Some Results in Universal Algebra

By

Alexander Wires

Dissertation

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

DOCTOR OF PHILOSOPHY

in

Mathematics

August, 2013

Nashville, Tennessee

Approved:

Professor Ralph McKenzie

Professor Constantine Tsinakis

Professor Steve Tschantz

Professor Mark Ellingham

Professor Yaqiong Xu

To my love, my wife,

Fengying

ACKNOWLEDGMENTS

I would like to express my sincere gratitude to my advisor, Ralph McKenzie. Whereas I knew only how to wade through, he could part the waters. The mathematical waters break around him and where he moves it is always high tide.

I would also like to thank him for his suggestion on how to prove that connectedness is a definable property(III.3.4). With this result in hand, much of the later results in that chapter could be simplified and concluded. If they still remain painful to slog through, I have only myself to blame.

I would like to thank the Vanderbilt math department and those at the helm during my stay here for providing an excellent research environment.

Finally, my eternal thanks goes to my family who have never failed in their support of me.

TABLE OF CONTENTS

LIST OF FIGURES

Figure | Page

CHAPTER I

DEFINABILITY IN SUBSTRUCTURE ORDERING

I.1 Introduction

In the paper [JM09a], Jaroslav Ježek and Ralph McKenzie introduce the following general situation. Let K be a fixed T-class of structures over a finite signature where T denotes a fixed type of axioms such as equations, quasi-equations, or universal sentences. Let $\mathcal{L}_K$ denote the collection of T-subclasses of K which usually forms a complete lattice ordered by inclusion. We may investigate if any of the following conditions are met by the first-order structure $\mathcal{L}_K$:

1. the finitely generated T-subclasses are definable in the language of lattices, and each such T-subclass is definable up to the automorphisms of $\mathcal{L}_K$;
2. the finitely axiomatizable T-subclasses are definable in the language of lattices, and each such T-subclass is definable up to the automorphisms of $\mathcal{L}_K$;
3. the classes axiomatizable by a single T-axiom is a definable subset of $\mathcal{L}_K$
4. the only automorphisms of $\mathcal{L}_K$ are the "obvious" ones.

If all of the above conditions are met we say that the T-theories of K has *positive definability*.

For an example of an "obvious" automorphism, let P be a poset. Reversing the direction of the ordering produces a new partial order over the same universe, denoted by P^{op}, where

$$a <_P b \quad \text{iff} \quad a >_{P^{op}} b.$$

The map $\mathbf{op} : P \to P^{op}$ takes substructures to substructures and can be seen to preserve the relation of embedding among posets. This means **op** induces a non-trivial automorphism of the partially ordered set of finite isomorphism types ordered by isomorphic substructures(embeddability); moreover, it is a non-trivial automorphism when restricted to lattices and distributive lattices. In a series of papers [JM09a], [JM10], [JM09b], [JM09c], Ježek and McKenzie investigated first-order definability in the substructure relation restricted to the finite isomorphism types of various subclasses of ordered sets. One of the principal results gathered from the separate papers is the following.

Theorem I.1.1. ([JM09a], [JM10], [JM09b], [JM09c]) Let $\mathcal{U}$ denote either the class of posets, lattices, or distributive lattices. Let $\langle \mathcal{PU}, \leq \rangle$ denote the poset of finite isomorphism types in the class $\mathcal{U}$ ordered by embeddability. Then there exists a single type $\mathbf{c} \in \mathcal{PU}$, such that every element of $\langle \mathcal{PU}, \leq, \mathbf{c} \rangle$ is first-order definable; moreover, **op** is the only non-trivial automorphism of $\mathcal{PU}$.

If *Sem* denotes the poset of finite isomorphism types of meet-semilattices ordered by embeddability, then every type is first-order definable in the order relation of *Sem*; in particular, *Sem* has no non-trivial automorphisms.

For meet-semilattices, posets, and distributive lattices, the Birkhoff dual of the substructure poset of finite isomorphism types is isomorphic to the lattice of universal subclasses. Using this isomorphism and the previous result, Ježek and McKenzie established positive definability for these universal theories. Lattices do not form a locally finite class, and so the question of positive definability for the universal theory of lattices remains tantalizingly open.

The case of ordered sets presents an additional remarkable result. In [JM10], it was shown that the connected finite isomorphism types of ordered sets were first-order definable in the substructure relation. It is a classical result of finite model theory, that among finite posets, the property of being connected is not first-order finitely axiomatizable; that is, there is no first-order sentence ψ in the language of ordered sets such that the finite models of ψ are precisely the connected ordered sets. In this regard, the substructure theory may capture strictly second-order properties. Indeed, it was shownby Jezek and McKenzie[JM10, Thm. 3.8] that in a certain sense, second-order finite axiomatizability among finite posets is equivalent to first order definability in the poset of finite isomorphism types ordered by substructure.

Whereas the previous work analyzed sublasses of ordered sets, the present work extends the theory of positive definability and definability in the substructure relation to the unordered structures of simple graphs (which are irreflexive symmetric digraphs) in Chapter III and equivalence relations in Chapter IV. In both cases, positive definability of the universal theories is achieved by establishing an analogue of Theorem I.1.1. We go further and characterize the expressive power of first-order definability in the substructure relation as equivalent to modeling full second-order properties when restricted to the finite members.

I.2 Outline of argument

In Chapter II, we establish the relationship between the lattice of universal subclasses and the poset of finite isomorphism types ordered by embeddability. The goal of this chapter is to provide a general template for the investigation of positive definability of universal theories, and so we pursue the material in an abstract setting and hope to be relatively thorough. The method does not work for arbitrary universal classes, and so the additional assumptions we must make are those which are necessary to preserve the approach established in [JM09a], [JM10], and [JM09b]. The chapter culminates with the proof of Theorem II.1.9 which guarantees positive definability of the universal theories provided we can first prove each finite type is definable in the poset of finite isomorphism types after adding a finite number of certain constants to the poset language.

An excellent and economical introduction to those aspects of model theory and first-order logic which are utilized in this dissertation can be found in the first two sections of chapter 5 in [BS81]. I will assume the basic familiarity with first-order logic and structures which can be found there.

In Section II.2, we briefly introduce the long-standing Reconstruction conjectures for finite ordered sets and simple graphs. We observe, as in [JM10], that the question of positive definability provides an excellent application for the affirmative resolution of these conjectures.

In Chapter III, we explore definability for finite isomorphism types of simple graphs under the substructure relation. Definability in this poset appears quite expansive, and we shall be able to conclude its elementary theory is undecidable and non-finitely axiomatizable. We establish the hypothesis required in Theorem II.1.9 and thereby conclude positive definability for the universal theories of simple graphs. This exploration culminates in Proposition III.7.1 of Section III.7 where it is shown each isomorphism type is definable. In the remaining part of the chapter, we establish the connection between definability in the poset of finite isomorphism types with the second-order language of graphs. The connection passes through the first-order language of a small category which is introduced and examined in Section III.9. This development follows very closely that of [JM10, Sec.3] which was the original inspiration for this work.

In Chapter IV, we turn our attention to the universal theory of equivalence relations. The poset of finite isomorphism types ordered by the substructure relation at first appears quite transparent, and so it is somewhat surprising, at least to the author, that we are able to interpret arithmetic and thereby conclude its elementary theory is both undecidable and non-finitely axiomatizable. Positive definability here follows more readily than in the case of simple graphs since Set Reconstruction holds for equivalence relations. In order to establish the connection between definability in the poset of finite isomorphism types and the second-order language, we must do more work. This culminates in Section IV.5.

Then in Chapter V, we look at some questions which arise from the work of the previous chapters.

CHAPTER II

UNIVERSAL CLASSES

In this chapter, we are concerned with classes of structures modeled by universal sentences. We will establish some conditions on a universal class for which positive definability of the universal theories is reduced to investigating definability in the isomorphic substructure relation of its finite members. When we refer to a class of structures $\mathcal{K}$ over a fixed signature, we always assume it is closed under taking isomorphic structures.

II.1 Definability in $\mathcal{L}_{\mathcal{U}}$

For a fixed signature, a first-order formula $\phi(x_1,...,x_n)$ is said to be open if it contains no quantifiers. A formula is in *prenex* form if it looks like

$$Q_1 y_1 \cdots Q_m y_m \psi(x_1,...,x_n)$$

where each Q_i is a quantifier, some of the y_i's may refer to the variables x_j, and $\psi(x_1,...,x_n)$ is an open formula. A standard result guarantees that every formula is logically equivalent to some formula in prenex form, which provides a canonical description for choosing interesting species of formulas. We may define one such species by saying that a formula is *universal* if it is logically equivalent to a prenex formula with only universal quantifiers. Recall, for any positive integer n there is a first-order sentence $\exists x_1 \cdots \exists x_n \Psi_{\geq n}(x_1,...,x_n)$ using only existential quantifiers such that for any structure B in any signature, $B \vDash \exists x_1 \cdots \exists x_n \Psi_{\geq n}$ iff $|B| \geq n$; for example, we can take for Ψ the following open formula

$$\Psi(x_1,...,x_n) := \bigwedge_{i<j} x_i \not\approx x_j.$$

Then $\forall x_1 \cdots \forall x_n \neg \Psi_{\geq n}(x_1,...,x_n)$ is a universal sentence which asserts a structure has at most $n-1$ elements.

For a set of first-order sentences Θ in some fixed signature, $\mathrm{Mod}(\Theta)$ is the class of structures in the same signature which satisfy every sentence of Θ. We say $\mathcal{U}$ is a *universal class* if $\mathcal{U} = \mathrm{Mod}(\Theta)$ for some set of universal sentences Θ. Universal classes can be described in an alternate manner. For a class $\mathcal{R}$, $S(\mathcal{R})$ will denote the class of structures isomorphic to substructures of structures in $\mathcal{R}$. The class $P_U(\mathcal{R})$ will consist of those structures isomorphic to ultraproducts of structures from $\mathcal{R}$. We shall make use the following characterization of universal classes.

Theorem II.1.1. [BS81, Thm 2.20] For any class of structures over a fixed signature, the following are equivalent:

1. $\mathcal{K}$ is a universal class
2. $\mathcal{K}$ is closed under S and P_U

3. $\mathcal{K} = SP_U(\mathcal{K}^*)$ for some class $\mathcal{K}^*$

For a non-empty class of structures $\mathcal{K}$, the universal class generated by $\mathcal{K}$, denoted by $\mathcal{U}(\mathcal{K})$, is by definition the smallest universal class containing $\mathcal{K}$. The previous theorem implies we may take $\mathcal{U}(\mathcal{K}) = SP_U(\mathcal{K})$.

A class $\mathcal{K}$ is said to be *locally finite* if all finitely generated substructures of structures in $\mathcal{K}$ are finite. For $\mathcal{R} \subseteq \mathcal{K}$, we say $\mathcal{U}(\mathcal{R}) \cap \mathcal{K}$ is the universal class generated by $\mathcal{R}$ relative to $\mathcal{K}$. If $\mathcal{R}$ =Mod(Θ) $\cap$ $\mathcal{K}$ for some set of universal sentences Θ, then $\mathcal{R}$ is called a universal class relative to $\mathcal{K}$. It is easy to see that both simple graphs and ordered sets form locally finite universal classes. As an ordered structure in a single binary relational signature, lattices do not form a universal class since they are not closed under substructure. In the signature of two binary operations, meet and join, lattices do indeed form a universal class, but it is not a locally finite.

Lemma II.1.2. For a locally finite class $\mathcal{K}$ closed under substructures, the relative universal subclasses are determined by their finite members.

`Proof`: Let $\mathcal{R}$, $\mathcal{P}$ be relative universal subclasses of $\mathcal{K}$. If they do not have the same finite members, then clearly $\mathcal{R} \neq \mathcal{P}$.

Suppose $\mathcal{R} \neq \mathcal{P}$. Then without loss of generality, there exists $P \in \mathcal{P}$ (and also in $\mathcal{K}$) such that $P \notin \mathcal{R}$, and so there exists a universal sentence ϕ such that $P \nvDash \phi$, but ϕ is satisfied by every structure in $\mathcal{R}$. Since only finitely many variables appear in ϕ, there exist $\{a_1, ..., a_n\} \subseteq P$ such that $D = Sg^P(a_1, ..., a_n) \leq P$ is a finite structure and $D \nvDash \phi$. This implies $D \notin \mathcal{R}$. •

Let $\mathcal{K}_{fin}$ denote the class of finite structures in $\mathcal{K}$.

Lemma II.1.3. For a class $\mathcal{K}$ of structures of finite signature, we have $\mathcal{U}(\mathcal{K})_{fin} \subseteq S(\mathcal{K})_{fin}$.

`Proof`: From Theorem II.1.1, we have $\mathcal{U}(\mathcal{K}) = SP_U(\mathcal{K})$. Let $B \leq \prod_U P_i$ be a finite substructure of an ultraproduct from $\mathcal{K}$. For all $\left\{\overline{a}_1, ..., \overline{a}_n, \overline{a}_{n+1}, \overline{b}_1, ..., \overline{b}_m\right\} \subseteq B$, operation symbol f and relation symbol R, if in B we have

$$f(\overline{a}_1/U, ..., \overline{a}_n/U) = \overline{a}_{n+1}/U \text{ and } R(\overline{b}_1/U, ..., \overline{b}_n/U),$$

then by definition of the ultraproduct the sets

$$[\![f(\overline{a}_1, ..., \overline{a}_n) = \overline{a}_{n+1}]\!] = \{i : P_i \vDash f(\overline{a}_1(i), ..., \overline{a}_n(i)) = \overline{a}_{n+1}(i)\}$$

$$[\![R(\overline{b}_1, ..., \overline{b}_n)]\!] = \left\{i : P_i \vDash R(\overline{b}_1(i), ..., \overline{b}_n(i))\right\}$$

all belong to the ultrafilter U. If $B = \{\overline{a}_1/U, ..., \overline{a}_p/U\}$, then since the signature is finite the intersection of all these sets for B together with $[\![\Psi_{\geq p}(\overline{a}_1, ..., \overline{a}_p)]\!]$ is non-empty and belongs to U. If k is a coordinate in this intersection, then it is straightforward to see that B embeds into P_k; therefore, $SP_U(\mathcal{K})_{fin} \subseteq S(\mathcal{K})_{fin}$. •

For a universal class $\mathcal{U}$, the universal classes contained in $\mathcal{U}$ may be ordered by containment; moreover, the order is a lattice order with meet given by intersection and the join of subclasses $\mathcal{K}$ and $\mathcal{V}$ given as

$\mathcal{K} \vee \mathcal{V} = SP_U(\mathcal{K} \cup \mathcal{V})$ by Theorem II.1.1. The lattice of universal subclasses of $\mathcal{U}$ is denoted by $\mathcal{L}_{\mathcal{U}}$. When we refer to definability in $\mathcal{L}_{\mathcal{U}}$, we refer to relations of the lattice $\mathcal{L}_{\mathcal{U}}$ definable by first-order formulas in the language of lattice theory.

For two structures A and B in the same signature, we write $A \leq B$ if A is embeddable into B; i.e., A is isomorphic to a substructure of B. For a non-empty class $\mathcal{K}$ in a finite signature, we may consider the relation $\leq$ of embeddability among the finite members. This naturally defines a quasi-order, and so by passing to the natural quotient by symmetric pairs, we arrive at the poset of finite isomorphism types $\langle \mathcal{PK}, \leq \rangle$ often denoted simply as $\mathcal{PK}$. Since elements of $\mathcal{PK}$ denote equivalence classes of isomorphic finite structures, we often refer to the order of $\mathcal{PK}$ as the substructure relation.

Recall for a poset P, the order ideals of P form a distributive lattice under the operations of set intersection and union. This lattice is the Birkhoff dual of P and is denoted by $\mathcal{O}(P)$.

Corollary II.1.4. For a locally finite universal class $\mathcal{U}$ of finite signature, we have the isomorphism $\mathcal{L}_{\mathcal{U}} \approx \mathcal{O}(\mathcal{PU})$.

`Proof`: For a universal subclass $\mathcal{K}$, we map $\phi(\mathcal{K}) = \mathcal{K}_{fin}$. Since $\mathcal{K}$ is closed under substructures, $\mathcal{K}_{fin}$ is an order ideal under $\leq$. Using Lemma II.1.3, for any universal sublcasses $\mathcal{K}$ and $\mathcal{R}$ we see that the join in the lattice of universal classes $\mathcal{K} \vee \mathcal{R} = SP_U(\mathcal{K} \cup \mathcal{R})$ implies $(\mathcal{K} \vee \mathcal{R})_{fin} = SP_U(\mathcal{K} \cup \mathcal{R})_{fin} \subseteq S(\mathcal{K} \cup \mathcal{R})_{fin} = S(\mathcal{K})_{fin} \cup S(\mathcal{R})_{fin}$; therefore, the map ϕ preserves joins. As the meet operation is just intersection in both lattices, ϕ is easily seen to preserve meets. Finally, Lemma II.1.2 guarantees ϕ is a bijection. •

In particular, the lattice of universal subclasses of $\mathcal{U}$ is complete. An element a in a lattice L is said to be strictly join-irreducible if whenever $a = \bigvee X$ for some set of lattice elements X, then $a \in X$. For any poset P, the strictly join-irreducible elements of $\mathcal{O}(P)$ are the principal order ideals. The property of an element being strictly join-irreducible is preserved under isomorphism.

Lemma II.1.5. In a complete lattice, an element x is strictly join-irreducible iff x has a unique lower cover. The set of strictly join-irreducible elements is first-order definable.

`Proof`: Let L be a complete lattice. Suppose x is strictly join-irreducible and set $Y = \{z \in L : z < x\}$. If $x = \bigvee Y = b$, then $x \in Y$; a contradiction. So $b < x$, and if $y < x$, then by definition we have $y \leq b$; therefore, b is the unique lower cover of x .

Suppose x has a unique lower cover $x^* \prec x$. If $x = \bigvee B$, then there exists $y \in B$ such that $y > x^*$; otherwise, $\bigvee B \leq x^* \prec x$, a contradiction. This means $y = x$ and so, x is strictly join-irreducible.

The property that x has a unique lower cover can be given a first order description; for example,

$$(\exists z)[(z < x) \wedge (x \not\leq z) \wedge ((\forall y)((y < x) \wedge (x \not\leq y) \rightarrow y \leq z))]$$

•

A class $\mathcal{R}$ of structures has the *finite embedding property* if for any finite set of finite structures $\{A_i\}_1^n \subseteq \mathcal{R}$, there is a finite structure $B \in \mathcal{R}$ such that $A_i \leq B$ for $1 \leq i \leq n$. Many classes have this property; for example, simple graphs, groups, rings, lattices, posets, and equivalence relations.

A universal class $\mathcal{R}$ is *finitely generated* if there is a finite set $\mathcal{J}$ of finite structures such that $\mathcal{R} = \mathcal{U}(\mathcal{J})$. A universal class $\mathcal{R}$ is *finitely axiomatizable* if there is a finite set Θ of universal sentences such that $\mathcal{R} = \mathrm{Mod}(\Theta)$. For a positive integer N, let $\mathcal{R}_N$ denote the subclass of N-generated substructures. In general, for a locally finite class and fixed N, all the N-generated substructures will be finite, but there is no a priori reason that their cardinalities should have a finite bound.

From this point on, we shall assume the universal class $\mathcal{U}$ is locally finite, of finite signature, has the finitely embedding property, and $\mathcal{U}_N$ is finite up to isomorphism for each N.

Lemma II.1.6. A universal subclass $\mathcal{K} \subseteq \mathcal{U}$ is finitely generated iff $\mathcal{K}$ is contained in a strictly join-irreducible universal subclass.

`Proof`: If $\mathcal{K} = SP_U(P_1,\ldots,P_n)$, let $P \in \mathcal{U}_{fin}$ such that each $P_i \hookrightarrow P$. Then $\mathcal{K} \subseteq SP_U(\{P\})$ and so, $\phi(K) = \mathcal{K}_{fin} \subseteq S(\{P\}) = (P]$; therefore, $\mathcal{K} \subseteq \phi^{-1}((P])$ which is strictly join-irreducible.

Likewise, if $\mathcal{K} \subseteq \mathcal{R}$ where $\mathcal{R}$ is strictly join-irreducible, then $\phi(\mathcal{R}) = (P]$ for some $P \in \mathcal{U}_{fin}$. This implies $\mathcal{K}_{fin} \subseteq \mathcal{R}_{fin} = S(P)$, and so $\mathcal{K}_{fin}$ is finite up to isomorphism. This implies $\mathcal{K}$ is finitely generated by Lemma II.1.2. $\bullet$

Lemma II.1.7. Let $\mathcal{U}$ be a finitely axiomatizable universal class. A universal subclass $\mathcal{K} \subseteq \mathcal{U}$ is finitely axiomatizable iff up to isomorphism there are only finitely many finite structures minimal under embeddability among structures of $\mathcal{U}$ outside of $\mathcal{K}$.

`Proof`: Let Θ be a finite set of universal sentences such that $\mathcal{U} = \mathrm{Mod}(\Theta)$. Suppose there exist finitely many finite structures $P_1,\ldots,P_n \notin \mathcal{K}$ such that for all $B \notin \mathcal{K}$ some $P_i \leq B$. Let $\boldsymbol{\Psi}$ be the sentence "I have a substructure isomorphic to some $P_1,\ldots,P_n$". Notice that $\boldsymbol{\Psi}$ can be taken to be a disjunction of existential sentences.

If $A \vDash \neg\boldsymbol{\Psi}$, then A does not embed any P_i and so by definition of the $P_1,\ldots,P_n$, we must have $A \in \mathcal{K}$.

If $A \in \mathcal{K}$, but $A \vDash \boldsymbol{\Psi}$, then some $P_i \leq A$ which implies $P_i \in \mathcal{K}$, a contradiction; thus, $A \vDash \neg\boldsymbol{\Psi}$.

Altogether we have that $\mathcal{K} = \mathrm{Mod}(\{\neg\boldsymbol{\Psi},\Theta\})$ and $\neg\boldsymbol{\Psi}$ is universal.

Conversely, suppose $\mathcal{K} = \mathrm{Mod}(\boldsymbol{\Sigma})$ where $\boldsymbol{\Sigma}$ is a finite set of universal sentences. We may assume $\mathcal{K}$ is a proper universal subclass. Then $\mathcal{K}^c \cap \mathcal{U} = \mathrm{Mod}(\bigvee\{\neg\phi : \phi \in \Sigma\}) \cap \mathcal{U}$. Let N denote the maximum number of variables used in $\neg\phi$ for all $\phi \in \boldsymbol{\Sigma}$. Then for any structure $A \in \mathcal{U}$, we have that $A \notin \mathcal{K}$ iff $A \vDash \neg\phi$ for some $\phi \in \boldsymbol{\Sigma}$ which implies there exists a finite substructure $B \leq A$ such that $B \vDash \neg\phi$ and B is at most N-generated.

We have shown that $A \notin \mathcal{K}$ iff $S(A) \cap (\mathcal{U}_N \cap \mathcal{K}^c) \neq \emptyset$. This shows the minimal structures of $\mathcal{U}$ outside of $\mathcal{K}$ are contained in $\mathcal{U}_N$ which, by hypothesis, is finite up to isomorphism. $\bullet$

Proposition II.1.8. Let $\mathcal{U}$ be a finitely axiomatizable universal class. The finitely generated and finitely axiomatizable universal subclasses of $\mathcal{U}$ are first-order definable in $\mathcal{L}_{\mathcal{U}}$.

`Proof`: By Lemma II.1.6 and Lemma II.1.5, "$\mathcal{K}$ such that $\mathcal{K} \subseteq \mathcal{R}$ for some $\mathcal{R}$ which has a unique lower cover" yields a first-order definition for the finitely generated universal subclasses. The proposition will be complete with the proof of the following claim.

`Claim`: $\mathcal{K}$ is finitely axiomatizable iff there exists a finitely generated class $\mathcal{N}$ such that $\forall \mathcal{M}$, $\mathcal{M} \nleq \mathcal{K} \Rightarrow \mathcal{M} \wedge \mathcal{N} \nleq \mathcal{K}$.

`Proof`: Suppose $\mathcal{K}$ is finitely axiomatizable. By Lemma II.1.7 let $P_1,...,P_n$ be a representative list of the finite minimal structures outside of $\mathcal{K}$. Then $\mathcal{N} = SP_U(P_1,...,P_n)$ is finitely generated. If $\mathcal{M} \nleq \mathcal{K}$, then there exists $R \in \mathcal{M}$, but $R \notin \mathcal{K}$ which implies some $P_i \in \mathcal{N} \cap \mathcal{M}$; therefore, the intersection is not empty and in particular, $\mathcal{N} \wedge \mathcal{M} \nleq \mathcal{K}$.

For the converse, suppose that $\mathcal{N}$ is is finitely generated by some $P_1,...,P_n$. Then $\mathcal{N} = SP_U(P_1,...,P_n) = S(P_1,...,P_n)$, and without loss of generality we may further assume $S(P_1,...,P_n) \subseteq \{P_1,...,P_n\}$. Let $\mathcal{R}$ be the class of structures of $(\mathcal{U} \cap \mathcal{K}^c)_{fin}$ minimal under substructure. For each $B \in \mathcal{R}$, set $\mathcal{K}_B = SP_U(\{B\})$. Then $\mathcal{K}_B \nleq \mathcal{K}$ implies $\mathcal{K}_B \cap \mathcal{N} \nleq \mathcal{K}$, and so there exists $F \in (\mathcal{K}_B \cap \mathcal{N})_{fin}$ such that $F \notin \mathcal{K}$. This means $F \in \{P_1,...,P_n\}$ up to isomorphism, and so by Lemma II.1.7, $\mathcal{K}$ is finitely axiomatizable.

This finishes the claim and the proposition. •

The above proposition characterizes the finitely axiomatizable universal classes as those classes $\mathcal{U}(F^c)$ where F is a finitely generated order filter in $\langle \mathcal{PU}, \leq \rangle$. The finitely generated universal classes are those classes $\mathcal{U}(I)$ where I is a finitely generated order ideal.

For any poset P, Aut$P$ will denote the group of automorphisms of the poset P. For any automorphism $\phi \in \mathrm{Aut}P$ and $J \in \mathcal{O}(P)$, the set $\phi(J) = \{\phi(a) : a \in J\}$ is an order ideal and so ϕ naturally determines an automorphism of $\mathcal{O}(P)$. In this way, the automorphisms of the poset $\mathcal{PU}$ induce automorphisms of the down-set lattice $\mathcal{O}(\mathcal{PU})$, and by using Corollary II.1.4, automorphisms of the lattice $\mathcal{L}_{\mathcal{U}}$. The next theorem provides conditions which guarantee these automorphisms are the only automorphisms of $\mathcal{L}_{\mathcal{U}}$.

Whenever a group G acts on a set X there is an induced action on n-tuples of elements $(x_1,...,x_n)^g = (x_1^g,...,x_n^g)$. To fix notation, let $\{(x_1,...,x_n)\}^G = \{(x_1,...,x_n)^g : g \in G\}$ denote the orbit of $(x_1,...,x_n)$ under G.

Theorem II.1.9. [JM10, Thm 2.35] Let $\mathcal{U}$ be finitely axiomatizable. Suppose there exist distinct finite isomorphism types $\mathbf{c}_1,...,\mathbf{c}_m$ such that each element of $\mathcal{PU}$ is definable in the pointed poset $\langle \mathcal{PU}, \leq, \mathbf{c}_1,...,\mathbf{c}_m \rangle$. Suppose the set

$$\{(\mathbf{c}_1,...,\mathbf{c}_m)\}^{\mathrm{Aut}(\mathcal{PU})}$$

is definable in $\mathcal{PU}$ without constants. Then each finitely generated universal subclass and finitely axiomatizable universal subclass of $\mathcal{U}$ is definable up to the automorphisms of $\mathcal{L}_{\mathcal{U}}$ induced by the automorphisms of $\mathcal{PU}$; moreover, the automorphisms induced by $\mathcal{PU}$ are the only automorphisms of $\mathcal{L}_{\mathcal{U}}$.

`Proof`: By Corollary II.1.4 and Proposition II.1.8, we may attain the result by showing individual definability of those order ideals $I \in \mathcal{O}(\mathcal{PU})$ which are finitely generated ideals or the complements of finitely generated filters.

Let I be an order ideal which is finitely generated or the complement of a finitely generated filter. We need to show that $\{I\}^{\mathrm{Aut}}$ is first order definable in $\mathcal{O}(\mathcal{PU})$. There are finitely many finite structures $P_1,...,P_n$ such that

$$I = \left\{B \in \mathcal{U}_{fin} : B \leq P_i \ \text{ for some } \ 1 \leq i \leq n\right\}$$

or

$$I = \left\{ B \in \mathcal{U}_{fin} : B \not\geq P_i \text{ for } 1 \leq i \leq n \right\}.$$

It is convenient to denote the principal order ideal generated by B as $B\downarrow$. Since the map $\Psi : \mathcal{PU} \to \mathcal{O}(\mathcal{PU})$ taking $\mathbf{c}$ to $\mathbf{c}\downarrow$ is an order-embedding, Lemma II.1.5 implies each $P_1\downarrow, \ldots, P_n\downarrow$ is definable in the pointed lattice $\mathcal{O}(\mathcal{PU})$ with constants $\mathbf{c}_1\downarrow, \ldots, \mathbf{c}_m\downarrow$. This means for each $1 \leq i \leq n$, there is a first-order formula $\psi_i(x, y_1, \ldots, y_m)$ in the language of lattices such that $P_i\downarrow$ is the unique element a_0 such that $\mathcal{O}(\mathcal{PU}) \models \psi_i(a_0, \mathbf{c}_1\downarrow, \ldots, \mathbf{c}_m\downarrow)$. Also, by hypothesis there is a first-order formula $\varepsilon(x_1, .., x_m)$ which defines the set $\{(\mathbf{c}_1\downarrow, \ldots, \mathbf{c}_m\downarrow)\}^{\mathrm{Aut}}$.

For ease of notation, let $+$ denote the lattice join operation in $\mathcal{O}(\mathcal{PU})$. Let $\Phi(x)$ be the formula

$$(\exists y \cdots \exists y_m)(\exists x_1 \cdots \exists x_n)[\varepsilon(y_1, \ldots, y_m) \wedge \bigwedge_{1 \leq i \leq n} \psi_i(x_i, y_1, \ldots, y_m) \wedge x = x_1 + \cdots + x_n]$$

and $\Theta(x)$ be the formula

$$(\exists y \cdots \exists y_m)(\exists x_1, \ldots, x_n)[\varepsilon(y_1, \ldots, y_m) \wedge \bigwedge_{1 \leq i \leq n} \psi_i(x_i, y_1, \ldots, y_m)] \wedge (\forall z)[z \leq x \leftrightarrow \bigwedge_{1 \leq i \leq n} x_i \not\leq z].$$

The claim is that $\mathcal{O}(\mathcal{PU}) \models \Phi(A)$ iff $A = \phi(I)$ for some $\phi \in \mathrm{Aut}\,\mathcal{O}(\mathcal{PU})$ where I is the order ideal generated by $P_1, \ldots, P_n$; and $\mathcal{O}(\mathcal{PU}) \models \Theta(B)$ iff $B = \phi(J)$ for some $\phi \in \mathrm{Aut}\,\mathcal{O}(\mathcal{PU})$ where J is the largest order ideal omitting each $P_1, \ldots, P_n$.

To verify the claim, suppose that $\mathcal{O}(\mathcal{PU}) \models \Theta(B)$. Let $Y_1, \ldots, Y_m$ and $X_1, \ldots, X_n$ be the elements which witness the satisfaction of $\Theta(B)$. Then we have that

$$\mathcal{O}(\mathcal{PU}) \models \varepsilon(Y_1, \ldots, Y_m)$$

$$\mathcal{O}(\mathcal{PU}) \models \psi_i(X_i, Y_1, \ldots, Y_m)$$

for $i = 1, \ldots, n$. It follows that $(Y_1, \ldots, Y_m) = (\mathbf{c}_1\downarrow, \ldots, \mathbf{c}_m\downarrow)^{\phi}$ for some automorphism ϕ.

Consider the order ideal $\phi^{-1}(B)$. Then $\mathcal{O}(\mathcal{PU}) \models \Theta(\phi^{-1}(B))$ with witnesses $\phi^{-1}(Y_i) = \mathbf{c}_i\downarrow$ and $\phi^{-1}(X_j)$ for $i = 1, .., m$ and $j = 1, \ldots, n$. The fact that $\mathcal{O}(\mathcal{PU}) \models \Theta(\phi^{-1}(B))$ implies $\phi^{-1}(B)$ is the largest element which fails to be above any element of the set $\{P_1\downarrow, \ldots, P_n\downarrow\}$; i.e., $\phi^{-1}(B)$ avoids each P_i and so $\phi^{-1}(B) = J$. So we have $B = \phi(J)$. Since it is straightforward to see that $\mathcal{O}(\mathcal{PU}) \models \Theta(J)$ and $\mathcal{O}(\mathcal{PU}) \models \Theta(\phi(J))$, we have shown J is definable up to the automorphisms.

Suppose that $\mathcal{O}(\mathcal{PU}) \models \Phi(B)$. Then as before, we have witnesses $Y_1, \ldots, Y_m$ and $X_1, \ldots, X_n$ such that $(Y_1, \ldots, Y_m) = (\mathbf{c}_1\downarrow, \ldots, \mathbf{c}_m\downarrow)^{\phi}$ for some automorphism ϕ.

We see that $\mathcal{O}(\mathcal{PU}) \models \Phi(\phi^{-1}(B))$ with witnesses $\phi^{-1}(Y_i) = \mathbf{c}_i\downarrow$ and $\phi^{-1}(X_j)$ for $i = 1, \ldots, m$ and $j = 1, \ldots, n$. Then $\mathcal{O}(\mathcal{PU}) \models \psi_j(\phi^{-1}(X_i), \mathbf{c}_1\downarrow, \ldots, \mathbf{c}_m\downarrow)$ implies $\phi^{-1}(X_j) = P_i\downarrow$ for $i = 1, \ldots, n$. We can also see that $\mathcal{O}(\mathcal{PU}) \models \Phi(\phi^{-1}(B))$ implies $\phi^{-1}(B)$ is the join of the principal ideals generated by the $\phi^{-1}(X_1), \ldots, \phi^{-1}(X_n)$; thus, $\phi^{-1}(B) = I$ which implies $B = \phi(I)$. Again, since it is clear that $\mathcal{O}(\mathcal{PU}) \models \Phi(I)$ and $\mathcal{O}(\mathcal{PU}) \models \Phi(\phi(I))$, we have shown the I is definable up to the automorphisms.

To finish the theorem, suppose σ is an automorphism of $\mathcal{O}(\mathcal{PU})$. Since the relation $\{(\mathbf{c}_1, \ldots, \mathbf{c}_m)\}^{\mathrm{Aut}(\mathcal{PU})}$

is definable, then by the above there exists some $\phi \in \text{Aut}(\mathcal{PU})$ such that $\sigma(\mathbf{c}_i\downarrow) = \bar{\phi}(\mathbf{c}_i\downarrow)$ for $i = 1,...,m$ where $\bar{\phi}$ is the unique automorphism of $\mathcal{O}(\mathcal{PU})$ induced by ϕ. Then the automorphism $\bar{\phi}^{-1}\circ\sigma$ fixes each $\mathbf{c}_i\downarrow$.

`Claim`: Any automorphism of $\mathcal{O}(\mathcal{PU})$ which fixes each $\mathbf{c}_i\downarrow$ is the identity.
`Proof`: Let τ be an automorphism of $\mathcal{O}(\mathcal{PU})$ which fixes each $\mathbf{c}_i\downarrow$. By the above, for each principal order ideal B there is a first-order formula $\psi(x, y_1, ..., y_m)$ in the language of lattices so that B is the unique element such that $\mathcal{O}(\mathcal{PU}) \vDash \psi(B, \mathbf{c}_1\downarrow, ..., \mathbf{c}_m\downarrow)$. Then $\mathcal{O}(\mathcal{PU}) \vDash \psi(B, \mathbf{c}_1\downarrow, ..., \mathbf{c}_m\downarrow)$ implies $\mathcal{O}(\mathcal{PU}) \vDash \psi(\tau(B), \tau(\mathbf{c}_1\downarrow), ..., \tau(\mathbf{c}_m\downarrow))$ and therefore, $\mathcal{O}(\mathcal{PU}) \vDash \psi(\tau(B), \mathbf{c}_1\downarrow, ..., \mathbf{c}_m\downarrow)$. This forces $\tau(B) = B$ and so τ fixes every principal order ideal. Since any order ideal is the complete join over the pincipal order ideals generated by its members, it must be the case that τ fixes every ideal, and so is the identity. This finishes the claim.

It follows from the claim that $\bar{\phi}^{-1}\circ\sigma = id$ and therefore, $\sigma = \bar{\phi}$. It now follows that the automorphisms of $\mathcal{L}_{\mathcal{U}}$ are just the induced automorphisms of $\mathcal{PU}$. •

In Chapters III and IV we will explore first-order definability in the substructure relation for simple graphs and equivalence relations. Each universal class is locally finite, of finite signature, has the finite embedding property, and has up to isomorphism only finitely many N-generated structures for each N. Furthermore, in the respective partially ordered sets of finite isomorphism types we will show that each isomorphism type is definable by adding an additional constant to the language. Since in each case, the orbits of the constants under automorphism will be definable without constants, the conditions of Theorem II.1.9 will be satisfied, and so we will have established positive definablity for each class.

II.2 Reconstruction

Let $\mathcal{D}$ be the class of digraphs; that is, structures in a single binary relation. For a finite digraph A and x a vertex of A, we let $A - x$ denote the substructure restricted on the remaining elements. If the digraph A has vertex set $V(A)$ and edge set $E(A)$, then $A - x$ is the digraph with vertex set $V(A)\setminus\{x\}$ and the edges are exactly thoses edges of $E(A)$ which are not incident with x. In the case of digraphs, our general notion of substructure is often referred to as an induced subdigraph. The *deck* of a digraph A is the multiset defined as $D(A) = \{A - x : x \in V(A)\}$ where we allow repeated elements in the case $A - x \approx A - y$ for $x \neq y$; in particular, the cardinality of $D(A)$ is equal to the cardinality of A. In the deck, we only consider the isomorphism types of the one vertex-deleted induced subdigraphs, so if A and B are isomorphic finite digraphs, then they will have the same decks. The general Reconstruction question asks if A and B have the same decks, then is it true that they are isomorphic?

If we consider a subclass $\mathcal{K} \subseteq \mathcal{D}$ closed under one-element deletions, then the reconstruction question can be posed for the finite digraphs in $\mathcal{K}$. It may be that reconstruction can be answered in the affirmative for all pairs of finite digraphs from $\mathcal{K}$ excluding a finite list of structures. In this case, we will slightly abuse terminology and say reconstruction holds for $\mathcal{K}$ if it is true for all structures larger than a predetermined cardinality; consequently, it is said to fail in $\mathcal{K}$ if it fails for an infinite family of pairs of structures.

The question of reconstruction appears in print for simple graphs in [Kel57] and for posets in [San85] (though in this paper it is mentioned the problem may stretch back almost a decade). In the case for graphs, it is easy to see that the path on three vertices and the disjoint union of a path on two vertices with an additional isolated vertex have the same decks, but are not isomorphic. Similarly, the three element poset with a unique bottom covered by two incomparable elements and its opposite have the same decks, but are not isomorphic. The reconstruction conjecture for graphs and posets states that these are the only exceptions; that is, reconstruction holds for simple graphs and posets of cardinality at least four. The truth of this conjecture is known for many subclasses of posets and graphs, but the full conjecture remains quite open. A good place to begin this topic is in the survey [Ram05].

In [Sto77] it was shown that reconstruction fails for digraphs. In fact, Stockmeyer showed the failure in the particular universal class of tournaments and then modified his construction in [Sto81] to give several lists of counterexamples for other digraphs.

There is a closely related question which appears in [Har64] which we describe in the general setting of this chapter. For any locally finite universal class $\mathcal{U}$ of finite signature, take $A \in \mathcal{PU}$ and consider the set $L_A = \{B \in \mathcal{PU} : B \prec A\}$; that is, L_A is the set of lower covers in the poset of finite isomorphism types ordered by the substructure relation. The Set Reconstruction question asks if

$$L_A = L_B \quad \Rightarrow \quad A = B \tag{II.1}$$

Again, by an abuse of terminology, set reconstruction is said to hold in a class if (II.1) holds after excluding finitely many counterexamples. One can see that set reconstruction implies reconstruction for those digraphs closed under one-element deletions since digraphs with the same deck have the same lower covers; posets, simple graphs, and tournamants are particular examples. As a result, it is unknown if set reconstruction holds for simple graphs and posets. Set reconstruction is known to hold for equivalence relations [PS04].

In the case of posets and graphs, the truth of set reconstruction may yield a quick proof of individual definability. For this section, a countable poset P is said to be *graded* if there exists a sequence of subsets $P_1, P_2, \ldots.$ which partition P such that each P_i is a non-empty set of incomparable elements, and every element of P covered by some element of P_{i+1} belongs to P_i. If $\mathcal{PU}$ is graded and each P_k consists of precisely those structures with cardinality k, then we say that $\mathcal{PU}$ is *graded by cardinality*.

Lemma II.2.1. Suppose Set Reconstruction holds in $\mathcal{PU}$ for structures of cardinality at least $m > 1$, and that $\mathcal{PU}$ is graded by cardinality. If all structures of cardinality at most $m-1$ are individually definable, then every structure is individually definable.

`Proof`: Since the poset is graded by cardinality, we may use induction on the cardinality of types to show each element at a given height is definable. Those structures of cardinality $m-1$ serve as the base of the induction. If each structure at height $n > m-1$ is definable, then by Set Reconstruction each structure A at height $n+1$ is uniquely determined by the set of lower covers L_A, and so inductively, is the unique element covering that particular first-order definable set L_A of elements at height n. ●

CHAPTER III

SIMPLE GRAPHS

We consider the class of simple graphs $\mathcal{G}$ which are those digraphs for which the binary edge relation is irreflexive and symmetric. Simple graphs form a locally finite universal class. For two vertices $u, v \in G$, the edge relation for u and v is denoted as $u \sim v$ and it is said that u and v are adjacent, or are connected by an edge.

We are concerned with the notion of embedding; that is, an injective map $\phi : G \to H$ such that for all vertices $u, v \in G$, $u \sim v$ iff $\phi(u) \sim \phi(v)$. Following the notation of Chapter II, we write $G \leq H$ if there is an embedding ϕ of G into H, and in this case, it is immediate that $\phi(G)$ is a substructure of H such that $G \approx \phi(G)$. For graphs, the general model-theoretic notion of substructure corresponds exactly to the definition of induced subgraph. If $\mathrm{QG}\textsc{r}$ denotes the finite simple graphs with vertex sets over the positive integers, then the embedding relation $\leq$ restricted to these graphs forms a quasi-ordered set $\langle \mathrm{QG}\textsc{r}, \leq \rangle$. The poset of finite isomorphism types $\mathcal{PG}$ ordered by substructure is then just the quotient of $\langle \mathrm{QG}\textsc{r}, \leq \rangle$ by the equivalence determined by isomorphism.

While we are interested in definability in the poset $\mathcal{PG}$, it will be more convenient to work within the quasi-ordered set $\langle \mathrm{QG}\textsc{r}, \leq \rangle$ where we will speak of graphs definable up to isomorphism rather than definable isomorphism types.

If $G < H$, but there does not exist $G < F < H$, then we write $G \prec H$ and say H covers G. It is easy to see that $G \prec H$ iff $G \leq H$ and $|H| = |G| + 1$. It follows that the poset $\mathcal{PG}$ is naturally graded according to cardinality, and so for each fixed positive integer n those graphs at height n (having cardinality n) are definable as having a maximal n-element chain in its principal order ideal. Notice this definition requires a fixed n, and so those graph with cardinality $n+1$ require a different package of formulas to define them. We shall see later in Section III.6 how to capture the cardinality of arbitrary graphs in a uniform manner.

The graph with a single vertex is the unique bottom element in $\mathcal{PG}$ below every other element. We do not consider the graph on an empty set of vertices.

When $A \leq B$, we will often without mention identify A with a particular induced subgraph U of B such that $U \approx A$. For example, if $A \prec B$, then we will say that B is formed from A by adding an additional vertex v to A and possibly some additional edges connecting v to vertices of A. If v is a vertex of G, then $G - v$ will denote the induced subgraph on the vertices of G omitting v; that is, the induced subgraph on the vertex set $V(G) - \{v\}$.

There is an obvious automorphism of $\langle \mathrm{QG}\textsc{r}, \leq \rangle$ (and of $\mathcal{PG}$) which is defined by edge complementation and denoted by σ; that is, $\sigma(G)$ is the graph over the same set of vertices as G, but $u \sim v$ in $\sigma(G)$ iff $u \not\sim v$ in G. We shall see in Section III.7 that this is the only non-trivial automorphism of $\mathcal{PG}$.

The *complete graph*, or *clique*, on m vertices is denoted as K_m and is characterized as the unique graph having every possible edge. The *empty graph*, or *trivial* graph, on m vertices is denoted as N_m and is characterized as the unique graph having no edges. It is easy to see that both $K_m \downarrow$ and $N_m \downarrow$ are chains. The *path* on n vertices is denoted by P_n and is a graph isomorphic to the graph $v_1 \sim v_2 \sim \cdots \sim v_n$ with no additional edges other than the ones specified. The *circuit*(or *cycle*) C_n is formed from the path P_n by adding only one additional edge $v_n \sim v_1$.

The set of complete or empty graphs is definable.

Lemma III.0.2. The set $\{K_m : m \geq 1\} \cup \{N_m : m \geq 1\}$ is definable in $\mathcal{PG}$.

Proof: A graph G is in the above set iff $G\downarrow$ is a chain.

To see this, suppose $G \not\approx K_n, N_k$ for any $n, k \geq 1$. Then $|G| \geq 3$ and so there exist vertices u, v, x, y such that $|\{u, v, x, y\}| \geq 3$ and $u \sim v$ and $x \not\sim y$. Then $K_2 \leq G$ and $N_2 \leq G$, but K_2 and N_2 are incomparable. •

For two graphs G and H, we form a new graph $G+H$ called the disjoint sum of G and H by taking the disjoint union of the two sets of vertices and allowing only those edges coming from G and H. This can be visualised as placing the graph G by the side of H. By construction, if $A \approx G$ and $B \approx H$, then $A+B \approx G+H$. We may consider the sum of more than two graphs, and so when taking many factors $\{G_i\}$ we can write the sum as $\sum G_i$. This will yield a convenient general notation for simple graphs.

Two vertices a and b in a graph G are *path-connected* if there is a path in G starting from a and ending at b; explicitly, if there is a sequence of vertices $a = x_1, \ldots, x_n = b$ in G such that $x_i \sim x_{i+1}$ for $i = 1, \ldots, n-1$. The graph G is *connected* if every two vertices of G are path-connected. We say an induced subgraph H of G is a connected component of G if H is connected, but no vertex of G outside of H is connected to any vertex of H; in this case we can write G as a disjoint sum of H and the induced subgraph on the remaining vertices. Naturally, any simple graph may be represented as $G \approx \sum G_i$ where G_i are the connected components of G.

Given two graphs G and H, we may construct a new graph $G \bigvee H$ called the join of G and H by taking $G+H$ and adding every possible edge of the form $u \sim v$ where $u \in G$ and $v \in H$. For example, $K_{p+q} \approx K_p \bigvee K_q$. Again, it is easy to see that if $A \approx G$ and $B \approx H$, then $A \bigvee B \approx G \bigvee H$.

The graph on four vertices with only the edges $u \sim v \sim x \sim y \sim u$ and $v \sim y$ will be denoted as $\mathbf{B}$.

The graph on four vertices with only the edges $u \sim v \sim x$ and $v \sim y$ will be denoted as $K_{1,3}$. This graph is often referred to as the *claw*.

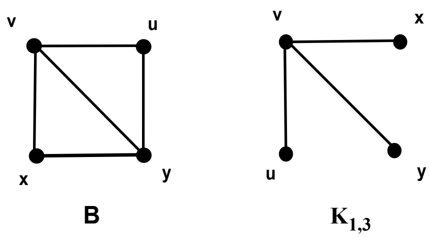

Figure III.1: Graphs $\mathbf{B}$ and $K_{1,3}$

At this point, we will add the constant P_3 representing the path on three vertices to form the pointed quasi-ordered structure $\langle \mathrm{QG}\textsc{r}, \leq, P_3 \rangle$. Unless otherwise specified, definability will refer to the language $\{\leq, P_3\}$.

Definition III.0.3. For $0 \leq k \leq n$, let $K_n +_k N_1$ denote the graph constructed from $K_n + N_1$ by arbitrarily adding k new edges. It is easy to see the graphs are isomorphic no matter how the k new edges are added, and so this produces a well-defined construction on isomorphism types. For example, $K_n +_n N_1 \approx K_{n+1}$.

If we denote the single vertex of N_1 by v, then we may define $N_n +_k N_1$ in a similar manner by adding k distinct edges $u \sim v$ where u is a vertex of N_n.

Remark III.0.4. For example, $\mathbf{B} \approx K_3 +_2 N_1$ and $K_{1,3} \approx N_3 +_3 N_1$.

Lemma III.0.5. For $m \geq 1$, the only covers of K_m are $\{K_m +_k N_1\}_{0 \leq k \leq m}$, and the only covers of N_m are $\{N_m +_k N_1\}_{0 \leq k \leq m}$

Proposition III.0.6. Every graph on at most four vertices is definable.

`Proof`: By Lemma III.0.2, the set $\{K_3, N_3\}$ is definable, and so then $\{P_3, P_2 + N_1\}$ is definable as the remaining graphs at height three; therefore, $P_2 + N_1$ is definable since we can call up the constant P_3.

Now P_4 is definable as the unique element covering P_3 and $P_2 + N_1$, but not covering anything in $\{K_3, N_3\}$.

The circuit C_4 is the element with the unique lower cover P_3.

We would like to define $P_3 + N_1$. Notice $P_3 + N_1$ and $K_3 +_1 N_1$ both cover P_3, $P_2 + N_1$, and some graph in $\{K_3, N_3\}$. We shall have recourse to uniquely define a graph which covers C_4 which will help separate $P_3 + N_1$ from $K_3 +_1 N_1$. First, we see that $C_4 + N_1$ can be defined as the unique cover of C_4 which has exactly two subcovers A and B, and if $A \approx C_4$, then $B \not\approx P_4$ and covers both P_3 and $P_2 + N_1$. We can then recover $P_3 + N_1$ as the unique subcover of $C_4 + N_1$ which is not isomorphic to C_4.

K_3 is then defined as the complete or empty graph at height three which is not N_3; consequently, K_4 is also definable. It follows that K_2 and N_2 are separately definable.

The graph $K_3 +_1 N_1$ is definable as the unique element with lower covers $P_3, P_2 + N_1$, and K_3.

$K_3 + N_1$ is the only cover of $P_2 + N_1$ which also covers K_3 and is not isomorphic to $K_3 +_1 N_1$.

$K_2 + K_2$ is the element with unique lower cover $P_2 + N_1$.

$K_2 + N_2$ is the unique element with only $P_2 + N_1$ and N_3 as lower covers.

The graph $\mathbf{B}$ is definable as the unique cover of P_3 not equal to $K_3 +_1 N_1$, but which also covers K_3.

The graph $K_{1,3}$ is definable as the element with P_3 and N_3 as the only lower covers.

This accounts for every type in $\mathcal{PG}$ of height at most four. •

Since K_3 and N_3 are separately definable, Lemma III.0.2 implies the sets of complete and trivial graphs are separately definable.

Definition III.0.7. For a graph A, let $cl(A) = n$ where K_n is largest clique which embeds in A, and $i(A) = m$ where N_m is the maximal trivial graph which embeds in A. A copy of N_m in A is called an independent set of A; thus, $i(A)$ is the size of a maximal independent set in A.

From the previous comments, $cl(\Gamma)$ and $i(\Gamma)$ are definable properties of Γ. It is also easy to see that $i(G+H) = i(G) + i(H)$.

To finish this section, we will show how to interpret the arithmetic of positive integers using disconnected cliques. Corollary III.0.9 references results of Section IV.2.

Proposition III.0.8. The set $\{\Gamma : \Gamma \text{ is a disjoint sum of cliques}\}$ is definable.

`Proof`: It is easy to see that Γ is a disjoint sum of complete graphs iff $P_3 \not\leq \Gamma$. $\bullet$

There is an obvious way to associate a disjoint sum of cliques to any equivalence relation; for the partition $\pi = (n_1, ..., n_t)$ consider the graph $\Gamma_\pi = \sum_{i=1}^{t} K_{n_i}$. It is easy to see that $\pi \leq \sigma$ iff $\Gamma_\pi \leq \Gamma_\sigma$; therefore, the above proposition implies $\langle \mathcal{PE}, \leq \rangle$ is definably present in $\mathcal{PG}$. Corollaries IV.2.15 and IV.2.16 yield the following.

Corollary III.0.9. The elementary theory of $\langle \mathcal{PG}, \leq \rangle$ is undecidable and not finitely axiomatizable.

III.1 Circuits, Paths and Trees

A graph H is said to contain a graph G as a subgraph if there is a subset of vertices V and some subset of edges E with vertices in V such that the graph with vertices V and edges E is isomorphic to G. This does not imply $G \leq H$. For example, **B** contains C_4 as a subgraph, but $C_4 \not\leq B$; that is, **B** does not contain C_4 as an induced subgraph.

A graph is said to be *acyclic* (or a *forest*) if it does not contain a circuit as a subgraph. The first result says we need only consider induced subgraphs.

Lemma III.1.1. Γ contains a circuit as a subgraph iff Γ contains a circuit as an induced subgraph.

`Proof`: Suppose $C \subseteq G$ is a circuit in G and write $C = x_1 \sim \cdots \sim x_n \sim x_1$. If $C \not\leq G$, then there exists $i+1 < j$, such that $x_i \sim x_j$, and so $C^{'} = x_1 \sim \cdots \sim x_i \sim x_j \sim \cdots \sim x_n \sim x_1$ is a smaller circuit. If we let $D \subseteq G$ be a circuit of minimal cardinality, then it follows that $D \leq G$. $\bullet$

Lemma III.1.2. $\{A : A \approx \sum C, \text{ C is a circuit }, |C| > 3\}$ is definable.

`Proof`: The claim is that A is in this set iff $|A| \geq 4$, and $P_3 \leq A$, $K_{1,3} \not\leq A$, $K_3 \not\leq A$, and A has a unique lower cover.

It is straightforward to see that these conditions are necessary. We must show that they are sufficient.

Suppose A is a graph which satisfies the conditions. If A is a circuit, then we are done. Assume A is not a circuit. Note that $P_3 \leq A$ implies A is not an empty graph.

`Claim`: A is the disjoint sum of circuits and paths.

`Proof`: First, note the maximum degree of every vertex in A is two. Suppose not. Let v be a vertex of A such that its neighborhood $N(v) = \{u : v \sim v\}$ has at least three vertices. If any two vertices of $N(v)$ are adjacent, then $K_3 \leq A$; a contradiction. Since $|N(v)| \geq 3$ and no vertices are adjacent, the induced subgraph on the vertices $N(v) \cup \{v\}$ embeds a copy of $K_{1,3}$; another contradiction.

Since every vertex has maximum degree two, it is not too hard to see that A must be the disjoint sum of circuits and paths. This finishes the claim.

Suppose A is the disjoint sum of k circuits and r paths. Let v be a vertex of some circuit and u a vertex of some path. Then $A - v$ has $k-1$ circuits, but $A - u$ still has k circuits. This contradictions the fact that A

has a unique lower cover. Suppose A is only the disjoint sum of t paths. Since $P_3 \leq A$, some path P in the sum of A has at least three vertices. Let x be a terminal vertex in P and let y be a vertex of degree two in P. Then $A-x$ is the disjoint sum of t paths, but $A-y$ is the disjoint sum of $t+1$ paths; a contradiction.

It must be the case that A is a disjoint sum of circuits. Since A has a unique lower cover, all the circuits must have the same length. •

For any $n \neq m$, we see that C_n and C_m are incomparable, and therefore, any two disjoint sums in the previous relation are comparable iff they are dsjoint sums over the same isomorphic circuit. In this case, the set of such sums over the same circuit are naturally linearly ordered according to the number of components; of course, the minimal elements are just the circuits.

Proposition III.1.3. The set of circuits is definable.

`Proof`: If we denote the definable set in Lemma III.1.2 as $CSUM$, then we have the following definable relation $R = \{(A,B) : A,B \in CSUM, A \leq B\}$. It follows from the above discussion that C is a circuit iff

$C \approx K_3$, or

$|C| > 3$, $C \in CSUM$, and $\forall B[(B,C) \in R \to B \approx C]$. •

We record the following corollary for use in the next section.

Corollary III.1.4. $\{(C,\Gamma) :$ C is circuit and $\Gamma \approx C+C\}$ is definable

`Proof`: (C,Γ) is in the relation iff C is a circuit and

$C \approx K_3$ and Γ is a disjoint sum of cliques, $cl(\Gamma) = 3$ and $i(\Gamma) = 2$, or

$C \approx C_4$ and C_4 is the unique circuit strictly below Γ, $K_{1,3} \not\leq \Gamma$, Γ has a unique lower cover, and $i(\Gamma) = 4$, or

$|C| > 4$, $\Gamma \in CSUM$, $C_m < \Gamma$, and there does not exist $R \in CSUM$ such that $C_m < R < \Gamma$.

The case where $C \approx C_4$ requires some explanation. The first two conditions imply Γ is a disjoint sum of copies of C_4 and possibly of some paths. Since Γ has a unique lower cover, there cannot be any paths present in the disjoint sum. The condition $i(\Gamma) = 4$ implies there are only two copies of C_4 in the sum, that is, $\Gamma \approx C_4 + C_4$. •

Corollary III.1.5. The set of forests is definable.

`Proof`: By definition, F is a forest iff it avoids every circuit. By Lemma III.1.1 and Proposition III.1.2, this is a definable condition. •

Proposition III.1.6. The set of paths is definable. In addition, the following two sets are definable:

1) $\{(A,P) : A \approx N_m$ for some m and $P \approx P_{2m}\}$

2) $\{(A,P) : A \approx N_m$ for some m and $P \approx P_{2m+1}\}$

`Proof`: The set of paths are just those elements which are the unique lower covers of a circuit.

1) (A,P) is in this relation iff $A \approx N_m$ for some m, P is a path, $i(P) = m$, and P covers a path L

such that $i(L) = m$.

2) (A,P) is in this relation iff $A \approx N_m$ for some m, P is a path, $i(P) = m$, and P covers a path L such that $i(L) = m-1$. •

By considering circuits which cover the appropriate paths, those circuits of even and odd cardinality are also separately definable.

Corollary III.1.7. $\{\Gamma : \Gamma$ is a disjoint sum of paths$\}$ is definable.

`Proof`: It is easy to see that a graph is strictly below some circuit iff it a disjoint sum of paths. The result now follows from Proposition III.1.3. •

Proposition III.1.8. The lattice of universal subclasses of the universal class generated by finite graphs has cardinality the continuum.

`Proof`: The set of isomorphism types of finite circuits constitutes a denumerably infinite antichain in $\mathcal{PG}$, and so by Corollary II.1.4 the collection of universal classes generated by arbitrary subsets of finite cycles forms a pair-wise distinct collection. Since universal classes of graphs are model classes over a finite signature, the cardinality is precisely the continuum. •

Proposition III.1.9. The class of non-finitely generated universal subclasses of simple graphs is equal to the union of the principal filters generated by $\mathcal{U}(\{K_m : m < \omega\})$ and $\mathcal{U}(\{N_m : m < \omega\})$.

`Proof`: Clearly, $\mathcal{U}(\{K_m : m < \omega\})$ and $\mathcal{U}(\{N_m : m < \omega\})$ are not finitely generated subclasses. Since simple graphs form a locally finite universal class, a universal subclass is non-finitely generated iff it contains infinitely many non-isomorphic finite simple graphs. The result is now just a direct application of Ramsey's Theorem. •

Definition III.1.10. For a,b in the same connected component of Γ, let $d(a,b)$ equal one less than the cardinality of the shortest path in Γ connecting a to b. The diameter of Γ is then taken to be $d(\Gamma) = \max\{d(a,b) : a,b$ in the same connected component of $\Gamma\}$. When a,b are in different connected components, set $d(a,b) = \infty$.

The distance between any two vertices a and b in the same connected component is always realized by some path, say $P = a \sim x_1 \sim \cdots \sim x_n \sim b$. If $x_i \sim x_j$ for some $j > i+1$, then we may construct a shorter path from a to b, contradicting the minimality of P; thus, the distance is always realized by an embedded path $P \leq \Gamma$. Since Γ is finite, the diameter is always realized by some path, and thus, by an embedded path. While it is not always true that $A \prec B$ implies $d(A) \leq d(B)$ - consider $A \approx P_4$ and $B \approx C_5$ - it is true for the class of forests.

A *tree* is a connected forest. In the case of forests, any induced path between two vertices is unique, and so the diameter of a forest is just the length of the largest induced subpath.

Proposition III.1.11. The set of trees is definable.

`Proof`: The claim is that T is a tree iff T is a forest and for any forest D such that $T \prec D$ we have $d(D) \leq d(T)+1$.

Suppose T is a tree, then whenever D is a forest such that $T \prec D$, we can construct D from T by adding a new vertex x and at most a single new edge $u \sim x$ where $u \in T$. Let $P \leq D$ realize the diameter of D. If $x \notin P$, then $d(T) \geq d(D)$. If $x \in P$, then x is a terminal vertex in the path P since it has degree one; therefore, we have a path $\overline{P} \leq T$ such that P is equal to adjoining x to the end of $\overline{P} \leq T$. Then $|\overline{P}| \leq d(T)$ which implies $d(D) = |P| \leq d(T)+1$.

Conversely, if F satisfies the conditions, then we may write $F = \sum_{i=1}^{m} F_i$ where each F_i is a tree. Note there exists $1 \leq k \leq m$ such that $d(F) = d(F_k)$. If $m > 1$, then choose $j \neq k$ and construct R, a cover for F, in the following manner: take F and a new vertex $v \notin F$, and add two new edges $a \sim v$ and $b \sim v$ where $a \in F_j$ and $b \in F_k$ such that b is an end-vertex of a path $P \leq F_k$ which realizes $d(F_k)$. It is easy to see that $d(R) > d(F_k)+1$; thus, we must have $m = 1$ which implies F is a connected forest. •

The above argument utilized the fact that the diameter for acyclic graphs was definable. If in general, the diameter of a graph was a definable property, then one would hope an argument similar to that of Proposition III.1.11 would yield the definability of the set of connected graphs. Explicitly, one would need the result that Γ is connected iff every upper cover increases the diameter by at most one. Unfortunately, this not true; one can find counterexamples among trees and their covers which are not forests. We will have to take a different approach to capture connected graphs in Section III.3.

III.2 Addition of Paths

It will be useful to do addition with paths instead of with cliques. The starting point is to observe that the lower covers of a path P_k are precisely the path P_{k-1} and the disjoint sums $P_r + P_t$ where $r+t = k-1$.

Lemma III.2.1. $\{(P,G) : P \approx P_m, G \approx P_m + P_m\}$ is definable.

`Proof`: The claim is that $G \approx P_m + P_m$ iff

$m = 1$ and $G \approx N_2$, or

$m \geq 2$ and $G \prec E \prec C_{m+1} + C_{m+1}$ for some E such that G is acyclic.

To see this one merely has to observe that any acyclic G such that $G \prec E \prec C_{m+1} + C_{m+1}$ must come from deleting a single vertex from each of the components in the sum. That $C_{m+1} + C_{m+1}$ is definable is precisely Corollary III.1.4. •

Proposition III.2.2. $\{(A,B,P) : A \approx P_n, B \approx P_m, P \approx P_k \text{ where } k = n+m\}$ is definable.

`Proof`: The claim is that $P \approx P_{n+m}$ iff P is a path, there exists a path R such that $P \prec R$, and there exists $G \prec R$ such that

(1) G is not a path

(2) $P_n \leq G$ and $P_m \leq G$

(3) If Q is a path such that $Q \leq G$, then $Q \leq P_n$ or $Q \leq P_m$

(4) If $P_m \leq P_n$ we have $P_m + P_m \leq G$, and if $P_n \leq P_m$ we have $P_n + P_n \leq G$

and for any path E such that $F \prec E$ and F satisfies (1) - (4), then $R \leq E$.

To show these conditions are sufficient, suppose P satisfies the conditions and is covered by a path R and $G \prec R$. Assume $R \approx P_s$ and so, $G \approx P_r + P_t$ where $r + t = s - 1$. Without loss of generality we may take $n \geq m$. If $r > n$ or $t > n$, then $P_{n+1} \leq G$, but $P_{n+1} \nleq P_n$ and $P_{n+1} \nleq P_m$ which contradicts (3); thus, $r, t \leq n$. By (2), $r = n$ or $t = n$, and so we may assume $r = n$. By (4), $t \geq m$ and so we can conclude that $n + m \leq s - 1 \leq n + n$. For $f - 1$ in the interval $[n + m, n + n]$, each P_f has a lower cover which satisfies (1) - (4), and so we must have $s - 1 = n + m$ by the requirement of minimality. This implies $P \approx P_{n+m}$.

Clearly, $P_n + P_m \prec P_{n+m+1}$ satisfies (1) - (4), and by the above argument any disjoint sum of two paths which satisfies (1) - (4) must be covered by a path P_s with $n + m + 1 \leq s \leq n + n + 1$ where $n \geq m$. This establishes the conditions are necessary, and completes the proof of the proposition. •

As a corollary we may establish the definability of the disjoint sum of two paths.

Corollary III.2.3. $\{(A, B, P) : A, B$ are paths and $P \approx A + B\}$ is definable.

`Proof`: $P \approx P_n + P_m$ iff $P \prec P_{n+m+1}$, and

(1) $P_n \leq P$ and $P_m \leq P$

(2) If Q is a path such that $Q \leq P$, then $Q \leq P_n$ or $Q \leq P_m$.

If P satisfies the conditions, then $P \approx P_{n+m}$ or $P \approx P_r + P_t$ where $r + t = n + m$. Without loss of generality, assume $n \geq m$. Since $P_{n+1} \leq P_{n+m}$, by (2) we see that $P \approx P_r + P_t$. Condition (1) implies $r \geq n$ or $t \geq n$. If $r > n$, or $t > n$, then $P_{n+1} \leq P$ and we arrive at a contradiction of (2); thus, $r = n$ or $t = n$ which implies $t = m$ or $r = m$, respectively.

As the necessity of the conditions is immediate, we have established the result. •

Since paths are the unique lower covers of circuits, we can also accomplish addition with the definable set of circuits in the obvious way.

Corollary III.2.4. $\{(A, B, C) : A \approx C_n, B \approx C_m, C \approx C_{n+m}\}$ is a definable relation.

III.3 Connectedness

In this section we will show the set of connected graphs is definable.

Lemma III.3.1. $\{(C, E) : C \approx C_m$ and $E \approx C_m + N_1\}$ is definable.

`Proof`: The claim is that $E \approx C_m + N_1$ iff

$m = 3$ and $E \approx K_3 + N_1$, or

$m > 3$ and $C_m \prec E$, $K_{1,3} \nleq E$, $K_3 \nleq E$.

The necessity of the conditions is immediate.

For sufficiency, suppose E satisfies the conditions and $E \not\approx K_3 + N_1$. Then we may form E from C_m by adding an additional vertex v and possibly some new edges connecting C_m to v. Suppose there exist $a, b \in C_m$ such that $a \sim v$ and $b \sim v$. If $a \sim b$, then the induced subgraph on the vertices $\{a, b, v\}$ is isomorphic to K_3, a contradiction. It must be the case that $a \not\sim b$, but then $K_{1,3} \leq E$ when we consider the induced subraph on the vertices $\{z, a, w, v\}$ where $z \sim a \sim w$ and $z, w \in C_m$, another contradiction. So there can be at most one new edge. Since $m > 3$ we see that $K_{1,3} \leq E$ if there is just one additional edge; therefore, $E \approx C_m + N_1$. •

Let $Path_{\geq 2}$ denote the set of graphs which are disjoint sums of paths with no isolated vertices.

Lemma III.3.2. *$Path_{\geq 2}$ is definable.*

`Proof`: We will show the set of graphs which are disjoint sums of paths with isolated vertices is definable, then the lemma will follow. Recall that the set of circuits forms an antichain under the substructure ordering; however, their unique subcovers, the set of paths, is linearly well-ordered. This implies there is a first-order definable well-ordering $\leq_*$ on circuits defined by

$$C \leq_* D \quad \text{if and only if} \quad P \leq Q$$

for circuits C and D and $P \prec C$ and $Q \prec D$.

The claim is that G is a disjoint sum of paths with isolated vertices iff

$|G| = 1$ and $G \approx N_1$, or

$|G| = 2$ and $G \approx N_2$, or

$|G| = 3$ and $G \approx N_3$, or $G \approx K_2 + N_1$, or

$|G| > 3$ and

(1) G is a disjoint sum of paths

(2) If C is a circuit such that $G \leq C$, and C is the smallest circuit D under $\leq_*$ such that $G \leq D$, then there exist circuits E and F such that $E \prec_* F \prec_* C$ and $G \leq E + N_1$.

The preceeding observations and Corollary III.1.7 guarantee these conditions are definable. For necessity, assume G is a disjoint sum of paths with isolated vertices and write $G \approx N_1 + \sum_{i=1}^{r} P_i$. If $n = \sum_{i=1}^{r} |P_i|$, then C_{n+r+2} is the circuit of smallest cardinality which embeds G. Then $\sum_{i=1}^{r} P_i \leq C_{n+r}$ and we set $E \approx C_{n+r}$ for condition (2).

For sufficiency, assume G satisfies the conditions and that $|G| > 3$. By (1), G is a disjoint sum of paths. It is easy to see that if G has no isolated vertices, then $G \leq C_k$ iff $G \leq C_k + N_1$. Let C be the smallest circuit under $\leq_*$ such that $G \leq C$. If G has no isolated vertices then using (2), $G \leq E$ where E is a circuit $E <_* C$; a contradiction. It must be the case that G has isolated vertices. •

Remark III.3.3. The first-order conditions for definability and the proof of the following proposition was suggested by Ralph McKenzie. This will provide for a rather simple way to capture connected graphs in Proposition III.3.6.

Proposition III.3.4. $\{(X, N, G) : N \approx N_m \ \text{ and } \ G \approx X + N_m\}$ is definable.

`Proof`: The claim is that $G \approx X + N_m$ iff

G is trivial, X is trivial and $i(G) = i(X) + m$, or

G is not trivial, and

(1) X is not trivial and $X < G$

(2) $i(G) = i(X) + m$

(3) For every circuit $\overline{C}$ there exists a circuit C such that

(a) $\overline{C} <_* C$

(b) There exists Γ such that

(i) $X \leq \Gamma$ and $C \leq \Gamma$, and for all $H \leq \Gamma$, $X \leq H$ and $C \leq H$ implies $H \approx \Gamma$

(ii) $G \leq \Gamma$

(iii) For all $R \leq \Gamma$ such that $C \prec R$, $C + N_1 \approx R$

(4) For all $B \in Path_{\geq 2}$, $B \leq G$ implies $B \leq X$

We first tackle the argument for sufficiency. Suppose G satisfies the conditions. We may assume G is not trivial; otherwise, definability follows from addition with trivial graphs which is provided by Proposition III.0.8 and the results of Section IV.2. We can represent G as $G \approx E + P$ where E is the disjoint sum of connected components which are not paths and P is the disjoint sum of all the connected components which are paths. Also, write $X \approx A + Q$ where A the disjoint sum of connected components which are not paths and Q is the disjoint sum of all the connected components which are paths. Suppose $n = |Q|$ and Q has r components (all of which are paths). Take $C_{n+r} = \overline{C}$ in (3). Let C be the circuit with $C_{n+r} <_* C$ given by (3a).

Let Γ be the graph whose existence is guaranteed by (3b). Condition (iii) implies that any copy of the circuit C in Γ must appear as a connected component, and since $C \leq \Gamma$, we can write $\Gamma \approx C + K$ for some sum of connected graphs K. Since $X \leq \Gamma$, we must have $A \leq K$. Because C^* was chosen large enough such that $Q \leq C$, we have $X \leq C + A \leq C + K \approx \Gamma$; thus, by (i) we have $\Gamma \approx A + C$. Note that

$$A + Q \approx X < G \approx E + P \leq \Gamma \approx A + C$$

implies $E \approx A$ and so, $G \approx A + P$ where $Q \leq P \leq C$.

We can further write $Q \approx F + N_t$ and $P \approx H + N_r$ where $F, H \in Path_{\geq 2}$. It is easy to see that whenever K is maximal among those graphs $\Phi \in Path_{\geq 2}$ such that $\Phi \leq X \approx A + Q \approx A + F + N_t$, then $K \approx J + F$ where $J \in Path_{\geq 2}$ and is maximal for $J \leq A$. Take $\overline{J} \in Path_{\geq 2}$ such that $\overline{J} \leq A$ and is of maximum cardinality. The condition $Q \leq P$ implies $F \leq H$. If $F < H$, then $\overline{J} + H \in Path_{\geq 2}$ and $\overline{J} + H \leq G$, and so by (4) we must have $\overline{J} + H \leq A + F$ which contradicts the choice of $\overline{J}$. It must be the case that $F \approx H$. Condition (2) implies $N_r \approx N_t + N_m$ and so,

$$G \approx A + P \approx A + H + N_r \approx A + F + N_t + N_m \approx A + Q + N_m \approx X + N_m.$$

To prove these conditions are necessary, assume X is not trivial and write $X \approx A + Q$ as before with $n = |Q|$ and r such that Q has r components (all of which are paths). Then for any $C_{n+r+2m} <_* C$ notice that $Q \leq C$. We may then take $\Gamma \approx A + C$ and it is straightforward to check conditions (3) and (4) are satisfied. •

The previous proposition actually yields more than is explicitly stated. What we have shown is that there is a first-order formula $\Psi(x,y,z,w)$ in the language of $\langle \mathrm{QG}\textsc{r}, \leq\rangle$ such that $\langle \mathrm{QG}\textsc{r}, \leq\rangle \vDash \Psi(A,N,G,P_3)$ iff N is trivial and $G \approx A+N$. If we apply the complementation automorphism σ we see that $\langle \mathrm{QG}\textsc{r}, \leq\rangle \vDash \Psi(B,K,H,\sigma(P_3))$ iff $\sigma(N) = K$ is complete and $H \approx \sigma(A+N) = \sigma(A) \bigvee \sigma(N) = B \bigvee K$. Since $\sigma(P_3) = K_2 + N_1$ is definable in $\langle \mathrm{QG}\textsc{r}, \leq, P_3\rangle$, there is a first-order formula $\gamma(x,w)$ such that $\langle \mathrm{QG}\textsc{r}, \leq\rangle \vDash \gamma(A,P_3)$ iff $A \approx K_2 + N_1$. We can then take the formula

$$\exists w \Psi(x,y,z,w) \wedge \gamma(w,P_3)$$

in order to define the join $X \bigvee K$ where K is complete.

Corollary III.3.5. $\{(X,K,G) : K \approx K_m \text{ and } G \approx X \bigvee K_m\}$ is definable.

An induced subgraph $A \leq G$ is called a *maximal connected component* iff A is connected and if $A < B \leq G$, then B is disconnected; in particular, a maximal connected component is a connected component. For example, if A and B are connected with $A \leq B$, then $G \approx A+B$ has only B as a maximal connected component.

Proposition III.3.6. The set of connected graphs is definable.

`Proof`: The claim is that G is connected iff there does not exist $B < G$ such that for all E, $B \prec E \leq G$ implies $E \approx B + N_1$.

Clearly, if G is disconnected with $G \approx B+H$ where B is a maximal connected component, then every cover F of B in G is of the form $F \approx B+N_1$.

If G is connected, then for every $B < G$ there exists $x \in G$ with $x \notin B$ but is adjacent to the connected component of B with largest cardinality. Then the induced subgraph on $B \cup \{x\}$ is certainly not isomorphic to $B+N_1$. ●

Since the property of being connected is definable, we can recognize the maximal connected components.

Lemma III.3.7. $\{(A,G) : A \text{ is a maximal connected component of } G\}$ is definable.

`Proof`: From the previous proposition and by the definition of maximal connected component. ●

The following lemma is the first step in showing the definability of the disjoint sum operation; however, it is such a specialized instance of a sum that we must do a little more preparation before we tackle the general case in Section III.5.

Lemma III.3.8. $\{(A,B,G) : G \approx A+B,$ A,B connected and incomparable $\}$ is definable.

`Proof`: The claim is that (A,B,G) is in the relation iff

(1) A and B are connected and incomparable

(2) A and B are maximal connected components of G

and G is smallest under $\leq$ among graphs satisfying (2). ●

The following sum will be useful in Section III.7.

Lemma III.3.9. $\{(C,D,\Gamma) : C \approx C_m, D \approx C_n$ for $n > m > 5,\ \Gamma \approx \sum_{k=m}^{n} C_k\}$ is definable.

`Proof`: The claim is that $\Gamma \approx \sum_{k=m}^{n} C_k$ iff every circuit C such that $C_m \leq_* C \leq_* C_n$ is a maximal connected component of Γ, and Γ is the smallest under $\leq$ with this property.

Since distinct circuits are incomparable, the argument from LemmaIII.3.8 can be applied here to establish the result. ●

III.4 Martians and Other Useful Graphs

For a circuit C_n we may construct the graph $C_n \rightarrow_1$ by adding only one new edge $u \sim x$ where x is some new vertex and u is an arbitrary vertex of C_n. Different choices of u result in isomorphic graphs, and so the construction is well-defined on isomorphism types.

The same construction for K_n in place of C_n yields a special case of definition 4.2; in this case we have $K_n \rightarrow_1 = K_n +_1 N_1$.

Lemma III.4.1. $\{(K,\Gamma) : K \approx K_n$ and $\Gamma \approx K_n \rightarrow_1\}$ is definable.

`Proof`: The claim is that $\Gamma \approx K_n \rightarrow_1$ iff

$n = 1$ and $\Gamma \approx K_2$, or

$n > 1$ and Γ covers K_n, Γ is connected, but Γ has a disconnected subcover.

To see this, suppose Γ satisfies the conditions and $n > 1$. Then $K_n \prec \Gamma$ implies $\Gamma \approx K_n +_k N_1$ for some $1 \leq k \leq n$. Since Γ is connected we must have $k \geq 1$. If $k \geq 2$, then every lower cover of Γ is connected which yields a contradiction; therefore, $k = 1$. That the conditions are necessary is immediate. ●

In the same way we have the definability of the graphs $C_n \rightarrow_1$.

Lemma III.4.2. $\{(C,\Gamma) : C \approx C_n$ and $\Gamma \approx C_n \rightarrow_1\}$ is definable.

`Proof`: $\Gamma \approx C_n \rightarrow_1$ iff $C_n \prec \Gamma$, Γ is connected, but Γ has a disconnected subcover. ●

Lemma III.4.3. $\{(C,\Gamma) :$ C is a circuit and $\Gamma \approx C + K_2\}$ is definable.

`Proof`: The claim is that $\Gamma \approx C_m + K_2$ iff

$m = 3$ and $\Gamma \approx K_3 + K_2$, or

$m > 3$ and there exists F such that $C_m \prec F \prec \Gamma$, $i(\Gamma) = i(C_m) + 1$, and C_m is a maximal component of Γ.

It is straightforward to see that $C_m + K_2$ satisfies the criteria.

For sufficiency, suppose Γ satsifies the conditions and $\Gamma \not\approx K_3 + K_2$. We may construct Γ from C_m by adding two new vertices u and v, and possibly some new edges. Since C_m is maximal component of a subcover, u and v are not connected to any vertices of C_m. The condition $i(\Gamma) = i(C_m) + 1$ implies $u \sim v$; threfore, $\Gamma \approx C_m + K_2$. $\bullet$

The graph $C \rightarrow_2$ refers to the cover of $C \rightarrow_1$ formed by adding an additional vertex and only one additional edge joining the new vertex to the unique vertex of degree one in $C \rightarrow_1$. The graph $K_n \rightarrow_2$ is defined in a similar manner.

Lemma III.4.4. $\{(C, \Gamma) : C$ is a circuit and $\Gamma \approx C \rightarrow_2\}$ is definable.

`Proof`: The claim is that $\Gamma \approx C \rightarrow_2$ where C is a circuit iff

(1) $C + N_1 \prec \Gamma$

(2) Γ is connected

(3) Γ has a disconnected acyclic subcover

We shall only verify sufficiency. Suppose Γ satisfies conditions (1) - (3). Then Γ may be constructed from $C + N_1$ by adding an additional vertex v and possibly some new edges joining v to C. If u denotes the isolated vertex of $C + N_1$, then condition (2) implies we have edges $v \sim u$ and $v \sim x$ for some $x \in C$. If v is adjacent to any other vertices of C, then every acyclic subcover is connected, a contradiction of (3); thus, v is adjacent to only one vertex of C which implies $\Gamma \approx C \rightarrow_2$. $\bullet$

Proposition III.4.5. $\{(K, \Gamma) : K \approx K_n$ and $\Gamma \approx K_n \rightarrow_2\}$ is definable.

`Proof`: The claim is that $\Gamma \approx K_n \rightarrow_2$ iff

$n = 1$ and $\Gamma \approx P_3$, or

$n = 2$ and $\Gamma \approx P_4$, or

$n \geq 3$ and $K_n \rightarrow_1 \prec \Gamma$, $K_n + N_1 \prec \Gamma$, $P_4 \leq \Gamma$, $K_{n+1} \not\leq \Gamma$, and $C_4 \not\leq \Gamma$.

It is easy to see these conditions are satisfied by $K_n \rightarrow_2$.

Suppose Γ meets these conditions, and we may assume $n \geq 3$. Since $K_n + N_1 \prec \Gamma$, we can construct Γ from $K_n + N_1$ by adding a new vertex v and possibly new edges of the form $x \sim v$ where $x \in K_n + N_1$. Since $K_{n+1} \not\leq \Gamma$, there exists $y \in K_n$ such that $y \not\sim v$. Since $K_n \rightarrow_1 \prec \Gamma$, there exists $u \in K_n$ such that $u \sim v$, and v is not adjacent to any other vertex of K_n. If z denotes the solitary vertex of N_1, then the condition $P_4 \leq \Gamma$ implies $v \sim z$, and this demonstrates that $\Gamma \approx K_n \rightarrow_2$. $\bullet$

We describe a certain cover of the sum $C + D$ when C and D are incomparable circuits. Let $\overline{C + D}$ denote the cover of $C + D$ formed by adding an additional vertex v and only two new edges connecting v to C and v to D. The choices of vertices in C and D which are adjacent to v is immaterial since every choice results in isomorphic graphs.

Lemma III.4.6. $\{(C, D, \Gamma) : C, D$ are incomparable circuits and $\Gamma \approx \overline{C + D}\}$ is definable.

`Proof`: The claim is that $\Gamma \approx \overline{C_n + C_m}$ iff

(1) $C_n + C_m \prec \Gamma$

(2) Γ is connected

(3) Γ has a disconnected subcover $E \prec \Gamma$ such that $C_n \not\leq E$

(4) Γ has a disconnected subcover $F \prec \Gamma$ such that $C_m \not\leq E$

Suppose Γ satisfies conditions (1) - (4). Then (1) implies Γ may be constructed by addiing an additional vertex v to $C_n + C_m$ and some new edges incident with v. Since Γ is connected, there must be edges connecting v to C_n and v to C_m. Take $n \geq m$ and suppose there are at least two vertices in C_n adjacent to v, then no mater how v is connected to C_m, every subcover which avoids C_n is connected; a contradiction with (3). A similar argument shows v is adjacent to only one vertex of C_m; therefore, $\Gamma \approx \overline{C_n + C_m}$.

For necessity, it is easy to see the only subcovers of $\overline{C_n + C_m}$ which avoid C_n or C_m are $P_{n-1} + C_m$ or $C_n + P_{m-1}$, respectively. •

Lemma III.4.7. $\{(C,D,G) : C,D$ are incomparable circuits and $G \approx C \rightarrow_1 +D\}$ is definable.

`Proof`: The definability of this relation follows from Lemma III.3.8 and Lemma III.4.2. •

Lemma III.4.8. The relation

$$\{(C,D,G) : \text{C and D are incomparable circuits and } G \approx C \rightarrow_1 +D \rightarrow_1\}$$

is definable.

`Proof`: Notice that when C and D are incomparable circuits, $C \rightarrow_1$ and $D \rightarrow_1$ must also be incomparable. The result now follows from Lemma III.3.8. •

Let $\gamma(n,m)$ denote the graph formed by adding a single new edge connecting the two unique vertices of $C_n \rightarrow_1 +C_m \rightarrow_1$ which have degree one.

Proposition III.4.9. The relation

$$\{(C,D,\Gamma) : C,D \text{ are incomparable circuits and } \Gamma \approx \gamma(|C|,|D|)\}$$

is definable.

`Proof`: The claim is that $\Gamma \approx \gamma(n,m)$ iff

(1)$C_n \rightarrow_1 +C_m \prec \Gamma$

(2) Γ is connected

(3) If $C_m \prec R \leq \Gamma$, then $R \approx C_m \rightarrow_1$ or $R \approx C_m + N_1$

(4) If $C_n \rightarrow_1 \prec R \leq \Gamma$, then $R \approx C_n \rightarrow_2$, or $R \approx C_n \rightarrow_1 +N_1$.

We only verify sufficiency. Suppose Γ satisfies conditions (1) - (4). From condition (1), we can form G by adding an additional vertex v to $C_n \rightarrow_1 +C_m$ and possibly some new edges incident with v. Since Γ is connected by (2), there is at least one new edge connecting v to the copy of C_m in $C_n \rightarrow_1 +C_m$, but by condition (3) there can be exactly one such edge. Again by (2), there is at least one edge connecting v to

$C_n \rightarrow_1$; however, by condition (4), $C_n \rightarrow_2$ is the only possibility for a connected induced subgraph on the vertices of $C_m \rightarrow_1 \cup \{v\}$. We have shown $\Gamma \approx \gamma(n,m)$. •

Definition III.4.10. For $n \geq 1$, a *martian* $M(n)$ is constructed from the two graphs K_n and $K_{1,3}$ by identifying a single vertex of K_n with a single vertex of $K_{1,3}$ which has degree one. Note the choice of the vertex in K_n is immaterial and so the construction is well-defined on isomorphism types. A *p-martian*, denoted by $pM(n)$, is constructed from $M(n)$ by connecting the remaining two vertices of degree one in $K_{1,3}$; thus connecting the "antennae".

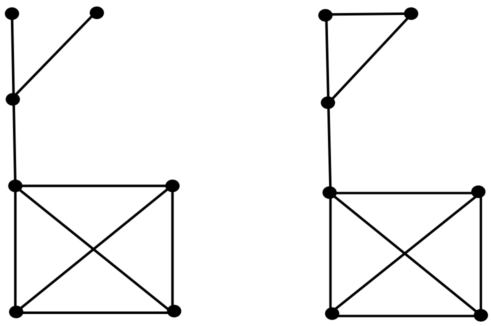

M(4) **pM(4)**

Figure III.2: A martian and p-martian

Before we show the definability of martians and p-martians, we need to show the definability of two auxiliary families of graphs.

Lemma III.4.11. $\{(K,R,G) : K \approx K_n, R \approx K_m$ and $G \approx K_n + K_m\}$ is definable.

`Proof`: We take $n \geq m$. The claim is that $G \approx K_n + K_m$ iff G is a disjoint sum of cliques with $cl(G) = n$, $i(G) = 2$, and

(1) If $n = m$, then G has a unique lower cover, or

(2) If $n > m$, then there exists $B < G$ such that B has a unique lower cover with $cl(B) = m$ and $i(B) = 2$, and whenever $A < G$ such that $i(A) = 2$ and A has a unique lower cover, then $cl(A) \leq m$.

If G satisfies the conditions, then $G \approx K_n + K_r$ for some $r \leq n$. Notice G has a unique lower cover iff $n = r$. This is the content of condition (1). Notice $K_t + K_t < G$ iff $t \leq m$. This is the content of condition (2). •

Lemma III.4.12. $\{(K,G) : K \approx K_n$ and $G \approx K_n + P_3\}$ is definable.

`Proof`: The claim is that $G \approx K_n + P_3$ iff

$n = 1$ and $G \approx P_3 + N_1$, or

$n = 2$ and $G \approx K_2 + P_3$, or

$n > 2$ and (K_n, P_3, G) is in the relation of Lemma III.3.8. •

Proposition III.4.13. $\{(K,M) : K \approx K_n \text{ and } M \approx M(n)\}$ is definable.

`Proof`: The claim is that $M \approx M(n)$ for $n \geq 2$ iff

$n = 1$ and $M \approx K_{1,3}$

$n = 2$ and M is the unique cover of $K_{1,3}$ which is acyclic and embeds P_4, or

$n \geq 3$ and

(1) $K_n + N_2 \prec M$, $K_n \rightarrow_2 \prec M$, and these are the only lower covers of M which embed K_n

(2) $P_4 \leq M$

(3) $K_{n-1} + P_3 \prec M$

To show these conditions characterize martians we need only check the cases for $n > 1$.

Suppose M satisfies the conditions and $n = 2$. We may construct M by adding a new vertex v to $K_{1,3}$ and possibly new edges incident with v. Let x be the unique vertex of $K_{1,3}$ with degree three, then $v \sim x$ implies $P_4 \not\leq K_{1,3}$ no matter what other edges are present; thus, $v \not\sim x$. Again, since $P_4 \leq M$, v is adjacent to a vertex of $K_{1,3}$ with degree one, but v must be adjacent to exactly one such vertex since M is acyclic. This finishes the demonstration that $M \approx M(2)$.

Suppose $n \geq 3$ and M satsifies (1) - (3). Since $K_n + N_2 \prec M$, we may construct M by adding an additional vertex v to $K_n + N_2$ and possibly some new edges incident with v. Let w and x be the vertices comprising this copy of N_2. By (2), we must have a new edge $a \sim v$ for some $a \in K_n$ and, without loss of generality, an edge $v \sim x$. If there exists $b \in K_n$ such that $b \neq a$ and $b \sim v$, there is no possibility for $K_{n-1} + P_3 \prec M$, a contradiction of (3); thus, the induced subgraph on $K_n \cup \{v\}$ is isomorphic to $K_n \rightarrow_1$. Suppose $w \not\sim v$. Then M has the three lower covers $K_n \rightarrow_2$, $K_n + N_2$, and $K_n \rightarrow_1 + N_1$ which all embed K_n, a contradiction of (1); therefore, we must have an edge $w \sim v$ and conclude that $M \approx M(n)$.

For necessity, it is easy to see that $M(2)$ is an acyclic cover of $K_{1,3}$ which embeds P_4.

Since $M(n)$ has a unique copy of K_n for $n \geq 3$, the only lower covers of $M(n)$ which embed K_n must come from deleting the vertices not included in K_n. In this case, $K_n \rightarrow_2$ and $K_n + N_2$ are the only such covers. Conditions (2) and (3) are immediate. Altogether we have shown (1) - (3) characterizes these martians. •

Proposition III.4.14. $\{(K,M) : K \approx K_n \text{ and } M \approx pM(n)\}$ is definable.

`Proof`: The claim is that $M \approx pM(n)$ iff

$n = 1$ and $M \approx K_3 \rightarrow_1$, or

$n = 2$ and $K_3 \rightarrow_2$, or

$n \geq 3$ and $K_n + K_2 \prec M$, $P_4 \leq M$, $P_5 \not\leq M$, and $\mathbf{B} \not\leq M$.

It is straightforward to check that each of these conditions must hold for the appropriate p-martian, so we shall concentrate on demonstrating that they are sufficient to characterize these graphs. Suppose M satisfies the conditions, and we may assume $n > 2$.

Then $K_n + K_2 \prec M$ implies we may construct M by adding an additional vertex x to $K_n + K_2$ and possibly some new edges. Since $P_4 \leq M$, but $P_5 \not\leq M$, we must have that x is adjacent to each vertex of K_2, and that x is adjacent to at least one vertex $a \in K_n$ and there exists $b \in K_n$ such that $x \not\sim b$. If x is adjacent to an additional vertex $c \in K_n$ distinct from a, then the induced subgraph on vertices $\{c,b,a,x\}$ is isomorphic to **B**, a contradiction; thus, there are no additional edges and we see that $M \approx pM(n)$. This finishes the proposition. ●

III.5 G + H

In this section we will prove the definability of the operation $G+H$ where G and H are arbitrary graphs. The starting point for this development will be the construction and definability of special connected graphs called pointed sums.

Definition III.5.1. Given two connected graphs A and B, the *pointed sum* is a graph $A+_p B$ formed by adding a new vertex v to $A+B$ and two new edges incident to v; one edge connects v to a vertex in A, and the other edge connects v to a vertex in B. Different choices of vertices in A and B lead to non-isomorphic graphs which are still considered as pointed sums; therefore, the notation $A+_p B$ will refer to the finite family of pointed sums for the different choices of vertices in A and B which are adjacent to the added vertex which has degree two.

We can see that when both A and B are complete, or both are circuits, then the choices of vertices in the definition is immaterial, and in these cases the family of pointed sums collapses to a unique graph. To make the connection with previous notation, for circuits C and D, $C+_p D \approx \overline{C+D}$. While we do not have definability of general pointed sums, we do have definability in a very specific and useful case.

Lemma III.5.2. The relation

$$\{(A,K,G) : A \text{ connected and not a clique}, K \approx K_n,\ n > cl(A)+1,\ G \in A+_p K\}$$

is definable.

`Proof`: Note that A is not complete and $n > cl(A)+1$ implies A and K_n are incomparable, and so by Lemma III.3.8, $A+K_n$ is definable.

The claim is that (A,K,G) is in the relation iff

(1) A is connected and not a clique

(2) $K \approx K_n$ for some n such that $n > cl(A)+1$

(3) $A+K_n \prec G$

(4) G is connected

(5) If $K_n \prec R \leq G$, then $R \approx K_n + N_1$ or $R \approx K_n \rightarrow_1$.

(6) $M(n) \not\leq G$ and $pM(n) \not\leq G$

The proof of necessity is straightforward and so we will establish sufficiency.

Suppose G satisfies the criteria. By (3), G may be constructed from $A+K_n$ by adding a new vertex p and perhaps some new edges incident with p. Note $n \geq 3$. Let V be the induced subgraph on the vertices of K_n together with p. Since G is connected, the vertex p is connected to some vertex of this copy of K_n; thus, $R \approx K_n \rightarrow_1$. Again, since G is connected there is at least one new edge $p \sim a$ with $a \in A$. Suppose

$p \sim b$ where $b \in A$ and $b \neq a$. If $b \sim a$, then the induced subgraph on $K_n \cup \{p,a,b\}$ is isomorphic to $M(n)$, a contradiction. If $b \not\sim a$, then we have a copy of $pM(n) \leq G$, another contradiction; therefore, p is adjacent to at most one vertex of A, and so we conclude that $G \in A +_p K_n$. •

Proposition III.5.3. $\{(A,\Gamma) : A$ is connected and $\Gamma \approx A+A\}$ is definable.

`Proof`: If A is a clique, then the definability of $A+A$ is guaranteed by Lemma III.4.11. We may suppose A is not a clique, and thus, $|A| > 2$. Let $m = cl(A)+3$.

Since A and K_m are connected and incomparable, $A+K_m$ is definable by Lemma III.3.8, and from Corollary III.3.5, $(A+K_m)\bigvee N_1$ is definable. Then by Lemma III.5.2 we have that $A+_p K_{m+1}$ is definable where, in a slight abuse of notation, $A+_p K_{m+1}$ will refer to any of the graphs in the set represented by the pointed sum $A+_p K_{m+1}$. By choice of m it is easy to see that $(A+K_m)\bigvee N_1$ and $A+_p K_{m+1}$ must be incomparable; therefore, by Lemma III.3.8 the disjoint union $(A+K_m)\bigvee N_1 + (A+_p K_{m+1})$ is definable.

We claim that $A+A+K_{m+1}+K_m$ is the unique graph Γ so that there exists E such that

(1) $\Gamma \prec E \prec (A+K_m)\bigvee N_1 + (A+_p K_{m+1})$

(2) $K_m \to_1 \not\leq E$ and K_{m+1} is a maximal connected component of E

(3) A and K_{m+1} are the only maximal connected components of Γ

To see this, set $G \approx (A+K_m)\bigvee N_1 + (A+_p K_{m+1})$. By (1), there exist vertices z and w such that $\Gamma \approx G-z-w$. In the construction of $A+_p K_{m+1}$, a new vertex v was added to the sum $A+K_{m+1}$ and an edge connecting v to a vertex of K_{m+1}. Let a denote this vertex of K_{m+1}. If $a \in \{z,w\}$, then K_{m+1} cannot appear as a maximal connected component; a contradiction of (3). If $v \notin \{z,w\}$, then $K_m \to_1$ embeds in every subcover of G, or K_{m+1} is not a maximal connected component; a contradiction of (2). Without loss of generality, we may take $v = z$. If we let q denote the unique vertex of N_1 in the construction of $(A+K_m)\bigvee N_1$ which is connected to every vertex of $A+K_m$, then (3) implies we must have $q = w$. Then $G-v-q \approx A+A+K_{m+1}+K_m$.

We will now see how to recover $A+A$ from $A+A+K_{m+1}+K_m$. This is the purpose of the following claim which will complete the proposition.

`Claim`: Consider the following property for a graph H:

$$(**)\ H+N_2 \leq A+A+K_{m+1}+K_m \quad \text{but} \quad H+N_3 \not\leq A+A+K_{m+1}+K_m$$

The graph $A+A$ is the unique graph among those maximal under $\leq$ for property $(**)$, which have A as the only maximal connected component.

`Proof`: Let X be maximal for property $(**)$ and having A as the only maximal connected component. We may write $X \approx A+X_2+\cdots+X_n$ where X_i are the connected components of X. Since $X+N_2 \leq A+A+K_{m+1}+K_m$, it must be the case that $X_2+\cdots+X_n+N_2 \leq A+K_{m+1}+K_m$. Let $G_2+\cdots+G_n+U$ be an induced subgraph of $A+K_{m+1}+K_m$ such that each $G_i \approx X_i$ and $N_2 \approx U$, and $X_2+\cdots+X_n+N_2 \approx G_2+\cdots+G_n+U \subseteq A+K_{m+1}+K_m$. The graph $G_2+\cdots+G_n+U$ fixes a copy of $X_2+\cdots+X_n+N_2$ in $A+K_{m+1}+K_m$.

If $U \subseteq K_{m+1}+K_m$, then $G_2+\cdots+G_n \subseteq A$ which implies $n=2$ by maximality; therefore, $G_2 \approx A$ and so $X \approx A+A$. We show this is the only possible case.

If $U \subseteq A + K_{m+1}$ but $U \nsubseteq A$, then $G_2 + \cdots + G_n \subseteq A + K_m$. Only a single connected component G_i can be in an induced subgraph in K_m, and so by maximality there is a component, say G_2, isomorphic to K_m. This implies $G_3 + \cdots + G_n \subseteq A$. But $m = cl(A) + 3$ implies K_m is a maximal connected component of X; a contradiction. The same argument for $U \subseteq A + K_m$ shows there exists some $G_i \approx K_{m+1}$ which yields another contradiction.

If $N_2 \subseteq A$, then some $G_i \subseteq K_{m+1}$ and maximality again shows we must have some $G_i \approx K_{m+1}$; a contradiction.

This finishes the claim and the proposition. $\bullet$

Lemma III.5.4. $\{(A,B,\Gamma) : A,B$ connected and $A < B,\ \Gamma \approx A + B\}$ is definable.

`Proof`: If B is a clique, then so is A and we already have definability of their sum. Assume B is not a clique and set $m = cl(B) + 3$. Then any two graphs from $A +_p K_{m+1}$ and $B +_p K_m$ are incomparable. By the same argument as in the previous lemma, we have that the sum $\Gamma \approx A + B + K_{m+1} + K_m$ is definable. We can then recover $A + B$ from Γ in a similar way, as well.

$A + B$ is the unique graph H which has B as the only maximal connected component and is maximal under $\leq$ for the property that $H + N_2 \leq \Gamma$ and $H + N_3 \nleq \Gamma$. To see this, let $H \approx H_1 + \cdots + H_r$ be such a graph with connected components H_i. We may take $H_1 \approx B$ and consider $B + H_2 + \cdots + H_r + N_2 \leq A + B + K_{m+1} + K_m$. Let $V + G_2 + \cdots + G_n + U$ be an induced subgraph of $A + B + K_{m+1} + K_m$ such that each $G_i \approx H_i$, $N_2 \approx U$, and $V \approx B$, and $B + H_2 + \cdots + H_n + N_2 \approx V + G_2 + \cdots + G_n + U \subseteq A + B + K_{m+1} + K_m$.

If $N_2 \approx U \nsubseteq K_{m+1} + K_m$, then by maximality some component G_i must intersect $K_{m+1} + K_m$. Again by maximality, we can conclude that $H_i \approx K_m$ or $H_i \approx K_{m+1}$; a contradiction. It must be that $U \subseteq K_{m+1} + K_m$ and intersects both cliques which implies $G_2 + \cdots + G_r \leq A$; thus, maximality implies $r = 2$ and so $H_2 \approx A$. Altogether, this shows $H \approx A + B$. $\bullet$

Putting these last two results together we can conclude the definability of a sum of two connected graphs..

Proposition III.5.5. $\{(A,B,\Gamma) : A,B$ connected and $\Gamma \approx A + B\}$ is definable.

We should note a useful property of the join construction. If V is a disconnected graph, then $V \bigvee N_1$ is connected and has a unique disconnected subcover; namely, if $U \prec V \bigvee N_1$ is disconnected, then $U \approx V$.

Lemma III.5.6. The relation

$$\{(U,V,\Gamma) : V \text{ disconnected},\ U \text{ a maximal connected component of } \Gamma,\ \Gamma \approx U + V\}$$

is definable.

`Proof`: (U,V,Γ) is in the relation iff

(1) V is disconnected and $V \leq \Gamma$

(2) U is a maximal connected component of Γ

(3) If M is a maximal connected component of Γ, then $M \approx U$ or $M \leq V$

(4) $\Gamma \prec U+V \bigvee N_1$

(5) Γ is not isomorphic to the disjoint sum of two connected graphs.

Since necessity is straightforward to check, we only prove sufficiency.

Suppose Γ satisfies conditions (1) - (5). By condition (4), $\Gamma \approx \Gamma'$ where Γ' is an induced subgraph of $U+V \bigvee \{v\}$, and there exists a vertex $x \in U+V \bigvee \{v\}$ such that $U+V \bigvee \{v\} - x = \Gamma'$. We will show Γ' is isomorphic to $U+V$ by considering the possible choices for the vertex x.

Suppose $x = v$. Then $V \bigvee \{v\} - x = V$ and so $\Gamma' = U+V \bigvee \{v\} - x = U+V$.

If $x \in V$, then $V \bigvee \{v\} - x$ is connected;. This implies $\Gamma' = U+V \bigvee \{v\} - x$ is the disjoint sum of two connected graphs, a contradiction of condition (5); therefore, $x \notin V$.

Suppose $x \in U$. Since $V \bigvee \{v\}$ is a connected subset of Γ', we must have that $V \bigvee \{v\} \leq M$ for some maximal connected component of $\Gamma' \approx \Gamma$. By condition (3), we have $V \bigvee \{v\} \leq M \approx U$ or $V \bigvee \{v\} \leq M \leq V$. Since $|V \bigvee \{v\}| > |V|$, we must have $V \bigvee \{v\} \leq U$. But $\Gamma' = (U-x)+V \bigvee \{v\}$ and condition (5) implies $U-x$ is disconnected. Since $|U-x| < |U|$, condition (2) implies $U \leq V \bigvee \{v\}$ which then yields $U \approx V \bigvee \{v\}$.

Since $U-x$ is disconnected, then $U-x \approx V \bigvee \{v\} - x \approx V$ implies $\Gamma' = (U-x)+(V \bigvee \{v\}) \approx V+U$. ●

Lemma III.5.7. $\{(A,P,\Gamma) : P \text{ a path, } P \nleq A, \text{ and } \Gamma \approx A+P\}$ is definable.

`Proof`: If A is a path or is just connected, then we already have the definability of $A+P$. If A is disconnected, then the definability of $A+P$ follows from Lemma III.5.6 since the condition $P \nleq A$ implies P is a maximal connected component of $A+P$. ●

Proposition III.5.8. $\{(A,B,\Gamma) : \Gamma \approx A+B\}$ is definable.

`Proof`: Let P be a path such that $P \nleq A$ and $P \nleq B$ and $|P| > 3$. Set $H = (A+P) \bigvee N_1 + (B+P) \bigvee N_1$ which is definable from A and B using Lemma III.5.7, Proposition III.5.5, and Corollary III.3.5. The claim is that $A+B+P+P$ is the unique graph G such that

(1) $G \prec E \prec H$ for some E

(2) P is a maximal connected component of G

(3) $P+P \leq G$

(4) $P \bigvee N_1 \nleq G$

To see this, assume G satisfies conditions (1) - (4). We can write $H = A' \cup B' \cup P' \cup P'' \cup \{p,q\}$ where $A' \approx A$, $B' \approx B$, $P' \approx P'' \approx P$, and $A' \cup P' \cup \{p\} \approx (A'+P') \bigvee N_1$ and $B' \cup P'' \cup \{q\} \approx (B'+P'') \bigvee N_1$. Then by (1), $G = H - \{u,v\}$ for some vertices u,v. If neither p nor q is in $\{u,v\}$, then P cannot be a maximal connected component of G; a contradiction of condition (2). We may assume, without loss of generality, $q = v$. Then $G = (A'+P') \bigvee N_1 + B' + P'' - u$. If $u \in B'$ or $u \in P''$ or $u \in A'$, then G we have $P \bigvee N_1 \nleq G$ which contradicts condition (4). If $u \in P'$, then $P+P \nleq G$ since any copy of P in $(A'+P') \bigvee N_1 - u$ will contain p and two consecutive vertices of P' which induce a copy of K_3; a contradiction of condition (3). It must be the case that $u = p$ which implies

$$G = (A' + P') \bigvee N_1 + B' + P'' - p = A' + P' + B' + P'' \approx A + B + P + P.$$

We may then use Lemma III.5.7 to capture $A + B$ as the unique graph F such that $F + P + P \approx G$. ●

III.6 Cardinality

In this section, we establish the cardinality of graphs as a definable property. The first step is to generalize the construction $\overline{C + D}$ where C and D are incomparable circuits. The graph $\overline{\sum_{k=m}^{n} C_k}$ is the cover of the sum $\sum_{k=m}^{n} C_k$ with one new vertex v and a new edge connecting v to each circuit.

Lemma III.6.1. $\left\{(C, D, \Gamma) : C \approx C_m, D \approx C_n \text{ for } n > m > 5,\ \Gamma \approx \overline{\sum_{k=m}^{n} C_k}\right\}$ is definable.

`Proof`: The claim is that $\Gamma \approx \overline{\sum_{k=m}^{n} C_k}$ iff

(1) $\sum_{k=m}^{n} C_k \prec \Gamma$

(2) If $C_i \prec R \leq \Gamma$, then $R \approx C_i + N_1$, or $R \approx C_i \rightarrow_1$ for $m \leq i \leq n$

(3) Γ is connected

Necessity is immediate. For sufficiency, notice (1) and (3) imply Γ is formed by adding a vertex v to $\sum_{k=m}^{n} C_k$ and at least one new edge $v \sim x_i$ with $x_i \in C_{m+i}$ for $m \leq i \leq n$. Condition (2) implies exactly one new edge $v \sim x_i$ is added. ●

Proposition III.6.2. $\{(K, A) : K \approx K_n \text{ and } |A| = n\}$ is definable.

`Proof`: It suffices to characterize when $|A| \geq n$. The claim is that $|A| \geq n$ iff A is a clique and $K_n \leq A$, or A is not a clique, $cl(A) = m$, and for every graph P with the following properties we must have $K_{n+1} \leq P$:

(1) $A \bigvee K_1 \leq P$

(2) For every clique K and graph Q such that $K_m < Q \leq A$ and $Q \bigvee K \leq P$, there exists Q' such that $K_m \leq Q' \prec Q$ and a clique $\overline{K}$ with $K \prec \overline{K}$ such that $Q' \bigvee \overline{K} \leq P$.

Assume first that A is not a clique and that $|A| \geq n$. Assume P satisfies conditions (1)-(3). By induction on k for $0 \leq k \leq |A| - m$, we argue that $Q \vee K_{k+1} \leq P$ for some graph Q such that $K_m \leq Q \leq A$ with $k = |A| - |Q|$. We see that (1) yields the base case $k = 0$ with $Q = A$. Condition (3) is applied at the inductive step for $1 \leq k < |A| - m$ to show $Q' \vee K_{k+1} \leq P$ for some $K_m \leq Q'$. At the step $k = |A| - m = |A| - |Q|$, we see that $K_m \leq Q$ implies $K_m \approx Q$ which yields

$$P \geq Q \vee K_{k+1} \approx K_m \vee K_{|A|-m+1} \approx K_{|A|+1} \geq K_{n+1}.$$

Now assume A is not a clique and $|A| < n$. Let $Q_1, \ldots, Q_p$ be a full list, up to isomorphism, of all graphs

Q such that $K_m \leq Q \leq A$. Let $r_i = |A| - |Q_i|$. We see that $cl(Q_i) = m$. Set

$$P = \sum_{i=1}^{p} Q_i \vee K_{r_i+1}.$$

Note $cl(Q_i \vee K_{r_i+1}) = m + r_i + 1$ and $cl(P) = max\{cl(Q_i \vee K_{r_i+1}) : i = 1, ..., p\}$. The clique size of P is determined by the component of maximum clique size which occurs when r_i is largest; that is, when $|Q_i| = |K_m| \Rightarrow Q_i \approx K_m$. For simplicity, let this occur at $i = 1$ and so we have

$$Q_1 \vee K_{r_1+1} \approx K_m \vee K_{|A|-m+1} \approx K_{|A|+1}.$$

Thus, $cl(P) = |A| + 1 \leq n \Rightarrow K_{n+1} \not\leq P$. It remains to show that P satisfies (1) and (2).

Since A is not a clique, condition (1) is immediately seen to hold by construction. For (2), suppose $K_m < Q \leq A$ and K is a clique such that $Q \vee K \leq P$. Since $Q \vee K$ is connected, we must have $Q \vee K \leq Q_i \vee K_{r_i+1}$ for some $i \in [p]$. Let $K \approx K_l$. Then

$$cl(Q) + l = cl(Q \vee K) \leq cl(Q_i \vee K_{r_i+1}) = m + r_i + 1$$

implies $l \leq r_i + 1$. We can assume $Q \vee K \subseteq Q_i \vee K_{r_i+1}$.

Since $K_m < Q \leq A$, Q is not a clique, and so there exist $U \subseteq Q$ and $q_0, q_1 \in Q$ such that $K_m \approx U$ and $q_0 \not\sim q_1$. We may take, without loss of generality, $q_0 \notin U$. We must have $\{q_0, q_1\} \subseteq Q_i - K_{r_i+1}$. So there exists $j \neq i$ such that $Q_j \approx Q_i - q_0$. Put $Q' = Q - q_0$. Then $K_m \leq Q'$ and $Q' \vee K_l \subseteq (Q_i - q_0) \vee K_{r_i+1} \approx Q_j \vee K_{r_i+1}$. Clearly, we have $K_{r_j+1} \approx K_{r_i+1} \vee K_1$. Set $K_l \prec K_l \vee K_1 = \overline{K}$. To finish the proposition we see that

$$Q' \vee \overline{K} \approx Q' \vee (K_l \vee K_1) \leq Q_j \vee (K_{r_i+1} \vee K_1) \approx Q_j \vee K_{r_j+1}.$$

•

Since we can do addition with cliques, we can use the previous proposition to define the ternary relation $\{(A, B, G) : |A| + |B| = |G|\}$ which allows us to do addition with the cardinality of arbitrary graphs. As a consequence, we have the definability of the n-step cover $\prec_n$ defined as $A \prec_n B$ if there exists a chain of covers $A \prec F_1 \prec \cdots \prec F_n \approx B$.

Lemma III.6.3. $\{(A, B, C) : C \approx C_n$ and $A \prec_n B\}$ is definable.

`Proof`: The claim is that $A \prec_n B$ iff $A \leq B$, and $C_n \approx C_{|B|-|A|}$. •

III.7 Individual Definability

In this section we will give a proof of the following proposition. Let $\mathbf{p}_3$ denote the isomorphism type of P_3.

Proposition III.7.1. Every element of $\langle \mathcal{PG}, \leq, \mathbf{p}_3 \rangle$ is definable. The complementation map is the only non-trivial automorphism of $\langle \mathcal{PG}, \leq \rangle$.

If we can show every graph is definable, then the following lemma will completely characterize the automorphisms of $\langle \mathcal{PG}, \leq \rangle$.

Lemma III.7.2. Suppose ϕ is a non-identity automorphism of the structure B of signature σ. Suppose $\mathbf{c}_o \in B$ such that $\{\mathbf{c}_0, \phi(\mathbf{c}_0)\}$ is definable in the signature σ. Suppose that each element of P is definable in the signature σ with the additional constant $\mathbf{c}_0$; that is, definable in the structure $\langle B, \sigma, \mathbf{c}_0 \rangle$. Then ϕ is the only non-identity automorphism of the structure B.

`Proof`: Suppose τ is an automorphism which fixes $\mathbf{c}_0$. By assumption, for every element $a \in B$, there is a formula $\Psi_a(x,y)$ such that $B \models \Psi_a(b, \mathbf{c}_0)$ iff $a = b$. Then $B \models \Psi_a(a, \mathbf{c}_0)$ implies $B \models \Psi_a(\tau(a), \tau(\mathbf{c}_0))$ which implies $B \models \Psi_a(\tau(a), \mathbf{c}_0)$; therefore, $\tau(a) = a$ for every $a \in B$, and so we conclude $\tau = id$.

If $\tau \neq id$, then because $\{\mathbf{c}_0, \phi(\mathbf{c}_0)\}$ is definable without constants, we must have $\tau(\mathbf{c}_0) = \phi(\mathbf{c}_o)$. Then $\phi^{-1} \circ \tau$ fixes $\mathbf{c}_0$, and so by the above, $\phi^{-1} \circ \tau = id$ which implies $\tau = \phi$. •

Definition III.7.3. Let A be any element of QGR with $|A| = n$. Let B be a graph with vertex set over the positive integers $\{1, ..., n\} = [n]$ such that $B \approx A$. Construct a finite graph denoted by $P_n(A,B)$ in the following way:

First, take the graph $B + \sum_{i=1}^{n} C_{n+2+i}$. Next, for each vertex k of B, add an edge connecting k to some vertex of C_{n+2+k}. In the end only n new edges are added. The resulting graph is called an *o-presentation* of A. The o-presentation $P_n(A,B)$ should look like the graph A with an edge leading out of each vertex to a circuit uniquely determined by cardinality. The figure below shows $P_4(A,B)$ where A is the ismorphism type of $K_3 \rightarrow_1$ and B is the isomorphic copy over the positive integers labeled as shown.

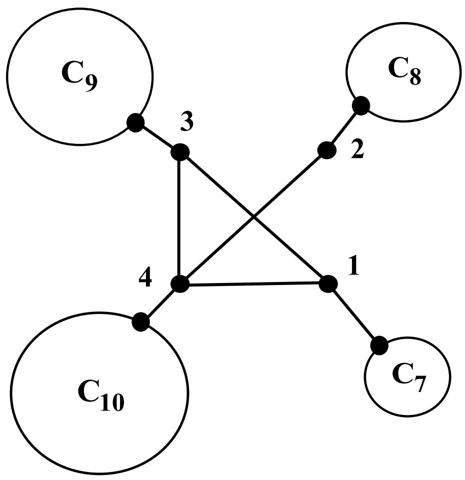

Figure III.3: An o-presentation for $P_4(K_3 \rightarrow_1, B)$

Proposition III.7.4. For a particular $A \in$ QGR, each o-presentation $P_n(A,B)$ is definable.

`Proof`: The idea is to use specific information of B as a graph on the vertices $[n]$ to write down first-order properties which capture $P_n(A,B)$. First we introduce a little simplifying notation; for $i,j \in [n]$, $i \neq j$, let

$$B(i,j) = \begin{cases} C_{n+2+i} \to_1 + C_{n+2+j} \to_1 & \text{if } i \not\sim j \text{ in } B \\ \gamma(n+2+i, n+2+j) & \text{if } i \sim j \text{ in } B \end{cases}$$

The claim is that $\Gamma \approx P_n(A,B)$ iff

(1) $\sum_{i=1}^n C_{n+2+i} \prec_n \Gamma$

(2) If $C_{n+2+i} \prec R \leq \Gamma$, then $R \approx C_{n+2+i} + N_1$ or $R \approx C_{n+2+i} \to_1$

(3) (for each $i,j \in [n]$, $i \neq j$) $B(i,j) \leq \Gamma$

(4) If $\sum_{i=1}^n C_{n+2+i} \prec R \leq \Gamma$, then there exists $j \in [n]$ such that $C_{n+2+j} \to_1 \leq R$.

(5) (for each $i,j \in [n]$, $i \neq j$) If $C_{n+2+i} + C_{n+2+j} \prec R \leq \Gamma$, then $\overline{C_{n+2+i} + C_{n+2+j}} \not\approx R$

(6) For $i \in [n]$, if $C_{n+2+i} \prec_2 R \leq \Gamma$, then $R \approx C_{n+2+i} \to_2$, or $R \approx C_{n+2+i} + K_2$, or $R \approx C_{n+2+i} \to_1 + N_1$, or $R \approx C_{n+2+i} + N_2$

It is easy to see that because the cardinality of C_{n+2+i} exceeds n and is connected uniquely to vertex i of B, $P_n(A,B)$ contains a unique copy of each C_{n+2+i} for $i \in [n]$; therefore, $P_n(A,B)$ also contains a unique copy of $\sum_{i=1}^n C_{n+2+i}$. By construction, each C_{n+2+i} is connected to a unique vertex. These facts make it straightforward to check that $P_n(A,B)$ satisfies the stated conditions.

For sufficiency, assume Γ satisfies conditions (1) - (6). From (1), we can assume, after passing to isomorphic induced subgraphs, that there exist vertices $\{v_1, \ldots v_n\}$ of Γ such that $\sum_{i=1}^n C_{n+2+i} = \Gamma - v_1 - \cdots - v_n$. Suppose there exist $k,j \in [n]$ such that v_k is adjacent to more than one vertex of C_{n+2+j}. Then the induced subgraph on the vertices $C_{n+2+j} \cup \{v_k\}$ is not isomorphic to $C_{n+2+j} + N_1$ nor to $C_{n+2+j} \to_1$; a contradiction of (2). Each vertex v_k is adjacent to at most one vertex of any C_{n+2+i} for $i = 1, \ldots, n$.

Suppose there exist $\{k,i,j\} \subseteq [n]$, $i \neq j$, such that the induced graph on the vertices of $C_{n+2+j} \cup \{v_k\}$ is isomorphic to $C_{n+2+j} \to_1$, and the induced subgraph on the vertices $C_{n+2+i} \cup \{v_k\}$ is isomorphic to $C_{n+2+i} \to_1$. Then the induced subgraph on $C_{n+2+i} \cup C_{n+2+j} \cup \{v_k\}$ is isomorphic to $\overline{C_{n+2+i} + C_{n+2+j}}$; a contradiction of (5). We see that if v_k is adjacent to some C_{n+2+j}, then it cannot be adjacent to any other circuit of $\sum_{i=1}^n C_{n+2+i}$.

If for each $k \in [n]$, we consider the induced subgraph on the vertices of $\sum_{i=1}^n C_{n+2+i} \cup \{v_k\}$, then condition (4) implies v_k is adjacent to some C_{n+2+j}. Altogether we have shown there is a function $\phi : [n] \to [n]$ such that for each $k \in [n]$, $C_{n+2+\phi(k)}$ is the unique circuit of $\sum_{i=1}^n C_{n+2+i}$ adjacent to v_k; moreover, the induced subgraph on the vertices yields a copy of $C_{n+2+\phi(k)} \to_1$.

We show ϕ is bijective. Suppose $\phi(i) = \phi(j) = k$. Then the induced subgraph on the vertices of $C_{n+2+k} \cup \{v_i, v_j\}$ is a graph which cannot be isomorphic to any of the four types of graphs listed in condition (6), a contradiction; therefore, ϕ is injective and so a bijection.

Condition (1) then implies there is a unique copy of each C_{n+2+i}. Since each $C_{n+2+\phi(k)}$ is uniquely connected to a single v_k, we see that $B(\phi(i),\phi(j)) \leq \Gamma$ if and only if $B(\phi(i),\phi(j))$ is isomorphic to the induced subgraph on the vertices of $C_{n+2+\phi(i)} \cup C_{n+2+\phi(j)} \cup \{v_i, v_k\}$. This implies $v_i \sim v_j$ if and only if $B(\phi(i),\phi(j)) \leq \Gamma$. Condition (3) then implies $v_i \sim v_j$ if and only if $\phi(i) \sim \phi(j)$ in B.

If F is the graph Γ induces on the vertices $\{v_1, \ldots, v_n\}$, then what we have shown is that the map $v_i \mapsto \phi(i)$

for $i = 1, \ldots, n$ yields an isomorphism $F \approx B$. Since $C_{n+2+\phi(i)}$ is uniquely connected to v_i by a single edge, the isomorphism ϕ extends to an isomorphism $\Gamma \approx P_n(A,B)$ in the natural way. •

The next task is to find a way to "read off" the copy of A sitting inside an o-presentation $P_n(A,B)$. The first step is to return to the topic of paths and isolate particular covers. When attaching a new vertex v to a path P, the choice of $u \in P$ for $u \sim v$ makes a difference. We will use the notation $P_n \to_1^t$ to denote the covers of P_n which are formed by adding a single new edge $u \sim v$ where v is a new vertex and $u \in P_n$ such that u has degree two. Different choices of u lead to non-isomorphic graphs, so the notation $P_n \to_1^t$ refers to the finite family of such graphs for a fixed n.

Lemma III.7.5. $\{(P,\Gamma) : P \approx P_n \text{ and } \Gamma \in P_n \to_1^t\}$ is definable.

`Proof`: The claim is that $\Gamma \in P_n \to_1^t$ iff $P_n \prec \Gamma$, $K_{1,3} \leq \Gamma$, $P_{n+1} \not\leq \Gamma$, and Γ is acyclic.

If Γ satisfies the conditions, then Γ is formed from P_n by adding a new vertex v and possibly new edges of the form $u \sim v$ for $u \in P_n$. If at least two new edges are added, then a circuit must be formed, and so there is at most one new edge $x \sim v$ with $x \in P_n$. Since $K_{1,3} \leq \Gamma$, there is exactly one new edge. Since $P_{n+1} \not\leq \Gamma$, the degree of x cannot be one. This establishes the lemma. •

Lemma III.7.6. $\{(C,\Gamma) : P \approx P_n \text{ and } \Gamma \approx \sum_{i=1}^n P_{n+1+i}\}$ is definable.

`Proof`: The claim is that $\Gamma \approx \sum_{i=1}^n P_{n+1+i}$ iff $\Gamma \prec_n \sum_{i=1}^n C_{n+2+i}$ and Γ is acyclic. •

We now have all the ingredients to finish the proof of Proposition III.7.1. From an o-presentation $P_n(A,B)$ we see that $A + \sum_{i=1}^n P_{n+1+i}$ is the unique graph G such that

(1) $G \prec_n P_n(A,B)$

(2) For all $k \in [n]$, $\quad C_{n+2+k} \not\leq G$

(3) For all $k \in [n]$, $\quad G$ embeds no element of $P_{n+1+k} \to_1^t$

(4) For all $k \in [n]$, $\quad$ each P_{n+1+k} is a connected component of G

This follows since conditions (1) and (2) imply G is obtained precisely by deleting exactly one vertex from each C_{n+2+i} for $i \in [n]$. Condition (3) and (4) imply each of those vertices must have degree three. We can then recover A as the unique graph H such that $G \approx H + \sum_{i=1}^n P_{n+1+i}$.

Remark III.7.7. Using a particular $P_3 \in$ QGR as a constant we have shown every finite graph is definable up to isomorphism in $\langle \text{QGR}, \leq, P_3 \rangle$. The same result could be achieved, but perhaps with greater difficulty, if we had chosen another graph C as the constant, provided C is not self-complementary. To see this, notice from the proof of Proposition III.0.6, there is a formula $\beta(x)$ in the language of $\leq$ such that $\langle \text{QGR}, \leq \rangle \vDash \beta(E)$ iff $E \approx P_3$ or $E \approx K_2 + N_1$. By what we have shown, for any $G \in$ QGR there is a formula $\phi_G(x,y)$ in the language of $\leq$ such that $\langle \text{QGR}, \leq \rangle \vDash \phi_G(E, P_3)$ iff $E \approx G$. It is not hard to see that the unary formula

$$(\exists y)\phi_G(x,y) \wedge \phi_C(C,y) \wedge \beta(y)$$

uniquely defines G.

In the next proposition, we shall see how to capture the pair (A,P) where P is isomorphic to some o-presentation of A. The proof of Proposition III.7.4 relied on the fact that we had a fixed graph on hand, and so we could "encode" the edge relation of this fixed graph with a certain packet of formulas. This means different graphs require different package of formulas to define the edge relations in the o-presentations. Since we have definable access to the cardinality of a graph, and can do arithmetic with circuits, we shall be able to describe a uniform packet of formulas which "encode" the edge relations of some graph in an o-presentation.

Proposition III.7.8. We have the following:

(1) $\{(A,P_n(A,B)):$ for some $B \approx A$ with $n = |A|\}$ is definable.

(2) If B is a graph over the vertices $[n]$ with $B \approx A$ and B' is a graph over the vertices $[m]$ with $B' \approx A'$, then $P_n(A,B) \approx P_m(A',B')$ if and only if $n = m$ and $B = B'$.

`Proof`: For part (1), the claim is that (A,P) is in the relation iff (where $|A| = n$ which is definable)

(1) $\sum_{i=1}^{n} C_{n+2+i} \prec_n P$

(2) If $C_{n+2+i} \prec R \leq P$, then $R \approx C_{n+2+i} + N_1$ or $R \approx C_{n+2+i} \rightarrow_1$

(3) (for each $i,j \in [n]$, $i \neq j$) If $C_{n+2+i} \rightarrow_1 + C_{n+2+j} \rightarrow_1 \not\leq P$, then $\gamma(n+2+i, n+2+j) \leq P$

(4) If $\sum_{i=1}^{n} C_{n+2+i} \prec R \leq P$, then there exists $j \in [n]$ such that $C_{n+2+j} \rightarrow_1 \leq R$.

(5)(for each $i,j \in [n]$, $i \neq j$) If $C_{n+2+i} + C_{n+2+j} \prec R \leq P$, then $\overline{C_{n+2+i} + C_{n+2+j}} \not\approx R$

(6) For $i \in [n]$, if $C_{n+2+i} \prec_2 R \leq P$, then $R \approx C_{n+2+i} \rightarrow_2$, or $R \approx C_{n+2+i} + K_2$, or $R \approx C_{n+2+i} \rightarrow_1 + N_1$, or $R \approx C_{n+2+i} + N_2$

(7) If G is the graph which satisfies the following properties

(a) $G \prec_n P$

(b) For all $k \in [n]$, $\quad C_{n+2+k} \not\leq G$

(c) For all $k \in [n]$, $\quad P_{n+1+k} \rightarrow_1^t \not\leq G$

(d) For all $k \in [n]$, $\quad$ each P_{n+1+k} is a connected component of G

then $G \approx A + \sum_{i=1}^{n} P_{n+1+i}$.

The proof of necessity and sufficiency exactly follows the proof of Proposition III.7.4. Any P which satisfies conditions (1) - (6) must be isomorphic to an o-presentation $P_n(E,F)$ for some F. Condition (7) then implies $E \approx A$. The details are left for the reader.

We establish part (2). Clearly, $n = m$ and $B = B'$ implies $P_n(A,B) \approx P_m(A',B')$. Suppose $P_n(A,B) \approx P_m(A',B')$. By using the definition of an o-presentation, $P_n(A,B)$ has $n + \sum_{i=1}^{n}(n+2+i)$ vertices. Since $P_n(A,B)$ and $P_m(A',B')$ have the same cardinality, we must have

$$n^2 + 3n + \frac{n(n+1)}{2} = m^2 + 3m + \frac{m(m+1)}{2}$$

which implies $n = m$, and so B and B' have the same vertices. Because $P_n(A,B)$ and $P_m(A',B')$ must then have a unique copy of each C_{n+2+i} and therefore, a unique copy of each $C_{n+2+i} \rightarrow_1$ for $i \in [n]$, the isomorphism of o-presentations restricts to the identity on B and B'. ●

III.8 Morphisms

We now turn to the task of encoding set functions. In the next section, we shall be interested in graph homomorphisms. We start with some auxiliary constructions.

Definition III.8.1. A *panda* is the graph $P(n)$ constructed from $C_n \to_1$ by adding two additional vertices x and y and only two new edges $x \sim u$ and $y \sim u$ where u is the unique vertex of $C_n \to_1$ with degree one. A *p-panda*, denoted by $pP(n)$ is formed from $P(n)$ be completing the triangle formed by the panda's arms; that is, by adding the edge $x \sim y$ to $P(n)$.

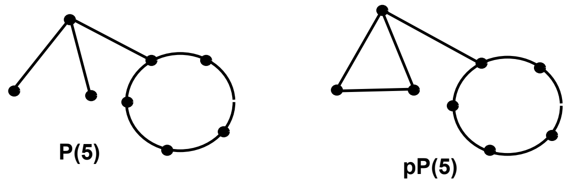

Figure III.4: A panda and p-panda

Lemma III.8.2. $\{(C,F) : C \approx C_n \text{ and } F \approx P(n)\}$ is definable.

`Proof`: The claim is that $F \approx P(n)$ iff

$n = 3$ and $F \approx M(3)$, or

$n > 3$ and $C_n + N_2 \prec F$, F is connected, and if $C_n \prec R \leq F$, then $R \approx C_n \to_1$ or $R \approx C_n + N_1$.

Suppose F satisfies the criteria and $F \not\approx M(3)$. We may construct F from $C_n + N_2$ by adding a new vertex v and some additional edges connecting v to $C_n + N_2$. Let N_2 be composed of the vertices a and b. Since F is connected, we must have edges $v \sim a$ and $v \sim b$, and at least one edge $v \sim x$ where $x \in C_n$. Since the induced subgraph on the vertices of $C_n \cup \{v\}$ is connected, we must have exactly one edge connecting v to C_n.

Since necessity is immediate, the proposition is established. •

Lemma III.8.3. $\{(C,F) : C \approx C_m \text{ and } F \approx pP(m)\}$ is definable.

`Proof`: The criteria is that $F \approx pP(m)$ iff $C_m + K_2 \prec F$, and if

$m = 3$, then $F \approx pM(3)$, or if

$m > 3$, F is connected, $K_3 \leq F$, and if $C_m \prec R \leq F$, then $R \approx C_m \to_1$ or $R \approx C_m + N_1$.

It is straightfoward to check the necessity of the criteria, and so we will establish their sufficiency.

Suppose F satisfies the conditions. Let F be constructed from $C_m + K_2$ by adding a new vertex v and possibly new edges incident with v. Let a and b be the two vertices of K_2. We may assume $m > 3$.

Let $m > 3$. Since F is connected, we have at least one edge $v \sim x$ for $x \in C_m$. Then we must have exactly one edge since the induced subgraph on $C_m \cup \{v\}$ is connected. Since $K_3 \leq F$, we must have edges $v \sim a$ and $v \sim b$. •

Let $Tuttle$ denote the cover of P_4 formed by adding a new vertex to P_4 and exactly two new edges connecting the new vertex to the two vertices of P_4 which have degree 2. In light of Proposition III.7.4, $Tuttle$ is definable.

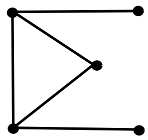

Figure III.5: Tuttle

Here is how to encode a function $f:[n]\to[m]$. Define the graph $\sigma(n,f,m)$ over the vertex set

$$\sum_{i=1}^{n} C_{3+i} \;+\; \sum_{i=1}^{m} K_{3+i} \;+\; N_n$$

with the following edge relations: Let $\{v_1,\ldots,v_n\}$ be the vertices of N_n; choose vertices $x_i \in C_{3+i}$ and $u_j \in K_{3+j}$ for $i=1,\ldots,n$ and $j=1,\ldots,m$; take all the edges of $\sum_{i=1}^{n} C_{3+i} \;+\; \sum_{i=1}^{m} K_{3+i}$ together with edges $x_i \sim v_i$ for $i=1,\ldots,n$, and edges $v_i \sim u_{f(i)}$ for $i=1,\ldots,n$; these are the only edges.

Notice the choices of u_i and x_i are immaterial. Below is $\sigma(n,f,m)$ for $f(1)=f(2)=1,\ f(3)=2$.

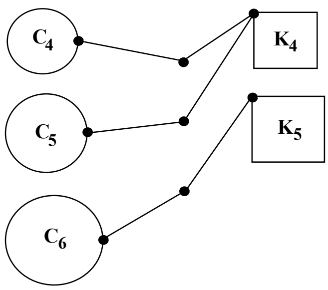

Figure III.6: $\sigma(3,f,2)$

Remark III.8.4. It may be possible to replace the conditions listed in the next proposition by a more efficient or elegent set of conditions, but the advantage of the list is that it makes the proof of sufficiency straightforward to verify.

Proposition III.8.5. We have the following:

(1) $\sigma(n,f,m)\approx\sigma(n',f',m')$ iff $n=n'$, $m=m'$ and $f=f'$.

(2) $\{(C_n,K_m,F): n,m>0,$ and $F\approx\sigma(n,f,m)$ for some $f:[n]\to[m]\}$ is definable.

`Proof`: We tackle statement (1). Observe that $\sigma(n,f,m)$ contains a unique copy of $\sum_{i=1}^{n} C_{3+i}$ and of $\sum_{i=1}^{m} K_{3+i}$. This implies we must have $n=n'$ and $m=m'$. For each $i\in[n]$, C_{3+i} appears in exactly one pointed sum $K_{3+f(i)} +_p C_{3+i}$ in $\sigma(n,f,m)$, and thus, also in $\sigma(n',f',m')$. This implies $f(i)=f'(i)$ for each $i\in[n]$, and so, $f=f'$.

For the second statement, the claim is that $F\approx\sigma(n,f,m)$ for some $f:[n]\to[m]$ iff

(1) $\sum_{i=1}^{n} C_{3+i} \;+\; \sum_{i=1}^{m} K_{3+i} \prec_n F$

(2) For $i \in [n]$, C_{3+i} is not a connected component of F
(3) For $i \in [n]$, if $C_{3+i} \prec R \leq F$, then $R \approx C_{3+i} \rightarrow_1$ or $R \approx C_{3+i} + N_1$.
(4) For $i, j \in [n]$, $i \neq j$, if $C_{3+i} + C_{3+j} \prec R \leq F$, then $R \not\approx \overline{C_{3+i} + C_{3+j}}$
(5) For each $i \in [n]$, $C_{3+i} \rightarrow_2 \leq F$.
(6) For each $i, j \in [n]$, $i \neq j$, $\gamma(3+i, 3+j) \nleq F$
(7) For $j \in [m]$, if $K_{3+j} \prec R \leq F$, then $R \approx K_{3+j} \rightarrow_1$ or $R \approx K_{3+i} + N_1$.
(8) For $i \in [n]$, $P(i) \nleq F$ and $pP(i) \nleq F$
(9) $Tuttle \nleq F$

Suppose F satisfies the conditions. Condition (1) implies we can construct F by adding the vertices $\{v_1, ..., v_n\}$ to $\sum_{i=1}^{n} C_{3+i} + \sum_{i=1}^{m} K_{3+i}$ and possibly some new edges incident with the vertices $\{v_1, ..., v_n\}$. Condition (2) implies each C_{3+i} is adjacent to at least one vertex of $\{v_1, \ldots, v_n\}$. Condition (3) implies that if C_{3+i} is adjacent to some v_k, then there is a unique vertex $x_i \in C_{3+i}$ such that $x_i \sim v_k$. Condition (4) implies no two distinct C_{3+i} and C_{3+j} are adjacent to the same vertex of $\{v_1, \ldots, v_n\}$, and therefore, no C_{3+i} can be adjacent to more than one vertex of $\{v_1, \ldots, v_n\}$. We may reorder the vertices of $\{v_1, \ldots, v_n\}$ so that for each $i \in [n]$, v_i is the unique vertex adjacent to C_{3+i} by a unique edge. Condition (6) implies $v_i \nsim v_j$ for $i \neq j$.

Condition (5) implies each v_i is adjacent to a vertex of some K_{3+j}, and condition (8) implies it is not connected to any other vertex of K_{3+j}, nor to any vertex of K_{3+r} for $r \neq j$. For each $i \in [n]$, let $f(i) = j$ where K_{3+j} is adjacent to v_i. Let $u_{f(i)} \in K_{3+f(i)}$ such that $v_i \sim u_{f(i)}$. Condition (8) asserts that whenever $f(i) = f(j)$ for $i \neq j$, then $u_{f(i)} = u_{f(j)}$; that is, a unique vertex is chosen in $K_{3+f(i)}$ so that whenever v_k is connected to the clique $K_{3+f(i)}$, it is connected by that vertex. The map $f : [n] \rightarrow [m]$ is the function we are after, and altogether we have shown $F \approx \sigma(n, f, m)$. •

III.9 A Small Category

We define a small category $\mathcal{CG}$. The objects are simple graphs whose vertex sets are initial segments of positive integers. The morphisms $\mathcal{CG}(A,B)$ are the graph homomorphisms from A to B which we write as a triple $F = (A, f, B)$ where $f : [n] \rightarrow [m]$ with $n = |A|$ and $m = |B|$. The category $\mathcal{CG}$ can naturally be thought of as a 2-sorted first-order structure, with one sort for objects and another sort for morphisms, together with a ternary relation over the sort of morphisms which reflects composition. The category structure is then described by the standard category axioms in this 2-sorted first-order language.

For a morphism, the property of being a monomorphism or an epimorphism is by definition first-order definable in the language of the category. In general categories we do not formally have access to the "inner" structure of the objects and so we don't expect to definably capture the property of injectivity or surjectivity; likewise, the property that $f \in \mathcal{CG}(A,B)$ is an embedding refers to the relational structure of A and B which is not included in the 2-sorted language of the category. In the case of simple graphs, for a morphism to be injective is equivalent to being a monomorphism and surjectivity is equivalent to being an epimorphism. So here, injectivity and surjectivity are definable properties. It will also be possible to capture embeddings, but we need to make an adjustment.

We enrich the small category by adding 4 new constants to the language and denote the resulting stucture by $\mathcal{CG}'$. We add the constants $\mathbf{K}_2$, $\mathbf{P}_3$, and the maps $\mathcal{CG}(N_1,\mathbf{K}_2) = \{\mathbf{t},\mathbf{b}\}$. Notice that N_1 is a terminal object in $\mathcal{CG}'$, and so is definable. We can use the first-order structure of the category and the constants to definably manipulate the edge relations of the object in the category. For any graph A and vertices $u,v \in A$, we see that $u \sim v$ in A iff where $x,y \in \mathcal{CG}(N_1,A)$ such that $x(0) = u$ and $y(0) = v$, then there exists $h \in \mathcal{CG}(\mathbf{K}_2,A)$ such that $h\mathbf{t} = x$ and $h\mathbf{b} = y$.

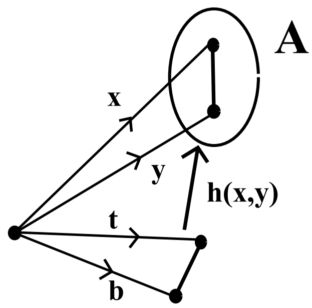

Figure III.7: Reading the edge relation

The first-order language of $\mathcal{CG}'$ is even more expressive. We see $f \in \mathcal{CG}(A,B)$ is an embedding iff f is a monomorphism and whenever there exist $x,y \in \mathcal{CG}(N_1,A)$ and $q \in \mathcal{CG}(\mathbf{K}_2,B)$ such that $g\mathbf{t} = fx$ and $g\mathbf{b} = fy$, then there exists $\phi \in \mathcal{CG}(\mathbf{K}_2,A)$ such that $x = \phi\mathbf{t}$ and $y = \phi\mathbf{b}$. This means that the substructure relation of $\langle \text{QG}\textsc{r}, \leq, P_3\rangle$ when restricted to the objects of the small category is first order-definable in $\mathcal{CG}'$.

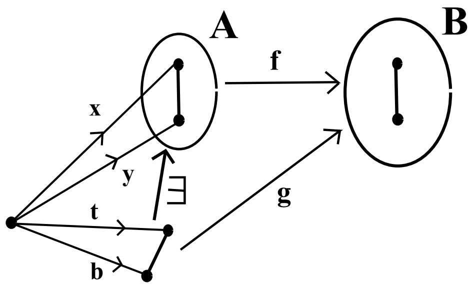

Figure III.8: Capturing embeddings

Using the bijection $G \ni u \leftrightarrow x \in \mathcal{CG}(N_1,G)$ such that $x(0) = u$, we see that given any $G \in \text{Obj } \mathcal{CG}$, we can construct an isomorphic graph

$$\hat{G} = \langle \mathcal{CG}(N_1,G), \hat{r} \subseteq \mathcal{CG}(N_1,G) \times \mathcal{CG}(N_1,G)\rangle \approx G \tag{III.1}$$

where both the set of vertices and the edge relation $\hat{r}$ have first-order definitions in the language of the small category. It is not difficult to see that the set of such graphs $\left\{\hat{G} : G \in \text{Obj}\,\mathcal{CG}\right\}$ is definable.

By following the procedure outlined in [JM10, Sec.3.1], we can use the first-order language of the category applied to the structures in $\left\{\hat{G} : G \in \text{Obj}\,\mathcal{CG}\right\}$ to parametrize arbitrary subsets of finitary cartesian products. To see this, take $\hat{G}_1,\dots,\hat{G}_m$ and $\hat{R}$ a subset of the cartesian product of their universes; that is, $\hat{R} \subseteq \mathcal{CG}(N_1,G_1) \times \cdots \times \mathcal{CG}(N_1,G_m)$. By the bijection in the previous paragraph there is a corresponding relation

$R \subseteq G_1 \times \cdots \times G_m$. If $|R| = k$, we shall use the maps in $\mathcal{CG}(N_1, N_k)$ to parametrize k-element subsets of the product in the same way as the maps of $\mathcal{CG}(N_1, A)$ parametrize the elements of A. If $\pi_i : G_1 \times \cdots \times G_m \to G_i$ denotes the i-th projection of the cartesian product as sets, then for any fixed bijection $p : [k] \to R$ there is a fixed sequence of morphisms $p_i \in \mathcal{CG}(N_k, G_i)$ given by $p_i = \pi_i \circ p$. This follows since any set map $\alpha : [k] \to G$ corresponds exactly to a morphism of the trivial graph N_k into G. An arbitrary tuple in R is then specified by $(p_1(s), ..., p_m(s))$ where $s \in [k]$. It is easy to see that with this choice of $(p_1, ..., p_m)$ we have

$$\hat{R} = \{(q_1, ..., q_m) \in \mathcal{CG}(N_1, G_1) \times \cdots \times \mathcal{CG}(N_1, G_m) : q_i = p_i \circ q \text{ for some } q \in \mathcal{CG}(N_1, N_k)\}$$

In this way, the first-order language of $\mathcal{CG}'$ when restricted to the structures $\left\{\hat{G} : G \in \mathrm{Obj}\,\mathcal{CG}\right\}$ is equivalent to a second-order language which has variables ranging over the elements of $\hat{G}$, variables for the morphisms between objects, and can express the edge relation in objects, application of morphisms to elements, composition of morphisms, and equality of elements and morphisms, and the apparatus to quantify over arbitrary subsets of finite products. Altogether, using the isomorphism in Eq.(III.1), we see that the first-order language of $\mathcal{CG}'$ when restricted to the objects of the category is equivalent in expressive power to a full second-order language of simple graphs over the same set of objects. What is surprising is that the isomorphism invariant relations definable in the first-order language of $\mathcal{CG}'$ (equivalently, a full second-order language), is up to isomorphism first-order definable in the theory of substructure. In order to establish this fact, we need to build a model of the small category in the definable relations of $\langle \mathrm{QG}\textsc{r}, \leq, P_3 \rangle$. The difficult part of this has already been accomplished.

Suppose we have graphs $G_i = \langle [m_i], r_i \rangle$ in the category and a morphism $F = (G_1, f, G_2)$ such that $f : [m_1] \to [m_2]$. We encode G_i as any graph isomorphic to $P_i = P_{m_i}(G_i, G_i)$ and encode F as any triple isomorphic to $M(F) = (P_1, \sigma(m_1, f, m_2), P_2)$.

In the next result, we see how to read off the values of a function f with statement (1), and how to capture that a f is a homomorphism with statement (2).

Lemma III.9.1. We have the following:

1. If $(U, V, W) \approx M(F)$ for $F = (G_1, f, G_2)$, then F(and f) are uniquely determined and for all $i \in [m_1]$ and $j \in [m_2]$, we have that $f(i) = j$ iff $K_{3+j} +_p C_{3+i} \leq V$.

2. $(U, V, W) \approx M(F)$ for some $F = (G_1, f, G_2)$ iff where $m_i = |G_i|$, we have $U \approx P_{m_1}(G_1, G_1)$, $W \approx P_{m_2}(G_2, G_2)$, and $V \approx \sigma(m_1, f, m_2)$ for some $f : [m_1] \to [m_2]$; and whenever we have $1 \leq i, i' \leq m_1$ and $1 \leq j, j' \leq m_2$, $j \neq j'$, and $K_{3+j} +_p C_{3+i} \leq V$ and $K_{3+j'} +_p C_{3+i'} \leq V$, then $\gamma(m_1 + 2 + i, m_1 + 2 + i') \leq U$ implies $\gamma(m_2 + 2 + j, m_2 + 2 + j') \leq W$.

`Proof`: For part (1), the first part of Proposition III.8.5 and the second part of Proposition III.7.8 guarantee that $(U, V, W) \approx M(F)$ iff $U \approx P_{m_1}(G_1, G_1)$, $W \approx P_{m_2}(G_2, G_2)$, and $V \approx \sigma(m_1, f, m_2)$. That $f(i) = j$ iff $K_{3+j} +_p C_{3+i} \leq V$ is explicit by construction.

For part (2), recall in the proof of Proposition III.7.4 that $\gamma(m_1 + 2 + i, m_1 + 2 + i') \leq P_{m_1}(B_1, B_1)$ iff $i \sim i'$ in B_1. •

We now account for the composition of morphisms.

Lemma III.9.2. Let $F = (G_1, f, G_2)$ and $H = (G_2, h, G_3)$ with $|G_i| = m_i$ for $i = 1,2,3$. Let $M(F) \approx (P_1, \sigma_1, P_2)$ and $M(H) \approx (P_2, \sigma_2, P_3)$. Then $M(HF) \approx (P_1, \sigma_3, P_3)$ iff $\sigma_3 \approx \sigma(m_1, p, m_3)$ for some p such that: for all $i \in [m_1]$, $j \in [m_2]$, $k \in [m_3]$, we have that $K_{3+j} +_p C_{3+i} \leq \sigma_1$ and $K_{3+k} +_p C_{3+j} \leq \sigma_2$ imply $K_{3+k} +_p C_{3+i} \leq \sigma_3$.

We have gathered together all the necessary definable relations. The analogue of Theorem 3.8([JM10]) below in the case of simple graphs goes through exactly word for word. Since we have strived for relative completeness in this dissertation, and since the corresponding result for equivalence relations in Section IV.5 requires a slight modification of the translation scheme, we will present the argument here.

Theorem III.9.3. Let R be an isomorphism invariant relation over QGR. Then R is first-order definable over $\langle \mathrm{QGR}, \leq, P_3 \rangle$ iff its restriction to Obj $\mathcal{CG}$ is first-order definable in the language of $\mathcal{CG}'$.

`Proof`: Let R be an isomorphism invariant N-ary relation over QGR. One direction has already been established in the beginning of this section.

Suppose $R \cap \mathrm{Obj}\,\mathcal{CG}^N$ is definable in the language of $\mathcal{CG}'$; that is, there exists a formula Φ in the language of $\mathcal{CG}'$ such that

$$R \cap \mathrm{Obj}\,\mathcal{CG}^N = \left\{ (A_1, \ldots, A_N) \in \mathrm{Obj}\,\mathcal{CG}^N : \mathcal{CG}' \vDash \Phi(A_1, \ldots, A_N) \right\}.$$

We need a formula $\Psi(x_1, \ldots, x_N)$ in the language of $\langle \mathrm{QGR}, \leq, P_3 \rangle$ such that

$$R = \left\{ (G_1, \ldots, G_N) \in \mathrm{QGR}^N : \langle \mathrm{QGR}, \leq, P_3 \rangle \vDash \Psi(G_1, \ldots, G_N) \right\}.$$

We will define by induction a formula $\hat{\Phi}(x_1, \ldots, x_N)$ so that whenever $A_i \approx B_i$ with $A_i \in \mathrm{QGR}$, and $|A_i| = k_i$ for $i = 1, \ldots, N$ we have

$$\mathcal{CG}' \vDash \Phi(B_1, \ldots, B_N) \quad \text{iff} \quad \langle \mathrm{QGR}, \leq, P_3 \rangle \vDash \hat{\Phi}(P_{k_1}(A_1, B_1), \ldots, P_{k_N}(A_N, B_N)).$$

We can then take $\Psi(x_1, \ldots, x_N)$ to be

$$(\exists u_1, \ldots, u_N)(\hat{\Phi}(u_1, \ldots, u_N) \wedge (\text{“there exist } v_i \text{ such that } k_i = |x_i| \text{ and } u_i \approx P_{k_i}(x_i, v_i) \text{ for } i = 1, \ldots, N))\text{”}$$

Let $X_1, \ldots, X_M$ be a list of all the object variables, both free and bound, which appear in Φ. Let $f_1, \ldots, f_T$ be a list of all the morphism variables which appear in Φ. Note that all the morphism variables must appear bound. We introduce variables $x_1, \ldots, x_M$ for $\hat{\Phi}$ which will correspond to the object variable $X_1, \ldots, X_M$, and $y_1, \ldots, y_T$ which will correspond to the morphism variables $f_1, \ldots, f_T$. By induction on the length of a formula, we define a correspondence from the subformulas of Φ to formulas in the substructure relation.

Our scheme for translating the atomic subformulas is the following:

1. If ϕ is $X_r = X_s$, then $\hat{\phi}$ is $x_r \leq x_s \wedge x_s \leq x_r$.

2. If ϕ is $f_s = f_r$, then $\hat{\phi}$ is $y_r \leq y_s \wedge y_s \leq y_r$

3. If ϕ is $f_s \in \mathcal{CG}(X_r, X_t)$, then $\hat{\phi}$ is

$$(\exists u_r, u_t)(\text{“there exist } v_r, v_t \text{ such that } k_r = |u_r|, k_t = |u_t|, \text{ and } x_r = P_{k_r}(u_r, v_r)$$
$$\text{and } x_t = P_{k_t}(u_t, v_t) \text{ and } (x_r, y_s, x_t) = M(F) \text{ for some } F \in \mathcal{CG}(v_r, v_t)\text{”})$$

4. If ϕ is $f_i \in \mathcal{CG}(X_r, X_t) \wedge f_j \in \mathcal{CG}(X_t, X_s) \wedge f_k = f_j \circ f_i$, then $\hat{\phi}$ is

$$(\exists u_r, u_t, u_s)(\text{“there exist } v_m \text{ such that } k_m = |v_m| \text{ and } x_m = P_{k_m}(u_m, v_m) \text{ for } m \in \{r, t, s\} \text{ and}$$
$$(x_r, y_i, x_t) = M(F) \text{ for some } F \in \mathcal{CG}(v_r, v_t) \text{ and } (x_t, y_j, x_s) = M(G)$$
$$\text{for some } G \in \mathcal{CG}(v_t, v_s) \text{ and } (x_r, y_k, x_s) = M(GF)\text{”})$$

5. If ϕ is $\neg\psi$, or $\psi \wedge \theta$, then $\hat{\phi}$ is $\neg\hat{\psi}$, or $\hat{\psi} \wedge \hat{\theta}$
6. If ϕ is $(\exists X_r)\psi$, then $\hat{\phi}$ is $(\exists x_r)[(\exists u_r)(\text{“there exists } v_r \text{ such that } k_r = |v_r| \text{ and } x_r = P_{k_r}(u_r, v_r)\text{”}) \wedge \hat{\psi}]$
7. If ϕ is $(\forall X_r)\psi$, then $\hat{\phi}$ is $(\forall x_r)[(\exists u_r)(\text{“there exists } v_r \text{ such that } k_r = |v_r| \text{ and } x_r = P_{k_r}(u_r, v_r)\text{”}) \rightarrow \hat{\psi}]$
8. If ϕ is $(\exists f_s \in \mathcal{CG}(X_r, X_t))\psi$, then $\hat{\phi}$ is

$$(\exists y_s)[(\exists u_r, u_t)(\text{“there exist } v_r, v_t \text{ such that } k_r = |v_r|, k_t = |v_t| \text{ and } x_r = P_{k_r}(u_r, v_r) \text{ and}$$
$$x_t = P_{k_t}(u_t, v_t) \text{ and } (x_r, y_s, x_t) = M(F) \text{ for some } F \in \mathcal{CG}(u_r, u_t)\text{”}) \wedge \hat{\psi}]$$

9. If ϕ is $(\forall f_s \in \mathcal{CG}(X_r, X_t))\psi$, then $\hat{\phi}$ is

$$(\forall y_s)[(\exists u_r, u_t)(\text{“there exist } v_r, v_t \text{ such that } k_r = |v_r|, k_t = |v_t| \text{ and } x_r = P_{k_r}(u_r, v_r) \text{ and}$$
$$x_t = P_{k_t}(u_t, v_t) \text{ and } (x_r, y_s, x_t) = M(F) \text{ for some } F \in \mathcal{CG}(v_r, v_t)\text{”}) \rightarrow \hat{\psi}]$$

It is now straightforward to prove by induction on the length of a formula that for all subformulas $\phi(X_1, ..., X_M; f_1, ..., f_T)$ of Φ, and for all $G_i \in \mathrm{Obj}\,\mathcal{CG}$, $F_j = (B_j, g_j, C_j) \in \mathcal{CG}(B_j, C_j)$ with $|G_i| = u_i$, $|B_j| = b_j$ and $|C_j| = c_j$ for $i \leq M$, $j \leq T$ it holds that

$$\mathcal{CG}' \vDash \phi(G_1, ..., G_M; f_1, ..., f_T)$$
$$\text{iff}$$
$$\langle \mathrm{QG{\scriptsize R}}, \leq, P_3 \rangle \vDash \hat{\phi}(P_{u_1}(G_1, G_1), ..., P_{u_M}(G_M, G_M); \sigma_1(b_1, g_1, c_1), ..., \sigma_T(b_M, g_M, c_M))$$

The theorem is then established when $\phi = \Phi$. •

Corollary III.9.4. For every sentence ϕ in the second-order language of simple graphs, there is a formula $\Phi(x)$ in the first-order language of the quasi-ordered set $\langle \mathrm{QG{\scriptsize R}}, \leq, P_3 \rangle$ such that a graph A in QGR models ϕ if and only if $\langle \mathrm{QG{\scriptsize R}}, \leq, P_3 \rangle \vDash \Phi(A)$.

CHAPTER IV

EQUIVALENCE RELATIONS

In this chapter, we consider the universal class of equivalence relations $\mathcal{E}$, and investigate definability in the partially ordered set $\langle \mathcal{PE}, \leq, \mathbf{n}_2 \rangle$ where $\mathbf{n}_2$ denotes the isomorphism type of the two element identity relation.

IV.1 Individual Definability

Recall, finite equivalence relations can be considered as arithmetic partitions. If π is a partition of an n element set into t blocks, then we can represent π as a sequence of positive integers $\pi = (n_1, ..., n_t)$ where each n_i represents the size of the i-th block and $n_{i+1} \leq n_i$. If σ embeds into π, then any two elements which are in separate blocks in σ must be mapped to separate blocks in π. If $\sigma = (s_1, ..., s_r)$, then it is easy to see that

$$\sigma \leq \pi \qquad \text{iff} \qquad r \leq t \;\text{ and }\; s_i \leq n_i \;\text{ for all }\; i \leq r. \tag{IV.1}$$

There is another way to represent the partition $\pi = (n_1, ..., n_t)$ called a Young diagram. This is a series of left-justified rows of boxes; the first row has n_1 number of boxes, the second row has n_2 boxes, the third row has n_3 boxes, etc. For example, the following is the Young diagram for the partitions $(4,2,2,1)$ and $(5,3,2,1)$, respectively:

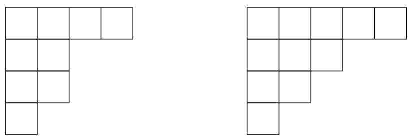

For any Young diagram, the transpose is defined by interchanging the rows and the columns in the same way that the transpose of a matrix is defined. For example, the Young diagram for $\pi = (n_1, ..., n_t)$ has first row with n_1 boxes, and second row with n_2 boxes, and third row with n_3 boxes, etc. Then the transpose Young diagram has the first column with n_1 boxes, the second column has n_2 boxes, the third column has n_3 boxes, etc. It is easy to see that the transpose of a Young diagram for π, is the Young diagram for another partition denoted π^∂. For example, $\pi = (4,2,2,1)$ and $\pi^\partial = (4,3,1,1)$ are pictured below:

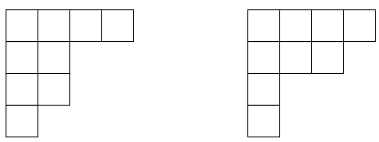

Using (IV.1), we see that $\pi \leq \sigma$ iff the Young diagram for π is contained in the Young diagram for σ. From this geometric picture it is easy to conclude:

Lemma IV.1.1. The transpose map $\pi \to \pi^\partial$ is an automorphism of $\langle \mathcal{PE}, \leq \rangle$.

If $\sigma < \pi$, but there does not exist ρ such that $\sigma < \rho < \pi$, then we write $\sigma \prec \pi$ and say π covers σ, or that σ is a subcover of π. For a partition π, $|\pi|$ will denote the cardinality of the underlying set. It is immediate from (IV.1) that $\pi \prec \rho$ implies $|\pi| + 1 = |\rho|$.

The identity, or *trivial*, relation on the set $\{1,\ldots,m\}$ will be denoted as N_m, and the set of identity relations as $\mathcal{N} = \{N_m : m < \omega\}$. The unique equivalence relation on $\{1,\ldots,m\}$ with only one block will be denoted as K_m. Such equivalence relations are said to be *complete*, or *total*, and the set of complete relations is denoted by $\mathcal{K} = \{K_m : m < \omega\}$.

Proposition IV.1.2. $\mathcal{K}$ and $\mathcal{N}$ are separately definable.

`Proof`: $\mathcal{K}$ consists of those partitions π which avoid N_2 together with N_1 which is clearly definable. $\mathcal{N}$ consists of N_1 together with those partitions π which are above N_2 such that $\pi\downarrow$ is linearly ordered. •

We have a nice characterization of those non-finitely generated universal subclasses of equivalence relations.

Proposition IV.1.3. The class of non-finitely generated universal subclasses of equivalence relations is equal to the union of the principal filters generated by $\mathcal{U}(\mathcal{K})$ and $\mathcal{U}(\mathcal{N})$.

`Proof`: Clearly, $\mathcal{U}(\mathcal{K})$ and $\mathcal{U}(\mathcal{N})$ are not finitely generated. Since equivalence relations form a locally finite universal class, a universal subclass is non-finitely generated iff it contains infinitely many non-isomorphic finite equivalence relations. If $\mathcal{R}$ is not finitely generated and avoids some N_m, then there are infinitely many finite partitions with at most m blocks; therefore, $\mathcal{R}$ contains finite partitions with blocks of arbitrarily large cardinality. •

It will be convenient to adjust the representation of a partition. Each total relation on an n-element set will correspond to a complete simple graph K_n on n vertices. Since each partition is a disjoint sum of blocks, and each block can be thought of as a total relation over that block, we would like to represent each partition as a disjoint sum of complete simple graphs. For the partition $\pi = (n_1,\ldots,n_t)$, this can be written as $\pi = \sum_{i=1}^{t} K_{n_i}$ where $n_{i+1} \leq n_i$.

We call a partition π *uniform* if it is the case that all the blocks are of the same cardinality which can be written as $\pi = \sum_{i=1}^{m} K_n \approx mK_n$. It is often useful to write the representation in the form $\pi = \sum_{i=1}^{r} m_i K_{n_i}$ where $n_{i+1} < n_i$ by grouping blocks of equal cardinality together. This will be referred to as the *canonical representation*. We have the following characterization of uniform partitions.

Lemma IV.1.4. π has a unique lower cover iff π is uniform; consequently, the set of uniform partitions is definable without constants.

`Proof`: For necessity, notice that if $\pi \approx mK_n$, then $\forall \rho \prec \pi, \rho \approx (m-1)K_n + K_{n-1}$.

If π is not uniform, then we can write $\pi \approx K_n + \sum_{i=1}^{r} K_{s_i}$ where $n \geq s_i$ and there exists s_k such that $n > s_k$. Then $K_{n-1} + \sum_{i=1}^{r} K_{s_i}$ and $K_n + K_{s_k-1} + \sum_{i \neq k} K_{s_i}$ are incomparable subcovers of π. •

At this point, we should note that those partitions at a fixed given height $n-1$ are definable in $\mathcal{PE}$ without constants, since they are precisely those partitions σ such that $\sigma\downarrow$ contains a chain of covers of length n, but no chain of greater length; for example, the set of isomorphism types $\{\mathbf{k}_2, \mathbf{n}_2\}$ are precisely those types at height 1.

With the proof of the next theorem, positive definability now follows from Theorem II.1.9.

Proposition IV.1.5. Every element of $\langle \mathcal{PE}, \leq, \mathbf{k}_2 \rangle$ is definable. The transpose map is the only non-trivial automorphism of $\langle \mathcal{PE}, \leq \rangle$.

`Proof`: Suppose every element of $\langle \mathcal{PE}, \leq, \mathbf{k}_2 \rangle$ is definable. Since $\{\mathbf{k}_2, \mathbf{n}_2\}$ is definable without constants and closed under the transpose map, Lemma III.7.2 completely characterizes the automorphisms of $\mathcal{PE}$. We now must show individual definability.

Note $K_2 + K_1$ is definable as the unique element with both N_2 and K_2 as lower covers. Together with Proposition IV.1.2, this shows every partition of cardinality at most three is definable. Using Lemma II.2.1, it is sufficient to show set reconstruction holds for partitions of cardinality at least four. While this is has been established as part of a slightly more general reconstruction result by Pretzel and Siemons[PS04], we will present an argument in our setting.

Assume $|\pi|,|\sigma| \geq 4$, and suppose $L_\pi = L_\sigma$. The goal is to show $\pi \approx \sigma$. Since $\rho \prec \pi \Rightarrow |\rho| + 1 = |\pi|$, we must have $|\pi| = |\sigma|$. Also, π and σ must have the same number of blocks in their canonical representations; i.e., if $\pi \approx \sum_{i=1}^{t} m_i K_{n_i}$ and $\sigma \approx \sum_{i=1}^{r} s_i K_{p_i}$, then $t = r$. This follows since $t = |L_\pi|$ and $r = |L_\sigma|$.

Assume $K_{n_i} = K_{p_i}$, but $m_i \neq s_i$. Without loss of generality, suppose $m_i < s_i$. Then every lower cover of σ has at least $s_i - 1$ blocks K_{p_i}, but π has a lower cover which does not. So in the canonical representation reading from the left, whenever the block sizes are equal, they must appear equally often.

If $t = 1$, then π and σ are uniform and so have unique lower covers. Here, $L_\pi = L_\sigma \Rightarrow \pi \approx \sigma$. Now, assume $t > 1$ and $\pi \not\approx \sigma$. Then let k be the first integer $j \leq t$ such that $K_{n_j} \neq K_{p_j}$; therefore, $K_{n_i} = K_{p_i}$ and (by the above), $m_i = s_i$, for $i < k$. Without loss of generality, we may assume $K_{n_k} < K_{p_k}$. If $k = 1$, then no lower cover of π has a block of K_{p_k}, but σ certainly does. If $k > 1$, then $\rho \approx (s_1 - 1)K_{p_1} + K_{p_1 - 1} + \sum_{i=2}^{r} s_i K_{p_i} \prec \sigma$ cannot be a lower cover of π. It must be the case that $\sigma \approx \pi$. •

IV.2 Arithmetic

For our purposes, it will be convenient to actually work with a closely related structure. Let QEQV denote the set of equivalence relations over finite sets of positive integers. We then establish our results on definability in the pointed quasi-ordered set $\langle \mathrm{QEQV}, \leq, K_2 \rangle$; consequently, whenever a particular equivalence relation is shown to be definable, it is definable up to the isomorphism of relations. Unless explicitly stated otherwise, definability refers to this pointed structure with the language $\{\leq, K_2\}$.

The poset $\langle \mathcal{PE}, \leq, \mathbf{k}_2 \rangle$ is then isomorphic to the quotient of $\langle \mathrm{QEQV}, \leq, K_2 \rangle$ by the equivalence determined by isomorphism. The transposition map is also an automorphism of $\langle \mathrm{QEQV}, \leq \rangle$.

For a partition ρ, let $l(\rho)$ equal the number of blocks in ρ. This will be referred to as the length of the partition.

Lemma IV.2.1. $\{(\pi, \rho) : \pi \approx N_m, l(\rho) = m, m \geq 1\}$ is definable.

`Proof`: We see that $l(\rho) = m$ iff $N_m \leq \rho$ but $N_{m+1} \not\leq \rho$. •

Lemma IV.2.2. $\{(k, \pi) : k \approx K_n$ and all blocks of π have at most n elements $\}$ is definable.

`Proof`: That every block of π has at most n elements is given by the condition $K_m \not\leq \pi$ for $K_m \succ K_n$. •

We set $b(\pi) = n$ if $K_n \leq \pi$, but $K_{n+1} \not\leq \pi$. While we have seen that uniform partitions are definable, the following lemmas will allow us to specify particular uniform partitions.

Lemma IV.2.3. $\{(k, n, \pi) : k \approx K_m, n \approx N_n, \pi \approx nK_m\}$ is definable.

`Proof`: We see that (k, n, π) is in this relation iff k is complete, n is trivial, $k \leq \pi$ but $\sigma \not\leq \pi$ when $k \prec \sigma$ and σ complete, $l(\pi) = |n|$, and π has a unique lower cover.

It is immediate that $\pi \approx nK_m$ satisfies the condition.

To see that they are sufficient, we must have $\pi \approx rK_t$ by uniformity, $l(\pi) = |n|$ implies $r = n$, and $t = |k|$ since k is the largest complete partition below π. •

Lemma IV.2.4. $\{(k, \pi) : k \approx K_n$ and all blocks of π are at least size n$\}$ is definable.

`Proof`: (k, π) is in the relation iff $k \approx K_n$ and whenever $l(\pi) = m$, $mK_n \leq \pi$. •

Proposition IV.2.5. $\{\pi :$ all blocks of π are distinct $\}$ is definable.

`Proof`: Let $b(\pi) = n$ and $l(\pi) = t$. Then all the blocks of π are distinct iff $\forall s < t$, there exists $K_{n_s} \leq K_n$ such that $sK_{n_s} \leq \pi$, $(s+1)K_{n_s} \not\leq \pi$, and $sK_p \not\leq \pi$ for $K_p > K_{n_s}$.

Suppose all blocks of π are distinct and order them as $K_{n_1} > K_{n_2} > \cdots > K_{n_t}$ where $n_1 = n$. For $s \leq t$, if $sK_{n_k} \leq \pi$, then $K_{n_k} \leq K_{n_s}$; therefore, K_{n_s} is the largest block of π such that $sK_{n_s} \leq \pi$.

Conversely, suppose π satisfies the conditions and consider the representation of π with $K_{n_1} \geq K_{n_2} \geq \cdots \geq K_{n_t}$ where $n_1 = n$. For $s = 1$, $K_{r_1} \leq \pi$, but $K_{r_1+1} \not\leq \pi$ implies $K_{r_1} \approx K_{n_1} \approx K_n$ since $2K_{r_1} \not\leq \pi$. This also implies $n_2 < n_1 = n$.

For $s = 2$ we have that $2K_{r_2} \leq \pi$ and so there exists $n_i < n_1$ such that $K_{r_2} \leq K_{n_i}$ which implies $K_{r_2} \leq K_{n_2}$. Since $2K_{r_2} \leq 2K_{n_2} \leq \pi$ but $sK_{r_2+1} \not\leq \pi$, we must have $K_{r_2} \approx K_{n_2}$; because $3K_{n_2} \approx 3K_{r_2} \not\leq \pi$, we have $n_3 < n_2$.

As we continue inductively, for $s < t$ we have $sK_{r_s} \leq \pi$ and so there exists $n_i < n_{s-1}$ such that $K_{r_s} \leq K_{n_i}$ which implies $K_{r_s} \leq K_{n_s}$. Since $sK_{r_s} \leq sK_{n_s} \leq \pi$ but $2K_{r_s+1} \not\leq \pi$, we must have $K_{r_s} \approx K_{n_s}$; because $(s+1)K_{n_s} \approx (s+1)K_{r_s} \not\leq \pi$, we have $n_{s+1} < n_s$.

For $s = t-1$, the conclusion $n_t < n_{t-1}$ finishes demonstrating that all n_i are distinct. •

We can now specify the existence of a particular block.

Proposition IV.2.6. $\{(k, \pi) : k \approx K_n$ and K_n is a block of $\pi\}$ is definable.

`Proof`: (k, π) is in this relation iff $k \approx K_n$ for some $n \geq 1$, $k \leq \pi$, and where $l(\pi) = t$ there exists $N_r \leq N_t$ such that $rK_n \leq \pi$, $(r+1)K_n \not\leq \pi$, and $rK_p \not\leq \pi$ for $K_p > K_n$.

To see that this characterizes the presence of a block K_r in π, notice that in the canonical representation for $\pi \approx \sum_{i=1}^{t} m_i K_{n_i}$, K_{n_k} appears m_k times, but the largest uniform partition with blocks of size n_k below π is $\left(\sum_{i=1}^{k} m_i\right) K_{n_k}$. For the block K_{n_k} the value of r we are after is then $r = \sum_{i=1}^{k} m_i$. •

Definition IV.2.7. For $n \geq 1$, a partition $\sigma \approx \sum_{i=1}^{n} K_i$ is called a *factorial* and will be denoted as $[n]!$.

Our approach to the definability of addition and multiplication is to first show that factorials are definable.

Proposition IV.2.8. $\{(k,\pi) : k \approx K_n,\ \pi \approx [n]!\ \}$ is definable.

`Proof`: Th claim is that π is isomorphic to the factorial $[n]!$ iff

(1) $b(\pi) = n$

(2) For all $K_r \leq K_n$ we have that K_r is a block of π

(3) All the blocks of π are distinct

If $\pi \approx [n]!$, then it is easy to see the conditions are satisfied.

Suppose π satisfies conditions (1) - (3). Conditions (1) and (2) imply $\pi \approx \sum_{i=1}^{n} m_i K_i$, and condition (3) implies each $m_i = 1$. •

We can now define the pairs of complete and trivial partitions which are at the same height.

Lemma IV.2.9. $\{(k,n) : k \text{ is complete},\ n \text{ is trivial}, |k| = |n|\}$ is definable.

`Proof`: (k,n) is in this relation iff $k \approx K_r$, $n \approx N_m$, and $l(\pi) = m$ where $\pi \approx [r]!$. •

With factorials, we don't have to start counting the components just from K_1; this means we can now do addition.

Proposition IV.2.10. $\{(k,r,\pi) : k,r,\pi \text{ are complete and } |k| + |r| = |\pi|\}$ is definable.

`Proof`: (k,r,π) is in this relation iff k,r,π are complete, $k,r < \pi$, and where ρ is the partition in which all the blocks are distinct, and K_m is a block of ρ iff $k < K_m \leq \pi$, we then have that $l(\rho) = |r|$. The last condition is definable by Lemma IV.2.9. •

It follows from the last two propositions that we can also do addition by considering the corresponding triplets of trivial relations.

We may refer to a partition of the form mK_n as n-uniform to denote the fact that all the blocks are of cardinality n. We will also say mK_n has *size* n. The *frequency* refers to m. We saw in Lemma IV.2.3 that the set of n-uniform partitions is definable; moreover, it is easy to see that the set of n-uniform partitions are linearly ordered. The next result allows us to pick out the uniform partitions which appear in a canonical representation.

Proposition IV.2.11. The relation

$$\{(k,n,\pi) : n \approx N_n \text{ and } k \text{ is a block of } \pi \text{ which appears exactly } n \text{ times }\}$$

is definable.

`Proof`: Let $k \approx K_r$ and $n \approx N_n$. We have that K_r is a block of π which appears exactly n times iff $\pi \approx nK_r$, or

(1)K_r is block of π,

(2) If $b(\pi) = r$, then nK_r is the maximal r-uniform partition below π.

(3) If $b(\pi) \neq r$ and there exists $K^* > K_r$ such that

(a) K^* is a block of π such that whenever K_s is a block of π above K_r, we have that $K_s \geq K^* > K_r$, and

(b) mK_r is the maximal r-uniform partition below π, and

(c) tK^* is the maximal $|K^*|$-uniform partition below π

then $n = m - t$.

If we examine the canonical representation of $\pi \approx \sum_{i=1}^{s} m_i K_{n_i}$, then for any block K_{n_r}, we see that $\left(\sum_{i=1}^{r} m_i\right) K_{n_r}$ is the largest n_r-uniform partition below π, and so the correctness of the above characterization follows since $m_r = \sum_{i=1}^{r} m_i - \sum_{i=1}^{r-1} m_i$. That the characterization is first-order is guaranteed by Lemma IV.2.4 and Lemma IV.2.9. •

Proposition IV.2.12. $\{(k,\pi) : k \text{ is complete and } |\pi| \geq |k|\}$ is definable.

`Proof`: The claim is that (k,π) is in the relation iff k is complete, and for any partition σ which satisfies the conditions below, we have $l(\sigma) \geq |k|$:

$(**)$ If $K_r \leq \pi$ and $mK_r \leq \pi$ but $(m+1)K_r \not\leq \pi$ for some m, then K_r is a block of σ which appears at least m times.

To check necessity, let $\pi \approx \sum_{i=1}^{t} m_i K_{n_i}$ with $|\pi| \geq n = |k|$, and suppose σ is a partition which satisfies the condition $(**)$. We wish to show $l(\sigma) \geq n$. Set $M_r = \sum_{i=1}^{r} m_i$ and note that $M_1 < M_2 < \cdots < M_t$. Since $M_i K_{n_i}$ is a maximal n_i-uniform partition below π, we must have that K_{n_i} appears as a block in σ M_i times. This implies $\sum_{i=1}^{t} M_i K_{n_i} \leq \sigma$. For an arbitrary block K_r such that $K_{n_{s+1}} < K_r < K_{n_s}$, $(**)$ implies K_r is a block of σ and must appear at least M_s times in σ; in particular, $[r]! \leq \sigma$ whenever K_r is a block of π. Altogether, it must be the case that $\sum_{i=1}^{t} m_i [n_i]! \leq \sigma$ and so

$$l(\sigma) \geq l\left(\sum_{i=1}^{t} m_i [n_i]!\right) = \sum_{i=1}^{t} l\left(m_i [n_i]!\right) = \sum_{i=1}^{t} m_i n_i = |\pi| \geq n.$$

To establish sufficiency, suppose (k,π) is in the relation, but $|\pi| < n = |k|$. Let $\pi \approx \sum_{i=1}^{t} m_i K_{n_i}$. Set $\rho \approx \sum_{i=1}^{t} m_i [n_i]!$ and observe that $l(\rho) = \sum_{i=1}^{t} m_i n_i < n$. Suppose $K_s \leq \pi$ and let k be the smallest number for which $K_s \leq K_{n_k}$. For the block K_{n_k}, we see that $\left(\sum_{i=1}^{k} m_i\right) K_{n_k} \leq \pi$ is maximal. If $rK_s \not\leq \pi$ for $r > \sum_{i=1}^{k} m_i$, then by definition of the canonical representation, we must have $K_s \not\leq K_{n_{k+1}}$ which contradicts the choice of

K_{n_k}; therefore, $\left(\sum_{i=1}^{k} m_i\right) K_s \leq \pi$ is maximal among s-uniform partitions. Notice that K_{n_k} appears in ρ for each factorial $[n_r]!$ where $n_r > n_k$; that is, K_{n_k} appears $\sum_{i=1}^{k} m_i$ times which is exactly how often K_s appears as a block in ρ. We have shown the partition ρ satsifies $(**)$, and so we arrive at a contradiction. It must be the case that $|\pi| \geq n$ •

By examining the above proof, it is interesting to note that we have essentially shown the definability of $\sigma \approx \sum_{i=1}^{t} m_i[n_i]!$ given $\pi \approx \sum_{i=1}^{t} m_i K_{n_i}$.

Proposition IV.2.13. $\{(k,\pi) : k$ is complete and $|\pi| = |k|\}$ is definable.

`Proof`: We would have $|\pi| \geq |k|$, but $|\pi| \not\geq |k|+1$. •

We can now interpret multiplication.

Proposition IV.2.14. $\{(k,\rho,\pi) : k,\rho,\pi$ are complete and $|\pi| = |k|\,|\rho|\}$ is definable.

`Proof`: (k,ρ,π) is in this relation iff k,ρ,π are complete and $|\pi| = |\sigma|$ where $\sigma \approx |k|K_{|\rho|}$. •

Let $\langle N_{>0}, +, \times\rangle$ denote the structure over the set of positive integers such that the operations of addition and multiplication have their usual meaning. From [TMR53, Thm 7], the elementary theory of this structure is undecidable. Propositions IV.1.2, IV.2.10, and IV.2.14 state that we can define the operations of addition and multiplication over the definable set of complete partitions, and so establish a first-order interpretation of the elementary theory of $\langle N_{>0}, +, \times\rangle$ into the elementary theory of $\langle \mathcal{PE}, \leq, \mathbf{n}_2\rangle$. According to [TMR53, Thm 7&10] this yields the following result.

Corollary IV.2.15. The elementary theory of $\langle \mathcal{PE}, \leq\rangle$ is undecidable.

Since the elementary theory of a fixed structure is complete, by [TMR53, Thm 1] we can conclude the following.

Corollary IV.2.16. The elementary theory of $\langle \mathcal{PE}, \leq\rangle$ is not finitely axiomatizable.

IV.3 Morphisms

Here is our scheme for encoding a function $f : [n] \to [m]$. We take a partition

$$\xi(n,f,m) \approx \sum_{i=1}^{n} m_i K_i + K_{n+1+m} \quad \text{where} \quad m_i = f(i).$$

In such a partition, K_n is the largest block of $\xi(n,f,m)$ such that K_{n+1} is not a block, but there does exist a strictly larger block. The next block K_r is the unique largest block, and $m = r-n-1$. This is how to read off the domain and range of f. Clearly, any such partition $\sum_{i=1}^{n} m_i K_i + K_{n+1+m}$ where $m_i \leq m$ defines a unique function $f : [n] \to [m]$ where we set $f(i) = m_i$.

We can realize a partition τ as $\xi(n,f,m)$ for some f in the following way:

(1) If m_rK_r is the uniform block of τ with largest size, then $m_r = 1$.

(2) If m_nK_n is the next largest uniform block of τ, then for all $1 \leq i < n$ we have a uniform block m_iK_i.

(3) $n+1 < r$

(4) If m_iK_i is a uniform block in τ, then $m_i \leq r-n-1$.

We have the following.

Proposition IV.3.1. We have the following:

(1) $\xi(n,f,m) \approx \xi(n',g,m')$ iff $n = n'$, $m = m'$ and $f = g$.

(2) $\{(K_n, K_m, \tau) : n, m > 0, \text{ and } \tau \approx \xi(n,f,m) \text{ for some } f : [n] \to [m]\}$ is definable.

IV.4 o-Presentations

Let σ be a concrete partition on the set $[n]$. We would like to encode in the isomorphism type of another partition information which records which elements of $[n]$ are in the same block in σ. Consider the set of partitions A_σ where $\pi = \sum_{i=1}^{n} m_iK_i \in A_\sigma$ if it satisfies the condition

$$m_i = m_j \quad \text{iff} \quad (i,j) \in \sigma. \tag{IV.2}$$

We can uniquely reconstruct σ from any $\pi \in A_\sigma$ by matching together into separate groups those uniform blocks which have the same frequency. For each such group, corresponding sizes of the uniform blocks precisely describes the elements which are in the same block of σ. Here, the actual frequency numbers are irrelevant since any member of A_σ will do. Since we would like our choice of a partition in A_σ to be definable, we must be more judicious.

For any $\pi \in \mathrm{QE}\textsc{qv}$, the isomorphism type $[\pi]$ contains a concrete partition we label π^* defined in the following manner. Let $\pi \approx \sum_{i=1}^{n} m_iK_{n_i}$ so that π is a partition of a set with $m = \sum_{i=1}^{n} m_in_i$ elements. Now π^* is a partition on $[m]$ where

- The first m_1n_1 integers are divided into m_1 blocks where the first block contains $1,\ldots,n_1$, the second block contains $n_1+1,\ldots.,2n_1$, and continuing in this manner the m_1-th block contains the integers $(m_1-1)n_1+1,\ldots.,m_in_i$
- The next m_2n_2 consecutive integers are partitioned in a similar manner.
- We continue partitioning consecutive intervals of m_in_i integers until we exhaust the uniform blocks of π.

Any element of A_{π^*} has a peculiar form - all the uniform blocks with the same frequencies appear as a consecutive interval in the sizes of the uniform blocks. We are now ready to define an o-presentation.

Definition IV.4.1. For any $\pi \in \mathcal{QE}$, choose $P(\pi) \in A_{\pi^*}$ such that whenever

$$\{mK_p, mK_{p+1}\ldots., mK_q\}$$

is a complete set of uniform blocks in $P(\pi)$ which have the same frequency, then $m = p$. We say $P(\pi)$ is an *o-presentation* for π.

Example IV.4.2. If we take $\pi = 3K_3 + 2K_5$, then

$$\begin{aligned} P(\pi) \quad = \quad & \underbrace{K_1 + K_2 + K_3} + \underbrace{4K_4 + 4K_5 + 4K_6} + \underbrace{7K_7 + 7K_8 + 7K_9} + \\ & \underbrace{10K_{10} + 10K_{11} + 10K_{12} + 10K_{13} + 10K_{14}} + \\ & \underbrace{15K_{15} + 15K_{16} + 15K_{17} + 15K_{18} + 15K_{19}} \end{aligned}$$

Given $\sigma \in \mathrm{QEQV}$, suppose $sK_s, \ldots, sK_r$ are all the uniform blocks in $P(\sigma)$ which have the same frequency. We call $sK_s + \cdots + sK_r$ a *pseudo-block* of $P(\sigma)$. The terminology comes from the fact that a pseudo-block in $P(\sigma)$ reflects the existence of a block in σ.

Lemma IV.4.3. The set $\{(K_s, \rho, K_r) : \rho \approx sK_s + \cdots + sK_r\}$ is definable.

`Proof`: We see that $\rho \approx sK_s + \cdots + sK_r$ iff whenever $s \leq i \leq r$, then sK_i is a uniform block of ρ and these are the only uniform blocks. •

Proposition IV.4.4. The following hold.

(1) The set $\{P(\pi) : \pi \in \mathrm{QEQV}\}$ is definable.
(2) The set $\{(\rho, P(\pi)) : \rho$ is a pseudo-block of $P(\pi)\}$ is definable.
(3) The set $\{(\sigma, P(\pi)) : \sigma \approx \pi\}$ is definable.
(4) $P(\pi) \approx P(\sigma)$ iff $\pi^* = \sigma^*$.

`Proof`:

(1) We see that $\sigma \approx P(\pi)$ for some π iff

- If mK_r is the uniform block of σ with the largest size, then for all $1 \leq i \leq r$ there exists n such that nK_i is a uniform block of σ.
- Whenever mK_s and mK_r are two uniform blocks of σ with $s < r$, then mK_i is a uniform block of σ for all $s \leq i \leq r$.
- Whenever mK_s and nK_r are two uniform blocks of σ with $m \neq n$ and $s < r$, then $m < n$.
- If mK_s is a uniform block which is smallest for all uniform blocks with the same frequency, then $m = s$.

The first three conditions imply $\sigma \in A_\pi$ for some π. The last condition guarantees σ is a disjoint union of pseudo-blocks and so is an o-presentation.

(2) $\rho \approx sK_s + \cdots + sK_r$ is a pseudo-block of $P(\pi)$ if sK_s and sK_r are uniform blocks of $P(\pi)$, and among all uniform blocks which have frequency s, sK_s is smallest and sK_r is the largest in size.

(3) From the canonical representation, the isomorphism type of a partition π is completely determined by the uniform blocks in π. The central question is, how do we read off the uniform blocks of π from the collection of pseudo-blocks of $P(\pi)$?

For a pseudo-block (K_s, ρ, K_r), we define the difference $d(\rho) = r + 1 - s$. Among all the pseudo-blocks of $P(\pi)$ which have the same difference, let $(K_{s_1}, \rho_1, K_{r_1})$ be the one in which s_1 is smallest, and let $(K_{s_2}, \rho_2, K_{r_2})$ be the one in which s_2 is largest. Then $\frac{r_2+1-s_1}{r_2+1-s_2} K_{r_2+1-s_2}$ is a uniform block of π.

Any such partition as $\frac{r_2+1-s_1}{r_2+1-s_2} K_{r_2+1-s_2}$ derived from $P(\pi)$ in this manner is called a *difference block* of $P(\pi)$. Then π is up to isomorphism the unique partition whose uniform blocks are precisely the difference blocks of $P(\pi)$.

(4) By construction, we have $\pi \approx \sigma$ iff $\pi^* = \sigma^*$, and from the above it follows that $\pi \approx \sigma$ iff $P(\pi) \approx P(\sigma)$. •

Example IV.4.5. If we take $P(\pi)$ from the previous example, then we can reconstruct the isomorphism type of π as

$$\frac{9-0}{3-0} K_{9+1-7} + \frac{19+1-10}{19+1-15} K_{19+1-15} = 3K_3 + 2K_5.$$

IV.5 A Small Category

In the same manner as in Section III.9, we define a small category $\mathcal{CE}$. The objects are precisely the concrete partitions π^* where $\pi \in \mathrm{QE}\textsc{qv}$. The morphisms $\mathcal{CE}(\pi^*, \sigma^*)$ are the relational homomorphisms from π^* to σ^* which we write as a triple $F = (\pi^*, f, \sigma^*)$ where $f : [n] \to [m]$ with $n = |\pi^*|$ and $m = |\sigma^*|$. The category $\mathcal{CE}$ can naturally be thought of as a 2-sorted first-order structure, with one sort for objects and another sort for morphisms, together with a ternary relation over the sort of morphisms which reflects composition. The category structure is then described by the standard category axioms in this 2-sorted first-order language.

We enrich the small category by adding 3 new constants to the language and denote the resulting structure by $\mathcal{CE}'$. We add the constants $\mathbf{K}_2^*$ and the maps $\mathcal{CE}(N_1^*, \mathbf{K}_2^*) = \{\mathbf{t}, \mathbf{b}\}$. Notice that K_1^* is a terminal object in $\mathcal{CE}'$, and so is definable. For any $\pi^* \in \mathrm{Obj}\,\mathcal{CE}$, we can use the maps in $\mathcal{CE}(\mathbf{K}_2^*, \pi^*)$ to parametrize the partition structure in the following way: for any i, j in the universe of $\pi^* = \langle [n], r \rangle$, we see that $(i, j) \in r$ iff where $x, y \in \mathcal{CE}(K_1^*, \pi^*)$ such that $x(0) = i$ and $y(0) = j$, then there exists $h \in \mathcal{CE}(\mathbf{K}_2^*, \pi^*)$ such that $h\mathbf{t} = x$ and $h\mathbf{b} = y$. This clearly yields a definable equivalence relation over $\mathcal{CE}(K_1^*, \pi^*)$ denoted by $\hat{r}$.

By a similar argument as in Section III.9, the property that a morphism is an embedding is definable in the first-order language of $\mathcal{CE}'$. We conclude that the embeddability relation of $\langle \mathrm{QE}\textsc{qv}, \leq, N_2 \rangle$ when restricted to the objects of the small category is first-order definable in $\mathcal{CE}'$.

Using the bijection $\pi^* \ni i \leftrightarrow x \in \mathcal{CE}(K_1^*, \pi^*)$ such that $x(0) = i$, we see that given any $\pi^* \in \mathrm{Obj}\,\mathcal{CE}$, we can construct an isomorphic equivalence relation

$$\Gamma(\pi^*) = \langle \mathcal{CE}(K_1^*, \pi^*), \hat{r} \subseteq \mathcal{CG}(K_1^*, \pi^*) \times \mathcal{CG}(K_1^*, \pi^*) \rangle \approx \pi^*$$

where both the set of elements and the relation $\hat{r}$ have first-order definitions in the language of the small category. The set of partitions in $\{\Gamma(\pi^*) : \pi^* \in \mathrm{Obj}\,\mathcal{CE}\}$ provide isomorphic copies of the objects of the category. For these ojects, we now have access to the “internal” stucture using the first-order theory of the category.

Putting all this together, the first-order language of $\mathcal{CE}'$ is equivalent in expressive power to a second-order language applied to the structures $\{\Gamma(\pi^*) : \pi^* \in \mathrm{Obj}\,\mathcal{CE}\}$. This new language has variables ranging over the elements of $\Gamma(\pi^*)$, variables for the morphisms between objects, can express the partition relation in objects, machinery to express the application of morphisms to elements, composition of morphisms, and equality of elements and morphisms. The procedure in [JM10] to parametrize arbitrary finitary relations works equally well in this small category of equivalence relations. Altogether, we see that the first-order language of $\mathcal{CE}'$ is equivalent in expressive power to a full second-order language over the structures $\{\Gamma(\pi^*) : \pi^* \in \mathrm{Obj}\,\mathcal{CE}\}$.

In order to establish the version of III.9.3 for partitions, we need to build a model of the small category $\mathcal{CE}'$ in the definable relations of $\langle \mathrm{QE}\textsc{qv}, \leq K_2\rangle$. With Propositions IV.4.4 and IV.3.1, almost all the work has been accomplished already. What remains is definably capturing the property that a function $f : n \to m$ encoded as $\xi(n, f, m)$ is a homomorphism between partitions π^* and σ^* encoded as $P(\pi)$ and $P(\sigma)$ where $|\pi| = n$ and $|\sigma| = m$.

Suppose we have objects π_i^* in the category with $m_i = |\pi_i^*|$ for $i = 1,2$, and a morphism $F = (\pi_1^*, f, \pi_2^*) \in \mathcal{CE}(\pi_1^*, \pi_2^*)$ where $f : [m_1] \to [m_2]$. We encode π_i^* as any graph isomorphic to $P_i = P(\pi_i^*)$ and encode F as any triple isomorphic to the ternary relation $M(F) = (P_1, \xi(m_1, f, m_2), P_2)$. Using Propositions IV.4.4 and IV.3.1 we can realize for $M(F) = (P_1, \xi(m_1, f, m_2), P_2)$ that P_1 and P_2 are o-presentations for some π and σ, respectively, and that $\xi(m_1, f, m_2)$ encodes some function from π^* to σ^* where $m_1 = |\pi^*|$ and $m_2 = |\sigma^*|$.

In the following lemma, we will see how to read off the values of the function f with statement (1), and how to capture the fact that f is a homomorphism with statement (2).

Lemma IV.5.1. We have the following:

1. If $(A,S,B) \approx M(F)$ for $F = (\pi_1^*, f, \pi_2^*)$, then f is uniquely determined and for all $i \in [m_1]$ and $j \in [m_2]$ where $|\pi_1^*| = m_1$ and $|\pi_2^*| = m_2$, we have that $f(i) = j$ iff jK_i is a uniform block of S.

2. $(A,S,B) \approx M(F)$ for some $F = (\pi_1^*, f, \pi_2^*) \in \mathcal{CE}(\pi_1^*, \pi_2^*)$ iff where $m_i = |\pi_i^*|$, we have $A \approx P(\pi_1^*)$, $B \approx P(\pi_2^*)$, and $S \approx \xi(m_1, f, m_2)$ for some $f : [m_1] \to [m_2]$; and whenever we have $1 \leq i, i' \leq m_1$ and $1 \leq j, j' \leq m_2$, $j \neq j'$, and $jK_i, j'K_{i'}$ are uniform blocks of S, then the blocks K_i and $K_{i'}$ have the same frequency in A implies the blocks K_j and $K_{j'}$ have the same frequency in B.

The next lemma captures the composition of morphisms.

Lemma IV.5.2. Consider morphisms $F = (\pi_1^*, f, \pi_2^*)$ and $G = (\pi_2^*, g, \pi_3^*)$ with $|\pi_i^*| = m_i$ for $i = 1,2,3$. Let $M(F) \approx (P_1, S_1, P_2)$ and $M(G) \approx (P_2, S_2, P_3)$. Then $M(GF) \approx (P_1, H, P_3)$ iff $H \approx \xi(m_1, h, m_3)$ for some h such that: for all $i \in [m_1]$, $j \in [m_2]$, $k \in [m_3]$, we have that jK_i is a uniform block of S_1 and kK_j is a uniform block of S_2 imply kK_i is a uniform block of H.

We have gathered together all the necessary definable relations. The proof of Theorem III.9.3 for simple graphs goes through exactly word for word in this setting with one slight change - we need only modify the translation of first-order formulas in the language of $\mathcal{CE}'$ taking into account our scheme for o-presentations(Theorem IV.4.4). The scheme for the translation of the atomic formulas in the case of equivalence relations is the following:

1. If ϕ is $X_r = X_s$, then $\hat{\phi}$ is $x_r \leq x_s \wedge x_s \leq x_r$.
2. If ϕ is $f_s = f_r$, then $\hat{\phi}$ is $y_r \leq y_s \wedge y_s \leq y_r$
3. If ϕ is $f_s \in \mathcal{CE}(X_r, X_t)$, then $\hat{\phi}$ is

$$(\exists u_r, u_t)(\text{“}x_r = P(u_r) \text{ and } x_t = P(u_t) \text{ and } (x_r, y_s, x_t) = M(F) \text{ for some } F \in \mathcal{CE}(u_r^*, u_t^*)\text{”})$$

4. If ϕ is $f_i \in \mathcal{CE}(X_r, X_t) \wedge f_j \in \mathcal{CE}(X_t, X_s) \wedge f_k = f_j \circ f_i$, then $\hat{\phi}$ is

$$(\exists u_r, u_t, u_s)(\text{“}x_m = P(u_m) \text{ for } m \in \{r,t,s\} \text{ and } (x_r, y_i, x_t) = M(F) \text{ for some } F \in \mathcal{CE}(u_r^*, u_t^*) \text{ and } (x_t, y_j, x_s) = M(G) \text{ for some } G \in \mathcal{CE}(u_t^*, u_s^*) \text{ and } (x_r, y_k, x_s) = M(GF)\text{”})$$

5. If ϕ is $\neg\psi$, or $\psi \wedge \theta$, then $\hat{\phi}$ is $\neg\hat{\psi}$, or $\hat{\psi} \wedge \hat{\theta}$
6. If ϕ is $(\exists X_r)\psi$, then $\hat{\phi}$ is $(\exists x_r)[(\exists u_r)(\text{“}x_r = P(u_r)\text{”}) \wedge \hat{\psi}]$
7. If ϕ is $(\forall X_r)\psi$, then $\hat{\phi}$ is $(\forall x_r)[(\exists u_r)(\text{“}x_r = P(u_r)\text{”}) \to \hat{\psi}]$
8. If ϕ is $(\exists f_s \in \mathcal{CE}(X_r, X_t))\psi$, then $\hat{\phi}$ is

$$(\exists y_s)[(\exists u_r, u_t)(\text{“}x_r = P(u_r) \text{ and } x_t = P(u_t) \text{ and } (x_r, y_s, x_t) = M(F) \text{ for some } F \in \mathcal{CE}(u_r^*, u_t^*)\text{”}) \wedge \hat{\psi}]$$

9. If ϕ is $(\forall f_s \in \mathcal{CE}(X_r, X_t))\psi$, then $\hat{\phi}$ is

$$(\forall y_s)[(\exists u_r, u_t)(\text{“}x_r = P(u_r) \text{ and } x_t = P(u_t) \text{ and } (x_r, y_s, x_t) = M(F) \text{ for some } F \in \mathcal{CE}(u_r^*, u_t^*)\text{”}) \to \hat{\psi}]$$

Now the argument of Theorem III.9.3 follows mutatis mutandis to conclude the following.

Theorem IV.5.3. Let R be an isomorphism invariant relation over QEQV. Then R is first-order definable over $\langle \text{QEQV}, \leq, K_2 \rangle$ iff its restriction to Obj $\mathcal{CE}$ is first-order definable in the language of $\mathcal{CE}'$.

Corollary IV.5.4. For every sentence ϕ in the second-order language of equivalence relations, there is a formula $\Phi(x)$ in the first-order language of the quasi-ordered set $\langle \text{QEQV}, \leq, K_2 \rangle$ such that an equivalence relation π in QEQV models ϕ if and only if $\langle \text{QEQV}, \leq, K_2 \rangle \vDash \Phi(\pi)$.

IV.6 The Number of Universal Classes

A quasi-ordered set Q is said to be *well-quasi-ordered*([Mil85]) if there are no infinite strictly descending chains, and no infinite anti-chains. Let $\mathcal{O}(Q)$ and $\mathcal{F}(Q)$ denote the set of order ideals and order filters, respectively ordered by inclusion. It is easy to see that if Q is well-quasi-ordered, then $\mathcal{O}(Q)$ has the descending chain condition and $\mathcal{F}(Q)$ has the ascending chain condition.

There is a nice connection between locally finite universal theories and well-quasi-orderings. The following two lemmas are essentially contained in a 1967 paper of A.I. Mal'cev[Mal67].

Lemma IV.6.1. Let $\mathcal{U}$ be a locally finite universal class of finite signature. $\mathcal{PU}$ is well quasi-ordered iff there are only countably many universal subclasses; otherwise, there are continuum many.

`Proof`: If $\mathcal{PU}$ is not well-quasi-ordered, then it must have an infinite anti-chain $\{A_i : i \in \omega\}$. This provides $2^{\aleph_0}$ many distinct universal subclass $\{\mathcal{U}(A_J) : J \subseteq \omega\}$.

Suppose $\mathcal{PU}$ is well-quasi-ordered. For $J \in \mathcal{O}(\mathcal{PU})$, let X_J denote the set of elements minimal in $\mathcal{PU} \setminus J$ which must form an anti-chain. Since $\mathcal{PU}$ is well-quasi-ordered, X_J is finite. Since X_J uniquely determines J, $|\mathcal{L}_{\mathcal{U}}| = |O(\mathcal{PU})| \leq \aleph_0$. •

Proposition IV.6.2. Suppose $\mathcal{U}$ is a finitely axiomatizable locally finite universal class of finite signature, and there are only finitely many N-generated structures up to isomorphism. Then every universal subclass is finitely axiomatizable iff $\mathcal{PU}$ is well quasi-ordered.

`Proof`: If $\mathcal{PU}$ is not well quasi-ordered, then there are continuum many universal subclasses. Since there are only countably many finite sets of sentences, there are non-finitely axiomatizable universal subclasses.

Suppose $\mathcal{PU}$ is well quasi-ordered. Let $\mathcal{K} \leq \mathcal{U}$ be a universal subclass. We may assume it is not finitely generated; thus, $\mathcal{K} = \mathcal{U}(I)$ for some infinite order ideal of $\mathcal{PU}$. For each n, let $H_n = \{A : |A| \leq n, A \notin I\}$. Let $F_n = H_n \uparrow$. Then each F_n is a finitely generated order filter and $I = (\bigcup F_n)^c$. Also, $F_1 \subseteq F_2 \subseteq \cdots$ is an ascending sequences of filters and so must converge; that is, there exists M such that $F_k = F_M$ for all $k \geq M$. Then $\mathcal{K} = \mathcal{U}(F_M^c)$ which implies $\mathcal{K}$ is finitely axiomatizable. •

A sequence $(a_n)_{n\in\omega}$ in a quasi-ordered set Q is *bad* if $a_i \not\leq a_j$ for $i < j$. It is not difficult to see that Q well-quasi-ordered(wqo) iff it has no bad sequences. The next result follows from a more general theorem of Higman [Hig52] on finite sequences, but we will provide a proof in our setting.

Proposition IV.6.3. The poset $\mathcal{PE}$ is well quasi-ordered; as a result, there are only countably many universal classes of equivalence relations every one of which is finitely axiomatizable.

`Proof`: It suffices to show $\mathcal{PE}$ has no bad sequences. For a contradiction, assume $(\pi_i)_{n\in\omega}$ is a bad sequence. Write each π_i as

$$\pi_i = \sum_{k=1}^{r_i} K^i_{t_k}$$

such that the blocks are non-increasing in cardinality. For each $i < j$, let $\alpha(i,j)$ be the smallest n such that

$$\sum_{k=1}^{n-1} K^i_{t_k} \leq \pi_j \quad \text{but} \quad \sum_{k=1}^{n} K^i_{t_k} \not\leq \pi_j.$$

For π_1, we have $\alpha(1,-) : \omega \smallsetminus \{1\} \to [r_1]$ and so there exists Δ_1 an infinite subset such that $\alpha(1,\Delta_1)$ is constant. Note

$$K^1_{t_{\alpha(1,j)}} > K^j_{t_{\alpha(1,j)}} \quad \text{for all} \quad j \in \Delta_1.$$

Choose $j_1^* \in \Delta_1$ such that $t_{\alpha(1,j_1^*)}$ is smallest.

Again we extract an infinite subset Δ_2 such that $\alpha(j_1^*,\Delta_2)$ is constant. Note

$$K^{j_1^*}_{t_{\alpha(j_1^*,j)}} > K^j_{t_{\alpha(j_1^*,j)}} \quad \text{for all} \quad j \in \Delta_2.$$

Choose $j_2^* \in \Delta_2$ such that $\alpha(j_1^*, j_2^*)$ is smallest.

Inductively, we find a sequence $\{j_1^*, j_2^*, \ldots\}$ such that

$$K^{j_p^*}_{t_{\alpha(j_p^*,j_{p+1}^*)}} > K^{j_{p+1}^*}_{t_{\alpha(j_p^*,j_{p+1}^*)}}$$

Now, suppose $\alpha(i,j)$ is bounded on the sequence $\{j_1^*, j_2^*, \ldots\}$. Then there exists an infinite subset $\theta \subseteq \{j_1^*, j_2^*, \ldots\}$ such that $\alpha(i,j)$ is constant. Say $\theta = \{k_1, k_2, \ldots\}$. This produces an infinite descending sequence

$$K^{k_1}_{t_{\alpha(k_1,k_2)}} > K^{k_2}_{t_{\alpha(k_2,k_3)}} > \cdots$$

which is a contradiction.

If $\alpha(i,j)$ is not bounded on $\{j_1^*, j_2^*, \ldots\}$, then there must exist a subsequence $\{s_1, s_2, \ldots\}$ such that

$$\alpha(s_i, s_{i+1}) < \alpha(s_{i+1}, s_{i+2}).$$

But then

$$K^{s_1}_{t_{\alpha(s_1,s_2)}} > K^{s_2}_{t_{\alpha(s_1,s_2)}} \geq K^{s_2}_{t_{\alpha(s_2,s_3)}} > K^{s_3}_{t_{\alpha(s_2,s_3)}} \geq \cdots$$

provides an infinite descending sequence; another contradiction. ●

CHAPTER V

QUESTIONS

In a review of substructure definability for distributive lattices([JM09b]), semilattices([JM09a]), lattices([JM09c]), posets([JM10]), and in extending these results to the unordered structures of equivalence relations and simple graphs, similiarities in constructions and arguments abound, but in each case there is enough distinction so that the approach almost starts over again each time. Having established strong definability results for these classes, can we abstract the combinatorial or model properties which may guarantee similar substructure definability in general universal classes? Even restricted to universal subclasses of digraphs, this appears to me a difficult question.

Does positive definability hold for the universal class of tournaments? Reconstruction([Sto77]), and thus set reconstruction, fails for tournaments, and so Lemma II.2.1 can offer no help in this case. The relationship between substructure definability and reconstruction is unclear, but perhaps we can offer another link with the following immediate corollary to Proposition III.7.1 and Theorem 2.29 in [JM10].

Corollary V.0.4. If P,Q are two counterexamples to Reconstruction(as posets or as simple graphs), then P and Q have distinct sets of upper covers.

`Proof`: If P and Q are two counterexamples to Reconstruction, then $|P|,|Q| > 4$ and must have the same lower covers in the substructure ordering since they have the same lower decks. If in addition they have the same upper covers, then P and Q must satisfy the same unary formulas in the language of $\{\leq, \mathbf{c}\}$ where $|\mathbf{c}| \leq 4$. This contradicts the fact that for simple graphs and posets each element is definable after adding a single constant of cardinality at most 3. ●

For tournaments, there is an obvious automorphism **rev** of the substructure ordering which comes from reversing the orientation of the edges. The counterexamples to the Reconstruction Conjecture for digraphs discoverd by Stockmeyer([Sto77]) appear in two infinite families (B_i, C_i) and (D_i, E_i). Interestingly, $\mathbf{rev}(B_i) = B_i$, $\mathbf{rev}(C_i) = C_i$, and $\mathbf{rev}(D_i) = E_i$. This is precisely what one must have if it is the case that the sets $\{B_i, C_i\}$ and $\{D_i, E_i\}$ are definable. This prompts the following two questions.

Question V.0.5. After adding a constant, every finite isomorphism type of tournaments is first-order definable in the poset of finite isomorphism types ordered by substructure; moreover, **rev** is the only non-trivial automorphism. Positive definability for universal theories holds.

Question V.0.6. Each pair of Stockmeyer's counterexamples $\{B_i, C_i\}$ and $\{D_i, E_i\}$ are definable in the poset of finite isomorphism types ordered by substructure without adding a constant to the language.

We saw that for posets, simple graphs, and equivalence relations the expressive power of first-order definability in the substructure relation was equivalent to modeling full second-order sentences when restricted to the finite members. Is this to be expected in general?

Question V.0.7. Let $\mathcal{U}$ be a finitely axiomatizable locally finite universal class in a finite signature. Let R be a class of finite structures of $\mathcal{U}$ closed under isomorphism. Suppose R is not the finite models of any first-order sentence, but the isomorphism types represented by R is a definable unary relation in $\langle \mathcal{PU}, \leq \mathbf{c}_1, \ldots, \mathbf{c}_k \rangle$ for some finite types $\mathbf{c}_1, \ldots, \mathbf{c}_k$. Must the expressive power of first-order definability in the substructure relation be equivalent to full second-order properties in the finite? Does it at least capture all first-order properties?

We attempt an application of Corollary III.9.4 which suggest possible answers to the previous question may be found among universal subclasses of simple graphs. We borrow the terminology of finite model theory from [EF95]. For two logics $\mathcal{L}_1$ and $\mathcal{L}_2$, we write $\mathcal{L}_1 \leq \mathcal{L}_2$ if for any signature τ and every sentence $\phi \in \mathcal{L}_1[\tau]$, there exists $\psi \in \mathcal{L}_2[\tau]$ such that finite models of ϕ are precisely the finite models of ψ. We write $\mathcal{L}_1 \equiv \mathcal{L}_2$ iff $\mathcal{L}_1 \leq \mathcal{L}_2$ and $\mathcal{L}_2 \leq \mathcal{L}_1$.

Recall, the following theorem concerning interpretation of structures.

Theorem V.0.8. ([Hod93, Thm. 5.5.1]) Let σ be a first-order language whose signature consists of the binary relation symbol R, and let L be a first-order language with finite signature. Then there is a first-order sentence χ in σ such that

(1) Every model of χ is an irreflexive symmetric graph.

(2) The class of models of χ is bi-interpretable with the class of all L-structures which have more than one element.

Moreover, both interpretations in (2) preserve embeddings.

With any FO-interpretation, there is the corresponding reduction theorem relating satisfiability by structures in the two signatures.

Proposition V.0.9. Let Π be an interpretation of σ in τ. For every FO(σ) sentence ψ there exists a FO(τ) ψ^Π such that for all τ-structures $\mathcal{A}$(with non-empty universe),

$$\mathcal{A} \models \psi^\Pi \text{ iff } \mathcal{A}^\Pi \models \psi.$$

As noted in ([EF95, Ex 11.2.4]), the logic FO can be replaced with some other logics $\mathcal{L}$ such that FO $\leq \mathcal{L} \leq$ SO like FO(IFP), FO(PFP), or SO.

Let τ be finite relational signature and let K_τ be the class of models in that signature with at least two elements. Then by Theorem V.0.8, Corollary III.9.4, and Prop. V.0.9 we can conclude that

- $\mathcal{PK}_\tau$ is definably present in $\mathcal{PG}$.
- For any of the logics $\mathcal{L}$ which satisfy the conclusion of Proposition V.0.9, the finite models of those finitely $\mathcal{L}$-axiomatizable subclasses of $\mathcal{K}_\tau$ are first-order definably present in $\mathcal{PG}$.
- If $\mathcal{L}_1 < \mathcal{L}_2 \leq$ SO satisfy the conclusion of Proposition V.0.9, then there exists a finite signature τ and sentence $\phi \in \mathcal{L}_2[\tau]$ such that the finite models of ϕ are not the finite models of any $\mathcal{L}_1[\tau]$-sentence, but the isomorphism types form a first-order definable set in $\mathcal{PG}$.

CHAPTER VI

A DISJUNCTIVE CHARACTERIZATION FOR QUASIVARIETIES

VI.1 Introduction

In [J6́8], Bjarni Jónsson established that all the algebras in a variety $\mathcal{V}$ have distributive congruence lattices iff the variety has ternary terms $p_0, \ldots, p_n$ which satisfy the identities

$$\begin{array}{rcll} p_0(xyz) & \approx & x & \\ p_n(xyz) & \approx & z & \\ p_i(xyx) & \approx & x & 0 \leq i \leq n \\ p_i(xxy) & \approx & p_{i+1}(xxy) & i \quad \text{even} \\ p_i(xyy) & \approx & p_{i+1}(xyy) & i \quad \text{odd} \end{array}$$

Kirby Baker noticed Jónsson's condition was equivalent to a closely related disjunction; namely, a variety $\mathcal{V}$ is congruence distributive iff there exist ternary terms $p_1, \ldots, p_n$ such that

$$\begin{array}{rcl} \mathcal{V} & \models & p_i(xux) \approx p_i(xvx) \qquad 0 \leq i \leq n \\ \mathcal{V} & \models & x \not\approx y \rightarrow \bigvee_{i=1}^{n-1} [p_i(xxy) \not\approx p_{i+1}(xyy)] \end{array}$$

Using the above characterization in an intricate analysis of principal congruence generation in congruence distributive varieties (a streamlined version of which can be found in [BS81]), Baker secured the following finite basis result:

Theorem VI.1.1. [Bak77] Let $\mathcal{V}$ be a variety of finite signature. If $\mathcal{V}$ is congruence distributive and has a finite residual bound, then $\mathcal{V}$ is finitely based.

In [Wil00], Ross Willard provided a new characterization for congruence meet-semidistributive varieties; a variety $\mathcal{V}$ is congruence meet-semidistributive iff there exist ternary terms $f_0, \ldots, f_n, g_1, \ldots, g_n$ such that

$$\begin{array}{rcl} \mathcal{V} & \models & f_i(xyx) \approx g_i(xyx) \qquad 0 \leq i \leq n \\ \mathcal{V} & \models & x \not\approx y \rightarrow \bigvee_{i=0}^{n} [f_i(xxy) \approx g_i(xxy) \leftrightarrow f_i(xyy) \not\approx g_i(xyy)] \end{array}$$

Ross Willard was able to use this disjunction characterization to provide an ingenious and involved analysis of principal congruences in congruence meet-semidistributive varieties which yields his generalization of Baker's finite basis theorem.

Theorem VI.1.2. [Wil00] Let $\mathcal{V}$ be a variety with finite signature. If $\mathcal{V}$ is congruence meet-semidistributive and has a finite residual bound, then $\mathcal{V}$ is finitely based.

The paper of Maroti and McKenzie[MM04] explores finite basis results in quasivarieties, and manages to provide a common framework which generalizes the finite basis results of Ross Willard for congruence meet-semidistributive varieties and Pigozzi's[Pig88] finite basis result for relatively congruence distributive quasivarieties. The start of their approach is the observation that Willard's disjunction can characterize quasivarieties satisfying a weaker condition than meet-semidistributivity, but which is equivalent to it in varieties.

In Section VI.2 we prove a similar disjunction for quasivarieties which is an amalgam of Willard's characterization of congruence meet-semidistributivity and Malcev's characterization of congruence permutability. For varieties, our characterization is equivalent to the existence of a weak difference term. Using the very nice Lemma VI.2.6 from [KS98, Lem.4.4], we will provide a relatively short proof of this fact.

In Section VI.3, we consider two applications of the characterization for Taylor varieties which will allow us to simplify the proof of an important result(Theorem VI.3.3), and provide an elementary generalization of another(Theorem VI.3.2). We will reference but not develop the algebraic framework recently developed to study the constraint satisfaction problem(CSP). The reader is directed to the papers of Bulatov, Jeavons, and Krokhin[BJK05] and Jeavons[Jea98]. For the required background in universal algebra consult Hoby and McKenzie[HM88].

VI.2 The Characterization

For a quasivariety $\mathcal{K}$, let $\alpha = \Theta(x,z)$, $\beta = \Theta(x,y)$, and $\gamma = \Theta(y,z)$ be the principle congruences determined in $F_{\mathcal{K}}(x,y,z)$ and make the definition

$$W_{\mathcal{K}}(x,y) := \bigvee_{(f,g)\in\alpha} [f(xxy) \approx g(xxy) \leftrightarrow f(xyy) \not\approx g(xyy)]\,.$$

For any algebra $A \in \mathcal{K}$, we have

$$A \models \forall x \forall y\, (\neg W_{\mathcal{K}}(x,y) \leftrightarrow \neg W_{\mathcal{K}}(y,x))\,.$$

To see this, for any term $f(xyz)$ define $f^*(xyz) = f(zyx)$. Then $f^{**} = f$ and $(f,g) \in \alpha$ implies $(f^*,g^*) \in \alpha$. Take $a,b \in A$ and assume $A \models \neg W_{\mathcal{K}}(a,b)$; that is, $f(aab) = g(aab) \leftrightarrow f(abb) = g(abb)$ holds in A for all terms f,g such that $f(xyx) = g(xyx)$. Then $f(baa) = g(baa)$ iff $f^*(aab) = g^*(aab)$ iff $f^*(abb) = g^*(abb)$ iff $f(bba) = g(bba)$; thus, $A \models \neg W_{\mathcal{K}}(b,a)$. A similar argument establishes the converse.

For a ternary term $c(xyz)$ in the signature of $\mathcal{K}$, define the formula $M_c(x,y)$ by

$$M_c(x,y) := [y \approx c(xxy) \wedge c(xxy) \approx c(yxx) \wedge c(yyx) \approx c(xyy) \wedge c(xyy) \approx x]\,.$$

When the context is clear, the subscript denoting the class will often be dropped, but in its place will be a positive integer to denote the disjunction is over a finite set of terms. For example,

$$W_n(x,y) := \bigvee_{i=1}^{n} [f_i(xxy) \approx g_i(xxy) \leftrightarrow f_i(xyy) \not\approx g_i(xyy)]$$

where each $f_i(xyx) \approx g_i(xyx)$.

For any $A \in \mathcal{K}$, let the set of $\mathcal{K}$-congruences be $\mathrm{Con}_{\mathcal{K}}(A) = \{\alpha \in \mathrm{Con}(A) : A/\alpha \in \mathcal{K}\}$. The set of $\mathcal{K}$-congruences is a complete lattice where the meet is the same as in $\mathrm{Con}(A)$, and the join denoted by $\vee^{\mathcal{K}}$ is the corresponding operation induced by the meet. Let $\alpha, \beta, \gamma \in \mathrm{Con}_{\mathcal{K}}(A)$, and define congruences β_m, $\gamma_m \in \mathrm{Con}_{\mathcal{K}}(A)$ inductively by $\beta_0 = \beta$, $\gamma_0 = \gamma$ and

$$\beta_{n+1} = \beta \vee^{\mathcal{K}} (\alpha \wedge \gamma_n) \quad \text{and} \quad \gamma_{n+1} = \gamma \vee^{\mathcal{K}} (\alpha \wedge \beta_n).$$

Notice $\beta \leq \beta_1 \leq \beta_2 \leq \cdots$ and $\gamma \leq \gamma_1 \leq \gamma_2 \leq \cdots$. Set

$$\beta_\infty = \bigcup_{n\in\omega} \beta_n \qquad \text{and} \qquad \gamma_\infty = \bigcup_{n\in\omega} \gamma_n$$

and note $\beta_\infty, \gamma_\infty \in \mathrm{Con}_{\mathcal{K}}(A)$.

We are now ready for the theorem.

Theorem VI.2.1. For any quasivariety $\mathcal{K}$ the following are equivalent.

(1) For any $A \in \mathcal{K}$ and $\alpha, \beta, \gamma \in \mathrm{Con}_{\mathcal{K}}(A)$, $\alpha \wedge \beta = \alpha \wedge \gamma = 0_A$ implies $\alpha \wedge (\beta \circ \gamma) \subseteq \gamma \circ \beta$.

(2) For any $A \in \mathcal{K}$ and $\alpha, \beta, \gamma \in \mathrm{Con}_{\mathcal{K}}(A)$, $\alpha \wedge (\beta \circ \gamma) \subseteq \gamma_\infty \circ \beta_\infty$.

(3) For the principle congruences $\alpha = \Theta(x,z)$, $\beta = \Theta(x,y)$, and $\gamma = \Theta(y,z)$ in $F_{\mathcal{K}}(x,y,z)$ there exists m such that $\alpha \cap (\beta \circ \gamma) \subseteq \gamma_m \circ \beta_m$.

(4) There exists a finite set of ternary terms $f_1, \ldots., f_n, g_1, \ldots., g_n, c$ such that $f_i(xyx) \approx g_i(xyx)$ for $i = 1, \ldots, n$, and $\mathcal{K}$ satisfies the sentence

$$\forall x \forall y\, [x \not\approx y \longrightarrow W_n(x,y) \vee M_c(x,y)]\,.$$

`Proof`: We show **(4)** $\Rightarrow$ **(1)**. Let $(a,b) \in \alpha \wedge (\beta \circ \gamma)$. Then $(a,d) \in \beta$ and $(d,b) \in \gamma$ for some $d \in A$. Suppose $a \neq b$ and $A \models W_n(a,b)$. Then take $1 \leq i \leq n$ such that $f_i(aab) = g_i(aab) \leftrightarrow f_i(abb) \neq g_i(abb)$. Without loss of generality, assume $f_i(aab) = g_i(aab)$. Then we have $f_i(adb) \neq g_i(adb)$ or $f_i(adb) = g_i(adb)$. Suppose we have $f_i(adb) = g_i(adb)$. Then

$$f_i(abb)\gamma f_i(adb) = g_i(adb)\gamma g_i(abb)$$

and

$$f_i(abb)\alpha f_i(aba) = g_i(aba)\alpha g_i(abb)$$

which shows $f_i(abb)\alpha \wedge \gamma g_i(abb)$. Since $\alpha \wedge \gamma = 0_A$, we arrive at the contradiction $f_i(abb) = g_i(abb)$. A similar argument for the case $f_i(adb) \neq g_i(adb)$ will show $f_i(adb)\alpha \wedge \beta g_i(adb)$, and so produce the contradiction $f_i(adb) = g_i(adb)$.

It must be the case that $A \models M_c(a,b)$. We have $a = c(abb)$ and $c(aab) = b$, and it is the case that $a = c(abb)\gamma c(adc)\beta c(aab) = b$. This shows $(a,b) \in \gamma \circ \beta$.

We show **(1)** $\Rightarrow$ **(2)**. Notice $\alpha \wedge \beta_\infty = \alpha \wedge \gamma_\infty$. Let $\delta = \alpha \wedge \beta_\infty$ and note $\delta \in \mathrm{Con}_{\mathcal{K}}(A)$. We also have $\delta \leq \alpha, \beta_\infty, \gamma_\infty$ and $\alpha/\delta, \beta_\infty/\delta, \gamma_\infty/\delta \in \mathrm{Con}_{\mathcal{K}}(A/\delta)$. Then $\alpha/\delta \wedge \beta_\infty/\delta = \alpha/\delta \wedge \gamma_\infty/\delta = 0_{A/\delta}$ which implies by (1),

$$\alpha/\delta \wedge (\beta_\infty/\delta \circ \gamma_\infty/\delta) \subseteq \gamma_\infty/\delta \circ \beta_\infty/\delta.$$

Let $(a,b) \in \alpha \wedge (\beta \circ \gamma) \subseteq \alpha \wedge (\beta_\infty \circ \gamma_\infty)$. If $(a,b) \in \delta$ the result is immediate. Suppose $(a,b) \notin \delta$, then

$$(a/\delta, b/\delta) \in \alpha/\delta \wedge (\beta_\infty/\delta \circ \gamma_\infty/\delta) \subseteq \gamma_\infty/\delta \circ \beta_\infty/\delta$$

and so there exist $m < \omega$ and $c \in A$ such that

$$(a/\delta, c/\delta) \in \gamma_m/\delta$$

$$(c/\delta, b/\delta) \in \beta_m/\delta.$$

Since δ refines γ_∞ and β_∞, we can conclude that

$$(a,c) \in \gamma_\infty$$

$$(c,b) \in \beta_\infty$$

which yields (2).

To show **(2)** $\Rightarrow$ **(3)**, notice that $(x,z) \in \alpha \cap (\beta \circ \gamma)$, and so by (2), there exists m such that $(x,z) \in \gamma_m \circ \beta_m$. If $(a,b) \in \alpha \cap (\beta \circ \gamma)$, then there exist n and $c \in F_{\mathcal{K}}(x,y,z)$ such that $(a,c) \in \gamma_n$ and $(c,b) \in \beta_n$. Take the endomorphism $\sigma : F_{\mathcal{K}}(x,y,z) \to F_{\mathcal{K}}(x,y,z)$ determined by $\sigma : (x,y,z) \to (a,c,b)$, and observe that $\alpha \leq \sigma^{-1}(\alpha)$, $\gamma \leq \sigma^{-1}(\gamma)$, and $\beta \leq \sigma^{-1}(\beta)$. It is not difficult to see that $\sigma(\gamma_m \circ \beta_m) \subseteq \gamma_m \circ \beta_m$ from which we conclude that $(a,b) = (\sigma(x), \sigma(z)) \in \gamma_m \circ \beta_m$.

We establish **(3)** $\Rightarrow$ **(4)**. Assume (3) holds. There exists m such that $(x,z) \in \alpha \cap (\beta \circ \gamma) \subseteq \gamma_m \circ \beta_m$. So there must be a ternary term $c(xyz) \in F_{\mathcal{K}}(x,y,z)$ such that $(x, c(xyz)) \in \gamma_m$ and $(c(xyz), z) \in \beta_m$.

We show $\mathcal{K}$ satisfies the infinite sentence

$$\forall x \forall y \left[x \not\approx y \longrightarrow W_{\mathcal{K}}(x,y) \vee M_c(x,y) \right].$$

A compactness argument applied to $\neg W_{\mathcal{K}}(x,y) \wedge \neg M_c(x,y) \longrightarrow x \approx y$ will then replace the formally infinite disjunction with a disjunction over a finite set of terms.

Suppose there exist $A \in \mathcal{K}$ and $a,b \in A$ such that $a \neq b$ and $A \models \neg W_{\mathcal{K}}(a,b)$. By the above remarks, we also have $A \models \neg W_{\mathcal{K}}(b,a)$. Altogether,

$$f(aab) = g(aab) \leftrightarrow f(abb) = g(abb)$$

and

$$f(bba) = g(bba) \leftrightarrow f(baa) = g(baa)$$

holds in A for all terms f,g such that $f(xyx) \approx g(xyx)$.

Consider the homomorphisms of $F_{\mathcal{K}}(x,y,z)$ into A given by $\pi_1 : (x,y,z) \to (a,a,b)$, $\pi_2 : (x,y,z) \to (a,b,b)$, $\sigma_1 : (x,y,z) \to (b,b,a)$, and $\sigma_2 : (x,y,z) \to (b,a,a)$. Then

$$\alpha \wedge \ker\pi_1 = \alpha \wedge \ker\pi_2 \quad \text{and} \quad \alpha \wedge \ker\sigma_1 = \alpha \wedge \ker\sigma_2$$

Also, $\beta \leq$ ker $\pi_1\wedge$ ker σ_1 and $\gamma \leq$ ker $\pi_2\wedge$ ker σ_2.

Inductively, we have $\beta_m \leq$ ker $\pi_1\wedge$ ker σ_1 and $\gamma_m \leq$ ker $\pi_2\wedge$ ker σ_2 for all m.

But for the term $c(xyz)$ we have

$$(x,c) \in \gamma_m \leq \ker\pi_2 \wedge \ker\sigma_2$$

$$(c,z) \in \beta_m \leq \ker\pi_1 \wedge \ker\sigma_1.$$

This implies $c(aab) = c(baa) = b \neq a = c(bba) = c(abb)$ and therefore,

$$A \models M_c(a,b).$$

•

The term $c(xyz)$ in $M_c(x,y)$ will be idempotent throughout $\mathcal{K}$. To this, take $B \in \mathcal{K}$ and $a \in B$. Consider the induced map $\sigma : F_{\mathcal{K}}(x,y,z) \to B$ defined by $\sigma(x) = \sigma(y) = \sigma(z) = a$. Then $\alpha, \beta, \gamma \leq \ker\sigma$ and so $\beta_m, \gamma_m \leq \ker\sigma$. Then $(x,c) \in \beta_m$ implies $c(aaa) = a$.

Remark VI.2.2. From the arguments (4) $\Rightarrow$ (1) and (1) $\Rightarrow$ (2) in Theorem VI.2.1 we have the following useful facts for any variety $\mathcal{V}$ satisfying condition (4). Let $A \in \mathcal{V}$, $\alpha, \beta, \gamma \in \mathrm{Con}(A)$, and $a, b \in A$ such that $a \neq b$:

• If $(a,b) \in \alpha \cap (\beta \vee \gamma)$ and $A \models W_{\mathcal{V}}(a,b)$, then

$$\alpha \wedge \beta \neq 0_A \quad \text{or} \quad \alpha \wedge \gamma \neq 0_A.$$

• If $(a,b) \in \alpha \cap (\beta \vee \gamma)$ and $\alpha \wedge \beta = \alpha \wedge \gamma = 0_A$, then $A \models M_c(a,b) \wedge \neg W_{\mathcal{V}}(a,b)$.

• If $(a,b) \in \alpha \cap (\beta \vee \gamma) \setminus \delta$ where $\delta = \alpha \wedge \beta_\infty = \alpha \wedge \gamma_\infty$, then

$$a\,\delta\,c(abb)\;\delta\,c(bba) \quad \text{and} \quad b\,\delta\,c(baa)\;\delta\,c(aab).$$

We say (a,b) is a Malcev pair if $A \models M_c(a,b)$, and a Willard pair if $A \models W_n(a,b)$.

By referring to Theorem 9.6 in [HM88], we have the following corollary.

Corollary VI.2.3. Let $\mathcal{V}$ be a locally finite variety. The following are equivalent.

(1) $\mathcal{V}$ omits type 1

(2) $\mathcal{V}$ has a Taylor term

(3) There exists an idempotent special variety $\mathcal{E} \leq \mathcal{V}$ such that $\mathcal{E} \not\leq Sets$.

(4) For the principle congruences $\alpha = \Theta(x,z)$, $\beta = \Theta(x,y)$, and $\gamma = \Theta(y,z)$ in $F_{\mathcal{V}}(x,y,z)$ there exists m such that $\alpha \cap (\beta \circ \gamma) \subseteq \gamma_m \circ \beta_m$.

(5) There exists a finite set of idempotent ternary terms $f_1, ..., f_n, g_1, ..., g_n, c$ such that $f_i(xyx) \approx g_i(xyx)$ for $i = 1, ..., n$, and $\mathcal{V}$ satisfies the sentence

$$\forall x \forall y \, (x \not\approx y \longrightarrow [W_n(x,y) \vee M_c(x,y)]) \, .$$

Remark VI.2.4. It is easy to see $(5) \Rightarrow (3)$ since the conditions in $W_n(x,y)$ and $M_c(x,y)$ cannot be satisfied by any interpretation by ternary projections. We shall see later that $(5) \Rightarrow (2)$ in the discussion of Sigger's([Sig10]) strong malcev condition, so the disjunctive characterization can imply in short order a 6-variable Taylor term for locally finite varieties. The equivalence of (4) and (2) passes through (1) and tame congruence theory. Is it possible to give a direct proof of $(2) \Rightarrow (4)$?

Definition VI.2.5. We say a quasivariety $\mathcal{Q}$ is *strictly active* if it satisfies condition (3)(equivalently, condition (4)) in Theorem VI.2.1 above; namely, for the principle congruences $\alpha = \Theta(x,z)$, $\beta = \Theta(x,y)$, and $\gamma = \Theta(y,z)$ in $F_{\mathcal{Q}}(x,y,z)$, there exists $m \in \omega$ such that $\alpha \cap (\beta \circ \gamma) \subseteq \gamma_m \circ \beta_m$.

In Theorem VI.2.7 we shall see that the term $c(xyz)$ in $M_c(x,y)$ will be a weak difference term for any strictly active variety. This will allow us to give an alternate proof of the result of Kearnes and Szendrei characterizing the malcev condition in (4) of Corollary VI.2.3 as the weakest malcev condition for varieties which guarantees that abelian algebras are affine. Our proof avoids the necessity of first developing the topic of quasi-affine varieties. In order to do so, we shall need a lemma from [KS98]. For $\alpha, \beta \in \mathrm{Con}(A)$, let $A(\alpha)$ denote the congruence α thought of as a subalgebra of A^2. Define the congruence in $A(\alpha)$,

$$\Delta_{\beta\alpha} = Cg^{A(\alpha)}(\{\langle (u,u),(v,v)\rangle : (u,v) \in \beta\}).$$

Lemma VI.2.6. (Lemma 4.4 [KS98]) Suppose $\mathcal{V}$ has a Taylor term. For $\gamma, \delta \in \mathrm{Con}(A)$, let $\gamma_i = \pi_i^{-1}(\gamma)$ for $i = 1,2$ where π_i denotes the coordinate projections of $A(\delta)$ onto A and $\eta_i = \ker \pi_i$. If $C(\gamma, \delta; 0)$ holds, then

$$\gamma_0 \wedge \eta_1 \wedge \Delta_{\gamma\delta} = 0 = \gamma_1 \wedge \eta_0 \wedge \Delta_{\gamma\delta}.$$

Theorem VI.2.7. (Theorem 4.8 [KS98]) For a variety $\mathcal{V}$, the following are equivalent:

(1) $\mathcal{V}$ has a weak difference term.

(2) $\mathcal{V}$ is strictly active.

(3) $\mathcal{V}$ has an idempotent term which interprets as a malcev operation in abelian algebras; consequently, abelian algebras are affine.

`Proof`: Assume (1) holds and let $c(xyz)$ be the weak difference term. Let $A \in \mathcal{V}$ and suppose $\alpha, \beta, \gamma \in \mathrm{Con}(A)$ such that $\alpha \wedge \beta = \alpha \wedge \gamma = 0$. Let $(a,b) \in \alpha \wedge (\beta \circ \gamma) = \theta$. Then there exists $e \in A$ such that $a\beta e\gamma b$. Then

$$a[\theta,\theta]c(bba)\gamma c(bea)\beta c(baa)[\theta,\theta]b.$$

Since $[\beta,\alpha] = [\gamma,\alpha] = 0$ we have $[\beta \vee \gamma, \alpha] = 0$. This implies $[\theta,\theta] \leq [\beta \vee \gamma, \alpha] = 0$ and so $(a,b) \in \gamma \circ \beta$. We have established

$$\alpha \wedge (\beta \circ \gamma) \subseteq \gamma \circ \beta.$$

This is precisely condition (1) in Theorem VI.2.1, and so we see that $\mathcal{V}$ is strictly active.

Now, assume $\mathcal{V}$ is strictly active. It is easy to see that neither $W_n(x,y)$ nor $M_c(x,y)$ can be satisfied by any interpretation by ternary projections, and so by a result of Walter Taylor[HM88, Lemma 9.4], $\mathcal{V}$ has a Taylor term and so we may make use of Lemma VI.2.6. Let $A \in \mathcal{V}$ be an abelian algebra. Then $[1_A, 1_A] = 0_A$ which implies $C(1_A, 1_A; 0_A)$ in A. In the notation of Lemma VI.2.6, $A(1_A) = A^2$ and $\pi_i^{-1}(1_A) = 1_{A^2}$ and therefore,

$$\Delta_{1_A,1_A} \wedge \eta_1 = 0 = \Delta_{1_A,1_A} \wedge \eta_0.$$

For any $a, b \in A$ we have $\langle (a,a),(b,b) \rangle \in \Delta_{1_A,1_A} \wedge (\eta_0 \circ \eta_1)$, and so by the remark proceeding Theorem VI.2.1 we see that $\langle (a,a),(b,b) \rangle$ is a Malcev pair in A^2. But this just means (a,b) is a Malcev pair in A. We have shown the term $c(xyz)$ in $M_c(x,y)$ is a malcev term for every abelian algebra in $\mathcal{V}$, and so each abelian algebra is affine.

Now, suppose $c(xyz)$ is an idempotent term which is a malcev operation on abelian algebras of $\mathcal{V}$. If we pass to the variety generated by the idempotent reducts of algebras in $\mathcal{V}$, then since $c(xyz)$ is idempotent, it will interpret as a malcev operation for the idempotent reducts which are abelian.

Let θ be a congruence of $A \in \mathcal{V}$ and $(a,b) \in \theta$. If $\theta = [\theta,\theta]$, then $c(bba)[\theta,\theta]a[\theta,\theta]c(abb)$. In case $[\theta,\theta] < \theta$, we factor by $[\theta,\theta]$ and observe that θ is abelian over $[\theta,\theta]$. Then each $\theta/[\theta,\theta]$-class is an abelian subalgebra of the idempotent reduct of $A/[\theta,\theta]$, and so $c(abb)/[\theta,\theta] = a/[\theta,\theta] = c(bba)/[\theta,\theta]$. It follows that $c(xyz)$ is a weak difference for $\mathcal{V}$. •

Remark VI.2.8. The implication (2) $\Rightarrow$ (3) required the use of a Taylor term. For general varieties, strictly active is a stronger condition than having a Taylor term [KS98, Ex.4.13]. Is it possible to use the disjunctive characterization to prove Lemma VI.2.6 directly?

VI.3 Applications

Here is an immediate applicaton of the disjunction characterization in Theorem VI.2.3. For a finite reflexive tournament T, CSP(T^c) will denote the constraint satisfaction problem over the structure which has all the singleton unary relations in addition to the edge relation of T. In [Lar06], Benoit Larose proves the following theorem.

Theorem VI.3.1. Let T be a finite reflexive tournament. Then T admits a Taylor operation if and only if T is transitive. If T is transitive, then the problem CSP(T^c) is in **P**, and it is **NP**-complete otherwise.

The second statement follows immediately from the first, since transitive tournaments are precisely linear orders, and so admit the lattice operations of max and min([JCC98]). The last statement follows since the relational structure T^c is a core, and it is known that if a core relational structure does not admit a Taylor operation, then CSP(T^c) is **NP**-complete [BKJ00].

The strategy of Larose's proof for the first statement is to consider a counter-example of minimal cardinality, and first show that it must have strictly more than three elements. The second step is to argue with the local combinatorics and produce a smaller counterexample, obtaining a contradiction. The first step is achieved by an application of a highly non-trivial and involved construction of a homotopy theory for finite reflexive binary structures developed in [LT04] and [Lar06]. Our first application of Theorem VI.2.3 is to provide an entirely elementary proof of this result; actually, a stronger result can be achieved.

Let $R \leq T^n$ be in the relational clone determined by T; equivalently, there is a primitive positive formula $\psi(x_1,...,x_n)$ in the edge relation of T such that $R = R_{\psi(x_1,...,x_n)} = \{(a_1,...,a_n) : T \models \psi(a_1,...,a_n)\}$. For any vertices $v_2,.,.,v_{n-1} \in T$, the unary relation $S_{\psi(x_1,v_1,...,v_n)} = \{a \in T : T \models \psi(a,v_1,...,v_n)\}$ may not be in the relational clone generated by T, but it will be closed under the idempotent polymorphisms of T. We call such unary relations $S_{\psi(x_1,v_1,...,v_n)}$ *inferred idempotent subalgebras*.

For a finite relational structure $\mathbb{X}$, $\mathbf{IdAlg}(\mathbb{X})$ denotes the non-indexed algebra with the same universe as $\mathbb{X}$ and whose operations are all the idempotent polymorphisms.

Theorem VI.3.2. Let T be a finite tournament(not neccessarily reflexive). If T contains a 3-cycle with at least two loops, then T is not closed under a Taylor polymorphism; consequently, CSP(T^c) is NP-complete.

`Proof`: For contradiction, assume there exists a finite tournament which contains a 3-cycle with at least two loops and is closed under a Taylor operation. Let T be such a tournament of minimal cardinality. Let $a \to b \to c \to a$ be a 3-cycle in T and without loss of generality, we may assume the vertices a and b have loops. By Corollary VI.2.3, $\mathcal{V}(\mathbf{IdAlg}(T))$ satisfies the disjunctive condition in (5). Suppose (a,b) is a Malcev pair. Then $b = c(aab) \to c(abb) = a$ which is a contradiction. It must be the case that (a,b) is a Willard pair, and so take $f(xyz), g(xyz)$ such that $f(xyx) \approx g(xyx)$ and

$$f(aab) = g(aab) \leftrightarrow f(abb) \neq g(abb).$$

If it were the case that $f(aab) \neq g(aab)$ and $f(abb) = g(abb)$, set $r(xyz) = f(zyx)$, $s(xyz) = g(zyx)$ and notice $r(xyx) \approx s(xyx)$. If $T^{\mathbf{rev}}$ denotes the tournament formed by reversing the orientation of the edges of T, then $T^{\mathbf{rev}}$ has the same polymorphisms as T. We then have a 3-cycle $b \to a \to c \to b$ in $T^{\mathbf{rev}}$ with loops at b and a where $r(bba) = s(bba)$ and $r(baa) \neq s(baa)$.

So we may assume $f(aab) = g(aab)$ and $f(abb) \neq g(abb)$. We claim that there exists a vertex w such that $a \to w \to b$. Let suppose this is not the case. We shall argue for a contradiction. Since f and g are polymorphisms of T, there is an obvious homomorphism of the digraph $\mathbb{G}$ (see Fig.VI.1a) into T which "fixes" the vertices $\{a,b,c\}$. It must be that $\{f(aab), g(aab), f(bab)\} \subseteq \{a,b\}$. There are two cases to consider.

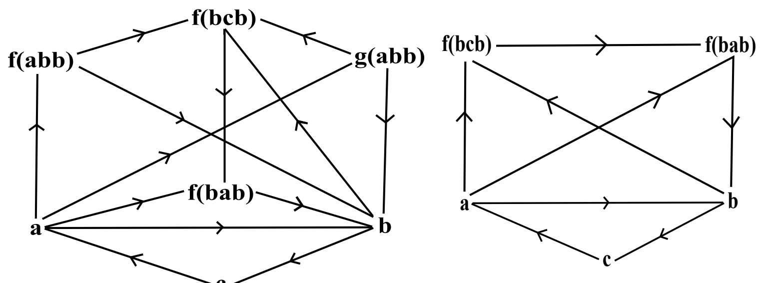

(a) $\mathbb{G}$ (b) $\mathbb{H}$

Figure VI.1: Digraphs $\mathbb{G}$ and $\mathbb{H}$

Suppose $f(abb) = a$. Then $g(abb) = b$ and $a \to f(aab) \to f(abb) = a$ implies $f(aab) = g(aab) = a$. This reduces to a homomorphic mapping of $\mathbb{H}$ (see Fig.VI.1b) into T. If $f(bab) = a$, then by $\mathbb{H}$ we must have $f(bcb) = a$, but $b \to f(bcb) = a$ yields a contradiction. If $f(bab) = b$, then $f(bcb) = b$. Now, $b = f(bcb) = g(bcb) \to g(cab) \to g(abb) = b$ implies $g(cab) = b$. But then $b = g(cab) \to g(aab) = a$ yields a contradiction.

Suppose $f(abb) = b$. Then we must have $g(abb) = a$ and $a \to f(aab) = g(aab) \to g(abb) = a$ implies $f(aab) = a$. Again we have reduced to a consideration of $\mathbb{H}$. If $f(bab) = a$, then we must have $f(bcb) = a$, but then $b \to f(bcb) = a$ is a contradiction. If $f(bab) = b$, then $b \to f(bcb) \to f(bab) = b$ implies $f(bcb) = b$. We have $b = f(bcb) \to f(cab) \to f(abb) = b$ which implies $f(cab) = b$. But then $b = f(cab) \to f(aab) = a$ is a contradiction.

So, there must exist a vertex w such that $a \to w \to b$. We may assume $w \to c$ to produce the configuration

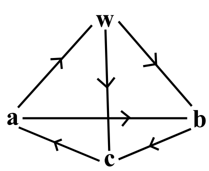

If $w \leftarrow c$, then we would consider $T^{\mathbf{rev}}$ and notice the induced subtournament on $\{a,b,c,w\}$ forms an isomorphic configuration.

Define the subalgebra $B = \{z : (\exists x)\,[(b \to x) \wedge (w \to x) \wedge (x \to z)]\}$. Then $\{a,b,c\} \subseteq B$ and so by minimality, $B = T$. But then $w \in B$ implies there exists x_0 such that $2 \to x_0$ and $w \to x_0 \to w$. Since $x_0 \neq 2, w$ and $w \neq b$, we arrive at the final contradiction which establishes the theorem. •

In [Sig10], Mark Siggers proved that omitting type 1 for locally finite varieties is equivalent to a strong malcev condition; namely, a locally finite variety omits type 1 if and only if it has a 6-variable Taylor term.

Sigger's startling short proof was based on the fact that a finite irreflexive symmetric graph with a triangle is not closed under a Taylor operation. This result is essentially established by Bulatov's reproof of the H-dichotomy conjecture when H is an irreflexive symmetric graph. Hell and Nešetřil[HN90] first established the dichotomy for H-coloring, and since then several authors in Bulatov[Bul05], Kun and Szegedy[KS09], Siggers[Sig09] and Barto and Kozik[BK12] have provided alternate proofs of this important result with varying levels of simplification. Our second application will be another proof that a finite irreflexive graph with an odd symmetric cycle is not closed under a Taylor operation. Since our proof is more algebraic, we begin in the same manner as Bulatov[Bul05] by taking a minimal counterexample and using primitive positive formulas to enforce special properties; however, the disjunctive characterization in Corollary VI.2.3 will allow us to short-cut the main argument entirely.

Theorem VI.3.3. Let $\mathbb{G}$ be a finite irreflexive digraph which contains an odd symmetric cycle. Then $\mathbb{G}$ is not closed under a Taylor polymorphism.

`Proof`: For a contradiction, suppose there exists a finite irreflexive digraph which contains a symmetric odd cycle and is closed under a Taylor operation. Let $\mathbb{G} = \langle V, E \rangle$ be such a digraph of minimal cardinality. By passing to the symmetric skeleton, we may assume $\mathbb{G}$ is symmetric. We may also assume $\mathbb{G}$ contains a triangle. If this is not so, then let $k \geq 5$ be the length of the smallest symmetric odd cycle in $\mathbb{G}$. The k-2-fold relational product E^{k-2} is a binary relation in the relational clone generated by the edge relation E, and so is closed under the polymorphisms of $\mathbb{G} = \langle V, E \rangle$. Since $k \geq 5$ is the length of the smallest odd cycle in $\mathbb{G}$, E^{k-2} is irreflexive. Let $x_1 \leftrightarrow x_2 \leftrightarrow \cdots \leftrightarrow x_k \leftrightarrow x_1$ be a cycle in $\mathbb{G}$. Then $(x_i, x_{i+2(mod k)}) \in E^{k-2}$ and $(x_i, x_{i+1(mod k)}) \in E^{k-2}$ because k is odd. Altogether, $\mathbb{H} = \left\langle V, E^{k-2} \right\rangle$ is an irreflexive symmetric digraph of the same cardinality as $\mathbb{G}$ with a triangle and closed under a Taylor operation. Let $1 \leftrightarrow 2 \leftrightarrow 3 \leftrightarrow 1$ be a symmetric triangle in $\mathbb{G}$.

Note that every vertex is part of a triangle. To see this, observe that aE^3a iff a is a vertex of a triangle. If we restrict to the subalgebra $S(x) = \left\{x \in \mathbb{G} : xE^3x\right\}$, then by minimality, we must have $S = \mathbb{G}$. Also, $\mathbb{G}$ cannot contain a complete graph on 4 vertices since the neighborhood of any of its vertices is a proper pp-definable subset whch contains a triangle.

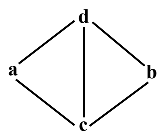

Figure VI.2: A rhombus

Claim: $\mathbb{G}$ cannot contain a proper rhombus(see Fig.VI.2).

`Proof`: : For a contradiction, $\mathbb{G}$ contains a rhombus where the two non-adjacent vertices are distinct. We will define a quotient graph from $\mathbb{G}$ which will have smaller cardinality, contain a triangle, and be closed under a Taylor operation. We start by defining a congruence. Consider the following primitive positive formula

$$R(x,y) := \exists u \exists v \left[E(x,u) \wedge E(x,v) \wedge E(u,v) \wedge E(u,y) \wedge E(v,y)\right].$$

Then $R(a,b)$ iff a and b can be connected as opposite vertices of a rhombus. Let θ be the transitive closure of R which is pp-definable since $\mathbb{G}$ is finite. Since every vertex belongs to some triangle, R is reflexive, and therefore, θ is a congruence. By assumption, θ is non-trivial.

We show θ does not contain an edge of $\mathbb{G}$. For a contradiction, suppose otherwise and choose $e\theta h$ where e and h are connected by a chain of n rhombii of minimal possible length. If $n=1$, then e and f are vertices of a complete graph on 4 vertices; therefore, $n>1$. There are two cases to consider.

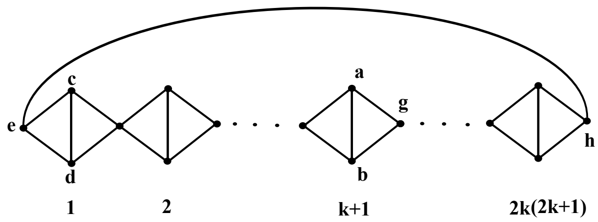

Figure VI.3: A chain of rhombii

If $n=2k$ is even, then e and h are connected by a chain of rhombii as in Fig.VI.3. Consider the subalgebra defined by the pp-formula

$$\begin{aligned} S(x) := \quad \exists \quad & z_1x_1,y_1z_2\cdots x_{k-1}y_{k-1}z_k[E(a,z_1)\wedge E(b,z_1)\wedge E(z_1,x_1)\wedge E(z_1,y_1) \\ \wedge \quad & E(x_1,y_1)\wedge E(x_1,z_2)\wedge E(y_1,z_2)\wedge\cdots\wedge E(x_{k-1},y_{k-1})\wedge E(x_{k-1},z_k) \\ \wedge \quad & E(y_{k-1},z_k)\wedge E(z_k,x)] \end{aligned}$$

This formula says that $u\in S$ if $E(u,v)$ for some v and v can be connected to the edge $E(a,b)$ by a chain of $k-1$ rhombii the same way h can be connected to $E(a,b)$ in Fig.VI.3. We see that S contains the triangle on $\{e,c,d\}$, and so by minimality, $S=\mathbb{G}$. But $h\in S$ implies there exists some vertex f connected to h by a chain of $n-1=2k-1$ rhombii. This contradicts the minimality of the chain.

If $n=2k+1$, then we can argue in the same manner using the subalgebra defined by the pp-formula

$$\begin{aligned} S(x) := \quad \exists \quad & x_1y_1z_1\cdots x_ky_kz_k[E(g,x_1)\wedge E(b,y_1)\wedge E(x_1,y_1)\wedge E(x_1,z_1)\wedge E(y_1,z_1) \\ \wedge \quad & E(z_1,x_2)\wedge E(z_1,y_2)\wedge\cdots\wedge E(x_k,y_k)\wedge E(x_k,z_k) \\ \wedge \quad & E(y_k,z_k)\wedge E(z_k,x)] \end{aligned}$$

Since both cases lead to a contradiction, θ cannot identify an edge of $\mathbb{G}$. This implies the vertices of any triangle are not identified by θ. The quotient graph $\mathbb{G}/\theta$ is defined with vertex set $\{a/\theta : a\in V\}$ and edge relation $\{(a/\theta,b/\theta) : E(a,b)\}$. If $t(x_1,\ldots,x_n)$ is a Taylor polymorphism of $\mathbb{G}$, then $\mathbb{G}/\theta$ admits a Taylor operation defined by $\hat{t}(a_1/\theta,\ldots,a_n/\theta)=t(a_1,\ldots,a_n)/\theta$. This is a polymorphism of $\mathbb{G}/\theta$ since θ is a congruence. Now, the quotient graph $\mathbb{G}/\theta$ is of smaller cardinality, has no loops, contains a triangle, and is closed under the corresponding Taylor operation - contradicting the minimality of $\mathbb{G}$. This establishes the claim.

By Corollary VI.2.3, $\mathcal{V}(\mathbf{IdAlg}(\mathbb{G}))$ satisfies condition (5). Suppose there are two Malcev pairs. Without loss of generality, we may assume $(1,2)$ and $(1,3)$ are Malcev. Then $1 = c(122) \leftrightarrow c(331) = 1$ which is a contradiction. We must have at least one Willard pair. Without loss of generality, we assume $(1,2)$ is a Willard pair and so there exist f, g such that $f(xyx) \approx g(xyx)$ with $f(122) = g(122)$ and $f(112) \neq g(112)$ (If it were the case that $f(112) \neq g(112)$ with $f(122) = g(122)$, set $r(xyz) = f(zyx)$, $s(xyz) = g(zyx)$ and notice r and s provide the required pattern).

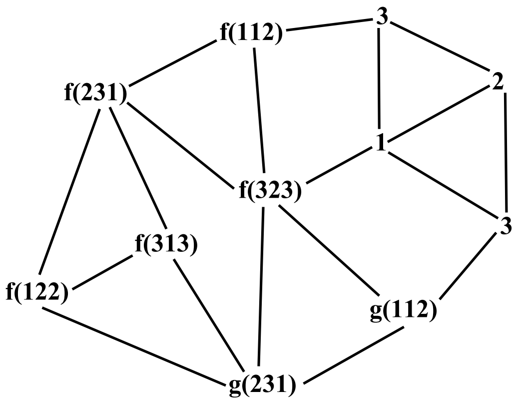

Figure VI.4: A leaf

Since f and g are polymorphisms, there is a homomorphism of the leaf graph in Fig.VI.4 which "fixes" the vertices $\{1,2,3\}$. By the previous claim, it must be the case that $g(231) = f(231)$ in $\mathbb{G}$. But this forms another rhombus which implies $g(112) = f(112)$, the final contradiction which establishes the theorem. •

VI.4 A Next Step

In studying polymorphisms of finite digraphs, it appears the new characterization for Taylor varieties may be useful. In the applications we considered, the arguments were more "local" and reduced to the analysis of a small and simple combinatorial configuration. Perhaps this approach can be generalized to more varied settings. The following would be an interesting first start.

- Characterize those finite tree digraphs which admit a Taylor operation
- Characterize all finite tournaments which admit a Taylor operation. In particular, if T is a finite smooth tournament with two directed cycles, is T closed under a Taylor operation?

BIBLIOGRAPHY

[Bak77] Kirby A. Baker. “Finite equational bases for finite algebras in a congruence-distributive equational class”. In: *Advances in Math.* 24.3 (1977), pp. 207–243.

[BK12] Libor Barto and Marcin Kozik. “Absorbing subalgebras, cyclic terms, and the constraint satisfaction problem”. In: *Log. Methods Comput. Sci.* 8.1 (2012), 27pp.

[Bul05] Andrei Bulatov. “H-coloring dichotomy revisited”. In: *Theoret. Comput. Sci.* 349.1 (2005), pp. 31–39.

[BJK05] Andrei Bulatov, Peter Jeavons, and Andrei Krokhin. “Classifying the complexity of constraints using finite algebras.” In: *SIAM J. Comput.* 34.3 (2005), pp. 720–742.

[BKJ00] Andreir Bulatov, Andrei Krokhin, and Peter. Jeavons. “Constraint satisfaction problems and finite algebras”. In: *Automata, languages and programming (Geneva, 2000)*. Springer, Berlin: Lecture Notes in Comput. Sci., 1853, 2000, pp. 272–282.

[BS81] Stanley Burris and H.P. Sankappanavar. *A Course in Universal Algebra*. 1st ed. Vol. 78. Graduate Texts in Mathematics. New York: Springer-Verlag, 1981, pp. xvi+276. ISBN: 0387905782.

[EF95] Heinz-Dieter Ebbinghaus and Jörg Flum. *Finite Model Theory*. 1st ed. Vol. 90. Perspectives in Mathematical Logic. Springer-Verlag, Berlin, 1995, pp. xvi+327. ISBN: 3-540-60149-X.

[Har64] Frank Harary. “On the reconstruction of a graph from a collection of subgraphs.” In: *Theory of Graphs and its Applications (Proc. Sympos. Smolenice, 1963)*. Prague: Publ. House Czechoslovak Acad. Sci., 1964, pp. 47–52.

[HN90] P. Hell and J. Nešetřil. “On the complexity of H-coloring”. In: *J. Combin. Theory B* 48.1 (1990), pp. 92–100.

[Hig52] Graham Higman. “Ordering by divisibility in abstract algebras”. In: *Proc. London Math. Soc. (3)* 2 (1952), pp. 326–336.

[HM88] David Hobby and Ralph McKenzie. *The Structure of Finite Algebras*. Vol. 76. Contemporary Mathematics. Providence, RI: American Mathematical Society, 1988, pp. xii+203. ISBN: 0-8218-5073-3.

[Hod93] Wilfrid Hodges. *Model Theory*. Vol. 42. Encyclopedia of Mathematics and its Applications. Cambridge University Press, Amsterdam, 1993, pp. xiv+772. ISBN: 0-521-30442-3.

[Jea98] Peter Jeavons. “On the algebraic structure of combinatorial problems”. In: *Theoretical Copmuter Science* 200.1–2 (1998), pp. 185–204.

[JCC98] P.G. Jeavons, D.A. Cohen, and M. Cooper. “Constraints consistency and closure”. In: *Artificial Intelligence* 101.1-2 (1998), pp. 251–265.

[JM09a] Jaroslav Ježek and Ralph McKenzie. “Definability in Substructure Orderings, I: Finite semilattices”. In: *Algebra Univers.* 61 (2009), pp. 59–75.

[JM09b] Jaroslav Ježek and Ralph McKenzie. “Definability in Substructure Orderings, III: Finite distributive lattices”. In: *Algebra Univers.* 61 (2009), pp. 283–300.

[JM09c] Jaroslav Ježek and Ralph McKenzie. “Definability in Substructure Orderings, IV: Finite lattices”. In: *Algebra Univers.* 61 (2009), pp. 301–312.

[JM10] Jaroslav Ježek and Ralph McKenzie. “Definability in Substructure Orderings, II: Finite ordered sets”. In: *Order* 27.2 (2010), pp. 115–145.

[J́68] Bjarni Jónsson. “Algebras whose congruence lattices are distributive”. In: *Math. Scand.* 21 (1968), pp. 110–121.

[KS98] Keith Kearnes and Agnes Szendrei. “The relationship between two commutators.” In: *Internat. J. Algebra Comput.* 8.4 (1998), pp. 497–531.

[Kel57] Paul J. Kelly. “A congruence theorem for trees”. In: *Pacific J. Math.* 7 (1957), pp. 961–968.

[KS09] G*á*bor Kun and Mario. Szegedy. “A new line of attack on the dichotomy conjecture [extended abstract]”. In: *Proceedings of the 2009 ACM International Symposium on Theory of Computing*. New York: ACM, 2009, pp. 725–734.

[Lar06] Benoit Larose. “Taylor operations on finite reflexive structures”. In: *Int. J. Math. Comput. Sci.* 1.1 (2006), pp. 1–21.

[LT04] Benoit Larose and Claude Tardif. “A discrete homotopy theory for binary reflexive structures”. In: *Adv. Math.* 189.2 (2004), pp. 268–300.

[Mal67] A.I. Mal’cev. “Universally axiomatizable subclasses of locally finite classes of models”. In: *Sibirsk. Mat.* . 8.3 (1967), pp. 1005–1014.

[MM04] Miklós Maróti and Ralph McKenzie. “Finite Basis Problems and Results for Quasivarieties”. In: *Studia Logica* 78.1–2 (2004), pp. 293–320.

[Mil85] E.C. Milner. “Basic WQO - and BQ-Theory”. In: *Graphs and order (Banff, Alta., 1984)*. Reidel, Dordrecht: NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci., 147, 1985, pp. 487–502.

[Pig88] Don Pigozzi. “Finite basis theorems for relatively congruence distributive quasivarieties”. In: *Trans. Amer. Math. Soc.* 310 (1988), pp. 499–533.

[PS04] Oliver Pretzel and Johannes Siemons. “Reconstruction of Partitions”. In: *Electronic J. Combin.* 11.2 (2004), Note 5(electronic).

[Ram05] Jean-Xavier. Rampon. “What is reconstruction for ordered sets”. In: *Discrete Math.* 291.1–3 (2005), pp. 191–233.

[San85] B. Sands. “Unsolved Problems”. In: *Order* 1.3 (1985), pp. 311–313.

[Sig09] Mark Siggers. “A new proof of the H-coloring dichotomy”. In: *SIAM J. Discrete Math.* 23.4 (2009), pp. 2204–2210.

[Sig10] Mark H. Siggers. “A Strong Mal’cev Condition for locally finite varieties omitting the unary type”. In: *Algebra Univers.* 64.1–2 (2010), pp. 15–20.

[Sto77] P.K. Stockmeyer. “The falsity of the reconstruction conjecture for tournaments”. In: *J. Graph Theory* 1.1 (1977), pp. 19–25.

[Sto81] P.K. Stockmeyer. “A census of reconstruction digraphs I: Six related families”. In: *J. Combin. Theory* 31.2 (1981), pp. 232–239.

[TMR53] Alfred Tarski, with Andrzej Mostowski, and Raphael Robinson. *Undecidable Theories*. Vol. 12. Studies in logic and the foundation of mathematics. Amsterdam: North Holland Pub. Co., 1953, pp. xi+ 98.

[Wil00] Ross Willard. “A finite basis theorem for residually finite, congruence meet-semidistributive varieties”. In: *J. Symbolic Logic* 65.1 (2000), pp. 187–200.