Descriptive Set Theory, Equivalence Relations, and Classification Problems in Analysis

by

John Daniel Clemens

A.B. (Princeton University) 1994
C.Phil (University of California at Berkeley) 1997

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Mathematics

in the

GRADUATE DIVISION
of the
UNIVERSITY of CALIFORNIA at BERKELEY

Committee in charge:

Professor John Steel, Chair
Professor Leo Harrington
Professor Chris Shannon

Spring 2001

The dissertation of John Daniel Clemens is approved:

Chair Date

Date

Date

University of California at Berkeley

Spring 2001

Descriptive Set Theory, Equivalence Relations, and Classification Problems in Analysis

Abstract

Descriptive Set Theory, Equivalence Relations, and Classification Problems in Analysis

by

John Daniel Clemens
Doctor of Philosophy in Mathematics

University of California at Berkeley

Professor John Steel, Chair

We apply descriptive set-theoretic techniques to analyze the complexity of various sets and relations arising in areas of analysis. In particular, we use the theory of Borel reducibility among definable equivalence relations to quantify the difficulty of classifying various mathematical objects up to some notion of equivalence.

We first analyze the group of Borel automorphisms of a Polish space, and show that its isomorphism relation is quite complicated: It is a $\mathbf{\Sigma}^1_2$-complete relation, and reduces the relation of equality on Borel sets.

Next we consider the notion of a weakly wandering sequence for a transformation and show that the set of sequences which are weakly wandering for some transformation is a $\mathbf{\Sigma}^1_1$-complete set, as are several related sets. We apply our techniques to produce specific sequences of interest, for example, a sequence which is exhaustive weakly wandering for some transformation but which is not weakly wandering for any ergodic transformation.

We then briefly consider equivalence relations which reduce all Borel equivalence relations, and ask whether there can be any minimal such relations. We show that the equality of Borel sets is not minimal, nor is it a universal $\mathbf{\Pi}^1_1$ equivalence relation.

Next we analyze Polish metric spaces. We first show that a set of non-negative reals is the set of distances of some Polish metric space if and only if it is either countable or it is analytic and has 0 as a limit point. We also characterize the distance sets for certain classes of metric spaces.

We then consider the equivalence relation of isometry of Polish metric spaces and present a technique for reducing the orbit equivalence relation of a Polish group action to this isometry relation. We also give lower bounds for the complexity of isometry restricted to certain classes of spaces.

Finally, we consider the isometry of spaces with large isometry groups. We prove several results about the isomorphism relation on various classes of countable structures which are of independent interest, in particular that the classification of countable vertex-transitive graphs up to isomorphism is Borel-complete.

Professor John Steel
Dissertation Committee Chair

Contents

Acknowledgements

I owe a great debt to my advisor, John Steel, without whom I might never have successfully navigated the road from unfocused problem-solver to mathematical researcher. In addition to providing advice on my research and patiently listening to my first feeble attempts to turn incoherent ideas into theorems, he has provided much needed prodding at critical points of my career.

I am also grateful to Alexander Kechris for his advice and support. Many of the results here arose from problems he suggested and discussed with me. His influence should be clear to anyone checking the citations.

Several other mathematicians here and elsewhere have influenced my development. John Addison first exposed me to the beauty of descriptive set theory and its classical foundations. Leo Harrington and Ted Slaman have shared their insights and given me a deeper understanding of recursion theory. I have had many useful conversations with Su Gao about descriptive set theory.

I would also like to thank Chris Shannon for agreeing to serve on my dissertation committee.

I am proud to consider several of my fellow graduate students both as colleagues and friends. These include Stuart Zoble, David Lippel, David Lippe, Andrés Caicedo, and Ben Miller. In addition to mathematical discussion, they have helped to keep me in touch with humanity.

Numerous other members of the department have been helpful. I must particularly thank Catalina Cordoba for help and knowledge of the inner workings of the department.

I could not have succeeded in graduate school or in the rest of my life were it not for the love and support of my parents. Their influence on me is immeasurable. I must also thank my girlfriend, Rebecca Evans, who has had to endure living with an often neurotic mathematician.

Lastly, I must thank all of the people at Au Coquelet who have provided me with coffee and conversation for many years while I sat and made incomprehensible marks on my paper.

Chapter 0

Introduction

Descriptive Set Theory is the branch of set theory concerned with the definability and classification of sets of real numbers and other Polish spaces (separable, completely metrizable spaces). The origins of the field lie in the point-set topology of the early part of the twentieth century. Two themes recur in this area. The first is that of *definability.* Given a particular set of reals, one can ask whether it can be given an explicit definition in a formal language, and how complicated this definition must be. This question is related to describing how the set may be "built from below," starting with some "simple" sets and applying set-theoretic operations to them. A classical example of this is the case of the Borel sets of real numbers. The Borel sets are those sets in the σ-algebra generated by the open sets, so one can stratify the Borel sets based on how many iterations of the operations of complementation and taking countable unions are necessary to produce a given set starting with some collection of open sets. Equivalently, one can ask for the simplest definition of the set using only number quantifiers (where "simplest" generally refers to the number of quantifier alternations necessary). This gives a measure of the complexity of the set. Similarly, one may consider the class of projective sets, those obtained from the open sets by taking complements and continuous images (or projections). Here again we have a stratification based on the number of iterations of these operations needed to produce the sets, and again this is equivalent to asking how complicated a definition we need using real quantifiers.

The second theme is that of *hierarchies of complexity.* This is by no means distinct from the first theme, but manifests itself in a different sort of result. In many situations one studies a notion of reducibility, in terms of which one object may be deemed to be "simpler" than another. A standard example of this is the notion of Wadge reducibility. One set of reals A is said to be Wadge reducible to another set B, $A \leq_W B$, if A is the continuous preimage of B, i.e. there is a continuous function f such that $A = f^{-1}[B]$. One thus obtains a partial pre-order (distinct elements may be equivalent) on the set of all sets of reals, and can study the structure of this ordering.

Implicit in these considerations is that the sets we consider should be in some sense *definable*, or at least well-behaved in some way. If we consider sets which are not explicitly definable (ones whose definition requires a necessary use of the Axiom of Choice, for instance), there is very little that can be said about general situations. One routinely produces "pathological sets," such as sets which are not Lebesgue-measurable. If we restrict our attention to more special sets, though, we do not find most of these pathologies. The most explicitly definable (and best-behaved) sets are the Borel sets, followed by the projective sets. Other classes of sets which are suitably well-behaved are the Baire-measurable sets and the universally-measurable sets (although these classes do not immediately suggest a notion of definability). If one adopts stronger set-theoretic hypotheses than the standard axioms of mathematics (ZFC), one obtains broader classes of sets which are well-behaved. Perhaps the most inclusive natural notion of definability is to consider the definable sets of reals to be those in the model $L(\mathbb{R})$, i.e. those sets which are constructible using real parameters. These include, in particular, all projective sets. Under the axiom $AD^{L(\mathbb{R})}$, which is the statement that every infinite two-player game of perfect information on the integers whose payoff set is constructible using real parameters is determined, these sets are well-behaved. For instance, they satisfy the regularity properties held by the Borel sets: They are all Lebesgue-measurable, have the property of Baire, and have the perfect set property. We shall restrict ourselves here primarily to results in the Borel and projective context, but many of them apply to broader notions of definability under variable determinacy hypotheses, which we will occasionally mention in passing but often leave unremarked.

In recent years, descriptive set theory has been used to study and classify *definable equivalence relations.* An equivalence relation is a reflexive, symmetric, and transitive binary relation on some set; as the name suggests, it is used to separate a set of objects into groups, called equivalence classes, which have something in common. A basic example is the relation on the integers of equivalence mod m for some natural number m. Here an equivalence class consists of all integers having the same remainder upon division by m. In the case that an equivalence relation E is defined on the reals or some other Polish space X, we may identify the relation with its graph, $E \subseteq X^2$. We can then discuss the definability of the relation in this space. For instance, E is said to be a Borel equivalence relation if its graph is a Borel subset of the space X^2 (in the product topology).

Here we have a natural notion of reducibility. We say that one equivalence relation E on a space X is *Borel reducible* to a second relation F on a space Y if there is a Borel-measurable function f from X to Y such that for all x_1 and x_2 in X we have

$$x_1 \; E \; x_2 \Longleftrightarrow f(x_1) \; F \; f(x_2)$$

We write this as $E \leq_B F$. This says that the relation E is essentially no more complicated than F: Given F (together with a reducing function f) we may recover E. Embeddibility may also be viewed from the standpoint of quotient spaces. If one considers the two quotient spaces X/E and Y/F, then having a definable reduction f of E to F implies that there is a definable injection $\widetilde{f} : X/E \hookrightarrow Y/F$. We can thus say that the "definable cardinality" of the quotient space X/E is less than or equal to that of Y/F.

We say that two relations are Borel bireducible, $E \sim_B F$, if $E \leq_B F$ and $F \leq_B E$. There are several strengthenings of these notions. We say that E is Borel embeddable in F, $E \sqsubseteq_B F$ if there is an injective function reducing E to F; we write $E \approx_B F$ for biembeddability. The strongest notion is that of Borel isomorphism: We say $E \cong_B F$ if there is a Borel bijection of the underlying spaces which is a reduction of E to F (so that its inverse is a reduction of F to E).

Instead of Borel-measurable maps, we could instead require continuous functions in our reductions. In many cases we can in fact find such reductions, but in general this gives a stronger notion of reducibility. The Borel context turns out to be a more useful notion for most of the situations we will consider, though. It allows us to work around features of an equivalence relation which are really features of the underlying space and not intrinsic to the relation, and allows us to be somewhat more flexible in the way we code various mathematical objects as elements of a Polish space as described below.

The reducibility hierarchy among Borel and other classes of definable equivalence relations has been studied extensively. Although many questions remain open, some of the broader structure has been determined. Several nice classes of equivalence relations have been isolated, and certain "benchmark" equivalence relations have been identified. These concrete examples allow us to test other equivalence relations to determine if they are simpler or more complicated than the known examples, and thus give some measure of where these other relations fall in the reducibility hierarchy. We will discuss several of these relations below.

This program has proved particularly fruitful in determining the complexity of various *classification problems* in analysis and other areas of mathematics. Given a collection of objects (such as countable groups) we can ask how difficult it is to classify them up to some notion of equivalence (such as group isomorphism). This is a somewhat vague question, but we can give a formulation which allows us to make this precise. We first want to represent the objects under consideration as elements of some Polish or standard Borel space. In the case of countable groups, for instance, we can view a group G as having as its underlying set the natural numbers, ω, and code the group as an element x of $2^{\omega\times\omega\times\omega}$ by setting $x(i,j,k) = 1$ if and only if $g_1 \cdot g_j = g_k$; that is, we code the group multiplication into a subset of $\omega \times \omega \times \omega$. The set of elements which legitimately code groups can then be topologized so as to be a Polish space. We can now define the isomorphism relation on this space by setting two codes equivalent if they code isomorphic groups. This relation turns out to be an analytic equivalence relation on the space of codes.

Having done this, we are free to apply the theory of reducibility of equivalence relations just described. This allows us to give relative comparisons between various isomorphism problems as well as other concrete equivalence relations. We can also view this analysis as giving us a gauge of how complicated a set of invariants must be in order to completely classify the given objects. An *invariant* for a classification problem is something we assign to each of our objects so that equivalent objects are assigned the same invariant. For instance, dimension is an invariant for linear isometry of Banach spaces, and entropy is an invariant for isomorphism of measure-preserving transformations (although neither of these is a complete invariant). To be useful, invariants should be able to be calculated in some fairly explicit fashion; it thus makes sense to require a Borel-measurable function sending an object to its invariants (or some more broadly definable function). These invariants should also be as concrete as possible if they are to be useful. For instance, one could always take the equivalence class of a given object as a complete invariant, but this tells us nothing new about the structure of the classification problem.

The simplest sort of invariant one can imagine is a real number (or perhaps an integer, but we will almost always be interested in situations with uncountably many isomorphism types). If we can assign a real number to each object (in a Borel manner) in such a way that two objects are equivalent if and only if they are assigned the same real, we say that the equivalence relation (and the associated classification problem) is *concretely classifiable* or *smooth*. This is equivalent to saying that the relation is Borel reducible to the identity relation on the reals. This is perhaps the best sort of classification we could hope for; see [16] or [32] for some considerations of what a concrete set of invariants should be.

There are many non-smooth equivalence relations. A canonical such relation is the equivalence relation E_0. This is defined on the space 2^ω by setting xE_0y if and only if x and y differ on only finitely many coordinates. Each equivalence class is then a countable set. In the realm of Borel equivalence relations, E_0 is a a minimum non-smooth equivalence relation: A theorem of Harrington, Kechris, and Louveau says that a Borel equivalence relation is non-smooth if and only if E_0 continuously embeds in it (see [14]). For non-Borel equivalence relations, the situation is slightly more complicated, but E_0 is still a good test relation.

Having a non-smooth classification problem means we can not in a reasonable manner assign single reals as complete invariants. So one tries to find other more complicated types of invariants. Some objects may be assigned countable sets of reals as complete invariants. We denote by F_2 the equivalence relation of equality of countable sets of reals; for definiteness this can be defined on the space $\mathbb{R}^\omega$ by setting one sequence $\langle x_n \rangle_{n\in\omega}$ equivalent to a second sequence $\langle y_n \rangle_{n\in\omega}$ if we have $\{x_n : n \in \omega\} = \{y_n : n \in \omega\}$ (in other words, we ignore the ordering and repetitions in the sequence). Then, being able to assign countable sets of reals as complete invariants is tantamount to saying that the classification problem is reducible to the relation F_2. These invariants are more complicated than single reals, but again there are relations which can not be classified by such invariants. We can continue along these lines and try countable sets of countable sets of reals and so forth. The most general sort of invariant we can consider of this form are countable structures of some type. For instance, we may try to assign to each object a countable graph so that two objects are equivalent if and only if they are assigned isomorphic graphs. Such an equivalence relation is said to be *classifiable by countable structures*. An equivalence relation turns out to be classifiable by countable structures precisely when it is reducible to an orbit equivalence relation of the infinite symmetric group S_∞. There are universal such relations, i.e. relations to which we may reduce any isomorphism relation on a class of countable structures. The isomorphism of countable graphs is an example of a relation which is universal for relations which are classifiable by countable structures. Such a universal relation is said to be *Borel-complete*.

Classification by countable structures is a more lenient interpretation of what good invariants should be, but once again we can find classification problems which are not classifiable by countable structures. Even in the case of equivalence relations induced by continuous actions of Polish groups on Polish spaces there are examples of relations which are not so classifiable. Hjorth's theory of *turbulence* provides a general framework for showing that certain relations are not classifiable by countable structures; this theory is set out in [16]. In the case of Borel equivalence relations induced by continuous actions of Polish groups, being classifiable by countable structures turns out to be equivalent to not reducing a turbulent action; for general Polish group actions the situation is essentially the same (but slightly more complicated). We shall not need any details of this theory here except that certain specific equivalence relations are not classifiable by countable structures. As a result, any equivalence relation to which we can reduce such a relation will also not be classifiable by countable structures.

There are two directions from which to attack a classification problem. On the one hand, we may try to give a set of complete invariants which is as simple as possible; equivalently, we may try to reduce the associated equivalence relation to some relatively simple known equivalence relation. On the other hand, we may try to show that the problem is difficult, and that certain simple invariants are not possible, by reducing some known complicated equivalence relation to it. The ideal situation is when we are able to do both: We find a known canonical equivalence relation which is bireducible with the classification problem we are interested in. This allows us to give a reasonably precise analysis of the difficulty of the problem, and also gives a natural example to improve our intuitions about the abstract reducibility hierarchy.

This subject has many connections to ergodic theory, where one is interested in studying dynamical properties of transformations and more complicated group actions. In fact, ergodicity is used in a fundamental way to show that the relation E_0 is non-smooth. Given a transformation of a space X, a subset of the space is said to be *invariant* if it is closed under the transformation. A measure on the space is said to be invariant for the transformation if the measure of a set is equal to the measure of its image under the transformation. The measure is said to be *ergodic* if, for any invariant measurable set A (equivalently, for any invariant Borel set), either A or $X \setminus A$ has measure 0. In the case of the equivalence relation E_0 introduced above, we can view this as the orbit equivalence relation of a single transformation. We can then define a version of Lebesgue measure on

2^ω which is invariant and ergodic for this transformation.

One can show that an equivalence relation E which admits an invariant, ergodic measure is necessarily non-smooth. For suppose that there were some Borel-measurable function f reducing E to the identity relation on $\mathbb{R}$. Fix a countable basis $\{B_n\}$ for $\mathbb{R}$. This basis separates points, so that for any $x \neq y$ in $\mathbb{R}$ there is some n with $x \in B_n$ and $y \notin B_n$. We then have:

$$xE_0y \iff f(x) = f(y) \iff (\forall n)[f(x) \in B_n \leftrightarrow f(y) \in B_n]$$

Now the sets $f^{-1}[B_n]$ must be invariant Borel sets, so for each n either $f^{-1}[B_n]$ or its complement must have full measure. Let A_n be whichever has full measure. By the calculation above, we must have that $\bigcap_n A_n$ can not contain two inequivalent elements, and is hence a single equivalence class of E_0. Each equivalence class of E_0 is a countable set, though, and the countable intersection of sets of full measure must itself have full measure and cannot then be countable. Hence we could not have had such a reducing function f.

This argument is very similar to what is used in Vitali's construction of a non-Lebesgue-measurable set. It provides the simplest non-classifiability result, which will be used several times in the sequel. The other non-classifiability results we will use are more complicated and we refer the reader, for instance, to [16] for results about non-classifiability by countable structures.

0.1 Preliminaries

We assume some familiarity with descriptive set theory. Our basic references are [26] and [21], and any undefined notation may be found in these. The following remarks are meant to be only the briefest explanation of terms common to much of this dissertation, and we will set out additional terminology as it is needed.

We shall denote the set of natural numbers, $\{0, 1, 2, \ldots\}$, by ω. A fundamental notion will be that of a *Polish space*, i.e. a separable, completely metrizable space. When we are interested primarily in the Borel structure of a space and not the topology, we will sometimes consider *standard Borel spaces*. A standard Borel space is a set equipped with a σ-algebra which is the algebra of Borel sets for some Polish topology on the space. In addition to the real numbers, $\mathbb{R}$, we will often work on the spaces 2^ω (Cantor space) and ω^ω (Baire space). These spaces are topologized using the product topology, where 2 (resp. ω)

is given the discrete topology. They can be given the following metric: $d(x, y) = \frac{1}{1+n(x,y)}$, where $n(x, y)$ is the least n such that $x(n) \neq y(n)$. We also use the space $[\omega]^\omega$ of infinite sequences from ω; this is a G_δ subspace of ω^ω and hence Polish in the relative topology. We shall generally refer to elements of all three of these spaces as "reals" (the Baire space is in fact homeomorphic to the set of irrational reals in the subspace topology, and elements of the Cantor space are associated to elements of the unit interval by viewing a sequence as a binary fractional expansion, i.e. $x \in 2^\omega \sim \sum_{n\in\omega} x(n)/2^{n+1}$). Note that all uncountable Polish spaces are Borel isomorphic, that is, there is a Borel-measurable bijection between any two of them. Hence, in terms of Borel structure, they are all the same, although the actual topologies may differ. Thus, when we deal with Borel reducibility, the actual Polish space we are working on will be inessential.

A *Polish group* is a topological group whose topology is Polish. Given a continuous (or Borel-measurable) action of a Polish group G on a Polish space X we define the *orbit equivalence relation* E_G^X by setting

$$x \; E_G^X \; y \Longleftrightarrow (\exists g \in G)[g \cdot x = y]$$

In the case that the action is continuous we call this a *Polish G-space.*

A *tree* on a set X (typically 2, ω, or some product of these) is subset of $X^{<\omega}$ (the set of finite sequences of elements of X) closed under initial segments. A *branch* through a tree T on X is an element of X^ω all of whose initial segments are in T. The set of branches through a tree T is denoted by $[T]$. Closed subsets of 2^ω and ω^ω may be represented as the set of branches of a tree on 2 or ω, respectively.

A set is *analytic* (or $\mathbf{\Sigma}^1_1$) if it is the continuous image of a Polish space; this is equivalent to being the continuous image of a Borel set. On the Baire space ω^ω, a set is analytic precisely when it is the projection (onto, say, the first coordinate) of a closed subset of $(\omega^\omega)^2$. This is equivalent to being the projection of a tree on $\omega \times \omega$, where the projection of a tree T is denoted by $p[T]$ and we have

$$\alpha \in p[T] \Longleftrightarrow (\exists \beta)(\forall n)[(\alpha \restriction n, \beta \restriction n) \in T]$$

i.e. $p[T]$ is the projection onto the first coordinate of the set of branches through T.

A set is *co-analytic* (or $\mathbf{\Pi}^1_1$) if its complement is analytic. There is a well-developed structure theory for such sets; we refer the reader to the above references, particularly for the notions of $\mathbf{\Pi}^1_1$-norms and boundedness. Further levels of the projective hierarchy are formed similarly: A set is $\mathbf{\Sigma}^1_2$ if it is the continuous image of a $\mathbf{\Pi}^1_1$ set; a set is $\mathbf{\Pi}^1_2$ if it is the complement of a $\mathbf{\Sigma}^1_2$ set, and so forth.

A pointclass is a collection of subsets of a space. Usually we shall want the class to have certain closure properties, for instance being closed under countable unions or under continuous preimages. Given a pointclass Γ (such as $\mathbf{\Pi}^1_1$), we say that a set A is Γ-*hard* if, for every set B in Γ, we have that B is the continuous preimage of A. We say that A is Γ-*complete* if A is in Γ and is Γ-hard. We say that a set is *true* Γ if it is in Γ but is not in any simpler pointclass; more precisely, this will generally mean the set is in Γ but its complement is not. In ZFC, being true Γ is generally weaker than being Γ-complete, although for reasonable pointclasses under some determinacy assumptions the notions coincide.

0.2 Outline of Results

In Chapter 1 we consider the relation of isomorphism of Borel automorphisms of a Polish space. We say that two automorphisms are isomorphic if they are conjugate via some Borel bijection of their underlying spaces. We show that this relation is quite complicated in two senses. First, the relation is complicated in a descriptive set-theoretic sense (when coded appropriately):

Theorem 0.1 *The relation of isomorphism of Borel automorphisms is $\mathbf{\Sigma}^1_2$-complete.*

Second, it is quite high in the Borel reducibility hierarchy of equivalence relations. It reduces the relation of equality of Borel sets, so in particular it reduces any Borel equivalence relation. We in fact prove a somewhat stronger statement. We say that one automorphism f embeds in a second one g if there is a Borel injection φ such that $\varphi \circ f = g \circ \varphi$; this is equivalent to saying that f is isomorphic to the restriction of g to some invariant Borel set. We then have:

Theorem 0.2 *There is a Borel map ψ from Borel subsets of the Baire space to Borel automorphisms of the Cantor space such that $A \subseteq B$ if and only if $\psi(A)$ embeds in $\psi(B)$.*

Since two automorphism will be isomorphic precisely when they embed in each other, this gives the reduction of equality of Borel sets to isomorphism of Borel automorphisms. In particular, this shows that isomorphism of Borel automorphisms in not concretely classifiable, answering a question of [5].

In Chapter 2 we consider the notion of a weakly wandering sequence for a transformation T on a space X. A sequence Ω of integers is called *weakly wandering* for T if there is a (non-trivial) set $A \subseteq X$ such that we have $T^n[A] \cap T^m[A] = \emptyset$ for all $n \neq m$ in Ω. The sequence is called *exhaustive weakly wandering* if there is such an A which also satisfies $X = \bigcup_{n \in \Omega} T^n[A]$. The set of (exhaustive) weakly wandering sets which a transformation admits is an invariant for isomorphism (and indeed is the driving force in the arguments of Chapter 1). We define the following four sets:

$$\begin{aligned}
\mathcal{WW} &= \{\Omega : \Omega \text{ is weakly wandering for some } T\} \\
\mathcal{WW}_0 &= \{\Omega : \Omega \text{ is weakly wandering for some ergodic } T\} \\
\mathcal{EWW} &= \{\Omega : \Omega \text{ is exhaustive weakly wandering for some } T\} \\
\mathcal{EWW}_0 &= \{\Omega : \Omega \text{ is exhaustive weakly wandering for some ergodic } T\}
\end{aligned}$$

Using characterizations of these sets due to Kamae, Eigen and Hajian, these sets are all easily seen to be $\mathbf{\Sigma}^1_1$. We show that they are also all $\mathbf{\Sigma}^1_1$-complete. In fact:

Theorem 0.3 *If X is any set with $\mathcal{EWW}_0 \subseteq X \subseteq \mathcal{WW}$, then X is $\mathbf{\Sigma}^1_1$-hard.*

As a corollary of the proof, we derive the following theorem of Mannsfield. Recall that a set of natural numbers A is said to be a *difference set* if there is a second set B such that $A = \{m - n : m, n \in B, m \geq n\}$.

Theorem 0.4 (Mannsfield) *The set $\{A : A$ contains an infinite difference set$\}$ is $\mathbf{\Sigma}^1_1$-complete.*

Using the techniques we develop, we are also able to construct particular sequences of interest, such as a sequence which is exhaustive weakly wandering for some transformation, but which is not weakly wandering for any ergodic transformation (thus answering a question of [6]). We prove:

Theorem 0.5 *The four sets: $\mathcal{WW} \setminus (\mathcal{WW}_0 \cup \mathcal{EWW})$, $\mathcal{EWW} \setminus \mathcal{WW}_0$, $\mathcal{WW}_0 \setminus \mathcal{EWW}$, and $(\mathcal{WW}_0 \cap \mathcal{EWW}) \setminus \mathcal{EWW}_0$ are all non-empty. In fact, they are all $\mathbf{\Sigma}^1_1$-hard.*

In Chapter 1 we used the relation of equality of Borel sets in order to show that isomorphism of Borel automorphisms is a complicated equivalence relation. This equality relation is an example of an equivalence relation which reduces all Borel equivalence relations. Such a relation must necessarily be non-Borel. In Chapter 3 we consider such relations and raise the question of whether there can be any minimal or minimum such relations. We start by considering the relation of equality of Borel sets. This is a $\mathbf{\Pi}^1_1$ relation which reduces all Borel equivalence relations; however, it is neither minimal nor universal among such:

Theorem 0.6 *The relation of equality of Borel sets is not a universal $\mathbf{\Pi}^1_1$ equivalence relation, nor is it minimal among those $\mathbf{\Pi}^1_1$ relations which reduce every Borel equivalence relation.*

It remains open whether there can be such a minimal relation, or even a minimum such. In the case of $\mathbf{\Sigma}^1_1$ relations, we show that there is no minimum by producing a "minimal pair":

Theorem 0.7 *There are two $\mathbf{\Sigma}^1_1$ equivalence relations, E_1 and E_2, such that for any equivalence relation E we have that E is Borel if and only if E is Borel reducible to both E_1 and E_2.*

In particular, E_1 and E_2 are incomparable and above all Borel equivalence relations; any minimum relation would have to be below both of them and hence Borel.

We next consider several aspects of Polish metric spaces. By a *Polish metric space* we mean a Polish space equipped with a complete, compatible metric. In Chapter 4 we determine what the set of distances of such a space may be. The *distance set* of a metric space (X, d) is defined to be $\{d(x, y) : x, y \in X\}$. We prove the following characterization:

Theorem 0.8 *A set of non-negative real numbers containing 0 is the distance set of some Polish metric space if and only if either it is countable or it is analytic and has 0 as a limit point.*

We also give characterizations of the possible distance sets of certain types of Polish metric spaces, such as zero-dimensional, locally compact, etc.

In Chapter 5 we analyze the equivalence relation of *isometry* of Polish metric spaces. In the case of compact metric spaces, this relation is concretely classifiable by a result of Gromov (see [12]). In the case of arbitrary Polish metric spaces, though, this relation turns out to be very complicated. Gao and Kechris in [11] have shown that this relation is bireducible with the universal Borel action of a Polish group on a Polish space. We give an independent proof of one direction of this:

Theorem 0.9 *Let E_G^X be the orbit equivalence relation induced by a Borel action of a Polish group G on a Polish space X. Then the relation E_G^X is Borel reducible to the isometry relation on Polish metric spaces.*

Our technique is "local" in the sense that the spaces we build share many properties with the Polish groups in question. This allows us to calculate lower bounds on the complexity of the isometry relation restricted to particular classes of metric spaces. For example:

Theorem 0.10 *If E_G^X is the orbit equivalence relation of a Polish group G which is zero-dimensional and has a complete left-invariant metric, then E_G^X is reducible to the isometry of zero-dimensional Polish metric spaces.*

We use Hjorth's theory of turbulence (see [16]) to conclude that the isometry problem for several classes of spaces is not classifiable by countable structures. We also show that several classes are not concretely classifiable by reducing the equivalence relation E_0 to the corresponding isometry relation.

In the final chapter we consider spaces with large isometry groups. A metric space is said to be *homogeneous* if its isometry group acts transitively. A space is said to be *ultra-homogeneous* if every partial isometry between finite sets extends to an isometry of the whole space. We prove the following two classifications:

Theorem 0.11 *Isometry of homogeneous discrete Polish metric spaces is bireducible with graph isomorphism. Isometry of ultra-homogeneous discrete or locally-compact Polish metric spaces is bireducible with the relation F_2.*

Along the way we consider the isomorphism problem for classes of countable structures with large automorphism groups. In particular, we show that classifying countable vertex-transitive graphs is as difficult as classifying arbitrary countable graphs:

Theorem 0.12 *The isomorphism relation on vertex-transitive countable connected graphs is Borel-complete.*

We also give a characterization of which countable languages have a Borel-complete isomorphism problem for their class of structures with transitive automorphism groups:

Theorem 0.13 *Let $\mathcal{L}$ be a countable first-order language and let $\mathcal{K}$ denote the class of countable $\mathcal{L}$-structures which have transitive automorphism groups. Then the isomorphism* *problem for $\mathcal{K}$ is Borel-complete if and only if the signature of $\mathcal{L}$ contains no constant symbols and contains either an n-ary relation or function symbol for some $n \geq 2$ or else contains at least two unary function symbols. In all other cases the isomorphism problem for $\mathcal{K}$ is concretely classifiable.*

Chapter 1

Borel Automorphisms

In this chapter we consider several complexity questions regarding *Borel automorphisms* of a Polish space. Recall that a Borel automorphism is a bijection of the space with itself whose graph is a Borel set (equivalently, the inverse image of any Borel set is Borel). Since the inverse of a Borel automorphism is another Borel automorphism, as is the composition of two Borel automorphisms, the set of Borel automorphisms of a given Polish space forms a group under the operation of composition. We can also consider the class of automorphisms of all Polish spaces. We will be primarily concerned here with the following notion of equivalence:

Definition 1.1 *Two Borel automorphisms f and g of the Polish spaces X and Y are said to be Borel isomorphic, $f \cong g$, if they are conjugate, i.e. there is a Borel bijection $\varphi : X \to Y$ such that $\varphi \circ f = g \circ \varphi$.*

We restrict ourselves to automorphisms of uncountable Polish spaces, as the Borel automorphisms of a countable space are simply the permutations of the space. Since any two uncountable Polish spaces are Borel isomorphic, any Borel automorphism is Borel isomorphic to some automorphism of a fixed space. Up to Borel isomorphism, then, we can fix a Polish space and represent any Borel automorphism as an automorphism of this space. We will use the Cantor space 2^ω as our representative space.

We may then represent a Borel automorphism by its graph, which is a subset of $(2^\omega)^2$. This graph is a Borel set, and may thus be coded as a real using a coding of Borel sets. The set of Borel automorphisms can then be viewed as a set of reals, and the relation of Borel isomorphism is then an equivalence relation on this set. This allows us to analyze the complexity of this relation using descriptive set-theoretic techniques. Two natural questions arise:

1. How complicated is this equivalence relation descriptively; i.e., where does it fall in the Wadge hierarchy?

2. How complicated is this relation in the hierarchy of equivalence relations under Borel reducibility?

We will be able to completely answer the first question by showing that the isomorphism relation is $\mathbf{\Sigma}^1_2$-complete. We will be able to give a partial answer to the second question by showing that the relation is quite complicated: The equivalence relation of equality of Borel sets is Borel reducible to the isomorphism relation. In particular, this shows that any Borel equivalence relation is reducible to the isomorphism relation.

In [17] Hjorth considers analogous questions for the group of measure-preserving transformations of a measure space. There is a notable distinction between the two situations, since the group of measure-preserving transformations can be made into a Polish group and hence has a $\mathbf{\Sigma}^1_1$ conjugacy relation. This relation turns out to be $\mathbf{\Sigma}^1_1$-complete and not classifiable by countable structures, so that it is also quite complicated; however, it is strictly simpler than the conjugacy relation of Borel automorphisms. Here we see an example of the differences between measure-theoretic and descriptive contexts.

1.1 Setup

Although we will be representing all Borel automorphisms by automorphisms of the Cantor space, it will be more convenient to define them on other spaces. Thus we will use (f, X) to indicate that f is an automorphism of the Polish space X. The Polish spaces we use will be sufficiently similar to each other that we will have a uniform way of producing isomorphic automorphisms of the Cantor space. We will discuss this later.

We will want a more general way of comparing automorphisms than isomorphism. We introduce the following partial order on the class of Borel automorphisms:

Definition 1.2 *We set* $(f, X) \preceq (g, Y)$ *if there is a Borel injection* $\varphi : X \hookrightarrow Y$ *such that* $\varphi \circ f = g \circ \varphi$.

Note that in the above definition, $\varphi[X]$ will be a g-invariant Borel set (since the injective image of a Borel set is Borel), so that we have the following equivalent form of the definition:

$$(f, X) \preceq (g, Y) \iff \text{ there is a } g\text{-invariant Borel set } B \subseteq Y \text{ such that } (f, X) \cong (g, B)$$

Also notice that, by a standard Shroeder-Bernstein argument, we have:

$$(f, X) \cong (g, Y) \iff (f, X) \preceq (g, Y) \text{ and } (g, Y) \preceq (f, X)$$

We now explain our coding of Borel automorphisms. We first fix a good parameterization of the Borel subsets of $(2^\omega)^2$ (see [21]). This consists of sets $D \subseteq 2^\omega$ and $P, S \subseteq (2^\omega)^3$ where D is $\mathbf{\Pi}^1_1$, P is $\mathbf{\Pi}^1_1$, and S is $\mathbf{\Sigma}^1_1$ such that:

1. $d \in D \Longrightarrow P_d = S_d$ (so that these sections are Borel subsets of $(2^\omega)^2$)

2. $\{P_d : d \in D\}$ contains all Borel subsets of $(2^\omega)^2$

3. For any Polish space X and Borel set $A \subseteq X \times (2^\omega)^2$, there is a Borel function $p : X \to 2^\omega$ such that for all $x \in X$ we have $p(x) \in D$ and $A_x = P_{p(x)}$

The properties of this parameterization will be necessary for the definability considerations later.

Definition 1.3 *Let $\mathcal{BA}$, the set of codes for Borel automorphisms, be:*

$$\mathcal{BA} = \{d \in D : P_d \textit{ is the graph of a Borel automorphism of } 2^\omega\}$$

Then, up to isomorphism, every Borel automorphism has a code in $\mathcal{BA}$. We can then define the isomorphism relation on $\mathcal{BA}$ by letting two elements be equivalent if they code isomorphic Borel automorphisms.

1.2 Reducing Equality of Borel Sets

In their paper [8], Eigen, Hajian and Weiss show how to construct a continuum of non-isomorphic Borel automorphisms. By extending their technique we will show the following:

Theorem 1.4 *There is a Borel map from the Borel subsets of the Baire space ω^ω to Borel automorphisms of the Cantor space, sending A to f_A, such that:*

$$A \subseteq B \iff f_A \preceq f_B$$

That is, the partial order of inclusion among the Borel sets embeds into the partial order of $\preceq$ *among Borel automorphisms.*

We will explain the Borel-ness of the map below. From the theorem we derive the following corollary:

Corollary 1.5 *The equivalence relation of equality of (codes for) Borel subsets of* ω^ω *is Borel reducible to the equivalence relation of isomorphism of Borel automorphisms of* 2^ω.

We will discuss some other consequences of these results in the next section.

Let us first set out some notation. Throughout this chapter, the transformation φ_0 will refer to the *odometer map* on 2^ω, given by adding 1 with carry (so that if $x = 1^k \frown 0 \frown \alpha$, then $\varphi_0(x) = 0^k \frown 1 \frown \alpha$, and $\varphi_0(1^\infty) = 0^\infty$). Except on the eventually constant sequences, φ_0 induces the equivalence relation E_0 (recall that two sequences are E_0 equivalent if they differ on only finitely many coordinates). A set $B \subseteq 2^\omega$ will be called *smooth* or φ_0-smooth if there is a Borel transversal for the φ_0-saturation of B, i.e., a Borel set S such that for each orbit of φ_0 which meets B, S meets the orbit in exactly one point. This is equivalent to saying that $E_0 \upharpoonright B$ is a smooth equivalence relation. Note that B is φ_0-smooth if and only if $B \in \mathcal{W}(\varphi_0)$, the σ-ideal generated by the wandering sets for φ_0 (a set B is *wandering* for a transformation T if $B \cap T^n[B] = \emptyset$ for all $n \in \omega$, $n \neq 0$). This is true since in the realm of countable Borel equivalence relations, being smooth is equivalent to having a Borel transversal, i.e. a set meeting each equivalence class in precisely one point. A transversal is a wandering set, and one easily constructs a transversal from a wandering set.

Finally, for $n \in \mathbb{N}$, we let $\mathrm{ord}_2(n)$ be the largest integer k such that 2^k divides n (so that $\mathrm{ord}_2(n)$ is the first non-zero coordinate in the binary expansion of n). Similarly, we can define $\mathrm{ord}_2(\alpha)$ for $\alpha \in 2^\omega$ to be the first non-zero coordinate of α. We let $\mathrm{ord}_2(0) = \infty$.

Proof of Theorem 1.4

Let $\{s_n : n \in \omega\}$ be a recursive enumeration of $\omega^{<\omega}$ such that if $s_n \sqsubseteq s_m$ then $n \leq m$. For $\alpha \in \omega^\omega$, let:

$$\widetilde{\alpha} = \{n : s_n \sqsubset \alpha\}$$

Thus, we can think of the nodes in the tree $\omega^{<\omega}$ as being labeled by natural numbers, and $\widetilde{\alpha}$ lists those nodes which the branch α passes through. Note that if $\alpha \neq \beta$, then $\widetilde{\alpha} \cap \widetilde{\beta}$ is finite.

Let $A \subseteq \omega^\omega$ be a given non-empty Borel set. We will define the Borel automorphism f_A to be the restriction of the odometer map φ_0 to some φ_0-invariant Borel set X_A (in the case of $A = \emptyset$ we will simply take the automorphism to be smooth and aperiodic, so it will embed in all the others produced here). Then $f_A = (\varphi_0, X_A)$ will be a Borel automorphism of the standard Borel space X_A. We will explain at the end how to uniformly represent this as a Borel automorphism of the Cantor space. We define X_A as follows:

$$X_A = \{x : x \text{ is infinite and there is } \alpha \in A \text{ such that } x \setminus \widetilde{\alpha} \text{ is finite}\}$$

We are here identifying subsets of ω with their characteristic functions in 2^ω. We check that X_A is a φ_0-invariant Borel set. If we let

$$W_A = \{x : x \text{ is infinite and } x \subseteq \widetilde{\alpha} \text{ for some } \alpha \in A\}$$

then we have:

$$X_A = \bigcup_{i \in \omega} \varphi_0^i[W_A]$$

So X_A is invariant, and it will suffice to check that W_A is Borel. Note that

$$x \in W_A \iff \exists \alpha(\alpha \in A \text{ and } x \subseteq \widetilde{\alpha} \text{ and } x \text{ is infinite})$$

so that W_A is the projection onto the first coordinate of the set

$$U = \{(x, \alpha) \in 2^\omega \times \omega^\omega : x \text{ is infinite and } \alpha \in A \text{ and } x \subseteq \widetilde{\alpha}\}$$

The set U is Borel since A is Borel, and for each x there is at most one α with $(x, \alpha) \in U$, since $\widetilde{\alpha} \cap \widetilde{\beta}$ is finite for $\alpha \neq \beta$. Thus W_A is the continuous injective image of a Borel set, and hence Borel.

It is clear that if $A \subseteq B$ then $X_A \subseteq X_B$, so taking φ to be the identity map on X_A shows that $(\varphi_0, X_A) \preceq (\varphi_0, X_B)$ in this case. Suppose on the other hand that $A \not\subseteq B$. We will show that $(\varphi_0, X_A) \not\preceq (\varphi_0, X_B)$.

If $A \not\subseteq B$, then there is some $\alpha \in A \setminus B$. Fix such an α, and suppose there were a Borel injection $\varphi : X_A \hookrightarrow X_B$ such that $\varphi \circ \varphi_0 = \varphi_0 \circ \varphi$. Let $Z = \varphi[X_A] \subseteq X_B$. Then Z is a φ_0-invariant Borel set. Let:

$$\begin{aligned} W_\alpha &= \{x : x \subseteq \widetilde{\alpha} \text{ and } x \text{ is infinite}\} \\ X_\alpha &= \bigcup_{i \in \omega} \varphi_0^i[W_\alpha] \end{aligned}$$

Then X_α is a φ_0-invariant subset of X_A. Let

$$W_Z = \{x \in Z : x \cap \widetilde{\alpha} = \emptyset\}$$

Then we have that $Z = \bigcup_i \varphi_0^i[W_Z]$ since any $x \in Z$ has finite intersection with $\widetilde{\alpha}$. Now, for $i, j \in \omega$, set:

$$V_{i,j} = \varphi^{-1} \circ \varphi_0^i[W_Z] \cap \varphi_0^j[W_\alpha]$$

Then each $V_{i,j}$ is a Borel set, and we have

$$\begin{aligned} X_\alpha &= X_\alpha \cap \varphi^{-1}[Z] \\ &= \left(\bigcup_j \varphi_0^j[W_\alpha]\right) \cap \varphi^{-1}\left[\bigcup_i \varphi_0^i[W_Z]\right] \\ &= \left(\bigcup_j \varphi_0^j[W_\alpha]\right) \cap \left(\bigcup_i \varphi^{-1} \circ \varphi_0^i[W_Z]\right) \\ &= \bigcup_{i,j} \left(\varphi^{-1} \circ \varphi_0^i[W_Z] \cap \varphi_0^j[W_\alpha]\right) \end{aligned}$$

Thus, $X_\alpha = \bigcup_{i,j} V_{i,j}$. We claim that each $V_{i,j}$ is a wandering set for φ_0. To see this, note that

$$\varphi_0^n[V_{i,j}] = \varphi_0^n[\varphi^{-1} \circ \varphi_0^i[W_Z] \cap \varphi_0^j[W_\alpha]] = \varphi^{-1} \circ \varphi_0^i[\varphi_0^n[W_Z]] \cap \varphi_0^j[\varphi_0^n[W_\alpha]]$$

so that

$$\begin{aligned} V_{i,j} \cap \varphi_0^n[V_{i,j}] &= \varphi^{-1} \circ \varphi_0^i[W_Z] \cap \varphi^{-1} \circ \varphi_0^i[\varphi_0^n[W_Z]] \cap \varphi_0^j[W_\alpha] \cap \varphi_0^j[\varphi_0^n[W_\alpha]] \\ &= \varphi^{-1} \circ \varphi_0^i[\varphi_0^n[W_Z] \cap W_Z] \cap \varphi_0^j[\varphi_0^n[W_\alpha] \cap W_\alpha] \end{aligned}$$

But now let $n \neq 0$ and set $k = \mathrm{ord}_2(n)$. Note that if $k \in \widetilde{\alpha}$ then every element of $\varphi_0^n[W_Z]$ has a 1 in the k coordinate. This is because applying φ_0^n can be thought of as adding the binary representation of n, which has its first k coordinates equal to 0 and the k coordinate equal to 1, and each element of W_Z has the k coordinate equal to 0 since $k \in \widetilde{\alpha}$ implies $k \notin x$ for $x \in W_Z$. Then, since every element of W_Z has its k coordinate equal to 0, we have that if $k \in \widetilde{\alpha}$, then $\varphi_0^n[W_Z] \cap W_Z = \emptyset$. Similarly, if $k \notin \widetilde{\alpha}$, we have that $\varphi_0^n[W_\alpha] \cap W_\alpha = \emptyset$. In any event, then, we have $V_{i,j} \cap \varphi_0^n[V_{i,j}] = \emptyset$, so $V_{i,j}$ is wandering for φ_0.

Thus, since X_α is a countable union of wandering sets, we have that X_α is φ_0-smooth, which is equivalent to saying $E_0 \restriction X_\alpha$ is a smooth equivalence relation. This, however, is false, which we will see by showing that $E_0 \sqsubseteq_c E_0 \upharpoonright X_\alpha$. Since E_0 is non-smooth, this gives a contradiction.

We define the embedding as follows. Let the elements of $\widetilde{\alpha}$ be enumerated in increasing order as $\{a_n : n \in \omega\}$. For $x \in 2^\omega$ we define $f(x) \in 2^\omega$ by:

$$f(x)(k) = \begin{cases} x(n) & \text{if } k = a_{(2n)} \\ 1 & \text{if } k = a_{(2n+1)} \\ 0 & \text{if } k \notin \widetilde{\alpha} \end{cases}$$

Then $f(x)$ is infinite and $f(x) \subseteq \widetilde{\alpha}$, so that $f(x) \in X_\alpha$. The map is clearly continuous and injective. Finally, it is clear that $xE_0y \iff f(x)E_0f(y)$, so that this is the desired embedding.

Note that the above argument also will show that for two distinct Borel sets A and B, the automorphisms produced will not be conjugate via a universally measurable map, since E_0 has no universally measurable selector. Stronger results of this form also follow from stronger set-theoretic hypotheses.

Let us explain the underlying principle of the argument here. For a set $S \subseteq \mathbb{N}$, let $\mathrm{IP}\{S\}$ be the set consisting of all sums of finite subsets of S. Now let $\Omega_\alpha = \mathrm{IP}\{2^n : n \notin \widetilde{\alpha}\}$. These can be viewed as those integers whose binary representations (viewed as a finite subset of ω) are disjoint from $\widetilde{\alpha}$. Note that for $n_1 \neq n_2$ in Ω_α, we have $\mathrm{ord}_2(n_1 - n_2) \notin \widetilde{\alpha}$. From this we can see that $X_\alpha = \bigsqcup_{n \in \Omega_\alpha} \varphi_0^n[W_\alpha]$, where the union is disjoint. Thus, W_α is an exhaustive weakly wandering set for the transformation (φ_0, X_α) and the exhaustive weakly wandering sequence Ω_α (see the next chapter for the definitions of these concepts). On the other hand, Ω_α can not be an exhaustive weakly wandering sequence for φ_0 restricted to any invariant subset of X_B, as would be necessary if (φ_0, X_α) were to embed in (φ_0, X_B), since exhaustive weakly wandering sequences are isomorphism invariants for Borel automorphisms. We will discuss various issues concerning weakly wandering sequences in Chapter 2.

Thus, our map $A \mapsto (\varphi_0, X_A)$ has the desired property that $A \subseteq B$ if and only if $(\varphi_0, X_A) \preceq (\varphi_0, X_B)$. We must lastly check that we can produce codes for the automorphisms (φ_0, X_A) in a Borel way, that is, find representatives for them as automorphisms of the Cantor space.

Recall the good parameterization of Borel subsets of $(2^\omega)^2$ introduced in the last section. Fix now a similar parameterization of Borel subsets of the Baire space ω^ω. For each Borel set coded by a parameter d we have just produced an automorphism which is the odometer map on a certain subset of the Cantor space, which we will denote X_d. Uniformly in the parameter d we can produce a bijection from 2^ω onto X_d (note that this bijection in general will depend on the exact code d and not merely on the Borel set it codes). Let σ_d denote this map. If we then define:

$$\varphi_d = \sigma_d^{-1} \circ \varphi_0 \circ \sigma_d$$

we will have that φ_d is a Borel automorphism of the Cantor space such that $(\varphi_0, X_d) \cong (\varphi_d, 2^\omega)$. Now we can produce a map p which takes a code d for a Borel subset of ω^ω and assigns a code $p(d)$ for a Borel automorphism of 2^ω, namely φ_d. Note that if d and e are distinct codes for the same Borel set, we do not necessarily have that φ_d and φ_e are the same automorphism, but we do have that they are isomorphic.

Since the automorphism coded by $p(d)$ is isomorphic to (φ_0, X_d), the map p is the desired Borel map claimed in the theorem. The proof of the corollary is immediate, as again p is a reduction of the equivalence relation of equality of Borel sets to the equivalence relation of isomorphism of Borel automorphisms. □

1.3 Consequences of the Reduction

Let us first note an important consequence of the above theorem.

Corollary 1.6 *Let E be a Borel equivalence relation. Then E is Borel reducible to isomorphism of Borel automorphisms.*

Proof: Given a Borel equivalence relation E, we have the map $x \mapsto [x]_E$ sending each element to its equivalence class, a Borel set. Given a Borel code for E, we define this map to uniformly produce Borel codes for the equivalence classes. Composing with the map p described above, we then have

$$xEy \iff [x]_E = [y]_E \iff f_{p([x])} \cong f_{p([y])}$$

giving a reduction of E to the relation of isomorphism of Borel automorphisms. □

Thus, the equivalence relation of isomorphism of Borel automorphisms, and the classification problem it represents, is quite complicated. For instance, one can not in any reasonably definable manner (universally measurable, e.g.) assign a real number as a complete invariant (the relation is not concretely classifiable). This answers a question from

[5]. In fact, this relation can not be reduced to the orbit equivalence relation of any Polish group action, since the equivalence relation E_1 is reducible to it (see [22]).

The exact complexity of the isomorphism relation is not known; however, Su Gao has improved the above result:

Theorem 1.7 (Gao [10]) *Let E be an analytic equivalence relation. Then E is Borel reducible to isomorphism of Borel automorphisms.*

This should be contrasted with the situation where we consider conjugacy in the group of Lipschitz automorphisms of the Cantor space. These are the isometries of the Cantor space under the standard metric

$$d(x,y) = \frac{1}{1+n(x,y)} \text{ where } n(x,y) \text{ is the least } n \text{ such that } x(n) \neq y(n)$$

Here, both the equivalence relation of conjugacy by Lipschitz automorphisms and the relation of conjugacy by arbitrary Borel automorphisms are concretely-classifiable (see [5]).

Lastly, we should contrast the notion of Borel isomorphism of two Borel automorphisms with that of orbit equivalence. Given a Borel automorphism f, let E_f denote the associated orbit equivalence relation,

$$x \; E_f \; y \Longleftrightarrow \exists n \in \mathbb{Z} \, [y = f^n(x)]$$

Then we say that two automorphisms f and g are orbit equivalent if their orbit equivalence relations are isomorphic, i.e., there is a bijection between the underlying spaces such that two elements are in the same f-orbit if and only if their images are in the same g-orbit. For aperiodic, non-smooth automorphisms, there are only countably many isomorphism types, corresponding to the number of invariant probability measures that the automorphism admits (see [5]). All of the automorphisms produced in the above argument turn out to be orbit equivalent, since they admit no invariant probability measures. To see this, let a Borel set $A \subseteq \omega^\omega$ be given, and suppose μ were an invariant probability measure for (φ_0, X_A). For α in A, define

$$\begin{aligned} W_\alpha &= \{x : x \text{ is infinite and } x \subseteq \widetilde{\alpha} \ \} \\ X_\alpha &= \bigcup_i \varphi_0^i[W_\alpha] \end{aligned}$$

so that $W_A = \bigsqcup_{\alpha \in A} W_\alpha$ and $X_A = \bigsqcup_{\alpha \in A} X_\alpha$. Since $X_A = \bigcup_i \varphi_0^i[W_A]$ and μ is φ_0-invariant, we must have $\mu(W_A) > 0$. But this implies that the conditional probability measure of

W_α on X_α must be positive for some α. This contradicts that (φ_0, X_α) has no invariant probability measure, as it admits a nontrivial weakly wandering sequence (see [8]). Thus, although all of the automorphisms produced are non-isomorphic, they are all orbit equivalent. It may be of interest to consider notions of equivalence intermediate between orbit equivalence and isomorphism (such as Kakutani equivalence) and determine the complexity of these equivalence relations.

1.4 Descriptive Complexity

We now will calculate the descriptive complexity of the relations of isomorphism and embeddability of Borel automorphisms of the Cantor space. Recall the set $\mathcal{BA}$ of codes for Borel automorphisms. We now define:

$$\begin{aligned} \mathcal{E}_\preceq &= \{(d,e) \in D^2 : d, e \in \mathcal{BA} \;\&\; f_d \preceq f_e\} \\ \mathcal{E}_\cong &= \{(d,e) \in D^2 : d, e \in \mathcal{BA} \;\&\; f_d \cong f_e\} \end{aligned}$$

where we recall that f_d is the automorphism whose graph is coded by d. We start with a few straightforward calculations.

Proposition 1.8 *The set $\mathcal{BA}$ is $\mathbf{\Pi}^1_1$.*

Proof: We have:

$$\begin{aligned} \mathcal{BA}(d) \iff\; & D(d) \;\&\; (\forall x)(\exists y)[P(d,x,y)] \;\&\; (\forall y)(\exists x)[P(d,x,y)] \;\& \\ & \neg(\exists x)(\exists y_1)(\exists y_2)[y_1 \neq y_2 \;\&\; S(d,x,y_1) \;\&\; S(d,x,y_2)] \;\& \\ & \neg(\exists x_1)(\exists x_2)(\exists y)[x_1 \neq x_2 \;\&\; S(d,x_1,y) \;\&\; S(d,x_2,y)] \end{aligned}$$

Notice that the existential witnesses in the first line must be unique, so that they can be calculated in a Δ^1_1 way. We then have:

$$\begin{aligned} \mathcal{BA}(d) \iff\; & D(d) \;\&\; (\forall x)(\exists y \in \Delta^1_1(d,x))[P(d,x,y)] \;\&\; (\forall y)(\exists x \in \Delta^1_1(d,y))[P(d,x,y)] \;\& \\ & \neg(\exists x)(\exists y_1)(\exists y_2)[y_1 \neq y_2 \;\&\; S(d,x,y_1) \;\&\; S(d,x,y_2)] \;\& \\ & \neg(\exists x_1)(\exists x_2)(\exists y)[x_1 \neq x_2 \;\&\; S(d,x_1,y) \;\&\; S(d,x_2,y)] \end{aligned}$$

Since $\mathbf{\Pi}^1_1$ is closed under Δ^1_1-quantification, this gives the desired result. □

Proposition 1.9 *The sets $\mathcal{E}_\preceq$ and $\mathcal{E}_\cong$ are both $\mathbf{\Sigma}^1_2$.*

Proof: A similar calculation to the previous one shows that the set $\mathcal{BI}$ of codes for Borel injections is also $\mathbf{\Pi}^1_1$. We then have

$$\begin{aligned} \mathcal{E}_{\preceq}(d,e) &\iff \mathcal{BA}(d) \ \& \ \mathcal{BA}(e) \ \& \ (\exists c)[\mathcal{BI}(c) \ \& \ f_c \circ f_d = f_e \circ f_c] \\ &\iff \mathcal{BA}(d) \ \& \ \mathcal{BA}(e) \ \& \ (\exists c)[\mathcal{BI}(c) \ \& \\ &\quad (\forall x)(\forall y)(\forall z)(\forall w)[S(d,x,y) \ \& \ S(c,x,w) \ \& \ S(c,y,z) \Longrightarrow P(e,w,z)]] \end{aligned}$$

So we see that $\mathcal{E}_{\preceq}$ is $\mathbf{\Sigma}^1_2$. Then, since

$$\mathcal{E}_{\cong}(d,e) \iff \mathcal{E}_{\preceq}(d,e) \ \& \ \mathcal{E}_{\preceq}(e,d)$$

we also have that $\mathcal{E}_{\cong}$ is $\mathbf{\Sigma}^1_2$. □

We now show that this calculation is optimal by showing that these sets are also $\mathbf{\Sigma}^1_2$-hard. The technique we use was introduced by Adams and Kechris in [1] to prove an analogous result about countable Borel equivalence relations. They first construct an equivalence relation for each element of ω^ω with certain properties; we will be able to use the automorphisms previously constructed in place of these.

Theorem 1.10 *The sets $\mathcal{E}_{\preceq}$ and $\mathcal{E}_{\cong}$ are both $\mathbf{\Sigma}^1_2$-complete.*

Proof: We will in fact show that there is a single automorphism (φ^*, X^*) such that the following two sets are $\mathbf{\Sigma}^1_2$-complete:

$$\begin{aligned} \mathcal{E}^*_{\preceq} &= \{d : d \in \mathcal{BA} \ \& \ (\varphi^*, X^*) \preceq f_d\} \\ \mathcal{E}^*_{\cong} &= \{d : d \in \mathcal{BA} \ \& \ (\varphi^*, X^*) \cong f_d\} \end{aligned}$$

As these two sets are clearly Wadge-reducible to $\mathcal{E}_{\preceq}$ and $\mathcal{E}_{\cong}$, respectively, this will establish the theorem.

Our proof relies on a result of Steel about trees with full Borel uniformizations. Given a tree T on $\omega \times \omega$, let $[T]$ denote its set of branches, which is then a closed subset of $(\omega^\omega)^2$. We say that T has a full Borel Uniformization if the set $[T]$ can be uniformized by a total Borel function, i.e. there is a total function f such that for all x in ω^ω we have $(x, f(x)) \in [T]$. We let $\mathcal{FBU}$ denote the set of trees with full Borel uniformizations. Then we have the following (for a proof, see [1]):

Theorem 1.11 (Steel) *The set $\mathcal{FBU}$ is $\mathbf{\Sigma}^1_2$-complete.*

We will first define (φ^*, X^*), and then define a a Borel map $T \mapsto (\varphi_T, X_T)$ sending trees to automorphisms such that

$$T \in \mathcal{FBU} \iff (\varphi^*, X^*) \preceq (\varphi_T, X_T) \iff (\varphi^*, X^*) \cong (\varphi_T, X_T)$$

This map thus gives a Borel reduction of the set $\mathcal{FBU}$ to the sets $\mathcal{E}_{\preceq}$ and $\mathcal{E}_{\cong}$, establishing the theorem. For technical reasons we will work only with trees having at least one branch; this presents no difficulty since we can uniformly transform a given tree to one having at least one branch without affecting whether or not the tree has a full Borel uniformization. We also note that the set of trees on $\omega \times \omega$ is easily topologized so as to be a Polish space. For $\alpha \in \omega^\omega$ recall the map $\alpha \mapsto \widetilde{\alpha}$ and the sets

$$\begin{aligned} W_\alpha &= \{x : x \text{ is infinite and } x \subseteq \widetilde{\alpha}\ \} \\ X_\alpha &= \bigcup_{i \in \omega} \varphi_0^i[W_\alpha] \end{aligned}$$

where φ_0 again denotes the odometer map on 2^ω. Let also $X_\alpha^* = \omega^\omega \times X_\alpha$. We now define $X^* \subseteq \omega^\omega \times \omega^\omega \times 2^\omega$ by:

$$X^* = \bigsqcup_{\alpha \in \omega^\omega} X_\alpha^*$$

i.e. $(\alpha, x, y) \in X^* \iff y \in X_\alpha$. We define φ^* by:

$$\varphi^*(\alpha, x, y) = (\alpha, x, \varphi_0(y))$$

Next, for a tree T on $\omega \times \omega$ we define $X_T \subseteq (\omega^\omega)^2 \times \omega^\omega \times 2^\omega$ by:

$$X_T = \bigsqcup_{(\alpha,\beta) \in [T]} X_\alpha^*$$

i.e. $(\alpha, \beta, x, y) \in X_T \iff (\alpha, \beta) \in [T]$ and $y \in X_\alpha$. This union is non-empty because we stipulated that $[T]$ be non-empty. We then set

$$\varphi_T(\alpha, \beta, x, y) = (\alpha, \beta, x, \varphi_0(y))$$

This completes the definitions of the various automorphisms. It will suffice now to establish the following four lemmas.

Lemma 1.12 *For all trees T we have $(\varphi_T, X_T) \preceq (\varphi^*, X^*)$.*

Lemma 1.13 *For $T \in \mathcal{FBU}$ we have $(\varphi^*, X^*) \preceq (\varphi_T, X_T)$.*

Lemma 1.14 *If* $(\varphi^*, X^*) \preceq (\varphi_T, X_T)$ *then* $T \in \mathcal{FBU}$.

Lemma 1.15 *There is a Borel function* p *sending a tree* T *to a code* $p(T) \in \mathcal{BA}$ *such that* $(\varphi_T, X_T) \cong (f_{p(T)}, 2^\omega)$.

Proof of lemma 1.12: Let $(x, y) \mapsto \langle x, y\rangle$ be a Borel bijection of $(\omega^\omega)^2$ with ω^ω and set

$$\psi(\alpha, \beta, x, y) = (\alpha, \langle \beta, x\rangle, y)$$

Then ψ is a Borel injection from X_T into X^* and it is easy to check that $\psi \circ \varphi_T = \varphi^* \circ \psi$, witnessing the embedding. □(lemma)

Proof of lemma 1.13: Let f be a total Borel function uniformizing $[T]$, so that for all α we have $(\alpha, f(\alpha)) \in [T]$. Set

$$\psi(\alpha, x, y) = (\alpha, f(\alpha), x, y)$$

Then ψ is an injection from X^* into X_T with $\psi \circ \varphi^* = \varphi_T \circ \psi$. □(lemma)

Proof of lemma 1.14: Let ψ be a Borel injection from X^* to X_T such that $\psi\circ\varphi^* = \varphi_T\circ\psi$. We want to exhibit a Borel function f such that for all α we have $(\alpha, f(\alpha)) \in [T]$. Let α be given and let φ_α be the injection from X_α into X_T given by

$$\varphi_\alpha(y) = \psi(\alpha, 0^\infty, y)$$

(where 0^∞, the constant 0 sequence, is chosen arbitrarily). Note that

$$\varphi_\alpha(\varphi_0(y)) = \psi(\alpha, 0^\infty, \varphi_0(y)) = \psi(\varphi^*(\alpha, 0^\infty, y)) = \varphi_T(\psi(\alpha, 0^\infty, y)) = \varphi_T(\varphi_\alpha(y))$$

so $\varphi_\alpha \circ \varphi_0 = \varphi_T \circ \varphi_\alpha$, i.e. φ_α witnesses that $(\varphi_0, X_\alpha) \preceq (\varphi_T, X_T)$. Let a_α, b_α, x_α, and y_α be the (unique) functions such that

$$\varphi_\alpha(y) = (a_\alpha(y), b_\alpha(y), x_\alpha(y), y_\alpha(y))$$

Then we have

$$\begin{aligned} \varphi_\alpha \circ \varphi_0(y) &= (a_\alpha(\varphi_0(y)), b_\alpha(\varphi_0(y)), x_\alpha(\varphi_0(y)), y_\alpha(\varphi_0(y))) \\ \varphi_T \circ \varphi_\alpha(y) &= (a_\alpha(y), b_\alpha(y), x_\alpha(y), \varphi_0(y_\alpha(y))) \end{aligned}$$

so that, comparing coordinates, we have

$$\begin{aligned} a_\alpha(\varphi_0(y)) &= a_\alpha(y) \\ b_\alpha(\varphi_0(y)) &= b_\alpha(y) \\ x_\alpha(\varphi_0(y)) &= x_\alpha(y) \\ y_\alpha(\varphi_0(y)) &= \varphi_0(y_\alpha(y)) \end{aligned}$$

Thus, the functions a_α, b_α, and x_α are all invariant under φ_0.

Recall from above that φ_0 on X_α induces the equivalence relation $E_0 \restriction X_\alpha$, and we have a uniform way of embedding E_0 into $E_0 \restriction X_\alpha$. We can then uniformly lift Lebesgue measure (which is E_0-invariant) from 2^ω to X_α and associate to each α a quasi-invariant, non-atomic, ergodic probability measure μ_α for φ_0 on X_α. The ergodicity of μ_α and the invariance of a_α, b_α, and x_α under φ_0 implies that these functions must be constant on a φ_0-invariant set $M_\alpha \subseteq X_\alpha$ of μ_α-measure 1. Call the (necessarily unique) constant values of these functions α_0, β_0, and x_0, respectively.

First, we claim that $\alpha_0 = \alpha$. Notice that we have $\varphi_\alpha(y) = (\alpha_0, \beta_0, x_0, y_\alpha(y))$ on M_α, so that $y_\alpha(y) \in X_{\alpha_0}$. Since $y_\alpha \circ \varphi_0 = \varphi_0 \circ y_\alpha$, this shows that in fact $(\varphi_0, M_\alpha) \preceq (\varphi_0, X_{\alpha_0})$. The set X_α here is the same as was considered earlier, and the same argument shows that if $\alpha_0 \neq \alpha$ then M_α would be a wandering set (smooth set) for φ_0. But now we have that μ_α restricted to M_α is still a quasi-invariant non-atomic probability measure for φ_0, so M_α can not be φ_0-smooth. Thus $\alpha_0 = \alpha$.

Next, notice that we must have $(\alpha_0, \beta_0) \in [T]$, so that if we define our function f by $f(\alpha) = \beta_0$, we will have that $(\alpha, f(\alpha)) \in [T]$ for all α in ω^ω, so that f gives a full uniformization of $[T]$. We need only to check that f so defined is Borel-measurable. For this, we check that its graph is a Borel set. Let $G_f \subseteq (\omega^\omega)^2$ denote this graph.

For each α, let $\langle \widetilde{\alpha}_n \rangle_{n\in\omega}$ be the enumeration of $\widetilde{\alpha}$ in increasing order. We now define a function $e : \omega^\omega \times 2^\omega \to 2^\omega$ by setting:

$$e(\alpha, x)(k) = \begin{cases} x(n) & \text{if } k = \widetilde{\alpha}_{(2n)} \text{ for some } n \\ 1 & \text{if } k = \widetilde{\alpha}_{(2n+1)} \text{ for some } n \\ 0 & \text{if } k \notin \widetilde{\alpha} \end{cases}$$

Thus, if we define functions e_α by $e_\alpha(x) = e(\alpha, x)$, we have that each e_α is an embedding of the equivalence relation E_0 into $E_0 \restriction X_\alpha$. The function e is Borel, in fact continuous. We can now uniformly define measures μ_α on X_α by setting

$$\mu_\alpha(A) = \mu_L(e_\alpha^{-1}[A])$$

where μ_L is Lebesgue measure on 2^ω. Let $\pi_1 : X_T \to \omega^\omega$ be projection onto the second coordinate. We then compute:

$$\begin{aligned}
(\alpha,\beta) \in G_f &\iff \mu_\alpha(\{y : b_\alpha(y) = \beta\}) = 1\\
&\iff \mu_L(\{x : b_\alpha(e_\alpha(x)) = \beta\}) = 1\\
&\iff \mu_L(\{x : \pi_1 \circ \varphi_\alpha(e_\alpha(x)) = \beta\}) = 1\\
&\iff \mu_L(\{x : \pi_1 \circ \psi(\alpha, 0^\infty, e(\alpha,x)) = \beta\}) = 1
\end{aligned}$$

Now, noting that the set

$$\{(x,\alpha,\beta) : \pi_1 \circ \psi(\alpha, 0^\infty, e(\alpha,x))\}$$

is Borel, we can apply Theorem 17.25 of [21] to conclude that G_f is Borel. $\square$(lemma)

Proof of lemma 1.15: First we define canonical Borel bijections σ_α between 2^ω and X_α. To do this, let

$$X_0 = \{y \in 2^\omega : \ y \text{ is infinite and } y \text{ has only finitely many non-zero odd coordinates}\}$$

and fix a Borel bijection σ_0 between 2^ω and X_0. Let $\langle \widetilde{\alpha}_n \rangle$ enumerate $\widetilde{\alpha}$ as before, and let $\langle \widetilde{\gamma}_n \rangle$ enumerate $\omega \setminus \widetilde{\alpha}$. Define a bijection $\sigma_{0,\alpha}$ between X_0 and X_α by

$$\sigma_{0,\alpha}(x)(k) = \begin{cases} x(2n) & \text{if } k = \widetilde{\alpha}_n \\ x(2n+1) & \text{if } k = \widetilde{\gamma}_n \end{cases}$$

We then let $\sigma_\alpha = \sigma_{0,\alpha} \circ \sigma_0$.

We also have that for a tree T with $[T] \neq \emptyset$ there is a canonical bijection between $[T] \times \omega^\omega$ and ω^ω. Combining all of these maps, we can produce a canonical bijection σ_T between 2^ω and X_T. Now let $\widetilde{\varphi}_T = \sigma_T^{-1} \circ \varphi_T \circ \sigma_T$, so that $\widetilde{\varphi}_T$ is an automorphism of 2^ω with $(\widetilde{\varphi}_T, 2^\omega) \cong (\varphi_T, X_T)$. We then have that the relation

$$G(T,x,y) \iff T \text{ is a tree and } y = \widetilde{\varphi}_T(x)$$

is Borel and G_T is the graph of $\widetilde{\varphi}_T$ on 2^ω. So, using property (3) of our good parameterization of Borel subsets of $(2^\omega)^2$, we have a Borel function p sending a tree to an element of

2^ω such that $p(T) \in D$ and $G_T = P_{p(T)}$. This implies that $p(T) \in \mathcal{BA}$, and that $f_{p(T)} = \widetilde{\varphi}_T$ as required. □(lemma)

This finishes the proof of the theorem. □

Again following Adams and Kechris, we can draw a further result. We say that $(f, X) \preceq_{\sigma(\mathbf{\Sigma}^1_1)} (g, Y)$ if there is an embedding of f into g which is measurable with respect to the σ-algebra generated by the $\mathbf{\Sigma}^1_1$ sets, and similarly for $\cong_{\sigma(\mathbf{\Sigma}^1_1)}$. The Jankov-von Neumann uniformization theorem says that a closed plane set has a full projection if and only if it admits a full $\sigma(\mathbf{\Sigma}^1_1)$-measurable uniformization. Let $\mathcal{FP}$ be the set of trees T on $\omega \times \omega$ such that $[T]$ has full projection. Inspecting the above proof, we see that we have the following:

$$(\varphi^*, X^*) \preceq_{\sigma(\mathbf{\Sigma}^1_1)} (\varphi_T, X_T) \iff (\varphi^*, X^*) \cong_{\sigma(\mathbf{\Sigma}^1_1)} (\varphi_T, X_T) \iff T \in \mathcal{FP}$$

Clearly $\mathcal{FBU} \subseteq \mathcal{FP}$; however, $\mathcal{FP}$ is easily seen to be $\mathbf{\Pi}^1_2$-complete (see [1]). From this we see that the relations $\preceq_{\sigma(\mathbf{\Sigma}^1_1)}$ and $\cong_{\sigma(\mathbf{\Sigma}^1_1)}$ are $\mathbf{\Pi}^1_2$-hard. We also know that the two sets $\mathcal{FP}$ and $\mathcal{FBBU}$ are not equal, so there is a tree T which has full projection but does not have a full Borel uniformization. Letting $f = (\varphi_T, X_T)$ for such a tree T, and $g = (\varphi^*, X^*)$, we conclude the following:

Corollary 1.16 *There is a pair of Borel automorphisms f and g which are conjugate via a $\sigma(\mathbf{\Sigma}^1_1)$-measurable automorphism, but are not conjugate via a Borel automorphism.*

Chapter 2

Weakly Wandering Sequences

In this chapter we consider descriptive aspects of *weakly wandering sequences*. These are invariants for measure-preserving transformations or Borel automorphisms introduced by Hajian and Kakutani in [13], and were in fact the driving force in the argument of the previous chapter for showing two automorphisms were not isomorphic. We first consider how difficult it is to determine whether some sequence can be a weakly wandering sequence for some transformation, an exhaustive weakly wandering sequence, etc. We show that these are all $\mathbf{\Sigma}^1_1$-complete questions. We then use the techniques developed to construct some particular sequences to show that all of these notions are distinct; for instance, we construct a sequence which is an exhaustive weakly wandering sequence for some transformation T, but for no *ergodic* T.

2.1 Notation

Weakly wandering sequences are usually considered in the context of measure-preserving transformations. Since we are primarily interested in descriptive issues here, i.e. those involving Borel structure, we will specialize our notation to the case of Borel-measurable transformations on a Polish space (or standard Borel space) and will usually not assume that we have a measure present. If we do have a measure, we will assume that it is a Borel measure, i.e. the associated σ-algebra is the algebra of Borel sets. By a *transformation* we will generally mean a Borel automorphism of a Polish space. The definitions given here will thus be slightly different from what is standard; we will try to indicate the differences when they occur. In particular, we will often take as a non-triviality condition on a set that it meets every orbit infinitely often, instead of the usual condition of having positive measure (this should also be distinguished from the notion of *fullness*, i.e. simply meeting every orbit). Similarly, we will usually demand that conditions hold everywhere, instead of almost everywhere.

Definition 2.1 *Let T be a transformation on a Polish space X. An increasing sequence of natural numbers Ω in $[\omega]^\omega$ is a (positive) weakly wandering sequence for T if there is a Borel set A which meets every orbit infinitely often, such that for all n and m in Ω with $n \neq m$, we have:*

$$T^n[A] \cap T^m[A] = \emptyset$$

We say Ω is an exhaustive weakly wandering sequence for T if there is such an A which also satisfies:

$$X = \bigcup_{n \in \Omega} T^n[A]$$

We can also consider weakly wandering sequences which are subsets of $\mathbb{Z}$ (where we demand that the sequence extend infinitely in both directions); we will restrict ourselves here to subsets of $\mathbb{N}$ since the results we obtain will also hold for subsets of $\mathbb{Z}$. We discuss the necessary modifications for bi-infinite sequences in Section 2.5.

A transformation is said to be *ergodic* if it admits an invariant ergodic measure μ. A measure is *invariant* if $\mu(T[A]) = \mu(A)$ for any Borel set A, and it is ergodic if for any T-invariant Borel set A, either $\mu(A) = 0$ or $\mu(X \setminus A) = 0$. This is equivalent to saying that the associated orbit equivalence relation generated by T is non-smooth. We can now define the key sets in this chapter.

Definition 2.2 *With the notation given, let:*

$$\begin{aligned}
\mathcal{WW} &= \{\Omega : \Omega \text{ is weakly wandering for some } T\} \\
\mathcal{WW}_0 &= \{\Omega : \Omega \text{ is weakly wandering for some ergodic } T\} \\
\mathcal{EWW} &= \{\Omega : \Omega \text{ is exhaustive weakly wandering for some } T\} \\
\mathcal{EWW}_0 &= \{\Omega : \Omega \text{ is exhaustive weakly wandering for some ergodic } T\}
\end{aligned}$$

We see that $\mathcal{WW}_0$ and $\mathcal{EWW}$ are subsets of $\mathcal{WW}$ and that $\mathcal{EWW}_0 \subseteq \mathcal{EWW} \cap \mathcal{WW}_0$. We will see in Section 2.4 that these inclusions are proper.

There are known characterizations of these sets which we will use in establishing their descriptive complexity. We must first introduce some terminology. We say that two sets of integers A and B have a *direct sum*, and denote this by $A+B=A\oplus B$, if for every $a_1 \neq a_2$ in A and every $b_1 \neq b_2$ in B we have $a_1+b_1 \neq a_2+b_2$. We say that A and B have direct sum $\mathbb{Z}$ if, in addition, every integer equals $a+b$ for some a in A and b in B. Thus:

$$A \oplus B = \mathbb{Z} \iff (\forall k \in \mathbb{Z})(\exists! a \in A)(\exists! b \in B)[k = a + b]$$

For a given set $A \subseteq \mathbb{Z}$ or $\mathbb{N}$, we say A is *complemented* if there is a set B with $A \oplus B = \mathbb{Z}$. Complemented sequences of natural numbers were first studied by DeBruijn in [4].

There is a useful fact about direct sums which we will exploit. For $A \subseteq \mathbb{N}$ or $\mathbb{Z}$, let:

$$\begin{aligned} D(A) &= \{a_2 - a_1 : a_1, a_2 \in A \ \&\ a_2 \geq a_1\} \\ S(A) &= \{a_1 + a_2 : a_1, a_2 \in A\} \end{aligned}$$

Then we have that $A+B=A\oplus B$ if and only if $D(A)\cap D(B)=\{0\}$. This is an immediate consequence of the fact that $a_1+b_1=a_2+b_2$ if and only if $a_1-a_2=b_2-b_1$. We call a set B a *difference set* (from $\mathbb{N}$, resp. $\mathbb{Z}$) if $B=D(A)$ for some $A\subseteq\mathbb{N}$ (resp. $\mathbb{Z}$). We will say more about these sets in Section 2.3. There are two further observations we will use. First, if we define the *shift* of a set A by an integer n to be: $A+n=\{a+n : a\in A\}$, then we see that a shift of a set has the same difference set as the original set: $D(A+n)=D(A)$. Thus, shifts of A and B have a direct sum if and only if A and B do. Second, if $A\oplus B=\mathbb{Z}$, then also $(A+n)\oplus B=\mathbb{Z}$.

We say that an increasing sequence $H=\langle \ldots h_{-1}, h_0, h_1, h_2, \ldots\rangle$ is a *hitting sequence* if for each finite consecutive subsequence $h_{-n},\ldots,h_0,\ldots h_n$, this subsequence occurs shifted to both the right and left in H. That is, there are positive numbers k_l, k_r, n_l and n_r in $\mathbb{N}$ such that

$$h_{i+k_r} = h_i + n_r \text{ and } h_{i-k_l} = h_i - n_l \text{ for } -n \leq i \leq n$$

We say that such an H has the *shift-repeat property.* We then have the following characterization, due to Eigen and Hajian ([6]), extending work of Kamae ([19]):

Theorem 2.3 (Eigen and Hajian, Kamae) *For Ω in $[\omega]^\omega$:*

1. *$\Omega \in \mathcal{WW} \iff$ there is $H \subseteq \mathbb{Z}$ such that $H+\Omega = H\oplus\Omega$.*

2. $\Omega \in \mathcal{WW}_0 \iff$ *there is a hitting sequence* $H \subseteq \mathbb{Z}$ *such that* $H + \Omega = H \oplus \Omega$.

3. $\Omega \in \mathcal{EWW} \iff$ *there is* $H \subseteq \mathbb{Z}$ *such that* $H \oplus \Omega = \mathbb{Z}$.

4. $\Omega \in \mathcal{EWW}_0 \iff$ *there is a hitting sequence* $H \subseteq \mathbb{Z}$ *such that* $H \oplus \Omega = \mathbb{Z}$.

Note that this gives $\mathbf{\Sigma}^1_1$ definitions for each of these sets. We shall now show that they are in fact $\mathbf{\Sigma}^1_1$-complete.

2.2 Complexity of the Set of Weakly Wandering Sequences

We let $\mathcal{T}$ denote the set of trees on ω, i.e. the set of all subsets of $\omega^{<\omega}$ which are closed under initial segments. Trees can be naturally coded as reals in the following way. Let $\{s_n : n \in \omega\}$ be a recursive enumeration of $\omega^{<\omega}$ such that if $s_n \sqsubset s_m$ then $n < m$. We can then set:

$$\mathcal{T} = \{T \in P(\omega) : (\forall n)(\forall k)[(n \in T \;\&\; s_k \sqsubset s_n) \Longrightarrow k \in T]\}$$

This is a closed subset of 2^ω and can thus be viewed as a Polish space in the relative topology. A tree is *well-founded* if it has no infinite branches. A tree which contains an infinite branch is called ill-founded. We let $\mathcal{WF} \subseteq \mathcal{T}$ be the set of well-founded trees. The following is a standard fact (see [21] or [26]):

Proposition 2.4 *The set* $\mathcal{WF}$ *is* $\mathbf{\Pi}^1_1$*-complete as a subset of* $\mathcal{T}$*. Hence,* $\mathcal{T} \setminus \mathcal{WF}$ *is* $\mathbf{\Sigma}^1_1$*-complete.*

We will now show that the four sets $\mathcal{WW}$, $\mathcal{WW}_0$, $\mathcal{EWW}$, and $\mathcal{EWW}_0$ are $\mathbf{\Sigma}^1_1$-complete by showing that $\mathcal{T} \setminus \mathcal{WF}$ is the continuous inverse image of each of them. We in fact will prove a more general fact. The following is the main technical lemma in this chapter:

Lemma 2.5 *There is a continuous function* $f : \mathcal{T} \to [\omega]^\omega$ *sending* T *to* Ω_T *such that:*

1. *If* T *is ill-founded, then there is a hitting sequence* H *such that* $H \oplus \Omega_T = \mathbb{Z}$

2. *If* T *is well-founded then* $D(\Omega_T) \setminus \{0\}$ *meets every infinite difference set (from* $\mathbb{N}$ *or from* $\mathbb{Z}$*)*

We defer the proof in order to give the consequence:

Theorem 2.6 *The sets $\mathcal{WW}$, $\mathcal{WW}_0$, $\mathcal{EWW}$, and $\mathcal{EWW}_0$ are all $\mathbf{\Sigma}^1_1$-complete. In fact, if X is any set with $\mathcal{EWW}_0 \subseteq X \subseteq \mathcal{WW}$, then X is $\mathbf{\Sigma}^1_1$-hard.*

Proof: We see that with f constructed as in Lemma 2.5, we have that $f[\mathcal{T} \setminus \mathcal{WF}] \subseteq \mathcal{EWW}_0$ and $f[\mathcal{WF}] \cap \mathcal{WW} = \emptyset$. Thus, if X is such that $\mathcal{EWW}_0 \subseteq X \subseteq \mathcal{WW}$, we have $\mathcal{T} \setminus \mathcal{WF} = f^{-1}[X]$. □

Proof of Lemma 2.5 We will define the continuous map $T \mapsto \Omega_T$. To make things combinatorially simpler, we first replace T with a new tree (which we will also denote by T) where we add the empty sequence (the root) and all sequences of length 1. The point is to make sure that T contains infinitely many nodes. This can be done continuously (in the codes for trees), and will not affect whether T contains an infinite branch. So we fix a tree T which we assume has infinitely many nodes. We let $\langle t_i \rangle_{i\in\omega}$ enumerate the nodes of T relative to the ordering on $\omega^{<\omega}$ introduced above. We denote that a node t_i is a predecessor of t_j in T by writing $t_i \prec t_j$. Given a node $t_j \in T$, we say that t_i is the *immediate predecessor* of t_j if $t_i \prec t_j$ and there is no node t_k with $t_i \prec t_k \prec t_j$. Every node other than the root has a unique immediate predecessor. Our enumeration of T is such that if $t_i \prec t_j$ then $i < j$. In particular, t_0 is the root of the tree. We say two nodes t_i and t_j are incompatible, $t_i \perp t_j$, if neither is a predecessor of the other. We let $|t|$ denote the length of t.

We will build our sequence $\Omega = \Omega_T$ as the union of finite sequences ω_i, and will simultaneously construct potential witnesses H_i which will give a complement to Ω precisely when T contains an infinite branch. We will proceed in stages, defining at stage n the finite sets $\omega_n \subseteq \mathbb{N}$ and $H_n \subseteq \mathbb{Z}$, and setting $A_n = D(\omega_n)$ and $D_n = D(H_n)$. We also set $\beta_n = \bigcup_{i \leq n} D_i$.

A few notational points: We are really building finite *sequences*, rather than subsets, so we will make decisions not only to add some integers to Ω, but also to keep some out. We will thus require, for instance, that ω_{n+1} is an *end-extension* of ω_n, meaning that ω_{n+1} adds no new integers less than the largest integer in ω_n. We will denote that B end-extends A by $B \sqsupseteq A$, and say $B \sqsupset A$ if B properly extends A. For sequences from $\mathbb{Z}$ instead of $\mathbb{N}$, end-extension will mean that the new sequence contains no new elements between the largest and the smallest element of the original sequence. We will use $B \supseteq A$ as usual to mean B contains all elements of A. The largest and smallest elements of a set A are denoted $\max(A)$ and $\min(A)$, respectively. So we can express end-extension as:

$$A \sqsubseteq B \iff B \cap [\min(A), \max(A)] = A$$

Also, given a set $A \subseteq \mathbb{Z}$ and an integer n, we let $A + n$ denote the shift of A by n: $\{a+n : a \in A\}$. If $n > 0$ we call this a shift to the right, and if $n < 0$ we call it a shift to the left. We say that one sequence *occurs* in another if some shift of the first sequence forms a consecutive subsequence of the second. We let $A + B$ denote the set $\{a + b : a \in A, b \in B\}$; we use $A \oplus B$ to indicate that the sum is direct. The set $A - B$ is defined similarly. To distinguish this from the set-theoretic difference of A and B, we will always use $A \setminus B$ to denote $\{n : n \in A \text{ and } n \notin B\}$.

We will inductively define our sequence to satisfy the following nine conditions at each stage n (for $i, j \leq n$):

1. If $i < j$ then $\omega_i \sqsubset \omega_j$ and $\beta_i \sqsubset \beta_j$
2. $A_i \cap \beta_i = \{0\}$
3. $A_i \cup \beta_i \supseteq [0, \max(\beta_i)]$
4. If $t_i \prec t_j$ then $H_i \sqsubset H_j$
5. Given $i < j$, let k be such that t_k is the maximal mutual predecessor of t_i and t_j. Then we require that $D_i \cap D_j \subseteq D_k$
6. Given $i < j$ with $t_i \perp t_j$, and $a \in D_i \setminus \beta_{i-1}$ and $b \in D_j \setminus \beta_{j-1}$, we require that $|b - a| \in A_j$ (i.e., $|b - a| \notin B_j$)
7. Given i, let m be such that t_m is the immediate predecessor of t_i. Then we require that shifts of H_m occur in H_i both to the left and the right as consecutive blocks
8. Let $\langle r_k \rangle_{k \in \omega}$ be the enumeration of $\mathbb{Z}$: $\langle 0, 1, -1, 2, -2, \ldots \rangle$, and let $k_i = |t_i|$. Then we require that $r_{k_i} \in \omega_i \oplus H_i$
9. For $i \leq j$ we require that $D(D_i)$ and $S(D_i)$ are disjoint from $\beta_j \setminus D_i$

At the end of the construction we set:

$$\begin{aligned} \Omega &= \bigcup_i \omega_i \\ A &= \bigcup_i A_i = D(\Omega) \\ \beta &= \bigcup_i \beta_i \end{aligned}$$

Note that we will have $A \cap \beta = \{0\}$ and $A \cup \beta = \mathbb{N}$.

The purpose of most of these conditions is fairly clear: (1), (2), and (3) ensure continuity; (2) and (4) ensure that there will be an H having direct sum with Ω in case T has an infinite branch; (7) guarantees that it will be a hitting sequence; (8) guarantees that this sum will be all of $\mathbb{Z}$; and (5) and (6) will guarantee that if T is well-founded then $A \setminus \{0\} = D(\Omega) \setminus \{0\}$ meets all infinite difference sets. The last condition, though, is a bit mysterious; this turns out to be precisely the inductive assumption necessary to extend the construction from one stage to the next.

Let us grant that the construction has been carried out, and see that it meets the requirements of the lemma. First, the map $T \mapsto \Omega_T$ is continuous, as knowing the first n elements of Ω requires knowing at most the first n nodes of T. Suppose that T has an infinite branch $\langle t_{m_i} \rangle_{i \in \omega}$, i.e. for all i we have $t_{m_i} \prec t_{m_{i+1}}$. Let $H = \bigcup_i H_{m_i}$. Then, since each H_{m_j} end-extends H_{m_i} for $i < j$ (by condition (4)), we will have:

$$D(H) = \bigcup_i D(H_{m_i}) = \bigcup_i D_{m_i} \subseteq \beta$$

Since $A \cap \beta = \{0\}$, we will have $D(\Omega) \cap D(H) = \{0\}$, so that Ω will have a direct sum with H. Condition (7) guarantees that H is a hitting sequence, since any finite subsequence of H must eventually occur as a subsequence of one of the H_{m_i}, and thus occurs shifted to the right and left in $H_{m_{i+1}}$ and hence in H. Finally, since each t_{m_i} has length i, each r_i will occur in $\omega_{m_i} \oplus H_{m_i}$, and hence $\Omega \oplus H = \mathbb{Z}$. Thus H gives us the desired witness for Ω_T.

Conversely, suppose that $D(\Omega_T) \setminus \{0\}$ is disjoint from some infinite difference set; we will show that T has an infinite branch. Any infinite difference set from $\mathbb{Z}$ contains an infinite difference set from $\mathbb{N}$, since if $B = D(C)$ with $C \subseteq \mathbb{Z}$, then one of the two sets:

$$\begin{aligned} C_+ &= \{n : n \in C \text{ and } n \geq 0\} \\ C_- &= \{-n : n \in C \text{ and } n \leq 0\} \end{aligned}$$

will be infinite, and both $D(C_+)$ and $D(C_-)$ are subsets of $D(C)$ and hence of B. Thus, we have some set $C \subseteq \mathbb{N}$ such that $D(C) \cap D(\Omega) = \{0\}$. Since $D(C - n) = D(C)$ for any n, we may shift C so that $0 \in C$. Thus we may suppose $C \subseteq \beta$. Let C_i be the subsequence containing the first i elements of C. Let m_i be the least m such that $C_i \subseteq \beta_m$. Then $m_i \leq m_j$ for $i < j$, and since each β_m is finite we will have that $m_i \to \infty$ as $i \to \infty$.

We claim that $t_{m_i} \preceq t_{m_j}$ for $i < j$. Suppose this fails for some $i < j$, and let m be such that t_m is the maximal predecessor of t_{m_i} and t_{m_j}. We thus have $m < m_i$, $m < m_j$, and $m_i < m_j$. By the minimality of m_i and m_j, there are elements $a \in C_i$ and $b \in C_j$, with $a \neq b$, such that:

$$\begin{aligned} a &\in \beta_{m_i} \setminus \beta_{m_i - 1} = D_{m_i} \setminus \beta_{m_i - 1} \\ b &\in \beta_{m_j} \setminus \beta_{m_j - 1} = D_{m_j} \setminus \beta_{m_j - 1} \end{aligned}$$

Then, by condition (6), we have $|b-a| \in A_{m_j}$, and we know $|b-a| \neq 0$. But $|b-a| \in D(C)$, contradicting that $D(C) \cap A = \{0\}$. Thus, each $t_{m_{i+1}}$ extends t_{m_i} in T, and since these have lengths approaching ∞, we get an infinite branch through T.

Our function f will thus be as desired, once we show that we can build Ω_T as described.

The construction of Ω_T

We begin at stage 0 by setting $\omega_0 = \{0\}$ and $H_0 = \{0\}$. Then $A_0 = D_0 = \beta_0 = \{0\}$. It is easy to see that we have satisfied all the conditions, since $|t_0| = 0$ and $r_0 = 0$. So now suppose the construction has been completed to stage n, with $\omega_0, \ldots, \omega_n$ and $H_0, \ldots, H_n$ defined so as to satisfy the given conditions. We now define ω_{n+1} and H_{n+1} so as to satisfy them at stage $n+1$.

Let m be such that t_m is the immediate predecessor of t_{n+1} in T. We will proceed in three steps:

Step I We end-extend H_m to $\widetilde{H}_{n+1}$ so as to satisfy condition (7) while not violating any of the other conditions (specifically, we need to ensure that the new differences introduced here to not violate conditions (1), (2), (5), (6), and (9)).

Step II If necessary, we end-extend ω_n to $\widetilde{\omega}_{n+1}$ and extend $\widetilde{H}_{n+1}$ to H_{n+1} in order to satisfy condition (8); otherwise we take $H_{n+1} = \widetilde{H}_{n+1}$ and $\widetilde{\omega}_{n+1} = \omega_n$. Again, we need to preserve conditions (1), (2), (5), (6), and (9).

Step III We end-extend $\widetilde{\omega}_{n+1}$ to ω_{n+1} in order to satisfy condition (3) (and make sure it is a proper extension of ω_n), while not violating condition (2).

Step I. We will choose two numbers, Δ_l and Δ_r, in $\mathbb{N}$, and define $\widetilde{H}_{n+1}$ to be the union of H_m and two shifts of H_m:

$$\widetilde{H}_{n+1} = (H_m - \Delta_l) \sqcup H_m \sqcup (H_m + \Delta_r)$$

We must choose the Δ's to preserve the necessary conditions. First, to guarantee that we end-extend H_m, we must make sure that the three blocks are disjoint. This is satisfied is we require:

$$\Delta_l, \Delta_r > (\max(H_m) - min(H_m)) = \max(D_m) \tag{2.1}$$

To preserve condition (1), we must make sure that any new differences produced are bigger than any of the elements of β_n. The differences we have in D_{n+1} at this point will be in one of the following sets:

$$\begin{gathered} D_m, \quad \Delta_l + D_m, \quad \Delta_r + D_m, \quad (\Delta_l + \Delta_r) + D_m, \\ \Delta_l - D_m, \quad \Delta_r - D_m, \quad (\Delta_l + \Delta_r) - D_m \end{gathered}$$

New differences, then, are only produced by pairs of elements from different blocks, so it will suffice to make

$$\min(H_m) - \max(H_m - \Delta_l) > \max(\beta_n)$$

and

$$\min(H_m + \Delta_r) - \max(H_m) > \max(\beta_n)$$

Both of these are satisfied if we require:

$$\Delta_l, \Delta_r > \max(\beta_n) + \max(D_m) \tag{2.2}$$

Preserving condition (2) simply requires making sure all new differences are bigger than the largest element of A_n. This is satisfied if:

$$\Delta_l, \Delta_r > \max(A_n) + \max(D_m) \tag{2.3}$$

To preserve condition (5), we need to ensure that, for $i \leq n$, $D_i \cap D_{n+1} \subseteq D_k$, where t_k is the maximal predecessor of t_i and t_{n+1}. Note that t_k is also the maximal predecessor of t_i and t_m, so we already know $D_i \cap D_m \subseteq D_k$. It thus suffices to make sure that the new differences are not in β_n, which is satisfied by our previous condition that

$$\Delta_l, \Delta_r > \max(\beta_n) + \max(D_m)$$

To preserve condition (6), we need to ensure the following: Given $i \leq n$ with $t_i \perp t_{n+1}$, for any $a \in D_i \setminus \beta_{i-1}$ and any new difference b, we need $|b - a| \notin \beta_{n+1}$. If we require any new differences to be bigger than $2 \cdot \max(\beta_n)$, we will have that any $|b - a| = b - a > \max(\beta_n)$ and so is not in β_n. We thus first require:

$$\Delta_l, \Delta_r > \max(D_m) + 2 \cdot \max(\beta_n) \tag{2.4}$$

We must still ensure that we do not have $b-a \in \beta_{n+1}$, which will only happen if $b-a \in D_{n+1}$, i.e. we add two new differences b_1 and b_2 such that there is an element a in $D_i \setminus \beta_{i-1}$ with $b_2 - a = b_1$. This amounts to saying that $b_2 - b_1 \in D_i \setminus \beta_{i-1}$. To prevent this, it suffices to make sure that $D(D_{n+1})$ is disjoint from $D_i \setminus \beta_{i-1}$ whenever $t_i \perp t_{n+1}$. Note that if we satisfy condition (9) at stage $n+1$, then (taking $i = j = n+1$), we will have that $D(D_{n+1})$ is disjoint from $\beta_{n+1} \setminus D_{n+1} = \beta_n \setminus D_{n+1}$. Thus we will have:

$$D(D_{n+1}) \cap D_i \subseteq D(D_{n+1}) \cap \beta_n \subseteq D_{n+1}$$

Then, since $t_i \perp t_{n+1}$ by assumption, we have $D_i \cap D_{n+1} \subseteq D_k$ where $k < i$ is the maximal mutual predecessor. But then $D_k \subseteq \beta_{i-1}$, so $D(D_{n+1})$ will be disjoint from $D_i \setminus \beta_{i-1}$. So we will preserve condition (6) if we preserve condition (9).

So we lastly check that we can choose Δ_l and Δ_r so as to preserve condition (9). This will amount to ensuring two things at this stage. First, for $i \leq n$, we need $D(D_i)$ and $S(D_i)$ to be disjoint from $\beta_{n+1} \setminus D_i$. Second, we need $D(D_{n+1})$ and $S(D_{n+1})$ to be disjoint from $\beta_{n+1} \setminus D_{n+1}$.

For the first requirement, note that we already know by the inductive assumption that $D(D_i)$ and $S(D_i)$ are disjoint from $\beta_n \setminus D_i$, so we need only ensure that new differences in D_{n+1} are disjoint from $D(D_i)$ and $S(D_i)$ for all $i \leq n$. This will hold if we require:

$$\Delta_l, \Delta_r > \max(D_m) + \max\left(\bigcup_{i \leq n} (D(D_i) \cup S(D_i))\right) \tag{2.5}$$

We now consider the second requirement. Note that $\beta_{n+1} \setminus D_{n+1} = \beta_n \setminus D_{n+1}$, and we know by assumption that $D(D_m)$ and $S(D_m)$ are disjoint from $\beta_n \setminus D_m$. Thus, we will have that $D(D_m)$ and $S(D_m)$ are disjoint from $\beta_{n+1} \setminus D_{n+1}$. To simplify computations, we will require:

$$\begin{aligned} \Delta_l &> \Delta_r + 2 \cdot \max(D_m) \\ \Delta_r &> 2 \cdot \max(D_m) \end{aligned} \tag{2.6}$$

Considering the elements of D_{n+1} added so far, we then see that elements of $D(D_{n+1})$ at this step will be in one of the following sets:

$$\begin{array}{lllll} D(D_m), & \Delta_l + D(D_m), & \Delta_r + D(D_m), & (\Delta_l - \Delta_r) + D(D_m), & (\Delta_l + \Delta_r) + D(D_m), \\ & \Delta_l - D(D_m), & \Delta_r - D(D_m), & (\Delta_l - \Delta_r) - D(D_m), & (\Delta_l + \Delta_r) - D(D_m), \\ S(D_m), & \Delta_l + S(D_m), & \Delta_r + S(D_m), & (\Delta_l - \Delta_r) + S(D_m), & (\Delta_l + \Delta_r) + S(D_m), \\ & \Delta_l - S(D_m), & \Delta_r - S(D_m), & (\Delta_l - \Delta_r) - S(D_m) & \end{array}$$

Similarly, elements of $S(D_{n+1})$ so far will be in one of the sets:

$$S(D_m), \text{ or } \left\{ \begin{array}{llll} \Delta_l, & 2\Delta_l, & 2\Delta_l+\Delta_r, & \Delta_l+\Delta_r, \\ \Delta_r, & 2\Delta_r, & \Delta_l+2\Delta_r, & 2\Delta_l+2\Delta_r \end{array} \right\} \pm \left\{ \begin{array}{l} D(D_m) \\ S(D_m) \end{array} \right\}$$

(with the exception of $\Delta_l - S(D_m)$ and $\Delta_r - S(D_m)$). By this we mean that we take one of the numbers in the first set and either add or subtract the elements of one of the sets $D(D_m)$ or $S(D_m)$. We already saw that $D(D_m)$ and $S(D_m)$ would not cause problems, and all of the other elements can be kept out of β_n (and hence out of $\beta_{n+1} \setminus D_{n+1}$) if we require

$$\Delta_l, \Delta_r, (\Delta_l - \Delta_r) > \max(\beta_n) + \max(S(D_m)) \tag{2.7}$$

Thus, if we pick Δ_l and Δ_r to satisfy all the requirements 2.1–2.7, we will be able to meet condition (7) without violating any of the other conditions at this step. So we can fix some enumeration of $\mathbb{N}^2$ and define:

$$(\Delta_l, \Delta_r) = \text{ the least pair satisfying requirements 2.1–2.7}$$

We then set:

$$\widetilde{H}_{n+1} = (H_m - \Delta_l) \sqcup H_m \sqcup (H_m + \Delta_r)$$

This finishes step I.

Step II. Now we must meet condition (8). So let $r = r_{|t_{n+1}|}$; we will ensure $r \in \omega_{n+1} \oplus H_{n+1}$. Let $\widetilde{H}_{n+1}$ be as produced in step I, and set:

$$\begin{aligned} \widetilde{D}_{n+1} &= D(\widetilde{H}_{n+1}) \\ \widetilde{\beta}_{n+1} &= \beta_n \cup \widetilde{D}_{n+1} \end{aligned}$$

If r is already in $\omega_n \oplus \widetilde{H}_{n+1}$ we need do nothing at this step; set $\widetilde{\omega}_{n+1} = \omega_n$ and $H_{n+1} = \widetilde{H}_{n+1}$.

Otherwise, we need to add r to the sum. We will add an element $h < 0$ to $\widetilde{H}_{n+1}$ and add $w = r - h$ to ω_n. We will set:

$$\begin{aligned} H_{n+1} &= \widetilde{H}_{n+1} \cup \{h\} \\ \widetilde{\omega}_{n+1} &= \omega_n \cup \{w\} \end{aligned}$$

We need to see that we can choose an h so as not to violate any of the other conditions. First, we require

$$h < \min(\widetilde{H}_{n+1}) \tag{2.8}$$

in order to make sure condition (4) holds. To satisfy the first part of condition (1), we must make sure that w is bigger than $\max(\omega_n)$. This can be achieved by requiring:

$$-h > \max(\omega_n) + r \tag{2.9}$$

(recall that h is to be negative). For the second part of (1), we require:

$$-h > \max(\widetilde{\beta}_{n+1}) - \min(\widetilde{H}_{n+1}) \tag{2.10}$$

Now we consider the new differences added, both to β_n and to A_n. Let:

$$\begin{aligned} \widehat{D} &= \{h' - h : h' \in \widetilde{H}_{n+1}\} \\ \widehat{A} &= \{w - w' : w' \in \omega_n\} \end{aligned}$$

We will then have:

$$\begin{aligned} D_{n+1} &= \widetilde{D}_{n+1} \cup \widehat{D} \\ \beta_{n+1} &= \widetilde{\beta}_{n+1} \cup \widehat{D} \\ \widetilde{A}_{n+1} &= A_n \cup \widehat{A} \end{aligned}$$

To meet condition (2), we must guarantee that the three sets

$$\widetilde{\beta}_{n+1} \cap \widehat{A},\ A_n \cap \widehat{D},\ \text{and } \widehat{D} \cap \widehat{A}$$

are all empty (since we guaranteed that $\widetilde{\beta}_{n+1} \cap A_n$ was empty in step I). The first of these will be empty if we make sure that $\min(\widehat{A}) > \max(\widetilde{\beta}_{n+1})$. So we require:

$$-h > \max(\widetilde{\beta}_{n+1}) + \max(\omega_n) - r \tag{2.11}$$

The second set will be empty if we ensure that $\min(\widehat{D}) > \max(A_n)$. So we require:

$$-h > \max(A_n) - \min(\widetilde{H}_{n+1}) \tag{2.12}$$

For the third set, we need to be sure that we do not have $h' - h = w - w'$ for some $h' \in \widetilde{H}_{n+1}$ and $w' \in \omega_n$. But this would mean that

$$w' + h' = w + h = r$$

so we would already have had $r \in \omega_n \oplus \widetilde{H}_{n+1}$, contrary to our assumption.

To satisfy condition (5), it suffices that $\widehat{D} \cap \beta_n \subseteq D_m$. We have already ensured this by making $\min(\widehat{D}) > \max(\widetilde{\beta}_{n+1})$. Preserving condition (6) at this step amounts to showing that if $a \in D_i \setminus \beta_{i-1}$ for some $i \leq n$ with $t_i \perp t_{n+1}$, and $b \in \widehat{D}$, then $b - a \notin \beta_{n+1}$. As in step I, this reduces to ensuring that $D(\widehat{D})$ is disjoint from $D_i \setminus \beta_{i-1}$ for such i. But here

$$\begin{aligned} D(\widehat{D}) &= \{(h'_1 - h) - (h'_2 - h) : h'_1 \geq h'_2 \in \widetilde{H}_{n+1}\} \\ &= \{h'_1 - h'_2 : h'_1 > h'_2 \in \widetilde{H}_{n+1}\} = D(\widetilde{D}_{n+1}) \end{aligned}$$

We already ensured in step I that $D(\widetilde{D}_{n+1}) \cap D_i \subseteq \beta_{i-1}$ for these i, so this is fine.

We must lastly preserve condition (9). Again we have two cases to check. For $i \leq n$ we need $\beta_{n+1} \setminus D_i$ disjoint from $D(D_i)$ and $S(D_i)$, and we need $D(D_{n+1})$ and $S(D_{n+1})$ disjoint from $\beta_{n+1} \setminus D_{n+1}$. We already know, from step I, that $\widetilde{\beta}_{n+1} \setminus D_i$ is disjoint from $D(D_i)$ and $S(D_i)$, and that $D(\widetilde{D}_{n+1})$ and $S(\widetilde{D}_{n+1})$ are disjoint from $\widetilde{\beta}_{n+1} \setminus \widetilde{D}_{n+1}$.

So for the first case, we just need to make sure that $\widehat{D}$ is disjoint from $D(D_i)$ and $S(D_i)$ for $i \leq n$. This is achieved if $\min(\widehat{D}) > \max\left(\bigcup_{i \leq n}(D(D_i) \cup S(D_i))\right)$, which is satisfied if we require:

$$-h > 2 \cdot \max(\beta_n) - \min(\widetilde{H}_{n+1}) \tag{2.13}$$

For the second case, note that $\beta_{n+1} \setminus D_{n-1} = \widetilde{\beta}_{n+1} \setminus \widetilde{D}_{n+1}$. We first consider $D(D_{n+1})$. We have:

$$D(D_{n+1}) = D(\widetilde{D}_{n+1} \cup \widehat{D}) = D(\widetilde{D}_{n+1}) \cup (\widehat{D} - \widetilde{D}_{n+1}) \cup D(\widehat{D})$$

So we need only ensure that $\widehat{D} - \widetilde{D}_{n+1}$ and $D(\widehat{D})$ are disjoint from $\widetilde{\beta}_{n+1} \setminus \widetilde{D}_{n+1}$. The first can be ensured if $\min(\widehat{D}) > \max(\widetilde{D}_{n+1}) + \max(\widetilde{\beta}_{n+1})$. For this we require:

$$-h > \max(\widetilde{D}_{n+1}) + \max(\widetilde{\beta}_{n+1}) - \min(\widetilde{H}_{n+1}) \tag{2.14}$$

For the second set, we observe that $D(\widehat{D}) = D(\widetilde{H}_{n+1}) = \widetilde{D}_{n+1}$, and hence is trivially disjoint from $\widetilde{\beta}_{n+1} \setminus \widetilde{D}_{n+1}$.

We next consider $S(D_{n+1})$. Here we have:

$$S(D_{n+1}) = S(\widetilde{D}_{n+1} \cup \widehat{D}) = S(\widetilde{D}_{n+1}) \cup (\widetilde{D}_{n+1} + \widehat{D}) \cup S(\widehat{D})$$

The first set we already know is okay. The second and third sets will also be disjoint from $\widetilde{\beta}_{n+1}$ by the previously imposed conditions.

Thus we can take h to be any number satisfying the requirements 2.8–2.14. So we set:

$$h = \text{ the greatest integer satisfying requirements 2.8–2.14}$$

and take H_{n+1} and $\widetilde{\omega}_{n+1}$ to be as defined above. This completes step II.

Step III. We must now add elements to $\widetilde{\omega}_{n+1}$ in order to satisfy condition (3) while not violating condition (2). Let $a_0, \ldots, a_{l-1}$ enumerate $\max(\beta_{n+1}) \setminus (\beta_{n+1} \cup D(\widetilde{\omega}_{n+1}))$. We successively pick pairs (w_i, w'_i) for $i < l$ such that $w'_i - w_i = a_i$. We will then let

$$\omega_{n+1} = \widetilde{\omega}_{n+1} \cup \{w_i, w'_i : i < l\}$$

We need to ensure that any new differences introduced are not in β_{n+1}. We can do this by making:

$$\begin{aligned} w_0 &> \max(\beta_{n+1}) + \max(\widetilde{\omega}_{n+1}) \qquad (2.15) \\ w_i &> \max(\beta_{n+1}) + w'_{i-1} \text{ for } 1 \leq i \leq l \end{aligned}$$

We thus successively define w_i to be the least number satisfying this, and $w'_i = w_i + a_i$.

This completes step III, and hence stage $n+1$ of the construction. We thus see that the construction can be continued from one stage to the next, and the lemma is established. □

2.3 Complemented Sets and Difference Sets

We can draw a few corollaries from Lemma 2.5 and its proof. Let us recall two definitions given earlier:

Definition 2.7 *A set A (of integers or natural numbers) is said to be complemented (in $\mathbb{Z}$) if there is a set $B \subseteq \mathbb{Z}$ such that $A \oplus B = \mathbb{Z}$.*

Definition 2.8 *A set $A \subseteq \mathbb{N}$ is said to be a difference set (from $\mathbb{N}$, resp. $\mathbb{Z}$) if there is a set B (of natural numbers, resp. integers) such that $A = D(B)$.*

Then Lemma 2.5 gives immediately:

Corollary 2.9 *The set $\{A : A \text{ is complemented in } \mathbb{Z}\}$ is $\mathbf{\Sigma}^1_1$-complete.*

There are several results characterizing when certain types of sets are complemented; see for instance [7]. This corollary, though, shows that in a descriptive context no simpler classification for arbitrary sets is possible than the definition itself.

In contrast to this, consider the case of sets A which are complemented in $\mathbb{N}$, i.e., where there is a $B \subseteq \mathbb{N}$ such that $A \oplus B = \mathbb{N}$. These were the types of complements originally studied by DeBruijn. Here the classification is much simpler, descriptively. For a given A, if we think of the tree of attempts to build a complement for A by finite approximations, we see that this tree can be chosen to be finitely branching since the elements of A and B that can produce a given sum k must be bounded between 0 and k. Thus the question of whether a set is complemented in $\mathbb{N}$ amounts to asking whether there is an infinite branch through a finitely branching tree, which is only a $\mathbf{\Pi}^0_2$ question. In the case of complements in $\mathbb{Z}$, the difficulty is that potential witnesses for a given sum are unbounded.

As for difference sets, the proof of Lemma 2.5 produces the following corollary, due originally to Mannsfield. Let:

$$\begin{aligned} \mathcal{DF} &= \{A \subseteq \mathbb{N} : \ A \text{ is a difference set}\} \\ \mathcal{DF}_\infty &= \{A \subseteq \mathbb{N} : \ A \text{ is an infinite difference set}\} \\ \mathcal{CDF} &= \{A \subseteq \mathbb{N} : \ A \text{ contains an infinite difference set}\} \end{aligned}$$

Corollary 2.10 (Mannsfield) *The set* $\mathcal{CDF}$ *is* $\mathbf{\Sigma}^1_1$*-complete.*

Proof: Note the function $f : T \mapsto \Omega_T$ constructed in Lemma 2.5 induces the function $g : T \mapsto \beta_T$. The conditions of the lemma ensure that g is also continuous. If T contains an infinite branch, then $D(H) \subseteq \beta_T$ is an infinite difference set. If T is well-founded, then every infinite difference set meets $D(\Omega_T) \setminus \{0\} = \mathbb{N} \setminus \beta_T$, and hence is not contained in β_T. □

A related theorem is due to Schmerl:

Theorem 2.11 (Schmerl [27]) *The set* $\mathcal{DF}$ *is* $\mathbf{\Sigma}^1_1$*-complete.*

From this one easily sees that $\mathcal{DF}_\infty$ is also $\mathbf{\Sigma}^1_1$-complete. There does not seem to be any way to derive one of these two theorems from the other. The construction given here, for instance, necessarily produces sets which are not difference sets in the case that a tree T has more than one branch. Schmerl ([28]) has raised the following question:

Question 2.12 *Is it the case that every set* X *with* $\mathcal{DF}_\infty \subseteq X \subseteq \mathcal{CDF}$ *is* $\mathbf{\Sigma}^1_1$*-hard?*

If true, this would of course imply both theorems.

Let us note one further interesting fact about difference sets. This is not a difficult result, but does not seem to appear in the literature. We need to pretend that sets do not contain 0 to avoid trivialities.

Proposition 2.13 *The set* $\{A \subseteq \mathbb{N} : A \setminus \{0\}$ *meets every infinite difference set*$\}$ *is a filter on* $\mathbb{N}$.

Proof: The only condition which is non-trivial to check is that for $A, B \subseteq \mathbb{N}$ (both containing 0), if $A \cup B$ contains an infinite difference set, then either A or B contains an infinite difference set. Let C be such that $D(C) \subseteq A \cup B$. Consider the partition of $[C]^2$:

$$P = \{(n, m) \in [C]^2 : |n - m| \in A\}$$

Applying the infinite Ramsey theorem to this partition, we see that there is an infinite set $H \subseteq C$ such that either $[H]^2 \subseteq P$ or $[H]^2 \cap P = \emptyset$. In the first case, $D(H) \subseteq A$, and in the second $D(H) \subseteq B$. □

2.4 Constructing Particular Sequences

We can use the techniques developed in the proof of Lemma 2.5 to show that all of the obvious inclusions among the sets $\mathcal{WW}$, $\mathcal{WW}_0$, $\mathcal{EWW}$, and $\mathcal{EWW}_0$ are proper inclusions. Let us start with the most interesting case. In their paper [6], Eigen and Hajian ask (essentially): If Ω is an exhaustive weakly wandering sequence for some transformation T, must Ω be an exhaustive weakly wandering sequence for some *ergodic* transformation T'? The answer is no:

Theorem 2.14 *There is a sequence* $\Omega_1 \in [\omega]^\omega$ *such that* $\Omega_1 \in \mathcal{EWW}$ *but* $\Omega_1 \notin \mathcal{WW}_0$ *(so* Ω_1 *is exhaustive weakly wandering for some transformation, but is not even weakly wandering for any ergodic transformation).*

Proof: This amounts to showing that we can construct a sequence Ω such that there is a sequence H with $\Omega \oplus H = \mathbb{Z}$, but there is no *hitting* sequence H' with $\Omega + H' = \Omega \oplus H'$. The trick will be to build Ω so that $(\mathbb{N} \setminus D(\Omega)) \cup \{0\}$ contains no arithmetic progressions of length 3 (as a subset, not necessarily as a subsequence). We claim that for such an Ω and any hitting sequence H', we have $D(\Omega) \cap D(H') \neq \{0\}$, so that H' does not have a direct sum with Ω. To see this, it suffices to show that for any hitting sequence H' (or even a sequence with the shift-repeat property in one direction), $D(H')$ contains an arithmetic progression of length 3.

By shifting if necessary, we may assume that $0 \in H'$. Let $h_1 > 0$ be the first positive element of H' (we can handle the case where all elements of H are negative in essentially the same way). Then, by the shift-repeat property, there are h_n and h_{n+1} in H' such that $h_{n+1} - h_n = h_1 - 0$. But then $D(H')$ contains the elements $h_n - h_1$, $h_n - 0$, and $h_{n+1} = h_n + h_1$, which form an arithmetic progression of length 3 with common difference h_1.

So we build Ω and H such that $\Omega \oplus H = \mathbb{Z}$ but $(\mathbb{N} \setminus D(\Omega)) \cup \{0\}$ contains no arithmetic progressions of length 3. We will make sure that H extends infinitely in both directions. We will again build Ω and H in stages, where we construct at stage n the finite sequences ω_n and H_n. At the end we set:

$$
\begin{aligned}
\Omega &= \bigcup_n \omega_n \\
H &= \bigcup_n H_n
\end{aligned}
$$

We set $A_n = D(\omega_n)$ and $D_n = D(H_n)$. Then, at stage n we require the following for $i, j \leq n$:

1. If $i < j$ then $\omega_i \sqsubset \omega_j$, $H_i \sqsubset H_j$, and $D_i \sqsubset D_j$ (where, for the H_i's, we require that we extend properly in both directions)
2. $A_i \cap D_i = \{0\}$
3. $A_i \cup D_i \supseteq [0, \max(D_i)]$
4. With $\langle r_k \rangle_{k \in \omega}$ enumerating $\mathbb{Z}$ (with $r_0 = 0$), we require that $r_i \in \omega_i \oplus H_i$
5. D_i does not contain any arithmetic progression of length 3

Noting the remarks above, it is evident that if we have carried out the construction, then $\Omega_1 = \Omega$ satisfies the conclusions of the theorem. So we proceed with the construction.

Let $\omega_0 = \{0\}$ and $H_0 = \{0\}$. The conditions are clearly satisfied at stage 0. So we assume that ω_n and H_n have been defined so that the conditions hold at stage n. We construct ω_{n+1} and H_{n+1} to continue to satisfy them. We proceed in three steps:

Step I We add an element $h_- < 0$ to H_n and (if necessary) an element w to ω_n such that $w + h_- = r_{n+1}$ to satisfy condition (4), while preserving conditions (1), (2), and (5).

Step II We add an element $h_+ > 0$ to H_n, again preserving conditions (1), (2), and (5).

Step III We add elements to ω_n in order to satisfy condition (3), while preserving conditions (1) and (2).

Step I. If it is already the case that $r_{n+1} \in \omega_n \oplus H_n$, then we will only add an element h_- to H_n at this step, making sure that h_- is to the left of all previous elements. Otherwise, we will also add the element $w = r_{n+1} - h_-$ to ω_n, which we want to be to the right of previous elements. We let $r = r_{n+1}$ If we are not adding a w, we ignore any of the requirements below related to w. The new differences we introduce will then be:

$$\begin{aligned} \widehat{D} &= \{h' - h_- : h' \in H_n\} \\ \widehat{A} &= \{w - w' : w' \in \omega_n\} \end{aligned}$$

Condition (1) is satisfied if we make $w > \max(\omega_n)$, $h_- < \min(H_n)$, and $\min(\widehat{D}) > \max(D_n)$. So we require:

$$\begin{aligned} -h_- &> \max(\omega_n) - r \qquad (2.16) \\ h_- &< \min(H_n) \\ -h_- &> \max(D_n) - min(H_n) \end{aligned}$$

Preserving condition (2) requires that

$$(D_n \cup \widehat{D}) \cap (A_n \cup \widehat{A}) = \{0\}$$

We know already that $D_n \cap A_n = \{0\}$, so will make sure that the sets $\widehat{D} \cap A_n$, $D_n \cap \widehat{A}$, and $\widehat{D} \cap \widehat{A}$ are all empty. For the first, it suffices that $\min(\widehat{D}) > \max(A_n)$, so we require:

$$-h_- > \max(A_n) - \min(H_n) \qquad (2.17)$$

For the second, it suffices that $\min(\widehat{A}) > \max(D_n)$, so we require:

$$-h_- > \max(D_n) + \max(\omega_n) - r \qquad (2.18)$$

For the third set, we need to ensure that there are not elements $h' \in H_n$ and $w' \in \omega_n$ with $h' - h_- = w - w'$. As in Lemma 2.5, this only happens if $h' + w' = r$, in which case we do not add w, so $\widehat{A}$ is empty.

To satisfy condition (5), we need to be sure that $D_n \cup \widehat{D}$ contains no arithmetic progression of length 3. We know that D_n does not, and we know that elements of $\widehat{D}$ are bigger than elements of D_n, so any arithmetic progression must be $\langle d_0, d_1, d_2 \rangle$ with $d_0 < d_1 < d_2$, $d_0 + d_2 = 2d_1$, and $d_2 \in \widehat{D}$. We can set $d_2 = h_2 - h_-$, with $h_2 \in H_n$. We consider three cases, depending on whether d_0 and d_1 are in D_n or in $\widehat{D}$.

1. If $d_0, d_1 \in D_n$, set $d_0 = h_0 - h'_0$ and $d_1 = h_1 - h'_1$, with h_0, h'_0, h_1, and h'_1 in H_n. We then have:

$$h_0 - h'_0 + h_2 - h_- = 2h_1 - 2h'_1$$

so that

$$-h_- = 2h_1 + h'_0 - 2h'_1 - h_0 - h_2$$

We can prevent this from happening by requiring:

$$-h_- > 3\max(H_n) - 4\min(H_n) \tag{2.19}$$

2. If $d_0 \in D_n$ and $d_1 \in \widehat{D}$, then $d_0 = h_0 - h'_0$, $d_1 = h_1 - h_-$ where h_0, h'_0, and h_1 are in D_n. We then have:

$$h_0 - h'_0 + h_2 - h_- = 2h_1 - 2h_-$$

so that

$$-h_- = h_0 + h_2 - h'_0 - 2h_1$$

This is prevented if

$$-h_- > 2\max(H_n) - 3\min(H_n)$$

which we have already ensured.

3. If d_0 and d_1 are in $\widehat{D}$, let $d_0 = h_0 - h_-$ and $d_1 = h_1 - h_-$, with h_0 and h_1 in H_n. Then:

$$h_0 - h_- + h_2 - h_- = 2h_1 - 2h_-$$

i.e. $h_0 + h_2 = 2h_1$. This would mean that $\langle h_0, h_1, h_2 \rangle$ was an arithmetic progression in H_n, so that $\langle 0, h_1 - h_0, h_2 - h_0 \rangle$ was an arithmetic progression in D_n, contradictory to our assumption.

So we can now safely choose h_-:

$$h_- = \text{ the greatest negative number satisfying requirements 2.16–2.19}$$

We let:

$$\begin{array}{rcl}\widetilde{H}_{n+1} & = & H_n \cup \{h_-\} \\ \widetilde{\omega}_{n+1} & = & \begin{cases} \omega_n & \text{if } r_{n+1} \in \omega_n \oplus H_n \\ \omega_n \cup \{w\} & \text{if } r_{n+1} \notin \omega_n \oplus H_n \end{cases}\end{array}$$

We also set $\widetilde{A}_{n+1} = A_n \cup \widehat{A}$ and $\widetilde{D}_{n+1} = D_n \cup \widehat{D}$. This finishes step I.

Step II. We now wish to add an element $h_+ > \max(H_n)$. Let:

$$D' = \{h_+ - h : h \in \widetilde{H}_{n+1}\}$$

For condition (1), we need $\min(D') > \max(\widetilde{D}_{n+1})$, so we require:

$$h_+ > \max(\widetilde{D}_{n+1}) + \max(\widetilde{H}_{n+1}) \tag{2.20}$$

For condition (2) we need to make sure $D' \cap \widetilde{A}_{n+1} = \emptyset$, so we require:

$$h_+ > \max(\widetilde{A}_{n+1}) + \max(\widetilde{H}_{n+1}) \tag{2.21}$$

For condition (5) we have three cases which are essentially the same as in step I. We need only require:

$$h_+ > 4\max(\widetilde{H}_{n+1}) - 3\min(\widetilde{H}_{n+1}) \tag{2.22}$$

If we now let:

$$h_+ = \text{ the least positive number satisfying requirements 2.20–2.22}$$

all the conditions will be preserved. We then set:

$$H_{n+1} = \widetilde{H}_{n+1} \cup \{h_+\}$$

This finishes step II.

Step III. We must now add elements to $\widetilde{\omega}_{n+1}$ in order to satisfy condition (3) while not violating conditions (1) and (2). Let $a_0, \ldots, a_{l-1}$ enumerate $\max(D_{n+1}) \setminus (D_{n+1} \cup D(\widetilde{\omega}_{n+1}))$. As in Lemma 2.5, we successively pick pairs (w_i, w_i') for $i < l$ such that $w_i' - w_i = a_1$ and then let

$$\omega_{n+1} = \widetilde{\omega}_{n+1} \cup \left(\bigcup_{i<l} \{w_i, w_i'\} \right)$$

We need to ensure that any new differences introduced are not in D_{n+1}. We can do this by making:

$$\begin{aligned} w_0 &> \max(D_{n+1}) + \max(\widetilde{\omega}_{n+1}) \\ w_i &> \max(D_{n+1}) + w'_{i-1} \text{ for } 1 \le i \le l \end{aligned} \tag{2.23}$$

We thus successively define w_i to be the least number satisfying this, and $w'_i = w_i + a_i$. This finishes step III, and the construction. Again, this allows us to proceed to the next stage of the construction, and so finishes the proof. □

We now proceed to show that there is a sequence which is weakly wandering for some ergodic transformation, but is not exhaustive weakly wandering for any transformation.

Theorem 2.15 *There is a sequence $\Omega_2 \in [\omega]^\omega$ such that $\Omega_2 \in \mathcal{WW}_0 \setminus \mathcal{EWW}$.*

Proof: We build $\Omega = \Omega_2$ and a hitting sequence H such that $\Omega + H = \Omega \oplus H$. This guarantees that $\Omega \in \mathcal{WW}_0$. We will prevent Ω being in $\mathcal{EWW}$ by requiring that for any $w \in \Omega$ we have $w + 1 \in D(\Omega)$. To see that this suffices, suppose there were an H' with $\Omega \oplus H' = \mathbb{Z}$. By shifting, we may assume that $0 \in H'$. Then we have some w in Ω and h in H' with $w + h = -1$. But now $-h = w + 1 > 0$ and $0 - h \in D(H')$, contradicting that $D(\Omega) \cap D(H') = \{0\}$.

So we will build Ω and H as before, satisfying the following conditions at stage n for $i, j \le n$:

1. If $i < j$ then $\omega_i \sqsubset \omega_j$, $H_i \sqsubset H_j$, and $D_i \sqsubset D_j$
2. $A_i \cap D_i = \{0\}$
3. $A_i \cup D_i \supseteq [0, \max(D_i)]$
4. Shifts of H_i occurs to the left and right in H_{i+1} as consecutive blocks
5. For all $w \in \omega_i$, we have $w + 1 \in A_{i+1}$ (so $w \notin D_i$)

We initially set $\omega_0 = H_0 = \{0\}$, satisfying the conditions at stage 0. We now assume they have been satisfied at stage n and proceed to construct ω_{n+1} and H_{n+1}. We have two steps:

Step I Add shifts of H_n to satisfy condition (4) while respecting (1), (2), and (5).

Step II Add elements to ω_n to satisfy conditions (3) and (5), while respecting (1) and (2).

Step I. We will again pick numbers Δ_l and Δ_r in $\mathbb{N}$ and let:

$$H_{n+1} = (H_n - \Delta_l) \sqcup H_n \sqcup (H_n + \Delta_r)$$

We will then have:

$$D_{n+1} = D_n \cup (\Delta_l \pm D_n) \cup (\Delta_r \pm D_n) \cup ((\Delta_l + \Delta_r) \pm D_n)$$

We know that D_n is okay, so we can satisfy the rest of the conditions by making sure that elements of the remaining sets are bigger than $\max(D_n)$, bigger than $\max(A_n)$, and bigger than $\max(\omega_n) + 1$. The following requirement suffices:

$$\Delta_l, \Delta_r > \max(D_n) + \max(A_n) + 1 \tag{2.24}$$

Step II. As in the previous constructions, we can now form ω_{n+1} by successively adding pairs to add the necessary differences to A_{n+1} while avoiding D_{n+1}. This will finish the construction. □

We continue by producing a sequence which is weakly wandering for some transformation, but for no ergodic one, and which is not exhaustively weakly wandering for any transformation.

Theorem 2.16 *There is a sequence $\Omega_3 \in \mathcal{WW} \setminus (\mathcal{WW}_0 \cup \mathcal{EWW})$.*

Proof: We build $\Omega = \Omega_3$ and H in stages such that $\Omega + H = \Omega \oplus H$. We use previously discussed conditions to ensure that there is no such hitting sequence H, and also no H' with $\Omega \oplus H' = \mathbb{Z}$. At stage n we require the following, for $i, j \leq n$:

1. If $i < j$ then $\omega_i \sqsubset \omega_j$, $H_i \sqsubset H_j$ (in both directions), and $D_i \sqsubset D_j$
2. $A_i \cap D_i = \{0\}$
3. $A_i \cup D_i \supseteq [0, \max(D_i)]$
4. D_i does not contain any arithmetic progression of length 3
5. For all $w \in \omega_i$, we have $w \in A_{i+1}$ (so $w \notin D_n$)

We set $\omega_0 = H_0 = \{0\}$. Then we assume the construction is completed to stage n, and construct stage $n+1$. There are three steps:

Step I Add an element $h_- < \min(H_n)$, preserving the other conditions.

Step II Add an element $h_+ > \max(H_n)$.

Step III Extend ω_n to include necessary differences.

Step I. We want to add $h_- < \min(H_n)$. Set:

$$\widehat{D} = \{h - h_- : h \in H_n\}$$

We want to make $\widehat{D}$ disjoint from A_n and D_n, and bigger than $\omega_n + 1$, so we require:

$$-h_- > \max(A_n) + 1 + \max(D_n) - \min(H_n) \tag{2.25}$$

We also need to make sure that $D_n \cup \widehat{D}$ contains no arithmetic progressions of length 3. As in the previous construction, the following suffices:

$$-h_- > 3\max(H_n) - 4\min(H_n) \tag{2.26}$$

If we now take h_- to be the greatest negative number satisfying these two requirements, we will be fine. We let $\widetilde{H}_{n+1} = H_n \cup \{h_-\}$ and $\widetilde{D}_{n+1} = D_n \cup \widehat{D}$.

Step II. We now add an element $h_+ > \max(H_n)$, much like in step I. The following conditions will suffice to preserve the other conditions:

$$\begin{aligned} h_+ &> \max(A_n) + 1 + \max(\widetilde{D}_{n+1}) + \max(\widetilde{H}_{n+1}) \\ h_+ &> 4\max(\widetilde{H}_{n+1}) - 3\min(\widetilde{H}_{n+1}) \end{aligned} \tag{2.27}$$

We take h_+ to be the least such number, and let $H_{n+1} = H_n \cup \{h_+\}$.

Step III. We once again add pairs to ω_n to produce the needed differences. This completes the construction. □

The previous three constructions can be viewed as performing the construction in Lemma 2.5 along a tree with a single infinite branch, with some additional requirements. We now complete the picture of the inclusions among our four sets $\mathcal{WW}$, $\mathcal{WW}_0$, $\mathcal{EWW}$, and $\mathcal{EWW}$ by producing a sequence which is weakly wandering for some ergodic transformation, exhaustive weakly wandering for some other transformation, but exhaustive weakly wandering for no ergodic transformation. This time we will be building two additional sequences as witnesses, and the construction can be viewed as occurring along a tree with two infinite branches.

Theorem 2.17 *There is a sequence* $\Omega_4 \in [\omega]^\omega$ *such that* $\Omega_4 \in (\mathcal{WW}_0 \cap \mathcal{EWW}) \setminus \mathcal{EWW}_0$.

Proof: This time we will build $\Omega = \bigcup_i \omega_i$ as well as two other sequences, $H^o = \bigcup H_i^o$ and $H^x = \bigcup H_i^x$. H^o will be a hitting sequence with $\Omega + H^o = \Omega \oplus H^o$, and H^x will be such that $\Omega \oplus H^x = \mathbb{Z}$. We will prevent any hitting sequence giving a direct sum of $\mathbb{Z}$. We set, at each stage, $A_n = D(\omega_n)$, $D_n^o = D(H_n^o)$, $D_n^x = D(H_n^x)$, and $\beta_n = \bigcup_{k \leq n}(D_k^o \cup D_k^x)$. At each stage we satisfy the following ten conditions:

1. If $i < j$ then $\omega_i \sqsubset \omega_j$, $H_i^o \sqsubset H_j^o$ and $H_i^x \sqsubset H_j^x$ (in both directions), and $D_i^o \sqsubset D_j^o$ and $D_i^x \sqsubset D_j^x$
2. $A_i \cap \beta_i = \{0\}$
3. $D_i^o \cap D_i^x = \{0\}$
4. $A_i \cup \beta_i \supseteq [0, \max(\beta_i)]$
5. H_{i+1}^o contains left and right shifts of H_i^o as consecutive blocks
6. D_i^x contains no arithmetic progression of length 3
7. With $\langle r_k \rangle_{k \in \omega}$ enumerating $\mathbb{Z}$ (with $r_0 = 0$), we have $r_i \in \omega_i \oplus H_i^x$
8. For all $w \in \omega_i$, we have $w + 1 \notin D_i^o$
9. If $a \in D_i^o \setminus D_i^x$ and $b \in D_i^x \setminus D_i^o$, then $|b - a| \in A_i$
10. $D(D_i^o)$ and $S(D_i^0)$ are disjoint from $D_i^x \setminus \{0\}$;

    $D(D_i^x)$ and $S(D_i^x)$ are disjoint from $D_i^o \setminus \{0\}$

It is clear that H^o will be a hitting sequence with $\Omega + H^o = \Omega \oplus H^o$, and that we have $\Omega \oplus H^x = \mathbb{Z}$. We need to check that the conditions guarantee there is no hitting sequence H with $\Omega \oplus H = \mathbb{Z}$.

Suppose we have such an H, so $D(\Omega) \cap D(H) = \{0\}$. We may assume $0 \in H$ by shifting if necessary. Let $H_- = \{-h : h \in H, h < 0\}$. Then $H_- \subseteq D(H) \subseteq D(H^o) \cup D(H^x)$. We claim that either $H_- \subseteq D(H^o)$ or $H_- \subseteq D(H^x)$. If not, there are a and b in H_- with $a \in D(H^o) \setminus D(H^x)$ and $b \in D(H^x) \setminus D(H^0)$. But then there is an n with $a \in D_n^o \setminus D_n^x$ and $b \in D_n^x \setminus D_n^o$, so that $|b - a| \in A_n$ by condition (9). But $a \neq b$, and $|b - a| \in D(H_-) \subseteq D(H)$, a contradiction.

If $H_- \subseteq D(H^o)$, we claim $\Omega + H \neq \mathbb{Z}$. For if the sum were $\mathbb{Z}$, there would be $w \in \Omega$ and $h \in H$ with $w + h = -1$. So we have that $h < 0$, and so $-h = w + 1 \in H_- \subseteq D(H^o)$. This contradicts condition (8). On the other hand, if $H_- \subseteq D(H^x)$, we claim that H is not hitting. We have $D(H_-) \subseteq D(D(H^x)) = \bigcup_n D(D_n^x)$. We also have that $D(D(H^x))$ is disjoint from $D(H^o) \setminus D(H^x)$ by condition (10) and so $D(D(H^x)) \cap D(H^o) = \{0\}$ by condition (3). Thus $D(H_-) \cap D(H^o) = \{0\}$ and $D(H_-) \cap D(\Omega) = \{0\}$, so that we have $D(H_-) \subseteq D(H^x)$ by condition (4). But now, as in the previous argument, if H is hitting, H_- must contain an arithmetic progression of length 3, which is impossible since $D(H^x)$ does not, by condition (6).

Thus, our construction will suffice once we carry it out. We start with $\omega_0 = H_0^o = H_0^x = \{0\}$. Assume we have finished stage n; we now construct stage $n+1$. We have four steps:

Step I We extend H_n^0 to meet condition (5), while respecting conditions (1), (2), (3), (8), (9), and (10).

Step II We extend H^x to the left and extend ω_n (if necessary) to meet condition (7), respecting conditions (1), (2), (3), (6), (9), and (10).

Step III We extend H^x to the right, preserving the same conditions.

Step IV We extend ω_n to meet condition (4), while preserving (1) and (2).

Step I. As usual, we will have

$$H_{n+1}^o = (H_n^o - \Delta_l) \sqcup H_n^o \sqcup (H_n^o + \Delta_r)$$

where we pick Δ_l and Δ_r large enough to make these sets disjoint and to make any new differences bigger than $\max(A_n)+1$, $\max(D_n^x)$, and $\max(D_n^o)$. The new differences will then be:

$$\widehat{D} = (\Delta_l \pm D_n^o) \cup (\Delta_r \pm D_n^o) \cup ((\Delta_l + \Delta_r) \pm D_n^o)$$

We must still preserve conditions (9) and (10).

For condition (9), we must ensure that for $b \in \widehat{D}$ and $a \in D_n^x \setminus D_n^o$, we have $b - a \notin \beta_n \cup \widehat{D}$. Keeping this out of β_n can be done by making Δ_l and Δ_r sufficiently large. For $b - a$ to be in $\widehat{D}$, we would have $D(\widehat{D})$ meeting $D_n^x \setminus D_n^o$. Preserving condition (10) will thus suffice to preserve condition (9).

To preserve condition (10) we must do two things. We must make sure that $\widehat{D}$ is disjoint from $D(D_n^x)$ and $S(D_n^x)$, and we must make sure that $D(D_n^o \cup \widehat{D})$ and $S(D_n^0 \cup \widehat{D})$ are disjoint from $D_n^x \setminus \{0\}$. The first is achieved by making Δ_l and Δ_r sufficiently large. For the second, we already know that differences and sums from D_n^o are okay. A difference or sum involving one element of D_n^o and one element of $\widehat{D}$ can be kept out of D_n^x by making Δ_l and Δ_r sufficiently large. As in the proof of Lemma 2.5, the inductive assumption allows us to satisfy the condition by making

$$\Delta_l, \Delta_r, (\Delta_l - \Delta_r) > \max(D_n^x) + \max(S(D_n^o))$$

Thus, taking Δ_l and Δ_r to be the least pair satisfying the requirements will be sufficient.

Step II. We add an element $h_- < \min(H_n^x)$ to H_n^x. If $r_{n+1} \notin \omega_n \oplus H_n^x$, we also add an element $w > \max(\omega_n)$ to ω_n; otherwise we add nothing to ω_n at this stage. We let the new differences be:

$$\begin{aligned} \widehat{D} &= \{h' - h_- : h' \in H_n^x\} \\ \widehat{A} &= \{w - w; : w' \in \omega_n\} \end{aligned}$$

The following conditions can be met by making $-h_-$ and w sufficiently large:

- $\widehat{D}$ is disjoint from A_n, D_{n+1}^o, $D(D_{n+1}^o)$ and $S(D_{n+1}^o)$
- $\widehat{A}$ is disjoint from D_n^x
- $(D_n^x \cup \widehat{D})$ contains no arithmetic progression of length 3 with at least one element in D_n^x
- For $b \in \widehat{D}$ and $a \in D_{n+1}^o \setminus \{0\}$, we have $|b - a| \notin (D_n^x \cup D_n^o)$
- $\widehat{D} - D_n^x$, $\widehat{D} + D_n^x$, and $S(\widehat{D})$ are disjoint from D_{n+1}^o

The only things left to preserve are the following:

- $\widehat{D}$ is disjoint from $\widehat{A}$
- $\widehat{D}$ contains no arithmetic progression of length 3
- $D(\widehat{D})$ is disjoint from $D_{n+1}^o \setminus \{0\}$

As before, the first of these is guaranteed by the assumption that $r_{n+1} \notin \omega_n \oplus H^x_n$. The second is guaranteed since an arithmetic progression in $\widehat{D}$ would imply one in H^x_n, and the third is guaranteed because $D(\widehat{D}) = D(H^x_n) = D^x_n$ which we know to be disjoint from $D^o_{n+1} \setminus \{0\}$. So we can take h_- to be the greatest negative number satisfying the appropriate conditions and add it to H^x_n. If necessary, we also add $w = r_{n+1} - h_-$ to ω_n. This completes step II.

Step III. This is handled like step II (without adding to ω_{n+1}).

Step IV. We again add pairs to ω_{n+1} to include the necessary differences in A_{n+1}. This will complete the construction at stage $n+1$. □

2.5 Modifications

In the four constructions in the previous section, we have been able to show that the four sets $\mathcal{WW} \setminus (\mathcal{WW}_0 \cup \mathcal{EWW})$, $\mathcal{EWW} \setminus \mathcal{WW}_0$, $\mathcal{WW}_0 \setminus \mathcal{EWW}$, and $(\mathcal{WW}_0 \cap \mathcal{EWW}) \setminus \mathcal{EWW}_0$ are non-empty. By combining the construction in Lemma 2.5 with the new conditions used here, we can in fact show that these four sets are all $\mathbf{\Sigma}^1_1$-hard. We will briefly sketch the modifications necessary.

We will want to use conditions (1)–(6) and condition (9) from the original construction; however, conditions (7) and (8) will sometimes be replaced by complementary conditions. Let us define the two alternatives:

(7′) Each D_i contains no arithmetic progression of length 3

(8′) If $w \in \omega_i$, then $w + 1 \notin \beta_i$

Then, to show $\mathcal{EWW} \setminus \mathcal{WW}_0$ is $\mathbf{\Sigma}^1_1$-hard, we would use conditions (7′) and (8), for $\mathcal{WW}_0 \setminus \mathcal{EWW}$ we use (7) and (8′), and for $\mathcal{WW} \setminus (\mathcal{EWW} \cup \mathcal{WW}_0)$ we use (7′) and (8′). For $(\mathcal{EWW} \cap \mathcal{WW}_0) \setminus \mathcal{EWW}_0$ we build H^o_i and H^x_i and define D^o_i, D^x_i, β^o_i, and β^x_i accordingly. We include the corresponding conditions for D^o_i and D^x_i, and also require $\beta^o_i \cap \beta^x_i = \{0\}$. As in the original argument, we can check that if there is an H which has a direct sum with Ω then there is an infinite branch through T; moreover, H will have to have the properties we wish of it.

In all of the constructions, we have built our sequences Ω as subsets of $\mathbb{N}$. We can also get the same results for sequences which are subsets of $\mathbb{Z}$ unbounded in both directions. There is no difficulty in adding negative elements to ω_n in the last step of the construction; we simply must make them small enough so that new differences avoid any previously constructed sets. The only modification necessary is that we should replace the condition "$w + 1 \notin \beta_i$" by "$|w + 1| \notin \beta_i$" for $w \in \omega_n$. This will change our requirements slightly, but causes no difficulty, and will again establish that there is no H with $H \oplus \Omega = \mathbb{Z}$ when necessary.

There is one question we should mention. We have been able to determine the complexity of the set of sequences which are exhaustive weakly wandering for some transformation T and so forth. It is not clear though, what the set of sequences which are, say, weakly wandering for a particular transformation can look like.

Definition 2.18 *For a given transformation T, set:*

$$\mathcal{WW}(T) = \{\Omega : \ \Omega \text{ is weakly wandering for } T\}$$

The set $\mathcal{EWW}(T)$ is defined similarly.

Question 2.19 *Which sets can be $\mathcal{WW}(T)$ or $\mathcal{EWW}(T)$ for some transformation T? For some ergodic T? In particular, what are the possible complexities of these sets?*

Chapter 3

Equivalence Relations Which Reduce All Borel Ones

In Chapter 1, we showed that the equivalence relation of isomorphism of Borel automorphisms was complicated by showing that we could reduce to it the equivalence relation of equality of codes for Borel sets. This implies in particular that any Borel equivalence relation is reducible to the isomorphism relation. In this chapter we try to obtain a more precise analysis of the complexity of the equivalence relation of equality of Borel sets (which would give more information about the isomorphism problem, as well as other problems to which it can be reduced). We also consider several questions about equivalence relations to which any Borel equivalence relation can be reduced.

3.1 Equality of Borel Sets

Recall the equivalence relation of equality of codes for Borel sets which we introduced in Chapter 1. We will use a slightly different formulation here which is more convenient to our purposes, but we can uniformly transfer between the two coding schemes, so this will not affect the notions.

We will mean by a Borel code on ω^ω a pair (T, f) where T is a well-founded tree on ω and f is a map from the terminal nodes of T to ω. The code indicates how a Borel set is built from basic open sets by taking countable intersections and negations (we could alternately consider countable unions). The Borel set coded by (T, f) is defined recursively on rank. We let $B(t)$ be the Borel set coded by a node t in T. Then $B(\langle\rangle)$, the set coded by the root of T will be the Borel set coded by (T, f). We fix an enumeration of the basic clopen intervals of ω^ω, $\langle C_i \rangle_{i\in\omega}$. Then:

1. A terminal node t codes the clopen set $B(t) = C_{f(t)}$.
2. A non-terminal node t codes the set $B(t) = \bigcap_{i\in\omega}(\omega^\omega \setminus B(t \frown i))$, i.e. the intersection of the complements of the Borel sets coded by the nodes immediately below t.

Since T is well-founded, this produces a well-defined Borel set coded by (T, f). Note that our coding produces only certain representations of sets (G_δ but not F_σ, for instance); we could introduce a more complicated coding to remedy this, but since we are working in a completely regular space this is not a real problem; it simply affects slightly the ranks assigned to some Borel sets.

Let $\mathcal{BC}$ denote the set of Borel codes. Note that (appropriately coded) this is a $\mathbf{\Pi}^1_1$ set, since it is a $\mathbf{\Pi}^1_1$ property of a tree to be well-founded. For a given countable ordinal α we will let $\mathcal{BC}_\alpha$ denote the set of Borel codes (T, f) where T is well-founded of rank less than α. These correspond to the Borel sets of Borel rank (roughly) less than α. All the $\mathcal{BC}_\alpha$'s are Borel sets, since for a fixed α being well-founded of rank less than α is a Borel property. Note that a given Borel set will have codes of arbitrarily high rank, so one should remember that $\mathcal{BC}_\alpha$ refers to all codes of rank less than α and not to all codes for Borel sets of rank less than α. We will denote the Borel set coded by (T, f) as $B(T, f)$; often we will use x to denote a Borel code, in which case we denote the set it codes as $B(x)$. We are now ready to define the relevant equivalence relations.

Definition 3.1 *Let E_B be the equivalence relation of equality of Borel sets, defined on the space $\{(T, f) : T$ is a tree and f is a function from the terminal nodes to $\omega\}$ by*

$$x \ E_B \ y \iff x = y \ \vee \ (x, y \in \mathcal{BC} \ \& \ B(x) = B(y))$$

For each countable ordinal α we also define the relation of equality of Borel sets of rank less than α, given by

$$x \ E_\alpha \ y \iff x = y \ \vee \ (x, y \in \mathcal{BC}_\alpha \ \& \ B(x) = B(y))$$

All of these are $\mathbf{\Pi}^1_1$ equivalence relations. It is Borel to ask whether a given element $z \in \omega^\omega$ is an element of the Borel set coded by a given parameter x, so

$$x \ E_B \ y \iff (\forall z)[z \in B(x) \iff z \in B(y)]$$

We have already noted that any Borel equivalence relation is Borel reducible to E_B. Let us record several other properties of these relations:

- For $\alpha < \omega_1$, any $\mathbf{\Pi}^0_\alpha$ equivalence relation E satisfies $E \leq_B E_\alpha$.
- For $\alpha < \beta < \omega_1$, we have $E_\alpha <_B E_\beta$. This is proved by Hjorth ([15]) using Stern absoluteness.
- For $\alpha < \omega_1$, $E_\alpha <_B E_B$. This follows from the previous result, since every E_β is reducible to E_B.

Let us also note that the Borel coding we use does not have a substantial effect; any standard system of Borel codes should produce an equality relation which is bireducible with the one given here.

Our first question is whether E_B is as complicated as possible, namely, whether every $\mathbf{\Pi}^1_1$ equivalence relation is Borel reducible to it. This is ostensibly possible, since Hjorth ([18]) has shown that there is a universal $\mathbf{\Pi}^1_1$ equivalence relation. This turns out not to be the case, though.

Theorem 3.2 *The relation E_B is not a universal $\mathbf{\Pi}^1_1$ equivalence relation.*

We will show this by exhibiting a $\mathbf{\Pi}^1_1$ equivalence relation which is not reducible to E_B. We begin with a lemma.

Lemma 3.3 *There is a $\mathbf{\Pi}^1_1$ equivalence relation E such that for each x, its equivalence class $[x]_E$ is a $\mathbf{\Pi}^1_1$-complete set.*

Proof: Let $G \subseteq (\omega^\omega)^2$ be a universal $\mathbf{\Pi}^0_3$ set, specifically

$$G(z, x) \iff (\forall n)(\exists j)(\forall k)[x \notin N(s_{z(\langle n,j,k\rangle)})]$$

where $\langle s_i \rangle_{i\in\omega}$ is an enumeration of $\omega^{<\omega}$, $N(s)$ is the basic open interval determined by s, and $(n, j, k) \mapsto \langle n, j, k\rangle$ is a recursive bijection between ω^3 and ω. We now set xEy if and only if $G_x = G_y$, i.e.

$$xEy \iff (\forall z)[G(x, z) \iff G(y, x)]$$

This is clearly $\mathbf{\Pi}^1_1$. Now let $x_0 \in \omega^\omega$ and a $\mathbf{\Pi}^1_1$ set $A \subseteq \omega^\omega$ be given. We will show that $A \leq_W [x_0]_E$, i.e. A is the continuous preimage of $[x_0]_E$. Fix a tree T to represent A:

$$y \in A \iff \neg(\exists z)[(z, y) \in [T]]$$

It now suffices to find a continuous function f such that, for each y:

$$(\exists z)[(z, y) \in [T]] \Longleftrightarrow G_{f(y)} \neq G_{x_0}$$

Since G_{x_0} is a Borel set, it either contains or is disjoint from a perfect set, and hence there is a continuous embedding of ω^ω into either this set or its complement. We consider the two cases.

If G_{x_0} is disjoint from a perfect set, let $\varphi : s \mapsto t_s$ be a mapping of finite sequences inducing a continuous embedding of ω^ω into $\omega^\omega \setminus G_{x_0}$ in the sense that x is sent to $\varphi(x) = \bigcup_n t_{x \restriction n}$. We can choose this map so that if s_1 and s_2 are incomparable sequences, then so are t_{s_1} and t_{s_2}. Let $\langle G_n \rangle_{n\in\omega}$ be the sequence of F_σ sets coded by x_0, where each $G_n = \bigcup_k G_{n,k}$ with $G_{n,k}$ closed. Thus, $G_{x_0} = \bigcap_n G_n$. We wish to express G_{x_0} as an intersection of nested sets, so we let $\widetilde{G}_n = \bigcap_{k \leq n} G_k$. These are still F_σ, and we can effectively write them in the coding scheme for our universal set (since there are recursive maps which, given indices for two F_σ sets, produce indices for their intersection and union). We still have $G_{x_0} = \bigcap_n \widetilde{G}_n$.

Now for $n \in \omega$ let

$$H_n = \widetilde{G}_n \cup \bigcup \{N(t_s) : |s| = n \ \& \ (s, y \restriction n) \in T\}$$

and let $\langle H_{n,k} \rangle_{k\in\omega}$ enumerate the closed sets necessary to express this as an F_σ code (we can inter-weave the closed sets making up G_n with codes for the $N(t_s)$'s in increasing order of their indexing). We let $f(y)(\langle n, k\rangle) = H_{n,k}$ (i.e. the index for $H_{n,k}$, so that $G_{f(y)} = \bigcap_n H_n$). Note that each $H_n \supseteq \widetilde{G}_n$ so that $G_{f(y)} \supseteq G_{x_0}$. Thus, we need only check that

$$(\exists z)[(z, y) \in [T]] \Longleftrightarrow \exists w \in G_{f(y)} \setminus G_{x_0}$$

If there is such a z, then each H_n contains $N(t_{z \restriction n})$, so that $\varphi(z) \in G_{f(y)}$. We know $\varphi(z) \notin G_{x_0}$, so this suffices.

Conversely, suppose there is a $w \in G_{f(y)} \setminus G_{x_0}$ There must then be some n_0 such that for all $n \geq n_0$ we have $w \notin \widetilde{G}_n$. Thus, for these n there must be some s_n with $(s_n, y \restriction n) \in T$ and $w \in N(t_{s_n})$ (so $w \sqsupseteq t_{s_n}$). Thus, all of the t_{s_n} must be compatible, so we must also have that the s_n's are all comparable. But then we can set $z = \bigcup_n s_n$ and we have that $(z, y) \in [T]$ as desired.

The case where G_{x_0} contains a perfect set is similar, except now we take elements away from G_{x_0} rather than adding them. Let φ be as before, but embedding ω^ω in G_{x_0}, and again let $G_{x_0} = \bigcap_n \widetilde{G}_n$ be given as the intersection of a descending sequence of F_σ sets. Let $H_{0,k}$ be the open set

$$\bigcup\{N(t_s) : |s| = k \;\&\; (s, y \restriction k) \in T\}$$

and let H_0 be the F_σ set

$$H_0 = \omega^\omega \setminus \bigcap_{k\in\omega} H_{0,k}$$

We now let $H_{n+1} = \widetilde{G}_n$ for $n \in \omega$ and let $f(y)$ code the intersection of these sets. Clearly $G_{f(y)} \subseteq G_{x_0}$, so it suffices to check that

$$(\exists z)[(z, y) \in [T]] \iff \exists w \in G_{x_0} \setminus G_{f(y)}$$

If there is such a z, then $\varphi(Z)$ is in each of the $H_{0,k}$'s, and hence not in H_0 so not in $G_{f(y)}$, whereas we know $\varphi(z) \in G_{x_0}$. Conversely, an element of G_{x_0} not in $G_{f(y)}$ must be in the complement of H_0 and hence in each of the $H_{0,k}$'s. As before, this gives us a z such that $(z, y) \in [T]$. □

We can now show that E_B is not universal.

Proposition 3.4 *Let E be the equivalence relation constructed in the previous lemma. Then $E_B <_B E \times E_B$.*

Proof: Clearly $E_B \leq_B E \times E_B$. Suppose there were a Borel function f reducing $E \times E_B$ to E_B. Since for all x and y we have

$$[(x, y)]_{E\times E_B} = [x]_E \times [y]_{E_B}$$

we have that all $E \times E_B$ equivalence classes are also $\mathbf{\Pi}^1_1$-complete. Thus, f must map each $E \times E_B$-class to a $\mathbf{\Pi}^1_1$-complete equivalence class of E_B, since the $E \times E_B$ classes are the inverse images under f of the E_B classes in the range. Thus, f can not map into any of the singleton classes of elements outside of $\mathcal{BC}$, and hence maps into $\mathcal{BC}$. Since $f[(\omega^\omega)^2]$ is now a $\mathbf{\Sigma}^1_1$ subset of $\mathcal{BC}$, we must have that $f[(\omega^\omega)^2] \subseteq \mathcal{BC}_\alpha$ for some $\alpha < \omega_1$ by boundedness (since the map sending a Borel code to its Borel rank is a $\mathbf{\Pi}^1_1$-rank on $\mathcal{BC}$). But then f would in fact witness that $E \times E_B \leq_B E_\alpha$ and hence we would have $E_B \leq_B E_\alpha$, a contradiction.

□

The proof actually gives us a more general theorem, that there is a type of "jump" for certain $\mathbf{\Pi}^1_1$ equivalence relations.

Theorem 3.5 *Let E be a $\mathbf{\Pi}^1_1$ equivalence relation none of whose classes is Borel, and let F be a $\mathbf{\Pi}^1_1$ equivalence relation such that F is not Borel reducible to $F \upharpoonright A$ for any $\mathbf{\Sigma}^1_1$ set $A \subseteq \{x : [x]_F \text{ is not Borel}\}$. Then we have that $F <_B E \times F$.*

We have seen that E_B is not a universal $\mathbf{\Pi}^1_1$ equivalence relation. In the next section we will consider whether it can be minimal among those $\mathbf{\Pi}^1_1$ equivalence relations which reduce all Borel equivalence relations.

3.2 Minimal $\mathbf{\Pi}^1_1$ Relations Above All Borel Ones

Equality of Borel sets is a canonical example of an equivalence relation to which we can reduce all Borel equivalence relations. We wish to analyze what equivalence relations of this form can look like, and see if this property applies any sort of additional complexity. The type of additional complexity we are seeking is some single complicated equivalence relation E such that, for all equivalence relations F of some given type, if every Borel equivalence relation is reducible to F, then E must also be reducible to F. An appropriate standard for complexity of E is that every Borel equivalence relation should be reducible to it. Note that, by a result of Friedman and Stanley (see [9]), no Borel equivalence relation can be universal for all Borel equivalence relations, so we should look to more complicated relations. This suggests we investigate the following notion.

Definition 3.6 *Given a collection $\mathcal{F}$ of equivalence relations, we say that an equivalence relation E is minimum above Borel for $\mathcal{F}$ if $E \in \mathcal{F}$, every Borel equivalence relation is reducible to E, and for every $F \in \mathcal{F}$ to which we can reduce all Borel equivalence relations we have $E \leq_B F$. We say that E is minimal above Borel for $\mathcal{F}$ if $E \in \mathcal{F}$, every Borel equivalence relation is reducible to E, and there is no equivalence relation $F \in \mathcal{F}$ with $F <_B E$ such that every Borel equivalence relation is reducible to F.*

We wish to investigate the possibility of there being minimum or minimal equivalence relations above Borel for classes of definable equivalence relations. In this section we consider the class of co-analytic relations and consider the class of analytic ones in the next section. We will first show that E_B is not minimum above Borel for the class of $\mathbf{\Pi}^1_1$ equivalence relations. We will in fact present two $\mathbf{\Pi}^1_1$ relations which reduce all Borel equivalence relations and to which we can not reduce E_B.

Let $D \subseteq \omega^\omega$ and $P, S \subseteq (\omega^\omega)^3$ be a parameterization of the Borel subsets of $(\omega^\omega)^2$, i.e. D is $\mathbf{\Pi}^1_1$, P is $\mathbf{\Pi}^1_1$, and S is $\mathbf{\Sigma}^1_1$ such that:

1. $d \in D \Longrightarrow P_d = S_d$ (and hence these are Borel subsets of $(\omega^\omega)^2$)

2. $\{P_d : d \in D\}$ contains all Borel subsets of $(2^\omega)^2$

We now define $B \subseteq D$ by

$$\begin{aligned} B(z) &\iff P_z = S_z \text{ is an equivalence relation} \\ &\iff (\forall x)(\forall y)(\forall w)[P_z(x,x) \ \& \ (S_z(x,y) \Longrightarrow P(y,z)) \ \& \\ &\qquad ((S_z(x,y) \ \& \ S_z(y,w)) \Longrightarrow P_z(x,w))] \end{aligned}$$

Thus, B (together with P and S) is a parameterization of all Borel equivalence relations. It is evidently $\mathbf{\Pi}^1_1$.

Lemma 3.7 *The set B is $\mathbf{\Pi}^1_1$-complete.*

Proof: Let A be a $\mathbf{\Pi}^1_1$ set, given by

$$A(x) \iff (\forall y)(\exists n)[(x \restriction n, y \restriction n) \notin T]$$

where T is a tree on ω. For a given x, we let $f(x)$ code the following binary relation R_x (as a Borel subset of $(\omega^\omega)^2$); we can easily produce a parameter for this set from x.

1. For $y_1, y_2 \neq x$, let $y_1 \ R_x \ y_2$.

2. For all y, let $x \ R_x \ y$.

3. For all y, set $y \ R_x \ x \iff (x,y) \notin [T]$.

Nothing else is related. A branch through T will prevent R_x from being symmetric; otherwise it is an equivalence relation. □

We now introduce our first relation.

Definition 3.8 *Let E_P be the $\mathbf{\Pi}^1_1$ equivalence relation defined on $(\omega^\omega)^2$ by*

$$(z_1, x_1) \ E_P \ (z_2, x_2) \iff z_1 = z_2 \ \& \ [x_1 = x_2 \ \vee \ (B(z_1) \ \& \ P(z_1, x_1, x_2))]$$

Thus, each z-slice of E_P either is the identity relation (if $z \notin B$), or it is the Borel equivalence relation parametrized by z for $z \in B$. Thus, it is immediate that every Borel equivalence relation is Borel reducible to E_P by mapping into the z-slice for some z parameterizing it. We also note that every equivalence class of E_P is Borel, so this is a relatively simple non-Borel $\mathbf{\Pi}^1_1$ equivalence relation in a descriptive sense. The following proposition is thus immediate:

Proposition 3.9 *If E is an equivalence relation such that $E \leq_B E_P$, then every equivalence class of E is Borel.*

Since E_B contains equivalence classes which are not Borel (such as the equivalence class of a code for the empty set), we have that E_B is not reducible to E_P. Thus:

Corollary 3.10 *The relation E_B is not minimum above Borel for the class of $\mathbf{\Pi}^1_1$ equivalence relations.*

Let us say a bit more about the relation E_P. Fix a $\mathbf{\Pi}^1_1$-norm ρ on B, and set

$$B_\alpha = \{x : x \in B \ \& \ \rho(x) < \alpha\}$$

Thus, each B_α is Borel, $B_\alpha \subseteq B_\beta$ for $\alpha < \beta$, and $B = \bigcup_{\alpha<\omega_1} B_\alpha$. For $\alpha < \omega_1$, define the following relation on $(\omega^\omega)^2$:

$$\begin{aligned} (z_1, x_1) \ E^\alpha_P \ (z_2, x_2) &\iff z_1 = z_2 \ \& \ [x_1 = x_2 \ \vee \ (B_\alpha(z_1) \ \& \ P(z_1, x_1, x_2))] \\ &\iff z_1 = z_2 \ \& \ [x_1 = x_2 \ \vee \ (B_\alpha(z_1) \ \& \ S(z_1, x_1, x_2))] \end{aligned}$$

Thus each E^α_P is a Borel equivalence relation, containing as slices all Borel equivalence relations with codes of rank less than α. We list several properties of these:

1. For $\alpha < \beta$, we have $E^\alpha_P \leq_B E^\beta_P <_B E_P$.
2. For $\alpha < \beta$, we have $E^\alpha_P \subseteq E^\beta_P \subseteq E_P$, and $E_P = \bigcup_{\alpha<\omega_1} E^\alpha_P$.
3. Every Borel equivalence relation is reducible to some E^α_P, and so, for an equivalence relation F we have that every Borel equivalence relation is Borel reducible to F if and only if every E^α_P is Borel reducible to F.

E_P seems to be a good candidate for a minimal above Borel $\mathbf{\Pi}^1_1$ equivalence relation, but we do not know whether this is true. As some evidence, we can see that $E_P \leq_B E_B$:

Proposition 3.11 *The relation E_P is Borel-reducible to E_B.*

Proof: Since the set B of codes for Borel equivalence relations is $\mathbf{\Pi}^1_1$ and the set $\mathcal{BC}$ of Borel codes is easily $\mathbf{\Pi}^1_1$-complete, we may fix a continuous function g such that for all z, $B(z) \iff g(z) \in \mathcal{BC}$. For $z \in B$, we have that the sets P_z and S_z are equal (hence $\mathbf{\Delta}^1_1$). We can thus apply the effective separation theorem (see [26]) to produce a Borel code for this set, the equivalence relation E_z. Both here and in the case of the function g we have a total function which produces a Borel code when it should (and produces garbage otherwise). So we will treat these objects as if they were Borel codes.

Now, given a pair (z, x) we can produce the Borel code $g(z)$, a Borel code for the singleton $\{z\}$, and a Borel code for the set $[x]_{E_z}$ (which we get from the Borel code for E_z obtained above). From these, we can form the Borel code for the following set:

$$\{0 \frown z\} \cup (1 \frown B(g(z))) \cup (2 \frown [x]_{E_z})$$

When z codes a Borel equivalence relation this procedure produces a legitimate Borel code, and when z does not code a Borel equivalence relation it will produce something which is not a Borel code. Call the function producing this putative code f. It is straightforward to check that for two of these objects $f(z_1, x_1)$ and $f(z_2, x_2)$ to be equivalent under E_B we must have had either that both z_1 and z_2 were not in B, or that $z_1 = z_2$ and x_1 is E_{z_1}-equivalent to x_2. Hence, f is a reduction from E_P to E_B. $\square$

Since we now have $E_P <_B E_B$, the following is immediate:

Corollary 3.12 *The relation E_B is not minimal above Borel for the class of bP^1_1 equivalence relations.*

We will say a bit more about E_P later. We now give a second example of a $\mathbf{\Pi}^1_1$ equivalence relation above all Borel relations to which we cannot reduce E_B. Recall that $\mathcal{WO}$ is the collection of codes for countable well-orders.

Definition 3.13 *Let E^*_B be the relation defined on $(\omega^\omega)^2$ by*

$$(x, \alpha)\ E^*_B\ (y, \beta) \iff (x = y \ \&\ \alpha = \beta) \ \vee\ (\alpha, \beta \in \mathcal{WO} \ \&\ |\alpha| = |\beta| \ \&\ x, y \in \mathcal{BC}_{|\alpha|} \ \&\ B(x) = B(y))$$

Thus, E^*_B is a spreading-out of E_B, demanding not only that two codes produce the same set, but also that they be of the same given rank. This is again a $\mathbf{\Pi}^1_1$ relation. Every Borel equivalence relation E is reducible to it because, having fixed a Borel code for E of rank α we can produce codes for its equivalence classed all of rank α. We also have that each of the relations E_α is reducible to E^*_B, since we can map a code in $\mathcal{BC}_\alpha$ to the pair $(x, \widetilde{\alpha})$, where $\widetilde{\alpha} \in \mathcal{WO}$ is some fixed representation of α.

Proposition 3.14 *The relation E_B is not Borel reducible to E^*_B.*

Proof: Suppose $\varphi : x \mapsto (x^*, \alpha_x)$ were such a reduction. Then we have that if $x \in \mathcal{BC}$ then $x^* \in \mathcal{BC}_{|\alpha_x|}$, and if $B(x) = B(y)$ then $|\alpha_x| = |\alpha_y|$. For $\gamma < \omega_1$, we have that $\{\alpha_x : x \in \mathcal{BC}_\gamma\}$ is thus a $\mathbf{\Sigma}^1_1$ subset of $\mathcal{WO}$, so there is some countable ordinal ξ_γ such that

$$x \in \mathcal{BC}_\gamma \Longrightarrow |\alpha_x| < \xi_\gamma$$

We choose ξ_γ to be the least possible such that $\xi_\gamma > \gamma$. Now let

$$\begin{aligned} \gamma_0 &= 1 \\ \gamma_{n+1} &= \xi_{\gamma_n} \text{ for } n \in \omega \\ \gamma_\infty &= \sup_{n \in \omega} \gamma_n \end{aligned}$$

Since the sequence of γ_n's is increasing, γ_∞ will be a limit ordinal, and if $x \in \mathcal{BC}_\delta$ for some $\delta < \gamma_\infty$, we will have $|\alpha_x| < \gamma_\infty$.

Now consider $x \in \mathcal{BC}_{(\gamma_\infty+1)}$. If $B(x) \in \mathbf{\Pi}^0_\delta$ for some $\delta < \gamma_\infty$, then $B(x)$ has a code $x' \in \mathcal{BC}_\delta$, so we have $|\alpha_{x'}| = |\alpha_x| < \gamma_\infty$. Conversely, if $|\alpha_x| < \gamma_\infty$, then $B(x) \in \mathbf{\Pi}^0_\delta$ for some $\delta < \gamma_\infty$. Hence, for $x \in \mathcal{BC}_{(\gamma_\infty+1)}$ we have

$$B(x) \text{ is true } \mathbf{\Pi}^0_{\gamma_\infty} \Longleftrightarrow |\alpha_x| \geq \gamma_\infty$$

This would imply that the set

$$C = \{x \in \mathcal{BC}_{(\gamma_\infty+1)} : B(x) \text{ is true } \mathbf{\Pi}^0_{\gamma_\infty}\}$$

is a Borel set. However, we claim that C is $\mathbf{\Sigma}^1_1$-hard.

Let A be a true $\mathbf{\Pi}^0_{\gamma_\infty}$ set, and fix a Borel code in $\mathcal{BC}_{(\gamma_\infty+1)}$ for the set $A \times \omega^\omega$ (which is also true $\mathbf{\Pi}^0_{\gamma_\infty}$). Here we are identifying $(\omega^\omega)^2$ with ω^ω by some fixed homeomorphism, so that basic neighborhoods referred to below will really be neighborhoods in $(\omega^\omega)^2$. We will describe a continuous reduction of the $\mathbf{\Sigma}^1_1$-complete set of ill-founded trees to the set C.

Given a tree T, we will map T to a Borel code x_T (in $\mathcal{BC}_{(\gamma_\infty+1)}$) for the set $A \times [T]$. We already have a Borel code for the set $A \times \omega^\omega$, so we need only find a code for the intersection of this set with $\omega^\omega \times [T]$. We can easily produce a code for the latter set (of rank less than γ_∞) in a continuous fashion; we now simply hang a copy of this code below a new node added in the first level of the tree of the original Borel code for $A \times \omega^\omega$. This new code will then be a Borel code for the intersection of these two sets. Let x_T be this code. Then, we will have that the set coded by x_T is empty if T is well-founded, and hence not true $\mathbf{\Pi}^0_{\gamma_\infty}$. On the other hand, if T is ill-founded then [T] will be a non-empty closed set, so clearly $A \leq_W A \times [T]$. We thus have that T is ill-founded if and only if $x_T \in C$. $\square$

So we know that E_B is not minimum or even minimal above Borel for $\mathbf{\Pi}^1_1$ equivalence relations. We do not know if there can be any minimal such equivalence relations, though E_P or E_B^* seem reasonable candidates.

Question 3.15 *Is there a $\mathbf{\Pi}^1_1$ equivalence relation which is minimum or minimal above Borel for the set of $\mathbf{\Pi}^1_1$ equivalence relations? What if we allow more general reductions, such as $\sigma(\mathbf{\Sigma}^1_1)$-measurable ones?*

Some simple test questions for E_P and E_B^* are:

Question 3.16 *Is $E_B^* \leq_B E_B$ or $E_B^* \leq_{\sigma(\mathbf{\Sigma}^1_1)} E_B$? Is $E_P \leq_B E_B^*$? Is $E_B^* \leq_B E_P$?*

We also do not know if E_B can be minimal under more general reductions. Here good test questions are the following:

Question 3.17 *Is $E_B \leq_{\sigma(\mathbf{\Sigma}^1_1)} E_P$? Is $E_B \leq_{\sigma(\mathbf{\Sigma}^1_1)} E_B^*$?*

A question along similar lines to minimality above Borel is this:

Question 3.18 *Is there a minimum $\mathbf{\Pi}^1_1$ equivalence relation above all of the E_α's?*

3.3 The $\mathbf{\Sigma}^1_1$ Case

Let us now consider the analogous questions for $\mathbf{\Sigma}^1_1$ equivalence relations. Here we can show that there is no $\mathbf{\Sigma}^1_1$ relation which is a minimum above Borel for the class of $\mathbf{\Sigma}^1_1$ equivalence relations. Recall the set B introduced earlier parameterizing Borel equivalence relations.

Definition 3.19 *Let E_S and E'_S be the equivalence relations defined on $(\omega^\omega)^2$ by*

$$(z_1, x_1)\ E_S\ (z_2, x_2) \iff z_1 = z_2\ \&\ [\neg B(z_1)\ \vee\ S(z_1, x_1, x_2)]$$
$$(z_1, x_1)\ E'_S\ (z_2, x_2) \iff [\neg B(z_1)\ \&\ \neg B(z_2)]\ \vee\ [z_1 = z_2\ \&\ S(z_1, x_1, x_2)]$$

These are both $\mathbf{\Sigma}^1_1$ equivalence relations. They are similar to E_P except that, instead of making slices outside of B into the identity relation, in the case of E_S we make each slice into a single equivalence class and in the case of E'_S we lump all of these slices together into a single class. We thus have $E_p \subseteq E_S \subseteq E'_S$. Again, we see that every Borel equivalence relation is Borel reducible to both E_S and E'_S. Also note that we can not have any Borel equivalence relation E with $E_P \subseteq E \subseteq E'_S$, since it would also have to be universal for all Borel equivalence relations. Every equivalence class of E_S is a Borel set, whereas E'_S contains one $\mathbf{\Sigma}^1_1$-complete class.

As we did with E_P, we can define E^α_S for $\alpha < \omega_1$ by

$$\begin{aligned}(z_1, x_1)\ E^\alpha_S\ (Z_2, x_2) &\iff z_1 = z_2\ \&\ [\neg B_\alpha(z_1)\ \vee\ P(z_1, x_1, x_2)]\\ &\iff z_1 = z_2\ \&\ [\neg B_\alpha(z_1)\ \vee\ S(z_1, x_1, x_2)]\end{aligned}$$

so that each E^α_S is a Borel equivalence relation, $E^\alpha_S \supseteq E^\beta_S$ for $\alpha < \beta < \omega_1$, and $E_S = \bigcap_{\alpha<\omega_1} E^\alpha_S$. Note that again every Borel equivalence relation is Borel reducible to some E^S_α. In fact, we have that for each α, E^α_S is Borel bireducible with E^α_P. We can also express E'_S in a similar manner.

Our main result in this section is the following:

Proposition 3.20 *For any equivalence relation E on a Polish space, E is Borel if and only if $E \leq_B E_S$ and $E \leq_B E'_S$.*

Proof: One direction is immediate. For the other, let f be a Borel function reducing E to E'_S. Since E'_S has only one non-Borel equivalence class, at most one E class can be mapped to it via f. Let A be the inverse image of the complement of this class under f, so A is E-invariant. Since E is reducible to E_S, all of its classes must be Borel sets, and since A omits at most one equivalence class, A must be Borel. Therefore, $f[A]$ is the Borel image of a Borel set and hence $\mathbf{\Sigma}^1_1$. Letting π_0 be projection onto the first coordinate, we have that $\pi_0 \circ f[A]$ is also $\mathbf{\Sigma}^1_1$. Since this is a subset of B, we must in fact have that $\pi_0 \circ f[A] \subseteq B_\alpha$ for some α. Thus, we in fact have that $E \restriction A \leq_B E'_S \restriction (B_\alpha \times \omega^\omega)$. Since

$E'_S \restriction (B_\alpha \times \omega^\omega) = E_S^\alpha \restriction (B_\alpha \times \omega^\omega)$, we actually have that $E \restriction A \leq_B E_S^\alpha$. This is a Borel equivalence relation, so $E \restriction A$ must be also. Then E itself must be Borel. □

Since each of these equivalence relations is reducible to itself, and neither is Borel, we then have:

Corollary 3.21 *The relations E_S and E'_S are $\leq_B$-incomparable.*

We also immediately have the promised:

Theorem 3.22 *There is no $\mathbf{\Sigma}^1_1$ equivalence relation which is a minimum above Borel for $\mathbf{\Sigma}^1_1$ equivalence relations.*

We do not know, however, if there can be minimal such relations.

Question 3.23 *Is there a $\mathbf{\Sigma}^1_1$ equivalence relation which is minimal above Borel for the class of $\mathbf{\Sigma}^1_1$ equivalence relations? What if we allow more general reductions?*

Although we have seen that E_S and E'_S are $\leq_B$-incomparable, they are in fact bireducible if we allow more general reductions. In fact:

Proposition 3.24 $E_s \sim_{\sigma(\mathbf{\Sigma}^1_1)} E'_S \sim_{\sigma(\mathbf{\Sigma}^1_1)} E_P$

Proof: First, to see that $E'_S \leq_{\sigma(\mathbf{\Sigma}^1_1)} E_S$, let z_0 be a point not in B. Now define the reduction f_1 by

$$f_1(z,x) = \begin{cases} (z_0, x) & \text{if } z \notin B \\ (z, x) & \text{if } z \in B \end{cases}$$

Next, to see that $E_S \leq_{\sigma(\mathbf{\Sigma}^1_1)} E_P$, define f_2 by

$$f_2(z,x) = \begin{cases} (z, z) & \text{if } z \notin B \\ (z, x) & \text{if } z \in B \end{cases}$$

Finally, to see $E_P \leq_{\sigma(\mathbf{\Sigma}^1_1)} E'_S$, fix a $z_1 \in B$ such that $P_{z_1} = S_{z_1} = \Delta(\omega^\omega)$, the identity relation. Now let f_3 be given by

$$f_3(z,x) = \begin{cases} (z_1, 0 \frown x) & \text{if } z \notin B \\ (z_1, 1 \frown x) & \text{if } z = z_1 \\ (z, x) & \text{if } z \in B \text{ and } z \neq z_1 \end{cases}$$

It is easy to verify that these functions are reductions and are $\sigma(\mathbf{\Sigma}^1_1)$-measurable. Compositions give the rest of the needed reductions (and we note that these compositions will be $\sigma(\mathbf{\Sigma}_1^1)$-measurable). Note that although the given functions are not injective, one can modify them to make them embeddings. □

The case of E_S and E'_S presents a strong argument that we should allow more general reductions at this level than Borel-measurable ones. These two relations are essentially the same, they are bireducible under $\sigma(\mathbf{\Sigma}_1^1)$-measurable reductions, and yet they are $\leq_B$-incomparable.

Chapter 4

Distances in Polish Metric Spaces

In this and the next two chapters we consider descriptive aspects of Polish metric spaces. By a *Polish metric space* we mean a pair (X, d), where X is a Polish space (a separable, completely-metrizable space) and d is a complete, compatible metric for X. We will consider two aspects. First, we will characterize which sets of reals can be the *set of distances* in a Polish metric space. In the following chapters, we will consider the complexity of the equivalence relation of *isometry* of Polish metric spaces. In both instances we will also obtain results about specific classes of metric spaces, e.g. compact, locally compact, zero-dimensional. In the last section of this chapter we will briefly consider the question of which sets of triangles (and larger configurations) can occur in a Polish metric space. Our interest in distance sets is related to the isometry question, since distance sets form an isometry invariant, although generally not a complete invariant.

4.1 The Set of Distances

Given a metric space (X, d), we will define its set of distances, Dist(X,d):

Definition 4.1 *Let Dist*$(X, d) = \{d(x, y) : x, y \in X\}$.

Distance sets have been studied in several contexts. Much of the work has been on the distance sets of subsets of the spaces $\mathbb{Z}$, $\mathbb{R}$, or $\mathbb{R}^n$, with the usual metrics. In the case of subsets of $\mathbb{Z}$ or $\mathbb{N}$ these are often called difference sets. One of the earliest results about them was Steinhaus's theorem ([31]) that the difference set of a subset of $\mathbb{R}$ of positive measure contains a (right-) neighborhood of 0. Sierpiński ([29]) showed that difference sets could be more complicated than the set itself by producing a G_δ subset of $\mathbb{R}$ whose difference set is $\mathbf{\Sigma}^1_1$-complete.

There are also several results characterizing which sets can be the set of distances in some metric space. Several results characterize which sets can be distance sets for subsets of $\mathbb{R}^n$ (see for instance [23]). More generally, in [24] the authors characterize which sets can be the set of distances of some separable metric space:

Theorem 4.2 (Kelly and Nordhaus) *A non-negative set of reals (containing 0) is the set of distances of some separable metric space if and only if it is either countable or has 0 as a limit point.*

Here we will be concerned with with the case of *Polish* metric spaces so we will need to produce complete metrics. As a result, our arguments will be significantly different.

Two properties of the distance set of a Polish metric space are clear. First, since d is a continuous map from the Polish space X^2 to $\mathbb{R}$, $\mathrm{Dist}(X,d)$ is an analytic set (of non-negative reals, containing 0). Second, if $\mathrm{Dist}(X,d)$ is uncountable, then X must also be uncountable, and so contains a perfect subset. Thus, if $\mathrm{Dist}(X,d)$ is uncountable, it must contain distances arbitrarily close to 0.

It turns out that these two conditions are sufficient for a non-negative set of reals to be the set of distances for some Polish metric space (X,d). In the next section we will consider some special classes of metric spaces which have more restrictive conditions on their sets of distances. Here we prove:

Theorem 4.3 *Given a set of non-negative reals, $A \subseteq [0,\infty)$, A is the set of distances for some Polish metric space if and only if A is an analytic set containing 0 and either A is countable or 0 is a limit point of A.*

Proof: First we consider the case where A is countable. If $A = \{0\}$ then we can take X to be the one-point space; otherwise, let $A \setminus \{0\} = \{a_i\}_{i\in\omega}$ (if A is finite we allow the a_i's to repeat). We will construct (X,d) so that $\mathrm{Dist}(X,d) = A$. Our space X will have as underlying set $\{x^*\} \cup \{x_i : i \in \omega\}$. We let:

$$\begin{aligned} d(x^*, x_i) &= a_i \\ d(x_i, x_j) &= \max(a_i, a_j) \text{ for } i \neq j \end{aligned}$$

We see that d is a metric (in fact an ultrametric), since the two largest sides in any triangle will have equal length. We can see that d is complete by noting that if $\langle x_{i_n}\rangle_{n\in\omega}$ is a Cauchy sequence (we may assume x^* does not occur) then $\lim a_{i_n} = 0$, so that the sequence converges to x^*. Since X is countable, it is separable, and so (X, d) as constructed is a Polish metric space, and it is clear that $\mathrm{Dist}(X, d) = \{0\} \cup \{a_i : i \in \omega\} = A$.

The more difficult case

Now we consider the case that A is analytic with 0 a limit point of A. We will first assume that $A \subseteq [0, 1)$ to simplify the proof, and handle the general case at the end. We may identify sequences from $\{0, 1\}$ with reals in $[0, 1]$ via the map

$$\alpha \mapsto \sum_{i\in\omega} \frac{\alpha(i)}{2^{i+1}}$$

This is a continuous map, so the pull-back of an analytic set of reals is an analytic subset of the Cantor space 2^ω. This is a one-to-one map except for those points with eventually constant binary representations. By removing the eventually 1 sequences, we may assume each point is uniquely represented (although this is not necessary); doing so will not affect the complexity of the set. We may thus represent A as an analytic subset of the Cantor space which we will identify with A. We will thus consider elements of A both as real numbers and as binary sequences in the sequel. We can then express A as the projection of a closed subset of $\mathcal{C} \times \mathcal{N}$ (where $\mathcal{N}$ is the Baire space ω^ω). We may thus represent A as the projection of a tree on $2 \times \omega$, which we may take to be pruned, i.e. without any terminal nodes. So, let T be a pruned tree on $2 \times \omega$ such that

$$\alpha \in A \Longleftrightarrow (\exists\beta)(\forall n)[(\alpha \restriction n, \beta \restriction n) \in T]$$

Now let

$$T^* = \{s \in 2^{<\omega} : (\exists a \in \omega^{<\omega})[(s, a) \in T]\}$$

Then T^* is a pruned tree on 2 with $[T^*] = \bar{A}$, the closure of A (in 2^ω). For each $s \in T^*$, we can thus pick $d_s \in A$ with $s \sqsubset d_s$, subject to the requirement that if $s = 0^k$ for some $k \in \omega$, then $2d_s \geq d_t$ for any $t \sqsupset s$. This can be achieved by having d_s "go right as soon as possible," i.e., if j is least such that there is a $t \in T^*$ with $t \sqsupset s$ and $t(j) = 1$, then require $d_s(j) = 1$. Note that this will ensure the following condition for all d_s's:

$$\text{If } s, t \in T^* \text{ and } s \sqsubseteq t \text{ then } 2d_s \geq d_t \qquad (*)$$

Choose finally a decreasing sequence $\langle \epsilon_i \rangle_{i \in \omega}$ with $\epsilon_i \in A$ an $\epsilon_i < \frac{1}{2^i}$, which is possible by the hypotheses on A. These d_s's and ϵ_i's will be used in constructing our space (X, d).

We will define our space (X, d) by first defining d on a countable set $\mathcal{D}$ and then taking (X, d) to be the completion of $(\mathcal{D}, d)$ (so that $\mathcal{D}$ will be a countable dense subset of X). Let $\mathcal{D}$ have underlying set:

$$\mathcal{D} = \{x^*\} \cup \{x_{s,a} : (s, a) \in T\}$$

We will construct d so that Cauchy sequences from $\mathcal{D}$ correspond to branches through T, and so that the associated limit points will have their distances from x^* being the associated values in $A = p[T]$.

Defining d on $\mathcal{D}$

First, for $(s, a) \in T$, let $d(x^*, x_{s,a}) = d_s$. To define $d(x_{s,a}, x_{t,b})$ for $(s, a), (t, b) \in T$, we consider three cases:

case 1: (t, b) is an extension of (s, a) in T ($(t, b) \sqsupset (s, a)$)

(a) If $s = 0^k$ for some k, let $d(x_{s,a}, x_{t,b}) = d_s$.

(b) If $s \neq 0^k$ and there is an ϵ_i with $\frac{1}{2^{|s|}} \leq \epsilon_i < d_s$, then let

$$d(x_{s,a}, x_{t,b}) = \text{ the least such } \epsilon_i$$

(c) Otherwise let $d(x_{s,a}, x_{t,b}) = d_s$.

case 2: (s, a) is an extension of (t, b). Same as case 1, interchanging roles.

case 3: (s, a) and (t, b) are incomparable in T.

Then let (r, c) be the maximal T-predecessor of (s, a) and (t, b).

(a) If $r = 0^k$ for some k, then let $d(x_{s,a}, x_{t,b}) = \max(d_s, d_t)$.

(b) If $r \neq 0^k$ then let $d(x_{s,a}, x_{t,b}) = \max(d(x_{r,c}, x_{s,a}), d(x_{r,c}, x_{t,b}))$.

This inductively defines the metric on d. The construction is motivated in part by considering metrics on the Cantor space of the following form: Let $\langle d_n \rangle_{n \in \omega}$ be a sequence of positive reals and set $d(x, y) = d_{n(x,y)}$, where $n(x, y)$ is the least n such that $x(n) \neq y(n)$, for $x \neq y$. Then d will give a separable complete metric on $\mathcal{C}$ precisely when $\langle d_n \rangle \to 0$ and $d_m \leq 2d_n$ whenever $n \leq m$. The metric will be somewhat like an ultrametric: two sides of a triangle will be equal, although the third may be longer. The present construction attempts to mimic this as much as possible (but with distances varying from branch to branch) since the "ultrametric-ness" will prevent unwanted Cauchy sequences.

Note that all distances between points in $\mathcal{D}$ are elements of A.

Verifying d is a metric on $\mathcal{D}$

We need only check that the triangle inequality holds for each trio of points in $\mathcal{D}$. Note that if the three sides have lengths d_1, d_2, and d_3, it suffices to check that:

$$|d_1 - d_2| \leq d_3 \leq d_1 + d_2$$

(this will frequently occur in the form: $d_1 = d_2$, $d_3 \leq 2d_1$). Also notice that, from the definition of d, whenever $s \neq 0^k$ and $(t, b) \sqsupset (s, a)$ we have

$$|d_t - d_s| \leq \frac{1}{2^{|s|}} \leq d(x_{s,a}, x_{t,b}) \leq d_s \tag{**}$$

We have two types of triangles to consider: those including the point x^* and those without it.

A. Triangles including x^*

Consider a triangle with vertices x^*, $x_{s,a}$, and $x_{t,b}$. Set $\delta = d(x_{s,a}, x_{t,b})$ and recall that $d(x^*, x_{s,a}) = d_s$ and $d(x^*, x_{t,b}) = d_t$. We have two cases:

1. $(s, a) \sqsubset (t, b)$ (or vice versa)

   (a) If $s = 0^k$ then $\delta = d_s$, and $d_t \leq 2d_s$ by our choice of the d_s's, so this triangle is legal as noted above.

   (b) If $s \neq 0^k$ then $|d_t - d_s| \leq \delta \leq d_s \leq d_s + d_t$ by (**), so this is ok.

2. $(s, a) \perp (t, b)$

   Let (r, c) be the maximal predecessor of (s, a) and (t, b).

   (a) If $r = 0^k$ then $\delta = \max(d_s, d_t)$ so this is ok.

   (b) If $r \neq 0^k$ then $\delta = \max(d(x_{r,c}, x_{s,a}), d(x_{r,c}, x_{t,b})) \geq \frac{1}{2^{|r|}}$, since both of these distances are at least $\frac{1}{2^{|r|}}$. They are both also at most d_r, so we have $\delta \leq d_r$. Since s and t both extend r we have $|d_s - d_t| \leq \frac{1}{2^{|r|}}$, and $|d_r - d_s| \leq \frac{1}{2^{|r|}} \leq d_t$ so that $d_r \leq d_s + d_t$. Putting this all together we have that $|d_s - d_t| \leq \delta \leq d_s + d_t$, so this type of triangle is ok.

B. Triangles without x^*

Let (r,a), (s,b) and (t,c) in T be given. We may assume $|r| \leq |s| \leq |t|$. The following lemma will be useful:

Lemma 4.4 *If $(r,a) \sqsubset (s,b) \sqsubset (t,c)$ in T, then $d(x_{s,b}, x_{t,c}) \leq 2d(x_{r,a}, x_{s,b})$.*

Proof of lemma:. There are three cases:

1. If $r = 0^k$ then $d(x_{r,a}, x_{s,b}) = d_r$ and $d(x_{s,b}, x_{t,c}) = d_s$, and $d_s \leq 2d_r$ by the construction of the d_s's, since $s \sqsupset r$.

2. If $r \neq 0^k$ and $d(x_{r,a}, x_{s,b}) = d_r$ then, since $d(x_{s,b}, x_{t,c}) \leq d_s$ and $d_s \leq 2d_r$, we have $d(x_{s,b}, x_{t,c}) \leq 2d(x_{r,a}, x_{s,b})$.

3. Otherwise $r \neq 0^k$ and $d(x_{r,a}, x_{s,b}) = \epsilon_i$ for one of the ϵ_i with $\frac{1}{2^{|r|}} \leq \epsilon_i < d_r$. Note that if $d_s \geq \epsilon_i$ then $\frac{1}{2^{|t|}} \leq \epsilon_i \leq d_s$, so $d(x_{s,b}, x_{t,c})$ will equal either ϵ_i or some $\epsilon_j < \epsilon_i$, in which case $d(x_{s,b}, x_{t,c}) \leq d(x_{r,a}, x_{s,b})$. Otherwise $d_s < \epsilon_i$, and so $d(x_{s,b}, x_{t,c}) \leq d_s < \epsilon_i = d(x_{r,a}, x_{s,b})$. □(lemma)

We will now consider six cases based on possible relations of the three sequences (note that the two ostensible cases: $(r,a) \perp (s,b)$, $(r,a) \sqsubset (t,b)$, $(s,b) \sqsubset (t,c)$; and $(r,a) \sqsubset (s,b)$, $(s,b) \sqsubset (t,c)$, $(r,a) \perp (t,c)$ are both impossible).

1. $(r,a) \perp (s,b), (s,b) \perp (t,c), (r,a) \perp (t,c)$

   Either all three sequences differ from each other for the first time at the same level, or one pair agrees for a longer initial segment than the third:

   (a) If they differ at the same level, let (q,d) be the mutual maximal predecessor of the three sequences.

      i. If $q = 0^k$ for some k, then

$$\begin{aligned} d(x_{r,a}, x_{s,b}) &= \max(d_r, d_s) \\ d(x_{s,b}, x_{t,c}) &= \max(d_s, d_t) \\ d(x_{r,a}, x_{t,c}) &= \max(d_r, d_t) \end{aligned}$$

      Thus we have an ultrametric triangle, so this case is ok.

ii. If $q \neq 0^k$ then

$$\begin{aligned} d(x_{r,a}, x_{s,b}) &= \max(d(x_{q,d}, x_{r,a}), d(x_{q,d}, x_{s,b})) \\ d(x_{s,b}, x_{t,c}) &= \max(d(x_{q,d}, x_{s,b}), d(x_{q,d}, x_{t,c})) \\ d(x_{r,a}, x_{t,c}) &= \max(d(x_{q,d}, x_{r,a}), d(x_{q,d}, x_{t,c})) \end{aligned}$$

So again this is ultrametric and hence legal.

(b) Otherwise one pair agrees longer; we may assume that (r, a) and (s, b) do so, since the relative lengths will not matter in this case. Let (p, e) be the maximal predecessor of (r, a) and (s, b), and let (q, d) be the maximal predecessor of (p, e) and (t, c). Note that $(q, d) \sqsubset (p, e)$.

i. If $q = 0^j$ and $p = 0^k$ (with $j < k$), then we have:

$$\begin{aligned} d(x_{r,a}, x_{s,b}) &= \max(d_r, d_s) \\ d(x_{r,a}, x_{t,c}) &= \max(d_r, d_t) \\ d(x_{s,b}, x_{t,c}) &= \max(d_s, d_t) \end{aligned}$$

So this is fine.

ii. If $q = 0^j$ but $p \neq 0^k$ then:

$$\begin{aligned} d(x_{r,a}, x_{s,b}) &= \max(d(x_{p,e}, x_{r,a}), d(x_{p,e}, x_{s,b})) \\ d(x_{s,b}, x_{t,c}) &= \max(d_s, d_t) \\ d(x_{r,a}, x_{t,c}) &= \max(d_r, d_t) \end{aligned}$$

Note that in all cases we have

$$|d(x_{r,a}, x_{t,c}) - d(x_{s,b}, x_{t,c})| \leq |d_r - d_s|$$

We also have

$$|d_r - d_s| \leq \frac{1}{2^{|p|}} \leq \max(d(x_{p,e}, x_{r,a}), d(x_{p,e}, x_{s,b})) = d(x_{r,a}, x_{s,b})$$

and $d(x_{r,a}, x_{s,b}) \leq d_r + d_s$ as in case (A.2.b) above. Finally, as

$$d_r + d_s \leq d(x_{r,a}, x_{t,c}) + d(x_{s,b}, x_{t,c})$$

we have

$$\begin{aligned} |d(x_{r,a}, x_{t,c}) - d(x_{s,b}, x_{t,c})| &\leq d(x_{r,a}, x_{s,b}) \\ &\leq d(x_{r,a}, x_{t,c}) + d(x_{s,b}, x_{t,c}) \end{aligned}$$

so that this case is legal.

iii. If $q \neq 0^j$ (so necessarily $p \neq 0^k$) then we have

$$d(x_{q,d}, x_{r,a}) = d(x_{q,d}, x_{s,b}) = d(x_{q,d}, x_{p,e})$$

so that $d(x_{r,a}, x_{t,c}) = d(x_{s,b}, x_{t,c})$. Since $(r, a) \sqsupset (p, e) \sqsupset (q, d)$, we have, by the lemma above:

$$d(x_{p,e}, x_{r,a}) \leq 2d(x_{q,d}, x_{p,e})$$

Similarly:

$$d(x_{p,e}, x_{s,b}) \leq 2d(x_{q,d}, x_{p,e})$$

so that

$$\begin{aligned} d(x_{r,a}, x_{s,b}) &= \max(d(x_{p,e}, x_{r,a}), d(x_{p,e}, x_{s,b})) \\ &\leq 2d(x_{q,d}, x_{p,e}) \leq 2d(x_{s,b}, x_{t,c}) \end{aligned}$$

so that this case is ok.

2. $(r, a) \sqsubset (s, b), (r, a) \perp (t, c), (s, b) \perp (t, c)$

Let (q, d) be the maximal predecessor of (r, a) and (t, c).

(a) If $q = 0^j$ and $r = 0^k$ (with $j < k$) then

$$\begin{aligned} d(x_{r,a}, x_{s,b}) &= d_r \\ d(x_{r,a}, x_{t,c}) &= \max(d_r, d_t) \\ d(x_{s,b}, x_{t,c}) &= \max(d_s, d_t) \end{aligned}$$

i. If $d_r \geq d_t$ then $d(x_{r,a}, x_{s,b}) = d(x_{r,a}, x_{t,c}) = d_r$, and since $d_s \leq 2d_r$ we have $d(x_{s,b}, x_{t,c}) \leq 2d(x_{r,a}, x_{s,b})$ and $d(x_{r,a}, x_{s,b}) = d(x_{r,a}, x_{t,c})$, so we are ok.

ii. Otherwise $d_r < d_t$, so

$$|d(x_{r,a}, x_{t,c}) - d(x_{r,a}, x_{s,b})| = d_t - d_r \leq d(x_{s,b}, x_{t,c})$$

and $d_s \leq 2d_r \leq d(x_{r,a}, x_{s,b}) + d(x_{r,a}, x_{t,c})$, so we have

$$\begin{aligned} |d(x_{r,a}, x_{t,c}) - d(x_{r,a}, x_{s,b})| &\leq d(x_{s,b}, x_{t,c}) \\ &\leq d(x_{r,a}, x_{t,c}) + d(x_{r,a}, x_{s,b}) \end{aligned}$$

which is ok.

(b) If $q = 0^j$ but $r \neq 0^k$ then

$$d(x_{r,a}, x_{t,c}) = \max(d_r, d_t)$$
$$d(x_{s,b}, x_{t,c}) = \max(d_s, d_t)$$

So in all cases we have:

$$|d(x_{s,b}, x_{t,c}) - d(x_{r,a}, x_{t,c})| \leq |d_s - d_r|$$

By (**),we know $|d_s - d_r| \leq d(x_{r,a}, x_{s,b}) \leq d_r$, and we also know

$$d_r \leq d(x_{r,a}, x_{t,c}) \leq d(x_{s,b}, x_{t,c}) + d(x_{r,a}, x_{t,c})$$

so that we have

$$\begin{aligned} |d(x_{s,b}, x_{t,c}) - d(x_{r,a}, x_{t,c})| &\leq d(x_{r,a}, x_{s,b}) \\ &\leq d(x_{s,b}, x_{t,c}) + d(x_{r,a}, x_{t,c}) \end{aligned}$$

and this case is fine.

(c) If $q \neq 0^j$ (so that $r \neq 0^k$) then we have $d(x_{q,d}, x_{s,b}) = d(x_{q,d}, x_{r,a})$ so that

$$\begin{aligned} d(x_{s,b}, x_{t,c}) &= d(x_{r,a}, x_{t,c}) \\ &= \max(d(x_{q,d}, x_{t,c}), d(x_{q,d}, x_{r,a})) \end{aligned}$$

Applying the lemma to the nodes $(q, d) \sqsubset (r, a) \sqsubset (s, b)$ we have

$$d(x_{r,a}, x_{s,b}) \leq 2d(x_{q,d}, x_{r,a}) \leq 2d(x_{s,b}, x_{t,c})$$

so these triangles are ok.

3. $(r, a) \sqsubset (t, c), (r, a) \perp (s, b), (s, b) \perp (t, c)$

This is similar to case(2), since only the inclusions, and not the relative lengths of the sequences, affect the argument.

4. $(s,b) \sqsubset (t,c), (r,a) \perp (s,b), (r,a) \perp (t,c)$

   This is also like case (2).

5. $(r,a) \sqsubset (s,b), (r,a) \sqsubset (t,c), (s,b) \perp (t,c)$

   Let (q,d) be the maximal predecessor of (s,b) and (t,c). Then $(r,a) \sqsubseteq (q,d)$ (we may have $(r,a) = (q,d)$).

   (a) If $r = 0^j$ and $q = 0^k$ with $j \leq k$, then

   $$\begin{aligned} d(x_{r,a}, x_{s,b}) &= d(x_{r,a}, x_{t,c}) = d_r \\ d(x_{s,b}, x_{t,c}) &= \max(d_s, d_t) \end{aligned}$$

   Since $s \sqsupset r$ and $t \sqsupset r$ we have both $d_s \leq 2d_r$ and $d_t \leq 2d_r$, so that $d(x_{s,b}, x_{t,c}) \leq 2d(x_{r,a}, x_{s,b})$ and this case is fine.

   (b) If $r = 0^j$ but $q \neq 0^k$ (so that $r \sqsubset q$), then:

   $$\begin{aligned} d(x_{r,a}, x_{s,b}) &= d(x_{r,a}, x_{t,c}) = d_r \\ d(x_{s,b}, x_{t,c}) &= \max(d(x_{q,d}, x_{s,b}), d(x_{q,d}, x_{t,c})) \end{aligned}$$

   By the lemma, we have

   $$d(x_{q,d}, x_{s,b}) \leq 2d(x_{r,a}, x_{q,d}) = 2d(x_{r,a}, x_{s,b})$$

   and also $d(x_{q,d}, x_{t,c}) \leq 2d(x_{r,a}, x_{t,c})$, so that $d(x_{s,b}, x_{t,c}) \leq 2d(x_{r,a}, x_{s,b})$ and this is ok.

   (c) If $r \neq 0^j$ (so $q \neq 0^k$) then:

   i. If $r = q$ we have $d(x_{s,b}, x_{t,c}) = \max(d(x_{r,a}, x_{s,b}), d(x_{r,a}, x_{t,c}))$ and we are done.

   ii. If $r \neq q$ then

   $$\begin{aligned} d(x_{r,a}, x_{s,b}) &= d(x_{r,a}, x_{t,c}) = d(x_{r,a}, x_{q,d}) \\ d(x_{s,b}, x_{t,c}) &= \max(d(x_{q,d}, x_{s,b}), d(x_{q,d}, x_{t,c})) \end{aligned}$$

   By the lemma, we have

   $$d(x_{q,d}, x_{s,b}) \text{ and } d(x_{q,d}, x_{t,c}) \leq 2d(x_{r,a}, x_{q,d}) = 2d(x_{r,a}, x_{s,b})$$

   so that $d(x_{s,b}, x_{t,c}) \leq 2d(x_{r,a}, x_{s,b})$ and we are done.

6. $(r, a) \sqsubset (s, b) \sqsubset (t, c)$

   We have $d(x_{r,a}, x_{s,b}) = d(x_{r,a}, x_{t,c})$, and we have $d(x_{s,b}, x_{t,c}) \leq 2d(x_{r,a}, x_{s,b})$ by the lemma, so this case is ok.

This completes the verification that d is a metric (on $\mathcal{D}$).

Verifying that Dist$(X, d) = A$

Recall that (X, d) is the completion of $(\mathcal{D}, d)$, so that d will continue to be a metric on X. We first show that $A \subseteq \text{Dist}(X, d)$.

Let $d_0 \in A$, $d_0 \neq 0$, be given. We will show that there is a $y \in X$ with $d(x^*, y) = d_0$. Let α witness that $d_0 \in A$, i.e. $\forall n (d_0 \upharpoonright n, \alpha \upharpoonright n) \in T$. Let $s_n = d_0 \upharpoonright n$, $a_n = \alpha \upharpoonright n$, and $y_n = x_{s_n, a_n}$. Since $(s_n, a_n) \in T$ for all n, we have each $y_n \in \mathcal{D}$ and $d(x^*, y_n) = d_{s_n}$.
claim: $\langle y_n \rangle$ is d-Cauchy.
proof: Since $(s_n, a_n) \sqsubset (s_m, a_m)$ for $n < m$, we have $d(y_n, y_m) = d(y_n, y_{n+1})$ for $n < m$. It thus suffices to show $d(y_n, y_{n+1}) \to 0$ as $n \to \infty$. Since $d_0 \neq 0$ there is a k_0 such that $s_{k_0} \neq 0^{k_0}$. Then for all $k \geq k_0$ we have $d_{s_k} \geq \frac{1}{2^{k_0}}$. We then have that there is some i_0 such that for all $i \geq i_0$ we have $\epsilon_i < \frac{1}{2^{k_0}}$. There is also a $k_1 > k_0$ with $\epsilon_{i_0} \geq \frac{1}{2^{k_1}}$. Now note that for $k \geq k_1$,

$$d(y_k, y_{k+1}) = \text{ the least } \epsilon_i \text{ such that } \epsilon_i \geq \frac{1}{2^k}$$

Thus, as $k \to \infty$ we have $\frac{1}{2^k} \to 0$, and since the $\epsilon_i \to 0$ we will have $d(y_k, y_{k+1}) \to 0$ as desired. $\square$(claim)

Thus there is some $y \in X$ with $\langle y_n \rangle \to y$, since we have taken X to be the completion of $\mathcal{D}$ in d. By the continuity of d we will have:

$$d(x^*, y) = \lim_{n \to \infty} d(x^*, y_n) = \lim_{n \to \infty} d_{s_n} = d_0$$

and thus $d_0 \in \text{Dist}(X, d)$.

We now show that $\text{Dist}(X, d) \subseteq A$. Let $d_0 \in \text{Dist}(X, d)$. We have two possibilities: either $d_0 = d(x^*, y)$ where $y = \lim y_i$ with $y_i \in \mathcal{D}$, or $d_0 = d(x, y)$ where $x = \lim x_i$, $y = \lim y_i$ for $x_i, y_i \in \mathcal{D}$.

1. $d_0 = d(x^*, y)$

   Let $y = \lim y_i$ where $y_i = x_{s_i, a_i}$ with $(s_i, a_i) \in T$. By passing to a subsequence, if necessary, we may assume that the s_i are increasing in length, since otherwise some

s_i occurs infinitely often and we have $d(x^*, y) = d_{s_i}$ for this s_i. So we may assume $|s_i| \geq i$. Since we have $d_0 = \lim d(x^*, y_i) = \lim d_{s_i}$, it must be the case that the sequence $\langle d_{s_i} \rangle$ converges, so the sequence $\langle s_i \rangle$ must converge in T. Thus there is some $\beta \in [T^*]$ with $\langle s_i \rangle \to \beta$. Also, since $d_0 \neq 0$, there is a k_0 with $\beta(k_0) = 1$, so for all but finitely many i, $d_{s_i} \geq \frac{1}{2^{k_0}}$.

claim: There is an $\alpha \in \omega^\omega$ such that $\langle a_i \rangle \to \alpha$.

proof: This amounts to showing that for each k, the sequence $\langle a_i \restriction k \rangle$ is eventually constant. We proceed by induction on k, with the case of $k = 0$ being immediate. Suppose $\langle a_i \restriction k \rangle$ is eventually constant. We may choose $i_0 \geq k$ and b such that $a_i \restriction k = b$ whenever $i \geq i_0$, and $s_i \restriction (k+1) = \beta \restriction (k+1)$ for $i \geq i_0$. If $\langle a_i \restriction (k+1) \rangle$ is not eventually constant, then there are infinitely many i with $a_i \restriction (k+1) \neq a_{i+1} \restriction (k+1)$. Thus we have infinitely many values of i for which $(s_{i+1}, a_{i+1}) \perp (s_i, a_i)$ with maximal predecessor $(\beta \restriction k, b)$.

If $\beta \restriction k = 0^k$ (i.e. $k < k_0$), then for such i with $i \geq k_0$ we have

$$d(y_i, y_{i+1}) = \max(d_{s_i}, d_{s_{i+1}}) \geq \frac{1}{2^{k_0}}$$

contradicting that $\langle y_i \rangle$ is Cauchy.

If $\beta \restriction k \neq 0^k$ (i.e. $k \geq k_0$), then for such i we have

$$d(y_i, y_{i+1}) = \max(d(x_{\beta\restriction k, b}, x_{s_i, a_i}), d(x_{\beta\restriction k, b}, x_{s_{i+1}, a_{a+1}})$$

Both of these are at least $\frac{1}{2^k}$, so $d(y_i, y_{i+1}) \geq \frac{1}{2^k}$ for infinitely many i, again contradicting that $\langle y_i \rangle$ is Cauchy. $\square$(claim)

Thus, we will have that for all n, $(\beta \restriction n, \alpha \restriction n) \in T$, so that $\beta \in p[T]$. Recalling that $A = p[T]$ and $\beta = \lim s_i = \lim d_{s_i} = d_0$ (identifying branches through T^* with reals), we have $d_0 \in A$ as desired, completing this case.

2. $d_0 = d(x, y)$

Let $x = \lim x_i$ with $x_i = x_{s_i, a_i}$ and $y = \lim y_i$ with $y_i = x_{t_i, b_i}$. As in the previous case, the s_i's converge to some $\beta_x \in [T^*]$, the t_i's converge to some β_y, the a_i's converge to some $\alpha_x \in \omega^\omega$ and the b_i's converge to some α_y. We may assume that β_x and β_y are not eventually 0, since this would imply that either $\langle x_i \rangle \to x^*$ or $\langle y_i \rangle \to x^*$. We may also assume $(\beta_x, \alpha_x) \neq (\beta_y, \alpha_y)$, since otherwise we would have $d_0 = 0$. Let k be least such that

$$(\beta_x, \alpha_x) \restriction (k+1) \neq (\beta_y, \alpha_y) \restriction (k+1)$$

Let $r = \beta_x \restriction k = \beta_y \restriction k$ and $c = \alpha_x \restriction k = \alpha_y \restriction k$. Then for all but finitely many i we have $(s_i, a_i) \perp (t_i, b_i)$ with maximal predecessor (r, c). Let

$$\begin{aligned} (s, a) &= (\beta_x \restriction (k+1), \alpha_x \restriction (k+1)) \\ (t, b) &= (\beta_y \restriction (k+1), \alpha_y \restriction (k+1)) \end{aligned}$$

Note that $(r, c) \sqsubset (s, a)$ and $(r, c) \sqsubset (t, b)$. Also, for all but finitely many i, $(s, a) \sqsubset (s_i, a_i)$ and $(t, b) \sqsubset (t_i, b_i)$. For these i,

$$\begin{aligned} d(x_{r,c}, x_{s_i,a_i}) &= d(x_{r,c}, x_{s,a}) \\ d(x_{r,c}, x_{t_i,b_i}) &= d(x_{r,c}, x_{t,b}) \end{aligned}$$

Thus, if $r \neq 0^k$ then for all but finitely many i,

$$\begin{aligned} d(x_i, y_i) &= \max(d(x_{r,c}, x_{s_i,a_i}), d(x_{r,c}, x_{t_i,b_i})) \\ &= \max(d(x_{r,c}, x_{s,a}), d(x_{r,c}, x_{t,b})) \end{aligned}$$

so $d_0 = \max(d(x_{r,c}, x_{s,a}), d(x_{r,c}, x_{t,b})) \in A$ as desired.

Otherwise, if $r = 0^k$ then $d(x_i, y_i) = \max(d_{s_i}, d_{t_i})$. So there will either be infinitely many i with $d(x_i, y_i) = d_{s_i}$, or infinitely many with $d(x_i, y_i) = d_{t_i}$ (according as β_x is to the right or to the left of β_y in the lexicographical order on 2^ω). But then we have that either $d_0 = \beta_x$ or $d_0 = \beta_y$, both of which are in $p[T] = A$. So again $d_0 \in A$.

This finishes the construction of (X, d) with Dist$(X, d) = A$ in the case that $A \subseteq [0, 1)$.

The general case of $A \subseteq [0, \infty)$

Now for $A \subseteq [0, \infty)$, let $A_n = A \cap [0, n)$. Since the A_n's satisfy the hypotheses of the theorem, we can construct (X_n, d_n) as above (stretching by n) such that Dist$(X_n, d_n) = A_n$. Now choose a sequence $\langle \delta_n \rangle$ with $\delta_n \in A_n$, $\delta_n \geq \frac{1}{2} \sup A_n$, and $\delta_n \leq \delta_{n+1}$. We let X be the disjoint union $\bigsqcup_{n \geq 1} X_n$. We set $d = d_n$ on each X_n, and for $x \in X_n$, $y \in X_m$ with $n < m$ we let $d(x, y) = \delta_m$. The conditions on the δ_n's guarantee that this is a metric and adds no distances other than those in the A_n's, so that Dist(X,d) $= \bigcup_{n \geq 1} A_n = A$. □

We should note that there is a lack of uniformity in the above construction. For a given analytic set A, we needed to pick a tree representation for it, as well as a sequence approaching 0, in order to build the metric space. Different choices of trees and sequences will produce different (non-isometric) spaces. We do not then have a map from analytic sets to metric spaces, unless we use the Axiom of Choice to pick representatives. Another way to say this is that we do not have a reduction of the relation of equality of (codes for) analytic sets to the isometry relation on Polish metric spaces. In fact, there can be no such "definable" map (e.g. Borel, or in $L(R)$ assuming $\mathrm{AD}^{L(R)}$), since the isometry relation is reducible to a Polish group action (by a result of [11]), whereas equality of analytic sets is not. We will discuss this point further in Chapter 5.

Along these lines, we can consider two questions. A set of distances is said to be *metrically rigid* for a class of metric spaces if it is the distance set for a unique space in the class (up to isometry). The authors of [24] consider this notion for certain classes of metric spaces. We can here ask:

Question 4.5 *Is there an analytic set which is metrically rigid for the class of Polish metric spaces?*

More generally, we can ask how far the distance set is from being a complete invariant:

Question 4.6 *For a given analytic set A, how complicated is it to classify up to isometry those Polish metric spaces having A as their set of distances?*

The precise meaning of this will be explained in the next chapter. We should note that the complexity will depend drastically on the set A; for instance, classifying spaces whose distance set is all of $[0, \infty)$ will be as complicated as classifying all Polish metric spaces.

4.2 Special Cases of Distances

We now consider various special classes of Polish metric spaces, and ask what the set of distances can look like. In general, the set of distances will retain some of the properties of the original space, but not all. The cases we will look at are: compact, locally compact, σ-compact, ultrametric, discrete, zero-dimensional, connected and path-connected. In the next chapter we will see that there is some analogy between the complexity of distance sets and the complexity of the isometry problem for a given class of spaces. For brevity, we will always assume that 0 is contained in a putative set of distances.

Theorem 4.7 *A set $A \subseteq [0,\infty)$ is the set of distances of some compact metric space if and only if either A is finite or A is compact and 0 is a limit point of A.*

Proof: $\Longrightarrow$ Since X^2 is compact, and $d : X^2 \to \mathbb{R}$ is continuous, we must have that $\mathrm{Dist}(X,d)$ is compact (the continuous image of a compact set is compact). Since the only discrete compact spaces are finite, we must have 0 as a limit point of $\mathrm{Dist}(X,d)$ if X is infinite.

$\Longleftarrow$ The case where A is finite is handled in the same way that the countable case of Theorem 4.3 was, except that we only include finitely many points in the space.

For the uncountable case, we may assume $A \subseteq [0,1]$ since A is bounded (and we may simply multiply the resulting metric by a constant). We can represent the set A as the branches through a tree on $\{0,1\}$, i.e. a closed subset of the Cantor space, using as before the fact that the map $\alpha \mapsto \sum_{i\in\omega} \frac{\alpha(i)}{2^{i+1}}$ is a continuous surjection of $\mathcal{C}$ onto $[0,1]$. We here must allow certain points to be represented twice, namely those with eventually constant binary expansions; this will not affect the set of distances, though. The construction will resemble the original construction, using a tree on $\{0,1\}$ rather than on $\{0,1\}\times\omega$ (in fact, it will essentially be that construction with a tree all of whose second coordinates are 0).

Let T be a pruned tree with $A = [T]$, and choose $\langle \epsilon_i \rangle$ and $\{d_s : s \in T\}$ as in Theorem 4.3, with T here taking the place of T^*. Our space will have a countable dense set

$$\mathcal{D} = \{x^*\} \cup \{x_s : s \in T\}$$

We set $d(x^*, x_s) = d_s$ for $s \in T$. For $s,t \in T$, we consider two cases:

case 1: $s \sqsubset t$ (or vice versa)

(a) If $s = 0^k$ for some k then let $d(x_s, x_t) = d_s$.

(b) If $s \neq 0^k$ and there is an ϵ_i with $\frac{1}{2^{|s|}} \leq \epsilon_i \leq d_s$ then let

$$d(x_s, x_t) = \text{ the least such } \epsilon_i$$

(c) Otherwise let $d(x_s, x_t) = d_s$.

case 2: $s \perp t$

Let r be the maximal predecessor of s and t.

(a) If $r = 0^k$ let $d(x_s, x_t) = \max(d_s, d_t)$.

(b) Otherwise, let $d(x_s, x_t) = \max(d(x_r, x_s), d(x_r, x_t))$.

Checking that d is a metric on $\mathcal{D}$ and that $\mathrm{Dist}(X,d) = A$ is exactly as before. It remains to show that (X,d) is compact.

We will show that (X, d) is totally bounded. Fix $\epsilon > 0$. Let ϵ_{i_0} be chosen from our sequence such that $\epsilon_{i_0} < \epsilon$, and let n_0 be such that $\frac{1}{2^{n_0}} \leq \epsilon_{i_0}$. Set:

$$D_\epsilon = \{d_s : s \in T \text{ and } |s| \leq n_0\}$$

Then D_ϵ is a finite set, and we claim that every point in X is within distance ϵ of some point in D_ϵ. First, note that since $0^{n_0} \in T$ and $d_{0^{n_0}} \leq \frac{1}{2^{n_0}} < \epsilon$, this holds for x^*. Any other point $x \in X$ (outside of D_ϵ) corresponds either to some x_t with $|t| > n_0$, or to a branch through T, i.e. there is a $\beta \in [T]$ such that $x = \lim_{n\to\infty} x_{\beta \upharpoonright n}$. Let $b = t \upharpoonright n_0$ or $b = \beta \upharpoonright n_0$, respectively, so $x_b \in D_\epsilon$. If $b = 0^{n_0}$ then $d(x, x_b) = d_{0^{n_0}} \leq \epsilon_{i_0}$. If $b \neq 0^{n_0}$ then we have $d(x, x_b)$ equal to one of: ϵ_{i_0}, some $\epsilon_j < \epsilon_{i_0}$, or $d_b < \epsilon_{i_0}$. In all cases $d(x, x_b) \leq \epsilon_{i_0} < \epsilon$ and we are done. □

Theorem 4.8 *The following are equivalent:*

1. $A \subseteq [0, \infty)$ *is the set of distances of some locally compact Polish metric space*

2. A *is the set of distances of some* σ*-compact Polish metric space*

3. A *is either countable or* A *is* K_σ *with 0 as a limit point.*

Proof: (1) $\implies$ (2) follows from the fact that any locally compact Polish space is σ-compact.

(2) $\implies$ (3): Let $X = \bigcup_{i\in\omega} K_i$ where each K_i is compact and $K_i \subseteq K_{i+1}$. Then $(K_i, d \upharpoonright K_i)$ is a compact metric space. Set $A_i = \text{Dist}(K_i, d \upharpoonright K_i)$, so that each A_i is compact. Notice that if $d_0 = d(x, y)$ for some $x, y \in X$, then there is some K_i such that $x, y \in K_i$, so that $d_0 \in A_i$. Thus, we get that $\text{Dist}(X, d) = \bigcup_{i\in\omega} A_i$, so that $\text{Dist}(X, d)$ is K_σ, and is either countable or contains 0 as a limit point since (X, d) is a Polish metric space.

(3) $\implies$ (1): Let $A = \bigcup_{i\in\omega} A_i$ where each A_i is compact, $A_i \subseteq A_{i+1}$, and $0 \in A_0$. If 0 is a limit point of A, then we can choose these sets such that 0 is a limit point of A_0: Simply add 0 to A_0 along with a sequence approaching 0; this sequence, together with the point 0, is a compact set. We will also assume that if A contains more than one point, then A_0 contains a non-zero point. Then each A_i satisfies the hypotheses of Theorem 4.7, and so we can build compact metric spaces (X_i, d_i) with $\text{Dist}(X_i, d_i) = A_i$. Let $M_i = \sup A_i < \infty$; this is an element of A_i by the compactness of A_i. Note that $M_i \leq M_{i+1}$, so for all i we

have $M_i \geq M_0 > 0$. We now define our space to be a disjoint union of the X_i. Let the underlying set be $X = \bigsqcup_{i\in\omega} X_i$. To define d, let

$$d(x,y) = \begin{cases} d_i(x,y) & \text{if } x, y \in X_i \text{ for some } i \\ M_j & \text{if } x \in X_i,\ y \in X_j \text{ for } i < j \end{cases}$$

This will define a Polish metric space, and it will be locally compact because for each $x \in X$, the open ball of radius M_0 centered at x is contained entirely within some X_i and hence has compact closure. □

Recall that a metric d is said to be an *ultrametric* if for all x, y, z, we have

$$d(x,z) \leq \max(d(x,y), d(y,z))$$

This is equivalent to saying that the longest two sides in any triangle have the same length.

Theorem 4.9 *A set $A \subseteq [0,\infty)$ is the set of distances of some ultrametric Polish metric space if and only if A is countable.*

Proof: $\Longleftarrow$ Recall that the construction given in Theorem 4.3 for the case where A is countable in fact produces an ultrametric space.

$\Longrightarrow$ Let (X,d) be a given ultrametric Polish metric space. Fix a countable dense set $\mathcal{D} \subseteq X$, $\mathcal{D} = \{x_i : i \in \omega\}$. We claim that

$$\text{Dist}(X,d) \ = \{d(x_i, x_j) : i, j \in \omega\}$$

which is a countable set. To see this, let $x \neq y \in X$ with $x = \lim x_{i_n}$, $y = \lim x_{j_n}$ where $\langle x_{i_n}\rangle$ and $\langle x_{j_n}\rangle$ are Cauchy sequences. Then $d(x,y) = \lim_{n\to\infty} d(x_{i_n}, x_{j_n})$. Set $\delta = d(x,y)$. Now, since the two sequences are Cauchy and approach x and y respectively, there is an N such that

$$(\forall n \geq N)\left(d(x_{i_n}, x) < \frac{\delta}{2} \ \&\ d(x_{j_n}, y) < \frac{\delta}{2}\right)$$

and

$$(\forall n, m \geq N)\left(d(x_{i_n}, x_{i_m}) < \frac{\delta}{2} \ \&\ d(x_{j_n}, x_{j_m}) < \frac{\delta}{2}\right)$$

Now, for $n \geq N$, consider the arrangement of x, y, x_{i_n} and x_{j_n}. Because d is an ultrametric, we must have

$$d(x, x_{j_n}) = \max(d(x,y), d(y, x_{j_n})) = \delta$$

since $d(y, x_{j_n}) < \delta$. Then, since $d(x, x_{i_n}) < \delta$, we have

$$d(x_{i_n}, x_{j_n}) = \max(d(x, x_{i_n}), d(x, x_{j_n})) = \delta$$

From this we get that

$$d(x, y) = d(x_{i_N}, x_{j_N}) \in \{d(x_i, x_j) : i, j \in \omega\}$$

as desired. □

Theorem 4.10 *A set $A \subseteq [0, \infty)$ is the set of distances of some discrete metric space if and only if A is countable.*

Proof: A discrete Polish space is of course countable and so has a countable set of distances. So, for the other direction, let A be given, and let $A \setminus \{0\} = \{a_i : i \in \omega\}$, where we allow repetitions in the case A is finite. Now let the space have underlying set

$$X = \{x_i : i \in \omega\} \cup \{y_i : i \in \omega\}$$

and define d by:

$$\begin{aligned} d(x_i, y_i) &= a_i \\ d(x_i, x_j) &= d(y_i, y_j) = d(x_i, y_j) = \max(a_0, a_i, a_j) \text{ for } i \neq j \end{aligned}$$

It is straightforward to check that this defines a metric. It is complete since all distances are at least a_0, so there are no non-trivial Cauchy sequences, and it is discrete since for each x_i or y_i, no point has distance less than a_i from it. □

A *zero-dimensional* space is one in which there is a basis consisting of clopen sets. Although zero-dimensional spaces are a special class of Polish metric spaces, the following result shows that their distance sets can be as complicated as those of arbitrary Polish metric spaces.

Theorem 4.11 *A set $A \subseteq [0, \infty)$ is the set of distances of some zero-dimensional Polish metric space if and only if either A is countable or A is analytic with 0 as a limit point.*

Proof: Let A be given. If A is countable, then the construction in Theorem 4.3 produces an ultrametric space, and every ultrametric space is zero-dimensional.

So suppose that 0 is a limit point of A. A key point here is the following: If (X, d) is a Polish metric space and $\mathrm{Dist}(X, d)$ is disjoint from some sequence $\langle \epsilon_i \rangle$ approaching 0, then (X, d) is zero-dimensional, since we can take as a basis the open balls of radius ϵ_i, which will in fact be clopen. Thus, if A is disjoint from some sequence approaching 0, the construction in Theorem 4.3 will necessarily produce a zero-dimensional space.

If A is not disjoint from such a sequence (i.e. A contains a half-neighborhood of 0), we proceed as follows. Let $A_0 = A \cap (\{0\} \cup \{\frac{1}{n} : n \geq 1\})$ and let $A_1 = \{0\} \cup (A \setminus A_0)$. We can now produce spaces (X_0, d_0) and (X_1, d_1) with $\mathrm{Dist}(X_0, d_0) = A_0$ and $\mathrm{Dist}(X_1, d_1) = A_1$. Both these spaces will be zero-dimensional as noted above. We now let $X = X_0 \sqcup X_1$. To define d, fix a point $x_1 \in X_1$ and let $\delta_0 = supA_0 \in A_0$. Then let

$$d(x, y) = \begin{cases} d_0(x, y) & \text{if } x, y \in X_0 \\ d_1(x, y) & \text{if } x, y \in X_1 \\ \delta_0 + d_1(x, x_1) & \text{if } x \in X_1,\ y \in X_0 \end{cases}$$

This defines a metric, and the resulting space is zero-dimensional since the union of bases for each component will give a basis for the whole space. □

Theorem 4.12 *The following are equivalent:*

1. *A is the set of distances of a path-connected Polish metric space*
2. *A is the set of distances of a connected Polish metric space*
3. *A is an interval of the form $[0, r)$, $[0, r]$, or $[0, \infty)$*

Proof: (1) $\Longrightarrow$ (2) is immediate.
(2) $\Longrightarrow$ (3): The continuous image of a connected set is connected, since the pull-back of a clopen partition would give a clopen partition of the original space. Thus $\mathrm{Dist}(X, d)$ must be a connected set containing 0, and hence an interval.
(3) $\Longrightarrow$ (1): For the first type of interval $[0, r)$ we can take the space to be $X = \mathbb{R}$ with the metric $d(x, y) = r\frac{|x-y|}{1+|x-y|}$. For the second, we can take $X = [0, r]$ with the standard metric, and for the third $X = \mathbb{R}$ with the standard metric. □

4.3 The Set of Triangles and Beyond

Having characterized the possible sets of distances in a Polish metric space, one can ask what sets of triangles are possible, or in general what sets of n-point configurations are possible for a given n.

Definition 4.13 *For (X,d) a Polish metric space and $n \geq 2$, let the n-point spectrum be*

$$Spec_n(X,d) = \{\langle d_{i,j}\rangle_{i<j<n} : (\exists x_0,\ldots,x_{n-1} \in X)(\forall i < j < n)[d_{i,j} = d(x_i,x_j)]\}$$

Note that $\mathrm{Spec}_2(X,d) = \mathrm{Dist}(X,d)$. In general, for $m < n$, $\mathrm{Spec}_n(X,d)$ completely determines $\mathrm{Spec}_m(X,d)$, and these must all be analytic sets. There is not such a simple characterization of the possible n-point spectra as there was in the case $n = 2$. Let us note one additional necessary condition in the case $n = 3$, the set of triangles.

Suppose the space contains a triangle with one side of length d_0, and suppose $d_1 > d_0$ is any distance occurring in the space. Then, by considering points x_1 and x_2 with $d(x_1,x_2) = d_0$ and points y_1 and y_2 with $d(y_1,y_2) = d_1$ (where one of the y's may be the same as one of the x's), we see

$$\begin{aligned} d(x_1,y_1) + d(x_1,y_2) &\geq d_1 \\ d(x_2,y_1) + d(x_2,y_2) &\geq d_1 \end{aligned}$$

Thus $d(x_1,y_1) + d(x_1,y_2) + d(x_2,y_1) + d(x_2,y_2) \geq 2d_1$, from which we get that at least one of these four distances is at least $\frac{d_1}{2}$. In other words: If $d_1 > d_0$ are any two distances in (X,d), then (X,d) contains a triangle with one side of length d_0 and another side of length at least $\frac{d_1}{2}$. So the behavior of the set of triangles is not "local"; adding one triangle will necessitate adding a number of others. So we can ask the following:

Question 4.14 *For $n \geq 3$, what sets can be $Spec_n(X,d)$ for some Polish metric space (X,d)?*

Although Polish metric spaces are not characterized up to isometry by their sets of distances, or even by the sequence $\langle \mathrm{Spec}_n(X,d)\rangle_{n\in\omega}$, there are two cases in which this is true. One example is the case of compact metric spaces: Two compact metric spaces are isometric if and only if they have the same n-point spectra for all $n \geq 2$. In this case, the spectra are fairly concrete: $\mathrm{Spec}_n(X,d)$ is a compact subset of $\mathbb{R}^{\frac{n(n-1)}{2}}$. Thus, we can take a sequence of compact sets as a complete invariant for isometry. This gives a proof of the result of Gromov (see [12]) that the isometry relation on compact metric spaces is concretely classifiable, as discussed in Chapter 5.

Another case where the spectra form complete invariants is that of *ultra-homogeneous* spaces, those in which any isometry between finite subsets of the space extends to an isometry of the whole space. Here the spectra are no longer compact, so we do not get a concrete classification. It would be interesting to know the possible spectra in this case, as an indication of the complexity of the isometry relation of ultra-homogeneous Polish metric spaces.

Chapter 5

Isometry of Polish Metric Spaces

In this chapter we consider the equivalence relation of *isometry of Polish metric spaces* and ask how complicated it is. In order to make this precise, we will view it as an equivalence relation on the space of codes for Polish metric spaces. Since we are dealing with separable, complete metric spaces, they are completely determined by their metrics restricted to a countable dense set. More precisely, we may fix a countable dense set $\{x_i : i \in \omega\}$ in a space, and then code it by an array $\langle d_{i,j} \rangle_{i,j\in\omega}$, where $d_{i,j}$ is the distance between the points x_i and x_j. The array will satisfy the conditions of a metric on the set $\{x_i : i \in \omega\}$. We will use such arrays as codes for Polish metric spaces, although we will permit two points to be at distance 0 to allow for the coding of finite spaces (in which case we identify the two points in the resulting space). Such an array then codes the Polish metric space which is the completion of the given countable space (which will be dense in the resulting space).

Definition 5.1 *Let the space of (codes for) Polish metric spaces be:*

$$\begin{aligned}\mathcal{M} \quad = \quad & \{\langle d_{i,j} \rangle_{i,j\in\omega} \textit{ such that} \\ & \textit{(1) } (\forall i,j)(d_{i,j} \in \mathbb{R} \textit{ and } d_{i,j} \geq 0) \\ & \textit{(2) } (\forall i)(d_{i,i} = 0) \\ & \textit{(3) } (\forall i,j)(d_{i,j} = d_{j,i}) \\ & \textit{(4) } (\forall i,j,k)(d_{i,k} \leq d_{i,j} + d_{j,k})\}\end{aligned}$$

Note that this is a closed subspace of $\mathbb{R}^{\omega\times\omega}$ and is hence a Polish space in the relative topology. We now define the isometry relation, $\cong_i$, on this space.

Definition 5.2 *For $\langle d_{i,j}\rangle$ and $\langle \widetilde{d}_{i,j}\rangle$ in $\mathcal{M}$, we set $\langle d_{i,j}\rangle \cong_i \langle \widetilde{d}_{i,j}\rangle$ if and only if the metric spaces coded by the two arrays are isometric.*

We note that this is a $\mathbf{\Sigma}^1_1$ equivalence relation on $\mathcal{M}$:

$$\begin{aligned}\langle d_{i,j}\rangle \cong_i \langle \widetilde{d}_{i,j}\rangle \iff & (\exists f : \omega \to \omega^\omega \text{ sending } n \mapsto f_n) \text{ such that:}\\ & (1)\ (\forall n)(\forall k)(\exists N)(\forall i,j \geq N)\left(\widetilde{d}_{f_n(i),f_n(j)} < \frac{1}{k}\right)\\ & (2)\ (\forall n,m)(\forall k)(\exists N)(\forall i \geq N)\left(\left|d_{n,m} - \widetilde{d}_{f_n(i),f_m(i)}\right| < \frac{1}{k}\right)\\ & (3)\ (\forall n)(\forall k)(\exists m)(\exists N)(\forall i \geq N)\left(\widetilde{d}_{n,f_m(i)} < \frac{1}{k}\right)\end{aligned}$$

That is, f induces a map sending each x_i to a $\hat{d}$-Cauchy sequence (by condition (1)), condition (2) guarantees that this is an isometric embedding, and condition (3) guarantees that it is surjective. It will turn out that $\cong_i$ is in fact a $\mathbf{\Sigma}^1_1$-complete equivalence relation.

In attempting to classify Polish metric spaces (or subclasses thereof) up to isometry, one can characterize the difficulty of the classification in terms of the complexity of the corresponding equivalence relation (in terms of Borel reducibility, $\leq_B$). As discussed in the introduction, this gives a measure of how complicated a set of complete invariants must be. In the case of compact metric spaces, the classification is quite simple:

Theorem 5.3 (Gromov [12]) *Isometry of compact metric spaces is concretely classifiable.*

Gromov's proof involves introducing a metric, the *Gromov-Hausdorff metric*, on the space of compact metric spaces. An alternate proof of this is to note that two compact metric spaces are isometric if and only if they have the same sets of n-point configurations, Spec_n, for all $n \geq 2$. These sets will be compact when the metric spaces are compact, and so the sequence of Spec_n's can be coded by a real.

The general problem of classifying Polish metric spaces up to isometry, on the other hand, is quite difficult:

Theorem 5.4 (Gao-Kechris [11]) *Isometry of Polish metric spaces, $\cong_i$, is Borel bireducible with the universal equivalence relation of a Borel action of a Polish group on a Polish space.*

Let us note one consequence of the fact that the isometry relation is reducible to the orbit equivalence relation of a Polish group. In the previous chapter we gave a construction which produced a Polish metric space whose set of distances was an arbitrary analytic set of non-negative reals (including 0, and having 0 as a limit point if uncountable). The result of Gao and Kechris shows that we can not do this uniformly (in any definable way), even for $\mathbf{\Sigma}^0_2$ sets. The reason for this is that the equivalence relation E_1, defined on $2^{\omega\times\omega}$ by

$$\langle x_n\rangle E_1 \langle y_n\rangle \iff \forall^\infty n(x_n = y_n)$$

is known not to be reducible to the orbit equivalence relation of any Polish group (Kechris and Louveau [22]). Since this is a $\mathbf{\Sigma}^0_2$ equivalence relation, it is easily reducible to the relation of equality of (codes for) $\mathbf{\Sigma}^0_2$ sets, which then can not be reducible to the isometry relation.

We present here a technique for reducing equivalence relations induced by Borel actions of Polish groups to the isometry relation, discovered independently. This method allows us to recover one direction of the above theorem, and also allows us to find new lower bounds for the complexity of isometry of some particular classes of Polish metric spaces.

The first section is devoted to setting out this technique. In the remaining sections we will use it, and certain refinements, to give lower bounds for the complexity of isometry of particular classes.

5.1 Reducing a Polish Group Action to Isometry

Let G be a Polish group, and let E^X_G be a *Borel G-space*, i.e. the equivalence relation on a Polish space X induced by a Borel action of G on X. We will show that $E^X_G \leq_B \cong_i$. The following theorem will simplify matters:

Theorem 5.5 (Becker and Kechris [2]) *For any Borel G-space E^X_G, there is a Polish G-space E^Y_G (i.e. the action is continuous) with Y a compact Polish space, such that $E^X_G \leq_B E^Y_G$.*

So it will suffice to reduce E^Y_G to isometry, where G acts continuously on a compact space Y. We will define, for each $z \in Y$, a Polish metric space (X_z, d_z) such that

$$z_1 E^Y_G z_2 \iff (X_{z_1}, d_{z_1}) \cong_i (X_{z_2}, d_{z_2})$$

Let d_Y be a complete metric on Y compatible with the topology, with $d_Y \leq 1$, and let d_G be a compatible left-invariant metric on G with $d_G \leq 1$. Recall that a metric on a topological group is left-invariant if

$$(\forall g_1, g_2 \in G)(\forall h \in G)[d(hg_1, hg_2) = d(g_1, g_2)]$$

Any Polish group G admits a compatible left-invariant metric, although it does not necessarily admit a complete left-invariant metric. Also, given any metric d, we can form the new metric d' given by

$$d'(x,y) = \frac{d(x,y)}{1 + d(x,y)}$$

This will have values less than 1, and will preserve most of the properties of the original metric. In particular, it will continue to be left-invariant when d is.

Before we begin, let us give the naive idea behind the construction. Suppose for the moment that G a has a complete left-invariant metric, that $Y = [0,1]$, and that the action is well-behaved in the sense that for all y in Y and all g_1 and g_2 in G we have

$$\left|g_1^{-1} \cdot y - g_2^{-1} \cdot y\right| \leq 2d_G(g_1, g_2)$$

In this case, we can define a map $z \mapsto (X_z, d_z)$ to directly encode the orbit of z into the set of distances in X_z. Let the underlying set be

$$X_z = \{x^*\} \cup \{x_g : g \in G\}$$

and set

$$\begin{aligned} d_z(x_{g_1}, x_{g_2}) &= d_G(g_1, g_2) \\ d_z(x^*, x_g) &= \tfrac{3}{2} + \tfrac{1}{2} g^{-1} \cdot z \end{aligned}$$

Thus, X_z consists of an isometric copy of G together with a distinguished point x^*, and we use the distances from x^* to directly code the orbit of z. Checking the triangle inequality only requires noting that, by the assumption above,

$$\begin{aligned} |d_z(x^*, x_{g_1}) - d_z(x^*, x_{g_2})| &= \frac{1}{2}\left|g_1^{-1} \cdot y - g_2^{-1} \cdot y\right| \\ &\leq d_G(g_1, g_2) \\ &= d_z(x_{g_1}, x_{g_2}) \end{aligned}$$

It is immediate that if we have two points z_1 and z_2 in distinct orbits, then the sets of distances from x^* will be different in X_{z_1} than it is in X_{z_2}. Hence $(X_{z_1}, d_{z_1}) \not\cong_i (X_{z_2}, d_{z_2})$, since we can see that any isometry between the two spaces would have to send the x^* of the first to that of the second. One could also note that if f is an isometry of (X_{z_1}, d_{z_1}) and (X_{z_2}, d_{z_2}) which sends x_{1_G} to x_h, then we must have $z_1 = h^{-1} \cdot z_2$.

Conversely, if z_1 and z_2 are in the same orbit, say $z_2 = h \cdot z_1$, then we can define the map:

$$f : \begin{cases} x^* & \mapsto & x^* \\ x_g & \mapsto & x_{hg} \end{cases}$$

The left-invariance of d_G ensures this is an isometry on the copy of G, and we check:

$$\begin{aligned} d_{z_2}(f(x^*), f(x_g)) &= d_{z_2}(x^*, x_{hg}) \\ &= \frac{3}{2} + \frac{1}{2}(hg)^{-1} \cdot z_2 \\ &= \frac{3}{2} + \frac{1}{2}g^{-1} \cdot (h^{-1} \cdot z_2) \\ &= \frac{3}{2} + \frac{1}{2}g^{-1} \cdot z_1 \\ &= d_{z_1}(x^*, x_g) \end{aligned}$$

Thus, we can think of the two spaces as encoding the orbit relative to the parameters z_1 and z_2, respectively, and the isometry simply changes the parameterization.

There are, of course, three problems with implementing this strategy in general. First, G may not have a complete left-invariant metric. We will resolve this problem by using, instead, copies of the Polish space $\widehat{G}$, which will be the completion of G with respect to a left-invariant metric. G will be dense in this space, and this will turn out to be sufficient to distinguish orbits. The second problem is that our space Y will not in general be the unit interval (or even continuously embeddable in the unit interval), so we can not literally encode the points in an orbit as distances. Instead, we will we will fix a countable dense subset of Y and encode the distances from a given point to each element of the dense set, which will be elements of $[0, 1]$. Two points with identical distances from all points in a dense set must in fact be the same. This will require that we have countably many places to do the encoding, rather than being able to use a single point as we did before. Finally, the action will not generally satisfy the distance condition we assumed. We can fix this by modifying the metric on G to make this condition hold. We could now do the construction using a copy of $\widehat{G}$ along with a countable set of distinguished points, but we will give a somewhat more complicated construction, using countably many copies of $\widehat{G}$ instead, which is more easily generalized.

We will first define a new metric, d'_G, on G which has better behavior with respect to the action of G on Y.

Definition 5.6 *For g_1 and g_2 in G, let*

$$d'_G(g_1, g_2) = \frac{1}{2} d_G(g_1, g_2) + \frac{1}{2} \sup\{d_Y(g_1^{-1} \cdot y, g_2^{-1} \cdot y) : y \in Y\}$$

Lemma 5.7 *d'_G is a compatible left-invariant metric on G with $d'_G \leq 1$, and d'_G is complete if d_G is.*

Proof: Clearly $d'_G \leq 1$. Left-invariance is also easy to check:

$$\begin{aligned} d'_G(hg_1, hg_2) &= \frac{1}{2} d_G(hg_1, hg_2) + \frac{1}{2} \sup\{d_Y((hg_1)^{-1} \cdot y, (hg_2)^{-1} \cdot y) : y \in Y\} \\ &= \frac{1}{2} d_G(g_1, g_2) + \frac{1}{2} \sup\{d_Y(g_1^{-1} \cdot (h^{-1} \cdot y), g_2^{-1} \cdot (h^{-1} \cdot y)) : y \in Y\} \\ &= \frac{1}{2} d_G(g_1, g_2) + \frac{1}{2} \sup\{d_Y(g_1^{-1} \cdot y, g_2^{-1} \cdot y) : y \in Y\} \\ &= d_G(g_1, g_2) \end{aligned}$$

We now show that a sequence $\langle g_n \rangle_{n \in \omega}$ converges with respect to d'_G if and only if it converges with respect to d_G, which shows that the identity map on G is a homeomorphism. This will give that d'_G is compatible with the topology, and also show that d'_G is complete if and only if d_G is complete. Since $d'_G \geq \frac{1}{2} d_G$, any d'_G-convergent sequence is d_G-convergent. So we need only check the converse. By left-invariance, we need only check this for sequences converging to 1_G.

Let $\langle g_n \rangle_{n \in \omega} \to 1_G$ with respect to d_G. Since the action of G on Y is continuous from $G \times Y \to Y$,

$$(\forall y_0)(\forall \epsilon)(\exists \delta)(\forall g)(\forall y)\, [d_g \times d_y((g, y), (1_G, y_0)) < \delta \Longrightarrow d_Y(g \cdot y, y_0) < \epsilon]$$

But Y is compact, so in fact we get

$$(\forall \epsilon)(\exists \delta)(\forall g)\, [d_G(g, 1_G) < \delta \Longrightarrow (\forall y)(d_Y(g \cdot y, y) < \epsilon)]$$

Thus, as $\langle g_n \rangle \to 1_G$ in d_G, we have $\sup\{d_Y(g_n^{-1} \cdot y, y) : y \in Y\} \to 0$, which shows that the sequence will converge with respect to d'_G. □

Note that d'_G now has the property that for any $z \in Y$ and $g_1, g_2 \in G$, we have

$$d'_G(g_1, g_2) \geq \frac{1}{2} d_Y(g_1^{-1} \cdot z, g_2^{-1} \cdot z) \qquad (*)$$

Now fix a countable dense subset of Y, enumerated as $\{y_n\}_{n\in\mathbb{Z}}$. Let $\widehat{G}$ be the completion of G in d'_G. Note that $\widehat{G}$ is no longer a Polish group if d_G is not complete, but it is a Polish space. Thus, since G is a Polish subspace of $\widehat{G}$ in the relative topology, we have that G is a dense G_δ subset of $\widehat{G}$, and hence comeager in $\widehat{G}$.

Defining (X_z, d_z)

Fix $z \in Y$. We will now define the Polish metric space (X_z, d_z). We will let X_z have as its underlying set $\widehat{G} \times \mathbb{Z} \times \mathbb{Z}_2$. We will define d_z on the dense subset $G \times \mathbb{Z} \times \mathbb{Z}_2$ (if we fix a countable dense set $G_0 \subseteq G$ then in fact $G_0 \times \mathbb{Z} \times \mathbb{Z}_2$ will be a countable dense subset of X_z, but it will be simpler to define d_z on the given subset). So let

$$X_z = \{x_{\hat{g}}^{i,n}\}_{\hat{g}\in\widehat{G}, n\in\mathbb{Z}, i\in\{0,1\}}$$

Let $\pi : \mathbb{Z} \leftrightarrow \omega$ be the bijection given by

$$\pi(n) = \begin{cases} 2n & \text{if } n \geq 0 \\ -1-2n & \text{if n ¡ 0} \end{cases}$$

Definition 5.8 *For* $g_0, g_1 \in G, n_0, n_1 \in \mathbb{Z}, i_0, i_1 \in \{0,1\}$, *let* $d_z(x_{g_0}^{i_0,n_0}, x_{g_1}^{i_1,n_1}) =$

$$\begin{cases} d'_G(g_0, g_1) & \text{if } i_0 = i_1 \text{ and } n_0 = n_1 \\ \frac{3}{2} + 4^{-|n_0-n_1|}\left[1 + d'_G(g_0, g_1)\right] & \text{if } i_0 = i_1 \text{ and } n_0 \neq n_1 \\ 1 + 4^{-1-\pi(n_0-n_1)}\left[1 + d_Y(y_{(n_0-n_1)}, g_1^{-1} \cdot z)\right] & \text{if } i_0 = 0 \text{ and } i_1 = 1 \\ 1 + 4^{-1-\pi(n_1-n_0)}\left[1 + d_Y(y_{(n_1-n_0)}, g_0^{-1} \cdot z)\right] & \text{if } i_0 = 1 \text{ and } i_1 = 0 \end{cases}$$

Our space thus consists of two $\mathbb{Z}$-chains of isometric copies of $\widehat{G}$. Each copy is separated from the others, and the metric is set up to establish some rigidity of the chains. Distances between points in different chains are then used to encode the orbit of z under the G-action.

Lemma 5.9 *The definition gives a metric on* $G \times \mathbb{Z} \times \mathbb{Z}_2$ *whose completion is a metric on* X_z.

Proof: Note that the definition is symmetric, so we need only check the triangle inequality. If we have three points all in the same G-block, then d_z behaves like d'_G, so we are fine. If all three points are in distinct blocks, then each distance is in the interval $[1, 2]$, and any such triangle is legal. So we need only consider the case where two points are in the same

G-block, and the third is in a distinct block. Let our points be $x_g^{i,n}$, $x_{h_1}^{j,m}$, and $x_{h_2}^{j,m}$, where $(i,n) \neq (j,m)$. Set:

$$\begin{aligned}\delta_0 &= d_z(x_{h_1}^{j,m}, x_{h_2}^{j,m}) \\ \delta_1 &= d_z(x_g^{i,n}, x_{h_1}^{j,m}) \\ \delta_2 &= d_z(x_g^{i,n}, x_{h_2}^{j,m})\end{aligned}$$

Note that $\delta_0 = d'_G(h_1, h_2)$ is necessarily the shortest of these, so it will suffice to show that $|\delta_1 - \delta_2| \leq \delta_0$. We have three cases:

1. If $i = j$, then $n \neq m$, so δ_1 and δ_2 are both defined by the second case of the definition of d_z, and we have:

$$\begin{aligned}|\delta_1 - \delta_2| &= 4^{-|n-m|} \left| d'_G(g, h_1) - d'_G(g, h_2) \right| \\ &\leq \frac{1}{4} d'_G(h_1, h_2) \text{ since } d'_G \text{ is a metric} \\ &\leq d'_G(h_1, h_2) = \delta_0\end{aligned}$$

2. If $i = 0$ and $j = 1$ then δ_1 and δ_2 are both defined by the third case, and we get:

$$\begin{aligned}|\delta_1 - \delta_2| &= 4^{-1-\pi(n-m)} \left| d_Y(y_{(n-m)}, h_1^{-1} \cdot z) - d_Y(y_{(n-m)}, h_2^{-1} \cdot z) \right| \\ &\leq 4^{-1-\pi(n-m)} d_Y(h_1^{-1} \cdot z, h_2^{-1} \cdot z) \\ &\leq \frac{1}{4} d_Y(h_1^{-1} \cdot z, h_2^{-1} \cdot z) \\ &\leq d'_G(h_1, h_2) = \delta_0 \text{ by (*)}\end{aligned}$$

3. If $i = 1$ and $j = 0$ then δ_1 and δ_2 are defined by the fourth case and we get:

$$\delta_1 = 4^{-1-\pi(m-n)} \left[1 + d_Y(y_{(m-n)}, g^{-1} \cdot z)\right] = \delta_2$$

   so this case is fine.

Finally, we note that any d_z-Cauchy sequence must eventually be contained entirely within a single G-block, and is hence Cauchy with respect to d'_G. Since the completion of G under d'_G is $\widehat{G}$, we get that the completion of $G \times \mathbb{Z} \times \mathbb{Z}_2$ under d_z is $\widehat{G} \times \mathbb{Z} \times \mathbb{Z}_2 = X_z$. □

Now we check that the map $z \mapsto (X_z, d_z)$ is a reduction of E_G^Y to $\cong_i$.

Lemma 5.10 *If* $z_1 E_G^Y z_2$ *then* $(X_{z_1}, d_{z_1}) \cong_i (X_{z_2}, d_{z_2})$.

Proof: Let $h \in G$ witness that $z_1 E^Y_G z_2$, i.e. $z_2 = h \cdot z_1$. Consider the map $m_h : g \mapsto hg$. This is an isometry of (G, d'_G) since d'_G is left-invariant. Moreover, since G is dense in $\widehat{G}$, this map extends uniquely to an isometry of $(\widehat{G}, d'_G)$. We now define $f : X_{z_1} \to X_{z_2}$. We will use $x^{i,n}_g$ to denote a point is X_{z_1} and $y^{i,n}_g$ to denote a point in X_{z_2}. It will suffice to define f on $G \times \mathbb{Z} \times \mathbb{Z}_2$ since this is dense and our function will be continuous. So let

$$\begin{aligned} f(x^{0,n}_g) &= y^{0,n}_g \\ f(x^{1,n}_g) &= y^{1,n}_{hg} \end{aligned}$$

Since f restricted to a given G-block is either the identity or the map m_h, this is a continuous bijection. We check that it is an isometry. Let $x_1 = x^{i_1,n_1}_{g_1}$ and $x_2 = x^{i_2,n_2}_{g_2}$. Note that the first two cases in the definition of d_z are independent of z, so if $i_1 = i_2$ isometry will be guaranteed by the left-invariance of d'_G. So we may restrict attention to the case where $i_1 \neq i_2$; without loss of generality we may assume that $i_1 = 0$ and $i_2 = 1$. Then let $y_1 = f(x_1) = y^{0,n_1}_{g_1}$ and $y_2 = f(x_2) = y^{1,n_2}_{hg_2}$. We compute:

$$\begin{aligned} d_{z_2}(y_1, y_2) &= d_{z_2}(y^{0,n_1}_{g_1}, y^{1,n_1}_{hg_2}) \\ &= 1 + 4^{-1-\pi(n_1-n_2)} \left[1 + d_Y(y_{(n_1-n_2)}, (hg_2)^{-1} \cdot z_2)\right] \\ &= 1 + 4^{-1-\pi(n_1-n_2)} \left[1 + d_Y(y_{(n_1-n_2)}, g_2^{-1} \cdot h^{-1} \cdot z_2)\right] \\ &= 1 + 4^{-1-\pi(n_1-n_2)} \left[1 + d_Y(y_{(n_1-n_2)}, g_2^{-1} \cdot z_1)\right] \\ &= d_{z_1}(x^{0,n_1}_{g_1}, x^{1,n_2}_{g_2}) = d_{z_1}(x_1, x_2) \end{aligned}$$

Thus f is the required isometry from (X_{z_1}, d_{z_1}) to (X_{z_2}, d_{z_2}). □

Lemma 5.11 *If* $(X_{z_1}, d_{z_1}) \cong_i (X_{z_2}, d_{z_2})$ *then* $z_1 E^Y_G z_2$.

Proof: Suppose $f : X_{z_1} \to X_{z_2}$ is an isometry. We again use $x^{i,n}_g$ to denote a point in X_{z_1} and $y^{i,n}_g$ to denote a point in X_{z_2} (where $g \in \widehat{G}$ here). Note that since distances within $\widehat{G}$-blocks are at most 1, and distances between elements in distinct blocks are greater than 1, the map f must carry each block to another block, i.e., if $f(x^{i,n}_{g_1}) = y^{j_1,m_1}_{h_1}$ and $f(x^{i,n}_{g_2}) = y^{j_2,m_2}_{h_2}$, then $j_1 = j_2$ and $m_1 = m_2$. Moreover, f must respect the ordering of the chains as follows. Note that the cases in the definition produce distances in the intervals: $[0, 1]$, $[\frac{3}{2} + 4^{-|n_0-n_1|}, \frac{3}{2} + 2 \cdot 4^{-|n_0-n_1|}]$, or $[1 + 4^{-1-\pi(n_0-n_1)}, 1 + 2 \cdot 4^{-1-\pi(n_0-n_1)}]$. These intervals are disjoint, so f must send consecutive blocks to consecutive blocks and hence must send each $\mathbb{Z}$-chain to one of the $\mathbb{Z}$-chains, although the order may be reversed and the two chains may be interchanged. The chains must also move in the same way, in the sense that if $x_g^{0,n} \mapsto y_h^{i,m}$, then also $x_g^{1,n} \mapsto y_{h'}^{1-i,m}$.

Suppose f maps $\widehat{G} \times \{0\} \times \{0\}$ to $\widehat{G} \times \{j_0\} \times \{m_0\}$ for some j_0 and m_0, and thus induces an isometry $\hat{f}_0 : \widehat{G} \to \widehat{G}$ given by $f(x_g^{0,0}) = y_{\hat{f}_0(g)}^{j_0,m_0}$. Since G is comeager in $\widehat{G}$, its inverse image $\hat{f}_0^{-1}[G]$ is also comeager in $\widehat{G}$ so we have $G \cap \hat{f}_0^{-1}[G] \neq \emptyset$. So there are $g_0, h_0 \in G$ such that $\hat{f}_0(g_0) = h_0$. Thus

$$f(x_{g_0}^{0,0}) = y_{h_0}^{j_0,m_0}$$

By the same argument, we can find j_1, m_1 and $g_1, h_1 \in G$ such that

$$f(x_{g_1}^{1,0}) = y_{h_1}^{j_1,m_1}$$

Notice that we must have $j_1 = 1 - j_0$ and $m_0 = m_1$. Also, we must have either

$$f(x_{g_0}^{0,n}) = y_{h_0}^{j_0,m_0+n} \text{ and } f(x_{g_1}^{1,n}) = y_{h_1}^{1-j_0,m_0+n} \text{ for all } n$$

or

$$f(x_{g_0}^{0,n}) = y_{h_0}^{j_0,m_0-n} \text{ and } f(x_{g_1}^{1,n}) = y_{h_1}^{1-j_0,m_0-n} \text{ for all } n$$

By considering, say, the blocks $\widehat{G} \times \{0\} \times \{0\}$ and $\widehat{G} \times \{1\} \times \{0\}$, we see that the first case must hold if $j_0 = 0$, and the second case if $j_0 = 1$. That is, if the $\mathbb{Z}$-chains are interchanged, then their ordering is flipped. We can summarize this as follows, where $\oplus$ denotes addition in $\mathbb{Z}_2$:

$$(\forall n \in \mathbb{Z})(\forall i \in \mathbb{Z}_2)\left[f(x_{g_i}^{i,n}) = y_{h_i}^{j_0 \oplus i, m_0 + (-1)^{j_0} \cdot n}\right]$$

(Note that this suggests the semi-direct product $\mathbb{Z} \ltimes \mathbb{Z}_2$ given by the action $j \cdot n = (-1)^j \cdot n$ of $\mathbb{Z}_2$ on $\mathbb{Z}$; we will return to this in the next section). Now, for each $n \in \mathbb{Z}$ we must have

$$d_{z_2}(f(x_{g_0}^{0,n}), f(x_{g_1}^{1,0})) = d_{z_1}(x_{g_0}^{0,n}, x_{g_1}^{1,0})$$

From the definition,

$$d_{z_1}(x_{g_0}^{0,n}, x_{g_1}^{1,0}) = 1 + 4^{-1-\pi(n)}\left[1 + d_Y(y_n, g_1^{-1} \cdot z_1)\right]$$

Because of the order flipping in the case $j_0 = 1$, we also have:

$$d_{z_2}(f(x_{g_0}^{0,n}), f(x_{g_1}^{1,0})) = 1 + 4^{-1-\pi(n)}\left[1 + d_Y(y_n, h_{1-j_0}^{-1} \cdot z_2)\right]$$

But then, setting $h_2 = h_{1-j_0}$, we have

$$(\forall n \in \mathbb{Z}) \left(d_Y(y_n, g_1^{-1} \cdot z_1) = d_Y(y_n, h_2^{-1} \cdot z_2)\right)$$

Then, since the y_n's are dense, for any $\epsilon > 0$ we can find an n with $d_Y(y_n, g_1^{-1} \cdot z_1) < \frac{\epsilon}{2}$, so also $d_Y(y_n, h_2^{-1} \cdot z_2) < \frac{\epsilon}{2}$, giving $d_Y(g_1^{-1} \cdot z_1, h_2^{-1} \cdot z_2) < \epsilon$. We thus have

$$g_1^{-1} \cdot z_1 = h_2^{-1} \cdot z_2$$

So, letting $g_2 = h_2 g_1^{-1}$, we have $z_2 = g_2 \cdot z_1$, so that $z_1 E_G^Y z_2$, which was our goal. □

Thus, $z \mapsto (X_z, d_z)$ is a reduction as desired. We lastly check that we can produce codes for these spaces in a Borel way.

Lemma 5.12 *The map $z \mapsto (X_z, d_z)$ is Borel-measurable as a map from Y to $\mathcal{M}$.*

Proof: Fix a countable dense subset of G (and thus of $\widehat{G}$),

$$G_0 = \{g_k : k \in \omega\}$$

and fix a bijection $(i, n, k) \mapsto \langle i, n, k\rangle$ of ω^3 with ω. We code (X_z, d_z) as an array in $\mathcal{M}$ by sending $z \mapsto \langle d_{i,j}\rangle_{i,j\in\omega}$ where

$$d_{\langle i,n.k\rangle,\langle j,m,l\rangle} = d_z(x_{g_k}^{i,n}, x_{g_l}^{j,m})$$

If we have a sequence $\langle z_i\rangle \to z$, then we have that $\langle g_k \cdot z_i\rangle \to g_k \cdot z$ for each k by the continuity of the action, and so $\langle d_Y(y_n, g_k \cdot z_i)\rangle \to d_Y(y_n, g_k \cdot z)$ for each n and k. Thus the codes for the (X_{z_i}, d_{z_i})'s will approach the code for (X_z, d_z). This shows that the map is in fact a continuous reduction. □

We thus have proved:

Theorem 5.13 *Let G be a Polish group and E_G^X a Borel G-space. Then $E_G^X \leq_B \cong_i$.*

Notice that for a given E_G^X, all the metric spaces produced by the embedding are homeomorphic to $\widehat{G} \times \mathbb{Z} \times \mathbb{Z}_2$. Non-isometry was established by producing different distance sets. We needed to use two $\mathbb{Z}$-chains in order to encode a point $z \in Y$ via its distances from the points y_n. If we had an action on a space Y which can be embedded continuously into $\mathbb{R}$, we could have done the construction using only two copies of $\widehat{G}$ by encoding z directly into the metric, rather than z's distances from the y_n's.

5.2 Isometry Restricted to Special Classes

We will use the construction from the previous section to get lower bounds on the complexity of isometry restricted to various special classes of metric spaces. Several other special cases are considered in [11]; for an overview of these and other cases see [3]. Two types of lower bounds will be of interest. First, we will be able to show that isometry of certain classes is not classifiable by countable structures by embedding turbulent group actions (see the introduction for the definition of turbulence and classifiability by countable structures), and thus show that these problems are strictly more difficult than graph isomorphism. Second, we will show that isometry of some other classes is not concretely classifiable by embedding the equivalence relation E_0.

In their paper [11], Gao and Kechris show that isometry of discrete Polish metric spaces, ultrametric Polish metric spaces, and zero-dimensional locally compact Polish metric spaces are all bireducible with graph isomorphism. They ask about general zero-dimensional spaces, and homogeneous locally compact spaces. We will be able to show that isometry of zero-dimensional spaces is strictly more complex than graph isomorphism, and that isometry of homogeneous locally compact spaces is at least as complicated as graph isomorphism (this will be proved in the next chapter).

The main idea used here is that the spaces we constructed above for a given G-action have many of the topological properties of the group G. First, note that in the case that G has a complete left-invariant metric (for instance, if G is locally compact or abelian), then the metric space constructed will have the topology of the Polish group $G \times \mathbb{Z} \times \mathbb{Z}_2$. We thus get:

Theorem 5.14 *Let E_G^X be a Borel G-space, where G has a complete left-invariant metric. Then $E_G^X \leq_B \cong_i \restriction$ (Polish group topologies).*

Noting that there are abelian Polish groups with turbulent actions (for instance the density ideal (I_d, Δ) discussed below), we get

Corollary 5.15 *Isometry of Polish group metrics is not classifiable by countable structures.*

The metrics constructed in this case will not in general be left-invariant (although this group will have a complete left-invariant metric). We will consider invariant metrics below.

Suppose, in addition, that G is a zero-dimensional Polish group. Then the resulting space will also be zero-dimensional, so we have

Theorem 5.16 *Let E_G^X be a Borel G-space, where G is zero-dimensional and has a complete left-invariant metric. Then $E_G^X \leq_B \cong_i\restriction$ (zero-dimensional Polish group topologies).*

Again taking $G = (I_d, \Delta)$, which is zero-dimensional (see below), we get

Corollary 5.17 *Isometry of zero-dimensional Polish metric spaces is not classifiable by countable structures.*

We will now consider metric spaces with rich isometry groups. Recall that a metric space is said to be *homogeneous* if its isometry group acts transitively, i.e., for any two points x and y, there is an isometry carrying x to y. A stronger property is the following:

Definition 5.18 *A metric space is ultra-homogeneous if any partial isometry between finite subsets of it can be extended to an isometry of the whole space.*

This is equivalent to saying that if $\langle x_i \rangle_{i \leq n}$ and $\langle y_i \rangle_{i \leq n}$ are two finite sets of points such that for all $i, j \leq n$ we have $d(x_i, x_j) = d(y_i, y_j)$, then there is an isometry f of the space such that for all $i \leq n$ we have $f(x_i) = y_i$. Such spaces are determined up to isometry by their n-point spectra as defined in Chapter 4.

One way to establish homogeneity is to produce a Polish group with a left-invariant metric, for then left-multiplication by any group element will be an isometry. Thus, given any g_1 and g_2 in the group, left multiplication by $g_2 g_1^{-1}$ will be an isometry sending g_1 to g_2. We will first modify our construction to produce (for certain groups) metric spaces which are in fact Polish groups with left-invariant and two-sided invariant metrics. Producing ultra-homogeneous spaces will require further modification and restrictions on G.

5.3 Producing invariant metrics

Let us turn now to the case of abelian Polish groups. Any abelian Polish group has a complete two-sided invariant metric (in fact, any compatible left-invariant metric will necessarily be two-sided invariant, and will be complete). Some of what follows should be applicable to groups with complete invariant metrics, but we seem to actually require commutativity for part of the argument. Let us fix a Polish G-space E_G^Y as before, where again we may assume Y is compact, d_G is an invariant metric for G with $d_G \leq 1$, and d_Y is a metric on Y with $d_Y \leq 1$. Let d'_G be again defined as in Definition 5.6. Fix a point $z \in Y$. We will now define a Polish metric space (X_z, d_z) such that X_z is in fact the Polish group $(G \times \mathbb{Z}) \ltimes \mathbb{Z}_2$, the semi-direct product given by the action $(g,n) \mapsto (g^{-1}, -n)$ of $\mathbb{Z}_2$ on $G \times \mathbb{Z}$, and such that the metric d_z is a complete, compatible left-invariant metric. This will guarantee that the metric space produced is homogeneous.

As before, X_z will have the underlying set

$$x_z = \{x_g^{i,n} : g \in G, i \in \mathbb{Z}_2, n \in \mathbb{Z}\}$$

Definition 5.19 *For $g_0, g_1 \in G, n_0, n_1 \in \mathbb{Z}, i_0, i_1 \in \{0,1\}$, let $d_z(x_{g_0}^{i_0,n_0}, x_{g_1}^{i_1,n_1}) =$*

$$\begin{cases} d'_G(g_0,g_1) & \text{if } i_0 = i_1 \text{ and } n_0 = n_1 \\ \frac{3}{2} + 4^{-|n_0-n_1|}\left[1 + d'_G(g_0,g_1)\right] & \text{if } i_0 = i_1 \text{ and } n_0 \neq n_1 \\ 1 + 4^{-1-\pi(n_0-n_1)}\left[1 + d_Y(y_{(n_0-n_1)}, (g_0 g_1^{-1})\cdot z)\right] & \text{if } i_0 = 0 \text{ and } i_1 = 1 \\ 1 + 4^{-1-\pi(n_1-n_0)}\left[1 + d_Y(y_{(n_1-n_0)}, (g_1 g_0^{-1})\cdot z)\right] & \text{if } i_0 = 1 \text{ and } i_1 = 0 \end{cases}$$

Verifying that this is a metric is similar to the previous case, noting that for invariant metrics we have $d_G(g_1^{-1}, g_2^{-1}) = d_G(g_1, g_2)$. Verifying that the map $z \mapsto (X_z, d_z)$ is an embedding of E_G^Y is also the same as before. The difference now is in the left-invariance of the metric. Consider the group $(G \times \mathbb{Z}) \ltimes \mathbb{Z}_2$ where multiplication is given by

$$(g,n,i)\cdot(h,m,j) = (g \cdot h^{(-1)^i}, n + (-1)^i \cdot m, i \oplus j)$$

We will check that d_z is left-invariant under this multiplication.

Lemma 5.20 *Let (X_z, d_z) be as above. Then for all $h, g_1, g_2 \in G$, for all $m, n_1, n_2 \in \mathbb{Z}$ and for all $j, i_1, i_2 \in \mathbb{Z}_2$ we have*

$$d_z(x_{hg_1^{(-1)^j}}^{j\oplus i_1, m+(-1)^j\cdot n_1}, x_{hg_2^{(-1)^j}}^{j\oplus i_2, m+(-1)^j\cdot n_2}) = d_z(x_{g_1}^{i_1,n_1}, x_{g_2}^{i_2,n_2})$$

Proof: For simplicity, let

$$(\tilde{g}_k, \tilde{n}_k, \tilde{i}_k) = (hg_k^{(-1)^j}, m + (-1)^j \cdot n_k, j \oplus i_k)$$

for $k = 1, 2$. We consider three possible cases in the definition of d_z.

1. If $i_1 = i_2$ and $n_1 = n_2$ then $\tilde{i}_1 = \tilde{i}_2$ and $\tilde{n}_1 = \tilde{n}_2$, so

$$\begin{aligned} d_z(x_{\tilde{g}_1}^{\tilde{i}_1,\tilde{n}_1}, x_{\tilde{g}_2}^{\tilde{i}_2,\tilde{n}_2}) &= d'_G(hg_1^{(-1)^j}, hg_2^{(-1)^j}) \\ &= d'_G(g_1^{(-1)^j}, g_2^{(-1)^j}) \\ &= d'_G(g_1,g_2) = d_z(x_{g_1}^{i_1,n_1}, x_{g_2}^{i_2,n_2}) \end{aligned}$$

2. If $i_1 = i_2$ and $n_1 \neq n_2$ then we have $\tilde{\imath}_1 = \tilde{\imath}_2$ and $|\tilde{n}_1 - \tilde{n}_2| = |n_1 - n_2|$. So we get $d_z(x_{\tilde{g}_1}^{\tilde{\imath}_1,\tilde{n}_1}, x_{\tilde{g}_2}^{\tilde{\imath}_2,\tilde{n}_2}) = d_z(x_{g_1}^{i_1,i_2}, x_{g_2}^{i_2,n_2})$ as in the previous case.

3. If $i_1 \neq i_2$ then $\tilde{\imath}_1 \neq \tilde{\imath}_2$. We may assume $i_0 = 0$.

   (a) If $j = 0$ then we have $\tilde{n}_1 - \tilde{n}_2 = n_1 - n_2$ so we get

$$\begin{aligned} d_z(x_{\tilde{g}_1}^{\tilde{\imath}_1,\tilde{n}_1}, x_{\tilde{g}_2}^{\tilde{\imath}_2,\tilde{n}_2}) &= 1 + 4^{-1-\pi(\tilde{n}_1-\tilde{n}_2)} \left[1 + d_Y(y_{\tilde{n}_1-\tilde{n}_2}, \tilde{g}_1\tilde{g}_2^{-1} \cdot z)\right] \\ &= 1 + 4^{-1-\pi(n_1-n_2)} \left[1 + d_Y(y_{n_1-n_2}, hg_1g_2^{-1}h^{-1} \cdot z)\right] \\ &= 1 + 4^{-1-\pi(n_1-n_2)} \left[1 + d_Y(y_{n_1-n_2}, g_1g_2^{-1} \cdot z)\right] \\ &= d_z(x_{g_1}^{i_1,n_1}, x_{g_2}^{i_2,n_2}) \end{aligned}$$

   (b) If $j = 1$ then $\tilde{n}_2 - \tilde{n}_1 = n_1 - n_2$ so we get

$$\begin{aligned} d_z(x_{\tilde{g}_1}^{\tilde{\imath}_1,\tilde{n}_1}, x_{\tilde{g}_2}^{\tilde{\imath}_2,\tilde{n}_2}) &= 1 + 4^{-1-\pi(\tilde{n}_2-\tilde{n}_1)} \left[1 + d_Y(y_{\tilde{n}_2-\tilde{n}_1}, \tilde{g}_2\tilde{g}_1^{-1} \cdot z)\right] \\ &= 1 + 4^{-1-\pi(n_1-n_2)} \left[1 + d_Y(y_{n_1-n_2}, hg_2^{-1}g_1h^{-1} \cdot z)\right] \\ &= 1 + 4^{-1-\pi(n_1-n_2)} \left[1 + d_Y(y_{n_1-n_2}, g_1g_2^{-1} \cdot z)\right] \\ &= d_z(x_{g_1}^{i_1,n_1}, x_{g_2}^{i_2,n_2}) \end{aligned}$$

So we get left-invariance in all three cases. □

As before, the embedding is clearly Borel, and so we get:

Theorem 5.21 *Let G be an abelian Polish group and E_G^X a Borel G-space. Then we have $E_G^X \leq_B \cong_i \restriction$ (left-invariant metrics for Polish groups).*

Taking $G = (I_d, \Delta)$ again, we get:

Corollary 5.22 *Isometry of homogeneous Polish metric spaces is not classifiable by countable structures.*

Also, if we take $G = (I_{\mathrm{FIN}}, \Delta)$, which generates the equivalence relation E_0, and note that the spaces produced are discrete (and hence locally compact), we get:

Corollary 5.23 *Isometry of homogeneous discrete Polish metric spaces (and hence of homogeneous locally compact Polish metric spaces) is not concretely classifiable.*

In the next chapter we will improve this result to show that isometry of homogeneous discrete Polish metric spaces is bireducible with graph isomorphism.

This should be contrasted with Corollary 5.8 of [11], that the isometry of *pseudo-connected* homogeneous locally compact Polish metric spaces is concretely classifiable. We can in fact strengthen the contrast. Let E_0 be represented (essentially) as the orbit equivalence relation of a $\mathbb{Z}$-action (we can use the odometer map on 2^ω), and let $G = \mathbb{Z}$ be given the metric $d_G(n, m) = \frac{1}{2}\left[1 + \frac{|n-m|}{1+|n-m|}\right]$ for $n \neq m$. Then (G, d_G) will be pseudo-connected. If we now let our space (X_z, d_Z) consist of two copies of $\mathbb{Z}$,

$$X_z = \{x_n^i : i \in \{0,1\}, n \in \mathbb{Z}\}$$

and define

$$\begin{aligned} d_z(x_n^i, x_m^i) &= d_G(n, m) \\ d_z(x_n^0, x_m^1) &= 1 + \varphi((n-m)\cdot z) \end{aligned}$$

where $\varphi(y) = \sum_{k\in\omega} \frac{y(k)}{2^{k+2}}$, then the resulting space will be homogeneous locally compact with two pseudo-connected components. A similar argument to that above shows that this is a reduction of E_0, so we thus get:

Corollary 5.24 *Isometry of homogeneous locally compact Polish metric spaces with two pseudo-connected components is not concretely classifiable.*

So Corollary 5.8, and hence Theorem 5.7, of [11] cannot be extended to the case of finitely many pseudo-components. The above construction can easily be modified to show that any countable abelian group action can be reduced to isometry of homogeneous discrete spaces with two pseudo-components. Louveau has been able to show this for ultra-homogeneous spaces:

Theorem 5.25 (Louveau) *Any countable abelian group action is reducible to isometry of ultra-homogeneous locally compact spaces with two pseudo-components.*

Here we have only been able to give lower bounds on the complexity in the given cases; the exact classification remains open. So we can restate two of the questions from [11]:

Question 5.26 *What is the exact complexity of the isometry of zero-dimensional Polish metric spaces?*

Based on the distance sets considered in the previous chapter, it seems reasonable to conjecture that this is as complicated as the isometry of arbitrary Polish metric spaces.

Question 5.27 *What is the exact complexity of the isometry of homogeneous locally compact Polish metric spaces?*

As noted, we will be able to show graph isomorphism is a lower bound, and we suspect that this is also an upper bound.

Towards the solution of Question 5.26 one may ask how complicated the action of a zero-dimensional Polish group may be. There are several known examples of universal Polish group actions, but none of these is given by a zero-dimensional group. So we may ask:

Question 5.28 *Is there a universal Polish group action given by the action of a zero-dimensional Polish group?*

There are two improvements to the above techniques that we we will briefly sketch. First, if G is a group such that every element has order 2, then we can produce two-sided invariant metrics for the group $G \times \mathbb{Z} \times \mathbb{Z}_2$ by replacing "$n_0 - n_1$" by "$|n_0 - n_1|$" throughout the definition of d_z. Second, if X is a space which embeds in $[0, 1]$, then we can produce metrics for the space $G \times \mathbb{Z}_2$. The representative case here is the action of (I, Δ) on 2^ω, where I is a Polishable ideal. From this we can get:

Corollary 5.29 *If G is a Polish group in which every element has order 2, then any E_G^X is reducible to isometry of invariant metrics for abelian Polish groups.*

Further slight modifications can be used to apply to other classes of metric spaces which we will not list here. We now turn to the case of ultra-homogeneous spaces.

5.4 Producing ultra-homogeneous spaces

Recall the definition of ultra-homogeneous spaces given in the first section. One might suspect that these spaces have a relatively simple classification because of their strong uniformity. In a certain sense they do: two ultra-homogeneous spaces are isometric if and only if they have the same set of n-point configurations for all $n \geq 2$. However, these sets are (potentially) quite complicated analytic sets, so they do not provide nice invariants from a descriptive set-theoretic point of view. We will in fact be able to show that the classification is complicated, in that it can not be classified by countable structures.

An ultra-homogeneous space is clearly homogeneous, so we will attempt to modify the homogeneous construction. Notice that in that construction, two isometric configurations of points must agree on the relative ordering of their G-blocks. The chief difficulty lies with G itself. For the spaces produced above, G itself must be ultra-homogeneous (in d'_G) for the resulting space to be ultra-homogeneous. So our first task will be to produce an ultra-homogeneous metric for a sufficiently complicated Polish group.

We will focus on the *density ideal* I_d which we have already used above. As the action of I_d on 2^ω via symmetric difference is turbulent, it will suffice to reduce this action to isometry of ultra-homogeneous Polish metric spaces in order to rule out classifiability by countable structures. This is a free action, which will be crucial to our method. Recall that:

$$I_d = \left\{ x \subseteq \omega : \lim_{n\to\infty} \frac{|x \cap n|}{n} = 0 \right\}$$

This is a *Polishable* ideal, i.e. it can be given a Polish topology compatible with the Borel structure it inherits as a subset of 2^ω such that it is a Polish group (with group addition being symmetric difference) under this topology. We will use a few facts from the theory of Polishable ideals, primarily their representation by *lower semi-continuous submeasures on* ω.

Definition 5.30 *A submeasure on* ω *is a map* $\varphi : P(\omega) \to [0, \infty]$ *such that:*

1. $\varphi(\emptyset) = 0$
2. $0 < \varphi(\{n\}) < \infty$ *for all* $n \in \omega$
3. $x \subseteq y \Longrightarrow \varphi(x) \leq \varphi(y)$
4. $\varphi(x \cup y) \leq \varphi(x) + \varphi(y)$

φ *is said to be lower semi-continuous (l.s.c.) if*

$$\varphi(x) = \lim_{n\to\infty} \varphi(x \cap n) \text{ for all } x$$

Given an l.s.c. submeasure φ, we define its *exhaustive ideal*

$$\mathrm{Exh}(\varphi) = \{x : \lim_{n\to\infty} \varphi(x \setminus n) = 0\}$$

A submeasure is said to be *finite* if $\varphi(x) < \infty$ for all x, and it is said to be *exhaustive* if $\mathrm{Exh}(\varphi) = P(\omega)$. We use the following representation from [30]:

Theorem 5.31 (Solecki) *An ideal I is Polishable if and only if there is a finite l.s.c submeasure φ such that $I = Exh(\varphi)$. In this case, the metric d_φ given by $d_\varphi(x, y) = \varphi(x\Delta y)$ is a complete invariant metric on I compatible with its (unique) Polish topology.*

We will now define such a submeasure for I_d. We will identify subsets of ω with their characteristic functions in 2^ω when this is convenient. Let $\rho : 2^\omega \to 2^\omega$ be given by:

$$\rho(x)(0) = 0$$

and for $k, m \geq 0$:

$$\rho(x)(2^m(2k + 1)) = x(m)$$

That is,

$$\rho(x) = \langle 0, x(0), x(1), x(0), x(2), x(0), x(1), x(0), x(3), \ldots\rangle$$

The main point is that if $x \neq y$ then $\rho(x)\Delta\rho(y)$ is infinite.

Definition 5.32 *For $x \subseteq \omega$, let*

$$\varphi(x) = \sup\left\{\frac{|x \cap 2^n|}{2^{n+1}} : n \in \omega\right\} + \sum_{n\in\omega} \frac{\rho(x)(n)}{2^{n+1}}$$

Note that $\varphi(x) \in [0, 1]$.

Lemma 5.33 *The function φ is a l.s.c. submeasure with $I_d = Exh(\varphi)$.*

Proof: First, note that the map ρ respects Boolean operations, so that the function

$$\varphi_1(x) = \sum_{n\in\omega} \frac{\rho(x)(n)}{2^{n+1}}$$

is an exhaustive submeasure (in fact, a measure). Thus it will suffice to show that the function

$$\varphi_0(x) = \sup\left\{\frac{|x \cap 2^n|}{2^{n+1}} : n \in \omega\right\}$$

is a l.s.c submeasure such that $I_d = \mathrm{Exh}(\varphi_0)$. It is easily seen to be a submeasure, and is lower semi-continuous because

$$(\forall\epsilon > 0)(\exists k)\left[\frac{|x \cap 2^k|}{2^{k+1}} \geq \sup\left\{\frac{|x \cap 2^n|}{2^{n+1}} : n \in \omega\right\} - \epsilon\right]$$

To check chat $I_d = \mathrm{Exh}(\varphi_0)$, we first note that if $x \in I_d$, then $\frac{|x\cap n|}{n} \to 0$, so if we fix an $\epsilon > 0$, then there is an N such that for all $k \geq 2^N$ we have $\frac{|x\cap k|}{2k} < \epsilon$. Consider $x \setminus k$ for

$k \geq 2^N$. We have, for $m < N$, that $\frac{|(x \setminus k) \cap 2^m|}{2^{m+1}} = 0$, whereas for $m \geq N$, this is less than ϵ. Taking the supremum we then have $\varphi_0(x \setminus k) < \epsilon$ whenever $k \geq 2^N$, so $x \in \mathrm{Exh}(\varphi_0)$.

Conversely, if $x \notin I_d$, there is an $\epsilon > 0$ such that there are infinitely many n with $\frac{|x \cap n|}{n} \geq \epsilon$. Fix such an ϵ. Now fix any m and let n_0 be large enough that $n_0 \cdot \epsilon \geq 2m$ and such that $\frac{|x \cap n_0|}{n_0} \geq \epsilon$. Let k be such that $2^{k-1} \leq n_0 < 2^k$. Then:

$$\frac{|(x \setminus m) \cap 2^k|}{2^{k+1}} \geq \frac{|(x \setminus m) \cap n_0|}{2^{k+1}} \geq \frac{n_0 \cdot \epsilon - m}{2^{k+1}} \geq \frac{n_0 \cdot \epsilon}{8n_0} \geq \frac{\epsilon}{8}$$

Thus $\varphi_0(x \setminus m) \geq \frac{\epsilon}{8}$ for any m, so $x \notin \mathrm{Exh}(\varphi_0)$. Hence $I_d = \mathrm{Exh}(\varphi_0) = \mathrm{Exh}(\varphi)$. □

Note that for $x \in I_d$, we have that $\varphi_0(x)$ defined above achieves its supremum at some value of n. This shows that φ_0 achieves only countably many values on I_d, and since this gives a metric on I_d we can see that $I - d$ is zero-dimensional.

We now show that the metric induced by φ will make I_d into an ultra-homogeneous space. First we see that φ will limit the number of isometric configurations:

Lemma 5.34 *If x and y are in I_d and $x \neq y$, then $\varphi(x) \neq \varphi(y)$.*

Proof: As just noted, for $x \in I_d$, since $\frac{|x \cap 2^n|}{2^{n+1}} \to 0$, the supremum in the definition of φ is actually achieved by some value of m. Let $m(x)$ be the least such m, so that we have

$$\varphi(x) = \frac{|x \cap 2^{m(x)}|}{2^{m(x)+1}} + \sum_{n \in \omega} \frac{\rho(x)(n)}{2^{n+1}}$$

Let $\sigma(x) \in 2^{<\omega}$ be the "binary expansion" of $\frac{|x \cap 2^{m(x)}|}{2^{m(x)+1}} \leq \frac{1}{2}$, i.e.

$$\frac{|x \cap 2^{m(x)}|}{2^{m(x)+1}} = \sum_{n < |\sigma(x)|} \frac{\sigma(x)(n)}{2^{n+1}}$$

Let $\tau(x)$ be defined as the coordinate-wise sum with left carry of $\sigma(x)$ and $\rho(x)$ (here we mean that the first $|\sigma(x)|$ digits of $\rho(x)$ should be listed, left-to-right, and considered as a binary number, and the same with $\sigma(x)$, and these added. The rest of $\rho(x)$ is concatenated and left unchanged, e.g. $\langle 011 \rangle + \langle 010 \ldots \rangle = \langle 101 \ldots \rangle$). This is well defined since $\rho(x)(0) = 0$ always, and if $\sigma(x)(0) = 1$ then all other digits of $\sigma(x)$ are 0. We thus have:

$$\varphi(x) = \sum_{n \in \omega} \frac{\tau(x)(n)}{2^{n+1}}$$

Now, note that for $x, y \in I_d$, x and y are not eventually 1, and so $\rho(x)$ and $\rho(y)$ are not eventually 1. Also, if $x \neq y$, then $\rho(x)$ and $\rho(y)$ differ on infinitely many coordinates. Since

$\sigma(x)$ and $\sigma(y)$ are finite strings, they affect only finitely-many coordinates in $\tau(x)$ and $\tau(y)$, respectively, so that we will have $\tau(x)$ differing from $\tau(y)$ on infinitely many coordinates as well; in particular $\tau(x) \neq \tau(y)$. This will then ensure (since neither of these is eventually 1) that $\varphi(x) \neq \varphi(y)$. □

Basically, we ensure ultra-homogeneity by simply limiting the possible repetitions of n-point configurations.

Lemma 5.35 *The Polish metric space* (I_d, d_φ) *is ultra-homogeneous, where* $d_\varphi(x, y) = \varphi(x\Delta y)$.

Proof: Suppose that $x_0, \ldots, x_n$ and $y_0, \ldots, y_n$ in I_d are such that

$$(\forall i < j \leq n)(d_\varphi(x_i, x_j) = d_\varphi(y_i, y_j))$$

Thus $\varphi(x_i\Delta x_j) = \varphi(y_i\Delta y_j)$. By the previous lemma, this implies that for all $i, j \leq n$ we have $x_i\Delta x_j = y_i\Delta y_j$.
Let $w = x_0\Delta y_0 \in I_d$ and define

$$f(x) = x\Delta w$$

This is an isometry of I_d since the metric is invariant, and we check:

$$f(x_i) = x_i\Delta w = x_i\Delta(x_o\Delta y_0) = (x_i\Delta x_0)\Delta y_0 = (y_i\Delta y_0)\Delta y_0 = y_i$$

So we have an isometry carrying each x_i to y_i as required. □

We are now ready to define the reduction. Recall that we are trying to reduce the action of I_d on 2^ω via symmetric difference. So we fix $z \in I_d$ and define (X_z, d_z). We will let X_z consist of two copies of I_d:

$$X_z = \{x^i_\alpha : i \in \{0, 1\}, \alpha \in I_d\}$$

and we set:

$$d_z(x^i_\alpha, x^j_\beta) = \begin{cases} \varphi(\alpha\Delta\beta) & \text{if } i = j \\ \frac{3}{2} + \sum_{n\in\omega} \frac{\rho(z\Delta\alpha\Delta\beta)(n)}{2^{n+1}} & \text{if } i \neq j \end{cases}$$

The proof that this defines a metric is much like before: The relationship between φ and

φ_0 is that of d'_G to d_G, and the fact that ρ respects Boolean operations ensures that:

$$\begin{aligned}\left|\sum_{n\in\omega}\frac{\rho(z\Delta\alpha\Delta\beta_1)(n)}{2^{n+1}}-\sum_{n\in\omega}\frac{\rho(z\Delta\alpha\Delta\beta_2)(n)}{2^{n+1}}\right| &\leq \sum_{n\in\omega}\frac{|\rho(z\Delta\alpha\Delta\beta_1)(n)-\rho(z\Delta\alpha\Delta\beta_2)(n)|}{2^{n+1}}\\ &= \sum_{n\in\omega}\frac{\left(\rho(z\Delta\alpha\Delta\beta_1)\Delta\rho(z\Delta\alpha\Delta\beta_2)\right)(n)}{2^{n+1}}\\ &= \sum_{n\in\omega}\frac{\rho(\beta_1\Delta\beta_2)(n)}{2^{n+1}}\end{aligned}$$

It is also easy to see that $z_1 \not E_{I_d} z_2 \implies (X_{z_1}, d_{z_1}) \not\cong_i (X_{z_2}, d_{z_2})$, since the orbits are coded more or less directly into the set of distances by ρ. There is a slight difficulty here, namely the case of eventually constant z's, for which the function ρ is not necessarily one-to-one. However, since $I_{\text{FIN}} \subseteq I_d$, these fall into only two orbits which can have overlapping distance sets. Since there are only a finite number of problematic orbits, we can simply redefine (X, d) in these two cases to be two other distinct ultra-homogeneous spaces.

We can also check that the map

$$f : \begin{cases} x^0_\alpha \mapsto x^0_\alpha \\ x^1_\alpha \mapsto x^1_{\alpha\Delta z_1\Delta z_2} \end{cases}$$

is an isometry of (X_{z_1}, d_{z_1}) and (X_{z_2}, d_{z_2}) when $z_1\Delta z_2 \in I_d$. Note that d_z is an invariant metric, viewing X_z as the Polish group $(I_d, \Delta)\times\mathbb{Z}_2$, so the space is homogeneous. It remains to check ultra-homogeneity.

Lemma 5.36 *The space (X_z, d_z) constructed above is ultra-homogeneous.*

Proof: Let $x_0, \ldots, x_n$ and $y_0, \ldots, y_n$ be given such that for all $i < j \leq n$ we have $d_z(x_i, x_j) = d_z(y_i, y_j)$. We may assume $x_0 = y_0$ by homogeneity. If all of the x_i lie in the same copy of I_d, then all of the y_i's must also, and in fact we must have $x_i = y_i$ for all $i \leq n$ as we saw in Lemma 5.35. In general, the x_i's and y_i's fall into two sets, those in the same copy of I_d as x_0 and those in the other copy. Here again the two sets in the same copy as x_0 must in fact be identical. Let $x_0 = x^j_\alpha$. For two points $x_i = x^{1-j}_{\beta_x}$ and $y_i = x^{1-j}_{\beta_y}$ in the other copy, we must have

$$\frac{3}{2}+\sum_{n\in\omega}\frac{\rho(z\Delta\alpha\Delta\beta_x)(n)}{2^{n+1}} = \frac{3}{2}+\sum_{n\in\omega}\frac{\rho(z\Delta\alpha\Delta\beta_y)(n)}{2^{n+1}}$$

This will ensure that $z\Delta\alpha\Delta\beta_x = z\Delta\alpha\Delta\beta_y$ (since if one of these is eventually 1, so is the other), and thus that $\beta_x = \beta_y$. Thus, both collections of points must be the same. □

We have thus proved:

Theorem 5.37 *Isometry of ultra-homogeneous Polish metric spaces is not classifiable by countable structures.*

We do not know an upper bound, though.

Question 5.38 *What is the exact complexity of the isometry of ultra-homogeneous Polish metric spaces?*

The technique and modifications presented in this chapter doubtless admit further refinements. It would be of interest to see how sharp the lower bounds produced actually are. For instance, it would be very interesting to see if there is any connection between the topological properties of a class of metric spaces and the properties of the Polish groups whose actions can be reduced to isometry of that class of spaces. We have no ideas for what such a relationship might be, if any.

Our technique, for instance, seems inapplicable to the question of the complexity of locally compact Polish metric spaces, which is perhaps the most interesting open question in this area. As it stands, our technique can only produce locally compact spaces in the case that G is locally compact. Since locally compact groups do not admit turbulent actions, this prevents a non-classifiability result. It also prevents us from using this technique to reduce actions of S_∞, which is a notable weakness in that we know such actions are reducible to isometry of locally compact spaces. It would certainly be reassuring to at least be able to close this gap.

Chapter 6

Homogeneous Spaces and Structures

In this chapter we continue our investigation of the complexity of the isometry relation for metric spaces with large isometry groups, specifically homogeneous and ultrahomogeneous spaces. We consider in particular the case of discrete and locally compact spaces and provide several exact classifications and other lower bounds. Our principal approach is to connect these isometry problems to the problem of classifying various classes of countable structures up to isomorphism. The classes of countable structures we will be interested in are those whose automorphism groups act transitively. We will also address several questions about such structures which are of independent interest, such as determining which classes of countable structures with transitive automorphism groups have classification problems of maximal complexity.

6.1 Homogeneous Discrete and Locally Compact Spaces

Recall that a metric space is said to *homogeneous* if its isometry group acts transitively on points. This usage should be distinguished from the model-theoretic usage which is a generally stronger property. When we refer to model-theoretic structures we shall use the term *transitive* to indicate that the automorphism group acts transitively on the underlying set of the structure.

We start by relating the isometry of homogeneous discrete metric spaces to the isomorphism of countable graphs with transitive automorphism groups.

Theorem 6.1 *The isomorphism relation on countable transitive connected graphs is Borel reducible to the isometry relation on homogeneous discrete metric spaces.*

Proof: The proof is essentially the same as showing that graph isomorphism is reducible to isometry of discrete metric spaces. Given a countable connected graph, we form the discrete metric space whose elements are the vertices of the graph and put on it the graph metric (i.e. the distance between two points is the length of the shortest path connecting them in the graph). Now we simply note that automorphisms of the graph induce isometries in the graph metric space, so that when the automorphism group of the original graph acts transitively, so too does the isometry group of the graph metric space. □

In section 6.3 we shall prove that isomorphism of transitive graphs is bireducible with graph isomorphism (Theorem 6.10). Since isometry of general discrete metric spaces is Borel reducible to graph isomorphism, we thus have:

Corollary 6.2 *Isometry of homogeneous discrete metric spaces is bireducible with graph isomorphism.*

This yields an exact classification in the case of homogeneous discrete spaces. Since discrete spaces are locally compact, we have the following lower bound:

Corollary 6.3 *Graph isomorphism is Borel reducible to isometry of homogeneous locally compact Polish metric spaces.*

This bound is probably sharp, but again as in the case of general locally compact spaces we do not have an exact upper bound.

Most of the work will be in the proof of Theorem 6.10 in Section 6.3. We defer this in order to consider the cases of ultra-homogeneous discrete and locally compact metric spaces.

6.2 Ultra-homogeneous Spaces

The techniques in this section will not involve countable structures, but will rely directly on metric space techniques. We begin with the discrete case. We first recall the equivalence relation F_2, equality of countable sets of reals, which is defined on the space $\mathbb{R}^\omega$ by

$$\langle x_n \rangle \ F_2 \ \langle y_n \rangle \Longleftrightarrow \{x_n : n \in \omega\} = \{y_n : n \in \omega\}$$

Our characterization is then:

Theorem 6.4 *Isometry of ultra-homogeneous discrete metric spaces is bireducible with F_2.*

Proof: The reduction of isometry of ultra-homogeneous discrete spaces to F_2 is simple. Observe that two ultra-homogeneous Polish metric spaces are isometric precisely when they have the same sets of n-point configurations for all $n \geq 2$. A discrete metric space is countable, so it contains only countable many n-point configurations for each n. These configurations are easily coded as reals, so that each set of n-point configurations can be coded by a countable set of reals. Then, the sequence of these codes for $n \geq 2$ can be coded by a countable set of reals, so that two spaces are isometric if and only if these two countable sets are equal.

To reduce F_2 to isometry of ultra-homogeneous discrete spaces we modify Katětov's construction of the Urysohn space (see [20]). The Urysohn space is an ultra-homogeneous space into which every Polish metric space can be embedded isometrically. First, we fix a homeomorphism ρ of $\mathbb{R}$ with the open interval (1,2):

$$\rho(x) = \frac{3}{2} + \frac{1}{2} \cdot \frac{x}{1+|x|}$$

Now let A be a countable set of reals. We will define the metric space (X_A, d_A). First, we set

$$A' = \{1\} \cup \rho[A] \subseteq [1,2)$$

We now define a sequence of metric spaces. We let (X_0, d_0) be the one-point space. Then, having defined (X_n, d_n) for some $n \in \omega$, we define (X_{n+1}, d_{n+1}) as follows. First, we set

$$X_{n+1} = X_n \sqcup E_A(X_n)$$

where

$$E_A(X) = \{f : X \to A' \text{ such that for all but finitely many } x \in X \text{ we have } f(x) = 1\ \}$$

We then define d_{n+1} by

$$\begin{aligned} d_{n+1}(x_1, x_2) &= d_n(x_1, x_2) \text{ for } x_1, x_2 \in X_n \\ d_{n+1}(f, x) &= f(x) \text{ for } f \in E_A(X_n), x \in X_n \\ d_{n+1}(f_1, f_2) &= 1 \text{ for } f_1 \neq f_2 \in E_A(X_n) \end{aligned}$$

As is the construction of the Urysohn space, this defines a metric space, where verification of the triangle inequality is immediate because all distances are in the interval $[1, 2)$. Since each of the functions in $E_A(X_n)$ has finite support and A' is countable, we will have that X_{n+1} is countable (and hence separable). Moreover, since all the distances are in the interval $[1,2)$, we have that the space (X_{n+1}, d_{n+1}) is discrete (hence complete). We also have that for each n, (X_n, d_n) is a subspace of (X_{n+1}, d_{n+1}). We lastly set

$$(X_A, d_A) = \bigcup_{n \in \omega} (X_n, d_n)$$

This is then a discrete Polish metric space. Note that the construction (up to isometry) is independent of the enumeration of A, so that the mapping $A \mapsto (X_A, d_A)$ is well-defined. That is, if $A_1 = A_2$ then we also have $(X_{A_1}, d_{A_1}) \cong_i (X_{A_2}, d_{A_2})$. For the converse, note that the set of distances in (X_A, d_A) will be equal to $\{0\} \cup A'$, and that $A'_1 = A'_2$ if and only if $A_1 = A_2$. Hence, if $A_1 \neq A_2$ then the distance sets of the two spaces will be different, and hence $(X_{A_1}, d_{A_1}) \not\cong_i (X_{A_2}, d_{A_2})$. Thus, our map is a reduction of F_2 to isometry, as desired.

We must lastly check that the spaces produced are in fact ultra-homogeneous. For this, we will show that the spaces have the *one-point extension property*:

Definition 6.5 *A metric space is said to have the one-point extension property if, whenever we are given two finite sets $\{x_1, \ldots x_n\}$ and $\{y_1, \ldots, y_n\}$, a partial isometry φ between them such than $\varphi(x_i) = y_i$ for $1 \leq i \leq n$, and another point x_{n+1}, there is a point y_{n+1} such that φ extends to a partial isometry with $\varphi(x_{n+1}) = y_{n+1}$.*

Ultra-homogeneity clearly implies the one-point extension property for Polish metric spaces, and a straightforward back-and-forth argument shows that if a space has this property then it is ultra-homogeneous.

The construction of (X_A, d_A) makes this property easy to verify. Given points $x_1, \ldots, x_{n+1}$ and $y_1, \ldots, y_n$ and a partial isometry, there will be some k with all of these points in X_k. There will then be an f in X_{k+1} which has the same distances relative to the y_n's as x_{n+1} does to the x_n's. □

Once again, we have that F_2 is a lower bound for the isometry relation on locally compact ultra-homogeneous Polish metric spaces. Here we are able to show that this is a precise characterization, by showing that F_2 is also an upper bound in the locally compact case. We begin with some preliminaries.

Let us recall from [11] the definition of a *pseudo-component* of a locally compact space. For a point x in a locally compact space X, we let $\rho(x)$ denote the radius of compactness of x, i.e.

$$\rho(x) = \sup\{r : B_r^{\mathrm{cl}}(x) \text{ is compact}\}$$

where $B_r^{\mathrm{cl}}(x)$ is the closed ball of radius r around the point x. Since the space is locally compact, we have $\rho(x) > 0$ for all x. Note that in a homogeneous space (and hence in an ultra-homogeneous space) the radius of compactness must be the same for all points, so that it makes sense here to refer to the radius of compactness of the space X as $\rho(X)$ (although we will not need to use this in what follows). We now define the binary relation R on X by

$$x \; R \; y \iff d(x, y) < \rho(x)$$

and let R^* be the transitive closure of R.. We then define the equivalence relation E on X by

$$x \; E \; y \iff x = y \;\vee\; (x \; R^* \; y \;\&\; y \; R^* \; x)$$

The pseudo-components of X are then the equivalence classes of E. As shown in [11], the map $x \mapsto \rho(x)$ is Lipschitz, and each pseudo-component is clopen, so there are at most countably many pseudo-components. A space with only one pseudo-component is said to be pseudo-connected. We also observe that

$$\rho(x) = \sup\{r : \overline{B_r(x)} \text{ is compact}\}$$

To see this, note that $\overline{B_r(x)} \subseteq B_r^{\mathrm{cl}}(x)$ so that if $B_r^{\mathrm{cl}}(x)$ is compact then so is $\overline{B_r(x)}$. On the other hand, if $\overline{B_r(x)}$ is compact, then for each $\epsilon > 0$ we have that $B_{r-\epsilon}^{\mathrm{cl}}$ is compact, so that the two suprema will be the same. This allows us to make the following observation:

$$\rho(x) > r \iff (\exists \delta > r)[\overline{B_\delta(x)} \text{ is compact}]$$

Also note that if $D \subseteq X$ is dense, then $\overline{B_\delta(x)} = \overline{D \cap B_\delta(x)}$, since points in $B_\delta(x)$ will have arbitrarily close points in $D \cap B_\delta(x)$. These observations will be useful to the calculations below.

Let the array $\langle d_{i,j} \rangle$ code the Polish metric space $\overline{\{x_i : i \in \omega\}}$, where $d(x_i, x_j) = d_{i,j}$ as before. We assume that this space is locally compact.

Lemma 6.6 *For $\delta > 0$ and $i \in \omega$, the set $\overline{B_\delta(x_i)}$ is compact if and only if the following holds:*

$$(\forall q \in \mathbb{Q}^+)(\exists s \in [\omega]^{<\omega})[(\forall k < |s|)[d_{i,s(k)} < \delta] \;\&\; (\forall j)[d_{i,j} < \delta \Longrightarrow (\exists k < |s|)[d_{j,s(k)} < q]]]$$

Proof: Fix a $\delta > 0$ and first suppose that $\overline{B_\delta(x_i)}$ is compact. Given $q \in \mathbb{Q}^+$, by total boundedness there are $y_0, \ldots y_{n-1}$ in $\overline{B_\delta(x_i)}$ such that

$$(\forall y \in \overline{B_\delta(x_i)})(\exists k < n)[d(y, y_k) < \frac{q}{2}$$

Also, for each y_k there is an x_{i_k} in $B_\delta(x_i)$ such that $d(y_k, x_{i_k}) < \frac{q}{2}$. Now let s be a sequence of length n such that $s(k) = i_k$ for $k < n$. We thus have $d_{i,s(k)} < \delta$. If j is such that $d_{i,j} < \delta$, then $x_j \in B_\delta(x_i)$ so there must be some k with $d(x_j, y_k) < \frac{q}{2}$, and hence $d_{j,s(k)} < q$.

Conversely, suppose the given property holds. We will show that $\overline{B_\delta(x_i)}$ is totally bounded. Given $\epsilon > 0$, let $q \in \mathbb{Q}^+$ be such that $q < \frac{\epsilon}{2}$, and let $s \in [\omega]^{<\omega}$ be a witness for q, so that for all $k < |s|$ we have $d_{i,s(k)} < \delta$ and for all j with $d_{i,j} < \delta$ we have some $k < |s|$ with $d_{j,s(k)} < q$. Then for any $y \in \overline{B_\delta(x_i)}$ there is an x_j with $d(y, x_j) < q$, and there is a $k < |s|$ with $d(x_j, x_{s(k)}) < q$, so that $d(y, x_{s(k)}) < \epsilon$. Thus, the set $\{x_{s(0)}, \ldots, x_{s(|s|-1)}\}$ witnesses total boundedness for ϵ. □

Lemma 6.7 *We have that x_i and x_j are in the same pseudo-component if and only if the following holds:*

$$(\exists i_0, \ldots, i_n)[i_0 = i \;\&\; i_n = j \;\&\; (\forall k < n)[d_{i_k, i_{k+1}} < \rho(x_{i_k})]] \;\&$$
$$(\exists j_0, \ldots, j_m)[j_0 = j \;\&\; j_m = i \;\&\; (\forall k < m)[d_{j_k, j_{k+1}} < \rho(x_{j_k})]]$$

Proof: If this condition holds, x_i and x_j are clearly in the same pseudo-component. Suppose conversely that x_i and x_j are in the same pseudo-component. Since $x_i \; R^* \; x_j$, we have a sequence of points $y_0, \ldots, y_n$ in the space with $y_0 = x_i$, $y_n = x_j$, and $d(y_k, y_{k+1}) < \rho(y_k)$ for each $k < n$. We wish to replace this sequence by a similar sequence where we use only x_i's. We can of course set $i_0 = i$ and $i_n = j$. Let:

$$\delta_0 = \rho(y_0) - d(y_0, y_1) > 0$$
$$\delta_1 = \rho(y_1) - d(y_1, y_2) > 0$$

Choose an ϵ such that $\epsilon < \min(\delta_0, \frac{\delta_1}{2})$ and pick i_1 such that $d(y_1, x_{i_1}) < \epsilon$. We will then have that

$$d(x_{i_0}, x_{i_1}) < \rho(x_{i_0})$$
$$d(x_{i_1}, y_2) < \rho(x_{i_1})$$

so that we may replace y_1 by x_{i_1} in our sequence. We may similarly find $i_2, \dots i_{n-1}$ as needed. The same argument handles the witnesses that $x_j \; R^* \; x_i$. □

We are now ready to prove the main definability lemma we will need.

Lemma 6.8 *There is a Borel-measurable function mapping an array $\langle d_{i,j} \rangle_{i,j}$ to another array $\langle d_{\langle n,i \rangle, \langle m,j \rangle} \rangle_{n,i,m,j}$ such that if $\langle d_{i,j} \rangle$ codes the space $X = \overline{\{x_i : i \in \omega\}}$ then $\langle d_{\langle n,i \rangle, \langle m,j \rangle} \rangle$ also codes this space $X = \overline{\{x_{n,i} : n, i \in \omega\}}$ and for each n we have that $X_n = \overline{\{x_{n,i} : i \in \omega\}}$ is a pseudo-component of X. In the case that X has infinitely many pseudo-components, we can also require that each one is enumerated only once.*

Proof: This follows directly from the two previous lemmas, which show that we can calculate the radius of compactness and determine when two elements are in the same pseudo-component in a Borel manner, along with the observation that $d_{i,j} < \rho(x_k)$ if and only if there is a $q \in \mathbb{Q}^+$ such that $d_{i,j} < q$ and $\overline{B_q(x_k)}$ is compact. It is then simply a matter of rearranging the indices to group together elements which are in the same pseudo-components. This suffices for spaces with infinitely many pseudo-components; otherwise we enumerate one of them infinitely often. □

We are now ready to prove our characterization.

Theorem 6.9 *Isometry of ultra-homogeneous locally compact Polish metric spaces is bireducible with F_2.*

Proof: We need to show that isometry is reducible to F_2. By a result of Hjorth (see [11]), isometry of locally compact Polish metric spaces with only finitely many pseudo-components is essentially countable, that is, it is reducible to a countable Borel equivalence relation. Every countable Borel equivalence relation is clearly reducible to F_2 by sending an element to its equivalence class, a countable set. We can thus fix a sequence of functions $\langle \rho_n \rangle_{n \in \omega}$ such that ρ_n reduces isometry of spaces with n pseudo-components to F_2 and moreover satisfies

$$\begin{aligned} X_1 \cong_i X_2 \quad &\iff \quad \rho_n(X_1) = \rho_n(X_2) \\ &\iff \quad \rho_n(X_1) \cap \rho_n(X_2) \neq \emptyset \end{aligned}$$

for X_1 and X_2 with n pseudo-components (where we also use $\rho_n(X)$ to denote the countable set it codes). For convenience, we also choose the sequence so that each ρ_n produces a subset of $[n, n+1)$.

Now, given an ultra-homogeneous locally compact space coded by the array $\langle d_{i,j}\rangle$, let $\langle X_n\rangle_{n\in\omega}$ be its pseudo-components as enumerated by the function from lemma 6.8. We will assume that X has infinitely many pseudo-components; this can be determined in a Borel way and it is straightforward to modify our reduction to also handle spaces with only finitely many pseudo-components. Now, for $n \in \omega$ let

$$\sigma_n(X) = \bigcup\{\rho_{n+1}(X_{i_0} \sqcup X_{i_1} \sqcup \cdots \sqcup X_{i_n}) : i_0 < i_1 < \cdots < i_n\}$$

where we are again identifying a countable sequence with the countable set it enumerates. By inter-weaving sequences we can produce a sequence enumerating the elements of $\sigma_n(X)$. Note that $\sigma_n(X)$ is a countable subset of $[n+1, n+2)$ and contains codes for all possible subspaces of X with $n+1$ pseudo-components. We then define our reducing function f by

$$f(X) = \bigcup_{n\in\omega} \sigma_n(X)$$

So $f(X)$ is a countable set of reals, and again we can produce a countable sequence rather than the countable set we have described. We claim that $X \cong_i Y$ if and only if $f(X) = f(Y)$, which establishes the theorem.

If $X \cong_i Y$, then (up to isometry and permutation of indexing) X and Y have the same subspaces with finitely many pseudo-components, so we have $\sigma_n(X) = \sigma_n(Y)$ for each n, and hence $f(X) = f(Y)$. Suppose conversely that $f(X) = f(Y)$. Since the ranges of the σ_n's are disjoint, we have that $\sigma_n(X) = \sigma_n(Y)$ for each n. Thus:

$$\bigcup\{\rho_{n+1}(X_{i_0} \sqcup \cdots \sqcup X_{i_n}) : i_0 < \cdots < i_n\} = \bigcup\{\rho_{n+1}(Y_{i_0} \sqcup \cdots \sqcup Y_{i_n}) : i_0 < \cdots < i_n\}$$

But recall that our functions ρ_n have the property that if

$$\rho_{n+1}(X_{i_0} \sqcup X_{i_1} \sqcup \cdots \sqcup X_{i_n}) \cap \rho_{n+1}(Y_{j_0} \sqcup Y_{j_1} \sqcup \cdots \sqcup Y_{j_n}) \neq \emptyset$$

then in fact

$$\rho_{n+1}(X_{i_0} \sqcup X_{i_1} \sqcup \cdots \sqcup X_{i_n}) = \rho_{n+1}(Y_{j_0} \sqcup Y_{j_1} \sqcup \cdots \sqcup Y_{j_n})$$

We therefore have that for each n,

$$\{\rho_{n+1}(X_{i_0} \sqcup \cdots \sqcup X_{i_n}) : i_0 < \cdots < i_n\} = \{\rho_{n+1}(Y_{i_0} \sqcup \cdots \sqcup Y_{i_n}) : i_0 < \cdots < i_n\}$$

Thus, in particular, for each n there are $i_0^n, \ldots, i_n^n$ and $j_0^n, \ldots, j_n^n$ such that

$$\begin{aligned} \rho_n(X_0 \sqcup \cdots \sqcup X_n) &= \rho_n(Y_{i_0^n} \sqcup \cdots \sqcup Y_{i_n^n}) \\ \rho_n(Y_0 \sqcup \cdots \sqcup Y_n) &= \rho_n(X_{j_0^n} \sqcup \cdots \sqcup X_{j_n^n}) \end{aligned}$$

Thus:

$$\begin{aligned} X_0 \sqcup \cdots \sqcup X_n &\cong_i Y_{i_0^n} \sqcup \cdots \sqcup Y_{i_n^n} \\ Y_0 \sqcup \cdots \sqcup Y_n &\cong_i X_{j_0^n} \sqcup \cdots \sqcup X_{j_n^n} \end{aligned}$$

Since each finite configuration of points in X (resp. Y) will occur in some $X_0 \sqcup \cdots \sqcup X_n$ (resp. $Y_0 \sqcup \cdots \sqcup Y_n$), we see that the same configuration occurs (up to isometry) in Y (resp. X). Thus, X and Y have the same n-point configurations for each n, and since they are ultra-homogeneous this suffices to establish that they are isometric. □

6.3 Isomorphism of Symmetric Graphs

We now return to our consideration of spaces with transitive automorphism groups. We begin with an analysis of graphs. In the theory of graphs there are two common notions of transitivity: vertex-transitivity and edge-transitivity. A vertex-transitive graph is one in which the automorphism group acts transitively on vertices, whereas an edge-transitive graph is one in which the automorphism group acts transitively on edges. It is more usual in model theory to axiomatize graphs so that the underlying set is the set of vertices and a symmetric binary relation tells which vertices are connected by an edge. To have a transitive automorphism group in this setting corresponds then to vertex-transitivity. Alternately, if we let the underlying set be the set of edges and use relations to indicate when two edges meet at a common vertex (which is technically more complicated), then having a transitive automorphism group corresponds to edge-transitivity. When we refer to graphs we will always assume they are axiomatized in the first manner. We say that two vertices are adjacent if they are joined by an edge.

We shall first concern ourselves with the case of vertex-transitive connected graphs. Recall that a first-order theory (or $\mathcal{L}_{\omega_1\omega}$ sentence) is said to be *Borel-complete* if its isomorphism problem reduces every isomorphism problem for first-order theories; this is equivalent to being universal for equivalence relations induced by Borel actions of the infinite symmetric group S_∞. Although the classes of countable structures we consider will not generally be axiomatizable by an $\mathcal{L}_{\omega_1\omega}$ sentence, we shall use the same term. We will use the well-known fact that the empty theory in the language whose signature consists of a single binary relation is Borel complete (see [9]).

Theorem 6.10 *Isomorphism of countable connected graphs with (vertex-) transitive automorphism groups is Borel-complete.*

Proof: Let $\mathcal{L}_0$ be the language whose signature contains a single binary relation symbol. We shall reduce isomorphism of countable $\mathcal{L}_0$-structures to isomorphism of connected vertex-transitive graphs. The main idea of the proof will be that Cayley graphs for countable groups are canonical vertex-transitive graphs. Our construction is based closely on Mekler's proof that the theory of nilpotent class 2 groups of prime exponent is Borel-complete (see [25]). Mekler begins by constructing a Borel map assigning to each $\mathcal{L}_0$-structure $\mathcal{A}$ a graph $G(\mathcal{A})$ in an isomorphism-preserving way, so that if $\mathcal{A}_1 \cong \mathcal{A}_2$ then $G(\mathcal{A}_1) \cong G(\mathcal{A}_2)$. This graph has the following three properties (of which we will need only the first here):

1. If $v_1 \neq v_2$ are two vertices, then there is a vertex v_3 which is adjacent to v_1 but not to v_2.

2. Any two vertices have at most one common adjacent vertex.

3. If two vertices are adjacent then they have no common adjacent vertex.

We can also require that this graph be infinite.

We start with an $\mathcal{L}_0$-structure $\mathcal{A}$ and let $\langle v_i \rangle_{i\in\omega}$ enumerate the vertices of $G(\mathcal{A})$. Let H be the group freely generated by the vertices of $G(\mathcal{A})$, except that we let adjacent vertices commute. That is, if $\langle g_i \rangle_{i\in\omega}$ are generators of the free group $\mathbb{F}_\omega$, then

$$H = \mathbb{F}_\omega / \{g_i g_j g_i^{-1} g_j^{-1} : \ v_i \text{ is adjacent to } v_j \text{ in } G(\mathcal{A})\}$$

Let $\mathcal{G}$ be the Cayley graph of H with the generators $\langle g_i \rangle$. Specifically, let N be the normal subgroup of $\mathbb{F}_\omega$ generated by $\{g_i g_j g_i^{-1} g_j^{-1} : \ v_i \text{ is adjacent to } v_j \text{ in } G(\mathcal{A})\}$. Nodes of $\mathcal{G}$ are left cosets of N in $\mathbb{F}_\omega$, and two nodes $w_1 N$ and $w_2 N$ are adjacent if there is a generator g_i such that $g_i w_1 N = w_2 N$ or $g_i w_2 N = w_1 N$. We can definably produce a code for this structure (that is, represent it as a structure with underlying set ω) in the following manner. First, fix an enumeration of $\langle w_i \rangle_{i\in\omega}$ of the words in $\mathbb{F}_\omega$ with the generators $\langle g_i \rangle$. For each coset of N, we can then pick the least i such that w_i is in this coset and take this as a representative of the coset (although it may be undecidable whether two integers index words in the same coset, that is irrelevant here). We then can enumerate these representatives, and define our binary relation of ω according to whether the corresponding cosets are adjacent in $\mathcal{G}$. Call the code for this graph $\mathcal{G}(\mathcal{A})$.

Let us note that the generators g_i are all in distinct cosets. Also, for later use note that we could instead form the directed Cayley graph, where an edge points from a vertex $w_1 N$ to another $w_2 N$ if there is a generator g_1 with $g_i w_1 N = w_2 N$.

We claim that the map $\mathcal{A} \mapsto \mathcal{G}(\mathcal{A})$ is the desired reduction. First, it is easy to check that each graph is vertex-transitive, for if we have two vertices $w_1 N$ and $w_2 N$ then the mapping φ defined by

$$\varphi(wN) = wNw_1^{-1}w_2 = ww_1^{-1}w_2N$$

will be an automorphism of $\mathcal{G}(\mathcal{A})$ sending $w_1 N$ to $w_2 N$.

Next, suppose that we have $\mathcal{A}_1 \cong \mathcal{A}_2$. Then the graphs $G(\mathcal{A}_1)$ and $G(\mathcal{A}_2)$ given by Mekler's construction are also isomorphic, so let f be an isomorphism between these two graphs. Then f induces a partial map φ from $\mathcal{G}(\mathcal{A}_1)$ to $\mathcal{G}(\mathcal{A}_2)$ given by $\varphi(g_i) = g_{f(i)}$ (more precisely, the cosets of these elements). We want to extend this to an isomorphism of the two graphs. Let N_1 and N_2 be the respective normal subgroups in the constructions of $\mathcal{G}(\mathcal{A}_1)$ and $\mathcal{G}(\mathcal{A}_2)$. We then let

$$\varphi(wN_1) = \widetilde{w}N_2$$

where, if $w = g_{i_n}^{\sigma_n} \cdots g_{i_0}^{\sigma_0}$ then $\widetilde{w} = g_{f(i_n)}^{\sigma_n} \cdots g_{f(i_0)}^{\sigma_0}$.

We see that φ is a bijection, and we check that it is well-defined. Note that the map $w \mapsto \widetilde{w}$ is an automorphism of $\mathbb{F}_\omega$ sending N_1 to N_2. Thus,

$$w_1 w_2^{-1} \in N_1 \Longleftrightarrow \widetilde{w_1}\widetilde{w_2}^{-1} \in N_2$$

To see that φ is an isomorphism, suppose that $w_1 N_1$ and $w_2 N_1$ are adjacent in $\mathcal{G}(\mathcal{A}_1)$, say $g_k w_1 N_1 = w_2 N_2$. We then have that $\widetilde{(g_k w_1)} N_2 = \widetilde{w_2} N_2$. But $\widetilde{(g_k w_1)} = g_{f(k)} \widetilde{w_1}$ so we have that $g_{f(k)} \varphi(w_1 N_1) = \varphi(w_2 N_1)$. The reverse direction is identical, so that we have the vertices $w_1 N_1$ and $w_2 N_1$ adjacent in $\mathcal{G}(\mathcal{A}_1)$ if and only if the vertices $\varphi(w_1 N_1)$ and $\varphi(w_2 N_1)$ are adjacent in $\mathcal{G}(\mathcal{A}_2)$.

Finally, suppose that $\mathcal{G}(\mathcal{A}_1) \cong \mathcal{G}(\mathcal{A}_2)$. We will show that $\mathcal{A}_1 \cong \mathcal{A}_2$ by showing that $G(\mathcal{A}_1) \cong G(\mathcal{A}_2)$. To see this, it will suffice to see how to recover $G(\mathcal{A})$ (up to isomorphism) from the isomorphism class of $\mathcal{G}(\mathcal{A})$. Fix a vertex in $\mathcal{G}(\mathcal{A})$. By transitivity it does not really matter which vertex we use, so we may as well assume that it is the vertex corresponding to N. We then identify the vertices adjacent to the fixed node (which will be the $g_k^{\pm 1} N$, which are all distinct, although we will not be able to identify which is which). Let these vertices be enumerated as $\langle u_i \rangle_{i\in\omega}$. Consider the binary relation R on this set which says that two vertices are related if they are opposite corners in a square (i.e. a cycle of length 4) in $\mathcal{G}(\mathcal{A})$, i.e.

$$u_i \; R \; u_j \iff u_i \neq u_j \;\&\; (\exists a)(\exists b)[a \neq b \;\&\; (u_i \text{ and } u_j \text{ are each adjacent to both } a \text{ and } b)]$$

This relationship can be determined entirely from the isomorphism class of $\mathcal{G}(\mathcal{A})$. We claim that $u_i \; R \; u_j$ if and only if there are some k_1 and k_2 with $u_i = g_{k_1}^{\sigma_1} N$ and $u_j = g_{k_2}^{\sigma_2} N$ such that v_{k_1} is adjacent to v_{k_2} in $G(\mathcal{A})$ (although again we are not claiming to be able to reconstruct $G(\mathcal{A})$).

First, if there are such a k_1 and k_2 then g_{k_1} and g_{k_2} commute in H, so that u_i and u_j are opposite vertices in the square which also includes N and $g_{k_1}^{\sigma_1} g_{k_2}^{\sigma_2} N = g_{k_2}^{\sigma_2} g_{k_1}^{\sigma_1} N$. Suppose conversely that $u_i \; R \; u_j$. Let $u_i = g_{k_1}^{\sigma_1} N$ and $u_j = g_{k_2}^{\sigma_2}$. Let a and b be the other two vertices of the square. There are thus generators $g_{n_1}^{\tau_1}$, $g_{n_2}^{\tau_2}$, $g_{m_1}^{\rho_1}$, and $g_{m_2}^{\rho_2}$ witnessing this, i.e.

$$\begin{aligned} a &= g_{n_1}^{\tau_1} g_{k_1}^{\sigma_1} N = g_{n_2}^{\tau_2} g_{k_2}^{\sigma_2} N \\ b &= g_{m_1}^{\rho_1} g_{k_1}^{\sigma_1} N = g_{m_2}^{\rho_2} g_{k_2}^{\sigma_2} N \end{aligned}$$

We therefore have

$$g_{k_1}^{-\sigma_1} g_{n_1}^{-\tau_1} g_{n_2}^{\tau_2} g_{k_2}^{\sigma_2} \in N \text{ and } g_{k_1}^{-\sigma_1} g_{m_1}^{-\rho_1} g_{m_2}^{\rho_2} g_{k_2}^{\sigma_2} \in N$$

Words in N must have the sum of the exponents of each generator equal to 0, so in particular we must have $k_1 = k_2$, $k_1 = n_1$, or $k_1 = n_2$. If $k_1 = k_2$ we must have $\sigma_2 = -\sigma_1$, since otherwise we would have $u_i = u_j$. This would require that $n_1 = n_2 = k_1 = k_2$ and that $\sigma_2 - \sigma_1 + \tau_2 - \tau_1 = 0$, from which we conclude that $\tau_1 = -\sigma_1$, so that $a = N$. Similarly, if $k_1 \neq k_2$ and $k_1 = n_1$, then we must have $\tau_1 = -\sigma_1$ and again we have $a = N$.

The last possibility is that $k_1 \neq k_2$ and $k_1 = n_2$. Then we also have $k_2 = n_1$, $\sigma_1 = \tau_2$, and $\sigma_2 = \tau_1$. Making these substitutions, we find

$$g_{k_1}^{-\sigma_1} g_{k_2}^{-\sigma_2} g_{k_1}^{\sigma_1} g_{k_2}^{\sigma_2} \in N$$

From the definition of N, this implies that g_{k_1} and g_{k_2} commute in H, which means that v_{k_1} was adjacent to v_{k_2} in $G(\mathcal{A})$.

A similar argument for the case of b shows that either $b = N$ or v_{k_1} is adjacent to v_{k_2} in $G(\mathcal{A})$. Since we have that $a \neq b$, they can not both be equal to N so that v_{k_1} and v_{k_2} must be adjacent as we wished to show.

We now identify pairs $\{u_i, u_j\}$ of elements such that u_i is R-related to the same elements as u_j. This will identify pairs of the form $\{g_k N, g_k^{-1} N\}$, and will not identify any other pairs because property (1) of $G(\mathcal{A})$ ensures that for distinct vertices there will be a vertex adjacent to the first but not to the second (and vice-versa). We then form the graph whose vertices are the pairs just described, and we set two pairs adjacent to one another if each of the elements of the first is R-related to each of the elements of the second. Our analysis of the relation R then shows that the graph we have just formed will be isomorphic to $G(\mathcal{A})$. □

The same proof works for the case of directed graphs (digraphs) if instead of forming the Cayley graph of H we instead form the directed Cayley graph as described in the proof. We thus get:

Theorem 6.11 *Isomorphism of countable (weakly-) connected directed graphs with vertex-transitive automorphism groups is Borel-complete.*

We now consider graphs with even larger automorphism groups. We consider the following property of a graph which implies both vertex-transitivity and edge-transitivity.

Definition 6.12 *We say that a graph G is symmetric if, for any two edges (u_1, u_2) and (v_1, v_2), there is an automorphism φ of G such that $\varphi(u_1) = v_1$ and $\varphi(u_2) = v_2$.*

Thus, not only can every edge be mapped to any other edge by an automorphism, but we can pick the orientation. This property is in general stronger than either vertex-transitivity or edge-transitivity. The following theorem shows that the isomorphism problem is no simpler, though. This theorem will also be useful to us in the next section.

Theorem 6.13 *Isomorphism of countable symmetric connected graphs is Borel-complete.*

Proof: We will reduce isomorphism of the vertex-transitive graphs produced in the proof of Theorem 6.10 to isomorphism of symmetric graphs. We will in fact re-use part of the embedding produced there. Recall that given a countable $\mathcal{L}_0$-structure $\mathcal{A}$ we produced a vertex-transitive graph $\mathcal{G}(\mathcal{A})$ which was in fact the Cayley graph of a countable group. Note that these graphs continue to have property (1) of Mekler's graphs: if v_1 and v_2 are distinct vertices then there is a vertex v_3 adjacent to v_1 but not to v_2. Thus, we can apply the embedding which sent the intermediate graph $G(\mathcal{A})$ to the vertex-transitive graph $\mathcal{G}(\mathcal{A})$ to these resulting graphs. If we let $G \mapsto \mathcal{G}$ be the result of applying this embedding to one of our vertex-transitive graphs G, we will thus have that

$$G_1 \cong G_2 \Longleftrightarrow \mathcal{G}_1 \cong \mathcal{G}_2$$

It thus suffices to show that whenever G is one of our vertex-transitive graphs then its image $\mathcal{G}$ is symmetric.

We have that $\mathcal{G}$ is vertex-transitive as before, so to verify symmetry it will suffice to show the following:

- If v_0 is some fixed vertex (say the coset N) and v_1 and v_2 are two vertices adjacent to v_0 in $\mathcal{G}$, then there is an automorphism π of $\mathcal{G}$ such that $\pi(v_0) = v_0$ and $\pi(v_1) = v_2$.

Let $v_0 = N$. The vertices adjacent to v_0 will then be of the form $g_k^{\pm 1}N$ where g_k is a generator of $\mathbb{F}_\omega$.

We first consider the case where $v_1 = g_k N$ and $v_2 = g_k^{-1}N$ and produce an automorphism π_1 fixing v_0 and interchanging v_1 and v_2. Let π_1 be defined by

$$\pi_1(wN) = \widetilde{w}N$$

where, if $w = g_{i_n}^{\sigma_n} \cdots g_{i_0}^{\sigma_0}$, then $\widetilde{w} = g_{i_n}^{-\sigma_n} \cdots g_{i_0}^{-\sigma_0}$. We check that this is well-defined. If $w_1 N = w_2 N$ then $w_1^{-1} w_2 \in N$, so $w_1^{-1} w_2$ is a product of conjugates of words of the form $g_i g_j g_1^{-1} g_j^{-1}$. Then $\widetilde{w_1}^{-1}\widetilde{w_2}$ will be of the same form, so that $\widetilde{w_1}N = \widetilde{w_2}N$. The map is clearly a bijection fixing $v_0 = N$ and interchanging v_1 and v_2. Finally, we see that it is a graph automorphism, since if $g_i w_1 N = w_2 N$, then $g_i^{-1}\widetilde{w_1}N = \widetilde{w_2}N$.

We next exhibit an automorphism π_2 fixing v_0 and sending $v_1 = g_i N$ to $v_2 = g_j N$. Since the graph G was vertex-transitive, there is an automorphism φ of G sending v_i to v_j. We think of φ as a permutation of the indices of the vertices of G. We then define π_1 by

$$\pi_1(wN) = \widetilde{w}N$$

where, if $w = g_{i_n}^{\sigma_n} \cdots g_{i_0}^{\sigma_0}$, then $\widetilde{w} = g_{\varphi(i_n)}^{\sigma_n} \cdots g_{\varphi(i_0)}^{\sigma_0}$. As before, it is straightforward to check that π_2 is an automorphism of $\mathcal{G}$ fixing v_0 and sending v_1 to v_2. Last, we can combine automorphism of the previous two types to produce an automorphism fixing v_0 and sending any v_1 adjacent to it to any other v_2 adjacent to it, so $\mathcal{G}$ is symmetric. □

Once again we could form the directed Cayley graph with edges from wN to $g_k wN$ in our construction. Symmetry in the case of directed graphs only requires that we move similarly oriented edges to one another. A similar proof works here also, since we need only produce automorphisms fixing N and sending g_iN to g_jN, and do not need to interchange g_kN and $g_k^{-1}N$. We thus have:

Theorem 6.14 *Isomorphism of symmetric (weakly-) connected countable directed graphs is Borel-complete.*

Let us also note that if we continue to iterate this embedding, we can get Borel-completeness for classes of graphs with even greater symmetry, for instance those in which every square (4-cycle) can be mapped to any other square by an automorphism. As discussed in the last section, there is an upper limit to the amount of symmetry we can demand while still having a complicated isomorphism problem.

These results have a trivial consequence which we find interesting enough to state:

Corollary 6.15 *Classifying countable connected symmetric graphs up to isomorphism is as complicated as classifying arbitrary countable graphs.*

6.4 Other Transitive Countable Structures

Continuing along this path, one would like to know other examples of theories whose class of transitive countable models has a Borel-complete isomorphism problem. In this section we analyze the simplest theories possible, namely the empty theory in languages with various signatures, and determine when these have a Borel-complete isomorphism problem for their classes of countable models with transitive automorphism groups.

We have already seen one case for which this is true, the language $\mathcal{L}_0$ whose signature contains a single binary relation symbol. This is because the theory of graphs can be axiomatized with a single binary relation symbol and so the class of $\mathcal{L}_0$-structures with transitive automorphism groups contains the class of vertex-transitive graphs, whose isomorphism problem we saw to be Borel-complete in theorem 6.10. We thus get

Corollary 6.16 *The isomorphism problem for transitive $\mathcal{L}_0$-structures is Borel-complete.*

We can conclude more from this. Before proceeding, let us note that we should only consider signatures without constant symbols. Since a constant symbol must be interpreted by a single element, it immediately produces a definable element. A definable element is fixed by every automorphism, so the structure cannot have a transitive automorphism group (unless it contains only that one element). So unless stated otherwise, we shall assume our signatures contain no constant symbols.

First, notice that if we add relation or function symbols to a language whose transitive models are Borel-complete, we will still have a Borel-complete isomorphism problem, because we can look at those structures where the new symbols have trivial interpretations (for instance, nothing is related under new relation symbols, and function symbols uniformly map to the first coordinate). These structures will then have the same automorphism groups as their reducts to the original language.

Next, notice that a binary relation can be coded into an n-ary relation for $n \geq 3$ by simply having the relation depend only on the first two coordinates. This will not affect the automorphism group. Likewise, an irreflexive (or reflexive) binary relation can be coded in a binary function so as to preserve automorphism: Let $f(x, x) = x$, and for $x \neq y$ let $f(x, y) = x$ if xRy and let $f(x, y) = y$ if $x \not R y$. Since the binary relation for adjacency in graphs is irreflexive, we can thus code transitive graphs into transitive structures for a binary function symbol. We can also encode a binary function in an n-ary function for $n \geq 3$ in an isomorphism-preserving way by again letting the function depend only on the first two coordinates. Summarizing this, we have:

Corollary 6.17 *If $\mathcal{L}$ is a language whose signature contains an n-ary relation or function symbol for some $n \geq 2$, then the isomorphism problem for the class of $\mathcal{L}$-structures with transitive automorphism groups is Borel-complete.*

On the other hand, all that we can code in a transitive structure with only unary relations is a real, sine each relation is either satisfied by everything or by nothing (we are assuming a countable language; in general we can encode an element of $2^{|\mathcal{L}|}$). Similarly, if we allow only a single unary function symbol there are only countably many isomorphism types for transitive structures, according to whether the function splits into some number of finite cycles or whether it splits into some number of uniformly branching bi-infinite trees. Thus, if we have a language whose signature contains only unary relation symbols and a single unary function symbol, the isomorphism problem for its transitive countable models is concretely-classifiable.

This leaves only the case where we have at least two unary function symbols. We shall show that this is enough to produce a Borel-complete isomorphism problem for the transitive models. Let $\mathcal{L}_{u2}$ be the language whose signature contains only two unary function symbols, u_0 and u_1. We now prove:

Proposition 6.18 *The isomorphism problem for countable $\mathcal{L}_{u2}$-structures with transitive automorphism groups is Borel-complete.*

Proof: We shall reduce isomorphism of the symmetric graphs produced in the proof of theorem 6.13 to isomorphism of transitive $\mathcal{L}_{u2}$-structures. By the result of theorem 6.13 this will be sufficient. Given a symmetric graph G we produce an $\mathcal{L}_{u2}$-structure $\mathcal{A} = \mathcal{A}(G)$, and let f_0 and f_1 denote the interpretations of u_0 and u_1 in $\mathcal{A}$. Recall that the symmetric graph G is connected, infinite, and each vertex has infinite degree.

We first set out an indexing for the underlying set of $\mathcal{A}$ and define f_0. This function will be defined so that each point has countably many preimages and there are countably many connected components, so that the structure is partitioned into countably many bi-infinite countably-branching trees. We refer to these as *components*. To each point we associate the countable set of its preimages, which we refer to as the *block* below the point. Thus, two elements x and y are in the same block if $f_0(x) = f_0(y)$, and they are in the same component if there are $n, m \in \omega$ with $f_0^n(x) = f_0^m(y)$.

If we distinguish a node a_0 in a given component, we can enumerate the elements of the component in the following manner. If we look at the preimages of any node, the preimages of these, and so forth, we have essentially a copy of the Baire space ω^ω below the node. Relative to a_0, then we can label points in the component by pairs $(n, s) \in \omega \times \omega^{<\omega}$, where n refers to to how far "up" we start from a_0 (i.e. start from $f^n(a_0)$), and s determines a point in the copy of Baire space below this point, with the understanding that a_0 is along the leftmost branch (0-branch). This gives some points multiple labels; we identify $(n + 1, 0 \frown s)$ with (n, s). The node a_0 is then indexed by $(0, \langle\rangle)$ (as well as other labels). The function f_0 is thus defined in this component as:

$$f_0(n, s) = \begin{cases} (n, s \upharpoonright (k-1)) & \text{if } |s| = k > 0 \\ (n+1, \langle\rangle) & \text{if } s = \langle\rangle \end{cases}$$

Then, starting with a distinguished component, we associate to each node $a_0 = (n_0, s_0)$ (and hence to the block below the node) countably many components which we index $\langle a_0, n\rangle$ for $n \in \omega$, where $\langle a_0, 0\rangle$ is the initial component. Nodes in these new components are then labeled $\langle a_0, n, a_1\rangle$ for $n \neq 0$ and $a_1 = (n_1, s_1)$ as in the initial component. We continue in a similar manner: For each node $\langle a_0, n_0, a_1\rangle$ in one of these new components, except for the nodes with $a_1 = (0, \langle\rangle)$ (which are already associated to a countable set of components), we associate countably many new components and so forth. All of the components are distinct, each is a connected component of f_0 with f_0 behaving as in the initial component, and the underlying set of our structure $\mathcal{A}$ then consists of all the nodes enumerated in this fashion. Thus, points correspond to sequences of the form

$$\langle a_0, n_0, a_1, n_1, \ldots, a_{l-1}, n_{l-1}, a_l \rangle$$

where each a_i is a pair (k_i, s_i), each $n_i > 0$, and $a_i \neq (0, \langle\rangle)$ for $0 < i < l$. Again, we identify sequences where two a_i's label the same point. Two nodes are thus in the same component if their sequences agree (modulo this identification) up to n_{l-1}.

We next define the *index* of a node w, $\text{ind}(w)$. A node has index 0 if it is of the form $\langle a_0 \rangle$ or of the form $\langle a_0, n_0, \ldots, a_l \rangle$ with $a_l \neq (0, \langle\rangle)$. These are the nodes from which we formed new components; we call these *initial nodes*. For a node of the form $\langle a_0, \ldots, n_{l-1}, a_l \rangle$ with $l \geq 1$ and $a_l = (0, \langle\rangle)$ we let the index be n_{l-1}. Note that the initial component has all of its indices equal to 0, whereas each other component has a single node with non-zero index. This will not affect transitivity because we will be unable to determine these indices within the structure.

To each non-initial node we associate the initial node $\langle a_0, \ldots, a_{l-1} \rangle$, and associate each initial node to itself. We let $I(w)$ be the initial node associated to a node w. We refer to the set of nodes associated to a given initial node as a *group*. We also say that two blocks are in the same group if the nodes above them are in the same group. We will use the blocks below the nodes in a group to code the graph G into the structure $\mathcal{A}$ using f_1. Up to this point our construction has been independent of G.

Let $\langle v_i \rangle_{i \in \omega}$ enumerate the vertices in our symmetric graph G (according to its coding). Let $\langle k^i_n \rangle_{n \in \omega}$ enumerate in increasing order the indices of the vertices adjacent to to v_i in G, and let $\langle m^i_n \rangle_{n \in \omega}$ indicate where v_i occurs in $v_{k^i_n}$'s enumeration, i.e.

$$k^{k^i_n}_{m^i_n} = i \text{ for each } i \text{ and } n$$

We then also have

$$m^{k^i_n}_{m^i_n} = n \text{ for each } i \text{ and } n$$

This indexing will not have an essential effect because of edge-transitivity.

For a node $w = \langle a_0, \ldots, n_{l-1}, a_l \rangle$ in $\mathcal{A}$ with $a_l = (n, s)$ we write $w \frown j$ to denote the node $\langle a_0, \ldots, n_{l-1}, a'_l \rangle$ where $a'_l = (n, s \frown j)$, so that $w \frown j$ is the j-th node in the block below w. We now define f_1:

$$f_1(w \frown j) = \begin{cases} \left\langle I(w), k_j^{\text{ind}(w)}, \left(0, \left\langle m_j^{\text{ind}(w)} \right\rangle\right) \right\rangle & \text{if } k_j^{\text{ind}(w)} \neq 0 \\ I(w) \frown m_j^{\text{ind}(w)} & \text{if } k_j^{\text{ind}(w)} = 0 \end{cases}$$

This serves to define f_1 everywhere, since each node is in the block below some unique node w. For simplicity, we shall write

$$f_1(w \frown j) = \left\langle I(w), k_j^{\text{ind}(w)}, (0, \langle\rangle) \frown m_j^{\text{ind}(w)} \right\rangle$$

with the understanding that this collapses to $I(w) \frown m_j^{\text{ind}(w)}$ if $k_j^{\text{ind}(w)} = 0$. Note that f_1 is an involution:

$$\begin{aligned} f_1(f_1(w \frown j)) &= f_1\left(\left\langle I(w), k_j^{\text{ind}(w)}, (0, \langle\rangle) \frown m_j^{\text{ind}(w)} \right\rangle\right) \\ &= f_1\left(\left\langle I(w), k_j^{\text{ind}(w)}, (0, \langle\rangle) \right\rangle \frown m_j^{\text{ind}(w)}\right) \\ &= \left\langle I\left(\left\langle I(w), k_j^{\text{ind}(w)}, (0, \langle\rangle) \right\rangle\right), k_{m_j^{ind(w)}}^{ind(\langle I(w), k_j^{ind(w)}, (0, \langle\rangle)\rangle)}, \right. \\ &\qquad \left. (0, \langle\rangle) \frown m_{m_j^{ind(w)}}^{ind(\langle I(w), k_j^{ind(w)}, (0, \langle\rangle)\rangle)} \right\rangle \\ &= \left\langle I(w), k_{m_j^{ind(w)}}^{k_j^{ind(w)}}, (0, \langle\rangle) \frown m_{m_j^{ind(w)}}^{k_j^{ind(w)}} \right\rangle \\ &= \langle I(w), \text{ind}(w), (0, \langle\rangle) \frown j \rangle \\ &= w \frown j \end{aligned}$$

Let us clarify how f_1 behaves. In each group as defined above we have nodes with indices in ω; let the group have nodes $\langle w_i \rangle_{i \in \omega}$ with $\text{ind}(w_i) = i$. If we look at the blocks below these nodes, we will then have that f_1 connects some element in the block below the node w_i to some element in the block below the node w_j if and only if the vertex v_i is adjacent to the vertex v_j in the graph G. The k_n^i's and m_n^i's determine which elements in each block are connected (the n-th element in the i-th block is connected to the m_n^i-th element of the k_n^i-th block), but this is primarily a matter of bookkeeping and not an essential feature of the structure.

This defines f_1 and completes the construction of the $\mathcal{L}_{u2}$-structure $\mathcal{A}(G)$. We now check that this works, i.e. that $\mathcal{A}(G)$ has a transitive automorphism group and that $G_1 \cong G_2$ if and only if $\mathcal{A}(G_1) \cong \mathcal{A}(G_2)$.

First, suppose that we have two graphs $G_1 \cong G_2$. The key feature of the structure $\mathcal{A}(G)$ is that the only interactions between f_0 and f_1 occur within groups. Aside from this, $\mathcal{A}(G)$ is "freely generated" by f_0 and f_1; we could have progressively defined f_0 and f_1 starting from an initial node in such a way so as to never revisit components. Thus, so long as we define a mapping which is an isomorphism between groups we will have no problems, and we can progressively define an isomorphism iπ from $\mathcal{A}(G_1)$ to $\mathcal{A}(G_2)$ in the same manner.

We start by setting $\pi(0, \langle\rangle) = (0, \langle\rangle)$, mapping the distinguished node of $\mathcal{A}(G_1)$ to that of $\mathcal{A}(G_2)$. We shall define π in pieces. There are two important types of extensions we will need to make:

1. If π is defined on a node w, we must extend π to the block containing that node and to the other blocks in the same group as this one.

2. If π is defined on a node w, then we must extend π to the block below this node and to the other blocks in the same group.

Then, as long as we ensure that π respects f_0 (which will be true if we map blocks to blocks and the node above a given block to the node above the image of that block) and that π respects f_1 within groups, we can continue to extend π to an isomorphism.

We first consider extensions of type (1). Suppose we have $w_1 \in \mathcal{A}(G_1)$ with $\pi(w_1) = w_2$. We must then have $\pi(f_0(w_1)) = f_0(w_2)$. Let i_1 be the index of $f_0(w_1)$ and i_2 the index of $f_0(w_2)$. Let n_1 be such that w_1 is the n_1-th node below $f_0(w_1)$, i.e. $w_1 = f_0(w_1) \frown n_1$, and let n_2 be such that $w_2 = f_0(w_2) \frown n_2$. We use labels (i, n) to refer to nodes in the group of blocks containing w_1, where i is the index of the node's block and n is the nodes position within its block, so that for instance w_1 is labeled (i_1, n_1). We similarly label the nodes in the group of blocks containing w_2.

We now want to ensure that $\pi(f_1(i, n)) = f_1(\pi(i, n))$. We know that $f_1(i, n) = (k_n^i, m_n^i)$ and that $\pi(i_1, n_1) = (i_2, n_2)$. By the symmetry of G_1 and G_2 we can pick an isomorphism φ from G_1 to G_2 sending v_i to $\widetilde{v}_{\varphi(i)}$ with $\varphi(i_1) = i_2$ and $\varphi(k_{n_1}^{i_1}) = \widetilde{k}_{n_2}^{i_2}$ (we use v, k, and m to refer to G_1 and $\widetilde{v}$, $\widetilde{k}$, and $\widetilde{m}$ to refer to G_2). We now define

$$\pi(i, n) = (\varphi(i), \rho(i, n))$$

where $\rho(i, n)$ is the unique j such that $\widetilde{k}_j^{\varphi(i)} = \varphi(k_n^i)$ (such a j exists since $\widetilde{v}_{\varphi(i)}$ is adjacent to $\widetilde{v}_{\varphi(k_n^i)}$ in G_2, as v_i is adjacent to $v_{k_n^i}$ in G_1). In particular, $\rho(i_1, n_1) = n_2$ since $\widetilde{k}_{n_2}^{\varphi(i_1)} =$

$\widetilde{k}^{i_2}_{n_2} = \varphi(k^{i_1}_{n_1})$ by our choice of φ, so that $\pi(i_1, n_1) = (i_2, n_2)$ as required. We also have:

$$\begin{aligned}\pi(f_1(i,n)) &= \left(\varphi\left(k^i_n\right), \rho\left(k^i_n, m^i_n\right)\right)\\ f_1(\pi(i,n)) &= \left(\widetilde{k}^{\varphi(i)}_{\rho(i,n)}, \widetilde{m}^{\varphi(i)}_{\rho(i,n)}\right)\end{aligned}$$

We already know $\varphi\left(k^i_n\right) = \widetilde{k}^{\varphi(i)}_{\rho(i,n)}$ by our definition of ρ, so we need only check that $\rho\left(k^i_n, m^i_n\right) = \widetilde{m}^{\varphi(i)}_{\rho(i,n)}$, which amounts to showing that

$$\widetilde{k}^{\varphi\left(k^i_n\right)}_{\widetilde{m}^{\varphi(i)}_{\rho(i,n)}} = \varphi\left(k^{k^i_n}_{m^i_n}\right)$$

The right-hand side is equal to $\varphi(i)$ from the definitions of k and m. But our definition of ρ implies that the left-hand side is equal to $\widetilde{k}^{\widetilde{k}^{\varphi(i)}_{\rho(i,n)}}_{\widetilde{m}^{\varphi(i)}_{\rho(i,n)}} = \varphi(i)$. Thus our extension of π respects f_1.

For extensions of type (2) we proceed in a similar manner but with more flexibility. Suppose that $\pi(u_1) = u_2$; we then need only ensure that the block below u_1 maps to the block below u_2 and that the rest of the blocks in the same group are mapped appropriately. If we set $w_1 = u_1 \frown 0$ and $w_2 = u_2 \frown 0$ we may then proceed exactly as in the first type of extension.

We now explain the global construction of our isomorphism. Starting with our initial $\pi((0, \langle\rangle)) = (0, \langle\rangle)$, we successively extend π to all blocks and corresponding groups in the initial component of $\mathcal{A}(G_1)$. If we then consider the group of some block in the initial component, and consider the component of another block in that group, we can extend π to this new component as we did in the initial component. Since we always extend π a group at a time we are ensured of respecting f_1, and our extensions also respect f_0. Continuing in this manner we will eventually reach all components (as the structure is generated from an initial node by f_0 and f_1), so that the domain of π will be all of $\mathcal{A}(G_1)$. The same is true for the range of π, since as we extend the domain to a component of a node already in the domain, the range is extended to the component of the image of that node, and similarly for groups and blocks. Thus, π will be an isomorphism from $\mathcal{A}(G_1)$ to $\mathcal{A}(G_2)$.

For the converse, we explain how to recover G (up to isomorphism) from the isomorphism type of $\mathcal{A}(G)$. We start by picking a node in $\mathcal{A}(G)$; because it has transitive automorphism group, the choice of node will have no effect. By looking at the behavior of f_0 we are able to determine which nodes are in the same blocks within the structure. We can also identify which nodes are in the same group: since the graph G is connected, two nodes

u and w are in the same group if and only if there is a sequence $\langle a_0, b_0, a_1, b_1, \ldots a_n, b_n \rangle$ where $a_0 = u$, $b_n = w$, a_i and b_i are in the same block for each i, and $f_1(a_i) = b_{i+1}$.

We can thus identify the group of our chosen node and form the graph whose vertices are blocks in this group. We set the vertices corresponding to two of these blocks adjacent if there is an element in the first block which is mapped to an element of the second block by f_1. It is clear from the construction of $\mathcal{A}(G)$ that this graph will be isomorphic to G.

We lastly check that the structure $\mathcal{A}(G)$ has a transitive automorphism group. This is similar to the verification that for $G_1 \cong G_2$ we have $\mathcal{A}(G_1) \cong \mathcal{A}(G_2)$. Fix two nodes w_1 and w_2 of $\mathcal{A}(G)$; we will produce an automorphism π of $\mathcal{A}(G)$ such that $\pi(w_1) = w_2$.

We start by setting $\pi(w_1) = w_2$. We will then progressively extend π so that it respects f_0 and f_1 at all stages. As before we must see how to extend π from a node to the block containing it and the group of this block (as well as the nodes above), and how to extend π from a node to the block and group below it. Looking at the earlier verification, we see that although we started by mapping the distinguished node of $\mathcal{A}(G_1)$ to that of $\mathcal{A}(G_2)$, nowhere did we rely of this fact; we could have initialized π by mapping any node of $\mathcal{A}(G_1)$ to any other node. If we thus take $G_1 = G_2 = G$ in that argument, we can extend π to an automorphism of $\mathcal{A}(G)$ as desired. □

We have thus examined all possible signatures for a countable first-order language. The following theorem summarizes the results of this section.

Theorem 6.19 *Let $\mathcal{L}$ be a countable first-order language and let $\mathcal{K}$ denote the class of countable $\mathcal{L}$-structures which have transitive automorphism groups. Then the isomorphism problem for $\mathcal{K}$ is Borel-complete if and only if the signature of $\mathcal{L}$ contains no constant symbols and contains either an n-ary relation or function symbol for some $n \geq 2$ or contains at least two unary function symbols. In all other cases the isomorphism problem for $\mathcal{K}$ is concretely classifiable.*

6.5 Additional Comments and Questions

We should note a few differences between the problem we have just considered and the question of whether a given first-order language is Borel-complete (when we consider all countable structures, not just the transitive ones). First, in that case having constant symbols in the signature has no effect on the complexity. Second, unary relations have more power. Although finitely many unary relations still do not allow us to code more than a real into the structure, countably many do. With countably many unary relations $\langle R_i \rangle_{i \in \omega}$ we can code a real $x \in 2^\omega$ into an element a of the structure by setting

$$R_i(a) \Longleftrightarrow x(i) = 1$$

Our structure thus codes a countable set of reals, one for each element in the structure. The isomorphism problem then turns out to be bireducible with F_2.

The most striking difference is in the case of a single unary relation. Friedman and Stanley show (in [9]) that the isomorphism problem for the countable structures in the language with a single unary function symbol is Borel-complete, by showing that the theory of trees (which can be axiomatized with a single unary function symbol) is Borel-complete. For the class of transitive structures, though, we saw that the isomorphism problem was concretely-classifiable. This allows us to draw the following conclusion: The theory of graphs can not be axiomatized in a language with only one unary function symbol in a way that preserves automorphism groups.

Another observation we should make is that is necessary to produce graphs with infinite degree for each vertex in the proof of Theorem 6.10. This is the case because the isomorphism problem for countable connected locally-finite vertex-transitive graphs is in fact concretely-classifiable. This can be shown by a direct argument, but it is also a simple consequence of Corollary 5.8 of [11], which says that isometry of homogeneous pseudo-connected locally compact Polish metric spaces is concretely-classifiable. A locally-finite graph when given the graph metric becomes a pseudo-connected locally compact Polish metric space, and its isometry group is the automorphism group of the graph.

It seems an interesting problem to determine which theories, like that of graphs, continue to have complicated isomorphism problems when we restrict to the class of transitive models. We can ask:

Question 6.20 *What other first-order theories have an isomorphism problem for their transitive models which is as complicated as that for all of their countable models? Are there other natural examples where the isomorphism problem for transitive models is Borel-complete? Can this happen for a complete theory T?*

We should note that many natural theories are immediately ruled out because their structures have definable sets or elements. As noted earlier, having any non-trivial definable sets prevents a structure from having a transitive automorphism group. Thus, trees, groups, and most algebraic structures with complicated isomorphism problems are eliminated.

Another question concerns structures with larger automorphism groups. A structure is said to be *n-transitive* if its automorphism group acts transitively on n-tuples of distinct elements (so being 1-transitive is the same as having a transitive automorphism group). We can then ask the analogous question to theorem 6.19 for n-transitive structures:

Question 6.21 *For which countable first-order languages is the isomorphism problem for the class of n-transitive structures Borel-complete, for a given n?*

The strongest property we could consider along these lines would be having an n-transitive automorphism group for all $n \in \omega$. Here, though, we note that a structure having this property has an $\aleph_0$-categorical theory, since one formulation of Ryll-Nardzewski's Theorem tells us that a theory is $\aleph_0$-categorical if and only if its countable models have oligomorphic automorphism groups, i.e. for each n there are only finitely many orbits on n-tuples. Isomorphism of such structures is thus concretely classifiable, since the first-order theory of the structure will completely determine it up to isomorphism, and this theory may be coded as a real. An alternative type of symmetry we could consider is that of n-homogeneity (in the model-theoretic sense), as opposed to transitivity. Let us note that structures with strong homogeneity will be easy to classify, though, since their isomorphism class will be determined by a countable set of reals.

Bibliography

[1] S. Adams and A.S. Kechris, *Linear algebraic groups and countable Borel equivalence relations*, **Journal of the American Mathematical Society**, vol. 13 (2000), no. 4, pp. 909-943.

[2] H. Becker and A.S. Kechris, **The descriptive set theory of Polish group actions**, London Mathematical Society Lecture Notes Series, Cambridge, 1996.

[3] J.D. Clemens, S. Gao, and A.S. Kechris, *Polish metric spaces: Their classification and isometry groups*, to appear in **Bulletin of Symbolic Logic**.

[4] N.G. De Bruijn, *On bases for the set of integers*, **Publicationes Mathematicae Debrecen** , vol. 1 (1950), pp. 232-242.

[5] R. Dougherty, S. Jackson, and A.S. Kechris, *The structure of hyperfinite Borel equivalence relations*, **Transactions of the American Mathematical Society**, vol. 341 (1994), no. 1, pp. 193-225.

[6] S. Eigen and A. Hajian, *A characterization of exhaustive weakly wandering sequences for nonsingular transformations*, **Commentarii Mathematici Universitatis Sancti Pauli**, vol. 36 (1987), no. 2, pp. 227-233.

[7] S. Eigen, A. Hajian, and S. Kakutani, *Complementing sets of integers – a result from ergodic theory*, **Japanese Journal of Mathematics**, vol. 18 (1992), no. 1, pp. 205-211.

[8] S. Eigen, A. Hajian, and B. Weiss, *Borel automorphisms with no finite invariant measure*, **Proceedings of the American Mathematical Society**, vol. 126 (1998), no. 12, pp. 3619-3623.

[9] H. Friedman and L. Stanley, *A Borel reducibility theory for classes of countable structures*, **Journal of Symbolic Logic,** vol. 54 (1989), pp. 894-914.

[10] S. Gao, *Some applications of the Adams-Kechris technique*, to appear in **Proceedings of the American Mathematical Society**.

[11] S. Gao and A.S. Kechris, *On the classification of Polish metric spaces up to isometry*, preprint, 2000.

[12] M. Gromov, **Metric structures for Riemannian and non-Riemannian spaces**, Progress in Mathematics, vol. 152, Birkhäuser, 1999.

[13] A. Hajian and S. Kakutani, *Weakly wandering sets and invariant measures*, **Transactions of the American Mathematical Society**, vol. 110 (1964), pp. 136-151.

[14] L.A. Harrington, A.S. Kechris, and A. Louveau, *A Glimm-Effros dichotomy for Borel equivalence relations*, **Journal of the American Mathematical Society**, vol. 4 (1990), pp. 903-927.

[15] G. Hjorth, *An absoluteness principle for Borel sets*, **Journal of Symbolic Logic**, vol. 63 (1998), no. 2, pp. 663-693.

[16] G. Hjorth, **Classification and orbit equivalence relations**, Mathematical Surveys and Monographs, vol. 75, American Mathematical Society, 2000.

[17] G. Hjorth, *On invariants for measure preserving transformations*, preprint, 1998.

[18] G. Hjorth, *Universal co-analytic sets*, **Proceedings of the American Mathematical Society**, vol. 124 (1996), no. 12, pp. 3867-3873.

[19] T. Kamae, *A characterization of weakly wandering sequences for non-singular transformations*, **Commentarii Mathematici Universitatis Sancti Pauli**, vol. 32 (1983), pp. 55-59.

[20] M. Katětov, *On universal metric spaces*, **General topology and its relations to modern analysis and algebra, VI (Prague, 1986)**, Research and Exposition in Mathematics, vol. 16, Heldermann, Berlin, 1988, pp. 323-330.

[21] A.S. Kechris, **Classical descriptive set theory**, Springer-Verlag Graduate Texts in Mathematics, Berlin, 1995.

[22] A.S. Kechris and A. Louveau, *The classification of hypersmooth Borel equivalence relations*, **Journal of the American Mathematical Society**, vol. 10 (1997), no.1, pp. 215-242.

[23] L.M. Kelley, *Distance sets*, **Canadian Journal of Mathematics**, vol. 3 (1951), pp. 187-194.

[24] L.M. Kelly and E.A. Nordhaus, *Distance sets in metric spaces*, **Transactions of the American Mathematical Society**, vol.71 (1951), pp. 440-456.

[25] A. Mekler, *Stability of nilpotent groups of class 2 and prime exponent*, **Journal of Symbolic Logic**, vol. 46 (1981), pp. 781-788.

[26] Y.N. Moschovakis, **Descriptive set theory**, North-Holland Press, Amsterdam, 1980.

[27] J. Schmerl, *What's the difference?*, **Annals of Pure and Applied Logic**, vol. 93 (1998), pp. 255-261.

[28] J. Schmerl, personal communication.

[29] W. Sierpiński, *Sur l'ensemble de distances entre les points d'un ensemble*, **Fundamenta Mathematicae**, vol. 7 (1925), pp. 144-148.

[30] S. Solecki, *Analytic ideals*, **Bulletin of Symbolic Logic**, vol. 2 (1996), no. 3, pp. 339-348.

[31] H. Steinhaus, *Sur les distances des points des ensembles de mesure positive*, **Fundamenta Mathematicae**, vol. 1 (1920), pp. 93-104.

[32] B. Weiss, *The isomorphism problem in ergodic theory*, **Bulletin of the American Mathematical Society**, vol. 78 (1972), pp. 668-84.