Notre Dame Journal of Formal Logic
Volume 50, Number 1, 2009

Isomorphism of Homogeneous Structures

John D. Clemens

Abstract We consider the complexity of the isomorphism relation on countable first-order structures with transitive automorphism groups. We use the theory of Borel reducibility of equivalence relations to show that the isomorphism problem for vertex-transitive graphs is as complicated as the isomorphism problem for arbitrary graphs and determine for which first-order languages the isomorphism problem for transitive countable structures is as complicated as it is for arbitrary countable structures. We then use these results to characterize the complexity of the isometry relation for certain classes of homogeneous and ultrahomogeneous metric spaces.

1 Introduction

In their article [4], Friedman and Stanley considered the question of how difficult it is to classify a collection of countable first-order structures up to isomorphism. To make this precise, they define the space of countable models of a given first-order theory and consider the isomorphism relation as an equivalence relation on this space. They then use the relation of Borel reducibility of equivalence relations to compare such isomorphism relations, thus characterizing the difficulty of the corresponding isomorphism problem.

Certain first-order languages and theories have an isomorphism problem of maximal complexity in the sense that any other such isomorphism relation can be reduced to them. Such theories are called Borel-complete. Many of the techniques for showing that a given theory is Borel-complete involve coding other structures into models of the given theory, and this generally involves the use of distinguished points or definable subsets in the models produced. The aim of this article is to consider the extent to which distinguished points can be eliminated, that is, to consider the complexity of the isomorphism problem for structures with no distinguished points.

Received January 27, 2008; accepted November 13, 2008; printed December 18, 2008
2000 Mathematics Subject Classification: Primary, 03E15; Secondary, 03C15, 03C50
Keywords: countable structures, homogeneous, isomorphism
 10.1215/00294527-2008-024

To that end, we consider structures whose automorphism group acts transitively, so that there are no nontrivial definable subsets. We first show that the isomorphism problem for countable vertex-transitive graphs is Borel-complete; that is, it is as complicated as the isomorphism problem for arbitrary countable graphs. We then use this result to show that the collection of countable $\mathcal{L}$-structures with transitive automorphism groups for a given first-order language $\mathcal{L}$ is Borel-complete precisely when $\mathcal{L}$ contains a relation or function symbol of arity at least 2, or contains at least two unary function symbols. We then use the result about vertex-transitive graphs in order to determine the complexity of the isometry relation on certain classes of homogeneous and ultrahomogeneous metric spaces.

In Section 2 we review the coding of countable models and the notion of Borel-completeness. Section 3 presents the proof that the collection of vertex-transitive graphs is Borel-complete, as well as several variants. In Section 4 we characterize the languages whose isomorphism problem for transitive structures is Borel-complete, and we discuss some related results and questions in Section 5. We then use these results to classify the complexity of the isometry relation for homogeneous discrete and locally compact metric spaces in Section 6, and we consider ultrahomogeneous discrete and locally compact metric spaces in Section 7.

2 The Space of Countable Models

We begin by defining the space of countable models for a given first-order language $\mathcal{L}$. Definitions of any undefined model-theoretic terms may be found, for example, in Hodges [7]. Results about spaces of countable structures may be found in [4] and Hjorth [6].

Definition 2.1 Let $\mathcal{L} = \{R_i : i \leq N\}$ be a finite relational language, where R_i has arity n_i. The *space of countable models of* $\mathcal{L}$, $\mathrm{Mod}(\mathcal{L})$, is the set

$$\prod_{i \leq N} \mathcal{P}(\mathbb{N}^{n_i}),$$

where $\mathcal{P}(\mathbb{N}^{n_i})$ is the set of all subsets of $\mathbb{N}^{n_i}$. This space is equipped with the product topology obtained by identifying $\mathcal{P}(\mathbb{N}^{n_i})$ with $2^{\mathbb{N}^{n_i}}$.

Thus, a point in the space codes a countable structure $\mathcal{M}$ whose underlying set is $\mathbb{N}$, and where the interpretation of R_i in $\mathcal{M}$ is given by the corresponding subset of $\mathbb{N}^{n_i}$. We can extend this coding to handle countably infinite languages and languages with constant or function symbols in a straightforward manner.

Definition 2.2 The *isomorphism relation on* $\mathrm{Mod}(\mathcal{L})$, $\cong_{\mathcal{L}}$, is defined by setting two points equivalent if they code isomorphic $\mathcal{L}$-structures.

We can also consider the collection of models of some first-order theory T (or $\mathcal{L}_{\omega_1,\omega}$-sentence) in the language $\mathcal{L}$, denoted $\mathrm{Mod}(T)$; this will be a Borel subset of $\mathrm{Mod}(\mathcal{L})$. We then identify the isomorphism problem for models of T with the isomorphism relation $\cong_{\mathcal{L}}$ restricted to $\mathrm{Mod}(T)$, and we use the same terminology for other collections of $\mathcal{L}$-structures. In order to compare the complexity of two isomorphism problems, we use the notion of Borel reducibility of equivalence relations.

Definition 2.3 Let E and F be equivalence relations on the standard Borel spaces X and Y. We say that E is *Borel reducible* to F, $E \leq_B F$, if there is a Borel function $f : X \to Y$ such that for all $x_1, x_2 \in X$ we have $x_1 \, E \, x_2$ if and only if $f(x_1) \, F \, f(x_2)$.

Definition 2.4 Let T be a first-order theory. We say that T is *Borel-complete* if any isomorphism relation $\cong_{\mathcal{L}'}$ is Borel reducible to $\cong_{\mathcal{L}} \restriction \mathrm{Mod}(T)$.

This is equivalent to saying that any orbit equivalence relation induced by an action of the infinite symmetric group S_∞ is Borel reducible to the isomorphism relation for T (see Theorem 2.7.3 of [1]). We similarly say that a language is Borel-complete when the empty theory in that language is Borel-complete, and we say a given class of $\mathcal{L}$-structures is Borel-complete when any other isomorphism relation is reducible to the isomorphism relation on that class of structures.

To show that a theory is Borel-complete, it suffices to show that some other Borel-complete theory is Borel reducible to it. Friedman and Stanley show, for instance, that the theory of graphs is Borel-complete, so we can show that a theory is Borel-complete by reducing to it the relation of graph isomorphism.

3 Isomorphism of Symmetric Graphs

We begin by considering isomorphism of transitive graphs. In the theory of graphs there are two common notions of transitivity: vertex-transitivity and edge-transitivity.

Definition 3.1 A graph G is *vertex-transitive* if the automorphism group of G acts transitively on the set of vertices. A graph is *edge-transitive* if the automorphism group acts transitively on the set of edges.

It is more usual in model theory to axiomatize graphs so that the underlying set of the structure is the set of vertices of the graph and a symmetric binary relation determines which vertices are connected by an edge. Having a transitive automorphism group in this setting then corresponds to vertex-transitivity. Alternately, if we let the underlying set of the structure correspond to 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 countable connected vertex-transitive graphs and show that their isomorphism problem is Borel-complete. 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 terminology. We will use the fact that the empty theory in the language whose signature consists of a single binary relation is Borel complete (see [4]).

Theorem 3.2 *Isomorphism of countable connected graphs having 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 countable connected vertex-transitive graphs. The main idea of the proof will be that Cayley graphs for countable groups provide canonical vertex-transitive graphs. In fact, any vertex-transitive graph is close to being the Cayley graph of some group; see Sabidussi [10]. Our construction is based closely on Mekler's proof that the theory of nilpotent class 2 groups of prime exponent is Borel-complete (see Mekler [9]). Mekler begins by constructing a Borel map which assigns to each $\mathcal{L}_0$-structure

$\mathcal{A}$ a graph $G(\mathcal{A})$ in an isomorphism-preserving way so that $\mathcal{A}_1 \cong \mathcal{A}_2$ if and only if $G(\mathcal{A}_1) \cong G(\mathcal{A}_2)$. The graph $G(\mathcal{A})$ 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 adjacent 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 $\langle v_i \rangle_{i\in\mathbb{N}}$ 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 $\langle g_i \rangle_{i\in\mathbb{N}}$ are generators of the free group on countably many generators, $\mathbb{F}_\omega$, 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 $\langle g_i \rangle_{i\in\mathbb{N}}$. Specifically, let N be the normal subgroup of $\mathbb{F}_\omega$ generated by

$$\{g_i g_j g_i^{-1} g_j^{-1} : v_i \text{ is adjacent to } v_j \text{ in } G(\mathcal{A})\}.$$

Vertices of $\mathcal{G}$ are left cosets of N in $\mathbb{F}_\omega$, and two vertices $w_1 N$ and $w_2 N$ are adjacent in $\mathcal{G}$ 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 $\mathbb{N}$) in the following manner. First, fix an enumeration $\langle w_i \rangle_{i\in\mathbb{N}}$ of the words in $\mathbb{F}_\omega$ with the generators $\langle g_i \rangle_{i\in\mathbb{N}}$. For each coset of N, we can then pick the least i such that w_i is in the given coset and take this element w_i as a representative of the coset. Note that it may be undecidable to determine whether two integers index words in the same coset, but that will be irrelevant here. We then can enumerate these representatives and define the binary relation on $\mathbb{N}$ which encodes this graph according to whether the corresponding cosets are adjacent in $\mathcal{G}$. Call the code for this graph $\mathcal{G}(\mathcal{A})$.

Observe that the generators $\langle g_i \rangle_{i\in\mathbb{N}}$ 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 vertex $w_2 N$ if there is a generator g_i 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 of $\cong_{\mathcal{L}_0}$ to the isomorphism relation on vertex-transitive graphs. First, it is easy to check that each graph $\mathcal{G}(\mathcal{A})$ is vertex-transitive, for if we have two vertices $w_1 N$ and $w_2 N$ in $\mathcal{G}(\mathcal{A})$ then the map 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 $\mathcal{L}_0$-structures $\mathcal{A}_1$ and $\mathcal{A}_2$ with $\mathcal{A}_1 \cong \mathcal{A}_2$. 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, φ acts on the cosets of these elements). We want to extend this map to an isomorphism of the whole 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 $\widetilde{w} = g_{f(i_n)}^{\sigma_n} \cdots g_{f(i_0)}^{\sigma_0}$ for $w = g_{i_n}^{\sigma_n} \cdots g_{i_0}^{\sigma_0}$.

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,

$$w_1 w_2^{-1} \in N_1 \iff \widetilde{w_1}\widetilde{w_2}^{-1} \in N_2.$$

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 vertex-transitivity of $\mathcal{G}(\mathcal{A})$ it does not matter which vertex we use, so we may assume that it is the vertex corresponding to N. We can then identify the vertices adjacent to this fixed vertex, which will be the vertices $g_k^{\pm 1} N$. These vertices are all distinct, although we will not be able to identify which is which. Let these vertices be enumerated as $\langle u_i \rangle_{i \in \mathbb{N}}$. Consider the binary relation R on this set, where two vertices are R-related if they are at opposite corners of a square (i.e., a cycle of length 4) in $\mathcal{G}(\mathcal{A})$. That is,

$$u_i \; R \; u_j \iff u_i \neq u_j \; \wedge \; \exists a \exists b \, [a \neq b \; \wedge \\ (u_i \text{ and } u_j \text{ are each adjacent to both } a \text{ and } b)].$$

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 k_1 and k_2 in $\mathbb{N}$ and σ_1 and σ_2 in $\{1, -1\}$ 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 $u_j = g_{k_2}^{\sigma_2} N$. Let a and b be the other two vertices of the square. There are thus generators g_{n_1}, g_{n_2}, g_{m_1}, and g_{m_2} and $\tau_1, \tau_2, \rho_1, \rho_2 \in \{1, -1\}$ witnessing this; that is,

$$a = g_{n_1}^{\tau_1} g_{k_1}^{\sigma_1} N = g_{n_2}^{\tau_2} g_{k_2}^{\sigma_2} N$$
$$b = g_{m_1}^{\rho_1} g_{k_1}^{\sigma_1} N = g_{m_2}^{\rho_2} g_{k_2}^{\sigma_2} N.$$

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 $k_1 = k_2$, then 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 that

$$g_{k_1}^{-\sigma_1} g_{k_2}^{-\sigma_2} g_{k_1}^{\sigma_1} g_{k_2}^{\sigma_2} \in N.$$

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 applied to b shows that either $b = N$ or v_{k_1} is adjacent to v_{k_2} in $G(\mathcal{A})$. Since we know that $a \neq b$, they cannot both be equal to N so that v_{k_1} and v_{k_2} must be adjacent as we wished to show.

We can 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 $G(\mathcal{A})$. □

We should note that the groups whose Cayley graphs are constructed here are different than the groups used in Mekler's result (since the groups here are not nilpotent); it is unclear whether Mekler's groups can be used directly. The above proof also 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 above proof. We thus get the following theorem.

Theorem 3.3 *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 3.4 We say that a graph G is *symmetric* if for any two edges (u_1, u_2) and (v_1, v_2) in G 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 3.5 *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 3.2 to isomorphism of symmetric graphs. We will in fact reuse 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 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

$$G_1 \cong G_2 \iff \mathcal{G}_1 \cong \mathcal{G}_2.$$

It thus suffices to show that whenever G is one of our earlier 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 $\widetilde{w} = g_{i_n}^{-\sigma_n} \cdots g_{i_0}^{-\sigma_0}$ for $w = g_{i_n}^{\sigma_n} \cdots g_{i_0}^{\sigma_0}$. 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 $g_i g_j g_i^{-1} g_j^{-1}$. 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 is 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_2 by letting

$$\pi_2(wN) = \widetilde{w}N,$$

where $\widetilde{w} = g_{\varphi(i_n)}^{\sigma_n} \cdots g_{\varphi(i_0)}^{\sigma_0}$ for $w = g_{i_n}^{\sigma_n} \cdots g_{i_0}^{\sigma_0}$. 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.

Finally, 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 instead 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 $g_i N$ to $g_j N$, and we do not need to interchange $g_k N$ and $g_k^{-1} N$. We thus have the following.

Theorem 3.6 *Isomorphism of symmetric weakly-connected countable directed graphs is Borel-complete.*

Let us also note that if we continue to iterate this embedding then we can get Borel-completeness for classes of graphs with even greater symmetry, for instance, graphs in which every square (4-cycle) can be mapped to any other square by an automorphism. As we will discuss below, there is an upper limit to the amount of symmetry we can demand while still having a complicated isomorphism problem.

To emphasize the complexity retained by transitive graphs, we can restate our main result as follows.

Corollary 3.7 *Classifying countable connected symmetric graphs up to isomorphism is as complicated as classifying arbitrary countable graphs.*

4 Other Transitive Countable Structures

Besides the theory of graphs, 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 they 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 3.2. We thus get the following corollary.

Corollary 4.1 *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 of a structure, 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.

Now, notice that if we add relation or function symbols to a language whose collection of transitive models is Borel-complete then we will still have a Borel-complete isomorphism problem because we can restrict our attention to those structures where the new symbols have trivial interpretations (for instance, nothing is related under new relation symbols, and new 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 into a binary function so as to preserve automorphisms. To do this, interpret f so that $f(x, x) = x$ and so that for $x \neq y$ we have $f(x, y) = x$ if $x \ R \ y$ and $f(x, y) = y$ 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 language with 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 the following.

Corollary 4.2 *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 for a language with only unary relations is an element of $2^{\mathbb{N}}$, since each relation must either be satisfied by everything or by nothing (we are assuming a countably infinite language; in general, we can encode an element of $2^{|\mathcal{L}|}$). Similarly, if the language contains only a single unary function symbol then there are only countably many isomorphism types for transitive structures, with the isomorphism type only depending on whether the function splits into some number of finite cycles (and the corresponding cycle size) or whether it splits into some number of uniformly branching bi-infinite trees (and the branching number, that is, the size of the preimage of a point). The isomorphism relation for transitive models of such languages will then be simple in the following sense.

Definition 4.3 An equivalence relation E on X is *concretely classifiable* if it is Borel reducible to the identity relation on some Polish space; that is, there is a Borel function $f : X \to Y$ for some Polish space Y such that for all $x_1, x_2 \in X$ we have $x_1 \; E \; x_2$ if and only if $f(x_1) = f(x_2)$.

We then get the following proposition.

Proposition 4.4 *If $\mathcal{L}$ is a language whose signature contains only unary relation symbols and a single unary function symbol then the isomorphism problem for the transitive countable models of $\mathcal{L}$ is concretely classifiable.*

Proof Let $\mathcal{L}$ have the unary function symbol f and the unary relation symbols R_i for $i \in \mathbb{N}$. For a transitive $\mathcal{L}$-structure $\mathcal{M}$, define $\varphi(\mathcal{M}) \in \mathbb{N}^{\mathbb{N}}$ by

$$\varphi(\mathcal{M})(0) = \begin{cases} 0 & \text{if the } f\text{-orbits in } \mathcal{M} \text{ are bi-infinite trees} \\ & \quad \text{with infinite branching} \\ 2n & \text{if the } f\text{-orbits in } \mathcal{M} \text{ are bi-infinite trees} \\ & \quad \text{with branching number } n \\ 2n-1 & \text{if the } f\text{-orbits in } \mathcal{M} \text{ are cycles of size } n \end{cases}$$

$$\varphi(\mathcal{M})(1) = \begin{cases} 0 & \text{if there are infinitely many } f\text{-orbits in } \mathcal{M} \\ m & \text{if there are } m\text{-many } f\text{-orbits in } \mathcal{M} \end{cases}$$

$$\varphi(\mathcal{M})(2+i) = \begin{cases} 0 & \text{if } R_i \text{ is satisfied by no element of } \mathcal{M} \\ 1 & \text{if } R_i \text{ is satisfied by every element of } \mathcal{M}. \end{cases}$$

Then the above discussion shows that two transitive $\mathcal{L}$-structures $\mathcal{M}_1$ and $\mathcal{M}_2$ will be isomorphic if and only if $\varphi(\mathcal{M}_1) = \varphi(\mathcal{M}_2)$, so φ witnesses that the isomorphism relation is concretely classifiable. □

This leaves only the case where we have a language with 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 the following.

Proposition 4.5 *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 3.5 to isomorphism of transitive $\mathcal{L}_{u2}$-structures. By the result of Theorem 3.5 this will be sufficient. Given a symmetric graph G we will produce an $\mathcal{L}_{u2}$-structure $\mathcal{A} = \mathcal{A}(G)$, where f_0 and f_1 will 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 f_0 will be defined so that each point has countably many preimages and there are countably many connected components in the graph it induces, 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 $n, m \in \mathbb{N}$ 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 preimages, and so forth, we have essentially a copy of the Baire space $\mathbb{N}^{\mathbb{N}}$ below this distinguished node. Relative to a_0, we can then label points in the component of a_0 by pairs $(n, s) \in \mathbb{N} \times \mathbb{N}^{<\mathbb{N}}$, where $\mathbb{N}^{<\mathbb{N}}$ is the set of finite sequences from $\mathbb{N}$. Here n indicates how far "up" we start from a_0 (i.e., we start from $f^n(a_0)$), and s determines a point in the copy of Baire space below this point $f^n(a_0)$, with the understanding that a_0 is along the leftmost branch (the branch with all coordinates 0). This gives some points multiple labels; we identify $(n+1, 0 \frown s)$ with (n, s) (where $s \frown t$ is the concatenation of two sequences s and t). The node a_0 is then indexed by $(0, \langle\rangle)$ (as well as by other labels). The function f_0 is then 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 $n \in \mathbb{N}$, 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 $a_1 = (0, \langle\rangle)$, we associate countably many new components and so forth. All of the components are distinct and each is a connected component of f_0 with f_0 behaving as in the initial component.

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_is label the same point. Two nodes are thus in the same component if their sequences agree up to n_{l-1} (modulo this identification). The function f_0 acts on the final pair a_l of a sequence, as indicated above.

To define the function f_1, we first 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 w of the form $\langle a_0, n_0, \ldots, a_{l-1}, n_{l-1}, a_l\rangle$ with $l \geq 1$ and $a_l = (0, \langle\rangle)$ we let the index of w 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 nonzero index. This will not affect the transitivity of the structure, though, because we will be unable to determine these indices within the structure.

To each noninitial node $\langle a_0, n_0, \dots, a_{l-1}, n_{l-1}, a_l\rangle$ we associate the initial node $\langle a_0, n_0 \dots, a_{l-1}\rangle$ 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 $\langle v_i\rangle_{i\in\mathbb{N}}$ enumerate the vertices in the given symmetric graph G (according to its coding). For each $i \in \mathbb{N}$, let $\langle k^i_n\rangle_{n\in\mathbb{N}}$ enumerate in increasing order the indices of the vertices adjacent to v_i in G and let $\langle m^i_n\rangle_{n\in\mathbb{N}}$ indicate where v_i occurs in $v_{k^i_n}$'s enumeration; that is, the m^i_ns are such that

$$k^{k^i_n}_{m^i_n} = i \text{ for each } i \text{ and } n.$$

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 $w = \langle a_0, n_0, \dots, n_{l-1}, a_l\rangle$ in $\mathcal{A}$ with $a_l = (n, s)$ we write $w \frown j$ to denote the node $\langle a_0, n_0, \dots, n_{l-1}, a'_l\rangle$ where $a'_l = (n, s \frown j)$, so that $w \frown j$ is the jth node in the block below w. We now define f_1:

$$f_1(w \frown j) = \begin{cases} \left\langle I(w), k^{\mathrm{ind}(w)}_j, \left(0, \left\langle m^{\mathrm{ind}(w)}_j\right\rangle\right)\right\rangle & \text{if } k^{\mathrm{ind}(w)}_j \neq 0 \\ I(w) \frown m^{\mathrm{ind}(w)}_j & \text{if } k^{\mathrm{ind}(w)}_j = 0. \end{cases}$$

This serves to define f_1 everywhere, since each node is in the block below some unique node w. For simplicity, we shall write

$$f_1(w \frown j) = \left\langle I(w), k^{\mathrm{ind}(w)}_j, (0, \langle\rangle) \frown m^{\mathrm{ind}(w)}_j\right\rangle,$$

with the understanding that this collapses to $I(w) \frown m^{\mathrm{ind}(w)}_j$ when $k^{\mathrm{ind}(w)}_j = 0$. Note that f_1 is an involution:

$$\begin{aligned} f_1(f_1(w \frown j)) &= f_1\left(\left\langle I(w), k^{\mathrm{ind}(w)}_j, (0, \langle\rangle) \frown m^{\mathrm{ind}(w)}_j\right\rangle\right) \\ &= f_1\left(\left\langle I(w), k^{\mathrm{ind}(w)}_j, (0, \langle\rangle)\right\rangle \frown m^{\mathrm{ind}(w)}_j\right) \\ &= \left\langle I\left(\left\langle I(w), k^{\mathrm{ind}(w)}_j, (0, \langle\rangle)\right\rangle\right), k^{\mathrm{ind}(\langle I(w), k^{\mathrm{ind}(w)}_j, (0,\langle\rangle)\rangle)}_{m^{\mathrm{ind}(w)}_j}, \right. \\ &\qquad \left. (0, \langle\rangle) \frown m^{\mathrm{ind}(\langle I(w), k^{\mathrm{ind}(w)}_j, (0,\langle\rangle)\rangle)}_{m^{\mathrm{ind}(w)}_j}\right\rangle \\ &= \left\langle I(w), k^{k^{\mathrm{ind}(w)}_j}_{m^{\mathrm{ind}(w)}_j}, (0, \langle\rangle) \frown m^{k^{\mathrm{ind}(w)}_j}_{m^{\mathrm{ind}(w)}_j}\right\rangle \\ &= \langle I(w), \mathrm{ind}(w), (0, \langle\rangle) \frown j\rangle \\ &= w \frown j. \end{aligned}$$

Let us clarify how f_1 behaves. In each group as defined above we have nodes with indices in $\mathbb{N}$; let the given group have nodes $\langle w_i\rangle_{i\in\mathbb{N}}$ with $\mathrm{ind}(w_i) = i$. 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^i_ns and m^i_ns determine which elements in each block are connected (the nth element in the ith block is connected to the (m^i_n)th element of the (k^i_n)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, that is, that $\mathcal{A}(G)$ has a transitive automorphism group and that $\mathcal{A}(G_1) \cong \mathcal{A}(G_2)$ if and only if $G_1 \cong G_2$.

First, suppose that we have two graphs G_1 and G_2 with $G_1 \cong G_2$. 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 from $\mathcal{A}(G_1)$ to $\mathcal{A}(G_2)$ which is an isomorphism between groups we will have no problems in extending it progressively to define an isomorphism π from $\mathcal{A}(G_1)$ to $\mathcal{A}(G_2)$ in the same manner.

We start by setting $\pi(0, \langle\rangle) = (0, \langle\rangle)$, thus 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 w 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 w 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 ensure 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 $f_0(w_1)$, that is, $w_1 = f_0(w_1) \frown n_1$, 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 node's 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 $f_1(i,n) = (k^i_n, m^i_n)$ 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 $\varphi(k^{i_1}_{n_1}) = \widetilde{k}^{i_2}_{n_2}$, where 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 $\widetilde{k}^{\varphi(i)}_j = \varphi(k^i_n)$ (such a j exists since $\widetilde{v}_{\varphi(i)}$ is adjacent to $\widetilde{v}_{\varphi(k^i_n)}$ in G_2, as v_i is adjacent to $v_{k^i_n}$ in G_1). In particular, $\rho(i_1, n_1) = n_2$ since $\widetilde{k}^{\varphi(i_1)}_{n_2} = \widetilde{k}^{i_2}_{n_2} = \varphi(k^{i_1}_{n_1})$ by our choice of φ, so that $\pi(i_1, n_1) = (i_2, n_2)$ as required. We also have

$$\pi(f_1(i,n)) = \left(\varphi\left(k^i_n\right), \rho\left(k^i_n, m^i_n\right)\right)$$

and

$$f_1(\pi(i,n)) = \left(\widetilde{k}^{\varphi(i)}_{\rho(i,n)}, \widetilde{m}^{\varphi(i)}_{\rho(i,n)}\right).$$

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 the k^i_ns and m^i_ns. But our definition of ρ implies that the left-hand side is equal to $\widetilde{k}^{\widetilde{k}^{\varphi(i)}_{\rho(i,n)}}_{\widetilde{m}^{\varphi(i)}_{\rho(i,n)}} = \varphi(i)$ as well. Thus our extension of π respects f_1.

For extensions of type (2) we proceed in a similar manner, but we have 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 the definition of π at our initial point, $\pi((0, \langle\rangle)) = (0, \langle\rangle)$, we successively extend π to all blocks and corresponding groups in the initial component of $\mathcal{A}(G_1)$. If we then take 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 (since 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 $\mathcal{A}(G)$ has a transitive automorphism group (which we will show below), 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 the 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; note that this will not require the above result that the map $G \mapsto \mathcal{A}(G)$ is a reduction (and hence introduces no circularity). This is similar to the verification that we have

$\mathcal{A}(G_1) \cong \mathcal{A}(G_2)$ when $G_1 \cong G_2$. 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 this node and to the group of this block (as well as to 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 on this fact; we could have initialized π by mapping any node of $\mathcal{A}(G_1)$ to any node of $\mathcal{A}(G_2)$. 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 4.6 *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.*

5 Additional Comments on Transitive Structures

We 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 and 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 $\langle R_i \rangle_{i \in \mathbb{N}}$ we can code a sequence $x \in 2^{\mathbb{N}}$ into an element a of the structure by setting

$$R_i(a) \Longleftrightarrow x(i) = 1.$$

Our structure can thus code a countable set of reals, one for each element in the structure. The isomorphism problem then turns out to be bireducible with the equivalence relation F_2 of equality of countable sets of reals (which we will define in Section 7 below).

The most striking difference is in the case of a single unary function symbol. Friedman and Stanley show (in [4]) that the isomorphism problem for 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 collection of transitive structures for a language with a single unary function symbol, though, we saw that the isomorphism problem is concretely classifiable. This allows us to draw the following conclusion: The theory of graphs cannot be axiomatized in a language with only one unary function symbol in a way that preserves automorphism groups (in the sense that the automorphism group when considered as an $\mathcal{L}$-structure is the same as for the original graph).

Another observation we should make is that it is necessary to produce graphs with infinite degree for each vertex in the proof of Theorem 3.2. 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 Gao and Kechris [5], which says that isometry of homogeneous pseudoconnected locally compact Polish metric spaces is concretely-classifiable. A locally-finite graph when given the graph metric becomes a pseudoconnected 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 them to the collection of transitive models. We can ask this question.

Question 5.1 *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?*

Many natural theories are immediately ruled out because their structures have definable sets or elements. As noted earlier, having any nontrivial definable sets prevents a structure from having a transitive automorphism group. Thus structures such as trees, groups, and most algebraic structures with complicated isomorphism problems are eliminated. The theory of linear orders, on the other hand, avoids this problem and seems a natural candidate for this question.

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 4.6 for n-transitive structures.

Question 5.2 *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 $n \in \mathbb{N}$. Here, though, we note that a structure having this property has an $\aleph_0$-categorical theory, since Ryll-Nardzewski's Theorem tells us that a theory is $\aleph_0$-categorical if and only if its countable models have oligomorphic automorphism groups; that is, 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 easier to classify, though, since their isomorphism class will be determined by a countable set of reals.

6 Homogeneous Locally Compact Spaces

We now use the above results about isomorphism of vertex-transitive graphs to derive some corollaries concerning the complexity of the isometry relation on certain classes of metric spaces. The relevant definitions may be found in Clemens [2], Gao and Kechris [5], or Clemens, Gao, and Kechris [3] (where several of the following results were announced).

Recall that a metric space is said to be *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 continue to 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 vertex-transitive automorphism groups.

Theorem 6.1 *The isomorphism relation on countable vertex-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 equip it with the graph metric, where 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. □

We showed in Section 3 that isomorphism of countable vertex-transitive graphs is bireducible with graph isomorphism (Theorem 3.2). Since isometry of general discrete metric spaces is Borel reducible to graph isomorphism, we thus have the following corollary.

Corollary 6.2 *Isometry of homogeneous discrete metric spaces is Borel 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 as with the case of general locally compact spaces we do not have an exact upper bound.

7 Ultrahomogeneous Locally Compact Spaces

We end by considering discrete and locally compact metric spaces with even richer isometry groups. The techniques in this section will not involve countable structures but will rely directly on metric space techniques. Recall that a metric space is *ultrahomogeneous* if any partial isometry between finite subsets of the space can be extended to an isometry of the whole space. We will use the following alternate characterization.

Definition 7.1 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 that $\varphi(x_i) = y_i$ 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 $\varphi(x_{n+1}) = y_{n+1}$.

Ultrahomogeneity 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 ultrahomogeneous.

We begin with the collection of discrete spaces. We first recall the equivalence relation F_2 of equality of countable sets of reals, which is defined on the space $\mathbb{R}^{\mathbb{N}}$ by setting

$$\langle x_n \rangle_{n\in\mathbb{N}} \; F_2 \; \langle y_n \rangle_{n\in\mathbb{N}} \iff \{x_n : n \in \mathbb{N}\} = \{y_n : n \in \mathbb{N}\}.$$

This equivalence relation is strictly simpler than graph isomorphism in terms of Borel reducibility. The complexity characterization here is then as follows.

Theorem 7.2 *Isometry of ultrahomogeneous discrete metric spaces is bireducible with F_2.*

Proof The reduction of isometry of ultrahomogeneous discrete spaces to F_2 is simple. Observe that two ultrahomogeneous Polish metric spaces are isometric precisely when they have the same sets of n-point distance configurations for all $n \geq 2$. A discrete metric space is countable, so it contains only countably many n-point distance 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 ultrahomogeneous discrete spaces we modify Katětov's construction of the Urysohn space in Katětov [8]. The Urysohn space is an ultrahomogeneous Polish metric 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. Given (X_n, d_n) for some $n \in \mathbb{N}$, we define (X_{n+1}, d_{n+1}) as follows. First, we set

$$X_{n+1} = X_n \sqcup E_A(X_n),$$

where

$$E_A(X) = \{f : X \to A' \text{ such that for all but finitely many } x \in X \text{ we have } f(x) = 1\},$$

where we can omit the usual condition that f satisfy a triangle inequality since the range of f is contained in $[1, 2]$. We then define d_{n+1} by setting

$$\begin{aligned} d_{n+1}(x_1, x_2) &= d_n(x_1, x_2) && \text{for } x_1, x_2 \in X_n \\ d_{n+1}(f, x) &= f(x) && \text{for } f \in E_A(X_n) \text{ and } x \in X_n \\ d_{n+1}(f_1, f_2) &= 1 && \text{for } f_1, f_2 \in E_A(X_n) \text{ with } f_1 \neq f_2. \end{aligned}$$

As is the construction of the Urysohn space, this defines a metric space; 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 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 (X_n, d_n) is a subspace of (X_{n+1}, d_{n+1}) for each n. We conclude by setting

$$(X_A, d_A) = \bigcup_{n \in \mathbb{N}} (X_n, d_n).$$

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 $(X_{A_1}, d_{A_1}) \cong_i (X_{A_2}, d_{A_2})$. For the converse, note that the set of distances in (X_A, d_A) is 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 ultrahomogeneous. For this, we will show that the spaces have the one-point extension property. The construction of (X_A, d_A) makes this property easy to verify. Given points $x_1, \ldots, x_n, x_{n+1}$ 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_ns as x_{n+1} does to the x_ns, and we can take y_{n+1} to be such an f. □

Once again, we have that F_2 is a lower bound for the isometry relation on locally compact ultrahomogeneous 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.

We recall from [5] the definition of a *pseudocomponent* 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; that is,

$$\rho(x) = \sup\{r : B_r^{\mathrm{cl}}(x) \text{ is compact}\},$$

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 ultrahomogeneous 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

$$x \; E \; y \iff x = y \; \vee \; (x \; R^* \; y \; \wedge \; y \; R^* \; x).$$

The pseudocomponents of X are then the equivalence classes of E. As shown in [5], the map $x \mapsto \rho(x)$ is Lipschitz and each pseudocomponent is clopen, so there are at most countably many pseudocomponents. A space with only one pseudocomponent is said to be *pseudoconnected*. 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 $\langle d_{i,j}\rangle_{i,j\in\mathbb{N}}$ code the Polish metric space $\overline{\{x_i : i \in \mathbb{N}\}}$, where $\{x_i : i \in \mathbb{N}\}$ is a countable dense subset and $d(x_i, x_j) = d_{i,j}$ for $i, j \in \mathbb{N}$. We assume that this space is locally compact.

Lemma 7.3 *For $\delta > 0$ and $i \in \mathbb{N}$, the set $\overline{B_\delta(x_i)}$ is compact if and only if the following holds:*

$$(\forall q \in \mathbb{Q}^+)(\exists s \in [\mathbb{N}]^{<\mathbb{N}})\,[(\forall k < |s|)\,[d_{i,s(k)} < \delta] \wedge \\ \forall j\,[d_{i,j} < \delta \implies (\exists k < |s|)\,[d_{j,s(k)} < q]\,]\,],$$

where $\mathbb{Q}^+$ is the set of positive rationals and $[\mathbb{N}]^{<\mathbb{N}}$ is the set of increasing finite sequences from $\mathbb{N}$.

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

$$(\forall y \in \overline{B_\delta(x_i)})(\exists k < n)\left[d(y, y_k) < \frac{q}{2}\right].$$

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 $s \in [\mathbb{N}]^{<\mathbb{N}}$ 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 7.4 *We have that x_i and x_j are in the same pseudocomponent if and only if the following holds:*

$$(\exists i_0, \ldots, i_n)\,[i_0 = i \;\wedge\; i_n = j \;\wedge\; (\forall k < n)\,[d_{i_k,i_{k+1}} < \rho(x_{i_k})]] \wedge \\ (\exists j_0, \ldots, j_m)\,[j_0 = j \;\wedge\; j_m = i \;\wedge\; (\forall k < m)\,[d_{j_k,j_{k+1}} < \rho(x_{j_k})]].$$

Proof If this condition holds then x_i and x_j are clearly in the same pseudocomponent. Suppose conversely that x_i and x_j are in the same pseudocomponent. 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_is. We can set $i_0 = i$ and $i_n = j$. Then let

$$\delta_0 = \rho(y_0) - d(y_0, y_1) > 0$$
$$\delta_1 = \rho(y_1) - d(y_1, y_2) > 0.$$

Choose $\epsilon < \min(\delta_0, \frac{\delta_1}{2})$ and choose 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 7.5 *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 $X = \overline{\{x_i : i \in \mathbb{N}\}}$ then $\langle d_{\langle n,i\rangle,\langle m,j\rangle}\rangle$ also codes this space, $X = \overline{\{x_{n,i} : n, i \in \mathbb{N}\}}$, and for each n we have that the space $X_n = \overline{\{x_{n,i} : i \in \mathbb{N}\}}$ is a pseudocomponent of X. In the case that X has infinitely many pseudocomponents, 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 pseudocomponent 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 pseudocomponents. This suffices for spaces with infinitely many pseudocomponents; otherwise, we enumerate one of them infinitely often. □

We are now ready to prove our characterization.

Theorem 7.6 *Isometry of ultrahomogeneous 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 [5]), isometry of locally compact Polish metric spaces with only finitely many pseudocomponents is essentially countable; that is, it is reducible to a countable Borel equivalence relation. Every countable Borel equivalence relation is reducible to F_2 by sending an element to its equivalence class, which is a countable set. We can thus fix a sequence of functions $\langle\rho_n\rangle_{n\in\mathbb{N}}$ such that ρ_n reduces isometry of spaces with n pseudocomponents to F_2 and, moreover, satisfies

$$X_1 \cong_i X_2 \iff \rho_n(X_1) = \rho_n(X_2)$$
$$\iff \rho_n(X_1) \cap \rho_n(X_2) \neq \varnothing$$

for X_1 and X_2 with n pseudocomponents (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 the interval $[n, n+1)$.

Now, given an ultrahomogeneous locally compact space coded by the array $\langle d_{i,j}\rangle_{i,j\in\mathbb{N}}$, let $\langle X_n\rangle_{n\in\mathbb{N}}$ be its pseudocomponents as enumerated by the function from Lemma 7.5. We will assume that X has infinitely many pseudocomponents; this can be determined in a Borel way and it is straightforward to handle spaces with only finitely many pseudocomponents. For $n \in \mathbb{N}$ let

$$\sigma_n(X) = \bigcup \left\{\rho_{n+1}(X_{i_0} \sqcup X_{i_1} \sqcup \dots \sqcup X_{i_n}) : i_0 < i_1 < \dots < i_n\right\},$$

where we are again identifying a countable sequence with the countable set it enumerates. By interweaving 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 $n+1$ pseudocomponents. We then define our reducing function f by setting

$$f(X) = \bigcup_{n\in\mathbb{N}} \sigma_n(X).$$

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 set of subspaces with finitely many pseudocomponents, 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_ns are disjoint, we have that $\sigma_n(X) = \sigma_n(Y)$ for each n. Thus,

$$\bigcup \left\{\rho_{n+1}(X_{i_0} \sqcup \cdots \sqcup X_{i_n}) : i_0 < \cdots < i_n\right\} = \bigcup \left\{\rho_{n+1}(Y_{i_0} \sqcup \cdots \sqcup Y_{i_n}) : i_0 < \cdots < i_n\right\}.$$

But recall that our functions ρ_n have the property that if

$$\rho_{n+1}(X_{i_0} \sqcup X_{i_1} \sqcup \cdots \sqcup X_{i_n}) \cap \rho_{n+1}(Y_{j_0} \sqcup Y_{j_1} \sqcup \cdots \sqcup Y_{j_n}) \neq \varnothing,$$

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

$$\rho_n(X_0 \sqcup \cdots \sqcup X_n) = \rho_n(Y_{i_0^n} \sqcup \cdots \sqcup Y_{i_n^n})$$
$$\rho_n(Y_0 \sqcup \cdots \sqcup Y_n) = \rho_n(X_{j_0^n} \sqcup \cdots \sqcup X_{j_n^n}).$$

Hence,

$$X_0 \sqcup \cdots \sqcup X_n \cong_i Y_{i_0^n} \sqcup \cdots \sqcup Y_{i_n^n}$$
$$Y_0 \sqcup \cdots \sqcup Y_n \cong_i X_{j_0^n} \sqcup \cdots \sqcup X_{j_n^n}.$$

Since each finite configuration of points in X (respectively, Y) will occur in some $X_0 \sqcup \cdots \sqcup X_n$ (respectively, $Y_0 \sqcup \cdots \sqcup Y_n$), we see that the same configuration occurs (up to isometry) in Y (respectively, X). Thus, X and Y have the same n-point distance configurations for each n, and since they are ultrahomogeneous this suffices to establish that they are isometric. □

References

[1] Becker, H., and A. S. Kechris, *The Descriptive Set Theory of Polish Group Actions*, vol. 232 of *London Mathematical Society Lecture Note Series*, Cambridge University Press, Cambridge, 1996. Zbl 0949.54052. MR 1425877. 3

[2] Clemens, J. D., "Isometry of Polish metric spaces," forthcoming in *Annals of Pure and Applied Logic*. 15

[3] Clemens, J. D., S. Gao, and A. S. Kechris, "Polish metric spaces: Their classification and isometry groups," *Bulletin of Symbolic Logic*, vol. 7 (2001), pp. 361–75. Zbl 0994.54037. MR 1860610. 16

[4] Friedman, H., and L. Stanley, "A Borel reducibility theory for classes of countable structures," *The Journal of Symbolic Logic*, vol. 54 (1989), pp. 894–914. Zbl 0692.03022. MR 1011177. 1, 2, 3, 14

[5] Gao, S., and A. S. Kechris, "On the classification of Polish metric spaces up to isometry," *Memoirs of the American Mathematical Society*, no. 766 (2003). Zbl 1012.54038. MR 1950332. 15, 16, 18, 20

[6] Hjorth, G., *Classification and Orbit Equivalence Relations*, vol. 75 of *Mathematical Surveys and Monographs*, American Mathematical Society, Providence, 2000. Zbl 0942.03056. MR 1725642. 2

[7] Hodges, W., *Model Theory*, vol. 42 of *Encyclopedia of Mathematics and its Applications*, Cambridge University Press, Cambridge, 1993. Zbl 0789.03031. MR 1221741. 2

[8] Katětov, M., "On universal metric spaces," pp. 323–30 in *General Topology and Its Relations to Modern Analysis and Algebra, VI (Prague, 1986)*, vol. 16 of *Research and Exposition in Mathematics*, Heldermann, Berlin, 1988. Zbl 0642.54021. MR 952617. 17

[9] Mekler, A. H., "Stability of nilpotent groups of class 2 and prime exponent," *The Journal of Symbolic Logic*, vol. 46 (1981), pp. 781–88. Zbl 0482.03014. MR 641491. 3

[10] Sabidussi, G., "Vertex-transitive graphs," *Monatshefte für Mathematik*, vol. 68 (1964), pp. 426–38. Zbl 0136.44608. MR 0175815. 3

Acknowledgments

I would like to thank the referee for suggesting several helpful clarifications and questions.

Department of Mathematics
The Pennsylvania State University
University Park PA 16802
USA
clemens@math.psu.edu
http://www.math.psu.edu/clemens