Contributions to Discrete Mathematics

Volume 12, Number 1, Pages 62–73
ISSN 1715-0868

THE CONJUGACY PROBLEM FOR AUTOMORPHISM GROUPS OF HOMOGENEOUS DIGRAPHS

SAMUEL COSKEY AND PAUL ELLIS

Abstract. We decide the Borel complexity of the conjugacy problem for automorphism groups of countable homogeneous digraphs. Many of the homogeneous digraphs, as well as several other homogeneous structures, have already been addressed in [5] and [4]. In this article we complete the program, and establish a dichotomy theorem that this complexity is either the minimum or the maximum among relations which are classifiable by countable structures. We also discuss the possibility of extending our results beyond graphs to more general classes of countable homogeneous structures.

1. Introduction

This article is a contribution to the study of the automorphism groups of countable homogeneous digraphs. We use the term *digraph* in the model-theoretic sense to mean an oriented simple graph. A countable digraph is said to be *homogeneous* if every finite partial automorphism extends to a total automorphism. A survey of the study of countable homogeneous structures can be found in [9]; the countable homogeneous digraphs are classified in [3].

Our main result will be stated in terms of the Borel complexity theory of equivalence relations. We recall that if E, F are equivalence relations on standard Borel spaces X, Y then we say that E is *Borel reducible* to F if there is a Borel function $f\colon X \to Y$ such that $x \mathrel{E} x'$ if and only if $f(x) \mathrel{F} f(x')$. An equivalence relation E is said to be *smooth* if it is Borel reducible to the equality relation on $\mathbb{R}$. An equivalence relation E is said to be *Borel complete* if it is Borel reducible to an isomorphism relation on a class of countable structures, and conversely any isomorphism relation on a class of countable structures is Borel reducible to E. (Here the countable structures are coded by a sequence of relations on $\mathbb{N}$.) We will use the standard fact that the isomorphism relations on the classes of linear orders and partial orders are Borel complete [6]. For a resource on Borel complexity theory we refer the reader to [7].

Received by the editors September 14, 2015, and in revised form December 9, 2015.
2000 *Mathematics Subject Classification.* 05C35, 05C12.
Key words and phrases. Metric dimension, resolving set, circulant graph.

In [5] and [4], we sought to compute the Borel complexity of the *conjugacy relation* on automorphism groups of numerous countable homogeneous structures. For countable homogeneous digraphs, we were able to decide this complexity in all but three cases, which turned out to be more difficult than the rest. For every digraph that we did analyze, the complexity turned out to be either smooth or Borel complete. In this article we show that in the three remaining cases the conjugacy problems are all Borel complete as well. This completes the proof of the dichotomy:

Theorem 1.1. *If G is a countable homogeneous digraph then the conjugacy problem for* $\mathrm{Aut}(G)$ *is either smooth or Borel complete.*

As we have said, to complete the proof it remains to consider just the three remaining digraphs. These are the generic partial order $\mathcal{P}$ (Section 2), the generic shuffled partial order $\mathcal{P}(3)$ (Section 3), and the semigeneric complete multipartite digraph $\infty \hat{*} I_\infty$ (Section 4). These three proofs showcase many of the tools used in [4] together with some additional tricks.

The conclusion of the proof of Theorem 1.1 suggests one should next ask about the classification of automorphisms of other countable homogeneous structures. In Section 5, we introduce the class of homogeneous structures with the (n-ary) *Borel amalgamation property*, and show how to generalize our methods to apply to such structures.

2. The generic partial order

Let $\mathcal{P}$ denote the generic countable homogeneous partial order. We refer the reader to [10] for the classification of all homogeneous partial orders. In this section we show that the conjugacy problem for $\mathrm{Aut}(\mathcal{P})$ is Borel complete. We begin with the following lemma giving a strong form of the amalgamation property for the class of partial orders.

Lemma 2.1. *Let P be a partial order and for each i let $P \cup \{a_i\}$ be a partial order extending P. Then the transitive closure P' of $P \cup \{a_0, a_1, \ldots\}$ is again a partial order. Also, P' adds no new relations to P.*

Proof. First note that the transitive closure is obtained in just one step by adding a relation $a_i \leq a_j$ whenever there is $p \in P$ such that $a_i \leq p \leq a_j$. To verify transitivity holds after performing this step, suppose that $a_i \leq a_j \leq a_k$. Then there exist $p, q \in P$ such that $a_i \leq p \leq a_j \leq q \leq a_k$. Since $P \cup \{a_j\}$ is a partial order, we must have that $p \leq q$. Therefore $a_i \leq p \leq a_k$ and it follows that $a_i \leq a_k$.

Now suppose towards a contradiction that antisymmetry fails in P'. Since we have added no new relations within P or between elements of P and elements of $\{a_i\}$, there must be distinct i, j such that $a_i \leq a_j \leq a_i$. Then there are $p, q \in P$ such that $a_i \leq p \leq a_j \leq q \leq a_i$. It follows that $p \leq q \leq p$, so by antisymmetry in P, we have $p = q$. But now $a_i \leq p \leq a_i$, which contradicts antisymmetry in $P \cup \{a_i\}$. $\square$

Theorem 2.2. *The isomorphism relation for countable partial orders is Borel reducible to the conjugacy relation on* $\mathrm{Aut}(\mathcal{P})$*. Hence the conjugacy relation on* $\mathrm{Aut}(\mathcal{P})$ *is Borel complete.*

Proof. Given a countable partial order P, we build a copy Q_P of $\mathcal{P}$ and an automorphism ϕ_P of Q_P in such a way that $P \cong P'$ if and only if ϕ_P is conjugate to $\phi_{P'}$. This will be done in stages Q_P^n, ϕ_P^n. To begin, let Q_P^0 consist of $\mathbb{Z}$ many incomparable copies of P together with $\mathbb{N}$ many copies of $\mathbb{Z}$. Denote these latter copies $\mathbb{Z}^{(i)} = \{m^{(i)} \mid m \in \mathbb{Z}\}$ for each $i \in \mathbb{N}$. Let ϕ_P^0 act on Q_P^0 by sending the ith copy of P to the $(i+1)$st copy of P, and sending $m^{(i)}$ to $(m+1)^{(i)}$.

Assume that Q_P^n and ϕ_P^n have been constructed, and construct Q_P^{n+1} and ϕ_P^{n+1} as follows. Begin by considering each set of constraints of the form $\bar{a} < x < \bar{b}$ and $x \perp \bar{c}$ which are consistent with the axioms of a partial order and the diagram of Q_P^n. (For example, we assume that $\bar{a} < \bar{b}$, $\bar{c} \not\leq \bar{a}$, and so forth.) Additionally assume

($\star$) the constraints contain relations of the form $m^{(i)} < x < (m+1)^{(i)}$.

For each such set of constraints, we add a new realization x to Q_P^{n+1}. We do not add any relations involving x except those implied by the constraints and transitivity. Then, we close Q_P^{n+1} under transitivity. By Lemma 2.1, we have added no new edges to Q_P^n.

We also extend ϕ_P^n to an automorphism ϕ_P^{n+1} of Q_P^n in the obvious way: If x is the point corresponding to the parameters $\bar{a}, \bar{b}, \bar{c}$, then we let $\phi_P^{n+1}(x)$ be the point corresponding to the parameters $\phi_P^n(\bar{a}), \phi_P^n(\bar{b}), \phi_P^n(\bar{c})$.

We claim that any element of Q_P is related to just finitely many copies of $\mathbb{Z}$ in Q_P^0. Clearly this claim holds for elements of Q_P^0 itself. Next assume that the claim holds for elements of Q_P^n and consider constraints of the form $\bar{a} < x < \bar{b}$ and $x \perp \bar{c}$ with parameters from Q_P^n. By hypothesis the elements $\bar{a}, \bar{b}, \bar{c}$ are related to just finitely many copies of $\mathbb{Z}$ in Q_P^0 among them all. Adding an element x satisfying this constraint (as done above) and closing under transitivity does not result in x being related to any additional copies of $\mathbb{Z}$. This completes the proof of the claim.

Now to see that Q_P satisfies the one-point extension property, let $\bar{a} < x < \bar{b}$ and $x \perp \bar{c}$ be an arbitrary set of constraints consistent with the axioms of a partial order. Let Q_P^n be the least level containing all of the parameters $\bar{a}, \bar{b}, \bar{c}$. By the claim, these parameters are related to just finitely many copies of $\mathbb{Z}$ in Q_P^0 among them all. By the argument of the claim, it is possible to add a realization x to Q_P^n which is related just to these finitely many copies of $\mathbb{Z}$ in Q_P^0. Let $\mathbb{Z}^{(i)}$ be the first copy of $\mathbb{Z}$ in Q_P^0 not related to x. We now consider the constraints $\bar{a} < y < \bar{b}$ and $y \perp \bar{c}$ and $0^{(i)} < y < 1^{(i)}$. This extended set of constraints is consistent and of the form ($\star$). Hence in the construction we have placed a realization y of the extended constraints into Q_P^{n+1}. This completes the verification that Q_P satisfies the one-point extension property.

Towards a conclusion, observe that for every x in a copy of P in Q_P^0 we have that the ϕ_P-orbit of x is an antichain. On the other hand for every other element x we have that the ϕ_P-orbit of x is a chain. This is because we have some constraint of the form $0^{(i)} < x < 1^{(i)}$. This implies $1^{(i)} = \phi(0^{(i)}) < \phi(x)$, and the two together imply that $x < \phi(x)$.

Thus we can recover the copies of P in Q_P^0 as the set of points whose ϕ_P-orbit is an antichain, and then further recover P. Hence if $\alpha\colon Q_P \cong Q_{P'}$ and $\alpha\phi_P = \phi_{P'}\alpha$ it follows that α restricts to an isomorphism $Q_P^0 \cong Q_{P'}^0$ that sends ϕ_P-orbits to $\phi_{P'}$-orbits. Therefore by passing to the quotient graphs of $Q_P^0, Q_{P'}^0$ by the ϕ_P and $\phi_{P'}$-orbit equivalence relations, we see that α induces an isomorphism $P \cong P'$.

To conclude, we remark briefly on how the construction can be exhibited in a Borel fashion. We fix the underlying sets of $P, Q_P, \mathcal{P}$ to be $\mathbb{N}$. The construction of Q_P can be made Borel by reserving an infinite subset $I_n \subset \mathbb{N}$ for each Q_P^n, and using a previously fixed enumeration of the finite subsets $S \subset I_k$. This immediately implies that the construction of ϕ_P is Borel also. Finally we can regard ϕ_P as an automorphism of $\mathcal{P}$ using a back-and-forth construction between Q_P and $\mathcal{P}$, where each choice in the construction is resolved by choosing the least available witness. $\square$

3. The generic shuffled partial order

The generic shuffled partial order, denoted $\mathcal{P}(3)$, is a graph obtained by "shuffling" three disjoint dense subsets of $\mathcal{P}$ in the following fashion.

Definition 3.1. *Let $\mathcal{P} = P_0 \sqcup P_1 \sqcup P_2$ be a partition of $\mathcal{P}$ into three dense subsets (that is, whenever $a < b$ there exist p_1, p_2, p_3 such that $a < p_i < b$). Define the shuffled graph $\mathcal{P}(3)$ on the underlying set of $\mathcal{P}$ as follows. First, if $x, y \in P_i$, then set $x \to_{\mathcal{P}(3)} y$ if and only if $x <_{\mathcal{P}} y$. Next for each $i \in \mathbb{Z}/3\mathbb{Z}$, if $x \in P_i$ and $y \in P_{i+1}$ then set*

$$\begin{aligned} x \to_{\mathcal{P}(3)} y &\iff x >_{\mathcal{P}} y \\ x \leftarrow_{\mathcal{P}(3)} y &\iff x \perp_{\mathcal{P}} y \\ x \perp_{\mathcal{P}(3)} y &\iff x <_{\mathcal{P}} y \end{aligned}$$

See Chapter 5 of [3] for the proof that $\mathcal{P}(3)$ is homogeneous. We remark that the construction of $\mathcal{P}(3)$ is similar to that of the digraph $S(3)$, which is obtained by shuffling three disjoint dense subsets of $\mathbb{Q}$. See Section 2.2 of [4] for our treatment of $S(3)$.

The argument of the previous section can be modified to show that the conjugacy problem for $\mathrm{Aut}(\mathcal{P}(3))$ is again Borel complete. In the proof we will let $\mathcal{P}_3$ denote the generic three-colored partial order. (Finite three-colored partial orders form an amalgamation class.) The structure $\mathcal{P}_3$ can be viewed simply as a copy of $\mathcal{P}$ partitioned into three distinguished dense subsets P_0, P_1, P_2.

Theorem 3.2. *The isomorphism relation for countable partial orders is Borel reducible to the conjugacy relation on* $\mathrm{Aut}(\mathcal{P}(3))$*. Hence the conjugacy relation on* $\mathrm{Aut}(\mathcal{P}(3))$ *is Borel complete.*

Proof. Given a partial order P we modify the proof of Theorem 2.2 to build a copy Q_P of $\mathcal{P}_3$ as follows. The first level Q_P^0 is constructed as before, with all vertices of Q_P^0 of color 0. When constructing Q_P^{n+1} we again add elements x satisfying each admissible constraint; in this way elements x of all three colors will be added. Continuing the construction as before, we obtain a structure Q_P which is a copy of $\mathcal{P}_3$, and an automorphism ϕ_P of Q_P. It is again the case that $x \in Q_P^0$ if and only if the ϕ_P-orbit of x is an antichain.

The structure Q_P gives rise to a corresponding copy $Q_P(3)$ of $\mathcal{P}(3)$ obtained by shuffling the colors according to the rules in Definition 3.1. Since ϕ_P preserves the colors of Q_P, it is easy to see that ϕ_P is an automorphism of $Q_P(3)$ too.

For the forward implication of the Borel reduction, beginning with an isomorphism $\alpha\colon P \cong P'$ it gives rise to an automorphism of the first level Q_P^0 (which we recall was monochromatic in color 0). This then extends naturally to a color-preserving automorphism of all Q_P, and hence to an automorphism of $Q_P(3)$. The resulting extension $\bar{\alpha}$ conjugates ϕ_P to $\phi_{P'}$.

For the converse implication, we note that since all the ϕ_P-orbits are monochromatic, they retain their structure even after the shuffle. That is, antichain orbits remain antichain orbits, and non-antichain orbits remain non-antichain orbits. Thus given ϕ_P we can recover a copy of Q_P^0 as the set of antichain orbits and conclude as in the proof of Theorem 2.2. $\square$

4. The semigeneric complete multipartite digraph

We say a digraph is *complete multipartite* if there is a partition of the digraph into maximal antichains (independent sets) A_i such that whenever $i \neq j$, $x \in A_i$, and $y \in A_j$ there is an edge between x and y. (In other words, the non-edge relation $\perp$ is an equivalence relation.) There is a generic complete multipartite digraph, denoted $\infty * I_\infty$. In this section we consider the following variant of $\infty * I_\infty$.

We say that a complete multipartite digraph satisfies the *parity property* if for every two maximal antichains A and B and every distinct $a, a' \in A$ and distinct $b, b' \in B$ there is an even number of edges pointing from the set $\{a, a'\}$ to the set $\{b, b'\}$. By [2, p. 75] there exists a generic complete mulpipartite digraph with the parity property. That is, there exists such a digraph which is homogeneous and universal among all such digraphs. In [2] this digraph is called "semigeneric"; we denote it by $\infty \mathbin{\hat{*}} I_\infty$.

In [4] we showed that the conjugacy problem for $\mathrm{Aut}(\infty * I_\infty)$ is Borel complete. In this section we show that the conjugacy problem for $\mathrm{Aut}(\infty \mathbin{\hat{*}} I_\infty)$ is again Borel complete. We remark that the argument for completeness of $\mathrm{Aut}(\infty * I_\infty)$ cannot be used directly for $\mathrm{Aut}(\infty \mathbin{\hat{*}} I_\infty)$, since the widget used in the proof did not have the parity property (see [4, Figure 2]). The following lemma will help us understand the structure $\infty \mathbin{\hat{*}} I_\infty$ and its automorphisms.

Lemma 4.1. *Let G be a complete multipartite digraph with the parity property, and suppose A, B are two maximal antichains. There exists subsets $R_{AB} \subset A$ and $S_{AB} \subset B$ such that for $x \in A$ and $y \in B$, we have $x \to y$ if and only if we have that $x \in R_{AB} \iff y \in S_{AB}$. Refer to Figure 1 for a diagram of these relationships.*

$$\begin{tikzcd} R_{AB} \arrow[r, Rightarrow] & S_{AB} \arrow[dl] \\ R^c_{AB} \arrow[r, Rightarrow] & S^c_{AB} \arrow[ul] \end{tikzcd}$$

FIGURE 1. The edge relationships between the sets R_{AB}, S_{AB}, and their complements $R^c_{AB} = A \setminus R_{AB}$ and $S^c_{AB} = B \setminus S_{AB}$.

Proof. Let A, B be two maximal antichains and, if it is possible, fix $x_0 \in A$ and $y_0 \in B$ such that $x_0 \to y_0$. We then let $R_{AB} = \{\, x \in A \mid x \to y_0 \,\}$ and let $S_{AB} = \{\, y \in B \mid x_0 \to y \,\}$. We claim these sets have the desired properties. Indeed, applying the parity property to the pairs $\{x_0, x\}$ and $\{y_0, y\}$ one can conclude that $x \to y$ in the cases when $x \in R_{AB}$ and $y \in S_{AB}$ or else $x \notin R_{AB}$ and $y \notin S_{AB}$. Similarly one can conclude that $x \leftarrow y$ in the cases when $x \in R_{AB}$ and $y \notin S_{AB}$ or else $x \notin R_{AB}$ and $y \in S_{AB}$, as desired.

If it is not possible to pick x_0, y_0 above, then instead pick $x_0 \in A$ and $y_0 \in B$ with $x_0 \leftarrow y_0$ and proceed similarly to find R_{BA} and S_{BA}. Then simply let $R_{AB} := S_{BA}$ and $S_{AB} := R^c_{BA}$. □

Theorem 4.2. *The isomorphism relation for countable linear orders is Borel reducible to the conjugacy problem for* $\mathrm{Aut}(\infty \mathbin{\hat{*}} I_\infty)$*. Hence the conjugacy problem for* $\mathrm{Aut}(\infty \mathbin{\hat{*}} I_\infty)$ *is Borel complete.*

Proof. Given a countable linear order L we will recursively construct a copy G_L of $\infty \mathbin{\hat{*}} I_\infty$ together with an automorphism ϕ_L of G_L such that $L \cong L'$ if and only if ϕ_L is conjugate to $\phi_{L'}$. To begin we let $G^0_L = L$ itself, and set ϕ^0_L to be the identity on G^0_L. We remark that G^0_L is multipartite with parts of size 1, and therefore it has the parity property. In the construction we will also require a linear ordering $<^n_L$ on G^n_L, and we initially set $<^0_L$ equal to the given ordering of L.

For the remainder of the proof we fix an enumeration $\tau_k(x, \bar{y})$ of the types consistent with the theory of digraphs.

Assuming G^n_L, ϕ^n_L, $<^n_L$ have been constructed, we construct G^{n+1}_L, ϕ^{n+1}_L, $<^{n+1}_L$ as follows. We consider in turn each pair $k \in \mathbb{N}$ and $S \in (G^n_L)^{<\infty}$

such that the parameterized type $\tau_k(x, S)$ does not contradict the parity property. For each such pair, we put three new points $a_k(S), b_k(S), c_k(S)$ into G_L^{n+1} satisfying $\tau_k(x, S)$.

If $\tau_k(x, S)$ forces x to be in a maximal antichain A of G_L^n, then we will place these new points into A. Otherwise we will create three new antichains with one point each. Formally, if $\tau_k(x, S)$ contains a formula of the form $x \perp s$, we set $a_k(S) \perp b_k(S) \perp c_k(S) \perp a_k(S)$. For future reference, let us say that the type $\tau_k(x, S)$ and the points $a_k(S), b_k(S), c_k(S)$ added in this manner are of *Class 1.* On the other hand, if $\tau_k(x, S)$ does not contain a formula of the form $x \perp s$, then we set $a_k(S) \to b_k(S) \to c_k(S) \to a_k(S)$. We say that the type $\tau_k(x, S)$ and the points $a_k(S), b_k(S), c_k(S)$ added in this manner are of *Class 2.*

Before we describe the rest of the edges of G_L^{n+1}, let us use Lemma 4.1 to define the sets R_{AB} and S_{AB} for every pair of maximal antichains A, B of G_L^n. Note that there is an ambiguity in the lemma, since R_{AB} and S_{AB} may be swapped with their complements. To make the definition uniform throughout our construction, note that each maximal antichain has an element e_A which was added *earliest*. We always choose R_{AB} so that it contains this element e_A.

Now given any element a of Class 1, a was added to some maximal antichain A. For each other maximal antichain B, we additionally specify that a lies in R_{AB} unless its type $\tau_K(S)$ explicitly forces us to put a into R_{AB}^c. This specification determines the remaining edges between elements of Class 1 and elements of G_L^n, and also between any two elements of Class 1.

Next given an element a of Class 2, a lies in its own maximal antichain $A = \{a\}$. For each other element b of G_L^n or of Class 1 we set $a \to b$ unless the type of a explicitly forces us to set $a \leftarrow b$.

It remains only to define the edges between two elements of Class 2. For this we will need to define the linear ordering $<_{n+1}$. First let $\prec_n$ be the lexicographic ordering of $(G_L^n)^{<\infty}$ inherited from $<_n$. We then make the definitions:

- If $d \in G_L^n$ and $d' \in G_L^{n+1} \smallsetminus G_L^n$, set $d <_{n+1} d'$.
- If $k < k'$, set $a_k(S), b_k(S), c_k(S) <_{n+1} a_{k'}(S), b_{k'}(S) c_{k'}(S)$.
- If $k \in \mathbb{N}$ and $S \prec_n S'$, set $a_k(S), b_k(S), c_k(S) <_{n+1} a_k(S'), b_k(S'), c_k(S')$.
- If $k \in \mathbb{N}$ and $S \in (G_L^n)^{<\infty}$, set $a_k(S) <_{n+1} b_k(S) <_{n+1} c_k(S)$.

Now if $\tau_k(S)$ and $\tau_{k'}(S')$ are two types of Class 2, we set $a_k(S), b_k(S), c_k(S) \to a_{k'}(S'), b_{k'}(S'), c_{k'}(S')$ precisely when $a_k(S) <_{n+1} a_{k'}(S')$.

Finally we extend the automorphism ϕ_L^n of G_L^n to an automorphism of G_L^{n+1}. Given a type $\tau_k(x, S)$ of the form considered above, we let $S' = \phi_n(S)$. Then we let ϕ_L^{n+1} map $a_k(S) \mapsto b_k(S')$, $b_k(S) \mapsto c_k(S')$, and $c_k(S) \mapsto a_k(S')$. This completes the construction.

Letting $G_L = \bigcup G_L^n$ we have that G_L is complete multipartite, satisfies the parity property, and has the one-point extension property with respect to this class. Thus G_L is a copy of $\infty \mathbin{\hat{*}} I_\infty$. Letting $\phi_L = \bigcup \phi_L^n$ we have that ϕ_L is an automorphism of G_L. As in our previous arguments, it is not hard to see that an isomorphism $L \cong L'$ gives rise to a conjugacy between ϕ_L and $\phi_{L'}$. Moreover we can recover L as the set of fixed points of ϕ_L, which guarantees that a conjugacy between ϕ_L and $\phi_{L'}$ gives rise to an isomorphism $L \cong L'$. $\square$

We conclude this section with a note motivating the need for the involved proof of the previous theorem. The "bowtie" structure of Figure 1 has an automorphism swapping R_{AB} with R^c_{AB} and S_{AB} with S^c_{AB}. One might expect that this symmetry can be used to build automorphisms of $\infty \mathbin{\hat{*}} I_\infty$. However the following result shows that these partial automorphisms do not extend to $\infty \mathbin{\hat{*}} I_\infty$, necessitating the more complicated construction above.

Proposition 4.3. *If ϕ is an automorphism of $\infty \mathbin{\hat{*}} I_\infty$ which fixes each maximal antichain setwise then ϕ is the identity.*

Proof. Suppose ϕ fixes each antichain setwise, and assume towards a contradiction that $a, \phi(a)$ are unequal and lie in the same maximal antichain A.

Case 1: $\phi^2(a) \neq a$.

By genericity there exists a maximal antichain B such that $a, \phi(a) \in R_{AB}$ and $\phi^2(a) \in R^c_{AB}$. Then $a, \phi(a) \in R_{AB}$ implies that $\phi(R_{AB}) = R_{AB}$. But on the other hand $\phi(a) \in R_{AB}, \phi^2(a) \in R^c_{AB}$ implies that $\phi(R_{AB}) = R^c_{AB}$. This is a contradiction.

Case 2: $\phi^2(a) = a$ and there exists another element $a' \in A$ besides $a, \phi(a)$ such that $\phi(a') \neq a'$.

There exists a maximal antichain B such that $a, \phi(a), a'$ lie in R_{AB} and $\phi(a')$ lies in R^c_{AB}. Then we reach a contradiction similarly to Case 1.

Case 3: $\phi^2(a) = a$ and some other element $a' \in A$ besides $a, \phi(a)$ is fixed by ϕ.

There exists a maximal antichain B such that a lies in R_{AB} and $\phi(a), a'$ lie in R^c_{AB}. We again reach a contradiction similarly to Case 1.

$\square$

5. Toward a more general theorem

In this final section we discuss the possibility of generalizing the methods of [5, 4] and the present article to classes of homogeneous structures other than graphs. More specifically, observe that many of our results are established by finding a Borel reduction from the isomorphism relation on the class of countable substructures of some countable homogeneous structure M to the conjugacy relation on $\mathrm{Aut}(M)$. It is natural to ask whether it is always possible to find such a reduction.

Recently, Bilge and Melleray showed that the answer is yes in the case when M has a property which is stronger than homogeneity. In order to describe this result, let us recall some notions from Fraïssé theory. In the following let $\mathcal{K}$ be a class of finite structures which is closed under substructures, and assume for simplicity that the structures are in a finite, relational language. We say that $\mathcal{K}$ has the *amalgamation property* (AP) if for any $A, B_1, B_2 \in \mathcal{K}$ and embeddings $g_i \colon A \to B_i$, there exists $C \in \mathcal{K}$ and embeddings $f_i \colon B_i \to C$ satisfying

$$f_1 \circ g_1 = f_2 \circ g_2.$$

Let $\mathcal{K}_\omega$ be the class of countable structures, all of whose finite substructures lie in $\mathcal{K}$. The main result of Fraïssé's theory states that $\mathcal{K}$ has the amalgamation property if and only if there exists a homogeneous structure $M \in \mathcal{K}_\omega$ such that the class of finite substructures of M is exactly $\mathcal{K}$. Such an M is necessarily unique up to isomorphism, and it is called the *Fraïssé limit* of $\mathcal{K}$.

The AP can be strengthened in a variety of ways. We say that $\mathcal{K}$ has the *strong amalgamation property* (SAP) if in the definition of AP we additionally have $f_1(B_1) \cap f_2(B_2) = f_1 \circ g_1(A) = f_2 \circ g_2(A)$. Finally we say that $\mathcal{K}$ has the *free amalgamation property* (FAP) if in the definition of SAP we can additionally assume there are no nontrivial relations between the sets $f_1(B_1) \smallsetminus f_1 \circ g_1(A)$ and $f_2(B_2) \smallsetminus f_2 \circ g_2(A)$. It is well-known that if $\mathcal{K}$ has the SAP or FAP, then $\mathcal{K}_\omega$ does too.

Theorem 5.1 ([1])**.** *Let $\mathcal{K}$ be a class of finite structures with the FAP, and let M be the Fraïssé limit. Then the isomorphism relation on $\mathcal{K}_\omega$ is Borel reducible to the conjugacy relation on* $\mathrm{Aut}(M)$*.*

The three structures considered in this article, as well as some of those considered in [4], do not have the FAP. Thus it is natural to seek a property that is weaker than the FAP but still strong enough to establish the conclusion of Theorem 5.1. In the rest of this section, we begin to do just that.

In the proof of Theorem 5.1 the authors do not use the FAP directly, but rather an extension construction: given a structure $A \in \mathcal{K}_\omega$, they build an extension $E(A) \in \mathcal{K}_\omega$ which witnesses all admissible finitary types over A. Here a quantifier-free type $\tau(x, \bar{a})$ with finitely many parameters from A is said to be an *admissible finitary type* over A (with respect to $\mathcal{K}_\omega$) if there exists $B \in \mathcal{K}_\omega$ such that A is a substructure of B and B contains a witness for τ. (This definition is written in slightly more general terms than the one given in [1].)

Note that there exists a map that assigns to each $A \in K_\omega$ a structure $E(A) \in K_\omega$ which contains witnesses for all admissible finitary types over A if and only if $\mathcal{K}_\omega$ has the SAP. However we have not been able to use the SAP alone to establish the conclusion of Theorem 5.1. To establish the desired Borel reduction, one should at least be able to construct $E(A)$ in an explicit fashion.

Definition 5.2. *Let $\mathcal{K}$, $\mathcal{K}_\omega$ be as above. We say that K_ω has the* Borel amalgamation property *(BAP) if there is a Borel assignment $E \colon \mathcal{K}_\omega \to \mathcal{K}_\omega$*

such that $E(A)$ contains A and $E(A)$ contains witnesses for all admissible finite types over A.

Thus with this definition, we have that

$$\text{FAP} \implies \text{BAP} \implies \text{SAP}.$$

The first implication is not reversible, since for example we will shortly check that the classes of countable partial orders and countable tournaments have the BAP but not the FAP. We conjecture that the second implication is also not reversible.

In order to establish a generalization of Theorem 5.1, we still require more than just the BAP. For lack of a better name, let us say that $\mathcal{K}_\omega$ has the *Automorphic BAP* (ABAP) if it has the BAP and the following properties:

(a) For any $A \in \mathcal{K}_\omega$ and automorphism $\phi\colon A \to A$, there is an automorphism $\tilde{\phi}\colon E(A) \to E(A)$ extending ϕ with no fixed points in $E(A) \smallsetminus A$. Moreover $\tilde{\phi}$ can be selected from the parameters A and ϕ in a Borel fashion.

(b) For any $A, A' \in \mathcal{K}_\omega$ and isomorphism $\alpha\colon A \to A'$, there is an extension $\hat{\alpha}\colon E(A) \to E(A')$ with the following property: given $A, A' \in \mathcal{K}_\omega$ and automorphisms ϕ, ϕ' of A, A' respectively, if α conjugates ϕ to ϕ' then $\hat{\alpha}$ conjugates $\tilde{\phi}$ to $\tilde{\phi}'$.

While these items may seem somewhat specialized, we will shortly give examples where (a) and (b) are satisfied in a natural way. One might find it natural to to write such a definition using the language of category theory, and indeed something like this has been done in the recent article [8]. However, in that article the authors exchange Borel definability of the extension mappings for a functorial condition on the isomorphism extensions. For our purposes, the ABAP is finely tuned to yield the following generalization of Theorem 5.1.

Theorem 5.3. *Let $\mathcal{K}$, $\mathcal{K}_\omega$ be as above, let M be the Fraïssé limit of $\mathcal{K}$, and suppose $\mathcal{K}_\omega$ satisfies the ABAP. Then the isomorphism relation on $\mathcal{K}_\omega$ is Borel reducible to the conjugacy relation on* $\operatorname{Aut}(M)$.

Proof. The details of the proof are straightforward and we provide just a sketch. Given $A \in K_\omega$, we let $A_0 = A$, $A_{n+1} = E(A_n)$, and $A_\infty = \bigcup A_n$. Meanwhile we let $\phi_A^0 = \mathrm{id}_A$, $\phi_A^{n+1} = \widetilde{\phi_A^n}$, and $\phi_A = \bigcup \phi_A^n$. By the BAP A_∞ is indeed a copy of M, and clearly ϕ_A is an isomorphism of A_∞.

We claim that the assignment $A \mapsto \phi_A$ is the desired Borel reduction. Indeed, if $\alpha\colon A \to A'$ is an isomorphism, then use condition (b) repeatedly to obtain an isomorphism $\alpha_\infty\colon A_\infty \to A'_\infty$. Now condition (b) ensures that α_∞ conjugates ϕ_A to $\phi_{A'}$.

For the reverse implication, note that for any $A \in K_\omega$, condition (a) together with the definition of ϕ_A implies that the set of fixed points of ϕ_A is precisely $A_0 = A$. Thus if $A, A' \in K_\omega$ and ϕ_A is conjugate to $\phi_{A'}$, then the conjugator witnesses that A is isomorphic to A'. $\square$

Turning to examples, the class of countable partial orders satisfies the ABAP and not the FAP. Given a countable partial order P we let $E(P)$ consist of P together with two new witnesses for every admissible finitary type over P. By Lemma 2.1 after closing under transitivity, $E(P)$ is a partial order extending P. Now any automorphism $\phi\colon P \to P$ can be extended to $E(P)$ as follows. If b_1, b_2 are the witnesses for the admissible finite type $\tau(x, \bar{a})$, and c_1, c_2 are the witnesses for the type $\tau(x, \phi(\bar{a}))$, then we let $\tilde{\phi}(b_1) = c_2$ and $\hat{\alpha}(b_2) = c_1$. Moreover any isomorphism $\alpha\colon P \to P'$ naturally extends to $\hat{\alpha}\colon E(P) \to E(P')$ by mapping the witnesses of a given admissible finitary type over P to the witnesses of the corresponding type over P'. It is easy to see that this satisfies the requirements of the ABAP. Of course the FAP fails because in any construction of $E(P)$ it is necessary to close the order relation under transitivity.

For a second example, the class of countable tournaments also satisfies the ABAP and (clearly) not the FAP. To establish the ABAP, we adapt the construction of [8, Example 2.10]. Given a countable tournament T we let $E(T)$ consist of T together with elements with names (s, n) where s is a finite sequence from T and $n \in \mathbb{Z}$. For an old vertex $v \in T$ and a new vertex (s, n), set $v \to (s, n)$ if v appears in s and $(s, n) \to v$ otherwise. The edges between the new vertices are defined lexicographically: if $|s| < |s'|$ we let $(s, n) \to (s', n')$; if $|s| = |s'|$, i is the first coordinate where they differ, and $s_i \to s'_i$, we let $(s, n) \to (s', n')$. Finally we let $(s, n) \to (s, n + k)$ for all $k > 0$. Now any automorphism ϕ of T can be extended to $E(T)$ by letting $\tilde{\phi}(s, n) = (\phi(s), n + 1)$, where $\phi(s)$ denotes the coordinatewise application of ϕ. Moreover any isomorphism $\alpha\colon T \to T'$ naturally extends to $E(T) \to E(T')$ by mapping $\hat{\alpha}(s, n) = (\alpha(s), n)$. Once again it is not difficult to check that this satisfies our requirements.

Thus Theorem 5.3 is strong enough to subsume some results from [4], provided one does the extra work of verifying that the ABAP holds in such cases. On the other hand, in some cases we are able to establish the same conlusion even when we don't know that the BAP holds. For example in the case of the semigeneric multipartite digraph, we don't how to define $E(A)$ for every A in the class. Instead we restricted to the subclass of linear orders, and made certain arbitrary choices according to an inherited linear order on the types.

To conclude, we remark that while the ABAP is somewhat specialized, the BAP and related notions may be of broader interest. As was mentioned above, it is known that if $\mathcal{K}$ has the SAP then $\mathcal{K}_\omega$ has the SAP too. Indeed this can be established by either iteratively applying the finitary amalgamation property, or using a compactness argument. In either proof, one does not arrive at an explicit and uniquely determined amalgam: it relies on the choice of the enumeration of the finite substructures, or on the weak form of the Axiom of Choice that is the compactness theorem. A Borel form of SAP such as BAP remedies this use of AC by isolating classes where countable amalgamation can be done in an explicit way.

Acknowledgement

We would like to thank Gregory Cherlin for helpful discussions regarding this material.

References

1. D. Bilge and J. Melleray, *Elements of finite order in automorphism groups of homogeneous structures*, Contrib. Discrete Math. **8** (2013), no. 2, 88–119.
2. G. L. Cherlin, *Homogeneous directed graphs. The imprimitive case*, Logic colloquium '85 (Orsay, 1985), Stud. Logic Found. Math., vol. 122, North-Holland, Amsterdam, 1987, pp. 67–88, `doi`.
3. ______, *The classification of countable homogeneous directed graphs and countable homogeneous n-tournaments*, vol. 131, Mem. Am. Math. Soc., no. 621, American Mathematical Society, Providence, RI, 1998, `doi`.
4. S. Coskey and P. Ellis, *The conjugacy problem for automorphism groups of countable homogeneous structures*, Preprint (2015).
5. S. Coskey, P. Ellis, and S. Schneider, *The conjugacy problem for the automorphism group of the random graph*, Arch. Math. Logic **50** (2011), no. 1-2, 215–221, `doi`.
6. H. Friedman and L. Stanley, *A Borel reducibility theory for classes of countable structures*, J. Symbolic Logic **54** (1989), no. 3, 894–914.
7. S. Gao, *Invariant descriptive set theory*, Pure and Applied Mathematics (Boca Raton), vol. 293, CRC Press, Boca Raton, FL, 2009.
8. W. Kubiś and D. Mašulović, *Katětov functors*, Preprint (2015).
9. D. Macpherson, *A survey of homogeneous structures*, Discrete Math. **311** (2011), no. 15, 1599–1634, `doi`.
10. J. H. Schmerl, *Countable homogeneous partially ordered sets*, Algebra Universalis **9** (1979), no. 3, 317–321, `doi`.

Samuel Coskey, Department of Mathematics, Boise State University, 1910 University Drive, Boise, ID, 83725
E-mail address: `scoskey@nylogic.org`
URL: `scoskey.github.io`

Paul Ellis, Department of Mathematics and Computer Science, Manhattanville College, 2900 Purchase Street, Purchase, NY, 10577
E-mail address: `paulellis@paulellis.org`
URL: `paulellis.org`