

ALL REDUCTS OF THE RANDOM GRAPH ARE MODEL-COMPLETE

MANUEL BODIRSKY AND MICHAEL PINSKER

ABSTRACT. We study locally closed transformation monoids which contain the automorphism group of the random graph. We show that such a transformation monoid is locally generated by the permutations in the monoid, or contains a constant operation, or contains an operation that maps the random graph injectively to an induced subgraph which is a clique or an independent set.

As a corollary, our result yields a new proof of Simon Thomas' classification of the five closed supergroups of the automorphism group of the random graph; our proof uses different Ramsey-theoretic tools than the one given by Thomas, and is perhaps more straightforward.

Since the monoids under consideration are endomorphism monoids of relational structures definable in the random graph, we are able to draw several model-theoretic corollaries: One consequence of our result is that all structures with a first-order definition in the random graph are model-complete. Moreover, we obtain a classification of these structures up to existential interdefinability.

1. INTRODUCTION

The random graph (also called the Rado graph) is the graph $G = (V; E)$ defined uniquely up to isomorphism by the property that for all finite disjoint subsets U, U' of the countably infinite vertex set V there exists a vertex $v \in V \setminus (U \cup U')$ such that v is in G adjacent to all vertices in U and to no vertex in U' ; we will refer to this property of the random graph as the *extension property*. For the many remarkable properties of this graph and its automorphism group, and various connections to many branches of mathematics, see e.g. [7, 8].

Simon Thomas has classified the five locally closed supergroups of the automorphism group of G in [18]. In this paper we more generally investigate locally closed transformation monoids that contain the automorphism group of G . We show that every such monoid is either a disguised group in the sense that it is generated by the largest permutation group which it contains, or it contains a constant operation, or an injective operation which either deletes all edges or all non-edges of the random graph. As a by-product of our proof, we obtain a new proof of Thomas' classification.

Not surprisingly, insights on the behavior of functions on G have consequences for model-theoretic questions concerning the random graph. Every closed supergroup of the automorphism group of G is the automorphism group of a relational structure definable in G ; such structures are called *reducts* of G . Moreover, two reducts Γ_1, Γ_2 have the same automorphism group iff they are first-order interdefinable, i.e., iff every relation of Γ_1 has a first-order definition in Γ_2 , and vice-versa. Thus, Thomas' theorem is the classification of the reducts of G up to first-order interdefinability. By considering monoids of self-embeddings instead of automorphism groups, we obtain a finer classification of these reducts, namely up to *existential* interdefinability, i.e., we do not distinguish between two structures Γ_1, Γ_2 whenever every relation of Γ_1 is definable in Γ_2 by an existential first-order formula, and vice versa.

Another consequence of our results is that all reducts Γ of the random graph are *model-complete*, i.e., all embeddings between models of the first-order theory of Γ preserve all first-order formulas. The analogous statement for the reducts of $(\mathbb{Q}; <)$, the dense linear order of the rationals, has been observed in [3]. Model-completeness is a central concept in model theory; see e.g. [11]. For

2000 *Mathematics Subject Classification*. Primary 03C40; secondary 08A40; 08A70; 03D15.

Key words and phrases. random graph, reduct, automorphism, endomorphism, model-completeness, invariant relation, existential positive definition, Galois connection, locally closed monoid.

The second author is grateful for support through Erwin Schrödinger Fellowship N2742-N18 of the Austrian Science Fund.

example, model-completeness plays an important role when establishing quantifier-elimination results. Whether or not a structure is model-complete is usually not preserved by first-order interdefinability. In this light, the result that all reducts of the random graph are model-complete might be surprising.

The results presented are also relevant for the study of the *constraint satisfaction problem* for structures with a first-order definition in the random graph. When Γ is a structure with a finite relational signature τ , then the constraint satisfaction problem for Γ ($\text{CSP}(\Gamma)$) is the computational problem of deciding whether a given primitive positive sentence over τ is true in Γ . A formula is called *primitive positive* iff it is of the form $\exists x_1, \dots, x_n. \psi_1 \wedge \dots \wedge \psi_m$ where ψ_1, \dots, ψ_m are atomic. The complexity of $\text{CSP}(\Gamma)$ does not change when Γ is expanded by finitely many relations with a primitive positive definition in Γ . Even though expansions by relations with an existential positive definition might increase the complexity of the constraint satisfaction problem, the classification of the reducts of $(\mathbb{Q}; <)$ up to existential positive interdefinability was an important ingredient in a recent complexity classification for the CSP of such reducts [3]. The results in this paper pave the way for a similar classification for reducts of the random graph.

2. RESULTS

We now present our main results, formulated in terms of transformation monoids and permutation groups; the proofs of these results will be of purely combinatorial nature. The model-theoretic corollaries of the results presented here will be drawn in Section 3.

A monoid \mathcal{M} of mappings from a set D to D is called *(locally) closed* iff the following holds: whenever $f : D \rightarrow D$ is such that for every finite $A \subseteq D$ there exists $e \in \mathcal{M}$ such that $e(x) = f(x)$ for all $x \in A$, then f is an element of \mathcal{M} . Equivalently, the monoid is a closed set in the product topology of D^D , where D is taken to be discrete. For the purposes of this paper, we call the smallest closed transformation monoid that contains a set of operations F from V to V and the automorphism group $\text{Aut}(G)$ of the random graph the monoid *generated* by F .

Similarly, a permutation group acting on D is called *(locally) closed* iff it is locally closed as a monoid. As before, we call the smallest closed group containing a set of permutations F on V as well as $\text{Aut}(G)$ the group *generated* by F .

The random graph contains all countable graphs as induced subgraphs. In particular, it contains an infinite complete subgraph, denoted by K_ω . It is clear that all injective operations from V to V whose image induces K_ω in G locally generate the same monoid. Let e_E be one such injective operation whose image induces K_ω in G . Similarly, G contains an infinite independent set, denoted by I_ω . Let e_N be an injective operation from $V \rightarrow V$ whose image induces I_ω in G .

Our main result is the following. It states that all closed monoids containing $\text{Aut}(G)$ either contain a quite primitive function, or are generated by their permutations. As it turns out, the permutations in such a monoid form a closed group.

Theorem 1. *For any closed monoid \mathcal{M} containing $\text{Aut}(G)$, one of the following cases applies.*

- (1) \mathcal{M} contains a constant operation.
- (2) \mathcal{M} contains e_E .
- (3) \mathcal{M} contains e_N .
- (4) \mathcal{M} is generated by its permutations.

The last case splits into five sub-cases, corresponding to the five locally closed permutation groups that contain $\text{Aut}(G)$. These groups have already been exhibited by Thomas [18]. In our proof of Theorem 1, we will be forced to re-derive this result. While our proof of that classification, being the proof of a more general result, is longer than the one in [18], it might be more canonical (in the sense of Definition 23). We now define the five groups.

It is clear that the complement graph of G is isomorphic to G . Note that any isomorphism between G and its complement locally generates the same transformation monoid (group). Let $-$ be one such isomorphism. We write function applications of $-$ without braces.

For any finite subset S of V , if we flip edges and non-edges between S and $V \setminus S$ in G , then the resulting graph is isomorphic to G (it is straightforward to verify the extension property). Let

i_S be such an isomorphism for each non-empty finite S . Every such operation generates the same transformation monoid (group). We also write sw for $i_{\{0\}}$, where $0 \in V$ is a fixed element for the rest of the paper, and refer to this operation as the *switch*.

Theorem 2 (of [18]). *Let \mathcal{G} be a closed permutation group containing $\text{Aut}(G)$. Then exactly one out of the following five cases is true.*

- (1) \mathcal{G} equals $\text{Aut}(G)$.
- (2) \mathcal{G} is the group generated by $-$.
- (3) \mathcal{G} is the group generated by sw .
- (4) \mathcal{G} is the group generated by $\{-, sw\}$.
- (5) \mathcal{G} is the group of all permutations on V .

The arguments given in [18] use a Ramsey-theoretic result by Nešetřil [15], namely that the class of all finite graphs excluding finite cliques of a fixed size forms a *Ramsey class* (in the sense of [16]). We also use a Ramsey-theoretic result, shown by Rödl and Nešetřil [13, 14] (and independently by [1]), which is different: we need the fact that finite ordered vertex-colored graphs form a Ramsey class. We believe that our approach is canonical, and that the proof techniques could very well be adapted to show similar classifications for supergroups of automorphism groups of other infinite structures Γ which have the property that the class of all finite structures that embed into Γ (possibly equipped with a linear order on the vertices) is a Ramsey class.

3. MODEL-THEORETIC COROLLARIES

We now discuss the results of the preceding section in a model-theoretic setting and establish some corollaries in this language.

One easily verifies that the endomorphism monoid $\text{End}(\Delta)$ (automorphism group $\text{Aut}(\Delta)$) of a structure Δ with domain D is a closed monoid (group) on D , and that every closed monoid (group) is of this form for an adequate structure Δ . Moreover, the automorphism group of a *reduct* Γ of a structure Δ , i.e., of a structure Γ which is first-order definable in Δ , clearly contains $\text{Aut}(\Delta)$. The following is Theorem 1, restated in terms of structures.

Theorem 3. *Let Γ be first-order definable in the random graph. Then one of the following cases applies.*

- (1) Γ has a constant endomorphism.
- (2) Γ has the endomorphism e_E .
- (3) Γ has the endomorphism e_N .
- (4) The endomorphisms of Γ are locally generated by the automorphisms of Γ .

For automorphism groups of reducts of the random graph, we have even more. It is well-known that the random graph is (*ultra-*) *homogeneous*, i.e., every isomorphism between two finite induced substructures of G can be extended to an automorphism of G . For structures with a finite signature, homogeneity implies ω -*categoricity*: all countable models of the first-order theory of G are isomorphic. Reducts of ω -categorical structures are ω -categorical.

Now, the theorem of Engeler, Ryll-Nardzewski, and Svenonius (see e.g. [9, 11]) states that a relation R is first-order definable in an ω -categorical structure Δ if and only if R is preserved by all automorphisms of Δ . As a consequence, the reducts of an ω -categorical structure Δ are, up to first-order interdefinability, in one-to-one correspondence with the locally closed permutation groups containing $\text{Aut}(\Delta)$. To illustrate this, we restate Theorem 2 by means of this connection.

On the random graph, let $R^{(k)}$ be the k -ary relation that holds on $x_1, \dots, x_k \in V$ if x_1, \dots, x_k are pairwise distinct, and the number of edges between these k vertices is odd. Note that $R^{(4)}$ is preserved by $-$, $R^{(3)}$ is preserved by sw , and that $R^{(5)}$ is preserved by $-$ and by sw , but not by all permutations of V .

Theorem 4 (of [18]). *Let Γ be a structure with a first-order definition in the random graph $(V; E)$. Then exactly one out of the following five cases is true.*

- (1) Γ is first-order interdefinable with $(V; E)$.
- (2) Γ is first-order interdefinable with $(V; R^{(4)})$.

- (3) Γ is first-order interdefinable with $(V; R^{(3)})$.
- (4) Γ is first-order interdefinable with $(V; R^{(3)}, R^{(4)})$.
- (5) Γ is first-order interdefinable with $(V; =)$.

For any reduct Γ , a case of Theorem 4 applies iff the case with the same number applies for $\text{Aut}(\Gamma)$ in Theorem 2. We will not prove this relational description in this paper; however, given Theorem 2 and the discussion above, verifying the equivalence is merely an exercise.

In the same way as automorphisms can be used to characterize first-order definability, self-embeddings can be used to characterize existential definability, and endomorphisms can be used to characterize existential positive definability in ω -categorical structures. This is the content of the following theorem.

Theorem 5. *A relation R has an existential positive (existential) definition in an ω -categorical structure Γ if and only if R is preserved by the endomorphisms (self-embeddings) of Γ .*

Proof. It is easy to verify that existential positive formulas are preserved by endomorphisms, and existential formulas are preserved by self-embeddings of Γ .

For the other direction, note that the endomorphisms and self-embeddings of Γ contain the automorphisms of Γ , and hence the theorem of Ryll-Nardzewski shows that R has a first-order definition in Γ ; let ϕ be a formula defining R . Suppose for contradiction that R is a relation preserved by all endomorphisms of Γ but has no existential positive definition in Γ . We use the well-known homomorphism preservation theorem in model theory (see e.g. [9, 11]), which states that a first-order formula ϕ is equivalent to an existential positive formula modulo a first-order theory T if and only if ϕ is preserved by all homomorphisms between models of T . Since by assumption ϕ is not equivalent to an existential positive formula in Γ , there are models Γ_1 and Γ_2 of the first-order theory of Γ and a homomorphism h from Γ_1 to Γ_2 that violates ϕ . By the Theorem of Löwenheim-Skolem the first-order theory of the two-sorted structure (Γ_1, Γ_2, h) has a countable model $(\Gamma'_1, \Gamma'_2, h')$. Since both Γ'_1 and Γ'_2 must be countably infinite, and because Γ is ω -categorical, we have that Γ'_1 and Γ'_2 are isomorphic to Γ , and h' can be seen as an endomorphism of Γ that violates ϕ ; a contradiction.

The argument for existential definitions and self-embeddings is similar, but we use the preservation theorem of Los-Tarski (see e.g. [9, 11]) instead of the homomorphism preservation theorem. \square

The following proposition, which links the operational generating process with preservation of relations of structures, is easy to prove; see e.g. [17].

Proposition 6. *Let F, H be sets of mappings from V to V . Then the monoid generated by F contains H iff every relation definable in G and preserved by F is also preserved by H .*

Using Theorem 5, we obtain an interesting and perhaps surprising consequence of our main result. A theory T is called *model-complete* if every embedding between models of T is elementary, i.e., preserves all first-order formulas. It is well-known that a theory T is model-complete if and only if every first-order formula is modulo T equivalent to an existential formula. A structure is said to be model-complete iff its first-order theory is model-complete. From the definition of model-completeness and ω -categoricity it is easy to see that an ω -categorical structure Γ is model-complete if all embeddings of Γ into itself preserve all first-order formulas. It has been observed in [3] (based on a proof of [12] of a result by Cameron [6]) that all reducts of the linear order of the rationals $(\mathbb{Q}; <)$ are model-complete. We now see that the same is true for the random graph.

Corollary 7. *Every structure Γ with a first-order definition in the random graph is model-complete.*

Proof. An ω -categorical structure Γ is model-complete if and only if all embeddings of Γ into itself are locally generated by the automorphisms of Γ . To see this, first assume that the automorphisms of Γ locally generate the self-embeddings of Γ , and let ϕ be a first-order formula. By the equivalent characterization of model-completeness mentioned above it suffices to show that ϕ is equivalent to an existential formula. Since ϕ is preserved by automorphisms of Γ , it is by Proposition 6 also

preserved by self-embeddings of Γ . Then Theorem 5 implies that ϕ is equivalent to an existential formula. Conversely, suppose that all first-order formulas are equivalent to an existential formula in Γ . Since existential formulas are preserved by self-embeddings of Γ , also the first-order formulas are preserved by self-embeddings of Γ . By Proposition 6, the self-embeddings are locally generated by the automorphisms of Γ .

We thus show that the embeddings of Γ are generated by its automorphisms. Note that when we expand Γ by \neq and by $\neg R$ for every relation in Γ , then the resulting structure has the same set of self-embeddings. Hence, we assume in the following that Γ contains \neq and $\neg R$ for all relations R , and hence that all endomorphisms of Γ are embeddings. We apply Theorem 1. If Case (4) of the theorem holds, we are done. Note that Γ cannot have a constant endomorphism since Γ contains \neq . So suppose that Γ is preserved by e_N . The substructure of Γ induced by $e_N(V)$ has a first-order definition in $(e_N(V); =)$; but then, since e_N is an embedding, Γ has a first-order definition in $(V; =)$. The set of self-embeddings of Γ is then the set of all injective mappings from V to V . It follows that Γ is model-complete, because every injective mapping from V to V is locally generated by the permutations of V (i.e., the automorphisms of Γ). The argument for e_E is analogous. \square

In ω -categorical structures, homogeneity is equivalent to having *quantifier-elimination*: every first-order formula is in G equivalent to a quantifier-free first-order formula; hence, the random graph has quantifier elimination. The same is not true for its reducts. For example, any two 2-element substructures of the structure

$$\Gamma = (V; \{(x, y, z) \mid E(x, y) \wedge \neg E(y, z)\})$$

are isomorphic. But since there is a first-order definition of G in Γ , an isomorphism between a 2-element substructure with an edge and a 2-element substructure without an edge cannot be extended to an automorphism of Γ . However, our results imply that a structure Γ with a first-order definition in the random graph is homogeneous when Γ is expanded by all relations with an existential definition in Γ .

Corollary 8. *Every structure Γ with a first-order definition in the random graph has quantifier-elimination if it is expanded by all relations with an existential definition in Γ .*

Proof. This follows directly from the fact mentioned above that in model-complete structures first-order formulas are equivalent to existential formulas (see e.g. [11]). \square

As another application of our main theorem, we refine Theorem 4 by giving a finer (at least in theory) classification of the reducts of the random graph.

Corollary 9. *Up to existential interdefinability, there are exactly five different structures with a first-order definition in the random graph.*

Proof. In the same way as in the proof of Corollary 9, we can use Theorem 1 to show that either the self-embeddings of a reduct Γ are generated by the automorphisms, and Γ is existentially interdefinable with one of the structures described in Theorem 2; or otherwise Γ has an existential definition in $(V; =)$, which is again one of the five cases from Theorem 2. \square

The endomorphism monoid $\text{End}(G)$ of the random graph has been studied in [4, 5, 10]. By Theorem 5, studying closed transformation monoids containing $\text{End}(G)$ is equivalent to studying structures with a first-order definition in G up to existential positive interdefinability. A complete classification of all locally closed transformation monoids that contain all permutations of V , and hence of the reducts of $(V; =)$ up to existential positive interdefinability, has been given in [2]; there is only a countable number of such monoids. The results of the present paper are far from providing a full classification of the locally closed transformation monoids that contain the automorphisms of the random graph — this is left for future investigation.

4. ADDITIONAL NOTIONS AND NOTATION

We will write $E(x, y)$ or $(x, y) \in E$ to express that two vertices $x, y \in V$ are adjacent in the random graph. The binary relation $N(x, y)$ is defined by $\neg E(x, y) \wedge x \neq y$. Pairs $\{x, y\}$ with $N(x, y)$ are referred to as *non-edges*.

Often when we have a graph $\mathcal{P} = (P; D)$, and $S \subseteq P$, then for notational simplicity we write $(S; D)$ for the subgraph of \mathcal{P} induced by S , i.e., we ignore the fact that D would have to be restricted to S^2 .

We say that an operation $e : V \rightarrow V$ (a set F of operations from V to V) is *generated* by a set of operations H from V to V iff it is contained in the monoid generated by H .

5. RAMSEY-THEORETIC PRELIMINARIES

We prepare the proof of our main theorem by recalling some Ramsey-type theorems and extending these theorems for our purposes. The notions and results of this section are of an abstract Ramsey-theoretic nature and do not refer to concrete structures such as the random graph.

We start by recalling a theorem on ordered structures due to Nešetřil and Rödl [13] which we will make heavy use of. Let $\tau = \tau' \cup \{\prec\}$ be a relational signature, and let $\mathcal{C}(\tau)$ be the class of all finite τ -structures \mathcal{S} where \prec denotes a linear order on the domain of \mathcal{S} . For τ -structures \mathcal{A}, \mathcal{B} , let $\binom{\mathcal{A}}{\mathcal{B}}$ be the set of all substructures of \mathcal{A} that are isomorphic to \mathcal{B} (we also refer to members of $\binom{\mathcal{A}}{\mathcal{B}}$ as *copies of \mathcal{B} in \mathcal{A}*). For a finite number $k \geq 1$, a *k-coloring* of the copies of \mathcal{B} in \mathcal{A} is simply a mapping χ from $\binom{\mathcal{A}}{\mathcal{B}}$ into a set of size k .

Definition 10. For $\mathcal{S}, \mathcal{H}, \mathcal{P} \in \mathcal{C}(\tau)$ and $k \geq 1$, we write $\mathcal{S} \rightarrow (\mathcal{H})_k^{\mathcal{P}}$ iff for every k -coloring χ of the copies of \mathcal{P} in \mathcal{S} there exists a copy \mathcal{H}' of \mathcal{H} in \mathcal{S} such that all copies of \mathcal{P} in \mathcal{H}' have the same color (under χ).

Theorem 11 (of [1,13,14]). *The class $\mathcal{C}(\tau)$ of all finite relational ordered τ -structures is a Ramsey class, i.e., for all $\mathcal{H}, \mathcal{P} \in \mathcal{C}(\tau)$ and $k \geq 1$ there exists $\mathcal{S} \in \mathcal{C}(\tau)$ such that $\mathcal{S} \rightarrow (\mathcal{H})_k^{\mathcal{P}}$.*

Corollary 12. *For every finite graph \mathcal{H} and for all colorings χ_E and χ_N of the edges and the non-edges of G , respectively, by finitely many colors, there exists an isomorphic copy of \mathcal{H} in G which both colorings are constant on.*

Proof. Let k be the number of colors used altogether by χ_E and χ_N . Let \prec be any total order on the domain of \mathcal{H} , and denote the structure obtained from \mathcal{H} by adding the order \prec to the signature by \mathcal{H} . Consider the complete graph \mathcal{K}_2 on two vertices, and order its two vertices anyhow to arrive at a structure \mathcal{K}_2 . Then the coloring χ_E of the edges of \mathcal{H} can be viewed as a coloring of the copies of \mathcal{K}_2 in \mathcal{H} . Let \mathcal{T} with $\mathcal{T} \rightarrow (\mathcal{H})_k^{\mathcal{K}_2}$ be provided by the preceding theorem, and let \mathcal{S} be \mathcal{T} without the order. Then \mathcal{S} is a graph with the property that whenever we color its edges with k colors, then there is a copy of \mathcal{H} in \mathcal{S} all of whose edges have the same color. Now we repeat the argument for the non-edges, starting from \mathcal{S} instead of \mathcal{H} . We then arrive at a graph \mathcal{T} with the property that whenever we color its edges and non-edges by k colors, then there is a copy \mathcal{H}' of \mathcal{H} in \mathcal{T} such that all edges of \mathcal{H}' have the same color, and such that non-edges of \mathcal{H}' have the same color. \mathcal{T} has a copy in G , proving the claim. \square

We will not only need to color edges of graphs, but also of graphs equipped with additional structure.

Definition 13. An *n-partitioned graph* is a structure $\mathcal{U} = (U; F, U_1, \dots, U_n)$, where $(U; F)$ is a graph and each U_i is a subset of U such that the U_i form a partition of U .

Definition 14. Let $\mathcal{U} = (U; F)$ be a graph, and let S_1, S_2 be disjoint subsets of U . Let χ be a coloring of the two-element subsets of U . We say that χ is *canonical on S_1* iff the color of a two-element subset of S_1 depends only on whether this set is an edge or a non-edge. Similarly, we say that χ is *canonical between S_1 and S_2* iff the color of every pair $\{s_1, s_2\}$, where $s_1 \in S_1$ and $s_2 \in S_2$, depends only on whether or not this pair is an edge.

Definition 15. Let $\mathcal{U} = (U; F, U_1, \dots, U_n)$ be an n -partitioned graph. We say that a coloring of the two-element subsets of U is *canonical on \mathcal{U}* iff it is canonical on all U_i and between all distinct U_i, U_j .

Lemma 16 (The n -partitioned graph Ramsey lemma). *Let $n, k \geq 1$. For any finite n -partitioned graph $\mathcal{U} = (U; F, U_1, \dots, U_n)$ there exists a finite n -partitioned graph $\mathcal{Q} = (Q; D, Q_1, \dots, Q_n)$ with the property that for all colorings of the two-element subsets of Q with k colors, there exists a copy of \mathcal{U} in \mathcal{Q} on which the coloring is canonical.*

Proof. We show the lemma for $n = 2$; the generalization to larger n is straightforward. For $n = 2$, we apply Theorem 11 six times: Once for the edges in U_1 , once for the edges in U_2 , once for the edges between U_1 and U_2 , and then the same for all three kinds of non-edges.

In general, we would have to apply the theorem 2 ($n + \binom{n}{2}$) times: Once for the edges of each part U_i , once for the edges between any two distinct parts U_i, U_j , and then the same for all non-edges on and between parts.

So assume $n = 2$. We exhibit the idea in detail for the edges between U_1 and U_2 . Let \prec be any total order on U with the property that $u_1 \prec u_2$ for all $u_1 \in U_1, u_2 \in U_2$. Consider the 2-partitioned graph $\mathcal{L}^1 = (\{a, b\}; \{(a, b), (b, a)\}, \{a\}, \{b\})$ and order its vertices by setting $a \prec b$; so \mathcal{L}^1 consists of two adjacent vertices which are ordered somehow, and which lie in different parts. By Theorem 11, there exists an ordered partitioned graph $\mathcal{Q}^1 = (Q^1; D^1, Q_1^1, Q_2^1, \prec)$ such that $\mathcal{Q}^1 \rightarrow (\mathcal{U})_k^{\mathcal{L}^1}$.

Now, if we change the order on \mathcal{Q}^1 in such a way that $r \prec s$ for all $r \in Q_1^1$ and all $s \in Q_2^1$ and such that the order within the parts Q_1^1, Q_2^1 remains unaltered, then the statement $\mathcal{Q}^1 \rightarrow (\mathcal{U})_k^{\mathcal{L}^1}$ still holds: For, given a coloring of the copies of \mathcal{L}^1 with respect to the new ordering, we obtain a coloring of (possibly fewer) copies of \mathcal{L}^1 with respect to the old ordering. There, we obtain a copy \mathcal{U}' of \mathcal{U} such that all copies of \mathcal{L}^1 in \mathcal{U}' have the same color. But in this copy, by the choice of the order on \mathcal{U} , we have that $r \prec s$ for all $r \in U_1'$ and all $s \in U_2'$. Therefore, this copy is also a substructure of \mathcal{Q}^1 with respect to the new ordering.

Since we can change the ordering on \mathcal{Q}^1 in the way described above, the colorings of the copies of \mathcal{L}^1 are just colorings of those pairs $\{r, s\}$, with $r \in Q_1^1$ and $s \in Q_2^1$, which are edges.

Now we repeat the process with the structure $\mathcal{L}^2 = (\{a, b\}; \{(a, b), (b, a)\}, \{a, b\}, \emptyset)$, ordered again by setting $a \prec b$, starting with \mathcal{Q}^1 . We then obtain a structure \mathcal{Q}^2 ; this step takes care of the edges which lie within U_1 . After that we proceed with $\mathcal{L}^3 = (\{a, b\}; \{(a, b), (b, a)\}, \emptyset, \{a, b\})$, thereby taking care of the edges within U_2 . We then apply Theorem 11 three more times with the structures $\mathcal{L}^4 = (\{a, b\}; \emptyset, \{a\}, \{b\})$, $\mathcal{L}^5 = (\{a, b\}; \emptyset, \{a, b\}, \emptyset)$, and $\mathcal{L}^6 = (\{a, b\}; \emptyset, \emptyset, \{a, b\})$, in order to ensure homogeneous non-edges. \square

The preceding lemma on partitioned graphs was an auxiliary tool to cope with graphs which have some distinguished vertices, as defined in the following.

Definition 17. An n -constant graph is a structure $\mathcal{U} = (U; F, u_1, \dots, u_n)$, where $\mathcal{U} = (U; F)$ is a graph, and $u_i \in U$ are distinct.

Observe that n -constant graphs are no relational structures; therefore, in order to apply Theorem 11, we have to make them relational: To every n -constant graph $\mathcal{U} = (U; F, u_1, \dots, u_n)$ we can assign an $n + 2^n$ -partitioned graph $\tilde{\mathcal{U}} = (U; F, \{u_1\}, \dots, \{u_n\}, U_1, \dots, U_{2^n})$ in which the u_i belong to singleton sets, and in which for every possible relative position (edge or non-edge) to the u_i we have a set U_j of all elements in $U \setminus \{u_1, \dots, u_n\}$ having this position. (In the language of model theory, every of the $n + 2^n$ sets corresponds to a maximal quantifier-free 1-type over the structure \mathcal{U} .) We call the parts U_i the *proper parts* of $\tilde{\mathcal{U}}$.

Definition 18. Let $\mathcal{U} = (U; F, u_1, \dots, u_n)$ be an n -constant graph. We say that a coloring of the two-element subsets of U is *canonical on \mathcal{U}* iff it is canonical on the corresponding $n + 2^n$ -partitioned graph.

We now arrive at the goal of this section, namely the following lemma, which we are going to apply to mappings on the random graph numerous times in the sections to come.

Lemma 19 (The n -constant graph Ramsey lemma). *Let $n, k \geq 1$. For any finite n -constant graph $\mathcal{U} = (U; F, u_1, \dots, u_n)$ there exists a finite n -constant graph $\mathcal{Q} = (Q; D, q_1, \dots, q_n)$ with the property that for all colorings of the two-element subsets of Q with k colors, there exists a copy of \mathcal{U} in \mathcal{Q} on which the coloring is canonical.*

Proof. Let $\tilde{\mathcal{U}} := (U; F, \{u_1\}, \dots, \{u_n\}, U_1, \dots, U_{2^n})$ be the partitioned graph associated with \mathcal{U} . We would like to use the partitioned graph Ramsey lemma (Lemma 16) in order to obtain \mathcal{Q} ; but we want the singleton sets $\{u_i\}$ of the partition to remain singletons, which is not guaranteed by that lemma.

So consider the 2^n -partitioned graph $\mathcal{R} := (U \setminus \{u_1, \dots, u_n\}; F, U_1, \dots, U_{2^n})$, and apply the partitioned graph Ramsey lemma to this graph to obtain a partitioned graph \mathcal{R}^0 .

Equip \mathcal{R}^0 with any linear order. Now consider the ordered 2^n -partitioned graph \mathcal{L}^1 which has just one vertex, and whose first part contains this single vertex. Apply Theorem 11 in order to obtain an ordered partitioned graph \mathcal{R}^1 such that $\mathcal{R}^1 \rightarrow (\mathcal{R}^0)_{k^n}^{\mathcal{L}^1}$.

Next, consider the ordered 2^n -partitioned graph \mathcal{L}^2 which has just one vertex, and whose second part contains this single vertex. Apply Theorem 11 in order to obtain an ordered partitioned graph \mathcal{R}^2 such that $\mathcal{R}^2 \rightarrow (\mathcal{R}^1)_{k^n}^{\mathcal{L}^2}$.

Repeat this procedure with the ordered 2^n -partitioned graphs $\mathcal{L}^3, \dots, \mathcal{L}^{2^n}$; \mathcal{L}^i has its single vertex in its i -th part. We end up with an ordered partitioned graph \mathcal{R}^{2^n} . We now forget its order and denote the resulting structure by $\mathcal{T} = (T; C, T_1, \dots, T_{2^n})$.

\mathcal{T} has the following property: Whenever we color its vertices with k^n colors, then we find a copy of \mathcal{R}^0 in \mathcal{T} such that the coloring is constant on each part of this copy. Hence, it has the property that if we color its two-element subsets and its vertices with k and k^n colors, respectively, then we find in it a copy of \mathcal{R} on which the first coloring is canonical, and such that the color of the vertices depends only on the part the vertex lies in.

Now consider the structure $\mathcal{S} := (T \cup \{u_1, \dots, u_n\}; B, \{u_1\}, \dots, \{u_n\}, T_1, \dots, T_{2^n})$, where B consists of the edges of \mathcal{T} , plus edges connecting the u_i with the vertices of some parts T_i , depending on whether u_i was in \mathcal{U} connected to the vertices in U_i or not. Clearly, \mathcal{S} is the partitioned graph of the n -constant graph $\mathcal{Q} := (T \cup \{u_1, \dots, u_n\}; B, u_1, \dots, u_n)$. We claim that \mathcal{Q} has the property we want to prove. Assume that we color the two-element subsets of $T \cup \{u_1, \dots, u_n\}$ with k colors. We must find a copy of \mathcal{U} in \mathcal{Q} on which the coloring is canonical. Divide the coloring into two colorings, namely the coloring restricted to two-element subsets of T , and the coloring of two-element subsets which contain at least one element u_i outside T . The color of the sets $\{u_i, u_j\}$ completely outside T is irrelevant for what we want to prove, so forget about these.

Now the coloring of those sets which have exactly one element outside T can be encoded in a coloring of the vertices of T : Each vertex is given one of k^n colors, depending on the colors of its edges leading to u_1, \dots, u_n . So we have encoded the original coloring into a coloring of two-element subsets of T and a coloring of the vertices of T . With our observation above, this proves the lemma. \square

6. FINDING STRUCTURE IN MAPPINGS ON THE RANDOM GRAPH

In this section we show how to use the Ramsey-theoretic results from the last section in our context. To warm up, we prove a simple observation (Proposition 22) applying Corollary 12. The proposition states that any mapping on the random graph behaves quite simple on arbitrarily large finite subgraphs.

Definition 20. Let $e, f : V \rightarrow V$. We say that e behaves as f on $F \subseteq V$ iff there is an automorphism α of G such that $f(x) = \alpha(e(x))$ for all $x \in F$. We say that e interpolates f modulo automorphisms iff for every finite $F \subseteq V$ there is an automorphism β of G such that $e(\beta(x))$ behaves as f on F ; so this is the case iff there exist automorphisms α, β such that $\alpha(e(\beta(x))) = f(x)$ for all $x \in F$.

Note that if e interpolates f modulo automorphisms, then it also generates f . We now want to make precise what it means that arbitrarily large structures have a certain property.

Definition 21. Let τ be any signature and let $\mathcal{C}(\tau)$ be a class of finite τ -structures closed under substructures and with the property that for any two structures in $\mathcal{C}(\tau)$ there exists a structure in $\mathcal{C}(\tau)$ containing both structures. We order $\mathcal{C}(\tau)$ by the embedding relation \subseteq . Let $P(w)$ be any property. We say that P holds for arbitrarily large elements of $\mathcal{C}(\tau)$ iff for any $\mathcal{F} \in \mathcal{C}(\tau)$ there exists $\mathcal{H} \in \mathcal{C}(\tau)$ such that $\mathcal{F} \subseteq \mathcal{H}$ and $P(\mathcal{H})$ holds. We say that P holds for all sufficiently large elements of $\mathcal{C}(\tau)$ iff there is an element \mathcal{F} of $\mathcal{C}(\tau)$ such that P holds for \mathcal{H} whenever \mathcal{F} embeds into \mathcal{H} .

Our properties $P(w)$ will be such that if $P(\mathcal{H})$ holds, then P also holds for all substructures of \mathcal{H} . The definition then says that P holds for arbitrarily large elements of $\mathcal{C}(\tau)$ iff for any $\mathcal{F} \in \mathcal{C}(\tau)$ there is $\mathcal{F}' \in \mathcal{C}(\tau)$ isomorphic to \mathcal{F} such that $P(\mathcal{F}')$ holds.

Observe also that if arbitrarily large structures in $\mathcal{C}(\tau)$ have one of finitely many properties, then one property holds for arbitrarily large elements of $\mathcal{C}(\tau)$.

Proposition 22. *Let $e : V \rightarrow V$ be a mapping on the random graph. Then e interpolates either the identity, e_E , e_N , a constant function, or $-$ modulo automorphisms.*

Proof. We show that arbitrarily large finite subgraphs of G have the property that e behaves on them like one of the operations of the proposition. Since there are finitely many operations to choose from, e then behaves like one fixed operation p from the list on arbitrarily large finite subgraphs of the random graph. By the homogeneity of the random graph, we can freely move finite graphs around by automorphisms, proving that e interpolates p .

So let \mathcal{F} be any finite graph; we have to find a copy \mathcal{F}' of \mathcal{F} in G such that e behaves like one of the mentioned operations on this copy.

We color all pairs $\{x, y\}$ of distinct vertices of G

- by 1 if $e(x) = e(y)$,
- by 2 if $E(e(x), e(y))$,
- by 3 if $N(e(x), e(y))$.

By Corollary 12 there exists a copy \mathcal{F}' of \mathcal{F} in G such that all edges and all non-edges of \mathcal{F}' have the same color χ_E and χ_N , respectively. If $(\chi_E, \chi_N) = (1, 1)$, then e behaves like the constant function on \mathcal{F}' . If $(\chi_E, \chi_N) = (2, 3)$, then it behaves like the identity, and if $(\chi_E, \chi_N) = (3, 2)$, then e behaves like $-$. If $(\chi_E, \chi_N) = (2, 2)$ or $(\chi_E, \chi_N) = (3, 3)$, then e behaves like e_E or e_N , respectively. Finally, it is easy to see that $(\chi_E, \chi_N) = (1, q)$ or $(\chi_E, \chi_N) = (q, 1)$, where $q \in \{2, 3\}$, is impossible if \mathcal{F} contains the two three-element graphs with one and two edges, respectively. \square

Definition 23. Let $\mathcal{U} = (U; F)$ be a graph, and let $f : U \rightarrow U$. Let S_1, S_2 be disjoint subsets of U . We say that f is *canonical on S_1* iff it behaves the same way on all edges and on all non-edges, respectively: This is to say that if f collapses one edge in S_1 , then it collapses all edges; if it makes an edge a non-edge, then it does so for all edges; etc. Similarly, we say that f is *canonical between S_1 and S_2* iff the same holds for all edges and non-edges between S_1 and S_2 .

We will often view \mathcal{U} as a subgraph of the random graph, and f will be injective. In this situation, f is canonical on S_1 and between S_1, S_2 iff it behaves like the identity, $-$, e_E , or e_N on S_1 and between S_1, S_2 , respectively. Observe that what we really proved in Proposition 22 is that any $e : V \rightarrow V$ is canonical on arbitrarily large subgraphs of the random graph.

Definition 24. Let $\mathcal{U} = (U; F, U_1, \dots, U_n)$ be a partitioned graph, and let $f : U \rightarrow U$. We say that f is *canonical on \mathcal{U}* iff it is canonical on all U_i and between all distinct U_i, U_j . If $\mathcal{U} = (U; F, u_1, \dots, u_n)$ is an n -constant graph, and $f : U \rightarrow U$, then f is *canonical on \mathcal{U}* iff it is canonical on the corresponding $n + 2^n$ -partitioned graph.

Definition 25. We call a countable structure \aleph_0 -*universal* iff it embeds all finite structures of the same signature.

Lemma 26 (The n -partite graph interpolation lemma). *Let $\mathcal{U} = (U; C, U_1, \dots, U_n)$ be an \aleph_0 -universal partitioned graph, and let $f : U \rightarrow U$. Then every finite partitioned graph has a copy in \mathcal{U} on which f is canonical.*

Proof. This is immediate from the n -partitioned graph Ramsey lemma (Lemma 16): Just like in the proof of Proposition 22, we color the edges and non-edges of \mathcal{U} according to what f does to them. \square

Lemma 27 (The n -constant graph interpolation lemma). *Let $\mathcal{U} = (U; C, u_1, \dots, u_n)$ be an \aleph_0 -universal n -constant graph, and let $f : U \rightarrow U$. Then every finite n -constant graph has a copy in \mathcal{U} on which f is canonical.*

Proof. This is immediate from the n -constant graph Ramsey lemma (Lemma 19). \square

7. PROOF OF THE MAIN THEOREM

We will apply Lemma 27 to prove

Lemma 28. *Let $e : V \rightarrow V$ be so that it preserves N but not E . Then e generates e_N .*

Proof. We prove that for every finite subset F of V , e produces an operation which behaves like e_N on F . We first claim that there are adjacent vertices $a, b \in V$ such that $N(e(a), e(b))$. Since e does not preserve E , there exist u, v with $(u, v) \in E$ such that $(e(u), e(v)) \notin E$. If $N(e(u), e(v))$, then we are done. If $e(u) = e(v)$, then choose w such that $E(w, u)$ and $N(w, v)$. We have $(e(w), e(u)) = (e(w), e(v)) \in N$, so u, w prove the claim.

Now, $\mathcal{U} := (V; E, a, b)$ is an \aleph_0 -universal 2-constant graph. Therefore, by Lemma 27, e is canonical on arbitrarily large substructures of \mathcal{U} . Since e preserves N , it is easy to see that if e is canonical on a 2-constant graph which is large enough, then e must be injective. (For example, if e is canonical on a graph which contains the three-element graph with two edges, then e cannot collapse any edges of that graph.) Hence, e is canonical and injective on arbitrarily large 2-constant subgraphs of \mathcal{U} . Since e preserves N , we have that for arbitrarily large substructures of \mathcal{U} , it behaves like the identity or like e_N on and between the parts of these structures; in particular, it does not add any edges. Hence, for any finite 2-constant graph, we can delete the edge between the two constants without adding any other edges. But that means that starting from any finite graph, we can delete all edges by repeating this process, choosing any edge we want to get rid of in each step. This proves the lemma. \square

The following is just the dual statement.

Corollary 29. *Let $e : V \rightarrow V$ be so that it preserves E but not N . Then e generates e_E .*

Lemma 30. *Let $e : V \rightarrow V$ be so that it preserves neither E nor N . If e is not injective, then e generates a constant operation.*

Proof. We must show that for any finite subset F of V , e generates an operation which is constant on F . Using the fact that e preserves neither E nor N , it is easy to check that e generates operations g, h which collapse an edge and a non-edge, respectively.

Having this, one proceeds inductively to collapse all the vertices of F , shifting F around with automorphisms accordingly and applying g and h . After at most $|F|$ steps, the whole of F is collapsed to a single vertex. \square

The following theorem states that there exist five functions on the random graph which are minimal in the sense that every function which is not an automorphism of G generates one of these five functions.

Theorem 31. *Let Γ be first-order definable in the random graph. Then one of the following cases applies.*

- (1) Γ has a constant endomorphism.
- (2) Γ has e_E as an endomorphism.
- (3) Γ has e_N as an endomorphism.
- (4) Γ has $-$ as an automorphism.
- (5) Γ has sw as an automorphism.
- (6) All endomorphisms of Γ are locally generated by the automorphisms of the random graph.

Proof. If Γ has an endomorphism e which preserves E but not N or N but not E , then we can refer to Lemma 28 and Corollary 29. If all of its endomorphisms preserve both N and E , then they are all generated by the automorphisms of G . We thus assume henceforth that Γ has an endomorphism e which violates both E and N .

If e is not injective, then it generates a constant operation, by Lemma 30. So suppose that e is injective. Fix distinct x, y such that $E(x, y)$ and $N(e(x), e(y))$.

By Proposition 22, e is canonical on arbitrarily large finite subgraphs of G . If e interpolates $-$, e_E , or e_N modulo automorphisms, then we are done. So assume this is not the case, i.e., there is a finite graph \mathcal{F}_0 with the property that on all copies of \mathcal{F}_0 in G , e does not behave like any of these operations. Observe that e then behaves like the identity on arbitrarily large subgraphs of G . Moreover, this assumption implies that if only a finite subgraph \mathcal{F} of G is sufficiently large (i.e., if it embeds \mathcal{F}_0), and e is canonical on \mathcal{F} , then e behaves like the identity on \mathcal{F} .

We now make a series of observations which rule out bad behavior of e between subsets of the random graph, and which follow from our assumptions of the preceding paragraph; the easily verifiable details are left to the reader.

- If e behaves like $-$ between the parts of arbitrarily large finite 2-partitioned subgraphs of G , then it generates sw .
- If e behaves like e_N between the parts of arbitrarily large finite 2-partitioned subgraphs of G , then it generates e_N .
- If e behaves like e_E between the parts of arbitrarily large finite 2-partitioned subgraphs of G , then it generates e_E .

We assume therefore that for sufficiently large finite 2-partitioned subgraphs of G , if e is canonical on such a graph, then e behaves like the identity on and between the parts.

Now observe that $\mathcal{Q} := (V; E, x, y)$ is an \aleph_0 -universal 2-constant graph. Let $\mathcal{F} = (F; D, f_1, f_2)$ be any finite 2-constant graph. By the n -constant interpolation lemma (Lemma 27), there is a copy \mathcal{F}' of \mathcal{F} in \mathcal{Q} on which e is canonical. By our assumption above, if only \mathcal{F} is large enough, then being canonical on a proper part F'_i of the 6-partitioned graph $\tilde{\mathcal{F}}' = (F'; E, \{x\}, \{y\}, F'_1, \dots, F'_4)$ corresponding to \mathcal{F}' means behaving like the identity thereon, and being canonical between proper parts means behaving like the identity between these parts. Therefore, all 2-constant graphs \mathcal{F} have a copy $\mathcal{F}' = (F'; E, x, y)$ in \mathcal{Q} such that e behaves like the identity on and between all of the parts F'_i, F'_j of the corresponding partitioned graph $\tilde{\mathcal{F}}' = (F'; E, \{x\}, \{y\}, F'_1, \dots, F'_4)$.

Of a two-constant graph \mathcal{F} , consider the reduct $\mathcal{H} = (F; D, f_1)$. This reduct has a copy \mathcal{H}' in $\mathcal{Q}^x = (V; E, x)$ on which e is canonical. The corresponding partitioned graph has two parts H'_1, H'_2 , and x is connected to, say, all vertices in H'_1 and to none in H'_2 . Since e is canonical on \mathcal{H}' , either all edges leading to H'_1 are kept or deleted. Similarly with the non-edges between x and H'_2 . If all edges are deleted and all non-edges kept for arbitrarily large \mathcal{H} , then e generates e_N . If all edges are deleted and all non-edges edged for arbitrarily large \mathcal{H} , then e interpolates sw modulo automorphisms. If all edges are kept and all non-edges edged for arbitrarily large \mathcal{H} , then e generates e_E . So we assume that if only \mathcal{H} is large enough, then all edges and non-edges are kept by e on those copies of \mathcal{H} on which e is canonical.

We use the same argument with the reduct $(F; D, f_2)$ and $\mathcal{Q}^y = (V; E, y)$, and arrive at the conclusion that if the two-constant graph \mathcal{F} is large enough, then on every copy of \mathcal{F} in \mathcal{Q} which e is canonical on, the edges and non-edges leading from x and y to the other vertices of the copy are kept.

Combining this with what we have established before, we conclude that if only \mathcal{F} is large enough, and \mathcal{F}' is a copy of \mathcal{F} in \mathcal{Q} which e is canonical on, then e behaves like the identity on \mathcal{F}' except between x and y , where it deletes the edge. Hence, for any finite \mathcal{F} we can find a copy in \mathcal{Q} on which e behaves that way. But this implies that starting from any finite graph $\mathcal{S} := (F; D)$, we can pick any edge in \mathcal{S} , say between vertices f_1, f_2 , and then find a copy of $\mathcal{F} := (F; D, f_1, f_2)$ in \mathcal{Q} such that e deletes exactly that edge from the copy without changing the rest. Hence, by shifting finite graphs around with automorphisms, we can delete a single edge from an arbitrary finite subgraph of G without changing the rest of the graph. Applying this successively, we can remove all edges from arbitrary finite graphs, proving that e generates e_N . \square

Proving Theorem 1 now amounts to showing that if cases (1),(2),(3), and (6) of Theorem 31 do not apply for a structure Γ , and hence if (4) or (5) of that theorem hold, then its endomorphisms are generated by its automorphisms. This will be accomplished in the three propositions to come.

Proposition 32. *Let Γ be first-order definable in the random graph, and suppose Γ is preserved by $-$ but not by e_N, e_E , or a constant operation. Then the endomorphisms of Γ are locally generated by $\{-\} \cup \text{Aut}(G)$, or Γ is preserved by sw .*

Proof. Suppose the endomorphisms of Γ are not generated by $\{-\} \cup \text{Aut}(G)$. Then, by Proposition 6, there is a relation R invariant under $\{-\} \cup \text{Aut}(G)$ and an endomorphism e of Γ which violates R ; that is, there exists a tuple $a := (a_1, \dots, a_n) \in R$ such that $e(a) = (e(a_1), \dots, e(a_n)) \notin R$.

Since R is definable in the random graph, e violates either an edge or a non-edge. Hence, as in the proof of Theorem 31, the assumption that e does not generate e_N, e_E , or a constant operation implies that e is injective.

Let $\mathcal{F} = (F; D, f_1, \dots, f_n)$ be any finite n -constant graph. By the n -constant interpolation lemma (Lemma 27), there is a copy \mathcal{F}' of \mathcal{F} in the \aleph_0 -universal n -constant graph $\mathcal{Q} := (V; E, a_1, \dots, a_n)$ such that e is canonical on this copy.

We now make a series of observations on the behavior of e on and between subsets of V where it is canonical.

- Since by assumption, e does not interpolate e_E, e_N , or a constant operation modulo automorphisms, it behaves like $-$ or the identity on sufficiently large finite subgraphs of G where it is canonical.
- Suppose that for arbitrarily large finite 2-partitioned subgraphs of G , e behaves like the identity on the parts and like $-$ between the parts. Then e generates sw .
- Suppose that for arbitrarily large finite 2-partitioned subgraphs of G , e behaves like the identity on the parts and like e_N (like e_E) between the parts. Then e generates e_N (e_E).
- Suppose that for arbitrarily large finite 2-partitioned subgraphs of G , e behaves like $-$ on the parts and like the identity $/ e_N / e_E$ between the parts. Then e and $-$ together generate $sw / e_E / e_N$. This is because we can apply the preceding two observations to $-e$.
- Suppose that for arbitrarily large finite 2-partitioned subgraphs of G which e is canonical on, e behaves like $-$ on one part and like the identity on the other part. Then e and $-$ together generate e_N .

To see the last assertion for the case where e behaves like the identity between the parts, select an edge within one of the parts that is mapped to a non-edge. For arbitrary finite $A \subseteq V$ we can now use the operation e to get rid of one edge in the graph induced by A in G and preserve all other edges, and so eventually generate an operation that behaves like e_N on A . For the case where e behaves like $-$ between the parts, we can apply the same argument to $-e$. If e behaves like e_N between the parts, then we can all the more delete edges. If it behaves like e_E between the parts, then $-e$ behaves like e_N and we are back in the preceding case.

Summarizing our observations, we can assume that for an arbitrary finite n -constant graph \mathcal{F} there is a copy of \mathcal{F} in \mathcal{Q} such that e behaves like the identity on and between all proper parts F'_i, F'_j of the corresponding partitioned graph, or like $-$ on and between all of its parts. If only the second case holds for arbitrarily large n -constant graphs \mathcal{F} , then we simply proceed our argument with $-e$ instead of e . We can do that since also $-e(a) \notin R$: For otherwise, picking an automorphism α of G such that $\alpha(-(-x)) = x$ for all $x \in V$, we would have $\alpha(-(-e(a))) = e(a) \in R$, contrary to our choice of a . Thus we assume that for arbitrary finite n -constant graphs \mathcal{F} there is a copy of \mathcal{F} in \mathcal{Q} such that e behaves like the identity on and between all proper parts of that copy.

As in the proof of Theorem 31, we may assume that if a copy $\mathcal{F}' = (F'; E, a_1, \dots, a_n)$ of \mathcal{F} in \mathcal{Q} is large enough and e is canonical on \mathcal{F}' and behaves like the identity on and between all proper parts F'_i, F'_j of the corresponding n -partitioned graph $\hat{\mathcal{F}}'$, then it leaves the edges and non-edges between the a_i and the vertices in $F' \setminus \{a_1, \dots, a_n\}$ unaltered. It follows that for arbitrary finite n -constant graphs \mathcal{F} there is a copy of \mathcal{F} in \mathcal{Q} such that the only edges or non-edges changed by e on this copy are those between the a_i .

Finally, note that since R is definable in the random graph and $e(a) \notin R$, e destroys at least one edge or one non-edge on $\{a_1, \dots, a_n\}$. Without loss of generality, say that a_1, a_2 are adjacent but their values under e are not. We have shown that for arbitrarily large 2-constant graphs \mathcal{H} , there is a copy of \mathcal{H} in $(V; E, a_1, a_2)$ such that e behaves like the identity on this copy, except for the edge between a_1 and a_2 , which is destroyed. This clearly implies that e generates e_N . \square

Proposition 33. *Let Γ be first-order definable in the random graph, and suppose Γ is preserved by sw but not by e_N, e_E , or a constant operation. Then the endomorphisms of Γ are locally generated by $\{sw\} \cup \text{Aut}(G)$, or Γ is preserved by $-$.*

Proof. The proof is very similar to the proof of the preceding proposition. This time we know that unless the endomorphisms are locally generated by $\{sw\} \cup \text{Aut}(G)$, there exists an endomorphism e that violates a relation R which is preserved by $\{sw\} \cup \text{Aut}(G)$. Fix a tuple a as before.

As in the preceding proof, we may assume that e is injective. If e interpolates $-$ modulo automorphisms, we are done. Suppose therefore that if e is canonical on a finite partitioned graph large enough, then it must behave like the identity on its parts.

If e behaves like e_N (e_E) between the parts of arbitrarily large finite 2-partitioned subgraphs of G , then it generates e_N (e_E). Thus we may assume that it behaves like the identity or $-$ between such parts.

Suppose that for arbitrarily large finite 3-partitioned subgraphs $\mathcal{F} = (F; E, F_1, F_2, F_3)$ of G which e is canonical on, e behaves like the $-$ between exactly two of the parts, say between F_1, F_2 , and like the identity between F_2, F_3 and F_1, F_3 . Then e is easily seen to generate both e_N and e_E . Indeed, if we want to delete¹ any edge from a finite graph, then we can view the vertices of the edge as two parts of a 3-partitioned graph, where the third part contains all the other vertices. If e behaves like $-$ between the two vertices whose edge we want to delete, and like the identity on and between the other parts, what happens is exactly that the edge is deleted.

If for arbitrarily large finite 3-partitioned subgraphs \mathcal{F} of G which e is canonical on, e behaves like $-$ between, say, F_1, F_2 and F_1, F_3 , and like the identity between F_2, F_3 , then by applying a suitable switch operation i_A to e we are back in the preceding case. Note here that there is an automorphism α of G such that $i_A(\alpha(i_A(x))) = x$ for all $x \in V$. Therefore, $i_A(e(a)) \notin R$; for otherwise, we would have $i_A(\alpha(i_A(e(a)))) = e(a) \in R$, a contradiction.

The latter argument works also if e behaves like $-$ between all three parts. Summarizing, we may assume that if e is canonical on a finite n -partitioned graph which is large enough, where $n \geq 3$, then it behaves like the identity on and between all of the parts.

As for n -constant graphs which e is canonical on, e might flip edges and non-edges between some parts and the constants. However, this situation can easily be repaired by a single application of sw .

Finally, observe that at least one edge or one non-edge on a_1, \dots, a_n is destroyed, and that we therefore can generate either e_N or e_E . \square

Proposition 34. *Let Γ be first-order definable in the random graph, and suppose Γ is preserved by sw and by $-$, but not by e_N, e_E , or a constant operation. Then the endomorphisms of Γ are locally generated by $\{-, sw\} \cup \text{Aut}(G)$, or Γ is preserved by all permutations.*

Proof. The argument goes as in the preceding two propositions; we leave the details to the reader. \square

Theorem 1 now is a direct consequence of Theorem 31, and Propositions 32, 33, 34: If a reduct Γ of G does not have e_E, e_N , or a constant operation as an endomorphism, and if its endomorphisms are not generated by the automorphisms of G , then Theorem 31 implies that it has either $-$ or sw as an endomorphism. Since $\text{Aut}(\Gamma)$ contains $\text{Aut}(G)$, once Γ has $-$ or sw as an endomorphism, it also has its inverse as an endomorphism; thus it has $-$ or sw as an automorphism. But then by the preceding three propositions, either $\text{End}(\Gamma)$ is generated by $\text{Aut}(\Gamma)$, or Γ is preserved by all permutations. The latter case, however, is impossible, as this would imply that e_E and e_N are among its endomorphisms, which we excluded already.

¹for the purposes of the proof, we identify ourselves with the personalized endomorphism monoid

Observe also how Thomas' classification of closed permutation groups containing $\text{Aut}(G)$ (Theorem 2) follows from our results: If a group properly contains $\text{Aut}(G)$, then it contains $-$ or sw , by Theorem 31. If it contains $-$ but is not generated by $-$, then it contains sw by Proposition 32. Similarly, if it contains sw but is not generated by sw , then it contains $-$ by Proposition 33. If it contains both $-$ and sw , but is not generated by these operations, then it must already contain all permutations (Proposition 34).

REFERENCES

- [1] F. G. Abramson and L. Harrington. Models without indiscernibles. *Journal of Symbolic Logic*, 43(3):572–600, 1978.
- [2] M. Bodirsky, H. Chen, and M. Pinsker. The reducts of equality up to primitive positive interdefinability. Preprint available from <http://arxiv.org/abs/0810.2270>.
- [3] M. Bodirsky and J. Kára. The complexity of temporal constraint satisfaction problems. In *Proceedings of STOC'08*, pages 29–38, 2008.
- [4] A. Bonato and D. Delić. The monoid of the random graph. *Semigroup Forum*, 61:138–148, 2000.
- [5] A. Bonato, D. Delić, and I. Dolinka. All countable monoids embed into the monoid of the infinite random graph. *Discrete Mathematics*. To appear.
- [6] P. J. Cameron. Transitivity of permutation groups on unordered sets. *Math. Z.*, 148:127–139, 1976.
- [7] P. J. Cameron. The random graph. *Algorithms and Combinatorics*, 14:333–351, 1997.
- [8] P. J. Cameron. The random graph revisited. In *Proceedings of the European Congress of Mathematics*, volume 201, pages 267–274. Birkhäuser, 2001.
- [9] C. C. Chang and H. J. Keisler. *Model theory*. Princeton University Press, 1977.
- [10] D. Delić and I. Dolinka. The endomorphism monoid of the random graph has uncountably many ideals. *Semigroup Forum*, 69:75–79, 2004.
- [11] W. Hodges. *Model theory*. Cambridge University Press, 1993.
- [12] M. Junker and M. Ziegler. The 116 reducts of $(\mathbb{Q}, <, a)$. *Journal of Symbolic Logic*, 73(3):861–884, 2008.
- [13] J. Nešetřil and V. Rödl. Ramsey classes of set systems. *J. Comb. Theory, Ser. A*, 34(2):183–201, 1983.
- [14] J. Nešetřil and V. Rödl. The partite construction and Ramsey set systems. *Discrete Mathematics*, 75(1-3):327–334, 1989.
- [15] J. Nešetřil. Partitions of finite relational and set systems. *J. Comb. Theory, Ser. A*, 22(3):289–312, 1977.
- [16] J. Nešetřil. Ramsey theory. *Handbook of Combinatorics*, pages 1331–1403, 1995.
- [17] Á. Szendrei. *Clones in Universal Algebra*. Seminaire de mathematiques superieures. Les Presses de L'Université de Montreal, 1986.
- [18] S. Thomas. Reducts of the random graph. *Journal of Symbolic Logic*, 56(1):176–181, 1991.

LABORATOIRE D'INFORMATIQUE (LIX), CNRS UMR 7161, ÉCOLE POLYTECHNIQUE, 91128 PALAISEAU, FRANCE
E-mail address: bodirsky@lix.polytechnique.fr
URL: <http://www.lix.polytechnique.fr/~bodirsky/>

LABORATOIRE DE MATHÉMATIQUES NICOLAS ORESME, CNRS UMR 6139, UNIVERSITÉ DE CAEN, 14032 CAEN CEDEX, FRANCE
E-mail address: marula@gmx.at
URL: <http://dmg.tuwien.ac.at/pinsker/>