Article:

Agarwal, L (2016) Reducts of the Generic Digraph. Annals of Pure and Applied Logic, 167 (3). pp. 370-391. ISSN 0168-0072

https://doi.org/10.1016/j.apal.2015.12.006

Reducts of the Generic Digraph

Lovkush Agarwal

School of Mathematics, University of Leeds, Leeds LS2 9JT, United Kingdom

Abstract

The generic digraph (D, E) is the unique countable homogeneous digraph that embeds all finite digraphs. In this paper, we determine the lattice of reducts of (D, E), where a structure $\mathcal{M}$ is a reduct of (D, E) if it has domain D and all its $\emptyset$-definable relations are $\emptyset$-definable relations of (D, E). As (D, E) is $\aleph_0$-categorical, this is equivalent to determining the lattice of closed groups that lie in between $\mathrm{Aut}(D, E)$ and $\mathrm{Sym}(D)$.

Keywords: Reduct, digraph, homogeneous structure, permutation group, closed group, canonical function.

This paper is a part of a large body of work concerning reducts of first-order structures, where $\mathcal{N}$ is said to be a reduct of $\mathcal{M}$ if all $\emptyset$-definable relations in $\mathcal{N}$ are $\emptyset$-definable in $\mathcal{M}$. A common set-up is that one studies the reducts of some given structure $\mathcal{M}$, where two reducts which are interdefinable are considered to be equal. These reducts form a lattice and when the structure is $\aleph_0$-categorical this is equivalent to studying the lattice of closed subgroups lying between $\mathrm{Aut}(\mathcal{M})$ and $\mathrm{Sym}(M)$.

The first results in this area were the classification of the reducts of $(\mathbb{Q}, <)$ ([1]) and of the random graph Γ ([2]). In [3], Thomas conjectured that all homogeneous structures in a finite relational language have only finitely many reducts. This question remains unsolved and continues to provide motivation for study. More recent results include the classification of the reducts of $(\mathbb{Q}, <, 0)$ ([4]) and of the affine and projective spaces over $\mathbb{Q}$ ([5]).

A surprising development in this area is the connection with constraint satisfaction in complexity theory, by Bodirsky and Pinsker. This connection is made via clone theory in universal algebra. In order to analyse certain closed clones they developed a Ramsey-theoretic tool, named 'canonical functions'. With further developments ([6], [7]), canonical functions now provide a powerful tool in studying reducts, for example, they were used to classify the reducts of the generic partial order ([8]) and of the generic ordered graph ([9]).

In this paper, we determine the lattice of reducts of the generic directed graph, which we denote by (D, E). For us, a directed graph (or digraph) means a set of vertices with directed edges between them, where we do *not* allow an edge going in both directions. The generic digraph is the unique countable homogeneous digraph that embeds all finite digraphs. 'Homogeneous' means that every isomorphism $f : A \to B$, where $A, B \subset D$ are finite, can be extended to an automorphism of (D, E).

We outline the structure of the paper. In Section 1, we provide the necessary preliminary definitions and facts about the generic digraph and about reducts. We also comment on some notational conventions that we use. In Section 2, we define the reducts of the generic graph and provide the lattice, $\mathcal{L}$, that these reducts form. The main theorem is that this lattice $\mathcal{L}$ is the lattice of all the reducts of the generic digraph. In Section 3, we describe the reducts in some detail, establishing notation and important lemmas that are used in the rest of the paper. In Section 4, we show that $\mathcal{L}$ is indeed a sublattice of the lattice of reducts. In Section 5, we prove that $\mathcal{L}$ does contain all the reducts of (D, E). The section starts by describing the information that is obtained from the known classifications of the random graph and the random tournament ([10]). We then give the background definitions and results on canonical functions at the start of Section 5.2, and we also carry out the the combinatorial analysis of the canonical functions in this section. Section 5 ends by using the analysis to complete the proof of the main theorem. In Section 6, we provide a summary and some open questions.

We thank the referee for their thorough and detailed comments and corrections.

1. Preliminaries

1.1. Notational Conventions

We sometimes write 'ab' as an abbreviation for (a, b), e.g., we may write "Let ab be an edge of the digraph D". Structures are denoted by $\mathcal{M}, \mathcal{N}$, and their domains are M and N respectively. If $A \subseteq M$,

A^c denotes the complement of A. $\mathrm{Sym}(M)$ is the set of all bijections $M \to M$ and $\mathrm{Aut}(\mathcal{M})$ is the set of all automorphisms of $\mathcal{M}$. Given a formula $\phi(x, y)$, we use $\phi^*(x, y)$ to denote the formula $\phi(y, x)$. $S(\mathcal{M})$ denotes the space of types of the theory of $\mathcal{M}$. If f has domain A and $(a_1, \dots, a_n) \in A^n$, then $f(a_1, \dots, a_n) := (f(a_1), \dots, f(a_n))$. For $\bar{a}, \bar{b} \in M^n$, we say $\bar{a}$ and $\bar{b}$ are isomorphic, and write $\bar{a} \cong \bar{b}$, to mean that the function $a_i \mapsto b_i$ for all i such that $1 \leq i \leq n$ is an isomorphism.

There will be instances where we do not adhere to strictly correct notational usage, however, the meaning will be clear from the context. For example, we may write '$a \in (a_1, \dots, a_n)$' instead of '$a = a_i$ for some i such that $1 \leq i \leq n$'. Another example is that we sometimes use c to represent the singleton set $\{c\}$ containing it. A third example is we may write '$\bar{a} \in A$' instead of '$\bar{a} \in A^n$ for some n'.

1.2. The Generic Digraph

Definition 1.1. (i) A directed graph (V, E) consists of a set V and an irreflexive, antisymmetric relation $E \subseteq V^2$. V represents the set of vertices and E represents the set of directed edges, so if $(a, b) \in E$, we visualise it as an edge going out of a and into b. We abbreviate 'directed graph' by 'digraph'.

(ii) By an empty digraph we mean a digraph whose edge set is empty.

(iii) We say that a structure $\mathcal{M}$ is homogeneous if every isomorphism $f : A \to B$, where A, B are finite substructures of $\mathcal{M}$, can be extended to an automorphism of $\mathcal{M}$.

(iv) The generic digraph, which we denote by (D, E), is the unique (up to isomorphism) countable homogeneous digraph that embeds all finite digraphs.

(v) $N(x, y)$ will denote the non-edge relation of (D, E), so $N(x, y) := \neg E(x, y) \wedge \neg E^*(x, y)$.

The fact that the generic digraph exists and is unique follows from the theory of Fraïssé limits and amalgamation classes, originally described in [11]. Details and proofs can be found in [12].

The following lemma collects several useful properties of the generic digraph.

Lemma 1.2. *(i) $Th((D, E))$ is $\aleph_0$-categorical and has quantifier elimination.*

(ii) Let $\bar{a}, \bar{b} \in D$. If $tp(\bar{a}) = tp(\bar{b})$, then there exists an automorphism mapping $\bar{a}$ to $\bar{b}$.

(iii) The generic digraph (D, E) is the unique, up to isomorphism, countable digraph satisfying the following extension property: for all finite pairwise disjoint subsets $U, V, W \subset D$ there exists $x \in D \backslash (U \cup V \cup W)$ such that $(\forall u \in U) E(x, u)$, $(\forall v \in V) E(v, x)$ and $(\forall w \in W) N(x, w)$.

(iv) All countable digraphs can be embedded into the generic digraph.

(v) Let $A \subseteq D$ and $B = A^c$. Then $(A, E|_A)$ or $(B, E|_B)$ is isomorphic to the generic digraph.

(vi) For all $a, b \in D$ there is $c \in D$ such that ac and cb are edges.

Remark. Due to the importance of the property in (iii), we give it the name 'the extension property'.

Remark. As a result of (ii), there is bijective correspondence between n-types and orbits of n-tuples. Given a type $p(\bar{x})$ one obtains the orbit $\{\bar{x} \in D \;\; \mathrm{tp}(\bar{x}) = p\}$, and given an orbit $A \subset D^n$ one obtains the type $p(\bar{a})$, where $\bar{a} \in A$. In this light, and as has become customary in modern model theory, we sometimes blur the distinction between a type and the set of tuples that realise that type.

Proof. (i) This is an instance of the more general statement that any countable homogeneous structure in a finite relational language is $\aleph_0$-categorical and has quantifier elimination. See [12] for details.

(ii)-(iv) These are standard results and left to the reader. We note that a back and forth argument is used for (iii) and a forth argument is used for (iv).

(v) Suppose for contradiction that both (A, E_A) and (B, E_B) fail the extension property. Let $U_1, V_1, W_1 \subset A$ and $U_2, V_2, W_2 \subset B$ witness this failure. Now let $U = U_1 \cup U_2, V = V_1 \cup V_2$ and $W = W_1 \cup W_2$. Let $x \in D$ be a witness of the extension property for U, V, W in D. But since x has to be in A or in B, this contradicts how U_i, V_i and W_i were initially chosen.

(vi) Follows straightforwardly from homogeneity. $\square$

1.3. Reducts

Let $\mathcal{M}$ be a structure. A relation $P \subseteq M^k$ is $\emptyset$-definable in $\mathcal{M}$ if there exists a formula $\phi(x_1, \ldots, x_k)$ in the language of $\mathcal{M}$ such that $P = \{(x_1, \ldots, x_k) \in M^k : \mathcal{M} \models \phi(x_1, \ldots, x_k)\}$.

Let $\mathcal{M}$ and $\mathcal{N}$ be two structures on the same domain M. We say that $\mathcal{N}$ is a reduct of $\mathcal{M}$ if for all $k \in \mathbb{N}$ and all relations $P \subset M^k$, if P is $\emptyset$-definable in $\mathcal{N}$ then P is $\emptyset$-definable in $\mathcal{M}$. We say $\mathcal{N}$ is a proper reduct of $\mathcal{M}$ if $\mathcal{N}$ is a reduct of $\mathcal{M}$ but $\mathcal{M}$ is not a reduct of $\mathcal{N}$.

In this article we determine all the reducts of the generic digraph, with the caveat that if two structures are reducts of each other (which implies they are $\emptyset$-interdefinable) we regard them as being equal. For the sake of conciseness, we choose not include the phrase "up to $\emptyset$-interdefinability" where we strictly ought to, with the understanding that we always consider two reducts that are first-order $\emptyset$-interdefinable to be equal.

An important fact about the reducts of a fixed structure $\mathcal{M}$ is that they form a lattice, where $\mathcal{N} \leq \mathcal{N}'$ if $\mathcal{N}$ is a reduct of $\mathcal{N}'$. The top element is always the original structure $\mathcal{M}$ and the bottom element is the trivial structure $(M, =)$. The meet (respectively join) of two structures $\mathcal{N}$ and $\mathcal{N}'$ will be the structure whose named relations are precisely the $\emptyset$-definable relations that are definable in both (respectively in at least one of) $\mathcal{N}$ and $\mathcal{N}'$. In addition to determining what the reducts of the generic digraph are, we also determine how they relate in this lattice.

There is a second, closely related notion of a reduct known as a group reduct. We say that $\mathcal{N}$ is a group reduct of $\mathcal{M}$ if $\mathrm{Aut}(\mathcal{N}) \geq \mathrm{Aut}(\mathcal{M})$. The group reducts of a fixed structure $\mathcal{M}$ form a lattice via the usual inclusion operation; the bottom element is $\mathrm{Aut}(\mathcal{M})$ and the top element is always $\mathrm{Sym}(M)$.

As a consequence of the Engeler–Ryll-Nardzewski–Svenonius theorem (see [12]), if $\mathcal{M}$ is $\aleph_0$-categorical then the lattice of reducts is anti-isomorphic to the lattice of group-reducts. In one direction, a reduct $\mathcal{N}$ is mapped to its automorphism group $\mathrm{Aut}(\mathcal{N})$. In the other direction, given a group reduct G one lets $\mathcal{N}$ be the structure whose n-ary relations are the orbits of the action of G on M^n (where for all $g \in G, \bar{x} \in M^n, g \cdot \bar{x} = g(\bar{x})$). In this light, we use the word 'reduct' to refer to either notion, with the meaning being clear from the context.

Furthermore, the group reducts of $\mathcal{M}$ are exactly the closed groups $G \leq \mathrm{Sym}(M)$ that contain $\mathrm{Aut}(\mathcal{M})$. We recall the topology on $\mathrm{Sym}(M)$ by describing what it means for a set to be closed: We say that $g \in \mathrm{Sym}(M)$ is in the closure of $F \subset \mathrm{Sym}(M)$ if for all finite $A \subset M$, there exists $f \in F$ such that $f(a) = g(a)$ for all $a \in A$. Then, F is closed if F is equal to the closure of itself.

From the above discussion, since (D, E) is $\aleph_0$-categorical, the task of determining its reducts is the same as determining its group reducts, which in turn is the same as determining the closed groups G where $\mathrm{Aut}(D, E) \leq G \leq \mathrm{Sym}(D)$.

2. Defining the Reducts

One can define a reduct by adding a function $f \in \mathrm{Sym}(D)$ to $\mathrm{Aut}(D, E)$, then closing under group operations and closing under the topology. Alternatively, one first defines a relation, P say, and then defines the reduct to be the automorphism group of (D, P). In view of this, we establish some notation:

(i) Let G be a topological group (e.g. $\mathrm{Sym}(D)$). For $F \subseteq G$, let $\langle F \rangle$ denote the smallest closed subgroup of G containing F. For brevity, when it is clear we are discussing reducts of (D, E), we may abuse notation and write $\langle F \rangle$ to mean $\langle F \cup \mathrm{Aut}(D, E) \rangle$.

(ii) Let G be a group. For $F \subseteq G$, we let $\mathrm{cl_g}(F)$, the group closure of F, denote the smallest subgroup of G containing F. As above, we may abuse notation where it is clear we are discussing supergroups of $\mathrm{Aut}(D, E)$.

We begin by showing that three particular functions $-$, sw and rot exist. These functions will give us the three reducts $\langle - \rangle$, $\langle sw \rangle$ and $\langle rot \rangle$.

Lemma 2.1. *There exists $f \in Sym(D)$ such that for all $x, y \in D$, $E(f(x), f(y))$ iff $E(y, x)$.*

Remark. For the rest of this article, we fix such a bijection and denote it by $-$.

Proof. The idea is to define a structure (D, E') that is isomorphic to (D, E) so that any isomorphism

$f : (D, E') \to (D, E)$ has the desired property. For this lemma, we let $E'(x, y) = E^*(x, y)$. The fact that $(D, E') \cong (D, E)$ follows straightforwardly from the extension property. $\square$

Lemma 2.2. *Let $a \in D$. Then there exists $f \in Sym(D)$ such that*

$$E(f(x), f(y)) \text{ if and only if } \begin{cases} E(x, y) \text{ and } x, y \neq a, \text{ OR,} \\ \\ E^*(x, y) \text{ and } x = a \vee y = a \end{cases}$$

Remark. For the rest of this article, we fix such a bijection and denote it by sw.

Proof. Use the same strategy as the previous lemma, but define $E'(x, y)$ as follows:

$$E'(x, y) := \begin{cases} E(x, y), & \text{if } x, y \neq a \\ \\ E^*(x, y), & \text{otherwise} \end{cases}$$

$\square$

Lemma 2.3. *Let $a \in D$. Then there exists $f \in Sym(D)$ such that*

$$E(f(x), f(y)) \text{ if and only if } \begin{cases} x, y \neq a \text{ and } E(x, y) \\ \\ x = a \text{ and } N(x, y) \\ \\ y = a \text{ and } E^*(x, y) \end{cases}$$

Remark. For the rest of this article, we fix such a bijection and denote it by rot.

Remark. In words, rot sends edges going out of a, to edges going into a, to non-edges, to edges going out of a.

Proof. Use the same strategy as for $-$ and sw. $\square$

Definition 2.4. (i) We let $\Gamma = (D, E_\Gamma)$, where $E_\Gamma := E(x, y) \vee E^*(x, y)$. Γ is a graph and, as will be proved later, is in fact (isomorphic to) the random graph.

(ii) We let $-_\Gamma \in \mathrm{Sym}(D)$ be a function which interchanges the sets of edges and non-edges in Γ.

(iii) Let $a \in D$. We let $sw_\Gamma \in \mathrm{Sym}(D)$ be a function which interchanges the sets of edges and non-edges adjacent to a, and preserves all other edges and non-edges.

Remarks. (D, E_Γ) is $\emptyset$-interdefinable with (D, N). The existence of $-_\Gamma$ and sw_Γ follows from the same argument as was used for $-$, sw and rot.

We now have all the background definitions necessary to state the main theorem:

Theorem 2.5. *The reducts of (D, E) are given by the following lattice, which we call $\mathcal{L}$:*

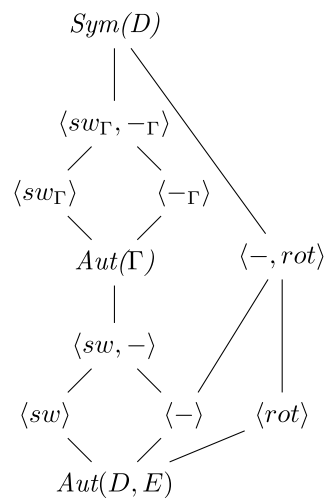

This theorem can be split into two main claims: that $\mathcal{L}$ is a sublattice of the reducts of (D, E) (so for example one needs to show that the meets and joins are correct), and that $\mathcal{L}$ contains all the reducts of (D, E).

3. Understanding the reducts

In this section we establish several useful lemmas that are used throughout the rest of the article. The first few lemmas will provide a concrete description of the three groups $\langle sw \rangle, \langle - \rangle$ and $\langle rot \rangle$. We do this by comparing how two functions behave, via the following definition.

Definition 3.1. Let $f, g : D \to D$ and $A \subset D$. We say f *behaves like* g *on* A if for all finite tuples $\bar{a} \in A$, $f(\bar{a})$ is isomorphic (as a finite digraph) to $g(\bar{a})$. If $A = D$, we simply say f *behaves like* g.

Example. All automorphisms of (D, E) behave like the identity $id : D \to D$. Conversely, all $f \in Sym(D)$ which behave like id are automorphisms.

Useful Remark. If $f : D \to D$ is any function and $g \in \mathrm{Aut}(D, E)$, then $h := g \circ f$ behaves like f. The converse is also true if f, h are bijections: if $h \in \mathrm{Sym}(D)$ behaves like $f \in \mathrm{Sym}(D)$, then there is $g \in \mathrm{Aut}(D, E)$ such that $h = g \circ f$.

We start with the simplest of the three groups, $\langle - \rangle$.

Lemma 3.2. *Let $f \in Sym(D)$. Then $f \in \langle - \rangle \setminus Aut(D, E) \Leftrightarrow f$ behaves like $-$.*

Proof. "$\Leftarrow$". Suppose f behaves like $-$. Then observe that $g := - \circ f$ behaves like id so $g \in \mathrm{Aut}(D, E)$. Hence, $f = -^{-1} \circ g \in \langle - \rangle$.

"$\Rightarrow$". Follows from the observation that all elements of $\mathrm{cl_g}(-) \setminus \mathrm{Aut}(D, E)$ behave like $-$. □

Next we look at $\langle sw \rangle$. For $A \subset D$, we let $sw_A : D \to D$ denote a function that behaves like id on A and A^c, and that switches the direction of all edges between A and A^c. For example, $sw = sw_a$ for some $a \in D$ and $sw_\emptyset$ is just an automorphism. The fact that sw_A exists for all $A \subseteq D$ follows from the fact that all countable digraphs are embeddable in the generic digraph (Lemma 1.2). If possible, we choose sw_A to be a bijection. Note that there are cases where sw_A cannot be a bijection, namely when the image of sw_A is not isomorphic to the generic digraph. For example, let $A = \{x \in D : E(a, x)\}$ where a is some element of D, then $sw_A(a)$ will not have any outward edges in its image.

Lastly, we define a 3-ary relation P_{sw} as follows:

$$\begin{aligned} P_{sw}(x, y, z) :=& (E(x, y) \wedge E(y, z) \wedge E(x, z)) \\ & \vee (E^*(x, y) \wedge E^*(y, z) \wedge E(x, z)) \\ & \vee (E^*(x, y) \wedge E(y, z) \wedge E^*(x, z)) \\ & \vee (E(x, y) \wedge E^*(y, z) \wedge E^*(x, z)) \end{aligned}$$

In words, a function f preserves P_{sw} if f preserves non-edges and if for all tournaments on three vertices f changes the direction of exactly 0 or 2 edges. Note that sw_A preserves P_{sw} for all A.

Lemma 3.3. *(i)* $cl_g(sw) = \{f \in Sym(D) : f$ *behaves like* sw_A*, for some finite* $A \subset D\}$.

(ii) For all proper non-empty $A \subset D$*, if* $sw_A \in Sym(D)$ *then* $\langle sw_A \rangle = \langle sw \rangle$.

(iii) $\langle sw \rangle = \{f \in Sym(D) : f$ *behaves like* sw_A*, for some* $A \subseteq D\} = \{f \in Sym(D) : f$ *preserves* $P_{sw}(x, y, z)\}$

Proof. For all of this proof, let $a \in D$ be the point such that $sw = sw_a$.

(i) RHS $\subseteq$ LHS. Let f behave like sw_A. We prove that $f \in \mathrm{cl_g}(sw)$ by induction on $|A|$. For the base case, let $A = \{a'\}$ and observe that $sw \circ h$ behaves like $sw_{a'}$ where $h \in \mathrm{Aut}(D, E)$ maps a' to a. By the useful remark above, this completes the base case. Now let $n > 1$ and let $A = \{a_1, \ldots, a_n\} \subset D$. Let $A' = \{a_2, \ldots, a_n\}$ and let $a' = sw_{A'}(a_1)$. Then consider the function $sw_{a'} \circ sw_{A'}$. By the inductive hypothesis this function is in $\mathrm{cl_g}(sw)$. Also, this function behaves like sw_A, as required.

LHS $\subseteq$ RHS. Follows from the following observation: If f behaves like sw_A and $a' = f^{-1}(a)$, then $sw \circ f$ behaves like $sw_{A \triangle \{a'\}}$. ('$\triangle$' denotes symmetric difference.)

(ii) Let A be a proper non-empty subset of D such that sw_A is a bijection. The fact that $sw_A \in \langle sw \rangle$ follows straightforwardly from the definitions and (i). To show the converse, it suffices to prove that for all $a_1, \ldots a_n \in D$ there exist $b_1, \ldots, b_n \in D$ such that $\bar{a} \cong \bar{b}$ and $A \cap \bar{b} = \{b_1\}$ or $\{b_2, \ldots, b_n\}$. This is immediate by homogeneity if A or A^c is finite, so suppose that A is infinite and coinfinite.

We complete the proof by induction on the length n of the tuple $\bar{a}$. The base case $n = 1$ is trivial so let $(a_1, \ldots, a_{n+1})$ be any tuple of length $n + 1$. By the inductive hypothesis, we can find $(b_1, \ldots, b_n)$ isomorphic to $(a_1, \ldots, a_n)$ where $A \cap \bar{b} = \{b_1\}$ or $\{b_2, \ldots, b_n\}$. Without loss we may assume that $A \cap \bar{b} = \{b_1\}$. If we can find $x \in A^c$ such that $(b_1, \ldots, b_n, x) \cong \bar{a}$, then we are done, so from now on assume that $(b_1, \ldots, b_n, x) \cong \bar{a}$ implies $x \in A$.

Now consider a tuple $(c_1, \ldots, c_{n+1})$ satisfying the following three conditions: c_1 is any element of $A^c \backslash \{b_2, \ldots, b_n\}$, $\bar{c} \cong \bar{a}$, and for each $2 \leq i \leq n + 1$, $(b_1, \ldots, b_n, c_i) \cong \bar{a}$. The first condition can be satisfied as A^c is infinite. The latter two conditions can be satisfied by the extension property. By the third condition, $c_2, \ldots, c_{n+1} \in A$. So $(c_1, \ldots, c_{n+1})$ satisfies all the conditions that we want, completing the induction and hence the proof.

(iii) If f behaves like sw_A for some $A \subseteq D$ then $f \in \langle sw \rangle$, by part (ii). Also, since sw and sw^{-1} preserve P_{sw} then every element of $\langle sw \rangle$ does too. It remains to be shown that if f preserves P_{sw} then f behaves like sw_A for some $A \subseteq D$.

To this end suppose that f preserves P_{sw}. Let $a' \in D$ be any element. Then let $A \subset D$ be the set of vertices v such that there exists u such that $a'u$ is an edge, uv is an edge and f switches the direction of exactly one of these two edges.

We claim that f behaves like sw_A. This amounts to case checking which we leave to the reader. Note that Lemma 1.2 (vi) is implicitly used in this proof. We provide one case as an example.

Case 1: We need to show that f behaves like id on A, so let v_1, v_2 be any edge in A. Then let $u_1, b = u_2$ be the vertices for v_1, v_2, respectively, given by the definition of A; without loss assume that f switches the edges $a'u_1$ and $a'u_2$. Now let d be a sixth vertex which is adjacent to a', a_1, a_2, b_1 and b_2. But now because f preserves P_{sw}, it is possible to see that f switches a_1d if and only if f switches a_2d. But then again because f preserves P_{sw}, this implies that f must not switch the direction of a_1a_2. Thus f fixes the direction of all the edges in A, as required. $\square$

The next reduct we analyse is $\langle rot \rangle$. For what follows, $A, B \subseteq D$ are disjoint and $C := (A \cup B)^c$. For the ordered pair (A, B), an *outward edge* is an edge going from A to B and an *inward edge* is one going from B to A. We say $f : D \to D$ behaves like rot between (A, B), or say 'between A and B', if f maps outward edges to inward edges to non-edges to outward edges. We let $rot_{A,B,C}$ be a function $D \to D$ which behaves like id on A, B and C and behaves like rot between $(A, B), (B, C)$ and (C, A). We often omit C from the subscript as its role is implicitly determined by A and B. If $C = \emptyset$, so that $B = A^c$, we just write rot_A. If possible we choose $rot_{A,B,C}$ to be a bijection.

Simple observations. $rot = rot_a$ for some $a \in D$. $rot_{B,C,A}$ and $rot_{C,A,B}$ both behave like $rot_{A,B,C}$. If f behaves like $rot_{A,B,C}$ and f fixes A, B and C setwise, then f^2 and f^{-1} behave like $rot_{C,B,A}$, and f^3 behaves like id.

Lastly, we define 3-ary relations for rot. These relations correspond to the orbits of $\mathrm{cl_g}(rot)$ acting on (D, E). We describe the orbits diagrammatically:

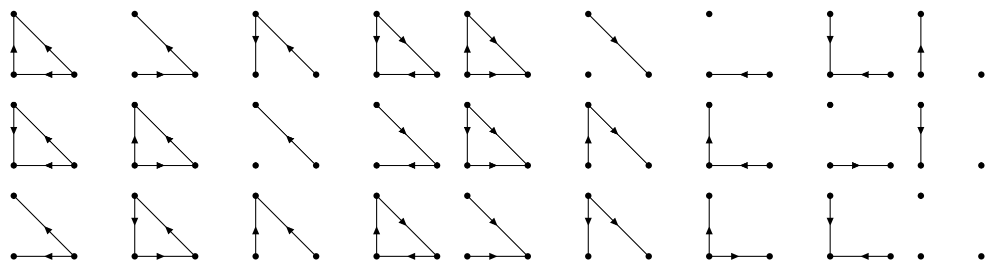

This diagram contains all the possible digraphs on three vertices. Each row of the diagram represents one of the orbits and hence, one of the relations that $\langle rot \rangle$ preserves. Let $P_{rot,1}, P_{rot,2}$ and $P_{rot,3}$ be the relations corresponding to the top, middle and bottom row respectively. One feature worth noting is that given any triple in D, making a change to exactly one (non-)edge of that triple causes a change in the orbit the triple is in.

Lemma 3.4. *(i) $cl_g(rot) = \{f \in Sym(D) : f$ behaves like $rot_{A,B}$ where A, B are finite$\}$.*

(ii) Let A, B be proper disjoint subsets of D such that at least one of A or B is non-empty. If $rot_{A,B} \in Sym(D)$, then $\langle rot_{A,B} \rangle = \langle rot \rangle$.

(iii) $\langle rot \rangle = \{f \in Sym(D) : f$ behaves like $rot_{A,B}$ where A, B are disjoint subsets of $D\} = \{f \in Sym(D) : f$ preserves $P_{rot,i}, i = 1, 2, 3\}$.

Proof. For this proof, let $a \in D$ be the point such that $rot = rot_a$.

(i) RHS $\subseteq$ LHS. It suffices to show that $rot_{A,B} \in \mathrm{cl_g}(rot)$. We start by showing that $rot_{a'} \in \mathrm{cl_g}(rot)$ for all $a' \in D$. This is easy: let $h \in \mathrm{Aut}(D, E)$ map a' to a then consider $rot \circ h$.

For the general case, let $A = \{a_1, \ldots, a_n\}$ and $B = \{b_1, \ldots, b_m\}$. The idea is to rotate twice about each element of A and rotate once about each element of B - we leave the details to the reader.

LHS $\subseteq$ RHS. Any $f \in \mathrm{cl_g}(rot)$ can be written in the form $g_n rot^{\epsilon_n} \ldots g_1 rot^{\epsilon_1} g_0$ where for all $i, \epsilon_i \in \{1, -1\}$. Since rot^{-1} behaves like rot^2, we can assume that $\epsilon_i = 1$ for all i. We prove by induction on n that there exist finite disjoint $A, B \subset D$ such that f behaves like $rot_{A,B}$.

The base case $n = 0$ is trivial, so assume that we know $h = g_{n-1}rot \ldots g_1 rot g_0$ behaves like $rot_{A,B}$ for finite A, B, and we consider $f = g_n rot\, h$. There are three cases depending on $a' := h^{-1}(a)$. If $a' \notin A \cup B$, then f behaves like $rot_{A,B\cup\{a'\}}$. If $a' \in B$, then f behaves like $rot_{A\cup\{a'\},B\setminus\{a'\}}$. Lastly, if $a' \in A$, then f behaves like $rot_{A\setminus\{a'\},B}$. This completes the induction and hence the proof.

(ii) Let $A, B \subseteq D$ be as described in the lemma, and let $C = (A \cup B)^c$. Using (i) it is straightforward to show that $rot_{A,B} \in \langle rot \rangle$. To show the converse, it suffices to show that rot or $rot^{-1} \in \langle rot_{A,B} \rangle$.

If one of A, B or C is empty, then we are done by imitating the corresponding argument for $\langle sw \rangle$. So assume A, B and C are all non-empty. If $(B \cup C, E_{B\cup C})$ is isomorphic to the generic graph, then we can ignore A and again imitate the argument from the switching case to get the result.

Hence, assume that $B \cup C$ is not isomorphic to the generic digraph. This means there exist finite, pairwise disjoint $U, V, W \subset B \cup C$ such that if $x \in D$ satisfies $\phi(x) := (\forall u \in U E(u, x)) \wedge (\forall v \in V E^*(v, x)) \wedge (\forall w \in W N(w, x))$, then $x \in A$.

Suppose that there exists $c \in C \cap (U \cup V \cup W)$. We will show that for all $(d_1, \ldots, d_n) \in D$, there exists $a_2, \ldots, a_n \in A$ such that $(c, a_2, \ldots, a_n) \cong (d_1, \ldots, d_n)$; this is sufficient to show that $rot \in \langle rot_{A,B} \rangle$. So, let $(d_1, \ldots, d_n) \in D$. Then let $(a_2, \ldots, a_n) \in D$ be such that $D \models \phi(a_2), \ldots, \phi(a_n)$ and $(c, a_2, \ldots, a_n) \cong (d_1, \ldots, d_n)$. Such a_i exist by the homogeneity of (D, E). Since $\phi(a_i)$ for all i, $(a_2, \ldots, a_n)$ has to be in A, as required, so $rot \in \langle rot_{A,B} \rangle$.

Now suppose that $C \cap (U \cup W \cup V) = \emptyset$, so there must be $b \in B \cap (U \cup V \cup W)$. By repeating the argument above, we can show that $rot^{-1} \in \langle rot_{A,B} \rangle$, so we are done.

(iii) If f behaves like $rot_{A,B}$ for some disjoint $A, B \subseteq D$ then by part (ii) we know that $f \in \langle rot \rangle$. Also, since rot and rot^{-1} preserve $P_{rot,i}$ then every element of $\langle rot \rangle$ does too. It remains to be shown that if f preserves $P_{rot,i}$ for $i = 1, 2, 3$ then f behaves like $rot_{A,B}$ for some disjoint $A, B \subseteq D$.

We find A and B as follows. Pick any $a \in D$. Let $A = \{a\} \cup \{x \in D : E(a, x) \wedge E(f(a, x))$ or $E^*(a, x) \wedge E^*(f(a, x))$ or $N(a, x) \wedge N(f(a, x))\}$. Let $B = \{x \in D : E(a, x) \wedge E^*(f(a, x))$ or $E^*(a, x) \wedge N(f(a, x))$ or $N(a, x) \wedge E(f(a, x))\}$.

We claim that f behaves like $rot_{A,B}$. This amounts to case checking, which we leave to the reader. We provide one case as an example.

Case 1. We need to show that f behaves like id on A. Suppose not, and let $a_1, a_2 \in A$ witness this fact. Then we have (a, a_1, a_2) such that f only changes what happens between a_1 and a_2, contradicting that f preserves $P_{rot,i}$. $\square$

The descriptions of $\langle -, sw\rangle$ and $\langle -, rot\rangle$ are straightforward:

Lemma 3.5. *(i)* $\langle -, sw\rangle = \{f \in \mathit{Sym}(D) : f = g \textit{ or } - \circ g \textit{ for some } g \in \langle sw\rangle\}$.

(ii) $\langle -, rot\rangle = \{f \in \mathit{Sym}(D) : f = g \textit{ or } - \circ g \textit{ for some } g \in \langle rot\rangle\}$.

Proof. (i) $\langle -, sw\rangle$ preserves the 6-ary relation $P_{sw,w} := P_{sw}(\bar{x}) \leftrightarrow P_{sw}(\bar{y})$. Now let $f \in \langle -, sw\rangle$. Then f must preserve $P_{sw,w}$, which implies that f or $- \circ f$ preserves P_{sw}, so we are done by Lemma 3.3.

(ii) $\langle -, rot\rangle$ preserves the 6-ary relation $P_{rot,w} := (P_{rot,1}(\bar{x}) \wedge P_{rot,1}(\bar{y})) \vee (P_{rot,2}(\bar{x}) \wedge P_{rot,2}(\bar{y}))$. Then continue as in (i). $\square$

The next lemmas will give us conditions on a group G to be equal to $\mathrm{Sym}(D)$ or to contain $\mathrm{Aut}(\Gamma)$.

Lemma 3.6. *Let* $G \leq \mathit{Sym}(D)$ *be a closed supergroup of* $\mathit{Aut}(D, E)$.

(i) If G *is* n*-transitive for all* $n \in \mathbb{N}$*, then* $G = \mathit{Sym}(D)$*. Note that* G *is* n*-transitive if for all pairs of tuples* $\bar{x}, \bar{y} \in D^n$*, there exists* $g \in G$ *such that* $g(\bar{x}) = \bar{y}$*.*

(ii) If G *is* n*-homogeneous for all* $n \in \mathbb{N}$*, then* $G = \mathit{Sym}(D)$*. Note that* G *is* n*-homogeneous if for all subsets* $A, B \subset D$ *of size* n*, there exists* $g \in G$ *such that* $g(A) = B$*.*

(iii) Suppose that whenever $A \subset D$ *is finite and has edges, there exists* $g \in G$ *such that* $g(A)$ *has less edges than* A *has. Then* $G = \mathit{Sym}(D)$*.*

(iv) Suppose that there exists a finite $A \subset D$ *and* $g \in G$ *such that* g *behaves like* id *on* $D \setminus A$*,* g *behaves like* id *between* A *and* $D \setminus A$*, and* g *deletes at least one edge in* A*. Then* $G = \mathit{Sym}(D)$*.*

Proof. (i) Well-known and easy result in permutation group theory.

(ii) We will show that G is n-transitive, so let $\bar{a}, \bar{b} \in D$ be tuples of length n. Then by n-homogeneity we can map $\bar{a}$ to the empty digraph, and again by n-homogeneity we can map the empty digraph to $\bar{b}$.

(iii) By repeatedly using the assumptions given in the lemma, we can map any finite set A to the empty digraph. This implies that G is n-homogeneous so we are done by (ii).

(iv) Let A and g be as in the lemma. We will show that for all finite $B \subset D$, if B contains edges then there is $f \in G$ such that $f(B)$ has less edges than B - this suffices by (iii). So let $B \subset D$ be finite. Let bb' be an edge in B, and let $aa' \in A$ be an edge that is deleted by g. Let h be an automorphism mapping bb' to aa' and all other elements of B to elements of $D \backslash A$. Then $gh \in G$ and $gh(B)$ contains less edges than in B, as required. ☐

We know look at the reduct $\mathrm{Aut}(\Gamma)$.

Lemma 3.7. Γ *is isomorphic to the random graph.*

Proof. It suffices to show that $\Gamma := (D, E_\Gamma)$ satisfies the extension property of the random graph. This follows immediately from the extension property of the digraph (Lemma 1.2). ☐

Terminology. Let $a_1, \ldots, a_n, b_1, \ldots, b_n \in D$. We say $\bar{a}$ and $\bar{b}$ are isomorphic as graphs if $N(a_i, a_j) \leftrightarrow N(b_i, b_j)$ for all i, j.

Lemma 3.8. *Let $G \leq Sym(D)$ be a closed supergroup of $Aut(D, E)$.*

(i) Suppose that whenever $\bar{a}$ and $\bar{b}$ are isomorphic as graphs, there exists $g \in G$ such that $g(\bar{a}) = \bar{b}$. Then $G \geq Aut(\Gamma)$.

(ii) Suppose that for all $A = \{a_1, \ldots, a_n\} \subset D$, there exists $g \in G$ such that for all edges $a_i a_j$ in A, $E(g(a_i), g(a_j))$ iff $i < j$. (Intuitively, such a g is switching the edges so they all point in the same direction.) Then, $G \geq Aut(\Gamma)$.

(iii) Suppose that for all finite $A \subset D$ and all edges $aa' \in A$ there is $g \in G$ such that g changes the direction of aa' and behaves like id on all other edges and non-edges of A. Then $G \geq Aut(\Gamma)$.

(iv) Suppose there is a finite $A \subset D$ and a $g \in G$ such that g behaves like id on $D \backslash A$, g behaves like id between A and $D \backslash A$, and g switches the direction of some edge in A. Then, $G \geq Aut(\Gamma)$.

Proof. (i) Follows immediately from the definitions.

(ii) Let $\bar{a_1}, \bar{a_2} \in D$ be isomorphic as graphs and let $g_1, g_2 \in G$ be the functions as described in the lemma for these tuples. Then $g_1(\bar{a}_1)$ is isomorphic to $g_2(\bar{a}_2)$ as digraphs. Thus we can get from $\bar{a}_1$ to $\bar{a}_2$ using functions in G, so we are done by (i).

(iii) Let $\bar{a}, \bar{b} \in D$ be isomorphic as graphs. Then by repeatedly using the condition in the lemma, we can switch the appropriate edges in $\bar{a}$ to end up with $\bar{b}$. Thus we are done by (i).

(iv) Let A and g be as stated in the lemma, and let $aa' \in A$ be an edge whose direction is switched by g. Now let $B \in D$ be finite and let bb' be any edge in $\bar{b}$. By homogeneity there is an automorphism h mapping bb' to aa' and all other elements of B to $D \backslash A$. Then applying $g \circ h$ to B switches the edge bb' and behaves like id everywhere else. Thus we are done by (iii). □

4. $\mathcal{L}$ is a sublattice of the reducts of (D, E)

Please note a convention that we will use for the remainder of the article. There will be proofs where we map some digraph A to a digraph B, that involve composing a sequence of functions $f_1, f_2, \ldots$, where the definition of each one depends on those defined earlier. For example, we may have defined f_1 and f_2, and f_3 is going to be a switching function. The convention is that we write 'Let f_3 be $sw_{A'}$' (where A' will be a particular subset of A), instead of the strictly correct 'Let f be $sw_{f_2 f_1(A')}$'.

There are two main benefits. First, the proofs will be easier to follow and will better match the intuition behind the argument. Second, we can avoid naming the functions altogether. We can use phrases like 'First switch about the subset A_1, then apply rot about the point a', whereas without the convention we would have to say '...then apply rot about the point which is the current image of a'.

Lemma 4.1. *(i)* $\langle - \rangle, \langle sw \rangle$ *and* $\langle rot \rangle$ *are proper reducts of* $Aut(D, E)$*.*

(ii) $\langle - \rangle, \langle sw \rangle$ *and* $\langle rot \rangle$ *are not reducts of each other.*

(iii) $\langle -, sw \rangle$ *is a proper reduct of* $\langle - \rangle$ *and* $\langle sw \rangle$*, and is not equal to* $Sym(D)$*.*

(iv) Γ *is a proper reduct of* $\langle -, sw \rangle$

(v) $\langle -, rot \rangle$ *is a proper reduct of* $\langle - \rangle$ *and* $\langle rot \rangle$*, and is not equal to* $Sym(D)$*.*

(vi) The join of $\langle rot \rangle$ *and* $\langle sw \rangle$ *is* $Sym(D)$*.*

(vii) The meet of $\langle sw \rangle$ *and* $\langle - \rangle$ *is* $Aut(D)$*.*

(viii) The meet of $\langle rot \rangle$ *and* $\langle sw_\Gamma, -_\Gamma \rangle$ *is* $Aut(D)$*.*

(ix) The meet of $\langle -, rot \rangle$ *and* $\langle sw_\Gamma, -_\Gamma \rangle$ *is* $\langle - \rangle$*.*

Proof. (i) This is immediate from the definitions.

(ii) We need to identify for each reduct a relation that it preserves but which the other two do not preserve. For $\langle - \rangle$ the relation is E_w, for $\langle sw \rangle$ the relation is P_{sw} and for $\langle rot \rangle$ we use $P_{rot,1}$.

(iii) By (ii), $\langle -, sw \rangle$ is a proper reduct of $\langle - \rangle$ and $\langle sw \rangle$. It preserves $P_{sw,w}$, so $\langle -, sw \rangle \neq \mathrm{Sym}(D)$.

(iv) Both $-$ and sw preserve $N(x, y)$, so $\langle -, sw \rangle \subseteq \mathrm{Aut}(D, N) = \Gamma$. Γ is a proper reduct because $\langle -, sw \rangle$ preserves $P_{sw,w}$ but Γ does not.

(v) By (ii), $\langle -, rot \rangle$ is a proper reduct of $\langle - \rangle$ and $\langle rot \rangle$. It preserves $P_{rot,w}$, so $\langle -, rot \rangle \neq \mathrm{Sym}(D)$.

(vi) By Lemma 3.6 (iii), it suffices to show that for all finite $A \subset D$ we can delete edges from A, if it has any, by applying functions from $\langle sw, rot \rangle$. So let $A \subset D$ be finite, let $a \in A$ be a point adjacent to at least one edge, and let $A_1 = \{a' \in A : E(a, a')\}, A_2 = \{a' \in A : E(a', a)\}$ and $A_3 = \{a' \in A : N(a, a')\}$.

First, switch about the subset A_1, so all the edges adjacent to a are now inward edges. Then apply rot_a^2: the edges between a and $A_1 \cup A_2$ become outward edges, and the non-edges between a and A_3 become inward edges. Then apply $sw_{A_1 \cup A_2}$: between a and $A \backslash \{a\}$ we now only have inward edges.

Applying rot_a for the last time results in all these edges becoming non-edges. The number of edges within $A\backslash\{a\}$ are unchanged, so we have reduced the number of edges in A, as required.

(vii) Let $f \in \langle - \rangle \cap \langle sw \rangle$. By Lemma 3.3, f behaves like sw_A for some $A \subseteq D$, so there exists an edge whose direction f does not switch, so f does not behave like $-$. Hence, by Lemma 3.2, $f \in \mathrm{Aut}(D, E)$.

(viii) For $A \subset D$, we say $f : D \to D$ behaves like $sw_{\Gamma,A}$ if f preserves edges and non-edges on A and A^c and if f swaps all edges and non-edges between A and A^c. Then by folklore (or by [2]), $\langle sw_\Gamma \rangle = \{f \in \mathrm{Sym}(D) : f$ behaves like $sw_{\Gamma,A}$ for some $A \subseteq D\}$, and $\langle -_\Gamma, sw_\Gamma \rangle = \{f \in \mathrm{Sym}(D) : \exists g \in \langle sw_\Gamma \rangle$ such that $f = g$ or $f = -_\Gamma \circ g\}$.

With this established, the proof consists of straightforward case checking, which we leave to the reader.

(ix) Let $f \in \langle -, rot \rangle \cap \langle sw_\Gamma, -_\Gamma \rangle$. By Lemma 3.5, and by composing with $-$ if necessary, we may assume that $f \in \langle rot \rangle$, so we are done by (viii). $\square$

5. $\mathcal{L}$ contains all the reducts

The task of showing that $\mathcal{L}$ contains all the reducts is split up into these lemmas:

Lemma 5.1. *Let G be a reduct of $Aut(D, E)$. Then either G contains $Aut(\Gamma)$, is contained in $Aut(\Gamma)$, or contains $\langle rot \rangle$.*

Lemma 5.2. *Let G be a reduct of $Aut(D, E)$ that contains $Aut(\Gamma)$. Then $G = \Gamma, \langle sw_\Gamma \rangle, \langle -_\Gamma \rangle, \langle sw_\Gamma, -_\Gamma \rangle$ or $Sym(D)$.*

Lemma 5.3. *Let G be a reduct of $Aut(D, E)$ that is contained in $Aut(\Gamma)$. Then $G = Aut(D, E)$, $\langle sw \rangle$, $\langle - \rangle, \langle sw, - \rangle$ or $Aut(\Gamma)$.*

Lemma 5.4. *Let G be a reduct of $Aut(D, E)$ that contains $\langle rot \rangle$. Then $G = \langle rot \rangle, \langle rot, - \rangle$ or $Sym(D)$.*

The main tool that will be used to prove these lemmas will be that of canonical functions, as developed by Bodirsky and Pinsker in [6] and [7]. However, before delving into the use of canonical functions, the next subsection describes the details that are obtained by other means.

5.1. Using the classification of the reducts of the random graph and of the random tournament

Knowing the reducts of the random graph is evidently necessary for this result, but it is also helpful to know the reducts of the random tournament. We begin by stating these two classifications.

Notation.

(i) We let $\mathcal{T} = (T, E_T)$ denote the random tournament. This can be defined as the countable homogeneous tournament which embeds all finite tournaments. Note that a tournament is a digraph that does not contain any non-edges.

(ii) Let $-_\mathcal{T}$ denote a function which switches the direction of all edges in the random tournament.

(iii) Let $sw_\mathcal{T}$ denote a function which switches the direction of only those edges that are adjacent to a particular fixed vertex.

Theorem 5.5. *(i) (Thomas [2].) The reducts of the random graph are:* $\Gamma, \langle sw_\Gamma \rangle$, $\langle -_\Gamma \rangle, \langle sw_\Gamma, -_\Gamma \rangle$ *and the full symmetric group.*

(ii) (Bennett, [10].) The reducts of the random tournament are: $Aut(T, E_T), \langle sw_\mathcal{T} \rangle, \langle -_\mathcal{T} \rangle$, $\langle sw_\mathcal{T}, -_\mathcal{T} \rangle$ *and the full symmetric group Sym(T).*

We immediately get:

Proof of Lemma 5.2. This is exactly the statement of Theorem 5.5 (i). $\square$

Theorem 5.5 (ii) contributes to the proof of Lemma 5.3, via the following construction.

Definition 5.6. (i) Let $f : \mathcal{M} \to \mathcal{N}$ be a function between two structures. The *behaviour* of f, $be_{\mathcal{M},\mathcal{N}}(f)$, is the relation $\{(p, q) \in S(\mathcal{M}) \times S(\mathcal{N}) : \exists \bar{a} \in M, \bar{b} \in N$ such that $\mathrm{tp}(\bar{a}) = p$, $\mathrm{tp}(\bar{b}) = q$ and $f(\bar{a}) = \bar{b}\}$. If $\mathcal{M} = \mathcal{N}$ we write $be_\mathcal{M}(f)$. We omit the subscript altogether if it is clear which structures we are considering.

(ii) Let $\mathcal{M}$ be any structure and $F \subseteq \mathrm{Sym}(M)$. The behaviour of F, $be_\mathcal{M}(F)$, is defined to be $\bigcup_{f \in F} be_\mathcal{M}(f)$.

(iii) Let G be a reduct of (D, E). We define $\Theta(G)$ to be $\{f \in \mathrm{Sym}(T) : be_{\mathcal{T}}(f) \subset be_{(D,E)}(G)\}$.

In words, $\Theta(G)$ contains those functions whose action on finite sets can be replicated by functions in G. The intuition is that $\Theta(G)$ tells us what G can do to tournaments. The idea behind this concept is as follows: $\Theta(G)$ is a reduct of $\mathcal{T}$, so by Theorem 5.5 $\Theta(G)$ has five different possibilities. Now if we assume that G fixes non-edges, G can only change the direction of edges. From this, one might suspect that G is determined by how it behaves on tournaments, i.e., that G is determined by $\Theta(G)$.

The following lemma captures two of the key properties that make this a useful construction.

Lemma 5.7. *Let G be a reduct of (D, E). Then:*

- *$\Theta(G)$ is a reduct of $\mathcal{T}$.*
- *$be(\Theta(G)) = be(G) \cap (\mathcal{S}(\mathcal{T}) \times \mathcal{S}(\mathcal{T}))$.*

Proof. Follows straightforwardly from the definitions and a back-and-forth argument, and is left as an exercise for the reader. □

Lemma 5.8. *Let G be a reduct of (D, E) contained in $Aut(\Gamma)$. Then:*

(i) $G = Aut(D, E) \Leftrightarrow \Theta(G) = Aut(T, E_T)$.

(ii) $G = \langle sw \rangle \Leftrightarrow \Theta(G) = \langle sw_{\mathcal{T}} \rangle$.

(iii) $G = \langle - \rangle \Leftrightarrow \Theta(G) = \langle -_{\mathcal{T}} \rangle$.

(iv) $G = \langle sw, - \rangle \Leftrightarrow \Theta(G) = \langle sw_{\mathcal{T}}, -_{\mathcal{T}} \rangle$.

Proof. By Lemma 5.7 and Theorem 5.5, $\Theta(G) = \mathrm{Aut}(T, E_T), \langle sw_{\mathcal{T}} \rangle, \langle -_{\mathcal{T}} \rangle$, $\langle sw_{\mathcal{T}}, -_{\mathcal{T}} \rangle$ or $\mathrm{Sym}(T)$.

(i) Follows immediately by considering the contrapositive statements.

(ii) "$\Rightarrow$". By (i), $\Theta(\langle sw \rangle) \neq \mathrm{Aut}(T, E_T)$. Suppose for contradiction that $\Theta(\langle sw \rangle)$ contains $\langle -_{\mathcal{T}} \rangle$. This implies that there is $g \in \langle sw \rangle$ which switches the direction of all the edges of some 3-element tournament in D. So g does not preserve P_{sw}, contradicting Lemma 3.3.

“$\Leftarrow$”. Suppose $\Theta(G) = \langle sw_T \rangle$. Since $\langle sw_T \rangle$ preserves P_{sw}, G also preserves P_{sw}. By Lemma 3.3, we get that $G = \mathrm{Aut}(D, E)$ or $\langle sw \rangle$. But it cannot be the former option, so $G = \langle sw \rangle$.

(iii) Same arguments as for part (ii).

(iv) “$\Rightarrow$”. This is proved similarly to previous cases.

“$\Leftarrow$”. Suppose $\Theta(G) = \langle sw_T, -_T \rangle$. This implies G preserves $P_{sw,w}$, which implies that $G \leq \langle sw, - \rangle$. In $\langle sw_T, -_T \rangle$, there is a function that does not preserve sw, so there is a function $g \in G$ which does not preserve sw. Hence, by Lemma 3.5, $g = - \circ g'$ where $g' \in sw$. Then g^2 will be in $\langle sw \rangle \backslash \mathrm{Aut}(D, E)$. Hence, by Lemma 3.3, $G \geq \langle sw \rangle$. By composing g with an appropriate element of $\langle sw \rangle$, we get that $- \in G$. Hence, we have that $G \geq \langle sw, - \rangle$. Thus, $G = \langle sw, - \rangle$, as required. $\square$

This lemma almost completes the proof of Lemma 5.2. What is left to prove is that if $\langle -, sw \rangle < G \leq \mathrm{Aut}(\Gamma)$, then $G = \mathrm{Aut}(\Gamma)$. We believe that this can be proved using $\Theta(G)$, without the need of canonical functions, but the combinatorics involved were just out of our reach.

5.2. Canonical functions

Definition 5.9. Let $\mathcal{M}, \mathcal{N}$ be any structures and let $f : M \to N$ be any function.

(i) If the behaviour of f is a function $S(M) \to S(N)$, then we say f is *canonical*. Rephrased, we say f is canonical if for all $\bar{a}, \bar{a}' \in M$, $\mathrm{tp}(\bar{a}) = \mathrm{tp}(\bar{a}') \Rightarrow \mathrm{tp}(f(\bar{a})) = \mathrm{tp}(f(\bar{a}'))$.

(ii) If f is canonical, we use the same symbol f to denote its behaviour.

For example, any $f \in \mathrm{Aut}(D, E)$ is a canonical function, and for all types p, $f(p) = p$. The benefit of canonical functions is that they are particularly well-behaved and can be easily manipulated and analysed. The next theorem will be treated as a ‘black-box’ for this article - a proof can be found in [7]. In order to state the theorem, we need to give a couple of definitions.

Definition 5.10. Let $F \subseteq D^D$. We let $\mathrm{cl}_{\mathrm{tm}}(F)$, the topological monoid closure of F, denote the smallest closed monoid in D^D containing F. We may abuse notation and write $\mathrm{cl}_{\mathrm{tm}}(F)$ for

$\mathrm{cl}_{\mathrm{tm}}(\mathrm{Aut}(D, E) \cup F)$.

Definition 5.11. We let $(D, E, <)$ denote the countable (linearly) ordered homogeneous digraph that embeds all finite ordered digraphs.

The theorem that follows is an application of the theorem in [7] to the structure $(D, E, <)$. For this to be valid, we need to know that $(D, E, <)$ is a Ramsey structure. The definition of a Ramsey structure can be found in [7]. The fact that $(D, E, <)$ is Ramsey follows from the main theorem of [13].

Theorem 5.12. *Let $f \in Sym(D)$ and $\bar{c} \in D$. Then there exists a function $g : D \to D$ such that*

(i) $g \in cl_{tm}(Aut(D, E) \cup \{f\})$.

(ii) $g(\bar{c}) = f(\bar{c})$.

(iii) When regarded as a function from $(D, E, <, \bar{c})$ to (D, E), g is a canonical function.

How is this theorem used? We illustrate by sketching how we will complete the proof of Lemma 5.3: G is a closed group such that $\langle -, sw \rangle < G \leq \mathrm{Aut}(\Gamma)$. Thus, G does not preserve $P_{sw,w}$; we let $f \in G$ and $c_1, \ldots, c_6 \in D$ witness this fact. We now use Theorem 5.12 to obtain the canonical g as in the theorem. We then examine the possibilities for g's behaviour, which boils down to some finite combinatorics. Using Lemma 3.8 we show that in all the possible behaviours, G must contain $\mathrm{Aut}(\Gamma)$.

The task of examining the possible behaviours is greatly simplified because the behaviour of a canonical function $f : (D, E, <, \bar{c}) \to (D, E)$ is determined by the restriction of the behaviour to 2-types. This follows because the $(D, E, <, \bar{c})$ has quantifier elimination and because the arities of the named relations are two or less.

5.2.1. Canonical functions from $(D, E, <)$

We start our analysis with the simplest situation, which is when no constants are added. As per the discussion above, it suffices to analyse the possible behaviours restricted to 2-types. To do this, we first need to describe what the possible 2-types of $(D, E, <)$ and (D, E) are.

Notation. Let $\phi_1(x,y), \ldots, \phi_n(x,y)$ be formulas. We let $p_{\phi_1,\ldots,\phi_n}(x,y)$ denote the (partial) type determined by the formula $\phi_1(x,y) \wedge \ldots \wedge \phi_n(x,y)$.

For example, let $a, b \in (D, E, <)$ be such that $a < b$ and $E(a,b)$. Then $p_{<,E}(x,y) = \mathrm{tp}(a,b)$. We will often omit the free variables x and y and write, for example, $p_{<,E}$.

With this notation in place, it is easy to state what the 2-types of $(D, E, <)$ and (D, E) are.

- There are three 2-types in (D, E): p_E, p_{E^*} and p_N.
- There are six 2-types in $(D, E, <)$: $p_{<,E}, p_{<,E^*}, p_{<,N}$, $p_{>,E}, p_{>,E^*}$ and $p_{>,N}$.

We can now analyse the behaviours.

Lemma 5.13. *Let G be a closed supergroup of $Aut(D, E)$ and let $f \in cl_{tm}(G)$ be a canonical function from $(D, E, <)$ to (D, E). Then (at least) one of the following is true:*

- *f behaves like id.*
- *f behaves like $-$.*
- *G contains $Aut(\Gamma)$.*

Proof. When we consider f as a function $(D, E, <) \to (D, E, <)$, we may assume without loss that f preserves the linear order. We can do this because if we let $f' : D \to D$ be a function with the same behaviour as f and which in addition preserves the linear order, then $f' \in \mathrm{cl_{tm}}(G)$, so f can be replaced by f'.

We split up the task according to the behaviour of f.

<u>Case 1</u>. $f(p_{<,N}) = p_N$.

Case 1a. $f(p_{<,E}) = p_E$ and $f(p_{<,E^*}) = p_{E^*}$, in which case f behaves like id.

Case 1b. $f(p_{<,E}) = p_{E^*}$ and $f(p_{<,E^*}) = p_E$, in which case f behaves like $-$.

Case 1c. $f(p_{<,E}) = p_E$ and $f(p_{<,E^*}) = p_E$. We will use Lemma 3.8 (ii) to show that G contains $\mathrm{Aut}(\Gamma)$, so let $\bar{a} \in D$. Then let $\bar{b} \in D$ such that $\bar{a} \cong \bar{b}$ *as digraphs*, and $b_i < b_j$ for all $i < j$. By homogeneity of (D, E) there is $g_1 \in G$ such that $g_1(\bar{a}) = \bar{b}$ and since $f \in \mathrm{cl}_{\mathrm{tm}}(G)$ there is $g_2 \in G$ such that $g_2(\bar{b}) = f(\bar{b})$. Then observe that $g = g_2 g_1$ satisfies $E(g(a_i), g(a_j)) \leftrightarrow i < j$. Thus we have satisfied the assumptions of Lemma 3.8 (ii), so $G \geq \mathrm{Aut}(\Gamma)$.

The case where $f(p_{<,E}) = p_{E^*}$ and $f(p_{<,E^*}) = p_{E^*}$ is symmetric to this case.

Case 1d. $f(p_{<,E}) = p_N$ or $f(p_{<,E^*}) = p_N$. Without loss suppose the first is true, the latter case is symmetric. We will use Lemma 3.6 (iii) and show that $G = \mathrm{Sym}(D)$, so in particular G contains $\mathrm{Aut}(\Gamma)$. Let $\bar{a} \in D$ contain an edge $a_i a_j$. By homogeneity of (D, E) there is $g_1 \in G$ such that $g_1(\bar{a}) \cong \bar{a}$ and $g_1(a_i) < g_1(a_j)$. Now let $g_2 \in G$ equal f on $g_1(\bar{a})$. Observe that g_2 deletes the edge $g_1(a_i)g_1(a_j)$ (and possibly others too) and we also know that g_2 preserves edges. Hence, $g_2 g_1(\bar{a})$ contains less edges than $\bar{a}$, so by Lemma 3.6, we are done.

<u>Case 2</u> $f(p_{<,N}) = p_E$.

Case 2a. Neither $f(p_{<,E}) = p_N$ nor $f(p_{<,E^*}) = p_N$. We will use Lemma 3.6 (iii) and show that $G = \mathrm{Sym}(D)$. Let $\bar{a} \in D$ contain an edge $a_i a_j$. By homogeneity of (D, E), map $\bar{a}$ to an isomorphic (as digraphs) tuple $\bar{b}$ where b_i, resp. b_j, is the least, resp. second least, element of $\bar{b}$. By assumption, f maps $\{b_i, b_j\}$ to an edge but we do not know its direction. This splits into two cases.

Subcase (i). Suppose we have $E(f(b_i), f(b_j))$. Now let $\bar{b}'$ be an ordered digraph which is the same as $\bar{b}$ except that b_i, b_j is changed to a non-edge. Observe that $f(\bar{b}') \cong f(\bar{b})$, because $f(p_{<,N}) = p_E$.

But now we are done: we can find mappings in G to get from $\bar{a}$ to $\bar{b}$ to $f(\bar{b})$ to $f(\bar{b}')$ to $\bar{b}'$ (noting that though f may not be invertible, the function in G which agrees with f on $\bar{b}'$ is invertible), and $\bar{b}'$ has less edges than in $\bar{a}$.

Subcase (ii). Suppose we have $E^*(f(b_i, b_j))$. The previous argument does not work as stated, because the edges $f(b_i, b_j)$ and $f(b'_i, b'_j)$ will not be in the same direction. To fix this, we modify $\bar{b}'$ by swapping b'_i and b'_j with respect to the linear order. Now, $f(b_i, b_j)$ and $f(b'_i, b'_j)$ will be in the same direction.

Furthermore, because we earlier specified that b_i and b_j should be the two least elements of $\bar{b}$, this swapping only affects the type of the pair b_ib_j - all other pairs' types are unaffected. This ensures that we have $f(\bar{b}') \cong f(\bar{b})$. The rest of the proof continues as in the previous case.

Case 2b. $f(p_{<,E}) = p_N$ and $f(p_{<,E^*}) = p_N$. Since we assumed that f preserves the linear order, f^2 is a canonical function in $\mathrm{cl}_{\mathrm{tm}}(G)$ where $f^2(p_{<,N}) = p_N$, $f(p_{<,E}) = p_E$ and $f(p_{<,E^*}) = p_E$. Hence, by Case 1c, G contains $\mathrm{Aut}(\Gamma)$.

Case 2c. $f(p_{<,E}) = p_N$ and $f(p_{<,E^*}) = p_E$. By considering f^2, this case is reduced to Case 1d, so G contains $\mathrm{Aut}(\Gamma)$.

Case 2d. $f(p_{<,E}) = p_N$ and $f(p_{<,E^*}) = p_{E^*}$. We will use Lemma 3.6 (iii). Let $\bar{a} \in D$ contain an edge $E(a_i, a_j)$. By composing with an element of $\mathrm{Aut}(D, E)$ if necessary, we may assume that a_j is the least element, and a_i is the second least element. In particular, we have $\mathrm{tp}(a_j, a_i) = p_{<,E^*}$. Let $\bar{b}$ be an ordered digraph such that $f(\bar{b}) = \bar{a}$. Now let $\bar{b}'$ be an ordered digraph which is the same as $\bar{b}$ except we swap the position of b'_i and b'_j in the linear order. This means that $\mathrm{tp}(b'_j, b'_i) = p_{>,E^*}$ and because b_i, b_j are the least elements, the types of all the other pairs are unaffected. Hence, $f(\bar{b}')$ is the same digraph as $\bar{a}$ but the edge a_ia_j is replaced by a non-edge. Hence, we are done.

Case 2e. $f(p_{<,E}) = p_E$ and $f(p_{<,E^*}) = p_N$. Considering f^2 reduces us to Case 2a.

Case 2f. $f(p_{<,E}) = p_{E^*}$ and $f(p_{<,E^*}) = p_N$. Imitate the argument in Case 2d to show that $G = \mathrm{Sym}(D)$.

Case 3 $f(p_{<,N}) = p_{E^*}$. This is symmetric to Case 2. □

5.2.2. Canonical functions from $(D, E, <, \bar{c})$

We now move on to the general situation where constants $\bar{c} \in D$ are added to the structure. For convenience, we assume that $c_i < c_j$ for all $i < j$. As before, n-types of $(D, E, <, \bar{c})$ correspond to orbits of $\mathrm{Aut}(D, E, <, \bar{c})$ acting on n-tuples. Consequently, we often conflate the notions of types and orbits.

$(D, E, <, \bar{c})$ has two kinds of 1-types, i.e. two kinds of orbits: singleton orbits, which are of the form $\{c_i\}$, and infinite orbits, which are determined by how their elements are related to the c_i. For example, $\{x \in D : x < c_1 \wedge \bigwedge_i E(x, c_i)\}$ is an infinite orbit. In order to describe the 2-types, we extend the notation from the previous section.

Notation Let A, B be definable subsets of $(D, E, <, \bar{c})$ and let $\phi_1(x, y), \ldots, \phi_n(x, y)$ be formulas. We let $p_{A,B,\phi_1,\ldots,\phi_n}(x, y)$ denote the (partial) type determined by the formula $x \in A \wedge y \in B \wedge \phi_1(x, y) \wedge \ldots \wedge \phi_n(x, y)$.

Now let X and Y be orbits, $\phi \in \{<, >\}$ and $\psi \in \{E, E^*, N\}$. Then all the 2-types of $(D, E, <, \bar{c})$ are of the form $p_{X,Y,\phi,\psi} = \{(a, b) \in D : a \in X, b \in Y, \phi(a, b) \text{ and } \psi(a, b)\}$.

Our task now is to analyse the possibilities for $f(p_{X,Y,\phi,\psi})$, where f is a canonical function. The analysis is split into cases depending on how the orbits X and Y relate. The first lemma deals with the situation when $X = Y$.

Lemma 5.14. *Let G be a closed supergroup of $Aut(D, E)$, let $f \in cl_{tm}(G)$ be a canonical function from $(D, E, <, \bar{c})$ to (D, E) and let X be an infinite orbit of $Aut(D, E, \bar{c})$. Then (at least) one of the following holds:*

- *f behaves like id on X.*
- *f behaves like $-$ on X.*
- *G contains $Aut(\Gamma)$.*

Proof. By noting that $(X, E|_X)$ is isomorphic to (D, E), unravelling the definitions will show that this lemma has exactly the same mathematical content as Lemma 5.13. □

Next, we look at how f can behave between two infinite orbits. For this we need to look at how two infinite orbits can relate to each other with respect to the linear order.

Facts and Notation There are two ways that two infinite orbits X and Y of $\mathrm{Aut}(D, E, <, \bar{c})$ can relate to each other with respect to the linear order $<$:

- All of the elements of one orbit, X say, are smaller than all of the elements of Y. This is abbreviated by $X < Y$

- X and Y are interdense: $\forall x < x' \in X, \exists y \in Y$ such that $x < y < x'$ and vice versa.

We deal with these possibilities separately, starting with the case where one orbit is below the other.

Lemma 5.15. *Let G be a closed supergroup of $Aut(D, E)$, let $f \in cl_{tm}(G)$ be a canonical function from $(D, E, <, \bar{c})$ to (D, E) and let X and Y be infinite orbits of $Aut(D, E, \bar{c})$ such that f behaves like id on X and $X < Y$. Then (at least) one of the following holds:*

- *f behaves like id, sw, rot or rot^{-1} between X and Y.*

- *G contains $Aut(\Gamma)$.*

Proof. Let $y_0 \in Y$ be fixed. We emphasise now an important feature of this proof, which is that our arguments only depend on how f behaves on $X \cup \{y_0\}$. This is done intentionally so that these arguments can be used unaltered in later lemmas.

As the arguments are similar to that of Lemma 5.14, the proofs are more sketchy, and we leave the details to the reader.

<u>Case 1</u> $f(p_{X,Y,N}) = p_N$. (Note that because $X < Y$, $p_{X,Y,\psi} = p_{X,Y,<,\psi}$ for any formula ψ.)

Case 1a. $f(p_{X,Y,E}) = p_E$ and $f(p_{X,Y,E^*}) = p_{E^*}$. Then f behaves like id between X and Y.

Case 1b. $f(p_{X,Y,E}) = p_{E^*}$ and $f(p_{X,Y,E^*}) = p_E$. Then f behaves like sw between X and Y.

Case 1c. $f(p_{X,Y,E}) = p_E$ and $f(p_{X,Y,E^*}) = p_E$. We will use Lemma 3.8 (ii) to show that G contains Γ. Let $\bar{a} = (a_1, \ldots, a_n) \in D$. We want to show that by using elements of G, we can switch the direction of the edges of $\bar{a}$ so they are all pointing in the same direction. We do this by induction on n. The base case $n = 1$ is trivial so let $n > 1$. By the inductive hypothesis, we can assume that for $1 \le i, j \le n-1$, if $a_i a_j$ is an edge, then $E(i, j) \leftrightarrow i < j$. By homogeneity, map a_n to y_0 and the other a_i's into X. Then applying f switches the edges adjacent to a_n so they are all directed into a_n and furthermore

f does not alter any of the other edges. Thus, the resulting digraph has all edges going in the same direction, as required.

The case where $f(p_{X,Y,E}) = p_{E^*}$ and $f(p_{X,Y,E^*}) = p_{E^*}$ can be dealt with with the same argument.

Case 1d. $f(p_{X,Y,E}) = p_N$ or $f(p_{X,Y,E^*}) = p_N$. Given any $\bar{a} \in D$ which contains edges, we use f to delete edges from it, so by Lemma 3.6 (iii), $G = \mathrm{Sym}(D)$. See Case 1d of Lemma 5.14 for more detail.

<u>Case 2</u> $f(p_{X,Y,N}) = p_E$.

Case 2a. Neither $f(p_{X,Y,E}) = p_N$ nor $f(p_{X,Y,E^*}) = p_N$. $G = \mathrm{Sym}(D)$, by using the same argument as in Case 2a of Lemma 5.14.

Case 2b. $f(p_{X,Y,E}) = p_N$ and $f(p_{X,Y,E^*}) = p_N$. By considering f^2 we reduce this to Case 1c.

Case 2c. $f(p_{X,Y,E}) = p_N$ and $f(p_{X,Y,E^*}) = p_E$. By considering f^2 we reduce this to Case 1d.

Case 2d. $f(p_{X,Y,E}) = p_N$ and $f(p_{X,Y,E^*}) = p_{E^*}$. We show that $G = \mathrm{Sym}(D)$ using the same idea as Case 2d in Lemma 5.14. Let $\bar{b} \in D$ contain an edge $E(b_i, b_j)$. Let $\bar{b}'$ be obtained by mapping b_i to y_0, applying f, then mapping b_j to y_0 and applying f again. Note that we have $N(b_i', b_j')$. Let $\bar{b}''$ be obtained from $\bar{b}$ in the same way, except we map b_j to y_0 first, and then b_i second. In this case, we have $E^*(b_i'', b_j'')$. Furthermore, $\bar{b}_1$ and $\bar{b}_2$ are otherwise the same. Now suppose $\bar{a}$ is given and has an edge. Find a $\bar{b}$ such that its corresponding $\bar{b}''$ is isomorphic to $\bar{a}$. Then, we can get from $\bar{a}$ to $\bar{b}''$ to $\bar{b}$ to $\bar{b}'$, i.e., we can delete an edge from $\bar{a}$. Thus, $G = \mathrm{Sym}(D)$.

Case 2e. $f(p_{X,Y,E}) = p_E$ and $f(p_{X,Y,E^*}) = p_N$. Considering f^2 reduces us to Case 2a.

Case 2f. $f(p_{X,Y,E}) = p_{E^*}$ and $f(p_{X,Y,E^*}) = p_E$. Then f behaves like *rot* between X and Y.

Case 3. $f(p_{X,Y,N}) = p_{E^*}$. This case is symmetric to Case 2. □

As mentioned at the start of the proof, what was relevant is how f behaved on $X \cup \{y_0\}$. More specifically, what was sufficient to make these arguments work was the following: For all finite digraphs $\bar{a} \in D$ and all points $a \in \bar{a}$, we can find a copy of $\bar{a}$ in $X \cup \{y_0\}$ such that $\bar{a} \cap \{y_0\} = \{a\}$.

This condition is satisfied in the remaining situations that need to be analysed, so their corresponding results are immediate corollaries of Lemma 5.15. The statement for interdense orbits has to be modified and will perhaps appear confusing. Clarification will be provided after the statement.

Corollary 5.16. *Let G be a closed supergroup of $Aut(D, E)$, let $f \in cl_{tm}(G)$ be a canonical function from $(D, E, <, \bar{c})$ to (D, E) and let X and Y be interdense infinite orbits of $Aut(D, E, \bar{c})$ such that f behaves like id on X . Suppose that G does not contain $Aut(\Gamma)$. Then both of the following hold:*

- *f behaves like id, sw, rot or rot^{-1} between increasing tuples from X to Y.*
- *f behaves like id, sw, rot or rot^{-1} between decreasing tuples from X to Y.*

For example, f behaves like sw between increasing tuples from X to Y means that $f(p_{X,Y,<,E}) = p_{E^*}, (p_{X,Y,<,E^*}) = p_E$ and $(p_{X,Y,<,N}) = p_N$.

Lastly, we look at how f can behave between the constants and the infinite orbits.

Definition 5.17. : Let c be one of the named constants of $(D, E, <, \bar{c})$ and let X_1, X_2 and X_3 be infinite orbits. If it is the case that we have outward edges from c to X_1, inward edges from c to X_2 and non-edges between c and X_3, we called the triple $\bar{X} = (X_1, X_2, X_3)$ a c-generic triple.

The reason for introducing this definition is that there is nothing to be analysed about how f behaves between c and a *single* orbit X. It is only useful to ask how f behaves between c and several infinite orbits, in particular, a c-generic triple. Note that if $\bar{X}$ is a c-generic triple, then $c \cup X_1 \cup X_2 \cup X_3$ is isomorphic to the generic digraph.

Corollary 5.18. *Let G be a closed supergroup of $Aut(D, E)$, let $f \in cl_{tm}(G)$ be a canonical function from $(D, E, <, \bar{c})$ to (D, E), let c be one of the named constants and let $\bar{X}$ be a c-generic triple such that f behaves like id on $X_1 \cup X_2 \cup X_3$. Then (at least) one of the following holds:*

- *f behaves like id, sw, rot or rot^{-1} between c and $\bigcup \bar{X}$.*
- *G contains $Aut(\Gamma)$.*

With this analysis, we are now in a position to prove the remaining lemmas.

Proof of Lemma 5.1. We recall that we want to show that if G is a reduct of (D, E), then G either contains $\mathrm{Aut}(\Gamma)$, is contained in $\mathrm{Aut}(\Gamma)$, or contains $\langle rot \rangle$. Suppose none of these are true - we will derive a contradiction.

G is not contained in $\mathrm{Aut}(\Gamma)$, which means that G does not preserve non-edges. Hence, there is $f \in G$ and an edge $c_1c_2 \in D$ such that $f(c_1c_2)$ is a non-edge. We apply Theorem 5.12 to obtain a canonical $g : (D, E, <, c_1, c_2) \to (D, E)$ which agrees with f on c_1 and c_2.

By Lemma 5.14, for any infinite orbit X, g behaves like id or $-$ on it - otherwise, G would contain $\mathrm{Aut}(\Gamma)$, contradicting our assumptions. By Lemma 5.15 and its corollaries, we have that g behaves like id or sw between orbits - otherwise, G would contain either $\mathrm{Aut}(\Gamma)$ or $\langle rot \rangle$, contradicting our assumptions.

But now we have a function $g \in \mathrm{cl_{tm}}(G)$ which deletes an edge (namely,c_1c_2) and maps all non-edges to non-edges. By imitating the proof of part (iv) of Lemma 3.6, we conclude that G equals $\mathrm{Sym}(D)$, contradicting that G does not contain $\mathrm{Aut}(\Gamma)$. □

Proof of Lemma 5.3. By Lemma 5.7 and Lemma 5.8, it remains to be proved that if $\langle sw, - \rangle < G \leq \mathrm{Aut}(\Gamma)$ is a closed group, then $G = \mathrm{Aut}(\Gamma)$.

So let G be such a closed group and suppose, for contradiction, that $G \neq \mathrm{Aut}(\Gamma)$. Since $\langle sw, - \rangle < G$, G does not preserve $P_{sw,w}$, so there exist $f \in G$ and $\bar{c} \in D$ such that $P_{sw,w}(\bar{c})$ and $\neg P_{sw,w}(f(\bar{c}))$. Then Theorem 5.12 gives $g \in \mathrm{cl_{tm}}(G)$ which is canonical from $(D, E, <, \bar{c})$ and which agrees with f on $\bar{c}$.

As in the previous proof, we use Lemma 5.14 to conclude that for any infinite orbit X, g behaves like id or $-$ on X, and we use Lemma 5.15 and its corollaries to conclude that g behaves like id or sw between orbits.

Claim 1. g behaves like id on all infinite orbits, or, g behaves like $-$ on all infinite orbits.

Proof of Claim 1. Suppose not, so there exists infinite orbits X and Y such that g behaves like id on X and like $-$ on Y. There are now two cases. The first case is if g behaves like id between X and Y. By using part (iii) of Lemma 3.8, we can conclude that $G \geq \mathrm{Aut}(\Gamma)$ - contradiction. The second case is if g behaves like sw between X and Y. This reduces to the previous case by considering $- \circ g$.

We can take the claim one step further: by considering $- \circ g$ if necessary, we may now assume that g behaves like id on all infinite orbits.

Enumerate the (finite number of) infinite orbits as $X_1, X_2, X_3, \ldots$.

Claim 2. We may assume g behaves like id between all pairs of orbits from X_1, X_2, X_3.

Proof of Claim 2. If g behaves like id between all three orbits, we are done. If g behaves like sw between precisely two of the pairs - without loss g behaves like sw between X_1 and X_2, and between X_1 and X_3 - then by switching about X_1 we may assume g behaves like id, so again we are done. If g behaves like sw between precisely one pair of infinite orbits, then by using Lemma 3.8 (iii), we get that $G \geq \mathrm{Aut}(\Gamma)$, which is a contradiction. The final possibility is that g behaves like sw between all pairs: this reduces to the third case by switching about X_1, so we again get a contradiction.

Claim 3. We may assume g behaves like id between all infinite orbits.

Proof of Claim 3. First consider how g behaves between X_4 and the first three X_i's. If g behaves like id between X_4 and all the previous X_i's, we move on. If g behaves like sw between X_4 and all the previous X_i's, we switch about X_4, reducing to the first case. The last case is if, without loss, g behaves like id between X_4 and X_1 and like sw between X_4 and X_2. But this is exactly the same as the contradictory case in the proof Claim 2, so this is not possible. Hence, we have shown that g must behave like id between all the pairs in $X_1, \ldots, X_4$.

One then moves on to X_5 and repeats this argument to show that we may assume g behaves like id between $X_1, \ldots X_5$. Continuing in this fashion proves the claim.

Claim 4. We may assume that g behaves like id between $\bar{c}$ and the infinite orbits.

Let $c_1 \in \bar{c}$. Suppose there are $\bar{c_1}$-generic triples $\bar{X}$ and $\bar{Y}$ such that g behaves like id, resp. sw, between c and $\bigcup \bar{X}$, resp. $\bigcup \bar{Y}$. Then for any finite digraph A and edge $aa' \in A$, we can map aa' to an edge between c_1 and $\bar{Y}$ and map the remaining vertices of A into $\bar{X}$. Then applying g will switch precisely the single edge aa'. Thus, by Lemma 3.8 (iii), we get that $G \geq \mathrm{Aut}(\Gamma)$, a contradiction. Hence, g behaves like id between c_1 and every c_1-generic triple, in which case we are done, or, g behaves like sw between c_1 and every c_1-generic triples, in which case we apply sw_{c_1}.

Repeating this for the other c_i will complete the proof of the claim.

Observe that all the manipulations we (may) have used on g have been applications of $-$ or sw. This ensures that $g(\bar{c}) \notin P_{sw,w}$. In particular, there is at least one edge in $\bar{c}$ whose direction g switches. Combining this observation with all the claims tells us that we are in the situation of Lemma 3.8 (iii): g behaves like id everywhere except on the finite set $\bar{c}$. Hence, $G \geq \mathrm{Aut}(\Gamma)$, giving us a contradiction, thus completing the proof. $\square$

Proof of Lemma 5.4. We recall the statement of the lemma: If G contains $\langle rot \rangle$, then G equals $\langle rot \rangle, \langle -, rot \rangle$ or $\mathrm{Sym}(D)$.

To prove the statement, it suffices to prove:

(i) If $G > \langle rot \rangle$ and $G \not\geq \langle -, rot \rangle$, then $G = \mathrm{Sym}(D)$.

(ii) If $G > \langle -, rot \rangle$, then $G = \mathrm{Sym}(D)$.

Recall that $\langle sw, rot \rangle = \mathrm{Sym}(D)$. Hence, we may assume in all that follows that $sw \notin G$.

(i) Suppose $G > \langle rot \rangle$ and $G \not\geq \langle -, rot \rangle$. The latter assumption implies that $- \notin G$. Then there exists $\bar{c}$ and $f \in G$ which witness the fact that G does not preserve $P_{rot,1}$. Then use Theorem 5.12 to obtain a canonical $g : (D, E, <, \bar{c}) \to (D, E)$ which agrees with f on $\bar{c}$.

By Lemma 5.14, g behaves like id on all infinite orbits, as otherwise G contains sw or $-$. Similarly, by Lemma 5.15 and its corollaries, we have that g behaves like id or rot between orbits.

We proceed in a similar fashion to the proof of Lemma 5.3.

Claim 1. We may assume that g behaves like the id between all infinite orbits.

Proof of Claim 1. Let $X_1, X_2, \ldots$ enumerate the infinite orbits. If g behaves like id between X_1 and X_2 then move on. Otherwise, g behaves like rot or rot^{-1} between X_1 and X_2. Hence, by composing with a rotation about X_1 or X_2 as appropriate, we can assume g behaves like id between X_1 and X_2.

Now consider X_3. Again by composing with a rotation if necessary, we may assume that g behaves like id between X_1 and X_3. Suppose that g does not behave like id between X_2 and X_3 - we will show that G must equal $\mathrm{Sym}(D)$. Without loss, g behaves like rot between X_2 and X_3. Given any finite digraph A and an edge aa' in A, we can find a copy of A in D such that aa' is an inward edge from X_2 to X_3 and such that the other vertices of A all lie in X_1. Applying g to this copy results in the edge aa' being deleted, with the rest of A being the same. Hence, by Lemma 3.6 (iii), we conclude that $G = \mathrm{Sym}(D)$. Therefore, we may assume that g behaves like id between X_1, X_2 and X_3.

Now consider X_4 and repeat this argument. Continuing in this fashion proves the claim.

Claim 2. We may assume that g behaves like id between $\bar{c}$ and all the infinite orbits.

Proof of Claim 2. First work with c_1. Let $\bar{X}$ be a c_1-generic triple. We know that g behaves like id, rot or rot^{-1} between c and $\bar{X}$, so by composing with a rotation if necessary, we may assume that g behaves like id. Now consider another c_1-generic triple $\bar{Y}$. If g does not behave like id between c_1 and $\bar{Y}$, then we use the same argument as in the proof of Claim 1 to show that G equals $\mathrm{Sym}(D)$. So we may assume that g behaves like id between c_1 and all infinite orbits.

Repeating this for c_2 and c_3 completes the proof of the claim.

Because all the modifications of g were compositions with rotations, we still have $\neg P_{rot,1}(g(\bar{c}))$. This means g acts like id everywhere except on $\bar{c}$. Then using Lemma 3.8 or Lemma 3.6, as appropriate, we get that $G \geq \mathrm{Aut}(\Gamma)$, so G contains sw, so $G = \mathrm{Sym}(D)$. This completes the proof of (i).

(ii) Suppose $G > \langle rot, - \rangle$. Then there exists $\bar{c}$ and $f \in G$ which witness the fact that G does not preserve $P_{rot,w}$. Then use Theorem 5.12 to obtain a canonical $g : (D, E, <, \bar{c}) \to (D, E)$ which agrees with f on $\bar{c}$.

Let X be an infinite orbit. Then by Lemma 5.14 and by replacing with $- \circ g$ if necessary, we may assume that g behaves like id on X.

Claim 1'. We may assume that g acts like id on all infinite orbits.

Proof of Claim 1'. Suppose not, so there is an infinite orbit Y such that g acts like $-$ on Y. Then g can behave like id, rot or rot^{-1} between X and Y. If g behaves like id, then use Lemma 3.8 (iii) to show that $G \geq \text{Aut}(\Gamma)$, so then G must equal $\text{Sym}(D)$. If g behaves like rot or rot^{-1} between X and Y, composing with rot_X or rot_X^{-1} reduces to the id behaviour between X and Y. This completes the proof of the claim.

Now by Lemma 5.15 and its corollaries, we have that g behaves like id or rot between the infinite orbits. Now continue in exactly the same way as in part (i) to complete the proof. □

6. Summary and Open Questions

We summarise the structure of the proof of the main theorem, Theorem 2.5, which states that $\mathcal{L}$ is the lattice of the reducts of the generic digraph. The first task is to show that $\mathcal{L}$ is a sublattice of the reducts of the generic digraph, which was done in Lemma 4.1.

The second task is to show that $\mathcal{L}$ contains all the reducts. By Lemma 5.1, which was proved using canonical functions at the start of Section 5.3, the task is split up into three regions of $\mathcal{L}$: The region above $\text{Aut}(\Gamma)$, the region below $\text{Aut}(\Gamma)$, and the rest. The region above $\text{Aut}(\Gamma)$ is immediately dealt with by Thomas' classification of the reducts of Γ. The proof of the region below $\text{Aut}(\Gamma)$, Lemma 5.3, has two parts. The first part is in Section 5.1, where we use the function $\Theta(G)$ and the classification for the random tournament, and the second part is in Section 5.3. The final region, Lemma 5.4, is proved using canonical functions at the end of Section 5.3.

We end by stating some problems of interest in this area. There is the obvious task of determining the reducts of your favourite structure(s), but some more specific questions are:

- (Thomas' Conjecture): If a structure is homogeneous in a finite relational language, then it only has finitely many reducts.

- Which lattices can be realised as the lattice of reducts of some structure?

- Is there always a maximal closed group between a closed group G and $\mathrm{Sym}(M)$ (where M is countable)?

The answer to the first question may be related to a question in structural Ramsey theory ([7]): Given a homogeneous structure, can you finitely extend its language so that the structure becomes Ramsey? It may be easier to prove Thomas' Conjecture for Ramsey structures and this could be sufficient, or alternatively, a counterexample for one question may lead to a counterexample for the other.

Another angle on the second question could be to consider whether there is any relationship between structures which have the same lattice of reducts. For example, it is curious that $(\mathbb{Q}, <)$, the random graph, the random tournament and the generic partial order have the same 5-element lattice as their lattice of reducts.

For clarification of the third question, we say that a closed group $F < \mathrm{Sym}(M)$ is maximal if there are no closed groups F' such that $F < F' < \mathrm{Sym}(M)$. To find a counterexample to this question, it would be sufficient to find an $\aleph_0$-categorical countable structure such that all of its non-trivial reducts have infinitely many reducts.

[1] P. Cameron, Transitivity of permutation groups on unordered sets, Mathematische Zeitschrift 148 (1976) 127–139.

[2] S. Thomas, Reducts of the random graph, Journal of Symbolic Logic 56 (1991) 176–181.

[3] S. Thomas, Reducts of random hypergraphs, Annals of Pure and Applied Logic 80 (1996) 165–193.

[4] M. Junker, M. Ziegler, The 116 reducts of $(\mathbb{Q}, <, 0)$, Journal of Symbolic Logic 74 (2008) 861–884.

[5] I. Kaplan, P. Simon, The affine and projective groups are maximal, to appear in the Transactions of the American Mathematical Society.

[6] M. Bodirsky, M. Pinsker, Reducts of ramsey structures, in: M. Grohe, J. Makowsky (Eds.), Contemporary Mathematics: Model Theorectic Methods in Finite Combinatorics, Vol. 558, American Mathematical Society, 2011, pp. 489–519.

[7] M. Bodirsky, M. Pinsker, T. Tsankov, Decidability of definability, Journal of Symbolic Logic 78 (2013) 1036–1054.

[8] P. Pach, M. Pinsker, G. Pluhár, A. Pongrácz, C. Szabó, Reducts of the random partial order, Advances in Mathematics 267 (2014) 94–120.

[9] M. Bodirsky, M. Pinsker, A. Pongrácz, The 42 reducts of the random ordered graph, Proc. Lond. Math. Soc. (3) 111 (3) (2015) 591–632.

[10] J. Bennett, The reducts of some infinite homogeneous graphs and tournaments, Ph.D. thesis, Rutgers University (1997).

[11] R. Fraïssé, Sur certaines relations généralisent l'ordre des nombres rationnels, Comptes Rendus d' l'Académie des Sciences de Paris 237 (1953) 540–542.

[12] Hodges, A Shorter Model Theory, Cambridge University Press, Cambridge, 1997.

[13] J. Nešetřil, V. Rödl, Partitions of relational and set systems, Journal of Combinatorial Theory (A) 22 (1977) 289–312.