

CORES OF COUNTABLY CATEGORICAL STRUCTURES

MANUEL BODIRSKY

Institut für Informatik, Unter den Linden 6, Humboldt-Universität zu Berlin
e-mail address: bodirsky@informatik.hu-berlin.de

ABSTRACT. A relational structure is a *core*, if all its endomorphisms are embeddings. This notion is important for computational complexity classification of constraint satisfaction problems. It is a fundamental fact that every finite structure has a core, i.e., an endomorphism such that the structure induced by its image is a core; moreover, the core is unique up to isomorphism.

We prove that every ω -categorical structure has a core. Moreover, every ω -categorical structure is homomorphically equivalent to a model-complete core, which is unique up to isomorphism, and which is finite or ω -categorical. We discuss consequences for constraint satisfaction with ω -categorical templates.

1. INTRODUCTION

The notion of the *core* of a finite structure is a central concept in structural combinatorics [22]. It can be defined in many equivalent ways, one of which is as follows (for detailed definitions see Section 2). A finite structure S is a *core* if every endomorphism of S is an automorphism. A structure C is called a *core of S* if C is a core and is isomorphic to a substructure induced by the image of an endomorphism of S .

For finite structures S it is well-known and easy to prove that a core of S always exists, and is unique up to isomorphism; see e.g. [22]. We therefore speak of *the core of* a finite relational structure S . For infinite structures various core-like properties were studied by Bauslaugh [5, 6]. In general, infinite structures might not have a core in the sense introduced above, see [5, 6].

An important application of the concept of a core is the theory of *constraint satisfaction* in computer science. Roughly speaking, a constraint satisfaction problem is a computational problem where the input consists of a given set of variables and a set of constraints that are imposed on these variables, and where the task is to assign values to the variables such that all the constraints are satisfied. Such problems appear in numerous areas of computer science, and constraint satisfaction problems attracted considerable attention in artificial intelligence.

2000 ACM Subject Classification: F.4.1.

Key words and phrases: Constraint satisfaction, cores, ω -categorical structures.

This is an extended and corrected version of a conference paper titled “The Core of a Countably Categorical Structure” [8].

Many constraint satisfaction problems can be formalized as follows, using the concept of structure homomorphisms; again, for a detailed introduction of all the involved concepts see Section 2. Let Γ be a finite or infinite structure with relational signature τ . The *constraint satisfaction problem (CSP)* for the so-called *template* Γ is the following computational problem.

CSP(Γ)

INSTANCE: A finite structure S of the same relational signature τ as the template Γ .

QUESTION: Is there a homomorphism $h : S \rightarrow \Gamma$?

We want to stress that Γ is not part of the input; each Γ defines a computational problem. Note that if two structures Γ and Δ are *homomorphically equivalent*, i.e., there is a homomorphism from Γ to Δ and from Δ to Γ , then these structures have the same constraint satisfaction problem. In particular, a structure S and its core C have the same constraint satisfaction problem.

For a finite template Γ , the computational problem CSP(Γ) is clearly contained in NP. A classification of tractable and NP-hard constraint satisfaction problems with a finite template is intensively studied, but still not complete. See [12,17,21], just to mention a few highlights on that subject. In all these approaches, the authors make use of the assumption that the templates of the constraint satisfaction problems under consideration are cores.

The class of constraint satisfaction problems with an *infinite* template was not yet studied systematically. It turns out that many interesting computational problems can be formulated with templates that are ω -categorical. The concept of ω -categoricity is central in classical model theory, and will be introduced carefully in Section 3. The following list of well-known computational problems can all be formulated with a countably infinite template that is ω -categorical.

- Allen’s interval algebra, and all its fragments [4,23,27,31]
- Problems in phylogenetic analysis [20,32]
- Tree description constraints in computational linguistics [10,11,15]
- Computational problems in the theory of relation algebras [16,24,30]
- All CSPs in monotone monadic SNP without inequality [9,17]

Moreover, every constraint satisfaction problem with a finite template Δ can also be formulated with an ω -categorical template.

In this article we study how the notion of the core of a finite relational structure can be generalized to ω -categorical structures. We say that an infinite structure Γ is a *core*, if every endomorphism e of Γ is an embedding, i.e., e is an isomorphism between Γ and the structure induced by the image of e in Γ . This is indeed a generalization of the previous definition for finite structures, since every embedding of a finite structure Δ in Δ is an automorphism. As in the finite case, we say that Γ *has a core* Δ , if Δ is a core and Δ is isomorphic to the structure induced in Γ by the image of some endomorphism of Γ .

Note that our definition of a core is fundamentally different from the notion of a *core structure* as introduced in [28] in model theory. There, a structure Γ is called a *core structure of a first-order theory T* if Γ is isomorphic to exactly one substructure of every model of T .

Results. We will show that every ω -categorical structure Γ has a core. Moreover, every ω -categorical τ -structure Γ is homomorphically equivalent to a model-complete core Γ^c , which is unique up to isomorphism, and which is finite or again ω -categorical.

If Γ is expanded by all primitive positive definable relations, then Γ is homomorphically equivalent to a homogeneous core Γ^c , which is unique up to isomorphism, and which is finite or ω -categorical. The condition that Γ is expanded by all primitive positive definable relations is natural in the context of constraint satisfaction, since a relational structure and its expansion by all primitive positive definable relations have the same computational complexity (see e.g. [26]). Since Γ^c is in this case homogeneous, it admits quantified elimination and has a $\forall\exists$ -axiomatization.

In Section 7 we also prove the following result, which has consequences for the theory of constraint satisfaction. If we expand Γ^c by a singleton relation, then the resulting constraint satisfaction problem has the same computational complexity as $\text{CSP}(\Gamma^c)$. This was shown for finite templates in [12], and is of fundamental importance for the so-called *algebraic approach* to constraint satisfaction.

2. CORES

Let Γ and Δ be relational structures with the same relational signature τ . We use the same symbols for the relation symbols from τ and for the respective relations in Γ , and the same symbols to denote a structure and its universe. A mapping $f : \Gamma \rightarrow \Delta$ is called a *homomorphism*, if for all relation symbols $R \in \tau$ and $x_1, \dots, x_n \in \Gamma$ the relation $R(f(x_1), \dots, f(x_n))$ holds in Δ whenever $R(x_1, \dots, x_n)$ holds in Γ . A homomorphism is called *strong*, if $R(x_1, \dots, x_n)$ holds in Γ *if and only if* $R(f(x_1), \dots, f(x_n))$ holds in Δ . An injective strong homomorphism from Γ to Δ is also called an *embedding*; if there exists an embedding of Γ into Δ , then Δ is called an *extension* of Γ . A homomorphism from Γ to Γ is called an *endomorphism* of Γ , and a bijective strong endomorphism of Γ is called an *automorphism* of Γ . Two structures Γ and Δ are called *homomorphically equivalent*, if there is a homomorphism from Γ to Δ and a homomorphism from Δ to Γ .

Definition 1. A (finite or infinite) structure Γ is a *core* if every endomorphism of Γ is an embedding. A core Γ is called a *core of* Δ if Γ is the image of an endomorphism of Δ .

Homomorphisms that are not embeddings are called *strict*. The above definition says that cores do not have strict endomorphisms. For *finite* cores it clearly holds that every endomorphism is an automorphism, hence our definition is a generalization of the notion of a core for finite structures.

The following proposition illustrates the relevance of cores for constraint satisfaction problems.

Proposition 2. *Let Γ be a relational structure. Then $\text{CSP}(\Gamma)$ can be formulated as a constraint satisfaction problem with a finite template if and only if Γ has a finite core.*

Proof. Clearly, if Γ has a finite core Δ , then $\text{CSP}(\Gamma)$ is equivalent to $\text{CSP}(\Delta)$. Conversely, suppose $\text{CSP}(\Gamma)$ equals $\text{CSP}(\Delta)$ for a finite template Δ . This implies that every finite substructure of Γ homomorphically maps to Δ . A standard compactness argument shows that there is a homomorphism from Γ to Δ , and therefore Γ has a finite core. \square

3. COUNTABLY CATEGORICAL STRUCTURES

Finite structures are up to isomorphism determined by their first-order theory. We can not expect this for infinite structures: by the theorem of Löwenheim-Skolem, every consistent theory with a model of cardinality λ has models of arbitrary cardinality $\geq \lambda$. However, it might still be the case that all models of a certain cardinality are isomorphic. If this is the case for the countably infinite models, we call the theory ω -categorical. A countably infinite structure Γ is called ω -categorical, if its first-order theory $Th(\Gamma)$ (i.e., the set of all first-order sentences that hold in Γ , where the atomic formulas are built from the symbols in τ and equality) is ω -categorical. Throughout the paper we only consider relational and at most countable structures and signatures. Despite the powerful theorems quoted below, the class of ω -categorical structures remains somewhat mysterious, and all classification results require some additional properties (stability in e.g. [29], or homogeneity in [14]).

Let G be a permutation group, and let G act on X . For all $k \geq 1$ there is a natural action of G on the set X^k of k -tuples over X , defined by $(x_1, \dots, x_k)\pi = (x_1\pi, \dots, x_k\pi)$ for permutations π from G . An orbit of k -tuples in G (with respect to this action) is a smallest subset S of X^k such that $\bar{x} \in S$ implies that $\bar{x}\pi \in S$ for all $\pi \in G$. All notions used here are standard and can be found e.g. in [25].

Theorem 3 (Engeler, Ryll-Nardzewski, Svenonius). *The following properties of a countably infinite structure Γ are equivalent:*

- (1) *the structure Γ is ω -categorical;*
- (2) *for each $n \geq 1$, the automorphism group of Γ contains finitely many orbits of n -tuples;*
- (3) *for each $n \geq 1$, Γ admits finitely many inequivalent formulas with n free variables. \square*

Permutation groups with the second property in Theorem 3 are called *oligomorphic* [13]. A famous example of an ω -categorical structure is $(\mathbb{Q}, <)$, the dense linear order of the rational numbers. A famous structure that is not ω -categorical is $(\mathbb{N}, <)$. This structure has an infinite number of orbits of pairs, and hence can not be ω -categorical, because whenever $x_1 - x_2 \neq y_1 - y_2$, then (x_1, x_2) can not be in the same orbit as (y_1, y_2) .

Constraint satisfaction with ω -categorical templates is a strict extension of constraint satisfaction with finite templates. Let Δ be a finite template for a constraint satisfaction problem. To formulate $CSP(\Delta)$ with an ω -categorical template, add for each vertex v in Δ a countably infinite number of copies v_1, v_2, \dots , such that for all $i \geq 1$ the relation $R(\dots, v_i, \dots)$ holds in the resulting structure Γ if and only if $R(\dots, v, \dots)$ holds in Δ . It is not hard to see that the structure Γ is ω -categorical, and that the core of Γ is isomorphic to the core of Δ . Clearly, there are constraint satisfaction problems with ω -categorical templates that can not be formulated with finite templates: all the classes of computational problems mentioned in the introduction contain examples of such problems.

To illustrate the concepts we have seen so far, we formulate several well-known computational problems as constraint satisfaction problems with ω -categorical templates. In all these examples, it is fairly easy to check that the chosen template is a core. We also do not always prove ω -categoricity of these structures, and that they indeed have the specified constraint satisfaction problem. But references with further illustration and proofs will be given. Three more examples follow at the end of Section 3, and for these examples the verification that the structures are indeed cores is more interesting. We do not present these examples here, because we need the concept of amalgamation to define them conveniently, which will be introduced in Section 4.

Betweenness. The hardness of many fragments of Allen’s Interval Algebra [4, 27] can be proven easily by reduction from Betweenness, an important NP-hard problem that can be found in Garey and Johnson [19]. Given a finite set V , and a collection C of ordered triples (x, y, z) of distinct elements from V , the computational question is whether there is an injective function $f : V \rightarrow \{1, \dots, |V|\}$ such that, for each $(a, b, c) \in C$, we have either $f(a) < f(b) < f(c)$ or $f(c) < f(b) < f(a)$. The formulation as a constraint satisfaction problem is straightforward, using for instance the rational numbers as the base set of the template.

Switching-Acyclicity. Given a digraph $D = (V; E)$, can we partition the vertices V into two parts, such that the graph that arises from D by switching all arcs between the two parts is acyclic? To formulate this as a constraint satisfaction problem with an ω -categorical template, consider a dense subset X of \mathbb{Q} , and switch the order $<$ between the elements of X and $\mathbb{Q} - X$, and leave the edges within X and within $\mathbb{Q} - X$ unchanged. The resulting structure is called $S(2)$ and is isomorphic for all choices of dense sets X , see e.g. [14]. The constraint satisfaction problem of $S(2)$ is the problem described above [11]. For equivalent definitions of $S(2)$ and an hardness-proof of its constraint satisfaction problem, see [7, 11].

Partial tree descriptions. Our next example is a computational problem that was studied in computational linguistics [15]. A polynomial time algorithm can be found in [10]. Let D be a digraph with two types of arcs, called *ancestorship* and *non-ancestorship* arcs. The question is whether D is a *consistent partial tree description*, i.e., whether we can find a forest with oriented edges on the vertex set of D , such that for every ancestor arc in D there is a directed path in the forest, and for every non-ancestor arc there is no directed path in the forest. As shown in [11], we can formulate this problem as a constraint satisfaction problem with an ω -categorical template.

Non-cores. Of course, there are plenty of ω -categorical structures that are not cores, for instance the Random graph \mathbf{R} [13, 25], whose core is the complete graph K_ω on a countably infinite set of vertices (the constraint satisfaction problem of \mathbf{R} and K_ω is trivial).

4. HOMOGENEOUS STRUCTURES AND AMALGAMATION CLASSES

We need another concept, which is of a more combinatorial nature, and links ω -categoricity via *homogeneity* and Fraïssé’s theorem to *amalgamation classes*. A structure is *homogeneous* (sometimes also called *ultra-homogeneous* [25]) if every isomorphism between finite substructures of Γ can be extended to an automorphism (in this paper, substructure always means *induced substructure*, as in [25]). A structure Γ admits *quantifier elimination*, if every first-order formula has in Γ a quantifier-free definition.

Proposition 4 (see e.g. 2.22 in [13], and [25]). *An ω -categorical structure has quantifier elimination if and only if it is homogeneous. A countable homogeneous structure Γ is ω -categorical if Γ contains finitely many relations of arity k , for all $k \geq 1$. \square*

For an example of a homogeneous structure that is not ω -categorical, consider the expansion of a countably infinite structure Γ by unary singleton predicates for each element in Γ . This structure is homogeneous, since there are no distinct isomorphic substructures in Γ , and it is not ω -categorical, since the number of orbits in the automorphism group of Γ is infinite.

The next theorem asserts that a countable homogeneous structure is up to isomorphism characterized by its *age*. The *age* of a relational structure Γ is the set of finite structures that embed into Γ (this is terminology that goes back to Fraïssé [18]). A class of finite relational structures \mathcal{C} is an *amalgamation class* if \mathcal{C} is nonempty, closed under isomorphisms and taking substructures, and has the *amalgamation property*, which says that for all $A, B_1, B_2 \in \mathcal{C}$ and embeddings $e_1 : A \rightarrow B_1$ and $e_2 : A \rightarrow B_2$ there exists $C \in \mathcal{C}$ and embeddings $f_1 : B_1 \rightarrow C$ and $f_2 : B_2 \rightarrow C$ such that $f_1 e_1 = f_2 e_2$.

Theorem 5 (Fraïssé [18]). *A countable class \mathcal{C} of finite relational structures with countable signature is the age of a countable homogeneous structure if and only if \mathcal{C} is an amalgamation class. In this case the homogeneous structure is up to isomorphism unique and called the Fraïssé-limit of \mathcal{C} . \square*

The following templates of well-known constraint satisfaction problems are easily defined with amalgamation classes.

Triangle-freeness. Given a graph G , is G triangle-free? Clearly, this problem can be solved in polynomial time. However, it can not be formulated as a constraint satisfaction problem with a *finite* template. To formulate this problem as a constraint satisfaction problem with an ω -categorical template, note that the class of all triangle-free graphs is an amalgamation class. Let us denote its Fraïssé-limit by \triangleleft . It is well-known and not hard to see [13] that this graph is up to isomorphism uniquely determined by the fact that it is triangle-free and has the following extension property: for all finite subsets A, B of vertices of \triangleleft such that all vertices in A are pairwise not adjacent there exists a vertex z in \triangleleft that is not in A , not in B , adjacent to all vertices in A , and not adjacent to all vertices in B . Clearly, $\text{CSP}(\triangleleft)$ is the computational problem described above.

We claim that the structure \triangleleft is a core. Suppose otherwise that there is a strict endomorphism e . If $e(u) = e(v)$, then u and v can not be connected in \triangleleft . We apply the extension property twice to derive that there must be adjacent vertices w and w' in \triangleleft such that w' is connected to u and w is connected to v . But then, $e(u) = e(v), e(w), e(w')$ form a triangle, a contradiction. Hence, since e is strict, there must be non-adjacent vertices a and b such that $e(a)$ is adjacent to $e(b)$. Again by the extension property \triangleleft contains a vertex w that is adjacent to both a and b . But then $e(a), e(w), e(b)$ form a triangle, again a contradiction. Therefore every endomorphism of \triangleleft is an embedding.

No-mono-tri. The structure $[\triangleleft, \triangleleft]$, i.e., the structure that consists of two copies of \triangleleft , where all vertices between the two copies are linked, has an interesting constraint satisfaction problem, which can be formulated as follows: Given a graph, can we partition its vertices into two parts such that both parts do not contain a triangle? This problem is a rather typical example from the class *monotone monadic SNP without inequality (MMSNP)*, a fragment of existential second-order logic introduced in [17] in the context of constraint satisfaction. A general result on so-called *G-free colorability* implies its NP-hardness [1]. Every constraint satisfaction problem in MMSNP can be formulated with an ω -categorical template [9]. The construction given in [9] also shows that the above structure $[\triangleleft, \triangleleft]$ is

ω -categorical. Similarly as in the previous example for the graph \mathcal{A} , it is not hard to show that $[\mathcal{A}, \mathcal{A}]$ is a core.

Quartet compatibility. The next example is an important structure in the theory of infinite permutation groups [13]. A *boron tree* is a finite tree in which all vertices have degree one (*hydrogen atoms*) or degree three (*boron atoms*). On the hydrogen atoms of a boron tree we can define a quaternary relation $xy|uv$ that holds when the paths joining x to y and u to v are disjoint. The class of all structures \mathcal{D} with a quaternary relation that stem from a boron tree as defined above is an amalgamation class [2]. Let D be the Fraïssé-limit of \mathcal{D} . Then $\text{CSP}(D)$ is a well-known NP-hard problem [32] that was independently studied in phylogenetic analysis (without any reference to constraint satisfaction), and is called *quartet-compatibility*: Given a collection C of quartets $xy|uv$ over a set X , is there some tree with leaf set X such that for each quadruple $xy|uv$ in C the paths from x to y and from u to v do not have common vertices?

Rooted triple consistency. The next problem is studied in phylogenetic analysis, again without notice that the problem can be stated as a constraint satisfaction problem. If we fix a point a in the previous structure D and consider the ternary relation ‘:’ defined by $x : yz \Leftrightarrow ax|yz$, we again obtain an ω -categorical structure (this is a *C-set* in [2]). The age of this structure now contains the finite structures T that come from finite *rooted trees*, and the relation $x : yz$ says that the least common ancestor of y and z is strictly below the least common ancestor of x, y , and z in the tree T . The corresponding constraint satisfaction problem is known as the rooted triple consistency problem [32], and tractable. The first polynomial time algorithm for this problem goes back to [3], motivated by a question in database theory.

5. EXISTENTIAL POSITIVE EXPANSIONS

In this section we study various syntactic restrictions of first-order logic. Recall that if relations are added to a given τ -structure Γ then the resulting structure Γ' is called an *expansion* of Γ , and Γ is called the τ -*reduct* of Γ' .

A first-order formula ϕ is *primitive (primitive positive)*, if it is of the form

$$\exists \bar{x}. \psi_1 \wedge \cdots \wedge \psi_k$$

where ψ_i are literals (atomic formulas) that might include the equality relation. It is called *existential (existential positive)*, if it is of the form $\exists \bar{x}. \Psi$ where Ψ is quantifier-free (and negation-free). The strongest of these four syntactic restrictions, *primitive positivity*, is important for constraint satisfaction, since the expansion of a template with primitive positive definable (short, *pp-definable*) relations does not change the complexity of the corresponding constraint satisfaction problem. This is an easy observation, see e.g. [26].

A mapping of a τ -structure Γ to a τ -structure Δ is called *elementary*, if it preserves all first-order τ -formulas (this is a standard notion in model-theory [25]). A theory T is called *model-complete*, if all embeddings between models of T are elementary. In the case that T is the theory of a structure Γ , we also say that Γ is model-complete, as usual. If T is a first-order theory, then Γ is called an *existentially complete* model of T if for every existential formula ϕ , every tuple \bar{a} in Γ , and every model Δ of T such that Γ is a substructure of Δ and $\phi(\bar{a})$ holds in Δ , the formula $\phi(\bar{a})$ also holds in Γ .

Model-completeness can be characterized in many different ways. We say that two formulas ϕ, ψ are *equivalent with respect to* a theory T if in every model of T the formulas ϕ and ψ define the same relations.

Proposition 6 (Theorem 7.3.1 in [25]). *Let T be a first-order theory. Then the following are equivalent.*

- T is model-complete
- Every model of T is an existentially complete model of T
- Every first-order formula is equivalent to an existential formula with respect to T . \square

An ω -categorical structure is model-complete if and only if its first-order theory is equivalent to a set of $\forall\exists$ -sentences, i.e., sentences of the form $\forall\bar{x}\exists\bar{y}\phi(\bar{x},\bar{y})$, where ϕ is quantifier-free (see e.g. Theorem 7.3.3 and 7.3.4 in [25]). The following proposition follows directly from Theorem 7.2.1 in [25].

Proposition 7. *Let Γ be a relational structure, and let T be the set of all $\forall\exists$ -sentences that hold in Γ . Then there exists an extension of Γ that is an existentially complete model of T . \square*

The case that all existential formulas are equivalent to existential positive formulas can be characterized in a different way.

Lemma 8. *Let T be a first-order theory such that every homomorphism between models of T is an embedding. Then every existential formula is equivalent to an existential positive formula with respect to T .*

Proof. A formula ϕ is equivalent to an existential positive formula with respect to a theory T if ϕ is preserved by all homomorphisms¹ between models of T ; this fact is stated as Exercise 2 in Section 5.5 in [25]. Let f be a homomorphism between two models of T . By assumption, f is an embedding, and therefore clearly preserves all existential formulas. Hence, all existential formulas are equivalent to an existential positive formula with respect to T . \square

The following lemma will be useful to construct homogeneous models.

Lemma 9. *If Γ is a structure that has been expanded by all primitive definable relations, then there is a homogeneous structure with the same age as Γ .*

Proof. Let $\bar{a} = (a_1, \dots, a_k)$ be a tuple of elements from Γ , let B_1, B_2 be finite induced substructures of Γ and $e_1 : \bar{a} \rightarrow B_1$ and $e_2 : \bar{a} \rightarrow B_2$ be embeddings. Since there are relation symbols for every primitive formula in the signature, there is a relation R_1 that holds on the tuple $e_1(\bar{a})$ in Γ , and is defined by the following primitive formula ϕ . Let b_1, \dots, b_m be the elements of B_1 . The formula ϕ has k free variables x_1, \dots, x_k and has the form $\phi := \exists x_{k+1}, \dots, x_m \psi$. The formula ψ is a conjunction of literals defined as follows. If there is an l -ary relation symbol R in the signature of Γ such that R holds on elements b_{i_1}, \dots, b_{i_l} in B , $1 \leq i_1, \dots, i_l \leq m$, then ψ contains the conjunct $R(x_{i_1}, \dots, x_{i_l})$. If the relation R does not hold on b_{i_1}, \dots, b_{i_l} in B , then ψ contains the conjunct $\neg R(x_{i_1}, \dots, x_{i_l})$. Moreover, ψ contains conjuncts of the form $x_i \neq x_j$ for all distinct indices i, j from $1, \dots, m$.

¹We say that a formula ϕ is *preserved by all homomorphisms between models of T* if whenever h is a homomorphism from M_1 to M_2 that maps a tuple \bar{a} of elements from M_1 pointwise to a tuple \bar{b} of elements from M_2 , where M_1 and M_2 are models of T , and \bar{a} satisfies ϕ in M_1 , then \bar{b} satisfies ϕ in M_2 .

We also have a relation R_2 corresponding in an analogous way to B_2 where the points from $B_2 - e_2(\bar{a})$ are existentially quantified, and which holds on $e_2(\bar{a})$ in Γ . Since e_1 and e_2 are embeddings, these relations also hold on \bar{a} in Γ . They assert that we can find an extension C of the structure induced by \bar{a} and embeddings $f_1 : B_1 \rightarrow C$, $f_2 : B_2 \rightarrow C$ such that $f_1 e_1 = f_2 e_2$. Thus, the age of Γ has the amalgamation property, and Theorem 5 implies that there is a homogeneous structure with the same age. \square

We can combine Lemma 8 and Lemma 9 and obtain the following.

Corollary 10. *Let Γ be a structure that has been expanded by all existential positive definable relations. If all homomorphisms between models of $\text{Th}(\Gamma)$ are embeddings, then there is a homogeneous structure with the same age as Γ .*

Proof. Proposition 8 shows that every existential, and in particular every primitive formula is in Γ equivalent to an existential positive formula. Hence, Γ is also expanded by all primitive definable relations, and Proposition 9 shows that there is a homogeneous structure that has the same age as Γ . \square

6. CORES OF COUNTABLY CATEGORICAL STRUCTURES

In this section we state and prove the main results of the paper. We start with a proposition on the existence of a ‘youngest’ endomorphic image of an ω -categorical structure. The proof employs a typical technique for ω -categorical structures.

Proposition 11. *Let Γ be an ω -categorical relational τ -structure. Then there exists an endomorphism c of Γ such that for every other endomorphism g , all finite substructures of $c(\Gamma)$ embed into $g(\Gamma)$. This is, there exists an endomorphic image of Γ of smallest age.*

Proof. Let \mathcal{S} be the set of all finite τ -structures S such that there is an endomorphism g of Γ so that S does not embed into $g(\Gamma)$. We have to show that there is an endomorphism c such that no structure from \mathcal{S} embeds into $c(\Gamma)$. For the construction of c we consider the following tree. Let a_1, a_2, \dots be an enumeration of Γ . The vertices on level n of the tree are labeled with equivalence classes of good homomorphisms from the structure induced by $\{a_1, \dots, a_n\}$ to Γ . A homomorphism h is good, if no structure from \mathcal{S} embeds into the structure induced by $h(\{a_1, \dots, a_n\})$ in Γ . Two homomorphisms g_1 and g_2 are equivalent, if there exists an automorphism α of Γ such that $g_1 = g_2 \alpha$. Clearly, if a homomorphism is good, then all equivalent homomorphisms and all restrictions are also good. A vertex u on level $n + 1$ in the tree is connected to a vertex v on level n , if some homomorphism from u is the restriction of some homomorphism from v . Because of ω -categoricity, the tree is finitely branching. We want to show that the tree has vertices on each level n , and iteratively construct a sequence h_1, h_2, \dots, h_k of homomorphisms from $\{a_1, \dots, a_n\}$ to Γ , where the last endomorphism h_k induces a good homomorphism. Initially, if no structure from \mathcal{S} embeds into the structure induced by $\{a_1, \dots, a_n\}$, we can choose the identity as a good homomorphism. Otherwise, there is a structure $S \in \mathcal{S}$ that embeds into the structure induced on $\{a_1, \dots, a_n\}$, and an endomorphism e such that $e(\Gamma)$ does not contain S . The mapping h_1 restricted to $\{a_1, \dots, a_n\}$ is a strict homomorphism, because if it was an embedding, S embeds into the structure induced by the image of $\{a_1, \dots, a_n\}$ under this mapping, which is by assumption impossible.

In step i , if no structure in \mathcal{S} embeds into the structure induced by $h_i(\{a_1, \dots, a_n\})$, then h_i is a good homomorphism, and we are again done. Otherwise there is an endomorphism e of Γ and a structure $S \in \mathcal{S}$ that embeds into the structure induced by $h_i(\{a_1, \dots, a_n\})$, such that S does not embed into $e(\Gamma)$. We can then define a homomorphism $h_{i+1} : \{a_1, \dots, a_n\} \rightarrow \Gamma$ by $h_{i+1}(x) := e(h_i(x))$, which is again a strict homomorphism. Since in the sequence of structures induced by $h_1(\{a_1, \dots, a_n\})$, $h_2(\{a_1, \dots, a_n\})$, \dots either the number of vertices decreases or the number of tuples in relations increases, and since Γ is ω -categorical, the sequence has to be finite. Hence, there exists a good homomorphism from $\{a_1, \dots, a_n\}$ to Γ , for all $n \geq 0$. By König's tree lemma, there exists an infinite path in the tree. Since adjacency in the tree was defined by restriction between homomorphisms, this path defines an endomorphism c of Γ . By construction, no structure in \mathcal{S} embeds into $c(\Gamma)$. \square

In the following, $c(\Gamma)$ denotes the structure induced by the image of the endomorphism c in Γ that was constructed in Proposition 11. Note that Proposition 11 says that all cores of Γ have the same age as $c(\Gamma)$. We will use the following well-known lemma several times (see e.g. Section 2.6 in [13]). It can be shown by a similar application of König's lemma as in the proof of the previous proposition.

Lemma 12. *Let Γ be a relational structure whose age is contained in the age of an ω -categorical structure Δ . Then Γ embeds into Δ .* \square

The following lemma is a central step in our arguments.

Lemma 13. *Let Γ be an ω -categorical τ -structure that contains all existential positive definable relations. Then every homomorphism between two structures Γ_1 and Γ_2 of the same age as $c(\Gamma)$ is an embedding.*

Proof. Since Γ is ω -categorical and because the age of Γ_1 and Γ_2 is contained in the age of Γ , Lemma 12 implies that both Γ_1 and Γ_2 embed into Γ . For simplicity of notation, we assume that Γ_1 and Γ_2 are substructures of Γ . Now suppose for contradiction that $f : \Gamma_1 \rightarrow \Gamma_2$ is a homomorphism that is not an embedding, this is, for some tuple $\bar{u} = (u_1, \dots, u_k)$ of elements and some k -ary relation R in Γ_1 the mapping f does not preserve the formula $\neg R(\bar{u})$ or does not preserve the formula $u_1 \neq u_2$. We will then construct an endomorphism h of Γ such that $h(\Gamma)$ does not contain a copy of the substructure S induced by \bar{u} in Γ_1 . This is a contradiction: On the one hand S is a substructure of $c(\Gamma)$, since Γ_1 has the same theory and thus the same age as $c(\Gamma)$. On the other hand, since S is not a substructure of $h(\Gamma)$, Proposition 11 says that S is not a substructure of $c(\Gamma)$.

To construct this homomorphism h we consider an infinite but finitely branching tree. The vertices on level n in this tree will be labeled by equivalence classes of *good* homomorphisms from $\{a_1, a_2, \dots, a_n\}$ to Γ , where a_1, a_2, \dots is an enumeration of Γ . A homomorphism g on level n is *good*, if the structure induced by $g(\{a_1, \dots, a_n\})$ does not contain an induced copy of S . Two homomorphisms g_1 and g_2 are equivalent if there exists an automorphism α of Γ such that $g_1 = g_2\alpha$. Adjacency is defined by restriction; this is, two nodes on level n and $n+1$ are adjacent in the tree if there are representatives g_1 and g_2 from these nodes such that g_1 is a restriction of g_2 . Clearly, all restrictions of a good homomorphism are again good homomorphisms, and all homomorphisms in an equivalence class are good, or all are not good. By ω -categoricity of Γ , the tree has only a finite number of vertices on level n , and in particular it follows that the tree is finitely branching. The crucial step is that the tree contains vertices on every level, i.e., there exists a good homomorphism

$h_n : \{a_1, \dots, a_n\} \rightarrow \Gamma$ for each $n \geq 1$. We show this in the following; and here we use the assumption that Γ contains all existential positive definable relations.

To find h_n for each $n \geq 1$, we consider a sequence $(h_n^i)_{i \geq 0}$ of homomorphisms to Γ , where the domain of h_n^0 is $\{a_1, \dots, a_n\}$, and the domain of h_n^{i+1} equals the image of h_n^i . Hence, we can define the following composed homomorphism $h_n^{(i)} : \{a_1, \dots, a_n\} \rightarrow \Gamma$ by $h_n^{(i)}(x) := h_n^i(\dots h_n^1(h_n^0(x)) \dots)$. We now define the sequence $(h_n^i)_{i \geq 0}$. The homomorphism h_n^0 is the identity. For $i > 0$, if the structure induced by the domain of h_n^i does not contain an induced copy of S , we are done, because then $h_n^{(i-1)}$ is a good homomorphism from $\{a_1, \dots, a_n\}$ to Γ . Otherwise, there are elements (b_1^i, \dots, b_k^i) in the domain of h_n^i that induce in Γ a structure isomorphic to S . We now define $h_n^i(b_j^i) := f(u_j)$ for $1 \leq j \leq k$, and want to extend this mapping to a (strict) homomorphism h_n^i on the other elements b_{k+1}^i, \dots, b_m^i , $m \leq n$, in the domain of h_n^i . Consider the formula $\phi := \exists x_{k+1}^i, \dots, x_m^i \psi$ with free variables x_1^i, \dots, x_k^i , where ψ is a conjunction of atomic formulas defined as follows. We use the structure induced by b_1^i, \dots, b_m^i in Γ to define ψ . For all $j_1, \dots, j_l \leq m$, the formula ψ contains a conjunct $R(x_{j_1}^i, \dots, x_{j_l}^i)$ iff R holds on $b_{j_1}^i, \dots, b_{j_l}^i$ in Γ . It is then clear that the formula ϕ holds for b_1^i, \dots, b_k^i .

Since Γ contains all existential positive definable relations, the existential positive formula ϕ also holds on u_1, \dots, u_k , since these vertices induce the same structure as b_1^i, \dots, b_k^i . Since f preserves existential positive formulas, ϕ also holds on $f(u_1), \dots, f(u_k)$ in Γ . We thus can find witnesses r_{k+1}, \dots, r_m in Γ for the variables x_{k+1}^i, \dots, x_m^i of the existential quantifiers in ϕ , and extend h_n^i by $h_n^i(b_j^i) := r_j$ for $k+1 \leq j \leq m$. Then h_n^i clearly is a homomorphism to Γ , which is also strict, because it does not preserve some of the inequalities or negated relations that hold on b_1^i, \dots, b_k^i . Therefore the sequence $(h_n^i)_i$ of homomorphisms must be finite, because ω -categoricity of Γ implies that there are only finitely many non-isomorphic homomorphic images of the structure induced by $\{a_1, \dots, a_n\}$ in Γ . Let $h_n^{(i_0)}$ be the last homomorphism in this sequence. By construction, this mapping is a good homomorphism h_n for every $n \geq k$.

Therefore, the constructed tree contains vertices on all levels, and König's tree lemma asserts that the tree contains an infinite path. Since adjacency is defined by restriction, this path defines an infinite endomorphism h of Γ . The image $h(\Gamma)$ does not contain an induced copy of S . This contradicts the minimality property of $c(\Gamma)$ formulated in Proposition 11. Hence, every homomorphism from Γ_1 to Γ_2 is an embedding. \square

Corollary 14. *Every ω -categorical τ -structure Γ has a core.*

Proof. Let Δ be the expansion of Γ by all existential positive definable relations, and let $c(\Delta)$ be the structure induced by the endomorphism c constructed in Proposition 11 for Δ . Lemma 13 shows in particular that endomorphisms of $c(\Delta)$ are cores. Let Γ_0 be the τ -reduct of $c(\Delta)$. Clearly, since $c(\Delta)$ is a core, Γ_0 must be a core as well. Because c is also an endomorphism of Γ , we have that Γ_0 is a core of Γ . \square

We already mentioned that the core of a finite structure is unique up to isomorphism. As we will see now, an ω -categorical structure might have non-isomorphic cores.

Example. Consider the following ω -categorical structure Γ . Let \mathbb{Q} be the set of rational numbers, and let X be a disjoint countably infinite set. The domain of Γ is $\mathbb{Q} \cup X$, and the signature contains two binary relation symbols $<$ and \neq , and one unary relation symbol P .

For two elements x, y in Γ the relation $x < y$ holds iff $x, y \in \mathbb{Q}$ and x is a strictly smaller number than y . The relation $x \neq y$ holds iff x and y are distinct elements in Γ . Finally, $P(x)$ holds iff $x > 0$ holds. It is easy to verify that Γ is ω -categorical. One way to see this is by Proposition 4, because the expansion of Γ by the unary relation $N := \{x \in \mathbb{Q} \mid x \leq 0\}$ is homogeneous.

The structure Γ has a model-complete core, which is unique up to isomorphism, namely the structure $\Delta := (\mathbb{Q}^+, <, \neq, P)$, where \mathbb{Q}^+ is the set of all positive rational numbers, where the interpretation of $<$ and \neq is as before, and where P denotes the trivial unary relation that contains all elements of the domain. The structure Δ is indeed the image of an endomorphism c of Γ , for instance the endomorphism that maps \mathbb{Q} in an order-preserving way to the set $\{x \in \mathbb{Q} \mid x > 3\}$, and maps X bijectively onto $\{x \in \mathbb{Q} \mid 0 < x \leq 3\}$ in an arbitrary way. Let Δ be the structure that is induced by the image \mathbb{Q}^+ of this endomorphism. It is easy to verify that Δ is a core, and hence Δ is a core of Γ . It will follow from Theorem 16 below that all other model-complete cores of Γ are isomorphic to Δ .

However, we claim that $(V, <, \neq, P)$ has other cores that are not isomorphic to Δ , and not model-complete. For instance, let Γ' be the substructure of Γ induced by $\mathbb{Q}^+ \setminus (1, 2)$, i.e., the set of positive rational numbers without the open interval $(1, 2)$. Then Γ' is the image of an endomorphism c' constructed similarly as above, with the only difference that c' maps X bijectively onto $\{x \in \mathbb{Q} \mid 0 < x \leq 1 \text{ or } 2 \leq x \leq 3\}$. It can be verified easily that the structure Γ' is a core. To see that Γ' is not model-complete, consider the restriction d of the mapping c' to the elements of Γ' . Since Γ' is a core, the mapping d is an embedding of Γ' into Γ' . However, it is not an elementary embedding, since the formula $\exists z. x < z < y$ holds for $x = d(1)$ and $y = d(2)$ in Γ' , but does not hold for $x = d(1)$ and $y = d(2)$ in the structure induced by $d(\Gamma')$.

We will show that every ω -categorical structure Γ is always homomorphically equivalent to a *model-complete* core, which is unique up to isomorphism. For that, we first prove a stronger result for structures that are expanded by all existential positive definable relations.

Proposition 15. *Let Γ be an ω -categorical τ -structure. If Γ contains all existential positive definable relations, then Γ is homomorphically equivalent to a homogeneous core Γ^c , which is unique up to isomorphism. Moreover, Γ^c is finite or ω -categorical.*

Proof. Let T be the set of all universal sentences that are true in the structure $c(\Gamma)$ constructed in Proposition 11. Proposition 7 shows that there is an extension Δ of $c(\Gamma)$ that is an existentially complete model of T . Since Δ satisfies the same universal sentences as $c(\Gamma)$, the two structures have the same age. Because Γ is ω -categorical, and because the age of Δ is contained in the age of Γ , we can apply Lemma 12 to show that Δ embeds into Γ . Assume for simplicity of notation that Δ is a substructure of Γ .

Let ϕ be an existential formula, and let \bar{t} be a tuple from Δ . We claim that ϕ holds on \bar{t} in Γ if and only if ϕ holds on \bar{t} in Δ . One direction is clear, because Δ embeds into Γ . Now, consider the restriction d of c to Δ . Because $d(\Delta)$ is a substructure of $c(\Gamma)$, and $c(\Gamma)$ is a substructure of Δ , d is an endomorphism of Δ . Because Δ and $c(\Gamma)$ have the same age, Lemma 13 implies that d is an embedding. Hence, the structure induced by $d(\Delta)$ is isomorphic to Δ . The embedding d preserves the existential formula ϕ , and therefore ϕ holds on $d(\bar{t})$ in Δ . Because Δ and thus also $d(\Delta)$ are existentially complete models of

T , the formula ϕ also holds on $d(\bar{t})$ in the structure $d(\Delta)$ (not only in Δ). Since d is an embedding, we then know that ϕ holds on \bar{t} in Δ .

Therefore, because the structure Γ is expanded by all existential positive relations, the structure Δ is also expanded by all existential positive relations. All models of $\text{Th}(\Delta)$ must have the same age as Δ and $c(\Gamma)$, and hence Lemma 13 shows that every homomorphism between two models of $\text{Th}(\Delta)$ is an embedding. By Corollary 10 there is a homogeneous structure Γ^c with the same age as Δ and $c(\Gamma)$. Theorem 5 shows that this structure is unique up to isomorphism. The structure Γ^c might be finite. If Γ^c is infinite, and since Γ^c contains only finitely many k -ary relations for each $k \geq 1$, Proposition 4 shows that Γ^c is ω -categorical.

To conclude the proof, we have to show that Γ^c is homomorphically equivalent to Γ . But this is clear since Γ homomorphically maps to $c(\Gamma)$, which embeds into Γ^c by Lemma 12. Lemma 12 also shows that Γ^c embeds into Γ , because it has a smaller age than Γ and because Γ is ω -categorical. \square

We can now prove one of the main results.

Theorem 16. *Every ω -categorical τ -structure Γ is homomorphically equivalent to a model-complete core Γ^c , which is unique up to isomorphism. The core Γ^c is ω -categorical or finite. All orbits of k -tuples in Γ^c are pp-definable in Γ^c .*

Proof. Let Δ be the expansion of Γ by all existential positive definable relations. By the previous Proposition, Δ has a homogeneous core Δ^c , which is finite or ω -categorical. By Proposition 4, the structure Δ^c has quantifier elimination, and therefore every first-order formula is equivalent to a quantifier-free formula (in the expanded signature of Δ and Δ^c). Moreover, Δ^c is ω -categorical. Let Γ^c be the τ -reduct of Δ^c . It is well-known (and follows from Theorem 3; see [25]) that Γ^c is ω -categorical as well.

We want to show that Γ^c is a model-complete core. Suppose that e is an endomorphism of Γ^c . We have to show that e is an elementary embedding. Let ϕ be a first-order formula. Then ϕ is in Δ^c equivalent to a quantifier-free formula ψ . Note that e preserves existential positive formulas, and because Δ^c is a core, e preserves ψ as well. Therefore, e is elementary.

Now suppose that Γ' is some model-complete core of Γ . We show that Γ' is isomorphic to Γ^c . In fact, we will show that Δ^c and the expansion Δ' of Γ' by all existential positive definable relations have the same age and are both homogeneous. Then Theorem 5 implies that Δ^c and Δ' , and hence also Γ^c and Γ' are isomorphic.

Similarly as in the proof of Proposition 15 we prove that every existential formula that holds on tuples from Δ' in Δ also hold in Δ' . This also shows that Δ' is ω -categorical. To show that Δ' is homogeneous, by Proposition 4 it suffices to show that Δ' has quantifier elimination. Let ϕ be a first-order formula. Since Γ' is model-complete, Proposition 6 shows that ϕ is in Γ' and therefore also in Δ' equivalent to an existential formula. Since Δ' is ω -categorical and a core, we can then apply Lemma 8 to show that ϕ is equivalent to an existential positive formula in Δ' . Since Δ' contains all existential positive formulas, we have shown that Δ' admits quantifier elimination.

Finally, let R be an orbit R of k -tuples in the automorphism group of Γ^c . All k -tuples in R induce isomorphic substructures S in Δ^c . Because Δ^c is homogeneous, all k -tuples in Δ^c that are isomorphic to S are contained in R . Thus, R has a definition as a conjunction φ of atomic formulas. We replace all relation symbols in φ that are contained in the signature of Δ^c but not in the signature of Γ^c by their existential positive definition. Then the resulting formula can be re-written as a disjunction of primitive positive formulas. Every disjunct is

either false or already defines R , because R is an orbit of k -tuples in $\text{Aut}(\Gamma^c)$. Hence, the formula is equivalent to a single disjunct, and R is pp-definable in Γ^c . \square

Corollary 17. *If Γ is an ω -categorical structure that is expanded by all pp-definable relations. Then Γ is homomorphically equivalent to a homogeneous core Γ^c (which is again unique up to isomorphism).*

Proof. Let Γ^c be the ω -categorical model-complete core that is homomorphically equivalent to Γ . Theorem 16 shows that every orbit of k -tuples in the automorphism group of Γ^c has a pp-definition in Γ^c , and the proof of Theorem 16 shows that because Γ is expanded by all pp-definable relations, Γ^c contains all pp-definable relations as well. We have also seen that in Γ^c , every first-order definable relation has a quantifier-free definition in the expanded structure that contains all existential positive formulas. Every existential positive formula is a disjunction of primitive positive formulas, and therefore the structure Γ^c has quantifier-elimination as well. Proposition 4 implies that Γ^c is homogeneous. \square

We also obtain alternative characterizations of when a model-complete ω -categorical structure is a core. A set of functions F from Γ to Γ *locally generates* a function g , if for every finite subset A of Γ there is a function $f \in F$ such that $g(a) = f(a)$ for all $a \in A$ (this is a standard notion in universal algebra [33]).

Theorem 18. *Let Γ be a model-complete ω -categorical structure. Then the following are equivalent.*

- (1) Γ is a core.
- (2) Every first-order formula is in Γ equivalent to an existential positive formula.
- (3) Every endomorphism of Γ is an elementary embedding.
- (4) The automorphism group of Γ locally generates the endomorphism monoid of Γ .

Proof. Every first-order definable k -ary relation R in Γ is the union of a finite number of orbits of k -tuples of $\text{Aut}(\Gamma)$. Theorem 16 shows that each of these orbits has a pp-definition. A disjunction of primitive positive formulas can be equivalently written as an existential positive formula, and we obtain an existential positive definition for R , showing that 1 implies 2.

Since endomorphisms clearly preserve existential positive formulas, 2 implies that every endomorphism is elementary.

A function from Γ to Γ is in the local closure of the automorphism group of Γ if and only if it preserves all first-order definable relations. If we assume that every endomorphism is elementary, then the automorphism group locally generates all endomorphisms of Γ , and therefore 3 implies 4.

To prove that 4 implies 1, assume that Γ has a strict endomorphism f , i.e., there is a tuple (u_1, \dots, u_k) in Γ such that $f(u_1) = f(u_2)$, or $R(f(u_1), \dots, f(u_k))$ and not $R(u_1, \dots, u_k)$. Clearly, such a function f cannot be locally generated by automorphisms. \square

7. ADDING CONSTANTS TO THE SIGNATURE

In relational structures Γ we can use singleton relations, i.e., relations of the form $\{c\}$ for an element c from the domain of Γ , to model the concept of *constants* in first-order structures [25]. One of the main results in [12] says that if Γ is a finite core, then adding a singleton-relation does not increase the complexity of the corresponding constraint

satisfaction problem. We show that the same holds for constraint satisfaction problems where the template is an ω -categorical model-complete core. Note that this directly applies to all the computational problems presented in the introduction and Sections 3 and 4.

Theorem 19. *Let Γ be a model-complete ω -categorical core, and let Γ' be the expansion of Γ by a unary singleton relation $C = \{c\}$. If $\text{CSP}(\Gamma)$ is tractable, then so is $\text{CSP}(\Gamma')$. (If $\text{CSP}(\Gamma')$ is NP-hard, then so is $\text{CSP}(\Gamma)$.)*

Proof. We show how to solve $\text{CSP}(\Gamma')$ in polynomial time, under the assumption that $\text{CSP}(\Gamma)$ can be solved in polynomial time. Let S' be an instance of $\text{CSP}(\Gamma')$. Let P be the orbit of c in the automorphism group of Γ . By Theorem 16, P is primitive positive definable in Γ . Thus we can assume without loss of generality that Γ and Γ' contain the relation P . Replace all occurrences of the relation C in S' by the relation P . Solve the resulting instance S of $\text{CSP}(\Gamma)$; by assumption this is possible in polynomial time. If S is not satisfiable, then in particular S' was not satisfiable. On the other hand, if there is a homomorphism h from S to Γ , we claim that there is a homomorphism from S' to Γ' . Since P is the orbit of the element c , there is an automorphism α of Γ such that $h\alpha$ is a solution of the instance S' of $\text{CSP}(\Gamma')$. \square

If an ω -categorical model-complete core Γ is expanded by a singleton relation, the resulting structure Γ' is again an ω -categorical model-complete core. The fact that ω -categoricity is preserved by such expansions is well-known, see e.g. [25]. Suppose that Γ' is not a model-complete core, i.e., there is a non-elementary endomorphism e of Γ' . Then e is also a non-elementary endomorphism of Γ , a contradiction. Hence, we can apply Theorem 19 several times, and obtain that Theorem 19 also holds for expansions by a *finite* number of singleton relations.

8. ACKNOWLEDGEMENTS

I am grateful to Julia Böttcher, Peter Cameron, Gregory Cherlin, Daniel Král, Martin Ziegler for discussions. Special thanks to anonymous referees, who greatly helped with their comments.

REFERENCES

- [1] D. Achlioptas. The complexity of G -free colourability. *Discrete Mathematics*, 165:21–30, 1997.
- [2] S. Adeleke and P. M. Neumann. Structure of partially ordered sets with transitive automorphism groups. *AMS Memoir*, 57(334), 1985.
- [3] A. Aho, Y. Sagiv, T. Szymanski, and J. Ullman. Inferring a tree from lowest common ancestors with an application to the optimization of relational expressions. *SIAM Journal on Computing*, 10(3):405–421, 1981.
- [4] J. F. Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843, 1983.
- [5] B. Bauslaugh. Core-like properties of infinite graphs and structures. *Disc. Math.*, 138(1):101–111, 1995.
- [6] B. Bauslaugh. Cores and compactness of infinite directed graphs. *Journal of Combinatorial Theory, Series B*, 68(2):255–276, 1996.
- [7] M. Bodirsky. Constraint satisfaction with infinite domains. Dissertation an der Humboldt-Universität zu Berlin, 2004.
- [8] M. Bodirsky. The core of a countably categorical structure. In *Proceedings of the 22nd Annual Symposium on Theoretical Aspects of Computer Science (STACS'05)*, LNCS 3404, pages 100–110, Springer-Verlag, 2005.

- [9] M. Bodirsky and V. Dalmau. Datalog and constraint satisfaction with infinite templates. In *Proceedings of STACS'06*, pages 646–659. Springer Verlag, 2006.
- [10] M. Bodirsky and M. Kutz. Pure dominance constraints. In *Proceedings of STACS'02*, pages 287–298. Springer Verlag, 2002.
- [11] M. Bodirsky and J. Nešetřil. Constraint satisfaction with countable homogeneous templates. In *Proceedings of CSL'03*, pages 44–57, Vienna, 2003.
- [12] A. Bulatov, A. Krokhin, and P. G. Jeavons. Classifying the complexity of constraints using finite algebras. *SIAM Journal on Computing*, 34:720–742, 2005.
- [13] P. J. Cameron. *Oligomorphic Permutation Groups*. Cambridge Univ. Press, 1990.
- [14] G. Cherlin. The classification of countable homogeneous directed graphs and countable homogeneous n -tournaments. *AMS Memoir*, 131(621), January 1998.
- [15] T. Cornell. On determining the consistency of partial descriptions of trees. In *Proceedings of the ACL*, pages 163–170, 1994.
- [16] I. Düntsch. Relation algebras and their application in temporal and spatial reasoning. *Artificial Intelligence Review*, 23:315–357, 2005.
- [17] T. Feder and M. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through Datalog and group theory. *SIAM Journal on Computing*, 28:57–104, 1999.
- [18] R. Fraïssé. *Theory of Relations*. North-Holland, 1986.
- [19] M. Garey and D. Johnson. *A guide to NP-completeness*. CSLI Press, 1978.
- [20] D. Gusfield. *Algorithms on strings, trees, and sequences. Computer Science and Computational Biology*. Cambridge University Press, New York, 1997.
- [21] P. Hell and J. Nešetřil. On the complexity of H-coloring. *Journal of Combinatorial Theory, Series B*, 48:92–110, 1990.
- [22] P. Hell and J. Nešetřil. *Graphs and Homomorphisms*. Oxford University Press, 2004.
- [23] R. Hirsch. Relation algebras of intervals. *Artificial Intelligence Journal*, 83:1–29, 1996.
- [24] R. Hirsch. Expressive power and complexity in algebraic logic. *Journal of Logic and Computation*, 7(3):309 – 351, 1997.
- [25] W. Hodges. *A shorter model theory*. Cambridge University Press, 1997.
- [26] P. Jeavons, D. Cohen, and M. Gyssens. Closure properties of constraints. *JACM*, 44(4):527–548, 1997.
- [27] P. Jeavons, P. Jonsson, and A. A. Krokhin. Reasoning about temporal relations: The tractable subalgebras of Allen’s interval algebra. *JACM*, 50(5):591–640, 2003.
- [28] D. W. Kueker. Core structures for theories. *Fundamenta Mathematicae*, 89.2:155–171, 1975.
- [29] A. H. Lachlan. Stable finitely homogeneous structures: A survey. In *Algebraic Model Theory, NATO ASI Series*, volume 496, pages 145–159, 1996.
- [30] P. B. Ladkin and R. D. Maddux. On binary constraint problems. *Journal of the Association for Computing Machinery*, 41(3):435–469, 1994.
- [31] B. Nebel and H.-J. Bürckert. Reasoning about temporal relations: A maximal tractable subclass of Allen’s interval algebra. *JACM*, 42(1):43–66, 1995.
- [32] M. Steel. The complexity of reconstructing trees from qualitative characters and subtrees. *Journal of Classification*, 9:91–116, 1992.
- [33] A. Szendrei. *Clones in universal Algebra*. Seminaire de mathematiques superieures. Les Presses de L’Universite de Montreal, 1986.