Constraint Satisfaction with Countable Homogeneous Templates

Manuel Bodirsky[1] and Jaroslav Nešetřil[2]

[1] Humboldt Universität zu Berlin, Germany
bodirsky@informatik.hu-berlin.de
[2] Institute for Theoretical Computer Science (ITI)
and Department of Applied Mathematics (KAM),
Charles University, Prague, Czech Republic
nesetril@kam.mff.cuni.cz

Abstract. For a fixed countable homogeneous relational structure Γ we study the computational problem whether a given finite structure of the same signature homomorphically maps to Γ. This problem is known as the *constraint satisfaction problem* CSP(Γ) for the *template* Γ and was intensively studied for finite Γ. We show that – as in the case of finite Γ – the computational complexity of CSP(Γ) for countable homogeneous Γ is determinded by the clone of polymorphisms of Γ. To this end we prove the following theorem, which is of independent interest: The primitive positive definable relations over an ω-categorical structure Γ are precisely the relations that are preserved by the polymorphisms of Γ. If the age of Γ is given by a finite number of finite forbidden induced substructures, then CSP(Γ) is in NP. We use a classification result by Cherlin and prove that in this case every constraint satisfaction problem for a countable homogeneous digraph is either tractable or NP-complete.
Keywords: complexity of constraint satisfaction, homogeneous digraphs, graph homomorphisms, ω-categorical structures, polymorphism preservation theorem

1 Introduction

For a fixed relational structure Γ (called the *template*), the constraint satisfaction problem CSP(Γ) is the following computational problem: Given

[2] Supported by grants LN00A056 and 1M0021620808 of the Czech Ministry of Education.

a finite structure S of the same signature as Γ, is there a homomorphism from S to Γ?

Constraint satisfaction problems frequently occur in theoretical computer science, and have attracted much attention for finite templates Γ. It is conjectured that $\mathrm{CSP}(\Gamma)$ has a *dichotomy* in the sense that every constraint satisfaction problem $\mathrm{CSP}(\Gamma)$ for finite structure Γ is either tractable or NP-complete. This is true for templates that are undirected graphs [25], for two element templates [39], and three element templates [8]. It is known that every constraint satisfaction problem with a finite template is polynomial time equivalent to a digraph homomorphism problem [20]. There are powerful classes of algorithms solving the known tractable constraint satisfaction problems, namely group theoretic algorithms and local-consistency based algorithms [17, 20, 27, 30].

But many constraint satisfaction problems in the literature can not be formulated as a constraint satisfaction problem with a finite template. One example is Allen's interval algebra [1] that has applications in temporal reasoning in artificial intelligence. The classification of the tractable and hard subalgebras of Allen's algebra was completed only recently [31, 35], and they also exhibit a complexity dichotomy. Other examples are tree description languages that were introduced in computational linguistics [3, 4, 15], and problems from phylogenetic analysis [7, 36, 41]. Even digraph acyclicity can not be formulated as a constraint satisfaction problem with a finite template.

It was already remarked in [20] that arbitrary infinite templates allow to describe all queries that are closed under disjoint unions and whose inverse is closed under homomorphisms. However, it turns out that many constraint satisfaction problems that can not be formulated with a finite template can be formulated with an infinite *well-behaved* template. We propose to study constraint satisfaction with countable templates that are *homogeneous*, a well-studied concept in model theory. Constraint satisfaction with such templates is a strict generalization of constraint satisfaction with finite templates, since every constraint satisfaction problem with a finite template is polynomial-time equivalent to a constraint satisfaction problem with a homogeneous template (see Section 3).

Countable homogeneous structures are intensively studied by model theorists, and they have many remarkable properties, for instance they allow quantifier elimination. If the signature contains only finitely many relation symbols for each arity they are ω-*categorical*, i.e., their first-order theories have only one countable model up to isomorphism. Countable homogeneous structures have been classified for all digraphs [14]. We use this result to determine the complexity of the constraint satisfaction problems where the template is a countable homogeneous digraph, and prove a dichotomy for the case that the structure is described by a finite number of finite forbidden induced subgraphs.

Adding *primitive positive definable* relations to a template Γ does not change the computational complexity of CSP(Γ). An equivalent characterization of primitive positive definable relations was independently found by [23] and [6]. They proved that a relation is primitive positive definable over a finite relational structure Γ if and only if it is preserved by the polymorphisms of Γ. This was first used in the context of constraint satisfaction by Jeavons et al. [30], and initiated the algebraic approach to constraint satisfaction, which has successfully been carried further e.g. in [9, 10, 16]. We generalize this result to ω-categorical structures Γ, and prove that a relation is p.p.-definable in Γ if and only if it is preserved by the polymorphisms of Γ.

Outline. We first give some background on relational homogeneous structures. In Section 3 we explain the rôle of *primitive positive definability* in constraint satisfaction. We give a characterization of primitive positive definability on countably categorical structures in Section 5 after introducing the necessary tools from universal algebra in Section 4. We close with a catalog of homogeneous digraphs and a discussion of their constraint satisfaction problems. An extended abstract of this paper appeared in [5].

2 Background

A *relational signature* τ is a (in this paper always at most countable) set of *relation symbols* R_i, each associated with an *arity* k_i. A *(relational) structure Γ over relational signature* τ (also called τ*-structure*) is a set D_Γ (the *domain*) together with a relation $R_i \subseteq D_\Gamma^{k_i}$ for each relation symbol of arity k_i. For simplicity we denote both a relation symbol and its corresponding relation with the same symbol. For a τ-structure Γ and $R \in \tau$ it will also be convenient to say that $R(u_1, \ldots, u_k)$ *holds in* Γ if $(u_1, \ldots, u_k) \in R$. We sometimes use the shortened notation $\overline{x}$ for a vector $x_1, \ldots, x_n$ of any length.

Let Γ and Γ' be τ-structures. A *homomorphism* from Γ to Γ' is a function f from D_Γ to $D_{\Gamma'}$ such that for each n-ary relation symbol in τ and each n-tuple $\overline{a}$, if $\overline{a} \in R^\Gamma$, then $(f(a_1), \ldots, f(a_n)) \in R^{\Gamma'}$. In this case we say that the map f *preserves* the relation R. A *strong homomorphism* f satisfies the stronger condition that for each n-ary relation symbol in τ and each n-tuple $\overline{a}$, $\overline{a} \in R^{\Gamma}$ if and only if $(f(a_1), \ldots, f(a_n)) \in R^{\Gamma'}$. An *embedding* of a Γ in Γ' is an injective strong homomorphism, and an *isomorphism* is a surjective embedding. Isomorphisms from Γ to Γ are called *automorphisms*. The set of all automorphisms of a structure Γ is a group with respect to composition, and denoted by $\mathrm{Aut}(\Gamma)$.

A first-order formula φ over the signature τ is said to be *primitive positive* (we say φ is a *p.p.-formula*, for short) if it is of the form

$$\exists \overline{x}(\varphi_1(\overline{x}) \wedge \cdots \wedge \varphi_k(\overline{x})) \ .$$

where $\varphi_1, \ldots, \varphi_k$ are atomic formulas. (For an introduction to first-order logic and model theory see [29].) Let Γ be a relational structure of signature τ. Then a p.p.-formula φ over τ with k free variables defines a k-ary relation $R \subseteq D_{\Gamma}^{k}$: the relation R is the set of all tuples satisfying the formula φ in Γ. Equivalently, there is a p.p.-formula defining a relation R if and only if there exists a finite relational τ-structure S containing k designated vertices $x_1, \ldots, x_k$ such that

$$R = \big\{ \big(f(x_1), \ldots, f(x_k)\big) \bigm| f \colon S \to \Gamma \text{ homomorphism} \big\} \ .$$

We call these relations *p.p.-definable*, and denote the relational structure that contains all such relations for a given Γ by $\langle \Gamma \rangle_{pp}$.

A relational structure Γ is called *homogeneous* (in the literature sometimes *ultrahomogeneous*) if every isomorphism between two finite induced substructures can be extended to an automorphism of Γ. Prominent examples of countable homogeneous structures are the *Rado graph* $\mathbf{R}$ and the dense linear order $(\mathbb{Q}, <)$. The Rado graph can be defined as the unique (up to isomorphism) model of the almost-sure theory of finite graphs. Homogeneous structures have been classified for graphs [33], tournaments [32], posets [40], and finally for digraphs [14] (there are continuum many homogeneous digraphs). For homogeneous structures with arbitrary relational signatures a classification is not yet known; even for signatures that consist of a single ternary relation this is open.

The *age of a relational τ-structure* Γ, denoted by $\mathrm{Age}(\Gamma)$, is the set of all finite *induced* substructures of Γ, i.e., the set of finite τ-structures that isomorphically embed in Γ. For a set of finite τ-structures $\mathcal{N}$, the class of all finite τ-structures without a substructure from $\mathcal{N}$ is denoted by $\mathrm{Forb}(\mathcal{N})$.

An important property of countable homogeneous structures is their characterization by their age. A class of finite structures $\mathcal{C}$ is an *amalgamation class* if $\mathcal{C}$ is nonempty, closed under isomorphism and taking induced substructures, and has the *amalgamation property*. The amalgamation property says that for all $A, B_1, B_2 \in \mathcal{C}$ and embeddings $e : A \to B_1$ and $f : A \to B_2$ there exists $C \in \mathcal{C}$ and embeddings $g : B_1 \to C$ and $h : B_2 \to C$ such that $ge = hf$.

Theorem 1 (Fraïssé [21]). *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. If this is the case, the countable structure is unique up to isomorphism, and called the* Fraïssé-limit *of $\mathcal{C}$.*

If the signature τ of a countable homogeneous structure Γ contains only finitely many relation symbols of each arity, then Γ is *ω-categorical*, i.e., every countable structure satisfying the same first-order formulas as Γ is isomorphic to Γ. On the other hand, every ω-categorical structure can be made homogeneous by expanding the signature by first-order definable relations. A permutation group over an infinite set D is called *oligomorphic* iff there is only a finite number of orbits on the set of n-tuples of D. The following theorem is essential (see e.g. [29]):

Theorem 2 (Engeler, Ryll-Nardzewski, Svenonius). *A countable structure Γ is ω-categorical if and only if $Aut(\Gamma)$ is oligomorphic.*

3 Combinatorial Constraint Satisfaction

Let Γ be an arbitrary structure with relational signature τ - also called the *template*. Then the constraint satisfaction problem CSP(Γ) is the following computational problem:

CSP(Γ)
INSTANCE: A finite τ-structure S.
QUESTION: Is there some homomorphism from S to Γ?

Formally, we denote by CSP(Γ) the set of all finite τ-structures that homomorphically map to Γ. All constraint satisfaction problems with finite Γ are clearly contained in NP.

Sometimes the age of a countable τ-structure Γ can be described by a finite set of forbidden induced substructures, i.e., $\mathrm{Age}(\Gamma) = \mathrm{Forb}(\mathcal{N})$ for a finite set $\mathcal{N}$ of finite τ-structures. We call such templates Γ *finitely constrained*. For finitely constrained Γ the constraint satisfaction problem $\mathrm{CSP}(\Gamma)$ is contained in NP. To see that, suppose we are given an instance S of $\mathrm{CSP}(\Gamma)$. A nondeterministic algorithm can then guess the image of S under a homomorphism, and verify in polynomial time that the image belongs to the age of Γ by checking the absence of a structure from $\mathcal{N}$.

Proposition 1. *Let Γ be a finitely constrained countable homogeneous relational structure. Then CSP(Γ) is in NP.*

Note that if Γ is not finitely constrained $\mathrm{CSP}(\Gamma)$ might be undecidable, see Section 7. In analogy with the dichotomy conjecture of Feder and Vardi [20], we ask the following.

Question [Dichotomy]. *Let Γ be a finitely constrained countable homogeneous relational structure. Is CSP(Γ) either NP-complete or tractable?*

For finite Γ we can assume without loss of generality that Γ is a *core*, i.e., all endomorphisms of Γ are embeddings [24]. The reason is that every finite relational structure has an endomorphism e such that the image of e induces a structure that is a core (one can prove that this core is unique up to isomorphism). This fact is not true in general for infinite structures, but remains valid for ω-categorical structures [2]. If Γ is a core, adding a singleton relation to the signature of Γ does not change the complexity of $\mathrm{CSP}(\Gamma)$ (due to [11]; see [2] for a different proof that works for ω-categorical structures as well). If Γ is a finite core, then we can add a singleton relation for every element in Γ, with the effect that the resulting structure is clearly homogeneous. Therefore constraint satisfaction with homogeneous templates can be seen as a generalization of constraint satisfaction with finite templates.

For both finite and infinite Γ, the following lemma explains the relevance of p.p.-definable relations in constraint satisfaction [30]. Suppose we extend a relational structure Γ by a p.p.-definable relation R. This does not change the computational complexity of the corresponding constraint satisfaction problem, since we can replace every occurrence of R in an instance of $\mathrm{CSP}(\Gamma)$ by the τ-structure that defines R.

Lemma 1. *Let Γ be a τ-structure and let Γ' be the extension of this structure by a relation R that is p.p.-definable over Γ. Then* $\mathrm{CSP}(\Gamma)$ *is polynomial-time equivalent to* $\mathrm{CSP}(\Gamma')$.

In the next section we introduce the algebraic notions needed to characterize p.p.-definability.

4 The Clone of Polymorphisms

Let D be a countable set, and O be the set of *finitary operations* on D, i.e., functions from D^k to D for finite k. We say that $f \in O$ *preserves* an m-ary relation $R \subseteq D^m$ if R is a subalgebra of the product algebra $(D, f)^m$. An operation that preserves all relations of a relational structure Γ is called a *polymorphism* of Γ. The set of all k-ary polymorphisms of Γ is denoted by $Pol^{(k)}(\Gamma)$, and we write $Pol(\Gamma)$ for the set of all finitary polymorphisms $Pol(\Gamma) = \bigcup_{i=1} Pol^{(i)}(\Gamma)$.

The notion of a *product* of relational structures allows an equivalent definition of polymorphisms, relating polymorphisms to homomorphisms. The *(categorical- or cross-) product* $\Gamma_1 \times \Gamma_2$ of two relational τ-structures Γ_1 and Γ_2 is a τ-structure on the domain $D_{\Gamma_1} \times D_{\Gamma_2}$. For all m-ary relations $R \in \tau$ the relation $R\big((x_1, y_2), \ldots, (x_m, y_m)\big)$ holds in $\Gamma_1 \times \Gamma_2$ iff $R(x_1, \ldots, x_m)$ holds in Γ_1 and $R(y_1, \ldots, y_m)$ holds in Γ_2. Then a k-ary polymorphism f of a relational structure corresponds to a homomorphism from $\Gamma^k = \Gamma \times \ldots \times \Gamma$ to Γ, i.e., for an m-ary relation R in τ, if $R(x_1, \ldots, x_m)$ holds in Γ^k then $R\big(f(x_1), \ldots, f(x_m)\big)$ holds in Γ.

An operation π is a *projection* (or a *trivial polymorphism*) if for all n-tuples, $\pi(x_1, \ldots, x_n) = x_i$ for some fixed $i \in \{1, \ldots, n\}$. The *composition* of a k-ary operation f and k operations $g_1, \ldots, g_k$ of arity n is an n-ary operation defined by

$$f(g_1, \ldots, g_k)(x_1, \ldots, x_n) = f\big(g_1(x_1, \ldots, x_n), \ldots, g_k(x_1, \ldots, x_n)\big) \quad .$$

A *clone* F is a set of operations from O that is closed under composition and that contains all projections. We write D_F for the *domain* D of the clone F. For a set of operations F from O we write $\langle F \rangle$ for the smallest clone containing all operations in F (the clone *generated* by F). Observe that $Pol(\Gamma)$ is a clone with the domain D_Γ.

Moreover, $Pol(\Gamma)$ is also closed under interpolation: We say that an operation $f \in O$ is *interpolated* by a set $F \subseteq O$ if for every finite subset B

of D there is some operation $g \in \langle F \rangle$ such that $f|_B = g|_B$ (f restricted to B equals g restricted to B, i.e., $f(a) = g(a)$ for every $a \in B^k$). The set of operations that are interpolated by F is called the *local closure* of F. The following is a well-known fact.

Proposition 2. *A set $F \subseteq O$ of operations is locally closed if and only if F is the set of polymorphisms of Γ for some relational structure Γ.*

We now define several important concepts for operations and clones. A k-ary operation f is *idempotent* iff $f(x, \ldots, x) = x$ for all $x \in D$. An operation f is called *essentially unary* iff there is a unary operation f_0 such that $f(x_1, \ldots, x_k) = f_0(x_i)$ for some $i \in \{1, \ldots, k\}$. A relational structure Γ is called *projective*, iff all idempotent polymorphisms of Γ are projections, and *strongly projective*, iff all polymorphisms of Γ are projections [38].

Let F be a clone with domain D. Then $R \subseteq D^m$ *is invariant under* F, if every $f \in F$ preserves R. We denote by $Inv(F)$ the relational structure containing the set of all relations that are invariant under F. A fundamental result that was independently found by Bodnarčuk et al. [6] and Geiger [23] (also see [37]) says that for arbitrary finite relational structures Γ the p.p.-definable relations are precisely the relations preserved by the polymorphisms of Γ.

Theorem 3 (Geiger [23], Bodnarčuk et al. [6]). *Let Γ be a finite relational structure. Then*

$$\langle \Gamma \rangle_{pp} \; = \; Inv(Pol(\Gamma)) \; .$$

To demonstrate the power of this theorem we prove that almost all constraint satisfaction problems are NP-complete. Nešetřil and Łuczak [34] showed that almost all structures with a fixed finite signature, chosen uniformly at at random, are strongly projective. Theorem 3 then implies that almost all such structures have a p.p.-definition of every relation, and in particular a p.p.-definition of the inequality relation. Therefore, and using Lemma 1, their constraint satisfaction problem is NP-hard on a domain of size $k \geq 3$, by reduction of k-colorability.

For arbitrary infinite structures, in general only one inclusion of Theorem 3 stays valid, which is easy to prove.

Proposition 3 (see e.g. [37]). *Let Γ be a relational structure. Then*

$$\langle \Gamma \rangle_{pp} \; \subseteq \; Inv(Pol(\Gamma)) \; .$$

We present an example that shows that the inclusion of Proposition 3 might be strict, communicated to the authors by Ferdinand Börner. Consider the relational structure $\Gamma := (\mathbb{N};\ R_1, R_2, R_3)$ on the natural numbers, where

$$\begin{aligned} R_1 &:= \{(a,b,c,d) \mid a = b \text{ or } c = d,\ a,b,c,d \in \mathbb{N}\} \\ R_2 &:= \{(0)\} \\ R_3 &:= \{(a, a+1) \mid a \in \mathbb{N}\} \end{aligned}$$

Every function preserving R_1 is essentially unary [37]. If f is unary and preserves R_2 then $f(0) = 0$. Furthermore, if f preserves R_3 we have $f(a+1) = f(a) + 1$ for all a, and inductively follows $f(a) = a$. Therefore $Pol(\Gamma)$ only contains the projections. Every projection preserves all relations. There are uncountably many relations over $\mathbb{N}$, but only countably many p.p.-formulas. Thus, $Inv(Pol(\Gamma))$ contains relations that are not p.p.-definable. A concrete example for such a relation is the predicate Odd that holds on all odd natural numbers. The reason is that for every p.p.-definable k-ary relation R in Γ the *diagonal set* $R^* := R \cap \{(x,\dots,x) | x \in \Gamma\}$ is either finite or cofinite, i.e., either R^* or $\Gamma - R^*$ contains finitely many elements. We call this property of a relation $(*)$, and observe that the relations R_1, R_2, and R_3 have property $(*)$, since the diagonal sets are full, of size 1, and empty, respectively. Any projection and all finite intersections preserve the property $(*)$. Since the relation Odd is neither finite nor cofinite, it cannot have a primitive positive definition in Γ.

For ω-categorical structures Γ the first-order definable relations are precisely the relations that are preserved by the automorphisms of Γ, i.e. $\langle \Gamma \rangle_{fo} = Inv(Aut(\Gamma))$ (see e.g. [12, 29]). We prove a corresponding theorem for primitive positive definability in the next section.

5 A Characterization of Primitive Positive Definability

We characterize the primitive positive first-order definable relations over an ω-categorical structure Γ by the polymorphisms of Γ.

Theorem 4. *Let Γ be an ω-categorical structure with relational signature τ. Then a relation R on Γ is preserved by the polymophisms of Γ if and only if R is p.p.-definable, i.e.,*

$$\langle \Gamma \rangle_{pp} = Inv(Pol(\Gamma)).$$

Proof. We already stated in Proposition 3 that the p.p.-definable relations over Γ are invariant under the polymophisms of Γ.

For the converse, let R be a k-ary relation from $Inv(Pol(\Gamma))$. Note that R is first-order definable in Γ: By ω-categoricity and Theorem 2, and since Γ and $Inv(Pol(\Gamma))$ have the same automorphism group, the relation R is a union of *finitely* many orbits of the automorphism group of Γ, and it can be defined by a disjunction φ of τ-formulas that define these orbits. Let $M_1, \ldots, M_w$ be the satisfiable monomials in this disjunction, and let $x_1, \ldots, x_k$ be the variables of these monomials. We have to construct a finite τ-structure Q with designated vertices $v_1, \ldots, v_k$ such that

$$R = \big\{ \big(f(v_1), \ldots, f(v_k)\big) \bigm| f \colon Q \to \Gamma \text{ homomorphism} \big\}.$$

The idea is to first consider an *infinite* τ-structure, namely the categorical product Γ^w, and then to apply a compactness argument to prove the existence of a suitable finite substructure.

For each monomial $M_j \in M_1, \ldots, M_w$ of φ we find a substructure $a_1^j, \ldots, a_k^j$ of Γ, such that $a_1^j, \ldots, a_k^j$ satisfies M_j in Γ. Let $b_1, b_2, \ldots$ be an enumeration of the w-tuples in D_Γ^w, starting with $b_i = (a_1^i, \ldots, a_w^i)$ for $1 \leq i \leq k$. Let us call a partial mapping from Γ^w to Γ a *bad* mapping if it maps $b_1, \ldots, b_k$ to a tuple not satisfying φ. Since R is preserved by all polymorphisms, no homomorphism from Γ^w to Γ is bad.

We now claim that there is a finite substructure Q of Γ^w such that no homomorphism from Q to Γ is bad. Assume for contradiction that all finite substructures of Γ^w containing $b_1, \ldots, b_k$ have a homomorphism to Γ mapping $b_1, \ldots, b_k$ to a tuple not satisfying φ. We shall construct a bad homomorphism from Γ^w to Γ, i.e. the images of $b_1, \ldots, b_k$ do not satisfy φ. This will contradict the fact that R is preserved by all polymorphisms. To this end, consider the following infinite but finitely branching tree. The nodes on level n in the tree are the equivalence classes of the bad homomorphisms from Γ^w restricted to $\{b_1, \ldots, b_n\}$ to Γ, where two homomorphisms f_1 and f_2 are equivalent if $f_1 = gf_2$ for some $g \in Aut(\Gamma)$. Adjacency between nodes on consecutive levels is defined by restriction. By our assumption, for each finite substructure of Γ^w there is a bad homomorphism, and thus the tree contains a node on each level. By Theorem 2, there are only finitely many nodes on each level. Hence, König's Lemma asserts the existence of an infinite path in the tree. This path defines a bad homomorphism from Γ^w to Γ.

We proved by contradiction that there must be a finite substructure Q containing the vertices $b_1, \ldots, b_k$ of Γ^w such that all homomorphisms from

Q to Γ map $b_1, \ldots, b_k$ to a tuple satisfying φ. Conversely, every mapping $f : Q \rightarrow \Gamma$ such that the tuple $(f(b_1), \ldots, f(b_k))$ satisfies in Γ the monomial M_j can be extended to a homomorphism $f : \Gamma^w \rightarrow \Gamma$. To see this note that both $a_1^j, \ldots, a_k^j$ and $(f(b_1), \ldots, f(b_k))$ satisfy M_j and thus both lie in the same orbit of $Aut(\Gamma)$. Thus we can choose f to be the jth projection combined with the automorphism sending $(a_1^j, \ldots, a_k^j)$ to $(f(b_1), \ldots, f(b_k))$. This completes the proof. $\square$

6 Examples of Countably Categorical Templates

In this section we demonstrate that several problems and classes of problems studied in the literature can be formulated as constraint satisfaction problems with ω-categorical templates. In fact, in all cases the templates can be expanded by a finite number of first-order definable relations, such that the resulting structure is homogeneous. Moreover, in all cases this homogeneous expansion is described by a finite number of finite forbidden induced substructures.

6.1 Tree Descriptions

We start with a computational problem that was first posed in [15], and that is motivated by questions in computational linguistics. There is a polynomial time graph-theoretic algorithm solving this problem [4]. We remark that this algorithm is neither group-theoretic nor uses Datalog or local consistency methods, in contrast to all efficient algorithms that are known for constraint satisfaction with finite templates [20].

PARTIAL-TREE-DESCRIPTION
INSTANCE: A finite structure S over the signature $\tau = \{\longrightarrow, \perp\}$ containing two binary relation symbols.
QUESTION: Can we find a rooted forest F on the vertices of S such that every edge from $\longrightarrow$ lies in the transitive closure of F, and every edge $\perp$ does not?

Using *a countable semi-linear order* (see [18]) we can formulate this problem (and related problems) as a constraint satisfaction problem CSP(Λ). In fact, we can find an ω-categorical structure Λ such that CSP(Λ) contains precisely the solvable instances. The signature of the structure Λ is $\{\longrightarrow, \perp\}$, and the domain is the set of all non-empty finite sequences of rational numbers. For $a = (q_1, q_1, \ldots, q_n), b = (q'_1, q'_1, \ldots, q'_m)$, $n \leq m$, we write $a \longrightarrow b$ if one of the following conditions holds:

– a is a proper initial subsequence of b, i.e., $q_i = q'_i$ for $1 \leq i \leq n$;
– $q_i = q'_i$ for $1 \leq i < n$, and $q_n < q'_n$.

The relation $\perp$ is the set of all unordered pairs of distinct points not in the relation $\longrightarrow$. We already mentioned that every ω-categorical structure can be made homogeneous by expanding the signature by some first-order definable relations. In the case of Λ this is possible with a single ternary relation [19], namely the relation x:yz defined by the following primitive positive formula.

$$\exists u.\, y \perp z \wedge u \longrightarrow y \wedge u \longrightarrow z \wedge u \perp x$$

It is also easy to find a finite list of forbidden induced substructures that characterizes the age of this structure.

6.2 Allen's Interval Algebra and its Fragments

We briefly introduce Allen's interval algebra [1], which is a famous framework for reasoning about temporal constraints. The satisfiability problem for all subclasses of Allen's interval algebra can be formulated as a constraint satisfaction problem with an ω-categorical template.

Consider as a base set D the closed intervals on the rational numbers, and the following binary relations on these intervals: Let $x = [x^-, x^+]$ and $y = [y^-, y^+]$ be closed intervals. We define

– The interval x *precedes* y, $x \;\mathsf{p}\; y$, iff $x^+ < y^-$.
– The interval x *overlaps* y, $x \;\mathsf{o}\; y$, iff $x^- < y^- < x^+$ and $x^+ < y^+$.
– The interval x *is during* y, $x \;\mathsf{d}\; y$, iff $y^- < x^-$ and $x^+ < y^+$.
– The interval x *starts* y, $x \;\mathsf{s}\; y$, iff $x^- = y^-$ and $x^+ > y^-$.
– The interval x *finishes* y, $x \;\mathsf{f}\; y$, iff $x^+ = y^+$ and $x^- > y^-$.
– The interval x *meets* y, $x \;\mathsf{m}\; y$, iff $x^+ = y^-$.
– The interval x *equals* y, $x \equiv y$, iff $x^- = y^-$ and $x^+ = y^+$.

For any set of relations derived from $\mathsf{p}, \mathsf{o}, \mathsf{d}, \mathsf{s}, \mathsf{m}, \mathsf{f}$ and $\equiv$ by union and complementation the corresponding countable relational structure is ω-categorical. The constraint satisfaction problems for these structures have a dichotomy [31, 35]. In contrast to constraint satisfaction with finite templates, where the problem one-in-three-sat is usually the most natural candidate to prove NP-hardness, here it is natural to use the problem *Betweennness* [22] to prove hardness.

BETWEENNESS
INSTANCE: A finite set V, and a collection C of ordered triples (x, y, z) of distinct elements from V.
QUESTION: Is there a one-to-one function $f : V \to \{1, \ldots, |V|\}$ such that, for each $(a, b, c) \in C$, we have either $f(a) < f(b) < f(c)$ or $f(a) < f(c) < f(b)$.

This problem can itself be formulated as a constraint satisfaction problem with an ω-categorical template. The domain is again the set of all rational numbers, and we have one relation symbol in the signature for the ternary relation $\{(x, y, z) \subseteq \mathbb{Q}^3 \mid x < y < z \text{ or } z < y < x\}$.

7 A Catalog of Homogeneous Templates

In this section we study the constraint satisfaction problems for homogeneous digraphs. We start with the homogeneous tournaments, which have been classified by Lachlan [32]. Up to isomorphism there are only five: the isolated vertex, which we denote by K_1, the oriented cycle C_3, the dense linear order $(\mathbb{Q}, <)$, the dense local order $S(2)$, and the generic tournament for the set of all finite tournaments.

The problem $\mathrm{CSP}(K_1)$ is trivial; and $\mathrm{CSP}(C_3)$ is known to be tractable (see e.g. [26]). The constraint satisfaction problem of the dense linear order $(\mathbb{Q}, <)$ is computationally equivalent to the problem whether a given digraph D is acyclic. Clearly, this tractable problem can not be formulated as a constraint satisfaction problem with a finite template. Note that the relational structure $(\mathbb{Q}, <)$ is not projective, e.g. the map $x, y \mapsto \max(x, y)$ is a polymorphism. The homogeneous tournament that is the Fraïssé-limit of all finite tournaments has a trivial constraint satisfaction problem: Every oriented graph maps to it. Thus the only interesting remaining case is the *dense local order* $S(2)$ (see [14]).

To define $S(2)$, consider a partition of the rational numbers $\mathbb{Q}$ in two disjoint dense subsets X and Y of (i.e., for every rational number a we can find sequences in X and in Y that converge against this number a). Then the relation $\prec$ of $S(2)$ is defined as the dense linear order of $\mathbb{Q}$ on $X \cup Y$,

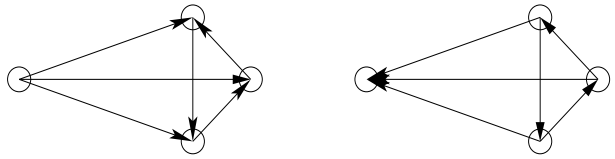

Fig. 1. The forbidden induced subgraphs $[I_1, C_3]$ and $[C_3, I_1]$ of $S(2)$.

where the edges between the sets X and Y are reversed. Formally, we define $u \prec v$ iff either $u < v$ and $u, v \in X$, $u < v$ and $u, v \in Y$, $v < u$ and $u \in X$, $v \in Y$, or $v < u$ and $v \in Y$, $u \in Y$. The tournament $S(2)$ is finitely constrained, and the two forbidden induced subtournaments are shown in Figure 1.

Proposition 4. *CSP($S(2)$) is NP-complete.*

Proof. Proposition 1 shows that CSP($S(2)$) is contained in NP, because $S(2)$ is finitely constrained. To prove hardness we reduce the problem Betweenness to CSP($S(2)$). Let $(V; C)$ be an instance of Betweenness. We define a polynomial size instance S of CSP($S(2)$) that is satisfiable if and only if $(V; C)$ is a yes-instance of Betweenness. The vertices of S consist of the vertices V and some additional vertices. We first introduce a new vertex u, and add $u < x$ to S for all $x \in V$. Then introduce for each triple x, y, z from C two new vertices v, w and add the constraints $v > x, v > y, v < z, w < x, w > y, w > z$ to the instance S of CSP($S(2)$). Also see Figure 2.

If there is a solution to the Betweenness instance $(V; C)$, there is also a homomorphism from S to $S(2)$: We map the vertex $u \in S$ to *some* vertex $f(u)$ in $S(2)$. Then the linearly ordered set $\{w \mid f(u) < w\}$ is isomorphic to the linear order of the rational numbers, and we map the vertices in V to this set in the same way as the solution of the Betweenness instance maps them to the rational line. Finally, we can map the existentially quantified variables to $S(2)$ in either of the two ways displayed in Figure 2: this is, if $x < y < z$ in the solution to the Betweenness instance we let $x < v$, $v < y$, $v < z$, $z < w$, $y < w$, $w < x$. Otherwise, if $z < y < x$ we let $z < w$, $y < w$, $w < x$, $y < v$, $x < v$, $v < z$.

Conversely, if there is a homomorphism f from S to $S(2)$, we also have a solution to the Betweenness instance $(V; C)$. The homomorphism f maps all vertices in V except u to the linearly ordered set $\{w \mid f(u) < w\}$. We claim that in the corresponding linear order of the vertices V, for each triple in C either $x < y < z$ or $z < y < x$. Assume otherwise that $y < x$ and $y < z$. Since x and y have to be comparable, either $x > z$ or $x < z$. In the first case we find the oriented three-cycle x, v, z, x, in the second the cycle x, z, w, x. In both cases there are arcs from y towards the vertices of that cycle, and we found the subgraph $[I_1, C_3]$, which is forbidden in $S(2)$. The case where $y > x$ and $y > z$ is analogous with the forbidden constellation $[C_3, I_1]$. Thus we have a contradiction to the assumption that f is a homomorphism from S to $S(2)$. $\square$

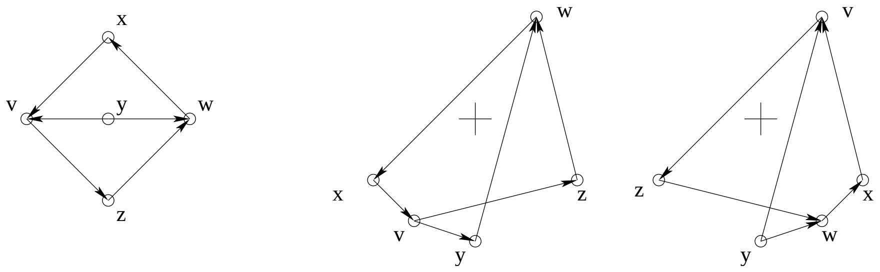

Fig. 2. On the left we find the gadget to simulate with $\mathrm{CSP}(S(2))$ the betweenness relation on x, y, z, where an arc from a to b means that there is the constraint $a < b$. On the right we find two possible ways to map these vertices homomorphically to $S(2)$, where the cross marks the origin in the plane.

Countable Homogeneous Digraphs. There are uncountably many countable homogeneous digraphs. But they have been classified by Cherlin [14]; the classification shows that the age of all but a countable number of well-understood homogeneous digraphs Γ has the property, that all subgraphs of Γ are also induced subgraphs of Γ. Henson showed that the set of isomorphism types of finite tournaments, partially ordered by embeddability, contains an infinite antichain [28]. For all distinct subsets $\mathcal{T}$ of tournaments in this antichain the classes $\mathrm{Forb}(\mathcal{T})$ are distinct amalgamation classes. Thus, there is an uncountable number of such classes, and by Fraïssé's theorem there is an uncountable number of non-isomorphic homogeneous digraphs.

Proposition 5. *There is a homogeneous digraph Γ such that CSP(Γ) is undecidable.*

Proof. Each of the uncountably many homogeneous digraphs from Henson's construction has a different constraint satisfaction problem, because the sets $\mathrm{CSP}(\Gamma)$ and $\mathrm{Age}(\Gamma)$ are equal for the constructed digraphs. Since there is a countable number of algorithms, undecidable constraint satisfaction problems exist, with templates that are countable homogeneous digraphs. □

However, if the age is described by a finite set $\mathcal{N}$ of finite forbidden induced subgraphs, then the constraint satisfaction problem for these templates is simple, since $\mathrm{CSP}(\Gamma) = \mathrm{Age}(\Gamma) = \mathrm{Forb}(\mathcal{N})$. Hence it suffices to check whether the input contains a forbidden induced subgraph, which can be done in polynomial time. If not, the input digraph is in the age of the template and is a yes-instance.

As in [14], we divide the remaining homogeneous digraphs into three classes. The first class contains the countable homogenous tournaments that we already know from the last paragraph in this section, and the empty graph on a countable number of vertices.

The second class contains *imprimitive* structures, i.e., structures that have a first-order definable nontrivial equivalence relation. In all cases except the first the definable equivalence relation corresponds to the pairs of vertices that are not connected. There are four types, classified in [13]. We use Cherlin's notation; $G[H]$ denotes the *composition* or *wreath product* of G and H: each vertex of G is replaced by a copy of H, and arcs between distinct copies of H are controlled by the edges of G.

1. "Wreathed (Composite)": $I_n[\mathbb{Q}]$, $I_n[S(2)]$, $I_n[T^\infty]$ and $\mathbb{Q}[I_n]$, $S(2)[I_n]$, $T^\infty[I_n]$, $C_3[I_\infty]$, where n might be infinite.
2. "Twisted": $\hat{\mathbb{Q}}$, $\hat{T^\infty}$. The digraph $\hat{\mathbb{Q}}$ is a variant of $S(2)$ where every point has an unconnected antipodal. The digraph $\hat{T^\infty}$ is universal for the class of all digraphs where the pairs of disconnected vertices form an equivalence relation with classes of size two, and the union of two such classes is a copy of C_4.
3. "Generified n-partite" (denoted by $n * I_\infty$ in [13]): Starting from $K_n[I_\infty]$ we orient the undirected edges randomly.
4. "Semigeneric" (denoted by $\infty * I_\infty$ in [13]): Starting from the undirected homogeneous digraph $K_\infty[I_\infty]$ we orient the edges arbitrarily such that for two distinct parts U and V and distinct vertices $u_1, u_2 \in U$ and $v_1, v_2 \in V$ the number of edges from u_1 or u_2 to v_1 or v_2 is even.

Finally we have the *exceptional class* containing the digraphs $S(3)$, $\mathbb{P}$, and $P(3)$. The digraph $S(3)$ denotes the set of points lying at a rational angle ϕ on the unit circle, and two points a, b are joined by and arc $a \to b$ iff the angle from a to b is in the range $(0, 2\pi/3)$. Equivalently we can produce the digraph similarly to $S(2)$, starting from a partition of $\mathbb{Q}$ into three dense sets Q_1, Q_2, Q_3. We identify the two possible orientations of an edge with $+1$ and -1, and 0 represents the absence of an edge. We then cyclically shift the edges between Q_i and Q_j by $j - i$ (the indices are integers modulo 3).

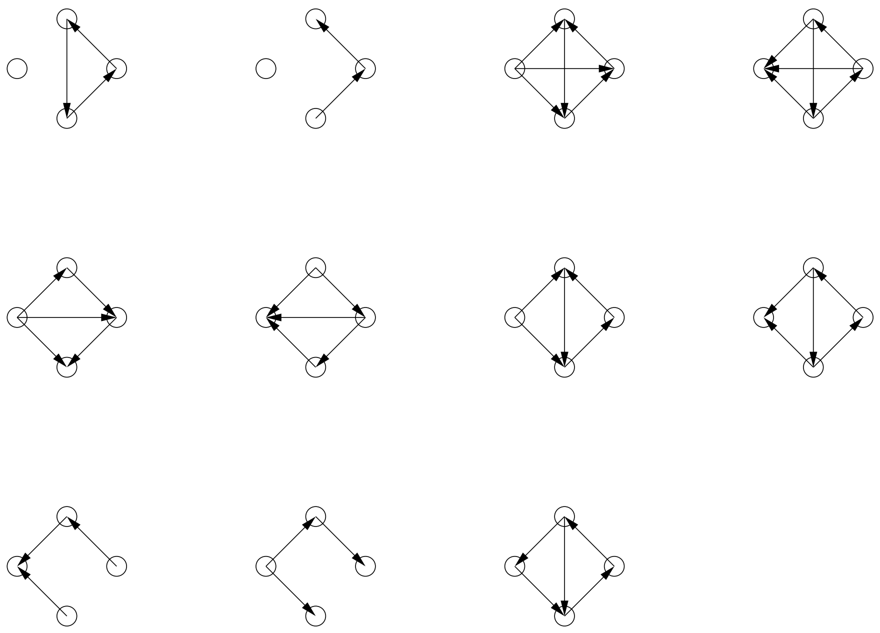

Fig. 3. The forbidden induced subgraphs of $P(3)$.

The countable homogeneous partial order $\mathbb{P}$ is the Fraïssé-limit of the class of all finite partial orders $\mathcal{P}$. We call a subset P of $\mathbb{P}$ *dense*, if for any pair a, b in $\mathbb{P}$ with $a \to b$ we have an element $c \in P$ such that $a \to c \to b$. $P(3)$ is the analog of $S(3)$, based on $\mathbb{P}$ instead of $\mathbb{Q}$. Let P_0, P_1, P_2 be three dense subsets of $\mathbb{P}$. The structure $P(3)$ is defined on $P_0 \uplus P_1 \uplus P_2$. The relation $\to$ on $P(3)$ restricted to P_i is defined as in $\mathbb{P}$. Again we identify the complete 2-types between elements of distinct parts with $\{-1, 0, +1\}$, and cyclically shift these types between P_i and P_j by $j-i$. Cherlin specified the age of $P(3)$ by the forbidden induced substructures shown in Figure 3 (their description can be extracted from the proof of Proposition 24 on page 126f in [14]).

Since $S(2)$ naturally embeds into $\hat{\mathbb{Q}}$, and $\hat{\mathbb{Q}}$ homomorphically maps to $S(2)$, these problems have the same constraint satisfaction problem. For the hardness of CSP($S(3)$) we can use the same gadget as above to simulate Betweenness on a suitable subset of vertices. Similarly to Proposition 4 we can prove the NP-hardness of CSP($P(3)$). Before we present a hardness proof, we reformulate the computational problem CSP($P(3)$); in this form we call the problem *switching-trigraph-transitivity.*

SWITCHING-TRIGRAPH-TRANSITIVITY
INSTANCE: A digraph $D = (V; E)$.
QUESTION: Can we partition the vertices V into three parts P_1, P_2, P_3, such that the digraph that arises from D by the following three operations is transitive?
i) deleting the edges from P_1 to P_2, P_2 to P_3, or P_3 to P_1;
ii) reversing the arcs from P_2 to P_1, P_3 to P_2, or P_1 to P_3;
iiii) adding arcs between unconnected pairs from P_2 to P_1, P_3 to P_2, or P_1 to P_3.

Proposition 6. *CSP($P(3)$) is NP-complete.*

Proof. As in the proof of Proposition 4 we use Proposition 1 to observe that CSP($P(3)$) is contained in NP, because $P(3)$ is finitely constrained. Again we prove hardness by reducing Betweenness to CSP($P(3)$), and we construct an instance S of CSP($P(3)$) from an instance $(V; C)$ of Betweenness in the same way as in the proof of Proposition 4. Now, we additionally introduce new vertices a, b for all pairs of vertices $x, y \in V$, and impose the constraint $a < b, x > b, b > y, y > a, a > x$; see Figure 4.

First we show that given a solution to the Betweenness instance $(V; C)$, we can construct a homomorphism from S to $P(3)$. Note that $P(3)$ contains copies of $S(2)$ as subgraphs; fix such a copy C of $S(2)$. We map the vertex $u \in S$ to some vertex $f(u)$ in C. The set $\{x \in C \mid f(u) < x\}$ induces the linear order of the rational numbers in $P(3)$. To construct a solution for

S, we map the vertices in V to this set in the same way as the solution of the Betweenness instance maps them to the rational line. As in the proof of Proposition 4, we can map the remaining existentially quantified variables to C in either of the two ways displayed in Figure 2.

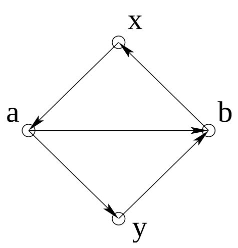

Fig. 4. We use this gadget to ensure that the vertices x and y are mapped to comparable vertices in $P(3)$.

Conversely, if there is a homomorphism f from the constructed instance S to $P(3)$, we also have a solution to the Betweenness instance $(V; C)$. Every solution to S maps pairs of distinct vertices $x, y \in V$ to *comparable* vertices in $P(3)$. To see that, consider the corresponding subgraph x, y, a, b in S. Since this structure is a forbidden induced subgraph of $P(3)$ (see Figure 3 on page 17), any homomorphism of this structure to $P(3)$ has to map x and y to comparable vertices in $P(3)$. Hence, the set $f(V)$ induces in $P(3)$ a tournament. Because $f(u) < x$ for all x in $f(V)$, and because $[I_1, C_3]$ is a forbidden induced subgraph in $P(3)$, the set $f(V)$ is linearly ordered. We claim for each triple in C either $x < y < z$ or $z < y < x$. This can be shown exactly as in the proof of Proposition 4. $\square$

Combining all the above results we get that if a countable homogeneous digraph is finitely constrained, then its constraint satisfaction problem is NP-complete or tractable.

References

1. J. F. Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843, 1983.

2. M. Bodirsky. The core of a countably categorical structure. In V. Diekert and B. Durand, editors, *Proceedings of the 22nd Annual Symposium on Theoretical Aspects of Computer Science (STACS'05), Stuttgart (Germany)*, LNCS 3404, Springer-Verlag Berlin Heidelberg, 2005.
3. M. Bodirsky, D. Duchier, J. Niehren, and S. Miele. A new algorithm for normal dominance constraints. In *12th ACM-SIAM Symposium on Discrete Algorithms (SODA'04)*, pages 59–67, 2003.
4. M. Bodirsky and M. Kutz. Pure dominance constraints. In *Proceedings of the 19th Annual Symposium on Theoretical Aspects of Computer Science (STACS'02)*, pages 287–298, Antibes - Juan le Pins, 2002.
5. M. Bodirsky and J. Nešetřil. Constraint satisfaction with countable homogeneous templates. In *Proceedings of Computer Science Logic (CSL'03), LNCS 2803*, pages 44–57, Vienna, 2003.
6. V. G. Bodnarčuk, L. A. Kalužnin, V. N. Kotov, and B. A. Romov. Galois theory for post algebras, part I and II. *Cybernetics*, 5:243–539, 1969.
7. D. Bryant. Building trees, hunting for trees, and comparing trees. PhD-thesis at the University of Canterbury, 1997.
8. A. Bulatov. A dichotomy theorem for constraints on a three-element set. In *Proceedings of the Symposium on Foundations of Comupter Science (FOCS'02)*, pages 649–658, 2002.
9. A. Bulatov. Tractable conservative constraint satisfaction problems. In *Proceedings of 18th Symposium on Logig in Computer Science (LICS'03)*, pages 321–330, 2003.
10. A. Bulatov, A. Krokhin, and P. Jeavons. The complexity of maximal constraint languages. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC'01)*, pages 667–674, 2001.
11. A. Bulatov, A. Krokhin, and P. G. Jeavons. Classifying the complexity of constraints using finite algebras. *To appear in SIAM Journal on Computing*, 2004.
12. P. J. Cameron. *Oligomorphic Permutation Groups.* Cambridge University Press, 1990.
13. G. Cherlin. Homogeneous directed graphs. The imprimitive case. *Logic Colloquium 1985*, pages 67–88, 1987.
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 *32nd Annual Meeting of the Association for Computational Linguistics (ACL'94)*, pages 163–170, 1994.
16. V. Dalmau. A new tractable class of constraint satisfaction problems. To appear in Annals of Mathematics and Artificial Intelligence. Presented at the 6th International Symposium on Mathematics and Artificial Intelligence, 2000.

17. R. Dechter. Local and global relational consistency. *Theoretical Computer Science*, 173(1):283–308, 1997.
18. M. Droste. Structure of partially ordered sets with transitive automorphism groups. *AMS Memoir*, 57(334), September 1985.
19. M. Droste, W. Holland, and D. Macpherson. Automorphism groups of infinite semilinear orders (I). *Proc. London Math. Soc.*, 58:454–478, 1989.
20. T. Feder and M. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through datalog and group theory. *SIAM J. Comput.*, 28:57–104, 1999.
21. R. Fraïssé. *Theory of Relations*. North-Holland, 1986.
22. Garey and Johnson. *A Guide to NP-completeness*. CSLI Press, Stanford, 1978.
23. D. Geiger. Closed systems of functions and predicates. *Pacific Journal of Mathematics*, 27:95–100, 1968.
24. P. Hell and J. Nesetril. The core of a graph. *Discrete Math.*, 109:117–126, 1992.
25. P. Hell and J. Nešetřil. On the complexity of H-coloring. *Journal of Combinatorial Theory, Series B*, 48:92–110, 1990.
26. P. Hell and J. Nešetřil. *Graphs and Homomorphisms*. Oxford University Press, 2004.
27. P. Hell, J. Nešetřil, and X. Zhu. Duality and polynomial testing of tree homomorphisms. *Trans. Amer. Math. Soc.*, 348(4):1281–1297, 1996.
28. C. W. Henson. Countable homogeneous relational systems and categorical theories. *Journal of Symbolic Logic*, 37:494–500, 1972.
29. W. Hodges. *A shorter model theory*. Cambridge University Press, 1997.
30. P. Jeavons, D. Cohen, and M. Gyssens. Closure properties of constraints. *Journal of the ACM*, 44(4):527–548, 1997.
31. P. Jeavons, P. Jonsson, and A. A. Krokhin. Reasoning about temporal relations: The tractable subalgebras of Allen's interval algebra. *Journal of the ACM*, 50(5):591–640, 2003.
32. A. H. Lachlan. Countable homogeneous tournaments. *Trans. Amer. Math. Soc.*, 284:431–461, 1984.
33. A. H. Lachlan and R. Woodrow. Countable ultrahomogeneous undirected graphs. *Trans. Amer. Math. Soc.*, 262(1):51–94, 1980.
34. T. Łuczak and J. Nešetřil. A note on projective graphs. *J. Graph Theory*, (47):81–86, 2004.
35. B. Nebel and H.-J. Bürckert. Reasoning about temporal relations: A maximal tractable subclass of Allen's interval algebra. *Journal of the ACM*, 42(1):43–66, 1995.
36. M. P. Ng, M. Steel, and N. C. Wormald. The difficulty of constructing a leaf-labelled tree including or avoiding given subtrees. *Discrete Applied Mathematics*, 98:227–235, 2000.

37. R. Pöschel and L. A. Kalužnin. *Funktionen- und Relationenalgebren.* DVW, 1979.
38. I. G. Rosenberg. Strongly rigid relations. *Rocky Mountain Journal of Mathematics*, 3(4):631–639, 1973.
39. T. J. Schaeffer. The complexity of satisfiability problems. *In Proc. 10th ACM Symp. on Theory of Computing*, pages 216–226, 1978.
40. J. H. Schmerl. Countable homogeneous partially ordered sets. *Algebra Universalis*, 9:317–321, 1979.
41. M. Steel. The complexity of reconstructing trees from qualitative charaters and subtrees. *Journal of Classification*, 9:91–116, 1992.