Qualitative Temporal and Spatial Reasoning Revisited*

Manuel Bodirsky and Hubie Chen

[1] Department Algorithms and Complexity, Humboldt-Universität zu Berlin,
bodirsky@informatik.hu-berlin.de
[2] Departament de Tecnologies de la Informació i les Comunicacions,
Universitat Pompeu Fabra, Barcelona, hubie.chen@upf.edu

Abstract. Establishing local consistency is one of the main algorithmic techniques in temporal and spatial reasoning. A central question for the various proposed temporal and spatial constraint languages is whether local consistency implies global consistency. Showing that a constraint language Γ has this "local-to-global" property implies polynomial-time tractability of the constraint language, and has further pleasant algorithmic consequences.
In the present paper, we study the "local-to-global" property by making use of a recently established connection of this property with universal algebra. Roughly speaking, the connection shows that this property is equivalent to the presence of a so-called quasi near-unanimity polymorphism of the constraint language. We obtain new algorithmic results and give very concise proofs of previously known theorems. Our results concern well-known and heavily studied formalisms such as the point algebra, Allen's interval algebra, and the spatial reasoning language RCC-5.

1 Introduction

Temporal and spatial reasoning is a subdiscipline in Artificial Intelligence that developed in the 1990s, and has many applications, for instance in natural language processing, geographic information systems, computational biology, and document interpretation; for references, see the monograph [FGV05] and the survey [RN07]. A common reasoning task in this field is to decide, given a set of relationships concerning temporal events or spatial regions, whether or not there exists a model fulfilling all of the relationships. It is well-acknowledged that instances of this reasoning task may be modelled using the *constraint satisfaction problem (CSP)*, a computational problem in which the input consists of a set of constraints on variables, and the question is whether or not there is an assignment to the variables satisfying all of the constraints. In this vein, a famous example from temporal reasoning is the CSP for *Allen's Interval Algebra*, where the variables denote intervals in time, and the constraints talk about relationships between intervals such as containment, overlap, and so forth [All83].

CSPs for temporal and spatial reasoning have been studied from a computational complexity perspective by restricting the sets of relationships that may be used. In the CSP formalism, this amounts to restricting the set of predicates that may be used to form constraints; we call such a restricted predicate set a *constraint language*, following the CSP literature.

A primary algorithmic technique for solving CSPs in spatial and temporal reasoning is the process of establishing k-consistency, an inferential process that yields a problem that is k-consistent: any partial solution on $(k-1)$ variables can be extended to any other variable. The notions of consistency that we employ in this paper are due to [Mac77, Fre82, Dec92]; formal definitions are provided later (see Section 2). For some constraint languages, it is known that (for some constant k) establishing k-consistency implies *global consistency*, the property of being i-consistent for all i. We refer to this property of constraint languages as the "local-to-global" property. Showing that a constraint language possesses this property implies that it is polynomial-time tractable, and has further desirable algorithmic consequences; for instance, we demonstrate a connection to a quantifier elimination algorithm for the *quantified* constraint satisfaction problem over the constraint

* An extended abstract of this article appeared in the 16th EACSL Annual Conference on Computer Science and Logic (CSL 2007).

language. One of the central questions for the various temporal and spatial constraint languages is to understand which such languages enjoy the "local-to-global" property.

In this paper, we study this question by making use of algebraic techniques for studying the complexity of constraint languages that have recently come into focus (see for instance the surveys [BJK05, CJ06]). Specifically, a fundamental result [BKJ05] associates to every constraint language an algebra in a way that permits the use of universal algebraic concepts and methods in the study of the complexity and algorithmic behavior of constraint languages. Utilizing this algebraic perspective, we both derive new algorithmic results and give very concise proofs of previously known theorems. We thus establish connections between two research areas–temporal and spatial reasoning, and algebraic techniques for constraint languages–that have up to the present seen little interaction, but indeed, as evidenced by our results, unite quite naturally. We hope that the present work serves to stimulate further interaction between these two areas.

Before giving more detail on our contributions, we describe how exactly we view temporal and spatial reasoning problems within the CSP framework, and which tools from CSP theory we employ. First, it should be pointed out that much of the work on CSP complexity has focused on constraint languages over *finite* domains. For such constraint languages, an exact algebraic characterization of the "local-to-global" property is known, namely, a constraint language has this property if and only if it has a so-called *near-unanimity polymorphism*. This characterization was presented in [JCC98] and in [FV99], and in part can be seen as reformulation of a classical result in universal algebra [BP74].

On the other hand, many temporal and spatial reasoning problems are naturally formulated as the CSP over constraint languages with *infinite* domains. Fortunately, a number of such problems can be formulated as the CSP over a language from the class of *ω-categorical constraint languages*, a class of languages that is known to be relatively manageable from the algebraic and logical viewpoints [BN06, HN92]. Formulability in this class is well-known for the Point Algebra, and for Allen's interval algebra and all its fragments [Hir96]. Moreover, it has recently been shown that the mentioned algebraic characterization of the "local-to-global" property essentially remains valid for ω-categorical structures [BD06]. Precisely, it has been shown that an ω-categorical constraint language of bounded maximal arity has this property if and only if it has a so-called *oligopotent quasi near-unanimity polymorphism*.

Contributions and outline. Sections 2 and 3 recall fundamental facts from the theory of constraint satisfaction that allow us to study temporal and spatial constraint languages algebraically. In Section 4 we prove results concerning the existence and properties of quasi near-unanimity polymorphisms for infinite posets. One of the applications of this result is a new and concise proof of the result of Koubarakis that $(2k+1)$-consistency (but not $2k$-consistency) implies global tractability for the Point Algebra with disjunctions of disequalities on at most k variables [Kou97] (discussed in Section 5). In Section 5, we provide characterizations of the fragments of the point algebra in terms of quasi near-unanimity polymorphisms. We also show that if we extend the constraint language for the Point Algebra to contain disjunctions of disequalities, the corresponding (uniform) *quantified CSP* can be solved in NL, giving a strict extension of the result of Koubarakis [Kou97].

In Section 6, we study the spatial reasoning formalism RCC-5. We first formulate the corresponding CSP with an ω-categorical constraint language, and then show that the so-called *basic relations* of RCC-5 possess the "local-to-global" property. In fact, in analogy to Koubarakis' result on the Point Algebra, we show that $(2k+1)$-consistency (but not $2k$-consistency) implies global consistency for the basic relations of RCC-5 with disjunctions of disequalities on at most k variables.

Finally, we present a general technique for establishing the "local-to-global" property for various temporal and spatial constraint languages, which is based on the model-theoretic concept of primitive positive interpretations (Section 7). We apply this technique to the pointizable fragment of Allen's interval algebra: we show that the basic relations of the rectangular algebra have a 5-ary quasi near-unanimity polymorphism, and hence that they have the "local-to-global" property.

2 Preliminaries

Relational structures. A *relational signature* τ is a set of *relation symbols* R_i, each of which has an associated finite *arity* k_i. A *relational structure* Γ over the 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 from τ. For simplicity, we use the same symbol for a relation symbol and the corresponding relation. If necessary, we write R^Γ to indicate that we are talking about the relation R belonging to the structure Γ. If a relational structure Δ can be obtained from a relational structure Γ by removing some of the relations from the structure and the signature of Γ, then Δ is called a *reduct* or *fragment* of Γ, and Γ is called an *expansion* of Δ.

A *homomorphism* from a structure Γ over signature τ to another structure Γ' over the same signature τ is a mapping h from D_Γ to $D_{\Gamma'}$ such that for each relation symbol $R \in \tau$, and for every tuple $(a_1, \ldots, a_k) \in R^\Gamma$, it holds that $(h(a_1), \ldots, h(a_k)) \in R^{\Gamma'}$. A homomorphism h from Γ to Γ' is called *strong* if for each relation symbol R and for every tuple $(a_1, \ldots, a_k) \in (D_\Gamma)^k$, it holds that: $(a_1, \ldots, a_k) \in R^\Gamma$ if and only if $(h(a_1), \ldots, h(a_k)) \in R^{\Gamma'}$. An *isomorphism* from a structure Γ to a structure Γ' is a strong homomorphism that is bijective, and an *automorphism* is an isomorphism from a structure Γ to itself. Note that in this paper, particularly in Section 6, we will sometime use the tuple notation $(D, R_1, \ldots, R_k)$ to denote a relational structure; the first element D denotes the universe and the R_i denote relations over D.

Let Γ be a structure with relational signature τ. If $E \subseteq (D_\Gamma)^2$ is an equivalence relation, we write Γ/E for the *factor structure* Δ with signature τ, defined as follows. The domain of Δ consists of the equivalence classes of E, and a tuple $(C_1, \ldots, C_k)$ is in a k-ary relation R^Δ iff there is $(c_1, \ldots, c_k) \in R^\Gamma$ with $c_1 \in C_1, \ldots, c_k \in C_k$.

For a subset S of D_Γ, we write $\Gamma[S]$ for the substructure of Γ induced by S. In this paper, substructure always means *induced substructure*, as in [Hod97]. An *embedding* of a τ-structure Γ in a τ-structure Δ is a mapping $f : D_\Gamma \to D_\Delta$ that is an isomorphism between Γ and $\Delta[f(D_\Gamma)]$.

The constraint satisfaction problem. A *constraint language* is simply a relational structure; we typically refer to a relational structure Γ as a constraint language when we are interested in the constraint satisfaction or the quantified constraint satisfaction problem for Γ, which are defined below.

A first-order formula ϕ is called a τ-*formula* if all symbols in ϕ are either the standard logical symbols $\{\exists, \forall, \wedge, \vee, \neg, =, \mathit{false}\}$, variable symbols, or from τ. A first-order τ-formula ϕ is called a τ-sentence if ϕ has no free variables. A first-order τ-formula is called *primitive positive* (for short, *pp*) if it is of the form

$$\exists x_1, \ldots, x_n . \psi_1 \wedge \cdots \wedge \psi_m \quad ,$$

where each ψ_i is an atomic τ-formula (a formula *false*, $x = y$, or $R(x_1, \ldots, x_k)$ for $R \in \tau$) that can contain both free variables and quantified variables from $\{x_1, \ldots, x_n\}$. The *constraint satisfaction problem (CSP) for Γ* is the computational problem to determine for a given primitive positive τ-sentence Φ whether Φ is true in Γ. We sometimes refer to the atomic formulas ψ_i in an instance of a CSP as *constraints*. It is well-known and easy to see (we refer to [KV98]) that the constraint satisfaction problem for Γ can also be formalized as the problem to determine for a given finite τ-structure whether there exists a homomorphism from the structure to Γ. We mostly use the logic formulation, but might sometimes also use the homomorphism formulation when this is convenient.

A first-order τ-formula is *conjunctive positive* if it has the form

$$Q_1 v_1 \ldots Q_n v_n (\psi_1 \wedge \ldots \wedge \psi_m),$$

where each Q_i is a quantifier from $\{\forall, \exists\}$, and each ψ_i is an atomic τ-formula that can contain both free variables and quantified variables from $\{v_1, \ldots, v_n\}$. The *quantified constraint satisfaction problem* for Γ, denoted by $QCSP(\Gamma)$, is the problem of deciding for a given conjunctive positive τ-sentence whether or not the formula is true in Γ. Note that both the universal and existential quantification is understood to take place over the entire universe of Γ.

Let R be a k-ary relation and let Γ be a τ-structure. We say that R has a *pp-definition* (or is *pp-definable*) in Γ if there exists a pp-formula ϕ with free variables $x_1, \dots, x_k$ such that $R(x_1, \dots, x_k) = \phi(x_1, \dots, x_k)$. Analogously, we define the concept of a *cp-definition.* The constraint language that contains all pp-definable relations in Γ is denoted by $\langle \Gamma \rangle$.

Amalgamation. As we have already mentioned in the introduction, it turns out that many CSPs in temporal and spatial reasoning (in particular, if they concern so-called *qualitative* formalisms [RN07]) can be formulated with constraint languages that are ω*-categorical.* A relational structure Γ is called ω*-categorical* if all countable models of the first-order theory[1] of Γ are isomorphic to Γ.

To formulate computational problems for temporal and spatial calculi as constraint satisfaction problems with ω-categorical constraint languages, the following concept is very powerful. 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 [Fra86]). 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 $e_1 : A \to B_1$ and $e_2 : A \to B_2$ there exists $C \in \mathcal{C}$ and embeddings $f_1 : B_1 \to C$ and $f_2 : B_2 \to C$ such that $f_1 e_1 = f_2 e_2$.

A structure is *homogeneous* (sometimes called *ultra-homogeneous* [Hod97]) if every isomorphism between finite substructures of Γ can be extended to an automorphism.

Theorem 1 (Fraïssé [Fra86]). *A countable class $\mathcal{C}$ of finite relational structures with countable signature is the age of a countable homogeneous structure Γ if and only if $\mathcal{C}$ is an amalgamation class. In this case Γ is up to isomorphism unique and called the* Fraïssé-limit *of $\mathcal{C}$.*

Homogeneous structures provide a rich source of ω-categorical structures.

Proposition 2 (see e.g. [Hod97]). *A countable homogeneous structure Γ over a finite relational signature is ω-categorical.*

The following is well-known; a proof can be found in [BD06]. The second part of the proposition is not proven there, but can be shown analogously.

Proposition 3. *Let Δ be a countable structure and let Γ be ω-categorical. Then Δ homomorphically maps to Γ if and only if all finite (induced, or equivalently weak) substructures of Δ homomorphically map to Γ. The structure Δ injectively homomorphically maps to Γ if and only if all finite induced substructures of Δ injectively homomorphically map to Γ.*

Polymorphisms. The *(direct-, 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 relations $R \in \tau$ the relation $R\big((x_1, y_1), \dots, (x_k, y_k)\big)$ holds in $\Gamma_1 \times \Gamma_2$ iff $R(x_1, \dots, x_k)$ holds in Γ_1 and $R(y_1, \dots, y_k)$ holds in Γ_2. Homomorphisms from $\Gamma^k = \Gamma \times \dots \times \Gamma$ to Γ are called *polymorphisms* of Γ. If $f : D^k \to D$ is a polymorphism of a relational structure (D, R), we also say that f *preserves* the relation R (and otherwise f *violates* the relation R).

The set of all polymorphisms of Γ gives rise to an algebra $\mathrm{Al}(\Gamma)$, defined as follows. The domain of the algebra equals the domain of Γ, and the algebra has a function (and an associated function symbol) for each polymorphisms of Γ. For finite domain constraint languages, Bulatov et al. [BKJ05] give a detailed exposition of this concept. A property of the algebra $\mathrm{Al}(\Gamma)$ is that it is *locally closed*, i.e., if f is a k-ary operation such that for every finite subset A of the domain there is a k-ary operation g in $\mathrm{Al}(\Gamma)$ such that $f(x) = g(x)$ for all $x \in A^k$, then f is also an operation in $\mathrm{Al}(\Gamma)$. (Note that this property is non-trivial only in the case that Γ has an infinite domain.)

We say that a polymorphism f of an ω-categorical structure Γ is *oligopotent* if the diagonal of f, that is, the function $f(x, \dots, x)$, is contained in the locally closed clone generated by the automorphisms of Γ.

The importance of polymorphisms stems from the following powerful preservation theorem that characterize primitive positive definability over ω-categorical structures.

[1] the first-order theory of a τ-structure Γ is the set of all τ-sentences that are true in Γ.

Theorem 4 (of [BN06]). *Let Γ be an ω-categorical structure. Then a relation R is pp-definable in Γ if and only if it is preserved by all polymorphisms of Γ.*

3 Consistency and QNUFs

Consistency. Establishing 2- and 3-consistency (defined below) are the most prominent algorithmic techniques for constraint satisfaction, due to their wide applicability in practical applications. On the other hand, the question whether a (quantified) constraint satisfaction problem can be solved in polynomial time by consistency techniques leads to challenging theoretical problems. We first introduce the basic definitions, and then present the mentioned connection to universal algebra.

A *projective homomorphism* from A to Γ is a mapping f from a subset S of D_A to Γ such that for every $(t_1, \ldots, t_k) \in R^A$ there exists a tuple $(b_1, \ldots, b_k) \in R^\Gamma$ such that $f(t_i) = b_i$ for all $t_i \in S$.

Definition 5. *Let Γ be a relational structure. An instance A of CSP(Γ) is called* k-consistent *if for every size k subset $S = \{v_1, \ldots, v_k\}$ of the elements of A and every projective homomorphism h from A to Γ defined on $\{v_1, \ldots, v_{k-1}\}$ there exists an extension of h defined on S that is a projective homomorphism from A to Γ. An instance A of CSP(Γ) is called* strongly k-consistent *if it is l-consistent for all $l \leq k$.*

An important feature of strong k-consistency is that for every fixed k and for every finite or ω-categorical structure Γ there is an algorithm that *establishes strong k-consistency* for a given instance A of CSP(Γ), i.e., computes a strongly k-consistent instance B that is logically equivalent to A; to formalize this idea, we need the following definition.

Definition 6. *Let Γ be constraint language, and let Δ be an expansion of Γ by finitely many primitive positive definable relations of Γ (in particular, Γ and Δ are defined on the same domain D). We say that an instance B of CSP(Δ) is a* strongly k-consistent variant of an instance *A of CSP(Γ) if B is strongly k-consistent, has the same set of variables $V(A)$ as A, and every mapping from $V(A)$ to D is a solution for B if and only if it is a solution for A.*

In Proposition 7 below, we state a fact that is well-known for constraint languages over a finite domain; for ω-categorical constraint languages a proof can be found in [BD06]. We would like to remark that the algorithms that are used in the proof of Proposition 7 can be formulated as Datalog programs (Datalog programs can be seen as Prolog programs without function symbols, and are a well-studied concept in Database theory and finite model theory; see e.g. [AHV95,EF99]).

Proposition 7. *Let Γ be a finite or an ω-categorical structure over a finite relational signature. Then for every k there is a polynomial-time algorithm that computes a strongly k-consistent variant of a given instance A of CSP(Γ).*

Proof. This is a well-known fact for constraint languages Γ over a finite domain. For ω-categorical structures Γ, the statement follows from [BD06], since the evaluation of the *canonical (l, k)-Datalog program* on A produces a k-consistent variant of A, where k is the maximal arity of the relations from Γ. □

We would like to remark that in artificial intelligence and in the theory of relation algebras, 3-consistency is usually called *path-consistency.* Note that if a strongly k-consistent variant B of an instance A of CSP(Γ) contains a constraint with a relation symbol that denotes an empty relation in Γ, then A does not have a solution. The converse need not be true in general.

Definition 8. *A constraint language Γ has* width k *when: a strongly $k+1$-consistent instance A of CSP(Γ) has a solution if and only if A does not contain a constraint with a relation symbol that denotes an empty relation in Γ.*

Global consistency. Some constraint languages Γ have the strong property that every strongly k-consistent instance of CSP(Γ) is automatically *globally consistent* (from [Fre82], see Definition 9 below).

Definition 9. *An instance A of CSP(Γ) is called* globally consistent *iff it is k-consistent for all $1 \leq k \leq |A|$.*

It follows easily from Proposition 7 that if Γ has the property that every strongly k-consistent instance of CSP(Γ) is globally consistent, then CSP(Γ) can be solved in polynomial time.

But note that Γ might have width k, while at the same time strong k-consistency does not establish global consistency. A well-known example over a two-element domain are boolean constraint languages that are preserved by the maximum operation (the relations in such a constraint language can be defined by Horn clauses). A well-known example over an infinite domain is CSP($\mathbb{Q}, \leq, \neq$). Vilain, Kautz and van Beek [MVvB89] have shown that this CSP has width 2, but establishing strong 3-consistency does not imply global consistency (however, establishing strong 5-consistency implies global consistency, due to a result by Koubarakis [Kou97]).

Quasi near-unanimity functions. We now present the connection of the "local-to-global" property to universal algebra mentioned in the introduction.

Definition 10. *A k-ary function $f : D^k \rightarrow D$, for $k \geq 3$, is called a* quasi near-unanimity function *(short, a* QNUF*), if it satisfies $f(x, \ldots, x, y) = f(x, \ldots, x, y, x) = \cdots = f(y, x, \ldots, x) = f(x, \ldots, x)$ for all $x, y \in D$.*

As an example, consider the structure ($\mathbb{Q}, \leq$), and the operation *median*, which is the ternary function that returns the median of its three arguments. More precisely, for three elements x, y, z from $\mathbb{Q}$, suppose that $\{x, y, z\} = \{a, b, c\}$, where $a \leq b \leq c$. Then $median(x, y, z)$ is defined to have value b. It is easy to verify that *median* is a ternary quasi near-unanimity function, and that it is a polymorphism of ($\mathbb{Q}, \leq, <$).

The following two results are of central importance in this paper, and they generalize well-known facts for finite structures Γ [BP74, FV99, JCC98]. The statement of item (3) in Theorem 8 in [BD06] mistakenly misses the word *oligopotent*; this has been corrected in the journal version [BD08]. Also note that CSP(Γ) has *strict width* k in the terminology of [BD06, BD08] if and only if every stronlgy k-consistent instance of CSP(Γ) is globally consistent.

Theorem 11 (Theorem 8 of [BD06], Theorem 13 of [BD08]). *Let $k \geq 3$. An ω-categorical structure Γ of bounded maximal arity has a k-ary oligopotent QNU-polymorphism if and only if every strongly k-consistent instance of CSP(Γ) is globally consistent.*

There is yet another characterization of constraint languages having a QNU-polymorphism. We say that a constraint language Γ is *k-decomposable* if every relation in Γ can be defined by a conjunction of at most k-ary primitive positive definable relations in Γ.

Theorem 12 (of [BC07a]). *Let $k \geq 3$. An ω-categorical structure Γ has a k-ary oligopotent QNU-polymorphism if and only if Γ is $(k-1)$-decomposable.*

Innermost quantifier elimination. In this section, we show that if an ω-categorical constraint language has a surjective oligopotent QNU polymorphism, then $QCSP(\Gamma)$ can be solved in polynomial time. This was already known for constraint languages over finite domains, see e.g. [Che]. The same idea that was applied there can be applied for ω-categorical constraint languages.

Theorem 13. *Let Γ be an ω-categorical constraint language with a surjective oligopotent QNU polymorphism. Then $QCSP(\Gamma)$ is in P.*

Proof. Let r be the maximal arity of the relations in Γ. In this proof we assume that Γ contains all cp-definable relations over Γ of arity at most $r-1$. This assumption is justified because expanding Γ by cp-definable relations clearly preserves the set of surjective polymorphisms.

The algorithm eliminates the variables of a given instance ϕ of $QCSP(\Gamma)$, starting from the innermost variable v_n, and computes a conjunctive-positive sentence ϕ' that is equivalent to ϕ, without the variable v_n. To do so, the algorithm first establishes strong k-consistency on the quantifier-free part of ϕ (viewing this part as an instance of the CSP; here we use Proposition 7 and the assumption that Γ is ω-categorical). If a constraint for the empty relation was derived during establishing strong k-consistency, then the algorithm reports that the instance is false over Γ. If v_n is existentially quantified, the algorithm replaces each constraint containing the variable v_n by the constraint obtained by projecting away the coordinate corresponding to v_n. This leads to an equivalent instance, because the instance was strongly $n-1$-consistent. If v_n is universally quantified, then the universal quantifier distributes over all the conjunctively combined constraints ψ. We replace each constraint ψ containing v_n by the constraint defined by $\forall v_n.\psi$. Clearly, we can then proceed with the next variable v_{n-1} in the same fashion. This process can be iterated, and if the algorithm never report that the instance is false, it eventually shows that the sentence is true. □

4 Posets

In this section, we develop some general results on QNU polymorphisms of posets, which we will utilize in the following sections. We use $[k]$ to denote the first k natural numbers, $\{1, \ldots, k\}$.

Definition 14. *Relative to a poset $(D, \leq)$, we say that $b \in D$ is the* middle value *of a tuple $t = (t_1, \ldots, t_k)$ (with $k \geq 3$) if there exists a permutation $\pi : [k] \to [k]$ such that $t_{\pi(1)} \leq \cdots \leq t_{\pi(k)}$ and $b = t_{\pi(2)} = \cdots = t_{\pi(k-1)}$. In this situation, we call $t_{\pi(1)}$ the* low value *of t, and $t_{\pi(k)}$ the* high value.

Note that not every tuple has a middle value, but when a tuple does have a middle value, it has a unique middle value.

Definition 15. *Relative to a poset $(D, \leq)$, we say that $b \in D$ is the* main value *of a tuple $t = (t_1, \ldots, t_k)$ (with $k \geq 3$) if either b is the middle value of the tuple t, or b occurs in (at least) $k-1$ coordinates of t.*

Again, not every tuple has a main value, but when a tuple has a main value, it is unique.

Using the definition of main value, we can now define an equivalence relation on the set of all tuples of length k. Let $(D, \leq)$ be a poset and let t, t' be two tuples of length k. We write $t \equiv_m t'$ if either $t = t'$ or t and t' have the same main value. We show that any QNUF on a poset must map equivalent tuples to the same point.

Proposition 16. *Let $\Gamma = (D, \leq)$ be a poset, $k \geq 3$, and let $f : \Gamma^k \to \Gamma$ be a QNU polymorphism of Γ. For two tuples t, t' of length k, if $t \equiv_m t'$, then $f(t) = f(t')$.*

Proof. Let t, t' be two tuples of length k such that $t \equiv_m t'$. If $t = t'$, then clearly $f(t) = f(t')$. If t and t' both have the same main value $b \in D$, we show that $f(t) = f(b, \ldots, b) = f(t')$. We prove that $f(t) = f(b, \ldots, b)$; the proof that $f(t') = f(b, \ldots, b)$ is identical. If t has $k-1$ coordinates equal to b, then $f(t) = f(b, \ldots, b)$ by the QNUF identities. So suppose that t has middle value b, let $\pi : [k] \to [k]$ be a witnessing permutation, and let $a = t_{\pi(1)}$ and $c = t_{\pi(k)}$, so that we have $a = t_{\pi(1)} \leq \cdots \leq t_{\pi(k)} = c$. For the sake of notation, we assume that π is the identity; the proof in the general case is identical, but with the coordinates permuted according to π. We have $t = (a, b, \ldots, b, c)$. Since f is a polymorphism of Γ, we have $f(t) \leq f(b, b, \ldots, b, c)$ and $f(a, b, \ldots, b, b) \leq f(t)$. On the other hand, since f is a QNUF we have $f(b, b, \ldots, b, c) = f(b, b, \ldots, b, b) = f(a, b, \ldots, b, b)$. It follows that $f(t) \leq f(b, b, \ldots, b, c) = f(b, b, \ldots, b, b) = f(a, b, \ldots, b, b) \leq f(t)$ and since Γ is a poset, we have $f(t) = f(b, b, \ldots, b, b)$. □

We now show that, under some mild assumptions, a poset has QNU polymorphisms of all arities which only identify together tuples that are equivalent under our equivalence relation $\equiv_m$.

Theorem 17. *Let $\Gamma = (D, \leq)$ be an ω-categorical poset such that there is an injective homomorphism from $(\mathbb{N}, \leq)$ to Γ. For all $k \geq 3$, there exists a k-ary QNU polymorphism $f : \Gamma^k \to \Gamma$ such that for all tuples $t, t' \in D^k$, it holds that $t \equiv_m t'$ if and only if $f(t) = f(t')$. Moreover, for any relation $R \subseteq D^m$ definable by a disjunction of disequalities over m variables and $k \geq 2m + 1$, the resulting polymorphism $f : \Gamma^k \to \Gamma$ preserves R.*

We say that a relation $R \subseteq D^m$ is definable by a disjunction of disequalities over m variables if there exists a disjunction of disequalities ϕ over variables $\{v_1, \ldots, v_m\}$ such that $R(v_1, \ldots, v_m) \equiv \phi(v_1, \ldots, v_m)$. As an example, consider the relation R such that $R(v_1, v_2, v_3, v_4, v_5) \equiv (v_1 \neq v_2 \vee v_2 \neq v_3 \vee v_4 \neq v_5)$. Here, we assume that each of the variables $v_1, \ldots, v_m$ appears in some disequality. Disjunctions of disequalities were studied in the context of temporal reasoning by Koubarakis [Kou01], and we may mention also that they appear in the classification result of [BC07b].

We need to present some intermediate results before proving Theorem 17. The following definition will be wieldy.

Definition 18. *Relative to a poset $(D, \leq)$, we say that a tuple $(t_1, \ldots, t_k)$ is j-*superior *to $d \in D$ if the tuple contains j coordinates that are greater than or equal to d in the poset, that is, there is a subset $S \subseteq [k]$ of size j such that $t_i \geq d$ for all $i \in S$. Similarly, we say that a tuple $(t_1, \ldots, t_k)$ is j-*inferior *to $d \in D$ if there is a subset $S \subseteq [k]$ of size j such that $t_i \leq d$ for all $i \in S$.*

Proposition 19. *Let $(D, \leq)$ be a poset, let $k \geq 3$ and let $j \in [k] \setminus \{1, k\}$. Let $t, t' \in D^k$ be two tuples such that $t \equiv_m t'$ and let $d \in D$ be a value. The tuple t is j-superior to d if and only if the tuple t' is j-superior to d. Similarly, the tuple t is j-inferior to d if and only if the tuple t' is j-inferior to d.*

Proof. Suppose that $t \equiv_m t'$. By symmetry, it suffices to show that if t is j-superior to d, then t' is j-superior to d. If $t = t'$, the claim is clear. So suppose that t and t' both have main value b. It suffices to show that $b \geq d$, since (in virtue of having main value b) t' is $(k-1)$-superior to b, and will thus be j-superior to b. Let $S \subseteq [k]$ be a set of size j such that $t_i \geq d$ for all $i \in S$. Let $T \subseteq [k]$ be the coordinates where t takes on its main value. If $S \cap T$ is non-empty, then for any $i \in S \cap T$ we have $b = t_i \geq d$. Now consider the case where $S \cap T$ is empty. We know $|S| \geq 2$ and $|T| \geq k - 2$, so in fact we have $|S| = 2$ and $|T| = k - 2$. We have that $\{t_s : s \in S\}$ contains both the high value and low value a of t, implying that $a \geq d$. Since $b \geq a$, we conclude that $b \geq d$. $\square$

Lemma 20. *Let $(D, \leq)$ be a poset and define $(E, \preceq)$ as $(D, \leq)^k / \equiv_m$ (for some $k \geq 3$). The relation $\preceq$ has no non-trivial cycles, that is, there are no distinct elements $T_1, \ldots, T_n \in E$ such that $T_1 \preceq \cdots \preceq T_n \preceq T_1$ and $n \geq 2$.*

Notice that, in light of Proposition 19, we may speak of an element $T \in E$ being j-superior to d (for $j \in [k] \setminus \{1, k\}$): by that proposition, either all tuples in T are j-superior to d, or none of them are. This same fact holds for the notion of j-inferiority.

Proof. We prove this by contradiction. Assume that $T_1 \preceq \cdots \preceq T_n \preceq T_1$ and $n \geq 2$, with the T_i pairwise distinct. If none of the T_i have a main value, then we have a contradiction since $(D, \leq)^k$ has no non-trivial cycles. So assume that T_1 has main value b. Every tuple in T_1 is $(k-1)$-superior to b. Since $T_1 \leq T_2$, every tuple in T_2 is $(k-1)$-superior to b. Similarly, since $T_2 \preceq \cdots \preceq T_n \preceq T_1$, every tuple in T_2 is $(k-1)$-inferior to b. Let t be any tuple in T_2, let $S \subseteq [k]$ be a $k-1$ size subset such that $t_j \geq b$ for all $j \in S$, and let $I \subseteq [k]$ be a $k-1$ size subset such that $t_j \leq b$ for all $j \in I$. Observe that for all $j \in (S \cap I)$, we have $t_j = b$. If $S = I$, then $t_i = b$ for all $j \in S = I$ and t_i has b as its main value, contradicting the distinctness of T_1 and T_2. If $S \neq I$, then let $s, i \in [k]$ be coordinates such that $S = \{s\} \cup (S \cap I)$, and $I = \{i\} \cup (S \cap I)$. Notice that $\{s\}$, $(S \cap I)$, and $\{i\}$ form a partition of $[k]$. We have $t_i \leq b \leq t_s$ and $b = t_j$ for all $j \in (S \cap I)$, so t has b as its middle value and hence as its main value, contradicting the distinctness of T_1 and T_2. $\square$

Proof (Theorem 17). We want to show that there is a homomorphism $f : \Gamma^k \to \Gamma$ that is a QNUF such that $t \equiv_m t'$ if and only if $f(t) = f(t')$. Let $(E, \preceq)$ be the structure $(D, \leq)^k / \equiv_m$. It suffices to show that there is an injective homomorphism from $(E, \preceq)$ to $(D, \leq)$; notice that $\equiv_m$ equates tuples that are related by the QNUF identities. By Proposition 3, it suffices to show that every finite induced substructure of $(E, \preceq)$ has an injective homomorphism to $(D, \leq)$. Let E' be a finite subset of E. We have shown in Lemma 20 that $(E, \preceq)$ contains no non-trivial cycles, implying that $(E', \preceq)$ contains no non-trivial cycles. Thus, there is an injective homomorphism h from $(E', \preceq)$ to $(\mathbb{N}, \leq)$. Composing this with the injective homomorphism from $(\mathbb{N}, \leq)$ to Γ, we obtain an injective homomorphism from $(E', \preceq)$ to Γ.

For the claim concerning disjunctions of disequalities, let $R \subseteq D^m$ be definable by a disjunction of disequalities. Assume $k \geq 2m + 1$ and let $t_1, \ldots, t_k \in R$ be tuples. We want to show that $f(t_1, \ldots, t_k) \in R$. Since R is definable by a disjunction of disequalities, it suffices to show that there exists a tuple t_l (with $l \in [k]$) such that for all pairs of coordinates $i, j \in [m]$, $f(t_{1i}, \ldots, t_{ki}) = f(t_{1j}, \ldots, t_{kj})$ implies $t_{li} = t_{lj}$. For each $i \in [m]$, if $(t_{1i}, \ldots, t_{ki})$ has a main value, let $S_i \subseteq [k]$ be coordinates where that tuple takes on the main value, and let $S_i = [k]$ otherwise. Notice that for each $i \in [m]$ we have $|S_i| \geq k - 2$. Combining this with the assumption that $k \geq 2m + 1$, we obtain that $\cap_{i \in [m]} S_i$ is non-empty. Let l be any value in $\cap_{i \in [m]} S_i$. Now, let $i, j \in [m]$ be coordinates such that $f(t_{1i}, \ldots, t_{ki}) = f(t_{1j}, \ldots, t_{kj})$. By the first part of this theorem, we have $(t_{1i}, \ldots, t_{ki}) \equiv_m (t_{1j}, \ldots, t_{kj})$. If the these two tuples are equal, it clearly holds that $t_{li} = t_{lj}$. If these two tuples have the same main value, we obtain $t_{li} = t_{lj}$ from our choice of l. □

The following proposition complements Theorem 17, showing that the given arity bound for preserving disjunctions of disequalities is optimal.

Proposition 21. *Let $(D, \leq)$ be a poset with three distinct values $a, b, c \in D$ such that $a \leq b \leq c$. Let $R \subseteq D^m$ be a relation definable by a disjunction of disequalities. There is no QNUF of arity $2m$ preserving R.*

Recall that in speaking of a relation definable by a disjunction of disequalities, we assume that each variable appears in at least one disequality.

Proof. We define $2m$ tuples $s_1, t_1, \ldots, s_m, t_m \in D^m$ in the following way. For all $i \in [m]$, define s_i to be the tuple equal to a in coordinate i, and b otherwise; and, define t_i to be the tuple equal to c in coordinate i, and b otherwise. All of these tuples are in R: in the tuples s_i and t_i, the ith coordinate is different from the rest, so the disequality involving the ith coordinate is satisfied. On the other hand, looking at any coordinate i, we have that $(s_{1i}, t_{1i}, \ldots, s_{mi}, t_{mi})$ is equal to b everywhere except at s_{ii} and t_{ii}, where it is a and c, respectively; thus $(s_{1i}, t_{1i}, \ldots, s_{mi}, t_{mi})$ has b as middle value, and by Proposition 16, $f(s_{1i}, t_{1i}, \ldots, s_{mi}, t_{mi})$ is equal to $f(b, \ldots, b)$. It follows that $f(s_1, t_1, \ldots, s_m, t_m) = (f(b, \ldots, b), \ldots, f(b, \ldots, b))$ which is not in R. □

5 The Point Algebra, fragments, and expansions

The Point Algebra is one of the most fundamental formalisms for temporal reasoning. The corresponding constraint language is

$$\Gamma^{PA} = (\mathbb{Q}; \leq, <, >, \geq, \neq, =, \mathbb{Q}^2, \emptyset)$$

with the obvious interpretation to the eight binary relation symbols of this structure. It is well-known that this structure is homogeneous and ω-categorical [Hod97]. We would like to remark that formally, the Point Algebra is a relation algebra [LM94,Due05], which has a natural representation by the relational structure Γ^{PA}. The notion of relation algebra and representations of relation algebras are not of importance in this paper, and it suffices to study Γ^{PA} for our purposes [LM94, Due05].

In this section, we study Γ^{PA}, its fragments, and its expansions from the algebraic viewpoint. We begin with Γ^{PA} itself. Koubarakis [Kou97] showed that strong 5-consistency implies global consistency for CSP(Γ^{PA}). In combination with Theorem 11, this implies the following.

Theorem 22. *The point algebra has an oligopotent QNU polymorphism of arity 5, but no oligopotent QNU polymorphism of arity 4.*

We observe that we may obtain an alternative proof of this theorem from results in the previous section. Note that via Theorem 11, this also leads to an alternative proof of the theorem of Koubarakis.

Proof (Theorem 22). It is clear that all relations of Γ^{PA} have a primitive positive definition in $(\mathbb{Q}, \leq, \neq)$, and we therefore work with this constraint language instead of Γ^{PA}. The existence of a 5-ary QNU polymorphism f then follows from Theorem 17. The polymorphism f must be oligomorphic, because the unary operation defined by $f(x, \ldots, x)$ preserves $\leq$ and $\neq$, and it is easy to see that such operations are interpolated by the automorphisms of $(\mathbb{Q}, \leq, \neq)$. The non-existence of a 4-ary QNU polymorphism follows from Proposition 21. □

It can be shown that the image of the constructed polymorphism is a dense subset of the rational numbers. It is then not hard to see from local closure arguments and the homogeneity of Γ^{PA} that Γ^{PA} also has a *surjective* 5-ary QNU polymorphism. Therefore, Theorem 13 shows that $QCSP(\Gamma^{PA})$ is in P. We will see a stronger algorithmic result for the QCSP over Γ^{PA} at the end of this section.

5.1 Fragments

In this subsection we give algebraic characterizations of fragments of the point algebra. Thereby we provide illustrations of simple but fundamental QNU operations. We believe that the operational descriptions of these fragments are important for future research, for example for a systematic investigation of the class of all temporal constraint languages up to pp-definability.

If $f_1, f_2, \ldots$ are operations from $\mathbb{Q}^k \to \mathbb{Q}$, then $Inv_{\mathbb{Q}}(f_1, f_2, \ldots)$ denotes the set of relations with a first-order definition in $(\mathbb{Q}, <)$ that are preserved by all operations $f_1, f_2, \ldots$. The three binary relations $\mathbb{Q}^2$, $\emptyset$, and $=$ are trivial from our standpoint, as we study pp-definability; other than these relations, there are exactly five binary first-order definable relations: $>$, $<$, $\geq$, $\leq$, and $\neq$ (this follows from the fact that every first-order definable relation over a relational structure Γ is preserved by the automorphisms of Γ, and that the 8 relations mentioned above are precisely the binary relations that are preserved by the automorphisms of $(\mathbb{Q}, <)$).

We first consider two fragments with two binary relations.

Theorem 23. $\langle(\mathbb{Q}, <, \leq)\rangle = Inv_{\mathbb{Q}}(median)$.

Recall that *median* (defined in Section 3) is a QNUF and preserves $<$ and $\leq$.

Proof. The $\subseteq$ direction is straightforward, so we prove the $\supseteq$ direction. Because $Inv_{\mathbb{Q}}(median)$ only contains relations that are first-order definable in $(\mathbb{Q}, <)$, it is clearly ω-categorical. By Theorem 12, every relation in $Inv_{\mathbb{Q}}(median)$ is 2-decomposable, and it suffices to show that for every binary relation, if the relation is preserved by *median*, then it is pp-definable over $(\mathbb{Q}, <, \leq)$. This is obvious for all relations except for $\neq$. We show that $\neq$ is not preserved by *median*. Let $a < b < c$ be three values in $\mathbb{Q}$. We have $a \neq b$, $b \neq c$, and $c \neq a$, but $median(a, b, c) = median(b, c, a)$. □

Theorem 24. *There exists a QNUF* $f : \mathbb{Q}^3 \to \mathbb{Q}$ *of arity* 3 *such that* $\langle(\mathbb{Q}, <, \neq)\rangle = Inv_{\mathbb{Q}}(f)$.

For the proof of this theorem, we will use the equivalence relation $\equiv$ where two tuples t, t' of arity $k \geq 3$ are equal ($t \equiv t'$) if either $t = t'$ or there exists a value b that occurs in at least $k - 1$ coordinates for each of t, t'. Intuitively, $t \equiv t'$ if $f(t) = f(t')$ for any QNUF f.

Proof. Set $(E, \preceq)$ to be the structure $(\mathbb{Q}, <)^3/\equiv$, and define $\preceq'$ to be the binary relation $\preceq \cup\{((2, 3, 4), (1, 3, 4))\}$. We first show that $\preceq'$ has no cycles. Notice that the property of being 2-superior to a value b applies to either all or none of the tuples in an $\equiv$-equivalence class, and so we apply this property to $\equiv$-equivalence classes. Notice also that $(2, 3, 4)$ is 2-superior to a value b

if and only if $(1,3,4)$ is 2-superior to a value b. Assume that $T_1 \preceq \cdots \preceq T_n \preceq T_1$ and $n \geq 2$, with the T_i pairwise distinct.

First suppose that for all T_i we have $|T_i| = 1$. It is clear that only one of the two tuples $(2,3,4)$, $(1,3,4)$ occurs in the sequence $T_1, \ldots, T_n$, since we do not have $3 < 3$ nor $4 < 4$. Hence, our cycle yields a cycle in $(\mathbb{Q}, <)^3$, a contradiction. Now assume $|T_1| > 1$. Let $b \in \mathbb{Q}$ be the value such that $(b, \ldots, b) \in T_1$. We have that T_1 is $(k-1)$-superior to b; it follows that all other T_i are $(k-1)$-superior to b. Select tuples $t_n \in T_n$ and $t_1 \in T_1$ such that $t_n < t_1$. By definition of T_1, the tuple t_n is strictly less than b in at least $(k-1)$ coordinates, contradicting that T_n is $(k-1)$-superior to b. We have thus shown that $\preceq'$ has no cycles.

We claim that there is an injective homomorphism f from $(E, \preceq')$ to $(\mathbb{Q}, <)$. Every finite induced substructure of $(E, \preceq')$ has an injective homomorphism to $(\mathbb{Q}, <)$, as $\preceq'$ has no cycles, so the claim follows from Proposition 3.

It is straightforward to verify that f preserves $\neq$. It preserves $<$ (and hence $>$) by construction, and so we have $\langle(\mathbb{Q}, <, \neq)\rangle \subseteq Inv_{\mathbb{Q}}(f)$. On the other hand, f but does not preserve $\leq$ as $f(1,3,4) > f(2,3,4)$. Thus by invoking 2-decomposability of the relations in $Inv_{\mathbb{Q}}(f)$, we obtain the $\supseteq$ direction. □

Theorem 25. *$\langle(\mathbb{Q}, \leq)\rangle = Inv_{\mathbb{Q}}(median, 0)$. Here, 0 denotes the unary operation that is always equal to the constant 0.*

Proof. The $\subseteq$ direction is straightforward. For the $\supseteq$ direction, by 2-decomposability of $Inv_{\mathbb{Q}}(median, 0)$, we need to show that $\leq$ and $\geq$ are the only (non-trivial) binary relations preserved by both *median* and 0. It is clear that none of $<$, $>$, $\neq$ are preserved by 0. □

Theorem 26. *$\langle(\mathbb{Q}, <)\rangle = Inv_{\mathbb{Q}}(median, f)$ where f is the operation from Theorem 24.*

Proof. The $\subseteq$ direction is straightforward. For the $\supseteq$ direction, by 2-decomposability of $Inv_{\mathbb{Q}}(median, f)$, we need to show that $<$ and $>$ are the only (non-trivial) binary relations preserved by both *median* and f. We refer the reader to the proofs of Theorems 23 and 24. □

Theorem 27. *$\langle(\mathbb{Q}, \neq)\rangle = Inv_{\mathbb{Q}}(-f)$ where f is the operation from Theorem 24.*

Proof. The $\subseteq$ direction follows from the fact that f preserves $\neq$. For the $\supseteq$ direction, by 2-decomposability of $Inv_{\mathbb{Q}}(-f)$, we need to show that $\neq$ is the only non-trivial binary relation preserved by $-f$. Since f preserves $<$, we have $f(1,2,3) < f(4,5,6)$. It follows that $-f(1,2,3) > -f(4,5,6)$ and $-f$ does not preserve $<$ nor $\leq$. □

5.2 Expansions

Koubarakis [Kou97] showed that establishing strong $2k+1$-consistency implies global consistency for the expansion of Γ^{PA} by disjunctions of disequalities on at most k variables, but strong $2k$ consistency is not sufficient. By Theorem 17 and Proposition 21 of Section 4 in combination with Theorem 11, we have a new proof of this theorem.

Now let us consider the constraint language Γ that is the expansion of the point algebra with all disjunctions of disequalities. This is a constraint language with infinitely many relations. What is the complexity of Γ? We can *not* derive the tractability of Γ using a QNUF, as it is immediate from Proposition 21 that Γ has no QNUF polymorphism. We show here that this expansion of the point algebra is in fact tractable[2]; indeed, we show that the QCSP over this constraint language is solvable in NL.

Theorem 28. *Let Γ be the expansion of the point algebra having arbitrary disjunctions of disequalities. The problem QCSP(Γ) is in NL. (Note that we assume that constraints are presented syntactically, for instance, as "$x \leq y$" or "$x \neq y \vee y \neq z$".)*

[2] In the constraint satisfaction literature, this form of tractability of the CSP where the instances might contain arbitrary constraints built from an infinite constraint language is called *uniform* or *global* tractability.

Let Φ be an instance of the described QCSP with variable set V. We let $G^{\leq}_{\Phi}$ denote the directed graph with vertex set V and with edge set containing all ordered pairs (v, v') such that $(v \leq v')$ is a constraint in Φ. Let $G^{=}_{\Phi}$ denote the undirected graph with vertex set V and with edge set containing all pairs $\{v, v'\}$ such that there exists a path from v to v' in $G^{\leq}_{\Phi}$ and there exists a path from v' to v in $G^{\leq}_{\Phi}$ (that is, v and v' are in the same strongly connected component of $G^{\leq}_{\Phi}$). When ψ is a disjunction of disequalities appearing in Φ, we use G^{ψ}_{Φ} to denote the undirected graph with vertex set V and with edge set containing all pairs $\{v, v'\}$ such that

– the disequality $v \neq v'$ appears in ψ, or
– $\{v, v'\}$ is an edge in $G^{=}_{\Phi}$.

The key to Theorem 28 is the following characterization of the false instances of the named QCSP. In the quantifier prefix of a QCSP instance, when the variable v' does not occur before the variable v (reading from left to right), we say that v is *earlier than* v', and also that v' is *later than* v. We use the modifier "strictly" to indicate that $v \neq v'$.

It is readily verified that this characterization can be checked in NL, yielding Theorem 28.

Theorem 29. *Let Γ be the expansion of the point algebra having arbitrary disjunctions of disequalities, and let Φ be an instance of QCSP(Γ). The formula Φ is false if and only if one of the following two conditions holds:*

1. *There exists a variable v and a universally quantified variable y strictly later than v such that in $G^{\leq}_{\Phi}$ there is a path from v to y or a path from y to v*
2. *There exists a disjunction of disequalities $\psi(z_1, \ldots, z_k)$ appearing in Φ such that for all existentially quantified variables $x \in \{z_1, \ldots, z_k\}$, either*
   (a) x is connected in $G^{=}_{\Phi}$ to a strictly earlier variable in $\{z_1, \ldots, z_k\}$, or
   (b) x is the earliest variable of the variables in $\{z_1, \ldots, z_k\}$ that are connected to x in G^{ψ}_{Φ}.

Proof. We conceive of a QCSP instance as a two-player game between a universal and existential player; the universal (existential) player sets the universally (respectively, existentially) quantified variables. Variables are set in the order dictated by the quantifier prefix. The existential player wins (equivalently, the formula is true) if and only if all constraints are satisfied after the variables have been set.

First, we show that the universal player wins if one of the two conditions hold. If condition (1) holds, the universal player sets y to any value so that $y < v$ in the case that there is a path from v to y, and so that $v < y$ in the case that there is a path from y to v. If condition (2) holds, then the universal player sets all universal variables y in ψ to the same value as the first variable in the component of y in G^{ψ}_{Φ}; if y is itself the first variable in its component in G^{ψ}_{Φ}, she sets it to an arbitrary value. We claim that in this manner the universal player wins.

By condition (2), all existential variables are either the first variable in their component in G^{ψ}_{Φ}, or have to be connected in $G^{=}_{\Phi}$ to a strictly earlier variable of ψ. In the latter case, unless some of the constraints of the form $x = y$ is violated, the existential player has to set each variable x in ψ to the same value as some variable from ψ that is strictly earlier; if this variable is universal, it has the same value as the first variable in the connected component of G^{ψ}_{Φ}, by the specified strategy of the universal player. If this variable is existential, it follows by induction that this variable must also be set to the value of the first variable in the connected component of G^{ψ}_{Φ}. Therefore, the universal player guarantees that in each connected component of G^{ψ}_{Φ} all variables are set to the earliest variable (assuming that the equalities in $G^{=}_{\Phi}$ are satisfied), in which case none of the disequalities in the disjunction ψ holds, and ψ is falsified.

Now, we prove that if neither of the two conditions hold, the existential player can win. We show that the existential player can preserve the following invariants during the course of the game:

I1 If two variables v, v' have been set (according to f) and there is a path from v to v' in $G^{\leq}_{\Phi}$, then $f(v) \leq f(v')$.

I2 If an existentially quantified variable x has been set (according to f) and it is the earliest variable among all variables connected to x in $G^{=}_{\Phi}$, then $f(x)$ is different from $f(v)$ for all variables v that are strictly earlier than x.

This suffices, as invariant I1 ensures that all constraints over $\leq$ are satisfied, and invariant I2 ensures that all disjunctions of disequalities are satisfied, since the negation of condition (2) is that for all disjunctions of disequalities ψ there is some connected component C in G^{ψ}_{Φ} containing an existentially quantified variable that is not connected to an earlier variable in $G^{=}_{\Phi}$ nor the earliest variable in C, and thus the new value selected by the existential player will satisfy one or more disequalities in ψ.

We consider two cases. The first case is that a universally quantified variable is set. Invariant I2 is preserved trivially. Invariant I1 is preserved because condition (1) does not apply.

Now consider the case where an existentially quantified variable x needs to be set. If the variable x is connected in $G^{=}_{\Phi}$ to a previously set variable v, we set x equal to the value of v. We clearly preserve I1, and trivially preserve I2. If the variable x is not connected in $G^{=}_{\Phi}$ to a previously set variable, let A be the set of all strictly earlier variables with paths in $G^{\leq}_{\Phi}$ to x, and let B be the set of all strictly earlier variables with paths in $G^{\leq}_{\Phi}$ from x. Let $a \in \mathbb{Q}$ be the maximum value of all variables in A, and $b \in \mathbb{Q}$ be the minimum value of all variables in B. Observe that $a \leq b$ by I1. If the values a and b originate from existentially quantified variables x_a, x_b, we have $a \neq b$; for if not, by I2 we would have that x_a, x_b are connected in $G^{=}_{\Phi}$. But because there is a path from x_a to x and a path from x to x_b, the three variables x_a, x_b, and x are all connected in $G^{=}_{\Phi}$, contradicting our assumption that x is not connected in $G^{=}_{\Phi}$ to a previously set variable.

By the negation of condition (1), at most one of the values a, b may originate from a universally quantified variable and when this occurs, again we have $a \neq b$ by the negation of condition (1) and invariant I2. We thus have that $a < b$ and we set x to any value in the interval (a, b). We clearly preserve both invariants: the first because the value selected for x is between a and b, the second because x is strictly between a and b. □

6 Spatial Reasoning

In this section, we study constraint satisfaction problems that arise from the spatial calculus known under the name *RCC-5* [RCC92, Ben94, JD97, RN99].

We first introduce the constraint satisfaction problems as in [JD97], and later discuss how to formulate the same problems with ω-categorical templates. Many different but equivalent ways to introduce these problems appeared in the literature, see e.g. [JD97, LM94, Due05, RN99, Ben94]; the formulation with an ω-categorical template presented here appears to be new.

Let X be a countably infinite set. Consider the relational structure

$$\mathbb{B}_1 = (2^X \setminus \{\emptyset\}, \mathtt{DR}, \mathtt{PO}, \mathtt{PP}, \mathtt{PPI}, \mathtt{EQ})$$

whose elements are the non-empty subsets of X, and whose relations are defined as follows.

$\mathtt{DR}(x, y)$ iff $x \cap y = \emptyset$	'x and y are disjoint'
$\mathtt{PP}(x, y)$ iff $x \subset y$	'y properly contains x'
$\mathtt{PPI}(x, y)$ iff $x \supset y$	'x properly contains y'
$\mathtt{EQ}(x, y)$ iff $x = y$	'x equals y'
$\mathtt{PO}(x, y)$ iff $\exists a, b, c\ (a \in x \setminus y, b \in x \cap y, c \in y \setminus x)$	'x and y properly overlap'

Note that for each pair (x, y) of elements of $\mathbb{B}_1$ exactly one of the relations $\mathtt{DR}, \mathtt{PO}, \mathtt{PP}, \mathtt{PPI}, \mathtt{EQ}$ holds. The *RCC-5 relations* are the binary boolean combinations of those relations. The constraint satisfaction problem for the expansion of $\mathbb{B}_1$ by the RCC-5 relations is also known as the *general network satisfaction problem for RCC-5*, and is NP-complete [RN99].

The structure $\mathbb{B}_1$ is not ω-categorical. To formulate the problem CSP($\mathbb{B}_1$) (and also the CSP for the RCC-5 relations) with a countably infinite ω-categorical structure, we use Fraïssé's theorem. The resulting structure is known in model theory as the ω-categorical countably infinite atomless boolean ring without 1 [Abi72]. The approach to define this structure with Fraïssé-amalgamation is also not new (this is mentioned, for example, in [Eva94]). However, we give the amalgamation argument here, because it allows us to prove that the CSP of the resulting structure is indeed the same problem as CSP($\mathbb{B}_1$); moreover, it shows an interesting connection between amalgamation problems and solving certain CSPs (Lemma 31).

Let $\mathcal{C}$ be the set of all finite induced substructures of $\mathbb{B}_1$, considered up to isomorphism.

Proposition 30. *The class $\mathcal{C}$ is an amalgamation class.*

Proof. By definition, $\mathcal{C}$ is closed under isomorphisms and substructures. The amalgamation property can in fact be formulated in constraint satisfaction terminology. We therefore first show the following.

Lemma 31. *Let A_1, A_2 be instances of CSP($\mathbb{B}_1$) that share some variables, where for each A_i and each pair of variables (x, y) in A_i there is a constraint on (x, y). If A_1, A_2 are both satisfiable (over $2^X \setminus \{\emptyset\}$), then there exists a mapping from the variables in A_1 and A_2 to $2^X \setminus \emptyset$ that satisfies all the constraints in A_1 and A_2. (Note that if A_1, A_2 are both satisfiable and x, y are both shared variables, A_1 and A_2 must have the same constraint on (x, y) as the relations are pairwise disjoint.)*

Proof. Let V_i denote the variables of A_i. Let $f_i : V_i \to 2^X \setminus \emptyset$ denote a solution to the instance A_i. We assume that $\cup_{v_1 \in V_1} f_1(v_1)$ and $\cup_{v_2 \in V_2} f_2(v_2)$ are disjoint. (One way to attain this is to first view f_1 and f_2 as mappings into $2^{X_1} \setminus \emptyset$ and $2^{X_2} \setminus \emptyset$ for two disjoint countably infinite sets X_1, X_2, and then view X as the disjoint union of X_1 and X_2.) Now consider the directed graph G with vertex set $V_1 \cup V_2$ and edge set containing all pairs (v, v') such that one of the constraints $\texttt{PP}(v, v')$, $\texttt{PPI}(v', v)$, $\texttt{EQ}(v, v')$, $\texttt{EQ}(v', v)$ is present in A_1 or A_2. Note that for any solution h to A_1, if v, v' are variables in V_1 and there is a path from v to v' in G, we must have $h(v) \subseteq h(v')$, and similarly for A_2.

We define a mapping $g : (V_1 \cup V_2) \to 2^X \setminus \emptyset$ on all of the variables as follows. For all $v_1 \in V_1$, we define $g(v_1) = f_1(v_1) \cup (\cup_{v_2} f_2(v_2))$ where the union is over all $v_2 \in V_2$ having a path to v_1 in G. Similarly, for all $v_2 \in V_2$, we define $g(v_2) = f_2(v_2) \cup (\cup_{v_1} f_1(v_1))$ where the union is over all $v_1 \in V_1$ having a path to v_2 in G. Here, we assume that for every variable v there appears a constraint $\texttt{EQ}(v, v)$. It is straightforward to verify that g is well-defined on variables in $V_1 \cap V_2$.

We claim that g satisfies all constraints in A_1 and A_2. Let v, v' be variables in V_1 (the case of V_2 is symmetric). If the constraint on (v, v') is $\texttt{PO}$, $\texttt{PP}$, $\texttt{PPI}$, or $\texttt{EQ}$, it is straightforward to verify from the definition of g that the constraint is satisfied. Now suppose that the constraint on (v, v') is $\texttt{DR}$. We prove by contradiction that it is satisfied by g. Suppose not. Then there exists an element $x \in g(v) \cap g(v')$. Since f_1 satisfied $\texttt{DR}(v, v')$ we must have that x appears in $\cup_{v_2 \in V_2} f(v_2)$. By definition of g, there exist variables $v_2, v_2' \in V_2$ such that $x \in f_2(v_2)$, $x \in f_2(v_2')$, there is a path from v_2 to v in G, and there is a path from v_2' to v' in G. The paths must pass through $V_1 \cap V_2$, so we may in fact assume that $v_2, v_2' \in V_1 \cap V_2$ (note that, again, by definition of G when traversing a directed path we may only gain elements). Since $x \in f_2(v_2)$ and $x \in f_2(v_2')$, the constraint on (v_2, v_2') cannot be $\texttt{DR}$. Since A_1 and A_2 have the same constraint on (v_2, v_2') this means that in A_1 we have variables v_2, v_2' which must have non-empty intersection in order to be satisfied, as well as paths in V_1 from v_2 to v and from v_2' to v' for variables v, v' with $\texttt{DR}(v, v')$ a constraint. But A_1 is satisfiable, a contradiction. □

Now, suppose that U, V_1, V_2 are structures in $\mathcal{C}$, e_1 an embedding of U into V_1, and e_2 an embedding of U into V_2. Both V_1 and V_2 can be considered as instances of CSP($\mathbb{B}_1$) that contain a constraint for each pair (x, y) of variables (recall that for every pair of elements of $\mathbb{B}_1$ exactly one of the relations $\texttt{DR}, \texttt{PO}, \texttt{PP}, \texttt{PPI}, \texttt{EQ}$ holds). Clearly, these instances are satisfiable. For all elements u in U, identify the element $e_1(u)$ of V_1 with the element $e_2(u)$ of V_2, and consider the resulting structure as an instance of CSP($\mathbb{B}_1$). Then Lemma 31 implies that this instance has a solution W, and the solution provides the embeddings f_1 and f_2 showing that $\mathcal{C}$ has the amalgamation property. $\square$

Corollary 32. *There exists an ω-categorical structure $\mathbb{B}_0$ such that CSP($\mathbb{B}_0$) equals CSP($\mathbb{B}_1$).*

Proof. We choose $\mathbb{B}_0$ to be the Fraïssé-limit of the class $\mathcal{C}$ described above, which exists and is ω-categorical due to Theorem 1. We have to show that a finite structure A homomorphically maps to $\mathbb{B}_0$ if and only if A homomorphically maps to $\mathbb{B}_1$. Suppose there is a homomorphism f from A to $\mathbb{B}_0$. Let A' be the structure induced by the image of f in $\mathbb{B}_0$. The structure A' is a member of $\mathcal{C}$, and in fact is an isomorphism class of substructures of $\mathbb{B}_1$, and f is a homomorphism to each of the members of this isomorphism class. Therefore, f gives rise to a homomorphism from A' to $\mathbb{B}_1$. In fact all the implications in this argument can be reversed, and this shows the statement. $\square$

We can easily obtain an ω-categorical constraint language for all of the RCC-5 relations by expanding the structure $\mathbb{B}_0$ by all binary relations that can be defined by boolean combinations of $\mathtt{DR}, \mathtt{PO}, \mathtt{PP}$, and $\mathtt{EQ}$. It is well-known (again, see [Eva94]) that there exists a countable set such that the elements of $\mathbb{B}_0$ can be seen as subsets of this set, and such that DR denotes disjointness, PO denotes proper overlap, and PP denotes strict containment between two subsets.

We now show that $\mathbb{B}_0$ has a QNU polymorphism. Observe that the the relation PP defines a partial order on the domain B_0 of $\mathbb{B}_0$, and we use infix notation $x\mathtt{PP}y$ when $(x, y) \in \mathtt{PP}$. We also define the relation $\equiv_m$ on tuples of $\mathbb{B}_0$ as in section 4 with respect to the partial order PP. The equivalence class of a tuple t with respect to $\equiv_m$ is denoted by $[t]$.

Theorem 33. *For all $k \geq 5$, the structure $\mathbb{B}_0$ has an oligopotent QNU polymorphism of arity k such that for all $t, t' \in (B_0)^k$, it holds that $t \equiv_m t'$ if and only if $f(t) = f(t')$. The polymorphism of arity k preserves any disjunction of disequalities over m variables with $k \geq 2m + 1$.*

In order to prove this theorem, we first establish a lemma. We will use $\hat{\mathbb{B}}_0 = (B_0; \mathtt{PP}, \mathtt{DR})$ to denote the reduct of $\mathbb{B}_0$ to the relations PP and DR. The lemma gives a sufficient condition for a *finite* structure to have an injective homomorphism into $\hat{\mathbb{B}}_0$, which we will then employ to construct the desired QNU polymorphism.

For a binary relation $\leq$, we say that b is a $\leq$-ancestor of b' if $b = b'$ or there exist elements $b_1, \dots, b_k$ such that $b = b_1 \leq \dots \leq b_k = b'$.

Lemma 34. *Suppose that $\Delta = (A, \mathtt{PP}, \mathtt{DR})$ is a finite structure where PP is irreflexive and has no non-trivial cycles, and there are no elements $(a, a') \in \mathtt{DR}$ such that a and a' have a common $\leq$-ancestor. Then there is an injective homomorphism from Δ to $\hat{\mathbb{B}}_0$.*

Proof. We expand Δ by a relation PO, defined as follows. For any distinct elements $a, a' \in A$, if a is not an ancestor of a', a' is not an ancestor of a, and a and a' have a common ancestor, place (a, a') in the relation PO. (Throughout this proof, we speak of ancestry with respect to the relation PP.) In addition, add to the relation DR all pairs (a, a') such that a and a' do not have a common ancestor.

We first show that $(A, \mathtt{PP}, \mathtt{DR}, \mathtt{PO})$ has a homomorphism h to $(B_0, \mathtt{PP}, \mathtt{DR}, \mathtt{PO})$. By Corollary 32, it suffices to show that there is a homomorphism from Δ to $(2^X \setminus \emptyset, \mathtt{PP}, \mathtt{DR}, \mathtt{PO})$. First, let $f : A \to X$ be an arbitrary injective mapping. Then define $g : A \to 2^X \setminus \emptyset$ by $g(a) = \cup\{f(b)\}$ where the union is over all elements b that are ancestors of a. It is straightforward to verify that g is a homomorphism as desired.

We have established that there exists a homomorphism h from $(A, \mathtt{PP}, \mathtt{DR}, \mathtt{PO})$ to $(B_0, \mathtt{PP}, \mathtt{DR}, \mathtt{PO})$. We now prove that h is in fact injective. Let a, a' be distinct elements of A. First we consider the case that a and a' have a common ancestor (again, with respect to PP). If a is an ancestor of a', then $h(a)$ and $h(a')$ are clearly not equal by the definition of PP. If a' is an ancestor of a, the reasoning is similar. If neither of a, a' is an ancestor of the other, then we have $\mathtt{PO}(h(a), h(a'))$, from which it follows that $h(a)$ and $h(a')$ are distinct. Next we consider the case that a and a' have no common ancestor. In that case, we have $\mathtt{DR}(h(a), h(a'))$, and $h(a)$ and $h(a')$ must be distinct. The homomorphism h thus satisfies the properties desired by the theorem. $\square$

Using the previous lemma, we can now give the sought-after QNUF polymorphisms. We want to point out that this proves not only that the basic relations of the constraint language for RCC-5 have the "local-to-global" property, but also that the same applies for every expansion of this constraint language by disjunctions of disequalities on a bounded number of variables.

Proof (Proof of Theorem 33). First we prove that every polymorphism of $\mathbb{B}_0$ is oligopotent. By Theorem 1, the Fraïssé-limit $\mathbb{B}_0$ is homogeneous; it is well-known that homogeneous structures over a finite relational signature are *model-complete* (see Section 7.3 and Corollary 6.4.2 in [Hod97]). A structure is model-complete iff every first-order formula is equivalent to an existential formula over the structure (this follows from Theorem 7.3.1 in [Hod97]); now it is easy to verify that this implies that every embedding of $\mathbb{B}_0$ into $\mathbb{B}_0$ is interpolated by the automorphisms of $\mathbb{B}_0$. So we only have to verify that every unary polymorphism of $\mathbb{B}_0$ is an embedding. But this is straightforward to verify, because for every pair of elements of $\mathbb{B}_0$ exactly one of the relations $\mathtt{DR}, \mathtt{PO}, \mathtt{PP}, \mathtt{PPI}, \mathtt{EQ}$ holds, so any unary polymorphism that preserves all these relations must be injective and strong.

Let Δ_0 be a finite induced substructure of $(\mathbb{B}_0)^k/\equiv_m$. By Proposition 3, it suffices to show that there is an injective homomorphism from Δ_0 to $\mathbb{B}_0$. Since $\mathbb{B}_1$ and $\mathbb{B}_0$ have the same finite induced substructures up to isomorphism, and since the definition of $\equiv_m$ only depends on the relation $\mathtt{PP}$, Δ_0 is isomorphic to an induced substructure Δ_1 of $(\mathbb{B}_1)^k/\equiv_m$. Recall that the elements of Δ_1 are equivalence classes of tuples of non-empty subsets of X. In the following, the $\cap$ and $\setminus$ operators are applied to tuples coordinate-wise.

Let Δ_2 be a finite induced structure of $(\mathbb{B}_1)^k/\equiv_m$ whose domain contains the domain of Δ_1 and which additionally contains elements such that the following holds:

- for all elements T, T' of Δ_1 such that $\mathtt{PO}(T, T')$ holds in Δ_1, there exist k-tuples $t \in T, t' \in T'$ in Δ_2 with $t \cap t'$, $t \setminus t'$ and $t' \setminus t$ non-empty at all coordinates, such that $[t \cap t']$, $[t \setminus t']$ and $[t' \setminus t]$ are in the universe of Δ_2.

We will show that there is an injective homomorphism from Δ_2 to $\mathbb{B}_0$, which suffices.

Consider the reduct Δ_3 of Δ_2 with the relations $\mathtt{PP}$ and $\mathtt{DR}$. We use Lemma 34 to show that there is an injective homomorphism from Δ_3 to $\hat{\mathbb{B}}_0$. It is straightforward to verify that in Δ_3 the relation $\mathtt{PP}$ is irreflexive, and it follows from Lemma 20 that $\mathtt{PP}$ has no non-trivial cycles; thus, we need only show that there are no elements T, T' with $(T, T') \in \mathtt{DR}$ that have a common $\mathtt{PP}$-ancestor.

We prove that if $T_0 \mathtt{PP} \cdots \mathtt{PP} T_l$, then for all $t_0 \in T_0$ and $t_l \in T_l$, it holds that $t_0 \subseteq t_l$ in at least $(k-2)$ coordinates. This suffices, for then if two elements T, T' of Δ_3 have a common $\mathtt{PP}$-ancestor T_0, for all $t_0 \in T_0$, $t \in T$ and $t' \in T$ we have $t_0 \subseteq t$ in at least $(k-2)$ coordinates and $t_0 \subseteq t'$ in at least $(k-2)$ coordinates, implying that there are at least $(k-4) \geq 1$ coordinates where $t_0 \subseteq t$ and $t_0 \subseteq t'$, from which it follows that $(T, T') \notin \mathtt{DR}$.

Suppose that $T_0 \mathtt{PP} \cdots \mathtt{PP} T_l$. If none of the equivalence classes T_i have a main value, then taking $T_0 = \{t_0\}$ and $T_l = \{t_l\}$, we in fact have $t_0 \subseteq t_l$ at all k coordinates. So suppose that some of the equivalence classes T_i have main values. Let m denote the index of the first class with a main value, and n the index of the last class with a main value; we have $0 \leq m \leq n \leq l$. Let B_m be the main value of T_m, and B_n the main value of T_n. It is straightforward to verify that $B_m \subseteq B_n$. All tuples in T_m are $(k-1)$-inferior to B_m; since $T_0 \mathtt{PP} \cdots \mathtt{PP} T_m$, we have that all tuples in T_0 are $(k-1)$-inferior to B_m. Likewise, all tuples in T_n are $(k-1)$-superior to B_n, from which it follows that all tuples in T_l are $(k-1)$-superior to B_n. Let $t_0 \in T_0$ and $t_l \in T_l$ be arbitrary tuples. Let I_0 be the coordinates where $t_0 \subseteq B_m$, and let I_l be the coordinates where $t_l \supseteq B_n$. We have $|I_0| \geq (k-1)$ and $|I_l| \geq (k-1)$, from which it follows that $|I_0 \cap I_l| \geq (k-2)$. For all $i \in I_0 \cap I_l$ we have $t_{0i} \subseteq B_m \subseteq B_n \subseteq t_{li}$.

We now have an injective homomorphism h from Δ_3 to $\hat{\mathbb{B}}_0$. We claim in fact that the restriction of h to Δ_1 is an injective homomorphism from Δ_1 to $\mathbb{B}_1$. We need to verify that it is a homomorphism with respect to the relation $\mathtt{PO}$; it is a homomorphism with respect to $\mathtt{DR}$, $\mathtt{PP}$, and $\mathtt{PPI}$ by definition of the reduct Δ_3, and trivially a homomorphism with respect to $\mathtt{EQ}$. Let T, T' be any two Δ_1-elements such that $(T, T') \in \mathtt{PO}$. By definition of Δ_2, there exist tuples $t \in T, t' \in T'$ such that $t \cap t'$, $t \setminus t'$, $t' \setminus t$ are non-empty at all coordinates and $[t \cap t']$, $[t \setminus t']$, and $[t' \setminus t]$ are in Δ_2. Now, since h preserves $\mathtt{PP}$ and $\mathtt{DR}$ we have $\mathtt{PP}(h([t \cap t']), h([t]))$, $\mathtt{PP}(h([t \cap t']), h([t']))$, $\mathtt{PP}(h([t \setminus t']), h([t]))$,

$\mathtt{DR}(h([t \setminus t']), h([t']))$, $\mathtt{PP}(h([t' \setminus t]), h([t']))$, and $\mathtt{DR}(h([t' \setminus t]), h([t]))$. This implies that we have $\mathtt{PO}(h([t]), h([t']))$.

The claim on disjunctions of disequalities is proved as in Theorem 17. □

7 Primitive Positive Interpretations

First-order interpretations are a fundamental concept in model theory. Here we apply interpretations in two different ways. The first is that we can sometimes define a constraint language Δ by a first-order interpretation in another (usually simpler) constraint language Γ. It is a well-known and useful fact that in this case Δ is ω-categorical whenever Γ is ω-categorical. Second, we use a restricted form of interpretations as a general tool to show the existence of oligopotent QNU-polymorphisms for various constraint languages.

Definition 35. *A σ-structure Δ has a* first-order interpretation *in a τ-structure Γ iff there exists a natural number d, called the* dimension *of the interpretation, and*

- *a first-order τ-formula $\delta(x_1, \ldots, x_d)$ – called* domain formula*,*
- *for each m-ary relation symbol R in σ a first-order τ-formula $\phi_R(\overline{x}_1, \ldots, \overline{x}_m)$ where the $\overline{x}_i$ denote disjoint n-tuples of distinct variables – called the* defining formulas*, and*
- *a surjective map $h : \delta(\Gamma^d) \to D_\Delta$ – called* coordinate map*,*

such that for all relations R in Δ and all tuples $\overline{a}_i \in \delta(\Gamma^d)$

$$\Delta \models R(h(\overline{a}_1), \ldots, h(\overline{a}_m)) \Leftrightarrow \Gamma \models \phi_R(\overline{a}_1, \ldots, \overline{a}_m)\,.$$

Lemma 36 (see e.g. [Hod97]). *If Γ is ω-categorical, then every structure Δ with a first-order interpretation in Γ is ω-categorical as well.*

Example. Allen's interval algebra is a famous temporal reasoning formalism. It can be described as a binary constraint language that has the following first-order interpretation in $(\mathbb{Q}, <)$: the dimension is two, the domain formula $\delta(x_1, x_2)$ is $x_1 < x_2$, the coordinate map is the identity mapping, and there is a binary relation for each 4-ary first-order definable relation over $(\mathbb{Q}, <)$.

Allen's interval algebra is in general NP-complete. A *fragment* of this constraint language is simply a reduct; by Lemma 36 this language and all its fragments are ω-categorical (this was also observed by [Hir96]). A special fragment of Allen's interval algebra is the *pointizable fragment* Γ^{PIA}, which consists of all binary relations over intervals that can be defined as a four-ary relation over $(\mathbb{Q}, <)$ from purely conjunctive combinations of formulas of the form $x_1 < x_2$, $x_1 = x_2$, $x_1 \neq x_2$, and $x_1 \leq x_2$. The pointizable fragment in particular contains the so-called *basic relations* of Allen's interval algebra: these are the 13 binary relations of the interval algebra that describe that an interval precedes, is preceded by, meets, overlaps, starts, (is) during, finishes, or equals another interval, with the obvious interpretations [All83]. It is known that CSP(Γ^{PIA}) has width 2, and hence is in P [vBC90].

Example. The rectangle algebra is a spatial reasoning formalism introduced by [Gue89]; it is based on the temporal reasoning formalism of Allen. Again, the corresponding constraint language can be conveniently described by a first-order interpretation, this time by a first-order interpretation in the fragment of Allen's Interval Algebra that contains the basic relations. The dimension of the interpretation is again two; therefore, we can think of the elements of the rectangle constraint language as rectangles with axes-parallel sides in Euclidean space. The domain formula is this time just true, the coordinate map is again the identity mapping, and the constraint language contains a binary relation for each 4-ary relation that is first-order definable over the basic relations of the interval algebra. The CSP for the rectangle algebra is NP-complete. Again, tractable fragments (i.e., reducts) have been identified [BCdC99]. By Lemma 36, the constraint language for the rectangle algebra is ω-categorical.

The *basic relations* of the rectangle algebra are the binary relations defined by a conjunction that describes the relative position of the projections of two rectangles to the first coordinate (the relative position is described by a basic relation of Allen's interval algebra), and the relative position of the projections to the second coordinate. For instance, there is a basic relation that describes that one rectangle R_1 properly contains another rectangle R_2, expressed by the conjunction that forces that the projection of R_1 to the first coordinate contains the projection of R_2 to the first coordinate, and that the projection of R_1 to the second coordinate contains the projection of R_2 to the second coordinate.

Primitive positive interpretations. The concept of interpretations can also be used to study the computational complexity of constraint satisfaction problems. To this end, we need the following restricted form of first-order interpretations.

Definition 37. *Suppose Δ has a first-order interpretation in Γ. If the formula δ and for all R the formula ϕ_R in the interpretation are primitive positive, we say that Δ has a* primitive positive interpretation *in Γ.*

Let Γ be a homogeneous τ-structure, and let Δ be a structure with a primitive positive interpretation in Γ of dimension d. Then instances A of CSP(Δ) can be reduced to instances of B of CSP(Γ) such that A homomomorphically maps to Δ if and only if B homomorphically maps to Γ, as follows. For each element a in A, we create d new variables $a^1, \dots, a^d$ in B. If $(a_1, \dots, a_k)$ is a tuple from a k-ary relation R in A, if ϕ_R is the primitive positive τ-formula that defines R, and if $S(a_{i_1}^{j_1}, \dots, a_{i_l}^{j_l})$ is a conjunct in ϕ_R where $S \in \tau$ is l-ary, then we add the tuple $(a_{i_1}^{j_1}, \dots, a_{i_l}^{j_l})$ to the relation S in B. Finally, if ϕ_R contains a conjunct $a_{i_1}^{j_1} = a_{i_2}^{j_2}$, then we identify the elements $a_{i_1}^{j_1}$ and $b_{i_2}^{j_2}$ in B. We write $T(A)$ for the τ-structure B obtained in this way from A (with respect to the given interpretation).

Proposition 38. *Let Γ be an ω-categorical relational structure, and let Δ be a structure with a primitive positive interpretation in Γ. If strong k-consistency implies global consistency for CSP(Γ), then Δ has a k-ary QNU polymorphism.*

We remark that for constraint languages over finite domains, an analog of this proposition follows from a result of [BKJ05].

Proof. The proposition is a consequence of the more general observation that there exists an algebra $\mathbf{A}$ in the pseudo-variety[3] $\mathcal{V}$ generated by Al(Γ) such that the domain of $\mathbf{A}$ equals the domain of Δ and all operations of $\mathbf{A}$ preserve all relations in Δ (see [Bod08]). Moreover, it is well-known that the operations of $\mathbf{A}$ satisfy the same equations as the operations in Al(Γ).

Hence, if Γ has a k-ary QNU polymorphism, then Δ has a k-ary quasi near-unanimity operation as well. Theorem 11 implies the statement of the proposition. □

We want to use Proposition 38 in combination with Theorem 11 to show that certain constraint languages have the local-to-global property. However, to use Theorem 11 we have to show that the operation whose existence we prove with Proposition 38 is oligopotent. The following Lemma is a useful tool to this end.

A satisfiable instance A of CSP(Γ) is called *maximally satisfiable* if whenever a constraint $a_1 = a_2$, or $R(a_1, \dots, a_k)$ for elements $a_1, \dots, a_k$ of A that is not implied by A is added to A, then A no longer homomorphically maps to Γ.

Lemma 39. *Let Γ be a homogeneous structure, and let Δ be a structure with a primitive positive interpretation in Γ. Moreover, assume that for every finite induced substructure A of Δ the structure $T(A)$ (defined before Proposition 38) is a maximally satisfiable instance of CSP(Γ). Then Δ is also homogeneous, and every polymorphism of Δ is oligopotent.*

[3] A *pseudo-variety* is a class of algebras with the same (functional) signature that is closed under homomorphisms (and hence also isomorphisms), subalgebras, and *finite* direct products.

Proof. Let τ be the signature of Γ, σ be the signature of Δ, and d be the dimension of the interpretation of Δ in Γ. To show that Δ is homogeneous, let f be an isomorphism between two n-element substructures A and B of Δ. We have to show that f can be extended to an automorphism of Δ. Let $(a_1^1, \dots, a_1^d), \dots, (a_n^1, \dots, a_n^d)$ be an enumeration of the elements of A, and $(b_1^1, \dots, b_1^d), \dots, (b_n^1, \dots, b_n^d)$ be an enumeration of the elements of B. Observe that $T(A)$ is isomorphic to the structure A' induced by $\{a_1^1, \dots, a_1^d, \dots, a_n^1, \dots, a_n^d\}$ in Γ: the function that maps the d variables that have been introduced in $T(A)$ for an element a_i from A to $a_i^1, \dots, a_i^d$ is clearly a homomorphism (and also well-defined with respect to potential vertex contractions in $T(A)$). We claim that the homomorphism must be injective and strong. If it is not injective, there are two distinct variables x_1 and x_2 mapped to the same element, it thus it would be possible to add the constraint $x_1 = x_2$ to $T(A)$ and maintain satisfiability. Since $T(A)$ is maximally satisfiable, the constraint $x_1 = x_2$ must have been already in $T(A)$. But then, by the definition of $T(A)$, x_1 and x_2 would have been identified, in contradiction to the assumption that x_1 and x_2 are distinct variables. If the homomorphism is not strong, it would be possible to add a constraint of the form $R(x_1, \dots, x_k)$, again contradicting maximal satisfiability of $T(A)$.

Similarly, $T(B)$ is isomorphic to the structure B' induced by $\{b_1^1, \dots, b_1^d, \dots, b_n^1, \dots, b_n^d\}$ in Γ. Since A and B are isomorphic, $T(A)$ and $T(B)$ are also isomorphic, and hence there is also an isomorphism between A' and B', which by homogeneity of Γ can be extended to an automorphism α' of Γ. Then the mapping α defined by $\alpha((a^1, \dots, a^d)) = (\alpha'(a^1), \dots, \alpha'(a^d))$ strongly preserves the domain formula of the interpretation, and also strongly preserves the τ-formulas ϕ_R for all $R \in \sigma$. The restriction β of α to the domain of Δ is also injective and reaches each element in the domain of Δ, and thus β is an automorphism of Δ. Since β extends f, we conclude that Δ is homogeneous.

We now show that every endomorphism of Δ is an embedding. Suppose that e is an endomorphism of Δ that is not an embedding, that is, e is not injective or there is $R \in \sigma$ and a tuple $(a_1, \dots, a_k)$ that is not in R^Δ, but $(e(a_1), \dots, e(a_k))$ is in R^Δ. Let A be the structure induced by $\{a_1, \dots, a_k\}$ in Δ, and B be the structure induced by $e(\{a_1, \dots, a_k\})$ in Δ. We define an operation g from $T(A)$ to $T(B)$ as follows: if $(a_i^1, \dots, a_i^d)$ are the vertices in $T(A)$ that have been introduced for a vertex a_i in A, and $(b_i^1, \dots, b_i^d)$ are the vertices that have been introduced for $e(a_i)$ in $T(B)$, then g maps a_i^j to b_i^j for all $1 \leq j \leq d$ (this is a well-defined map with respect to the potential vertex contractions in $T(A)$). The mapping g must be a homomorphism, but since e is not an embedding, g is also not an embedding. We know that there is a homomorphism h from $T(B)$ to Γ. Then the homomorphism $h(g)$ from $T(A)$ to Γ is not an embedding as well. As we have seen before, this shows that $T(A)$ is not maximally satisfiable, in contradiction to our assumption. So every endomorphism of Δ is an embedding.

Since Δ is homogeneous, every embedding of Δ into itself is locally generated by the automorphisms of Δ, and so the first and the second part of the proof together show that every polymorphism of Δ is oligopotent. □

Example. We continue the example with the pointizable fragment of Allen's interval algebra, and show that for this constraint language establishing 5-consistency implies global consistency.

Recall that for CSP$(\mathbb{Q}, \leq, \neq)$ strong 5-consistency implies global consistency [vBC90]. We use Theorem 38, Lemma 39, and Theorem 11 to show that 5-consistency of the pointizable fragment of Allen's interval algebra implies global consistency. This answers a question posed in [Kou97].

Theorem 40. *Establishing 5-consistency for the pointizable fragment of Allen's Interval Algebra implies global consistency.*

Proof. By definition, the pointizable fragment of Allen's interval algebra has a primitive positive interpretation in the constraint language $(\mathbb{Q}, \leq, \neq)$. Theorem 38 shows that the pointizable fragment of Allen's Interval Algebra has a 5-ary QNUF f.

Let A be a finite substructure of Γ^{PIA}. On any pair of elements from A, one of the basic relations from the pointizable fragment holds in A. This implies that in the structure $T(A)$ on any pair of elements one of the relations $=$, $<$, or $>$ holds. Hence, if we add a tuple to any of the relations in $T(A)$ there is no homomorphism from $T(A)$ to Γ. Since $(\mathbb{Q}, \leq, \neq)$ is a homogeneous structure, Lemma 39 shows that f is oligopotent. Theorem 11 implies the statement of the Theorem.

Example. We continue the example with the rectangle algebra. By definition, the basic relations of the constraint language for the rectangle algebra have a primitive positive interpretation in the pointizable fragment Γ^{PIA} of Allen's interval algebra. As we have seen in the previous paragraph, for Γ^{PIA} strong 5-consistency implies global consistency. Again it is straighforward to verify the conditions of Lemma 39. As in the previous example, we can use Theorem 11 to show that strong 5-consistency implies global consistency for the basic relations of the rectangle algebra.

Acknowledgements

We would like to thank the anonymous referee for his work. Hubie Chen is supported by the MEC Program 'Ramon y Cajal'.

References

[Abi72] Alexander Abian. Categoricity of denumerable atomless boolean rings. *Studia Logica*, 30(1):63–67, 1972.

[AHV95] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison Wesley, 1995.

[All83] J. F. Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843, 1983.

[BC07a] Manuel Bodirsky and Hubert Chen. Oligomorphic clones. *Algebra Universalis*, 57(1):109–125, 2007.

[BC07b] Manuel Bodirsky and Hubie Chen. Quantified equality constraints. In *Proceedings of LICS'07*, pages 203–212, 2007.

[BCdC99] Philippe Balbiani, Jean-Francois Condotta, and Luis Farinas del Cerro. A new tractable subclass of the rectangle algebra. In *IJCAI*, pages 442–447, 1999.

[BD06] Manuel Bodirsky and Víctor Dalmau. Datalog and constraint satisfaction with infinite templates. In *Proceedings of STACS'06*, pages 646–659, 2006.

[BD08] Manuel Bodirsky and Víctor Dalmau. Datalog and constraint satisfaction with infinite templates. *An extended abstract appeared in the proceedings of STACS'06. The full version is available online at arXiv:0809.2386v1 [cs.LO]*, 2008.

[Ben94] Brandon Bennett. Spatial reasoning with propositional logics. In *International Conference on Knowledge Representation and Reasoning*. Morgan Kaufmann, 1994.

[BJK05] A. Bulatov, P. Jeavons, and A. Krokhin. The complexity of constraint satisfaction: An algebraic approach (a survey paper). *In: Structural Theory of Automata, Semigroups and Universal Algebra (Montreal, 2003), NATO Science Series II: Mathematics, Physics, Chemistry*, 207:181–213, 2005.

[BKJ05] A. Bulatov, A. Krokhin, and P. G. Jeavons. Classifying the complexity of constraints using finite algebras. *SIAM Journal on Computing*, 34:720–742, 2005.

[BN06] Manuel Bodirsky and Jaroslav Nešetřil. Constraint satisfaction with countable homogeneous templates. *Journal of Logic and Computation*, 16(3):359–373, 2006.

[Bod08] Manuel Bodirsky. Constraint satisfaction problems with infinite templates. In Heribert Vollmer, editor, *Complexity of Constraints (a collection of survey articles), to appear*. Springer, LNCS 5250, 2008.

[BP74] K. A. Baker and A. F. Pixley. Polynomial interpolation and the Chinese remainder theorem for algebraic systems. *Math. Z.*, 143:165–174, 1974.

[Che] Hubie Chen. The computational complexity of quantified constraint satisfaction. Ph.D. thesis, Cornell University, August 2004.

[CJ06] David Cohen and Peter Jeavons. The complexity of constraint languages. *Appears in: Handbook of Constraint Programming*, 2006.

[Dec92] Rina Dechter. From local to global consistency. *Artif. Intell.*, 55(1):87–108, 1992.

[Due05] Ivo Duentsch. Relation algebras and their application in temporal and spatial reasoning. *Artificial Intelligence Review*, 23:315–357, 2005.

[EF99] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer, 1999. 2nd edition.

[Eva94] David Evans. Examples of $\aleph_0$-categorical structures. *In 'Automorphisms of first-order structures', R. Kaye and H.D. Macpherson, Oxford University Press*, pages 33–72, 1994.

[FGV05] M. Fisher, D. Gabbay, and L. Vila, editors. *Handbook of Temporal Reasoning in Artificial Intelligence*. Elsevier, 2005.

[Fra86] R. Fraïssé. *Theory of Relations*. North-Holland, 1986.

[Fre82] E.C. Freuder. A sufficient condition for backtrack-free search. *Journal of the ACM*, 29(1):24–32, 1982.

[FV99] 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.

[Gue89] H. Guesgen. Spatial reasoning based on Allen's temporal logic. Report ICSI TR 89-049, International Computer Science Institute, 1989.

[Hir96] R. Hirsch. Relation algebras of intervals. *Artificial Intelligence Journal*, 83:1–29, 1996.

[HN92] P. Hell and J. Nesetril. The core of a graph. *Discrete Math.*, 109:117–126, 1992.

[Hod97] Wilfrid Hodges. *A shorter model theory*. Cambridge University Press, 1997.

[JCC98] Peter Jeavons, David Cohen, and Martin Cooper. Constraints, consistency and closure. *AI*, 101(1-2):251–265, 1998.

[JD97] Peter Jonsson and Thomas Drakengren. A complete classification of tractability in RCC-5. *J. Artif. Intell. Res.*, 6:211–221, 1997.

[Kou97] Manolis Koubarakis. From local to global consistency in temporal constraint networks. *Theor. Comput. Sci.*, 173(1):89–112, 1997.

[Kou01] M. Koubarakis. Tractable disjunctions of linear constraints: Basic results and applications to temporal reasoning. *Theoretical Computer Science*, 266:311–339, 2001.

[KV98] Ph. G. Kolaitis and M. Y. Vardi. Conjunctive-query containment and constraint satisfaction. In *Proceedings of PODS'98*, pages 205–213, 1998.

[LM94] Peter B. Ladkin and Rober D. Maddux. On binary constraint problems. *Journal of the Association for Computing Machinery*, 41(3):435–469, 1994.

[Mac77] A. K. Mackworth. Consistency in networks of relations. *AI*, 8:99–118, 1977.

[MVvB89] Henry Kautz Marc Vilain and Peter van Beek. Constraint propagation algorithms for temporal reasoning: A revised report. *Reading in Qualitative Reasoning about Physical Systems*, pages 373–381, 1989.

[RCC92] D. A. Randell, Z. Cui, and A. G. Cohn. A spatial logic based on regions and connection. In *Conference on Principles on Knowledge Representation and Reasoning (KR'92)*, pages 165–176, 1992.

[RN99] Jochen Renz and Bernhard Nebel. On the complexity of qualitative spatial reasoning: A maximal tractable fragment of the region connection calculus. *Artif. Intell.*, 108(1-2):69–123, 1999.

[RN07] J. Renz and B. Nebel. Qualitative spatial reasoning using constraint calculi. In M. Aiello, I. Pratt-Hartmann, and J. van Benthem, editors, *Handbook of Spatial Logics*. Springer Verlag, Berlin, 2007.

[vBC90] Peter van Beek and Robin Cohen. Exact and approximate reasoning about temporal relations. *Computational Intelligence*, 6:132–144, 1990.