Temporal Constraint Satisfaction Problems in Fixed-Point Logic

Manuel Bodirsky*
TU Dresden
manuel.bodirsky@tu-dresden.de

Wied Pakusa
RWTH Aachen
pakusa@logic.rwth-aachen.de

Jakub Rydval
TU Dresden
jakub.rydval@tu-dresden.de

Abstract

Finite-domain constraint satisfaction problems are either solvable by Datalog, or not even expressible in fixed-point logic with counting. The border between the two regimes can be described by a strong height-one Maltsev condition. For infinite-domain CSPs, the situation is more complicated even if the template structure of the CSP is model-theoretically tame. We prove that there is no Maltsev condition that characterizes Datalog already for the CSPs of first-order reducts of $(\mathbb{Q}; <)$; such CSPs are called *temporal CSPs* and are of fundamental importance in infinite-domain constraint satisfaction. Our main result is a complete classification of temporal CSPs that can be expressed in one of the following logical formalisms: Datalog, fixed-point logic (with or without counting), or fixed-point logic with the Boolean rank operator. The classification shows that many of the equivalent conditions in the finite fail to capture expressibility in Datalog or fixed-point logic already for temporal CSPs.

CCS Concepts: • **Theory of computation** → **Constraint and logic programming**; *Design and analysis of algorithms*; Randomness, geometry and discrete structures.

Keywords: temporal constraint satisfaction problems, fixed-point logic, Maltsev conditions

ACM Reference Format:
Manuel Bodirsky, Wied Pakusa, and Jakub Rydval. 2020. Temporal Constraint Satisfaction Problems in Fixed-Point Logic. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS '20), July 8–11, 2020, Saarbrücken, Germany*. ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3373718.3394750

*Manuel Bodirsky was supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant Agreement No 681988, CSP-Infinity).

LICS '20, July 8–11, 2020, Saarbrücken, Germany

ACM ISBN 978-1-4503-7104-9/20/07…$15.00
https://doi.org/10.1145/3373718.3394750

1 Introduction

The quest for finding a logic capturing Ptime is an ongoing challenge in the field of finite model theory originally motivated by questions from database theory [29]. Ever since its proposal, most candidates are based on various extensions of *fixed-point logic* (FP), for example by *counting* or by *rank operators*. Though not a candidate for capturing Ptime, *Datalog* is perhaps the most studied fragment of FP. Datalog is particularly well-suited for formulating various algorithms for solving *constraint satisfaction problems* (CSPs); examples of famous algorithms that can be formulated in Datalog are the *arc consistency* procedure and the *path consistency* procedure. In general, the expressive power of FP is limited as it fails to express counting properties of finite structures such as even cardinality. However, the combination of a mechanism for iteration and a mechanism for counting provided by *fixed-point logic with counting* (FPC) is strong enough to express most known algorithmic techniques leading to polynomial-time procedures [19, 28]. In fact, all known decision problems for finite structures that provably separate FPC from Ptime are at least as hard as deciding solvability of systems of linear equations over a non-trivial finite Abelian group [44]. If we extend FPC further to by the *Boolean rank operator* [28], we obtain the logic FPR_2 which is known to capture Ptime for Boolean CSPs [45].

Proving inexpressibility results for FPR_2 seems to be very difficult. The first inexpressibility result for FPC is due to Cai, Fürer and Immerman for systems of linear equations over $\mathbb{Z}_2$ [16]. In 2009, this result was extended to arbitrary non-trivial finite Abelian groups by Atserias, Bulatov and Dawar [1]; their work was formulated purely in the framework of CSPs. At around the same time, Barto and Kozik [4] settled the closely related bounded width conjecture of Larose and Zádori [37]. A combination of both works together with results from [35, 39] yields the following theorem.

Theorem 1.1 ([1, 4, 35, 39])**.** *For a finite structure $\mathfrak{B}$, the following six statements are equivalent.*

1. $\text{CSP}(\mathfrak{B})$ *is expressible in* Datalog.
2. $\text{CSP}(\mathfrak{B})$ *is expressible in* FP.
3. $\text{CSP}(\mathfrak{B})$ *is expressible in* FPC.
4. $\mathfrak{B}$ *does not pp-construct linear equations over any non-trivial finite Abelian group.*
5. $\mathfrak{B}$ *has weak near-unanimity polymorphisms for all but finitely many arities.*

6. $\mathfrak{B}$ *has weak near-unanimity polymorphisms* f, g *that satisfy the* 3-4 equation $g(x,x,y) \approx f(x,x,x,y)$.

In particular, Datalog, FP, and FPC are equally expressive when it comes to finite-domain CSPs. This observation raises the question whether there are natural classes of CSPs where the above-mentioned fragments and extensions of FP do not collapse. In fact, this question was already answered positively in 2007 by Bodirsky and Kára for the CSPs of first-order reducts of $(\mathbb{Q}; <)$, also known as (infinite-domain) temporal CSPs [11]; the decision problem $\mathrm{CSP}(\mathbb{Q}; \mathrm{R}_{\mathrm{MIN}})$, where

$$\mathrm{R}_{\mathrm{MIN}} := \{(x,y,z) \in \mathbb{Q}^3 \mid x > y \vee x > z\},$$

is provably not solvable by any Datalog program [12] but it is expressible in FP, as we will see later. Since every CSP formally represents a class of finite structures whose complement is closed under homomorphisms, this also yields an alternative proof of a result from [21] stating that the homomorphism preservation theorem fails for FP.

We present a complete classification of temporal CSPs that can be solved in Datalog, FP, FPC, or FPR_2. Several famous NP-hard problems such as the *Betweenness* problem or the *Cyclic Ordering* problem are temporal CSPs. Temporal CSPs have been studied, e.g., in artificial intelligence [40], Scheduling [12], and approximation [31]. Random instances of temporal CSPs have been studied in [25]. Temporal CSPs also play a particular role for the theory of infinite-domain CSPs since the important technique of reducing infinite-domain CSPs to finite-domain CSPs [13] cannot be used to prove polynomial-time tractability results for this class. The classification leads to the following sequence of inclusions for temporal CSPs:

$$\text{Datalog} \subsetneq \text{FP} = \text{FPC} \subsetneq \text{FPR}_2 = \text{Ptime}$$

Our results show that the expressibility of temporal CSPs in these logics can be characterized in terms of avoiding pp-constructibility of certain structures. If a structure can pp-construct the complete graph on three vertices, $\mathfrak{K}_3$, then its CSP is not expressible in any of the listed logics. If a structure can pp-construct $(\mathbb{Q}; \mathrm{R}_{\mathrm{MIN}})$ (see the paragraph below Theorem 1.1), then its CSP is not expressible in Datalog; conversely, if a temporal CSP can pp-construct neither $\mathfrak{K}_3$ nor $(\mathbb{Q}; \mathrm{R}_{\mathrm{MIN}})$, then it is contained in Datalog. We show that a temporal CSP is expressible in FP and in FPC iff it can pp-construct neither $\mathfrak{K}_3$ nor $(\mathbb{Q}; \mathrm{X})$ where

$$\mathrm{X} := \{(x,y,z) \in \mathbb{Q}^3 \mid x = y < z \vee x = z < y \vee y = z < x\}$$

Finally, we show that FPR_2 captures polynomial time on temporal CSPs (unless P=NP).

Our results also show that every temporal CSP with a template that pp-constructs $(\mathbb{Q}; \mathrm{X})$ but not all finite structures is solvable in polynomial time, is not expressible in FPC, and cannot encode linear equation constraints over any non-trivial finite Abelian group. Such temporal CSPs are Datalog equivalent to the following decision problem:

3-Ord-Xor-Sat

INPUT: A finite set of homogeneous linear Boolean equations of length 3.
QUESTION: Does every non-empty subset S of the equations have a solution where at least one variable occurring in an equation from S denotes the value 1.

We have eliminated the following candidates for general algebraic criteria for expressibility of CSPs in FP (Section 7):

- the inability to pp-construct linear equations over a non-trivial finite Abelian group [1],
- the 3-4 equation for weak near-unanimity polymorphisms modulo outer endomorphisms [13],
- the existence of weak near-unanimity polymorphisms modulo outer endomorphisms for all but finitely many arities [4].

We have good news and bad news regarding the existence of general algebraic criteria for expressibility of CSPs in fragments and/or extensions of FP. The bad news is that there is no Maltsev condition that would capture expressibility of temporal CSPs in Datalog (see Theorem 7.2) which carries over to CSPs of *reducts of finitely bounded homogeneous* structures and more generally to CSPs of ω*-categorical* templates. This is particularly striking because ω-categorical CSPs are otherwise well-behaved when it comes to expressibility in Datalog—every ω-categorical CSP expressible in Datalog admits a *canonical Datalog program* [9]. The question which ω-categorical CSPs are in Datalog is the central theme in the survey article [10]. The good news is that there is a strong height-one Maltsev condition that characterises the expressibility in FP for finite-domain and temporal CSPs (Theorem 7.8). It is based on a family $\mathcal{E}_{k,n}$ of strong height-one Maltsev conditions closely related to the *dissected weak near-unanimity* identities introduced in [3, 24]; the set of polymorphisms of every first-order reduct of a finitely bounded homogeneous structure known to the authors (in particular of the examples from Theorem 1.3 in [14]) satisfies $\mathcal{E}_{k,k+1}$ for all but finitely many k iff its CSP is in FP.

2 Preliminaries

We need various notions from model theory, constraint satisfaction, and universal algebra. The set $\{1,\ldots,n\}$ is denoted by $[n]$. We use the boldface notation for tuples and matrices; for a tuple $\boldsymbol{t}$ indexed by a set I, the value of $\boldsymbol{t}$ at the position $i \in I$ is denoted by $\boldsymbol{t}[i]$.

2.1 Structures

A (*relational*) *signature* τ is a set of *relation symbols*, each with an associated natural number called *arity*, and *constant symbols*. A (*relational*) τ*-structure* $\mathfrak{A}$ consists of a set A (the *domain*) together with the relations $R^{\mathfrak{A}} \subseteq A^k$ for each relation symbol $R \in \tau$ of arity k and the *constants* $c^{\mathfrak{A}} \in A$ for each constant symbol $c \in \tau$. We often describe structures by listing their domain, relations, and constants, that is, we write

$\mathfrak{A} = (A; R_1^{\mathfrak{A}}, \ldots, c_1^{\mathfrak{A}}, \ldots)$. An *expansion* of $\mathfrak{A}$ is a σ-structure $\mathfrak{B}$ with $A = B$ such that $\tau \subseteq \sigma$, $R^{\mathfrak{B}} = R^{\mathfrak{A}}$ for each relation symbol $R \in \tau$, and $c^{\mathfrak{B}} = c^{\mathfrak{A}}$ for each constant symbol $c \in \tau$. Conversely, we call $\mathfrak{A}$ a *reduct* of $\mathfrak{B}$.

A *homomorphism* $h\colon \mathfrak{A} \to \mathfrak{B}$ for τ-structures $\mathfrak{A}, \mathfrak{B}$ is a mapping $h\colon A \to B$ that *preserves* each constant and each relation of $\mathfrak{A}$, that is, $h(c^{\mathfrak{A}}) = c^{\mathfrak{B}}$ holds for every constant symbol $c \in \tau$ and if $\boldsymbol{t} \in R^{\mathfrak{A}}$ for some k-ary relation symbol $R \in \tau$, then $(h(\boldsymbol{t}[1]), \ldots, h(\boldsymbol{t}[k])) \in R^{\mathfrak{B}}$. We write $\mathfrak{A} \to \mathfrak{B}$ if $\mathfrak{A}$ homomorphically maps to $\mathfrak{B}$ and $\mathfrak{A} \not\to \mathfrak{B}$ otherwise. We say that $\mathfrak{A}$ and $\mathfrak{B}$ are *homomorphically equivalent* if $\mathfrak{A} \to \mathfrak{B}$ and $\mathfrak{B} \to \mathfrak{A}$. An *endomorphism* is a homomorphism from $\mathfrak{A}$ to $\mathfrak{A}$. By an *embedding* we mean an injective homomorphism $e\colon \mathfrak{A} \to \mathfrak{B}$ that additionally satisfies the following condition: for every k-ary relation symbol $R \in \tau$ and $\boldsymbol{t} \in A^k$ we have $(h(\boldsymbol{t}[1]), \ldots, h(\boldsymbol{t}[k])) \in R^{\mathfrak{B}}$ only if $\boldsymbol{t} \in R^{\mathfrak{A}}$. We write $\mathfrak{A} \hookrightarrow \mathfrak{B}$ if $\mathfrak{A}$ embeds to $\mathfrak{B}$. An *isomorphism* is a surjective embedding. Two structures $\mathfrak{A}$ and $\mathfrak{B}$ are *isomorphic* if there exists an isomorphism from $\mathfrak{A}$ to $\mathfrak{B}$. An *automorphism* is an isomorphism from $\mathfrak{A}$ to $\mathfrak{A}$. A *substructure* of $\mathfrak{A}$ is a structure $\mathfrak{B}$ over $B \subseteq A$ such that the inclusion map $i\colon B \to A$ is an embedding.

An n-ary *polymorphism* of a relational structure $\mathfrak{A}$ is a mapping $f\colon A^n \to A$ such that for every constant symbol $c \in \tau$ we have $f(c^{\mathfrak{A}}, \ldots, c^{\mathfrak{A}}) = c^{\mathfrak{A}}$, and for every k-ary relation symbol $R \in \tau$ and tuples $\boldsymbol{t}_1, \ldots, \boldsymbol{t}_n \in R^{\mathfrak{A}}$ we have $\big(f(\boldsymbol{t}_1[1], \ldots, \boldsymbol{t}_1[n]), \ldots, f(\boldsymbol{t}_k[1], \ldots, \boldsymbol{t}_k[n])\big) \in R^{\mathfrak{A}}$. We say that f *preserves* $\mathfrak{A}$ to indicate that f is a polymorphism of $\mathfrak{A}$. We might also say that an operation *preserves* a relation R over A if it is a polymorphism of $(A; R)$.

2.2 Model theory

A short and precise definition of the notion of a *logic* can be found in [29]. We assume that the reader is familiar with classical *first-order* logic (FO); we allow the first-order formulas $x = y$ and $\bot$. A first-order τ-formula ϕ is *primitive positive* (pp) if it is of the form $\exists x_1, \ldots, x_m(\phi_1 \wedge \cdots \wedge \phi_n)$, where each ϕ_i is *atomic*, that is, of the form $\bot$, $x_i = x_j$, or $R(x_{i_1}, \ldots, x_{i_\ell})$ for some $R \in \tau$. For a τ-structure $\mathfrak{A}$ and a set of FO τ-formulas Θ, we say that an n-ary relation has a Θ-*definition* in $\mathfrak{A}$ if it is of the form $\{\boldsymbol{t} \in A^n \mid \mathfrak{A} \models \phi(\boldsymbol{t})\}$ for some $\phi \in \Theta$. The following statement follows a well-known principle that connects logic and algebra.

Proposition 2.1 ([32]). *Let $\mathfrak{A}$ be a relational structure.*

1. *Every relation over A with a FO-definition in $\mathfrak{A}$ is preserved by all automorphisms of $\mathfrak{A}$.*
2. *Every relation over A with a pp-definition in $\mathfrak{A}$ is preserved by all polymorphisms of $\mathfrak{A}$.*

The set of all automorphisms of $\mathfrak{A}$, denoted by $\mathrm{Aut}(\mathfrak{A})$, forms a *permutation group* w.r.t. the map composition [32]. The *orbit* of a tuple $\boldsymbol{t} \in A^k$ under the *natural action* of $\mathrm{Aut}(\mathfrak{A})$ on A^k is the set $\{(g(\boldsymbol{t}[1]), \ldots, g(\boldsymbol{t}[k])) \mid g \in \mathrm{Aut}(\mathfrak{A})\}$. A structure is ω-*categorical* if its first-order theory has exactly one countable model up to isomorphism. The theorem of Engeler, Ryll-Nardzewski, and Svenonius (Theorem 6.3.1 in [32]) asserts that the following statements are equivalent for a countably infinite structure $\mathfrak{A}$ with countable signature:

- $\mathfrak{A}$ is ω-categorical.
- Every relation preserved by all automorphisms of $\mathfrak{A}$ has a FO-definition in $\mathfrak{A}$.
- For every $k \geq 1$, there are only finitely many orbits of k-tuples under the natural action of $\mathrm{Aut}(\mathfrak{A})$.

A structure $\mathfrak{A}$ is *homogeneous* if every isomorphism between finite substructures of $\mathfrak{A}$ extends to an automorphism of $\mathfrak{A}$. Every homogeneous structure with a finite relational signature is ω-categorical [32]. A structure $\mathfrak{A}$ is *finitely bounded* if there is a universal first-order sentence ϕ such that a finite structure embeds into $\mathfrak{A}$ iff it satisfies ϕ. A prime example of a finitely bounded homogeneous structure is $(\mathbb{Q}; <)$ [13].

Definition 2.2 (Counting, finite variable logics [2, 43])**.** By FOC we denote the extension of FO by the counting quantifiers $\exists^i$. If $\mathfrak{A}$ is a τ-structure and ϕ a τ-formula with a free variable x, then $\mathfrak{A} \models \exists^i x.\, \phi(x)$ iff there exist i distinct $a \in A$ such that $\mathfrak{A} \models \phi(a)$. While FOC is not more expressive than FO, the presence of counting quantifiers might affect the number of variables that are necessary to define a particular relation. We denote the fragment of FO in which every formula has at most k variables by L^k, and its existential positive fragment by $\exists^+ L^k$. We write $\mathfrak{A} \Rightarrow_k \mathfrak{B}$ if every sentence of $\exists^+ L^k$ that is true in $\mathfrak{A}$ is also true in $\mathfrak{B}$. The k-variable fragment of FOC is denoted by C^k. We write $\mathfrak{A} \equiv_{C^k} \mathfrak{B}$ if every sentence of C^k is true in $\mathfrak{A}$ iff it is true in $\mathfrak{B}$. The infinitary logic $L^k_{\infty\omega}$ extends L^k with infinite disjunctions and conjunctions. The extension of $L^k_{\infty\omega}$ by the counting quantifiers $\exists^i$ is denoted by $C^k_{\infty\omega}$. We write $\mathfrak{A} \equiv_{C^k_{\infty\omega}} \mathfrak{B}$ if every sentence of $C^k_{\infty\omega}$ is true in $\mathfrak{A}$ iff it is true in $\mathfrak{B}$.

2.3 Fixed-point logic

Inflationary fixed-point logic (IFP) is defined by adding formation rules to FO whose semantics is defined with inflationary fixed-points of arbitrary operators, and *least fixed-point logic* (LFP) is defined by adding formation rules to FO whose semantics is defined using least fixed-points of monotone operators. The logics LFP and IFP are *equivalent* in the sense that they define the same relations over the class of all structures [36]. For this reason, they are both commonly referred to as FP (see, e.g., [2]).

Datalog is usually understood as the existential positive fragment of LFP (see [21]). The existential positive fragments of LFP and IFP are equivalent, because the fixed-point operator induced by a formula from either of the fragments is monotone, which implies that its least and inflationary fixed-point coincide (see Proposition 10.3 in [38]). This allows us to define Datalog as the existential positive fragment of FP.

For the definitions of the counting extensions IFPC and LFPC we refer the reader to [26]. One important detail is that the equivalence LFP ≡ IFP extends to LFPC ≡ IFPC (see p. 189 in [26]). Again, we refer to both counting extensions simply as FPC. It is worth mentioning that the extension of Datalog with counting is also equivalent to FPC [27]. In the following, all we need to know about FPC is Theorem 2.3.

Theorem 2.3 (Immerman and Lander [19]). *For every* FPC *sentence* ϕ, $\mathfrak{A} \equiv_{C^k} \mathfrak{B}$ *implies* $\mathfrak{A} \models \phi \Leftrightarrow \mathfrak{B} \models \phi$ *for some* $k \in \mathbb{N}$.

This result follows from the fact that for every FPC formula ϕ there exists k such that, on structures with at most n elements, ϕ is equivalent to a formula of C^k whose quantifier depth is bounded by a polynomial function of n [19]. Clearly $\mathfrak{A} \equiv_{C^k_{\infty\omega}} \mathfrak{B}$ implies $\mathfrak{A} \equiv_{C^k} \mathfrak{B}$. The difference here is that every formula of FPC is actually equivalent to a formula of $C^k_{\infty\omega}$ for some k, that is, FPC forms a fragment of the infinitary logic $C^\omega_{\infty\omega} := \bigcup_{k\in\mathbb{N}} C^k_{\infty\omega}$ (Corollary 4.20 in [43]).

The logic FPR_2 extends FPC by the Boolean rank operator rk, making it the most expressive logic explicitly treated in this paper. Intuitively, rk is a logical constructor that can be used to form a rank term $[\text{rk}_{x,y}\phi(x,y)]$ from a given formula $\phi(x,y)$. The value of $[\text{rk}_{x,y}\phi(x,y)]$ in an input structure $\mathfrak{A}$ is the rank of a Boolean matrix specified by $\phi(x,y)$ through its evaluation in $\mathfrak{A}$. For instance, $[\text{rk}_{x,y}(x = y) \wedge \phi(x)]$ computes in an input structure $\mathfrak{A}$ the number of elements $a \in A$ such that $\mathfrak{A} \models \phi(a)$ for a given formula $\phi(x)$ [20]. The satisfiability of a suitably encoded system of Boolean linear equations $\boldsymbol{Ax} = \boldsymbol{b}$ can be tested in FPR_2 by comparing the rank of $\boldsymbol{A}$ with the rank of the extension of $\boldsymbol{A}$ by $\boldsymbol{b}$ as a last column. A thorough definition of FPR_2 can be found in [20, 28].

2.4 CSPs in general

Let $\mathfrak{B}$ be a structure with finite relational signature τ. The *constraint satisfaction problem* CSP($\mathfrak{B}$) is the computational problem of deciding whether a given finite τ-structure $\mathfrak{A}$ maps homomorphically to $\mathfrak{B}$. We call $\mathfrak{B}$ a *template* of CSP($\mathfrak{B}$). Formally, we denote by CSP($\mathfrak{B}$) the class of all finite τ-structures that homomorphically map to $\mathfrak{B}$. The CSP of a τ-structure $\mathfrak{B}$ is *expressible* in a logic $\mathcal{L}$ if there exists a sentence $\phi_{\mathfrak{B}}$ in $\mathcal{L}$ that defines the complementary class co-CSP($\mathfrak{B}$) of all finite τ-structures which do not homomorphically map to $\mathfrak{B}$. A *solution* for an instance $\mathfrak{A}$ of CSP($\mathfrak{B}$) is a homomorphism $\mathfrak{A} \to \mathfrak{B}$.

Definition 2.4 (Interpretation, reducibility). Let σ, τ be finite relational signatures and Θ a set of FPR_2 formulas with first-order free variables only. A Θ-*interpretation of* τ *in* σ is a tuple $\mathcal{I}$ of σ-formulas from Θ consisting of a distinguished d-ary *domain* formula $\delta_{\mathcal{I}}(x_1,\ldots,x_d)$ and, for each n-ary atomic τ-formula $\phi(x_1,\ldots,x_n)$, an $(n\cdot d)$-ary formula $\phi_{\mathcal{I}}(x_1,\ldots,x_{n\cdot d})$. A τ-structure $\mathfrak{B}$ has an Θ-*interpretation in a* σ-*structure* $\mathfrak{A}$ if there is an Θ-interpretation $\mathcal{I}$ of τ in σ and a surjective *coordinate map* $h\colon \{\boldsymbol{t} \in A^d \mid \mathfrak{A} \models \delta_{\mathcal{I}}(\boldsymbol{t})\} \to B$ such that, for every atomic τ-formula $\phi(x_1,\ldots,x_n)$ and all $\boldsymbol{t}_1,\ldots,\boldsymbol{t}_n \in A^d$, we have $\mathfrak{B} \models \phi(h(\boldsymbol{t}_1),\ldots,h(\boldsymbol{t}_n))$ iff $\mathfrak{A} \models \phi_{\mathcal{I}}(\boldsymbol{t}_1[1],\ldots,\boldsymbol{t}_1[d],\ldots,\boldsymbol{t}_n[1],\ldots,\boldsymbol{t}_n[d])$. If h is the identity map, then we write $\mathfrak{B} = \mathcal{I}(\mathfrak{A})$. Let $\mathfrak{B}$ be a σ-structure and $\mathfrak{A}$ a τ-structure. We write CSP($\mathfrak{B}$) $\leq_\Theta$ CSP($\mathfrak{A}$) and say that CSP($\mathfrak{B}$) *reduces to* CSP($\mathfrak{A}$) *under* Θ-*reducibility* if there exists a Θ-interpretation $\mathcal{I}$ of τ in $\sigma \cup \{c_1,\ldots,c_k\}$ for some constant symbols $c_1,\ldots,c_k$ fresh w.r.t. σ such that, for every σ-structure $\mathfrak{D}$ with $|D| \geq k$, we have $\mathfrak{D} \to \mathfrak{B}$ iff $\mathcal{I}(\mathfrak{C}) \to \mathfrak{A}$ for some $\sigma \cup \{c_1,\ldots,c_k\}$-expansion $\mathfrak{C}$ of $\mathfrak{D}$ by distinct constants iff $\mathcal{I}(\mathfrak{C}) \to \mathfrak{A}$ for every $\sigma \cup \{c_1,\ldots,c_k\}$-expansion $\mathfrak{C}$ of $\mathfrak{D}$ by distinct constants.

Both Θ-reducibility and Θ-interpretability, seen as binary relations, are transitive if Θ is any of the standard logical fragments or extensions of FO we have mentioned so far. The following reducibility result was obtained in [1] for finite-domain CSPs. A close inspection of the original proof reveals that the statement holds for infinite-domain CSPs as well.

Theorem 2.5 (Atserias, Bulatov, and Dawar [1]). *Let* $\mathfrak{A}$ *and* $\mathfrak{B}$ *be structures with finite relational signatures such that* $\mathfrak{B}$ *is pp-interpretable in* $\mathfrak{A}$. *Then* CSP($\mathfrak{B}$) $\leq_{\text{Datalog}}$ CSP($\mathfrak{A}$).

Clearly, the requirement of pp-interpretability in Theorem 2.5 can be replaced with the more general notion of pp-constructibility [5, 14]. A relational structure $\mathfrak{B}$ can be *pp-constructed* from $\mathfrak{A}$ if there exists a sequence $\mathfrak{C}_1, \mathfrak{C}_2, \ldots, \mathfrak{C}_k$ such that $\mathfrak{C}_1 = \mathfrak{A}$, $\mathfrak{C}_k = \mathfrak{B}$ and, for every $1 \leq i < k$, $\mathfrak{C}_i$ admits a pp-interpretation of $\mathfrak{C}_{i+1}$, or $\mathfrak{C}_i$ is homomorphically equivalent to $\mathfrak{C}_{i+1}$. What is not clear is whether $\leq_{\text{Datalog}}$ actually preserves the expressibility of CSPs in Datalog / FP / FPC / FPR_2, since [1] only states so for $C^\omega_{\infty\omega}$. This is indeed true, see Corollary 2.6; the proof of the second part of Corollary 2.6 is elementary in all cases except for Datalog where it relies on Theorem 2 from [23].

Corollary 2.6. *Let* $\mathfrak{A}$ *and* $\mathfrak{B}$ *be structures with finite relational signatures such that* $\mathfrak{B}$ *can be pp-constructed from* $\mathfrak{A}$. *Then* CSP($\mathfrak{B}$) $\leq_{\text{Datalog}}$ CSP($\mathfrak{A}$). *Moreover,* $\leq_{\text{Datalog}}$ *preserves the expressibility of CSPs in* Datalog, FP, FPC, *or* FPR_2.

We now introduce a formalism that simplifies the presentation of known algorithms for TCSPs.

Definition 2.7 (Projections of instances). Let $\mathfrak{B}$ be a structure with finite relational signature τ and $\mathfrak{A}$ an instance of CSP($\mathfrak{B}$). Let R be an n-ary symbol from τ. The *projection* of $R^{\mathfrak{B}}$ to $I \subseteq [n]$, denoted by $\text{pr}_I(R^{\mathfrak{B}})$, is the $|I|$-ary relation defined in $\mathfrak{B}$ by the pp-formula $\exists_{j\in[n]\setminus I} x_j.\, R(x_1,\ldots,x_n)$. We call it *proper* if $I \neq \emptyset$, and *trivial* if it represents the relation $B^{|I|}$. In this paper, we assume that the set of relations of $\mathfrak{B}$ is always closed under taking projections, that is, we assume that τ contains symbols $\text{pr}_I R$ for all possible projections $\text{pr}_I(R^{\mathfrak{B}})$. Note that this convention neither leads to a different set of polymorphisms for $\mathfrak{B}$ (Proposition 2.1) nor does it influence the expressibility of CSP($\mathfrak{B}$) in any of the logics Datalog, FP, FPC, or FPR_2 (Corollary 2.6). The *projection* of $\mathfrak{A}$ to $V \subseteq A$ is the τ-structure $\text{pr}_V(\mathfrak{A})$ obtained as follows.

For every n-ary symbol $R \in \tau$ and every tuple $\boldsymbol{t} \in R^{\mathfrak{A}}$, we remove $\boldsymbol{t}$ from $R^{\mathfrak{A}}$ and add the tuple $\mathrm{pr}_I(\boldsymbol{t})$ to $(\mathrm{pr}_I R)^{\mathfrak{A}}$ for $I := \{i \in [n] \mid \boldsymbol{t}[i] \in V\}$. Then we replace the domain with V.

2.5 Temporal CSPs

A *temporal constraint language* (TCL) is a structure $\mathfrak{B}$ with domain $\mathbb{Q}$ all of whose relations are FO-definable in $(\mathbb{Q}; <)$. As the structure $(\mathbb{Q}; <)$ is homogeneous, every order preserving map between two finite subsets of $\mathbb{Q}$ can be extended to an automorphism of all TCLs due to Proposition 2.1. The relations of a TCL are called *temporal*. The *dual* of a k-ary temporal relation R is defined as $\{(-\boldsymbol{t}[1], \dots, -\boldsymbol{t}[k]) \mid \boldsymbol{t} \in R\}$. The *dual* of a TCL is the TCL whose relations are precisely the duals of the relations of the original one. Note that every TCL is homomorphically equivalent to its dual via $x \mapsto -x$, which means that both structures have the same CSP. The CSP of a TCL is called a *temporal CSP* (TCSP).

Definition 2.8 (Min-sets and free sets)**.** We define the *min-indicator function* $\chi\colon \mathbb{Q}^k \to \{0,1\}^k$ by setting $\chi(\boldsymbol{t})[i] := 1$ iff $\boldsymbol{t}[i]$ is a minimal entry in $\boldsymbol{t}$; we call $\chi(\boldsymbol{t}) \in \{0,1\}^k$ the *min-tuple* of $\boldsymbol{t} \in \mathbb{Q}^k$. The *min-set* of a tuple $\boldsymbol{t} \in \mathbb{Q}^k$ is defined as the set $\{i \in [k] \mid \chi(\boldsymbol{t})[i] = 1\}$. Let $\mathfrak{B}$ be a TCL and $\mathfrak{A}$ an instance of CSP($\mathfrak{B}$). A *free set* of $\mathfrak{A}$ is a non-empty subset $F \subseteq A$ such that, for every k-ary tuple $\boldsymbol{t}' \in R^{\mathfrak{A}}$, either no entry of $\boldsymbol{t}'$ is contained in F, or there exists a tuple $\boldsymbol{t} \in R^{\mathfrak{B}}$ which has the set $\{i \in [k] \mid \boldsymbol{t}'[i] \in F\}$ as its min-set.

2.6 Clones

The set of all polymorphisms of a relational structure $\mathfrak{A}$, denoted by Pol($\mathfrak{A}$), forms an algebraic structure called a *clone* w.r.t. compositions for maps of all arities. For instance, the clone Pol($\{0,1\}$; 1IN3) where $1\mathrm{IN}k := \{\boldsymbol{t} \in \{0,1\}^k \mid \boldsymbol{t}[i] = 1$ for exactly one $i \in [k]\}$ consists of all projection maps on $\{0,1\}$, and is called the *projection clone* [3].

Definition 2.9. A map $\xi\colon \mathrm{Pol}(\mathfrak{A}) \to \mathrm{Pol}(\mathfrak{B})$ for structures $\mathfrak{A}$ and $\mathfrak{B}$ is called

- a *clone homomorphism* (or we say that ξ *preserves identities*) if it preserves arities, projections, and compositions, that is, $\xi(f(g_1,\dots,g_n)) = \xi(f)(\xi(g_1),\dots,\xi(g_n))$ holds for all n-ary f and m-ary $g_1,\dots,g_n$ from Pol($\mathfrak{A}$),
- a *h1 clone homomorphism* (or we say that ξ *preserves h1 identities*) if it preserves arities, projections and those compositions where $g_1,\dots,g_n$ are projections,
- *uniformly continuous* if for all finite $B' \subseteq B$ there exists a finite $A' \subseteq A$ such that if $f, g \in \mathrm{Pol}(\mathfrak{A})$ of the same arity agree on A', then $\xi(f)$ and $\xi(g)$ agree on B'.

In the language of clones, the recently closed finite-domain CSPs tractability conjecture can be reformulated as follows: the polymorphism clone of a finite structure $\mathfrak{A}$ either admits a height-one (h1) clone homomorphism to the projection clone in which case CSP($\mathfrak{A}$) is NP-complete, or it does not and CSP($\mathfrak{A}$) is polynomial-time tractable [5]. The former is the case iff $\mathfrak{A}$ pp-constructs all finite structures. For detailed information about clones and clone homomorphisms we refer the reader to [5, 14].

2.7 Polymorphisms of TCLs

The following notions were used in the P versus NP-complete complexity classification of TCSPs [11]. Let MIN denote the binary minimum operation on $\mathbb{Q}$. The *dual* of a k-ary operation f on $\mathbb{Q}$ is the map $(x_1,\dots,x_k) \mapsto -f(-x_1,\dots,-x_n)$. Let us fix any endomorphisms α, β, γ of $(\mathbb{Q}; <)$ such that $\alpha(x) < \beta(x) < \gamma(x) < \alpha(x+\varepsilon)$ for every $x \in \mathbb{Q}$ and every $\varepsilon \in \mathbb{Q}_{>0}$. Such unary operations can be constructed inductively, see the paragraph below Lemma 26 in [11]. Then MI is the binary operation on $\mathbb{Q}$ defined by

$$\textsc{mi}(x,y) := \begin{cases} \alpha(\textsc{min}(x,y)) & \text{if } x = y,\\ \beta(\textsc{min}(x,y)) & \text{if } x > y,\\ \gamma(\textsc{min}(x,y)) & \text{if } x < y,\end{cases}$$

and MX is the binary operations on $\mathbb{Q}$ defined by

$$\textsc{mx}(x,y) := \begin{cases} \alpha(\textsc{min}(x,y)) & \text{if } x \neq y,\\ \beta(\textsc{min}(x,y)) & \text{if } x = y.\end{cases}$$

The main property of the operations MI and MX is that they both refine the kernel of the binary minimum operation on $\mathbb{Q}$. Let LL be an arbitrary binary operation on $\mathbb{Q}$ such that $\textsc{ll}(a,b) < \textsc{ll}(a',b')$ if

- $a \leq 0$ and $a < a'$, or
- $a \leq 0$ and $a = a'$ and $b < b'$, or
- $a, a' > 0$ and $b < b'$, or
- $a > 0$ and $b = b'$ and $a < a'$.

Theorem 2.10 (Bodirsky and Kára [11, 15])**.** *Let $\mathfrak{B}$ be a TCL. Either $\mathfrak{B}$ is preserved by* MIN, MI, MX, LL, *the dual of one of these operations, or a constant operation and* CSP($\mathfrak{B}$) *is in P, or $\mathfrak{B}$ pp-constructs all finite structures and* CSP($\mathfrak{B}$) *is NP-complete.*

There are two additional operations that appear in correctness proofs of algorithms for TCSPs; PP is an arbitrary binary operation on $\mathbb{Q}$ that satisfies $\textsc{pp}(a,b) \leq \textsc{pp}(a',b')$ iff $a \leq 0$ and $a \leq a'$, or $0 < a$, $0 < a'$, and $b \leq b'$, and LEX is an arbitrary binary operation on $\mathbb{Q}$ that satisfies $\textsc{lex}(a,b) < \textsc{lex}(a',b')$ iff $a < a'$, or $a = a'$ and $b < b'$. If a TCL is preserved by MIN, MI, or MX, then it is preserved by PP, and if a TCL is preserved by LL, then it is preserved by LEX [11].

3 Fixed-point algorithms for TCSPs

In this section, we discuss the expressibility in FP for some particularly chosen TCSPs that are provably in P. By Theorem 2.10, a TCSP is polynomial-time tractable if its template is preserved by one of the operations MIN, MI, MX, or LL. In the case of MIN, the known algorithm from [11] can be formulated as an FP algorithm. In the case of MI, the known algorithm from [11] cannot be implemented in FP as it involves choices of arbitrary elements. We show that there exists a choiceless version that can be turned into an FP sentence. In the case of LL, the known algorithm from [12] cannot be implemented in FP for the same reason as in the case of MI. Again, we show that there exists a choiceless algorithm. In the case of MX, the known algorithm from [11] cannot be turned into an FP sentence because it relies on the use of linear algebra. We show in Section 4 that, in general, the CSP of a TCL preserved by MX cannot be expressed in FP but it can be expressed in the logic FPR_2.

We first describe a procedure for temporal languages preserved by PP as it appears in [11], and then the choiceless version that is necessary for translation into an FP sentence.

Let $\mathfrak{A}$ be an instance of $\mathrm{CSP}(\mathfrak{B})$. The original procedure searches for a non-empty set $S \subseteq A$ for which there exists a solution $\mathfrak{A} \to \mathfrak{B}$ under the assumption that the projection of $\mathfrak{A}$ to $A \setminus S$ has a solution as an instance of $\mathrm{CSP}(\mathfrak{B})$. It was shown in [11] that S has this property if it is a free set of $\mathfrak{A}$, and that $\mathfrak{A} \not\to \mathfrak{B}$ if no free set of $\mathfrak{A}$ exists. We improve the original result by showing that the same holds if we replace "a free set" in the statement above with "a non-empty union of free sets".

Proposition 3.1. *Let $\mathfrak{A}$ be an instance of* $\mathrm{CSP}(\mathfrak{B})$ *for some TCL $\mathfrak{B}$ preserved by* PP *and S a union of free sets of $\mathfrak{A}$. Then $\mathfrak{A}$ has a solution iff* $\mathrm{pr}_{A\setminus S}(\mathfrak{A})$ *has a solution.*

The above proposition leads to the desired choiceless version of the original algorithm. Suitable Ptime procedures for finding unions of free sets for TCSPs with a template preserved by MIN, MI, LL, or MX exist by the results of [11], and they generally exploit the algebraic structure of the CSP that is witnessed by one of these operations.

The following lemma in combination with Corollary 2.6 shows that instead of presenting an FP algorithm for each TCSP with a template preserved by MIN, it suffices to present one for $\mathrm{CSP}(\mathbb{Q}; \mathrm{R}^{\geq}_{\mathrm{MIN}}, >)$ where

$$\mathrm{R}^{\geq}_{\mathrm{MIN}} := \{(x, y, z) \in \mathbb{Q}^3 \mid x \geq y \vee x \geq z\}.$$

Lemma 3.2. $(\mathbb{Q}; \mathrm{R}^{\geq}_{\mathrm{MIN}}, >)$ *is preserved by* MIN *and has a pp-definition of every temporal relation preserved by* MIN.

In the case of $\mathrm{CSP}(\mathbb{Q}; \mathrm{R}^{\geq}_{\mathrm{MIN}}, >)$, the suitable procedure from [11] for finding free sets can clearly be implemented in FP.

Corollary 3.3. $\mathrm{CSP}(\mathbb{Q}; \mathrm{R}^{\geq}_{\mathrm{MIN}}, >)$ *is expressible in* FP.

The following lemma in combination with Corollary 2.6 shows that instead of presenting an FP algorithm for each TCSP with a template preserved by MI, it suffices to present one for $\mathrm{CSP}(\mathbb{Q}; \mathrm{R}_{\mathrm{MI}}, \mathrm{S}_{\mathrm{MI}}, \neq)$ where

$$\mathrm{R}_{\mathrm{MI}} := \{(x, y, z) \in \mathbb{Q}^3 \mid x > y \vee x \geq z\},$$
$$\mathrm{S}_{\mathrm{MI}} := \{(x, y, z) \in \mathbb{Q}^3 \mid x \neq y \vee x \geq z\}.$$

Lemma 3.4. $(\mathbb{Q}; \mathrm{R}_{\mathrm{MI}}, \mathrm{S}_{\mathrm{MI}}, \neq)$ *is preserved by* MI *and has a pp-definition of every temporal relation preserved by* MI.

In the case of $(\mathbb{Q}; \mathrm{R}_{\mathrm{MI}}, \mathrm{S}_{\mathrm{MI}}, \neq)$, the suitable procedure from [11] for finding free sets can be implemented in FP if it returns the union of all free sets computed during the main loop instead of a single free set.

Corollary 3.5. $\mathrm{CSP}(\mathbb{Q}; \mathrm{R}_{\mathrm{MI}}, \mathrm{S}_{\mathrm{MI}}, \neq)$ *is expressible in* FP.

If a TCL $\mathfrak{B}$ is preserved by LL, then it is also preserved by LEX, but not necessarily by PP. In general, the choiceless procedure based on Proposition 3.1 is not correct for $\mathrm{CSP}(\mathfrak{B})$. We show that there exists a modified version of this procedure, motivated by the approach of repeated contractions from [12] for TCSPs whose template is preserved by LEX, and this version is correct for $\mathrm{CSP}(\mathfrak{B})$.

Let $\mathfrak{A}$ be an instance of $\mathrm{CSP}(\mathfrak{B})$. We repeatedly simulate on $\mathfrak{A}$ the choiceless procedure based on Proposition 3.1 and, every time a union S of free sets is computed, we contract in $\mathfrak{A}$ all variables in every free set within S that is minimal w.r.t. set inclusion among all existing free sets in the current projection. This loop terminates when a fixed-point is reached where $\mathfrak{A}$ no longer changes in which case we accept.

The following lemma in combination with Corollary 2.6 shows that instead of presenting an FP algorithm for each TCSP with a template preserved by LL, it suffices to present one for $\mathrm{CSP}(\mathbb{Q}; \mathrm{R}_{\mathrm{LL}}, \mathrm{S}_{\mathrm{LL}}, \neq)$ where

$$\mathrm{R}_{\mathrm{LL}} := \{(x, y, z) \in \mathbb{Q}^3 \mid x > y \vee x > z \vee x = y = z\},$$
$$\mathrm{S}_{\mathrm{LL}} := \{(x, y, z, w) \in \mathbb{Q}^4 \mid x \neq y \vee z \geq w\}.$$

Lemma 3.6. $(\mathbb{Q}; \mathrm{R}_{\mathrm{LL}}, \mathrm{S}_{\mathrm{LL}}, \neq)$ *is preserved by* LL *and has a pp-definition of every temporal relation preserved by* LL.

In the case of $\mathrm{CSP}(\mathbb{Q}; \mathrm{R}_{\mathrm{LL}}, \mathrm{S}_{\mathrm{LL}}, \neq)$, we can use the same FP procedure for finding free sets from [11] that we use for instances of $\mathrm{CSP}(\mathbb{Q}; \mathrm{R}_{\mathrm{MI}}, \mathrm{S}_{\mathrm{MI}}, \neq)$. Clearly, repeated contraction of minimal free sets can be implemented in FP as well because minimality is a choiceless condition.

Corollary 3.7. $\mathrm{CSP}(\mathbb{Q}; \mathrm{R}_{\mathrm{LL}}, \mathrm{S}_{\mathrm{LL}}, \neq)$ *is expressible in* FP.

4 A TCSP in FPR_2 which is not in FP

Let X be the temporal relation as defined in the introduction. In this section, we show that $\mathrm{CSP}(\mathbb{Q}; \mathrm{X})$ is expressible in FPR_2 (Proposition 4.2) but inexpressible in FPC (Theorem 4.11). Moreover, the following lemma in combination with Corollary 2.6 shows that instead of presenting an FPR_2 algorithm for each TCSP with a template preserved by MX, it suffices to present one for $\mathrm{CSP}(\mathbb{Q}; \mathrm{X})$.

Lemma 4.1. *The structure* $(\mathbb{Q}; \mathrm{X})$ *is preserved by* MX *and has a pp-definition of every temporal relation preserved by* MX.

The polynomial-time algorithms from [11] for CSPs of TCLs preserved by MX and for CSPs of TCLs preserved by MI only differ in the way how one determines if a variable is contained in a free set in the current projection. Thus the expressibility of $\mathrm{CSP}(\mathbb{Q}; X)$ in FPR_2 can be shown using the same approach as in the first part of Section 3 via Proposition 3.1 if the suitable procedure from [11] for finding free sets can be implemented in FPR_2. This is possible by encoding systems of Boolean linear equations in FPR_2 similarly as it is done in the case of symmetric reachability in directed graphs in the paragraph above Corollary III.2. in [20].

Corollary 4.2. $\mathrm{CSP}(\mathbb{Q}; \mathrm{X})$ *is expressible in* FPR_2.

Interestingly, the inexpressibility of $\mathrm{CSP}(\mathbb{Q}; X)$ in FPC cannot be shown by giving a pp-construction of systems of Boolean linear equations and utilizing the inexpressibility result of Atserias, Bulatov, and Dawar [1] (see Corollary 7.14). For this reason we resort to the standard strategy of showing that $\mathrm{CSP}(\mathbb{Q}; \mathrm{X})$ has unbounded counting width and then applying Theorem 2.3 [22]. The *counting width* of $\mathrm{CSP}(\mathfrak{B})$ for a τ-structure $\mathfrak{B}$ is the function that assigns to each $n \in \mathbb{N}$ the minimum value k for which there is a τ-sentence ϕ in C^k such that, for every τ-structure $\mathfrak{A}$ with $|A| \leq n$, $\mathfrak{A} \models \phi$ iff $\mathfrak{A} \to \mathfrak{B}$. Recall the reflexive and transitive relation $\Rightarrow_k$ and the equivalence relation $\equiv_{C^k}$ from Definition 2.2. In our proof of inexpressibility, we utilize both the *existential k-pebble game* which characterizes $\Rightarrow_k$, and the *bijective k-pebble game* which characterizes $\equiv_{C^k}$. See [2] for details about the approach to $\Rightarrow_k$ and $\equiv_{C^k}$ via model-theoretic games. For easier understanding we reformulate $\mathrm{CSP}(\mathbb{Q}; \mathrm{X})$ as a certain decision problem for systems of linear equations over $\mathbb{Z}_2$.

The satisfiability problem for systems of linear equations $\boldsymbol{Ax} = \boldsymbol{b}$ over a finite Abelian group $\mathcal{G}$ with at most k variables per equation can be formulated as $\mathrm{CSP}(\mathfrak{E}_{\mathcal{G},k})$ where $\mathfrak{E}_{\mathcal{G},k}$ is the structure over the domain G of $\mathcal{G}$ with the relations $\{\boldsymbol{t} \in G^j \mid \sum_{i\in[j]} \boldsymbol{t}[i] = a\}$ for every $j \leq k$ and $a \in G$ [1]. In the present paper, we denote the instance of $\mathrm{CSP}(\mathfrak{E}_{\mathbb{Z}_2,k})$ that is derived from a system of Boolean linear equations $\boldsymbol{Ax} = \boldsymbol{b}$ by $\mathfrak{I}_{\boldsymbol{A},\boldsymbol{b}}$. Recall the decision problem 3-Ord-Xor-Sat defined in the introduction. Formally, we understand 3-Ord-Xor-Sat as a proper subset of $\mathrm{CSP}(\mathfrak{E}_{\mathbb{Z}_2,3})$. Now we turn our attention back to the original problem. The structure $(\mathbb{Q}; \mathrm{X})$ meets the sole requirement of being a TCL preserved by mx for applicability of Theorem 42 from [11]. This result provides us with a polynomial-time algorithm for $\mathrm{CSP}(\mathbb{Q}; \mathrm{X})$. An inspection of the algorithm reveals that $\mathrm{CSP}(\mathbb{Q}; \mathrm{X})$ and 3-Ord-Xor-Sat are the same problems up to renaming of symbols. We use the probabilistic construction of 3-multipedes from [6, 30] as a black box for extracting certain systems of Boolean linear equations that represent instances of $\mathrm{CSP}(\mathbb{Q}; \mathrm{X})$. More specifically, we use the reduction of the isomorphism problem for 3-multipedes to the satisfiability of a system of linear equations over $\mathbb{Z}_2$ with 3 variables per equation from the proof of Theorem 23 in [6].

The following concepts were introduced in [30]; we mostly follow the terminology in [6].

Definition 4.3. A 3-*multipede* is a finite relational structure $\mathfrak{M}$ with the signature $\{<, E, H\}$, where $<, E$ are binary symbols and H is a ternary symbol, such that $\mathfrak{M}$ satisfies the following axioms. The domain of $\mathfrak{M}$ has a partition into *segments* $\mathcal{S}(\mathfrak{M})$ and *feet* $\mathcal{F}(\mathfrak{M})$ such that $<^{\mathfrak{M}}$ is a linear order on $\mathcal{S}(\mathfrak{M})$, and $E^{\mathfrak{M}}$ is the graph of a surjective function $\mathrm{seg}\colon \mathcal{F}(\mathfrak{M}) \to \mathcal{S}(\mathfrak{M})$ with $|\mathrm{seg}^{-1}(x)| = 2$ for every $x \in \mathcal{S}(\mathfrak{M})$. For every $\boldsymbol{t} \in H^{\mathfrak{M}}$, either the entries of $\boldsymbol{t}$ are contained in $\mathcal{S}(\mathfrak{M})$ and we call $\boldsymbol{t}$ a *hyperedge*, or they are contained in $\mathcal{F}(\mathfrak{M})$ and we call $\boldsymbol{t}$ a *positive triple*. The relation $H^{\mathfrak{M}}$ is *fully symmetric*, that is, closed under all permutations of entries, and only contains triples with pairwise distinct entries. For every positive triple $\boldsymbol{t}$, the triple $(\mathrm{seg}(\boldsymbol{t}[1]), \mathrm{seg}(\boldsymbol{t}[2]), \mathrm{seg}(\boldsymbol{t}[3]))$ is a hyperedge. If $\boldsymbol{t} \in H^{\mathfrak{M}}$ is an hyperedge where $\mathrm{seg}^{-1}(\boldsymbol{t}[i]) = \{x_{i,0}, x_{i,1}\}$ for every $i \in [3]$, then we require that exactly 4 elements of the set $\{(x_{1,i}, x_{2,j}, x_{3,k}) \mid i, j, k \in \{0, 1\}\}$ are positive triples. We also require that, for each pair of triples $(x_{1,i}, x_{2,j}, x_{3,k}), (x_{1,i'}, x_{2,j'}, x_{3,k'})$ from the set above, we have $(i - i') + (j - j') + (k - k') = 0 \bmod 2$. A 3-multipede $\mathfrak{M}$ is:

- *odd* if for each $\emptyset \subsetneq X \subseteq \mathcal{S}(\mathfrak{M})$ there is a hyperedge $\boldsymbol{t} \in H^{\mathfrak{M}}$ such that $|\{\boldsymbol{t}[1], \boldsymbol{t}[2], \boldsymbol{t}[3]\} \cap X|$ is odd,
- *k-meager* if for each $\emptyset \subsetneq X \subseteq \mathcal{S}(\mathfrak{M})$ of size at most $2k$ we have $|X| > 2 \cdot |H^{\mathfrak{M}} \cap X^3|$.

Example 4.4. The 3-multipede $\mathfrak{M}$ from Figure 1 has segments $\mathcal{S}(\mathfrak{M}) = \mathbb{Z}_9$, feet $\mathcal{F}(\mathfrak{M}) = \mathbb{Z}_9 \times \mathbb{Z}_2$, $<^{\mathfrak{M}}$ is the linear order $0 < \cdots < 8$, $E^{\mathfrak{M}} = \{(\boldsymbol{t}, s) \in \mathcal{F}(\mathfrak{M}) \times \mathcal{S}(\mathfrak{M}) \mid \boldsymbol{t}[1] = s\}$, and $H^{\mathfrak{M}}$ consists of: all triples $\boldsymbol{s} \in \mathcal{S}(\mathfrak{M})^3$ with $\boldsymbol{s}[2] = \boldsymbol{s}[1] + 2 \bmod 9$ and $\boldsymbol{s}[3] = \boldsymbol{s}[1] + 5 \bmod 9$, and all triples $(\boldsymbol{t}_1, \boldsymbol{t}_2, \boldsymbol{t}_3) \in \mathcal{F}(\mathfrak{M})^3$ such that $\boldsymbol{s} := (\boldsymbol{t}_1[1], \boldsymbol{t}_2[1], \boldsymbol{t}_3[1])$ satisfies the previous condition and additionally $\boldsymbol{t}_1[2] + \boldsymbol{t}_2[2] + \boldsymbol{t}_3[2] = 0 \bmod 2$. Note that the hyperedges of $\mathfrak{M}$ do not overlap on more than one segment, because the minimal distances between two elements of an hyperedge are $2, 3$, or $4 \bmod 9$. This directly implies that $\mathfrak{M}$ is 2-meager. Using Gaussian elimination, one can check that the system of linear equations $\boldsymbol{Ax} = \boldsymbol{0}$ over $\mathbb{Z}_2$, where $\boldsymbol{A}$ is the incidence matrix of the hyperedges of $\mathfrak{M}$ on the segments of $\mathfrak{M}$, only admits the trivial solution. From this fact it follows that $\mathfrak{M}$ is odd. Otherwise, suppose that there exists a non-empty $X \subseteq \mathcal{S}(\mathfrak{M})$ witnessing that $\mathfrak{M}$ is not odd. Then $\boldsymbol{Ax} = \boldsymbol{0}$ is satisfied by the non-trivial assignment that maps $\boldsymbol{x}[s]$ to 1 iff $s \in X$, which yields a contradiction.

The following two statements are crucial for our application of 3-multipedes in the context of $\mathrm{CSP}(\mathbb{Q}; \mathrm{X})$. An automorphism is called *trivial* if it is an identity map.

Proposition 4.5 ([6], Proposition 17)**.** *Odd* 3*-multipedes have no non-trivial automorphisms.*

For a pair $\mathfrak{M}_1, \mathfrak{M}_2$ of 3-multipedes we say that $\mathfrak{M}_2$ is obtained from $\mathfrak{M}_1$ by *transposing the feet* of a segment s of $\mathfrak{M}_1$ ([30], p. 12) if the domains and relations of both 3-multipedes coincide up to the following property: a positive triple $\boldsymbol{t}$ of $\mathfrak{M}_1$ is a positive triple of $\mathfrak{M}_2$ iff $s \notin \{\mathrm{seg}(\boldsymbol{t}[1]), \mathrm{seg}(\boldsymbol{t}[2]), \mathrm{seg}(\boldsymbol{t}[3])\}$.

Lemma 4.6 ([30], Lemma 4.5)**.** *For any* $k \in \mathbb{N}_{>0}$*, let* $\mathfrak{M}_1, \mathfrak{M}_2$ *be two* $2k$*-meager* 3*-multipedes such that one is obtained from the other by transposing the feet of one segment. Then* $\mathfrak{M}_1 \equiv_{C^k_{\infty\omega}} \mathfrak{M}_2$*. The statement holds even if we extend the signature by means of individual constants for every segment.*

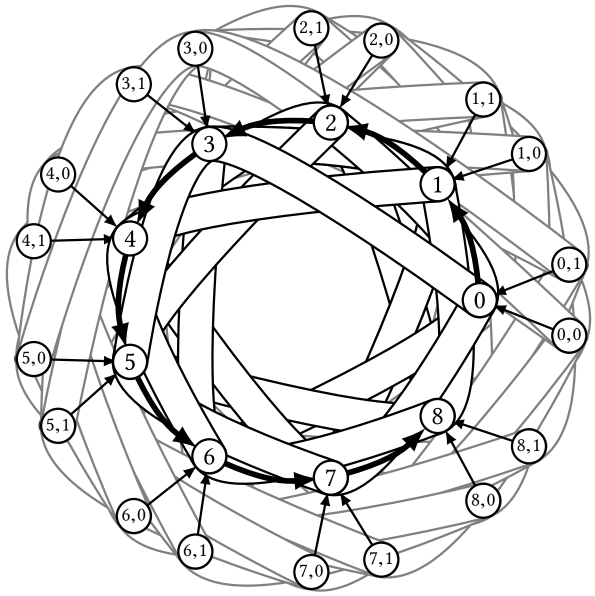

Figure 1. An odd 2-meager 3-multipede.

Proposition 4.7 ([6], Proposition 18)**.** *For every $k \in \mathbb{N}_{>0}$, there exists an odd k-meager 3-multipede.*

Let $\mathfrak{M}_1, \mathfrak{M}_2$ be any two 3-multipedes with $M_1 = M_2$, $<^{\mathfrak{M}_1} = <^{\mathfrak{M}_2}$, $E^{\mathfrak{M}_1} = E^{\mathfrak{M}_2}$, and $H^{\mathfrak{M}_1} \cap \mathcal{S}(\mathfrak{M}_1)^3 = H^{\mathfrak{M}_2} \cap \mathcal{S}(\mathfrak{M}_2)^3$. Fix an arbitrary bijection $f\colon M_1 \to M_2$ that preserves $<^{\mathfrak{M}_1}$ and $E^{\mathfrak{M}_1}$. Let $\boldsymbol{t}_1, \ldots, \boldsymbol{t}_m$ be an enumeration of $H^{\mathfrak{M}_1} \cap \mathcal{S}(\mathfrak{M}_1)^3$ and $s_1, \ldots, s_\ell$ be an enumeration of $\mathcal{S}(\mathfrak{M}_1)$. Consider the matrix $\boldsymbol{A}' \in \{0,1\}^{m \times \ell}$ with $\boldsymbol{A}'[i,j] = 1$ iff s_j occurs in some entry of $\boldsymbol{t}_i$, and the tuple $\boldsymbol{b}' \in \{0,1\}^m$ with $\boldsymbol{b}'[i] = 0$ iff f preserves positive triples of $\mathfrak{M}_1$ at $\boldsymbol{t}_i$. Each isomorphism $f_X\colon \mathfrak{M}_1 \to \mathfrak{M}_2$ can be obtained from f by transposing the images of the feet at every segment from a particular subset $X \subseteq \mathcal{S}(\mathfrak{M}_1)$, i.e.,

$$f_X(x) = \begin{cases} f(y) & \text{if } \mathrm{seg}^{-1}(s) = \{x, y\} \text{ for some } s \in X, \\ f(x) & \text{otherwise.} \end{cases}$$

For every $X \subseteq \mathcal{S}(\mathfrak{M}_1)$, let $\boldsymbol{t}_X \in \{0,1\}^\ell$ be the tuple defined by $\boldsymbol{t}_X[i] \coloneqq 1$ iff $s_i \in X$. The following lemma is a simple consequence of the definition of a 3-multipede.

Lemma 4.8 ([6], the proof of Theorem 23)**.** *For every $X \subseteq \mathcal{S}(\mathfrak{M}_1)$, the mapping f_X is an isomorphism from $\mathfrak{M}_1$ to $\mathfrak{M}_2$ iff $\boldsymbol{t}_X$ is a solution to the system $\boldsymbol{A}'\boldsymbol{x} = \boldsymbol{b}'$.*

Keeping this construction in mind, we can derive the following statement about systems of Boolean linear equations.

Proposition 4.9. *For every $k \geq 2$ there exist Boolean vectors $\boldsymbol{b}, \boldsymbol{c} \neq \boldsymbol{0}$ and a Boolean matrix $\boldsymbol{A}$ with 3 non-zero entries per row such that*

1. *$\boldsymbol{A}\boldsymbol{x} = \boldsymbol{0}$ only admits the trivial solution,*
2. *$\boldsymbol{A}\boldsymbol{x} = \boldsymbol{c}$ has a solution and $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$ has no solution,*
3. $\mathfrak{I}_{\boldsymbol{A},\boldsymbol{b}} \equiv_{C^{2k}} \mathfrak{I}_{\boldsymbol{A},\boldsymbol{c}}$.

Proof. For a given $k \geq 2$, let $\mathfrak{M}_1$ be an odd $12k$-meager 3-multipede whose existence follows from Proposition 4.7. We obtain a second 3-multipede $\mathfrak{M}_2$ from $\mathfrak{M}_1$ by transposing the feet of an arbitrary segment s of $\mathfrak{M}_1$. Fix any bijection $f\colon M_1 \to M_2$ that preserves $<^{\mathfrak{M}_1}$ and $E^{\mathfrak{M}_1}$, and let $\boldsymbol{A}'\boldsymbol{x} = \boldsymbol{b}'$ be the system of linear equations over $\mathbb{Z}_2$ derived from $\mathfrak{M}_1$, $\mathfrak{M}_2$ and f using the construction described in the paragraph above Lemma 4.8. By the definition of $\mathfrak{M}_2$, every isomorphism $\mathfrak{M}_1 \to \mathfrak{M}_2$ yields a non-trivial automorphism of $\mathfrak{M}_1$. Since $\mathfrak{M}_1$ is odd, it cannot have a non-trivial automorphism due to Proposition 4.5. Hence $\mathfrak{M}_1$ and $\mathfrak{M}_2$ are non-isomorphic. But then $\boldsymbol{A}'\boldsymbol{x} = \boldsymbol{b}'$ cannot have a solution due to Lemma 4.8. Now consider the situation where $\mathfrak{M}_2$ was a copy of $\mathfrak{M}_1$ instead, and we chose f to be the identity map. Then f would preserve all positive triples of $\mathfrak{M}_1$. Thus the system of linear equations over $\mathbb{Z}_2$ obtained from $\mathfrak{M}_1$, $\mathfrak{M}_2$ and f using the identical construction as in the non-isomorphic case would be precisely $\boldsymbol{A}'\boldsymbol{x} = \boldsymbol{0}$, the homogeneous companion of $\boldsymbol{A}'\boldsymbol{x} = \boldsymbol{b}'$. By an additional application of Lemma 4.8, $\boldsymbol{A}'\boldsymbol{x} = \boldsymbol{0}$ cannot have any non-trivial solution, because $\mathfrak{M}_1$ has no non-trivial automorphism.

Note that we have $\mathfrak{I}_{\boldsymbol{A}',1} \Rightarrow_{2k} \mathfrak{E}_{\mathbb{Z}_2,3}$ because Duplicator has the trivial winning strategy of placing all pebbles on 1 in the existential $2k$-pebble game played on $\mathfrak{I}_{\boldsymbol{A}',1}$ and $\mathfrak{E}_{\mathbb{Z}_2,3}$. We claim that also $\mathfrak{I}_{\boldsymbol{A}',\boldsymbol{b}'} \Rightarrow_{2k} \mathfrak{E}_{\mathbb{Z}_2,3}$. In the terminology of [2] we would say that $\mathfrak{I}_{\boldsymbol{A}',\boldsymbol{b}'}$ is $2k$-*locally satisfiable*. Without loss of generality, we may assume that $\mathfrak{M}_1, \mathfrak{M}_2$ have their signature expanded by constant symbols for every segment (see Lemma 4.6). For convenience, we fix an arbitrary linear order on $M_1 = M_2$, and say that x is a *left foot* and y a *right foot* of a segment s with $\mathrm{seg}^{-1}(s) = \{x, y\}$ if x is less than y w.r.t. this order. We know that Duplicator has a winning strategy in the bijective $6k$-pebble game played on $\mathfrak{M}_1$ and $\mathfrak{M}_2$. We use it to construct a winning strategy for Duplicator in the existential $2k$-pebble game played on $\mathfrak{I}_{\boldsymbol{A}',\boldsymbol{b}'}$ and $\mathfrak{E}_{\mathbb{Z}_2,3}$.

Suppose we have a position in the existential $2k$-pebble game with pebbles $1, \ldots, \ell$ placed on some $x_1, \ldots, x_\ell \in I_{\boldsymbol{A}',\boldsymbol{b}'}$ and $v_1, \ldots, v_\ell \in \{0,1\}$ for $\ell \leq 2k$. If Spoiler chooses a pebble $i > \ell$ and places it onto some $x_i \in I_{\boldsymbol{A}',\boldsymbol{b}'}$, then we consider the situation in the bijective $6k$-pebble game played on $\mathfrak{M}_1$ and $\mathfrak{M}_2$ where Spoiler places, in three succeeding rounds, a pebble i on the corresponding segment x_i of $\mathfrak{M}_1$, and two pebbles i_ℓ, i_r on its left and right foot. Since Duplicator has a winning strategy in this game, she can react by placing the pebbles i, i_ℓ, i_r on some elements y_i, y_{i_ℓ}, y_{i_r} of the second 3-multipede $\mathfrak{M}_2$. Since her placement corresponds to a partial isomorphism and the signature contains constant symbols for every segment, y_i must be a segment with the same number of predecessors as x_i with respect to the linear order on the segments, and y_{i_ℓ}, y_{i_r} must be its feet. Now if y_{i_ℓ} is the left and y_{i_r} the right foot of y_i, then Duplicator places v_i on 0 in the existential $2k$-pebble game, otherwise on 1. The case when $i \leq \ell$ corresponds to the situation when pebbles i, i_ℓ, i_r are lifted from both $\mathfrak{M}_1$ and $\mathfrak{M}_2$. Clearly Duplicator can maintain this condition, and her pebbling specifies a partial homomorphism by a local variant of Lemma 4.8.

The following construction originates from [2]. We define the system $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$ so that it contains, for each equation $x_{i_1} + x_{i_2} + x_{i_3} = b'_i$ of $\boldsymbol{A}'\boldsymbol{x} = \boldsymbol{b}'$ and all $a_1, a_2, a_3 \in \{0, 1\}$, the equation $x^{a_1}_{i_1} + x^{a_2}_{i_2} + x^{a_3}_{i_3} = b'_i + a_1 + a_2 + a_3$. Analogously we obtain $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{c}$ from $\boldsymbol{A}'\boldsymbol{x} = \boldsymbol{1}$, and $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{z}$ from $\boldsymbol{A}'\boldsymbol{x} = \boldsymbol{0}$. As a direct consequence of Lemma 2 in [2] we have

$$\mathfrak{I}_{\boldsymbol{A},\boldsymbol{b}} \equiv_{C^{2k}} \mathfrak{I}_{\boldsymbol{A},\boldsymbol{z}} \equiv_{C^{2k}} \mathfrak{I}_{\boldsymbol{A},\boldsymbol{c}}$$

while we also have $\mathfrak{I}_{\boldsymbol{A},\boldsymbol{c}} \to \mathfrak{E}_{\mathbb{Z}_2,3}$ and $\mathfrak{I}_{\boldsymbol{A},\boldsymbol{b}} \not\to \mathfrak{E}_{\mathbb{Z}_2,3}$ due to Lemma 3 in [2]. Note that $\mathfrak{I}_{\boldsymbol{A},\boldsymbol{0}}$ contains a copy of $\mathfrak{I}_{\boldsymbol{A}',\boldsymbol{0}}$ with variables x^a_i for both upper indices $a \in \{0, 1\}$, and thus $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{0}$ only admits the trivial solution. □

The following result can be derived from Proposition 4.9 by adding dummy variables based on $\boldsymbol{b}$ and $\boldsymbol{c}$ to the equations of $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{0}$ and subsequently reducing their length back to 3.

Corollary 4.10. *For every $k \geq 2$, there exist Boolean matrices $\boldsymbol{B}$ and $\boldsymbol{C}$ with 3 non-zero entries per row such that*
- *each non-empty subset F of the equations of $\boldsymbol{C}\boldsymbol{x} = \boldsymbol{0}$ has a non-trivial solution with respect to the variables that occur in F,*
- *$\boldsymbol{B}\boldsymbol{x} = \boldsymbol{0}$ has only the trivial solution,*
- *$\mathfrak{I}_{\boldsymbol{B},\boldsymbol{0}} \equiv_{C^k} \mathfrak{I}_{\boldsymbol{C},\boldsymbol{0}}$.*

Theorem 4.11. CSP($\mathbb{Q}$; X) *is inexpressible in* FPC.

Proof. We have 3-Ord-Xor-Sat = CSP($\mathbb{Q}$; X) up to renaming of symbols by inspection of the original algorithm for TCSPs with a template preserved by mx from [11]. It follows from Corollary 4.10 and Theorem 2.3 that 3-Ord-Xor-Sat is inexpressible in FPC, which completes the proof. □

5 Classification of TCSPs in FP

In this section, we classify CSPs of TCLs with respect to expressibility in fixed-point logic. We start with the case of a TCL $\mathfrak{B}$ that is not preserved by any operation mentioned in Theorem 2.10. Note that the NP-completeness of CSP($\mathfrak{B}$) is not sufficient for obtaining inexpressibility in FP. What is sufficient is the fact that $\mathfrak{B}$ pp-constructs all finite structures.

Lemma 5.1. *Let $\mathfrak{B}$ be a relational structure. If $\mathfrak{B}$ pp-constructs all finite structures, then* CSP($\mathfrak{B}$) *is inexpressible in FPC.*

Proof. In particular, $\mathfrak{B}$ pp-construct the structure $\mathfrak{E}_{\mathbb{Z}_2,3}$ whose CSP is inexpressible in FPC by Theorem 10 in [1]. Thus CSP($\mathfrak{B}$) is inexpressible in FPC by Corollary 2.6. □

We show in Theorem 5.2 that those TCLs preserved by mx for which we know that their CSP is expressible in FP by the results in Section 3 are precisely the ones unable to pp-define the relation X which we have studied in Section 4.

Theorem 5.2. *Let $\mathfrak{B}$ be a TCL preserved by* mx. *Then either $\mathfrak{B}$ admits a pp-definition of* X, *or one of the following is true:*
1. *$\mathfrak{B}$ is preserved by a constant operation,*
2. *$\mathfrak{B}$ is preserved by* min.

We are now ready for our first classification result.

Theorem 5.3. *Let $\mathfrak{B}$ be a TCL. The following are equivalent:*
1. CSP($\mathfrak{B}$) *is expressible in* FP.
2. CSP($\mathfrak{B}$) *is expressible in* FPC.
3. *$\mathfrak{B}$ does not pp-construct all finite structures and $\mathfrak{B}$ does not pp-construct* ($\mathbb{Q}$; X).
4. *$\mathfrak{B}$ is preserved by* min, mi, ll, *the dual of one of these operations, or by a constant operation.*

Proof. Let $\mathfrak{B}$ be a TCL.

(1)⇒(2): Trivial because FP is a fragment of FPC.

(2)⇒(3): Lemma 5.1 implies that $\mathfrak{B}$ does no pp-construct all finite structures; Theorem 4.11 and Corollary 2.6 show that $\mathfrak{B}$ does not pp-construct ($\mathbb{Q}$; X).

(3)⇒(4): Since $\mathfrak{B}$ does not pp-construct all finite structures, by Theorem 2.10, $\mathfrak{B}$ is preserved by min, mi, mx, ll, the dual of one of these operations, or by a constant operation. If $\mathfrak{B}$ is preserved by mx but neither by min nor by a constant operation, then $\mathfrak{B}$ pp-defines X by Theorem 5.2, a contradiction to (3). If $\mathfrak{B}$ is preserved by dual mx but neither by max nor by a constant operation, then $\mathfrak{B}$ pp-defines $-X$ by the dual version of Theorem 5.2. Since ($\mathbb{Q}$; X) and ($\mathbb{Q}$; $-X$) are homomorphically equivalent, we get a contradiction to (3) in this case as well. Thus (4) must hold for $\mathfrak{B}$.

(4)⇒(1): If $\mathfrak{B}$ has a constant polymorphism, then CSP($\mathfrak{B}$) is trivial and thus expressible in FP. If $\mathfrak{B}$ is preserved by min, mi, or ll, then every relation of $\mathfrak{B}$ is pp-definable in ($\mathbb{Q}$; $R^{\geq}_{\text{MIN}}$, >) by Lemma 3.2, or in ($\mathbb{Q}$; R_{MI}, S_{MI}, ≠) by Lemma 3.4, or in ($\mathbb{Q}$; R_{LL}, S_{LL}, ≠) by Lemma 3.6. Thus CSP($\mathfrak{B}$) is expressible in FP by Corollary 3.3, Corollary 3.5, or Corollary 3.7 combined with Corollary 2.6. Each of the previous statements can be dualized to obtain expressibility of CSP($\mathfrak{B}$) in FP if $\mathfrak{B}$ is preserved by max, dual mi, or dual ll. □

The following corollary can be derived from Theorem 5.3, Corollary 4.2, Lemma 4.1, and the fact that CSP($\mathfrak{E}_{\mathbb{Z}_3,3}$) is inexpressible in FPR_2 because 2, 3 are distinct primes [28].

Corollary 5.4. *The CSP of a TCL $\mathfrak{B}$ is expressible in* FPR_2 *iff $\mathfrak{B}$ does not pp-construct all finite structures.*

6 Classification of TCSPs in Datalog

In this section, we classify CSPs of TCLs with respect to expressibility in Datalog in several steps.

A formula over the signature {<} is *Ord-Horn* if it is a conjunction of clauses of the form $(x_1 \neq y_1) \vee \cdots \vee (x_m \neq y_m) \vee (x \leq y)$ where the last disjunct is optional (see [7]). Nebel and Bürckert [40] showed that testing satisfiability of Ord-Horn formulas can be done in polynomial time. It is easy to see that their algorithm can be formulated as a Datalog procedure. We show in Proposition 6.1 that Ord-Horn definability of the relations of a TCL preserved by one of the operations min, mi, mx, or ll can be characterized in terms of admitting certain polymorphisms. Interestingly, there is no such characterization in terms of identities for polymorphism clones (see Proposition 7.3). Proposition 6.1 is proved using the syntactic normal form for temporal relations preserved by PP from [8] and the syntactic normal form for temporal relations preserved by LL from [7].

Proposition 6.1. *A temporal relation preserved by* MIN*,* MI*,* MX*, or* LL *is Ord-Horn definable iff it is preserved by every binary injective operation on* $\mathbb{Q}$ *that preserves* $\leq$*.*

Let $\mathrm{R_{MIN}}$ be the temporal relation as defined in the introduction. Recall that $\mathrm{CSP}(\mathbb{Q}; \mathrm{R_{MIN}})$ is inexpressible in Datalog [12]. This time, the reason for inexpressibility is not unbounded counting width, but the combination of the two facts that $\mathrm{CSP}(\mathbb{Q}; \mathrm{R_{MIN}})$ admits unsatisfiable instances of arbitrarily high girth, and that all proper projections of $\mathrm{R_{MIN}}$ are trivial. We show in Theorem 6.2 that the inability of a TCL preserved by one of the operations MIN, MI, MX, or LL to pp-define $\mathrm{R_{MIN}}$ can be characterized in terms of being preserved by a constant operation, or by the operations from Proposition 6.1 which witness Ord-Horn definability.

Theorem 6.2. *Let* $\mathfrak{B}$ *be a TCL preserved by* MIN*,* MI*,* MX*, or* LL*. Then either* $\mathfrak{B}$ *admits a pp-definition of the relation* $\mathrm{R_{MIN}}$*, or one of the following is true:*

1. $\mathfrak{B}$ *is preserved by a constant operation, or*
2. $\mathfrak{B}$ *is preserved by every binary injective operation on* $\mathbb{Q}$ *that preserves* $\leq$*.*

We are ready for our second classification result; its proof is a straightforward combination of Corollary 2.6, Proposition 6.1, Theorem 6.2, and the results from previous sections.

Theorem 6.3. *Let* $\mathfrak{B}$ *be a TCL. The following are equivalent:*

1. $\mathrm{CSP}(\mathfrak{B})$ *is expressible in* Datalog.
2. $\mathfrak{B}$ *does not pp-construct all finite structures and* $\mathfrak{B}$ *does not pp-construct* $(\mathbb{Q}, \mathrm{R_{MIN}})$.
3. *Every relation of* $\mathfrak{B}$ *is Ord-Horn definable or* $\mathfrak{B}$ *has a constant polymorphism.*

7 Algebraic conditions for temporal CSPs

In this section, we consider several candidates for general algebraic criteria for expressibility of CSPs in FP and Datalog stemming from the well-developed theory of finite-domain CSPs. Our results imply that none of them can be used in the setting of ω-categorical CSPs. We also present a new simple algebraic condition which characterises expressibility of both finite-domain and temporal CSPs in FP.

h1 conditions. We call an at least binary operation f *weak near-unanimity* (WNU) if it satisfies $f(y, x, \dots, x) \approx \cdots \approx f(x, \dots, x, y)$. The requirement for existence of such an operation is an example of a height one condition; *height one* (h1) *conditions* are given by sets of identities of the form $f_1(x_1^1, \dots, x_{n_1}^1) \approx \cdots \approx f_k(x_k^1, \dots, x_{n_k}^k)$ and we have already encountered them in Theorem 1.1.

Failures of known h1 conditions. Despite their success in the setting of finite-domain CSPs, finite h1 conditions such as item (5) in Theorem 1.1 are insufficient for classification purposes in the context of ω-categorical CSPs.

Proposition 7.1. *Let* $\mathcal{L}$ *be any logic at least as expressive as the existential positive fragment of* FO*. There is no finite h1 condition capturing the expressibility of the CSPs of reducts of finitely bounded homogeneous structures in* $\mathcal{L}$*.*

Proposition 7.1 is a consequence of the proof of Theorem 1.3 in [14]. Both statements rely on a result from [18] which states that, for every finite family $\mathcal{F}$ of finite connected structures with a finite signature τ, there exists a τ-reduct $\mathfrak{U}_h(\mathcal{F})$ of a finitely bounded homogeneous structure such that $\mathfrak{U}_h(\mathcal{F})$ embeds precisely those finite τ-structures which do not contain a homomorphic image of any member of $\mathcal{F}$ (also see, e.g., the presentation in [33] and [14]).

Proof of Proposition 7.1. Suppose, on the contrary, that there exists such a condition $\mathcal{E}$. By the proof of Theorem 1.3 in [14], there exists a finite family $\mathcal{F}$ of finite connected structures with a finite signature τ such that $\mathrm{Pol}(\mathfrak{U}_h(\mathcal{F}))$ violates $\mathcal{E}$. By a standard result from database theory, $\mathfrak{A}$ homomorphically maps to $\mathfrak{B}$ iff the *canonical conjunctive query* $Q_{\mathfrak{A}}$ is true in $\mathfrak{B}$ [17]; $Q_{\mathfrak{A}}$ is the pp-sentence with existentially quantified variables x_a for every $a \in A$, and a conjunction of literals $R(x_{a_1}, \dots, x_{a_n})$ for every $(a_1, \dots, a_n) \in R^{\mathfrak{A}}$. Clearly the existential positive sentence $\phi_{\mathfrak{U}_h(\mathcal{F})} := \bigvee_{\mathfrak{A} \in \mathcal{F}} Q_{\mathfrak{A}}$ defines the complement of $\mathrm{CSP}(\mathfrak{U}_h(\mathcal{F}))$. But then $\mathrm{CSP}(\mathfrak{U}_h(\mathcal{F}))$ is expressible in $\mathcal{L}$, a contradiction. ◻

The satisfiability of h1 identities in polymorphism clones is preserved under h1 clone homomorphisms, and the satisfiability of arbitrary identities in polymorphism clones is preserved under clone homomorphisms [5]. We use the latter to show that, for Datalog, Proposition 7.1 can be strengthened to sets of arbitrary identities, see Theorem 7.2. We hereby give a negative answer to an open question from [14] concerning the existence of a fixed set of identities that would capture Datalog expressibility for ω-categorical CSPs. Recall the relation S_{LL} defined in Lemma 3.6.

Theorem 7.2. *There is no set of identities for polymorphism clones that would capture the expressibility of the CSPs of reducts of finitely bounded homogeneous structures in* Datalog.

Proof. By Proposition 7.3, every set of identities satisfiable in $\mathrm{Pol}(\mathbb{Q}; \neq, \mathrm{S_{LL}})$ is satisfiable in $\mathrm{Pol}(\mathbb{Q}; \mathrm{R_{MIN}})$. We also have that $\mathrm{CSP}(\mathbb{Q}; \neq, \mathrm{S_{LL}})$ is expressible in Datalog by Theorem 6.3 because $\neq$ and $\mathrm{S_{LL}}$ are Ord-Horn definable, and $\mathrm{CSP}(\mathbb{Q}; \mathrm{R_{MIN}})$ is inexpressible in Datalog by Theorem 5.2 in [12]. ◻

We say that an operation $f : B^n \to B$ *depends* on the i-th argument if there exist $b_1, \dots, b_n, b \in B$ with $b_i \neq b$ such that $f(b_1, \dots, b_n) \neq f(b_1, \dots, b_{i-1}, b, b_{i+1}, \dots, b_n)$.

Proposition 7.3. *There exists a uniformly continuous clone homomorphism* $\xi\colon \mathrm{Pol}(\mathbb{Q};\neq,\mathrm{S_{LL}}) \to \mathrm{Pol}(\mathbb{Q};\mathrm{R_{MIN}})$.

Proof. For an n-ary $f \in \mathrm{Pol}(\mathbb{Q};\neq,\mathrm{S_{LL}})$, let $\{i_1,\dots,i_m\} \subseteq [n]$ be the set of all indices $i \in [n]$ such that $f(x_1,\dots,x_n)$ depends on the i-th argument. Since $(\mathbb{Q};\neq,\mathrm{S_{LL}})$ has no constant polymorphism, we have $m > 0$. We define the *essential part* of f as the map $f^{ess}\colon \mathbb{Q}^m \to \mathbb{Q}$, $(x_1,\dots,x_m) \mapsto f(x_{\mu_f(1)},\dots,x_{\mu_f(n)})$ where $\mu_f\colon [n] \to [m]$ is any map that satisfies $\mu_f(i_\ell) = \ell$ for each $\ell \in [m]$. By Proposition 6.1.4 in [7], f^{ess} is injective because the relation $\{\boldsymbol{t} \in \mathbb{Q}^4 \mid \boldsymbol{t}[1] = \boldsymbol{t}[2] \Rightarrow \boldsymbol{t}[3] = \boldsymbol{t}[4]\}$ is pp-definable in $(\mathbb{Q};\neq,\mathrm{S_{LL}})$. We define $\xi\colon \mathrm{Pol}(\mathbb{Q};\neq,\mathrm{S_{LL}}) \to \mathrm{Pol}(\mathbb{Q};\mathrm{R_{MIN}})$ so that it sends an n-ary operation $f \in \mathrm{Pol}(\mathbb{Q};\neq,\mathrm{S_{LL}})$ to the map $(x_1,\dots,x_n) \mapsto \textsc{min}\{x_{\mu_f(1)},\dots,x_{\mu_f(n)}\}$. Clearly ξ preserves arities and projections. We claim that ξ is a clone homomorphism, that is, $\xi(f(g_1,\dots,g_n)) = \xi(f)(\xi(g_1),\dots,\xi(g_n))$ holds for all n-ary f and m-ary $g_1,\dots,g_n$ from $\mathrm{Pol}(\mathbb{Q};\neq,\mathrm{S_{LL}})$. By the injectivity of the essential parts, we have

$$f(g_1,\dots,g_n)^{ess} = f\big(g^{ess}_{\mu_f(1)},\dots,g^{ess}_{\mu_f(n)}\big).$$

Now the claim that ξ is a clone homomorphism follows from the simple fact that

$$\textsc{min}\{x_{1,1},\dots,x_{\ell,k_\ell}\} = \textsc{min}\{\textsc{min}\{x_{i,1},\dots,x_{i,k_i}\} \mid i \in [\ell]\}$$

holds for any $\ell, k_1,\dots,k_\ell \geq 1$. Finally, ξ defined this way is trivially uniformly continuous by choosing $A' := B'$. □

It is worth mentioning that there is no uniformly continuous clone homomorphism $\xi\colon \mathrm{Pol}(\mathbb{Q};\neq,\mathrm{S_{LL}}) \to \mathrm{Pol}(\mathbb{Q};\mathrm{R_{MIN}})$ such that the invertible unary operations of $\xi(\mathrm{Pol}(\mathbb{Q};\neq,\mathrm{S_{LL}}))$ act with finitely many orbits on $\mathbb{Q}$. Otherwise, $(\mathbb{Q};\mathrm{R_{MIN}})$ could be pp-constructed from $(\mathbb{Q};\neq,\mathrm{S_{LL}})$ by Corollary 6.10 in [5] which would, through Corollary 2.6, yield a contradiction to the inexpressibility of $\mathrm{CSP}(\mathbb{Q};\mathrm{R_{MIN}})$ in Datalog.

The main reason behind the failure of most equational conditions coming from finite-domain CSPs in the setting of TCSPs is the existence of TCLs with a tractable CSP whose polymorphisms have very small kernels (sometimes even injective). Instead of identifying output values of a single operation, one needs to identify orbits of tuples obtained by a row-wise application of several different operations to matrix-like schemes of variables. Such conditions were previously considered in the literature, e.g., the $(m+n)$*-terms* introduced in [41]. A condition $\mathcal{E}$ given by a set of identities is called *idempotent* if, for each operation symbol f appearing in the condition, $f(x,\dots,x) \approx x$ is a consequence of $\mathcal{E}$, and *trivial* if $\mathcal{E}$ can be satisfied by projections. By Theorem 1.7 in [41], the existence of $(3+n)$-terms for some n is an non-trivial idempotent condition that characterizes precisely those polymorphism clones over arbitrary sets whose *variety* is *congruence meet-semidistributive*, short $\mathrm{SD}(\wedge)$. All we need to know about $\mathrm{SD}(\wedge)$ is that, in the finite-domain setting, it is in 1-1 correspondence to the expressibility of CSPs in Datalog / FP / FPC, see, e.g., Theorem 1.7 in [42]. This correspondence fails for temporal CSPs, see Proposition 7.4.

Proposition 7.4. *The polymorphism clone of* $(\mathbb{Q};\neq,\mathrm{S_{LL}})$ *does not satisfy any non-trivial idempotent condition and hence its variety is not* $\mathrm{SD}(\wedge)$*, but* $\mathrm{CSP}(\mathbb{Q},\neq,S_{\mathrm{LL}})$ *is expressible in* FP.

Simply dropping idempotence does not bring us any further—Proposition 7.5 shows that, when working with TCSPs, we also lose the correspondence of expressibility of CSPs in FP to satisfiability of an arbitrary equational condition that is unsatisfiable by affine combinations over any field [42].

Proposition 7.5. *The polymorphism clone of* $(\mathbb{Q};\mathrm{X})$ *contains operations witnessing the non-idempotent part of* $(3+3)$*-terms.*

New h1 conditions. The expressibility of temporal and finite-domain CSPs in FP / FPC can still be characterized by a finite non-idempotent h1 condition (Theorem 7.8) which is implied by the 3-4 WNU equation and is closely related to the h1 conditions studied in [3, 24].

Definition 7.6 (Dissected WNUs)**.** Let $\mathfrak{A}$ be a finite structure with a single relation $R^{\mathfrak{A}}$ of arity k. For every $\boldsymbol{t} \in R^{\mathfrak{A}}$ we introduce an k-ary function symbol $g_{\boldsymbol{t}}$. We define the h1 condition $\mathcal{E}_{\mathfrak{A}}$ so that it contains, for each pair $\boldsymbol{t},\boldsymbol{t}' \in R^{\mathfrak{A}}$ with $\boldsymbol{t}[i] = \boldsymbol{t}'[j]$, the equation $g_{\boldsymbol{t}}(x,\dots x,y,x\dots x) \approx g_{\boldsymbol{t}'}(x,\dots x,y,x\dots x)$ where y appears in the i-th argument of $g_{\boldsymbol{t}}(\cdots)$ and in the j-th argument of $g_{\boldsymbol{t}'}(\cdots)$. We write $\mathcal{E}_{k,n}$ instead of $\mathcal{E}_{\mathfrak{A}}$ if the domain A of $\mathfrak{A}$ is $[n]$ and $R^{\mathfrak{A}} = \{\boldsymbol{t} \in [n]^k \mid \boldsymbol{t}[1] < \cdots < \boldsymbol{t}[k]\}$. The condition $\tilde{\mathcal{E}}_{k,n}$ is defined as the extension of $\mathcal{E}_{k,n}$ by the equations $g_{\boldsymbol{t}}(x,\dots x,y,x\dots x) \approx f(x,\dots x,y,x\dots x)$ where y appears in the i-th argument of $g_{\boldsymbol{t}}(\cdots)$ and in the $\boldsymbol{t}[i]$-th argument of $f(\cdots)$.

We first present a general fact about the conditions $\mathcal{E}_{\mathfrak{A}}$ and then we restrict our attention to the families $(\mathcal{E}_{k,n})$ and $(\tilde{\mathcal{E}}_{k,n})$. Recall the structure $(\{0,1\};1\mathrm{IN}k)$ defined in Section 2.

Lemma 7.7. $\mathcal{E}_{\mathfrak{A}}$ *is trivial iff* $\mathfrak{A} \to (\{0,1\};1\mathrm{IN}k)$.

The condition $\tilde{\mathcal{E}}_{k,n}$ first appeared in Lemma 4.3 in [3] for $n = 2k-1$ in a slightly extended form consisting of several intertwined copies of $\tilde{\mathcal{E}}_{k,2k-1}$. Later in [24], the n-ary operation in $\tilde{\mathcal{E}}_{k,n}$ was replaced by a set of binary operations. We propose to consider an even simpler version obtained by dropping the n-ary operation symbol completely while retaining the implied equalities for the remaining operation symbols. This leads to the identities $\mathcal{E}_{k,n}$. The reason is that, although $\tilde{\mathcal{E}}_{k,n}$ and $\mathcal{E}_{k,n}$ are not equivalent as finite h1 conditions as we demonstrate below, the satisfiability of one of these conditions for all but finitely many $n > k > 1$ yields two criteria which are equivalent within the scope of the present paper. Note that $\mathcal{E}_{k,n}$ is implied by the existence of a single k-ary WNU operation, whereas $\tilde{\mathcal{E}}_{k,n}$ is implied by the existence of two WNU operations f and g with arities k and n, respectively, which satisfy $f(y,x,\dots,x) \approx g(y,x,\dots,x)$. Also note that $\mathcal{E}_{k,n}$ implies $\mathcal{E}_{k,k+1}$ for all $n > k > 1$. Using

Lemma 7.7, it is easy to see that $\mathcal{E}_{k,k+1}$ is non-trivial for every $k \geq 2$. We can prove the following result for finite-domain and temporal CSPs, where $\tilde{\mathcal{E}}_{3,4}$ can be replaced with $\tilde{\mathcal{E}}_{k,k+1}$ for any odd $k > 1$.

Theorem 7.8. *Let $\mathfrak{B}$ be a finite structure or a TCL. Then* $\mathrm{CSP}(\mathfrak{B})$ *is expressible in* FP / FPC *iff* $\mathrm{Pol}(\mathfrak{B})$ *satisfies* $\tilde{\mathcal{E}}_{3,4}$.

It follows from Proposition 7.1 that Theorem 7.8 cannot hold for the CSPs of reducts of finitely bounded homogeneous structures in general. However, the polymorphism clone of each of the structures $\mathfrak{U}_h(\mathcal{F})$ used in the proof of Proposition 7.1 can only violate $\tilde{\mathcal{E}}_{k,n}$ for finitely many $n > k > 1$. Also, requiring the simpler condition $\mathcal{E}_{3,4}$ in Theorem 7.8 instead of $\tilde{\mathcal{E}}_{3,4}$ still captures the expressibility of TCSPs in FP. This is not the case for finite-domain CSPs because $\mathcal{E}_{3,4}$ is implied by the existence of a ternary WNU operation. However, Theorem 7.8 can still be reformulated using the simpler conditions if we instead require the satifiability of $\mathcal{E}_{k,k+1}$ for all but finitely many arities, which draws a parallel to Theorem 1.1. For the proof of this more general result, Theorem 7.11, which also covers the structures $\mathfrak{U}_h(\mathcal{F})$, we need the following definition.

Definition 7.9. Let $k \in \mathbb{N}_{\geq 2}$. We denote the k-ary minimum on $\mathbb{Q}$ by $\textsc{min}_k$. The operation $\textsc{mx}_k\colon \mathbb{Q}^k \to \mathbb{Q}$ is defined inductively as follows. In the base case $k = 2$, we set $\textsc{mx}_2(\boldsymbol{t}) \coloneqq \textsc{mx}(\boldsymbol{t})$. For $k > 2$, we set

$$\textsc{mx}_k(\boldsymbol{t}) \coloneqq \textsc{mx}\big(\textsc{mx}_{k-1}(\boldsymbol{t}[1], \ldots, \boldsymbol{t}[k-1]), \textsc{mx}_{k-1}(\boldsymbol{t}[2], \ldots, \boldsymbol{t}[k])\big).$$

The following definitions specify k-ary operations on $\mathbb{Q}$:

$$\begin{aligned}
\textsc{lex}_k(\boldsymbol{t}) &\coloneqq \textsc{lex}\big(\boldsymbol{t}[1], \textsc{lex}\big(\boldsymbol{t}[2], \ldots \textsc{lex}(\boldsymbol{t}[k-1], \boldsymbol{t}[k]) \ldots\big)\big),\\
\textsc{med}_k(\boldsymbol{t}) &\coloneqq \textsc{max}\Big\{\textsc{min}\{\boldsymbol{t}[i] \mid i \in I\} \;\Big|\; I \in \tbinom{[k]}{k-1}\Big\},\\
\textsc{mi}_k(\boldsymbol{t}) &\coloneqq \textsc{lex}_{k+2}\big(\textsc{min}_k(\boldsymbol{t}), \textsc{med}_k(-\chi(\boldsymbol{t})), -\chi(\boldsymbol{t})\big),\\
\textsc{ll}_k(\boldsymbol{t}) &\coloneqq \textsc{lex}_{k+2}\big(\textsc{min}_k(\boldsymbol{t}), \textsc{med}_k(\boldsymbol{t}), \boldsymbol{t}\big).
\end{aligned}$$

Proposition 7.10. *Let $\mathfrak{B}$ be a TCL that is preserved by* $\textsc{min}$ / $\textsc{mi}$ / $\textsc{mx}$ / $\textsc{ll}$. *Then $\mathfrak{B}$ is preserved by* $\textsc{min}_k$ / $\textsc{mi}_k$ / $\textsc{mx}_k$ / $\textsc{ll}_k$ *for all $k \geq 2$.*

Theorem 7.11. *Let $\mathfrak{B}$ be a TCL, a finite structure, or the structure $\mathfrak{U}_h(\mathcal{F})$ for some $\mathcal{F}$. Then* $\mathrm{CSP}(\mathfrak{B})$ *is expressible in* FP / FPC *iff* $\mathrm{Pol}(\mathfrak{B})$ *satisfies $\mathcal{E}_{k,k+1}$ for all but finitely many k.*

Proof. In the context of $\mathcal{E}_{k,k+1}$, we write g_i instead of $g_{\boldsymbol{t}}$ for $\boldsymbol{t}$ that omits i as an entry.

For $\mathfrak{B}$ a TCL, we proceed by a case distinction according to the proof of Theorem 5.3. Suppose that $\mathfrak{B}$ is a TCL that is neither preserved by $\textsc{min}$, $\textsc{mi}$, $\textsc{mx}$, $\textsc{ll}$, the dual of one of these operations, nor a constant operation. Then $\mathfrak{B}$ pp-constructs all finite structures by Theorem 2.10 and in particular the structure $(\{0,1\}; \mathrm{1IN3})$. By Theorem 1.8 in [5], there exists a uniformly continuous h1 clone homomorphism from $\mathrm{Pol}(\mathfrak{B})$ to $\mathrm{Pol}(\{0,1\}; \mathrm{1IN3})$, the projection clone. By Lemma 7.7, $\mathcal{E}_{k,k+1}$ is non-trivial for every k. Thus $\mathrm{Pol}(\mathfrak{B})$ cannot satisfy $\mathcal{E}_{k,k+1}$ due to Theorem 1.9 in [3].

Next, we distinguish the cases where $\mathfrak{B}$ is a TCL preserved by one of the operations listed above. *1)* If $\mathfrak{B}$ is preserved by a constant operation, then $\mathcal{E}_{k,k+1}$ is witnessed by a set of k-ary constant operations for every k. *2)* If $\mathfrak{B}$ is preserved by $\textsc{min}$, then $\mathcal{E}_{k,k+1}$ is witnessed by a set of k-ary minimum operations for every $k \geq 2$. *3)* If $\mathfrak{B}$ has $\textsc{mx}$ as a polymorphism, then, by Theorem 5.2, either $\mathfrak{B}$ is preserved by $\textsc{min}$ or a constant operation, which are cases that we have already treated, or otherwise $\mathfrak{B}$ admits a pp-definition of X. We claim that $\mathrm{Pol}(\mathbb{Q}; \mathrm{X})$ does not satisfy $\mathcal{E}_{k,k+1}$ for any odd $k > 1$. Suppose, on the contrary, that $\mathrm{Pol}(\mathbb{Q}; \mathrm{X})$ satisfies $\mathcal{E}_{k,k+1}$ for some odd $k > 1$. By Theorem 6 in [8], the temporal relation $R^k_{\textsc{mx}} \coloneqq \{\boldsymbol{t} \in \mathbb{Q}^k \mid \sum_{\ell=1}^{k} \chi(\boldsymbol{t})[\ell] = 0 \bmod 2\}$ is preserved by $\textsc{mx}$. By Lemma 4.1, $R^k_{\textsc{mx}}$ is pp-definable in $(\mathbb{Q}; \mathrm{X})$. Thus, by Proposition 2.1, $\mathrm{Pol}(\mathbb{Q}; R^k_{\textsc{mx}})$ satisfies $\mathcal{E}_{k,k+1}$ as well. Consider the inputs $x = 0$ and $y = 1$. Since the columns of the $k \times k$ unit matrix $\boldsymbol{U}_k$ are contained in $R^k_{\textsc{mx}}$, for every $i \in [k+1]$, the application of g_i to the rows of $\boldsymbol{U}_k$ produces a tuple $\boldsymbol{t}_i$ that is contained in $R^k_{\textsc{mx}}$. By the definition of $\mathcal{E}_{k,k+1}$, there exists a tuple $\boldsymbol{t} \in \mathbb{Q}^{k+1}$ such that $\mathrm{pr}_{[k+1]\setminus\{i\}}(\boldsymbol{t}) = \boldsymbol{t}_i$ for every $i \in [k+1]$. By the definition of $R^k_{\textsc{mx}}$, for every $i \in [k+1]$, an even number of entries in $\mathrm{pr}_{[k+1]\setminus\{i\}}(\boldsymbol{t})$ must be minimal. If $\boldsymbol{t}$ exists, then the homogeneous system $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{0}$ of $k+1$ Boolean linear equations of the form $\sum_{j \in [k+1]\setminus\{i\}} x_j = 0 \bmod 2$ has a non-trivial solution because some entry in $\boldsymbol{t}$ must be minimal. But $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{0}$ only has the trivial solution for odd $k > 1$, which completes the proof of our claim. *4)* If $\mathfrak{B}$ has $\textsc{mi}$ as a polymorphism, then we proceed similarly as in the proof of Proposition 4.10 in [3]. For every $k \geq 3$, we set $\tilde{f} \coloneqq \textsc{mi}_{k+1}$ and $\tilde{g}_i \coloneqq \textsc{mi}_k$ for every $i \in [k+1]$. By homogeneity of $(\mathbb{Q}; <)$, for every finite $S \subseteq \mathbb{Q}$, there exists $\alpha \in \mathrm{Aut}(\mathbb{Q}; <)$ such that $\alpha \circ \tilde{g}_1(y, x, \ldots, x) = \tilde{f}(x, y, x, \ldots, x), \ldots, \alpha \circ \tilde{g}_1(x, \ldots, x, y) = \tilde{f}(x, x, \ldots, x, y)$ for all $x, y \in S$. An analogous statement holds for $\tilde{g}_i$ with $i > 1$ where we appropriately shift the first arguments on the right hand side containing only the variable x to the i-th position. Then Lemma 4.4 in [3] yields functions $g_1, \ldots, g_{k+1}$ (and an auxiliary function f) which witness $\mathcal{E}_{k,k+1}$ for $\mathrm{Pol}(\mathfrak{B})$. *5)* If $\mathfrak{B}$ has $\textsc{ll}$ as a polymorphism, then we repeat the strategy above starting with $\tilde{f} \coloneqq \textsc{ll}_{k+1}$ and $\tilde{g}_i \coloneqq \textsc{ll}_k$ for every $i \in [k+1]$. Each of the statements *2-5)* can be dualized in order to obtain witnesses for $\mathcal{E}_{k,k+1}$ for $k \geq 3$ in the cases where $\mathfrak{B}$ is preserved by $\textsc{max}$, $\textsc{dual mi}$, $\textsc{dual ll}$, and show that $\mathrm{Pol}(\mathfrak{B})$ does not satisfy $\mathcal{E}_{k,k+1}$ for odd $k > 1$ if it admits a pp-definition of $-\mathrm{X}$.

To resolve the case where $\mathfrak{B}$ is a finite structure, we show that $\mathcal{E}_{k,k+1}$ can be satisfied by a set of operations given by affine combinations over $\mathbb{Z}_n$ iff $\gcd(k, n) = 1$. Then the claim follows from previously known results [1, 4, 35, 39]. If $\gcd(k, n) = 1$, then there exists $\lambda \in \mathbb{Z}_n$ such that $k\lambda = 1 \bmod n$. Then $\mathcal{E}_{k,k+1}$ is witnessed by a set of k-ary WNU operations $g_1, \ldots, g_{k+1}$ given by the affine combinations $g_j(x_1, \ldots, x_k) \coloneqq \sum_{i=1}^{k} \lambda x_i$. Conversely, suppose that $\mathcal{E}_{k,k+1}$ is witnessed by some operations $g_1, \ldots, g_{k+1}$ given by some affine combinations $g_j(x_1,\dots,x_k) := \sum_{i=1}^{k} \lambda_{j,i} x_i$, that is, $\sum_{i=1}^{k} \lambda_{j,i} = 1 \bmod n$ for every $j \in [k+1]$. If we apply g_j to the rows of the $k \times k$ unit matrix $\boldsymbol{U}_k$ for every $j \in [k+1]$, then $\mathcal{E}_{k,k+1}$ implies that there exists a particular partition $\{F_1,\dots,F_{k+1}\}$ of $[k+1]\times[k]$ such that the value of $\lambda_{j,i}$ is constant when (j,i) ranges over one of the sets F_ℓ, and, for every $j \in [k+1]$, the pairs $(j,1),\dots,(j,k)$ are contained in $\bigcup_{\ell\in[k+1]\setminus\{j\}} F_\ell$. Let λ_ℓ be any fixed representative of the coefficients $\lambda_{j,i}$ for $(j,i) \in F_\ell$. Since g_1 is given by an affine combination, one of the coefficients $\lambda_{1,1},\dots,\lambda_{1,k}$ is non-zero, say $\lambda_{1,1}$. W.l.o.g., $(1,1) \in F_{k+1}$. Then we have $k+1 = \sum_{j=1}^{k+1}\sum_{i=1}^{k} \lambda_{j,i} = \sum_{\ell=1}^{k+1} k\lambda_\ell = k\lambda_{1,1} + k(\sum_{i=1}^{k} \lambda_{k+1,i}) = k\lambda_{1,1} + k \bmod n$. Thus $k\lambda_{1,1} = 1 \bmod n$. But this can only be the case if $\gcd(k,n) = 1$.

If $\mathfrak{B}$ is one of the structures $\mathfrak{U}_h(\mathcal{F})$, then $\mathrm{CSP}(\mathfrak{U}_h(\mathcal{F}))$ is expressible in FP by the proof of Proposition 7.1, and $\mathrm{Pol}(\mathfrak{U}_h(\mathcal{F}))$ satisfies $\mathcal{E}_{k,k+1}$ for all but finitely many arities because it contains WNU operations for all but finitely many arities by Lemma 5.4 in [14]. □

Pseudo h1 conditions. In the context of infinite-domain ω-categorical CSPs, most classification results are formulated using *pseudo h1 conditions* [3] which extend h1 conditions by outer unary operations: $e_1 \circ f_1(x_1^1,\dots,x_{n_1}^1) \approx \cdots \approx e_k \circ f_k(x_k^1,\dots,x_{n_k}^k)$. For instance, the following generalization of a WNU operation was used in [7] to give an alternative classification of the computational complexity of TCSPs. An at least binary operation $f \in \mathrm{Pol}(\mathfrak{B})$ is called *pseudo weak near-unanimity* (PWNU) if there exist $e_1,\dots,e_n \in \mathrm{End}(\mathfrak{B})$ such that $e_1 \circ f(x,\dots,x,y) \approx \cdots \approx e_n \circ f(y,x,\dots,x)$.

Theorem 7.12 ([7]). *Let $\mathfrak{B}$ be a TCL. Then either $\mathfrak{B}$ has a PWNU polymorphism and* $\mathrm{CSP}(\mathfrak{B})$ *is in P, or $\mathfrak{B}$ pp-constructs all finite structures and* $\mathrm{CSP}(\mathfrak{B})$ *is NP-complete.*

Failures of known pseudo h1 conditions. It is natural to ask whether pseudo h1 conditions can be used to formulate a generalization of the 3-4 equation that would capture the expressibility in FP for the CSPs of reducts of finitely bounded homogeneous structures. One such generalization was considered in [13]. Proposition 7.13 shows that the criterion provided by Theorem 8 in [13] is insufficient in general.

Proposition 7.13. *There exist PWNU polymorphisms f, g of* $(\mathbb{Q};\mathrm{X})$ *that satisfy $g(x,x,y) \approx f(x,x,x,y)$.*

Proposition 7.13 has another important consequence, namely Corollary 7.14. The present proof of Corollary 7.14 is by an interesting application from [3] of the fact that the automorphism group of any *ordered homogeneous Ramsey structure* [15] is *extremely amenable* [34]. Consider the structure $\mathfrak{E}_{\mathcal{G},3}$ defined in Section 4 for every finite Abelian group $\mathcal{G}$.

Corollary 7.14. *The structure* $(\mathbb{Q};\mathrm{X})$ *does not pp-construct $\mathfrak{E}_{\mathcal{G},3}$ for any finite non-trivial Abelian group $\mathcal{G}$.*

Proof. Suppose, on the contrary, that this is the case. By Theorem 1.8 in [5], there exists a uniformly continuous h1 clone homomorphism $\xi\colon \mathrm{Pol}(\mathbb{Q};\mathrm{X}) \to \mathrm{Pol}(\mathfrak{E}_{\mathcal{G},3})$ that preserves all h1 conditions which hold in $\mathrm{Pol}(\mathbb{Q};\mathrm{X})$. Because $(\mathbb{Q};<)$ is an ordered homogeneous Ramsey structure and $\mathfrak{E}_{\mathcal{G},3}$ is finite, by the second proof of Theorem 1.9 in [3], there exists a uniformly continuous h1 clone homomorphism $\xi'\colon \mathrm{Pol}(\mathbb{Q};\mathrm{X}) \to \mathrm{Pol}(\mathfrak{E}_{\mathcal{G},3})$ that preserves all pseudo h1 conditions with outer embeddings which hold in $\mathrm{Pol}(\mathbb{Q};\mathrm{X})$ for at most 4-ary operations. Since every endomorphism of $(\mathbb{Q};\mathrm{X})$ is an embedding, the 3-4 equation for PWNUs from Proposition 7.13 is such a condition. Thus it must also be satisfied in $\mathrm{Pol}(\mathfrak{E}_{\mathcal{G},3})$. But the only endomorphism of $\mathfrak{E}_{\mathcal{G},3}$ is the identity, which means that $\mathrm{Pol}(\mathfrak{E}_{\mathcal{G},3})$ would have to satisfy the 3-4 equation for WNUs, which cannot be the case by Theorem 1.1 combined with Theorem 10 in [1]. □

Another possible criterion which also turns out to be insufficient is the existence of PWNU polymorphism for all but finitely many arities, see Proposition 7.15

Proposition 7.15. *For every $k \geq 3$,* $\textsc{min}_k$, $\textsc{mx}_k$, $\textsc{mi}_k$, *and* $\textsc{ll}_k$ *are PWNU operations.*

New pseudo h1 conditions. We present a new candidate for an algebraic condition given by pseudo h1 identities that could capture the expressibility in FP for the CSPs of reducts of finitely bounded homogeneous structures. Let $\mathcal{E}'_{k,k+1}$ be the pseudo h1 condition obtained from $\mathcal{E}_{k,k+1}$ by replacing each g_t in $\mathcal{E}_{k,k+1}$ with $e_t \circ g$ where e_t is unary and g is k-ary. For instance, $\mathcal{E}'_{3,4}$ is the following condition:

$$\begin{aligned}
e_2 \circ g(y,x,x) &\approx e_3 \circ g(y,x,x) \approx e_4 \circ g(y,x,x),\\
e_1 \circ g(y,x,x) &\approx e_3 \circ g(x,y,x) \approx e_4 \circ g(x,y,x),\\
e_1 \circ g(x,y,x) &\approx e_2 \circ g(x,y,x) \approx e_4 \circ g(x,x,y),\\
e_1 \circ g(x,x,y) &\approx e_2 \circ g(x,x,y) \approx e_3 \circ g(x,x,y).
\end{aligned}$$

Clearly $\mathcal{E}'_{k,k+1}$ is implied by $\mathcal{E}_{k,k+1}$, and also by the existence of a k-ary WNU operation. It follows from the proof of Theorem 7.11 that the statement of the theorem remains true if we replace $\mathcal{E}_{k,k+1}$ with $\mathcal{E}'_{k,k+1}$. Note that, although $\mathcal{E}'_{k,k+1}$ contains only a single k-ary operation symbol for $k \geq 2$, it is in general not implied by the existence of a k-ary PWNU operation by Theorem 7.11 combined with Proposition 7.15.

8 Open questions

Question 1. Can $\mathrm{CSP}(\mathfrak{A})$ for $\mathfrak{A}$ a reduct of a finitely bounded homogeneous structure be expressed in FP iff $\mathrm{Pol}(\mathfrak{A})$ satisfies $\mathcal{E}_{k,k+1}$ or $\mathcal{E}'_{k,k+1}$ for all but finitely many $k > 1$?

Question 2. Does the satisfiability of $\mathcal{E}_{k,k+1}$ for all but finitely many $k > 1$ imply the satisfiability of $\mathcal{E}_{k,n}$ for all but finitely many $n > k > 1$ in polymorphism clones of reducts of finitely bounded homogeneous structures?

Question 3. Is $\mathrm{CSP}(\mathbb{Q};\mathrm{X})$ expressible in *choiceless polynomial time* (CPT) with (or without) counting [6]?

Question 4. Do FPR or CPT+C capture PTime for CSPs of reducts of finitely bounded homogeneous structures?

Question 5. If a CSP is in FPC, is it is also in FP?

References

[1] A. Atserias, A. Bulatov, and A. Dawar. 2009. Affine systems of equations and counting infinitary logic. *Theoretical Computer Science* 410, 18 (2009), 1666–1683.

[2] A. Atserias and A. Dawar. 2019. Definable inapproximability: New challenges for duplicator. *Journal of Logic and Computation* 29, 8 (2019), 1185–1210. Available at arXiv:1806.11307.

[3] L. Barto, M. Kompatscher, M. Olšák, T. Van Pham, and M. Pinsker. 2019. Equations in oligomorphic clones and the Constraint Satisfaction Problem for ω-categorical structures. *Journal of Mathematical Logic* 19, 2 (2019), 1950010. available at arXiv:1612.07551.

[4] L. Barto and M. Kozik. 2009. Constraint Satisfaction Problems of Bounded Width. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, Atlanta, GA, USA, 595–603.

[5] L. Barto, J. Opršal, and M. Pinsker. 2018. The wonderland of reflections. *Israel Journal of Mathematics* 223, 1 (2018), 363–398.

[6] A. Blass, Y. Gurevich, and S. Shelah. 2002. On polynomial time computation over unordered structures. *The Journal of Symbolic Logic* 67, 3 (2002), 1093–1125.

[7] M. Bodirsky. 2012. Complexity Classification in Infinite-Domain Constraint Satisfaction. (2012). Mémoire d'habilitation à diriger des recherches, Université Diderot – Paris 7. Available at arXiv:1201.0856.

[8] M. Bodirsky, H. Chen, and M. Wrona. 2014. Tractability of quantified temporal constraints to the max. *International Journal of Algebra and Computation* 24, 08 (2014), 1141–1156.

[9] M. Bodirsky and V. Dalmau. 2013. Datalog and constraint satisfaction with infinite templates. *J. Comput. System Sci.* 79, 1 (2013), 79–100.

[10] M. Bodirsky and P. Jonsson. 2017. A model-theoretic view on qualitative constraint reasoning. *Journal of Artificial Intelligence Research* 58 (2017), 339–385.

[11] M. Bodirsky and J. Kára. 2010. The complexity of temporal constraint satisfaction problems. *Journal of the ACM (JACM)* 57, 2 (2010), 9.

[12] M. Bodirsky and J. Kára. 2010. A fast algorithm and Datalog Inexpressibility for temporal reasoning. *ACM Transactions on Computational Logic* 11, 3 (2010), 1–21.

[13] M. Bodirsky and A. Mottet. 2016. Reducts of finitely bounded homogeneous structures, and lifting tractability from finite-domain constraint satisfaction. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science*. ACM, IEEE, New York, NY, USA, 623–632.

[14] M. Bodirsky, A. Mottet, M. Olšák, J. Opršal, M. Pinsker, and R. Willard. 2019. Topology is relevant (in a dichotomy conjecture for infinite-domain constraint satisfaction problems). In *2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. IEEE, Vancouver, BC, Canada, 1–12. Available at arXiv:1901.04237.

[15] M. Bodirsky and M. Pinsker. 2011. Reducts of Ramsey structures. *AMS Contemporary Mathematics* 558 (2011), 489–519.

[16] J.-Y. Cai, M. Fürer, and N. Immerman. 1992. An optimal lower bound on the number of variables for graph identification. *Combinatorica* 12, 4 (1992), 389–410.

[17] A. K. Chandra and P. M. Merlin. 1977. Optimal implementation of conjunctive queries in relational databases. In *Proceedings of the ninth annual ACM symposium on Theory of computing*. ACM, New York, NY, United States, 77–90.

[18] G. Cherlin, S. Shelah, and N. Shi. 1999. Universal graphs with forbidden subgraphs and algebraic closure. *Advances in Applied Mathematics* 22, 4 (1999), 454–491.

[19] A. Dawar. 2015. The nature and power of fixed-point logic with counting. *ACM SIGLOG News* 2, 1 (2015), 8–21.

[20] A. Dawar, M. Grohe, B. Holm, and B. Laubner. 2009. Logics with rank operators. In *2009 24th Annual IEEE Symposium on Logic In Computer Science*. IEEE, Los Angeles, CA, USA, 113–122.

[21] A. Dawar and S. Kreutzer. 2008. On Datalog vs. LFP. In *International Colloquium on Automata, Languages and Programming*. Springer, Reykjavik, Iceland, 160–171.

[22] A. Dawar and P. Wang. 2017. Definability of semidefinite programming and Lasserre lower bounds for CSPs. In *2017 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. IEEE, Reykjavik, Iceland, 1–12.

[23] T. Feder and M. Y. Vardi. 2003. Homomorphism closed vs. existential positive. In *18th Annual IEEE Symposium of Logic in Computer Science, 2003. Proceedings.* IEEE, Ottawa, Ont., Canada, 311–320.

[24] P. Gillibert, J. Jonušas, M. Kompatscher, A. Mottet, and M. Pinsker. 2020. Hrushovski's encoding and ω-categorical CSP monsters. In *International Colloquium on Automata, Languages and Programming*. ACM, Saarbrücken, Germany, to appear. Available at arXiv:2002.07054.

[25] A. Goerdt. 2009. On random ordering constraints. In *International Computer Science Symposium in Russia*. Springer, Novosibirsk, Russia, 105–116.

[26] E. Grädel, P. G. Kolaitis, L. Libkin, M. Marx, J. Spencer, M. Y. Vardi, Y. Venema, and S. Weinstein. 2007. *Finite Model Theory and its applications*. Springer, Berlin, Heidelberg, New York.

[27] E. Grädel and M. Otto. 1993. Inductive Definability with Counting on Finite Structures. In *Computer Science Logic, 6th Workshop, CSL '92, San Miniato 1992, Selected Papers*, Vol. 702. Springer, San Miniato, Italy, 231–247.

[28] E. Grädel and W. Pakusa. 2019. Rank logic is dead, long live rank logic! *The Journal of Symbolic Logic* 84, 1 (2019), 54–87.

[29] M. Grohe. 2008. The quest for a logic capturing PTIME. In *2008 23rd Annual IEEE Symposium on Logic in Computer Science*. IEEE, Pittsburgh, PA, USA, 267–271.

[30] Y. Gurevich and S. Shelah. 1996. On finite rigid structures. *The Journal of Symbolic Logic* 61, 2 (1996), 549–562.

[31] V. Guruswami, J. Håstad, R. Manokaran, P. Raghavendra, and M. Charikar. 2011. Beating the random ordering is hard: Every ordering CSP is approximation resistant. *SIAM J. Comput.* 40, 3 (2011), 878–914.

[32] W. Hodges. 1997. *A shorter model theory*. Cambridge university press, Cambridge, United Kingdom.

[33] J. Hubička and J. Nešetřil. 2016. Homomorphism and Embedding Universal Structures for Restricted Classes. *Journal of Multiple-Valued Logic & Soft Computing* 27 (2016), 229–253. Available at arXiv:0909.4939.

[34] A. S. Kechris, V. G. Pestov, and S. Todorcevic. 2005. Fraïssé limits, Ramsey theory, and topological dynamics of automorphism groups. *Geometric and Functional Analysis* 15, 1 (2005), 106–189.

[35] M. Kozik, A. Krokhin, M. Valeriote, and R. Willard. 2015. Characterizations of several Maltsev conditions. *Algebra universalis* 73, 3-4 (2015), 205–224.

[36] S. Kreutzer. 2004. Expressive equivalence of least and inflationary fixed-point logic. *Annals of Pure and Applied Logic* 130, 1-3 (2004), 61–78.

[37] B. Larose and L. Zádori. 2007. Bounded width problems and algebras. *Algebra universalis* 56, 3-4 (2007), 439–466.

[38] L. Libkin. 2013. *Elements of finite model theory*. Springer, Berlin, Heidelberg, New York.

[39] M. Maróti and R. McKenzie. 2008. Existence theorems for weakly symmetric operations. *Algebra universalis* 59, 3 (2008), 463–489.

[40] B. Nebel and H.-J. Bürckert. 1995. Reasoning about temporal relations: a maximal tractable subclass of Allen's interval algebra. *Journal of the ACM (JACM)* 42, 1 (1995), 43–66.

[41] M. Olšák. 2018. Maltsev conditions for general congruence meet-semidistributive algebras. (2018). Available at arXiv:1810.03178.

[42] M. Olšák. 2019. *Generalizing CSP-related results to infinite algebras*. PhD dissertation. Univerzita Karlova, Matematicko-fyzikální fakulta.

[43] M. Otto. 2000. *Bounded Variable Logics and Counting. A Study in Finite Model Theory*. Springer, Berlin, Heidelberg, New York.

[44] W. Pakusa. 2015. *Linear Equation Systems and the Search for a Logical Characterisation of Polynomial Time*. Ph.D. Dissertation. PhD thesis, RWTH Aachen.

[45] P. Wang. 2018. *Descriptive complexity of constraint problems.* Ph.D. Dissertation. University of Cambridge.