Modular Counting CSP: Reductions and Algorithms

Amirhossein Kazeminia
Simon Fraser University, Burnaby, Canada

Andrei A. Bulatov
Simon Fraser University, Barnaby, Canada

Abstract

The Constraint Satisfaction Problem (CSP) is ubiquitous in various areas of mathematics and computer science. Many of its variations have been studied including the Counting CSP, where the goal is to find the number of solutions to a CSP instance. The complexity of finding the exact number of solutions of a CSP is well understood (Bulatov, 2013, and Dyer and Richerby, 2013) and the focus has shifted to other variations of the Counting CSP such as counting the number of solutions modulo an integer. This problem has attracted considerable attention recently. In the case of CSPs based on undirected graphs Bulatov and Kazeminia (STOC 2022) obtained a complexity classification for the problem of counting solutions modulo p for arbitrary prime p. In this paper we report on the progress made towards a similar classification for the general CSP, not necessarily based on graphs.

We identify several features that make the general case very different from the graph case such as a stronger form of rigidity and the structure of automorphisms of powers of relational structures. We provide a solution algorithm in the case $p = 2$ that works under some additional conditions and prove the hardness of the problem under some assumptions about automorphisms of the powers of the relational structure. We also reduce the general CSP to the case that only uses binary relations satisfying strong additional conditions.

2012 ACM Subject Classification Theory of computation → Problems, reductions and completeness; Theory of computation → Constraint and logic programming

Keywords and phrases Constraint Satisfaction Problem, Modular Counting

Digital Object Identifier 10.4230/LIPIcs.STACS.2025.60

Related Version *Full Version*: `https://arxiv.org/abs/2501.04224` [29]

Funding *Andrei A. Bulatov*: NSERC Discovery Grant.

1 Introduction

Counting problems in general have been intensively studied since the pioneering work by Valiant [34, 33]. One of the most interesting and well studied problems in this area is the Counting Constraint Satisfaction Problem (#CSP), which provides a generic framework for a wide variety of counting combinatorial problems that arise frequently in multiple disciplines such as logic, graph theory, and artificial intelligence. The counting CSP also allows for generalizations including weighted CSPs and partition functions [2, 7] that yield connections with areas such as statistical physics, see, e.g. [27, 32]. While the complexity of exact counting solutions of a CSP is now well-understood [13, 3, 14, 12], modular counting such as finding the parity of the number of solutions remains widely open.

Homomorphisms and the Constraint Satisfaction Problem. The most convenient way to introduce CSPs is through homomorphisms of relational structures. A *relational signature* σ is a collection of *relational symbols* each of which is assigned a positive integer, the *arity* of the symbol. A *relational structure* $\mathcal{H}$ with signature σ is a set H and an *interpretation* $\mathcal{R}^{\mathcal{H}}$

of each $\mathcal{R} \in \sigma$, where $\mathcal{R}^{\mathcal{H}}$ is a relation or a predicate on H whose arity equals that of $\mathcal{R}$. The set H is said to be the *base set* or the *universe* of $\mathcal{H}$. We will use the same letter for the base set as for the structure, only in the regular font. A structure with signature σ is often called a *σ-structure*. Structures with the same signature are called *similar*.

Let $\mathcal{G}, \mathcal{H}$ be similar structures with signature σ. A *homomorphism* from $\mathcal{G}$ to $\mathcal{H}$ is a mapping $\varphi : G \to H$ such that for any $\mathcal{R} \in \sigma$, say, of arity r, if $\mathcal{R}^{\mathcal{G}}(a_1, \ldots, a_r)$ is true for $a_1, \ldots, a_r \in G$, then $\mathcal{R}^{\mathcal{H}}(\varphi(a_1), \ldots, \varphi(a_r))$ is also true. The set of all homomorphisms from $\mathcal{G}$ to $\mathcal{H}$ is denoted $\mathsf{Hom}(\mathcal{G}, \mathcal{H})$. The cardinality of $\mathsf{Hom}(\mathcal{G}, \mathcal{H})$ is denoted by $\mathsf{hom}(\mathcal{G}, \mathcal{H})$. A homomorphism φ is an *isomorphism* if it is bijective and the inverse mapping φ^{-1} is a homomorphism from $\mathcal{H}$ to $\mathcal{G}$. A homomorphism of a structure to itself is called an *endomorphism*, and an isomorphism to itself is called an *automorphism*.

Following Feder and Vardi [17], in a CSP, the goal is, given similar relational structures $\mathcal{G}, \mathcal{H}$, to decide whether there is a homomorphism from $\mathcal{G}$ to $\mathcal{H}$. The restricted problem in which $\mathcal{H}$ is fixed and only $\mathcal{G}$ is given as an input is denoted by $\mathsf{CSP}(\mathcal{H})$. The complexity of such problems is well understood [4, 35].

Counting CSP. In the (exact) Counting CSP the goal is to find the number $\mathsf{hom}(\mathcal{G}, \mathcal{H})$ of homomorphisms from a structure $\mathcal{G}$ to a similar structure $\mathcal{H}$. Restricted versions of the Counting CSP can be introduced in the same way as for the decision one. In the counting version of $\mathsf{CSP}(\mathcal{H})$ denoted $\#\mathsf{CSP}(\mathcal{H})$ the goal is to find $\mathsf{hom}(\mathcal{G}, \mathcal{H})$ for a given structure $\mathcal{G}$as an input.

The complexity class the Counting CSP belongs to is #P, the class of problems of counting accepting paths of polynomial time nondeterministic Turing machines. There are several ways to define reductions between counting problems, but the most widely used ones are parsimonious reductions and Turing reductions. A *parsimonious reduction* from a counting problem $\#A$ to a counting problem $\#B$ is an algorithm that, given an instance I of $\#A$ produces (in polynomial time) an instance I' of $\#B$ such that the answers to I and I' are equal. A *Turing reduction* is a polynomial time algorithm that solves $\#A$ using $\#B$ as an oracle. Completeness in #P is then defined in the standard way. This paper and all the papers we cite predominantly use Turing reductions.

A complexity classification of counting CSPs of the form $\#\mathsf{CSP}(\mathcal{H})$ was obtained by Bulatov [3] and was further improved and simplified by Dyer and Richerby [14]. Bulatov's proof makes heavy use of techniques of universal algebra. Dyer and Richerby's proof, on the other hand, uses combinatorial and structural properties of relational structures. The more general version of the counting CSP, the weighted CSP, has also been thoroughly studied. Cai and Chen [10] obtained a complexity classification for weighted CSP, where each homomorphism has a complex weight. One of the main features of counting with complex weights is the phenomenon of cancellation, when complex weights of homomorphisms cancel each other rather than add up. This, of course, never happens in exact unweighted counting problems, but is frequently encountered in modular counting.

Modular Counting. Another natural variation of counting problems is counting modulo some integer. In this paper we consider the problem of computing the number of solutions of a CSP modulo a prime number p. If a relational structure $\mathcal{H}$ is fixed, this problem is denoted $\#_p\mathsf{CSP}(\mathcal{H})$. More precisely, in $\#_p\mathsf{CSP}(\mathcal{H})$ the objective is, given a relational structure $\mathcal{G}$, to find the number of homomorphisms from $\mathcal{G}$ to $\mathcal{H}$ modulo p.

There are several complexity classes related to modular counting. The more established type of classes is $\mathsf{Mod}_k\mathrm{P}$, the class of problems of deciding whether the number of accepting paths of a polynomial time nondeterministic Turing machine is *not* divisible by k, [11, 25].

In particular, if $k = 2$ then $\mathsf{Mod}_k\mathrm{P}$ is the class $\oplus\mathrm{P}$. However, problems of counting accepting paths naturally belong to classes of the form $\#_k\mathrm{P}$, introduced by Faben in [15] that contain problems of counting accepting paths modulo k. The standard notion of reduction is again Turing reduction. Faben in [15] studied the basic properties of such classes, in particular, he identified several $\#_k\mathrm{P}$-complete problems.

In the case of the CSP, the research has mostly been focused on graph homomorphisms. The only exceptions we are aware of are a result of Faben [15], who characterized the complexity of counting the solutions of a Generalized Satisfiability problem modulo an integer, and a generalization of [15] to problems with weights by Guo et al. [21]. The study of modular counting of graph homomorphisms has been much more vibrant.

Before discussing the existing results on modular counting and the results of this study, we need to mention some features of the automorphism group of a relational structure. The automorphisms of a relational structure $\mathcal{H}$ form a group with respect to composition denoted $\mathsf{Aut}(\mathcal{H})$. The order of an automorphism $\pi \in \mathsf{Aut}(\mathcal{H})$ is the smallest number k such that π^k is the identity permutation. An element $a \in H$ is a fixed point of $\pi \in \mathsf{Aut}(\mathcal{H})$ if $\pi(a) = a$. The set of all fixed points of π is denoted by $\mathsf{Fix}(\pi)$.

A systematic study of counting homomorphisms in graphs was initiated by Faben and Jerrum in [15]. They observed that the automorphism group $\mathsf{Aut}(\mathcal{H})$, particularly the automorphisms of order p, plays a crucial role in the complexity of the problem $\#_p\mathsf{Hom}(\mathcal{H})$. This insight extends to relational structures, as discussed in [9]. Specifically, for a homomorphism φ from a relational structure $\mathcal{G}$ to $\mathcal{H}$, composing φ with an automorphism from $\mathsf{Aut}(\mathcal{H})$ yields another homomorphism from $\mathcal{G}$ to $\mathcal{H}$. Thus, any automorphism of $\mathcal{H}$ acts on the set $\mathsf{Hom}(\mathcal{G}, \mathcal{H})$ of all homomorphisms from $\mathcal{G}$ to $\mathcal{H}$. If $\mathsf{Aut}(\mathcal{H})$ contains an automorphism π of order p, the size of the orbit of φ is divisible by p unless $\pi \circ \varphi = \varphi$, making this orbit contribute 0 modulo p to the total homomorphism count from $\mathcal{G}$ to $\mathcal{H}$. If $\pi \circ \varphi = \varphi$, the range of φ lies within the set of fixed points $\mathsf{Fix}(\pi)$ of π. This observation motivates the following construction: let $\mathcal{H}^\pi$ denote the substructure of $\mathcal{H}$ induced by $\mathsf{Fix}(\pi)$. We write $\mathcal{H} \to_p \mathcal{H}'$ there exists $\pi \in \mathsf{Aut}(\mathcal{H})$ such that $\mathcal{H}'$ is isomorphic to $\mathcal{H}^\pi$. Furthermore, we write $\mathcal{H} \to_p^* \mathcal{H}'$ if there exist structures $\mathcal{H}_1, \dots, \mathcal{H}_k$ such that $\mathcal{H}$ is isomorphic to $\mathcal{H}_1$, $\mathcal{H}'$ is isomorphic to $\mathcal{H}_k$, and $\mathcal{H}_1 \to_p \mathcal{H}_2 \to_p \cdots \to_p \mathcal{H}_k$.

Relational structures without order p automorphisms will be called *p-rigid*.

▶ **Lemma 1** ([9, 16]). *Let $\mathcal{H}$ be a relational structure and p a prime. Then*
(a) *Up to an isomorphism there exists a unique p-rigid structure $\mathcal{H}^{*p}$ such that $\mathcal{H} \to_p^* \mathcal{H}^{*p}$.*
(b) *For any relational structure $\mathcal{G}$ it holds that $\mathsf{hom}(\mathcal{G}, \mathcal{H}) \equiv \mathsf{hom}(\mathcal{G}, \mathcal{H}^{*p}) \pmod p$.*

By Lemma 1 it suffices to determine the complexity of $\#_p\mathsf{CSP}(\mathcal{H})$ for p-rigid structures $\mathcal{H}$.

The existing results on modular counting. As we mentioned before, the research in modular counting CSPs has mostly been aimed at counting graph homomorphisms. The complexity of the problem $\#_p\mathsf{Hom}(H)$ of counting homomorphism to a fixed graph H modulo a prime number p has received significant attention in the last ten years. Faben and Jerrum in [16] posed a conjecture that up to order p automorphism reduction $\to_p$ the complexity of this problem is the same as that for exact counting. More precisely, they conjectured that $\#_p\mathsf{Hom}(H)$ is solvable in polynomial time if and only if $\mathsf{Hom}(H^{*p})$ is. By the result of Dyer and Greenhill [13] $\#_p\mathsf{Hom}(H)$ is solvable in polynomial time if and only if every connected component of H^{*p} is a complete graph with all loops present or a complete bipartite graph. Therefore, proving that if a p-rigid H does not satisfy these conditions then $\#_p\mathsf{Hom}(H)$ is $\#_pP$-hard suffices to confirm the conjecture of Faben and Jerrum. Over several years the hardness of $\#_p\mathsf{Hom}(H)$ was established for various graph classes [16, 22, 19, 20, 28, 18, 31]. Finally, it was proved for arbitrary graphs in [9].

▶ **Theorem 2** ([9]). *For any prime p and any graph H the problem $\#_p\mathsf{Hom}(H)$ is solvable in polynomial time if and only if $\mathsf{Hom}(H^{*p})$ is solvable in polynomial time. Otherwise it is $\#_pP$-complete.*

Our Results. In this paper we begin a systematic study of the problem $\#_p\mathsf{CSP}(\mathcal{H})$ for general relational structures $\mathcal{H}$. Note that to the best of our knowledge, it is the first paper attempting at such a general modular counting problem. The ultimate goal is to obtain a complexity classification similar to Theorem 2 for arbitrary relational structures. The full version of the paper can be found in [29].

The contribution of the paper is twofold. First, we analyse the existing techniques such as those from [9], and the methods used in exact counting [3, 14, 10], and their applicability to the general case. We conclude that few of them work. More specifically, Theorem 2 asserts that the complexity of modular counting for p-rigid graphs is the same as of exact counting. We, however, suggest a relational structure, a digraph $\mathcal{T}_p$, that is p-rigid, its exact counting problem is hard, but modular counting is easy, see Example 18. Another important ingredient of the proof of Theorem 2 is a structural theorem on automorphisms of products of graphs [24]. No such result exists for products of relational structures. Moreover, in Example 18 in Section 6 we suggest an example (again, a digraph) showing that nothing similar to such a structural result can be true. Some of the standard techniques in counting CSPs involve properties of relations and relational structures such as rectangularity, permutability, balancedness, the presence of a Mal'tsev polymorphism. In the case of exact counting these concepts are closely related to each other and make efficient algorithms possible. Unfortunately, these concepts are of little help for modular counting. We introduce their modular equivalents, but then a series of examples show that no tight connections are possible in this case. This makes algorithm design very difficult.

On the positive side, we obtain some results to somewhat remedy the situation. The first step is to convert the problem into a richer framework of multi-sorted relational structures and CSPs. The main idea is, given a CSP instance $\mathcal{G}$, to try to identify the possible images of each vertex of $\mathcal{G}$, and then treat vertices with different ranges as having different types and map them to different disjoint domains. In Section 4 we call this process refinement and propose two types of refinement, one is based on local propagation techniques, and the other on solving the decision version of the problem. The main benefit of using multi-sorted structures and CSPs is the richer structure of their automorphisms. This often allows stronger reductions than single-sorted structures do. In particular, if the digraph $\mathcal{T}_p$ mentioned above is subjected to this process, it results in a multi-sorted structure that is no longer p-rigid, the corresponding reduced structure is very simple and can easily be solved. We are not aware of any examples of a structure whose multi-sorted refinement is p-rigid, but that would give rise to an easy modular counting problem.

In the next line of research we follow the approach of [9] to expand the relational structure $\mathcal{H}$ by adding relations to $\mathcal{H}$ that are primitive positive (pp-)definable in $\mathcal{H}$, that is, relations that can be derived from the relations of $\mathcal{H}$ using equality, conjunction and existential quantifiers. However, expanding the general relational structure by pp-definable relations does not work as well as for graphs. To overcome this obstacle, we introduce a new form of expansion which uses p-modular quantifiers instead of regular existential quantifiers. The semantics of a p-modular quantifier is "there are non-zero modulo p values of a variable" rather than "there exists at least one value of a variable" as the regular existential quantifier asserts. Every relational structure is associated with a *relational clone* $\langle\mathcal{H}\rangle$ that consists of all relations pp-definable in $\mathcal{H}$. The new concept gives rise to new definitions of pp-formulas and

relational clones. If regular existential quantifiers in pp-formulas are replaced with p-modular quantifiers, we obtain p-modular primitive positive formulas (p-mpp-formulas, for short). Then, similar to pp-definitions, a relation $\mathcal{R}$ is said to be p-mpp-definable in a structure $\mathcal{H}$ if there is a p-mpp-formula in $\mathcal{H}$ expressing $\mathcal{R}$. The p-modular clone $\langle\mathcal{H}\rangle_p$ associated with $\mathcal{H}$ is the set of all relations p-mpp-definable in $\mathcal{H}$. We show in Theorem 10 (see also Theorem 7) that, similar to the result of Bulatov and Dalmau [6], expanding a structure by a p-mpp-definable relation does not change the complexity of the problem $\#_p\mathsf{CSP}(\mathcal{H})$.

▶ **Theorem 3.** *Let p be a prime number and $\mathcal{H}$ a p-rigid relational structure. If $\mathcal{R}$ is p-mpp-definable in $\mathcal{H}$, then $\#_p\mathsf{CSP}(\mathcal{H}+\mathcal{R})$ is polynomial time reducible to $\#_p\mathsf{CSP}(\mathcal{H})$.*

In the remaining part of the paper we identify a number of conditions under which it is possible to design an algorithm or to prove the hardness of the problem. One such case is $\#_2\mathsf{CSP}(\mathcal{H})$ when $\mathcal{H}$ satisfies both 2-rectangularity and the usual strong rectangularity conditions (or, equivalently, has a Mal'tsev polymorphism). It starts with applying the known techniques [5, 14] to find a concise representation or a frame of the set of solutions of a given CSP. However, such a representation cannot be used directly to find the parity of the number of solutions. The algorithm performs multiple recomputations of the frame to exclude the parts that produce an even number of solutions. Unfortunately, this algorithm does not generalize to larger p even under very strong assumptions, because the structure of finite fields start playing a role.

While studying the structure of automorphisms of products of relational structures such as $\mathsf{Aut}(\mathcal{H}^n)$ may be a difficult problem, in Section 7 we make a step forward by reducing the class of structures $\mathcal{H}$ for which such structural results are required. More precisely, for any relational structure $\mathcal{H}=(H;\mathcal{R}_1,\dots,\mathcal{R}_k)$ we construct its binarization $b(\mathcal{H})$ whose relations are binary and rectangular. This makes such structures somewhat closer to graphs and the hope is that it will be easier to study the structure of $\mathsf{Aut}(b(\mathcal{H})^n)$ than $\mathsf{Aut}(\mathcal{H}^n)$ itself. We prove that $\mathcal{H}$ and $b(\mathcal{H})$ share many important properties.

▶ **Theorem 4.** *Let $\mathcal{H}$ be a relational structure. Then $\mathcal{H}$ is strongly rectangular (p-strongly rectangular, p-rigid, has a Mal'tsev polymorphism) if and only if $b(\mathcal{H})$ is strongly rectangular (p-strongly rectangular, p-rigid, has a Mal'tsev polymorphism).*

2 Preliminaries

Let $[n]$ denote the set $\{1,2,\dots,n\}$. Let H^n be the Cartesian product of the set H with itself n times and $H_1\times\cdots\times H_n$ the Cartesian product of sets $H_1,\dots,H_n$. We denote the members of H^n and $H_1\times\cdots\times H_n$ using bold font, $\mathbf{a}\in H^n$, $\mathbf{a}\in H_1\times\cdots\times H_n$. The i-th element of $\mathbf{a}$ is denoted by $\mathbf{a}[i]$ or $\mathbf{a}_i$.

For $I=\{i_1,\dots,i_k\}\subseteq[n]$ we use $\mathsf{pr}_I\mathbf{a}$ to denote $(\mathbf{a}_{i_1},\dots,\mathbf{a}_{i_k})$; and for $\mathcal{R}\subseteq H_1\times\cdots\times H_n$ we use $\mathsf{pr}_I\mathcal{R}$ to denote $\{\mathsf{pr}_I\mathbf{a}\mid\mathbf{a}\in\mathcal{R}\}$. For $\mathbf{a}\in H^s$ by $\#\mathsf{ext}_{\mathcal{R}}(\mathbf{a})$ we denote the number of assignments $\mathbf{b}\in H^{n-s}$ such that $(\mathbf{a},\mathbf{b})\in\mathcal{R}$. (Note that the order of elements in $\mathbf{a}$ and $\mathbf{b}$ and $\mathcal{R}$ might differ, hence we slightly abuse the notation here.) We denote the number of assignments mod p by $\#_p\mathsf{ext}_{\mathcal{R}}(\mathbf{a})$. Moreover, $\mathsf{pr}^p_I\mathcal{R}$ denotes the set $\{\mathsf{pr}_I\mathbf{a}|\mathbf{a}\in\mathcal{R}\text{ and }\#_p\mathsf{ext}_{\mathcal{R}}(\mathsf{pr}_I\mathbf{a})\neq 0\}$. Often, we treat relations $\mathcal{R}\subseteq H_1\times\cdots\times H_n$ as predicates $\mathcal{R}:H_1\times\cdots\times H_n\to\{0,1\}$.

2.1 Multi-Sorted Sets and Relational Structures

We begin with introducing *multi-sorted* or *typed* sets. Let $H=\{H_i\}_{i\in[k]}=\{H_1,\dots,H_k\}$ be a collection of sets. We will assume that the sets $H_1,\dots,H_k$ are disjoint.

The cardinality of a multi-sorted set H equals $|H| = \sum_{i\in[k]} |H_i|$. A mapping φ between two multi-sorted sets $G = \{G_i\}_{i\in[k]}$ and $H = \{H_i\}_{i\in[k]}$ is defined as a collection of mappings $\varphi = \{\varphi_i\}_{i\in[k]}$, where $\varphi_i : G_i \to H_i$, that is, φ_i maps elements of the sort i in G to elements of the sort i in H. A mapping $\varphi = \{\varphi_i\}_{i\in[k]}$ from $\{G_i\}_{i\in[k]}$ to $\{H_i\}_{i\in[k]}$ is injective (bijective), if for all $i \in [k]$, φ_i is injective (bijective).

A *multi-sorted relational signature* σ over a set of types $\{1,\dots,k\}$ is a collection of *relational symbols*, a symbol $\mathcal{R} \in \sigma$ is assigned a positive integer $\ell_{\mathcal{R}}$, the *arity* of the symbol, and a tuple $(i_1,\dots,i_{\ell_{\mathcal{R}}}) \subseteq [k]^{\ell_{\mathcal{R}}}$, the *type* of the symbol. A *multi-sorted relational structure* $\mathcal{H}$ with signature σ is a multi-sorted set $\{H_i\}_{i\in[k]}$ and an *interpretation* $\mathcal{R}^{\mathcal{H}}$ of each $\mathcal{R} \in \sigma$, where $\mathcal{R}^{\mathcal{H}}$ is a relation or a predicate on $H_{i_1} \times ... \times H_{i_{\ell_{\mathcal{R}}}}$. The multi-sorted structure $\mathcal{H}$ is finite if H and σ are finite. All structures in this paper are finite. The set H is said to be the *base set* or the *universe* of $\mathcal{H}$. For the base set we will use the same letter as for the structure, only in the regular font. Multi-sorted structures with the same signature and type are called *similar*.

Let $\mathcal{G},\mathcal{H}$ be similar multi-sorted structures with signature σ. A *homomorphism* φ from $\mathcal{G}$ to $\mathcal{H}$ is a collection of mappings $\varphi_i : G_i \to H_i$, $i \in [k]$, from G to H such that for any $\mathcal{R} \in \sigma$ with type $(i_1,\dots,i_{\ell_{\mathcal{R}}})$, if $\mathcal{R}^{\mathcal{G}}(a_1,\dots,a_{\ell_{\mathcal{R}}})$ is true for $(a_1,\dots,a_r) \in G_{i_1} \times ... \times G_{i_{\ell_{\mathcal{R}}}}$, then $\mathcal{R}^{\mathcal{H}}(\varphi_{i_1}(a_1),\dots,\varphi_{i_{\ell_{\mathcal{R}}}}(a_{\ell_{\mathcal{R}}}))$ is true as well. The set of all homomorphisms from $\mathcal{G}$ to $\mathcal{H}$ is denoted $\mathsf{Hom}(\mathcal{G},\mathcal{H})$. The cardinality of $\mathsf{Hom}(\mathcal{G},\mathcal{H})$ is denoted by $\mathsf{hom}(\mathcal{G},\mathcal{H})$. For a multi-sorted structure $\mathcal{H}$, the corresponding *counting CSP*, $\#_p\mathsf{CSP}(\mathcal{H})$, is the problem of computing $\mathsf{hom}(\mathcal{G},\mathcal{H})$ modulo a prime number p for a given structure $\mathcal{G}$. A homomorphism φ is an *isomorphism* if it is bijective and the inverse mapping φ^{-1} is a homomorphism from $\mathcal{H}$ to $\mathcal{G}$. If $\mathcal{H}$ and $\mathcal{G}$ are isomorphic, we write $\mathcal{H} \cong \mathcal{G}$. A homomorphism of a structure to itself is called an *endomorphism*, and an isomorphism to itself is called an *automorphism*. As is easily seen, automorphisms of a structure $\mathcal{H}$ form a group denoted $\mathsf{Aut}(\mathcal{H})$.

The *direct product* of multi-sorted σ-structures $\mathcal{H},\mathcal{G}$, denoted $\mathcal{H} \times \mathcal{G}$ is the multi-sorted σ-structure with the base set $H \times G = \{H_i \times G_i\}_{i\in[k]}$, the interpretation of $\mathcal{R} \in \sigma$ is given by $\mathcal{R}^{\mathcal{H}\times\mathcal{G}}((a_1,b_1),\dots,(a_k,b_k)) = 1$ if and only if $\mathcal{R}^{\mathcal{H}}(a_1,\dots,a_k) = 1$ and $\mathcal{R}^{\mathcal{G}}(b_1,\dots,b_k) = 1$. By $\mathcal{H}^\ell$ we will denote the *ℓth power* of $\mathcal{H}$, that is, the direct product of ℓ copies of $\mathcal{H}$.

For a prime number p we say that π has order p if π is not the identity in $\mathsf{Aut}(\mathcal{H})$ and has order p in $\mathsf{Aut}(\mathcal{H})$. In other words, each of the π_j's is either the identity mapping or has order p, and at least one of the π_j's is not the identity mapping. Structure $\mathcal{H}$ is said to be *p-rigid* if it has no automorphism of order p. Similar to regular relational structures we can introduce reductions of multi-sorted structures by their automorphisms of order p.

A *substructure* $\mathcal{H}'$ of $\mathcal{H}$ *induced* by a collection of sets $\{H'_i\}_{i\in[k]}$, where $H'_i \subseteq H_i$ is the relational structure given by $(\{H'_i\}_{i\in[k]};\mathcal{R}'_1,\dots,\mathcal{R}'_m)$, where $\mathcal{R}'_j = \mathcal{R}_j \cap (H'_{i_1} \times \cdots \times H'_{i_\ell})$ and $(i_1,\dots,i_\ell)$ is the type of $\mathcal{R}_j$. By $\mathsf{Fix}(\pi)$ we denote the collection $\{\mathsf{Fix}(\pi_i)\}_{j\in[k]}$ of sets of fixed points of the π_i's. Let $\mathcal{H}^\pi$ denote the substructure of $\mathcal{H}$ induced by $\mathsf{Fix}(\pi)$. We write $\mathcal{H} \to_p \mathcal{H}'$ if there is $\pi \in \mathsf{Aut}(\mathcal{H})$ of order p such that $\mathcal{H}'$ is isomorphic to $\mathcal{H}^\pi$. We also write $\mathcal{H} \to_p^* \mathcal{H}'$ if there are structures $\mathcal{H}_1,\dots,\mathcal{H}_t$ such that $\mathcal{H}$ is isomorphic to $\mathcal{H}_1$, $\mathcal{H}'$ is isomorphic to $\mathcal{H}_t$, and $\mathcal{H}_1 \to_p \mathcal{H}_2 \to_p \cdots \to_p \mathcal{H}_t$. Let also $\mathcal{H}^{*p}$ be a p-rigid structure such that $\mathcal{H} \to_p^* \mathcal{H}^{*p}$.

▶ **Proposition 5.** *Let $\mathcal{H}$ be a multi-sorted structure and p a prime. Then up to an isomorphism there exists a unique p-rigid multi-sorted structure $\mathcal{H}^{*p}$ such that $\mathcal{H} \to_p^* \mathcal{H}^{*p}$. Moreover, for any relational structure $\mathcal{G}$ it holds $\mathsf{hom}(\mathcal{G},\mathcal{H}) \equiv \mathsf{hom}(\mathcal{G},\mathcal{H}^{*p}) \pmod p$.*

The proof of Proposition 5 follows the same lines as the analogous statement in the single-sorted case.

We complete this section with a definition of polymorphisms. Let $\mathcal{R} \subseteq H^n$ be a relation over a set H. A k-ary *polymorphism* of $\mathcal{R}$ is a mapping $f : H^k \to H$ such that for any choice of $\mathbf{a}_1, \dots, \mathbf{a}_k \in \mathcal{R}$, it holds that $f(\mathbf{a}_1, \dots, \mathbf{a}_k) \in \mathcal{R}$ (computed component-wise). The mapping f is a polymorphism of a (single-sorted) relational structure $\mathcal{H} = (H, \mathcal{R}_1, \dots, \mathcal{R}_m)$ if it is a polymorphism of each relation $\mathcal{R}_1, \dots, \mathcal{R}_m$. In the multi-sorted case the definitions are a bit more complicated. Let $\mathcal{R} \subseteq H_1 \times \cdots \times H_n$ be a multi-sorted relation. Instead of a single mapping we consider a family of mappings $f = \{f_i\}_{i\in[n]}$, $f_i : H_i^k \to H_i$. The family f is said to be a polymorphism of $\mathcal{R}$ if it satisfies the same condition: for any choice of $\mathbf{a}_1, \dots, \mathbf{a}_k \in \mathcal{R}$, it holds that $f(\mathbf{a}_1, \dots, \mathbf{a}_k) \in \mathcal{R}$, only in this case the mapping f_i is applied in the ith coordinate position, $i \in [n]$. A polymorphism of a multi-sorted structure $\mathcal{H} = (\{H_i\}_{i\in[q]}; \mathcal{R}_1, \dots, \mathcal{R}_m)$ is again a family $f = \{f_i\}_{i\in q} : H_i^k \to H_i$ such that for each $j \in [m]$, f (or rather its appropriate subfamily) is a polymorphism of $\mathcal{R}_j$. For a complete introduction into polymorphisms the reader is referred to [1].

The following special type of polymorphisms often occurs in the study of counting CSPs. For a set H a mapping $\varphi : H^3 \to H$ is said to be a *Mal'tsev* operation if for any $a, b \in H$ it satisfies the conditions $f(a, a, b) = f(b, a, a) = b$. In the multi-sorted case a family $f = \{f_i\}_{i\in[n]}$ is Mal'tsev, if every f_i is Mal'tsev. A Mal'tsev operation (or a family of operations) that is a polymorphism of a relation $\mathcal{R}$ or a relational structure $\mathcal{H}$ is said to be a Mal'tsev polymorphism of $\mathcal{R}$ or $\mathcal{H}$.

2.2 Expansion of relational structures

One of the standard techniques when studying constraint problems is to identify ways to expand the target relational structure or constraint language with additional relations without changing the complexity of the problem.

Let $\mathcal{H}$ be a relational structure with signature σ and $\mathcal{H}^=$ its expansion by adding a binary relational symbol $=$ interpreted as $=_H$, the equality relation on H. The following reduction is straightforward.

▶ **Lemma 6** ([9]). *For any relational structure $\mathcal{H}$ and any prime p, $\#_p\mathsf{CSP}(\mathcal{H}^=) \leq_T \#_p\mathsf{CSP}(\mathcal{H})$.*

A constant relation over a set H is a unary relation $C_a = \{a\}$, $a \in H$. For a relational structure $\mathcal{H}$ by $\mathcal{H}^\mathsf{c}$ we denote the expansion of $\mathcal{H}$ by all the constant relations C_a, $a \in H$. Theorem 7 was proved for exact counting in [6], for modular counting of graph homomorphisms in [16, 19, 20], and for general modular counting CSP in [9].

▶ **Theorem 7** ([9]). *Let $\mathcal{H}$ be a p-rigid σ-structure. Then $\#_p\mathsf{CSP}(\mathcal{H}^\mathsf{c})$ is polynomial time reducible to $\#_p\mathsf{CSP}(\mathcal{H})$.*

Lemma 6 and Theorem 7 can be generalized to the multi-sorted case. Let $\mathcal{H} = \{\mathcal{H}_i\}_{i\in[k]}$ be a multi-sorted structure with signature σ and $\mathcal{H}^=$ its expansion by adding a family of binary relational symbols $=_{H_i}$ (one for each type) interpreted as the equality relation on H_i, $i \in [k]$.

A constant relation over a set $\{H_i\}_{i\in[k]}$ is a unary relation $C_a = \{a\}$, $a \in H_i, i \in [k]$ (such a predicate can only be applied to variables of type i). For a structure $\mathcal{H}$ by $\mathcal{H}^\mathsf{c}$ we denote the expansion of $\mathcal{H}$ by all the constant relations C_a, $a \in H_i, i \in [k]$.

▶ **Theorem 8.** *Let $\mathcal{H}$ be a multi-sorted relational structure and p prime.*
(1) *$\#_p\mathsf{CSP}(\mathcal{H}^=)$ is Turing reducible to $\#_p\mathsf{CSP}(\mathcal{H})$;*
(2) *Let $\mathcal{H}$ be p-rigid. Then $\#_p\mathsf{CSP}(\mathcal{H}^\mathsf{c})$ is Turing reducible to $\#_p\mathsf{CSP}(\mathcal{H})$.*

Yet another way to expand a relational structure is by primitive positive definable (pp-definable for short) relations. Primitive-positive definitions have played a major role in the study of the CSP. It has been proved in multiple circumstances that expanding a structure with pp-definable relations does not change the complexity of the corresponding CSP. This has been proved for the decision CSP in [26, 8] and the exact Counting CSP [6]. The reader is referred to [1] for details about pp-definitions and their use in the study of the CSP.

Conjunctive definitions are a special case of pp-definitions that do not use quantifiers. Let $\mathcal{H}$ be a structure with signature σ. A *conjunctive formula* Φ over variables $\{x_1, \dots, x_k\}$ is a conjunction of atomic formulas of the form $\mathcal{R}(y_1, \dots, y_\ell)$, where $\mathcal{R} \in \sigma$ is an (ℓ-ary) symbol and $y_1, \dots, y_\ell \in \{x_1, \dots, x_k\}$. A k-ary predicate $\mathcal{Q}$ is conjunctive definable by Φ if $(a_1, \dots, a_k) \in \mathcal{Q}$ if and only if $\Phi(a_1, \dots, a_k)$ is true.

▶ **Lemma 9** ([9]). *Let $\mathcal{H}$ be a relational structure with signature σ, $\mathcal{R}$ be conjunctive definable in $\mathcal{H}$, and $\mathcal{H} + \mathcal{R}$ denote the expansion of $\mathcal{H}$ by a predicate symbol $\mathcal{R}$ that is interpreted as the relation $\mathcal{R}$ in $\mathcal{H}$. Then $\#_p\mathsf{CSP}(\mathcal{H} + \mathcal{R}) \leq_T \#_p\mathsf{CSP}(\mathcal{H})$.*

We now extend the concept of primitive-positive definability to the multi-sorted case. Let $\mathcal{H}$ be a multi-sorted relational structure with the base set H. As before *primitive positive* (pp-) formula in $\mathcal{H}$ is a first-order formula $\exists y_1, \dots, y_s \Phi(x_1, \dots, x_k, y_1, \dots, y_s)$, where Φ is a conjunction of atomic formulas of the form $z_1 =_H z_2$ or $\mathcal{R}(z_1, \dots, z_\ell)$, $z_1, \dots, z_\ell \in \{x_1, \dots, x_k, y_1, \dots, y_s\}$, and $\mathcal{R}$ is a predicate of $\mathcal{H}$. However, every variable in Φ is now assigned a type $\tau(x_i), \tau(y_j)$ in such a way that for every atomic formula $z_1 =_H z_2$ it holds that $\tau(z_1) = \tau(z_2)$, and for any atomic formula $\mathcal{R}(z_1, \dots, z_\ell)$ the sequence $(\tau(z_1), \dots, \tau(z_\ell))$ matches the type of $\mathcal{R}$. We say that $\mathcal{H}$ *pp-defines* a predicate $\mathcal{Q}$ if there exists a pp-formula such that

$$\mathcal{Q}(x_1, \dots, x_k) = \exists y_1, \dots, y_s \Phi(x_1, \dots, x_k, y_1, \dots, y_s).$$

For $\mathbf{a} \in \mathcal{R}$ by $\#\mathsf{ext}_\Phi(\mathbf{a})$ we denote the number of assignments $\mathbf{b} \in H_{\tau(y_1)} \times \dots \times H_{\tau(y_s)}$ to $y_1, \dots, y_s$ such that $\Phi(\mathbf{a}, \mathbf{b})$ is true. We denote the number of such assignments mod p by $\#_\mathsf{p}\mathsf{ext}_\Phi(\mathbf{a})$.

While when $\mathcal{H}$ is a graph it is possible to prove a statement similar to Lemma 9 for pp-definable relations [9], we will later see that it is unlikely to be true for general relational structures.

Finally, for a relational structure $\mathcal{H}$ (single- or multi-sorted) $\langle\mathcal{H}\rangle$ denotes the relational clone of $\mathcal{H}$, that is, the set of all relations pp-definable in $\mathcal{H}$

2.3 Modular Expansion of Relational Structures

We follow the approach of [9] by expanding the relational structure $\mathcal{H}$ by adding pp-definable relations, but doing in a manner friendly to modular counting.

We introduce a new form of expansion which is using p-modular quantifiers instead of regular existential quantifiers. The semantics of a p-modular quantifier is "there are non-zero modulo p values of a variable" rather than "there exists at least one value of a variable" as the regular existential quantifier asserts. The new concept gives rise to new definitions of pp-formulas. If regular existential quantifiers in pp-formulas are replaced with p-modular quantifiers, we obtain p-modular primitive positive formulas (p-mpp-formulas, for short). The p-modular quantifier is denoted $\exists^{\equiv p}$, and so p-mpp-formulas have the form

$$\exists^{\equiv p} y_1, \dots, y_{\ell_1} \dots \exists^{\equiv p} y_{\ell_1 + \dots + \ell_{s-1}+1}, \dots, y_k \Phi(z_1, \dots, z_m).$$

Note the more complicated form of the quantifier prefix: modular quantification is not as robust as the regular one and quantifying variables away in groups or one-by-one may change the result. For example, let $\mathcal{R} = \{(1,0,0),(1,1,0),(1,1,1),(2,2,2)\}$ be a relation on $\{0,1,2\}$. Then formulas $\exists^{\equiv 3} y \exists^{\equiv 3} z \;\; \mathcal{R}(x,y,z)$ and $\exists^{\equiv 3} y, z \;\; \mathcal{R}(x,y,z)$ define sets $\{1,2\}$ and $\{2\}$, respectively.

Every relational structure is associated with a relational clone $\langle \mathcal{H} \rangle$ that consists of all relations pp-definable in $\mathcal{H}$. Then, similar to pp-definitions, a relation $\mathcal{R}$ is said to be p-mpp-definable in a structure $\mathcal{H}$ if there is a p-mpp-formula in $\mathcal{H}$ expressing $\mathcal{R}$. The p-modular clone $\langle \mathcal{H} \rangle_p$ associated with $\mathcal{H}$ is the set of all relations p-mpp-definable in $\mathcal{H}$. Similar to the result of Bulatov and Dalmau [6], expanding a structure by a p-mpp-definable relation does not change the complexity of the problem $\#_p\mathsf{CSP}(\mathcal{H})$.

▶ **Theorem 10.** *Let $\mathcal{H}$ be a be a σ-structure, and p a prime. Let $\mathcal{R}$ be a relation that is defined by $R(x_1,\dots,x_k) = \exists^{\equiv p} y_1,\dots,y_\ell \Phi(x_1,\dots,x_k,y)$, then $\#_p\mathsf{CSP}(\mathcal{H}+\mathcal{R})$ is polynomial time reducible to $\#_p\mathsf{CSP}(\mathcal{H})$.*

3 Structural Properties for Counting

3.1 Rectangularity, Permutability, and Friends

The key properties of relational structures heavily exploited in [3, 14, 10] to obtain characterizations of the complexity of exact counting are rectangularity, balancedness, congruence permutability, and the presence of a Mal'tsev polymorphism. Indeed, according to [14] $\#\mathsf{CSP}(\mathcal{H})$ is polynomial time solvable if and only if $\mathcal{H}$ is strongly balanced. However, in order for the solution algorithm to work, it requires a Mal'tsev polymorphism to be applied over and over again to construct a *frame*, that is, a compact representation of the set of solutions, [3, 14].

Using modular pp-definitions, we can modify these properties' definitions accordingly to obtain the properties of strong p-rectangularity, p-balancedness, and p-permutability. Modular-pp-definitions preserve the complexity of modular problems, however, they destroy the connections between the concepts above.

p-Rectangularity. Recall that a binary relation $\mathcal{R} \subseteq H_1 \times H_2$ is said to be *rectangular* if $(a,c),(a,d),(b,c) \in \mathcal{R}$ implies $(b,d) \in \mathcal{R}$ for any $a,b \in H_1, c,d \in H_2$. A relation $\mathcal{R} \subseteq H_1 \times \cdots \times H_n$ for $n \geq 2$ is rectangular if for every $I \subsetneq [n]$, the relation R is rectangular when viewed as a binary relation, a subset of $\mathsf{pr}_I\mathcal{R} \times \mathsf{pr}_{[n]-I}\mathcal{R}$. A relational structure $\mathcal{H}$ is *strongly rectangular* if every relation $\mathcal{R} \in \langle \mathcal{H} \rangle$ of arity at least 2 is rectangular. Finally, a relational structure $\mathcal{H}$ is said to be strongly *p-rectangular*, if every $\mathcal{R} \in \langle \mathcal{H} \rangle_p$ is rectangular.

p-Balancedness. An n-by-m matrix M is said to be a *rank-1 block matrix* if by permuting rows and columns it can be transformed to a block-diagonal matrix (not necessarily square), where every nonzero diagonal block has rank at most 1. Note that the rank can also be computed in $\mathbb{Z}_p$, in which case we use the term *rank-1 block matrix modulo p*.

Let $\mathcal{R}(x,y,z)$ be a ternary relation, a subset of $H_1 \times H_2 \times H_3$. We call $\mathcal{R}$ *balanced* if the matrix $\mathsf{M}_\mathcal{R} \in \mathbb{Z}^{|H_1|\times|H_2|}$, where $\mathsf{M}_\mathcal{R}[x,y] = |\{z \in H_3 : (x,y,z) \in \mathcal{R}\}|$ such that $x \in H_1$ and $y \in H_2$ is a rank-1 block matrix. It is *p-balanced* if $\mathsf{M}_\mathcal{R}$ is a rank-1 block matrix modulo p. A relation $\mathcal{R}$ of arity $n > 3$ is balanced if every representation of $\mathcal{R}$ as a ternary relation, a subset of $H^k \times H^\ell \times H^{n-k-\ell}$, is balanced. Similarly, A relation $\mathcal{R}$ on H of arity $n > 3$ is p-balanced if every representation of $\mathcal{R}$ as a ternary relation, a subset of $H^k \times H^\ell \times H^{n-k-\ell}$, is p-balanced.

A relational structure $\mathcal{H}$ is called *strongly balanced* if every relation $\mathcal{R} \in \langle\mathcal{H}\rangle$ is balanced. Similarly, a relational structure $\mathcal{H}$ is called *strongly p-balanced* if every relation $\mathcal{R} \in \langle\mathcal{H}\rangle_p$ is p-balanced.

p-Permutability. A *congruence* of a relational structure $\mathcal{H}$ is an equivalence relation θ on H that is pp-definable in $\mathcal{H}$. More generally, let $\mathcal{R} \in \langle\mathcal{H}\rangle$ be a k-ary relation. A congruence of $\mathcal{R}$ is a $2k$-ary relation pp-definable in $\mathcal{H}$ that is an equivalence relation on $\mathcal{R}$. We denote the set of all congruences of $\mathcal{H}$ (of $\mathcal{R}$) by $\mathsf{Con}(\mathcal{H})$ (respectively, by $\mathsf{Con}(\mathcal{R})$). By $\circ$ we denote the product of binary relations: $(a, b) \in \mathcal{R} \circ \mathcal{Q}$ if and only if there is c such that $(a, c) \in \mathcal{R}$ and $(c, b) \in \mathcal{Q}$. We say $\mathcal{H}$ is *congruence permutable* if for all $\alpha, \beta \in \mathsf{Con}(\mathcal{R})$, where $\mathcal{R} \in \langle\mathcal{H}\rangle$, it holds that $\alpha \circ \beta = \beta \circ \alpha$.

Congruence p-permutability is defined as follows: A *p-congruence* of $\mathcal{H}$ or of $\mathcal{R} \in \langle\mathcal{H}\rangle_p$ is an equivalence relation on H or $\mathcal{R}$, respectively, that is p-mpp-definable in $\mathcal{H}$. We denote the set of all p-congruences of $\mathcal{H}$ ($\mathcal{R}$) by $\mathsf{Con}_p(\mathcal{H})$ ($\mathsf{Con}_p(\mathcal{R})$). By $\textcircled{p}$ we denote the product of binary relations: $(a, b) \in \mathcal{R}\textcircled{p}\mathcal{Q}$ if and only if the number of elements c such that $(a, c) \in \mathcal{R}$ and $(c, b) \in \mathcal{Q}$ is not a multiple of p. We say that $\mathcal{H}$ is *congruence p-permutable* if for all $\alpha, \beta \in \mathsf{Con}_p(\mathcal{R})$, where $\mathcal{R} \in \langle\mathcal{H}\rangle_p$, we have $\alpha\textcircled{p}\beta = \beta\textcircled{p}\alpha$. Figure 1 demonstrates the connection between congruence permutabililty, strong rectangularity, the existence of a Mal'tsev polymorphism, strong balancedness and their modular counterparts. A collection of statements and examples proving these connections or lack thereof will be given in the full version of the paper. Below we give one such example showing that the existence of a Mal'tsev polymorphism does not guarantee 2-rectangularity.

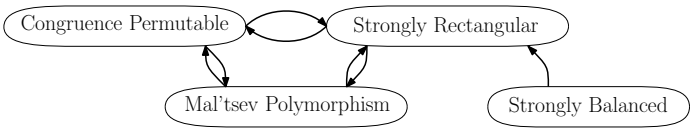

(a) Congruence Permutability, Strong Rectangularity, and Mal'tsev polymorphism are equavalent (See [23, 14, 3]). Also, Strong Balancedness implies Strong Rectangularity(See [14]).

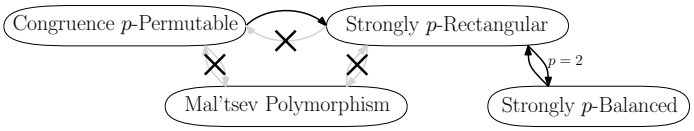

(b) The only connection that is preserved for the modular case is Congruence p-Permutability implies Strong p-rectangularity. Also, Strong 2-rectangularity is equivalence to Strong 2-Balancedness.

Figure 1 The connection between 4 properties is shown above. Figure (1a) shows the connection for the exact counting. Figure (1b) shows the the connection for the modular counterparts.

▶ **Example 11.** Let $H = \{0, 1, 2\}$, $p = 2$, and $\mathcal{H} = (H; \mathcal{R})$, where $\mathcal{R}$ is the following ternary relation, (triples are written vertically), $\mathcal{R} = \{(0,0,0), (0,1,1), (1,0,1), (1,0,2), (1,1,0)\}$. As is easily see, $\exists^{\equiv 2} z\mathcal{R}(x, y, z)$ is the relation $\{(0,0), (0,1), (1,1)\}$, which is not rectangular, and so, $\mathcal{H}$ is not strongly 2-rectangular. We now show that $\mathcal{H}$ is strongly rectangular by presenting a Mal'tsev polymorphism of $\mathcal{H}$. Let $f(x, y, z) = x + y + z$, where addition is modulo 2, is an operation on $\{0, 1\}$. For $\mathbf{a} = (a_1, a_2, a_3) \in H^3$, let also $\mathbf{a}' \in \{0, 1\}^3$ denote the triple obtained by replacing 2's with 1's. Then let $g(x, y, z)$ be given by

$$g(\mathbf{a}) = \begin{cases} 2, & \text{if } \mathbf{a} \in \{(2, a, a), (a, a, 2) \mid a \in H\}, \\ f(\mathbf{a}') & \text{otherwise.} \end{cases}$$

It is straightforward that g is a Mal'tsev operation and is a polymorphism of $\mathcal{H}$.

4 Rigitity and Multi-sorted Structures

We start with applying a well-known framework of multi-sorted relational structures to modular counting. While multi-sorted structures is a standard tool in the study of decision CSPs, it usually only used to simplify arguments and streamline algorithms. Here we use this framework in a more fundamental way, to strengthen the main (hypothetical) tractability condition.

The classification result in Theorem 2 asserts that $\#_p\mathsf{Hom}(H)$ for a p-rigid graph H is hard whenever the exact counting problem is hard. In the following example, we briefly show that this is no longer the case for general relational structures.

▶ **Example 12.** Let p be a prime number and $T_p = \{0, \dots, p-1, p, p+1\}$. A relational structure $\mathcal{T}_p$ has the base set T_p and predicates $\mathcal{R}, C_0, \dots, C_{p+1}$, where C_a is the constant relation $\{(a)\}$ and $\mathcal{R} = T_p^2 - \{(0,p), \dots, (p-1,p)\}$. The structure $\mathcal{T}_p$ is p-rigid as it contains all the constant relations and every automorphism must preserve them, implying that every element of T_p is a fixed point. By [4, 35] the decision $\mathsf{CSP}(\mathcal{T}_p)$ is solvable in polynomial time, while the exact counting problem $\#\mathsf{CSP}(\mathcal{T}_p)$ is #P-complete by [6], as it does not have a Mal'tsev polymorphism and $\mathcal{R}$ is not rectangular (see below). However, if $\mathcal{G}$ is a structure similar to $\mathcal{T}_p$ and there is a homomorphism $\varphi : \mathcal{G} \to \mathcal{T}_p$ such that some vertex v of $\mathcal{G}$ is mapped to $a \in \{0, \dots, p-1\}$ then, unless v is bound by C_a, the mapping that differs from φ only at v by sending it to any $b \in \{0, \dots, p-1\}$ is also a homomorphism. Since $|\{0, \dots, p-1\}| = p$, this means that the elements $0, \dots, p-1$ can be effectively eliminated from $\mathcal{T}_p$, and the resulting structure is somewhat trivial. Therefore $\#_p\mathsf{CSP}(\mathcal{T}_p)$ can be solved in polynomial time.

We introduce a different concept of rigidity that is stronger than the one used before. In particular, it will explain the tractability of the problem from Example 12.

While Proposition 5 allows one to reduce CSPs over non-p-rigid structures, it is sometimes possible to go further and reduce a CSP to one with a richer structure, and a richer set of automorphisms. Let $\mathcal{H} = (\{H_i\}_{i\in[k]}; \mathcal{R}_1, \dots, \mathcal{R}_m)$ be a multi-sorted relational structure. We say that a structure $\mathcal{G}$ is a *refinement* of $\mathcal{H}$ if it satisfies the following conditions:

(a) $\mathcal{G} = (\{G_i\}_{i\in[q]}; \mathcal{Q}_1, \dots, \mathcal{Q}_t)$, where the G_i' are pairwise disjoint,

(b) for every $i \in [q]$, there is an injective mapping $\xi_i : G_i \to H_{i'}$ for some $i' \in [k]$,

(c) for every $j \in [t]$ there is $j' \in [m]$ with $\mathcal{R}_{j'} \subseteq H_{i_1} \times \dots \times H_{i_\ell}$ and $\mathcal{Q}_j = \{(a_1, \dots, a_\ell) \in G_{i'_1} \times \dots \times G_{i'_\ell} \mid (\xi_{i'_1}(a_1), \dots, \xi_{i'_\ell}(a_\ell)) \in \mathcal{R}_{j'}\}$, where $G_{i'_r}$ is such that $\xi_{i'_r}(G_{i'_r}) \subseteq H_{i_r} \cap \mathsf{pr}_r\mathcal{R}_{j'}$.

Condition (a) is required because the domains of a multi-sorted structure have to be disjoint. Condition (b) basically says that G_i consists of copies of some elements of $H_{i'}$. Condition (c) amounts to saying that $\mathcal{Q}_j \subseteq \mathcal{R}_{j'}$, except it uses copies of the elements of $\mathcal{H}$. In the notation of item (c) we use $\xi(a_1, \dots, a_\ell)$ to denote $(\xi_{i'_1}(a_1), \dots, \xi_{i'_\ell}(a_\ell))$, we use $\xi(\mathcal{Q}_j)$ to denote $\{\xi(a_1, \dots, a_\ell) \mid (a_1, \dots, a_\ell) \in \mathcal{Q}_j\}$, and $\xi^{-1}(\mathcal{R}_{j'})$ for $\{(a_1, \dots, a_\ell) \in G_{i'_1} \times \dots \times G_{i'_\ell} \mid \xi(a_1, \dots, a_\ell) \in \mathcal{R}_{j'}\}$.

It is possible that while $\mathcal{H}$ is p-rigid, its refinement is not and Proposition 5.

In order to relate refinement structures with reductions between counting problems we introduce two special types of refinement. First of all, we will need an alternative approach to pp-definitions based on homomorphisms, see [17, 30].

▶ **Lemma 13** ([17, 30]). *A predicate $\mathcal{R}(x_1, \dots, x_k)$ is pp-definable in a multi-sorted structure $\mathcal{H}$ containing the equality predicate if and only if there exists a similar structure $\mathcal{G}_\mathcal{R}$ containing vertices $x_1, \dots, x_k$ such that for any $(a_1, \dots, a_k) \in \mathcal{R}$ it holds that $(a_1, \dots, a_k) \in \mathcal{R}$ if and only if there is a homomorphism from $\mathcal{G}_\mathcal{R}$ to $\mathcal{H}$ that maps x_i to a_i, $i \in [k]$. We will say that $\mathcal{G}_\mathcal{R}$* defines *R in $\mathcal{H}$.*

Let $\mathcal{H} = (\{H_i\}_{i\in[k]}; \mathcal{R}_1, \dots, \mathcal{R}_m)$ be a multi-sorted relational structure. The *Gaifman graph* of $\mathcal{H}$ is the graph $G(\mathcal{H}) = (V, E)$, where $V = \bigcup_{i\in[k]} H_i$ and $(a, b) \in E$ if and only if there is $j \in [m]$ and $\mathbf{a} \in \mathcal{R}_j$ such that $\mathbf{a}[s] = a, \mathbf{a}[t] = b$ for some coordinate positions s, t of $\mathcal{R}_j$. The structure $\mathcal{H}$ has *treewidth* d if $G(\mathcal{H})$ has treewidth d.

We say that a refinement $\mathcal{G} = (\{G_i\}_{i\in[q]}; \mathcal{Q}_1, \dots, \mathcal{Q}_t)$ of $\mathcal{H}$ is *width d definable* if for every $i \in [q]$ there is a structure $\mathcal{G}_i$ of treewidth at most d that defines $\xi_i(G_i)$ in $\mathcal{H}$. In a similar way, we say that $\mathcal{G}$ is a *definable* refinement if for every $i \in [q]$ the set $\xi_i(G_i)$ is pp-definable in $\mathcal{H}$. Finally, we say that $\mathcal{G}$ is *the full width d definable refinement* (respectively, *full definable refinement*) if it satisfies the following conditions.

(1) It is a width d definable refinement (respectively, a definable refinement).

(2) For every unary relation U definable in $\mathcal{H}$ by a structure of treewidth at most d (respectively, pp-definable unary relation) there is $i \in [q]$ such that $\xi_i(G_i) = U$.

(3) For every relation S obtained from some relation $\mathcal{R}_j$ by restricting it to domains definable by structures of width d (respectively by pp-definable domains), there is $\mathcal{Q}_{j'}$ such that $\xi(\mathcal{Q}_{j'}) = S$.

Since the original domains H_i, $i \in [k]$, have trivial pp-definitions, they (or rather their copies) are always among the G_j's, and copies of the original relations are also among the $\mathcal{Q}_j$'s, although they may be over different, smaller domains than the $\mathcal{R}_i$'s.

Next, we extend refinements to CSP instances. Let $\mathcal{G} = (\{G_i\}_{i\in[q]}; \mathcal{Q}_1, \dots, \mathcal{Q}_t)$ be a refinement of $\mathcal{H}$ and $\mathcal{P} = (V, \mathcal{C})$ an instance of $\mathsf{CSP}(\mathcal{H})$. Recall that every variable $v \in V$ is assigned a sort $\sigma(v) \in [k]$. Let $\sigma' : V \to [q]$ be such that $\xi_{\sigma'(v)} : G_{\sigma'(v)} \to H_{\sigma(v)}$ for each $v \in V$. The instance $\mathcal{P}^{\sigma'} = (V, \mathcal{C}^{\sigma'})$ is said to be a *refinement for $\mathcal{G}$* of $\mathcal{P}$ with the sort function σ' if it satisfies the following two conditions

(a) every $v \in V$ is assigned the sort $\sigma'(v)$;

(b) for every $C = \langle \mathbf{s}, \mathcal{R}\rangle \in \mathcal{C}$, $\mathbf{s} = (v_1, \dots, v_\ell)$, it holds that $\xi_{\sigma'(v)}(G_{\sigma'(v_i)}) \subseteq \mathsf{pr}_i\mathcal{R}$, $i \in [\ell]$, and there is $C' = \langle \mathbf{s}, \xi^{-1}(\mathcal{R})\rangle \in \mathcal{C}^{\sigma'}$.

The refinement $\mathcal{P}^{\sigma'}$ is *lossless* if for every solution φ of $\mathcal{P}$ the mapping $\xi^{-1} \circ \varphi$ is a solution of $\mathcal{P}^{\sigma'}$. Suppose $\mathcal{G}$ is full pp-definable. The refinement $\mathcal{P}^{\sigma'}$ is *minimal lossless* if it is lossless and for each $v \in V$, $\sigma'(v)$ is minimal (with respect to inclusion of $\xi_{\sigma'(v)}(G_{\sigma'(v)})$ in the original domain). If $\mathcal{G}$ is full of treewidth d, the definition is bit more complicated. The instance $\mathcal{P}$ can also be viewed as a structure $\mathcal{F}$ with vertex set V and such that the solutions of $\mathcal{P}$ are exactly the homomorphisms from $\mathcal{F}$ to $\mathcal{H}$, see [17]. Then $\mathcal{P}^{\sigma'}$ is *minimal lossless of width d* if it is lossless and for every $v \in V$, $\xi_{\sigma'(v)}(G_{\sigma'(v)})$ is minimal with respect to inclusion among unary relations defined by a structure $\mathcal{G}_v$ of treewidth d with a designated variable x and such that there is a homomorphism from $\mathcal{G}_i$ to $\mathcal{F}$ mapping x to v. In fact, a minimal lossless of width d structure $\mathcal{P}^{\sigma'}$ is produced from $\mathcal{P}$ by applying an appropriate local propagation algorithm [30].

▶ **Proposition 14.** *Let $\mathcal{H}$ be a multi-sorted relational structure.*

(1) *Let $\mathcal{G}$ be the full width d definable refinement of $\mathcal{H}$. For any instance $\mathcal{P}$ of $\mathsf{CSP}(\mathcal{H})$, its minimal lossless refinement of width d for $\mathcal{G}$ can be found in polynomial time.*

(2) *Let $\mathcal{H}$ be a relational structure containing all the constant relations and such that $\mathsf{CSP}(\mathcal{H})$ is solvable in polynomial time, and $\mathcal{G}$ the full definable refinement of $\mathcal{H}$. Then for any instance $\mathcal{P}$ of $\mathsf{CSP}(\mathcal{H})$, its minimal lossless refinement for $\mathcal{G}$ can be found in polynomial time.*

5 An Algorithm for Parity

In this section we outline an algorithm that solves $\#_2\mathsf{CSP}(\mathcal{H})$, that is, finds the parity of the number of homomorphisms to $\mathcal{H}$, provided the structure $\mathcal{H}$ satisfies some additional conditions.

▶ **Theorem 15.** *Let $\mathcal{H}$ be a 2-rigid, strongly 2-rectangular, and $\langle\mathcal{H}\rangle_2$ has a Mal'tsev polymorphism*[1]*. Then $\#_2\mathsf{CSP}(\mathcal{H})$ can be solved in time polynomial time.*

In order to prove Theorem 15 we apply some of the existing techniques such as compact representations of relations with a Mal'tsev polymorphism, but in a novel way that is very different from its original use.

Frames and Witness Functions. Suppose that $\mathcal{R}$ is an n-ary relation with a Mal'tsev polymorphism φ. For each $i \in [n]$ we define the following relation $\sim_i$ on $\mathsf{pr}_i\mathcal{R}$: $a \sim_i b$ if there exist tuples $\mathbf{x} \in H^{i-1}$ and $\mathbf{y}_a, \mathbf{y}_b \in H^{n-i}$ such that $(\mathbf{x}, a, \mathbf{y}_a) \in \mathcal{R}$ and $(\mathbf{x}, b, \mathbf{y}_b) \in \mathcal{R}$. For the case $i = 1$, we have $a \sim_1 b$ for all $a, b \in \mathsf{pr}_1\mathcal{R}$ because they share the common empty prefix ε. The relations $\sim_i$ will be called *frame relations*. The following observations are straightforward corollaries of the rectangularity of $\mathcal{R}$ and were used in [14].

▶ **Lemma 16** (Folklore)**.** *Let $\mathcal{R}$ be a relation with a Mal'tsev polymorphism.*
(1) *$\sim_i$ is an equivalence relation for all $i \in [n]$.*
(2) *If $a \sim_i b$ and $\mathbf{x} \in \mathcal{R}$ with $\mathbf{x}_i = a$, then there is a $\mathbf{y} \in \mathcal{R}$ with $\mathbf{y}_i = b$ and $\mathsf{pr}_{[i-1]}\mathbf{x} = \mathsf{pr}_{[i-1]}\mathbf{y}$.*

A mapping $\omega : [n] \times H \to H^n \cup \{\bot\}$ is called a witness function of $\mathcal{R}$ if
(i) For any $i \in [n]$ and $a \in H - \mathsf{pr}_i\mathcal{R}$, $\omega(i, a) = \bot$;
(ii) For any $i \in [n]$ and $a \in \mathsf{pr}_i\mathcal{R}$, $\omega(i, a) \in \mathcal{R}$ is a witness for (i, a), i.e., $\mathsf{pr}_i\omega(i, a) = a$;
(iii) For any $i \in [n]$ and $a, b \in \mathsf{pr}_i\mathcal{R}$ with $a \sim_i b$, we have $\mathsf{pr}_{[i-1]}\omega(i, a) = \mathsf{pr}_{[i-1]}\omega(i, b)$.

A witness function ω provides a concise representation of $\mathcal{R}$. Let $F = \{\omega(i, a) \mid i \in [n], a \in \mathsf{pr}_i\mathcal{R}\}$. Such a set of tuples is called a *frame* of $\mathcal{R}$. A witness function (or a frame) can be found in polynomial time given a conjunctive definition (i.e. a CSP instance) of a relation in a relational structure with a Mal'tsev polymorphism. This is the property that makes them essential for solving CSPs. The following transformations of frames can be performed in polynomial time.

▶ **Proposition 17** ([5, 14])**.** *Let $\mathcal{H}$ be a relational structure with a Mal'tsev polymorphism and a relation $\mathcal{R}$ has a* conjunctive *definition $\mathcal{R}(x_1, \dots, x_n) = \bigwedge_{i\in[m]} \mathcal{R}_i(x_{i_1}, \dots, x_{i_t})$ in $\mathcal{H}$.*
(1) *A frame and a witness function for $\mathcal{R}$ can be computed in $O(mn^4)$.*
(2) *Let $I \subseteq [n]$. Given a frame F for $\mathcal{R}$, a frame for $\mathcal{R}(x_1, x_2, ..., x_n) \wedge (\bigwedge_{s\in I} C_{a_s}(x_s))$ (i.e., x_s is the constant $a_s \in H$, $s \in I$) can be constructed in $O(n^3)$ time.*
(3) *Given a frame F for $\mathcal{R}$, a frame for $\mathcal{Q}(x_1, x_2..., x_{n-1}) = \exists y\mathcal{R}(x_1, ..., x_{n-1}, y)$, can be constructed in $O(n)$ time.*

Overview of the algorithm. The exact counting algorithm in [14] first finds a frame of a conjunctive defined relation (a.k.a. the set of solutions of a CSP instance), and then uses the frame and the condition of balancedness to compute the number of solutions. Unfortunately, this approach does not work in our case, because according to the results of Section 3.1 the property of 2-balancedness that we need here does not correlate well with the existence of a Mal'tsev polymorphism. We therefore choose a different approach.

[1] Note that the latter condition does not follow from $\mathcal{H}$ having a Mal'tsev polymorphism, because 2-mpp-definitions do not preserve polymorphisms.

Let $\mathcal{H}$ be a structure satisfying the conditions of Theorem 15 and $\mathcal{R} \in \langle\mathcal{H}\rangle_2$. Since $\langle\mathcal{H}\rangle_2$ has a Mal'tsev polymorphism, a frame for $\mathcal{R}$ can be computed in polynomial time by Proposition 17(1). Suppose that $\mathcal{R}$ is n-ary. It is not hard to see that

$$|\mathcal{R}| \equiv |\exists^{\equiv 2} y \mathcal{R}(x_1, \dots, x_{n-1}, y)| \pmod 2.$$

Therefore, if it is possible to find a frame of $\widetilde{\mathcal{R}} = \exists^{\equiv 2} y \mathcal{R}(x_1, \dots, x_{n-1}, y)$, we could repeat this process $n-1$ times eventually obtaining a unary relation $\mathcal{R}'$ such that $|\mathcal{R}'| \equiv |\mathcal{R}| \pmod 2$ and then just count the elements in $\mathcal{R}'$ using its frame. Unfortunately, it is not that straightforward, because Proposition 17(3) does not work for modular quantifiers. Instead, we use a more convoluted method.

Let $\mathsf{PAR}_{\mathcal{R}}(\mathbf{x}, y) = \mathcal{R}(\mathbf{x}, y) \wedge (\exists^{\equiv 2} z \ \ \mathcal{R}(\mathbf{x}, z))$. This relation contains essentially the same tuples as $\widetilde{\mathcal{R}}$, except it also keeps their extensions to the last coordinate position. This means that $\widetilde{\mathcal{R}} = \exists y \mathsf{PAR}_{\mathcal{R}}(\mathbf{x}, y)$, and if we know a frame of $\mathsf{PAR}_{\mathcal{R}}(\mathbf{x}, y)$, a frame of $\widetilde{\mathcal{R}}$ can be found by Proposition 17(3). Finding a frame for $\mathsf{PAR}_{\mathcal{R}}(\mathbf{x}, y)$ is the crux of the algorithm.

Finding a frame for $\mathsf{PAR}_{\mathcal{R}}(\mathbf{x}, y)$. Let $\sim_i$ and ω be frame relations and a witness function (a frame) of the relation $\mathcal{R}$ found in the previous step. Let $\mathcal{E}_i$ denote the collection of the equivalence classes of $\sim_i$, and $\mathcal{E}_i = \{\mathcal{E}_{i,1}, \dots, \mathcal{E}_{i,\ell_i}\}$, $\mathcal{E}_{i,j} \subseteq \mathsf{pr}_i\mathcal{R}$, where $j \in [\ell_i]$. We often refer to these classes as *frame classes*. By $\sim'_i, \omega'$, and $\mathcal{E}'_i$, $i \in [n]$, we denote yet unknown frame relations, witness function, and frame classes of $\mathsf{PAR}_{\mathcal{R}}$ (clearly, the number of classes in $\mathcal{E}'_i$ may differ from that of $\mathcal{E}_i$).

First, we observe that by definition a tuple $(\mathbf{x}, a) \in \mathcal{R}$ belongs to $\mathsf{PAR}_{\mathcal{R}}$ if and only if there is a class $\mathcal{E}_{n,s} \in \mathcal{E}_n$, $s \in [\ell_n]$, such that $a \in \mathcal{E}_{n,s}$ and $|\mathcal{E}_{n,s}| \equiv 1 \pmod 2$. This makes finding $\sim'_n$ and $\omega'(n, *)$ easy, they are just restrictions of $\sim_n$ and $\omega(n, *)$ on the union of odd classes from $\mathcal{E}_n$. For $k \in [n-1]$ and $a \in \mathsf{pr}_k\mathcal{R}$ we use Proposition 17(2) to find a witness function $\omega^{k \leftarrow a}$ and frame classes $\mathcal{E}_i^{k \leftarrow a}$ of $\mathcal{R}(x_1, x_2, ..., x_n) \wedge C_a(x_k)$. We examine $\mathcal{E}_{n,s}^{k \leftarrow a}$ for $s \in [|\mathcal{E}_n^{k \leftarrow a}|]$. We check whether there exists $s \in [|\mathcal{E}_n^{k \leftarrow a}|]$ such that $|\mathcal{E}_{n,s}^{k \leftarrow a}| \equiv 1 \mod 2$. If we find such an s, we select $b \in \mathcal{E}_{n,s}^{k \leftarrow a}$ and set $\omega'(k, a) = \omega^{k \leftarrow a}(n, b)$. By construction $\mathsf{pr}_k\omega'(k, a) = \mathsf{pr}_k\omega^{k \leftarrow a}(n, b) = a$, and by the observation above $\omega'(k, a) \in \mathsf{PAR}_{\mathcal{R}}$. Finally, in order to check whether for some $b \in \mathsf{pr}_k\mathcal{R}$ it holds that $a \sim'_k b$ it suffices to complete the following steps. Use use Proposition 17(2) to compute a witness function ω'' and frame classes of

$$\mathcal{R}(x_1, \dots, x_n) \wedge C_b(x_k) \wedge (\bigwedge_{s=1}^{k-1} C_{d_s}(x_s)),$$

where $(d_1, \dots, d_n) = \omega^{k \leftarrow a}(k, a)$. It can be shown that $a \sim'_k b$ if and only if there exists $t \in [|\mathcal{E}''_n|]$ such that $|\mathcal{E}''_{n,t}| \equiv 1 \pmod 2$.

This completes the outline of the algorithm.

6 Hardness and Automorphisms of Direct Products of Structures

Another crucial component of Theorem 2 is the structure of automorphisms of direct products of graphs [24]. It essentially asserts that every automorphism of a direct product $H_1 \times \cdots \times H_n$ can be thought of as a composition of a permutation of factors in the product and automorphisms of those factors. In the following example we show that this breaks down already for digraphs.

▶ **Example 18.** let $\mathcal{H} = (V, E)$ be a directed graph where $V = \{a, b, c, d\}$ with (directed) edge set $E = \{(b, a), (b, c), (c, d)\}$. This digraph is rigid. However, the automorphism group of $\mathcal{H}^2$, see Figure 2, has a complicated structure. As is easily seen $\mathcal{H}^2$ has a large number of automorphisms of order 2 and 3, not all of which have a simple representation mentioned above.

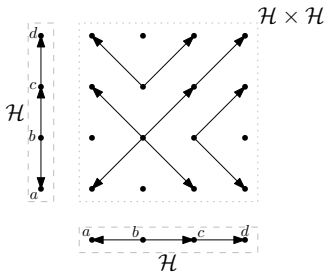

Figure 2 The structure of $\mathcal{H}$ and $\mathcal{H}^2$.

In [9] one of the important applications of the structural theorem for automorphisms of graph products is that it allows one to prove that $\#_p\mathsf{Hom}(\mathcal{H} + \mathcal{R})$, where $\mathcal{H} + \mathcal{R}$ denotes the expansion of $\mathcal{H}$ by a relation $\mathcal{R}$ pp-definable in $\mathcal{H}$, is polynomial time reducible to $\#_p\mathsf{Hom}(\mathcal{H})$. The example above indicates that this result may no longer be true for general relational structures.

Next, we explore what implications of a structural result similar to that in [24] that is used in [9] about $\mathsf{Aut}(\mathcal{H}^n)$ can be. A rectangularity obstruction is a violation of the rectangularity or p-rectangularity property, that is, a (n-ary) relation $\mathcal{R}$ pp- or p-mpp-definable in a structure $\mathcal{H}$, $k \in [n]$, and tuples $\mathbf{a}, \mathbf{b} \in \mathsf{pr}_{[k]}\mathcal{R}, \mathbf{c}, \mathbf{d} \in \mathsf{pr}_{[n]-[k]}\mathcal{R}$ such that $(\mathbf{a}, \mathbf{c}), (\mathbf{a}, \mathbf{d}), (\mathbf{b}, \mathbf{d}) \in \mathcal{R}$, but $(\mathbf{b}, \mathbf{c}) \notin \mathcal{R}$. A *generalized rectangularity obstruction* are the relation $\mathcal{R}$ and sets $A_{1,1}, A_{1,2} \subseteq \mathsf{pr}_{[k]}\mathcal{R}$, $A_{2,1}, A_{2,2} \subseteq \mathsf{pr}_{[n]-[k]}\mathcal{R}$ such that $A_{1,1} \cap A_{1,2} = \emptyset$, $A_{2,1} \cap A_{2,2} = \emptyset$, and any $\mathbf{a} \in A_{1,1}, \mathbf{b} \in A_{1,2}, \mathbf{c} \in A_{2,1}, \mathbf{d} \in A_{2,2}$ form a rectangularity obstruction.

At the first glance, if such an obstruction exists, it should be possible to prove the hardness of $\#_p\mathsf{CSP}(\mathcal{H})$. Indeed, $\mathcal{R}$ can be viewed as a bipartite graph $\mathsf{K}_{\mathcal{R}}$, whose parts of the bipartition are $\mathsf{pr}_{[k]}\mathcal{R}$ and $\mathsf{pr}_{[n]-[k]}\mathcal{R}$, and as the rectangularity obstruction shows, this graph is not complete bipartite implying that $\#\mathsf{Hom}(\mathsf{K}_{\mathcal{R}})$ is #P-complete. However, the hardness of $\#_p\mathsf{Hom}(\mathsf{K}_{\mathcal{R}})$ also involves the requirement that $\mathsf{K}_{\mathcal{R}}$ is p-rigid. p-rigidity is achieved by restricting the problem to induced subgraphs of $\mathsf{K}_{\mathcal{R}}$ as in Lemma 1. However, such a subgraph may avoid the (generalized) rectangularity obstruction rendering it useless.

The obstruction $A_{1,1}, A_{1,2} \subseteq \mathsf{pr}_{[k]}\mathcal{R}$, $A_{2,1}, A_{2,2} \subseteq \mathsf{pr}_{[n]-[k]}\mathcal{R}$ is said to be *protected* in $\mathcal{R}$ if, after applying a p-reduction to $\mathcal{R}$ under a sequence of p-automorphisms from $\mathsf{Aut}(\mathcal{R})$, for the resulting relation $\tilde{\mathcal{R}}$ it holds that $\mathsf{pr}_{[k]}\tilde{\mathcal{R}} \cap A_{1,1}, \mathsf{pr}_{[k]}\tilde{\mathcal{R}} \cap A_{1,2} \neq \emptyset$, $\mathsf{pr}_{[n]-[k]}\tilde{\mathcal{R}} \cap A_{2,1}, \mathsf{pr}_{[n]-[k]}\tilde{\mathcal{R}} \cap A_{2,2} \neq \emptyset$. In fact, Theorem 4.2 [9] implies, although implicitly, that any p-rigid graph that is not a complete bipartite graph contains a protected rectangularity obstruction. One case of a protected rectangularity obstruction is when it is protected in $\mathsf{K}_{\mathcal{R}}$, that is, survives p-reductions of $\mathsf{K}_{\mathcal{R}}$ itself. In this case we say that the obstruction is *graph-protected*.

▶ **Proposition 19.** *Let $\mathcal{H}$ be a (multi-sorted) relational structure and p a prime number. If $\mathcal{H}$ has a graph protected generalized rectangularity obstruction modulo p, $\#_p\mathsf{CSP}(\mathcal{H})$ is $\#_pP$-complete.*

We consider a special case of graph-protected generalized rectangularity obstructions, standard hardness gadgets, that have to satisfy the additional condition $A_{1,1} \cup A_{1,2} = \mathsf{pr}_{[k]}\mathcal{R}, A_{2,1} \cup A_{2,2} = \mathsf{pr}_{[n]-[k]}\mathcal{R}$. It is can be proved that a standard hardness gadget is indeed a graph-protected obstruction.

Standard hardness gadgets provide a fairly limited condition for the hardness of $\#_p\mathsf{CSP}(\mathcal{H})$. In fact, it is possible to prove that $\#_p\mathsf{CSP}(\mathcal{H})$ is $\#_pP$-complete whenever $\mathcal{H}$ has any protected rectangularity obstruction, not necessarily a standard gadget. However, it cannot be done using Theorem 2 as a black box, and is outside of the scope of this paper.

7 Binarization

While studying the structure of $\mathsf{Aut}(\mathcal{H}^k)$ for a relational structure $\mathcal{H}$ and an integer k may be a difficult problem, in this Section we make a step forward by reducing the class of structures $\mathcal{H}$ for which such a characterization is required. In particular, we show that it suffices to obtain a characterization for structures with only binary rectangular relations. More precisely, for any relational structure $\mathcal{H} = (H; \mathcal{R}_1, \dots, \mathcal{R}_k)$ we construct its binarization $b(\mathcal{H})$ as follows. The structure $b(\mathcal{H})$ is multi-sorted, and the domains are the relations $\mathcal{R}_1, \dots, \mathcal{R}_k$ viewed as sets of tuples, thus, $b(\mathcal{H})$ has k domains. For every pair $i, j \in [k]$ (i, j are allowed to be equal) and any $s \in [\ell_i], t \in [\ell_j]$, where ℓ_i, ℓ_j are the arities of $\mathcal{R}_i, \mathcal{R}_j$, the structure $b(\mathcal{H})$ contains a binary relation $\Omega^{ij}_{st} = \{(\mathbf{a}, \mathbf{b}) \mid \mathbf{a} \in \mathcal{R}_i, \mathbf{b} \in \mathcal{R}_j, \mathbf{a}[s] = \mathbf{b}[t]\}$. We show that $\mathcal{H}$ and $b(\mathcal{H})$ share many important properties.

▶ **Theorem 20.** *Let $\mathcal{H}$ be a relational structure. Then $\mathcal{H}$ is strongly rectangular (p-strongly rectangular, p-rigid, has a Mal'tsev polymorphism) if and only if $b(\mathcal{H})$ is strongly rectangular (p-strongly rectangular, p-rigid, has a Mal'tsev polymorphism).*

In addition to Theorem 20 every relation of $b(\mathcal{H})$ is binary and rectangular. This makes such structures somewhat closer to graphs and the hope is that it will be easier to study the structure of $\mathsf{Aut}(b(\mathcal{H})^n)$ than $\mathsf{Aut}(\mathcal{H}^n)$ itself.

References

1 Libor Barto, Andrei Krokhin, and Ross Willard. Polymorphisms, and how to use them. In *Dagstuhl Follow-Ups*, volume 7. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017. doi:10.4230/DFU.VOL7.15301.1.

2 Alexander I. Barvinok. *Combinatorics and Complexity of Partition Functions*, volume 30 of *Algorithms and combinatorics*. Springer, 2016. doi:10.1007/978-3-319-51829-9.

3 Andrei A. Bulatov. The complexity of the counting constraint satisfaction problem. *J. ACM*, 60(5):34:1–34:41, 2013. doi:10.1145/2528400.

4 Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. In Chris Umans, editor, *FOCS*, pages 319–330, 2017. doi:10.1109/FOCS.2017.37.

5 Andrei A. Bulatov and Víctor Dalmau. A simple algorithm for Mal'tsev constraints. *SIAM J. Comput.*, 36(1):16–27, 2006. doi:10.1137/050628957.

6 Andrei A. Bulatov and Víctor Dalmau. Towards a dichotomy theorem for the counting constraint satisfaction problem. *Information and Computation*, 205(5):651–678, 2007. doi:10.1016/J.IC.2006.09.005.

7 Andrei A. Bulatov and Martin Grohe. The complexity of partition functions. *Theor. Comput. Sci.*, 348(2-3):148–186, 2005. doi:10.1016/J.TCS.2005.09.011.

8 Andrei A. Bulatov, Peter Jeavons, and Andrei A. Krokhin. Classifying the complexity of constraints using finite algebras. *SIAM J. Comput.*, 34(3):720–742, 2005. doi:10.1137/S0097539700376676.

9 Andrei A. Bulatov and Amirhossein Kazeminia. Complexity classification of counting graph homomorphisms modulo a prime number. In *STOC*, pages 1024–1037. ACM, 2022. doi:10.1145/3519935.3520075.

10 Jin-Yi Cai and Xi Chen. Complexity of counting CSP with complex weights. *J. ACM*, 64(3), June 2017. doi:10.1145/2822891.

11 Jin-yi Cai and Lane A. Hemachandra. On the power of parity polynomial time. In Burkhard Monien and Robert Cori, editors, *STACS*, volume 349 of *Lecture Notes in Computer Science*, pages 229–239. Springer, 1989. doi:10.1007/BFB0028987.

12 Víctor Dalmau and Peter Jonsson. The complexity of counting homomorphisms seen from the other side. *Theor. Comput. Sci.*, 329(1-3):315–323, 2004. doi:10.1016/J.TCS.2004.08.008.

13 M. Dyer and C. Greenhill. The complexity of counting graph homomorphisms. *Random Structures and Algorithms*, 17:260–289, 2000. doi:10.1002/1098-2418(200010/12)17:3/4\%3C260::AID-RSA5\%3E3.0.CO;2-W.

14 Martin Dyer and David Richerby. An effective dichotomy for the counting constraint satisfaction problem. *SIAM J. on Comp*, 42(3):1245–1274, 2013. doi:10.1137/100811258.

15 John Faben. The complexity of counting solutions to generalised satisfiability problems modulo k, 2008. arXiv:0809.1836.

16 John Faben and Mark Jerrum. The complexity of parity graph homomorphism: an initial investigation. *arXiv preprint arXiv:1309.4033*, 2013. arXiv:1309.4033.

17 Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through datalog and group theory. *SIAM J. Comput.*, 28(1):57–104, 1998. doi:10.1137/S0097539794266766.

18 Jacob Focke, Leslie Ann Goldberg, Marc Roth, and Stanislav Zivný. Counting homomorphisms to k_4-minor-free graphs, modulo 2. In Dániel Marx, editor, *SODA*, pages 2303–2314. SIAM, 2021. doi:10.1137/1.9781611976465.137.

19 Andreas Göbel, Leslie Ann Goldberg, and David Richerby. Counting homomorphisms to square-free graphs, modulo 2. *ACM Transactions on Computation Theory (TOCT)*, 8(3):1–29, 2016. doi:10.1145/2898441.

20 Andreas Göbel, J. A. Gregor Lagodzinski, and Karen Seidel. Counting homomorphisms to trees modulo a prime. In *MFCS*, volume 117, pages 49:1–49:13, 2018. doi:10.4230/LIPICS.MFCS.2018.49.

21 Heng Guo, Sangxia Huang, Pinyan Lu, and Mingji Xia. The Complexity of Weighted Boolean #CSP Modulo k. In *STACS)*, volume 9, pages 249–260, 2011. doi:10.4230/LIPICS.STACS.2011.249.

22 Andreas Göbel, Leslie Ann Goldberg, and David Richerby. The complexity of counting homomorphisms to cactus graphs modulo 2. *ACM Trans on Comp Th*, 6(4):1–29, 2014. doi:10.1145/2635825.

23 J. Hagemann and A. Mitschke. On n-permutable congruences. *Algebra Universalis*, 3:8–12, 1973.

24 Richard Hammack, Wilfried Imrich, and Sandi Klavzar. *Handbook of Product Graphs, Second Edition*. CRC Press, Inc., USA, 2nd edition, 2011.

25 Ulrich Hertrampf. Relations among mod-classes. *Theor. Comput. Sci.*, 74(3):325–328, 1990. doi:10.1016/0304-3975(90)90081-R.

26 Peter Jeavons. On the algebraic structure of combinatorial problems. *Theoretical Computer Science*, 200(1-2):185–204, 1998. doi:10.1016/S0304-3975(97)00230-2.

27 Mark Jerrum and Alistair Sinclair. Polynomial-time approximation algorithms for the Ising model. *SIAM J. Comput.*, 22(5):1087–1116, 1993. doi:10.1137/0222066.

28 Amirhossein Kazeminia and Andrei A Bulatov. Counting homomorphisms modulo a prime number. In *MFCS*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019.

29 Amirhossein Kazeminia and Andrei A. Bulatov. Modular counting csp: Reductions and algorithms, 2025. arXiv:2501.04224.

30 Phokion G Kolaitis. Constraint satisfaction, complexity, and logic. In *Hellenic Conference on Artificial Intelligence*, pages 1–2. Springer, 2004. doi:10.1007/978-3-540-24674-9_1.

31 J. A. Gregor Lagodzinski, Andreas Göbel, Katrin Casel, and Tobias Friedrich. On counting (quantum-)graph homomorphisms in finite fields of prime order. *CoRR*, abs/2011.04827, 2021. arXiv:2011.04827.

32 E.H. Lieb and A.D. Sokal. A general Lee-Yang theorem for one-component and multicomponent ferromagnets. *Communications in Mathematical Physics*, 80(2):153–179, 1981.

33 L. Valiant. The complexity of computing the permanent. *Theoretical Computing Science*, 8:189–201, 1979. doi:10.1016/0304-3975(79)90044-6.

34 L. Valiant. The complexity of enumeration and reliability problems. *SIAM Journal on Computing*, 8(3):410–421, 1979. doi:10.1137/0208032.

35 Dmitriy Zhuk. A proof of the CSP dichotomy conjecture. *J. ACM*, 67(5):30:1–30:78, 2020. doi:10.1145/3402029.