



Cohomology in Constraint Satisfaction and Structure Isomorphism

Adam Ó Conghaile   

Computer Laboratory, Cambridge University, United Kingdom

The Alan Turing Institute, United Kingdom

Abstract

Constraint satisfaction (CSP) and structure isomorphism (SI) are among the most well-studied computational problems in Computer Science. While neither problem is thought to be in PTIME, much work is done on PTIME approximations to both problems. Two such historically important approximations are the k -consistency algorithm for CSP and the k -Weisfeiler-Leman algorithm for SI, both of which are based on propagating local partial solutions. The limitations of these algorithms are well-known – k -consistency can solve precisely those CSPs of bounded width and k -Weisfeiler-Leman can only distinguish structures which differ on properties definable in C^{rk} . In this paper, we introduce a novel sheaf-theoretic approach to CSP and SI and their approximations. We show that both problems can be viewed as deciding the existence of global sections of presheaves, $\mathcal{H}_k(A, B)$ and $\mathcal{I}_k(A, B)$ and that the success of the k -consistency and k -Weisfeiler-Leman algorithms correspond to the existence of certain efficiently computable subpresheaves of these. Furthermore, building on work of Abramsky and others in quantum foundations, we show how to use Čech cohomology in $\mathcal{H}_k(A, B)$ and $\mathcal{I}_k(A, B)$ to detect obstructions to the existence of the desired global sections and derive new efficient cohomological algorithms extending k -consistency and k -Weisfeiler-Leman. We show that cohomological k -consistency can solve systems of equations over all finite rings and that cohomological Weisfeiler-Leman can distinguish positive and negative instances of the Cai-Fürer-Immerman property over several important classes of structures.

2012 ACM Subject Classification Theory of computation \rightarrow Finite Model Theory

Keywords and phrases constraint satisfaction problems, finite model theory, descriptive complexity, rank logic, Weisfeiler-Leman algorithm, cohomology

Digital Object Identifier 10.4230/LIPIcs.MFCS.2022.75

Related Version *Full Version:* <https://arxiv.org/abs/2206.15253>

Funding Funded in part by EPSRC grants EP/R513180/1 and EP/S03238X/1 and supported by a PhD Enrichment Award from the Alan Turing Institute.

Acknowledgements I thank Anuj Dawar and Samson Abramsky for useful discussions in writing this paper and am especially gratefully to Samson Abramsky for permission to use results from [1].

1 Introduction

Constraint satisfaction problems (CSP) and structure isomorphism (SI) are two of the most well-studied problems in complexity theory. Mathematically speaking, an instance of one of these problems takes a pair of structures (A, B) as input and asks whether there is a homomorphism $A \rightarrow B$ for CSP or an isomorphism $A \cong B$ for SI. These problems are not in general thought to be tractable. Indeed the general case of CSP is NP-Complete and restricting our structures to graphs the best known algorithm for SI is Babai’s quasi-polynomial time algorithm [8]. As a result, it is common in complexity and finite model theory to study approximations of the relations \rightarrow and \cong .



© Adam Ó Conghaile;

licensed under Creative Commons License CC-BY 4.0

47th International Symposium on Mathematical Foundations of Computer Science (MFCS 2022).

Editors: Stefan Szeider, Robert Ganian, and Alexandra Silva; Article No. 75; pp. 75:1–75:16

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The k -consistency and k -Weisfeiler-Leman¹ algorithms efficiently determine two such approximations to \rightarrow and \cong which we call \rightarrow_k and \equiv_k . These relations have many characterisations in logic and finite model theory, for example in [18] and [13]. One that is particularly useful is that of the existence of winning strategies for Duplicator in certain Spoiler-Duplicator games with k pebbles [28, 25]. For both of these games Duplicator’s winning strategies can be represented as non-empty sets $S \subset \mathbf{Hom}_k(A, B)$ of k -local partial homomorphisms which satisfy some extension properties and connections between these games have been studied before. For example, a joint comonadic semantics is given by the pebbling comonad of Abramsky, Dawar and Wang [4].

The limitations of these approximations are well-known. In particular, it is known that k -consistency only solves CSPs of *bounded width* and k -Weisfeiler-Leman can only distinguish structures which differ on properties expressible in the infinitary counting logic C^k . Feder and Vardi [18] showed that CSP encoding linear equations over the finite fields do not have bounded width, while Cai, Fürer, and Immerman [13] demonstrated an efficiently decidable graph property which is not expressible in C^k for any k .

In the present paper, we introduce a novel approach to the CSP and SI problems based on presheaves of k -local partial homomorphisms and isomorphisms, showing that the problems can be reframed as deciding whether certain presheaves admit global sections. We show that the classic k -consistency and k -Weisfeiler-Leman algorithms can be derived by computing greatest fixpoints of presheaf operators which remove some efficiently computable obstacles to global sections. Furthermore, we show how invariants from sheaf cohomology can be used to find further obstacles to combining local homomorphisms and isomorphisms into global ones. We use these to construct new efficient extensions to the k -consistency and k -Weisfeiler-Leman algorithms computing relations $\rightarrow_k^{\mathbb{Z}}$ and $\equiv_k^{\mathbb{Z}}$ which refine \rightarrow_k and \equiv_k .

The application of presheaves has been particularly successful in computer science in recent decades with applications in semantics [32, 19], information theory [33] and quantum contextuality [3, 5, 2]. This work draws in particular on the application of sheaf theory to quantum contextuality, pioneered by Abramsky and Brandenburger [3] and developed by Abramsky and others for example in [5] and [2].

Using this work, we prove that these new cohomological algorithms are strictly stronger than k -consistency and k -Weisfeiler-Leman. In particular, we show that cohomological k -consistency decides solvability of linear equations with k variables per equation over all finite rings and that there is a fixed k such that $\equiv_k^{\mathbb{Z}}$ distinguishes structures which differ on Cai, Fürer and Immerman’s property.

It is also interesting to compare $\rightarrow_k^{\mathbb{Z}}$ and $\equiv_k^{\mathbb{Z}}$ with other well-studied refinements of \rightarrow_k and \equiv_k . For \rightarrow_k , such refinements include the algorithms of Bulatov [12] and Zhuk [35] which decide all tractable CSPs and the algorithms of Brakensiek, Guruswami, Wrochna and Živný [11] and Ciardo and Živný [14] for Promise CSPs. For \equiv_k , comparable approximations to \cong include linear Diophantine equation methods employed by Berkholz and Grohe [9] and the invertible-map equivalence of Dawar and Holm [17] which bounds the expressive power of rank logic. The latter was recently used by Lichter [30] to demonstrate a property which is decidable in PTIME but not expressible in rank logic. In our paper, we show that $\equiv_k^{\mathbb{Z}}$, for some fixed k , can distinguish structures which differ on this property. Comparing $\rightarrow_k^{\mathbb{Z}}$ to the Bulatov-Zhuk algorithm and algorithms for PCSPs remains a direction for future work.

¹ The algorithm we call “ k -Weisfeiler-Leman” is more commonly called “ $(k - 1)$ -Weisfeiler-Leman” in the literature, see for example [13]. We prefer “ k -Weisfeiler-Leman” to emphasise its relationship to k -variable logic and sets of k -local isomorphisms.

The rest of the paper proceeds as follows. Section 2 establishes some background and notation. Section 3 introduces the presheaf formulation of CSP and SI and new formulations of k -consistency and k -Weisfeiler-Leman in this framework. Section 4 demonstrates how to apply aspects of sheaf cohomology to CSP and SI and defines new algorithms along these lines. Section 5 surveys the strength of these new cohomological algorithms. Section 6 concludes with some open questions and directions for future work. Major proofs and additional background are left to the full version.

2 Background and definitions

In this section, we record some definitions and background which are necessary for our work.

2.1 Relational structures & finite model theory

Throughout this paper we use the word *structure* to mean a relational structure over some finite relational signature σ . A structure A consists of an underlying set (which we also call A) and for each relational symbol R of arity r in σ a subset $R^A \subset A^r$ or tuples related by R . A *homomorphism* of structures A, B over a common signature is a function between the underlying sets $f: A \rightarrow B$ which preserves related tuples. An *isomorphism* of structures is a bijection between the underlying sets which both preserves and reflects related tuples. A partial function $s: A \rightarrow B$ (seen as a set $s \subset A \times B$) is a *partial homomorphism* if it preserves the related tuples in $\mathbf{dom}(s)$. s is a *partial isomorphism* if it is a bijection onto its image and both preserves and reflects related tuples. A partial homomorphism or isomorphism is said to be *k -local* if $|\mathbf{dom}(s)| \leq k$. For two structures over the same signature we write $\mathbf{Hom}_k(A, B)$ and $\mathbf{Isom}_k(A, B)$ respectively for the sets of k -local homomorphisms and isomorphisms from A to B .

In the paper, we make reference to several important logics from finite model theory and descriptive complexity theory. The logics we make reference to in this paper are as follows.

- Fixed-point logic with counting (written FPC) is first-order logic extended with operators for inflationary fixed-points and counting, for example see [20].
- For any natural number k , C^k is infinitary first-order logic extended with counting quantifiers with at most k variables. This logic bounds the expressive power of FPC in the sense that, for each k' there exists k such that any FPC formula in k' variables is equivalent to one in C^k . We write C^ω for the union of these logics.
- Rank logic is first-order logic extended with operators for inflationary fixed-points and computing ranks of matrices over finite fields, see [34].
- Linear algebraic logic is first-order infinitary logic extended with quantifiers for computing *all* linear algebraic functions over finite fields, see [15]. This logic bounds rank logic in the sense described above.

At different points in the history of descriptive complexity theory, both FPC and rank logic were considered as candidates for “capturing PTIME” and thus refuting a well-known conjecture of Gurevich [23]. Each has since been proven not to capture PTIME, for FPC see Cai, Fürer and Immerman [13], for rank logic see Lichter [30]. Infinitary logics such as C^ω and linear algebraic logic are capable of expressing properties which are not decidable in PTIME but have been shown not to contain any logic which does not capture PTIME. For C^ω , see Cai, Fürer and Immerman [13] and for linear algebraic logic, see Dawar, Grädel, and Lichter [16].

2.2 Constraint satisfaction problems & Structure Isomorphism

Assuming a fixed relational signature σ , we write CSP for the set of all pairs of σ -structures (A, B) such that there is a homomorphism witnessing $A \rightarrow B$. We use $CSP(B)$ to denote the set of relational structures A such that $(A, B) \in CSP$. We also use CSP and $CSP(B)$ to denote the decision problem on these sets. For general B , $CSP(B)$ is well-known to be NP-complete. However for certain structures B the problem is in PTIME. Indeed, the Bulatov-Zhuk Dichotomy Theorem (formerly the Feder-Vardi Dichotomy Conjecture) states that, for any B , $CSP(B)$ is either NP-complete or it is PTIME. Working out efficient algorithms which decide $CSP(B)$ for larger and larger classes of B was an active area of research which culminated in Bulatov and Zhuk's exhaustive classes of algorithms [12, 35].

Similarly, we write SI for the set of all pairs of σ -structures (A, B) such that there is an isomorphism witnessing $A \cong B$. The decision problem for this set is also thought not to be in PTIME however there are no general hardness results known for this. The best known algorithm (in the case where σ is the signature of graphs) is Babai's [8] which is quasi-polynomial.

There are many efficient algorithms which approximate the decision problems of CSP and SI . Two such examples, which are of particular importance to this paper, are the k -consistency and k -Weisfeiler-Leman algorithms. Explicit modern presentations of these algorithms can be seen, for example, in [7] and [27]. We instead focus on equivalent formulations in terms of positional Duplicator winning strategies. These are given by Kolaitis and Vardi [28] for k -consistency and Hella [24] for k -Weisfeiler-Leman. In the case of k -consistency, a pair (A, B) is accepted by the algorithm if and only if there is a non-empty subset $S \subset \mathbf{Hom}_k(A, B)$ which is downward-closed and satisfies the so-called forth property. This means any $s \in S$ with $|\mathbf{dom}(s)| < k$ satisfies the property $\mathbf{Forth}(S, s)$ which is defined as

$$\forall a \in A, \exists b \in B \text{ s.t. } s \cup \{(a, b)\} \in S.$$

If such an S exists we write $A \rightarrow_k B$. The similar strategy-based characterisation of k -Weisfeiler-Leman is captured by non-empty downward-closed sets $S \subset \mathbf{Isom}_k(A, B)$ where each element satisfies the bijective forth property $\mathbf{BijForth}(S, s)$ defined by

$$\exists b_s: A \rightarrow B \text{ a bijection s.t. } \forall a \in A \ s \cup \{(a, b_s(a))\} \in S.$$

If such an S exists we write $A \equiv_k B$. For more details, see the full version of this paper.

2.3 Presheaves & cohomology

Here we give a brief account of the category-theoretic preliminaries for this paper. For a more comprehensive introduction to category theory we refer to Chapter 1 of Leinster's textbook [29] and for a complete account of presheaves we refer to Chapter 2 of MacLane and Moerdijk [31].

Given two categories \mathbf{C} and \mathbf{S} , an \mathbf{S} -valued presheaf over \mathbf{C} is a contravariant functor $\mathcal{F}: \mathbf{C}^{op} \rightarrow \mathbf{S}$. We will assume that \mathbf{C} is some subset of the powerset of some set X with subset inclusion as the morphisms. We call X the underlying space of \mathbf{C} . For this reason, when $U' \subset U$ in \mathbf{C} we write $(\cdot)_{|_{U'}}$ for the restriction map $\mathcal{F}(U' \subset U): \mathcal{F}(U) \rightarrow \mathcal{F}(U')$. We assume \mathbf{S} is either the category \mathbf{Set} of sets or the category \mathbf{AbGrp} of abelian groups. We call \mathbf{AbGrp} -valued presheaves, abelian presheaves. \mathbf{Set} -valued presheaves are just called presheaves or presheaves of sets where there is ambiguity.

For any \mathbf{C} and \mathbf{S} as above, the category of presheaves $\mathbf{PrSh}(\mathbf{C}, \mathbf{S})$ has as objects the presheaves $\mathcal{F}: \mathbf{C}^{op} \rightarrow \mathbf{S}$ and, as morphisms, natural transformations between these functors. If \mathbf{S} has a terminal object 1 (as both \mathbf{Set} and \mathbf{AbGp} do) then the presheaf $\mathbb{I} \in \mathbf{PrSh}(\mathbf{C}, \mathbf{S})$ which sends all elements of \mathbf{C} to 1 is a terminal object in $\mathbf{PrSh}(\mathbf{C}, \mathbf{S})$. For any $\mathcal{F} \in \mathbf{PrSh}(\mathbf{C}, \mathbf{S})$, a *global section* of \mathcal{F} is a natural transformation $S: \mathbb{I} \Rightarrow \mathcal{F}$.

3 Presheaves of local homomorphisms and isomorphisms

Some important efficient algorithms for CSP and SI involve working with sets of k -local homomorphisms between the two structures in a given instance. These sets of partial homomorphisms of domain size $\leq k$ are useful for constructing efficient algorithms because computing the sets $\mathbf{Hom}_k(A, B)$ and $\mathbf{Isom}_k(A, B)$ can be done in polynomial time in $|A| \cdot |B|$. In this section, we see that these sets can naturally be given the structure of sheaves, that the CSP and SI problems can be seen as the search for global sections of these sheaves and that the k -consistency and k -Weisfeiler-Leman algorithms can both be seen as determining the existence of certain special subpresheaves. The framework of considering sheaves of local homomorphisms and isomorphisms is novel in this work and essential for the main cohomological algorithms later. The results in Section 3.3 are from a technical report of Samson Abramsky [1] and we thank him for his permission to include them here.

3.1 Defining presheaves of homomorphisms and isomorphisms

Let A and B be relational structures over the same signature. A *partial homomorphism* is a partial function $s: A \rightarrow B$ that preserves related tuples in $\mathbf{dom}(s)$. A *partial isomorphism* is a partial homomorphism $s: A \rightarrow B$ which is injective and reflects related tuples from $\mathbf{im}(s)$. A *k -local homomorphism (resp. isomorphism)* is a partial homomorphism (resp. isomorphism) s such that $|\mathbf{dom}(s)| \leq k$. We write $\mathbf{Hom}_k(A, B)$ (resp. $\mathbf{Isom}_k(A, B)$) for the sets of k -local homomorphisms (resp. isomorphisms). We write $\mathbf{Hom}(A, B)$ for the union $\bigcup_{1 \leq k \leq |A|} \mathbf{Hom}_k(A, B)$ and $\mathbf{Isom}(A, B)$ for the union $\bigcup_{1 \leq k \leq |A|} \mathbf{Isom}_k(A, B)$.

It is not hard to see that these sets can be given the structure of presheaves on the underlying space A . Indeed, we define the *presheaf of homomorphisms from A to B* $\mathcal{H}(A, B): \mathbf{P}(A)^{op} \rightarrow \mathbf{Set}$ as $\mathcal{H}(A, B)(U) = \{s \in \mathbf{Hom}(A, B) \mid \mathbf{dom}(s) = U\}$ with restriction maps $\mathcal{H}(A, B)(U' \subset U)$ given by the restriction of partial homomorphisms $(\cdot)_{|U'}$. Similarly, let $\mathcal{I}(A, B)$ be the subpresheaf of $\mathcal{H}(A, B)$ containing only partial isomorphisms. Now, consider the cover of A by subsets of size at most k , written $A^{\leq k} \subset \mathbf{P}(A)$. We define the *presheaves of k -local homomorphisms and isomorphisms* $\mathcal{H}_k(A, B)$ and $\mathcal{I}_k(A, B)$ as the functors $\mathcal{H}(A, B)$ and $\mathcal{I}(A, B)$ restricted to the subcategory $(\mathbf{A}^{\leq k})^{op} \subset \mathbf{P}(A)^{op}$.

We now see how these presheaves and their global sections encode the CSP and SI problems for the instance (A, B) .

3.2 CSP and SI as search for global sections

Fix an instance (A, B) for the CSP or SI problem and let \mathcal{H} and \mathcal{I} stand for the presheaves of all partial homomorphisms and isomorphisms between A and B defined in the last section. For either of these presheaves \mathcal{S} a global section $s: \mathbb{I} \Rightarrow \mathcal{S}$ is a collection $\{s_U \in \mathcal{S}(U)\}_{U \in \mathbf{P}(A)}$ where naturality implies that for any subsets U and U' of A $(s_U)_{|U \cap U'} = (s_{U'})_{|U \cap U'}$. As the poset $\mathbf{P}(A)$ has a maximal element, namely A , any such global section is determined by a choice of $s_A \in \mathcal{S}(A)$. This leads us to the following observation.

► **Observation 1.** *Given a pair (A, B) relational structures over the same signature then*

$$(A, B) \in \text{CSP} \iff \mathcal{H} \text{ has a global section}$$

and if $|A| = |B|$ then

$$(A, B) \in \text{SI} \iff \mathcal{I} \text{ has a global section.}$$

This observation reframes the CSP and SI problems in terms of presheaves but algorithmically this not a particularly useful restating as computing the full objects \mathcal{H} and \mathcal{I} requires solving the CSP and SI problems for all subsets of A and B . A much more interesting equivalent condition is that for large enough k , whether or not a particular instance (A, B) is in CSP or SI is determined by the global sections of the presheaves of k -local homomorphisms and isomorphisms.

► **Lemma 2.** *For a pair (A, B) relational structures over the same signature, σ , and k at least the arity of sigma then*

$$(A, B) \in \text{CSP} \iff \mathcal{H}_k \text{ has a global section}$$

and if $|A| = |B|$ then

$$(A, B) \in \text{SI} \iff \mathcal{I}_k \text{ has a global section.}$$

Proof. See full version. ◀

This is more interesting than the previous observation as \mathcal{H}_k and \mathcal{I}_k can be computed for any relational structures A and B in $\mathcal{O}(\text{poly}(|A| \cdot |B|))$. Indeed, we can just list all $\mathcal{O}(|A|^k \cdot |B|^k)$ possible k -local functions and check which ones preserve (and reflect) related tuples. This also gives us an interesting starting point for designing efficient algorithms for approximating CSP and SI. In particular, any efficient algorithms which finds obstacles to the existence of global sections in \mathcal{H}_k and \mathcal{I}_k will provide a tractable approximation to CSP and SI. We now see how this approach can be used to capture some classical approximations of these problems.

3.3 Algorithms and games in terms of presheaves

In this section, we consider the approximations $A \rightarrow_k B$ and $A \equiv_k B$ to CSP and SI which are computed respectively by the k -consistency and k -Weisfeiler-Leman algorithms and we show that these algorithms can be seen as searching for certain obstructions to global sections in $\mathcal{H}_k(A, B)$ and $\mathcal{I}_k(A, B)$. In particular, we define efficiently computable monotone operators on subpresheaves of \mathcal{H}_k and \mathcal{I}_k and show that they have non-empty greatest fixpoints if and only if (A, B) are accepted by k -consistency and k -Weisfeiler-Leman respectively. Proposition 3 is reproduced with permission from an unpublished technical report of Samson Abramsky and the formulation of the fixpoint operators is inspired by the same report.

3.3.1 Flasque presheaves and k -consistency

Recall that $A \rightarrow_k B$ if and only if there is a positional winning strategy for Duplicator in the existential k -pebble game [18] and that a presheaf \mathcal{F} is flasque if all of the restriction maps $\mathcal{F}(U \subset U')$ are surjective. In a recent technical report, Abramsky [1] proves the following characterisation of these strategies in our presheaf setting.

- **Proposition 3.** *For A, B relational structures and any k there is a bijection between:*
- *positional strategies in the existential k -pebble game from A to B , and*
 - *non-empty flasque subpresheaves $\mathcal{S} \subset \mathcal{H}_k(A, B)$.*

This gives an alternative description of the k -consistency algorithm as constructing the largest flasque subpresheaf $\overline{\mathcal{H}}_k$ of \mathcal{H}_k and checking if it is empty. As pointed out by Abramsky [1], this is the process of coflasquification of the presheaf \mathcal{H}_k and can be seen as dual to the Godement construction [21], an important early construction in homological algebra. $\overline{\mathcal{H}}_k$ can be computed efficiently as the greatest fixpoint of the presheaf operator $(\cdot)^{\uparrow\downarrow}$ which computes the largest subpresheaf of a presheaf $\mathcal{S} \subset \mathcal{H}_k$ such that every $s \in \mathcal{S}^{\uparrow\downarrow}(C)$ satisfies the forth property **Forth**(\mathcal{S}, s). For further details see the full version of this paper.

3.3.2 Greatest fixpoints and k -Weisfeiler-Leman

In a similar way to the k -consistency algorithm, k -Weisfeiler-Leman can be formulated as determining the existence of a positional strategy for Duplicator in the k -pebble bijection game between A and B . This inspires the definition of another efficiently computable presheaf operator $(\cdot)^{\#\downarrow}$ which computes the largest subpresheaf of a presheaf $\mathcal{S} \subset \mathcal{I}_k$ such that for every $s \in \mathcal{S}^{\#\downarrow}(C)$ satisfies the bijective forth property **BijForth**(\mathcal{S}, s). We call the greatest fixpoint of this operator $\overline{\mathcal{S}}$ and we have that $A \equiv_k B$ if and only if $\overline{\mathcal{I}}_k$ is non-empty. For more details, see the full version of this paper.

To conclude, in this section, we have seen how to reformulate the search for homomorphisms and isomorphisms between relational structures A and B as the search for global sections in the presheaves $\mathcal{H}_k(A, B)$ and $\mathcal{I}_k(A, B)$. We have also seen that well-known approximations of homomorphism and isomorphism, \rightarrow_k and \equiv_k , can be computed as greatest fixpoints of presheaf operators which remove elements which cannot form part of any global section. In the next section, we look at sheaf-theoretic obstructions to forming a global section which come from cohomology and see how these can be used to define stronger approximations of homomorphism and isomorphism.

4 Cohomology of local homomorphisms and isomorphisms

As we showed in the previous section, an instance of CSP and SI with input (A, B) can be seen as determining the existence of a global section for the presheaf $\mathcal{H}_k(A, B)$ or $\mathcal{I}_k(A, B)$ respectively and that the classic k -consistency and k -Weisfeiler-Leman algorithms can be reformulated as computing greatest fixed points of presheaf operations which successively remove sections which are obstructed from being part of some global section. In this section, we extend these algorithms by considering further efficiently computable obstructions which arise naturally from presheaf cohomology. From this we derive new cohomological algorithms for CSP and SI.

4.1 Cohomology and local vs. global problems

The notion of computing cohomology valued in an **AbGp**-valued presheaf \mathcal{F} on a topological space X has a long history in algebraic geometry and algebraic topology which dates back to Grothendieck's seminal paper on the topic [22]. The cohomology valued in \mathcal{F} consists of a sequence of abelian groups $H^i(X, \mathcal{F})$ where $H^0(X, \mathcal{F})$ is the free \mathbb{Z} -module over global sections of \mathcal{F} . As seen in the previous section we may be interested in such global sections but their existence may be difficult to determine. This is where the functorial nature of cohomology is extremely useful. Indeed, any short exact sequence of presheaves

$$0 \rightarrow \mathcal{F}_L \rightarrow \mathcal{F} \rightarrow \mathcal{F}_R \rightarrow 0$$

lifts to a long exact sequence of cohomology groups

$$0 \rightarrow H^0(X, \mathcal{F}_L) \rightarrow H^0(X, \mathcal{F}) \rightarrow H^0(X, \mathcal{F}_R) \rightarrow H^1(X, \mathcal{F}_L) \rightarrow \dots$$

This tells us that the global sections of \mathcal{F}_R which are not images of global sections of \mathcal{F} are mapped to non-trivial elements of the group $H^1(X, \mathcal{F}_L)$ by the maps in this sequence. This means that these higher cohomology groups can be seen as a source of obstacles to lifting “local” solutions in \mathcal{F}_R to “global” solutions in \mathcal{F} .

An important recent example of such an application of cohomology to finite structures can be found in the work of Abramsky, Barbosa, Kishida, Lal and Mansfield [2] in quantum foundations. They show that cohomological obstructions of the type described above can be used to detect contextuality (locally consistent measurements which are globally inconsistent) in quantum systems which were earlier given a presheaf semantics by Abramsky and Brandenburger [3]. In the full version of this paper, we describe these obstructions in general and show how the presheaves we constructed in the last section admit the same cohomological obstructions. This similarity inspires the definitions and algorithms which follow in the next two sections.

4.2 \mathbb{Z} -local sections and \mathbb{Z} -extendability

Returning to presheaves of local homomorphisms and isomorphisms let \mathcal{S} be a subpresheaf of \mathcal{H}_k . Then we define the presheaf of \mathbb{Z} -linear local sections of \mathcal{S} to be the presheaf of formal \mathbb{Z} -linear sums of local sections of \mathcal{S} . This means that for any $C \in A^{\leq k}$

$$\mathbb{Z}\mathcal{S}(C) := \left\{ \sum_{s \in \mathcal{S}(C)} \alpha_s s \mid \alpha_s \in \mathbb{Z} \right\}.$$

This is an abelian presheaf on $A^{\leq k}$ and we call the global sections $\{r_U \in \mathbb{Z}\mathcal{S}(U)\}_{U \in A^{\leq k}}$ \mathbb{Z} -linear global sections of \mathcal{S} . We say that a local section $s \in \mathcal{S}(C)$ is \mathbb{Z} -extendable if there is a \mathbb{Z} -linear global section $\{r_U \in \mathbb{Z}\mathcal{S}(U)\}_{U \in A^{\leq k}}$ such that $r_C = s$. We write this condition as $\mathbf{Zext}(\mathcal{S}, s)$. As outlined by Abramsky, Barbosa and Mansfield [5], this condition corresponds to the absence of a cohomological obstruction to \mathcal{S} containing a global section involving s .

Importantly for our purposes, deciding the condition $\mathbf{Zext}(\mathcal{S}, s)$ for any $\mathcal{S} \subset \mathcal{H}_k(A, B)$ is computable in polynomial time in the sizes of A and B . This is because the compatibility conditions for a collection $\{r_U \in \mathbb{Z}\mathcal{S}(U)\}_{U \in A^{\leq k}}$ being a global section of $\mathbb{Z}\mathcal{S}$ can be expressed as a system of polynomially many linear equations in polynomially many variables. Indeed, we write each r_U as $\sum_{s \in \mathcal{S}(U)} \alpha_s s$ where α_s is a variable for each $s \in \mathcal{S}(U)$. This gives a total number of variables bounded by $\mathcal{O}(|A|^k \cdot |B|^k)$, the size of $\mathbf{Hom}_k(A, B)$. For each of the $\mathcal{O}(|A|^{2k})$ pairs of contexts $U, U' \in A^{\leq k}$, the compatibility condition $(r_U)|_{U \cap U'} = (r_{U'})|_{U \cap U'}$ yields a linear equation in the α_s variables for each $s' \in \mathcal{S}(U \cap U')$, leading to a total number of equations bounded by $\mathcal{O}(|A|^{2k} \cdot |B|^k)$. By an algorithm of Kannan and Bachem [26] can be solved in polynomial time in the sizes of A and B . This allows us to define the following efficient algorithms for CSP and SI based on removing cohomological obstructions.

4.3 Cohomological algorithms for CSP and SI

We saw in Section 3 that the k -consistency and k -Weisfeiler-Leman algorithms can be recovered as greatest fixpoints of presheaf operators removing local sections which fail the forth and bijective-forth properties respectively. Now that we have from cohomological considerations a new necessary condition $\mathbf{Zext}(\mathcal{S}, s)$ for a local section to feature in a global section of \mathcal{S} , we can define natural extensions to the k -consistency and k -Weisfeiler-Leman algorithms as follows.

4.3.1 Cohomological k -consistency

To define the cohomological k -consistency algorithm, we first define an operator which removes those local sections which admit a cohomological obstruction. Let $(\cdot)^{\mathbb{Z}\downarrow}$ be the operator which computes for a given presheaf $\mathcal{S} \subset \mathcal{H}_k$ the subpresheaf $\mathcal{S}^{\mathbb{Z}\downarrow}$ where $\mathcal{S}^{\mathbb{Z}\downarrow}(C)$ contains exactly those local sections $s \in \mathcal{S}(C)$ which satisfy both the forth property $\mathbf{Forth}(\mathcal{S}, s)$ and the \mathbb{Z} -extendability property $\mathbb{Z}\mathbf{ext}(\mathcal{S}, s)$. As this process may remove the local sections in \mathcal{S} which witness the extendability of other local sections we need to take a fixpoint of this operator to get a presheaf with the right extendability properties at every local section. So, we write $\overline{\mathcal{S}}^{\mathbb{Z}}$ for the greatest fixpoint of this operator starting from \mathcal{S} . As both $\mathbf{Forth}(\mathcal{S}, s)$ and $\mathbb{Z}\mathbf{ext}(\mathcal{S}, s)$ are both computable in polynomial time in the size of \mathcal{S} and $\overline{\mathcal{S}}^{\mathbb{Z}}$ has a global section if and only if \mathcal{S} has a global section, this allows us to define the following efficient algorithm for approximating CSP.

► **Definition 4.** *The cohomological k -consistency algorithm accepts an instance (A, B) if the greatest fixpoint $\overline{\mathcal{H}_k(A, B)}^{\mathbb{Z}}$ is non-empty and otherwise rejects. If (A, B) is accepted by this algorithm we write $A \rightarrow_k^{\mathbb{Z}} B$ and say that the instance (A, B) is cohomologically k -consistent.*

We conclude this section by showing that the relation $\rightarrow_k^{\mathbb{Z}}$ is transitive.

► **Proposition 5.** *For all k , given A, B and C structures over a common finite signature*

$$A \rightarrow_k^{\mathbb{Z}} B \rightarrow_k^{\mathbb{Z}} C \implies A \rightarrow_k^{\mathbb{Z}} C.$$

Proof. See full version. ◀

4.3.2 Cohomological k -Weisfeiler-Leman

We now define cohomological k -equivalence to generalise k -WL-equivalence in the same way as we did for cohomological k -consistency, by removing local sections which are not \mathbb{Z} -extendable. As \mathbb{Z} -extendability in $S \subset \mathbf{Isom}_k(A, B)$ is not *a priori* symmetric in A and B we need to check that both s is \mathbb{Z} -extendable in S and s^{-1} is \mathbb{Z} -extendable in $S^{-1} = \{t^{-1} \mid t \in S\}$. We call this s being \mathbb{Z} -bi-extendable in S and write it as $\mathbb{Z}\mathbf{bext}(\mathcal{S}, s)$. We incorporate this into a new presheaf operator $(\cdot)^{\mathbb{Z}\#}$ as follows. Given a presheaf $\mathcal{S} \subset \mathcal{I}_k$ let $\mathcal{S}^{\mathbb{Z}\#}$ be the largest subpresheaf of \mathcal{S} such that every $s \in \mathcal{S}^{\mathbb{Z}\#}(C)$ satisfies both the bijective forth property $\mathbf{BijForth}(\mathcal{S}, s)$ and the \mathbb{Z} -bi-extendability property $\mathbb{Z}\mathbf{bext}(\mathcal{S}, s)$. We write $\overline{\mathcal{S}}^{\mathbb{Z}\#}$ for the greatest fixpoint of this operator starting from \mathcal{S} . As both $\mathbf{BijForth}(\mathcal{S}, s)$ and $\mathbb{Z}\mathbf{bext}(\mathcal{S}, s)$ are computable in polynomial time in the size of \mathcal{S} and $\overline{\mathcal{S}}^{\mathbb{Z}\#}$ has a global section if and only if \mathcal{S} has a global section, this allows us to define the following efficient algorithm for approximating SI.

► **Definition 6.** *The cohomological k -Weisfeiler-Leman accepts an instance (A, B) if the greatest fixpoint $\overline{\overline{\mathcal{I}_k(A, B)}}^{\mathbb{Z}\#}$ is non-empty and otherwise rejects. If (A, B) is accepted by this algorithm we write $A \equiv_k^{\mathbb{Z}} B$ and say that the instance (A, B) is cohomologically k -equivalent.*

Finally, we observe that the existence of a non-empty subpresheaf of \mathcal{I}_k satisfying the $\mathbf{BijForth}$ and $\mathbb{Z}\mathbf{bext}$ properties also satisfies the conditions for witnessing cohomological k -consistency of the pairs (A, B) and (B, A) . Formally we have

► **Observation 7.** *For any two structures A and B , $A \equiv_k^{\mathbb{Z}} B$ implies that $A \rightarrow_k^{\mathbb{Z}} B$ and $B \rightarrow_k^{\mathbb{Z}} A$.*

In Section 5, we will demonstrate the power of these new algorithms by showing that both cohomological k -consistency and cohomological k -Weisfeiler-Leman can solve instances of CSP and SI on which the non-cohomological versions fail. Before doing this, we briefly review some other algorithms for CSP and SI which involve solving systems of linear equations and establish a possible connection to be explored in future work.

4.4 Other algorithms for CSP and SI

While the connections to cohomology in approximating CSP and SI are novel in this paper, the algorithms introduced here are not the first to use solving systems of linear equations to approximate these problems.

On the CSP side, some examples of such algorithms include the BLP+AIP [11] and CLAP [14] algorithms studied in the Promise CSP community. One difference here is that for an instance (A, B) the variables in BLP and AIP are indexed by valid assignments to each variable and to each related tuple instead of being indexed by valid k -local homomorphisms as in the algorithm derived above. This means that directly comparing these algorithms as stated is not straightforward and is beyond the scope of this paper. However, it seems likely that these algorithms can also be expressed in terms of appropriate presheaves. For example, let $\mathbf{C}(A)$ be the category whose objects are the elements of A and the related tuples of A and with maps for each projection from a related tuple to an element, and let the **Set**-valued presheaf $\mathcal{H}_C(A, B)$ on $\mathbf{C}(A)$ map any $a \in A$ to the set of all elements in B and any $\mathbf{a} \in R^A$ to the set of all related tuples R^B . Then, in a similar way to above, we can see that global sections of \mathcal{H}_C are homomorphisms from A to B . In future work, we will compare the fixpoints $\overline{\mathcal{H}_C}$ and $\overline{\mathcal{H}_C}^{\mathbb{Z}}$ with solutions to the BLP and AIP systems of equations and we will explore a possible presheaf representation for CLAP.

On the SI side, Berkholz and Grohe [9] have studied \mathbb{Z} -linear versions of the Sherali-Adams hierarchy of relaxations of the graph isomorphism problem. They establish that no level of this hierarchy decides the full isomorphism relation on graphs. Their algorithm for the k^{th} level of the hierarchy appears similar to checking the \mathbb{Z} -extendability in $\mathcal{H}_k^{A,B}$ of the empty solution $\epsilon \in \mathcal{H}_k^{A,B}(\emptyset)$. A full comparison of this algorithm and the algorithm described above is an interesting direction for future work.

5 The (unreasonable) effectiveness of cohomology in CSP and SI

In this section, we prove that the new algorithms arising from this cohomological approach to CSP and SI are substantially more powerful than the k -consistency and k -Weisfeiler-Leman algorithms. In particular, we show that cohomological k -consistency resolves CSP over all domains of arity less than or equal to k which admit a ring representation and that for a fixed small k cohomological k -Weisfeiler-Leman can distinguish structures which differ on a very general form of the CFI property, in particular, showing that cohomological k -Weisfeiler-Leman can distinguish a property which Lichter [30] claims not to be expressible in rank logic.

5.1 Cohomological k -consistency solves all affine CSPs

In this section, we demonstrate the power of the cohomological k -consistency algorithm by proving that it can decide the solvability of systems of equations over finite rings.

To express the main theorem of this section in terms of the finite relational structures on which our algorithm is defined, we first need to fix a notion of ring representation of a relational structure. Let A be a relational structure over signature σ with relations given by $\{R^A\}_{R \in \sigma}$. We say that A has a *ring representation* if we can give the set A a ring structure $(A, +, \cdot, 0, 1)$ such that for every relational symbol $R \in \sigma$ the set $R^A \subset A^m$ is an affine subset of the ring $(A, +, \cdot, 0, 1)$, meaning that there exists $b_1^R, \dots, b_m^R, a^R \in A$ such that

$$R^A = \{\mathbf{x} \in A^m \mid \sum_{i \in [m]} b_i^R \cdot x_i = a^R\}$$

With this necessary background we state the main theorem of this section.

► **Theorem 8.** *For any structure B with a ring representation, there is a k such that the cohomological k -consistency algorithm decides $\mathbf{CSP}(B)$.*

Alternatively stated, there exists a k such that for all σ -structures A

$$A \xrightarrow{\mathbb{Z}}_k B \iff A \rightarrow B$$

Proof. See full version. ◀

This theorem is notable because there are relational structures B with ring representations for which there are families of structures A_k such that $A_k \rightarrow_k B$ but $A_k \not\rightarrow B$, see for example the examples given by Feder and Vardi [18]. Furthermore, there exist pairs (A_k, B_k) where $A_k \equiv_k B_k$, $B_k \rightarrow B$ and $A_k \rightarrow_k B$ but $A_k \not\rightarrow B$, see for example the work of Atserias, Bulatov and Dawar [6]. As the sequence of relations \equiv_k bounds the expressive power of FPC, this effectively proves that the solvability of systems of linear equations over \mathbb{Z} , which is central to the cohomological k -consistency algorithm, is not expressible in FPC. This result was not previously known to the author.

5.2 Cohomological k -Weisfeiler-Leman decides the CFI property

The Cai-Fürer-Immerman construction [13] on ordered finite graphs is a very powerful tool for proving expressiveness lower bounds in descriptive complexity theory. While it was originally used to separate the infinitary k variable logic with counting from PTIME, it has since been used in adapted forms to prove bounds on invertible maps equivalence [15], computation on Turing machines with atoms [10] and rank logic [30]. In this section, we show that $\equiv_k^{\mathbb{Z}}$ separates a very general form of this

The version we consider in this paper is parameterised by a prime power q and takes any totally ordered graph $(G, <)$ and any map $g: E(G) \rightarrow \mathbb{Z}_q$ to a relational structure $\mathbf{CFI}_q(G, g)$. The construction effectively encodes a system of linear equations over \mathbb{Z}_q based on the edges of G and the “twists” introduced by the labels g . The result is the following well-known fact.

► **Fact 9.** *For any prime power, q , ordered graph G , and functions $g, h: E(G) \rightarrow \mathbb{Z}_q$,*

$$\mathbf{CFI}_q(G, g) \cong \mathbf{CFI}_q(G, h) \iff \sum g = \sum h$$

We say that the structure $\mathbf{CFI}_q(G, g)$ has the *CFI property* if $\sum g = 0$. For more details on this construction we refer to the recent paper of Lichter [30] whose presentation we follow in the full version of this paper.

We now recall the two major separation results based on this construction. The first is a landmark result of descriptive complexity from the early 1990’s.

► **Theorem 10** (Cai, Fürer, Immerman [13]). *There is a class of ordered (β -regular) graphs $\mathcal{G} = \{G_n\}_{n \in \mathbb{N}}$ such that in the respective class of CFI structures*

$$\mathcal{K} = \{\mathbf{CFI}_2(G, g) \mid G \in \mathcal{G}, g: V(G) \rightarrow \mathbb{Z}_2\}$$

the CFI property is decidable in polynomial-time but cannot be expressed in FPC.

The second is a recent breakthrough due to Moritz Lichter.

► **Theorem 11** (Lichter [30]). *There is a class of ordered graphs $\mathcal{G} = \{G_n\}_{n \in \mathbb{N}}$ such that in the respective class of CFI structures*

$$\mathcal{K} = \{\mathbf{CFI}_{2^k}(G, g) \mid G \in \mathcal{G}\}$$

the CFI property is decidable in polynomial-time (indeed, expressible in choiceless polynomial time) but cannot be expressed in rank logic.

Despite this CFI property proving to be inexpressible in both FPC and rank logic, we show that (perhaps surprisingly) there is a fixed k such that cohomological k -Weisfeiler-Leman algorithm can separate structures which differ on this property in the following general way. The proof of this theorem relies on showing that $\equiv_k^{\mathbb{Z}}$ behaves well with logical interpretations and the details are left to the full version of this paper.

► **Theorem 12.** *There is a fixed k such that for any q given $\mathbf{CFI}_q(G, g)$ and $\mathbf{CFI}_q(G, h)$ with $\sum g = 0$ we have*

$$\mathbf{CFI}_q(G, g) \equiv_k^{\mathbb{Z}} \mathbf{CFI}_q(G, h) \iff \mathbf{CFI}_q(G, g) \cong \mathbf{CFI}_q(G, h)$$

Proof. See full version. ◀

As a direct consequence of this result, there is some k such that the set of structures with the CFI property in Lichter's class \mathcal{K} from Theorem 11 is closed under $\equiv_k^{\mathbb{Z}}$. This means that, by the conclusion of Theorem 11, the equivalence relation $\equiv_k^{\mathbb{Z}}$ can distinguish structures which disagree on a property that is not expressible in rank logic. Indeed, Dawar, Grädel and Lichter [16] show further that this property is also inexpressible in linear algebraic logic. By the definition of our algorithm for $\equiv_k^{\mathbb{Z}}$ this implies that solvability of systems of \mathbb{Z} -linear equations is not definable in linear algebraic logic.

6 Conclusions & future work

In this paper, we have presented novel approach to CSP and SI in terms of presheaves and have used this to derive efficient generalisations of the k -consistency and k -Weisfeiler-Leman algorithms, based on natural considerations of presheaf cohomology. We have shown that the relations, $\rightarrow_k^{\mathbb{Z}}$ and $\equiv_k^{\mathbb{Z}}$, computed by these new algorithms are strict refinements of their well-studied classical counterparts \rightarrow_k and \equiv_k . In particular, we have shown in Theorem 8 that cohomological k -consistency suffices to solve linear equations over all finite rings and in Theorem 12 that cohomological k -Weisfeiler-Leman distinguishes positive and negative instances of the CFI property on the classes of structures studied by Cai, Fürer and Immerman [13] and more recently by Lichter [30]. These results have important consequences for descriptive complexity theory showing, in particular, that the solvability of systems of linear equations over \mathbb{Z} is not expressible in FPC, rank logic or linear algebraic logic. Furthermore, the results of this paper demonstrate the unexpected effectiveness of a cohomological approach to constraint satisfaction and structure isomorphism, analogous to that pioneered by Abramsky and others for the study of quantum contextuality.

The results of this paper suggest several directions for future work to establish the extent and limits of this cohomological approach. We ask the following questions which connect it to important themes in algorithms, logic and finite model theory.

Cohomology and constraint satisfaction. Firstly, Bulatov and Zhuk’s recent independent resolutions of the Feder-Vardi conjecture [12, 35], show that for all domains B either $\mathbf{CSP}(B)$ is NP-Complete or B admits a weak near-unanimity polymorphism and $\mathbf{CSP}(B)$ is tractable. As the cohomological k -consistency algorithm expands the power of the k -consistency algorithm which features as one case of Bulatov and Zhuk’s general efficient algorithms, we ask if it is sufficient to decide all tractable \mathbf{CSP} s.

► **Question 13.** *For all domains B which admit a weak near-unanimity polymorphism, does there exist a k such that for all A*

$$A \rightarrow B \iff A \xrightarrow{\mathbb{Z}}_k B?$$

Cohomology and structure isomorphism. Secondly, as cohomological k -Weisfeiler-Leman is an efficient algorithm for distinguishing some non-isomorphic relational structures we ask if it distinguishes all non-isomorphic structures. As the best known structure isomorphism algorithm is quasi-polynomial [8], we do not expect a positive answer to this question but expect that negative answers would aid our understanding of the hard cases of structure isomorphism in general.

► **Question 14.** *For every signature σ does there exist a k such that for all σ -structures A, B*

$$A \cong B \iff A \equiv_k^{\mathbb{Z}} B?$$

Cohomology and game comonads. Thirdly, as \rightarrow_k and \equiv_k have been shown by Abramsky, Dawar, and Wang [4] to correspond to the coKleisli morphisms and isomorphisms of a comonad \mathbb{P}_k , we ask whether a similar account can be given to $\xrightarrow{\mathbb{Z}}_k$ and $\equiv_k^{\mathbb{Z}}$. As the coalgebras of the \mathbb{P}_k comonad relate to the combinatorial notion of treewidth, an answer to this question could provide a new notion of “cohomological” treewidth.

► **Question 15.** *Does there exist a comonad \mathbb{C}_k for which the notion of morphism and isomorphism in the coKleisli category are $\xrightarrow{\mathbb{Z}}_k$ and $\equiv_k^{\mathbb{Z}}$?*

The search for a logic for PTIME. Finally, as the algorithms for $\xrightarrow{\mathbb{Z}}_k$ and $\equiv_k^{\mathbb{Z}}$ are likely expressible in rank logic extended with a quantifier for solving systems of linear equations over \mathbb{Z} and as $\equiv_k^{\mathbb{Z}}$ distinguishes all the best known family separating rank logic from PTIME, we ask if solving systems of equations over \mathbb{Z} is enough to capture all PTIME queries.

► **Question 16.** *Is there a logic $FPC+rk+\mathbb{Z}$ incorporating solvability of \mathbb{Z} -linear equations into rank logic which captures PTIME?*

References

- 1 Samson Abramsky. Notes on presheaf representations of strategies and cohomological refinements of k -consistency and k -equivalence, 2022. doi:10.48550/ARXIV.2206.12156.

- 2 Samson Abramsky, Rui Soares Barbosa, Kohei Kishida, Raymond Lal, and Shane Mansfield. Contextuality, Cohomology and Paradox. In Stephan Kreutzer, editor, *24th EACSL Annual Conference on Computer Science Logic (CSL 2015)*, volume 41 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 211–228, Dagstuhl, Germany, 2015. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.CSL.2015.211.
- 3 Samson Abramsky and Adam Brandenburger. The sheaf-theoretic structure of non-locality and contextuality. *New Journal of Physics*, 13(11):113036, November 2011. doi:10.1088/1367-2630/13/11/113036.
- 4 Samson Abramsky, Anuj Dawar, and Pengming Wang. The pebbling comonad in finite model theory. In *Proceedings of the 32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '17*. IEEE Press, 2017.
- 5 Samson Abramsky, Shane Mansfield, and Rui Soares Barbosa. The cohomology of non-locality and contextuality. In B. Jacobs, P. Selinger, and B. Spitters, editors, *Proceedings of 8th International Workshop on Quantum Physics and Logic (QPL 2011)*, volume 95, pages 1–14, 2011. doi:10.4204/EPTCS.95.1.
- 6 Albert Atserias, Andrei Bulatov, and Anuj Dawar. Affine systems of equations and counting infinitary logic. In *In ICALP'07, volume 4596 of LNCS*, pages 558–570, 2007.
- 7 Albert Atserias, Andrei A. Bulatov, and Víctor Dalmau. On the power of k -consistency. In Lars Arge, Christian Cachin, Tomasz Jurdzinski, and Andrzej Tarlecki, editors, *Automata, Languages and Programming, 34th International Colloquium, ICALP 2007, Wroclaw, Poland, July 9-13, 2007, Proceedings*, volume 4596 of *Lecture Notes in Computer Science*, pages 279–290. Springer, 2007. doi:10.1007/978-3-540-73420-8_26.
- 8 László Babai. Graph isomorphism in quasipolynomial time [extended abstract]. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing, STOC '16*, pages 684–697, New York, NY, USA, 2016. Association for Computing Machinery. doi:10.1145/2897518.2897542.
- 9 Christoph Berkholz and Martin Grohe. Linear diophantine equations, group CSPs, and graph isomorphism. In Philip N. Klein, editor, *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 327–339. SIAM, 2017. doi:10.1137/1.9781611974782.21.
- 10 Mikolaj Bojanczyk, Bartek Klin, Slawomir Lasota, and Szymon Torunczyk. Turing machines with atoms. In *2013 28th Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 183–192, 2013. doi:10.1109/LICS.2013.24.
- 11 Joshua Brakensiek, Venkatesan Guruswami, Marcin Wrochna, and Stanislav Živný. The power of the combined basic linear programming and affine relaxation for promise constraint satisfaction problems. *SIAM Journal on Computing*, 49(6):1232–1248, 2020. doi:10.1137/20M1312745.
- 12 Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 319–330, 2017. doi:10.1109/FOCS.2017.37.
- 13 Jin-Yi Cai, Martin Fürer, and Neil Immerman. An optimal lower bound on the number of variables for graph identification. *Combinatorica*, 12(4):389–410, December 1992. doi:10.1007/BF01305232.
- 14 Lorenzo Ciardo and Stanislav Živný. CLAP: A new algorithm for promise CSPs. In Joseph (Seffi) Naor and Niv Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022*, pages 1057–1068. SIAM, 2022. doi:10.1137/1.9781611977073.46.
- 15 Anuj Dawar, Erich Grädel, and Wied Pakusa. Approximations of Isomorphism and Logics with Linear-Algebraic Operators. In Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi, editors, *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, volume 132 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 112:1–112:14, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.ICALP.2019.112.

- 16 Anuj Dawar, Erich Grädel, and Moritz Lichter. Limitations of the invertible-map equivalences, 2021. [arXiv:2109.07218](https://arxiv.org/abs/2109.07218).
- 17 Anuj Dawar and Bjarki Holm. Pebble games with algebraic rules. In Artur Czumaj, Kurt Mehlhorn, Andrew Pitts, and Roger Wattenhofer, editors, *Automata, Languages, and Programming*, pages 251–262, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- 18 Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through Datalog and group theory. *SIAM Journal on Computing*, 28(1):57–104, 1998. doi:10.1137/S0097539794266766.
- 19 Marcelo Fiore and Sam Staton. Comparing operational models of name-passing process calculi. *Information and Computation*, 204(4):524–560, 2006. Seventh Workshop on Coalgebraic Methods in Computer Science 2004. doi:10.1016/j.ic.2005.08.004.
- 20 Jörg Flum and Martin Grohe. On fixed-point logic with counting. *The Journal of Symbolic Logic*, 65(2):777–787, 2000. URL: <http://www.jstor.org/stable/2586569>.
- 21 Roger Godement. Topologie algébrique et théorie des faisceaux. *The Mathematical Gazette*, 44:69, 1960.
- 22 Alexander Grothendieck. Sur quelques points d’algèbre homologique, I. *Tohoku Mathematical Journal*, 9(2):119–221, 1957. doi:10.2748/tmj/1178244839.
- 23 Yuri Gurevich. Logic and the challenge of computer science, 1988.
- 24 Lauri Hella. Logical hierarchies in PTIME. *Information and Computation*, 129(1):1–19, 1996. doi:10.1006/inco.1996.0070.
- 25 Neil Immerman and Eric Lander. *Describing Graphs: A First-Order Approach to Graph Canonization*, pages 59–81. Springer New York, New York, NY, 1990. doi:10.1007/978-1-4612-4478-3_5.
- 26 Ravindran Kannan and Achim Bachem. Polynomial algorithms for computing the smith and hermite normal forms of an integer matrix. *SIAM Journal on Computing*, 8(4):499–9, November 1979. Copyright - Copyright] © 1979 © Society for Industrial and Applied Mathematics; Last updated - 2021-09-11. URL: <https://ezp.lib.cam.ac.uk/login?url=https://www.proquest.com/scholarly-journals/polynomial-algorithms-computing-smith-hermite/docview/918468506/se-2?accountid=9851>.
- 27 Sandra Kiefer. *Power and limits of the Weisfeiler-Leman algorithm*. Dissertation, RWTH Aachen University, Aachen, 2020. Veröffentlicht auf dem Publikationsserver der RWTH Aachen University; Dissertation, RWTH Aachen University, 2020. doi:10.18154/RWTH-2020-03508.
- 28 P.G. Kolaitis and M.Y. Vardi. On the expressive power of Datalog: Tools and a case study. *Journal of Computer and System Sciences*, 51(1):110–134, 1995. doi:10.1006/jcss.1995.1055.
- 29 Tom Leinster. *Categories, functors and natural transformations*, pages 9–40. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2014. doi:10.1017/CB09781107360068.003.
- 30 Moritz Lichter. Separating rank logic from polynomial time. In *36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, Rome, Italy, June 29 - July 2, 2021*, pages 1–13. IEEE, 2021. doi:10.1109/LICS52264.2021.9470598.
- 31 Saunders Mac Lane and Ieke Moerdijk. *Sheaves of Sets*, pages 64–105. Springer New York, New York, NY, 1994. doi:10.1007/978-1-4612-0927-0_4.
- 32 Cristina Matache, Sean K. Moss, and Sam Staton. Recursion and sequentiality in categories of sheaves. In Naoki Kobayashi, editor, *6th International Conference on Formal Structures for Computation and Deduction, FSCD 2021, July 17-24, 2021, Buenos Aires, Argentina (Virtual Conference)*, volume 195 of *LIPICs*, pages 25:1–25:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.FSCD.2021.25.
- 33 Baudot P. and Bennequin D. The homological nature of entropy. *Entropy*, 17(5):3253–3318, 2015. Cited by: 27; All Open Access, Gold Open Access, Green Open Access. doi:10.3390/e17053253.

75:16 Cohomology in Constraint Satisfaction and Structure Isomorphism

- 34 Wied Pakusa. *Linear Equation Systems and the Search for a Logical Characterisation of Polynomial Time*. PhD thesis, RWTH Aachen University, 2016. URL: <https://logic.rwth-aachen.de/~pakusa/diss.pdf>.
- 35 Dmitry Zhuk. A proof of the CSP dichotomy conjecture. *J. ACM*, 67(5), August 2020. doi:10.1145/3402029.