Ineffectiveness for Search and Undecidability of PCSP Meta-Problems

Alberto Larrauri [*]

University of Zaragoza

Abstract

It is an open question whether the search and decision versions of promise CSPs are equivalent. Most known algorithms for PCSPs solve only their *decision* variant, and it is unknown whether they can be adapted to solve *search* as well. The main approaches, called BLP, AIP and BLP + AIP, handle a PCSP by finding a solution to a relaxation of some integer program. We prove that rounding those solutions to a proper search certificate can be as hard as any problem in the class TFNP. In other words, these algorithms are ineffective for search. Building on the algebraic approach to PCSPs, we find sufficient conditions that imply ineffectiveness for search. Our tools are tailored to algorithms that are characterized by minions in a suitable way, and can also be used to prove undecidability results for meta-problems. This way, we show that the families of templates solvable via BLP, AIP, and BLP + AIP are undecidable.

Using the same techniques we also analyze several algebraic conditions that are known to guarantee the tractability of finite-template CSPs. We prove that several meta-problems related to cyclic polymorphims and WNUs are undecidable for PCSPs. In particular, there is no algorithm deciding whether a finite PCSP template (1) admits cyclic a polymorphism, (2) admits a WNU.

---

[*] This work was supported by the UKRI grant EP/X024431/1, and the UK EPSRC grant EP/X033201/1. The research leading to these results was conducted at the University of Oxford between 2024 and 2025, and continued during the summer of 2025 at Durham University.

Contents

1 Introduction 2

2 Preliminaries 4

3 Main Results 8

4 Overview of the Proofs 9
4.1 Organization of the Paper 14

5 Sources of Undecidability, Non-Computability and Hardness 15

6 Main Definitions 17
6.1 Interpretations over Minions 17
6.2 Manifold Minions 18
6.3 Descriptions 18
6.4 Patterns 21

7 From Minions to Templates 23

8 Main Reductions 27

9 Proof of the Main Results 31
9.1 The AIP Algorithm 32
9.1.1 AIP: Interpreting the Super-Grid 32
9.1.2 AIP: Interpreting the Grid 35
9.2 The BLP Algorithm 38
9.2.1 BLP: Interpreting the Super-Grid 38
9.2.2 BLP: Interpreting the Grid 42
9.3 The BLP + AIP Algorithm 44
9.3.1 BLP + AIP: Interpreting the Super-Grid 45
9.4 Weak Near-Unanimity Polymorphisms 48
9.4.1 WNUs: Interpreting the Grid 49
9.4.2 WNUs: Interpreting Triangular Slices 53
9.5 Cyclic Polymorphisms 53
9.5.1 Cyclic Polymorphisms: Interpreting the Grid 54
9.5.2 Cyclic Polymorphisms: Interpreting Triangular Slices 58

10 Discussion 59

A Proofs of Section 5 67
A.1 Proof of Proposition 4.2-(2) 67
A.2 Proof of Proposition 4.2-(3) 68
A.3 Proof of Proposition 5.3 68
A.4 Hardness Proofs 69
A.4.1 Proof of Proposition 4.2-(1) 70
A.4.2 Proof of Proposition 5.2-(1) 70

B Characterizing Polymorphism Minions up to Isomorphism 72

1 Introduction

The Dichotomy Theorem [25,65] and, more generally, the algebraic approach to constraint satisfaction [15,49] show that finite-template constraint satisfaction problems (CSPs) form a particularly well-behaved class of NP problems. This class includes a wide range of natural problems relevant across many domains, such as variants of the Boolean satisfiability problem, graph coloring problems, and systems of equations over finite algebraic structures. Unconditionally, the search and decision variant of each finite-template CSPs are polynomial-time equivalent. Each problem in this class is either NP-complete or in P, and there is an explicit procedure that correctly assigns one of these two cases to each given CSP. This stands in contrast to the well-known fact that there are NP-intermediate problems if P≠ NP [52], and there are problems in FNP (the search analogue of NP) which we believe to be hard but whose decision variants are trivial, captured by the TFNP class [40]. Therefore, it is natural to ask how far (and in which directions) can finite-template CSPs be generalized while still keeping their nice properties.

Promise CSPs (PCSPs) are qualitative relaxations of CSPs generalizing the task of coloring a c-colorable graph using d colors for fixed integers $d \geqslant c$, called the *approximate graph coloring* (AGC) problem. A PCSP is given by a pair of relational structures $(\boldsymbol{A}, \boldsymbol{B})$, the *template*, where the first maps homomorphically into the second, denoted $\boldsymbol{A} \to \boldsymbol{B}$. In the decision variant of PCSP$(\boldsymbol{A}, \boldsymbol{B})$ the task is to distinguish input structures $\boldsymbol{I}$ satisfying $\boldsymbol{I} \to \boldsymbol{A}$ from those satisfying $\boldsymbol{I} \not\to \boldsymbol{B}$. In the search variant, the goal is to find an explicit homomorphism from $\boldsymbol{I}$ to $\boldsymbol{B}$ given the promise that $\boldsymbol{I} \to \boldsymbol{A}$. The algebraic approach, instrumental in much of the CSP theory and both proofs of the Dichotomy Theorem, has recently been extended to the PCSP framework [11]. Roughly, this approach studies the complexity of PCSP$(\boldsymbol{A}, \boldsymbol{B})$ by analyzing the algebraic properties of the set Pol$(\boldsymbol{A}, \boldsymbol{B})$ of homomorphisms $f : \boldsymbol{A}^n \to \boldsymbol{B}$ from some direct power of $\boldsymbol{A}$ to $\boldsymbol{B}$, called *polymorphisms*. This extension has motivated a surge of activity in the area that has yielded partial complexity classifications [10,18,19,36,54,58,59], efficient algorithms [11,21,27,30], and hardness conditions [7,14,22,51]. However, many basic questions remain unanswered. For example, the complexity classifications of Boolean and graph PCSPs (including the complexity of AGC) are still open, despite relevant progress in those directions [37,51]. Crucially, the relationship between search and decision PCSPs is not well understood, and it is unknown whether there is always an efficient way of solving a finite-template search PCSP whose decision variant is tractable.

In this work we analyze several PCSP algorithms, as well as some algebraic conditions that guarantee tractability in the CSP setting. It is known that, unless P=NP, there are tractable finite-template PCSPs which cannot be reduced via gadgets to a tractable finite-template CSP, as exemplified by "1-in-3-SAT vs not-all-equal-SAT" [2,9,57,61]. Therefore, there is a need to develop algorithms that go beyond the finite CSP case. Solving the decision variant of PCSP$(\boldsymbol{A}, \boldsymbol{B})$ involves solving a tractable relaxation of CSP$(\boldsymbol{A})$ which should be a restriction of CSP$(\boldsymbol{B})$. The main relaxations used for this purpose are direct relaxations of the basic integer programming formulation of CSP$(\boldsymbol{A})$ [11, 20, 21], the local consistency algorithm [4, 12], and algorithmic hierarchies built on top of these two previous approaches [27, 28, 30]. Virtually all these algorithms present two inherent limitations. The first is that they only solve the decision version of PCSP$(\boldsymbol{A}, \boldsymbol{B})$: when an instance $\boldsymbol{I}$ is accepted by the relaxation, we do not directly obtain a homomorphism $\boldsymbol{I} \to \boldsymbol{B}$. To obtain such a homomorphism we must *round* the solution to the relaxation, and there is no obvious efficient way of doing so except for in a few known cases [17,18]. The second limitation of these algorithms is that, despite some of them admitting nice algebraic characterizations, we do not know of a way to recognize the problems $\mathrm{PCSP}(\boldsymbol{A}, \boldsymbol{B})$ that they solve. This is known as the *meta-problem* related to these algorithms.

In our analysis of algorithms we focus on the *affine integer programming* (AIP) and the *basic linear programming* (BLP) relaxations for PCSPs, as well as on the BLP + AIP relaxation, which combines the power of the previous two. We present results of three kinds: hardness, undecidability, and non-computability. We show that rounding the output of these algorithms to search certificates can be as hard as any problem in the *total* FNP (TFNP) class, and that all these algorithms have undecidable meta-problems.

Theorem 1.1 (Main algorithmic result, informal)**.** *Let $\mathcal{Q} \in \{AIP, BLP, BLP + AIP\}$, and let $\mathcal{S}^{\mathcal{Q}}$ be the family of finite templates $(\boldsymbol{A}, \boldsymbol{B})$ such that $\mathcal{Q}$ solves* $\mathrm{PCSP}(\boldsymbol{A}, \boldsymbol{B})$*. Then, given any problem Λ in the TFNP class, there is a finite template $(\boldsymbol{A}_\Lambda, \boldsymbol{B}_\Lambda) \in \mathcal{S}^{\mathcal{Q}}$ such that Λ is many-one reducible in polynomial time to the problem of finding a homomorphism $\boldsymbol{I} \to \boldsymbol{B}_\Lambda$ for an input structure $\boldsymbol{I}$ accepted by $\mathcal{Q}$. Furthermore, $\mathcal{S}^{\mathcal{Q}}$ is undecidable.*

This is the first *hardness-of-rounding* result in the PCSPs literature. If TFNP contains some problem that has no polynomial-time solution (which is widely conjectured [45, 55]), then this gives a negative answer to the question of whether the output of these algorithms can always be rounded to a search solution efficiently, posed in [50] [1]. Furthermore, for AIP and BLP we show that there are finite-template PCSPs solved by those algorithms where no computable function is a valid rounding rule. This remains true even for *left-Boolean* templates $(\boldsymbol{A}, \boldsymbol{B})$ (i.e., those where $\boldsymbol{A}$'s domain contains only two elements). This means that directly rounding the output of these algorithms through consistent families of well-behaved polymorphisms such as the ones defined in [17] is not always possible. However, changing the domain of our relaxations to other *efficiently computable rings*, as proposed in [17], could be a way to bypass the issues presented in this paper. This is expanded upon in the discussion section (Section 10). We also point out that while our results include small bounds on $\boldsymbol{A}$'s domain size, the domain of the right structure $\boldsymbol{B}$ can grow quite large. Hence, our results do not rule out the possibility that BLP, AIP, and BLP + AIP could always be directly adapted for search in the Boolean setting.

Our techniques leverage the characterizations of the algorithms AIP, BLP, and BLP + AIP by means of objects called *minions* (or minor-closed classes [62]) that lie at the heart of the algebraic method. On a high level, our results follow from reducing tiling problems (e.g., [43]) to *promise minor condition (*PMC*) problems* [11]. We remark that, although the connection between the TFNP class and tiling problems is not difficult to prove, to our knowledge it has not been previously used to obtain hardness results within TFNP. To achieve our reductions we develop a new way of encoding PCSPs in PMC problems by means of a sheaf-like [23] construction on minions where elements represent partial homomorphisms, and minoring operations ensure consistency between the corresponding local homomorphisms.

We also study some well-known algebraic conditions following the same approach as in our analysis of algorithms. A polymorphism $f : \boldsymbol{A}^n \to \boldsymbol{B}$ is called *cyclic* if it is invariant under any cyclic permutation of its arguments, and *weak near-unanimity* (WNU) if the value of $f(x, \ldots, x, y, x, \ldots, x)$ is independent of the position of the lone y. The existence of a WNU and the existence of a cyclic

---

[1] A previous version of this work also claimed the results presented here yielded negative answers to the questions in [17, 21], without further clarification. This was incorrect and misleading.

polymorphism are equivalent conditions that characterize the tractability of finite-template CSPs if P≠ NP [25, 65]. There is a rich network of equivalences between algebraic conditions in this setting (e.g., [15, Theorems 41 and 47]) which were established using tools from universal algebra. These equivalences are a powerful aid in proving the decidability of various *meta-questions*. For instance, it is known that a finite structure $\boldsymbol{A}$ admits a cyclic polymorphism if and only if it admits a cyclic polymorphism of every prime arity larger than the domain size of $\boldsymbol{A}$ [13], and hence one can easily decide the existence of a WNU or a cyclic polymorphism. Another important algebraic condition is the presence of WNUs of all arities $k \geqslant 3$, which characterizes bounded width CSPs in the finite-template setting [12, 48]. This condition is equivalent to the presence of WNUs of arities 3 and 4 satisfying a particular relation [48] and is, hence, decidable. The complexity of these and other meta-questions has been examined in [26].

For PCSPs the picture is significantly less structured. We know that, unless P≠ NP, no finite family of polymorphisms can guarantee tractability [11]. Hence, all algebraic tractability conditions must involve an infinite family of polymorphisms. Even then, admitting WNUs of all arities $k \geqslant 3$ is no longer a sufficient condition for bounded width (but remains a necessary one) [4], and does not even guarantee tractability [29]. Since many of the tools from universal algebra no longer apply to the PCSP setting, few non-trivial implications between algebraic conditions have been shown (with some exceptions, e.g. [21, Theorem 4]). In particular, the decidability of most meta-questions related to natural algebraic conditions remains open for PCSPs. Our main algebraic result is the following.

Theorem 1.2. *The following problems are undecidable. Given an input finite template* $(\boldsymbol{A}, \boldsymbol{B})$ *with* $|A| \leqslant 3$, *determine whether* $\mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$ *contains:*

- *(1) cyclic polymorphisms*
  - *(i) for every prime arity p,*
  - *(ii) for at least one arity k,*
  - *(iii) for all but finitely many prime arities p,*
  - *(iv) for infinitely many prime arities p.*
- *(2) weak near-unanimity polymorphisms (WNUs)*
  - *(i) for every arity $k \geqslant 3$,*
  - *(ii) for at least one arity $k \geqslant 3$,*
  - *(iii) for all but finitely many arities $k \geqslant 3$,*
  - *(iv) for infinitely many arities $k \geqslant 3$.*

2 Preliminaries

The set of natural numbers $\mathbb{N}$ starts at 1, and $[k]$ is the set $\{1, 2, \ldots, k\}$. Given sets S, T, the set of maps $f : S \to T$ is denoted by T^S. A *partial function* from S to T is denoted as $f : S \rightharpoonup T$. We write id_X for the identity map on a set X. We identify tuples $\boldsymbol{t} = (t_1, \ldots, t_n) \in T^n$ with functions in $T^{[n]}$ in the natural way (i.e., $\boldsymbol{t}(i) = t_i$). Disjoint unions are denoted using $\sqcup$. We write U^* for the set of finite strings over a finite alphabet U, i.e. $U^* = \bigsqcup_{n \geqslant 0} U^n$. A partial map $f : U^* \rightharpoonup U^*$ is said to be *computable* if there is a Turing machine on an alphabet containing U that computes $f(\boldsymbol{x})$ for any input $\boldsymbol{x} \in U^*$ where f is defined, and runs forever on $\boldsymbol{x} \in U^*$ where f is undefined. We informally say that a partial map $f : S \rightharpoonup T$ between countable sets S, T is computable if f is computable under some implicit encoding. That is, we implicitly refer to a finite alphabet U, and to injective maps $\alpha : S \to U^*$ and $\beta : T \to U^*$, and mean that the partial map $\beta \circ f \circ \alpha^{-1}$ is computable [31].

Search Problems We refer to [60, Section 10.3] for an introduction to search complexity and the classes FP, FNP, and TFNP. Let $\mathfrak{R} \subseteq U^* \times V^*$ be a binary relation on finite words, and $\perp$ a special "reject" symbol outside the alphabets U, V. The *search problem* $\Lambda_{\mathfrak{R}}$, is defined as: given a string $\boldsymbol{x} \in U^*$, find some $\boldsymbol{y} \in V^*$ satisfying $(\boldsymbol{x}, \boldsymbol{y}) \in \mathfrak{R}$, or reject $\boldsymbol{x}$ and return $\perp$ if no such $\boldsymbol{y}$ exists. The class of *total functional NP* (TFNP) problems consists of all problems $\Lambda_{\mathfrak{R}}$ where $\mathfrak{R}$ is (1) a total relation, meaning that for all input strings $\boldsymbol{x}$ there is some $\boldsymbol{y}$ satisfying $(\boldsymbol{x}, \boldsymbol{y}) \in \mathfrak{R}$, (2) polynomially bounded, meaning that there is some polynomial $p(n)$ such that $(\boldsymbol{x}, \boldsymbol{y}) \in \mathfrak{R}$ implies $|\boldsymbol{y}| \leqslant p(|\boldsymbol{x}|)$, and (3) $\mathfrak{R}$ is recognizable in polynomial time. The class of *tally TFNP* problems, denoted TFNP_1, consists of all problems $\Lambda_{\mathfrak{R}} \in \text{TFNP}$ such that $\mathfrak{R} \subseteq U^* \times V^*$ is a relation where U is the unary alphabet $\{1\}$.

The class TFNP has been studied extensively (e.g., [40, 42, 45]) and contains several problems for which no efficient solution is expected such as integer factoring, or the problem of computing a Nash equilibrium [33]. There is compelling reason to believe that TFNP $\not\subseteq$ FP [45] (FP being the class of search problems solvable in polynomial time), which would imply that P $\neq$ NP. On the other hand, there are oracles with respect to which TFNP $\subseteq$ FP but the polynomial hierarchy is infinite [24]. To our knowledge, the class TFNP_1 has not been studied explicitly, but also contains problems that have no known polynomial-time solution, such as PRIME, where 1^n is given as input and the task is to find a prime number on n bits [39].

A (polynomial-time) *many-one reduction* from $\Lambda_{\mathfrak{R}_1}$ to $\Lambda_{\mathfrak{R}_2}$ consists of a pair (α, β) of polynomial-time computable functions satisfying that for all $\boldsymbol{x}$ (1) if $(\alpha(\boldsymbol{x}), \boldsymbol{z}) \in \mathfrak{R}_2$, then $(\boldsymbol{x}, \beta(\boldsymbol{x}, \boldsymbol{z})) \in \mathfrak{R}_1$, and (2) if $(\alpha(\boldsymbol{x}), \boldsymbol{z}) \not\in \mathfrak{R}_2$ for all $\boldsymbol{z}$, then $(\boldsymbol{x}, \boldsymbol{y}) \notin\in \mathfrak{R}_1$ for all $\boldsymbol{y}$. Similarly, a *generalized many-one reduction* from $\Lambda_{\mathfrak{R}_1}$ to $\Lambda_{\mathfrak{R}_2}$ is a pair of polynomial-time computable functions (α, β) satisfying that whenever $\boldsymbol{z}$ is a valid answer to $\alpha(\boldsymbol{x})$ in the problem $\Lambda_{\mathfrak{R}_2}$ (including the case where $\boldsymbol{z} = \perp$), then $\beta(\boldsymbol{x}, \boldsymbol{z})$ is a valid answer to $\boldsymbol{x}$ in the problem $\mathfrak{R}_1$ (including again the case where $\beta(\boldsymbol{x}, \boldsymbol{z}) = \perp$). Given a class $\mathcal{C}$ of search problems, and a set $\mathcal{F}$ of search problems, we say that $\mathcal{F}$ is $\mathcal{C}$-hard if for any $\mathfrak{R} \in \mathcal{C}$ there is some $\mathfrak{R}' \in \mathcal{F}$ such that $\Lambda_{\mathfrak{R}}$ has a many-one reduction to $\Lambda_{\mathfrak{R}'}$. A problem $\Lambda_{\mathfrak{R}}$ is called $\mathcal{C}$-hard if the family $\{\Lambda_{\mathfrak{R}}\}$ is $\mathcal{C}$-hard. The class TFNP is conjectured to contain no TFNP-hard problems [55], so our results show TFNP-hardness of families.

We also consider *promise* search problems, which can be seen as "partial" search problems. The notions in this section extend to promise problems in the natural way, along the lines of e.g., [35, 41].

Promise Constraint Satisfaction A *relational signature* Σ is a finite set of symbols where each $R \in \Sigma$ has some *arity* $\text{ar}(R) \in \mathbb{N}$. A Σ-*structure* $\boldsymbol{A}$ consists of: a set A called its *universe*, and a relation $R^A \subseteq A^{\text{ar}(R)}$ for each $R \in \Sigma$. Given two *similar* (i.e., with the same signature) structures $\boldsymbol{A}, \boldsymbol{B}$, a homomorphism $h : \boldsymbol{A} \to \boldsymbol{B}$ is a map from A to B satisfying $(h(e(1)), \ldots, h(e(\text{ar}(R)))) \in R^B$ for each $R \in \Sigma$ and each tuple $e \in R^A$. We write $\boldsymbol{A} \to \boldsymbol{B}$ to denote there is a homomorphism from $\boldsymbol{A}$ to $\boldsymbol{B}$. The *n-th power* of $\boldsymbol{A}$, denoted $\boldsymbol{A}^n$, is a structure similar to $\boldsymbol{A}$ whose universe is A^n, and where R^{A^n} consists of the tuples $(\boldsymbol{a}_1, \ldots, \boldsymbol{a}_{\text{ar}(R)})$, satisfying $(a_{1,j}, \ldots a_{\text{ar}(R),j}) \in R^A$ for all $j \in [n]$.

Templates are pairs of structures $(\mathbf{A}, \mathbf{B})$ satisfying $\mathbf{A} \to \mathbf{B}$. The template is *finite* if both $\mathbf{A}, \mathbf{B}$ are finite. The *decision promise constraint satisfaction problem* (PCSP) defined by a template $(\mathbf{A}, \mathbf{B})$, denoted $\operatorname{PCSP}(\mathbf{A}, \mathbf{B})$, is the problem of, given an input finite structure $\mathbf{I}$, to accept it if $\mathbf{I} \to \mathbf{A}$, and to reject it if $\mathbf{I} \not\to \mathbf{B}$. Similarly, in the *search PCSP* defined by $(\mathbf{A}, \mathbf{B})$, denoted $\operatorname{sPCSP}(\mathbf{A}, \mathbf{B})$, the goal is to find a homomorphism $F : \mathbf{I} \to \mathbf{B}$, if $\mathbf{I} \to \mathbf{A}$, or to reject $\mathbf{I}$ if $\mathbf{I} \not\to \mathbf{B}$. Observe that this only makes sense if $\mathbf{B}$ is finite or a suitable encoding is fixed. We define the problems $\operatorname{CSP}(\mathbf{A})$ and $\operatorname{sCSP}(\mathbf{A})$ as $\operatorname{PCSP}(\mathbf{A}, \mathbf{A})$ and $\operatorname{sPCSP}(\mathbf{A}, \mathbf{A})$ respectively.

Efficient Algorithms and Rounding Problems Let $\mathbf{A}$, $\mathbf{I}$ be finite Σ-structures. The following system of equations over the integers $\{0, 1\}$, denoted $\operatorname{IP}_{\mathbf{A}}(\mathbf{I})$, is satisfiable if and only if $\mathbf{I} \to \mathbf{A}$:

$$\begin{aligned} \text{Variables: } & \{x_{v,a} \mid v \in I, a \in A\} \quad \sqcup \\ & \{x_{r_A, r_I} \mid R \in \Sigma, r_I \in R^I, r_A \in R^A\}. \end{aligned}$$

$$\text{Equations: } \begin{array}{l} \sum_{a \in A} x_{v,a} = 1, \\ \sum_{r_A \in R^A} x_{r_I, r_A} = 1, \\ \sum_{r_A \in R^A, r_A(i) = \alpha} x_{r_I, r_A} = x_{r_I(i), \alpha}, \end{array} \tag{1}$$
$$\text{for each } v \in I,\ R \in \Sigma,\ r_I \in R^I,\ i \in [\operatorname{ar}(R)],\ \alpha \in A.$$

Solving this system is as difficult as solving $\operatorname{CSP}(\mathbf{A})$, but if we allow the variables to take arbitrary values in $\mathbb{Z}$ or in the rational interval $[0, 1]$ then the task can be carried out in polynomial time. We write $\operatorname{AIP}_{\mathbf{A}}(\mathbf{I})$ and $\operatorname{BLP}_{\mathbf{A}}(\mathbf{I})$ for the system $\operatorname{IP}_{\mathbf{A}}(\mathbf{I})$ when the domain of the variables is $\mathbb{Z}$ or $[0, 1] \subseteq \mathbb{Q}$ respectively.

Given an input structure $\mathbf{I}$ for $\operatorname{PCSP}(\mathbf{A}, \mathbf{B})$, where $(\mathbf{A}, \mathbf{B})$ is a finite template, the AIP algorithm (resp., BLP) [11] solves $\operatorname{AIP}_{\mathbf{A}}(\mathbf{I})$ (resp., $\operatorname{BLP}_{\mathbf{A}}(\mathbf{I})$) and accepts $\mathbf{I}$ if and only if this system is satisfiable. The algorithm BLP + AIP [21] combines the power of the previous two, and checks in polynomial time whether $\operatorname{AIP}_{\mathbf{A}}(\mathbf{I})$ and $\operatorname{BLP}_{\mathbf{A}}(\mathbf{I})$ have compatible solutions, and accepts $\mathbf{I}$ when this is the case. Compatibility here means that whenever a variable is assigned to 0 in $\operatorname{BLP}_{\mathbf{A}}(\mathbf{I})$, it is also assigned to 0 in $\operatorname{AIP}_{\mathbf{A}}(\mathbf{I})$. An algorithm $\mathcal{Q} \in \{\mathrm{AIP}, \mathrm{BLP}, \mathrm{BLP} + \mathrm{AIP}\}$ solves $\operatorname{PCSP}(\mathbf{A}, \mathbf{B})$ if it always outputs a correct answer in this problem, meaning that the following two implications hold: (1) if there is a homomorphism $\mathbf{I} \to \mathbf{A}$ then $\mathcal{Q}$ accepts $\mathbf{I}$, and (2) if $\mathcal{Q}$ accepts $\mathbf{I}$ then there is a homomorphism $\mathbf{I} \to \mathbf{B}$. In this situation we define the *rounding problem* $\operatorname{sPCSP}_{\mathcal{Q}}(\mathbf{A}, \mathbf{B})$ as the problem of, given an instance $\mathbf{I}$ accepted by $\mathcal{Q}$, to find a homomorphism from $\mathbf{I}$ to $\mathbf{B}$.

Minions A *minion* $\mathscr{M}$ consists of a collection of disjoint sets $\mathscr{M}(n)$ indexed by the natural numbers $n \in \mathbb{N}$, and a map $\pi^{\mathscr{M}} : \mathscr{M}(n) \to \mathscr{M}(m)$ for each pair of numbers $n, m \in \mathbb{N}$ and each function $\pi \in [m]^{[n]}$, satisfying that (1) $\pi^{\mathscr{M}} = \mathrm{id}_{\mathscr{M}(n)}$ if $\pi = \mathrm{id}_{[n]}$, and (2) $\pi^{\mathscr{M}} = \pi_1^{\mathscr{M}} \circ \pi_2^{\mathscr{M}}$ whenever $\pi = \pi_1 \circ \pi_2$. The elements $f \in \mathscr{M}(n)$ are called *n-ary*, the functions $\pi^{\mathscr{M}}$ are called *minoring operations*, and an element $\pi^{\mathscr{M}}(f)$ is called a *minor* of f. When $\mathscr{M}$ is clear from the context, we write f^{π} instead of $\pi^{\mathscr{M}}(f)$. The minion $\mathscr{M}$ is called *locally finite* if $\mathscr{M}(n)$ is a finite set for all $n \in \mathbb{N}$. An element $p \in \mathscr{M}(n)$ is called *cyclic* if $p = p^{(n,1,2,\ldots,n-1)}$ (we remind the reader that we represent maps $\pi \in [n]^{[n]}$ as tuples). A *weak near-unanimity* (WNU) element is some $p \in \mathscr{M}(n)$ satisfying that $p^{\sigma_i} = p^{\sigma_j}$ for all $i, j \in [n]$, where $\sigma_i \in [2]^{[n]}$ sends i to 1 and all other $o \in [n]$ to 2. Given a template $(\mathbf{A}, \mathbf{B})$, the *polymorphism minion* $\operatorname{Pol}(\mathbf{A}, \mathbf{B})$ is a minion whose n-ary elements are the homomorphisms $f : \boldsymbol{A}^n \to \boldsymbol{B}$, which are called *polymorphisms*. Given an n-ary polymorphism f, and a map $\pi : [n] \to [m]$, the minor f^π is defined by $f^\pi(\boldsymbol{a}) = f(\boldsymbol{a} \circ \pi)$ for every $f \in A^n$. Finally, a *minion homomorphism* $F : \mathscr{M} \to \mathscr{N}$ is a map from elements of $\mathscr{M}$ to elements of $\mathscr{N}$ that preserves arities and minoring operations, i.e., satisfying that $F(f)^\pi = F(f^\pi)$ for each suitable f, π. Similarly, given $h \in \mathbb{N}$, a *partial homomorphism* $F : \mathscr{M} \xrightarrow{h} \mathscr{N}$ *up to arity* h is a partial map defined on all elements $f \in \mathscr{M}$ of arity at most h that preserves arities and minoring operations.

We define three minions that characterize the power of the algorithms AIP, BLP, BLP + AIP. In the minion $\mathscr{M}_{\mathrm{AIP}}$ the n-ary elements are the tuples $f \in \mathbb{Z}^n$ satisfying $\sum_{i\in[n]} f(i) = 1$. Given $f \in \mathscr{M}_{\mathrm{AIP}}(n)$, and $\pi \in [m]^{[n]}$, minoring is defined by the identity $(f^\pi)(i) = \sum_{j\in\pi^{-1}(i)} f(j)$ for each $i \in [m]$. In the minion $\mathscr{M}_{\mathrm{BLP}}$ the n-ary elements are the the tuples $f \in [0,1]^n$ of rational numbers for which $\sum_{i\in[n]} f(i) = 1$. Minoring is defined as for $\mathscr{M}_{\mathrm{AIP}}$, i.e., by the identity $(f^\pi)(i) = \sum_{j\in\pi^{-1}(i)} f(j)$. Finally, in the minion $\mathscr{M}_{\mathrm{BLP+AIP}}$ the n-ary elements are pairs (f, g), where $f \in \mathscr{M}_{\mathrm{BLP}}(n), g \in \mathscr{M}_{\mathrm{AIP}}(n)$, and $f(i) = 0$ implies $g(i) = 0$ for each $i \in [n]$. Minoring in $\mathscr{M}_{\mathrm{BLP+AIP}}$ is defined component wise. That is, $(f, g)^\pi = (f^\pi, g^\pi)$, where $f^\pi = \pi^{\mathscr{M}_{\mathrm{BLP}}}(f)$, and $g^\pi = \pi^{\mathscr{M}_{\mathrm{AIP}}}(g)$. The following theorem has been shown in [11] for the algorithms AIP, BLP and in [21] for BLP + AIP.

Theorem 2.1. *Let* $\mathcal{Q} \in \{AIP, BLP, BLP + AIP\}$, *and let* $(\boldsymbol{A}, \boldsymbol{B})$ *be a finite template. Then the algorithm* $\mathcal{Q}$ *solves* $\mathrm{PCSP}(\boldsymbol{A}, \boldsymbol{B})$ *if and only if* $\mathscr{M}_{\mathcal{Q}} \to \mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$.

In (non-)computability results we consider the plain encoding of $\mathscr{M}_{\mathrm{BLP}}, \mathscr{M}_{\mathrm{AIP}}, \mathscr{M}_{\mathrm{BLP+AIP}}$, where tuples are represented as comma separated lists, delimited with parentheses, integers are represented with their decimal representations, and rational numbers are represented as irreducible fractions, written as two numbers separated by a forward slash (i.e., n/m).

Minor Conditions Another useful way of looking at minions is to consider them as multi-sorted structures [32], where the sort of each element $p \in \mathscr{M}$ is its arity $\mathrm{ar}(p)$, and $\cdot^\pi$ is a function from m-ary elements to n-ary elements for each map $\pi \in [n]^{[m]}$. We consider the multi-sorted first-order (FO) language $\mathcal{L}_{\mathrm{MC}}$ of minions. For some background of multi-sorted (or many-sorted) FO logic we refer to [38], but we require only the very basics. Formulas in $\mathcal{L}_{\mathrm{MC}}$ are built using variables, each of which has an arity (i.e., its sort), Boolean connectives, and function symbols $\cdot^\pi$ for each $n, m \in \mathbb{N}$ and each map $\pi \in [m]^{[n]}$. Variables represent elements of minions, and each symbol $\cdot^\pi$ represents the corresponding minoring operation. We write $\exists^n x$ rather than $\exists x$ to make explicit that x is a n-ary variable, and $\phi(x_1^{n_1}, \ldots, x_k^{n_k})$ to express that the free variables $x_1, \ldots, x_k$ of the formula ϕ have arities $n_1, \ldots, n_k$ respectively. For example, the formula $\exists^3 x \left(x = x^{(2,3,1)}\right)$ expresses the existence of a 3-ary cyclic element. The *maximum arity* of a formula $\phi \in \mathcal{L}_{\mathrm{MC}}$ is the largest maximum arity of any of its sub-terms. A *primitive positive formula* is one that does not include disjunction, negation, or universal quantification, and a *sentence* (or a *closed formula*) is a formula with no free variables. *Minor conditions* are closed pp-formulas in $\mathcal{L}_{\mathrm{MC}}$. We remark that more commonly the notion of minor condition is introduced using bipartite Label Cover instances (e.g., [11]).

Given a minion $\mathscr{M}$, a formula $\phi(x_1, \ldots, x_k) \in \mathcal{L}_{\mathrm{MC}}$, and elements $f_1, \ldots, f_k \in \mathscr{M}$ such that f_i has the same arity as x_i for all $i \in [k]$, we write $\mathscr{M} \models \phi(f_1, \ldots, f_k)$ to express that $\mathscr{M}$ satisfies ϕ when substituting each x_i with the element f_i. A *pp-definition* of a set $Q \subseteq \mathscr{M}$ is a pp-formula $\Phi(x)$ such that $\mathscr{M} \models \Phi(f)$ if and only if $f \in Q$. Suppose that $\mathscr{M} \models \phi$ for a minor condition ϕ. A *satisfying assignment* of ϕ in $\mathscr{M}$ maps each occurrence of the existential quantifier $\exists x \varphi(x)$ in ϕ to an element f of $\mathscr{M}$ of the same arity as x in such a way that ϕ is satisfied after substituting $\exists x \varphi(x)$ with $\varphi(f)$ in each sub-formula. When there is no ambiguity, we will simply treat assignments as maps from variables of ϕ to elements of $\mathscr{M}$. Given two minions $\mathscr{M} \to \mathscr{N}$ and a number $h \in \mathbb{N}$, in the *promise minor condition problem* $\mathrm{PMC}_h(\mathscr{M}, \mathscr{N})$ we consider an input minor condition ϕ whose maximum arity is at most h, and the task is to accept it if $\mathscr{M} \models \phi$ and reject it if $\mathscr{N} \not\models \phi$. Additionally, if $\mathscr{N}$ is locally finite, we define the *search promise minor condition problem* $\mathrm{sPMC}_h(\mathscr{M}, \mathscr{N})$ as the problem of either finding a satisfying assignment of ϕ in $\mathscr{N}$ or to reject ϕ if $\mathscr{M} \not\models \phi$.

3 Main Results

We are ready to state our main results about the algorithms AIP, BLP, and BLP + AIP. These are summarized in Figure 1. Given $\mathcal{Q} \in \{\mathrm{AIP}, \mathrm{BLP}, \mathrm{BLP} + \mathrm{AIP}\}$ and $k \in \mathbb{N}$, $\mathcal{S}_k^{\mathcal{Q}}$ denotes the family of finite templates $(\boldsymbol{A}, \boldsymbol{B})$ with $|A| \leqslant k$ such that $\mathcal{Q}$ solves $\mathrm{PCSP}(\boldsymbol{A}, \boldsymbol{B})$.

Theorem 3.1 (Main result for AIP)**.** *The following hold:*

(1) the family of rounding problems $\mathrm{sPCSP}_{AIP}(\boldsymbol{A}, \boldsymbol{B})$ *for* $(\boldsymbol{A}, \boldsymbol{B}) \in \mathcal{S}_4^{AIP}$ *is TFNP-hard,*

(2) the family $\mathrm{sPCSP}_{AIP}(\boldsymbol{A}, \boldsymbol{B})$ *for* $(\boldsymbol{A}, \boldsymbol{B}) \in \mathcal{S}_2^{AIP}$ *is* TFNP_1*-hard,*

(3) the family $\mathcal{S}_2^{AIP}$ *is undecidable, and*

(4) there is a template $(\boldsymbol{A}, \boldsymbol{B}) \in \mathcal{S}_2^{AIP}$ *for which there is no computable minion homomorphism from* $\mathscr{M}_{\mathrm{AIP}}$ *to* $\mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$*.*

Theorem 3.2 (Main result for BLP)**.** *The following hold:*

(1) the family of rounding problems $\mathrm{sPCSP}_{BLP}(\boldsymbol{A}, \boldsymbol{B})$ *for* $(\boldsymbol{A}, \boldsymbol{B}) \in \mathcal{S}_5^{BLP}$ *is TFNP-hard,*

(2) the family $\mathrm{sPCSP}_{BLP}(\boldsymbol{A}, \boldsymbol{B})$ *for* $(\boldsymbol{A}, \boldsymbol{B}) \in \mathcal{S}_2^{BLP}$ *is* TFNP_1*-hard,*

(3) the family $\mathcal{S}_2^{BLP}$ *is undecidable, and*

(4) there is a template $(\boldsymbol{A}, \boldsymbol{B}) \in \mathcal{S}_2^{BLP}$ *for which there is no computable minion homomorphism from* $\mathscr{M}_{\mathrm{BLP}}$ *to* $\mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$*.*

Theorem 3.3 (Main result for BLP + AIP)**.** *The following hold:*

(1) the family of rounding problems $\mathrm{sPCSP}_{BLP+AIP}(\boldsymbol{A}, \boldsymbol{B})$ *for* $(\boldsymbol{A}, \boldsymbol{B}) \in \mathcal{S}_5^{BLP+AIP}$ *is TFNP-hard, and*

(2) the family $\mathcal{S}_5^{BLP+AIP}$ *is undecidable.*

We also recall here our main result about cyclic polymorphisms and WNUs. We remark that $\mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$ admitting some cyclic polymorphism is equivalent to it admitting one of some prime arity.

Theorem 1.2. *The following problems are undecidable. Given an input finite template* $(\boldsymbol{A}, \boldsymbol{B})$ *with* $|A| \leqslant 3$*, determine whether* $\mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$ *contains:*

(1) cyclic polymorphisms

  (i) for every prime arity p*,*

   - *(ii) for at least one arity k,*
   - *(iii) for all but finitely many prime arities p,*
   - *(iv) for infinitely many prime arities p.*

- *(2) weak near-unanimity polymorphisms (WNUs)*
   - *(i) for every arity $k \geqslant 3$,*
   - *(ii) for at least one arity $k \geqslant 3$,*
   - *(iii) for all but finitely many arities $k \geqslant 3$,*
   - *(iv) for infinitely many arities $k \geqslant 3$.*

4 Overview of the Proofs

We sketch the proof of the following theorem while outlining the ideas used in our main results.

Theorem 4.1. *Let $\mathcal{S}$ be the set of finite templates $(\boldsymbol{A}, \boldsymbol{B})$ with $|A| \leqslant 3$ such that* $\mathrm{PCSP}(\boldsymbol{A}, \boldsymbol{B})$ *is solved by AIP. Then (1) the family of problems* $\mathrm{sPCSP}_{AIP}(\boldsymbol{A}, \boldsymbol{B})$ *where $(\boldsymbol{A}, \boldsymbol{B}) \in \mathcal{S}$ is TFNP$_1$-hard, (2) $\mathcal{S}$ is undecidable, and (3) there is a template $(\boldsymbol{A}, \boldsymbol{B}) \in \mathcal{S}$ for which there is no computable homomorphism $F : \mathscr{M}_{AIP} \to \mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$.*

The proof of this result is a reduction from tiling problems. The signature Σ_Γ consists of a unary symbol O and binary symbols E_1, E_2, and $\boldsymbol{\Gamma}$ is the Σ_Γ-structure whose universe is the upper-right quadrant $\mathbb{N}^2$, and where $O^\Gamma = \{(1,1)\}$ marks the origin, and $E_1^\Gamma = \{((m,n),(m+1,n)) \mid (m,n) \in \mathbb{N}^2\}$, $E_2^\Gamma = \{((m,n),(m,n+1)) \mid (m,n) \in \mathbb{N}^2\}$ are the horizontal and vertical adjacency relations. We write $\mathrm{Hom}(\boldsymbol{\Gamma}, \cdot)$ for the set of finite structures $\boldsymbol{T}$ satisfying $\boldsymbol{\Gamma} \to \boldsymbol{T}$. A Σ_Γ-structure $\boldsymbol{T}$ can be seen as a set of tiles, equipped with horizontal and vertical constraints, and a set of distinguished tiles that can occupy the initial position. This way, a homomorphism $F : \boldsymbol{\Gamma} \to \boldsymbol{T}$ represents a tiling of the upper-right quadrant of the plane that satisfies all the imposed restrictions.

Proposition 4.2. *The following hold:*

- *(1) the family of problems* $\mathrm{sPCSP}(\boldsymbol{\Gamma}, \boldsymbol{T})$ *where* $\boldsymbol{T} \in \mathrm{Hom}(\boldsymbol{\Gamma}, \cdot)$ *is TFNP$_1$-hard,*
- *(2)* $\mathrm{Hom}(\boldsymbol{\Gamma}, \cdot)$ *is undecidable, and*
- *(3) there exists* $\boldsymbol{T} \in \mathrm{Hom}(\boldsymbol{\Gamma}, \cdot)$ *for which there is no computable homomorphism $F : \boldsymbol{\Gamma} \to \boldsymbol{T}$.*

At a high-level, this is follows from the fact that, given a non-deterministic Turing machine M, we can construct a finite structure $\boldsymbol{T}_M$ for which the homomorphisms $F : \boldsymbol{\Gamma} \to \boldsymbol{T}_M$ represent non-halting runs of $\boldsymbol{T}_M$.

In order to prove Theorem 4.1 using this proposition we first represent $\boldsymbol{\Gamma}$ inside the minion $\mathscr{M}_{\mathrm{AIP}}$, and then we find a way of encoding finite Σ_Γ-structures in finite templates $(\boldsymbol{A}, \boldsymbol{B})$ in a suitable way. The first task is the simpler one. We can identify each pair $\boldsymbol{m} = (m_1, m_2)$ in $\mathbb{N}^2$ with the 3-ary element $f_{\boldsymbol{m}} = (m_1, m_2, 1 - m_1 - m_2) \in \mathscr{M}_{\mathrm{AIP}}$. Similarly, we can represent each pair $(\boldsymbol{m}, \boldsymbol{m}') \in E_i^\Gamma$ with the 4-ary element $g_{E_i(\boldsymbol{m},\boldsymbol{m}')} = (m_1, m_2, 1, -m_1 - m_2)$ with the idea that $f_{\boldsymbol{m}} = g^{(1,2,3,3)}_{E_i(\boldsymbol{m},\boldsymbol{m}')}$ and $f_{\boldsymbol{m}'} = g^{(1,2,i,3)}_{E_i(\boldsymbol{m},\boldsymbol{m}')}$ in mind. We call this way of representing a relational structure inside a minion an *interpretation* (Section 6.1) and we denote it by $\mathcal{I}$. We define $U^{\mathcal{I}} = \{f_{\boldsymbol{m}} \mid \boldsymbol{m} \in \mathbb{N}^2\}$, $O^{\mathcal{I}} = \{f_{(1,1)}\}$, and $E_i^{\mathcal{I}} = \{g_{E_i(\boldsymbol{m},\boldsymbol{m}')} \mid (\boldsymbol{m},\boldsymbol{m}') \in E_i^{\Gamma}\}$ for $i = 1, 2$.

The interpretation $\mathcal{I}$ is *almost* pp-definable in $\mathscr{M}_{\text{AIP}}$. Indeed, the element $(1,0) \in \mathscr{M}_{\text{AIP}}(2)$ is the only witness of the pp-formula $\phi_1(x^2) \equiv x = x^{(1,1)}$, so the 4-ary elements of the form $(m_1, m_2, 1, -m_1 - m_2) \in \mathscr{M}_{\text{AIP}}$ are precisely the witnesses of $\phi_E(y^4) \equiv \phi_1(y^{(2,2,1,2)})$. However, this does not take into account the restriction that $m_1, m_2 \in \mathbb{N}$. Under closer inspection it is not hard to see that the set of elements $(m, n) \in \mathscr{M}_{\text{AIP}}$ with $m > 0$ is not pp-definable, so, in fact, the interpretation $\mathcal{I}$ itself is not pp-definable either [2]. However, we get something almost as good: for any element $\boldsymbol{m} \in \mathbb{N}^2$ we can construct in polynomial time a pp-definition $\psi_{\boldsymbol{m}}(x^3)$ of $f_{\boldsymbol{m}}$ in $\mathscr{M}_{\text{AIP}}$. For $\boldsymbol{m} = (1,1)$, we simply define $\psi_{\boldsymbol{m}}(x^3) \equiv \phi_1(x^{(1,2,2)}) \wedge \phi_1(x^{(2,1,2)})$. Given $\boldsymbol{m} \in \mathbb{N}^2$, for $\boldsymbol{m}' = (m_1 + 1, m_2)$, we define $\psi_{\boldsymbol{m}'}(x^3)$ as the formula

$$\overset{4}{\exists} y \overset{3}{\exists} z \left(\psi_{\boldsymbol{m}}(z) \wedge \phi_E(y) \wedge z = y^{(1,2,3,3)} \wedge x = y^{(1,2,1,3)} \right), \tag{2}$$

and so on. This way, given a finite substructure $\boldsymbol{G} \subset \boldsymbol{\Gamma}$, we can construct in polynomial time a minor condition $\Psi_{\boldsymbol{G}}$ containing an existentially-quantified variable $x_{\boldsymbol{m}}$ for each vertex $\boldsymbol{m} \in G$ an existentially-quantified variable $y_{E_i(\boldsymbol{m},\boldsymbol{m}')}$ for each edge $(\boldsymbol{m},\boldsymbol{m}') \in E_i^G$, such that the satisfying assignments of $\Psi_{\boldsymbol{G}}$ on $\mathscr{M}_{\text{AIP}}$ must map each variable to the minion element that encodes the corresponding vertex or edge. That is, $x_{\boldsymbol{m}} \mapsto f_{\boldsymbol{m}}$ and $y_{E_i(\boldsymbol{m},\boldsymbol{m}')} \mapsto g_{E_i(\boldsymbol{m},\boldsymbol{m}')}$. The minor condition $\Psi_{\boldsymbol{G}}$ can be defined as

$$\mathop{\overset{3}{\exists}}_{\boldsymbol{m}\in G} x_{\boldsymbol{m}} \mathop{\overset{4}{\exists}}_{i\in[2],(\boldsymbol{m},\boldsymbol{m}')\in E_i^G} y_{E_i(\boldsymbol{m},\boldsymbol{m}')} \left(\bigwedge_{\boldsymbol{m}\in G} \psi_{\boldsymbol{m}}(x_{\boldsymbol{m}}) \right) \bigwedge \left(\bigwedge_{i\in[2],(\boldsymbol{m},\boldsymbol{m}')\in E_i^G} \psi_E(y_{E_i(\boldsymbol{m},\boldsymbol{m}')}) \wedge x_{\boldsymbol{m}} = y^{(1,2,3,3)}_{E_i(\boldsymbol{m},\boldsymbol{m}')} \wedge x_{\boldsymbol{m}'} = y^{(1,2,i,3)}_{E_i(\boldsymbol{m},\boldsymbol{m}')} \right). \tag{3}$$

This construction is called a *pattern* (Section 6.4). Having found a nice way to represent $\boldsymbol{\Gamma}$ inside $\mathscr{M}_{\text{AIP}}$, the next step is to develop a construction that, given a finite Σ_Γ-structure $\boldsymbol{T}$, builds a suitable finite template $(\boldsymbol{A_T}, \boldsymbol{B_T})$. We would like $\mathscr{M}_{\boldsymbol{T}} = \text{Pol}(\boldsymbol{A_T}, \boldsymbol{B_T})$ to satisfy the following.

Property 1 (Conditions that enable our reductions)**.**

(I) $\mathscr{M}_{\text{AIP}} \to \mathscr{M}_{\boldsymbol{T}}$ if and only if $\boldsymbol{\Gamma} \to \boldsymbol{T}$.

(II) A homomorphism $F : \boldsymbol{\Gamma} \to \boldsymbol{T}$ can be computed given oracle access to a homomorphism $H : \mathscr{M}_{\text{AIP}} \to \mathscr{M}_{\boldsymbol{T}}$.

(III) Given a finite substructure $\boldsymbol{G} \subset \boldsymbol{\Gamma}$, there is a polynomial-time reduction from the task of finding a homomorphism $F : \boldsymbol{G} \to \boldsymbol{T}$ to the task of finding a satisfying assignment of $\Psi_{\boldsymbol{G}}$ in $\mathscr{M}_{\boldsymbol{T}}$.

It is easy enough to see that the conditions (I) and (II) allow us to respectively reduce the undecidability and non-computability parts of Theorem 4.1 to those of Proposition 4.2. In to bridge

---

[2] This could be circumvented by having defined $\boldsymbol{\Gamma}$ on the whole integer plane $\mathbb{Z}^2$ instead of $\mathbb{N}^2$, but we want to illustrate the point that our reductions are not uniform pp-definitions. This will be useful in the more involved proofs.

the gap between the TFNP_1-hardness statements in those results, we use condition (III) together with the following additional result, which is a consequence of Theorem 8.5 (Section 8). We note that this is a standard modification of an analogous result, [11, Theorem 3.12], that has been applied widely in the PCSP literature.

Proposition 4.3. *Suppose that AIP solves* $\mathrm{PCSP}(\boldsymbol{A}, \boldsymbol{B})$ *for a finite template* $(\boldsymbol{A}, \boldsymbol{B})$*. Then* $\mathrm{sPCSP}_{AIP}(\boldsymbol{A}, \boldsymbol{B})$ *is log-space equivalent to* $\mathrm{sPMC}_N(\mathscr{M}_{\mathrm{AIP}}, \mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B}))$*, where N is at least as large as $|A|$ and $|R^A|$ for all relation symbols R.*

We remark that we do not always achieve the conditions from Property 1 in tandem: for example, sometimes we obtain templates that satisfy both (I) and (III), but not (II), meaning that we are able to obtain undecidability, and hardness-of-rounding results for a given algorithm, but not a related non-computability result. This is the case, for example, of BLP + AIP. With our methods, the condition (I) used to prove undecidability has the weakest requirements, while the conditions (II) and (III) used to show non-computability and hardness have non-comparable requirements.

Rather than constructing the template $(\boldsymbol{A_T}, \boldsymbol{B_T})$ directly, we focus on its polymorphism minion instead. We start by defining a minion $\mathscr{N}_{\boldsymbol{T}}$ that satisfies Property 1 of $\mathrm{Pol}(\boldsymbol{A_T}, \boldsymbol{B_T})$. The n-ary elements of $\mathscr{N}_{\boldsymbol{T}}$ are pairs (f, χ), where $f \in \mathscr{M}_{\mathrm{AIP}}(n)$, and $\chi : [3]^{[n]} \rightharpoonup T$ is a partial map defined on the elements $\gamma \in [3]^{[n]}$ such that $f^\gamma \in U^{\mathcal{I}}$ that satisfies the following properties: if $f_{(1,1)} = f^\gamma$, then $\chi(\gamma) \in O^T$, and if $g_{E_i(\boldsymbol{m},\boldsymbol{m}')} = f^\pi$, then $(\chi(\gamma), \chi(\gamma')) \in E_i^T$, where $\gamma = (1,2,3,3) \circ \pi$, and $\gamma' = (1,2,i,3) \circ \pi$. Intuitively, the element $f \in \mathscr{M}_{\mathrm{AIP}}$ *covers* some part of $\mathbf{\Gamma}$ (as interpreted by $\mathcal{I}$), which consists of the elements $f_{\boldsymbol{m}} \in U^{\mathcal{I}}$ that can be obtained by minoring f. Similarly, we also think of f as covering the edges $(\boldsymbol{m}, \boldsymbol{m}') \in E_i^{\Gamma}$ where $g_{E_i(\boldsymbol{m},\boldsymbol{m}')} \in E_i^{\mathcal{I}}$ is a minor of f. Then, the map χ represents a homomorphism from the substructure $\mathbf{S}_f \subseteq \mathbf{\Gamma}$ covered by f to $\boldsymbol{T}$. Given a map $\pi \in [m]^{[n]}$, we define the minor $(f, \chi)^\pi$ as the pair (f^π, χ^π), where $\chi^\pi : [3]^{[m]} \rightharpoonup T$ is the partial map defined by $\gamma \mapsto \chi(\gamma \circ \pi)$. The rationale behind this construction is that the structure $\mathbf{S}_{f^\pi}$ covered by f^π is a substructure of $\mathbf{S}_f$, so χ^π is defined so that it represents the restriction of χ to $\mathbf{S}_{f^\pi}$. If the substructures $\mathbf{S}_f$ formed a topology over $\mathbf{\Gamma}$ (which, we remark, is not the case), then the minion $\mathscr{N}_{\boldsymbol{T}}$ would encode a *sheaf* [23] over $\mathbf{\Gamma}$ whose *sections* correspond to partial homomorphisms to $\boldsymbol{T}$. We keep this topological intuition in mind throughout the proof. We call $\mathscr{N}_{\boldsymbol{T}}$ the *manifold minion* [3] given by the interpretation $\mathcal{I}$ and the structure $\boldsymbol{T}$ (Section 6.2).

Proposition 4.4. *The minion $\mathscr{N}_{\boldsymbol{T}}$ satisfies Property 1.*

Let us sketch the proof of this fact. It is not difficult to show that every element $f \in \mathscr{M}_{\mathrm{AIP}}$ is pp-definable, so the only minion homomorphism $F : \mathscr{M}_{\mathrm{AIP}} \to \mathscr{M}_{\mathrm{AIP}}$ is the identity. Hence, any homomorphism $F : \mathscr{M}_{\mathrm{AIP}} \to \mathscr{N}_{\boldsymbol{T}}$ composed with the left projection must yield the identity over $\mathscr{M}_{\mathrm{AIP}}$. In other words, F must be of the form $f \mapsto (f, \chi_f)$. If such homomorphism F exists, then the charts χ_f must be compatible with each other, yielding a global homomorphism from $\mathbf{\Gamma}$ to $\boldsymbol{T}$. Conversely, if there is a homomorphism $H : \mathbf{\Gamma} \to \boldsymbol{T}$, one can restrict it to each local structure $\mathbf{S}_f$ to obtain compatible charts χ_f and define a minion homomorphism F. This proves $\mathscr{N}_{\boldsymbol{T}}$ satisfies (I). To show (II), suppose that we have oracle access to a homomorphism $F : \mathscr{M}_{\mathrm{AIP}} \to \mathscr{N}_{\boldsymbol{T}}$. Then, by the previous reasoning, in order to compute a homomorphism $H : \mathbf{\Gamma} \to \boldsymbol{T}$, given an element

---

[3] In previous versions of this work, this was called the exponential minion. We changed the name to appease a category-theoretically inclined reader who would point out that these minions are not, in fact, exponential objects in the category of minions.

$\boldsymbol{m} \in \mathbb{N}^2$ we just need to query $F(f_{\boldsymbol{m}}) = (f_{\boldsymbol{m}}, \chi_{f_{\boldsymbol{m}}})$ and consider the chart $\chi_{f_{\boldsymbol{m}}}$. Finally, to see that (III) holds, consider a finite substructure $\boldsymbol{G} \subset \boldsymbol{\Gamma}$, and a satisfying assignment $x \mapsto (f_x, \chi_x)$ of $\Psi_{\boldsymbol{G}}$ in $\mathscr{N}_{\boldsymbol{T}}$. The map $x \mapsto f_x$ is a satisfying assignment of $\Psi_{\boldsymbol{G}}$ in $\mathscr{M}_{\mathrm{AIP}}$, so by construction of $\Psi_{\boldsymbol{G}}$ it must hold that $f_x = f_{\boldsymbol{m}}$ when $x = x_{\boldsymbol{m}}$ for each $\boldsymbol{m} \in G$, and $f_x = g_{E_i(\boldsymbol{m},\boldsymbol{m}')}$ when $x = y_{E_i(\boldsymbol{m},\boldsymbol{m}')}$ for each $i \in [2]$, $(\boldsymbol{m}, \boldsymbol{m}') \in E_i^G$. Hence, when x ranges over all variables $\{x_{\boldsymbol{m}}\}_{\boldsymbol{m}\in G} \cup \{y_{E_i(\boldsymbol{m},\boldsymbol{m}')}\}_{i\in[2],(\boldsymbol{m},\boldsymbol{m}')\in E_i}$, the elements f_x cover the whole structure $\boldsymbol{G}$ (as induced by the interpretation $\mathcal{I}$), and the charts χ_x must piece together a homomorphism from $\boldsymbol{G}$ to $\boldsymbol{T}$.

The main issue at this point is that, despite having the desired properties, the minion $\mathscr{N}_{\boldsymbol{T}}$ is not (isomorphic to) the polymorphism minion of any finite template $(\boldsymbol{A_T}, \boldsymbol{B_T})$. In fact, it is not even locally finite. The first step towards constructing $(\boldsymbol{A_T}, \boldsymbol{B_T})$ from the minion $\mathscr{N}_{\boldsymbol{T}}$ is to find a locally finite quotient $\mathscr{N}'_{\boldsymbol{T}}$ of $\mathscr{N}_{\boldsymbol{T}}$ that still satisfies Property 1. More precisely, $\mathscr{N}'_{\boldsymbol{T}}$ will be obtained by performing the manifold construction on a locally finite quotient of $\mathscr{M}_{\mathrm{AIP}}$ by some equivalence relation $\sim$. Let us spell this out. The equivalence $\sim$ defines a quotient minion $\mathscr{M}_{\mathrm{AIP}}/\sim$ where each element $\langle f \rangle$ is the $\sim$-class of an element f in the original minion $\mathscr{M}_{\mathrm{AIP}}$. Similarly, given a subset $S \subseteq \mathscr{M}_{\mathrm{AIP}}$, we write $\langle S \rangle$ for the subset of $\mathscr{M}_{\mathrm{AIP}}/\sim$ consisting of the elements $\langle f \rangle$ for each $f \in S$. The relation $\sim$ also induces an equivalence on $\boldsymbol{\Gamma}$ through the interpretation $\mathcal{I}$ defined as $\boldsymbol{m} \sim \boldsymbol{m}'$ whenever $f_{\boldsymbol{m}} \sim f_{\boldsymbol{m}'}$. Finally, we can also speak of the quotient interpretation $\mathcal{I}/\sim$ over the quotient minion $\mathscr{M}_{\mathrm{AIP}}/\sim$, which is given by $U^{\mathcal{I}/\sim} = \langle U^{\mathcal{I}} \rangle$, $O^{\mathcal{I}/\sim} = \langle O^{\mathcal{I}} \rangle$, and $E_i^{\mathcal{I}/\sim} = \langle E_i^{\mathcal{I}} \rangle$ for $i = 1, 2$. It can be shown that $\mathcal{I}/\sim$ induces the quotient structure $\Gamma/\sim$ on the quotient minion $\mathscr{M}_{\mathrm{AIP}}/\sim$. Then, we define $\mathscr{N}'_{\boldsymbol{T}}$ as the manifold minion given by the interpretation $\mathcal{I}/\sim$ and the structure $\boldsymbol{T}$, similarly to $\mathscr{N}_{\boldsymbol{T}}$.

The following conditions ensure that $\mathscr{N}'_{\boldsymbol{T}}$ still satisfies Property 1 and that this property can be further preserved when transforming $\mathscr{N}'_{\boldsymbol{T}}$ into a polymorphism minion in the final step.

Property 2 (Requirements for a good quotient)**.**

(A) *The interpretation $\mathcal{I}$ is $\sim$-stable.* It holds that $\langle f \rangle \in \langle U^{\mathcal{I}} \rangle$ if and only if $f \in U^{\mathcal{I}}$, $\langle f \rangle \in \langle O^{\mathcal{I}} \rangle$ if and only if $f \in O^{\mathcal{I}}$, and $\langle f \rangle \in \langle E_i^{\mathcal{I}} \rangle$ if and only if $f \in E_i^{\mathcal{I}}$ for all $i \in [2]$.

(B) *The interpretation $\mathcal{I}$ is internal with respect to $\sim$ at arity* 4. For each partial homomorphism $F : \mathscr{M}_{\mathrm{AIP}} \rightharpoonup \mathscr{M}_{\mathrm{AIP}}/\sim$ defined up to arity 4, it holds that $F(U^{\mathcal{I}}) \subseteq \langle U^{\mathcal{I}} \rangle$, $F(O^{\mathcal{I}}) \subseteq \langle O^{\mathcal{I}} \rangle$, and $F(E_i^{\mathcal{I}}) \subseteq \langle E_i^{\mathcal{I}} \rangle$ for each $i \in [2]$.

(C) *The map $\boldsymbol{G} \mapsto \Psi_{\boldsymbol{G}}$ defined* (3) *is an internal pattern with respect to $\sim$.* Given a finite substructure $\boldsymbol{G} \subset \boldsymbol{\Gamma}$, any satisfying assignment of $\Psi_{\boldsymbol{G}}$ in $\mathscr{M}_{\mathrm{AIP}}/\sim$ must correspond to a homomorphism from $\boldsymbol{G}$ into $\Gamma/\sim$[4].

Let us briefly go over these conditions. Condition (A) ensures that if there is a homomorphism $F : \Gamma \to \boldsymbol{T}$ then there is a minion homomorphism $H : \mathscr{M}_{\mathrm{AIP}} \to \mathscr{N}'$. Indeed, in this case we can define H as the map $f \mapsto (\langle f \rangle, \chi_f)$, where χ_f is obtained by looking at the restriction of F to the local structure induced by $\mathcal{I}$ on f, and then factoring the local homomorphism through the quotient. The reason this construction works is that the quotient map $\boldsymbol{\Gamma} \to \boldsymbol{\Gamma}/\sim$ preserves non-relations (i.e., sends non-edges to non-edges and sends non-origin elements to non-origin elements) so

---

[4] We warn the reader that, for the sake of exposition in this introductory section, the statement in (C) is not fully precise.

partial homomorphisms on local structures can be factored through the quotient [5]. Condition (B) implies that every minion homomorphism $H : \mathscr{M}_{\mathrm{AIP}} \to \mathscr{N}'_{\boldsymbol{T}}$ corresponds to a homomorphism $F : \mathbf{\Gamma} \to \boldsymbol{T}$. Indeed, intuitively, this condition means that every homomorphism $\mathscr{M}_{\mathrm{AIP}} \to \mathscr{M}_{\mathrm{AIP}}/\sim$ must induce a homomorphism between the structures $\mathbf{\Gamma}$ and $\mathbf{\Gamma}/\sim$ interpreted on those minions. So, if the homomorphism H maps $f \in \mathscr{M}_{\mathrm{AIP}}$ to some pair $(\langle g \rangle, \chi) \in \mathscr{N}'_{\boldsymbol{T}}$, then it must hold that the local structure $\boldsymbol{S}_{\langle g \rangle}$ given by $\langle g \rangle$ in $\mathscr{M}_{\mathrm{AIP}}/\sim$ is a homomorphic image of the local structure $\boldsymbol{S}_f$ of f in $\mathscr{M}_{\mathrm{AIP}}$, meaning that the local homomorphism $\chi : \boldsymbol{S}_{\langle g \rangle} \to \boldsymbol{T}$ can be lifted to another one $\tilde{\chi} : \boldsymbol{S}_f \to \boldsymbol{T}$. It can be seen that these lifted local homomorphisms form a consistent family, so they define a global homomorphism from $\mathbf{\Gamma}$ to $\boldsymbol{T}$. Finally, condition (C) requires that any assignment of $\Psi_{\boldsymbol{G}}$ in $\mathscr{M}_{\mathrm{AIP}}/\sim$ maps, for each $\boldsymbol{m} \in G$, the variable $x_{\boldsymbol{m}}$ to some element $\langle f_{\boldsymbol{m}'} \rangle$ in such a way that the map $\boldsymbol{m} \mapsto \langle \boldsymbol{m}' \rangle$ is a homomorphism from $\boldsymbol{G}$ to $\mathbf{\Gamma}/\sim$. Arguing again that you can factor local homomorphism through the quotient, this condition allows us to transfer (III) in Property 1 from the minion $\mathscr{N}_{\boldsymbol{T}}$ to $\mathscr{N}'_{\boldsymbol{T}}$.

Another important observation is that (B) would be weaker if we considered non-partial homomorphisms instead. Indeed, every homomorphism can be restricted to a partial homomorphism, but not every partial homomorphism can, in principle, be extended to a fully-defined homomorphism. We consider partial homomorphisms planning for a future step where, in order to obtain a polymorphism minion, we will have to let go of high-arity terms in a certain way. The choice of arity 4 corresponds to the fact that the minor conditions $\Psi_{\boldsymbol{G}}$ have maximum arity 4.

In order to construct the relation $\sim$ we select some relevant *predicates* (related to $\mathcal{I}$) containing the information that we would like to preserve in the quotient. These predicates will be sets of binary elements $D_1, D_{\mathbb{N}} \subseteq \mathscr{M}_{\mathrm{AIP}}(2)$ defined as $D_1 = \{(1,0)\}$, and $D_{\mathbb{N}} = \{(m, 1-m) \mid m \in \mathbb{N}\}$. We call the set of predicates $\mathcal{D} = \{D_1, D_{\mathbb{N}}\}$ a *description* (Section 6.3). Then, we write $f \sim_{\mathcal{D}} g$ for two elements $f, g \in \mathscr{M}_{\mathrm{AIP}}$ of the same arity n whenever for all $\pi \in [2]^{[n]}$ and all $P \in \mathcal{D}$ the inclusion $f^{\pi} \in P$ holds if and only if $g^{\pi} \in P$ as well. We shorten expressions of the form $\cdot/\sim_{\mathcal{D}}$ to $\cdot/\mathcal{D}$. Figure 2 displays the quotient $\mathbf{\Gamma}/\mathcal{D}$.

Fact 4.5. *The equivalence relation $\sim_{\mathcal{D}}$ satisfies Property 2.*

The fact that $\sim_{\mathcal{D}}$ meets condition (A) in Property 2 follows easily from the definition of this relation. Items (B) and (C), hold roughly, because the formula $\psi_{\boldsymbol{m}}(x)$ given in (2) is a pp-definition of $\langle f_{\boldsymbol{m}} \rangle$ in $\mathscr{M}_{\mathrm{AIP}}/\mathcal{D}$ for each $\boldsymbol{m} \in \mathbb{N}^2$ [6], and there are analogous pp-definitions for the elements $\langle g_{E_i(\boldsymbol{m}_1, \boldsymbol{m}_2)} \rangle$ representing each edge. Let us sketch this fact. First, we can see that $\phi_1(x^2) \equiv x = x^{(1,1)}$ defines the element $\langle (1,0) \rangle$. Indeed, if $\mathscr{M}_{\mathrm{AIP}}/\mathcal{D} \models \phi_1(\langle f \rangle)$ then $f \sim_{\mathcal{D}} f^{(1,1)}$. However, $f^{(1,1)} \in D_1$ for all $f \in \mathscr{M}_{\mathrm{AIP}}(2)$, so f must belong to D_1 as well, meaning that $f = (1,0)$. Following this argument, we see that $\psi_{(1,1)}(x)$ defines $\langle f_{(1,1)} \rangle$, and the other values of $\boldsymbol{m} \in \mathbb{N}^2$ can be handled by induction.

The fact that $\psi_{\boldsymbol{m}}(x)$ is a pp-definition of $\langle f_{\boldsymbol{m}} \rangle$ in $\mathscr{M}_{\mathrm{AIP}}/\mathcal{D}$ forces any partial homomorphism $H : \mathscr{M}_{\mathrm{AIP}} \rightharpoonup \mathscr{M}_{\mathrm{AIP}}/\mathcal{D}$ defined up to arity 4 (which is the maximum arity of $\phi_{\boldsymbol{m}}(x)$) to satisfy $H(f_{\boldsymbol{m}}) = \langle f_{\boldsymbol{m}} \rangle$ for all $\boldsymbol{m} \in \mathbb{N}^2$. Indeed, it holds that $\mathscr{M}_{\mathrm{AIP}} \models \psi_{\boldsymbol{m}}(f_{\boldsymbol{m}})$, and this is witnessed by elements of arity at most 4, so $\mathscr{M}_{\mathrm{AIP}}/\mathcal{D} \models \psi_{\boldsymbol{m}}(H(f_{\boldsymbol{m}})$, and it must be that $H(f_{\boldsymbol{m}}) = \langle f_{\boldsymbol{m}} \rangle$. This shows

---

[5] Making this statement formal requires being more careful in our definition of local structures. The direct approach, which we follow in this proof sketch, of defining the local structure $\boldsymbol{S}_f$ as the substructure of $\mathbf{\Gamma}$ induced on the minors of f does not work. The issue is that the quotient map $f \mapsto \langle f \rangle$ might identify minors of f, which prevents us from factoring partial homomorphisms through the quotient.

[6] This is a stronger statement than what we really need. What we will use is the existence of a kind of formulas which we call "internal references". These are introduced in Section 6.3.

that condition (B) holds. Finally, one can also show that (C) holds via a similar argument using the fact that the minor conditions Ψ_G are built-up from the formulas $\psi_m(x)$.

Having obtained a good equivalence relation, we construct the locally finite minion $\mathscr{N}'_T$ as outlined previously. To reiterate, we consider the quotient interpretation $\mathcal{I}/\mathcal{D}$ on the quotient minion $\mathscr{M}_{\text{AIP}}/\mathcal{D}$ and define $\mathscr{N}'_T$ as the manifold minion given by $\mathcal{I}/\mathcal{D}$ and the structure $\boldsymbol{T}$. The fact that $\sim_{\mathcal{D}}$ satisfies Property 2, means that $\mathscr{N}'_T$ manages to preserve Property 1, which $\mathscr{N}_T$ previously satisfied. However, we are still not done: The minion $\mathscr{N}'_T$ is locally finite, but it may not be isomorphic to a polymorphism minion. The reason is that we do not know whether it is finitizable in the sense of [18]. We circumvent this issue by proving the following.

Proposition 4.6. *There is a finite template* $(\boldsymbol{A}, \boldsymbol{B})$ *with* $|A| = 3$ *and* $|R^A| \leqslant 4$ *for any relation symbol* R *satisfying that (1) there is a partial minion isomorphism* $F : \mathscr{N}'_T \rightharpoonup \operatorname{Pol}(\boldsymbol{A}, \boldsymbol{B})$ *defined up to arity* 4*, and (2) for any minion* $\mathscr{M}$*, it holds that* $\mathscr{M} \to \operatorname{Pol}(\boldsymbol{A}, \boldsymbol{B})$ *if and only if there is a partial homomorphism* $H : \mathscr{M} \rightharpoonup \mathscr{N}'_T$ *defined up to arity* 4*.*

This is a consequence of the more general result Theorem 7.1 (Section 7). The template $(\boldsymbol{A}, \boldsymbol{B})$ is constructed as follows. We consider a single relation symbol R of arity 3^4. Then $\boldsymbol{A}$ is *the most general structure* on 3 elements containing a single relation R^A of size 4, following a construction given in [18]. The fact that $|A| = 3$ in this result has to do with the fact that any two elements $f, g \in \mathscr{N}'_T$ of the same arity are equal if and only if all their 3-ary minors are equal. Hence, $\mathscr{N}'_T$ can be seen as a function minion on a domain of size 3. The structure $\boldsymbol{B}$ is defined as the *free-structure* [11] of $\mathscr{N}'_T$ generated by $\boldsymbol{A}$.

The fact that we do not obtain a full isomorphism in Proposition 4.6 may seem concerning, but we already accounted for this by considering partial homomorphisms in Property 2. Crucially, the second item in this proposition, ensures that $\mathscr{M}_{\text{AIP}} \to \operatorname{Pol}(\boldsymbol{A}, \boldsymbol{B})$ exactly when there is a partial homomorphism $H : \mathscr{M}_{\text{AIP}} \rightharpoonup \mathscr{N}'_T$ defined up to arity 4. Putting together some of the implications we have shown so far shows that $\operatorname{Pol}(\boldsymbol{A}, \boldsymbol{B})$ satisfies condition (I) in Property 1, and the rest of the conditions follow similarly. This way, the following holds and finishes the proof sketch of Theorem 4.1.

Proposition 4.7. *Let* $(\boldsymbol{A}, \boldsymbol{B})$ *be a template satisfying the statement of Proposition 4.6. Then* $\operatorname{Pol}(\boldsymbol{A}, \boldsymbol{B})$ *satisfies Property 1.*

4.1 Organization of the Paper

The rest of the paper follows roughly the same structure as the proof sketch in this section. We begin in Section 5 by introducing the sources of hardness, undecidability and non-computability that are the starting point of our reductions. In Section 6 we properly introduce the notions required to prove the main results, including *interpretations*, *manifold minions*, *descriptions*, and *patterns*. Each subsection defines some concepts and includes proofs of related auxiliary results. In Section 7 we describe how to obtain polymorphism minions that are partially isomorphic to a given locally-finite minion. In Section 8 we show the reductions that allow us to transfer undecidability, non-computability, and hardness results from tiling problems to rounding problems. In Section 9 we put everything together and prove our main results. Finally Section 10 discusses in greater depth the link between our results and some open questions in the area, outlining some research directions.

Algorithm $\mathcal{Q}$ Result	$\mathcal{Q} = \text{AIP}$	$\mathcal{Q} = \text{BLP}$	$\mathcal{Q} = \text{BLP+AIP}$
Undecidability of $\mathscr{M}_{\mathcal{Q}} \to \operatorname{Pol}(\boldsymbol{A}, \boldsymbol{B})$	$\lvert A \rvert = 2$	$\lvert A \rvert = 2$	$\lvert A \rvert = 5$
Non-computability of homomorphisms $F : \mathscr{M}_{\mathcal{Q}} \to \operatorname{Pol}(\boldsymbol{A}, \boldsymbol{B})$	$\lvert A \rvert = 2$	$\lvert A \rvert = 2$	$--$
TFNP_1-hardness of $\text{sPCSP}_{\mathcal{Q}}(\boldsymbol{A}, \boldsymbol{B})$	$\lvert A \rvert = 2$	$\lvert A \rvert = 2$	$\lvert A \rvert = 5$
TFNP-hardness of $\text{sPCSP}_{\mathcal{Q}}(\boldsymbol{A}, \boldsymbol{B})$	$\lvert A \rvert = 4$	$\lvert A \rvert = 5$	$\lvert A \rvert = 5$

Figure 1: Main algorithmic results.

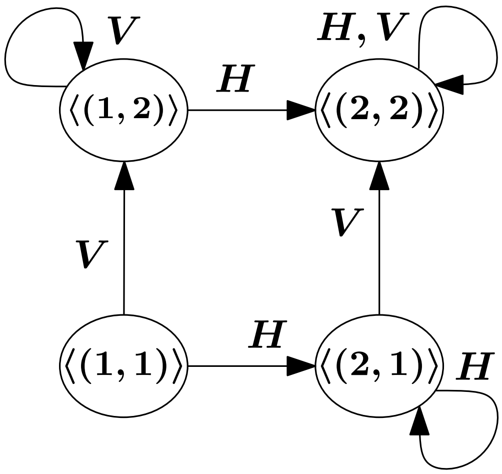

Figure 2: The quotient $\boldsymbol{\Gamma}/\sim_{\mathcal{D}}$

5 Sources of Undecidability, Non-Computability and Hardness

Given a family $\mathbb{S}$ of similar structures, $\operatorname{Hom}(\mathbb{S}, \cdot)$ denotes the set of finite structures satisfying $\boldsymbol{S} \to \boldsymbol{I}$ for some $\boldsymbol{S} \in \mathbb{S}$. When $\mathbb{S}$ is a singleton family $\{\boldsymbol{S}\}$, we write $\operatorname{Hom}(\boldsymbol{S}, \cdot)$ rather than $\operatorname{Hom}(\{\boldsymbol{S}\}, \cdot)$. We write $\operatorname{Hom}_{\text{eg}}(\mathbb{S}, \cdot)$, where eg stands for *eventually globally*, for the set of finite structures $\boldsymbol{I}$ satisfying $\boldsymbol{S} \to \boldsymbol{I}$ for all but finitely many $\boldsymbol{S} \in \mathbb{S}$. Finally, we define $\operatorname{Hom}_{\text{io}}(\mathbb{S}, \cdot)$, where io stands for *infinitely often*, for the set of finite structures $\boldsymbol{I}$ satisfying $\boldsymbol{S} \to \boldsymbol{I}$ holds for infinitely many $\boldsymbol{S} \in \mathbb{S}$.

When $\mathbb{S}$ is a finite family of finite structures, or $\boldsymbol{S}$ is a finite structure, all the sets described above can easily be recognized in polynomial time. However, when $\mathbb{S}$ is an infinite family, or $\boldsymbol{S}$ is an infinite structure, the previous problems can be undecidable. We deal mostly with these later cases.

Let $\boldsymbol{S}$ and $\boldsymbol{T}$ be two similar structures. We say that $\boldsymbol{S}$ and $\boldsymbol{T}$ are *finitely equivalent* if $\boldsymbol{I} \to \boldsymbol{S}$ and $\boldsymbol{I} \to \boldsymbol{T}$ are equivalent conditions for every finite structure $\boldsymbol{I}$. A standard argument shows that, given a finite structure $\boldsymbol{I}$, the fact that $\boldsymbol{S} \to \boldsymbol{I}$ is equivalent to $\boldsymbol{G} \to \boldsymbol{I}$ for every finite substructure

$\boldsymbol{G} \subseteq \boldsymbol{S}$. See [11, Remark 7.13] for a proof of this fact in the countable case, which is the only one we use. This yields the following result.

Lemma 5.1. *Let $\boldsymbol{S}$, $\boldsymbol{T}$ be other two finitely equivalent Σ-structures. Let $\boldsymbol{I}$ be a finite Σ-structure. Then $\boldsymbol{S} \to \boldsymbol{I}$ if and only if $\boldsymbol{T} \to \boldsymbol{I}$.*

Proof. We show that $\boldsymbol{S} \to \boldsymbol{I}$ implies that $\boldsymbol{T} \to \boldsymbol{I}$. The reverse implication follows analogously. For any finite substructure $\boldsymbol{G} \subseteq \boldsymbol{S}$ it holds that $\boldsymbol{G} \to \boldsymbol{S}$. By the definition of finitely equivalent structures, it must hold that $\boldsymbol{G} \to \boldsymbol{T}$ as well. Given that we have assumed that $\boldsymbol{T} \to \boldsymbol{I}$, composing homomorphisms we obtain that $\boldsymbol{G} \to \boldsymbol{I}$. This way, any finite substructure of $\boldsymbol{S}$ maps homomorphically to $\boldsymbol{I}$, meaning that $\boldsymbol{S} \to \boldsymbol{I}$ as well. ☐

The proofs of the remaining results in this section can be found in Section A. The techniques are standard, and essentially follow the idea from [64] that runs of a Turing machine M can be encoded in tilings of the plane by using consecutive horizontal lines to describe consecutive configurations of the machine M.

Recall the definition of the grid structure $\boldsymbol{\Gamma}$ from Section 4. We repeat here the main result about that structure for the sake of completeness.

Proposition 4.2. *The following hold:*

(1) the family of problems $\mathrm{sPCSP}(\boldsymbol{\Gamma}, \boldsymbol{T})$ *where* $\boldsymbol{T} \in \mathrm{Hom}(\boldsymbol{\Gamma}, \cdot)$ *is TFNP*$_1$*-hard,*

(2) $\mathrm{Hom}(\boldsymbol{\Gamma}, \cdot)$ *is undecidable, and*

(3) there exists $\boldsymbol{T} \in \mathrm{Hom}(\boldsymbol{\Gamma}, \cdot)$ *for which there is no computable homomorphism* $F : \boldsymbol{\Gamma} \to \boldsymbol{T}$*.*

Items (2) and (3) were shown in [64] and [43] respectively. In (3) we consider the plain encoding of $\boldsymbol{\Gamma}$, which represents pairs $(m_1, m_2) \in \mathbb{N}^2$ as comma-separated lists delimited by parentheses, where the integers are written in decimal notation.

It is worth remarking that the notion of finite equivalence introduced earlier preserves both the undecidability and the hardness parts in this result. Indeed, if $\boldsymbol{G}$ is finitely equivalent to $\boldsymbol{\Gamma}$ then $\mathrm{Hom}(\boldsymbol{G}, \cdot)$ and $\mathrm{Hom}(\boldsymbol{\Gamma}, \cdot)$ are the same family, and the problems $\mathrm{sPCSP}(\boldsymbol{G}, \boldsymbol{T})$ and $\mathrm{sPCSP}(\boldsymbol{\Gamma}, \boldsymbol{T})$ are the same whenever $\boldsymbol{G} \to \boldsymbol{T}$ for some finite structure $\boldsymbol{T}$. However, the non-computability part of this result is not preserved by finite equivalence. To see this, define $\boldsymbol{G}$ as a disjoint union of increasingly large grids. For instance, for each $n \in \mathbb{N}$, let $\boldsymbol{G}_n$ be the substructure of $\boldsymbol{\Gamma}$ induced on $[n]$, and let $\boldsymbol{G} = \bigsqcup_{n=1}^{\infty} \boldsymbol{G}_n$. It is not difficult to see that $\boldsymbol{G}$ is finitely equivalent to $\boldsymbol{\Gamma}$. Even so, it also holds that there exists a computable homomorphism $F : \boldsymbol{G} \to \boldsymbol{T}$ whenever $\boldsymbol{G} \to \boldsymbol{T}$ for a finite structure $\boldsymbol{T}$. To see this, observe that when querying the value of a global homomorphism on an element v, an algorithm just needs to compute a homomorphism from v's connected component, which is finite, to $\boldsymbol{T}$, and store it in memory. This yields a global homomorphism that can be computed on the fly.

In order to obtain TFNP-hardness instead of TFNP$_1$-hardness we utilize a three-dimensional grid with extra constraints corresponding to "doubling" each coordinate. The *super-grid* $\boldsymbol{\Gamma}^+$ has signature $\Sigma_{\Gamma^+} = \{O, E_1, E_2, E_3, \mathbb{E}_1, \mathbb{E}_2, \mathbb{E}_3\}$, where the symbol O is unary and all other symbols are binary. The universe Γ^+ is the set of triples $\mathbb{N}^3$. The relations of $\boldsymbol{\Gamma}^+$ are defined as follows. We have $O^{\Gamma^+} = \{(1,1,1)\}$. The relations $E_1^{\Gamma^+}, E_2^{\Gamma^+}, E_3^{\Gamma^+}$ describe unit increments in the first, second, and third coordinate respectively. I.e., $E_1^{\Gamma^+} = \{((m,n,o),(m+1,n,o)) \mid (m,n,o) \in \mathbb{N}^3\}$, and so on. Similarly, the relations $\mathbb{E}_1^{\Gamma^+}, \mathbb{E}_2^{\Gamma^+}, \mathbb{E}_3^{\Gamma^+}$ describe doubling increments in the first, second, and third coordinate respectively. That is, $\mathbb{E}_1^{\Gamma^+} = \{((m,n,o),(2m,n,o)) \mid (m,n,o) \in \mathbb{N}^3\}$, and so on.

Proposition 5.2. *The following hold:*

(1) the family of problems $\mathrm{sPCSP}(\mathbf{\Gamma}^+, \boldsymbol{T})$ *where* $\boldsymbol{T} \in \mathrm{Hom}(\mathbf{\Gamma}^+, \cdot)$ *is TFNP-hard,*

(2) $\mathrm{Hom}(\mathbf{\Gamma}^+, \cdot)$ *is undecidable, and*

(3) there exists $\boldsymbol{T} \in \mathrm{Hom}(\mathbf{\Gamma}^+, \cdot)$ *for which there is no computable homomorphism* $F : \mathbf{\Gamma}^+ \to \boldsymbol{T}$*.*

Again, in the non-computability result we consider the plain encoding of $\mathbf{\Gamma}^+$. Here we point out that the undecidability and non-computability parts of this result just follow from Proposition 4.2. Indeed, given a Σ_Γ-structure $\boldsymbol{T}$ it is easy to construct a Σ_{Γ^+}-structure $\boldsymbol{T}^+$ such that $\mathbf{\Gamma} \to \boldsymbol{T}$ if and only if $\mathbf{\Gamma}^+ \to \boldsymbol{T}^+$, and where a homomorphism $F : \mathbf{\Gamma} \to \boldsymbol{T}$ can be computed given oracle access to a homomorphism $H : \mathbf{\Gamma}^+ \to \boldsymbol{T}^+$. Indeed, $\boldsymbol{T}^+$ can be obtained extending $\boldsymbol{T}$ by interpreting each symbol $R \in \Sigma_{\Gamma^+} \backslash \Sigma_\Gamma$ as the total relation of arity $\mathrm{ar}(R)$ over T.

Finally, we need one last source of undecidability results, which will be given by a family of growing triangular slices of the two-dimensional grid. Given $m \in \mathbb{N}$, the structure $\boldsymbol{\nabla}_m$ has signature $\Sigma_\nabla = \{O, W, E_1, E_2\}$, where O, W are unary symbols and E_1, E_2 are binary. The universe ∇_m consists of all pairs $(n, o) \in \mathbb{N}^2$ satisfying $n + o \leqslant m$. The relations $O^{\nabla_m}, E_1^{\nabla_m}, E_2^{\nabla_m}$ are defined as in $\mathbf{\Gamma}$. That is, $O^{\nabla_m} = \{(1, 1)\}$, $E_1^{\nabla_m}$ consists of all pairs of the form $((n, o), (n + 1, o))$ and $E_2^{\nabla_m}$ contains the pairs $((m, o), (m, o + 1))$. Finally, the relation W^{∇_m} contains all pairs (n, o) satisfying $n + o = m$ (i.e., the upper-right boundary of the triangle).

Proposition 5.3. *Let* $(a_n)_{n\in\mathbb{N}}$ *be a strictly increasing sequence of natural numbers. Then the following families are undecidable (1)* $\mathrm{Hom}(\{\nabla_{a_n} \mid n \in \mathbb{N}\}, \cdot)$*, (2)* $\mathrm{Hom}_{\mathrm{eg}}(\{\nabla_{a_n} \mid n \in \mathbb{N}\}, \cdot)$*, and (3)* $\mathrm{Hom}_{\mathrm{io}}(\{\nabla_{a_n} \mid n \in \mathbb{N}\}, \cdot)$*.*

6 Main Definitions

6.1 Interpretations over Minions

We begin by introducing formally the notion of interpretation that we outlined in the proof sketch. Let $\mathscr{M}$ be a minion. Given a number $n \in \mathbb{N}$, a *n-ary predicate* over $\mathscr{M}$ is a subset $P \subseteq \mathscr{M}(n)$. We write $n = \mathrm{ar}(P)$. We write $2^{\mathscr{M}}$ for the set of predicates over $\mathscr{M}$ of arbitrary arity. Given a relational signature Σ, a Σ*-interpretation* $\mathcal{I}$ *over* $\mathscr{M}$ consists of (1) a predicate $U^{\mathcal{I}} \in 2^{\mathscr{M}}$, (2) a predicate $R^{\mathcal{I}} \in 2^{\mathscr{M}}$ for each symbol $R \in \Sigma$, and (3) a map $\Pi^{\mathcal{I}}_{R,i} : [\mathrm{ar}(R^{\mathcal{I}})] \mapsto [\mathrm{ar}(U^{\mathcal{I}})]$ for each symbol $R \in \Sigma$ and each index $i \in [\mathrm{ar}(R)]$. Interpretations over minions induce two kinds of structures, global and local, defined as follows.

- The *global structure induced by* $\mathcal{I}$, denoted $\mathbf{S} = \mathbf{S}_{\mathcal{I}}$, is a Σ-structure with universe $S = U^{\mathcal{I}}$, where for each symbol $R \in \Sigma$, a tuple $(f_1, \ldots, f_{\mathrm{ar}(R)}) \in (U^{\mathcal{I}})^{\mathrm{ar}(R)}$ belongs to $R^{\mathbf{S}}$ if there is an element $g \in R^{\mathcal{I}}$ satisfying that $f_i = g^{\pi_i}$ for each $i \in [\mathrm{ar}(R)]$, where $\pi_i = \Pi^{\mathcal{I}}_{R,i}$.
- Given an element $f \in \mathscr{M}$, we define the set $U^{\mathcal{I},f} \subseteq [\mathrm{ar}(U^{\mathcal{I}})]^{[\mathrm{ar}(f)]}$ as the subset of maps π for which $f^\pi \in U^{\mathcal{I}}$. The *local structure induced by* $\mathcal{I}$ *on f's minors*, denoted $\mathbf{S} = \mathbf{S}_{\mathcal{I},f}$, is a Σ-structure whose universe is $U^{\mathcal{I},f}$, where for each $R \in \Sigma$ the relation $R^{\mathbf{S}}$ consists of all the tuples of the form $(\sigma \circ \pi_1, \ldots, \sigma \circ \pi_{\mathrm{ar}(R)}) \in (U^{\mathcal{I},f})^{\mathrm{ar}(R)}$, where $f^\sigma \in R^{\mathcal{I}}$, and $\pi_i = \Pi^{\mathcal{I}}_{R,i}$ for each $i \in [\mathrm{ar}(R)]$.

It is important to remark that the definition of local structure appearing here is not exactly the same as the one given in the proof sketch. Indeed, the local structure $\boldsymbol{S}_{\mathcal{I},f}$ is *not* the induced substructure of $\boldsymbol{S}_{\mathcal{I}}$ on the minors of f. The reason is that it may be that $f^{\pi_1} = f^{\pi_2} \in U^{\mathcal{I}}$ for some suitable maps, whereas in the local structure $\boldsymbol{S}_{\mathcal{I},f}$ the maps π_1, π_2 represent different elements.

Although this is not relevant in the proofs of our algorithmic results, when we analyse cyclic and WNU polymorphisms we will deal with minions that correspond to disjoint unions of simpler minions. For instance, the minion representing the existence of a k-ary WNU polymorphism for each arity $k \geqslant 3$ is the disjoint union of the minion generated by an abstract 3-ary WNU element, with the one generated by a 4-ary WNU element, and so on. Hence, we also specialise some of our auxiliary results to disjoint unions of minions.

Given a subminion $\mathscr{N} \subseteq \mathscr{M}$, the *restricted interpretation* $\mathcal{J} = \mathcal{I}|_{\mathscr{N}}$ is defined by $U^{\mathcal{J}} = U^{\mathcal{I}} \cap \mathscr{N}$, $R^{\mathcal{J}} = R^{\mathcal{I}} \cap \mathscr{N}$ for each $R \in \Sigma$, and $\Pi^{\mathcal{J}}_{R,i} = \Pi^{\mathcal{I}}_{R,i}$ for each $R \in \Sigma$, $i \in [\mathrm{ar}(R)]$.

Observation 1. *Let J be a set, and $\mathcal{I}$ be a Σ interpretation over a disjoint union of minions $\mathscr{M} = \bigsqcup_{j \in J} \mathscr{M}_j$, and let $\mathcal{I}_j = \mathcal{I}|_{\mathscr{M}_j}$ for each $j \in J$. Then the following hold: (1) $\mathbf{S}_{\mathcal{I}} = \bigsqcup_{j \in J} \mathbf{S}_{\mathcal{I}_j}$, and (2) $\mathbf{S}_{\mathcal{I},f} = \mathbf{S}_{\mathcal{I}_j,f}$ for all $j \in J$, $f \in \mathscr{M}_j$.*

6.2 Manifold Minions

Now we are in position to introduce manifold minions, which is arguably the main construction that enables our proofs.

Let $\mathscr{M}$ be a minion, $\mathcal{I}$ be a Σ interpretation over it, and $\boldsymbol{C}$ be a Σ-structure. The *manifold minion* $\mathscr{N} = \boldsymbol{C}^{\mathcal{I}}$ is defined as follows. The elements in $\mathscr{N}(n)$ are pairs (f, χ), where $f \in \mathscr{M}(n)$, and $\chi : U^{\mathcal{I},f} \to C$ is a homomorphism from $\mathbf{S}_{\mathcal{I},f}$ to $\boldsymbol{C}$, which we call a *local homomorphism*. Given a map π, the minoring operation is given by $(f, \chi)^{\pi} = (f^{\pi}, \chi^{\pi})$, where χ^{π} is defined by $\gamma \mapsto \chi(\gamma \circ \pi)$. Observe that if $g = f^{\pi}$, then $U^{\mathcal{I},g} = \{\sigma \mid \sigma \circ \pi \in U^{\mathcal{I},f}\}$, so the homomorphism χ^{π} is well-defined. The *canonical projection* $\mathrm{pj} : \boldsymbol{C}^{\mathcal{I}} \rightharpoonup \boldsymbol{C}$ is the partial map defined on pairs (f, χ) such that $f \in U^{\mathcal{I}}$, which maps each such pair to the element $\chi(\mathrm{id})$, where $\mathrm{id} = \mathrm{id}_{[\mathrm{ar}(U^{\mathcal{I}})]}$ refers to the identity map over $[\mathrm{ar}(U^{\mathcal{I}})]$. Observe that $\chi(\mathrm{id})$ is well-defined. Indeed, the fact that $f \in U^{\mathcal{I}}$ implies that $\mathrm{id} \in \mathbf{S}_{\mathcal{I},f}$.

6.3 Descriptions

In this section we introduce a way to obtain quotient minions that preserve sufficient information about interpretations. Roughly, the main result here, Lemma 6.3, states that given a Σ-interpretation $\mathcal{I}$ over a minion $\mathscr{M}$, a Σ-structure $\boldsymbol{C}$, and a nice quotient $\mathscr{M}/\sim$ of $\mathscr{M}$, then $\mathscr{M} \to \boldsymbol{C}^{\mathcal{I}/\sim}$ if and only if the global structure $\boldsymbol{S}_{\mathcal{I}}$ maps homomorphically into $\boldsymbol{C}$.

A *description* of $\mathscr{M}$ is a set of predicates $\mathcal{D} \subseteq \mathscr{M}$. A description $\mathcal{D}$ induces an equivalence relation $\sim_{\mathcal{D}}$ on $\mathscr{M}$ as follows. Let $f_1, f_2 \in \mathcal{M}(n)$ for some $n \in \mathbb{N}$. Then $f_1 \sim_{\mathcal{D}} f_2$ means that $f_1^{\pi} \in P$ if and only if $f_2^{\pi} \in P$ for every $P \in \mathcal{D}$ and every map $\pi : [n] \to [\mathrm{ar}(P)]$. We shorten $\sim_{\mathcal{D}}$ to $\mathcal{D}$ when writing quotients to keep the notation light (i.e., we write $\cdot/\mathcal{D}$ instead of $\cdot/\sim_{\mathcal{D}}$). Given an element $f \in \mathscr{M}$, we write $\langle f \rangle$ to denote its equivalence class in $\mathscr{M}/\mathcal{D}$, and given a set $S \subseteq \mathscr{M}$ we write $\langle S \rangle$ for the set of equivalence classes $\langle f \rangle$ of elements $f \in S$. The quotient $\mathscr{M}/\mathcal{D}$ inherits a natural minion structure from $\mathscr{M}$: for every $n \in \mathbb{N}$, we define $(\mathscr{M}/\mathcal{D})(n) = \langle \mathscr{M}(n) \rangle$, and for every $\langle f \rangle \in \mathscr{M}/\mathcal{D}$ and every suitable map π, we define the minoring operation as $(\langle f \rangle)^{\pi} = \langle f^{\pi} \rangle$. Observe that this operation is well defined and does not depend on the chosen representative f. Moreover, if the description $\mathcal{D}$ is finite, then $\mathscr{M}/\mathcal{D}$ is locally finite.

Given a description $\mathcal{D} \subseteq 2^{\mathscr{M}}$, and a Σ-interpretation $\mathcal{I}$, the *quotient interpretation* $\mathcal{J} = \mathcal{I}/\mathcal{D}$ is the Σ interpretation over $\mathscr{M}/\mathcal{D}$ defined by $U^{\mathcal{J}} = \langle U^{\mathcal{I}} \rangle$, $R^{\mathcal{J}} = \langle R^{\mathcal{I}} \rangle$ for each $R \in \Sigma$, and $\Pi^{\mathcal{J}}_{R,i} = \Pi^{\mathcal{I}}_{R,i}$ for each $R \in \Sigma$, $i \in [\mathrm{ar}(R)]$.

We remark that by considering minion quotients corresponding to descriptions we are not loosing any generality. Indeed, if there is a surjective homomorphism $F : \mathscr{M} \to \mathscr{N}$ then $\mathscr{N}$ is isomorphic to $\mathscr{M}/\mathcal{D}$, where $\mathcal{D}$ is the description consisting of all the preimages $F^{-1}(g)$. So descriptions should be understood as a way to efficiently encode equivalence relations over minions. Indeed, given a description $\mathcal{D}$ of $\mathscr{M}$ consisting only of k-ary predicates, the equivalence relation $\sim_{\mathcal{D}}$ over $\mathscr{M}$ could have at most $k^n 2^{|\mathcal{D}|}$ classes of n-ary elements.

Internal Predicates The next step now is to formally introduce what it means for a description $\mathcal{D}$ of a minion $\mathscr{M}$ to preserve the meaningful information about an interpretation $\mathcal{I}$. More precisely, what we want is that the interpretation $\mathcal{I}$ is preserved, in some way, by all homomorphisms from $\mathscr{M}$ to the quotient $\mathscr{M}/\mathcal{D}$. The key definition here is that of an *internal predicate with respect to a description*. The description $\mathcal{D}$ will define a good quotient for $\mathcal{I}$ if all the predicates that form $\mathcal{I}$ are internal with respect to $\mathcal{D}$.

Let $\mathscr{M}$ be a minion and let $\mathcal{D} \subseteq 2^{\mathscr{M}}$ be a description. A predicate $Q \in 2^{\mathscr{M}}$ is called *$\mathcal{D}$-stable* if $f \sim_{\mathcal{D}} g$ together with $f \in Q$ imply that $g \in Q$ for any $f, g \in \mathscr{M}(\mathrm{ar}(Q))$. Given $h \geqslant \mathrm{ar}(Q)$, we say that Q is *internal at arity* h with respect to $\mathcal{D}$ if (1) Q is $\mathcal{D}$-stable, and (2) all partial homomorphisms $F : \mathscr{M} \xrightarrow{h} \mathscr{M}/\mathcal{D}$ satisfy $F(Q) \subseteq \langle Q \rangle$. Similarly, a Σ-interpretation $\mathcal{I}$ over $\mathscr{M}$ is $\mathcal{D}$-stable if all the predicates $U^{\mathcal{I}}$, $R^{\mathcal{I}}$ for $R \in \Sigma$ are $\mathcal{D}$-stable; and $\mathcal{I}$ is internal w.r.t. $\mathcal{D}$ at arty h if all the predicates $U^{\mathcal{I}}$, $R^{\mathcal{I}}$ for $R \in \Sigma$ are internal at arity h w.r.t. $\mathcal{D}$. In particular this requires that h is at least as large as $\mathrm{ar}(U^{\mathcal{I}})$ and $\mathrm{ar}(R^{\mathcal{I}})$ for each $R \in \Sigma$.

An *internal reference to* Q w.r.t. $\mathcal{D}$ is a pp-formula $\Phi(x) \in \mathcal{L}_{\mathrm{MC}}$ with one free variable x of arity $\mathrm{ar}(Q)$ satisfying

$$\mathscr{M}/\mathcal{D} \models \Phi(\langle f \rangle) \implies \langle f \rangle \in \langle Q \rangle$$

for all $\langle f \rangle \in \mathscr{M}(\mathrm{ar}(Q))/\mathcal{D}$. Our main tool for showing that a predicate is internal is the following result.

Lemma 6.1 (Main criterion for internal predicates)**.** *Let $\mathscr{M}$ be a minion $\mathcal{D} \subseteq 2^{\mathscr{M}}$ a description, $Q \in 2^{\mathscr{M}}$ a predicate that is $\mathcal{D}$-stable, and $h \in \mathbb{N}$ a number. Suppose that for each $f \in Q$ there is an internal reference $\phi_f(x)$ to Q w.r.t. $\mathcal{D}$ whose maximum arity is at most h, and such that $\mathscr{M} \models \phi_f(f)$. Then Q is internal at arity h w.r.t. $\mathcal{D}$*

Proof. Let $F : \mathscr{M} \xrightarrow{h} \mathscr{M}/\mathcal{D}$ be a partial homomorphism and $f \in Q$. Observe that $\mathscr{M} \models \phi_f(f)$ is witnessed by an assignment over elements of arity at most h, so this implies $\mathscr{M}/\mathcal{D} \models \phi_f(F(f))$. The fact that ϕ_f is an internal reference to Q w.r.t. $\mathcal{D}$ yields $F(f) \in \langle Q \rangle$. This proves the result. ☐

A situation in which this last criterion is especially easy to apply is that in which ϕ_f can be chosen to be the same for all $f \in Q$. This motivates the following notion. Let $\mathscr{M}$ be a minion, $\mathcal{D} \subseteq 2^{\mathscr{M}}$ a description and $Q \in 2^{\mathscr{M}}$ a predicate. An *internal definition of Q with respect to $\mathcal{D}$* is an internal reference $\Phi(x)$ to Q w.r.t. $\mathcal{D}$ that additionally satisfies

$$f \in Q \implies \mathscr{M} \models \Phi(f)$$

for all $f \in \mathscr{M}(\mathrm{ar}(Q))$. If such definition $\Phi(x)$ exists, its arity is bounded by a number $h \in \mathbb{N}$, and additionally Q is $\mathcal{D}$-stable, then Q is said to be *internally definable at arity h w.r.t. $\mathcal{D}$*. Observe that by Lemma 6.1, this implies that Q is internal at arity h w.r.t. $\mathcal{D}$.

Main Results About Quotient Interpretations The next two lemmas show that, given a minion $\mathscr{M}$, if $\mathcal{I}$ is an internal interpretation with respect to a description $\mathcal{D} \subseteq 2^{\mathscr{M}}$, then homomorphisms $\mathscr{M} \to \mathscr{M}/\mathcal{D}$ preserve the interpretation, and homomorphisms to power minions $\mathscr{M} \to \boldsymbol{C}^{\mathcal{I}/\mathcal{D}}$ correspond to homomorphisms from the global induced structure $\boldsymbol{S}_{\mathcal{I}}$ to $\boldsymbol{C}$.

Lemma 6.2 (Homomorphisms to good quotients preserve local structures)**.** *Let $\mathscr{M}$ be a minion, $\mathcal{I}$ a Σ-interpretation over $\mathscr{M}$, and $\mathcal{D} \subseteq 2^{\mathscr{M}}$ a description. Define $\mathcal{J} = \mathcal{I}/\mathcal{D}$. The following hold.*

(1) Suppose that $\mathcal{I}$ is $\mathcal{D}$-stable. Then $\mathbf{S}_{\mathcal{I},f} = \mathbf{S}_{\mathcal{J},\langle f\rangle}$ for all $f \in \mathscr{M}$.

(2) Suppose that $\mathcal{I}$ is internal at arity $h \in \mathbb{N}$ w.r.t. $\mathcal{D}$, and that $F : \mathscr{M} \xrightarrow{h} \mathscr{M}/\mathcal{D}$ is a partial homomorphism. Then for any $f \in \mathscr{M}$ whose arity is at most h, the local structure $\mathbf{S}_{\mathcal{I},f}$ is contained in $\mathbf{S}_{\mathcal{J},F(f)}$.

Proof. (1) This is a direct consequence of the stability condition. For any $f \in \mathscr{M}$ and any suitable map π it holds that $f^{\pi} \in U^{\mathcal{I}}$ if and only if $\langle f\rangle^{\pi} \in U^{\mathcal{J}}$, and $f^{\pi} \in R^{\mathcal{I}}$ if and only if $\langle f\rangle^{\pi} \in R^{\mathcal{J}}$ for each $R \in \Sigma$. This proves the statement.

(2) Let $\boldsymbol{A} = \mathbf{S}_{\mathcal{I},f}$, $\boldsymbol{B} = \mathbf{S}_{\mathcal{J},F(f)}$, and $n_U = \mathrm{ar}(U^{\mathcal{I}})$. First, we show that if a map $\sigma \in [n_U]^{[n]}$ belongs to A, then it also belongs to B. Indeed, the first condition is equivalent to $f^{\sigma} \in U^{\mathcal{I}}$. Because $\mathcal{I}$ is internal, it follows that $F(f)^{\sigma} \in \langle U^{\mathcal{I}}\rangle = U^{\mathcal{J}}$, which is equivalent to the second condition. Now, let $R \in \Sigma$, and $n_R = \mathrm{ar}(R^{\mathcal{I}})$. We show that if a tuple $\boldsymbol{\sigma}$ belongs to R^{A}, then $\boldsymbol{\sigma}$ belongs to R^{B} as well. The first condition means that there is a map $\gamma \in [n_R]^{[n]}$ such that $\sigma_i = \Pi^{\mathcal{I}}_{R,i}$ for each $i \in [\mathrm{ar}(R)]$, and $f^{\gamma} \in R^{\mathcal{I}}$. Because $\mathcal{I}$ is internal, it must hold that $F(f)^{\gamma} \in R^{\mathcal{J}}$, which shows that $\boldsymbol{\sigma}$ also belongs to R^{B}. □

We gently remind the reader that, given a manifold minion $\boldsymbol{C}^{\mathcal{J}}$ defined by a Σ-structure $\boldsymbol{C}$ and a Σ-interpretation $\mathcal{J}$, we write pj for the canonical projection from $\boldsymbol{C}^{\mathcal{J}}$ to $\boldsymbol{C}$, defined in Section 6.2.

Lemma 6.3 (Power minions on good quotients capture homomorphisms from the global structure)**.** *Let $\mathscr{M}$ be a minion, $\mathcal{I}$ a Σ-interpretation over $\mathscr{M}$, $\mathcal{D} \subseteq 2^{\mathscr{M}}$ a description, and $\boldsymbol{C}$ a Σ-structure. Define $\mathcal{J} = \mathcal{I}/\mathcal{D}$. Then the following hold.*

(1) Suppose that $\mathcal{I}$ is $\mathcal{D}$-stable and $\boldsymbol{S}_{\mathcal{I}} \to \boldsymbol{C}$. Then $\mathscr{M} \to \boldsymbol{C}^{\mathcal{J}}$.

(2) Suppose that $\mathcal{I}$ is internal w.r.t. $\mathcal{D}$ at arity $h \in \mathbb{N}$, and $F : \mathscr{M} \xrightarrow{h} \boldsymbol{C}^{\mathcal{J}}$ is a partial homomorphism. Then $\mathrm{pj} \circ F|_{U^{\mathcal{I}}}$ is a homomorphism from $\boldsymbol{S}$ to $\boldsymbol{C}$.

Proof. (1) Suppose there is a homomorphism $H : \boldsymbol{S}_{\mathcal{I}} \to \boldsymbol{C}$. Given an element $f \in \mathscr{M}$, the local homomorphism $\chi_f : \mathbf{S}_{\mathcal{I},f} \to \boldsymbol{C}$ is given by $\pi \mapsto H(f^{\pi})$ for each $\pi \in U^{\mathcal{I},f}$. Observe that if $g = f^{\pi}$, and $\sigma \in U^{\mathcal{I},g}$, then $\chi_g(\sigma) = \chi_f(\sigma \circ \pi)$, so the local homomorphisms we have defined are compatible with minoring. Now, by item (1) of Lemma 6.2, $\mathbf{S}_{\mathcal{I},f} = \mathbf{S}_{\mathcal{J},\langle f\rangle}$ for each $f \in \mathscr{M}$, so the map $f \mapsto (\langle f\rangle, \chi_f)$ is a minion homomorphism from $\mathscr{M}$ to $\boldsymbol{C}^{\mathcal{J}}$.

(2) Suppose there is a partial homomorphism $F : \mathscr{M} \xrightarrow{h} \boldsymbol{C}^{\mathcal{J}}$, given by $f \mapsto (p_f, \chi_f)$, and ρ be the canonical projection from $\boldsymbol{C}^{\mathcal{J}}$ to $\boldsymbol{C}$. We show that $H = \mathrm{pj} \circ F|_{U^{\mathcal{I}}}$ is a homomorphism from $\boldsymbol{S}_{\mathcal{I}}$

to $\boldsymbol{C}$. The map H sends every element $f \in U^{\mathcal{I}}$ to $\chi_f(\mathrm{id})$, where id denotes the identity map over $[\mathrm{ar}(U^{\mathcal{I}})]$. First, let us see that H is a well-defined map. Observe that the map $F' : \mathscr{M} \xrightarrow{h} \mathscr{M}/\mathcal{D}$ given by $f \mapsto p_f$ is a partial minion homomorphism. By item (2) of Lemma 6.2, $\mathbf{S}_{\mathcal{I},f} \subseteq \mathbf{S}_{\mathcal{J},p_f}$ for every $f \in \mathscr{M}(n)$, $n \leqslant h$. In particular, if $f \in U^{\mathcal{I}}$, then id belongs to $\mathbf{S}_{\mathcal{J},p_f}$, so $H(f) = \chi_f(\mathrm{id})$ is well-defined. Now let $R \in \Sigma$ and $(f_1, \dots, f_{\mathrm{ar}(R)}) \in R^{S_{\mathcal{I}}}$. We prove that $(H(f_1), \dots, H(f_{\mathrm{ar}(R)})) \in R^{C}$. Let $\pi_i = \Pi^{\mathcal{I}}_{R,i}$ for each $i \in [\mathrm{ar}(R)]$. By the definition of $\boldsymbol{S}_{\mathcal{I}}$, there must be an element $f_R \in R^{\mathcal{I}}$ such that $f_i = f_R^{\pi_i}$ for each $i \in [\mathrm{ar}(R)]$. In particular, this means that $(\pi_1, \dots, \pi_{\mathrm{ar}(R)}) \in R^{A}$, where $\boldsymbol{A} = \mathbf{S}_{\mathcal{I},f_R}$. Additionally, given $i \in [\mathrm{ar}(R)]$, the following chain of identities holds:

$$\chi_{f_R}(\pi_i) = \chi_{f_R}^{\pi_i}(\mathrm{id}) = \chi_{f_i}(\mathrm{id}) = H(f_i).$$

Hence, as χ_{f_R} is a homomorphism from $\mathbf{S}_{\mathcal{I},f_R}$ to $\boldsymbol{C}$, the tuple $(H(f_1), \dots, H(f_{\mathrm{ar}(R)}) = (\chi_{f_R}(\pi_1), \dots, \chi_{f_R}(\pi_{\mathrm{ar}(R)}))$ belongs to R^{C}, as we wanted to prove. ☐

Finally, we give another version of this last lemma that handles minions $\mathscr{M}$ that can be decomposed as disjoint unions of subminions.

Lemma 6.4. *Let K be a set, $\mathscr{M}$ a disjoint union $\bigsqcup_{k \in K} \mathscr{M}_k$ of subminions, $\mathcal{I}$ a Σ-interpretation over $\mathscr{M}$, $\mathcal{D} \subseteq 2^{\mathscr{M}}$ a description, and $\boldsymbol{C}$ a Σ-structure. Define $\mathcal{I}_k = \mathcal{I}|_{\mathscr{M}_k}$ for each $k \in K$, and $\mathcal{J} = \mathcal{I}/\mathcal{D}$. The following hold.*

1. *Suppose that $\mathcal{I}$ is $\mathcal{D}$-stable, and $\boldsymbol{S}_{\mathcal{I}_k} \to \boldsymbol{C}$ for some $k \in K$. Then $\mathscr{M}_k \to \boldsymbol{C}^{\mathcal{J}}$.*
2. *Suppose that $\mathcal{I}$ is internal at arity h w.r.t. $\mathcal{D}$, and $F : \mathscr{M}_k \xrightarrow{h} \boldsymbol{C}^{\mathcal{J}}$ is a partial homomorphism for some $k \in K$. Then $\mathrm{pj} \circ F|_{U^{\mathcal{I}_k}}$ is a homomorphism from $\boldsymbol{S}_{\mathcal{I}_k}$ to $\boldsymbol{C}$.*

Proof. The proof follows from the same arguments as the proof of Lemma 6.3, using Observation 1. The main insight is that, given a minion $\mathscr{N}$, partial homomorphisms $\mathscr{M} \xrightarrow{h} \mathscr{N}$ correspond to families $(F_k)_{k \in K}$ of partial homomorphisms $F_k : \mathscr{M}_k \xrightarrow{h} \mathscr{N}$ in a one-to-one fashion. If we let $\mathscr{N} = \mathscr{M}/\mathcal{D}$, the fact that $\mathcal{I}$ is internal at arity h w.r.t. $\mathcal{D}$, implies that for all $F_k : \mathscr{M}_k \xrightarrow{h} \mathscr{N}$ it must hold that $F_k(U^{\mathcal{I}_k}) \subseteq \langle U^{\mathcal{I}} \rangle$ and $F_k(R^{\mathcal{I}_k}) \subseteq \langle R^{\mathcal{I}} \rangle$ for all $R \in \Sigma$. ☐

6.4 Patterns

The ideas introduced in the previous sections are (almost) enough to prove our main undecidability and non-computability results. Indeed, consider a Σ-interpretation $\mathcal{I}$ over a minion $\mathscr{M}$ which is internal at arity k w.r.t. some description $\mathcal{D} \subseteq 2^{\mathscr{M}}$, and a Σ-structure $\boldsymbol{T}$. Then, Lemma 6.3 shows that there is a homomorphism from the global structure $\boldsymbol{S}_{\mathcal{I}}$ to $\boldsymbol{T}$ if and only if $\mathscr{M} \to \boldsymbol{T}^{\mathcal{I}/\mathcal{D}}$. Moreover, it is not difficult to show that it is possible to compute a homomorphism $F : \boldsymbol{S}_{\mathcal{I}} \to \boldsymbol{T}$ when given oracle access to a homomorphism $H : \mathscr{M} \to \boldsymbol{T}^{\mathcal{I}/\mathcal{D}}$. What is missing in this picture is a way to reduce, assuming that $\boldsymbol{S}_{\mathcal{I}} \to \boldsymbol{T}$, the problem $\mathrm{sPCSP}(\boldsymbol{S}_{\mathcal{I}}, \boldsymbol{T})$ to $\mathrm{sPMC}_h(\mathscr{M}, \boldsymbol{T}^{\mathcal{I}/\mathcal{D}})$ for some $h \in \mathbb{N}$. This reduction will be given by what we call a *pattern*, which is a polynomial-time algorithm that constructs minor conditions related to input Σ-structures, in a way similar to the formulas $\Psi_{\boldsymbol{G}}$ defined in (3) during our proof sketch. The formal definition is given below.

Let $\mathscr{M}$ a minion, $h \in \mathbb{N}$ an integer, $\mathcal{D} \subseteq 2^{\mathscr{M}}$ a description, and $\mathcal{I}$ a $\mathcal{D}$-stable Σ-interpretation over $\mathscr{M}$. A *pattern of h-ary internal references to $\mathcal{I}$ w.r.t. $\mathcal{D}$* is a function Ψ computable in polynomial time that sends every finite structure $\boldsymbol{I}$ satisfying $\boldsymbol{I} \to \boldsymbol{S}_{\mathcal{I}}$ to a minor condition $\Psi_{\boldsymbol{I}}$ of arity at most h, such that $\mathscr{M} \models \Psi_{\boldsymbol{I}}$, and that can be written as

$$\exists_{v \in I} x_v \exists_{R \in \Sigma, r \in R^I} x_r \left(\bigwedge_{v \in I} \psi_v(x_v) \right) \bigwedge \left(\bigwedge_{R \in \Sigma, r \in R^I} \psi_r(x_r) \bigwedge_{i \in [\mathrm{ar}(R)]} x_r^{\Pi^{\mathcal{I}}_{R,i}} = x_{r(i)} \right),$$

where each $v \in I$ the formula $\psi_v(x)$ is an internal reference to $U^{\mathcal{I}}$ w.r.t. $\mathcal{D}$, and for each $R \in \Sigma$ and each $r \in R^I$, the formula $\psi_r(x)$ is an internal reference to $R^{\mathcal{I}}$ w.r.t. $\mathcal{D}$. To abbreviate, we will refer to the tuple $(\mathcal{I}, \mathcal{D}, \Psi, h)$ as a Σ-*pattern* over $\mathscr{M}$.

Lemma 6.5. *Let $(\mathcal{I}, \mathcal{D}, \Psi, h)$ be a pattern over a minion $\mathscr{M}$, and $\boldsymbol{T}$ a finite structure satisfying $\boldsymbol{S}_{\mathcal{I}} \to \boldsymbol{T}$. Then there is a many-one reduction from* $\mathrm{sPCSP}(\boldsymbol{S}_{\mathcal{I}}, \boldsymbol{T})$ *to* $\mathrm{sPMC}_h(\mathscr{M}, \boldsymbol{T}^{\mathcal{I}/\mathcal{D}})$.

Proof. Let $\mathcal{J} = \mathcal{I}/\mathcal{D}$, and $\mathscr{N} = \boldsymbol{T}^{\mathcal{J}}$. Observe that by Lemma 6.3, the fact that $\mathcal{I}$ is $\mathcal{D}$-stable implies that $\mathscr{M} \to \boldsymbol{T}^{\mathcal{J}}$, so $\mathrm{sPMC}_h(\mathscr{M}, \mathscr{N})$ is well-defined.

We give a many-one reduction from $\mathrm{sPCSP}(\boldsymbol{G}, \boldsymbol{T})$ to $\mathrm{sPMC}_h(\mathscr{M}, \mathscr{N})$. This reduction consists of a pair of polynomial-time computable functions (α, β), where (1) α maps structures $\boldsymbol{I}$ satisfying $\boldsymbol{I} \to \boldsymbol{G}$ to minor conditions ψ of arity at most h that satisfy $\mathscr{M} \models \psi$, and (2) β maps pairs $(\boldsymbol{I}, F)$ to homomorphisms $H : \boldsymbol{I} \to \boldsymbol{T}$, where $\boldsymbol{I}$ is a structure satisfying $\boldsymbol{I} \to \boldsymbol{G}$ and F is an assignment satisfying $\alpha(\boldsymbol{I})$ over $\mathscr{N}$. We simply define the function α as the map $\boldsymbol{I} \mapsto \Psi_{\boldsymbol{I}}$. Now consider an assignment F that satisfies $\Psi_{\boldsymbol{I}}$ over $\mathscr{N}$, and let pj be the canonical projection form $\mathscr{N}$ to $\boldsymbol{T}$. Let $H_F : \boldsymbol{I} \to \boldsymbol{T}$ be the map that sends each $v \in I$ to $\mathrm{pj}(F(x_v))$, which is clearly computable in polynomial time. We claim that H_F is a homomorphism. This way, we can define β as $(\boldsymbol{I}, F) \mapsto H_F$. Let us show that H_F is indeed a homomorphism. Recall the structure of the minor condition $\Psi_{\boldsymbol{I}}$. We can write

$$\Psi_{\boldsymbol{I}} \equiv \exists_{v \in I} x_v \exists_{R \in \Sigma, r \in R^I} x_r \left(\bigwedge_{v \in I} \psi_v(x_v) \right) \bigwedge$$
$$\left(\bigwedge_{R \in \Sigma, r \in R^I} \psi_r(x_r) \bigwedge_{i \in [\mathrm{ar}(R)]} x_r^{\Pi^{\mathcal{I}}_{R,i}} = x_{r(i)} \right),$$

where $\psi_v(x)$ is an internal reference (w.r.t. $\mathcal{D}$) to $U^{\mathcal{I}}$ for each $v \in I$, and $\psi_r(x)$ is an internal reference to $R^{\mathcal{I}}$ for each $R \in \Sigma$, $r \in R^I$.

We suppose that F maps $v \mapsto (\langle p_v \rangle, \chi_v)$ for each $v \in I$, and $r \mapsto (\langle p_r \rangle, \chi_r)$ for each $R \in \Sigma$, and tuple $r \in R^I$. First we argue that H_F is a well-defined map. Let $v \in I$. Then $\mathscr{M}/\mathcal{D} \models \psi_v(\langle p_v \rangle)$, which means that $\langle p_v \rangle \in U^{\mathcal{J}}$, by the definition of internal reference. Hence $H_F(v) = \mathrm{pj}(F(x_v)) = \chi_v(\mathrm{id})$ is well defined. Now we show that H_F is a homomorphism. Let $R \in \Sigma$ and $r = (v_1, \ldots, v_{\mathrm{ar}(R)}) \in R^I$. We have that $\mathscr{M}/\mathcal{D} \models \psi_r(\langle p_r \rangle)$, so $\langle p_r \rangle \in R^{\mathcal{J}}$. It also holds that $F(x_r)^{\pi_i} = F(x_{v_i})$ for each $i \in [\mathrm{ar}(R)]$, where $\pi_i = \Pi^{\mathcal{J}}_{R,i}$. This means that $\chi_r(\pi_i) = \chi_{v_i}(\mathrm{id}) = H_F(v_i)$ for each $i \in [\mathrm{ar}(R)]$. Observe that $(\pi_1, \ldots, \pi_{\mathrm{ar}(R)}) \in R^{\boldsymbol{A}}$, where $\boldsymbol{A} = \mathbf{S}_{\mathcal{J}, \langle p_r \rangle}$. Indeed, we have established that $\langle p_r \rangle \in R^{\mathcal{J}}$, and it holds that $\pi_i = \pi_i \circ \mathrm{id}$ for each $i \in [\mathrm{ar}(R)]$. Using that χ_r is a homomorphism from $\mathbf{S}_{\mathcal{J}, \langle p_r \rangle}$ to $\boldsymbol{T}$, we obtain that $(\chi_r(\pi_1), \ldots, \chi_r(\pi_{\mathrm{ar}(R)})) = (H_F(v_1), \ldots, H_F(v_{\mathrm{ar}(R)}) \in R^{\boldsymbol{T}}$. This proves that H_F is a homomorphism. $\square$

7 From Minions to Templates

In this section, given a locally finite minion $\mathscr{M}$ and a number $h \in \mathbb{N}$, we obtain a finite template $(\boldsymbol{A}, \boldsymbol{B})$ satisfying that $\operatorname{Pol}(\boldsymbol{A}, \boldsymbol{B})$ has a partial isomorphism to $\mathscr{M}$ up to arity h and that $\mathscr{N} \xrightarrow{h} \mathscr{M}$ if and only if $\mathscr{N} \to \operatorname{Pol}(\boldsymbol{A}, \boldsymbol{B})$ for any other minion $\mathscr{N}$. Furthermore, we want to do so while keeping $\boldsymbol{A}$ small, controlling both the size of its universe A and the size of its relations.

In order to state the main result of this section we need to define one more notion. Let $\mathscr{M}$ be a minion. The *rank* of $\mathscr{M}$ is the smallest number $r \in \mathbb{N}$ satisfying that, for any $n \in \mathbb{N}$, whenever two elements $f_1, f_2 \in \mathscr{M}(n)$ have the same r-ary minors (i.e., $f_1^\pi = f_2^\pi$ for all $\pi \in [r]^{[n]}$), then $f_1 = f_2$. We say that $\mathscr{M}$ has infinite rank if no such r exists. Our goal is to prove the following.

Theorem 7.1. *Let $\mathscr{M}$ be a minion whose rank is at most r, and let $h \geqslant r$. Then there is a template $(\boldsymbol{A}, \boldsymbol{B})$ where $|A| = r$ and $|R^A| \leqslant h$ for all relations in $\boldsymbol{A}$ such that*

(1) There is a partial minion isomorphism $\mathscr{M} \xrightarrow{h} \operatorname{Pol}(\boldsymbol{A}, \boldsymbol{B})$ defined up to arity h.

(2) Given a minion $\mathscr{M}'$, it holds that $\mathscr{M}' \to \operatorname{Pol}(\boldsymbol{A}, \boldsymbol{B})$ if and only if $\mathscr{M}' \xrightarrow{h} \mathscr{M}$.

Moreover, if $\mathscr{M}$ is locally finite, then $\mathbf{B}$ is a finite structure.

Through similar ideas to the ones used to prove this result, we are also able to characterize abstract minions which are isomorphic to the polymorphism minion of some finite template. This characterization is not essential to show the main results of this paper, and we defer it to Section B.

In our reductions we apply Theorem 7.1 to manifold minions built using quotient interpretations. We will need the following bound on the rank of such minions.

Lemma 7.2. *Let $\boldsymbol{T}$ be a Σ-structure, $\mathscr{M}$ a minion, $\mathcal{D} \subseteq 2^{\mathscr{M}}$, $\mathcal{I}$ a Σ-interpretation over $\mathscr{M}$, and $r \in \mathbb{N}$. Suppose that*

(1) $r \geqslant \operatorname{ar}(Q)$ for all $Q \in \mathcal{D}$, and

(2) $r \geqslant \operatorname{ar}(U^{\mathcal{I}})$.

Then the rank of $\boldsymbol{T}^{\mathcal{I}/\mathcal{D}}$ is at most r.

Proof. Let $f, g \in \mathscr{M}(n)$ for some $n \in \mathbb{N}$, and suppose that $f^\pi \sim_{\mathcal{D}} g^\pi$ for all $\pi \in [r]^{[n]}$. We show this implies $f \sim_{\mathcal{D}} g$. By Lemma 7.3 this implies that the rank of $\mathscr{M}/\mathcal{D}$ is at most r. Indeed, suppose that $f \not\sim_{\mathcal{D}} g$. Then, without loss of generality we may assume there is some $Q \in \mathcal{D}$ and some $\pi \in [\operatorname{ar}(Q)]^{[n]}$ such that $f^\pi \in Q$ but $g^\pi \notin Q$. Let $\alpha : [\operatorname{ar}(Q)] \to [r]$ and $\beta : [r] \to [\operatorname{ar}(Q)]$ be such that $\beta \circ \alpha = \operatorname{id}_{[\operatorname{ar}(Q)]}$. Such maps exist because $\operatorname{ar}(Q) \leqslant r$. By assumption, $f^{\alpha\circ\pi} = g^{\alpha\circ\pi}$. However $f^\pi = (f^{\alpha\circ\pi})^\beta$, and $g^\pi = (g^{\alpha\circ\pi})^\beta$, a contradiction.

Now define $\mathcal{J} = \mathcal{I}/\mathcal{D}$, and $\mathscr{N} = \boldsymbol{T}^{\mathcal{I}/\mathcal{D}}$. Let $(\langle f\rangle, \chi_f), (\langle g\rangle, \chi_g) \in \mathscr{N}(n)$ for some $n \in \mathbb{N}$. Suppose that

$$(\langle f\rangle, \chi_f)^\pi = (\langle g\rangle, \chi_g)^\pi \quad \text{for all } \pi \in [r]^{[n]}.$$

Then by the previous arguments $\langle f\rangle = \langle g\rangle$, so the local structure $\mathbf{S}_{\mathcal{J},\langle f\rangle}$ and $\mathbf{S}_{\mathcal{J},\langle g\rangle}$ are the same. Now, suppose, for the sake of a contradiction, that there is a map $\sigma \in U^{\mathcal{J},\langle f\rangle}$ such that $\chi_f(\sigma) \neq$

$\chi_g(\sigma)$. Let $\alpha : [\mathrm{ar}(U^{\mathcal{I}})] \to [r]$ and $\beta : [r] \to [\mathrm{ar}(U^{\mathcal{I}})]$ be such that $\beta \circ \alpha = \mathrm{id}_{[\mathrm{ar}(U^{\mathcal{I}})]}$ (observe that $\mathrm{ar}(U^{\mathcal{I}}) = \mathrm{ar}(U^{\mathcal{J}})$). Then it must hold that

$$\chi_f^{\alpha \circ \sigma}(\beta) = \chi_f(\sigma) \neq \chi_g(\sigma) = \chi_g^{\alpha \circ \sigma}(\beta).$$

However, $\alpha \circ \sigma \in [r]^{[n]}$, yielding a contradiction. This completes the proof. $\square$

Finite rank minions are precisely those which are isomorphic to a function minion. Indeed, any function minion, as defined in [18], on a domain D must have rank at most $|D|$. In the other direction, if $\mathscr{M}$ has rank r, and $p \in \mathscr{M}$ is a n-ary element, then it can be seen as a function from $[r]^n$ to $\mathscr{M}(r)$ by letting $p(\pi) = p^\pi$ for each $\pi \in [r]^{[n]}$ (recall that we identify tuples and maps). The following is a useful auxiliary fact.

Lemma 7.3. *Let $\mathscr{M}$ be a minion, $n \in \mathbb{N}$, and $f, g \in \mathscr{M}(n)$. Let $r \leqslant h$ be natural numbers. Suppose that $f^\pi = g^\pi$ for every $\pi \in [h]^{[n]}$. Then $f^\pi = g^\pi$ for every $\pi \in [r]^{[n]}$.*

Proof. Let $\alpha : [r] \to [h]$ and $\beta : [h] \to [r]$ be such that $\beta \circ \alpha = \mathrm{id}_{[r]}$. Then for any $\pi \in [r]^{[n]}$ it holds that

$$f^\pi = f^{\beta \circ (\alpha \circ \pi)} = g^{\beta \circ (\alpha \circ \pi)} = g^\pi,$$

as we wanted to show. $\square$

Minion Closures Let A and B be finite sets, and consider an arbitrary minion $\mathscr{M}$ whose n-ary elements are functions of the form $f : A^n \to B$ where minoring is defined as usual. One of the insights of [18] is that a minion of this kind is a polymorphism minion of a finite template if and only if it has bounded *finitised arity*. In other words, there is a number $h \in \mathbb{N}$ such that for any $n \in \mathbb{N}$ and any function $f : A^n \to B$ it holds that $f \in \mathscr{M}$ if all h-ary minors of f belong to $\mathscr{M}$. Roughly, this means that every consistent family of h-ary elements in $\mathscr{M}$ occurs as the family of h-ary minors of a higher arity element. Conversely, in the case where $\mathscr{M}$ is not isomorphic to the polymorphism minion of a finite template, it must be that some consistent family of h-ary elements in $\mathscr{M}$ is not represented by a higher arity element. In a way, this means that $\mathscr{M}$ is missing some elements. To prove Theorem 7.1, we first define the h-closure of the minion $\mathscr{M}$ which results from adding to it an element representing each consistent family of h-ary minors. Then, we show how to construct a finite template $(\boldsymbol{A}, \boldsymbol{B})$ such that $\mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$ is isomorphic to this closure.

An *m-ary system of k-ary minors* over $\mathscr{M}$ is a map $\zeta : [k]^{[m]} \to \mathscr{M}(k)$ satisfying that for any pair of maps $\pi_1, \pi_2 \in [k]^{[m]}$ and any map $\sigma \in [k]^{[k]}$ for which $\pi_1 = \sigma \circ \pi_2$ it holds that $\zeta(\pi_2)^\sigma = \zeta(\pi_1)$. If ζ is an m-ary system, and $\sigma \in [n]^{[m]}$ is a map, we denote by ζ^σ the n-ary system corresponding to the map $\pi \mapsto \zeta(\pi \circ \sigma)$. The h-closure of a minion $\mathscr{M}$ is another minion $\mathscr{M}^{(h)}$ whose n-ary elements are the n-ary systems of h-ary minors over $\mathscr{M}$, and where minoring is given by the operation $\zeta \mapsto \zeta^\pi$.

Any element $p \in \mathscr{M}(n)$ defines a n-ary system $\zeta_p \in \mathscr{M}^{(h)}(n)$ in a natural way. That is, for any $\pi : [n] \to [h]$ we set $\zeta_p(\pi) = p^\pi$. This mapping $p \mapsto \zeta_p$ is, in fact, a minion homomorphism $\mathrm{Cl}_h : \mathscr{M} \to \mathscr{M}^{(h)}$, which we call the *canonical homomorphism* from $\mathscr{M}$ to $\mathscr{M}^{(h)}$.

Now we prove that the h-closure of a minion $\mathscr{M}$ satisfies both properties we require from $\mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$ in Theorem 7.1.

Proposition 7.4. *Let $\mathscr{M}$ be a minion and let $h \in \mathbb{N}$. Then the restriction of the canonical homomorphism $\mathrm{Cl}_h : \mathscr{M} \to \mathscr{M}^{(h)}$ to elements of arity at most h is a partial isomorphism.*

Proof. Let F be the restriction of Cl_h to elements of arity at most h. We define another partial homomorphism $H : \mathscr{M}^{(h)} \xrightarrow{h} \mathscr{M}$ and show that F and H are inverses. For each $k \leqslant h$ we fix maps $\sigma_k : [k] \to [h]$ and $\pi_k : [h] \to [k]$ satisfying $\pi_k \circ \sigma_k = \mathrm{id}_{[k]}$. Then, given a k-ary element $\zeta \in \mathscr{M}^{(h)}$, we define $H(\zeta)$ as $(\zeta(\sigma_k))^{\pi_k}$.

Given $n \leqslant h$ and an element $f \in \mathscr{M}(n)$, we define $F(f)$ as the n-ary system ζ_f of h-ary minors that maps each function $\pi \in [h]^{[n]}$ to f^{π}. The map F defined this way is clearly a h-partial minion homomorphism. Now let us show that F is injective. Given $n \leqslant h$, it is possible to find maps $\pi : [n] \to [h]$ and $\sigma : [h] \to [n]$ such that $\sigma \circ \pi = \mathrm{id}_{[n]}$. Hence $f^{\pi} = g^{\pi}$ implies $f = g$ for all pairs $f, g \in \mathscr{M}(n)$. In particular, this means that $\zeta_f = \zeta_g$ if and only if $f = g$. Finally, let us prove that F is surjective. Let $n \leqslant h$ and let π, σ be the same maps as before. Consider an arbitrary n-ary system ζ of h-ary minors over $\mathscr{M}$. We claim that $\zeta = \zeta_f$, where $f = \zeta(\pi)^{\sigma}$. To prove this we need to show that $\zeta(\pi') = f^{\pi'}$ for all $\pi : [n] \to [h]$. Observe that $\pi' = \pi' \circ \sigma \circ \pi$. Thus, by the definition of system, $\zeta(\pi') = \zeta(\pi)^{\pi' \circ \sigma}$. However, $\zeta(\pi)^{\pi' \circ \sigma} = f^{\pi'}$, proving that $\zeta = \zeta_f$. □

Proposition 7.5. *Let $\mathscr{M}$ and $\mathscr{N}$ be minions, and $h \in \mathbb{N}$ be a number. Then $\mathscr{M} \xrightarrow{h} \mathscr{N}$ if and only if $\mathscr{M} \to \mathscr{N}^{(h)}$.*

Proof. Suppose there is a minion homomorphism $F : \mathscr{M} \to \mathscr{N}^{(h)}$. Then the restriction of F to elements of arity at most h yields a partial homomorphism from $\mathscr{M}$ to $\mathscr{N}^{(h)}$ up to arity h. By Proposition 7.4, this implies there is a partial homomorphism from $\mathscr{M}$ to $\mathscr{N}$ up to arity h.

Now suppose there is a partial homomorphism $F : \mathscr{M} \xrightarrow{h} \mathscr{N}$. We use F to define a minion homomorphism $F' : \mathscr{M} \to \mathscr{N}^{(h)}$. Let $n \in \mathbb{N}$, and let $f \in \mathscr{M}(n)$. Then we define $F'(f)$ to be the system ζ_f that sends each map $\pi \in [h]^{[n]}$ to $F(f^{\pi})$. The system ζ_f is well-defined: if $\pi_1 = \sigma \circ \pi_2$ for some maps $\pi_1, \pi_2 \in [h]^{[n]}$, $\sigma \in [h]^{[h]}$, then $\zeta_f(\pi_1) = F(f^{\sigma \circ \pi_2}) = F(f^{\pi_2})^{\sigma} = \zeta_f(\pi_2)^{\sigma}$. Finally, the map F' is a minion homomorphism. Indeed, if $f = g^{\pi}$, then $\zeta_f = \zeta_g^{\pi}$ following the definition of minoring for systems.

□

Finite Templates We introduce two kinds of structures that will be used in our templates. Let $h \geqslant r$ be two natural numbers. Let $m = r^h$, and let $\pi_1, \dots \pi_m$ be the lexicographical ordering of $[r]^{[h]}$. The *complete structure* $\boldsymbol{K}_r^h$ is the relational structure whose signature consists of a single m-ary symbol R, whose universe is $[r]$, and where $R^{\boldsymbol{K}_r^h}$ is defined as the set of tuples $(\pi_1(i), \dots, \pi_m(i))$, where $i \in [h]$. This construction was given without a name in [18, Lemma 6.7]. This is the most general structure on r elements with relations of size at most h, in the sense that any other such structure is pp-definable on $\boldsymbol{K}_r^h$.

The second kind of structures we use are the so-called *free structures*, introduced in [11]. Let $\mathscr{M}$ be a minion and let $\boldsymbol{A}$ be a Σ-structure. Let $n = |A|$ and identify $A = [n]$ in some fixed way. Similarly, for each $R \in \Sigma$, let $m_R = |R^{\boldsymbol{A}}|$, and identify $R^{\boldsymbol{A}}$ with m_R in a fixed way. The *free structure of $\mathscr{M}$ generated by $\boldsymbol{A}$* is a Σ-structure, denoted $\boldsymbol{F} = \boldsymbol{F}_{\mathscr{M}}(\boldsymbol{A})$ has universe $F = \mathscr{M}(n)$, and for each symbol $R \in \Sigma$ the relation $R^{\boldsymbol{F}}$ is given by the set of tuples $(f_1, \dots, f_{\mathrm{ar}(R)})$ for which there is an element $g \in \mathscr{M}(m)$ satisfying $g^{\pi_i} = f_i$ for each $i \in [\mathrm{ar}(R)]$, where $\pi_i : [m] \to [n]$ is the i-th projection $r \mapsto r(i)$ (recall that $r \in R^{\boldsymbol{A}}$ is seen as an element of m, and $r(i) \in A$ as an element of $[n]$).

Having defined complete structures and free structures, we are in conditions to prove the main result of the section.

Proof of Theorem 7.1. The template witnessing the statement is given by $\boldsymbol{A} = \boldsymbol{K}_r^h$ and $\boldsymbol{B} = \boldsymbol{F}_{\mathscr{M}}(\boldsymbol{A})$. Observe that if $\mathscr{M}$ is locally finite, then $\boldsymbol{B}$ is a finite structure. We prove that $\operatorname{Pol}(\boldsymbol{A}, \boldsymbol{B})$ is isomorphic to the closure $\mathscr{M}^{(h)}$. Observe that this proves the theorem: the first item follows then from Proposition 7.4, and the second from Proposition 7.5.

Let $\mathscr{N} = \operatorname{Pol}(\boldsymbol{A}, \boldsymbol{B})$, $m = r^h$, and $\pi_1, \dots, \pi_m$ be the lexicographical ordering of $[r]^{[h]}$. We see elements $p \in \mathscr{N}(n)$ as maps from $[r]^{[n]}$ to $\mathscr{M}(r)$ by identifying $[r]^n$ with $[r]^{[n]}$ as usual. Given a map $\sigma \in [h]^{[n]}$, the elements $(p(\pi_1 \circ \sigma), \dots, p(\pi_m \circ \sigma))$ must belong to the relation R^B. In other words, there must be some element $f \in \mathscr{M}(h)$ such that $f^\pi = p(\pi \circ \sigma)$ for each $\pi \in [r]^{[h]}$. Moreover, as the rank of $\mathscr{M}$ is bounded by r, there can be only one element f with such property. We denote f as $p^*(\sigma)$.

Now we are able to define an isomorphism $H : \mathscr{N} \to \mathscr{M}^{(h)}$. Given $n \in \mathbb{N}$ and a polymorphism $p \in \mathscr{N}(n)$, we define $H(p)$ as the system ζ_p that sends each map $\sigma \in [h]^{[n]}$ to the element $p^*(\sigma) \in \mathscr{M}(h)$. Let us show that ζ_p is a well-defined system. Let $\sigma \in [h]^{[n]}$, $\gamma \in [h]^{[h]}$, $f = p^*(\sigma) = \zeta_p(\sigma)$, and $g = p^*(\gamma \circ \sigma) = \zeta_p(\gamma \circ \sigma)$. We need to show that $f^\gamma = g$. Given $\pi \in [r]^{[h]}$, using the definition of p^* and the fact that $\pi \circ \gamma \in [r]^{[h]}$ we obtain

$$(f^\gamma)^\pi = f^{\pi \circ \gamma} = p((\pi \circ \gamma) \circ \sigma) = p(\pi \circ (\gamma \circ \sigma)) = g^\pi.$$

As the rank of $\mathscr{M}$ is at most r, Lemma 7.3 implies $f^\gamma = g$, as we wanted.

Now let us prove that H is a minion homomorphism. Let $p \in \mathscr{N}(n_1)$ be a polymorphism, and let $q = p^\gamma$ where $\gamma \in [n_2]^{[n_1]}$. We need to show that $\zeta_q = \zeta_p^\gamma$. Let $\sigma \in [h]^{[n_2]}$, be an arbitrary map, $g = \zeta_q(\sigma)$, and $f = \zeta_p^\gamma(\sigma) = \zeta_p(\sigma \circ \gamma)$. It is enough to prove that $f = g$. For each $\pi \in [r]^{[h]}$ we have that $f^\pi = p(\pi \circ \sigma \circ \gamma) = q(\pi \circ \sigma) = g^\pi$. Using that the rank of $\mathscr{M}$ is at most r, this shows that $f = h$, completing the proof.

Now let us prove that H is a bijective map. This is enough because the inverse map of a bijective minion homomorphism is also a minion homomorphism. As $r \leqslant h$, there are two maps $\sigma \in [h]^{[r]}$ and $\gamma \in [r]^{[h]}$ such that $\mathrm{id}_{[r]} = \gamma \circ \sigma$. To see that H is injective, let $p, q \in \mathscr{N}(n)$ be two different polymorphisms. Using that the rank of $\mathscr{N}$ is at most r, Lemma 7.3 yields some $\tau \in [r]^{[n]}$ such that $p(\tau) \neq q(\tau)$. But $p(\tau) = \zeta_p(\sigma \circ \tau)^\gamma$, and $q(\tau) = \zeta_q(\sigma \circ \tau)^\gamma$, so $H(p) = \zeta_p \neq H(q) = \zeta_q$. Finally, we show that H is surjective. Let $\zeta \in \mathscr{M}^{(h)}(n)$ be an arbitrary system. We need to prove that there is some polymorphism $p \in \mathscr{N}(n)$ such that $\zeta = \zeta_p$. Define $p : [r]^{[n]} \to \mathscr{M}(r)$ as the function that maps each $\tau \in [r]^{[n]}$ to $\zeta(\sigma \circ \tau)^\gamma$. We need to show that p is a n-ary polymorphism in $\mathscr{N}$. For this it is enough to show that $(p(\pi_1 \circ \tau'), \dots, p(\pi_m \circ \tau')) \in R^B$ for each map $\tau' \in [h]^{[n]}$. Indeed, $p(\pi \circ \tau') = \zeta(\sigma \circ \pi \circ \tau')^\gamma$ for all $\pi \in [r]^{[h]}$. However, by the definition of system, this last element equals $\zeta(\tau')^{\gamma \circ (\sigma \circ \pi)} = \zeta(\tau')^\pi$, where the last equality uses the fact that $\gamma \circ \sigma = \mathrm{id}_{[r]}$. The fact that $p(\pi \circ \tau') = \zeta(\tau')^\pi$ for all $\pi \in [r]^{[h]}$ implies that $(p(\pi_1 \circ \tau'), \dots, p(\pi_m \circ \tau')) \in R^B$. To see that $\zeta = \zeta_p$, observe that $p^*(\tau')$ must equal $\zeta(\tau')$ for all $\tau' \in [h]^{[n]}$. ☐

Rather than referencing this result, it will be more convenient to use the following more concrete corollary, which simply follows from the fact that the template constructed in the last proof is precisely $(\boldsymbol{K}_r^h, \boldsymbol{F}_{\mathscr{M}}(\boldsymbol{K}_r^h))$.

Corollary 7.6. *Let $\mathscr{M}$ be a minion whose rank is at most r, and let $h \geqslant r$. Then the template $(\boldsymbol{A}, \boldsymbol{B}) = (\boldsymbol{K}_r^h, \boldsymbol{F}_{\mathscr{M}}(\boldsymbol{K}_r^h))$ satisfies the following.*

(1) *There is a partial minion isomorphism $\mathscr{M} \xrightarrow{h} \operatorname{Pol}(\boldsymbol{A}, \boldsymbol{B})$ defined up to arity h.*

(2) Given a minion $\mathscr{M}'$, it holds that $\mathscr{M}' \to \mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$ if and only if $\mathscr{M}' \xrightarrow{h} \mathscr{M}$.

Moreover, if $\mathscr{M}$ is locally finite, then $\boldsymbol{B}$ is a finite structure.

8 Main Reductions

In this section we show the reductions that enable our main results. Most proofs are short, as they are just a matter of putting together the pieces we have constructed up until now. As in the overview given in Section 4, the general strategy can be described as follows. We start with an interpretation $\mathcal{I}$ over a minion $\mathscr{M}$ which is internal at some arity h w.r.t. a description $\mathcal{D}$. Then we are able to reduce problems related to the global structure $\boldsymbol{S}_{\mathcal{I}}$ to problems related to $\mathscr{M}$ by considering manifold minions of the form $\boldsymbol{T}^{\mathcal{I}/\mathcal{D}}$ for finite structures $\boldsymbol{T}$ similar to $\boldsymbol{S}_{\mathcal{I}}$, and then obtaining finite templates from these manifold minions applying Corollary 7.6.

Reductions for Undecidability

Theorem 8.1. *Let $\boldsymbol{G}$ be a Σ-structure, let $\mathscr{M}$ be a minion, let $\mathcal{D} \subseteq 2^{\mathscr{M}}$ be a finite description, and let $\mathcal{I}$ be a Σ-interpretation over $\mathscr{M}$. Let $h, r \in \mathbb{N}$ with $h \geqslant r$. Suppose that*

(1) $\mathrm{Hom}(\boldsymbol{G}, \cdot)$ is undecidable,

(2) $\mathcal{I}$ is internal at arity h w.r.t. $\mathcal{D}$,

(3) $\boldsymbol{S}_{\mathcal{I}}$ is finitely equivalent to $\boldsymbol{G}$, and

(4) $r \geqslant \mathrm{ar}(Q)$ for every $Q \in \mathcal{D}$ and $r \geqslant \mathrm{ar}(U^{\mathcal{I}})$.

Then the family of finite templates of the form $(\boldsymbol{K}^h_r, \boldsymbol{B})$ satisfying $\mathscr{M} \to \mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$ is undecidable.

Proof of Theorem 8.1. We show that, given a finite Σ-structure $\boldsymbol{C}$, there is an algorithm that computes a finite template of the form $(\boldsymbol{K}^h_r, \boldsymbol{B})$ such that $\mathscr{M} \to \mathrm{Pol}(\boldsymbol{K}^h_r, \boldsymbol{B})$ if and only if $\boldsymbol{G} \to \boldsymbol{C}$. Let $\mathscr{N}$ be the manifold minion $\boldsymbol{C}^{\mathcal{I}/\mathcal{D}}$. We define $\boldsymbol{B} = \boldsymbol{F}_{\mathscr{N}}(\boldsymbol{K}^h_r)$. Let us reassure the reader that $\boldsymbol{B}$ is constructible from $\boldsymbol{C}$, even though this may not be immediately obvious. This follows from observing that $\boldsymbol{B}$ can be constructed when given access to the restriction of $\mathscr{M}/\mathcal{D}$ to elements of arity at most h and to the interpretation $\mathcal{I}/\mathcal{D}$, and these are finite objects that are independent from the input $\boldsymbol{C}$, so they can be precomputed.

Let us see that $\mathscr{M} \to \mathrm{Pol}(\boldsymbol{K}^h_r, \boldsymbol{B})$ if and only if $\boldsymbol{G} \to \boldsymbol{C}$. This is a consequence of the following chain of double implications:

$$\begin{array}{lr} \boldsymbol{G} \to \boldsymbol{C} & \\ & \xLeftrightarrow{\text{Lemma 6.3}} \\ \mathscr{M} \xrightarrow{h} \mathscr{N} & \\ & \xLeftrightarrow{\text{Corollary 7.6}} \\ \mathscr{M} \to \mathrm{Pol}(\boldsymbol{K}^h_r, \boldsymbol{B}). & \end{array}$$

Let us spell the arguments out. As $\boldsymbol{S}_{\mathcal{I}}$ is finitely equivalent to $\boldsymbol{G}$, $\boldsymbol{G} \to \boldsymbol{C}$ holds if and only if $\boldsymbol{S}_{\mathcal{I}} \to \boldsymbol{C}$. By Lemma 6.3, there is a homomorphism from $\boldsymbol{S}_{\mathcal{I}}$ to $\boldsymbol{C}$ if and only if there is a partial homomorphism $F : \mathscr{M} \xrightarrow{h} \mathbf{C}^{\mathcal{I}/\mathcal{D}} = \mathscr{N}$ By Lemma 7.2, the rank of $\mathscr{N}$ is at most r, so we can apply Corollary 7.6 to the minion $\mathscr{N}$ and the template $(\mathbf{K}_r^h, \mathbf{B}) = (\mathbf{K}_r^h, \mathbf{F}_{\mathscr{N}}(\mathbf{K}_r^h))$. In particular, $\mathscr{M} \xrightarrow{h} \mathscr{N}$ if and only if $\mathscr{M} \to \mathrm{Pol}(\mathbf{K}_r^h, \mathbf{B})$. This completes the proof. □

We also give an analogue of last result tailored for minions that can be expressed as disjoint unions of subminions. In this case, the global structure induced by an interpretation is a disjoint union of the global structures induced on each of the disjoint subminions (Observation 1), and we can apply the strategy from last proof to each of the parts.

Theorem 8.2. *Let $\mathbb{G}$ be a family of Σ-structures. Let $\mathscr{M} = \bigsqcup_{\mathbf{G} \in \mathbb{G}} \mathscr{M}_{\mathbf{G}}$ be a minion, let $\mathcal{D} \subseteq 2^{\mathscr{M}}$ be a finite description, and let $\mathcal{I}$ be a Σ-interpretation over $\mathscr{M}$. Let $h, r \in \mathbb{N}$ with $h \geqslant r$. Suppose that*

1. *(1) $\mathcal{I}$ is internal at arity h w.r.t. $\mathcal{D}$,*
2. *(2) $\mathbf{S}_{\mathcal{I}_{\mathbf{G}}}$ is finitely equivalent to $\mathbf{G}$ for each $\mathbf{G} \in \mathbb{G}$, where $\mathcal{I}_{\mathbf{G}} = \mathcal{I}|_{\mathscr{M}_{\mathbf{G}}}$, and*
3. *(3) $r \geqslant \mathrm{ar}(Q)$ for every $Q \in \mathcal{D}$ and $Q = U^{\mathcal{I}}$,*

Then there is an algorithm that, given a finite Σ-structure $\mathbf{C}$, yields a finite template $(\mathbf{K}_r^h, \mathbf{B})$ such that $\mathbf{G} \to \mathbf{T}$ if and only if $\mathscr{M}_{\mathbf{G}} \to \mathrm{Pol}(\mathbf{K}_r^h, \mathbf{B})$ for each $\mathbf{G} \in \mathbb{G}$. In particular, the following hold:

- *(i) The set $\mathrm{Hom}(\mathbb{G}, \cdot)$ is Turing-reducible to the family of finite templates $(\mathbf{K}_r^h, \mathbf{B})$ satisfying $\mathscr{M}_{\mathbf{G}} \to \mathrm{Pol}(\mathbf{K}_r^h, \mathbf{B})$ for some $\mathbf{G} \in \mathbb{G}$.*
- *(ii) The set $\mathrm{Hom}_{\mathrm{eg}}(\mathbb{G}, \cdot)$ is Turing-reducible to the family of finite templates $(\mathbf{K}_r^h, \mathbf{B})$ satisfying $\mathscr{M}_{\mathbf{G}} \to \mathrm{Pol}(\mathbf{K}_r^h, \mathbf{B})$ for all but finitely many $\mathbf{G} \in \mathbb{G}$.*
- *(iii) The set $\mathrm{Hom}_{\mathrm{io}}(\mathbb{G}, \cdot)$ is Turing-reducible to the family of finite templates $(\mathbf{K}_r^h, \mathbf{B})$ satisfying $\mathscr{M}_{\mathbf{G}} \to \mathrm{Pol}(\mathbf{K}_r^h, \mathbf{B})$ for infinitely many $\mathbf{G} \in \mathbb{G}$.*

Proof of Theorem 8.2. Given a finite Σ-structure $\mathbf{C}$, define $\mathscr{N} = \mathbf{C}^{\mathcal{I}/\mathcal{D}}$ and $\mathbf{B} = \mathbf{F}_{\mathscr{N}}(\mathbf{K}_r^h)$. As in the proof of Theorem 8.1, there is an algorithm that computes the template $(\mathbf{K}_r^h, \mathbf{B})$ in response to the input $\mathbf{C}$. In this case, Lemma 6.4 tells us that $\mathscr{M}_{\mathbf{G}} \xrightarrow{h} \mathscr{N}$ if and only if $\mathbf{G} \to \mathbf{C}$ for a given $\mathbf{G} \in \mathbb{G}$. By the same arguments as in the proof of Theorem 8.1, this is the case if and only if $\mathscr{M}_{\mathbf{G}} \to \mathrm{Pol}(\mathbf{K}_r^h, \mathbf{B})$, as we wanted to show.

Observe that the last three items in the statement of the theorem are a direct consequence of the facts that $\mathbf{G} \to \mathbf{T}$ if and only if $\mathscr{M}_{\mathbf{G}} \to \mathrm{Pol}(\mathbf{K}_r^h, \mathbf{B})$ for each $\mathbf{G} \in \mathbb{G}$, and that $(\mathbf{K}_r^h, \mathbf{B})$ can be computed in response to $\mathbf{C}$. □

Reduction for Non-Computability

Theorem 8.3. *Let $\mathbf{G}$ be a Σ-structure, $\mathscr{M}$ a minion, $\mathcal{D} \subseteq 2^{\mathscr{M}}$ a finite description, $\mathcal{I}$ a Σ-interpretation over $\mathscr{M}$, and $h \geqslant r$ natural numbers. Suppose that*

1. *(1) there is $\mathbf{T} \in \mathrm{Hom}(\mathbf{G}, \cdot)$ for which no homomorphism $H : \mathbf{G} \to \mathbf{T}$ is computable,*
2. *(2) $\mathcal{I}$ is internal at arity h w.r.t. $\mathcal{D}$,*
3. *(3) $\mathbf{S}_{\mathcal{I}}$ is finitely equivalent to $\mathbf{G}$,*

(4) there is a computable homomorphism $F : \boldsymbol{G} \to \boldsymbol{S}_{\mathcal{I}}$, and

(5) $r \geqslant \mathrm{ar}(Q)$ for every $Q \in \mathcal{D}$ and $Q = U^{\mathcal{I}}$.

Then there exists a finite template of the form $(\boldsymbol{K}_r^h, \boldsymbol{B})$ satisfying that $\mathscr{M} \to \mathrm{Pol}(\boldsymbol{K}_r^h, \boldsymbol{B})$, but there is no computable minion homomorphism $H : \mathscr{M} \to \mathrm{Pol}(\boldsymbol{K}_r^h, \boldsymbol{B})$.

Proof. Define $\mathcal{J} = \mathcal{I}/\mathcal{D}$, and $\mathscr{N} = \boldsymbol{T}^{\mathcal{J}}$. We claim that the finite template satisfying the theorem's statement is $(\boldsymbol{K}_r^h, \boldsymbol{B})$, where $\boldsymbol{B} = \boldsymbol{F}_{\mathscr{N}}(\boldsymbol{K}_r^h)$. By Lemma 7.2, we have that $\mathscr{N}$ has rank at most r, so $\mathscr{N}$ and the template $(\boldsymbol{K}_r^h, \boldsymbol{B})$ witness Corollary 7.6. Hence $\boldsymbol{G} \to \boldsymbol{T}$ implies $\mathscr{M} \to \mathrm{Pol}(\boldsymbol{K}_r^h, \boldsymbol{B})$. We show that there is no computable minion homomorphism from $\mathscr{M}$ to $\mathrm{Pol}(\boldsymbol{K}_r^h, \boldsymbol{B})$.

We proceed by contradiction. Suppose there is a computable minion homomorphism $H : \mathscr{M} \to \mathrm{Pol}(\boldsymbol{K}_r^h, \boldsymbol{B})$. We give a composition of computable partial maps that yields a homomorphism from $\boldsymbol{G}$ to $\boldsymbol{T}$. By Corollary 7.6, there is a partial homomorphism $H' : \mathrm{Pol}(\boldsymbol{K}_r^h, \boldsymbol{B}) \xrightarrow{h} \mathscr{N}$. Observe that H' is given by a finite table (i.e, it is defined on a finite set, and its co-domain, consisting of the elements $f \in \mathscr{N}$ of arity bounded by h, is finite). Hence, H' is computable. This way, $H' \circ H$ yields a computable partial homomorphism from $\mathscr{M}$ to $\mathscr{N}$ defined up to arity h. Let $\mathrm{pj} : \mathscr{N} \rightharpoonup \boldsymbol{T}$ be the canonical projection (recall the definition from Section 6.2). This, again, is a computable partial map. By Lemma 6.3, the map $\mathrm{pj} \circ H' \circ H$ restricted to $U^{\mathcal{I}}$ is a homomorphism from $\boldsymbol{S}_{\mathcal{I}}$ to $\boldsymbol{T}$ (we do not require the restriction to be computable; we just use that $\mathrm{pj} \circ H' \circ H$ is computable). By hypothesis, there exists a computable homomorphism $F : \boldsymbol{G} \to \boldsymbol{S}_{\mathcal{I}}$. Then $\mathrm{pj} \circ H' \circ H \circ F$ is a computable homomorphism from $\boldsymbol{G}$ to $\boldsymbol{T}$, yielding a contradiction. Thus, there cannot be a computable homomorphism from $\mathscr{M}$ to $\mathrm{Pol}(\boldsymbol{K}_r^h, \boldsymbol{B})$, as we wanted to prove. ☐

Reduction for Hardness

Let us take another look at the algorithms $\mathcal{Q} \in \{\mathrm{AIP}, \mathrm{BLP}, \mathrm{BLP} + \mathrm{AIP}\}$. Unrolling the definitions, it turns turns out that for a given instance $\boldsymbol{I}$ the following are equivalent: (1) $\boldsymbol{I}$ is accepted by $\mathcal{Q}$, and (2) $\boldsymbol{I} \to \boldsymbol{F}_{\mathscr{M}_{\mathcal{Q}}}(\boldsymbol{A})$. This is shown in [11] for BLP and AIP, and further discussed in [21] for the case of BLP + AIP. From this we obtain the following alternative formulation of rounding problems as left-infinite PCSPs.

Fact 8.4. *Let $\mathcal{Q} \in \{AIP, BLP, BLP + AIP\}$, and $(\boldsymbol{A}, \boldsymbol{B})$ be a finite template. Suppose that $\mathcal{Q}$ solves* $\mathrm{PCSP}(\boldsymbol{A}, \boldsymbol{B})$. *Then the problems* $\mathrm{sPCSP}_{\mathcal{Q}}(\boldsymbol{A}, \boldsymbol{B})$ *and* $\mathrm{sPCSP}(\boldsymbol{F}_{\mathscr{M}_{\mathcal{Q}}}(\boldsymbol{A}), \boldsymbol{B})$ *are the same.*

The *projection minion* $\mathscr{P}$ is defined as the minion where $\mathscr{P}(n) = [n]$ and where $f^{\pi} = \pi(f)$ for each $f \in \mathscr{P}(n)$ and each $\pi \in [m]^{[n]}$. If P≠ NP, the minion $\mathscr{P}$ is homomorphically equivalent to $\mathrm{Pol}(\boldsymbol{A}) = \mathrm{Pol}(\boldsymbol{A}, \boldsymbol{A})$ for each finite structure $\boldsymbol{A}$ for which $\mathrm{CSP}(\boldsymbol{A})$ is NP-hard. It is also not difficult to show that $\mathscr{P} \to \mathscr{M}$ for every non-empty minion $\mathscr{M}$.

One of the cornerstones of the algebraic approach to PCSPs is the result [11, Theorem 3.12] that, for a finite template $(\boldsymbol{A}, \boldsymbol{B})$ and h at least as large as $|A|$ and each relation of $\boldsymbol{A}$, the problems $\mathrm{sPCSP}(\boldsymbol{A}, \boldsymbol{B})$ and $\mathrm{sPMC}_h(\mathscr{P}, \mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B}))$ are log-space equivalent, and similarly for their decision variants. The same proof actually shows the following more general result [7].

---

[7] To see that this indeed generalizes [11, Theorem 3.12], choose $\mathscr{M}$ to be the projection minion $\mathscr{P}$, and use the fact that $\boldsymbol{F}_{\mathscr{P}}(\boldsymbol{A})$ is isomorphic to $\boldsymbol{A}$.

Theorem 8.5. *Let $\mathscr{M}$ be a minion, $N \in \mathbb{N}$ a number, and $(\boldsymbol{A}, \boldsymbol{B})$ a finite template satisfying that $\mathscr{M} \to \mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$. Then there is a log-space reduction from $\mathrm{sPMC}_N(\mathscr{M}, \mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B}))$ to $\mathrm{sPCSP}(\boldsymbol{F}_{\mathscr{M}}(\boldsymbol{A}), \boldsymbol{B})$. Further suppose that N is at least as large as $|A|$ and $|R^A|$ for any relation R in the signature of $\boldsymbol{A}$. Then there is a log-space reduction from $\mathrm{sPCSP}(\boldsymbol{F}_{\mathscr{M}}(\boldsymbol{A}), \boldsymbol{B})$ to $\mathrm{sPMC}_N(\mathscr{M}, \mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B}))$.*

Proof. The reductions are essentially the ones given in the proof of [11, Theorem 3.12]. We sketch them below. For the rest of the proof we identify A with $[|A|]$, and R^A with $[|R^A|]$ for each relation symbol R in some fixed way. Given a relation symbol R and an index $i \in [\mathrm{ar}(R)]$ we write $\Pi_{R,i} : R^A \to A$ for the map $\boldsymbol{a} \mapsto a_i$.

Reducing from sPMC *to* sPCSP. Let Φ be a minor condition of arity at most N. Without loss of generality, we may assume that Φ is of the form $\exists x_1, \ldots, x_k \varphi(x_1, \ldots, x_k)$, where φ is a conjunction of formulas of the form $x_i = x_j^{\pi}$. For each $i \in [k]$ we let n_i be the arity of x_i. The first part of the reduction is the construction of an structure $\boldsymbol{I}_\Phi$ satisfying $\boldsymbol{I}_\Phi \to \boldsymbol{F}_{\mathscr{M}}(\boldsymbol{A})$ if $\mathscr{M} \models \Phi$. First, define the set $S = \{x_i(\pi) \,|\, i \in [k], \pi \in A^{[n_i]}\}$. Now, we identify two elements $x_i(\pi)$ and $x_j(\sigma)$ if there is an identity in φ of the form $x_i = x_j^{\gamma}$, where $\sigma = \pi \circ \gamma$. We define the universe I_Φ as the result of performing all these identifications in S, and write $\langle x_i(\pi) \rangle$ to denote the equivalence class of $x_i(\pi)$ in I_Φ. Given a relation symbol R of arity m, we define R^{I_Φ} as the set of tuples of the form $(\langle x_i(\pi_1) \rangle, \ldots, \langle x_i(\pi_m) \rangle)$ for which there is a map $\sigma \in [R^A]^{[n_i]}$ satisfying $\Pi_{R,j} \circ \sigma = \pi_j$ for each $j \in [m]$. The structure $\boldsymbol{I}_\Phi$ can be built in log-space. Now, if $x_i \mapsto p_i$ is a satisfying assignment of Φ in $\mathscr{M}$, then the map $\langle x_i(\pi) \rangle \mapsto p_i^{\pi}$ is a well-defined homomorphism from $\boldsymbol{I}_\Phi$ to $\boldsymbol{F}_{\mathscr{M}}(\boldsymbol{A})$. Now, let $F : \boldsymbol{I}_\Phi \to \boldsymbol{B}$ be a homomorphism. The second part of the reduction computes an assignment of Φ over $\mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$ from F. For each $i \in [k]$ we define $f_i : \boldsymbol{A}^{n_i} \to \boldsymbol{B}$ as the map $\boldsymbol{a} \mapsto F(\langle x_i(\boldsymbol{a}) \rangle)$, where we recall that $\boldsymbol{a}$ can be seen as a map from $[n_i]$ to A. It is routine to verify that f_i is indeed a homomorphism and that the map $x_i \mapsto f_i$ is a satisfying assignment of Φ in $\mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$ that can be obtained in log-space.

Reducing from sPCSP *to* sPMC. Let $\boldsymbol{I}$ be an instance of $\mathrm{sPCSP}(\boldsymbol{F}_{\mathscr{M}}(\boldsymbol{A}), \boldsymbol{B})$. The first map of our reduction constructs in log-space a minor condition $\Phi_{\boldsymbol{I}}$ of arity at most N that satisfies $\mathscr{M} \models \Phi_{\boldsymbol{I}}$ if $\boldsymbol{I} \to \boldsymbol{F}_{\mathscr{M}}(\boldsymbol{A})$. Let Σ be the relational signature of all structures under consideration. Then the condition $\Phi_{\boldsymbol{I}}$ is defined as

$$\exists_{v \in I}^{|A|} x_v \ \exists_{R \in \Sigma, r \in R^I}^{|R^A|} y_r \left(\bigwedge_{R \in \Sigma, r \in R^I, i \in [\mathrm{ar}(R)]} x_{r(i)} = y_r^{\Pi_{R,i}} \right).$$

The arity of this minor condition is the maximum of $|A|$ and $|R^A|$ for all $R \in \Sigma$, so it is bounded by N by assumption. Now suppose that $F : \boldsymbol{I} \to \boldsymbol{F}_{\mathscr{M}}(\boldsymbol{A})$ is a homomorphism. Given $R \in \Sigma$, and $r \in R^I$ we write $\widehat{F}(r)$ for the element $p \in \mathscr{M}(R^A)$ witnessing that $(F(r(1)), \ldots, F(r(\mathrm{ar}(R)))) \in R^{F_{\mathscr{M}}(A)}$. Then the map $x_v \mapsto F(v)$ together with $y_r \mapsto \widehat{F}(r)$ is a satisfying assignment of $\Phi_{\boldsymbol{I}}$ in $\mathscr{M}$. Now, let H be a satisfying assignment of $\Phi_{\boldsymbol{I}}$ in $\mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$. The second map of the reduction needs to construct a homomorphism $F : \boldsymbol{I} \to \boldsymbol{B}$ in log-space using H. Given $v \in I$, we define $F(v)$ as $H(x_v)(\boldsymbol{a})$, where $\boldsymbol{a} \in A^{|A|}$ is our fixed enumeration of A. Now it is routine to check that F is indeed a homomorphism. ☐

By this point showing our main hardness reduction is just a matter of composing the reductions we have shown up until now.

Theorem 8.6. *Let* $\boldsymbol{G}$ *be a* Σ*-structure,* $\mathscr{M}$ *a minion,* $(\mathcal{I}, \mathcal{D}, \Psi, h)$ *a pattern over* $\mathscr{M}$ *and* $r \leqslant h$ *a natural number. Suppose that*

(1) $\boldsymbol{S}_{\mathcal{I}}$ *is finitely equivalent to* $\boldsymbol{G}$*, and*

(2) $r \geqslant \operatorname{ar}(Q)$ *for every* $Q \in \mathcal{D}$ *and* $Q = U^{\mathcal{I}}$*.*

Then for each $\boldsymbol{C} \in \operatorname{Hom}(\boldsymbol{G}, \cdot)$ *there is a finite template* $(\boldsymbol{K}_r^h, \boldsymbol{B})$ *such that* $\operatorname{sPCSP}(\boldsymbol{G}, C)$ *is many-one reducible to* $\operatorname{sPCSP}(\boldsymbol{F}_{\mathscr{M}}(\boldsymbol{K}_r^h), \boldsymbol{B})$*.*

Proof. Let $\boldsymbol{C} \in \operatorname{Hom}(\boldsymbol{G}, \cdot)$. Let $\mathscr{N} = \boldsymbol{C}^{\mathcal{I}/\mathcal{D}}$. By Lemma 6.5 there is a many-one reduction from $\operatorname{sPCSP}(\boldsymbol{G}, \boldsymbol{T})$ to $\operatorname{sPMC}_h(\mathscr{M}, \mathscr{N})$. By Lemma 7.2 the rank of $\mathscr{N}$ is at most r, so the minion $\mathscr{N}$ and the template $(\boldsymbol{K}_r^h, \boldsymbol{B}) = (\boldsymbol{K}_r^h, \boldsymbol{F}_{\mathscr{N}}(\boldsymbol{K}_r^h))$ satisfy Corollary 7.6. Hence, there is a partial isomorphism $F : \operatorname{Pol}(\boldsymbol{K}_r^h, \boldsymbol{B}) \xrightarrow{h} \mathscr{N}$. This yields a many-one reduction from $\operatorname{sPMC}_h(\mathscr{M}, \mathscr{N})$ to $\operatorname{sPMC}_h(\mathscr{M}, \operatorname{Pol}(\boldsymbol{K}_r^h, \boldsymbol{B}))$. Finally, by Theorem 8.5, the problem $\operatorname{sPMC}_h(\mathscr{M}, \operatorname{Pol}(\boldsymbol{K}_r^h, \boldsymbol{B}))$ has a many-one reduction to $\operatorname{sPCSP}(\boldsymbol{F}_{\mathscr{M}}(\boldsymbol{K}_r^h), \boldsymbol{B})$. ☐

9 Proof of the Main Results

In this section we present the proofs of Theorems 3.1 to 3.3, and Theorem 1.2. Each of these proofs consists of reducing one of the source problems described in Section 5 to our target problems through the reductions shown in Section 8. The lengthy lists of requirements for our reductions indicate that applying them requires some work: even after coming up with a suitable interpretation $\mathcal{I}$ and a suitable description $\mathcal{D}$, proving that $\mathcal{I}$ is internal w.r.t. $\mathcal{D}$ requires multiple applications of Lemma 6.1, which means constructing multiple logical formulas, and analysing their satisfying assignments on the original minion $\mathscr{M}$ and the quotient $\mathscr{M}/\mathcal{D}$.

To make matters worse, this section makes painfully apparent that there is, in principle, no way of transferring our results between the minions $\mathscr{M}_{\mathrm{AIP}}, \mathscr{M}_{\mathrm{BLP}}, \mathscr{M}_{\mathrm{BLP+AIP}}$ despite their similarities. This means that we need to write very similar proofs (with small, but significant changes) over and over. We remark that minion homomorphisms do not yield reductions for any of our results. Indeed, suppose that $\mathscr{M} \to \mathscr{N}$. Then, in general there is no relation between the decidability of the finite templates $(\boldsymbol{A}, \boldsymbol{B})$ such that $\mathscr{M} \to \operatorname{Pol}(\boldsymbol{A}, \boldsymbol{B})$, and the decidability of those satisfying $\mathscr{N} \to \operatorname{Pol}(\boldsymbol{A}, \boldsymbol{B})$. Similarly, in the absence of additional restrictions, it is possible that all homomorphisms $F : \mathscr{N} \to \operatorname{Pol}(\boldsymbol{A}, \boldsymbol{B})$ are non-computable, but there is a computable homomorphism $H : \mathscr{M} \to \operatorname{Pol}(\boldsymbol{A}, \boldsymbol{B})$, or vice versa. As for hardness, it is true that $\operatorname{sPCSP}(\boldsymbol{F}_{\mathscr{M}}(\boldsymbol{A}), \boldsymbol{B})$ has a straight-forward reduction to $\operatorname{sPCSP}(\boldsymbol{F}_{\mathscr{N}}(\boldsymbol{A}), \boldsymbol{B})$, so it may seem that the family of rounding problems arising from $\mathscr{M}$ is easier than the one resulting from $\mathscr{N}$. However, this is not true in general: There may be templates $(\boldsymbol{A}, \boldsymbol{B})$ for which $\mathscr{M} \to \operatorname{Pol}(\boldsymbol{A}, \boldsymbol{B})$ but $\mathscr{N} \not\to \operatorname{Pol}(\boldsymbol{A}, \boldsymbol{B})$, implying that there are *more* rounding problems for $\mathscr{M}$ than for $\mathscr{N}$, and these may be hard.

How to read this section

We try to avoid being meticulous and redundant to the point of exhausting the reader, and being so careless as to skip relevant details. The bulk of the following subsections is spent proving that several interpretations are internal with respect to corresponding descriptions. This is done through several claims. We believe that after getting acquainted with the typical arguments, the proofs of many of these claims are routine.

Sections 9.1 to 9.3 are conveniently ordered from less to more difficult, and they all employ different versions of the arguments in Section 9.1. These subsections are further divided to show interpretations of $\mathbf{\Gamma}^+$ and interpretations of $\mathbf{\Gamma}$. Our interpretations of $\mathbf{\Gamma}^+$ are relatively simpler than the ones of $\mathbf{\Gamma}$, because the interpretations in this second group are meant to produce Boolean templates. Hence, a good overview of the following subsections is the following.

- **Interpretations of the super-grid $\mathbf{\Gamma}^+$** (ordered from simpler to more complex):

  Section 9.1.1 (AIP), Section 9.2.1 (BLP), and Section 9.3.1 (BLP+AIP).

- **Interpretations of the grid $\mathbf{\Gamma}$** (ordered from simpler to more complex):

  Section 9.1.2 (AIP), and Section 9.2.2 (BLP).

The proofs in Section 9.4 and Section 9.5 handle WNU and cyclic polymorphisms respectively, and stand on their own. In both these cases we define minions characterizing the presence of those polymorphisms, and then find ways to interpret useful structures (e.g., grids) on them. However, these minions lack the arithmetical structure of $\mathscr{M}_{\text{AIP}}$, $\mathscr{M}_{\text{BLP}}$, and $\mathscr{M}_{\text{BLP+AIP}}$, so we need to use different arguments. There is ample difference between Section 9.4 and Section 9.5, but they handle minions constructed in a similar way, and some arguments are analogous. They are, again, ordered from simpler to more complex, so we recommend the reader to tackle Section 9.4 first.

9.1 The AIP Algorithm

We prove Theorem 3.1 in this section. To prove item (1) we give an interpretation of the super-grid $\mathbf{\Gamma}^+$ over $\mathscr{M}_{\text{AIP}}$, shown in Section 9.1.1. Items (2),(3), and (4) are proven similarly, by showing in Section 9.1.2 a suitable interpretation of the grid $\mathbf{\Gamma}$ over $\mathscr{M}_{\text{AIP}}$.

9.1.1 AIP: Interpreting the Super-Grid

The following interpretation $\mathcal{I}$ induces a global structure $\mathbf{S}_{\mathcal{I}}$ that is finitely equivalent to $\mathbf{\Gamma}^+$.

Interpretation 1. The Σ_{Γ^+}-interpretation $\mathcal{I}$ over $\mathscr{M}_{\text{AIP}}$ is given by

$$U^{\mathcal{I}} = \left\{(m_1, m_2, m_3, n) \in \mathscr{M}_{\text{AIP}}(4) \,\middle|\, m_i > 0 \text{ for all } i = 1,2,3\right\},$$

$$O^{\mathcal{I}} = \left\{(1,1,1,-4)\right\}, \quad \text{and} \quad \Pi^{\mathcal{I}}_{O,1} = \text{id},$$

$$E_i^{\mathcal{I}} = \left\{(m_1, m_2, m_3, 1, n) \in \mathscr{M}_{\text{AIP}}(5)\right\}, \quad \text{and}$$

$$\Pi^{\mathcal{I}}_{E_i,j} = \begin{cases} (1,2,3,4,4), \text{ when } j = 1, \\ (1,2,3,i,4) \text{ when } j = 2, \end{cases} \quad \text{for each } i \in [3],$$

$$\mathbb{E}_i^{\mathcal{I}} = \Big\{(m_1, m_2, m_3, m_i, n) \in \mathscr{M}_{\mathrm{AIP}}(5)\Big\}, \quad \text{and}$$

$$\Pi^{\mathcal{I}}_{\mathbb{E}_i,j} = \begin{cases} (1,2,3,4,4) \text{ for } j = 1, \\ (1,2,3,i,4) \text{ for } j = 2 \end{cases} \quad \text{for each } i \in [3].$$

This way, $\mathbf{\Gamma}^+$ is easily seen to be isomorphic to $\mathbf{S}_{\mathcal{I}}$ via the homomorphism $(m_1, m_2, m_3) \mapsto (m_1, m_2, m_3, 1 - m_1 - m_2 - m_3)$. Now we define a description $\mathcal{D} \subseteq 2^{\mathscr{M}_{\mathrm{AIP}}}$ so that $\mathcal{I}$ is internal at arity 5 w.r.t. $\mathcal{D}$.

Description 1. The description $\mathcal{D} \subseteq 2^{\mathscr{M}_{\mathrm{AIP}}}$ consists of the predicates

$$D_< = \Big\{(m_1, m_2, m_3) \in \mathscr{M}_{\mathrm{AIP}} \,\Big|\, m_1 < m_2\Big\},$$
$$D_1 = \Big\{(1, 0)\Big\}.$$

The interpretation $\mathcal{I}$ and the description $\mathcal{D}$ are the simplest in Section 9. The fact that we are not aiming to obtain Boolean templates means that we have virtually no restrictions on $\mathcal{I}$ and $\mathcal{D}$, and the minion $\mathscr{M}_{\mathrm{AIP}}$ already has the arithmetical structure required to express unit increments and multiplication times two. The following claims show that $\mathcal{I}$ is internal at arity 5 (w.r.t. $\mathcal{D}$):

Claim 1: $R^{\mathcal{I}}$ is $\mathcal{D}$-stable for each $R \in \Sigma_{\Gamma^+}$. This follows by a direct application of the definitions.

Claim 2: D_1 is internal at arity 5. Indeed, the minor condition $\phi_1(x) \equiv x = x^{(1,1)}$ is an internal definition of D_1.

Claim 3: $O^{\mathcal{I}}$ is internal at arity 5. The following is an internal definition of $O^{\mathcal{I}}$:

$$\phi_O(x) \equiv \phi_1(x^{(1,2,2,2)}) \wedge \phi_1(x^{(2,1,2,2)}) \wedge \phi_1(x^{(2,2,1,2)}).$$

Claim 4: $E_i^{\mathcal{I}}$ is internal at arity 5 for all $i \in [3]$. The following is an internal definition of $E_i^{\mathcal{I}}$:

$$\phi_{E_i}(x) \equiv \phi_1(x^{(2,2,2,1,2)}).$$

Claim 5: $\mathbb{E}_i^{\mathcal{I}}$ is internal at arity 5 for all $i \in [3]$. The following is an internal definition of $\mathbb{E}_i^{\mathcal{I}}$:

$$\phi_{\mathbb{E}_i}(x) \equiv x^{\sigma} = x^{\tau},$$

where $\sigma \in [3]^{[5]}$ is the map sending i to 1, 4 to 2, and the other elements to 3, and $\tau \in [3]^{[5]}$ sends i to 2, 4 to 1, and the other elements to 3. The key insight is that if $f^{\sigma} \sim_{\mathcal{D}} f^{\tau}$ for some element $f \in \mathscr{M}_{\mathrm{AIP}}(5)$, then it cannot be that $f(i) < f(4)$ or $f(i) > f(4)$, so it must be that $f(i) = f(4)$.

Claim 6: Let $i \in [3]$. Suppose that $\phi(x)$ is an internal reference to $U^{\mathcal{I}}$. Then the following formula is also an internal reference to $U^{\mathcal{I}}$:

$$\phi'(x) \equiv \exists y \exists z \Big(\phi_{E_i}(z) \wedge \phi(y) \wedge y = z^{(1,2,3,4,4)} \wedge x = z^{(1,2,3,i,4)} \Big).$$

Moreover, if f satisfies ϕ on $\mathscr{M}_{\mathrm{AIP}}$, then g satisfies ϕ' on $\mathscr{M}_{\mathrm{AIP}}$, where $g(i) = f(i) + 1$, $g(4) = f(4) - 1$, and $g(j) = f(j)$ for $j \neq i, 4$. Let us show that $\phi(x)$ is indeed an internal reference to $U^{\mathcal{I}}$. Suppose that $\mathscr{M}_{\mathrm{AIP}}/\mathcal{D} \models \phi(\langle f_x \rangle)$ with $\langle f_y \rangle, \langle f_z \rangle$ as existential witnesses for y, z. The formula ϕ is an internal reference to $U^{\mathcal{I}}$, so $f_y \in U^{\mathcal{I}}$. By a similar reasoning we also obtain that $f_z \in E_i^{\mathcal{I}}$. It must hold that $f_z^{(1,2,3,4,4)} \sim_{\mathcal{D}} f_y$, so $f_z^{(1,2,3,4,4)} \in U^{\mathcal{I}}$ (by Claim 1). Let $f_z = (m_1, m_2, m_3, 1, m_4)$. The fact that $f^{(1,2,3,4,4)} \in U^{\mathcal{I}}$ means that $m_1, m_2, m_3 > 0$. Hence, the fist three elements in $f^{(1,2,3,i,4)}$ must also be positive, so this tuple belongs to $U^{\mathcal{I}}$ as well. Using the fact that $f_z^{(1,2,3,i,4)} \sim_{\mathcal{D}} f_x$ we conclude that $f_x \in U^{\mathcal{I}}$, proving that ϕ is an internal reference to $U^{\mathcal{I}}$.

Finally, suppose that $\mathscr{M}_{\mathrm{AIP}} \models \phi(f)$ for some element f, and let g be defined as in the statement. Then g satisfies $\phi'(x)$ on $\mathscr{M}_{\mathrm{AIP}}$ with f as an existential witness for y and the tuple $(f(1), f(2), f(3), 1, f(4) - 1)$ as an existential witness for z.

Claim 7: Let $i \in [3]$. Suppose that $\phi(x)$ is an internal reference to $U^{\mathcal{I}}$. Then the following formula is also an internal reference to $U^{\mathcal{I}}$:

$$\phi'(x) \equiv \exists y \exists z \Big(\phi_{\mathbb{E}_i}(z) \wedge \phi(y) \wedge y = z^{(1,2,3,4,4)} \wedge x = z^{(1,2,3,i,4)} \Big).$$

Moreover, if f satisfies ϕ on $\mathscr{M}_{\mathrm{AIP}}$, then g satisfies ϕ' on $\mathscr{M}_{\mathrm{AIP}}$, where $g(i) = 2f(i)$, $g(4) = f(4) - f(i)$, and $g(j) = f(j)$ for $j \neq i, 4$. This follows analogously to the previous claim.

Claim 8: $U^{\mathcal{I}}$ is internal at arity 5. We prove this claim using Lemma 6.1. We define an internal reference $\phi_{n_1,n_2,n_3}(x)$ to $U^{\mathcal{I}}$ inductively for each $(n_1, n_2, n_3) \in \mathbb{N}^3$ following the lexicographical order, satisfying that $\mathscr{M}_{\mathrm{AIP}} \models \phi_{n_1,n_2,n_3}((n_1, n_2, n_3, 1 - n_1 - n_2 - n_3))$. We define $\phi_{1,1,1}(x) \equiv \phi_O(x)$. Now let $(n_1, n_2, n_3) \in \mathbb{N}^3$ be different from $(1, 1, 1)$, $i \in [3]$ be the largest index for which $n_i > 1$. We have two cases. Suppose that n_i is odd. Then we let $m_i = n_i - 1$, and $m_j = n_j$ for $j \in [3]$, $j \neq i$, and define

$$\phi_{n_1,n_2,n_3}(x) \equiv \exists y \exists z \Big(\phi_{m_1,m_2,m_3}(y) \wedge \phi_{E_i}(z) \wedge z^{(1,2,3,4,4)} = y \wedge z^{(1,2,3,i,4)} = x \Big).$$

Otherwise, if n_i is even, we let $m_i = n_i/2$, and $m_j = n_j$ for $j \in [3]$, $j \neq i$, and define

$$\phi_{n_1,n_2,n_3}(x) \equiv \exists y \exists z \Big(\phi_{m_1,m_2,m_3}(y) \wedge \phi_{\mathbb{E}_i}(z) \wedge z^{(1,2,3,4,4)} = y \wedge z^{(1,2,3,i,4)} = x \Big).$$

Now Claims 6 and 7 prove the statement.

Proof of item (1) in Theorem 3.1. The claims in this section show that $\mathcal{I}$ is internal w.r.t. $\mathcal{D}$ at arity 5. We also have that $\boldsymbol{S}_{\mathcal{I}}$ is isomorphic to $\boldsymbol{\Gamma}^+$. Hence, the result will follow from applying Theorem 8.6 after defining a 5-ary pattern Ψ of internal references to $\mathcal{I}$ w.r.t. $\mathcal{D}$. Given a structure $\boldsymbol{I}$ satisfying $\boldsymbol{I} \to \boldsymbol{\Gamma}^+$, we find a homomorphism $v \mapsto (m_v, n_v, o_v)$ from $\boldsymbol{I}$ to $\boldsymbol{\Gamma}^+$ in polynomial time. Moreover, we can assume that $\max_v \log_2(m_v n_v o_v) \leqslant |I|$. Observe that the map $v \mapsto (m_v, n_v, o_v, 1 - m_v - n_v - o_v)$ is a homomorphism from $\boldsymbol{I}$ to $\boldsymbol{S}_{\mathcal{I}}$. We define

$$\Psi_{\boldsymbol{I}} \equiv \exists_{v \in I} x_v \exists_{R \in \Sigma_{\Gamma+}, r \in R^I} x_r \left(\bigwedge_{v \in I} \phi_{m_v, n_v, o_v}(x_v) \right) \bigwedge \left(\bigwedge_{R \in \Sigma_{\Gamma}, r \in R^I} \phi_R(x_r) \bigwedge_{i \in [\mathrm{ar}(R)]} x_r^{\Pi^{\mathcal{I}}_{R,i}} = x_{r(i)} \right).$$

Here, the minor conditions $\phi_{m,n,o}(x)$, and $\phi_R(x)$ for $R \in \Sigma_{\Gamma+}$, are the internal references defined in the previous claims. To see that $\Psi_{\boldsymbol{I}}$ can be computed in polynomial time, observe that the formula $\phi_{m,n,o}(x_v)$ can be constructed inductively in time $O(\log_2(mno))$, and the maximum of $\log_2(m_v, n_v, o_v)$ is at most $|I|$ for $v \in I$. Now we only need to show that $\mathscr{M}_{\mathrm{AIP}} \models \Psi_{\boldsymbol{I}}$ in order to prove that $(\mathcal{I}, \mathcal{D}, \Psi)$ is a valid pattern. We give existential witnesses for the variables in $\Psi_{\boldsymbol{I}}$ to show that $\mathscr{M}_{\mathrm{AIP}} \models \Psi_{\boldsymbol{I}}$. For each $v \in I$, we choose

$$f_v = (m_v, n_v, o_v, 1 - m_v - n_v - o_v)$$

as the existential witness for x_v. By Claim 6, we know that f_v satisfies $\phi_{m_v, n_v, o_v}(x)$ on $\mathscr{M}_{\mathrm{AIP}}$. The fact that the map $v \mapsto f_v$ is a homomorphism from $\boldsymbol{I}$ to $\boldsymbol{S}_{\mathcal{I}}$ means that for each $R \in \Sigma_{\Gamma+}, r \in R^I$ there is some $f_r \in R^{\mathcal{I}}$ satisfying that

$$f_r^{\Pi^{\mathcal{I}}_{R,i}} = f_{r(i)} \qquad \text{for each } i \in \mathrm{ar}(R).$$

The element f_r must also satisfy $\phi_R(x)$ on $\mathscr{M}_{\mathrm{AIP}}$, because $\phi_R(x)$ is an internal definition of R. Hence, f_r is a valid existential witness for x_r. This shows that $\mathscr{M}_{\mathrm{AIP}} \models \Psi_{\boldsymbol{I}}$, and completes the proof that $(\mathcal{I}, \mathcal{D}, \Psi, 5)$ is a pattern over $\mathscr{M}_{\mathrm{AIP}}$.

Now Theorem 3.1-(1) follows from Theorem 8.6 together with Proposition 5.2-(1). Observe that $3 \geqslant \mathrm{ar}(U^{\mathcal{I}})$ and $3 \geqslant \mathrm{ar}(P)$ for all $P \in \mathcal{D}$, so it is enough to consider templates of the form $(\boldsymbol{K}_3^5, \boldsymbol{B})$ in Theorem 3.1-(1). $\square$

9.1.2 AIP: Interpreting the Grid

The following interpretation $\mathcal{I}$ induces a global structure which is finitely equivalent to $\boldsymbol{\Gamma}$.

Interpretation 2. The Σ_Γ-interpretation $\mathcal{I}$ over $\mathscr{M}_{\mathrm{AIP}}$ is given by

$$U^{\mathcal{I}} = \left\{ (2^m 3^n, y) \in \mathscr{M}_{\mathrm{AIP}}(2) \,\middle|\, m, n \text{ non-negative integers} \right\}$$

$$O^{\mathcal{I}} = \left\{ (1, 0) \right\}, \quad \text{and} \quad \Pi^{\mathcal{I}}_{O,1} = \mathrm{id},$$

$$E_1^{\mathcal{I}} = \left\{ (m, m, n, o) \in \mathscr{M}_{\mathrm{AIP}}(4) \,\middle|\, m + n = 1 \right\}, \quad \text{and} \quad \Pi^{\mathcal{I}}_{E_1, j} = \begin{cases} (1,2,2,2) \text{ for } j = 1, \text{ and} \\ (1,1,2,2) \text{ for } j = 2. \end{cases}$$

$$E_2^{\mathcal{I}} = \left\{ (m, m, m, n, o) \in \mathscr{M}_{\mathrm{AIP}}(5) \,\middle|\, m + n = 1 \right\}, \quad \text{and} \quad \Pi^{\mathcal{I}}_{E_2, j} = \begin{cases} (1,2,2,2,2) \text{ for } j = 1, \text{ and} \\ (1,1,1,2,2) \text{ for } j = 2. \end{cases}$$

Observe the grid structure $\mathbf{\Gamma}$ is isomorphic to $\boldsymbol{S}_{\mathcal{I}}$ via the bijection

$$(m,n) \mapsto (2^{m-1}3^{n-1}, 1 - 2^{m-1}3^{n-1}).$$

Moreover, this map is computable under the plain encoding for $\mathbf{\Gamma}$ and $\mathscr{M}_{\mathrm{AIP}}$. Next, we define a description $\mathcal{D} \subseteq 2^{\mathscr{M}_{\mathrm{AIP}}}$ so that $\mathcal{I}$ is internal w.r.t. $\mathcal{D}$ at arity 5.

Description 2. The description $\mathcal{D} \subseteq 2^{\mathscr{M}_{\mathrm{AIP}}}$ consists of the predicates $U^{\mathcal{I}}, O^{\mathcal{I}}$ defined in the previous interpretation.

We have defined the interpretation $\mathcal{I}$ and description $\mathcal{D}$ with the goal to obtain templates $(\boldsymbol{A}, \boldsymbol{B})$ where $\boldsymbol{A}$ is Boolean after applying Corollary 7.6. By Lemma 7.2, this requires $\mathrm{ar}(U^{\mathcal{I}}) \leqslant 2$ and $\mathrm{ar}(D) \leqslant 2$ for all predicates $D \in \mathcal{D}$. This is the reason behind the more awkward encoding of Γ compared to the previous subsection. An added difficulty is that because $\mathcal{D}$ cannot contain ternary predicates, we cannot speak about equality between coordinates in a straight-forward way. In the previous subsection, $\mathcal{D}$ contained the predicate $D_< = \{(m_1, m_2, m_3) \in \mathscr{M}_{\mathrm{AIP}} \mid m_1 < m_2\}$, so given an element $f = (n_1, n_2, n_3, n_4) \in \mathscr{M}_{\mathrm{AIP}}(4)$ we could tell whether its first two coordinates were equal by checking whether $\langle f \rangle = \langle f \rangle^{(2,1,3,4)}$ in the quotient $\mathscr{M}_{\mathrm{AIP}}/\mathcal{D}$ [8]. In this section we cannot compare coordinates in this direct way. But instead, we can still verify that the first two coordinates of f are equal as long as we have, for example, that $n_1 + n_3 = 1$ and $n_2 + n_3 = 1$. This motivates the definition of the predicates $E_1^{\mathcal{I}}$ and $E_2^{\mathcal{I}}$ in this subsection.

Let us show that $\mathcal{I}$ is indeed internal at arity 5 (w.r.t. $\mathcal{D}$). This is a consequence of the following claims.

Claim 1: The predicates $U^{\mathcal{I}}$, $O^{\mathcal{I}}$, $E_1^{\mathcal{I}}$, and $E_2^{\mathcal{I}}$ are all $\mathcal{D}$-stable. This follows directly from the definitions.

Claim 2: The predicate $O^{\mathcal{I}}$ is internal at arity 5. Indeed, the minor condition $\phi_O(x) \equiv x = x^{(1,1)}$ is an internal definition of $O^{\mathcal{I}}$.

Claim 3: The predicates $E_1^{\mathcal{I}}$ and $E_2^{\mathcal{I}}$ are internal at arity 5. Consider the following internal definitions for $E_1^{\mathcal{I}}$ and $E_2^{\mathcal{I}}$:

$$\phi_{E_1}(x) \equiv \phi_O(x^{(1,2,1,2)}) \wedge \phi_O(x^{(2,1,1,2)}),$$

$$\phi_{E_2}(x) \equiv \phi_O(x^{(1,2,2,1,2)}) \wedge \phi_O(x^{(2,1,2,1,2)}) \wedge \phi_O(x^{(2,2,1,1,2)}).$$

Claim 4: The following implications hold.

$$f^{(1,2,2,2)} \in U^{\mathcal{I}} \implies f^{(1,1,2,2)} \in U^{\mathcal{I}} \textbf{ for all } f \in E_1^{\mathcal{I}}$$
$$f^{(1,2,2,2,2)} \in U^{\mathcal{I}} \implies f^{(1,1,1,2,2)} \in U^{\mathcal{I}} \textbf{ for all } f \in E_2^{\mathcal{I}}.$$

We prove the statement for $f \in E_1^{\mathcal{I}}$. The case where $f \in E_2^{\mathcal{I}}$ follows analogously. Let $(m,n) = f^{(1,2,2,2)}$. By the definition of $E_1^{\mathcal{I}}$, it must hold that $f^{(1,1,2,2)} = (2m, n-m)$. Now, observe that $(m,n) \in U^{\mathcal{I}}$ implies that $(2m, m-n) \in U^{\mathcal{I}}$ as well, following the definition of $U^{\mathcal{I}}$.

---

[8] We do this when proving Claim 5 in Section 9.1.1.

Claim 5: Let $\phi(x)$ be an internal reference to $U^{\mathcal{I}}$ then the following formulas are also internal references to $U^{\mathcal{I}}$:

$$\phi_1(x) = \exists y \exists z \Big(\phi(y) \wedge \phi_{E_1}(z) \wedge y = z^{(1,2,2,2)} \wedge x = z^{(1,1,2,2)} \Big),$$

and

$$\phi_2(x) = \exists y \exists z \Big(\phi(y) \wedge \phi_{E_2}(z) \wedge y = z^{(1,2,2,2,2)} \wedge x = z^{(1,1,1,2,2)} \Big).$$

Indeed, let us argue the statement for ϕ_1. The case of ϕ_2 follows similarly. Suppose that $\mathscr{M}_{\mathrm{AIP}}/\mathcal{D} \models \phi_1(\langle f_x \rangle)$ for some $f \in \mathcal{Q}(2)$, and let $\langle f_y \rangle$ and $\langle f_z \rangle$ be existential witnesses for y and z respectively. In particular,

$$\mathscr{M}_{\mathrm{AIP}}/\mathcal{D} \models \phi(\langle f_y \rangle), \quad \text{and } \mathscr{M}_{\mathrm{AIP}}/\mathcal{D} \models \phi_{E_1}(\langle f_z \rangle).$$

The fact that $U^{\mathcal{I}}$, and $E_1^{\mathcal{I}}$ are $\mathcal{D}$-stable and $\phi(x)$, $\phi_{E_1}(x)$ are internal references to those predicates implies that $f_y \in U^{\mathcal{I}}$ and $f_z \in E_1^{\mathcal{I}}$. It must also hold that $f_y \sim_{\mathcal{D}} f_z^{(1,2,2,2)}$, so we can conclude that $f_z^{(1,2,2,2)} \in U^{\mathcal{I}}$. By Claim 5 this implies that $f_z^{(1,1,2,2)} \in U^{\mathcal{I}}$ as well. Using that $f_z^{(1,1,2,2)} \sim_{\mathcal{D}} f_x$ we finally obtain that $f_x \in U^{\mathcal{I}}$, proving that $\phi_1(x)$ is an internal reference to $U^{\mathcal{I}}$.

Claim 7: The predicate $U^{\mathcal{I}}$ is internal at arity 5. We use Lemma 6.1 to prove the claim. We define an internal reference $\phi_{m,n}(x)$ to $U^{\mathcal{I}}$ inductively for each $m, n \in \mathbb{N}$. Additionally, we keep the invariant that $\mathscr{M}_{\mathrm{AIP}} \models \phi_{m,n}(f)$, for each $m, n \in \mathbb{N}$ where $f = (2^{m-1}2^{n-1}, 1 - 2^{m-1}3^{n-1})$. We define $\phi_{1,1}$ as the minor condition ϕ_O. Now, given $m > 1$, we define

$$\phi_{m,1}(x) \equiv \exists y \exists z \Big(\phi_{m-1,1}(y) \wedge \phi_{E_1}(z) \wedge y = z^{(1,2,2,2)} \wedge x = z^{(1,1,2,2)} \Big).$$

By the previous Claim, the fact that $\phi_{m-1,1}$ is an internal reference to $U^{\mathcal{I}}$, means that so is $\phi_{m,1}$ as well. Moreover, the fact that $(2^{m-2}, 1 - 2^{m-2})$ satisfies $\phi_{m-1,1}(x)$ on $\mathscr{M}_{\mathrm{AIP}}$, means that $(2^{m-1}, 1 - 2^{m-1})$ satisfies $\phi_{m,1}(x)$ on $\mathscr{M}_{\mathrm{AIP}}$ by taking $(2^{m-2}, 1 - 2^{m-2})$ as an existential witness for y and

$$(2^{m-2}, 2^{m-2}, 1 - 2^{m-2}, -2^{m-2})$$

as an existential witness for z. Arguing in a similar way, if $n > 1$, we define

$$\phi_{m,n}(x) = \exists y \exists z \Big(\phi_{m,n-1}(y) \wedge \phi_{E_2}(z) \wedge y = z^{(1,2,2,2,2)} \wedge x = z^{(1,1,1,2,2)} \Big).$$

Again, using the previous claim we obtain that $\phi_{m,n}$ is an internal reference to $U^{\mathcal{I}}$. One can also see that $(2^{m-1}3^{n-1}, 1 - 2^{m-1}3^{n-1})$ satisfies $\phi_{m,n}(x)$ on $\mathscr{M}_{\mathrm{AIP}}$. This shows the claim.

Proof of items (2), (3), (4) of Theorem 3.1. In this section we have shown that $\mathbf{S}_{\mathcal{I}}$ is isomorphic to $\mathbf{\Gamma}$ through a computable homomorphism, and that $\mathcal{I}$ is internal at arity 5 w.r.t. $\mathcal{D}$. Additionally, $\mathrm{ar}(U^{\mathcal{I}}) = 2$, and $\mathrm{ar}(Q) = 2$ for all $Q \in \mathcal{D}$. Then items (3), (4) of Theorem 3.1 follow from Proposition 4.2 together via Theorem 8.1 and Theorem 8.3.

Finally, to prove Theorem 3.1-(2) we define a 5-ary pattern Ψ of internal references to $\mathcal{I}$ w.r.t. $\mathcal{D}$. We map each $\boldsymbol{I} \in \mathrm{Hom}(\mathbf{\Gamma}, \cdot)$ to a minor condition $\Psi_{\boldsymbol{I}}$ defined as follows. First, we find in polynomial time a homomorphism $v \mapsto (m_v, n_v)$ from $\boldsymbol{I}$ to $\boldsymbol{\Gamma}$. Moreover, this can be done in such a way that $\max_{v \in I} m_v + n_v \leqslant |I| + 1$. Then

$$\Psi_{\boldsymbol{I}} \equiv \exists_{v \in I} x_v \exists_{R \in \Sigma_\Gamma, r \in R^I} x_r \left(\bigwedge_{v \in I} \phi_{m_v, n_v}(x_v) \right) \bigwedge \left(\bigwedge_{R \in \Sigma_\Gamma, r \in R^I} \phi_R(x_r) \bigwedge_{i \in [\mathrm{ar}(R)]} x_r^{\Pi^{\mathcal{I}}_{R,i}} = x_{r(i)} \right).$$

Here, for each $R \in \Sigma_R$, $\phi_R(x)$ is the internal definition of $R^{\mathcal{I}}$ given in the previous claims, and $\phi_{m,n}(x)$ is the internal reference to $U^{\mathcal{I}}$ defined in Claim 7. To see that $\Psi_{\boldsymbol{I}}$ can be constructed in polynomial time, observe that $\phi_{m,n}$ takes $O(n+m)$ time to construct inductively, and we have that $\max_{v \in I} n_v + m_v \leqslant |I| + 1$. To prove that $(\mathcal{I}, \mathcal{D}, \Psi, 5)$ is a valid pattern we need to show that $\Psi_{\boldsymbol{I}}$ is satisfiable over $\mathscr{M}_{\mathrm{AIP}}$. Recall that $\mathscr{M}_{\mathrm{AIP}} \models \phi_{m,n}(f_{m,n})$, where $f_{m,n} = (2^{m-1}3^{n-1}, 1 - 2^{m-1}3^{n-1})$, and observe that the map $v \mapsto f_{m_v, n_v}$ is a homomorphism from $\boldsymbol{I}$ to $\boldsymbol{S}_{\mathcal{I}}$. Using this fact, $\mathscr{M}_{\mathrm{AIP}} \models \Psi_{\boldsymbol{I}}$ can be proven following the same reasoning as in the proof of Theorem 3.1-(1), in Section 9.1.1. ☐

9.2 The BLP Algorithm

We prove Theorem 3.2 in this section. To prove item (1) we give an interpretation of the super-grid $\boldsymbol{\Gamma}^+$ over $\mathscr{M}_{\mathrm{BLP}}$, shown in Section 9.2.1. Items (2),(3), and (4) are proven similarly, by showing in Section 9.2.2 a suitable interpretation of the grid $\boldsymbol{\Gamma}$ over $\mathscr{M}_{\mathrm{BLP}}$.

9.2.1 BLP: Interpreting the Super-Grid

The following interpretation $\mathcal{I}$ induces a global structure which is finitely equivalent to $\boldsymbol{\Gamma}^+$.

Interpretation 3. The Σ_{Γ^+}-interpretation $\mathcal{I}$ over $\mathscr{M}_{\mathrm{BLP}}$ is given by

$$U^{\mathcal{I}} = \left\{ (ms, ns, os, s, t) \in \mathscr{M}_{\mathrm{BLP}}(5) \,\middle|\, s = \frac{1}{2^j} \text{ for some } j \in \mathbb{N}, \text{ and } m, n, o \in \mathbb{N} \right\},$$

$$O^{\mathcal{I}} = \left\{ (s, s, s, s, t) \in \mathscr{M}_{\mathrm{BLP}}(5) \middle| s = 2^{-i} \text{ for some } i \in \mathbb{N} \right\}, \quad \text{and} \quad \Pi^{\mathcal{I}}_{O,1} = \mathrm{id},$$

$$E_i = \left\{ (s_1, s_2, s_3, s, s, t) \in \mathscr{M}_{\mathrm{BLP}}(6) \right\} \quad \text{and} \quad \Pi^{\mathcal{I}}_{E_i,j} = \begin{cases} (1,2,3,4,5,5) \text{ if } j = 1, \\ (1,2,3,i,4,5) \text{ if } j = 2 \end{cases} \quad \text{for all } i \in [3],$$

$$\mathbb{E}^{\mathcal{I}}_i = \{ (s_1, s_2, s_3, s_i, s, t) \in \mathscr{M}_{\mathrm{BLP}}(6) \} \quad \text{and} \quad \Pi^{\mathcal{I}}_{\mathbb{E}_i,j} = \begin{cases} (1,2,3,4,5,5) \text{ if } j = 1, \\ (1,2,3,i,4,5) \text{ if } j = 2 \end{cases} \quad \text{for each } i \in [3].$$

Here it is worth pointing out the main difference with respect to the interpretation of $\boldsymbol{\Gamma}^+$ over $\mathscr{M}_{\mathrm{AIP}}$ given in Section 9.1.1. Addition and, by extension, multiplication times two are easy to perform in the minion $\mathscr{M}_{\mathrm{BLP}}$, but we are missing a natural notion of "unit increment". To remediate this, we represent a triple $(m, n, o) \in \mathbb{N}^3$ in the 3-dimensional grid with elements of the form (ms, ns, os, s, t) in $\mathscr{M}_{\mathrm{BLP}}$, where the fourth coordinate represents the increment that we have chosen in our encoding. Of course, for a given increment s, the resulting grid in $\mathscr{M}_{\mathrm{BLP}}$ is finite, but we consider the disjoint union of these grids when s ranges over all values of the form $2^{-\ell}$ for $\ell \in \mathbb{N}$. The resulting global structure is not isomorphic to $\mathbf{\Gamma}^+$, but it is finitely equivalent to it, as indicated below.

Claim 1: The structure $S = S_{\mathcal{I}}$ induced by $\mathcal{I}$ is finitely equivalent to $\mathbf{\Gamma}^+$. Let us describe S. For each integer $i \geqslant 2$ we define S as the connected component of S_i containing the element

$$\left(\frac{1}{2^i}, \frac{1}{2^i}, \frac{1}{2^i}, \frac{1}{2^i} 1 - \frac{1}{2^{i-2}}\right) \in U_{\mathcal{I}}.$$

Observe that S is the disjoint union $\bigsqcup_{i \geqslant 2} S_i$. Given $i \geqslant 2$, we define $\mathbf{\Gamma}_i^+$ as the substructure of $\mathbf{\Gamma}^+$ induced on the set of elements $(m, n, o) \in \mathbb{N}^3$ with $m + n + o + 1 \leqslant 2^i$. The structure $\mathbf{\Gamma}_i^+$ is isomorphic to S_i through the bijection

$$(m, n, o) \mapsto \left(\frac{m}{2^i}, \frac{n}{2^i}, \frac{o}{2^i}, \frac{1}{2^i}, 1 - \frac{m + n + o + 1}{2^i}\right).$$

Finally, observe that $\mathbf{\Gamma}^+$ is finitely equivalent to the disjoint union $\bigsqcup_{i > 2} \mathbf{\Gamma}_i^+$. Indeed, if $\boldsymbol{I} \to \mathbf{\Gamma}^+$, for some finite $\boldsymbol{I}$, then it must be that $\boldsymbol{I} \to \mathbf{\Gamma}_i^+$ for some $i > 2$, and we also have that $\mathbf{\Gamma}_i^+ \to \mathbf{\Gamma}^+$ for all i. This proves the claim.

Now let us define a description $\mathcal{D} \subseteq 2^{\mathscr{M}_{\text{BLP}}}$ so that $\mathcal{I}$ is internal at arity 6 w.r.t. $\mathcal{D}$.

Description 3. The description $\mathcal{D} \subseteq 2^{\mathscr{M}_{\text{BLP}}}$ consists of the predicate $U_{\mathcal{I}}$ defined in the previous interpretation, and the predicates

$$D_p = \left\{ \left(\frac{1}{2^i}, 1 - \frac{1}{2^i}\right) \mid i \geqslant 0 \right\}, \text{ and}$$
$$D_< = \left\{ (s_1, s_2, s_3) \in \mathscr{M}_{\text{BLP}} \mid s_1 < s_2 \right\}.$$

The following claims establish that $\mathcal{I}$ is internal at arity 6 (w.r.t. $\mathcal{D}$).

Claim 2: The predicates $U^{\mathcal{I}}$, and $R^{\mathcal{I}}$ for $R \in \Sigma_{\Gamma+}$ are all $\mathcal{D}$-stable. This follows from the definitions.

Claim 3: The predicates $E_i^{\mathcal{I}}$ are internal at arity 6. The following is an internal definition of $E_i^{\mathcal{I}}$:

$$\phi_{E_i}(x) \equiv x^{(3,3,3,1,2,3)} = x^{(3,3,3,2,1,3)}.$$

Claim 4: The predicates $\mathbb{E}_i^{\mathcal{I}}$ are internal at arity 6. The following is an internal definition of $E_i^{\mathcal{I}}$:

$$\phi_{\mathbb{E}_i}(x) \equiv x^{\sigma} = x^{\tau},$$

where $\sigma : [6] \to [3]$ maps i to 1, 4 to 2 and the other elements to 3, and $\tau : [6] \to [3]$ maps i to 2, 4 to 1 and the other elements to 3.

Claim 5: If $\phi(x)$ is an internal reference to D_p, then the following formula is also an internal reference to D_p:

$$\phi'(x) = \exists y \exists^3 z \Big(\phi(y) \wedge z = z^{(2,1,3)} \wedge y = z^{(1,1,2)} \wedge x = z^{(1,2,2)}\Big).$$

Moreover, if $(1/2^i, -1/2^i)$ satisfies $\phi(x)$ (over $\mathscr{M}_{\mathbf{BLP}}$), then $(1/2^{i+1}, 1-1/2^{i+1})$ satisfies $\phi'(x)$. Let us begin with the second part of the statement. To see that $(1/2^{i+1}, 1-1/2^{i+1})$ satisfies $\phi'(x)$, observe that $(1/2^i, 1-1/2^i)$ and $(1/2^{i+1}, 1/2^{i+1}, 1-1/2^i)$ are valid existential witnesses for y and z.

Now let us show that ϕ' is an internal reference to $\mathcal{D}_{1/2}$. Suppose $\mathscr{M}_{\mathrm{BLP}}/\mathcal{D} \models \phi'(\langle f_x \rangle)$ and $\langle f_y \rangle, \langle f_z \rangle$ are existential witnesses for y and z. As ϕ is an internal reference to D_p, we have $f_y \in D_p$. Hence, the fact that $f_y \sim_{\mathcal{D}} f_z^{(1,1,2)}$ means that $f_z^{(1,1,2)} \in D_p$, so $f_z(1) + f_z(2) = 1/2^i$ for some $i \geqslant 0$. Also, observe that the fact that $f_z \sim_{\mathcal{D}} f_z^{(2,1,3)}$ implies $f_z(1) = f_z(2)$. Indeed, otherwise exactly one of f_z or $f_z^{(2,3,1)}$ would belong to $D_<$. Hence $f_z(1) = 1/2^{i+1}$, and $f_z^{(1,2,2)} \in D_p$. Finally, because $f_x \sim_{\mathcal{D}} f_z^{(1,2,2)}$, we must have $f_x \in D_p$, proving the claim.

Claim 6: The predicate D_p is internal at arity 6. We use Lemma 6.1 to prove the claim. We define an internal reference $\phi_i(x)$ to D_p inductively for each $i \geqslant 0$ in such a way that $(1/2^i, 1-1/2^i)$ satisfies $\phi_i(x)$ on $\mathscr{M}_{\mathrm{BLP}}$. We define

$$\phi_0(x) \equiv x = x^{(1,1)}.$$

Given $i > 0$, we define

$$\phi_i(x) \equiv \exists y \exists^3 z \Big(\phi_{i-1}(y) \wedge z = z^{(2,1,3)} \wedge y = z^{(1,1,2)} \wedge x = z^{(1,2,2)} \Big).$$

Now the previous claim proves that $\phi_i(x)$ is an internal reference to D_p and $(1/2^i, 1-1/2^i)$ satisfies it. This proves the statement.

Claim 7: The predicate $O^{\mathcal{I}}$ is internal at arity 6. Let $i \geqslant 2$. Then, by last claim, the following formula is an internal reference to $O^{\mathcal{I}}$:

$$\begin{aligned} \phi_{O,i}(x) \equiv \phi_i(x^{(1,2,2,2,2)}) &\wedge x^{(1,2,3,3,3)} = x^{(2,1,3,3,3)} \\ &\wedge x^{(1,3,2,3,3)} = x^{(2,3,1,3,3)} \\ &\wedge x^{(1,3,3,2,3)} = x^{(2,3,3,1,3)}. \end{aligned}$$

Additionally, $\phi_{O,i}(x)$ is satisfied on $\mathscr{M}_{\mathrm{BLP}}$ by the element

$$\left(\frac{1}{2^i}, \frac{1}{2^i}, \frac{1}{2^i}, \frac{1}{2^i}, 1 - \frac{1}{2^{i-2}} \right).$$

Now the claim follows from Lemma 6.1.

Claim 8: Let $i \in [3]$. Suppose $\phi(x)$ is an internal reference to $U^{\mathcal{I}}$. Then the following formula is also an internal reference to $U^{\mathcal{I}}$:

$$\phi'(x) \equiv \exists y \exists^6 z \Big(\phi(y) \wedge \phi_{E_i}(z) \wedge y = z^{(1,2,3,4,5,5)} \wedge x = z^{(1,2,3,4,i,5)} \Big).$$

Moreover, if f satisfies $\phi(x)$ (on $\mathscr{M}_{\mathbf{BLP}}$), then the tuple g satisfies $\phi'(x)$, where g is defined by $g(i) = f(i) + f(4)$, $g(5) = f(5) - f(4)$ and $g(j) = f(j)$, and $g(j) = f(j)$ for $j \neq i, 5$. We begin with the second part of the statement. To see that g satisfies $\phi'(x)$, observe that f and $f_z = (f(1), f(2), f(3), f(4), f(4), f(5) - f(4))$ are existential witnesses for y and z.

Now let us show that ϕ' is an internal reference to $U^{\mathcal{I}}$. Suppose $\mathscr{M}_{\mathrm{BLP}}/\mathcal{D} \models \phi'(\langle f_x \rangle)$ with $\langle f_y \rangle, \langle f_z \rangle$ as existential witnesses for y, z. As ϕ is an internal reference to $U^{\mathcal{I}}$, it must hold that

$f_y \in U^{\mathcal{I}}$. Moreover, $f_y \sim_{\mathcal{D}} f_z^{(1,2,3,4,5,5)}$, so $f_z^{(1,2,3,4,5,5)} \in U^{\mathcal{I}}$, meaning that $f_z(4) = \frac{1}{2^i}$ for some $i \geqslant 2$, and $f_z(j) = \frac{m_j}{2^i}$ for some $m_j \in \mathbb{N}$ for each $j \in [3]$. Now, the fact that $f_z \in E_i^{\mathcal{I}}$ means that $f_z(4) = f_z(5)$. Hence, we conclude that $f_z(5) = \frac{1}{2^i}$ as well. This implies that $f_z^{(1,2,3,4,i,5)}$ belongs to $U^{\mathcal{I}}$. Finally, using that $f_x \sim_{\mathcal{D}} f_z^{(1,2,3,4,i,5)}$, we obtain $f_x \in U^{\mathcal{I}}$, as we wanted.

Claim 9: Let $i \in [3]$. Suppose $\phi(x)$ is an internal reference to $U^{\mathcal{I}}$. Then the following formula is also an internal reference to $U^{\mathcal{I}}$:

$$\phi'(x) \equiv \exists y \exists^6 z \Big(\phi(y) \wedge \phi_{\mathbb{E}_i}(z) \wedge y = z^{(1,2,3,4,5,5)} \wedge x = z^{(1,2,3,4,i,5)}\Big)$$

Moreover, if f satisfies $\phi(x)$ (on $\mathscr{M}_{\mathrm{BLP}}$), then the tuple g satisfies $\phi'(x)$, where g is defined by $g(i) = 2f(i)$, $g(5) = f(5) - f(i)$ and $g(j) = f(j)$, and $g(j) = f(j)$ for $j \neq i, 5$. This can be shown analogously to the previous claim.

Claim 10: The predicate $U^{\mathcal{I}}$ is internal at arity 6. We prove the claim using Lemma 6.1, as usual. We define an internal reference $\phi_{m_1,m_2,m_3,m_4}(x)$ to $U^{\mathcal{I}}$ inductively, following the lexicographical order, for each $(m_1, m_2, m_3, m_4) \in \mathbb{N}^4$ such that $m_2 + m_3 + m_4 + 1 \leqslant 2^{m_1}$, in such a way that

$$\left(\frac{m_2}{2^{m_1}}, \frac{m_3}{2^{m_1}}, \frac{m_4}{2^{m_1}}, \frac{1}{2^{m_1}}, 1 - \frac{m_2 + m_3 + m_4 + 1}{2^{m_1}}\right)$$

satisfies $\phi_{m_1,m_2,m_3,m_4}(x)$ on $\mathscr{M}_{\mathrm{BLP}}$. For each $m \geqslant 2$ we define

$$\phi_{m,1,1,1}(x) \equiv \phi_{O,m}(x),$$

where $\phi_{O,m}(x)$ is the internal reference to $O^{\mathcal{I}}$ defined in Claim 7. Now suppose that $i > 1$ is the maximum index such that $m_i > 1$. We have two cases. Suppose that m_i is odd. Then we let $n_i = m_i - 1$, $n_j = m_j$ for all $j \neq i$, and define $\phi_{m_1,m_2,m_3,m_4}(x)$ as

$$\exists y \exists^6 z \Big(\phi_{n_1,n_2,n_3,n_4}(y) \wedge \phi_{E_i}(z) \wedge y = z^{(1,2,3,4,5,5)} \wedge x = z^{(1,2,3,4,i,5)}\Big).$$

Otherwise, suppose m_i is even. Then Then we let $n_i = m_i/2$, $n_j = m_j$ for all $j \neq i$, and define $\phi_{m_1,m_2,m_3,m_4}(x)$ as

$$\exists y \exists^6 z \Big(\phi_{n_1,n_2,n_3,n_4}(y) \wedge \phi_{\mathbb{E}_i}(z) \wedge y = z^{(1,2,3,4,5,5)} \wedge x = z^{(1,2,3,4,i,5)}\Big).$$

Now the statement follows from Claims 8 and 9.

Proof of item (1) of Theorem 3.1. We define a 6-ary pattern Ψ of internal references to $\mathcal{I}$ w.r.t. $\mathcal{D}$. Then the result will frollow from Theorem 8.6. Given $\boldsymbol{I} \in \mathrm{Hom}(\boldsymbol{\Gamma}^+, \cdot)$, we first compute a homomorphism $v \mapsto (m_v, n_v, o_v)$ from $\boldsymbol{I}$ to $\boldsymbol{\Gamma}^+$ in polynomial time in such a way that $\max_{v \in I} \log_2(m_v n_v o_v) \leqslant |I|$. Let $M = \max_{v \in I} m_v + n_v + o_v + 1$, and let $j = \lceil \log_2 M \rceil$. Then F is actually a homomorphism from $\boldsymbol{I}$ to $\boldsymbol{\Gamma}_j^+$, where $\boldsymbol{\Gamma}_j^+ \subseteq \boldsymbol{\Gamma}^+$ is the substructure defined in Claim 1. Then the minor condition $\Psi_{\boldsymbol{I}}$ is defined as

$$\begin{aligned} \exists_{v \in I} x_v \exists_{R \in \Sigma_{\Gamma^+}, r \in R^I} x_r \left(\bigwedge_{v \in I} \phi_{j,m_v,n_v,o_v}(x_v)\right) \bigwedge \left(\bigwedge_{r \in O^I} \phi_{O,j}(x_r) \wedge x_r = x_{r(1)}\right) \\ \bigwedge \left(\bigwedge_{R \in \Sigma_{\Gamma^+}, R \neq O, r \in R^I} \phi_R(x_r) \bigwedge_{i \in \mathrm{ar}(R)} x_r^{\Pi^{\mathcal{I}}_{R,i}} = x_{r(i)}\right). \end{aligned}$$

To see that $\Psi_{\boldsymbol{I}}$ can be computed in polynomial time, observe that ϕ_{j,m_v,n_v,o_v} can be constructed inductively in time $O(j+\log_2(m_v n_v o_v))$, $j \leqslant 1+\log_2 4|I|$, and $\max_{v\in I}\log_2(m_v n_v o_v) \leqslant |I|$. In order to prove that $(\mathcal{I},\mathcal{D},\Psi,6)$ is a valid pattern, we only need to show that $\mathscr{M}_{\mathrm{BLP}} \models \Psi_{\boldsymbol{I}}$. To do this, observe that the map

$$v \mapsto f_v = \left(\frac{m_v}{2^j}, \frac{n_v}{2^j}, \frac{o_v}{2^j}, \frac{1}{2^j}, 1-\frac{m_v+n_v+o_v+1}{2^j}\right),$$

is a homomorphism from $\boldsymbol{I}$ to the connected component of

$$\left(\frac{1}{2^j}, \frac{1}{2^j}, \frac{1}{2^j}, \frac{1}{2^j}, 1-\frac{1}{2^{j-2}}\right)$$

in $\boldsymbol{S}_{\mathcal{I}}$ (recall the isomorphisms from Claim 1). Now $\mathscr{M} \models \Psi_{\boldsymbol{I}}$ can be proven analogously to Theorem 3.1-(1) in Section 9.1.1. Observe that $5 \geqslant \mathrm{ar}(U^{\mathcal{I}})$ and $5 \geqslant \mathrm{ar}(P)$ for all $P \in \mathcal{D}$. Hence, we only need to consider templates of the form $(\boldsymbol{K}_5^6, \boldsymbol{B})$ to achieve the TFNP-hardness result. ☐

9.2.2 BLP: Interpreting the Grid

The following interpretation $\mathcal{I}$ induces a global structure that is finitely equivalent to $\boldsymbol{\Gamma}$.

Interpretation 4. The Σ_Γ-interpretation $\mathcal{I}$ over $\mathscr{M}_{\mathrm{BLP}}$ is given by

$$U^{\mathcal{I}} = \left\{\left(\frac{1}{2^m 3^n}, y\right) \in \mathscr{M}_{\mathrm{BLP}}(2) \mid m, n \text{ non-negative integers}\right\},$$

$$O^{\mathcal{I}} = \left\{(1,0)\right\}, \quad \text{and} \quad \Pi^{\mathcal{I}}_{O,1} = \mathrm{id},$$

$$E_1^{\mathcal{I}} = \left\{(x,x,y,y) \in \mathscr{M}_{\mathrm{BLP}}(4)\right\}, \quad \text{and} \quad \Pi^{\mathcal{I}}_{E_1,i} = \begin{cases}(1,1,2,2) \text{ for } i=1,\\ (1,2,2,2) \text{ for } i=2,\end{cases}$$

$$E_2^{\mathcal{I}} = \left\{(x,x,x,y,z) \in \mathscr{M}_{\mathrm{BLP}}(4) \mid x+y=1/3\right\}, \quad \text{and} \quad \Pi^{\mathcal{I}}_{E_2,i} = \begin{cases}(1,1,1,2,2) \text{ for } i=1,\\ (1,2,2,2,2) \text{ for } i=2.\end{cases}$$

This way, the grid structure $\boldsymbol{\Gamma}$ is isomorphic to $\boldsymbol{S}$ via the bijection

$$(m,n) \mapsto \left(\frac{1}{2^{m-1}3^{n-1}}, 1-\frac{1}{2^{m-1}3^{n-1}}\right).$$

Moreover, this map is computable under the plain encoding for $\boldsymbol{\Gamma}$ and $\mathscr{M}_{\mathrm{BLP}}$. Next, we define a description $\mathcal{D} \subseteq 2^{\mathscr{M}_{\mathrm{BLP}}}$ so that $\mathcal{I}$ is internal w.r.t. $\mathcal{D}$ at arity 5.

Description 4. The description $\mathcal{D} \subseteq 2^{\mathscr{M}_{\mathrm{BLP}}}$ consists of the binary predicates $U^{\mathcal{I}}$ and $O^{\mathcal{I}}$ defined in the previous interpretation $\mathcal{I}$, and the new predicates

$$D_{1/2} = \left\{(1/2, 1/2)\right\}, \quad \text{and} \quad D_{1/3} = \left\{(1/3, 2/3)\right\}.$$

It is worth to compare the interpretation $\mathcal{I}$ and description $\mathcal{D}$ with the ones from Section 9.1.2, used to encode $\mathbf{\Gamma}$ on $\mathscr{M}_{\text{AIP}}$. Both constructions are completely analogous, with the main difference that in that subsection we considered positive powers 2^m3^n and here we consider negative ones $2^{-m}3^{-n}$. Let us show that $\mathcal{I}$ is indeed internal at arity 5 w.r.t. $\mathcal{D}$. This is a consequence of the following claims.

Claim 1: The predicates $U^{\mathcal{I}}$, $O^{\mathcal{I}}$, $E_1^{\mathcal{I}}$, and $E_2^{\mathcal{I}}$ are all $\mathcal{D}$-stable. This is a routine check.

Claim 2: The predicate $O^{\mathcal{I}}$ is internal at arity 5 (w.r.t. $\mathcal{D}$). Indeed, the minor condition $\phi_O(x) \equiv x = x^{(1,1)}$ is an internal definition of $O^{\mathcal{I}}$.

Claim 3: The predicates $D_{1/2}$ and $D_{1/3}$ are internal at arity 5. The following are internal definitions for $D_{1/2}$ and $D_{1/3}$ w.r.t. $\mathcal{D}$:

$$\begin{aligned}
&\phi_{1/2}(x) \equiv x = x^{(2,1)},\\
&\phi_{1/3}(x) \equiv \exists^3 y\Big(y = y^{(3,1,2)} \wedge y = y^{(2,3,1)} \wedge x = y^{(1,2,2)}\Big).
\end{aligned}$$

Claim 4: The predicates $E_1^{\mathcal{I}}$ and $E_2^{\mathcal{I}}$ are internal at arity 5. Consider the following internal definitions for $E_1^{\mathcal{I}}$ and $E_2^{\mathcal{I}}$:

$$\begin{aligned}
&\phi_{E_1}(x) \equiv \phi_{1/2}(x^{(1,2,1,2)}) \wedge \phi_{1/2}(x^{(2,1,1,2)}), \quad \text{and}\\
&\phi_{E_2}(x) \equiv \phi_{1/3}(x^{(1,2,2,1,2)}) \wedge \phi_{1/3}(x^{(2,1,2,1,2)}) \wedge \phi_{1/3}(x^{(2,2,1,1,2)}).
\end{aligned}$$

Claim 5: The following implications hold.

$$\begin{aligned}
&f^{(1,1,2,2)} \in U^{\mathcal{I}} \implies f^{(1,2,2,2)} \in U^{\mathcal{I}} \textbf{ for all } f \in E_1^{\mathcal{I}}, \quad \textbf{and}\\
&f^{(1,1,1,2,2)} \in U^{\mathcal{I}} \implies f^{(1,2,2,2,2)} \in U^{\mathcal{I}} \textbf{ for all } f \in E_2^{\mathcal{I}}.
\end{aligned}$$

We prove the statement for $f \in E_1^{\mathcal{I}}$. The case where $f \in E_2^{\mathcal{I}}$ follows analogously. Let $(m, n) = f^{(1,2,2,2)}$. By the definition of $E_1^{\mathcal{I}}$, it must hold that $f^{(1,1,2,2)} = (2m, n - m)$. Now, observe that $(m, n) \in U^{\mathcal{I}}$ implies that $(2m, n - m) \in U^{\mathcal{I}}$ as well, following the definition of $U^{\mathcal{I}}$.

Claim 6: Let $\phi(x)$ be an internal reference to $U^{\mathcal{I}}$ then the following formulas are also internal references to $U^{\mathcal{I}}$:

$$\begin{aligned}
&\phi_1(x) = \exists y \exists z\Big(\phi(y) \wedge \phi_{E_1}(z) \wedge y = z^{(1,1,2,2)} \wedge x = z^{(1,2,2,2)}\Big), \quad \textbf{and}\\
&\phi_2(x) = \exists y \exists z\Big(\phi(y) \wedge \phi_{E_2}(z) \wedge y = z^{(1,1,1,2,2)} \wedge x = z^{(1,2,2,2,2)}\Big).
\end{aligned}$$

Indeed, let us argue the statement for ϕ_1. The case of ϕ_2 follows similarly. Suppose that $\mathscr{M}_{\text{BLP}}/\mathcal{D} \models \phi_1(\langle f_x \rangle)$ for some $f \in \mathcal{Q}(2)$, and let $\langle f_y \rangle$ and $\langle f_z \rangle$ be existential witnesses for y and z respectively. In particular,

$$\mathscr{M}_{\text{BLP}}/\mathcal{D} \models \phi(\langle f_y \rangle), \quad \text{and } \mathscr{M}_{\text{BLP}}/\mathcal{D} \models \phi_{E_1}(\langle f_z \rangle).$$

The fact that $U^{\mathcal{I}}$, and $E_1^{\mathcal{I}}$ are $\mathcal{D}$-stable and $\phi(x)$, $\phi_{E_1}(x)$ are internal references to those predicates implies that $f_y \in U^{\mathcal{I}}$ and $f_z \in E_1^{\mathcal{I}}$. It must also hold that $f_y \sim_{\mathcal{D}} f_z^{(1,1,2,2)}$, so we can conclude that $f_z^{(1,1,2,2)} \in U^{\mathcal{I}}$. By Claim 5 this implies that $f_z^{(1,2,2,2)} \in U^{\mathcal{I}}$ as well. Using that $f_z^{(1,2,2,2)} \sim_{\mathcal{D}} f_x$ we finally obtain that $f_x \in U^{\mathcal{I}}$, proving that $\phi_1(x)$ is an internal reference to $U^{\mathcal{I}}$.

Claim 7: The predicate $U^{\mathcal{I}}$ is internal at arity 5**.** We use Lemma 6.1 to prove the claim. We define an internal reference $\phi_{m,n}(x)$ to $U^{\mathcal{I}}$ inductively for each $m, n \in \mathbb{N}$. Additionally, we keep the invariant that $\mathscr{M}_{\mathrm{BLP}} \models \phi_{m,n}(f)$, for each $m, n \in \mathbb{N}$ where $f = (\frac{1}{2^{m-1}2^{n-1}}, 1 - \frac{1}{2^{m-1}3^{n-1}})$. We define $\phi_{1,1}$ as the minor condition ϕ_O. Now, given $m > 1$, we define as

$$\phi_{m,1}(x) \equiv \exists y \exists z \Big(\phi_{m-1,1}(y) \wedge \phi_{E_1}(z) \wedge y = z^{(1,1,2,2)} \wedge x = z^{(1,2,2,2)} \Big).$$

By the previous Claim, the fact that $\phi_{m-1,1}$ is an internal reference to $U^{\mathcal{I}}$, means that so is $\phi_{m,1}$ as well. Moreover, the fact that $(\frac{1}{2^{m-2}}, 1 - \frac{1}{2^{m-2}})$ satisfies $\phi_{m-1,1}(x)$ on $\mathscr{M}_{\mathrm{BLP}}$, means that $(\frac{1}{2^{m-1}}, 1 - \frac{1}{2^{m-1}})$ satisfies $\phi_{m,1}(x)$ on $\mathscr{M}_{\mathrm{BLP}}$ by taking $(\frac{1}{2^{m-2}}, 1 - \frac{1}{2^{m-2}})$ as an existential witness for y and

$$\left(\frac{1}{2^{m-2}}, \frac{1}{2^{m-2}}, \frac{1}{2} - \frac{1}{2^{m-2}}, \frac{1}{2} - \frac{1}{2^{m-2}} \right)$$

as an existential witness for z. Arguing in a similar way, if $n > 1$, we define

$$\begin{aligned} \phi_{m,n}(x) = \exists y \exists z \Big(& \phi_{m,n-1}(y) \wedge \phi_{E_2}(z) \wedge \\ & y = z^{(1,1,1,2,2)} \wedge x = z^{(1,2,2,2,2)} \Big). \end{aligned}$$

Again, using the previous claim we obtain that $\phi_{m,n}$ is an internal reference to $U^{\mathcal{I}}$. One can also see that $(\frac{1}{2^{m-1}3^{n-1}}, 1 - \frac{1}{2^{m-1}3^{n-1}})$ satisfies $\phi_{m,n}(x)$ on $\mathscr{M}_{\mathrm{BLP}}$. This shows the claim.

Proof of items (2),(3), and (4) of Theorem 3.2. Observe that $\mathrm{ar}(U^{\mathcal{I}}) = 2$, and $\mathrm{ar}(Q) = 2$ for all $Q \in \mathcal{D}$. Thus items (3) and (4) of Theorem 3.2 follow from Proposition 4.2 together with Theorem 8.1 and Theorem 8.3 using the interpretation $\mathcal{I}$ and the description $\mathcal{D}$ in this section.

Finally, to prove item (2) of Theorem 3.2 we define a 5-ary pattern Ψ of internal references to $\mathcal{I}$ w.r.t. $\mathcal{D}$. Given $\boldsymbol{I} \in \mathrm{Hom}(\boldsymbol{\Gamma}, \cdot)$ we construct the minor condition $\Psi_{\boldsymbol{I}}$ as follows. First, we find in polynomial time a homomorphism $v \mapsto (m_v, n_v)$ from $\boldsymbol{I}$ to $\boldsymbol{\Gamma}$ such that $\max_{v \in I} m_v + n_v \leqslant |I| + 1$. Then

$$\Psi_{\boldsymbol{I}} \equiv \exists_{v \in I} x_v \exists_{R \in \Sigma_\Gamma, r \in R^I} x_r \left(\bigwedge_{v \in I} \phi_{m_v, n_v}(x_v) \right) \bigwedge \left(\bigwedge_{R \in \Sigma_\Gamma, r \in R^I} \phi_R(x_r) \bigwedge_{i \in [\mathrm{ar}(R)]} x_r^{\Pi^{\mathcal{I}}_{R,i}} = x_{r(i)} \right).$$

Here, for each $R \in \Sigma_R$, $\phi_R(x)$ is the internal definition of $R^{\mathcal{I}}$ given in the previous claims, and $\phi_{m,n}(x)$ is the internal reference to $U^{\mathcal{I}}$ defined in Claim 7. To see that $\Psi_{\boldsymbol{I}}$ can be computed in polynomial time, observe that $\phi_{m,n}(x)$ takes $O(m+n)$ time to be constructed inductively, and $\max_{v \in I} m_v + n_v \leqslant |I| + 1$. The fact that $\mathscr{M}_{\mathrm{BLP}} \models \Psi_{\boldsymbol{I}}$ can be proven similarly to Theorem 3.1-(1) in Section 9.1.1, using that the map $v \mapsto (\frac{1}{2^{m-1}3^{n-1}}, 1 - \frac{1}{2^{m-1}3^{n-1}})$ is a homomorphism from $\mathcal{I}$ to $\boldsymbol{\Gamma}$. ☐

9.3 The BLP + AIP Algorithm

We prove Theorem 3.3 in this section. This result follows from a suitable interpretation of the super-grid $\boldsymbol{\Gamma}^+$ over $\mathscr{M}_{\mathrm{BLP+AIP}}$, shown in Section 9.3.1.

9.3.1 BLP + AIP: Interpreting the Super-Grid

Let $\mathscr{M} = \mathscr{M}_{\text{BLP+AIP}}$ for the rest of this section. Recall that elements in $\mathscr{M}(n)$ consists of pairs (f, g) where $f \in [0,1]^n$ and $g \in \mathbb{Z}^n$. We consider the lexicographical order on pairs $(s, m) \in [0,1] \times \mathbb{Z}$. We perform arithmetic on tuples coordinate-wise. We write $m \vdots n$ for $n, m \in \mathbb{Z}$ to denote that there exists an integer $o \in \mathbb{Z}$ such that $on = m$, and $m \not\vdots n$ for the negation of this statement [9]. Similarly, given $(s, m), (t, n) \in [0,1] \times \mathbb{Z}$, we write $(s, m) \vdots (t, n)$ if there is some integer $o \geqslant 0$ such that $o(t, n) = (s, m)$.

The constructions and arguments in this section are very similar to those in Section 9.2.1, so we will omit some details. The following interpretation $\mathcal{I}$ induces a global structure that is finitely equivalent to $\mathbf{\Gamma}^+$.

Interpretation 5. The Σ_{Γ^+}-interpretation $\mathcal{I}$ over $\mathscr{M}_{\text{BLP+AIP}}$ is given by

$$U^{\mathcal{I}} = \left\{ (f, g) \in \mathscr{M}(5) \,\middle|\, f(4) < 1/3, g(4) \not\vdots 2, \text{ and } (f(i), g(i)) \vdots (f(4), g(4)) \text{ for each } i \in [3] \right\}$$

$$O^{\mathcal{I}} = \left\{ (f, g) \in \mathscr{M}(5) \,\middle|\, f(1) < 1/3,\ g(1) \not\vdots 2, \text{ and } (f(i), g(i)) = (f(4), g(4)) \text{ for all } i \in [3] \right\}, \text{ and}$$

$$\Pi^{\mathcal{I}}_{O,1} = \mathrm{id}\,.$$

And, for all $i \in [3]$

$$E_i^{\mathcal{I}} = \left\{ (f, g) \in \mathscr{M}(6) \,\middle|\, (f(4), g(4)) = (f(5), g(5)) \right\}, \quad \text{and} \quad \Pi^{\mathcal{I}}_{E_i, j} = \begin{cases} (1,2,3,4,5,5) \text{ for } j = 1, \\ (1,2,3,i,4,5) \text{ for } j = 2, \end{cases}$$

$$\mathbb{E}_i^{\mathcal{I}} = \left\{ (f, g) \in \mathscr{M}(6) \,\middle|\, (f(4), g(4)) = (f(i), g(i)) \right\}, \quad \text{and} \quad \Pi^{\mathcal{I}}_{\mathbb{E}_i, j} = \begin{cases} (1,2,3,4,5,5) \text{ for } j = 1, \\ (1,2,3,i,4,5) \text{ for } j = 2. \end{cases}$$

As in Section 9.2.1, we follow the idea of interpreting increasingly large grids and keeping track of the increments in each of these grids. Similarly to before, if we represent a pair $(f, g) \in U^{\mathcal{I}}$ as $(f, g) = ((f(1), g(1)), \ldots, (f(5), g(5)))$, it holds that the first three coordinates of this vector (which are two-dimensional vectors themselves) are integer multiples of $(f(4), g(4))$, which takes the role of the unit increment. Again, the fact that we consider tuples (f, g) where $f(4)$ can be arbitrarily small but positive means that the global structure given by $\mathcal{I}$ contains arbitrarily large grids. This yields the following claim.

Claim 1: The structure $S = S_{\mathcal{I}}$ induced by $\mathcal{I}$ is finitely equivalent to $\mathbf{\Gamma}^+$. Given $j \in \mathbb{N}$, we define $\mathbf{\Gamma}_j^+$ as the substructure of $\mathbf{\Gamma}^+$ induced on the elements $(m, n, o) \in \mathbb{N}^3$ satisfying $m + n + o < j$. Consider an element $(f, g) \in U^{\mathcal{I}}$. Let $j = \lceil \frac{1}{f(4)} \rceil$. Then $\mathbf{\Gamma}_j^+$ is isomorphic to the connected component of (f, g) in S via the bijection $(m_1, m_2, m_3) \mapsto (f', g')$, where $(f'(i), g'(i)) = m_i(f(4), g(4))$ for $i \in [3]$, $(f'(4), g'(4)) = (f(4), g(4))$, and $(f'(5), g'(5))$ is defined so that both the elements in f' and the elements in g' add up to 1. We have that $\mathbf{\Gamma}^+$ is finitely equivalent to the disjoint union $\bigsqcup_{j \in \mathbb{N}} \mathbf{\Gamma}_j^+$, so this proves the result.

Let us define a description $\mathcal{D}$ so that $\mathcal{I}$ is internal at arity 6 w.r.t. $\mathcal{D}$.

---

[9] The more standard notation $n \mid m$ would quickly lead to readability issues.

Description 5. The description $\mathcal{D} \subseteq 2^{\mathscr{M}_{\mathrm{BLP+AIP}}}$ is given by the predicate $U^{\mathcal{I}}$ defined in the previous interpretation, and

$$D_{<} = \Big\{(f,g) \in \mathscr{M}(3) \,\Big|\, (f(1),g(1)) < (f(2),g(2))\Big\},$$
$$D_{\not\vdots 2} = \Big\{(f,g) \in \mathscr{M}(2) \,\Big|\, f(1) < 1/3,\ g(1) \not\vdots\, 2\Big\}.$$

Let us show that $\mathcal{I}$ is internal at arity 6 (w.r.t. $\mathcal{D}$). This is a consequence of the following claims.

Claim 2: The predicates $U^{\mathcal{I}}, O^{\mathcal{I}}$, and $E_i^{\mathcal{I}}$, $\mathbb{E}_i^{\mathcal{I}}$ for $i \in [3]$ are all $\mathcal{D}$-stable.

Claim 3: The predicate $E_i^{\mathcal{I}}$ is internal at arity 6 for each $i \in [3]$. This is shown in a similar way to Claim 3 in Section 9.2.1. The following is an internal definition of $E_i^{\mathcal{I}}$:

$$\phi_{E_i}(x) \equiv x^{(3,3,3,1,2,3)} = x^{(3,3,3,2,1,3)}.$$

Claim 4: The predicate $\mathbb{E}_i^{\mathcal{I}}$ is internal at arity 6 for each $i \in [3]$. This is shown in a similar way to Claim 4 in Section 9.2.1. The following is an internal definition of $E_i^{\mathcal{I}}$:

$$\phi_{E_i}(x) \equiv x^{\sigma} = x^{\tau},$$

where $\sigma : [6] \to [3]$ sends i to 1, 4 to 2 and the other elements to 3, and $\tau : [6] \to [3]$ sends i to 2, 4 to 1 and the other elements to 3.

Claim 5: The predicate $D_{\not\vdots 2}$ is internal at arity 6. Indeed, the following is an internal definition of $D_{\not\vdots 2}$:

$$\phi_{\not\vdots 2}(x) \equiv \exists^5 y\Big(y^{(1,2,3,3,3)} = y^{(2,1,3,3,3)} \wedge y^{(1,3,2,3,3)} = y^{(2,3,1,3,3)} \wedge y^{(3,3,3,1,2)} = y^{(3,3,3,2,1)} \wedge x = y^{(1,2,2,2,2)}\Big).$$

We need to show both that (f,g) satisfy $\phi_{\not\vdots 2}(x)$ on $\mathscr{M}$ for all $(f,g) \in D_{\not\vdots 2}$ and that $\phi_{\not\vdots 2}$ is an internal reference to $D_{\not\vdots 2}$. Let us begin with the first statement. Let $(f,g) \in D_{\not\vdots 2}$. Then let $(f_y,g_y) \in \mathscr{M}(5)$ be defined by $(f_y(i),g_y(i)) = (f(1),g(1))$, for all $i \in [3]$, and $(f_y(i),g_y(i)) = \frac{1}{2}(1-f(1),1-g(1))$ for $i = 4,5$. The fact that $(f_y,g_y) \in \mathscr{M}$ follows from the definition of $D_{\not\vdots 2}$. Observe that (f,g) satisfies $\phi_{\not\vdots 2}(x)$ on $\mathscr{M}$ with (f_y,g_y) as an existential witness for y.

Now let us show that $\phi_{\not\vdots 2}$ is an internal reference to $D_{\not\vdots 2}$. Suppose that $\mathscr{M}/\mathcal{D} \models \phi_{\not\vdots 2}(\langle f,g\rangle)$ with $\langle f_y,g_y\rangle$ as an existential witness for y. We claim that $(f_y(i),g_y(i)) = (f_y(1),g_y(1))$ for $i \in [3]$, and $(f_y(4),g_y(4)) = (f_y(5),g_y(5))$. Indeed, suppose that $(f_y(1),g_y(1)) = (f_y(2),g_y(2))$ Then exactly one of $(f_y,g_y)^{(1,2,3,3,3)}$ or $(f_y,g_y)^{(2,1,3,3,3)}$ would belong to $D_{<}$, contradicting $(f_y,g_y)^{(1,2,3,3,3)} \sim_{\mathcal{D}} (f_y,g_y)^{(2,1,3,3,3)}$. The other identities can be argued similarly. Hence, we have that $3g_y(1) + 2g_y(4) = 1$, where $g_y(1), g_y(4) \in \mathbb{Z}$. This forces $g_y(1) \not\vdots\, 2$ and $g_y(4) \not\vdots\, 3$. In particular $g_y(1), g_y(4) \neq 0$, so $f_y(1), f_y(4) > 0$ by the definition of $\mathscr{M}$. We also have that $3f_y(1) + 2f_y(4)$, so we obtain $f_y(1) < 1/3$ using that $f_y(4) > 0$. All of this together implies that $(f_y,g_y)^{(1,2,2,2,2)} \in D_{\not\vdots 2}$. Finally, the fact that $(f,g) \sim_{\mathcal{D}} (f_y,g_y)^{(1,2,2,2,2)}$ shows that $(f,g) \in D_{\not\vdots 2}$, as we wanted to prove.

Claim 6: The predicate $O^{\mathcal{I}}$ is internal at arity 6. The formula $\phi_O(x)$ defined below is an internal definition of $O^{\mathcal{I}}$.

$$\exists y\Big(\phi_{\not{\colon}2}(y) \wedge y = x^{(1,2,2,2,2)} \wedge x^{(1,2,3,3,3)} = x^{(2,1,3,3,3)} \\ \wedge x^{(1,3,2,3,3)} = x^{(2,3,1,3,3)} \wedge x^{(1,3,3,2,3)} = x^{(2,3,3,1,3)}\Big).$$

Claim 7: Let $i \in [3]$. Suppose $\phi(x)$ is an internal reference to $U^{\mathcal{I}}$. Then the following formula is also an internal reference to $U^{\mathcal{I}}$:

$$\phi'(x) \equiv \exists y \exists^6 z\Big(\phi(y) \wedge \phi_{E_i}(z) \wedge y = z^{(1,2,3,4,5,5)} \wedge x = z^{(1,2,3,4,i,5)}\Big).$$

Moreover, if (f,g) satisfies $\phi(x)$ (on $\mathscr{M}_{\mathbf{BLP}}$), then (f',g') satisfies $\phi'(x)$, where

- $(f'(i), g'(i)) = (f(i), g(i)) + (f(4), g(4))$
- $(f'(5), g'(5)) = (f(5), g(5)) - (f(4), g(4))$**, and**
- $(f'(j), g'(j)) = (f(j), g(j))$**, for $j \neq i, 5$.**

This is shown exactly as Claim 8 in Section 9.2.1.

Claim 8: Let $i \in [3]$. Suppose $\phi(x)$ is an internal reference to $U^{\mathcal{I}}$. Then the following formula is also an internal reference to $U^{\mathcal{I}}$:

$$\phi'(x) \equiv \exists y \exists^6 z\Big(\phi(y) \wedge \phi_{\mathbb{E}_i}(z) \wedge y = z^{(1,2,3,4,5,5)} \wedge x = z^{(1,2,3,4,i,5)}\Big).$$

Moreover, if (f,g) satisfies $\phi(x)$ (on $\mathscr{M}_{\mathbf{BLP}}$), then (f',g') satisfies $\phi'(x)$, where

$(f'(i), g'(i)) = (2f(i), 2g(i)),$
$(f'(5), g'(5)) = (f(5), g(5)) - (f(i), g(i)),$ **and**
$(f'(j), g'(j)) = (f(j), g(j)),$ **for $j \neq i, 5$.**

Again, this follows similarly to Claim 8 in Section 9.2.1.

Claim 9: The predicate $U^{\mathcal{I}}$ is internal at arity 6. We prove the claim using Lemma 6.1. For each $(m_1, m_2, m_3) \in \mathbb{N}^3$ we define an internal reference $\phi_{m_1,m_2,m_3}(x)$ to $U^{\mathcal{I}}$ inductively following the lexicographical order on $\mathbb{N}^3$ in such a way that $\mathscr{M} \models \phi_{m_1,m_2,m_3}((f,g))$ for all elements $(f,g) \in U^{\mathcal{I}}$ satisfying that $(f(i), g(i)) = m_i(f(4), g(4))$ for each $i \in [3]$. We define

$$\phi_{1,1,1}(x) \equiv \phi_O(x).$$

Now, let $(m_1, m_2, m_3) \in \mathbb{N}^3$ and suppose $i \in [3]$ is the greatest index for which $m_i > 1$. We have two cases. Suppose m_i is odd. Then we let $n_j = m_j$ for $j \neq i$, and $n_i = m_i - 1$, and define

$$\phi_{m_1,m_2,m_3}(x) \equiv \exists y \exists^6 z\left(\phi_{n_1,n_2,n_3}(y) \wedge \phi_{E_i}(z) \wedge y = z^{(1,2,3,4,5,5)} \wedge x = z^{(1,2,3,4,i,5)}\right).$$

Otherwise, suppose m_i is even. Then we let $n_j = m_j$ for $j \neq i$, and $n_i = m_i/2$, and define

$$\phi_{m_1,m_2,m_3}(x) \equiv \exists y \exists^6 z \left(\phi_{n_1,n_2,n_3}(y) \wedge \phi_{\mathbb{E}_i}(z) \wedge y = z^{(1,2,3,4,5,5)} \wedge x = z^{(1,2,3,4,i,5)} \right).$$

Now the statement follows from Claims 7 and 8.

Proof of Theorem 3.3. Observe that $5 \geqslant \mathrm{ar}(U^{\mathcal{I}})$ and $5 \geqslant \mathrm{ar}(P)$ for all $P \in \mathcal{D}$. Theorem 3.3-(2) follows from Theorem 8.3 and Proposition 5.2 using that $\boldsymbol{S}_{\mathcal{I}}$ is finitely equivalent to $\boldsymbol{\Gamma}^+$, and $\mathcal{I}$ is internal at arity 6 w.r.t. $\mathcal{D}$.

We prove Theorem 3.3-(1) using Theorem 8.6 and Proposition 5.2. To do this it is enough to define a 6-ary pattern Ψ of internal references to $\mathcal{I}$ w.r.t. $\mathcal{D}$. We map each finite structure $\boldsymbol{I} \in \mathrm{Hom}(\boldsymbol{\Gamma}^+, \cdot)$ to the minor condition $\Psi_{\boldsymbol{I}}$ defined as follows. First, we compute a homomorphism $v \mapsto (m_v, n_v, o_v)$ from $\boldsymbol{I}$ to $\boldsymbol{\Gamma}^+$ in polynomial time, in such a way that $\max_{v \in I} \log_2(m_v n_v o_v) \leqslant |I|$. Then we set

$$\Psi_I \equiv \exists_{v \in I} x_v \exists_{R \in \Sigma_{\Gamma^+}, r \in R^I} x_r \left(\bigwedge_{v \in I} \phi_{m_v,n_v,o_v}(x_v) \right) \bigwedge \left(\bigwedge_{R \in \Sigma_{\Gamma^+}, r \in R^I} \phi_R(x_r) \bigwedge_{i \in \mathrm{ar}(R)} x_r^{\Pi^{\mathcal{I}}_{R,i}} = x_{r(i)} \right).$$

To see that $\Psi_{\boldsymbol{I}}$ can be computed in polynomial time, observe that $\phi_{m,n,o}$ takes $O(\log_2(mno))$ time to be constructed inductively, and $\max_{v \in I} \log_2(m_v n_v o_v) \leqslant |I|$. In order to prove that $(\mathcal{I}, \mathcal{D}, \Psi, 6)$ is a valid pattern we just need to show that $\mathscr{M} \models \Psi_{\boldsymbol{I}}$. We give an explicit homomorphism from $\boldsymbol{I}$ to $\boldsymbol{S}_{\mathcal{I}}$. This map is defined by $v \mapsto (f_v, g_v)$, where

$$f_v = \left(\frac{m_v}{j}, \frac{n_v}{j}, \frac{o_v}{j}, \frac{1}{j}, 1 - \frac{m_v + n_v + o_v + 1}{j} \right),$$

where $j = \max_{v \in I} m_v + n_v + o_v + 1$, and

$$g_v = (m_v, n_v, o_v, 1, 1 - (m_v + n_v + o_v + 1)).$$

Now $\mathscr{M} \models \Psi_{\boldsymbol{I}}$ can be shown the same way as in the proof of Theorem 3.1-(1) in Section 9.1.1. □

9.4 Weak Near-Unanimity Polymorphisms

In this section we prove items (2)(i-iv) of Theorem 1.2. We begin by introducing a minion $\overline{\mathscr{W}}$ that characterizes the existence of WNUs suitably. Then, item (1)(i) will follow from interpreting the grid $\boldsymbol{\Gamma}$ over $\overline{\mathscr{W}}$ (Section 9.4.1), and items (1)(ii-iv) from interpreting growing triangular slices $\boldsymbol{\nabla}_n$ over $\overline{\mathscr{W}}$ (Section 9.4.2).

Given $k \geqslant 2$, we introduce a minion $\mathscr{W}_k$ that characterizes the existence of a n-ary w.n.u. We define an auxiliary minion first. Let $\mathscr{W}_k$ be the minion whose n-ary elements are the ordered partitions of [2]. That is,

$$\mathscr{W}_k(n) = \left\{ \gamma \in \left(2^{[k]}\right)^n \;\middle|\; \bigcup_{i \in [n]} \gamma(i) = [k], \text{ and } \gamma(i) \cap \gamma(j) = \varnothing \text{ for all } i \neq j \right\}.$$

Given a map $\pi : [n] \to [m]$ and elements $\omega \in \mathscr{W}_k(n)$ and $\gamma \in \mathscr{W}_k(m)$, the identity $\gamma = \omega^\pi$ holds if $\gamma(i) = \bigcup_{j \in \pi^{-1}(i)} \omega(j)$. We write $\gamma \sim_k \omega$ for two elements $\gamma, \omega \in \mathscr{W}_k(n)$ if for any $\pi : [n] \to [2]$ we have that $\gamma^\pi = \omega^\pi$, that $|\gamma^\pi(1)| = |\omega^\pi(1)| = 1$, or that $|\gamma^\pi(1)| = |\omega^\pi(1)| = k-1$. Clearly the equivalence relation $\sim_k$ is compatible with minoring (i.e. $\gamma \sim_k \omega$ implies $\gamma^\pi \sim_k \omega^\pi$ for all suitable maps π), so we can define $\overline{\mathscr{W}}_k = \mathscr{W}_k / \sim_k$. Given $\gamma \in \mathscr{W}_k$, we write $\overline{\gamma}$ to denote its $\sim_k$-class.

Lemma 9.1. *Let $\mathscr{M}$ be a minion and $p \in \mathbb{N}$ be a prime number. Then $\mathscr{M}$ contains a w.n.u. element of arity k if and only if $\overline{\mathscr{W}}_k \to \mathscr{M}$.*

Proof. Let $\omega \in \mathscr{W}_k(k)$ be the element defined by $\omega(i) = \{i\}$ for each $i \in [k]$. Suppose there is a homomorphism $\alpha : \overline{\mathscr{W}_k} \to \mathscr{M}$. Then $\alpha(\overline{\omega})$ must be a k-ary w.n.u.

In the other direction, suppose that $f \in \mathscr{M}(k)$ is a w.n.u. Then we define a homomorphism $\alpha : \overline{\mathscr{W}_k} \to \mathscr{M}$ by setting $\alpha(\overline{\omega}) = f$. This defines the homomorphism completely. Indeed, for any $\gamma \in \mathscr{W}_k(n)$ we have that $\gamma = \omega^{\pi_\gamma}$, where for each $i \in [k]$ we have $\pi_\gamma(i) = j$ if and only if $i \in \gamma(j)$. Hence, we define $\alpha(\overline{\gamma}) = f^{\pi_\gamma}$. In order to prove that α is well-defined we need to prove that $f^{\pi_{\gamma_1}} = f^{\pi_{\gamma_2}}$ for any $\gamma_1, \gamma_2 \in \mathscr{W}_k$ satisfying $\gamma_1 \sim \gamma_2$. Suppose that $\gamma_1, \gamma_2 \in \mathscr{W}_k(n)$ are elements satisfying $\gamma_1 \sim \gamma_2$ but $\gamma_1 \neq \gamma_2$. Then there must be indices $i, j \in [n]$ satisfying $|\gamma_1(i)| = |\gamma_2(i)| = 1$ and $|\gamma_1(j)| = |\gamma_2(j)| = k-1$. Let $\tau : [2] \to [n]$ be the map $1 \mapsto i$, $2 \mapsto j$. Then for $s = 1, 2$ we have $\pi^{\gamma_s} = \tau \circ \sigma_s$, where $\sigma_s : [k] \to [2]$ satisfies $|\sigma_s^{-1}(1)| = 1$ and $|\sigma_s^{-1}(2)| = k-1$. The fact that f is a w.n.u. implies that $f^{\sigma_1} = f^{\sigma_2}$, so $f^{\pi_{\gamma_1}} = f^{\pi_{\gamma_2}}$, as we wanted to prove. ☐

We define $\overline{\mathscr{W}}$ as the disjoint union $\bigsqcup_{k \geqslant 3} \overline{\mathscr{W}}_k$. A straight-forward corollary of last lemma is the following.

Corollary 9.2. *Let $\mathscr{M}$ be a minion. Then $\mathscr{M}$ contains a w.n.u. of each arity $k \geqslant 3$ if and only if $\overline{\mathscr{W}} \to \mathscr{M}$.*

9.4.1 WNUs: Interpreting the Grid

The following interpretation $\mathcal{I}$ induces a global structure that is finitely equivalent to $\mathbf{\Gamma}$

Interpretation 6. The Σ_Γ-interpretation $\mathcal{I}$ over $\overline{\mathscr{W}}$ is given by

$$U^{\mathcal{I}} = \overline{\mathscr{W}}(3),$$

$$O^{\mathcal{I}} = \left\{ \overline{\omega} \in \overline{\mathscr{W}}(3) \;\middle|\; \omega(1) = \omega(2) = \varnothing \right\}, \quad \text{and} \quad \Pi^{\mathcal{I}}_{O,1} = \text{id},$$

and, for each $i \in [2]$,

$$E_i^{\mathcal{I}} = \left\{ \overline{\omega} \in \overline{\mathscr{W}}(4) \;\middle|\; |\omega(3)| = 1 \right\}, \quad \text{and} \quad \Pi^{\mathcal{I}}_{E_i,j} = \begin{cases} (1,2,3,3) \text{ for } j = 1, \\ (1,2,i,3) \text{ for } j = 2. \end{cases}$$

Given an integer $k \geqslant 3$ we define $\mathcal{I}_k$ as the restriction $\mathcal{I}|_{\overline{\mathscr{W}}_k}$. Then we have that $\mathbf{S}_{\mathcal{I}} = \bigsqcup_{k \geqslant 3} \mathbf{S}_{\mathcal{I}_k}$.

The idea behind this interpretation relatively simple (the difficult part will be to show that its predicates are internal w.r.t. some suitable description). We represent grid elements $(n, m) \in \mathbb{N}$ as elements $\overline{\omega} \in \overline{\mathscr{W}}(3)$ where the partition ω satisfies $|\omega(1)| = n-1$ and $|\omega(2)| = m-1$. This is correspondence is far from a bijection, but we only need to show finite equivalence.

Claim 1: The structure $S_{\mathcal{I}}$ induced by $\mathcal{I}$ is finitely equivalent to $\mathbf{\Gamma}$. Given a number $m \in \mathbb{N}$, we define $\mathbf{\Gamma}_m$ to be the substructure of $\mathbf{\Gamma}$ induced on the elements $(n, o) \in \mathbb{N}^2$ satisfying $n + o \leqslant m$. Observe that $\mathbf{\Gamma}$ is finitely equivalent to the disjoint union $\bigsqcup_{i \in \mathbb{N}} \mathbf{\Gamma}_i$, and that $\mathbf{\Gamma}_i \to \mathbf{\Gamma}_j$ for each pair $i \leqslant j$. We prove that $\mathbf{\Gamma}_{k+2}$ is finitely equivalent to $\mathbf{S}_{\mathcal{I}_k}$ for each $k \geqslant 3$. Observe that this proves the claim. We define suitable homomorphisms. Let $F : \mathbf{S}_{\mathcal{I}_k} \to \mathbf{\Gamma}_{k+2}$ be the map $\overline{\omega} \mapsto (|\omega(1)| + 1, |\omega(2)| + 1)$. To see that F is well defined, observe that the relation $\sim_k$ preserves the size of sets. That is, if $\omega_1 \sim_k \omega_2$ for some $\omega_1, \omega_2 \in \mathscr{W}_k$, then $|\omega_1(i)| = |\omega_2(i)|$ for all i. The fact that F is indeed a homomorphism follows from the definition of $\mathbf{S}_{\mathcal{I}_i}$.

Now let $H : \mathbf{\Gamma}_{k+2} \to \mathbf{S}_{\mathcal{I}_k}$ be the map given by $(m, n) \mapsto \overline{\omega_{m,n}}$, where $\omega_{m,n} \in \mathscr{C}_p(3)$ is defined as $(X_m, Y_n, [k] \backslash (X_m \cup Y_n))$, where $X_m = [m-1]$ and $Y_n = \{k - n + 2, \dots k\}$. Observe that for $m = n = 1$ we have $X_m = Y_n = \varnothing$. Hence, $H(1,1) \in O^{S_{\mathcal{I}_k}}$. To see that H is a homomorphism we need to prove that H preserves E_1 and E_2. We show the statement for E_1, the other case is analogous. In other words, we need to prove that $(\overline{\omega_{m,n}}, \overline{\omega_{m+1,n}}) \in E_1^{S_{\mathcal{I}_k}}$ for all $((m,n),(m+1,n)) \in E_1^{\Gamma_{k+2}}$. Consider the element $\omega = (X_m, Y_n, \{m\}, [k] \backslash (X_m \cup Y_n \cup \{m\}))$. Then we have that $\overline{\omega} \in E_1^{\mathcal{I}}$, and

$$\overline{\omega}_{m,n} = \overline{\omega}^{\Pi^{\mathcal{I}}_{E_1,1}}, \quad \overline{\omega}_{m+1,n} = \overline{\omega}^{\Pi^{\mathcal{I}}_{E_1,2}},$$

as we wanted to prove. This completes the proof of the claim.

We define a description $\mathcal{D}$ so that $\mathcal{I}$ is internal at arity 4 (w.r.t. $\mathcal{D}$).

Description 6. The description $\mathcal{D} \subseteq 2^{\overline{\mathscr{W}}}$ consists of the predicates

$D_0 = \{\overline{\omega} \in \overline{\mathscr{C}}(2) \mid \omega(1) = \varnothing\}$,
$D_1 = \{\overline{\omega} \in \overline{\mathscr{C}}(2) \mid |\omega(1)| = 1\}$,
$D_* = \{\overline{\omega} \in \overline{\mathscr{C}}(2) \mid |\omega(1)|, |\omega(2)| \neq 1, \text{ and } 1 \in \omega(1)\}$.

Given $\square \in \{0, 1, *\}$ we also define the auxiliary binary predicate $C_\square \in 2^{\overline{\mathscr{W}}}$ which contains all elements $\overline{\omega}$ such that $\overline{\omega}^{(2,1)} \in D_\square$. This way,

$$\left(\bigsqcup_{\square \in \{0,1,*\}} D_\square \right) \bigsqcup \left(\bigsqcup_{\square \in \{0,1,*\}} C_\square \right)$$

is a partition of $\overline{\mathscr{W}}(2)$.

We warn the reader that we deal with two nested equivalence relations from now on: An element $\langle \overline{\omega} \rangle \in \overline{\mathscr{W}}/\mathcal{D}$ is a $\sim_{\mathcal{D}}$-class of some $\overline{\omega} \in \overline{\mathscr{W}}$, which is in turn a $\sim_k$-class of an element $\omega \in \mathscr{W}_k$ for some integer $k \geqslant 3$.

Now the task at hand is showing that $\mathcal{I}$ is internal with respect to $\mathcal{D}$. The crux is proving that the predicate D_1 is internal w.r.t. $\mathcal{D}$, as that predicate is what enables us to define unit increments. This seems intuitive: after all, the elements $\overline{\omega} \in \overline{\mathscr{W}}(2)$ with $|\omega(1)| = 1$ are the ones that explain the existence of WNU elements in $\overline{\mathscr{W}}$. The intuition behind the internal reference to D_1 shown in Claim 4 is that, starting from an element of the form $\overline{\omega} = \overline{([k], \varnothing')}$, the elements of D_1 are the only ones which allow us to transfer all the weight of $\overline{\omega}$ from its first coordinate to its second coordinate by taking $\mathcal{D}$-equivalent "scoops". We show that $\mathcal{I}$ is internal at arity 4 (w.r.t. $\mathcal{D}$) through the following claims.

Claim 2: The predicates $U^{\mathcal{I}}$, $O^{\mathcal{I}}$, $E_1^{\mathcal{I}}$ and $E_2^{\mathcal{I}}$ are all $\mathcal{D}$-stable. This follows from the definitions.

Claim 3: The predicate $O^{\mathcal{I}}$ is internal at arity 4**.** Indeed, the following formula is an internal definition of $O^{\mathcal{I}}$:

$$\phi_O(x) \equiv x = x^{(3,3,3)}.$$

Claim 4: Let $m \geqslant 1$ be an integer. Then the following is an internal reference to D_1:

$$\phi_{1,m}(x^2) \equiv \underset{i\in[m+2]}{\overset{2}{\exists}} y_i \underset{j\in[m]}{\overset{4}{\exists}} z_j \left(y_1 = y_1^{(1,1)} \wedge y_{m+1} = y_{m+1}^{(2,2)} \right) \bigwedge$$

$$\left(\bigwedge_{i\in[m]} x = z_i^{(2,1,2,2)} \wedge x = z_i^{(2,2,1,2)} \wedge y_i = z_i^{(1,1,1,2)} \wedge y_{i+1} = z_i^{(1,1,2,2)} \wedge y_{i+2} = z_i^{(1,2,2,2)} \right).$$

We recall the intuition given at the start of this series of claims. Consider an element $\overline{\omega_x} \in \overline{\mathscr{W}}$ such that $\overline{\mathscr{W}}/\mathcal{D} \models \phi_{1,m}(\langle \overline{\omega_x} \rangle)$ and consider a satisfying assignment witnessing this, given by $y_i \mapsto \langle \overline{\omega_{y_i}} \rangle$ and $z_i \mapsto \langle \overline{\omega_{z_i}} \rangle$ for all suitable indices i. What this formula indicates is that all the mass of ω_{y_1} is in its first coordinate, and all the mass of $\omega_{y_{m+1}}$ is in its second one. The sequence $\omega_{y_1}, \ldots, \omega_{y_{m+1}}$ is obtained by successively taking $\mathcal{D}$-equivalent scoops out of the first coordinate of ω_{y_i} and placing them into its second coordinate, obtaining $\omega_{y_{i+1}}$. The key idea is that this is only possible if the scoops are of size one. Otherwise, only one of them would contain the element 1 and would not be $\mathcal{D}$-equivalent to the rest. This motivates the predicate $D_* \in \mathcal{D}$. Now we give the formal argument.

Suppose that $\overline{\mathscr{C}}/\mathcal{D} \models \phi_{1,m}(\langle \overline{\omega} \rangle)$, with $\langle \overline{\omega_{y_i}} \rangle$ and $\langle \overline{\omega_{z_i}} \rangle$ as witnesses for each variable y_i, z_i. We prove that $\overline{\omega} \in D_1$. Suppose that $\overline{\omega} \notin D_1$ for the sake of contradiction. Then $\overline{\omega}$ must belong to either D_0, D_*, C_0, C_1, C_*. We rule out each of the possibilities by case analysis.

Suppose that $\overline{\omega} \in D_0$. Then, $\overline{\omega}_{z_i}^{(2,1,2,2)}, \overline{\omega}_{z_i}^{(2,2,1,2)} \in D_0$ for all $i \in [m]$. This means that for all $i \in [m]$ we have $\omega_{z_i}(2) = \omega_{z_i}(3) = \varnothing$, and $\overline{\omega}_{z_i}^{(1,1,1,2)} = \overline{\omega}_{z_i}^{(1,1,2,2)} = \overline{\omega}_{z_i}^{(1,2,2,2)}$. This yields that $\overline{\omega}_{y_i} \sim_{\mathcal{D}} \overline{\omega}_{y_{i+1}} \sim_{\mathcal{D}} \overline{\omega}_{y_{i+2}}$ for all $i \in [m]$, so in particular $\overline{\omega}_{y_1} \sim_{\mathcal{D}} \overline{\omega}_{y_{m+2}}$. This implies $\overline{\omega}_{y_1} \in D_0 \cap C_0$, a contradiction.

Suppose that $\overline{\omega} \in D_*$. This implies that $\overline{\omega_{z_i}}^{(2,1,2,2)}, \overline{\omega_{z_i}}^{(2,2,1,2)} \in D_*$, for all $i \in [m]$. However, this means that $1 \in \omega_{z_i}(2)$, and $1 \in \omega_{z_i}(3)$, contradicting the fact that ω_{z_i} is an ordered partition.

Suppose that $\overline{\omega} \in C_1$. Then, $\overline{\omega}_{z_i}^{(2,1,2,2)}, \overline{\omega}_{z_i}^{(2,2,1,2)} \in D_*$ for all $i \in [m]$. This means that

$$|\omega_{z_i}(1)| + |\omega_{z_i}(3)| + |\omega_{z_i}(4)| = |\omega_{z_i}(1)| + |\omega_{z_i}(2)| + |\omega_{z_i}(4)| = 1,$$

This implies that $|\omega_{z_i}(2)| = |\omega_{z_i}(3)| \leqslant 1$. However, these identities together with

$$|\omega_{z_i}(1)| + |\omega_{z_i}(2)| + |\omega_{z_i}(3)| + |\omega_{z_i}(4)| \geqslant 3$$

also imply that $|\omega_{z_i}(2)|, |\omega_{z_i}(3)| \geqslant 2$. This yields a contradiction.

Suppose that $\overline{\omega} \in C_0$. Then, $\overline{\omega_{z_i}}^{(2,1,2,2)}, \overline{\omega_{z_i}}^{(2,2,1,2)} \in C_0$ for all $i \in [m]$. The first inclusion means that $\omega_{z_i}(1) = \omega_{z_i}(3) = \omega_{z_i}(4) = \varnothing$, and the second means that $\omega_{z_i}(1) = \omega_{z_i}(2) = \omega_{z_i}(4) = \varnothing$. This means that all entries of ω_{z_i} contain the empty set, contradicting the fact that ω_{z_i} is an ordered partition of $[k]$ for some $k \geqslant 3$.

Suppose that $\overline{\omega} \in C_*$. This is the hardest case. We have that $\overline{\omega}_{z_i}^{(2,1,2,2)}, \overline{\omega}_{z_i}^{(2,2,1,2)} \in C_*$ for all $i \in [m]$. Using the facts that $\overline{\omega}_{y_1} \in C_0$ and $\overline{\omega}_{y_1} \sim_{\mathcal{D}} \overline{\omega}_{z_1}^{(1,1,1,2)}$ we obtain that $\omega_{z_1}(4) = \varnothing$. This way, $\overline{\omega}_{y_2} \sim_{\mathcal{D}} \overline{\omega}_{z_1}^{(1,1,2,2)} = \overline{\omega}_{z_1}^{(1,1,2,1)}$, and in particular $\overline{\omega}_{y_2} \in D_*$. Let $2 < j \leqslant m+2$ be the smallest index satisfying $\overline{\omega_{y_j}} \notin D_*$. Such index must exist because D_0 and D_* are disjoint, and $\overline{\omega}_{y_{m+2}} \in D_0$. We prove that $\overline{\omega y_j} \in D_1$. Indeed, we have both

$$\overline{\omega}_{y_{j-1}} \sim_{\mathcal{D}} \overline{\omega}_{z_{j-2}}^{(1,1,2,2)}, \text{ and } \overline{\omega}_{y_j} \sim_{\mathcal{D}} \overline{\omega}_{z_{j-2}}^{(1,2,2,2)}.$$

By our choice of j, it must be that $\overline{\omega}_{y_{j-1}}, \overline{\omega}_{z_{j-2}}^{(1,1,2,2)} \in D_*$, meaning that $1 \notin \omega_{z_{j-1}}(3) \sqcup \notin \omega_{z_{j-2}}(4)$. Additionally, the fact that $\overline{\omega}_{z_{j-2}}^{(2,1,2,2)} \in C_*$ implies that $1 \notin \omega_{z_{j-2}}(2)$ as well. Hence, $1 \in \omega_{z_{j-2}}(1)$. This implies that $\overline{\omega}_{z_{j-2}}^{(1,2,2,2)}$ belongs to either D_1 or D_*. By our choice of j the first case must hold, yielding $\overline{\omega}_{y_j} \in D_1$. Thus $j < m+2$, because $\overline{\omega}_{y_{m+2}} \in D_0$. Now consider the element $\omega_{z_{j-1}}$. The fact that $\overline{\omega}_{y_j} \sim_{\mathcal{D}} \overline{\omega}_{z_{j-1}}^{(1,1,2,2)}$ implies that $|\omega_{z_{j-1}}(1)| + |\omega_{z_{j-1}}(2)| = 1$. However, the fact that $\overline{\omega}_{z_{j-1}}^{(2,1,2,2)} \in D_*$ implies that $|\omega_{z_{j-1}}(2)| > 1$, a contradiction.

Claim 5: Let $m \geqslant 1$ be an integer and let $\overline{\omega} \in \overline{\mathscr{W}}_{m+2}(2) \cap D_1$. Then $\overline{\mathscr{W}} \models \phi_{1,m}(\overline{\omega})$. We find witnesses for each variable in $\phi_{1,m}$. For each $i \in [m+2]$, let $\omega_{y_i} \in \mathscr{W}_{m+2}(2)$ be defined as

$$\omega_{y_i}(1) = [m+2-i],$$

and

$$\omega_{y_i}(2)\{m+3-i, \ldots, m+2\}.$$

Similarly, given $i \in [m]$, we define $\omega_{z_i} \in \mathscr{W}_{m+2}(4)$ as

$$\begin{aligned} &\omega_{z_i}(1) = [m-i], \quad \omega_{z_i}(2) = \{m+1-i\}, \\ &\omega_{z_i}(3) = \{m+2-i\}, \quad \text{ and} \\ &\omega_{z_i}(4) = \{m+3-i, \ldots, m+2\}. \end{aligned}$$

Now it is routine to check that $\overline{\mathscr{W}} \models \phi_{1,m}(\overline{\omega})$ with $\overline{\omega}_{y_i}$ as a witness for y_i for each $i \in [m+2]$, and $\overline{\omega}_{z_i}$ as a witness for z_i for each $i \in [m]$.

Claim 6: Let $i \in [2]$. Then $E_i^{\mathcal{I}}$ is internal at arity 4. By Claim 4, given an integer $m \geqslant 1$, the following is an internal reference to $E_i^{\mathcal{I}}$:

$$\phi_{E,m}(x) \equiv \phi_{1,m}(x^{(2,2,1,2)}).$$

Moreover, by Claim 5, if $\overline{\omega} \in \overline{\mathscr{W}}_{m+2} \cap E_i^{\mathcal{I}}$, then $\overline{\mathscr{W}} \models \phi_{E,m}(\overline{\omega})$. Hence the claim follows from Lemma 6.1.

Proof of item (2)(i) of Theorem 1.2. Observe that $3 \geqslant \text{ar}(U^{\mathcal{I}})$ and $3 \geqslant \text{ar}(P)$ for all $P \in \mathcal{D}$. The claims in this section show that $\boldsymbol{S}_{\mathcal{I}}$ is finitely equivalent to $\boldsymbol{\Gamma}$, and $\mathcal{I}$ is internal at arity 4 w.r.t. $\mathcal{D}$. Hence, the result follows from Theorem 8.1 together with Proposition 4.2. It is enough to consider templates of the form $(\boldsymbol{K}_3^4, \boldsymbol{B})$ in the statement. □

9.4.2 WNUs: Interpreting Triangular Slices

The following interpretation $\mathcal{I}$ over $\overline{\mathscr{W}}$ induces structures $\mathbf{S}_{\mathcal{I}_m}$ that are homomorphically equivalent to $\mathbf{\nabla}_m$ for each integer $m \geqslant 3$, where $\mathcal{I}_m = \mathcal{I}|_{\overline{\mathscr{W}}_m}$.

Interpretation 7. The Σ_∇-interpretation $\mathcal{I}$ over $\overline{\mathscr{W}}$ is defined as in Interpretation 6, by adding

$$W^{\mathcal{I}} = \{\overline{\omega} \in \overline{\mathscr{W}}(3) \mid \omega(3) = \varnothing\}, \quad \text{and} \quad \Pi^{\mathcal{I}}_{W,1} = \mathrm{id}\,.$$

Claim 1: For each integer k, the structures $\mathbf{\nabla}_k$ and $S_{\mathcal{I}_k}$ are homomorphically equivalent. The maps defined in Claim 1 in Section 9.4.1 are homomorphisms in both directions between $\mathbf{\nabla}_k$ and $\mathbf{S}_{\mathcal{I}_k}$.

We define the description $\mathcal{D} \subseteq 2^{\overline{\mathscr{W}}}$ in the same way as in Description 6. We claim $\mathcal{I}$ is internal at arity 4 (w.r.t. $\mathcal{D}$). The following, together with the claims from Section 9.4.1, proves the statement.

Claim 2: The predicate $W^{\mathcal{I}}$ is internal at arity 4. Clearly $W^{\mathcal{I}}$ is $\mathcal{D}$-stable. Additionally, the following is an internal definition of $W^{\mathcal{I}}$:

$$\phi_W(x) \equiv x^{(1,1,2)} = x^{(1,1,1)}.$$

Proof of items (2)(ii-iv) of Theorem 1.2. Observe that $3 \geqslant \mathrm{ar}(U^{\mathcal{I}})$ and $3 \geqslant \mathrm{ar}(P)$ for all $P \in \mathcal{D}$. Let $\mathcal{I}$ and $\mathcal{D}$ be the Σ_∇-interpretation over $\overline{\mathscr{C}}$ and the description over $\overline{\mathscr{W}}$ given in this section. The claims in this section and the previous one show that for each integer $k \geqslant 3$, the structure $\mathbf{S}_{\mathcal{I}_k}$ is homomorphically equivalent to $\mathbf{\nabla}_{k+2}$, where $\mathcal{I}_k = \mathcal{I}|_{\overline{\mathscr{W}}_k}$, and $\mathcal{I}$ is internal at arity 4 w.r.t. the description $\mathcal{D}$. Then the result follows from Proposition 5.3 together with Theorem 8.2. It is enough to consider templates of the form $(\boldsymbol{K}_3^4, \boldsymbol{B})$ in the statement. $\square$

9.5 Cyclic Polymorphisms

In this section we prove items (1)(i-iv) of Theorem 1.2. We begin by introducing a minion $\overline{\mathscr{C}}$ that characterizes the existence of cyclic polymorphisms suitably. Then, item (1)(i) will follow from interpreting the grid $\mathbf{\Gamma}$ over $\overline{\mathscr{C}}$ (Section 9.5.1), and items (1)(ii-iv) from interpreting growing triangular slices $\mathbf{\nabla}_n$ over $\overline{\mathscr{C}}$ (Section 9.5.2).

Given prime arity. Given a prime number $p \in \mathbb{N}$, we define the minion $\mathscr{C}_p$ as follows. We let

$$\mathscr{C}(n) = \left\{ \gamma \in \left(2^{\mathbb{Z}_p}\right)^n \;\middle|\; \bigcup_{i \in [n]} \gamma(i) = \mathbb{Z}_p, \text{ and } \gamma(i) \cap \gamma(j) = \varnothing \text{ for all } i \neq j \right\}.$$

The minoring operation is defined as follows. Let $\gamma \in \mathscr{C}(n)$, $\pi : [n] \to [m]$. Then $\gamma^\pi = \omega$, where $\omega(j) = \bigcup_{i \in \pi^{-1}(j)} \gamma(i)$ for each $j \in [m]$, and empty unions yield the empty set.

Given a set $S \subseteq \mathbb{Z}_p$ and an element $m \in \mathbb{Z}_p$ we write $S + m$ for the set $\{n + m \mid n \in S\}$. Given two elements $\gamma, \omega \in \mathscr{C}_p(n)$, we write $\gamma \sim_p \omega$ if there is some element $m \in \mathbb{Z}_p$ such that $\gamma(i) = \omega(i) + m$

for all $i \in [n]$. Observe that $\sim_p$ is an equivalence relation and it is compatible with minoring, in the sense that $\gamma \sim_p \omega$ implies $\gamma^\pi \sim_p \omega^\pi$. We write $\overline{\gamma}$ to denote the $\sim_p$-class of an element γ, and define $\overline{\mathscr{C}_p}$ as the quotient minion $\mathscr{C}_p / \sim_p$.

Lemma 9.3. *Let $\mathscr{M}$ be a minion and $p \in \mathbb{N}$ be a prime number. Then $\mathscr{M}$ contains a cyclic element of arity p if and only if $\overline{\mathscr{C}}_p \to \mathscr{M}$.*

Proof. We show both directions. Suppose there is a homomorphism $\alpha : \overline{\mathscr{C}}_p \to \mathscr{M}$. Let $\gamma = (\{0\}, \{1\}, \dots, \{p-1\}) \in \mathscr{C}_p(p)$. Then the element $\overline{\gamma} \in \overline{\mathscr{C}}_p(p)$ is cyclic, so $\alpha(\overline{\gamma})$ must be cyclic as well.

In the other direction, suppose that $f \in \mathscr{M}(p)$ is a cyclic element. Then we define a homomorphism $\alpha : \overline{\mathscr{C}}_p \to \mathscr{M}$ by setting $\alpha(\overline{\gamma}) = f$. This defines the image of any element $\overline{\omega} \in \mathscr{C}_p(n)$. Indeed, we have that $\omega = \gamma^{\pi_\omega}$, where $\pi_\omega : [p] \to [n]$ is the map that sends $i \in [p]$ to $j \in [n]$ if the element $i-1 \in \mathbb{Z}_p$ belongs to $\omega(j)$. Hence, we can define $\alpha(\overline{\omega}) = f^{\pi_\omega}$. To see that this is well defined, we need to prove that whenever $\omega_1 \sim_p \omega_2$ then $f^{\pi_{\omega_1}} = f^{\pi_{\omega_2}}$. However, if $\omega_1 \sim_p \omega_2$, then there is some $m \in \mathbb{Z}_p$ such that $\omega_1(i) = \omega_2(i) + m$ for all i. This means that $\pi_{\omega_2} = \pi_{\omega_1} \circ \sigma^m$, where $\sigma = (p-1, 1, \dots, p-2)$ is the cyclic shift. Hence, as f is cyclic, we obtain that $f^{\pi_{\omega_2}} = (f^{\sigma^m})^{\pi_{\omega_1}} = f^{\omega_1}$, as we wanted. We have shown that α is a well-defined map. The fact that α is a minion homomorphism now follows from the fact that if $\omega_1 = \omega_2^\pi$, for some elements $\omega_1, \omega_2 \in \mathscr{C}_p$ then $\pi_{\omega_1} = \pi \circ \pi_{\omega_2}$. $\square$

We define $\overline{\mathscr{C}}$ as the disjoint union $\bigsqcup_{p \text{ prime}} \overline{\mathscr{C}}_p$. A straight-forward corollary of last lemma is the following.

Corollary 9.4. *Let $\mathscr{M}$ be a minion. Then $\mathscr{M}$ contains a cyclic element of each prime arity p if and only if $\overline{\mathscr{C}} \to \mathscr{M}$.*

9.5.1 Cyclic Polymorphisms: Interpreting the Grid

The following interpretation $\mathcal{I}$ that induces a global structure which is finitely equivalent to $\mathbf{\Gamma}$. We note that this interpretation is completely analogous to Interpretation 8, defined over $\overline{\mathscr{C}}$.

Interpretation 8. The Σ_Γ-interpretation $\mathcal{I}$ over $\overline{\mathscr{C}}$ is given by

$$U^{\mathcal{I}} = \overline{\mathscr{C}}(3),$$

$$O^{\mathcal{I}} = \left\{ \overline{\omega} \in \overline{\mathscr{C}}(3) \,\middle|\, \omega(1) = \omega(2) = \varnothing \right\}, \quad \text{and} \quad \Pi^{\mathcal{I}}_{O,1} = \mathrm{id},$$

and, for each $i \in [2]$,

$$E_i^{\mathcal{I}} = \left\{ \overline{\omega} \in \overline{\mathscr{C}}(4) \,\middle|\, |\omega(3)| = 1 \right\}, \quad \text{and} \quad \Pi^{\mathcal{I}}_{E_i,j} = \begin{cases} (1,2,3,3) \text{ for } j = 1, \\ (1,2,i,3) \text{ for } j = 2. \end{cases}$$

Given a prime number p we define $\mathcal{I}_p$ as the restriction $\mathcal{I}|_{\overline{\mathscr{C}}_p}$. Then we have that $S_{\mathcal{I}} = \bigsqcup_{p \text{ prime}} S_{\mathcal{I}_p}$.

Claim 1: The structure $S_{\mathcal{I}}$ induced by $\mathcal{I}$ is finitely equivalent to $\mathbf{\Gamma}$. This follows similarly to Claim 1 in Section 9.4.1. Given a number $m \in \mathbb{N}$, we define $\mathbf{\Gamma}_m$ to be the substructure of $\mathbf{\Gamma}$ induced on the elements $(n, o) \in \mathbb{N}^2$ satisfying $n + o \leqslant m$. Observe that $\mathbf{\Gamma}$ is finitely equivalent to the disjoint union $\bigsqcup_{i \in \mathbb{N}} \mathbf{\Gamma}_i$, and that $\mathbf{\Gamma}_i \to \mathbf{\Gamma}_j$ for each pair $i \leqslant j$. We prove that $\mathbf{\Gamma}_{p+2}$ is finitely equivalent to $S_{\mathcal{I}_p}$

for each prime number p. Observe that this proves the claim. We define suitable homomorphisms. Let $F : \mathbf{S}_{\mathcal{I}_p} \to \mathbf{\Gamma}_{p+2}$ be the map $\overline{\omega} \mapsto (|\omega(1)| + 1, |\omega(2)| + 1)$. To see that F is well defined, observe that the relation $\sim_p$ preserves the size of sets. That is, if $\omega_1 \sim_p \omega_2$ for some $\omega_1, \omega_2 \in \mathscr{C}_p$, then $|\omega_1(i)| = |\omega_2(i)|$ for all i. The fact that F is indeed a homomorphism follows from the definition of $\mathbf{S}_{\mathcal{I}_i}$.

Now let $H : \mathbf{\Gamma}_{p+2} \to \mathbf{S}_{\mathcal{I}_p}$ be the map given by $(m, n) \mapsto \overline{\omega_{m,n}}$, where $\omega_{m,n} \in \mathscr{C}_p(3)$ is defined as $(X_m, Y_n, \mathbb{Z}_p \backslash (X_m \cup Y_n))$, where $X_m = \{0, \ldots, m-1\}$ and $Y_n = \{p - n + 1, \ldots p - 1\}$. Observe that for $m = n = 1$ we have $X_m = Y_n = \varnothing$. Hence, $H(1,1) \in O^{S_{\mathcal{I}_p}}$. To see that H is a homomorphism we need to prove that H preserves E_1 and E_2. We show the statement for E_1, the other case is analogous. In other words, we need to prove that $(\overline{\omega_{m,n}}, \overline{\omega_{m+1,n}}) \in E_1^{S_{\mathcal{I}_p}}$ for all $((m,n),(m+1,n)) \in E_1^{\Gamma_{p+2}}$. Consider the element $\omega = (X_m, Y_n, \{m\}, \mathbb{Z}_p \backslash (X_m \cup Y_n \cup \{m\}))$. Then we have that $\overline{\omega} \in E_1^{\mathcal{I}}$, and

$$\overline{\omega}_{m,n} = \overline{\omega}^{\Pi^{\mathcal{I}}_{E_1,1}}, \quad \overline{\omega}_{m+1,n} = \overline{\omega}^{\Pi^{\mathcal{I}}_{E_1,2}},$$

as we wanted to prove. This completes the proof of the claim.

We define a description $\mathcal{D}$ so that $\mathcal{I}$ is internal at arity 5 w.r.t. $\mathcal{D}$.

Description 7. The description $\mathcal{D} \subseteq 2^{\overline{\mathscr{C}}}$ consists of the predicates

$$D_0 = \left\{ \overline{\omega} \in \overline{\mathscr{C}}(2) \,\middle|\, \omega(1) = \varnothing \right\},$$
$$D_1 = \left\{ \overline{\omega} \in \overline{\mathscr{C}}(2) \,\middle|\, |\omega(1)| = 1 \right\},$$
$$D_< = \left\{ \overline{\omega} \in \overline{\mathscr{C}}(3) \,\middle|\, |\omega(1)| < |\omega(2)| \right\},$$
$$D_\div = \left\{ \overline{\omega} \in \overline{\mathscr{C}}(3) \,\middle|\, |\omega(2)| = n|\omega(1)| \text{ for some integer } n \geqslant 0 \right\}.$$

We also consider the auxiliary predicate

$$D_= = \left\{ \overline{\omega} \in \overline{\mathscr{C}}(3) \,\middle|\, |\omega(1)| = |\omega(2)| \right\}.$$

As in Section 9.4, we warn the reader again that we deal with two nested equivalence relations from now on: An element $\langle \overline{\omega} \rangle \in \overline{\mathscr{C}}/\mathcal{D}$ is a $\sim_{\mathcal{D}}$-class of some $\overline{\omega} \in \overline{\mathscr{C}}$, which is in turn a $\sim_p$-class of an element $\omega \in \mathscr{C}_p$ for some prime p.

As in the previous section about WNU polymorphisms, here the more involved part is showing that $\mathcal{I}$ is internal with respect to $\mathcal{D}$. Here the important step is showing that the predicate D_1 is internal, because it allows us to define unit increments. We can show that D_1 is internal if $D_\div$ is internal, because the elements in D_1 are precisely the classes of tuples ω where $\omega(1)$ is a set whose size divides the size of $\omega(2)$ (here it is crucial that we are dealing with cyclic elements of prime arity). To speak about division, we need to speak about equality, so the most important part is proving that $D_=$ is internal, as shown in Claim 4. The intuition is that, if we represent $\mathbb{Z}_p$ in a circle in the usual way, an element ω with $\overline{\omega} \in D_=$ selects two points in the circle, given by $\omega(1) = \{a\}$ and $\omega(2) = \{b\}$, and we can think of the pair (a, b) as a "step". Then the idea is that in $\overline{\mathscr{C}}$ one can return to a from b taking $\mathcal{D}$-equivalent steps. We show that $\mathcal{I}$ is internal at arity 5 (w.r.t. $\mathcal{D}$) through the following claims.

Claim 2: The predicates $U^{\mathcal{I}}$, $O^{\mathcal{I}}$, $E_1^{\mathcal{I}}$ and $E_2^{\mathcal{I}}$ are all $\mathcal{D}$-stable. This is a routine check.

Claim 3: The predicate $O^{\mathcal{I}}$ is internal at arity 5**.** Indeed, the following formula is an internal definition of $O^{\mathcal{I}}$:

$$\phi_O(x) \equiv x = x^{(3,3,3)}.$$

Claim 4: Let p be a prime number. Then the following predicate is an internal reference to $D_=$:

$$\phi_{=,p}(x) \equiv x = x^{(2,1,3)}$$

if $p = 2$, and

$$\phi_{=,p}(x) \equiv \mathop{\exists}\limits_{i\in[p-2]}^{4} y_i \mathop{\exists}\limits_{i\in[p-3]}^{5} z_i \left(y_1^{(1,2,3,3)} = x \wedge y_1^{(3,1,2,3)} = x \wedge y_{p-2}^{(2,3,1,3)} = x \right) \bigwedge$$
$$\left(\bigwedge_{i\in[p-3]} z_i^{(1,2,3,3,3)} = z_i^{(3,3,1,2,3)} \wedge y_i = z_i^{(1,2,3,4,4)} \wedge y_{i+1} = z_i^{(1,2,4,3,4)} \right)$$

if $p > 2$. Let us show that $\phi_{=,p}(x)$ is an internal reference. The case $p = 2$ is straightforward. We assume that $p \geqslant 3$. Suppose that $\overline{\mathscr{C}}/\mathcal{D} \models \phi_{=,p}(\langle \overline{\omega_x} \rangle)$ with $\langle \overline{\omega_{y_i}} \rangle$ as the witness for y_i for each $i \in [p-2]$ and $\langle \overline{\omega_{z_i}} \rangle$ as the witness for z_i for each $i \in [p-3]$. We repeat the intuition given before this chain of claims. The idea is that ω_x fixes a step inside some cyclic group. I.e., $\omega_x = (\{a\}, \{b\}, \mathbb{Z}_p \backslash \{a, b\})$. Then, up to equivalence in $\overline{\mathscr{C}}/\mathcal{D}$, each element ω_{y_i} is of the form $(\{a\}, \{b\}, \{c_i\}, \mathbb{Z}_q \backslash \{a, b, c\})$, where $c_{i+1} = c_i + b - a$, and $c_1 = 2b - a$. Then the idea is that at the end we have come full circle, so to say, and obtain $c_{p-2} = 2a - b$. Now we proceed with the formal argument.

First, we show that $|\omega_x(1)| = |\omega_x(2)|$. Suppose that $|\omega_x(1)| < |\omega_x(2)|$ for a contradiction (the reverse inequality can be dealt with analogously). We prove that $|\omega_{y_i}(1)| < |\omega_{y_i}(2)| < |\omega_{y_i}(3)|$ for all $i \in [p-2]$ by induction on i. For $i = 1$, we have that $\overline{\omega_{y_1}}^{(1,2,3,3)} \sim_{\mathcal{D}} \overline{\omega_x}$, and $\overline{\omega_{y_1}}^{(3,1,2,3)} \sim_{\mathcal{D}} \overline{\omega_x}$, so necessarily $|\omega_{y_1}(1)| < |\omega_{y_1}(2)| < |\omega_{y_1}(3)|$. Now let $i > 1$ and suppose that

$$|\omega_{y_{i-1}}(1)| < |\omega_{y_{i-1}}(2)| < |\omega_{y_{i-1}}(3)|.$$

We have that $\overline{\omega_{y_{i-1}}} \sim_{\mathcal{D}} \overline{\omega_{z_{i-1}}}^{(1,2,3,4,4)}$, so

$$|\omega_{z_{i-1}}(1)| < |\omega_{z_{i-1}}(2)| < |\omega_{z_{i-1}}(3)|.$$

Additionally, $\overline{\omega_{z_{i-1}}}^{(1,2,3,3,3)} \sim_{\mathcal{D}} \overline{\omega_{z_{i-1}}}^{(3,3,1,2,3)}$, so

$$|\omega_{z_{i-1}}(3)| < |\omega_{z_{i-1}}(4)|.$$

Finally, $\overline{\omega_{y_i}} \sim_{\mathcal{D}} \overline{\omega_{z_{i-1}}}^{(1,2,4,3,4)}$, so we can conclude that

$$|\omega_{y_i}(1)| < |\omega_{y_i}(2)| < |\omega_{y_i}(3)|,$$

as we wanted to show. We have shown that

$$|\omega_{y_{p-2}}(1)| < |\omega_{y_{p-2}}(2)| < |\omega_{y_{p-2}}(3)|.$$

However, $\overline{\omega_x} \sim_{\mathcal{D}} \overline{\omega_{y_{p-2}}}^{(2,3,1,3)}$ implies that

$$|\omega_{y_{p-2}}(3)| < |\omega_{y_{p-2}}(1)|,$$

a contradiction. Hence, it must be that $|\omega_x(1)| = |\omega_x(2)|$ to begin with.

Claim 5: Let p be a prime, and let $\omega \in \mathscr{C}_p(3)$ be such that $|\omega(1)| = |\omega(2)| = 1$. Then $\overline{\mathscr{C}} \models \phi_{=,p}(\overline{\omega})$. Suppose that $p = 2$. Then $\omega \sim_p (\{0\}, \{1\}, \varnothing) \sim_p (\{1\}, \{0\}, \varnothing)$, proves the statement. Suppose that $p \geqslant 3$. Without loss of generality we can assume that $\omega = (\{0\}, \{m\}, \mathbb{Z}_p \backslash \{0, m\})$ for some $m \in \mathbb{Z}_p$. We find witnesses $\overline{\omega_{y_i}}, \overline{\omega_{z_i}}$, for every variable y_i, z_i. Given $i \in [p-2]$ we define

$$\omega_{y_i} = \big(\{0\},\ \{m\},\ \{(i+1)m\},\ \mathbb{Z}_p \backslash \{0, m, (i+1)m\}\big).$$

Given $i \in [p-3]$, we define

$$\omega_{z_i} = \big(\{0\},\ \{m\},\ \{(i+1)m\},\ \{(i+2)m\},\ \mathbb{Z}_p \backslash \{0, m, (i+1)m, (i+2)m\}\big).$$

Now it is routine to check that our choice of representatives satisfies $\phi_{=,p}(\overline{\omega})$.

Claim 6: Let p be a prime, and let $m \geqslant 0$ be an integer. Then the following is an internal reference to $D_{\div}$:

$$\phi_{\div,p,m}(x) \equiv \underset{i \in [m+1]}{\overset{3}{\exists}} y_i \underset{j \in [m]}{\overset{4}{\exists}} z_j \left(y_1^{(2,2,2)} = y_1^{(2,1,2)} \wedge x = y_{m+1} \right) \wedge$$
$$\left(\bigwedge_{i \in [m]} z_i^{(1,2,3,3)} = y_i \wedge z_i^{(1,2,2,3)} = y_{i+1} \wedge \phi_{=,p}(z_i^{(1,3,2,3)}) \right).$$

Suppose that $\overline{\mathscr{C}}/\mathcal{D} \models \phi_{\div,p,m}(\langle \overline{\omega} \rangle)$, with $\langle \overline{\omega_{y_i}} \rangle$ and $\langle \overline{\omega_{z_i}} \rangle$ as witnesses for each variable y_i, z_i. We prove that $\overline{\omega} \in D_{\div}$. In order to show this, we prove by induction on i that $\overline{\omega_{y_i}} \in D_{\div}$ for each $i \in [m+1]$. This proves the result because $\overline{\omega} \sim_{\mathcal{D}} \overline{\omega_{y_{m+1}}}$. For $i = 1$, we have that $\overline{\omega_{y_1}}^{(2,2,2)} \sim_{\mathcal{D}} \overline{\omega_{y_1}}^{(2,1,2)}$. Given that the first element must belong to D_0, so does the second one, and we obtain that $\omega_{y_1}(2) = \varnothing$. Hence $\overline{\omega_{y_1}} \in D_{\div}$ vacuously. Now let $i > 1$ and assume that $|\omega_{y_i}(1)|$ divides $|\omega_{y_i}(2)|$. Using that $\overline{\omega_{y_i}} \sim_{\mathcal{D}} \overline{\omega_{z_i}}^{(1,2,3,3)}$, we obtain that $|\omega_{z_i}(1)|$ divides $|\omega_{z_i}(2)|$. We also have that $|\omega_{z_i}(1)| = |\omega_{z_i}(3)|$ using that $\phi_{=,p}(x)$ is an internal reference to $D_{=}$. Hence $|\omega_{z_i}(1)|$ divides $|\omega_{z_i}(2)| + |\omega_{z_i}(3)|$. Finally, using that $\overline{\omega_{y_{i+1}}} \sim_{\mathcal{D}} \overline{\omega_{z_i}}^{(1,2,2,3)}$, we obtain that $\overline{\omega_{y_{i+1}}}$ belongs to $D_{\div}$, as we wanted to prove. This shows the claim.

Claim 7: Let p be a prime and let $0 \leqslant m \leqslant p-1$ be an integer. Let $\omega_{p,m} \in \mathscr{C}_p(3)$ be defined as $(\{0\}, \{1, \ldots, m\}, \{m+1, \ldots, p-1\})$. Then $\overline{\mathscr{C}} \models \phi_{\div,p,m}(\overline{\omega_{p,m}})$. We show the claim by defining suitable witnesses for each variable in $\phi_{\div,p,m}$. For each $i \in [m+1]$ we define $\omega_{y_i} = \omega_{p,i-1}$, and for each $i \in [m]$ we define

$$\omega_{z_i} = (\{0\}, \{1, \ldots, i-1\}, \{i\}, \{i+1, \ldots, p-1\}).$$

Now it is routine to check that the elements $\overline{\omega_{y_i}}, \overline{\omega_{z_i}}$ are witnesses for the variables y_i, z_i. The key observation is that $\overline{\omega_{z_i}}^{(1,3,2,3)}$ satisfies $\phi_{=,p}(x)$ on $\overline{\mathscr{C}}$ for all $i \in [m]$ by Claim 5.

Claim 8: The predicate D_1 is internal at arity 5. We apply Lemma 6.1. Let $\omega_p \in \mathscr{C}_p$ be the element $(\{0\}, \mathbb{Z}_p \backslash \{0\})$. We define an internal reference to D_1 that is satisfied by $\overline{\omega_p}$. This interpretation is the following.

$$\phi_{1,p}(x) \equiv \exists^3 y \left(y = x^{(1,2)} \wedge \phi_{\div,p,p-1}(y) \right).$$

Let us show that $\phi_{1,p}$ is an internal reference to D_1. Suppose that $\overline{\mathscr{C}}/\mathcal{D} \models \phi_{1,p}(\langle \overline{\omega_x} \rangle)$ with $\langle \overline{\omega_y} \rangle$ as a witness for y. Observe that $\omega_y(3)$ must be the empty set. Indeed, we have that $\overline{\omega_y} \sim_{\mathcal{D}} \overline{\omega_x}^{(1,2)}$, so $\overline{\omega_y}^{(2,2,1)} \sim_{\mathcal{D}} \overline{\omega_x}^{(2,2)}$, and $\overline{\omega_x}^{(2,2)} \in D_0$, so $\overline{\omega_y}^{(2,2,1)} \in D_0$ as well. Additionally, the formula $\phi_{\div,p,p-1}(x)$ is an internal reference to $D_{\div}$, so $|\omega_y(1)|$ must divide $|\omega_y(2)|$. Both these numbers must add up to some prime q, so the only possibility is that $|\omega_y(1)| = 1$ and $|\omega_y(2)| = q-1$. This means that $\overline{y}^{(1,2,2)}$ belongs to D_1. Finally, using again that $\overline{\omega_y} \sim_{\mathcal{D}} \overline{\omega_x}^{(1,2)}$, we obtain again that $\overline{x}$ belongs to D_1, as we wanted to prove.

Now, in order to see that $\overline{\omega_p}$ satisfies $\phi_{1,p}(x)$ on $\overline{\mathscr{C}}$, just consider $\overline{\omega_y}$ as a witness for y, where ω_y is the tuple $(\{0\}, \mathbb{Z}_p \backslash \{0\}, \varnothing) \in \mathscr{C}_p$, and apply the previous claim.

Claim 9: Let $i \in [2]$. Then the predicate $E_i^{\mathcal{I}}$ is internal at arity 5. We apply Lemma 6.1, as usual. Let $\overline{\omega} \in E_i^{\mathcal{I}}$ be such that $\omega \in \mathscr{C}_p$ for a prime p. Then, using the previous claim we obtain that

$$\phi_{E_i,p}(x) \equiv \phi_{1,p}(x^{(2,2,1,2)})$$

is an internal reference to $E_i^{\mathcal{I}}$ and is satisfied by $\overline{\omega}$ on $\overline{\mathscr{C}}$.

Claim 10: The predicate $U^{\mathcal{I}}$ is internal at arity 5. Trivially, the following is an internal definition of $U^{\mathcal{I}}$:

$$\phi_U(x) \equiv x = x.$$

Proof of Theorem 1.2-(1)(i). Observe that $3 \geqslant \mathrm{ar}(U^{\mathcal{I}})$ and $3 \geqslant \mathrm{ar}(P)$ for all $P \in \mathcal{D}$. The claims in this section show that $\mathbf{S}_{\mathcal{I}}$ is finitely equivalent to $\mathbf{\Gamma}$, and $\mathcal{I}$ is internal at arity 5 w.r.t. $\mathcal{D}$. Hence, the result follows from Theorem 8.1 together with Proposition 4.2. ☐

9.5.2 Cyclic Polymorphisms: Interpreting Triangular Slices

The following interpretation $\mathcal{I}$ over $\overline{\mathscr{C}}$ induces structures $\mathbf{S}_{\mathcal{I}_p}$ that are homomorphically equivalent to $\mathbf{\nabla}_{p+2}$ for each prime number p, where $\mathcal{I}_p = \mathcal{I}|_{\overline{\mathscr{C}_p}}$.

Interpretation 9. The Σ_∇-interpretation $\mathcal{I}$ over $\overline{\mathscr{C}}$ is defined as in Interpretation 8, by adding

$$W^{\mathcal{I}} = \left\{ \overline{\omega} \in \overline{\mathscr{C}}(3) \;\middle|\; \omega(3) = \varnothing \right\}, \quad \text{and} \quad \Pi_{W,1}^{\mathcal{I}} = \mathrm{id}.$$

Claim 1: For each prime p, the structures $\mathbf{\nabla}_{p+2}$ and $S_{\mathcal{I}_p}$ are homomorphically equivalent. The maps defined in Claim 1 in Section 9.5.1 are homomorphisms in both directions between $\mathbf{\nabla}_{p+2}$ and

$\boldsymbol{S}_{\mathcal{I}_p}$.

We define the description $\mathcal{D} \subseteq 2^{\overline{\mathscr{C}}}$ as in Description 7. The interpretation $\mathcal{I}$ is internal at arity 5 (w.r.t. $\mathcal{D}$). The following claim together with the claims from Section 9.5.1 prove the statement.

Claim 2: The predicate $W^{\mathcal{I}}$ is internal at arity 5. Clearly $W^{\mathcal{I}}$ is $\mathcal{D}$-stable. Additionally, the following is an internal definition of $W^{\mathcal{I}}$:

$$\phi_W(x) \equiv x^{(1,1,2)} = x^{(1,1,1)}.$$

Proof of items (1)(ii-iv) of theorem 1.2. Observe that $3 \geqslant \mathrm{ar}(U^{\mathcal{I}})$ and $3 \geqslant \mathrm{ar}(P)$ for all $P \in \mathcal{D}$. The claims in this section and the previous one show that for each prime number p, the structure $\boldsymbol{S}_{\mathcal{I}_p}$ is homomorphically equivalent to $\boldsymbol{\nabla}_{p+2}$, where $\mathcal{I}_p = \mathcal{I}|_{\overline{\mathscr{C}_p}}$, and $\mathcal{I}$ is internal at arity 5 w.r.t. the description $\mathcal{D}$. Then the result follows from Proposition 5.3 together with Theorem 8.2. It is enough to consider templates of the form $(\boldsymbol{K}_3^5, \boldsymbol{B})$ in the statement. □

10 Discussion

This work represents a step towards understanding the relationship between search and decision in promise constraint satisfaction. It is important to insist that the relationship between our results, the Search vs Decision question for PCSPs, and other related questions in the literature [17, 21] is nuanced. While we prove that rounding the output of the algorithms BLP, AIP, and BLP + AIP is hard (in the TFNP sense), it is still possible that, for instance, whenever BLP solves the decision version of PCSP($\boldsymbol{A}, \boldsymbol{B}$), another linear programming relaxation, such as the ones introduced in [17], can be used to solve the search variant of PCSP($\boldsymbol{A}, \boldsymbol{B}$). A counterintuitive possibility is that even if both algorithms BLP and BLP + AIP cannot be used for search separately, it could happen that BLP + AIP could be adapted for search in the templates where decision is solvable via BLP. In other words, it could be the case that $\mathrm{sPCSP}_{\mathrm{BLP+AIP}}(\boldsymbol{A}, \boldsymbol{B})$ is tractable whenever BLP solves PCSP($\boldsymbol{A}, \boldsymbol{B}$).

The Search vs Decision Question We have shown that, conditional to TFNP $\not\subseteq$ FP, not every efficient decision algorithm for PCSP($\boldsymbol{A}, \boldsymbol{B}$) can be turned into an efficient decision algorithm for sPCSP($\boldsymbol{A}, \boldsymbol{B}$) that accepts the same instances. So in this particular sense search PCSPs are more difficult to solve than decision PCSPs, although this is not a complexity-theoretic separation. We remark that TFNP $\not\subseteq$ FP is the weakest assumption under which these questions make sense. If $\mathcal{Q}$ is a polynomial-time algorithm solving PCSP($\boldsymbol{A}, \boldsymbol{B}$), then $\mathrm{sPCSP}_{\mathcal{Q}}(\boldsymbol{A}, \boldsymbol{B})$ can be seen as a problem in TFNP if one considers rejections by $\mathcal{Q}$ as proper search certificates. Hence, $\mathrm{sPCSP}_{\mathcal{Q}}(\boldsymbol{A}, \boldsymbol{B})$ can be solved in polynomial time if TFNP $\subseteq$ FP.

We have considered the problem of obtaining search algorithms from efficient decision algorithms, but standing above is the open question of whether, in the finite-template setting, sPCSP($\boldsymbol{A}, \boldsymbol{B}$) has a polynomial-time solution whenever PCSP($\boldsymbol{A}, \boldsymbol{B}$) does. We do not consider our results strong evidence to the contrary: There is some reason to believe that, say, the third level of the BLP + AIP hierarchy [28] could be used to solve sPCSP($\boldsymbol{A}, \boldsymbol{B}$) for all templates used to prove

Theorems 3.1 to 3.3. We sketch the argument here. Given $\mathcal{Q} \in \{\text{AIP}, \text{BLP}, \text{BLP} + \text{AIP}\}$, those results show hardness of $\text{sPCSP}_{\mathcal{Q}}(\boldsymbol{A}, \boldsymbol{B})$ by proving that $\text{sPMC}_N(\mathscr{M}_{\mathcal{Q}}, \mathscr{N})$ is hard, where $\mathscr{N}$ is some manifold minion built on top of a quotient of $\mathscr{M}_{\mathcal{Q}}$. If, instead, we want to prove the stronger result that $\text{sPCSP}(\boldsymbol{A}, \boldsymbol{B})$ is hard, then we need to show that $\text{sPMC}_N(\mathscr{P}, \mathscr{N})$ is hard, where $\mathscr{P}$ denotes the minion of projections. The reason our proof fails to show this result is that the minor conditions $\Phi_{\boldsymbol{G}}$ that arise in our patterns do not hold in $\mathscr{P}$. In fact, they can be ruled out by a 3-consistency check. Hence, there could be an efficient method of obtaining search certificates for minor conditions that are both accepted by $\mathcal{Q}$ and 3-consistency in a way that does not involve solving the tiling problem encoded in $\mathscr{N}$.

We further observe that a separation of search and decision PCSPs implies a non-dichotomy for search PCSPs unless the polynomial hierarchy collapses to its first level.

Theorem 10.1. *Let* $(\boldsymbol{A}, \boldsymbol{B})$ *be a finite template for which* $\text{PCSP}(\boldsymbol{A}, \boldsymbol{B})$ *has a polynomial-time solution. Suppose that NP≠ coNP. Then* $\text{sPCSP}(\boldsymbol{A}, \boldsymbol{B})$ *is not FNP-hard. In particular, if* $\text{sPCSP}(\boldsymbol{A}, \boldsymbol{B})$ *has no polynomial-time solution, then it must be FNP-intermediate.*

Proof. Let $\mathcal{Q}$ be a polynomial-time algorithm solving $\text{PCSP}(\boldsymbol{A}, \boldsymbol{B})$. Then, by our previous reasoning, $\text{sPCSP}_{\mathcal{Q}}(\boldsymbol{A}, \boldsymbol{B}) \in \text{TFNP}$. The problem $\text{sPCSP}(\boldsymbol{A}, \boldsymbol{B})$ is trivially reducible to $\text{sPCSP}_{\mathcal{Q}}(\boldsymbol{A}, \boldsymbol{B})$. We note that if we see $\text{sPCSP}_{\mathcal{Q}}(\boldsymbol{A}, \boldsymbol{B})$ as a total problem, then this reduction must be seen as a generalized many-one reduction, as it may need to map some answers in $\text{sPCSP}_{\mathcal{Q}}(\boldsymbol{A}, \boldsymbol{B})$ to rejections in $\text{sPCSP}(\boldsymbol{A}, \boldsymbol{B})$. It is known [55, Theorem 2.1] that, unless NP = coNP, there is not a problem in TFNP that is FNP-hard under generalized many-one reductions [10]. Hence $\text{sPCSP}(\boldsymbol{A}, \boldsymbol{B})$ cannot be FNP-hard. □

Small Templates Our techniques produce templates $(\boldsymbol{A}, \boldsymbol{B})$ where $\boldsymbol{B}$ can grow quite large. To construct $\boldsymbol{B}$ we start with a set of tiles $\boldsymbol{T}$ that exhibits some interesting behavior. For instance, if we drop the origin constraint, it is known that the smallest *aperiodic* tile set has 11 elements [46]. Or alternatively, we start with a problem $\Pi \in \text{TFNP}$, translate it into a problem about a particular Turing machine, and encode this machine in a tile set $\boldsymbol{T}$. Then, construct an manifold minion $\mathscr{N}$ using $\boldsymbol{T}$ and a suitable quotient of an interesting minion $\mathscr{M}$. Finally, $\boldsymbol{B}$ is obtained as some free structure of $\mathscr{N}$, although this last step does not have a significant impact on $\boldsymbol{B}$'s size. Therefore it seems safe to assume that our techniques have little to say about templates where the domains of both $\boldsymbol{A}$ and $\boldsymbol{B}$ are small. For instance, it is still possible that AIP, BLP, and BLP + AIP can always be adapted to solve search in Boolean PCSPs.

A somewhat unsatisfactory aspect of our reductions is that they are oblivious to the structure inside TFNP. For example, can we obtain an explicit and relatively small template $(\boldsymbol{A}, \boldsymbol{B})$ such that BLP solves $\text{PCSP}(\boldsymbol{A}, \boldsymbol{B})$ and the rounding problem $\text{sPCSP}_{\text{BLP}}(\boldsymbol{A}, \boldsymbol{B})$ is PPAD-hard (i.e., as hard as the problem of computing Nash equilibria [33])? What can we say about the templates $(\boldsymbol{A}, \boldsymbol{B})$ for which $\text{sPCSP}_{\text{BLP}}(\boldsymbol{A}, \boldsymbol{B})$ is PPAD-hard?

Infinitely Many Symmetric Polymorphisms It is known that the algorithm BLP + AIP solves (the decision variant of) $\text{PCSP}(\boldsymbol{A}, \boldsymbol{B})$ whenever $\text{Pol}(\boldsymbol{A}, \boldsymbol{B})$ contains symmetric polymorphisms of infinitely many arities. It is not known, however, whether $\text{sPCSP}(\boldsymbol{A}, \boldsymbol{B})$ is also tractable for these

---

[10] This result is often misquoted as referring only to many-one reductions, but notice that in that case the statement is vacuous: a non-total search problem cannot be reduced to a total problem via many-one reductions.

templates [17,21,50]. Of course, a negative answer would imply a separation between search and decision for PCSPs. One open possibility is that running the BLP algorithm on a different ring other than $\mathbb{Q}$, such as $\mathbb{Z}[\sqrt{2}]$, is enough to obtain an output that can be rounded in polynomial time and to solve $\mathrm{sPCSP}(\boldsymbol{A},\boldsymbol{B})$. The intuitive reason for this is that those rings avoid "rounding boundaries", which seemed to be the obstacle to widely-applicable rounding procedures. An insight from our results is that this difficulty is intrinsic: when we prove TFNP-hardness, the tiling problems are precisely encoded on the rounding boundaries of the corresponding PCSPs (for instance, on the values of the form $\frac{1}{2^n 3^m}$ in Section 9.2.2). If these boundaries are the root reason the rounding problems are difficult, then avoiding them may be key for obtaining efficient search algorithms.

The work [17] poses a question, later repeated in [21] for the Boolean setting, whose affirmative answer would yield efficient search algorithms for all templates with infinitely many symmetric polymorphisms. This question is whether these templates must also contain an infinite consistent family of so-called "regional-periodic" polymorphisms. We remark that our results do not rule out this possibility, but they restrict the ways in which it might hold. For example, let $(\boldsymbol{A},\boldsymbol{B})$ be a finite template satisfying that BLP solves $\mathrm{PCSP}(\boldsymbol{A},\boldsymbol{B})$ but all homomorphisms $\mathscr{M}_{\mathrm{BLP}} \to \mathrm{Pol}(\boldsymbol{A},\boldsymbol{B})$ are non-computable, as in Theorem 3.2. Then the minion $\mathrm{Pol}(\boldsymbol{A},\boldsymbol{B})$ cannot contain a consistent family of nicely-behaved symmetric polymorphisms of all (emphasis on the word *all*) arities, because they would entail a computable homomorphism $\mathscr{M}_{\mathrm{BLP}} \to \mathrm{Pol}(\boldsymbol{A},\boldsymbol{B})$. Still, it could be the case that one of these families exists, but only containing polymorphisms of *infinitely* many arities.

Other Algorithms A natural next step in light of our results is to study other PCSP algorithms from the same perspective. An obstacle is that we do not know of explicit minion characterizations for k-consistency for $k \geqslant 3$ and other algorithmic hierarchies [28] for any level after the second. Two cases where relatively tame minion characterizations can be obtained are the SDP algorithm [20,28], named after the *semi-definite programming relaxation* [63], and extensions of singleton arc-consistency [34] such as the CLAP algorithm, introduced in [27]. Out of these, SDP is the case most similar to the relaxations studied in this paper. In the minion $\mathscr{M}_{\mathrm{SDP}}$, introduced independently in [20,28], n-ary elements are n-tuples of finitely-supported orthogonal vectors in $\mathbb{R}^{\mathbb{N}}$ that add-up to $e_1 = (1,0,\ldots,0,\ldots)$. Minoring is defined by means of addition, as in the case of $\mathscr{M}_{\mathrm{AIP}}$ and $\mathscr{M}_{\mathrm{BLP}}$. We outline some intuition indicating that our methods may be difficult to apply to $\mathscr{M}_{\mathrm{SDP}}$. An observation is that our reductions exploit the *lack of symmetry* of the minions $\mathscr{M}_{\mathcal{Q}}$ for $\mathcal{Q} \in \{\mathrm{AIP},\mathrm{BLP},\mathrm{BLP}+\mathrm{AIP}\}$: an interpretation $\mathcal{I}$ over $\mathscr{M}_{\mathcal{Q}}$ that is internal (w.r.t. some description) must, in particular, be invariant under the endomorphisms of $\mathscr{M}_{\mathcal{Q}}$. This lack of symmetry also seems necessary to obtain internal references to interesting predicates. In the cases of $\mathscr{M}_{\mathrm{AIP}}$ and $\mathscr{M}_{\mathrm{BLP}}$ the only endomorphisms are the identity maps. However, any isometry of $\mathbb{R}^{\mathbb{N}}$ that fixes the origin and e_1 induces an endomorphism of $\mathscr{M}_{\mathrm{SDP}}$, making this minion extremely symmetric compared to the ones studied in this paper. This theme of lack of symmetry leading to hardness is recurrent in the theory of constraint satisfaction.

Other Meta-Problems We point out several meta-problems whose decidability is still open. All of these are well known in the area, but are most often posed as quests for characterizations. Maybe an equally promising direction would be to consider them as invitations to prove the *absence of effective characterizations*. We believe that resolving those questions could shed light on the ways in which PCSPs may be too expressive, or pinpoint some structure that aids in the further development of the theory.

The meta-problems for virtually all PCSP algorithms referred to in this paper other than AIP, BLP, BLP + AIP remain open. The exception is *arc-consistency*, but tractability through this algorithm is equivalent to pp-constructability from a fixed tractable finite-template CSP (Horn 3-SAT) which is decidable [11]. An important case is the one of k-consistency. Both [4] and [29] have found some sufficient conditions implying that $\mathrm{PCSP}(\boldsymbol{A}, \boldsymbol{B})$ has linear width (and hence is not solved by any fixed level of the local consistency algorithm), but we do not know of any non-trivial necessary conditions. In this context, is also worth going back to the CSP setting. There, the meta-problems for BLP, k-consistency, k-Sherali Adams, are decidable ([16], [8], and [3,5] respectively) and the relationship between those algorithms is well understood. However, for finite-template CSPs we do not have a good understanding of AIP and its derived algorithms (e.g., BLP + AIP, cohomological k-consistency [30], the BLP + AIP hierarchy [28], CLAP [27]), and the related meta-problems are open.

Other than meta-problems related to algorithms, an important open question is whether we can recognize the cases where $\mathrm{PCSP}(\boldsymbol{A}, \boldsymbol{B})$ is *finitely tractable* [2,9], meaning that there is a finite structure $\boldsymbol{C}$ such that $\boldsymbol{A} \to \boldsymbol{C} \to \boldsymbol{B}$ and $\mathrm{CSP}(\boldsymbol{C})$ can be solved in polynomial time. In [47] it has been shown that the size of the smallest witness $\boldsymbol{C}$ of finite tractability can grow quite large compared to $\boldsymbol{A}, \boldsymbol{B}$ (if P$\neq$ NP), suggesting that characterizing this phenomenon may be difficult. In the same direction, another open-problem is the one of recognizing the cases where $\mathrm{PCSP}(\boldsymbol{A}, \boldsymbol{B})$ is solvable in first-order logic. Recently [56] showed that this occurs precisely when $\boldsymbol{A} \to \boldsymbol{C} \to \boldsymbol{B}$ for some finite $\boldsymbol{C}$ such that $\mathrm{CSP}(\boldsymbol{C})$ is definable in first-order logic. The cases where $\mathrm{CSP}(\boldsymbol{C})$ is definable in first-order logic are decidable [53], but in [56] there is no obvious bound on the size of $\boldsymbol{C}$ in terms of $(\boldsymbol{A}, \boldsymbol{B})$, suggesting that this also may be a difficult meta-question.

Finally, the different notions of reductions between PCSPs are another source of interesting problems. Reductions between finite-template PCSPs by means of pp-constructions are characterized by the existence of a homomorphism between the corresponding polymorphism minions [11], which can be shown to be a decidable condition. However, these reductions are provably not enough to obtain all NP-hard finite-template PCSPs from, say, 3-SAT (see the discussion in [50]). Other proposed reductions are the ones given by so-called (d, r)*-homomorphisms* between minions [14], and the more general *local consistency* reductions [32]. These give rise to natural meta-problems: Given finite templates $(\boldsymbol{A}, \boldsymbol{B})$, $(\boldsymbol{A}', \boldsymbol{B}')$ can we decide (1) whether there is a (d, r)-homomorphism from $\mathrm{Pol}(\boldsymbol{A}', \boldsymbol{B}')$ to $\mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$ for any $d, r \in \mathbb{N}$? (2) whether $\mathrm{PCSP}(\boldsymbol{A}, \boldsymbol{B})$ reduces to $\mathrm{PCSP}(\boldsymbol{A}', \boldsymbol{B}')$ via the k-consistency reduction for any $k \in \mathbb{N}$? Another related, and perhaps more accessible question is: given a finite template $(\boldsymbol{A}, \boldsymbol{B})$, can we decide whether $\mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$ has *bounded essential arity*? By this we mean that there is some $d \in \mathbb{N}$ such that every polymorphism $f \in \mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$ depends on at most d variables [11]. This implies the existence of a (strong version of a) $(d, 1)$-homomorphism to the minion of projections $\mathscr{P}$, and is a condition that has been used to prove hardness of some PCSPs (e.g., [6]).

Acknowledgements

I am especially grateful to Lorenzo Ciardo for his comments on the manuscript, and our discussions on the minions characterizing CLAP and singleton arc-consistency. When I was trying to encode tiling problems in $\mathscr{M}_{\mathrm{BLP}}$, he also suggested that $\mathscr{M}_{\mathrm{AIP}}$ could possibly be an easier target, and that approach ultimately worked. I thank Andrei Krokhin for our later discussions on the literature surrounding the search-vs-decision question for PCSPs. I am also thankful to Silvia Butti, Lorenzo Ciardo, Tamio-Vesa Nakajima, and Standa Živný for our early discussions on rounding algorithms. I wish to thank Paul Goldberg for discussing with me the status of the class TFNP_1, and other notions related to TFNP. Thanks should also go to Antoine Mottet for pointing out to me that the class of templates admitting symmetric polymorphisms of all arities was not known to be decidable, and discussing this and other meta-questions with me. Finally, I wish to thank the CWC 2024 attendees for their feedback on a talk presenting this work.

References

[1] Sanjeev Arora and Boaz Barak. *Computational complexity: a modern approach*. Cambridge University Press, 2009.

[2] Kristina Asimi and Libor Barto. Finitely tractable promise constraint satisfaction problems. In *46th International Symposium on Mathematical Foundations of Computer Science (MFCS 2021)*, pages 11–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2021.

[3] Albert Atserias, Andrei Bulatov, and Anuj Dawar. Affine systems of equations and counting infinitary logic. *Theoretical Computer Science*, 410(18):1666–1683, 2009.

[4] Albert Atserias and Víctor Dalmau. Promise constraint satisfaction and width. In *Proc. 2022 ACM-SIAM Symposium on Discrete Algorithms (SODA'22)*, pages 1129–1153, 2022. `arXiv:2107.05886`, `doi:10.1137/1.9781611977073.48`.

[5] Albert Atserias and Elitza Maneva. Sherali-adams relaxations and indistinguishability in counting logics. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pages 367–379, 2012.

[6] Per Austrin, Venkatesan Guruswami, and Johan Håstad. $(2+\varepsilon)$-sat is np-hard. *SIAM Journal on Computing*, 46(5):1554–1573, 2017.

[7] Demian Banakh and Marcin Kozik. Injective hardness condition for pcsps. In *Proceedings of the 39th Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 1–10, 2024.

[8] Libor Barto. The collapse of the bounded width hierarchy. *Journal of Logic and Computation*, 26(3):923–943, 2014.

[9] Libor Barto. Promises make finite (constraint satisfaction) problems infinitary. In *2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–8. IEEE, 2019.

[10] Libor Barto, Diego Battistelli, and Kevin M. Berg. Symmetric Promise Constraint Satisfaction Problems: Beyond the Boolean Case. In *Proc. 38th International Symposium on Theoretical Aspects of Computer Science (STACS'21)*, volume 187 of *LIPIcs*, pages 10:1–10:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. `arXiv:2010.04623`, `doi:10.4230/LIPIcs.STACS.2021.10`.

[11] Libor Barto, Jakub Bulín, Andrei A. Krokhin, and Jakub Opršal. Algebraic approach to promise constraint satisfaction. *J. ACM*, 68(4):28:1–28:66, 2021. `arXiv:1811.00970`, `doi:10.1145/3457606`.

[12] Libor Barto and Marcin Kozik. Constraint satisfaction problems of bounded width. In *2009 50th Annual IEEE symposium on foundations of computer science*, pages 595–603. IEEE, 2009.

[13] Libor Barto and Marcin Kozik. Absorbing subalgebras, cyclic terms, and the constraint satisfaction problem. *Logical Methods in Computer Science*, 8, 2012.

[14] Libor Barto and Marcin Kozik. Combinatorial Gap Theorem and Reductions between Promise CSPs. In *Proc. 2022 ACM-SIAM Symposium on Discrete Algorithms (SODA'22)*, pages 1204–1220, 2022. arXiv:2107.09423, doi:10.1137/1.9781611977073.50.

[15] Libor Barto, Andrei Krokhin, and Ross Willard. Polymorphisms, and how to use them, 2017.

[16] Zarathustra Brady. Notes on csps and polymorphisms. *arXiv preprint arXiv:2210.07383*, 2022.

[17] Joshua Brakensiek and Venkatesan Guruswami. An algorithmic blend of lps and ring equations for promise csps. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 436–455. SIAM, 2019.

[18] Joshua Brakensiek and Venkatesan Guruswami. Promise Constraint Satisfaction: Algebraic Structure and a Symmetric Boolean Dichotomy. *SIAM J. Comput.*, 50(6):1663–1700, 2021. arXiv:1704.01937, doi:10.1137/19M128212X.

[19] Joshua Brakensiek, Venkatesan Guruswami, and Sai Sandeep. Conditional dichotomy of boolean ordered promise csps. *TheoretiCS*, 2, 2023.

[20] Joshua Brakensiek, Venkatesan Guruswami, and Sai Sandeep. SDPs and robust satisfiability of promise CSP. In *Proc. 55th Annual ACM Symposium on Theory of Computing (STOC'23)*, pages 609–622. ACM, 2023. arXiv:2211.08373, doi:10.1145/3564246.3585180.

[21] Joshua Brakensiek, Venkatesan Guruswami, Marcin Wrochna, and Stanislav Živný. The power of the combined basic LP and affine relaxation for promise CSPs. *SIAM J. Comput.*, 49:1232–1248, 2020. arXiv:1907.04383, doi:10.1137/20M1312745.

[22] Alex Brandts, Marcin Wrochna, and Stanislav Živnỳ. The complexity of promise sat on non-boolean domains. *ACM Transactions on Computation Theory (TOCT)*, 13(4):1–20, 2021.

[23] Glen E Bredon. *Sheaf theory*, volume 170. Springer Science & Business Media, 2012.

[24] Harry Buhrman, Lance Fortnow, Michal Kouckỳ, John D Rogers, and Nikolay Vereshchagin. Does the polynomial hierarchy collapse if onto functions are invertible? *Theory of Computing Systems*, 46:143–156, 2010.

[25] Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. In *Proc. 58th Annual IEEE Symposium on Foundations of Computer Science (FOCS'17)*, pages 319–330, 2017. arXiv:1703.03021, doi:10.1109/FOCS.2017.37.

[26] Hubie Chen and Benoit Larose. Asking the metaquestions in constraint tractability. *ACM Transactions on Computation Theory (TOCT)*, 9(3):1–27, 2017.

[27] Lorenzo Ciardo and Stanislav Živný. CLAP: A new algorithm for promise CSPs. *SIAM J. Comput.*, 52(1):1–37, 2023. arXiv:2107.05018, doi:10.1137/22M1476435.

[28] Lorenzo Ciardo and Stanislav Živný. Hierarchies of minion tests for PCSPs through tensors. In *Proc. 2023 ACM-SIAM Symposium on Discrete Algorithms (SODA'23)*, pages 568–580, 2023. arXiv:2207.02277, doi:10.1137/1.9781611977554.ch25.

[29] Lorenzo Ciardo and Stanislav Živnỳ. The periodic structure of local consistency. *arXiv preprint arXiv:2406.19685*, 2024.

[30] Adam Ó Conghaile. Cohomology in constraint satisfaction and structure isomorphism. In *47th International Symposium on Mathematical Foundations of Computer Science (MFCS 2022)*, pages 75–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2022.

[31] Nigel Cutland. *Computability: An introduction to recursive function theory*. Cambridge university press, 1980.

[32] Víctor Dalmau and Jakub Opršal. Local consistency as a reduction between constraint satisfaction problems. In *Proceedings of the 39th Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 1–15, 2024.

[33] Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. The complexity of computing a nash equilibrium. *Communications of the ACM*, 52(2):89–97, 2009.

[34] Romuald Debruyne and Christian Bessière. Some practicable filtering techniques for the constraint satisfaction problem. In *IJCAI (1)*, pages 412–417. Morgan Kaufmann, 1997. URL: `http://dblp.uni-trier.de/db/conf/ijcai/ijcai97.html#DebruyneB97`.

[35] Shimon Even, Alan L Selman, and Yacov Yacobi. The complexity of promise problems with applications to public-key cryptography. *Information and control*, 61(2):159–173, 1984.

[36] Miron Ficak, Marcin Kozik, Miroslav Olsák, and Szymon Stankiewicz. Dichotomy for symmetric boolean pcsps. *arXiv preprint arXiv:1904.12424*, 2019.

[37] Miron Ficak, Marcin Kozik, Miroslav Olšák, and Szymon Stankiewicz. Dichotomy for Symmetric Boolean PCSPs. In *Proc. 46th International Colloquium on Automata, Languages, and Programming (ICALP'19)*, volume 132, pages 57:1–57:12. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. `arXiv:1904.12424`, `doi:10.4230/LIPIcs.ICALP.2019.57`.

[38] Jean H Gallier. *Logic for computer science: foundations of automatic theorem proving*. Courier Dover Publications, 2015.

[39] Eran Gat and Shafi Goldwasser. Probabilistic search algorithms with unique answers and their cryptographic applications. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 18, pages 1–3, 2011.

[40] Paul W Goldberg and Christos H Papadimitriou. Tfnp: an update. In *International Conference on Algorithms and Complexity*, pages 3–9. Springer, 2017.

[41] Oded Goldreich. On promise problems: A survey. In *Theoretical Computer Science: Essays in Memory of Shimon Even*, pages 254–290. Springer, 2006.

[42] Mika Göös, Alexandros Hollender, Siddhartha Jain, Gilbert Maystre, William Pires, Robert Robere, and Ran Tao. Separations in proof complexity and tfnp. *Journal of the ACM*, 71(4):1–45, 2024.

[43] William Hanf. Nonrecursive tilings of the plane. i. *The Journal of Symbolic Logic*, 39(2):283–285, 1974.

[44] John E Hopcroft and Jeffrey D Ullman. *Formal languages and their relation to automata*. Addison-Wesley Longman Publishing Co., Inc., 1969.

[45] Pavel Hubácek, Moni Naor, and Eylon Yogev. The journey from np to tfnp hardness. In *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*, pages 60–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2017.

[46] Emmanuel Jeandel and Michael Rao. An aperiodic set of 11 wang tiles. *arXiv preprint arXiv:1506.06492*, 2015.

[47] Alexandr Kazda, Peter Mayr, and Dmitriy Zhuk. Small promise csps that reduce to large csps. *Logical Methods in Computer Science*, 18, 2022.

[48] Marcin Kozik, Andrei Krokhin, Matt Valeriote, and Ross Willard. Characterizations of several maltsev conditions. *Algebra universalis*, 73(3):205–224, 2015.

[49] Andrei Krokhin, Andrei Bulatov, and Peter Jeavons. The complexity of constraint satisfaction: an algebraic approach. *Structural Theory of Automata, Semigroups, and Universal Algebra: Proceedings of the NATO Advanced Study Institute on Structural Theory of Automata, Semigroups and Universal Algebra Montreal, Quebec, Canada 7–18 July 2003*, pages 181–213, 2005.

[50] Andrei Krokhin and Jakub Opršal. An invitation to the promise constraint satisfaction problem. *ACM SIGLOG News*, 9(3):30–59, 2022.

[51] Andrei Krokhin, Jakub Opršal, Marcin Wrochna, and Stanislav Živný. Topology and adjunction in promise constraint satisfaction. *SIAM J. Comput.*, 52(1):37–79, 2023. `arXiv:2003.11351`, `doi:10.1137/20M1378223`.

[52] Richard E Ladner. On the structure of polynomial time reducibility. *Journal of the ACM (JACM)*, 22(1):155–171, 1975.

[53] Benoit Larose, Cynthia Loten, and Claude Tardif. A characterisation of first-order constraint satisfaction problems. *Logical Methods in Computer Science*, 3, 2007.

[54] Alberto Larrauri and Stanislav Živnỳ. Solving promise equations over monoids and groups. *ACM Transactions on Computational Logic*, 26(1):1–24, 2024.

[55] Nimrod Megiddo and Christos H Papadimitriou. On total functions, existence theorems and computational complexity. *Theoretical Computer Science*, 81(2):317–324, 1991.

[56] Antoine Mottet. Promise and infinite-domain constraint satisfaction. In *32nd EACSL Annual Conference on Computer Science Logic (CSL 2024)*, pages 41–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024.

[57] Antoine Mottet. Algebraic and algorithmic synergies between promise and infinite-domain csps. *arXiv preprint arXiv:2501.13740*, 2025.

[58] Tamio-Vesa Nakajima and Stanislav Zivný. Boolean symmetric vs. functional PCSP dichotomy. In *Proc. 38th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS'23)*, 2023. `doi:10.1109/LICS56636.2023.10175746`.

[59] Tamio-Vesa Nakajima and Stanislav Živnỳ. On the complexity of symmetric vs. functional pcsps. *ACM Transactions on Algorithms*, 20(4):1–29, 2024.

[60] Christos H Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and system Sciences*, 48:498–532, 1994.

[61] Michael Pinsker, Jakub Rydval, Moritz Schöbi, and Christoph Spiess. Three meta-questions on infinite-domain constraint satisfaction problems. *arXiv preprint arXiv:2502.06621*, 2025.

[62] Nicholas Pippenger. Galois theory for minors of finite functions. *Discrete Mathematics*, 254(1-3):405–419, 2002.

[63] Prasad Raghavendra. *Approximating np-hard problems efficient algorithms and their limits*. University of Washington, 2009.

[64] Hao Wang. Dominoes and the aea case of the decision problem. *Computation, Logic, Philosophy: A Collection of Essays*, pages 218–245, 1990.

[65] Dmitriy Zhuk. A proof of the CSP dichotomy conjecture. *J. ACM*, 67(5):30:1–30:78, 2020. `arXiv:1704.01914`, `doi:10.1145/3402029`.

A Proofs of Section 5

We assume familiarity with Turing machines and computability-related notions during this section. We refer to e.g., [1] or [31]. We introduce some related notation below.

A non-deterministic Turing machine M is a tuple (Q, Ω, Δ), where (1) Q is a finite set of states, containing an initial state q_0, an accepting state q_+ and a rejecting state q_-, where $q_+ \neq q_-$, (2) Ω is a finite alphabet containing a distinguished "blank" symbol $\square$, and (3) $\Delta \subseteq (Q \times \Omega) \times (Q \times \Omega \times \{L, R\})$ is the transition relation. The machine M is deterministic if the transition function Δ contains at most one pair whose first element is (q, a) for each $(q, a) \in Q \times \Omega$. We adopt the convention that a Turing machine M is executed on a single right-infinite tape, and the initial position of the head is always the start of the tape. We say a Turing machine is *immortal* if it has an infinite run starting from any finite input word.

Let $M = (Q, \Omega, \Delta)$ be a (non-deterministic) Turing machine and # a fresh symbol. A *configuration* of M is given by a pair $(\boldsymbol{\omega}, \boldsymbol{q}) \in (\Omega \cup \{\#\})^{\mathbb{Z}_{\geqslant 0}} \times (Q \cup \{\#\})^{\mathbb{Z}_{\geqslant 0}}$. The 0-th position of the configuration symbolizes a blank spot to the left of the tape, which starts at position 1. Formally, (1) $\boldsymbol{\omega}$ describes the contents of M's tape from left to right, starting with the symbol # which marks the extra position before the leftmost end of the tape, and (2) $\boldsymbol{q}$ contains a state $q \in Q$ in a single entry and the # symbol in all the others, denoting that the head of M is at the given position in state q. Given $i \in \mathbb{N}$ (observe here indices start at 1), the i-th local description of $(\boldsymbol{\omega}, \boldsymbol{q})$ is the pair $((\omega_{i-1}, \omega_i, \omega_{i+1}), (q_{i-1}, q_i, q_{i+1}))$. We denote by $L_M \subseteq (\Omega \cup \{\#\})^3 \times (Q \cup \{\#\})^3$ the set of local descriptions of (configurations of) M. If $(\boldsymbol{\omega}, \boldsymbol{q})$ is a local restriction, we will often use the indices $-1, 0, 1$ to access its elements, rather than $1, 2, 3$. Hence, the sensible restrictions apply: # can only occupy the first position of $\boldsymbol{\omega}_t$, and at most one element in $\boldsymbol{q}_t$ can be different from #.

A.1 Proof of Proposition 4.2-(2)

In order to prove the results from Section 5 we need some insights from the proof of Proposition 4.2-(2) given in [64]. This work is in the context of the domino problem. In this problem we are given a finite set of square tiles T of the same size where each side of each tile is colored, together with an initial tile t_0. Our task is to decide whether it is possible to tile the infinite upper-right quadrant of the plane using the tiles from T (each tile can be use infinitely many times) in such a way that t_0 is placed at the origin and any adjoining edges have the same color. Observe this problem can be seen as a restriction of $\mathrm{Hom}(\Gamma, \cdot)$ where we only consider instances $\boldsymbol{T}$ where each element in T is a tuple $t = (t_U, t_D, t_L, t_R)$, corresponding to the edge colors of the tile t, the relation E_1^T consists of the pairs (t, t') such that $t_R = t'_L$, and the relation E_2^T consists of the pairs (t, t') such that $t_U = t'_D$. In [64] this problem is shown to be undecidable, so by extension $\mathrm{Hom}(\Gamma, \cdot)$ is undecidable as well.

The proof in [64] is a reduction from the Halting Problem. The idea is that any (non-deterministic) Turing Machine M can be encoded into a finite set of tiles T_M in such a way that a tiling of the upper-right quadrant of the plane corresponds to a non-halting execution of M starting from an empty tape. The intuition behind this encoding is that the i-th row of a tiling should represent a configuration of M (i.e., tape contents, head position, and machine state) at the i-th time step. This is fairly straightforward, but we sketch a construction that is slightly simpler from the one in [64], since we do not need to consider only "domino" tiles.

Fix a Turing machine $M = (Q, \Omega, \Delta)$. We a Σ_Γ-structure $\boldsymbol{T}_M$ as follows. We allow ourselves to be slightly informal in the description of $\boldsymbol{T}_M$ and leave some small gaps. Elements from T_M

are tuples $t = (\boldsymbol{\omega}_t, \boldsymbol{q}_t, b_t, c_t)$ where $(\boldsymbol{\omega}_t, \boldsymbol{q}_t) \in L_M$ is a local description of M, $b_t \in \{0, 1\}$ is a bit that equals 1 when t describes a local configuration to the right of M's head, and $c_t \in \{0, 1\}$ is a bit that equals 1 when t is describes part of M's initial configuration. Hence, $\boldsymbol{q}_t = (\#, \#, \#)$, when $b_t = 1$. Similarly, when $c_t = 1$, only blank symbols $\square$ and end-of-tape symbols # are allowed in $\boldsymbol{\omega}_t$, and the head of M can only be at the beginning of the tape. The relation O^{T_M} contains only the initial tile $((\#, \square, \square), (\#, q_0, \#), 0, 1)$. Two elements (t, t') belong to $E_1^{T_M}$ if they are two consecutive local descriptions of some configuration of M. More explicitly, the last two elements of $\boldsymbol{\omega}_t$ must equal the first two elements of $\boldsymbol{\omega}_{t'}$, the last two elements of $\boldsymbol{q}_t$ must equal the first two elements of $\boldsymbol{q}_{t'}$, and $c_t = c_{t'}$. If $\boldsymbol{q}_t = (q, \#, \#)$, for some $q \in Q$, then $\boldsymbol{q}_{t'} = (\#, \#, \#)$ and $b_{t'} = 1$. The relation $E_2^{T_M}$ consists of the pairs (t, t') that represent consistent local descriptions of two successive configurations of M. We apply the natural rules: (1) $c_{t'} = 0$, (2) if $\boldsymbol{q}_t = (\#, \#, \#)$ then $\boldsymbol{\omega}_t = \boldsymbol{\omega}_{t'}$, and $\boldsymbol{q}_{t'}$ contains # in the second position, and (3) if $\boldsymbol{q}_t = (\#, q, \#)$ for some $q \in Q$, then t' describes the evolution of t according to some transition in Δ.

Given our definition of $\boldsymbol{T}_M$ and a homomorphism $F : \boldsymbol{\Gamma} \to \boldsymbol{T}_M$, it is easy to see that there is an infinite run of M starting from an empty tape which is given by a sequence of configurations $(\boldsymbol{\omega}_1, \boldsymbol{q}_1), (\boldsymbol{\omega}_2, \boldsymbol{q}_2), \ldots$ such that the element $F(i, j)$ contains the j-th local description of $(\boldsymbol{\omega}_1, \boldsymbol{q}_1)$. Conversely, given an infinite run of M starting from an empty tape it is straight-forward to describe a homomorphism $F : \boldsymbol{\Gamma} \to \boldsymbol{T}_M$. This shows Proposition 4.2-(2).

A.2 Proof of Proposition 4.2-(3)

The same proof from [43] essentially shows our result. This work is also in the context of the domino problem. Their main result states that there are sets of domino tiles that can tile the whole plane starting with a fixed tile at the origin, but satisfying that any such tiling must be non-computable. In our setting, we only need to tile the upper-right quadrant of the plane and we consider more general sets of tiles, other than domino tiles, but the proof from [43] can be easily adapted. The starting point is the result that there exists a Turing machine M that does not halt for some input words, but all those input words are non-computable. This is shown in [43] for Turing machines on a two-way infinite tape, rather than just a semi-infinite tape as in our setting, but it is known that both models are equivalent. This can be seen, for example, by "folding" a two-way infinite tape into a semi-infinite tape as in [44]. Then Proposition 4.2-(3) follows from considering a Σ_Γ-structure $\boldsymbol{T}_M$ derived from M as in Section A.1, modified suitably so that arbitrary input words are allowed in the initial configuration.

A.3 Proof of Proposition 5.3

We build on the ideas from Section A.1 again. The idea is to build a Σ_∇-structure $\boldsymbol{S}_M$ starting from a Turing machine M in such a way that homomorphisms $F : \boldsymbol{\nabla}_n \to \boldsymbol{S}_M$ precisely encode accepting runs of M which start from the empty input, take at most n-steps, and where the head of M is placed left to the $(n - i + 1)$-th position at the i-th step. This way, $\boldsymbol{\nabla}_n \to \boldsymbol{S}_M$ for any $n \in \mathbb{N}$, if and only if $\boldsymbol{\nabla}_n \to \boldsymbol{S}_M$ for all but finitely many $n \in \mathbb{N}$, if and only if $\boldsymbol{\nabla}_n \to \boldsymbol{S}_M$ for infinitely many $n \in \mathbb{N}$, if and only if M has a halting run starting from the empty word. Then Proposition 5.3 follows from the fact that the problem of determining whether an input Turing machine accepts the empty word is undecidable.

To construct $\boldsymbol{S}_M$ we start from the tile set $\boldsymbol{T}_M$ described in Section A.1, and add an additional element $\blacksquare$. The new relation W^{S_M} consists only of this element $\{\blacksquare\}$. We also add the pair $(\blacksquare, \blacksquare)$ to both $E_1^{S_M}, E_2^{S_M}$. Given a tile $t = (\boldsymbol{\omega}_t, \boldsymbol{q}_t, b_t, c_t) \in T_M$, we add $(t, \blacksquare)$ to $E_2^{S_M}$ if $\boldsymbol{q}_t = (\#, \#, \#)$, or if the only state showing in $\boldsymbol{q}_t$ is the accepting state q_+. We also add $(t, \blacksquare)$ to $E_1^{S_M}$ if the last position of $\boldsymbol{q}_t$ contains #. Let us give some intuition for this construction. When searching for a homomorphism $F : \boldsymbol{\nabla}_n \to \boldsymbol{S}_M$, we proceed as in Section A.1, by tiling the plane in a way that describes a run of M on the empty word. Now, at any point we may chose to stop describing this run and start placing $\blacksquare$ tiles instead, with the conditions that (1) $\blacksquare$ tiles must propagate right and up, and (2) $\blacksquare$ tiles cannot replace local descriptions that contain the head of M, as long as M is not in the accepting state. It is routine to check that $\boldsymbol{S}_M$ has the desired properties described in the previous paragraphs.

A.4 Hardness Proofs

Before moving on to showing Proposition 4.2-(1) and Proposition 5.2-(1) we describe some TFNP$_1$-hard and TFNP-hard families. Given a immortal Turing machine $M = (Q, \Omega, \Delta)$, in the problem R_M we are given a pair $(\boldsymbol{x}, \boldsymbol{n})$ as an input, where $\boldsymbol{x}$ is an input word for M, and $\boldsymbol{n} \in \{1\}^*$ is a number in unary representation, and the task is to output a run of length $|\boldsymbol{n}|$ of M on the input $\boldsymbol{x}$. Here a run is given as a sequence $\delta_1, \ldots, \delta_n$ of transitions in Δ. Observe the problem R_M belongs to TFNP. We also consider the "tally" version of the problem R_M. Let M be a (non-deterministic) Turing machine that has an infinite run on the empty input. In the problem R_M^1 we are given a number $\boldsymbol{n} \in \{1\}^*$ in unary representation, and the task is to output a run of M on the empty input of length $|\boldsymbol{n}|$.

Lemma A.1. *Let $\mathcal{F}$ be the the family consisting of the problems R_M for each immortal Turing machine M on the binary alphabet $\{0, 1, \square\}$. Then $\mathcal{F}$ is TFNP-hard.*

Proof. We can ignore the restriction to the alphabet $\{0, 1, \square\}$ by noting that larger alphabets can be suitably encoded in binary by paying some small overhead [1]. Hence, it is enough to prove that the larger family consisting of the problems R_M for each immortal Turing machine M is TFNP-hard.

Let $\Lambda_{\mathfrak{R}}$ be a problem in TFNP, where $\mathfrak{R} \subseteq U^* \times V^*$. Let p be a polynomial, and let N be a polynomial-time Turing machine satisfying that for each $\boldsymbol{x} \in U^*$ there is some $\boldsymbol{y} \in V^*$ such that $(\boldsymbol{x}, \boldsymbol{y}) \in \mathfrak{R}$ and $|\boldsymbol{y}| \leqslant p(|\boldsymbol{x}|)$, and N decides $\mathfrak{R}$. We build a immortal Turing machine M, informally described as follows. The alphabet of M contains both U and V. The machine M loops forever on input words $\boldsymbol{x}$ that are not in U^*. Given an input $\boldsymbol{x} \in U^*$, the machine M guesses a word $\boldsymbol{y} \in V^*$ of length at most $p(|\boldsymbol{x}|)$ and then runs N on $(\boldsymbol{x}, \boldsymbol{y})$. If N rejects this pair, then M also rejects, and if N accepts the pair, then M loops forever. Observe that M must be immortal. Let p' be a polynomial such that, given $\boldsymbol{x} \in U^*$, it takes M at most $p'(|\boldsymbol{x}|)$ time steps to guess $\boldsymbol{y}$, simulate N, and continue its execution for one more step.

Now we claim that there is a polynomial-time many-one reduction from $\Lambda_{\mathfrak{R}}$ to R_M. The first part of this reduction sends the input $\boldsymbol{x} \in U^*$ to $\Lambda_{\mathfrak{R}}$ to the input $(\boldsymbol{x}, \boldsymbol{n})$ to R_M, where $\boldsymbol{n}$ is a unary representation of $p'(|\boldsymbol{x}|)$. Observe that a valid answer to $(\boldsymbol{x}, \boldsymbol{n})$ in R_M is a run of M on the input $\boldsymbol{x}$ that lasts for $p'(|\boldsymbol{x}|)$ steps. During such a run, M must guess correctly a word $\boldsymbol{y} \in V^*$ satisfying $(\boldsymbol{x}, \boldsymbol{y}) \in \mathfrak{R}$. Indeed, otherwise, N would reject, leading to M halting prematurely. The second part of the reduction extracts the word $\boldsymbol{y}$ from the description of the run. This completes the proof. $\square$

Lemma A.2. *Let $\mathcal{F}$ be the the family consisting of the problems R_M^1 for each Turing machine M on the binary alphabet $\{0, 1, \square\}$ that has an infinite run on the empty input. Then $\mathcal{F}$ is TFNP$_1$-hard.*

Proof. We ignore the restriction on the alphabet, as in the previous lemma. Let $\Lambda_{\mathfrak{R}}$ be a problem in TFNP_1, where $\mathfrak{R} \subseteq \{1\}^* \times V^*$. Let p be a polynomial satisfying that for any $n \in \mathbb{N}$ there is some $\boldsymbol{y} \in V^*$ satisfying $(1^n, \boldsymbol{y}) \in \mathfrak{R}$ and $|\boldsymbol{y}| \leqslant p(n)$, and let N be polynomial-time deterministic Turing machine deciding $\mathfrak{R}$.

We construct a Turing machine M that has an infinite run on the empty input which satisfies that $\Lambda_{\mathfrak{R}}$ has a many-one polynomial-time reduction to R^1_M. The machine M is informally described as follows. On the empty input, M keeps track of an integer counter k, whose value starts at one. Then, M guesses a word $\boldsymbol{y} \in V^*$ with $|\boldsymbol{y}| \leqslant p(k)$, and then runs N on the input $(1^k, \boldsymbol{y})$. If N rejects, then M rejects and halts. Otherwise, M increases the value of k by one and repeats the process again. Let q be a polynomial such that it takes M at most $p'(m)$ steps to set its counter k to the value $m+1$.

We describe a suitable reduction from $\Lambda_{\mathfrak{R}}$ to R^1_M. The first part of the reduction takes an input 1^n to $\Lambda_{\mathfrak{R}}$ and constructs the input $1^{p'(n)}$ to R^1_M. A valid response to $1^{p'(n)}$ in R^1_M is a run of M on the empty input that lasts for $p'(n)$ steps. During such a run M must guess correctly a word $\boldsymbol{y} \in V^*$ satisfying $(1^n, \boldsymbol{y}) \in \mathfrak{R}$ (otherwise M would halt prematurely). The second part of the reduction extracts the word $\boldsymbol{y}$ from the description of this run. □

A.4.1 Proof of Proposition 4.2-(1)

Let M be a Turing machine that has an infinite run on the empty word, and let $\boldsymbol{T}_M$ be the Σ_Γ-structure constructed in Section A.1. We show that the problem R^1_M has a polynomial-time many-one reduction to $\mathrm{sPCSP}(\boldsymbol{\Gamma}, \boldsymbol{T}_M)$. Observe that this reduction together with Lemma A.2 prove the result.

The first part of the reduction takes an input 1^n to R^1_M and produces an input $\boldsymbol{\Gamma}_n$ to $\mathrm{sPCSP}(\boldsymbol{\Gamma}, \boldsymbol{T}_M)$. The structure $\boldsymbol{\Gamma}_n$ is the substructure of $\boldsymbol{\Gamma}$ induced on $[n] \times [n]$, i.e., the $n \times n$ grid. A valid output to $\boldsymbol{\Gamma}_n$ in $\mathrm{sPCSP}(\boldsymbol{\Gamma}, \boldsymbol{T}_M)$ is a homomorphism $F : \boldsymbol{\Gamma}_n \to \boldsymbol{T}_M$. The horizontal lines $F(i,j)$ for each $i \in [n]$ describe successive configurations (truncated to the first n tape spaces) of M starting from the empty input. Therefore the homomorphism F contains the description of a n-step run of M on the empty input. The second part of the reduction simply extracts this run from F.

A.4.2 Proof of Proposition 5.2-(1)

This result can be shown similarly to Proposition 4.2-(1), with some work. Let M be an immortal Turing machine. We begin by modifying the construction from Section A.1 to obtain a Σ_{Γ^+}-structure $\boldsymbol{T}_M$ satisfying that any homomorphism $F : \boldsymbol{\Gamma}^+ \to \boldsymbol{T}_M$ admits the following description. We view the super-grid $\boldsymbol{\Gamma}^+$ as the positive quadrant of the 3-dimensional grid, where we call the dimensions vertical, horizontal, and normal respectively. Each normal slice (i.e., each set consisting all tuples (i,j,k) where k is fixed) is called a *floor*. Then the tiling $F(i,j,k)$ of the k-th floor describes an infinite run of M on the input given by the binary representation $\lfloor i \rfloor$ of i. We adopt the convention that binary representations start from the least-significant bit. I.e., $\lfloor 6 \rfloor = 011$.

We construct $\boldsymbol{T}_M$ as follows. Let $M = (Q, \Omega, \Delta)$ be a Turing machine on the binary alphabet $\Omega = \{0, 1, \square\}$. Similarly to the Σ_Γ-structure constructed in Section A.1, the elements in T_M are tuples $t = (\boldsymbol{\omega}_t, \boldsymbol{q}_t, b_t, c_t, d_t, e_t) \in \Omega^3 \times Q^3 \times \{0,1\}^4$, where $(\boldsymbol{\omega}_t, \boldsymbol{q}_t) \in L_M$ is a local description of M. The bit b_t keeps track of whether the tile t describes a local configuration to the right of M's head, bit c_t keeps track of whether the tile t is in the initial horizontal line of some floor (e.g., is in a position

(i,j,k) with $j = 1$), and the bit d_t keeps track of whether t is in the first floor (e.g., its position is (i,j,k) with $k = 1$). Finally, the bit e_t is a *carry-over* bit, used during addition operations described below. The origin relation O^{T_M} is the singleton containing the tile

$$((\#,1,\square),(\#,q_0,\#),0,1,1,0).$$

The vertical and horizontal relations $E_1^{T_M}$, $E_2^{T_M}$ are described analogously to the ones in Section A.1. The main difference is that we force the tape described by the first horizontal line in the first floor to be $(\#,1,\square,\square,\dots)$ (i.e., the input 1), rather than the blank tape. We also need to describe how $E_1^{T_M}$ and $E_2^{T_M}$ interact with the carry-over bit e_t. The vertical relation $E_2^{T_M}$ "forgets" about the carryover, meaning that if $(t,t') \in E_2^{T_M}$, then $e_{t'} = 0$. However, the horizontal relation $E_1^{T_M}$ propagates the carryover. This means that if $(t,t') \in E_2^{T_M}$, and $e_t = 0$, then $e_{t'} = 0$. However if $e_t = 1$ then we have two options. If $e_{t'} = 0$ then the rightmost element of $\boldsymbol{\omega}_{t'}$ must be 1, or, otherwise, if $e_{t'} = 1$ then the rightmost element of $\omega_{t'}$ is 0. The doubling relations $\mathbb{E}_1^{T_M}$, $\mathbb{E}_2^{T_M}$ do not impose any constraint. The normal relations $E_3^{T_M}$, $\mathbb{E}_3^{T_M}$ only constrain the tiles that lie in the first horizontal line of each floor, i.e., those $t \in T_M$ with $c_t = 1$. The relation $\mathbb{E}_3^{T_M}$ ensures that if the tape contains a binary representation $\lfloor n \rfloor$ in the horizontal line $(i,1,1)$, then the horizontal line at $(2i,1,1)$ displays $\lfloor 2n \rfloor$. This is achieved by by shifting right $\boldsymbol{\omega}_t$ in each local description $(\boldsymbol{\omega}_t, \boldsymbol{q}_t)$ and by adding a leading 0. In a similar way, the relation $E_3^{T_M}$ ensures that if the horizontal line $(i,1,1)$ displays $\lfloor n \rfloor$, then $(i+1,1,1)$ displays $\lfloor n+1 \rfloor$. This is requires addition to be performed on $\boldsymbol{\omega}_t$ from left to right in each local description $(\boldsymbol{\omega}_t, \boldsymbol{q}_t)$ following the standard *column addition* algorithm, and using the carry-over bit e_t as needed. This way, we obtain a structure $\boldsymbol{T}_M$ satisfying the high-level description given at the start of the section.

In order to complete the proof of Proposition 5.2-(1), we describe a reduction from R_M to sPCSP($\boldsymbol{\Gamma}^+, \boldsymbol{T}_M$). Let $(\lfloor n \rfloor, 1^m)$ be an input to R_M, where $\lfloor n \rfloor$ corresponds to the binary representation of $n \in \mathbb{N}$, and $m \in \mathbb{N}$. Without loss of generality we may assume that $m \geqslant n$: if $m < n$ we can act as if $m = n$. The first map of our reduction constructs in polynomial time a structure $\boldsymbol{I}_{n,m}$ satisfying $\boldsymbol{I}_{n,m} \to \boldsymbol{\Gamma}^+$. First, observe that there is a sequence $\spadesuit_1, \dots, \spadesuit_\ell$ of length $\ell \in O(\log n)$ consisting of the operations $+1$ and $\times 2$ that yields n starting from 1. Such a sequence $\boldsymbol{L}$ can be obtained in polynomial time from $\lfloor n \rfloor$. For each $i \in [\ell+1]$ we define the number n_i inductively as follows. We let $n_1 = 1$, and $n_i = n_{i-1}\spadesuit_{i-1}$ for $i > 1$, where we recall that $\spadesuit_{i-1} \in \{+1, \times 2\}$. This way, $n_{\ell+1} = n$. The universe $I_{n,m}$ is defined as

$$\{(i,j,n) \mid i,j \in [m]\} \bigcup \{(i,1,n_k) \mid i \in [m], k \in [\ell+1]\}.$$

In other words, this includes the initial $m \times m$ quadrant in the n-th floor of $\mathbb{N}^3$, plus several horizontal segments of length m that "climb" up to the n-th floor. The relations are defined so that $\boldsymbol{I}_{n,m}$ is the substructure of $\boldsymbol{\Gamma}^+$ induced on $I_{n,m}$. By construction $\boldsymbol{I}_{n,m} \to \boldsymbol{\Gamma}^+$. Now, for the second map of our reduction we need a polynomial-time procedure that computes a m-step run of M starting from the input $\lfloor n \rfloor$ by accessing a homomorphism $F : \boldsymbol{I}_{n,m} \to \boldsymbol{T}_M$. This is done in the intuitive way: By construction, image of the horizontal line $(i,1,n)$ through F must describe the initial configuration of M on the input $\lfloor n \rfloor$, and the image through F of the quadrant $\{(i,j,n) \mid i,j \in [m]\}$ must describe a m-step run of M starting from this input. Hence, by accessing F one can compute the desired run in polynomial time. This completes the proof.

B Characterizing Polymorphism Minions up to Isomorphism

We extend the ideas of [18, Section 6.2], where function minions corresponding to polymorphisms are characterized. The main result there is that a function minion is a polymorphism minion if and only if it has bounded *finitized arity*, as we described in Section 7. The following characterizes abstract minions that are isomorphic to the polymorphism minion of some finite template.

Theorem B.1. *Let $\mathscr{M}$ be a minion. Then there exists a finite template $(\boldsymbol{A}, \boldsymbol{B})$ such that $\mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$ is isomorphic to $\mathscr{M}$ if and only if $\mathscr{M}$ is locally finite and $\mathscr{M}$ is isomorphic to $\mathscr{M}^{(h)}$ for some $h \in \mathbb{N}$.*

We need the following short auxiliary results.

Lemma B.2. *Let $\mathscr{M}$ and $\mathscr{N}$ be minions, let $h \in \mathbb{N}$, and let $F : \mathscr{M} \xrightarrow{h} \mathscr{N}$ be a partial minion isomorphism defined up to arity h. Then $\mathscr{M}^{(h)}$ is isomorphic to $\mathscr{N}^{(h)}$.*

Proof. We define an isomorphism $H : \mathscr{M}^{(h)} \to \mathscr{N}^{(h)}$. Given a system $\xi \in \mathscr{M}^{(h)}(n)$, its image $H(\xi)$ is the system defined by $H(\xi)(\pi) = F(\xi(\pi))$ for all $\pi \in [h]^{[n]}$. Checking that this is a minion homomorphism is routine, and its inverse is given by $H^{-1}(\rho)(\pi) = F^{-1}(\rho(\pi))$ for all $n \in \mathbb{N}$ all $\rho \in \mathscr{N}^{(h)}(n)$ and all $\pi \in [h]^{[n]}$ ☐

As a corollary we obtain that the h-closure of a h-closure is isomorphic to the original h-closure.

Corollary B.3. *Let $\mathscr{M}$ be a minion, let $h \in \mathbb{N}$ and let $\mathscr{N} = \mathscr{M}^{(h)}$. Then, $\mathscr{N}$ is isomorphic to $\mathscr{N}^{(h)}$.*

Proof. By Proposition 7.4, $\mathscr{M}$ and $\mathscr{M}^{(h)} = \mathscr{N}$ are isomorphic up to arity h, so the statement follows from last lemma. ☐

Now we proceed with our main proof.

Proof of Theorem B.1. We show both directions. Let $(\boldsymbol{A}, \boldsymbol{B})$ be an arbitrary finite template. Let $r = |A|$, and let $h \in \mathbb{N}$ be at least as large as r and as large as the number of tuples in any relation of $\boldsymbol{A}$. We identify A with $[r]$. A consequence of the definition of polymorphism is that a function of the form $f : A^n \to B$ belongs to $\mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$ if and only if $f^\pi \in \mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$ for every function $\pi \in [h]^{[n]}$. Define $\mathscr{M} = \mathrm{Pol}(\boldsymbol{A}, \boldsymbol{B})$. Clearly, $\mathscr{M}$ is locally finite. We show that $\mathscr{M}$ is isomorphic to $\mathscr{M}^{(h)}$. The isomorphism is the canonical map $\mathrm{Cl}_h : \mathscr{M} \to \mathscr{M}^{(h)}$. We already know that Cl is a minion homomorphism, so all that is left is to show that it is bijective. An argument that is convenient to remember is that, under the usual identification of tuples from $[r]^n$ with maps in $[r]^{[n]}$, the evaluation of a function $f : [r]^n \to B$ on a tuple $\boldsymbol{a} \in [r]^n$ can be seen as the evaluation $f^{\boldsymbol{a}}(\mathrm{id}_{[r]})$.

Cl_h *is injective*. Let $n \in \mathbb{N}$, and let $f, g \in \mathscr{M}(n)$. Suppose that $f \neq g$. Then by the remark above, it holds that $f^\pi \neq g^\pi$ for some $\pi \in [r]^{[n]}$. Composing π with the inclusion $[r] \hookrightarrow [n]$ we obtain $\pi' \in [h]^{[n]}$ such that $f^{\pi'} \neq g^{\pi'}$. Thus, by the definition of the canonical map, it holds that

$$\mathrm{Cl}_h(f)(\pi') = f^{\pi'} \neq g^{\pi'} = \mathrm{Cl}_h(g)(\pi'),$$

meaning that $\mathrm{Cl}_h(f) \neq \mathrm{Cl}_h(h)$.

Cl_h *is surjective*. Let $n \in \mathbb{N}$ and let $\xi \in \mathscr{M}^{(h)}(n)$ be a system. The idea is that ξ is given by a family of consistent h-ary polymorphisms in $\mathscr{M}$, so there is a n-ary function $f_\xi : [r]^n \to B$ whose h-ary minors are given by $f_\xi^\pi = \xi(\pi)$. Constructing such f_ξ would show that Cl_h is surjective: Indeed, under this assumption all h-ary minors of f_ξ belong to $\mathscr{M}$ forcing $f_\xi \in \mathscr{M}$ as well. Finally, by definition, $\mathrm{Cl}_h(f_\xi) = \xi$, so this shows that Cl_h is surjective. We define the desired function $f_\xi : [r]^n \to B$ as follows. Fix two maps $\alpha : [r] \to [h]$ and $\beta : [h] \to [r]$ such that $\beta \circ \alpha = \mathrm{id}_{[r]}$. Given a tuple $\boldsymbol{a} \in [r]^n$, we define

$$f_\xi(\boldsymbol{a}) = \xi(\alpha \circ \boldsymbol{a})(\beta).$$

To parse this equality, recall that $\xi(\alpha \circ \boldsymbol{a})$ is a h-ary polymorphism in $\mathscr{M}$, which can be evaluated at a tuple in $[r]^h$. All that is left to show is that $f_\xi^\pi = \xi(\pi)$ for all $\pi \in [h]^{[n]}$. Indeed, in this case, given a tuple $\boldsymbol{a} \in [r]^{[h]}$ it holds that

$$f_\xi^\pi(\boldsymbol{a}) = f_\xi(\boldsymbol{a} \circ \pi) = \xi(\alpha \circ \boldsymbol{a} \circ \pi)(\beta) = (\xi(\pi))^{\alpha \circ \boldsymbol{a}}(\beta) = \xi(\pi)(\beta \circ \alpha \circ \boldsymbol{a}) = \xi(\pi)(\boldsymbol{a}).$$

This proves that Cl_h is a surjective homomorphism, and together with the fact that it is also injective, completes the proof that it is an isomorphism.

We have shown that if $\mathscr{M}$ is the polymorphism minion of some finite template, then it is locally finite and is isomorphic to $\mathscr{M}^{(h)}$ for some $h \in \mathbb{N}$. Now we prove the other direction. Let $\mathscr{M}$ be a locally finite minion, and suppose that it is isomorphic to $\mathscr{N} = \mathscr{M}^{(h)}$ for some $h \in \mathbb{N}$. An observation is that the rank of $\mathscr{N}$ is at most h. Indeed, let $n \in \mathbb{N}$ and let $\xi_1, \xi_2 \in \mathscr{N}(n)$ be two systems of h-ary minors. Suppose that $\xi_1^\pi = \xi_2^\pi$ for all $\pi \in [h]^{[n]}$. We also have that $\xi_1(\pi) = \xi_1^\pi(\mathrm{id}_{[h]})$ and analogously with ξ_2. This means that $\xi_1(\pi) = \xi_2(\pi)$ for all $\pi \in [h]^{[n]}$ and $\xi_1 = \xi_2$. Then, by Corollary 7.6, it holds that $\mathscr{N}^{(h)}$ is isomorphic to $\mathrm{Pol}(\boldsymbol{K}_h^h, \boldsymbol{F}_{\mathscr{N}}(\boldsymbol{K}_h^h))$. Also, $\mathscr{N}^{(h)}$ is isomorphic to $\mathscr{N} = \mathscr{M}^{(h)}$ by Corollary B.3, and $\mathscr{M}^{(h)}$ is, by assumption, isomorphic to $\mathscr{M}$. Putting together all these isomorphisms we obtain that $\mathscr{M}$ is isomorphic to $\mathrm{Pol}(\boldsymbol{K}_h^h, \boldsymbol{F}_{\mathscr{N}}(\boldsymbol{K}_h^h))$, completing the proof. □

An alternative, and perhaps better way to present the second part of this proof could have been to show, using essentially the same reasoning as in the proof of Theorem 7.1, that given a minion $\mathscr{M}$ and a number $h \in \mathbb{N}$, it holds in general that $\mathscr{M}^{(h)}$ is isomorphic to $\mathrm{Pol}(\boldsymbol{K}_h^h, \boldsymbol{F}_{\mathscr{M}}(\boldsymbol{K}_h^h))$.