

The Complexity of 3-Colouring H -Colourable Graphs

Andrei Krokhin

Department of Computer Science

Durham University

Durham, UK

Email: andrei.krokhin@durham.ac.uk

Jakub Opršal

Department of Computer Science

Durham University

Durham, UK

Email: jakub.oprsal@durham.ac.uk

Abstract—We study the complexity of approximation on satisfiable instances for graph homomorphism problems. For a fixed graph H , the H -colouring problem is to decide whether a given graph has a homomorphism to H . By a result of Hell and Nešetřil, this problem is NP-hard for any non-bipartite graph H . In the context of promise constraint satisfaction problems, Brakensiek and Guruswami conjectured that this hardness result extends to promise graph homomorphism as follows: fix any non-bipartite graph H and another graph G with a homomorphism from H to G , it is NP-hard to find a homomorphism to G from a given H -colourable graph. Arguably, the two most important special cases of this conjecture are when H is fixed to be the complete graph on 3 vertices (and G is any graph with a triangle) and when G is the complete graph on 3 vertices (and H is any 3-colourable graph). The former case is equivalent to the notoriously difficult approximate graph colouring problem. In this paper, we confirm the Brakensiek-Guruswami conjecture for the latter case. Our proofs rely on a novel combination of the universal-algebraic approach to promise constraint satisfaction, that was recently developed by Barto, Bulín and the authors, with some ideas from algebraic topology.

Index Terms—graph colouring; graph homomorphism problem; constraint satisfaction; polymorphism; promise problem;

I. INTRODUCTION

In this paper we investigate the complexity of finding an approximate solution to satisfiable instances. For example, for the problem of 3-colouring a graph, one natural approximation version is the *approximate graph colouring* problem: The goal is to find a c -colouring of a given 3-colourable graph. There is a huge gap in our understanding of the complexity of this problem. The best known efficient algorithm uses roughly $c = O(n^{0.199})$ colours where n is the number of vertices of the graph [1]. It has been long conjectured the problem is NP-hard for any fixed constant $c \geq 3$, but the state-of-the-art here has only recently been improved from $c = 4$ [2], [3] to $c = 5$ [4], [5].

Graph colouring problems naturally generalise to graph homomorphism problems and further to constraint satisfaction problems (CSPs). In a graph homomorphism problem, one is given two graphs and needs to decide whether there is a homomorphism (edge-preserving map) from the first graph to the second [6]. The CSP is generalisation of this that uses arbitrary relational structures in place of graphs. One particularly important case that attracted much attention is

when the second graph/structure is fixed, this is the so-called non-uniform CSP [7], [8]. For graph homomorphisms, this gives the so-called H -colouring problem: decide whether a given graph has a homomorphism to a fixed graph H [6]. The P vs. NP-complete dichotomy of H -colouring given in [9] was one of the base cases that supported the Feder-Vardi dichotomy conjecture for CSPs [8]. The study of the complexity of the (standard) CSP and the final resolution of the dichotomy conjecture [10], [11] was greatly influenced by the algebraic approach (see survey [7]). This approach has also made important contributions to the study of approximability of CSPs (e.g. [12]).

Brakensiek and Guruswami [13], [14] suggested that perhaps progress on approximate graph colouring and similar open problems can be made by looking at a broader picture, by extending it to *promise graph homomorphism* and further to the *promise constraint satisfaction problem* (PCSP). The promise graph homomorphism is an approximation version of the graph homomorphism problem in the following sense: we fix (not one but) two graphs H and G such that there is a homomorphism from H to G ; the goal is then to find a G -colouring for a given H -colourable graph. The *promise* is that the input graph is always H -colourable. The general promise CSP (PCSP, for short) is a natural generalisation of this to arbitrary relational structures.

Given the huge success of the algebraic approach to the CSP, it is natural to investigate what it can do for PCSPs. This investigation was started by Austrin, Håstad, and Guruswami [15], with an application to a promise version of SAT. It was further developed by Brakensiek and Guruswami [13], [14], [16] and applied to a range of problems, including versions of approximate graph and hypergraph colouring. A recent paper [4], [5] describes a general abstract algebraic theory for PCSPs. However, the algebraic theory of PCSP is still very young and much remains to be done both in further developing it and in applying it to specific problems. We note that the aforementioned NP-hardness of 5-colouring a given 3-colourable graph was proved in [4], [5] by applying this abstract theory.

In the present paper, we apply this general theory to prove NP-hardness for an important class of promise graph homomorphism problems.

A. Related Work

The notion of the PCSP has been coined in [15], but problems from the class has been around for a long time, e.g. the approximate graph colouring [17].

Most notable examples of PCSPs studied before are related to graph and hypergraph colouring. We already mentioned some results concerning colouring 3-colourable graphs with a constant number of colours. By additionally assuming non-standard (perfect-completeness) variants of the Unique Games Conjecture, NP-hardness was shown for all constant $c \geq 3$ [18]. Without additional complexity-theoretic assumptions, the strongest known NP-hardness results for colouring k -colourable graphs are as follows. For any $k \geq 3$, it is NP-hard to colour a given k -colourable graph with $2k - 1$ colours [4], [5]. For large enough k , it is NP-hard to colour a given k -colourable graph with $2^{\Omega(k^{1/3})}$ colours [19]. The only earlier result about promise graph homomorphisms (with $\mathbf{H} \neq \mathbf{G}$) that involves more than graph colouring is the NP-hardness of 3-colouring graphs that admit a homomorphism to \mathbf{C}_5 , the five-element cycle [5], which is the simplest problem within the scope of the main result of this paper.

A colouring of a hypergraph is an assignment of colours to its vertices that leaves no edge monochromatic. It is known that, for any constants $2 \leq k \leq c$, it is NP-hard to find a c -colouring of a given 3-uniform k -colourable hypergraph [20]. Further variants of approximate hypergraph colouring, e.g. relating to strong or rainbow colourings, were studied in [21], [13], [22], [23], but most complexity classifications related to them are still open in full generality. There are also hardness results concerning hypergraph colouring with a super-constant number of colours, e.g. [24], [25].

An accessible exposition of the algebraic approach to the CSP can be found in [7], where many ideas and results leading to (but not including) the resolution [10], [11] of the Feder-Vardi conjecture are presented. The volume [26] contains surveys concerning many aspects of the complexity and approximability of CSPs.

The first link between the algebraic approach and PCSPs was found by Austrin, Håstad, and Guruswami [15], where they studied a promise version of $(2k+1)$ -SAT called $(2+\varepsilon)$ -SAT. It was further developed by Brakensiek and Guruswami [13], [14], [16]. They use a notion of *polymorphism* (which is the central concept in the algebraic theory of CSP) suitable for PCSPs, and show that the complexity of a PCSP is fully determined by its polymorphisms — in the sense that that two PCSPs with the same set of polymorphisms have the same complexity. They also use polymorphisms to prove several hardness and tractability results. The algebraic theory of PCSP was lifted to an abstract level in [4], [5], where it was shown that abstract properties of polymorphisms determine the complexity of PCSP. The main result of this paper heavily relies on [4], [5].

B. Our Contribution

The approximate graph colouring problem is about finding a c -colouring of a given 3-colourable graph. In other words,

it relaxes the goal in 3-colouring. We can instead insist that we want to find a 3-colouring, but strengthen the promise, i.e., fix a 3-colourable graph \mathbf{H} , and ask how hard it is to find a 3-colouring of a given \mathbf{H} -colourable graph. We prove that this problem is NP-hard for any non-bipartite graph \mathbf{H} that is 3-colourable. Note that if \mathbf{H} is bipartite, then this problem is solvable in polynomial time, and therefore our result completes a dichotomy of this special case of the promise graph homomorphism problem.

The scope of our result can be seen as a certain dual of approximate graph colouring in the landscape of promise graph homomorphism, in the following sense. It is not hard to see that, in order to prove that promise graph homomorphism is NP-hard for any pair of non-bipartite graphs (\mathbf{H}, \mathbf{G}) , it enough to prove this for all pairs $(\mathbf{C}_k, \mathbf{K}_n)$, $k \geq 3$ odd and $n \geq 3$, where the first graph is an odd cycle and the second is a complete graph. This is because we have a chain of homomorphisms

$$\dots \rightarrow \mathbf{C}_k \rightarrow \dots \rightarrow \mathbf{C}_5 \rightarrow \mathbf{C}_3 = \mathbf{K}_3 \rightarrow \mathbf{K}_4 \rightarrow \dots \rightarrow \mathbf{K}_n \rightarrow \dots \quad (\text{I.1})$$

and, for each (\mathbf{H}, \mathbf{G}) with a homomorphism $\mathbf{H} \rightarrow \mathbf{G}$, the problem $\text{PCSP}(\mathbf{H}, \mathbf{G})$ admits a (trivial) reduction from $\text{PCSP}(\mathbf{C}_k, \mathbf{K}_n)$ where k is the size of an odd cycle in \mathbf{H} and n is the chromatic number of \mathbf{G} (so we have $\mathbf{C}_k \rightarrow \mathbf{H} \rightarrow \mathbf{G} \rightarrow \mathbf{K}_n$). The chain of homomorphisms (I.1) has a natural middle point \mathbf{K}_3 . From this middle point, the right half of the chain corresponds to approximate graph colouring and the left half is the scope of this paper.

The result of this paper can also be viewed as NP-hardness of colouring $(2 + \varepsilon)$ -colourable graphs with 3 colours. This statement can be made more formal using so-called *circular chromatic number* (see e.g. [27, p. 7]): one can check that indeed a graph maps homomorphically to \mathbf{C}_{2k+1} if and only if its circular chromatic number is at most $2 + 1/k$.

We remark that the promise graph homomorphism problem has not been studied much beyond the case of approximate graph colouring. This somewhat narrow focus of earlier research can probably be explained by the serious difficulties encountered already in this special case. However, broadening the scope is advantageous, as this brings new methods into the picture, which can potentially resolve the difficult special case too.

Our proofs rely on the universal-algebraic approach to promise constraint satisfaction, that was recently developed by Barto, Bulín, and the authors [4], [5], as well as on some ideas from algebraic topology. To the best of our knowledge, this is the first time when ideas from universal algebra and algebraic topology are applied together to analyse the complexity of approximation. We remark that three earlier results on the complexity of approximate hypergraph colouring [21], [25], [20] were based on results from topological combinatorics using the Borsuk-Ulam theorem or similar [28], [29]. Their use of topology seems different from ours, and it remains to be seen whether they are all occurrences of a common pattern.

C. Subsequent Work

In [27], Wrochna and Živný adapted and formalized the topological intuition of the present paper and generalized our result. In particular, they showed that the complexity of \mathbf{G} -colouring $(2 + \varepsilon)$ -colourable graphs depends only on an inherent topology of the graph \mathbf{G} . They showed that this implies NP-hardness of colouring $(2 + \varepsilon)$ -colourable graphs with $4 - \varepsilon$ colours, and obtained a similar result for ‘square-free’ graphs \mathbf{G} . Remarkably, using somewhat related methods, they also improved the state-of-the-art of the standard approximate graph colouring. More precisely, building on the results of Huang [19], they proved NP-hardness of colouring a k -colourable graph with $\binom{k}{\lfloor k/2 \rfloor} - 1$ colours for any $k \geq 4$, which improves both [19] (where the number of colours is smaller, and the result is proved only for large enough k) and the result of [4], [5] for any $k > 5$.

D. Overview of Key Technical Ideas

We prove the hardness via a reduction from Gap Label Cover. The general structure of the proof is similar to [15], where they first give a general sufficient condition for the existence of such a reduction for general PCSPs, and then apply it to specific PCSPs which they call $(2 + \varepsilon)$ -SAT. However, the structure of our problems is rather more complicated — in particular, our problems do not satisfy the sufficient condition from [15], so we need to do substantially more. It was shown in [5] that a reduction in the style of [15] works under a weaker structural assumption (than [15]), and the technical part of this paper shows that this weaker assumption is satisfied for our problems. Let us explain this in more detail.

The general reduction in [15] encodes an instance of Gap Label Cover as an instance of a PCSP instance by using a polymorphism gadget. (Roughly, polymorphisms are multivariate functions compatible with the constraint relations of the PCSP.) That is, solutions of this gadget are polymorphisms, one for each variable of the original label cover instance. In this encoding, the arity of polymorphisms corresponds to the size of label sets in the label cover instance, so an assignment of a label to a variable corresponds to choosing a coordinate in the corresponding polymorphism. The completeness of the reduction follows automatically from the structure of the gadget. The proof of soundness uses the assumption such that any polymorphism of the PCSP at hand essentially depends only on a bounded number of variables (i.e., is a *junta*) — this is the sufficient condition. It is not hard to prove that this is enough to provide a good-enough approximation for a label cover instance. One can assign to each variable of the label cover instance a label chosen uniformly at random from the bounded-size set of labels corresponding to these essential variables of the corresponding polymorphism.

This approach does not work directly for our problems, since we do not have the property that all polymorphisms are juntas. However, we can use a stronger version of the above mentioned result from [15] given in [5]. This stronger version weakens the assumption the all polymorphisms are

juntas — instead, we assume that we can map our polymorphisms to another set of multivariate functions that does have this property, and we can do it in a way that works well with the label cover constraints, so we can use it to identify the (bounded-size set of) important coordinates in our polymorphisms. Formally, such a map is called a minion homomorphism (see Definition II.10). This notion plays a very important role in the algebraic theory of PCSP (see [4], [5]). The construction of this map and the proof that it is a minion homomorphism is the technical content of the paper. Once this is done, our main result follows from [5].

In order to identify which coordinates in polymorphisms are important and which are ‘noise’, we need to analyse the structure of our polymorphisms. In our case, the polymorphisms are simply 3-colourings of direct powers of a fixed odd cycle. Since \mathbf{K}_3 is also an odd cycle, we have that both graphs defining our PCSPs are discretisations of a circle. Our analysis is inspired by ideas from algebraic topology. We assign to each coordinate an integer, called a *degree*. For a unary polymorphism, i.e., a homomorphism from the odd cycle to the 3-cycle, this degree has a precise intuitive meaning: it is the number of times the domain cycle wraps around the range cycle under the homomorphism. This corresponds to the topological degree of a continuous map between two copies of a circle. We further generalise this degree to higher arity polymorphisms — roughly, the degree at a certain coordinate is supposed to count how many times that coordinate wraps around the circle when ignoring all other coordinates. To define this number formally and consistently, we borrow a few notions from algebraic topology: To graphs and graph homomorphisms, we associate Abelian groups and group homomorphisms. These correspond to so-called groups of chains in topology, that are further used to define homology. However, we do not follow this theory that far, and define the degrees directly from these group homomorphisms. Finally, we show that only a bounded number of variables in a polymorphism can have a non-zero degree and use this fact to define our minion homomorphism.

E. Organisation of the Paper

In Section II, we introduce technical notions that we will need in our proof. Section III states our main result and the result from [5] that our proof relies on. In Section IV, we give an overview of the topological intuition of the proof. This section is useful for those who want to get a deeper understanding of the topological intuition — however, it is not required for checking the formal proof of the main result, which is presented in Section V.

II. PRELIMINARIES

A. Promise Graph Homomorphism Problems

The approximate graph colouring problem and promise graph homomorphism problem are special cases of the PCSP, and we use the theory of PCSPs. However, we will not need the general definitions, so we define everything only for graphs. For general definitions, see, e.g. [5]. All graphs in this paper are loopless (i.e. irreflexive).

Definition II.1. A *homomorphism* from a graph $\mathbf{H} = (V(H), E(H))$ to another graph $\mathbf{G} = (V(G), E(G))$ is a map $h: V(H) \rightarrow V(G)$ such that $(h(u), h(v)) \in E(G)$ for every $(u, v) \in E(H)$. In this case we write $h: \mathbf{H} \rightarrow \mathbf{G}$, and simply $\mathbf{H} \rightarrow \mathbf{G}$ to indicate that a homomorphism exists.

We now define formally the promise graph homomorphism problem.

Definition II.2. Fix two graphs \mathbf{H} and \mathbf{G} such that $\mathbf{H} \rightarrow \mathbf{G}$.

- The *search* variant of $\text{PCSP}(\mathbf{H}, \mathbf{G})$ is, given an input graph \mathbf{I} that maps homomorphically to \mathbf{H} , *find* a homomorphism $h: \mathbf{I} \rightarrow \mathbf{G}$.
- The *decision* variant of $\text{PCSP}(\mathbf{G}, \mathbf{H})$ requires, given an input graph \mathbf{I} such that either $\mathbf{I} \rightarrow \mathbf{H}$ or $\mathbf{I} \not\rightarrow \mathbf{G}$, to output YES in the former case, and NO in the latter case.

Note that there is an obvious reduction from the decision variant of each PCSP to the search variant, but it is not known whether the two variants are equivalent for each PCSP. The hardness results in this paper hold for the decision (and hence also for the search) version of $\text{PCSP}(\mathbf{H}, \mathbf{G})$.

It is obvious that if at least one of \mathbf{H}, \mathbf{G} is bipartite then the problem can be solved in polynomial time by using an algorithm for 2-colouring.

Conjecture II.3 ([14]). Let \mathbf{H} and \mathbf{G} be any non-bipartite graphs with $\mathbf{H} \rightarrow \mathbf{G}$. Then $\text{PCSP}(\mathbf{H}, \mathbf{G})$ is NP-hard.

The graphs that we will be working with in this paper are cycles and their direct powers. As usual, we denote by \mathbf{K}_k the complete graph on k vertices, and by \mathbf{C}_k the k -cycle. We will assume throughout that the set of vertices of both graphs is $\{0, 1, \dots, k-1\}$ and that the edges of the k -cycle are $(0, 1), \dots, (k-2, k-1), (k-1, 0)$.

Definition II.4. The n -th *direct (or tensor) power* of a graph \mathbf{G} is the graph \mathbf{G}^n whose vertices are all n -tuples of vertices of \mathbf{G} (i.e., $V(\mathbf{G}^n) = V(\mathbf{G})^n$), and whose edges are defined as follows: $((u_1, \dots, u_n), (v_1, \dots, v_n))$ is an edge of \mathbf{G}^n if and only if (u_i, v_i) is an edge of \mathbf{G} for all $i \in \{1, \dots, n\}$.

B. Polymorphisms

Although this paper does not use the general PCSPs, we will use the tools developed for analysis of these kind of problems. Namely, we use the notions of polymorphisms [15], [14], minions and minion homomorphisms [4], [5]. We introduce these notions in the special case of graphs below. The general definitions and more insights can be found in [5], [7].

Definition II.5. An n -ary *polymorphism* from a graph \mathbf{G} to a graph \mathbf{H} is a homomorphism from \mathbf{G}^n to \mathbf{H} . To spell this out, it is a mapping $f: V(\mathbf{G})^n \rightarrow V(\mathbf{H})$ such that, for all tuples $(u_1, v_1), \dots, (u_n, v_n)$ of edges of \mathbf{G} , we have

$$(f(u_1, \dots, u_n), f(v_1, \dots, v_n)) \in E(\mathbf{H}).$$

We denote the set of all polymorphisms from \mathbf{G} to \mathbf{H} by $\text{Pol}(\mathbf{G}, \mathbf{H})$.

Example II.6. The n -ary polymorphisms from a graph \mathbf{G} to the k -clique \mathbf{K}_k are the k -colourings of \mathbf{G}^n .

An important notion in our analysis of polymorphisms is that of an essential coordinate.

Definition II.7. A coordinate i of a function $f: A^n \rightarrow B$ is called *essential* if there exist a_1, \dots, a_n and b_i in A such that

$$f(a_1, \dots, a_{i-1}, a_i, a_{i+1}, \dots, a_n) \neq f(a_1, \dots, a_{i-1}, b_i, a_{i+1}, \dots, a_n).$$

A coordinate of f that is not essential is called *inessential* or *dummy*.

The set of polymorphisms between any two graphs is closed under the operation of taking a minor, that is, it is a minion. Let us formally define these notions.

Definition II.8. An n -ary function $f: A^n \rightarrow B$ is called a *minor* of an m -ary function $g: A^m \rightarrow B$ given by a map $\pi: \{1, \dots, m\} \rightarrow \{1, \dots, n\}$ if

$$f(x_1, \dots, x_n) = g(x_{\pi(1)}, \dots, x_{\pi(m)})$$

for all $x_1, \dots, x_n \in A$.

Alternatively, one can say that f is a minor of g if it is obtained from g by identifying variables, permuting variables, and introducing inessential variables.

Definition II.9. Let $\mathcal{O}(A, B) = \{f: A^n \rightarrow B \mid n \geq 1\}$. A (*function*) *minion* \mathcal{M} on a pair of sets (A, B) is a non-empty subset of $\mathcal{O}(A, B)$ that is closed under taking minors. For fixed $n \geq 1$, let $\mathcal{M}^{(n)}$ denote the set of n -ary functions from \mathcal{M} .

Definition II.10. Let \mathcal{M} and \mathcal{N} be two minions (not necessarily on the same pairs of sets). A mapping $\xi: \mathcal{M} \rightarrow \mathcal{N}$ is called a *minion homomorphism* if

- 1) it preserves arities, i.e., maps n -ary functions to n -ary functions for all n , and
- 2) it preserves taking minors, i.e., for each

$$\pi: \{1, \dots, m\} \rightarrow \{1, \dots, n\}$$

and each $g \in \mathcal{M}^{(m)}$ we have

$$\xi(g)(x_{\pi(1)}, \dots, x_{\pi(m)}) = \xi(g(x_{\pi(1)}, \dots, x_{\pi(m)})).$$

We refer to [5, Example 2.22] for an example of a minion homomorphism.

Definition II.11. A minion \mathcal{M} is said to have *essential arity at most k* , if each function $f \in \mathcal{M}$ has at most k essential variables. It is said to have *bounded essential arity* if it has essential arity at most k for some k .

Remark II.12. It is well known (see, e.g. [30]), and not hard to check, that the minion $\text{Pol}(\mathbf{K}_3, \mathbf{K}_3)$ has essential arity at most 1. However, it is easy to show that, for any odd $k > 3$, the minion $\text{Pol}(\mathbf{C}_k, \mathbf{K}_3)$ does not have bounded essential arity. Fix a homomorphism $h: \mathbf{C}_k \rightarrow \mathbf{K}_3$ such that $h(0) = h(2) =$

0 and $h(1) = 1$ and define the following function from \mathbf{C}_k^n to \mathbf{K}_3 :

$$f(x_1, \dots, x_n) = \begin{cases} 2 & \text{if } x_1 = \dots = x_n = 1, \\ h(x_1) & \text{otherwise.} \end{cases}$$

It is easy to check that $f \in \text{Pol}(\mathbf{C}_k, \mathbf{K}_3)$. By using Definition II.7 with $a_1 = \dots = a_n = 1$ and $b_i = 0$, one can verify that every coordinate i of f is essential.

Our proof will rely on the following theorem which is a special case of a result in [5] that generalised [15, Theorem 4.7]. We remark that the proof of this theorem is by a reduction from Gap Label Cover, which is a common source of inapproximability results.

Theorem II.13 ([5, Proposition 5.15]). *Let \mathbf{H}, \mathbf{G} be graphs such that $\mathbf{H} \rightarrow \mathbf{G}$. Assume that there exists a minion homomorphism $\xi: \text{Pol}(\mathbf{H}, \mathbf{G}) \rightarrow \mathcal{M}$ for some minion \mathcal{M} on a pair of (possibly infinite) sets such that \mathcal{M} has bounded essential arity and does not contain a constant function (i.e., a function without essential variables). Then $\text{PCSP}(\mathbf{H}, \mathbf{G})$ is NP-hard.*

C. Graph Homology

In this section we introduce a simple way to associate Abelian groups and group homomorphisms to graphs and graph homomorphisms. We will use this connection to find a minion homomorphism needed to apply Theorem II.13 to $\mathbf{H} = \mathbf{C}_k$, $k \geq 3$ odd, and $\mathbf{G} = \mathbf{K}_3$. What we describe here is a special case of standard notions in algebraic topology [31], but we do not assume any topology background.

For an edge (u, v) in a graph \mathbf{G} , let $[u, v]$ denote an orientation of the edge from u to v .

Definition II.14. Fix a graph $\mathbf{G} = (V(G), E(G))$. Let $\Delta_V(\mathbf{G})$ denote the free Abelian group with generators $[v]$, $v \in V(G)$. That is, the elements of this group are formal sums $\sum_{v \in V(G)} c_v [v]$, where $c_v \in \mathbb{Z}$ for all $v \in V(G)$, and the addition in this group is naturally defined as

$$\sum_{v \in V(G)} c_v [v] + \sum_{v \in V(G)} c'_v [v] = \sum_{v \in V(G)} (c_v + c'_v) [v].$$

Similarly, let $\Delta_E(\mathbf{G})$ denote the free Abelian group with generators $[u, v]$, $(u, v) \in E(G)$, where we additionally postulate that $[u, v] = -[v, u]$ for every edge. The elements of $\Delta_V(\mathbf{G})$ are called *vertex chains* and the elements of $\Delta_E(\mathbf{G})$ *edge chains* in \mathbf{G} .

Note that any multiset W of oriented edges in \mathbf{G} gives rise to the edge chain $\sum_{[u, v] \in W} [u, v]$, where each oriented edge appears in the sum with the corresponding multiplicity. With a slight abuse of notation, we will denote this edge chain also by W . For example, if W is a walk that uses some edge (u, v) the same number of times in each direction, then the corresponding coefficient in the edge chain of W will be 0.

Note also that one can consider both $\Delta_V(\mathbf{G})$ and $\Delta_E(\mathbf{G})$ not only as Abelian groups, but also as \mathbb{Z} -modules. That is, for any integer c and any vertex chain or edge chain W , one can

consider the chain $c \cdot W$ defined by multiplying all coefficients in W by c .

For any two graphs \mathbf{H} and \mathbf{G} , any homomorphism $f: \mathbf{H} \rightarrow \mathbf{G}$ naturally gives rise to group homomorphisms $f_V: \Delta_V(\mathbf{H}) \rightarrow \Delta_V(\mathbf{G})$ and $f_E: \Delta_E(\mathbf{H}) \rightarrow \Delta_E(\mathbf{G})$ defined by

$$f_V(\sum_i c_i [v_i]) = \sum_i c_i [f(v_i)],$$

and

$$f_E(\sum_i c_i [u_i, v_i]) = \sum_i c_i [f(u_i), f(v_i)].$$

Since f is a graph homomorphism, $[f(u_i), f(v_i)]$ is always an (orientation of an) edge in \mathbf{G} .

Definition II.15. For a graph \mathbf{G} , we define a map $\partial: \Delta_E(\mathbf{G}) \rightarrow \Delta_V(\mathbf{G})$ as the group homomorphism such that $[u, v] \mapsto [v] - [u]$, for every $[u, v] \in \Delta_E(\mathbf{G})$

Note that the above condition uniquely defines ∂ .

The map ∂ computes the ‘boundary’ of an edge chain. For example, the boundary of an edge chain corresponding to a walk from u to v in \mathbf{G} is $[v] - [u]$, and more generally, the boundary ∂W of an edge chain W counts for each vertex v the difference between how many times edges in W arrive to v and how many times they leave.

We will also use the following observation which is a generalization of the fact that mapping a walk from u to v by a homomorphism f results in a walk from $f(u)$ to $f(v)$.

Lemma II.16. *For each graph homomorphism $f: \mathbf{H} \rightarrow \mathbf{G}$ and each $P \in \Delta_E(\mathbf{H})$, we have $f_V(\partial P) = \partial f_E(P)$.*

Proof. Since all the involved maps are group homomorphisms, it is enough to check the required equality on the generators of $\Delta_E(\mathbf{H})$. Pick $[u, v]$ an oriented edge of \mathbf{H} , then

$$\begin{aligned} f_V(\partial[u, v]) &= f_V([v] - [u]) = [f(v)] - [f(u)] = \\ &= \partial[f(u), f(v)] = \partial f_E([u, v]) \end{aligned}$$

as required. \square

III. THE MAIN RESULT

Our main result is as follows.

Theorem III.1. *Let \mathbf{H} be a 3-colourable non-bipartite graph. Then $\text{PCSP}(\mathbf{H}, \mathbf{K}_3)$ is NP-hard.*

As we explained in the introduction, it is enough to prove this theorem for the case $\mathbf{H} = \mathbf{C}_k$, $k \geq 3$ odd. We do this by using Theorem II.13 for $\text{Pol}(\mathbf{C}_k, \mathbf{K}_3)$ and the minion \mathcal{M} defined as follows.

Definition III.2. Let N be an odd number, we define a minion $\mathcal{L}_{\leq N}$ to be the set of all functions $f: \mathbb{Z}^n \rightarrow \mathbb{Z}$ such that $f(x_1, \dots, x_n) = c_1 x_1 + \dots + c_n x_n$ for some $c_1, \dots, c_n \in \mathbb{Z}$ with $\sum_{i=1}^n |c_i| \leq N$ and $\sum_{i=1}^n c_i$ odd.

Alternatively, the set $\mathcal{L}_{\leq N}$ can be described as the set of all minors of the functions of the form $(x_1, \dots, x_N) \mapsto \pm x_1 \pm \dots \pm x_N$. It is clear that, for any fixed odd N , $\mathcal{L}_{\leq N}$

is a minion that has bounded essential arity and contains no constant function.

Theorem III.3. *Let $k \geq 3$ be odd and let N be the largest odd number such that $N \leq k/3$. Then there is a minion homomorphism from $\text{Pol}(\mathbf{C}_k, \mathbf{K}_3)$ to $\mathcal{L}_{\leq N}$.*

Remark III.4. If k is the size of an odd cycle in \mathbf{H} , there also exists a minion homomorphism $\xi: \text{Pol}(\mathbf{H}, \mathbf{K}_3) \rightarrow \text{Pol}(\mathbf{C}_k, \mathbf{K}_3)$, which can be composed with the minion homomorphism from Theorem III.3 to give a minion homomorphism from $\text{Pol}(\mathbf{H}, \mathbf{K}_3)$ to $\mathcal{L}_{\leq N}$. Given a graph homomorphism $h: \mathbf{C}_k \rightarrow \mathbf{H}$, we can define a map $\xi: \text{Pol}(\mathbf{H}, \mathbf{K}_3) \rightarrow \text{Pol}(\mathbf{C}_k, \mathbf{K}_3)$ by

$$\xi(f)(x_1, \dots, x_n) = f(h(x_1), \dots, h(x_n)).$$

It is easy to show that this map preserves minors and is therefore a minion homomorphism.

The bound on N given in the above theorem is tight. More precisely, one can show that there is also a minion homomorphism in the opposite direction, i.e., from $\mathcal{L}_{\leq N}$ to $\text{Pol}(\mathbf{C}_k, \mathbf{K}_3)$ (see the appendix). It is not hard to check that this in particular implies that [5, Corollary 5.19] cannot be used to provide NP-hardness of $\text{PCSP}(\mathbf{C}_k, \mathbf{K}_3)$ for any $k \geq 9$.

As mentioned above, Theorem III.1 follows immediately from Theorem II.13 and Theorem III.3.

IV. A TOPOLOGICAL DETOUR

The proof presented in Section V is heavily influenced by several topological observations, and even though they are not formally needed, we present them here to provide some intuition. The only intention of this section is to give an intuition about the combinatorial statements in the Section V, therefore we will omit any formal proofs or statements. We believe that an interested reader with an access to a book about algebraic topology (e.g. [31]) will be able to check correctness of our statements. Throughout this section, we add a few remarks intended for readers skilled with algebraic topology.

The analogy between our discrete setting and topology is based on the observation that both \mathbf{C}_k for $k \geq 3$ and \mathbf{C}_3 look from the topological perspective like the circle $S^1 = \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 = 1\}$. Any continuous mapping $f: S^1 \rightarrow S^1$ is assigned a topological invariant called *degree* of f , and denoted by $\deg f$. Intuitively, this number counts ‘how many times f loops around the circle’. A positive degree means it loops around counter-clockwise, a negative one means it loops around clockwise. A similar invariant can be used for graph homomorphisms between two cycles (see Definition V.2). The essence of our argument is to generalize this degree to polymorphisms, i.e., mappings that have multiple values on the input.

Remark IV.1. In algebraic topology, the degree is formally defined through the *fundamental group*. The fundamental group $\Pi_1(S^1)$ is isomorphic to the free cyclic group \mathbb{Z} , the generator of this group is the class of a loop that loops around once counter-clockwise. Any continuous mapping $f: S^1 \rightarrow S^1$




Fig. 1. A torus with two representatives of coordinate loops. The dashed line represents points with coordinates (x, x) for some x .

induces a group homomorphism between the fundamental groups, i.e., a group homomorphism $\Pi_1(f): \mathbb{Z} \rightarrow \mathbb{Z}$, and any such mapping is of the form $f(x) = cx$. This c is then defined as the degree of f .

Let us borrow the term ‘polymorphism’ to use for continuous mappings from a power of a topological space to another, i.e., a polymorphism of our circle S^1 is a continuous map from n -th power of a circle to S^1 (with the product topology). The n -th power of S^1 is an n -torus, usually denoted by T^n . The second power is the usual torus T^2 (surface of a doughnut) depicted on Figure 1. That is for n -ary polymorphisms, we are interested in continuous maps $f: T^n \rightarrow S^1$.

Such a mapping $f: T^n \rightarrow S^1$ is assigned n different degrees $\deg_1 f, \dots, \deg_n f$ each corresponding to one of the coordinates of f . A *degree of f at a coordinate i* is obtained by fixing all other coordinates to a point, and following the i -th coordinate around S^1 and counting how many times one loops around the circle in the image. For example, for $n = 2$, each of the two degrees are obtained by following one of the two loops depicted in Figure 1. A necessary observation is that degree assigned this way does not depend on the choice of values to which other coordinates are fixed. This is due to a simple fact that any two such choices of loops can be connected by a continuous transformation, continuously changing one loop into the other (such a continuous transformation is usually called a *homotopy*), this implies that the degree has to change continuously as well. But the degree can only attain discrete values, and therefore it has to remain constant.

This assigns a quantity $\deg_i f$, which is always an integer, to each of the coordinates of f . Intuitively, we can say that the higher the absolute value of this degree is, the more important the corresponding variable is. In particular, inessential variables have degree 0. This is in essence how we identify which variables of a function are important, and which should become inessential after applying a minion homomorphism.

Remark IV.2. Using the fundamental groups in the n -ary case can also bring a little more insight. In particular, as it is well-known that $\Pi_1(T^n)$ is isomorphic to the n -generated free Abelian group. The loops that we described in the above para-




Fig. 2. A homotopy.

graphs correspond to the n different generators. And similarly, as in the unary case, any continuous mapping $f: T^n \rightarrow S^1$ induces a group homomorphism between the fundamental groups, i.e., a group homomorphism $\Pi_1(f): \mathbb{Z}^n \rightarrow \mathbb{Z}$. Any such map is of the form

$$(x_1, \dots, x_n) \mapsto c_1 x_1 + \dots + c_n x_n,$$

and each of these coefficients c_i correspond to the degree $\deg_i f$.

To bound the number of ‘interesting’ coordinates, we need use the discrete structure of the graph. One easy observation is that a degree of a graph homomorphism from \mathbf{C}_k to \mathbf{C}_3 cannot be arbitrarily large: we can walk around the cycle \mathbf{C}_3 at most $k/3$ times in k steps. We need to bring this bound on a single degree of a unary map to bound the number of coordinates with non-zero degree. This is done by proving that if f is n -ary, and g is defined from f by identifying all variables, i.e., $g(x) = f(x, \dots, x)$, then $\deg g = \deg_1 f + \dots + \deg_n f$. This is not so easy to see, let us sketch the proof for $n = 2$. Let $f: T^2 \rightarrow S^1$, then g is defined as the restriction of f to the diagonal, i.e., points with coordinates (x, x) , see the dashed line in Figure 1.

We want to connect the degree of this restriction with the degrees of the two restrictions of f to the loops that define $\deg_1 f$ and $\deg_2 f$. This is again done by observing that walking along the two loops one after another can be continuously transformed to walking along the diagonal. This can be done by continuously shifting the walk along the lines shown in Figure 2. A similar argumentation works for higher dimensions as well. The last small technical obstacle is what to do with negative degrees as they could cancel out with positive ones. This is only a minor problem since we can simply reverse the corresponding coordinates to obtain a mapping that has only positive degrees that are identical to the original ones, up to a sign.

Remark IV.3. The above argumentation is an instance of a more general statement that says that the mapping $\Pi_1: \mathcal{S}^1 \rightarrow \mathcal{L}$ (here \mathcal{S}^1 denotes the minion of all continuous maps from T^n to S^1 and \mathcal{L} the minion of all group




Fig. 3. The graph \mathbf{C}_9^2 on a torus.

homomorphisms from \mathbb{Z}^n to \mathbb{Z}) is a minion homomorphism. In other words, if g is defined from f using $\pi: \{1, \dots, n\} \rightarrow \{1, \dots, m\}$ by $g(x_1, \dots, x_m) = f(x_{\pi(1)}, \dots, x_{\pi(n)})$, then the same identity holds for $\Pi_1(g)$ and $\Pi_1(f)$, i.e.,

$$\deg_1 g \cdot x_1 + \dots + \deg_m g \cdot x_m = \deg_1 f \cdot x_{\pi(1)} + \dots + \deg_n f \cdot x_{\pi(n)}.$$

The above is equivalent to the statement that for all $i \in \{1, \dots, m\}$ we have

$$\deg_i g = \sum_{j \in \pi^{-1}(i)} \deg_j f.$$

In Section V, we prove that the degrees we define for graph polymorphisms also have this property.

In our attempt to bring these topological considerations to proper statements about polymorphisms from \mathbf{C}_k to \mathbf{C}_3 , there are a few points where the analogy does not work nicely. We already mentioned one, that the degree of a graph homomorphism is bounded, but the degree of a continuous map $S^1 \rightarrow S^1$ is not. This is due to the fact that, unlike topological spaces which are sometimes described as ‘being made of rubber’, i.e., they can be infinitely stretched and folded, graphs are ‘made of sticks’, i.e., they can be folded but not stretched. This property works to our advantage. The second issue is that the second power of \mathbf{C}_k is not exactly topologically equivalent to a torus, rather it forms a certain mesh that can be drawn on a torus in some way (see Figure 3). This can be avoided by using a less naïve assignment of a topological space to a graph and a more robust theory [32] which is in line with Lovász’s approach [28]. This approach is described in detail in a subsequent paper of Wrochna and Živný [27, Appendix A]. In this paper, we stick to the naïve approach and present an ad-hoc (‘discrete continuity’) argument using an alternative definition of a degree that resembles the topological one.

V. PROOF OF THEOREM III.3

We prove the theorem by analysing the polymorphisms from \mathbf{C}_k to $\mathbf{C}_3 (= \mathbf{K}_3)$, where $k \geq 3$ is an odd number (which we assume to be fixed for the rest of this section).

A. Degree of a Homomorphism

Recall that, for $m \geq 3$, we define the graph \mathbf{C}_m to be the m -cycle with vertices $0, 1, \dots, m-1$. Here vertices are connected by an edge if they differ by exactly one modulo m .

We fix an orientation of any \mathbf{C}_m in the increasing order modulo m , and denote by O_m the edge chain $[0, 1] + [1, 2] + \dots + [m-1, 0]$ in $\Delta_{\mathbb{E}}(\mathbf{C}_m)$.

The *degree* of a homomorphism $f: \mathbf{C}_m \rightarrow \mathbf{C}_l$ is intuitively defined as the (possibly non-positive) number of times the image of \mathbf{C}_m under f walks around \mathbf{C}_l in a fixed direction (say, counter-clockwise). The formal definition is based on the following observation.

Lemma V.1. *Let $m, l \geq 3$, and let $f: \mathbf{C}_m \rightarrow \mathbf{C}_l$ be a homomorphism. Then there is an integer d such that $f_{\mathbb{E}}(O_m) = d \cdot O_l$.*

Proof. Clearly, we have $\partial O_m = 0$. Lemma II.16 then implies that $\partial(f_{\mathbb{E}}(O_m)) = 0$. We claim that the only edge chains W in $\Delta_{\mathbb{E}}(\mathbf{C}_l)$ such that $\partial W = 0$ are chains of the form $d \cdot O_l$, so $f_{\mathbb{E}}(O_m)$ is of this form. Indeed, observe that if $W = d_0[0, 1] + \dots + d_{l-1}[l-1, 0]$, then

$$\begin{aligned} \partial W &= d_0([1] - [0]) + d_1([2] - [1]) + \dots + d_{l-1}([0] - [l-1]) \\ &= (d_{l-1} - d_0)[0] + (d_0 - d_1)[1] + \dots + (d_{l-2} - d_{l-1})[l-1]. \end{aligned}$$

If $\partial W = 0$, all coefficients in the above sum are 0, and therefore $d_0 = d_1 = \dots = d_{l-1}$ concluding that $W = d \cdot O_l$ for $d = d_0$. \square

Definition V.2. *Let $m, l \geq 3$. The *degree* of a homomorphism $f: \mathbf{C}_m \rightarrow \mathbf{C}_l$, denoted by $\deg f$, is defined as the integer d from the above lemma, i.e., the number $\deg f$ such that*

$$f_{\mathbb{E}}(O_m) = \deg f \cdot O_l.$$

Lemma V.3. *Let $m, l \geq 3$, assume that l is odd, and let $f: \mathbf{C}_m \rightarrow \mathbf{C}_l$ be a homomorphism. Then*

- 1) $|\deg f| \leq m/l$,
- 2) *the parity of $\deg f$ is the same as the parity of m , and*
- 3) *if $m = 4$ then $\deg f = 0$.*

Proof. (1) We have that

$$\begin{aligned} \deg f \cdot O_l &= f_{\mathbb{E}}(O_m) = \\ &[f(0), f(1)] + [f(1), f(2)] + \dots + [f(m-1), f(0)]. \end{aligned}$$

It is clear that, for each $[i, i+1]$ in O_l , the last expression above contains at least $|\deg f|$ terms that are either $[i, i+1]$ or $[i+1, i]$. It follows that $|\deg f| \leq m/l$.

(2) This follows by similar considerations as above. For each $[i, i+1]$ in O_l , the parity of the number of terms in the above sum that are either $[i, i+1]$ or $[i+1, i]$ is the same as the parity of $\deg f$. Since l is odd, the result follows.

(3) From (1) and (2), we know that the degree of any homomorphism $f: \mathbf{C}_4 \rightarrow \mathbf{C}_l$ is an even integer with absolute value at most $4/l$. For $l > 2$, there is only one such number, namely 0. \square

Note that as a direct consequence of item (2) in the above lemma, we get that for m, l odd, any homomorphism from \mathbf{C}_m to \mathbf{C}_l has a non-zero degree.

B. Degrees of a Polymorphism

We generalise the notion of a degree of a homomorphism to polymorphisms between odd cycles. More precisely, for a polymorphism $f: \mathbf{C}_k^n \rightarrow \mathbf{C}_3$ and coordinate $i \in \{1, \dots, n\}$, we define a quantity that we will call ‘a degree of f at coordinate i ’. Since this quantity will be used to define a minion homomorphism, the main requirement here will be that the degree behaves nicely with respect to minors. Formally, we will need that if g is a minor of f defined by

$$g(x_1, \dots, x_m) = f(x_{\pi(1)}, \dots, x_{\pi(n)}),$$

then

$$\deg_i g = \sum_{j \in \pi^{-1}(i)} \deg_j f.$$

This property is equivalent to saying that the mapping that maps f to the function on \mathbb{Z} defined by $(x_1, \dots, x_n) \mapsto \deg_1 f \cdot x_1 + \dots + \deg_n f \cdot x_n$ is minor-preserving.

Intuitively, a degree of a unary function counts how many times one loops around the cycle \mathbf{C}_3 if one follows the values of the function. We would like to bring this intuition to the n -ary case, so that the degree of f at some coordinate would mean ‘number of times one loops around the circle if one follows edges going in the given direction at this coordinate’.

We will formalise this intuition and prove that the degree at a coordinate can be defined in two equivalent ways, one global and the other local. In what follows we fix $l = 3$, but all proofs work for any odd $3 \leq l \leq k$.

We denote by $O_{k,i}^n$ the set of all oriented edges of \mathbf{C}_k^n whose i -th coordinate is oriented as in O_k , i.e.,

$$\begin{aligned} O_{k,i}^n &= \{[(a_1, \dots, a_n), (b_1, \dots, b_n)] \mid \\ &(a_j, b_j) \in E(C_k) \text{ for all } j \in \{1, \dots, n\} \text{ and} \\ &[a_i, b_i] \in O_k\}. \end{aligned}$$

We will also view $O_{k,i}^n$ as an edge chain in $\Delta_{\mathbb{E}}(\mathbf{C}_k^n)$.

Definition V.4. *Let $f: \mathbf{C}_k^n \rightarrow \mathbf{C}_3$ be a polymorphism. We define the *degree of f at coordinate i* as the integer $\deg_i f$ such that*

$$f_{\mathbb{E}}(O_{k,i}^n) = (2k)^{n-1} \deg_i f \cdot O_3.$$

Note that $|O_{k,i}^n| = 2^{n-1}k^n$, and therefore the above definition agrees with the intuitive meaning. Also if $n = 1$, then $\deg_1 f$ coincides with $\deg f$ since $(2k)^{1-1} = 1$ and $O_{k,1}^1 = O_k$. For a general n , it is not even clear that such a number $\deg_i f$ always exists. It is easy to show that there is an integer d' such that $f_{\mathbb{E}}(O_{k,i}^n) = d' \cdot O_3$ since $\partial(O_{k,i}^n) = 0$. However, there is no obvious reason that this number is a multiple of $(2k)^{n-1}$. Let us show that this is the case. We need a technical definition first.




Fig. 4. Proof of Lemma V.6(2).

Definition V.5. For an unoriented edge $e = (\bar{u}, \bar{v})$ of \mathbf{C}_k^{n-1} , we define

$$e \times_i O_k = \{[(u'_1, \dots, u'_n), (v'_1, \dots, v'_n)] \mid \\ \{(\dots, u'_{i-1}, u'_{i+1}, \dots), (\dots, v'_{i-1}, v'_{i+1}, \dots)\} = \{\bar{u}, \bar{v}\}, \\ [u'_i, v'_i] \in O_k\}.$$

Note that $|e \times_i O_k| = 2k$ since for each $[u_i, v_i] \in O_k$ there are two edges in $e \times_i O_k$ (one for each orientation of e) whose i -th coordinate agree with $[u_i, v_i]$. We also note that $O_{k,i}^n = \bigcup_e e \times_i O_k$ where the union runs through all unoriented edges of \mathbf{C}_k^{n-1} . Again, we can view $e \times_i O_k$ as an edge chain in $\Delta_E(\mathbf{C}_k^n)$.

Lemma V.6. Let $n \geq 2$, $f: \mathbf{C}_k^n \rightarrow \mathbf{C}_3$ be a polymorphism, and let $i \in \{1, \dots, n\}$. Then

- 1) for each edge e of \mathbf{C}_k^{n-1} , there is an integer d such that $f_E(e \times_i O_k) = 2d \cdot O_3$;
- 2) the above d does not depend on the choice of e ;
- 3) $d = \deg_i f$.

Proof. Without loss of generality, assume that $i = 1$, and to simplify the notation, we will write \times instead of \times_1 .

(1) Observe that $e \times O_k$ is an oriented $2k$ -cycle in \mathbf{C}_k^n , and consider $g_e: \mathbf{C}_{2k} \rightarrow \mathbf{C}_3$ to be the restriction of f to this $2k$ -cycle. Then $\deg g_e$ is even from Lemma V.3(2), and therefore it is equal to $2d$ for some d .

(2) We first prove the claim for two incident edges e and e' . Let

$$e = ((u_2, \dots, u_n), (v_2, \dots, v_n)), \\ e' = ((u_2, \dots, u_n), (w_2, \dots, w_n)).$$

We want to prove that $f_E(e \times O_k) = f_E(e' \times O_k)$ which is equivalent to $f_E(e \times O_k - e' \times O_k) = 0$ since f_E is a group homomorphism. Note that $-e' \times O_k$ is obtained from $e' \times O_k$ by reversing edges. Our goal is then decompose these two oriented cycles into several 4-cycles and then apply Lemma V.3(3). The four cycles are defined on vertices

$$(j, u_2, \dots, u_n), (j+1, v_2, \dots, v_n), \\ (j+2, u_2, \dots, u_n), (j+1, w_2, \dots, w_n)$$

where the addition in the first coordinate is considered modulo k . We denote by S_j the sum of oriented edges of the above 4-cycle, with the orientation following the order above. Observe that indeed (see Figure 4)

$$\sum_{j < k} S_j = e \times_i O_k - e' \times_i O_k,$$

and therefore

$$f_E(e \times_i O_k - e' \times_i O_k) = f_E\left(\sum_{j < k} S_j\right) = \sum_{j < k} f_E(S_j) = 0$$

where the last equality follows from Lemma V.3(3). This implies that $f_E(e \times_i O_k) = f_E(e' \times_i O_k)$, as required. The general case is then obtained by transitivity, since one can move from any edge of \mathbf{C}_k^{n-1} to any other edge by following a sequence of incident edges.

(3) We have

$$f_E(O_{k,i}^n) = \sum_{e \in E(\mathbf{C}_k^{n-1})} f_E(e \times_i O_k) = \\ \sum_{e \in E(\mathbf{C}_k^{n-1})} 2d \cdot O_3 = (2k)^{n-1} d \cdot O_3$$

since $|E(\mathbf{C}_k^{n-1})| = (2k)^{n-1}/2$. \square

C. Minor Preservation

Let \mathcal{L} denote the minion of all linear maps over \mathbb{Z} , i.e., of the functions of the form $\sum c_i x_i$ where all $c_i \in \mathbb{Z}$. We define a mapping $\delta: \text{Pol}(\mathbf{C}_k, \mathbf{C}_3) \rightarrow \mathcal{L}$ by

$$\delta(f): (x_1, \dots, x_n) \mapsto \deg_1 f \cdot x_1 + \dots + \deg_n f \cdot x_n.$$

In this subsection we prove that δ is minor-preserving, and therefore a minion homomorphism. In the following one we show that the image of δ contains only functions of bounded essential arity (but no constant function).

Lemma V.7. The map δ is a minion homomorphism.

It is clear that δ preserves the arity, so we need to show it also preserves the operation of taking minors. We decompose this operation into a few steps: permuting variables, introducing new dummy variables, and identifying two variables. It is not hard to observe that δ preserves the operation of

permuting variables (this corresponds to the case when π in Definition II.10 is a bijection). We deal with the case of identifying two variables in Lemma V.8, and then consider the addition of dummy variables in Lemma V.11.

Lemma V.8. *Let $n \geq 2$. If $f \in \text{Pol}(\mathbf{C}_k, \mathbf{C}_3)$ is n -ary and g is obtained from f by identifying the first two variables, i.e.,*

$$g(y, x_3, \dots, x_n) = f(y, y, x_3, \dots, x_n)$$

then $\deg_1 g = \deg_1 f + \deg_2 f$.

Before, we get to the proof, we need some technical definitions and a simple technical lemma. Similarly to $O_{k,i}^n$, we denote by $O_{k,\{1,2\}}^n$ the set of all oriented edges of \mathbf{C}_k^n whose first and second coordinates are oriented as in O_k , i.e.,

$$O_{k,\{1,2\}}^n = O_{k,1}^n \cap O_{k,2}^n.$$

Note that, unlike $O_{k,i}^n$, $O_{k,\{1,2\}}^n$ does not contain all edges of \mathbf{C}_k^n in some orientation, e.g. neither the edge $[(0,1), (1,0)]$ nor $[(1,0), (0,1)]$ is contained in $O_{k,\{1,2\}}^n$. Also observe that when considering this set as an edge chain, we have

$$O_{k,\{1,2\}}^n = 1/2 \cdot (O_{k,1}^n + O_{k,2}^n)$$

which follows, since in the sum on the right-hand side the edges that disagree in the orientation in the first two coordinates cancel out, and those which agree count twice.

We define the joint degree of f at coordinates 1 and 2, which intuitively expresses ‘the average number of times one loops around O_3 when following k edges that increase in both coordinates 1 and 2’. For the formal definition below, note that $|O_{k,\{1,2\}}^n| = 2^{n-2}k^n$.

Definition V.9. Let $n \geq 2$, and let $f: \mathbf{C}_k^n \rightarrow \mathbf{C}_3$ be a polymorphism. We define a *joint degree* of f at coordinates 1 and 2 as the integer $\deg_{1,2} f$ such that

$$f_{\mathbb{E}}(O_{k,\{1,2\}}^n) = 2^{n-2}k^{n-1} \deg_{1,2} f \cdot O_3.$$

As in the case of degrees of polymorphisms, it is not obvious that such a number exists, but we prove this in the following lemma.

Lemma V.10. *For each $n \geq 2$ and a polymorphism $f: \mathbf{C}_k^n \rightarrow \mathbf{C}_3$, we have*

$$\deg_{1,2} f = \deg_1 f + \deg_2 f.$$

Proof. Since $O_{k,\{1,2\}}^n = 1/2 \cdot (O_{k,1}^n + O_{k,2}^n)$, we have

$$f_{\mathbb{E}}(O_{k,\{1,2\}}^n) = 1/2 \cdot f_{\mathbb{E}}(O_{k,1}^n + O_{k,2}^n),$$

and consequently,

$$2^{n-2}k^{n-1} \deg_{1,2} f = 1/2 \cdot ((2k)^{n-1} \deg_1 f + (2k)^{n-1} \deg_2 f).$$

Cancelling $2^{n-2}k^{n-1}$ on both sides gives the claim. \square

Proof of Lemma V.8. By the previous lemma, it is enough to prove that $\deg_{1,2} f = \deg g$. This is done in a similar way to proving that the degree of f at a coordinate is both a local and a global property of f (Lemma V.6). We prove this statement




Fig. 5. Proof of Case 1 of Lemma V.8.

separately for two cases: (1) f is binary and g is unary; and (2) f has arity at least 3 and g has arity at least 2. The two cases are very similar. We present them separately to ease some technical difficulties of the proof.

Case 1: f is binary. The assumption says that $g(x) = f(x, x)$, and we aim to prove that $\deg g = \deg_{1,2} f$. Note that

$$g_{\mathbb{E}}(O_k) = f_{\mathbb{E}}(O'_{k,0})$$

where $O'_{k,0}$ is the set of oriented edges of the k -cycle $(0,0), (1,1), \dots, (k-1, k-1)$ in \mathbf{C}_k^2 (in the increasing orientation). Note that $O'_{k,0} \subseteq O_{k,\{1,2\}}^2$. More generally, for $i \in \{1, \dots, k-1\}$, we denote by $O'_{k,i}$ the set of oriented edges of the k -cycle $(0,i), (1,i+1), \dots, (k-1, i-1)$ where in the second coordinate, 0 succeeds $k-1$. Thus, the set $O_{k,\{1,2\}}^2$ is a disjoint union of the cycles $O'_{k,i}$, $0 \leq i \leq k-1$. Since each $O'_{k,i}$ is a k -cycle, we know that there is d_i such that

$$f_{\mathbb{E}}(O'_{k,i}) = d_i \cdot O_3.$$

In a similar way as in Lemma V.6, we prove that d_i does not depend on i and is actually equal to $\deg_{1,2} f$. Also note that $d_0 = \deg g$.

First, we fix any $0 \leq i \leq k-1$ and show that we have $d_i = d_{i+2}$ where the addition is modulo k . For each $x = 0, \dots, k-1$, let S_x be the set of edges of the oriented 4-cycle

$$(x, x+i), (x+1, x+i+1), (x, x+i+2), (x-1, x+i+1)$$

where the addition is considered modulo k . Observe that $\sum_{x=0}^{k-1} S_x = O'_{k,i} - O'_{k,i+2}$ (see Figure 5). By Lemma V.3(3), we have $f_{\mathbb{E}}(S_x) = 0$ for each x . So we get

$$0 = f_{\mathbb{E}}\left(\sum_{x=0}^{k-1} S_x\right) = f_{\mathbb{E}}(O'_{k,i} - O'_{k,i+2}) = f_{\mathbb{E}}(O'_{k,i}) - f_{\mathbb{E}}(O'_{k,i+2})$$

and therefore $f_E(O'_{k,i}) = f_E(O'_{k,i+2})$ which implies that $d_i = d_{i+2}$. Since k is odd, it follows that d_i is the same for all i . This implies that $d_i = \deg g$, and therefore

$$f_E(O_{k,\{1,2\}}^2) = \sum_{i=0}^{k-1} f_E(O'_{k,i}) = \sum_{i=0}^{k-1} \deg g \cdot O_3 = k \deg g \cdot O_3.$$

Consequently, $\deg_{1,2} f = \deg g$ which concludes the proof of the first case.

Case 2: The polymorphism f is of arity $n > 2$. The proof is similar to the first case: We decompose the set $O_{k,\{1,2\}}$ into k disjoint sets each of which is a copy of $O_{k-1,1}^{n-1}$ used in the definition of $\deg_1 g$. One of these copies will exactly correspond to the edges of \mathbf{C}_k^n on vertices that have the first two coordinate identical. This is important since g is essentially the restriction of f to such vertices.

The sets P'_i are defined to consist of those oriented edges of \mathbf{C}_k^n whose projection to the first two coordinates is in $O'_{k,i}$. In other words, we put

$$P'_i = \{[(x, x+i, u_3, \dots, u_n), (x+1, x+i+1, v_3, \dots, v_n)] \mid x = 0, \dots, k-1, \text{ and } (u_j, v_j) \in E(C_k) \text{ for all } j \geq 3\}$$

(again, the addition in the first two coordinates is considered modulo k). It is easy to see that $\partial P'_i = 0$, $O_{k,\{1,2\}} = \bigcup_{i=0}^{k-1} P'_i$, and $g_E(O_{k,1}^{n-1}) = f_E(P'_0)$.

We prove $f_E(P'_i) = f_E(P'_{i+2})$. For each edge

$$e = [(x, x+i, u_3, \dots, u_n), (x+1, x+i+1, v_3, \dots, v_n)] \in P'_i,$$

let S_e denote the set of oriented edges of the 4-cycle

$$\begin{aligned} & (x, x+i, u_3, \dots, u_n), (x+1, x+i+1, v_3, \dots, v_n), \\ & (x, x+i+2, u_3, \dots, u_n), (x-1, x+i+1, v_3, \dots, v_n). \end{aligned}$$

Note that $e \in S_e$, and also $[(x, x+i+2, u_3, \dots, u_n), (x-1, x+i+1, v_3, \dots, v_n)] \in -P'_{i+2}$. We claim that $\sum_{e \in P'_i} S_e = P'_i - P'_{i+2}$. In other words, all edges of the S_e 's not contained in $P'_i - P'_{i+2}$ cancel out. Indeed, every edge of the form

$$d = [(x+1, x+i+1, v_3, \dots, v_n), (x, x+i+2, u_3, \dots, u_n)]$$

that appears in S_e , has its reverse in S_{e^+} for

$$e^+ = [(x+1, x+i+1, u_3, \dots, u_n), (x+2, x+i+2, v_3, \dots, v_n)].$$

Furthermore, the correspondence $e \leftrightarrow e^+$ is 1-to-1, and therefore we paired each $d \in S_e$ with a unique $-d \in S_{e^+}$.

We conclude the proof in a similar way as Case 1: From Lemma V.3(3), we have $f_E(S_e) = 0$ for each e , and therefore

$$0 = f_E\left(\sum_{e \in P'_i} S_e\right) = f_E(P'_i - P'_{i+2})$$

which shows that $f_E(P'_i) = f_E(P'_{i+2})$. Consequently, we have $P'_i = P'_j$ for all i, j , and

$$\begin{aligned} f_E(O_{k,\{1,2\}}^n) &= \sum_{i < k} f_e(P'_i) = k \cdot f_e(P'_0) = \\ & k \cdot g_E(O_{k,1}^{n-1}) = k(2k)^{n-2} \deg_1 g \cdot O_3. \end{aligned}$$

This implies that $\deg_{1,2} f = \deg_1 g$, as required. \square

The following lemma says that δ preserves the operation of adding a dummy variable.

Lemma V.11. *Let $f: \mathbf{C}_k^n \rightarrow \mathbf{C}_3$ be a polymorphism, and $i \in \{1, \dots, n\}$. If the coordinate i in f is dummy, then $\deg_i f = 0$.*

Proof. Loosely speaking, the degree of f at the i -th coordinate is determined by the image of $O_{k,i}^n$ under f_E . Since f does not depend on the i -th coordinate, neither does (in the corresponding meaning) f_E , and therefore, it cannot distinguish the orientation of the edges in $O_{k,i}^n$.

Formally, we use the local definition of $\deg_i f$ and prove that, for an edge $e \in E(C_k^{n-1})$, we have $f_E(e \times_i O_k) = 0$. Fix $e = ((u_2, \dots, u_n), (v_2, \dots, v_n))$ and assume without loss of generality that $i = 1$. As mentioned in the proof of Lemma V.6(1), $e \times_1 O_k$ is a $2k$ -cycle. The values of f on this cycle are all of the form $f(a, u_2, \dots, u_n)$ and $f(a, v_2, \dots, v_n)$ where $a \in V(C_k)$. Since the first coordinate in f is dummy, none of this values depends on a , so f attains only two possible values on this $2k$ -cycle. This implies that necessarily $f_E(e \times_1 O_k) = 0 \cdot O_3$, and consequently, by Lemma V.6, $\deg_1 f = 0$. \square

This concludes the proof of Lemma V.7.

D. Bounding the Essential Arity

To finish the analysis of polymorphisms from \mathbf{C}_k to \mathbf{C}_3 , we need to bound the essential arity of functions in the image of δ (defined in Section V-C) and show that none of these functions is a constant function. We now prove that the image of δ is contained in $\mathcal{Z}_{\leq N}$, where N is the largest odd number such that $N \leq k/3$. Recall that, for an odd number N , the minion $\mathcal{Z}_{\leq N}$ is defined to be the set of all functions $f: \mathbb{Z}^n \rightarrow \mathbb{Z}$ of the form $f(x_1, \dots, x_n) = c_1 x_1 + \dots + c_n x_n$ for some $c_1, \dots, c_n \in \mathbb{Z}$ with $\sum_{i=1}^n |c_i| \leq N$ and $\sum_{i=1}^n c_i$ odd. It is easy to see that the number of non-zero coefficients in any function from $\mathcal{Z}_{\leq N}$ is between 1 and N , giving us the required result. We remark that our bound on N is tight (for a proof of this fact, see Lemma A.1), but we will not need that.

Lemma V.12. *Let N be the largest odd number such that $N \leq k/3$. Then we have $\delta(f) \in \mathcal{Z}_{\leq N}$ for all $f \in \text{Pol}(\mathbf{C}_k, \mathbf{C}_3)$.*

Proof. We need to prove (1) $\sum_{i=1}^n |\deg_i f| \leq N$, and (2) $\sum_{i=1}^n \deg_i f$ is odd. Consider the unary minor of f , i.e., the mapping $g: \mathbf{C}_k \rightarrow \mathbf{C}_3$ defined by $g(x) = f(x, \dots, x)$. Parts (1) and (2) of Lemma V.3 imply that $\deg g$ is an odd number not greater than $k/3$, consequently $\deg g \leq N$. Further, Lemma V.7 implies that $\deg g = \sum_{i=1}^n \deg_i f$, and therefore we immediately get item (2). This argument also shows item (1) if $\deg_i f$ is non-negative for all i . We reduce the general case to this case. More precisely, for each f we will find a new polymorphism f' of the same arity such that $|\deg_i f| = \deg_i f'$ for all i .

This f' is constructed by a simple trick: ‘reversing’ all coordinates x_i with $\deg_i f$ negative. Let us do that for one coordinate at a time, and for simplicity, consider just the




Fig. 6. The graph D_9 with a 3-colouring h_9 .

first coordinate. Let θ be the automorphism of C_k such that $\theta(0) = 0$ and $\theta(i) = k - i$ for $i \neq 0$, and define f' as

$$f'(x_1, \dots, x_n) = f(\theta(x_1), x_2, \dots, x_n).$$

It is easy to see that f' is also a polymorphism. We claim that $\deg_1 f' = -\deg_1 f$ and $\deg_i f' = \deg_i f$ for all $i \neq 1$. This follows from the fact that applying θ in the first coordinate reverses the orientation of all edges in $O_{k,1}^n$ and it does not change the orientation of edges in $O_{k,j}^n$ for $j \neq 1$. (This is clear since the orientation of edges in $O_{k,i}^n$ is given by the orientation in the i -th coordinate.) In other words, we have $f'_E(O_{k,1}^n) = f_E(-O_{k,1}^n)$ and $f'_E(O_{k,j}^n) = f_E(O_{k,j}^n)$ for $j \neq 1$. This directly implies the claim about degrees.

Repeating this trick, we eventually obtain a polymorphism whose degrees are all positive and up to the sign same as degrees of f . This completes the proof. \square

Theorem III.3 now follows from Lemma V.7 and Lemma V.12.

APPENDIX

In this appendix, we prove that (for k odd and N the largest odd number smaller than or equal to $k/3$) the minion $\mathcal{L}_{\leq N}$ is homomorphically equivalent to $\text{Pol}(C_k, C_3)$, i.e., that there exists minion homomorphisms between these minions in both directions. A minion homomorphism $\delta: \text{Pol}(C_k, C_3) \rightarrow \mathcal{L}_{\leq N}$ was given in Section V, here we provide one from $\mathcal{L}_{\leq N}$ to $\text{Pol}(C_k, C_3)$.

Lemma A.1. *Let k and N be odd such that $N \leq k/3$. There exists a minion homomorphism $\eta: \mathcal{L}_{\leq N} \rightarrow \text{Pol}(C_k, C_3)$.*

Proof. For simplicity, let us assume that $k = 3N$. The general case is obtained by an easy observation that $\text{Pol}(C_k, C_3) \rightarrow \text{Pol}(C_{k'}, C_3)$ for any odd $k' \leq k$. As an intermediate step, we define a graph D_k with vertices $V(D_k) = V(C_k)$ such that $(u, v) \in E(D_k)$ if the distance of u and v in C_k is odd and at most N . Alternatively, we can say $(u, v) \in E(D_k)$ if they are connected in C_k by a walk of length exactly N . It is easy to see that these two are equivalent. (See also Figure 6.)

We claim:

- 1) there is a minion homomorphism η' from $\mathcal{L}_{\leq N}$ to $\text{Pol}(C_k, D_k)$, and
- 2) there is a graph homomorphism $h_k: D_k \rightarrow C_3$.

To prove the first claim, we define η' by

$$\eta'(f): (x_1, \dots, x_n) \mapsto f(x_1, \dots, x_n) \bmod k.$$

That is, we apply the function f to the n -tuple of vertices of C_k as they would be numbers in \mathbb{Z} , and then take the residue modulo k of the result to get a number between 0 and $k - 1$. From the definition, it is clear that η' is minor-preserving, therefore we only need to prove that each $\eta'(f)$ is a polymorphism from C_k to D_k . It is enough to prove the claim for f of the form: $f(x_1, \dots, x_N) = \pm x_1 \pm \dots \pm x_N$ since all other functions in $\mathcal{L}_{\leq N}$ are minors of some such f . Let $f' = \eta'(f)$, i.e.,

$$f'(x_1, \dots, x_N) = (\pm x_1 \pm \dots \pm x_N) \bmod k$$

Assume that $(u_i, v_i) \in E(C_k)$ for $i = 1, \dots, N$, and observe that

$$f'(u_1, \dots, u_N), f'(v_1, u_2, \dots, u_N), f'(v_1, v_2, u_3, \dots, u_N), \dots, f'(v_1, \dots, v_N)$$

is a walk in C_k from $f'(u_1, \dots, u_N)$ to $f'(v_1, \dots, v_N)$ of length exactly N , which implies that

$$(f'(u_1, \dots, u_N), f'(v_1, \dots, v_N)) \in E(D_k)$$

and $f' \in \text{Pol}(C_k, D_k)$.

For the second claim, consider $h_k: V(D_k) \rightarrow V(C_3)$ defined by

$$h_k(x) = \begin{cases} 0 & \text{if } x < N \text{ and } x \text{ is even, or} \\ & x > 2N \text{ and } x \text{ is odd,} \\ 1 & \text{if } x < 2N \text{ and } x \text{ is odd,} \\ 2 & \text{if } x > N \text{ and } x \text{ is even.} \end{cases}$$

It is easy to check that h is indeed a graph homomorphism. (Figure 6 shows h_9 .)

Finally, the minion homomorphism η is given by

$$\eta(f): (x_1, \dots, x_n) \mapsto h_k(\eta'(f)(x_1, \dots, x_n)).$$

It is straight-forward to check that η is minor-preserving, we also have that $\eta(f)$ is a polymorphism from C_k to C_3 since it is a composition of a polymorphism $\eta'(f): C_k^n \rightarrow D_k$ and a homomorphism $D_k \rightarrow C_3$. \square

We remark, without giving a proof, that the composition of minion homomorphisms $\delta \circ \eta$ is the identity on $\mathcal{L}_{\leq N}$, i.e., for $f(x_1, \dots, x_n) = \sum_{i=1}^n c_i x_i$ and $i \in \{1, \dots, n\}$, we have $\deg_i \eta(f) = c_i$.

As a simple corollary, we can obtain polymorphisms that forbid simple reductions from some other PCSPs.

Corollary A.2. *Let $k \geq 9$ and $n \leq k/3$ be both odd, $\text{Pol}(C_k, C_3)$ contains functions c, s , and o satisfying:*

- 1) $c_n(x_1, \dots, x_n) \approx c(x_2, \dots, x_n, x_1)$,
- 2) $s(x, y, x, z, y, z) \approx s(y, x, z, x, z, y)$, and

$$3) \ o(x, x, y, y, y, x) \approx o(x, y, x, y, x, y) \\ \approx o(y, x, x, x, y, y).$$

Proof. Let N be the largest odd number smaller than or equal to $k/3$, note that $n \leq N$ and $3 \leq N$. Since a minion homomorphism preserves satisfaction of the above identities, it is enough to find such functions in $\mathcal{L}_{\leq N}$. For that, we consider

$$c_n(x_1, \dots, x_n) = x_1 + \dots + x_n, \\ s(x_1, \dots, x_6) = x_1 + x_3 + x_5, \\ o(x_1, \dots, x_6) = x_1 + x_2 + x_3.$$

Clearly, $c_n, s, o \in \mathcal{L}_{\leq N}$, and it is easy to check that they satisfy the required identities. The claim is then given by Lemma A.1. \square

A function satisfying item (3) of the above corollary is called an *Olšák* function, and absence of such a polymorphism is a requirement for a reduction from approximate hypergraph colouring using [5, Corollary 6.3]. A function satisfying item (2) is called a *Siggers* function, and absence of such shows that [5, Theorem 3.1] cannot be used for a reduction from approximate graph colouring to PCSP($\mathbf{C}_k, \mathbf{C}_3$) (see also [5, Theorem 6.9]).

ACKNOWLEDGEMENTS

We would like to thank John Hunton for useful discussions about algebraic topology, and the reviewers for useful comments.

This paper was supported by the UK EPSRC grant EP/R034516/1.

REFERENCES

- [1] K. Kawarabayashi and M. Thorup, “Coloring 3-colorable graphs with less than $n^{1/5}$ colors,” *J. ACM*, vol. 64, no. 1, pp. 4:1–4:23, 2017.
- [2] S. Khanna, N. Linial, and S. Safra, “On the hardness of approximating the chromatic number,” *Combinatorica*, vol. 20, no. 3, pp. 393–415, Mar 2000.
- [3] V. Guruswami and S. Khanna, “On the hardness of 4-coloring a 3-colorable graph,” *SIAM Journal on Discrete Mathematics*, vol. 18, no. 1, pp. 30–40, 2004.
- [4] J. Bulín, A. Krokhnin, and J. Opršal, “Algebraic approach to promise constraint satisfaction,” in *Proceedings of the 51st Annual ACM SIGACT Symposium on the Theory of Computing (STOC ’19)*. New York, NY, USA: ACM, 2019.
- [5] L. Barto, J. Bulín, A. Krokhnin, and J. Opršal, “Algebraic approach to promise constraint satisfaction,” Jun. 2019, arXiv:1811.00970, version 3. [Online]. Available: <https://arxiv.org/abs/1811.00970v3>
- [6] P. Hell and J. Nešetřil, *Graphs and Homomorphisms*. Oxford University Press, 2004.
- [7] L. Barto, A. Krokhnin, and R. Willard, “Polymorphisms, and how to use them,” in *The Constraint Satisfaction Problem: Complexity and Approximability*, ser. Dagstuhl Follow-Ups, A. Krokhnin and S. Živný, Eds. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017, vol. 7, pp. 1–44.
- [8] T. Feder and M. Y. Vardi, “The computational structure of monotone monadic SNP and constraint satisfaction: A study through datalog and group theory,” *SIAM J. Comput.*, vol. 28, no. 1, pp. 57–104, Feb. 1998.
- [9] P. Hell and J. Nešetřil, “On the complexity of H -coloring,” *J. Combin. Theory Ser. B*, vol. 48, no. 1, pp. 92–110, 1990.
- [10] A. A. Bulatov, “A dichotomy theorem for nonuniform CSPs,” in *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, Oct. 2017, pp. 319–330.
- [11] D. Zhuk, “A proof of CSP dichotomy conjecture,” in *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, Oct 2017, pp. 331–342.
- [12] L. Barto and M. Kozik, “Robustly solvable constraint satisfaction problems,” *SIAM Journal on Computing*, vol. 45, no. 4, pp. 1646–1669, 2016.
- [13] J. Brakensiek and V. Guruswami, “New hardness results for graph and hypergraph colorings,” in *31st Conference on Computational Complexity (CCC 2016)*, ser. Leibniz International Proceedings in Informatics (LIPIcs), R. Raz, Ed., vol. 50. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2016, pp. 14:1–14:27.
- [14] —, “Promise constraint satisfaction: Structure theory and a symmetric boolean dichotomy,” in *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, ser. SODA’18. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2018, pp. 1782–1801.
- [15] P. Austrin, V. Guruswami, and J. Håstad, “ $(2+\epsilon)$ -Sat is NP-hard,” *SIAM J. Comput.*, vol. 46, no. 5, pp. 1554–1573, 2017.
- [16] J. Brakensiek and V. Guruswami, “An algorithmic blend of LPs and ring equations for promise CSPs,” in *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, 2019, pp. 436–455.
- [17] M. R. Garey and D. S. Johnson, “The complexity of near-optimal graph coloring,” *J. ACM*, vol. 23, no. 1, pp. 43–49, 1976.
- [18] I. Dinur, E. Mossel, and O. Regev, “Conditional hardness for approximate coloring,” *SIAM J. Comput.*, vol. 39, no. 3, pp. 843–873, 2009.
- [19] S. Huang, “Improved hardness of approximating chromatic number,” in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques: 16th International Workshop, APPROX 2013, and 17th International Workshop, RANDOM 2013, Berkeley, CA, USA, August 21-23, 2013. Proceedings*, P. Raghavendra, S. Raskhodnikova, K. Jansen, and J. D. P. Rolim, Eds. Berlin, Heidelberg: Springer, 2013, pp. 233–243.
- [20] I. Dinur, O. Regev, and C. Smyth, “The hardness of 3-uniform hypergraph coloring,” *Combinatorica*, vol. 25, no. 5, pp. 519–535, Sep. 2005.
- [21] P. Austrin, A. Bhangale, and A. Potukuchi, “Improved inapproximability of rainbow coloring,” Oct. 2018, arXiv:1810.02784. [Online]. Available: <http://arxiv.org/abs/1810.02784>
- [22] J. Brakensiek and V. Guruswami, “The quest for strong inapproximability results with perfect completeness,” in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2017, August 16–18, 2017, Berkeley, CA, USA, 2017*, pp. 4:1–4:20.
- [23] V. Guruswami and E. Lee, “Strong inapproximability results on balanced rainbow-colorable hypergraphs,” *Combinatorica*, Dec. 2017.
- [24] P. Austrin, A. Bhangale, and A. Potukuchi, “Simplified inapproximability of hypergraph coloring via t -agreeing families,” Apr. 2019, arXiv:1904.01163. [Online]. Available: <http://arxiv.org/abs/1904.01163>
- [25] A. Bhangale, “NP-hardness of coloring 2-colorable hypergraph with poly-logarithmically many colors,” in *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic, 2018*, pp. 15:1–15:11.
- [26] A. Krokhnin and S. Živný, Eds., *The Constraint Satisfaction Problem: Complexity and Approximability*, ser. Dagstuhl Follow-Ups. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017, vol. 7.
- [27] M. Wrochna and S. Živný, “Improved hardness for H -colourings of G -colourable graphs,” Jul. 2019, arXiv:1907.00872. [Online]. Available: <http://arxiv.org/abs/1907.00872>
- [28] L. Lovász, “Kneser’s conjecture, chromatic number, and homotopy,” *Journal of Combinatorial Theory, Series A*, vol. 25, no. 3, pp. 319–324, 1978.
- [29] J. Matoušek, *Using the Borsuk-Ulam Theorem*, 1st ed., ser. Lectures on Topological Methods in Combinatorics and Geometry. Springer-Verlag Berlin Heidelberg, 2003.
- [30] D. L. Greenwell and L. Lovász, “Applications of product colouring,” *Acta Mathematica Academiae Scientiarum Hungarica*, vol. 25, no. 3, pp. 335–340, Sep. 1974.
- [31] A. Hatcher, *Algebraic topology*. Cambridge: Cambridge University Press, 2001.
- [32] J. W. Walker, “From graphs to ortholattices and equivariant maps,” *Journal of Combinatorial Theory, Series B*, vol. 35, no. 2, pp. 171–192, 1983.