TOPOLOGY AND ADJUNCTION IN PROMISE CONSTRAINT SATISFACTION*

ANDREI KROKHIN[†], JAKUB OPRŠAL[‡], MARCIN WROCHNA[§], AND STANISLAV ŽIVNÝ[¶]

Abstract. The approximate graph coloring problem, whose complexity is unresolved in most cases, concerns finding a *c*-coloring of a graph that is promised to be *k*-colorable, where $c \ge k$. This problem naturally generalizes to promise graph homomorphism problems and further to promise constraint satisfaction problems. The complexity of these problems has recently been studied through an algebraic approach. In this paper, we introduce two new techniques to analyze the complexity of promise CSPs: one is based on topology and the other on adjunction. We apply these techniques, together with the previously introduced algebraic approach, to obtain new unconditional NP-hardness results for a significant class of approximate graph coloring and promise graph homomorphism problems.

Key words. graph coloring, graph homomorphism problem, constraint satisfaction, polymorphism, promise problem

MSC codes. 68Q17, 68Q25, 68R05, 05C15

DOI. 10.1137/20M1378223

1. Introduction. In this paper we investigate the complexity of finding an approximate solution to fully satisfiable instances of constraint satisfaction problems. For example, for the classical problem of k-coloring a graph, one natural approximation version is the approximate graph coloring problem: the goal is to find a c-coloring of a given k-colorable graph, where $c \ge k \ge 3$. There is a huge gap in our understanding of the complexity of this problem. For k = 3, the best known efficient algorithm uses roughly $c = O(n^{0.199})$ colors where n is the number of vertices of the graph [48]. It has been long conjectured the problem is NP-hard for any fixed constants $c \ge k \ge 3$, but, say for k = 3, the state-of-the-art has only recently been improved from c = 4 [38, 49] to c = 5 [5, 23].

Graph coloring problems naturally generalize 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

^{*}Received by the editors November 3, 2020; accepted for publication (in revised form) September 23, 2022; published electronically February 14, 2023. Preliminary versions of parts of this paper appeared in the Proceedings of the 60th Annual Symposium on Foundations of Computer Science (FOCS'19), IEEE, Piscataway, NJ, 2019, pp. 1227–1239 and Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'20), SIAM, Philadelphia, 2020, pp. 1426–1435.

https://doi.org/10.1137/20M1378223

Funding: This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement 714532). The first and second authors were supported by the UK EPSRC grant EP/R034516/1. The fourth author was supported by a Royal Society University Research Fellowship. The paper reflects only the authors' views and not the views of the ERC or the European Commission. The European Union is not liable for any use that may be made of the information contained therein.

Department of Computer Science, Durham University, UK (andrei.krokhin@durham.ac.uk).

[‡]Institute of Science and Technology Austria, Klosterneuburg, Austria (jakub.oprsal@ist.ac.at).

[§]Faculty of Mathematics, Informatics, and Mechanics, Unversity of Warsaw, Poland (m.wrochna@mimuw.edu.pl).

 $[\]P$ Department of Computer Science, University of Oxford, Oxford, UK (standa.zivny@cs.ox ac.uk).

homomorphism (an edge-preserving map) from the first graph to the second [45]. The CSP is a generalization 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 nonuniform CSP [7, 30]. This is also the only case we consider in this article. For graph homomorphisms, this is known as the **H**-coloring problem: decide whether a given graph has a homomorphism to a fixed graph **H** [45]. The P versus NP-complete dichotomy of **H**-coloring given in [44] was one of the base cases that supported the Feder–Vardi dichotomy conjecture for CSPs [30]. The study of the complexity of the CSP and the complete resolution of the CSP dichotomy conjecture [22, 75, 76] was greatly influenced by the algebraic approach [21] (see survey [7]). This approach has also made important contributions to the study of approximability of CSPs (e.g., [6]).

Brakensiek and Guruswami [13, 16] suggested that perhaps progress on approximate graph coloring 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 CSP (PCSP). Promise graph homomorphism is an approximation version of the graph homomorphism problem in the following sense: in PCSP(\mathbf{H}, \mathbf{G}), we fix (not one but) two graphs \mathbf{H} and \mathbf{G} such that there is a homomorphism from \mathbf{H} to \mathbf{G} (we write $\mathbf{H} \rightarrow \mathbf{G}$ to denote this). The goal is then to find a \mathbf{G} -coloring for a given graph when an \mathbf{H} -coloring is guaranteed to exist (but not given as part of input). The promise is that the input graph is always \mathbf{H} -colorable and hence \mathbf{G} -colorable as well. The PCSP is a natural generalization of this to arbitrary relational structures, or in other words, a generalization of the decision CSP to the promise setting. Brakensiek and Guruswami proposed a conjecture that PCSP(\mathbf{H}, \mathbf{G}) is NP-hard for all nonbipartite loopless graphs \mathbf{H} and \mathbf{G} such that $\mathbf{H} \rightarrow \mathbf{G}$. This would generalize the approximate graph coloring conjecture and greatly extend the Hell–Nešetřil dichotomy for \mathbf{H} -coloring [44].

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, Guruswami, and Håstad [3], with an application to a promise version of SAT. It was further developed by Brakensiek, Guruswami and others [13, 14, 15, 16, 18, 19] and applied to a range of problems, including versions of approximate graph and hypergraph coloring. A recent pair of papers [5, 23] describe a general abstract algebraic theory for PCSPs, which shows, in particular, how algebraic properties precisely capture the power of gadget reductions in PCSPs. However, the algebraic theory of PCSPs 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-coloring a given 3-colorable graph was proved in [5, 23] by applying this abstract theory.

The gist of the algebraic theory is that the complexity of $PCSP(\mathbf{H}, \mathbf{G})$ depends only on (certain properties of) *polymorphisms*, which are multivariable functions that can be defined as homomorphisms from direct powers \mathbf{H}^n into \mathbf{G} . However, the analysis of polymorphisms is in general a highly nontrivial task, and powerful tools are needed to conduct it. For resolving the CSP dichotomy conjecture, the structural theory of finite universal algebras provided such a tool. However, it is not clear how much this theory can be applied to the promise setting. In this paper, we show that algebraic topology gives a very useful tool to analyze polymorphisms and pinpoint the complexity of PCSPs. We do this by explaining how general PCSPs are naturally equipped with a topological structure, called homomorphism complexes, and how polymorphisms of a given PCSP can be understood through the continuous maps they induce. Homomorphism complexes (as well as several related constructions) have been actively studied in topological combinatorics [53, 59], though mainly to give obstructions to the existence of homomorphisms and mostly for the case of graphs. However, methods of algebraic topology can also be used to obtain important information about polymorphisms: for example, to identify "influential" variables. We demonstrate how this new methodology can be applied to resolve a significant part of the Brakensiek–Guruswami conjecture.

We also show that the simple notion of adjunction, which is a certain form of homomorphism duality, provides a powerful tool to reason about reductions between PCSPs. We observe that adjunctions always give rise to reductions between PC-SPs. Moreover, we prove that many reductions between PCSPs work because of the presence of adjunction. This includes, in particular, all gadget reductions (that are captured by the algebraic approach) and all reductions satisfying very mild technical conditions. We demonstrate how adjunction can be applied by significantly improving the state-of-the-art in approximate graph coloring—via reductions that provably cannot be explained via the algebraic approach from [5].

Related work. The notion of PCSP was coined in [3], though one of the main examples of problems of this form, approximate graph coloring, has been around for a long time [35]. The complexity landscape of PCSPs (beyond CSPs) is largely unknown, even for the Boolean case (see [16, 31]) or for graph coloring and homomorphisms.

Most notable examples of PCSPs studied before are related to graph and hypergraph coloring. We already mentioned some results concerning coloring 3-colorable graphs with a constant number of colors. Without additional complexity-theoretic assumptions, the strongest known NP-hardness results for coloring k-colorable graphs are as follows. For any $k \geq 3$, it is NP-hard to color a given k-colorable graph with 2k-1 colors [5, 23]. For large enough k, it is NP-hard to color a given k-colorable graph with $2^{\Omega(k^{1/3})}$ colors [46]. The only earlier result about promise graph homomorphisms (with $\mathbf{H} \neq \mathbf{G}$) that involves more than approximate graph coloring is the NP-hardness of 3-coloring for graphs that admit a homomorphism to \mathbf{C}_5 , the five-element cycle [5].

Under stronger assumptions (Khot's 2-to-1 conjecture [51] for $k \ge 4$ and its nonstandard variant for k = 3), Dinur, Mossel, and Regev showed that finding a *c*-coloring of a *k*-colorable graph is NP-hard for all constants $c \ge k \ge 3$ [27]. It was shown in [40] that the above result for k = 2d still holds if one assumes the *d*-to-1 conjecture of Khot [51] for any fixed $d \ge 2$ instead of the 2-to-1 conjecture (which is the strongest in the family of *d*-to-1 conjectures). A variant of Khot's 2-to-1 conjecture with imperfect completeness has recently been proved [26, 52], which implies hardness for approximate coloring variants for the weaker promise that most but not all of the graph is guaranteed to be *k*-colorable.

A coloring of a hypergraph is an assignment of colors to its vertices that leaves no edge monochromatic. It is known that, for any constants $c \ge k \ge 2$, it is NP-hard to find a *c*-coloring of a given 3-uniform *k*-colorable hypergraph [28]. Further variants of approximate hypergraph coloring, e.g., relating to strong or rainbow colorings, were studied in [2, 13, 17, 39, 41], but most complexity classifications related to them are still open in full generality.

Some results are also known for colorings with a superconstant number of colors. For graphs, conditional hardness can be found in [29], and for hypergraphs, NP-hardness results were obtained in [1, 9].

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

40

[22, 75, 76] of the Feder–Vardi conjecture are presented. The book [55] 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, Guruswami, and Håstad [3], where they studied a promise version of (2k + 1)-SAT called $(2 + \epsilon)$ -SAT. They use a notion of polymorphism (which is the central concept in the algebraic theory of CSP) suitable for PCSPs. In [16], it was shown that the complexity of a PCSP is fully determined by its polymorphisms—in the sense 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 [5, 23], where it was shown that abstract properties of polymorphisms determine the complexity of PCSP.

The topological methods that we develop in this paper originate in topological combinatorics, specifically in Lovász's celebrated proof [57] that gives a tight lower bound on the chromatic number of Kneser graphs. We refer to [59] for an approachable introduction, and to [53] for an in-depth technical reference. The modern view of this method is to assign a topological space to a graph in such a way that combinatorial properties of the graph (e.g., the chromatic number) are influenced by topological properties of the resulting space (e.g., topological connectivity). An intermediate step in the construction of the topological space is to assign a certain abstract simplicial complex to a graph (we introduce these below). In our proof, we use so-called *homomorphism complexes* that give a simplicial structure on the set of homomorphisms between two graphs (or other structures). We remark that restricting those complexes to vertices and edges (so called 1-skeletons) gives graphs of homomorphisms which have been used in CSP research before (see, e.g., [20, 56]).

We remark that three earlier results on the complexity of approximate hypergraph coloring [2, 9, 28] were based on results from topological combinatorics using the Borsuk–Ulam theorem or similar [57, 59]. Their use of topology seems different from ours, and it remains to be seen whether they are all occurrences of a common pattern.

Topological methods and adjunction (including some specific cases that we use in this paper) have also been actively used in research around Hedetniemi's conjecture about the chromatic number of graph products [34, 61, 68, 69, 70, 71] (recently disproved by Shitov [67]). A few ideas in this paper are inspired by this line of research. A survey on adjunction and graph homomorphisms can be found in [32] (see also [33]), which also discusses several specific constructions that we use in this paper.

Our contributions. We first describe our methodological contributions related to topology and adjunction and then specific applications to promise graph homomorphism and approximate graph coloring. For simplicity, we will present the general theory for the case of graphs, which is what our applications are about; nevertheless, it generalizes immediately to arbitrary relational structures. We do not assume that the reader is familiar with topological combinatorics or algebraic topology and provide the necessary definitions and explanations here and in later sections.

It will be clear to the reader familiar with category theory that much of what we do in this paper can be naturally expressed in category-theoretic language. However, we avoid using this language, for the benefit of the readers less familiar with category theory.

Topological analysis of polymorphisms. As we mentioned before, the complexity of a problem $PCSP(\mathbf{H}, \mathbf{G})$ is completely determined by certain abstract properties of polymorphisms from \mathbf{H} to \mathbf{G} . Our first contribution is the introduction of topology as a tool to analyze polymorphisms. The basis for such analysis is the fact that the set of all homomorphisms from a graph \mathbf{H} to another graph \mathbf{G} can be made into an abstract simplicial complex denoted by $\mathsf{Hom}(\mathbf{H}, \mathbf{G})$.

An abstract simplicial complex K is a downwards closed family of nonempty subsets of a vertex set V(K)—subsets in the family are called *faces* (or *simplices*), their elements are *vertices*. A simplical complex describes a topological space: the *geometric realization* of K, denoted |K|, is the subspace of $\mathbb{R}^{V(K)}$ obtained by identifying vertices with affinely independent points and, for each face, adding to the space the convex hull of the vertices in the face. Thus a pair $\{u, v\} \in K$ becomes an edge, a triple (i.e., 3-element face) becomes a filled triangle, a quadruple becomes a filled tetrahedron, and so on. The resulting space can be analyzed by using algebraic topology and the algebraic structures (groups, rings) that it associates with a topological space.

The vertex set of the complex $\mathsf{Hom}(\mathbf{H}, \mathbf{G})$ is the set of all homomorphisms from \mathbf{H} to \mathbf{G} . A nonempty set $\{h_1, \ldots, h_\ell\}$ of such homomorphisms is a face if every function $h: V(H) \to V(G)$ satisfying $h(v) \in \{h_1(v), \ldots, h_\ell(v)\}$ for all v is a homomorphism $\mathbf{H} \to \mathbf{G}$. For example, if two homomorphisms h_1, h_2 differ at only one vertex $v \in V(H)$, then they are connected by a line in $|\mathsf{Hom}(\mathbf{H}, \mathbf{G})|$. Note the definition generalizes in a straightforward way from graphs to arbitrary relational structures.

There are several ways to use this notion for analysis of polymorphisms. One is to directly use the topological structure of $|\mathsf{Hom}(\mathbf{H}^n, \mathbf{G})|$ —for example, by looking at various connectivity properties in this space and asking when polymorphisms (as points in this space) belong to the same component. Another one, and this is what we use in the paper, goes as follows. Any (say, *n*-ary) polymorphism *f* from **H** to **G**, i.e., a homomorphism from \mathbf{H}^n to **G**, induces in a natural way a continuous map \tilde{f} from the space $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{H}^n)|$ to $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})|$, where \mathbf{K}_2 is the two-element clique. One can then obtain information about *f* from algebraic invariants of \tilde{f} .

As an important example, suppose that H, G are (possibly different) odd cycles. It is well known and not hard to check that $|\mathsf{Hom}(\mathbf{K}_2,\mathbf{H})|$ is topologically equivalent to the circle S^1 (we do this later in Example 3.4) and $|\mathsf{Hom}(\mathbf{K}_2,\mathbf{H}^n)|$ to the *n*-torus $\mathcal{T}^n = \mathcal{S}^1 \times \cdots \times \mathcal{S}^1$. A homomorphism f from **H** to **G** induces a continuous map \tilde{f} from \mathcal{S}^1 to \mathcal{S}^1 , and the main algebraic invariant of such a map is its *degree*, or winding number, which is an integer that intuitively measures how many times f wraps the domain circle around the range circle (and in which direction). The degree of f will be bounded because it arises from a discrete map f. Similarly, when analyzing a homomorphism f from \mathbf{H}^n to **G**, we study \tilde{f} , which is now a continuous map from \mathcal{T}^n to \mathcal{S}^1 . Each of the *n* variables of \tilde{f} corresponds to a circle in \mathcal{T}^n and thus to a degree of \tilde{f} restricted to that circle. We show that the number of variables whose degrees is nonzero is bounded, again because f arises from a discrete function f. In this way, we obtain that each polymorphism f has a bounded number of coordinates (independent of n) that are "important", or influential, and we can then use this information, together with the previously developed algebraic theory [5], to show that $PCSP(\mathbf{H}, \mathbf{G})$ is NP-hard.

Adjunction. We use symbols Λ, Γ for functions from the class of all (finite) graphs to itself. It will be convenient to write $\Lambda \mathbf{H}$ instead of $\Lambda(\mathbf{H})$ for the image of \mathbf{H} under Λ . The definitions and general properties again extend to all relational structures. Adjunction is a certain duality property between functions, best introduced with a concrete example.

Example 1.1. For a graph **G** and an odd integer k one can consider the following functions: $\Lambda_k \mathbf{G}$ is defined to be the graph obtained by subdividing each edge of **G**

into a path of k edges, and $\Gamma_k \mathbf{G}$ is the graph obtained by taking the kth power of the adjacency matrix (with zeroes on the diagonal; equivalently, the vertex set remains unchanged and two vertices are adjacent if and only if there is a walk between them of length exactly k in \mathbf{G}). For example, $\Gamma_3 \mathbf{G}$ has loops if \mathbf{G} has triangles.

Two functions Λ, Γ are called *adjoint* if

 $\Lambda \mathbf{H} \rightarrow \mathbf{G}$ if and only if $\mathbf{H} \rightarrow \Gamma \mathbf{G}$

for all graphs \mathbf{G}, \mathbf{H} . In this case Λ is also called the *left adjoint to* Γ , and Γ is *the* right adjoint to Λ . For example, it is well known and easy to check that Λ_k, Γ_k are adjoint, for any fixed odd k [32]. Adjoint functions are always monotone with respect to the homomorphism preorder, i.e., $\mathbf{H} \to \mathbf{G}$ implies both $\Lambda \mathbf{H} \to \Lambda \mathbf{G}$ and $\Gamma \mathbf{H} \to \Gamma \mathbf{G}$ (see Lemma 4.5).

Adjoint functions give us a way to reduce one PCSP to another. Indeed, consider any function Λ . We can always attempt to use it as a reduction between *some* PCSPs: if an instance graph **I** is guaranteed to be **H**-colorable, then Λ **I** is guaranteed to be Λ **H**-colorable if Λ is monotone. On the other hand if we find Λ **I** to be **G**-colorable, this may imply that **I** is **X**-colorable for some graph **X**. In such a case Λ would be a reduction from PCSP(**H**, **X**) to PCSP(Λ **H**, **G**). What is the best possible **X**? It is a graph **X** such that for any instance **I**, Λ **I** \rightarrow **G** holds if and only if **I** \rightarrow **X**. If such an **X** exists, it is essentially unique (since we just defined what homomorphisms **X** admits). The function that assigns to a graph **G** this best possible **X** is exactly the right adjoint to Λ . In this way, adjoints help us identify the best possible reduction a function gives, even though the proof that the reduction works might not need to mention the right adjoint.

Applications. Our applications of the above methodologies aim towards resolving the Brakensiek–Guruswami conjecture mentioned earlier:

CONJECTURE 1.2 (Brakensiek and Guruswami [16]). Let \mathbf{H} and \mathbf{G} be any nonbipartite loopless graphs with $\mathbf{H} \to \mathbf{G}$. Then $\mathrm{PCSP}(\mathbf{H}, \mathbf{G})$ is NP-hard.

We remark that the Hell–Nešetřil theorem [44] confirms Conjecture 1.2 for the case $\mathbf{H} = \mathbf{G}$. We also remark that Conjecture 1.2 covers all graphs: as discussed in section 2, if either \mathbf{H} or \mathbf{G} is bipartite or contains a loop, then PCSP(\mathbf{H}, \mathbf{G}) can be easily solved in polynomial time.

It is not hard to see that the conjecture is equivalent to the statement that $PCSP(\mathbf{C}_k, \mathbf{K}_c)$ is NP-hard for all $k \geq 3$ odd and $c \geq 3$, where \mathbf{C}_k is a cycle on k vertices and \mathbf{K}_c is a clique on c vertices. This is because we have a chain of homomorphisms

(1.1)
$$\cdots \rightarrow \mathbf{C}_k \rightarrow \cdots \rightarrow \mathbf{C}_5 \rightarrow \mathbf{C}_3 = \mathbf{K}_3 \rightarrow \mathbf{K}_4 \rightarrow \cdots \rightarrow \mathbf{K}_c \rightarrow \cdots$$

and, for each (\mathbf{H}, \mathbf{G}) with a homomorphism $\mathbf{H} \to \mathbf{G}$, the problem PCSP (\mathbf{H}, \mathbf{G}) admits a trivial reduction from PCSP $(\mathbf{C}_k, \mathbf{K}_c)$, where the promise is strengthened by requiring the input graph to be \mathbf{C}_k -colorable, for an odd cycle \mathbf{C}_k in \mathbf{H} , and the goal is weakened to that of finding a \mathbf{K}_c -coloring, where c is the chromatic number of \mathbf{G} (so we have $\mathbf{C}_k \to \mathbf{H}$ and $\mathbf{G} \to \mathbf{K}_c$).

The chain (1.1) has a natural middle point \mathbf{K}_3 . The right half corresponds to the classical approximate graph coloring: find a *c*-coloring of a 3-colorable graph. Our applications make progress on the right half and show hardness for all of the left half.

For the left half, we use the topological analysis of polymorphisms, as described above, to confirm Conjecture 1.2 for $\mathbf{G} = \mathbf{K}_3$.



FIG. 1.1. Examples of circular cliques.

THEOREM 1.3. $PCSP(\mathbf{H}, \mathbf{K}_3)$ is NP-hard for every nonbipartite 3-colorable **H**.

Equivalently, $PCSP(\mathbf{C}_k, \mathbf{K}_3)$ is NP-hard for all odd $k \geq 3$. We in fact prove a more general result which covers other graphs \mathbf{G} with similar topological properties to \mathbf{K}_3 —namely that $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})|$ maps to the circle S^1 via a \mathbb{Z}_2 -map (see Definition 3.7).

THEOREM 1.4. Let \mathbf{H}, \mathbf{G} be nonbipartite loopless graphs such that $\mathbf{H} \to \mathbf{G}$ and there is a \mathbb{Z}_2 -map from $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})|$ to S^1 . Then $\mathrm{PCSP}(\mathbf{H}, \mathbf{G})$ is NP-hard.

We give two specific classes of graphs **G** satisfying the assumptions of Theorem 1.4: certain circular cliques and all square-free graphs.

For positive integers p, q such that $p \ge 2q$, the *circular clique* $\mathbf{K}_{p/q}$ is the graph that has the same vertex set as the cycle \mathbf{C}_p and two vertices in it are connected by an edge if and only if they are at distance at least q in \mathbf{C}_p (see Figure 1.1). It well known that $\mathbf{K}_{n/1}$ is isomorphic to \mathbf{K}_n , $\mathbf{K}_{(2n+1)/n}$ is isomorphic to \mathbf{C}_{2n+1} , and that $\mathbf{K}_{p/q} \to \mathbf{K}_{p'/q'}$ if and only if $p/q \le p'/q'$ (see, e.g., Theorem 6.3 in [45]), thus circular cliques refine the homomorphism order (1.1) on odd cycles and cliques described above. The *circular chromatic number* of $\mathbf{G}, \chi_c(\mathbf{G})$, is defined as $\inf\{p/q \mid \mathbf{G} \to \mathbf{K}_{p/q}\}$. Note that we always have $\chi(\mathbf{G}) = \lceil \chi_c(\mathbf{G}) \rceil$ and also $\chi_c(\mathbf{G}) \le 2 + \frac{1}{n}$ if and only if $\mathbf{G} \to \mathbf{C}_{2n+1}$.

The fact that circular cliques $\mathbf{K}_{p/q}$ with 2 < p/q < 4 satisfy the topological condition of Theorem 1.4 is folklore, though we prove it later for completeness. The theorem in this case can be viewed as NP-hardness of coloring $(2+\epsilon)$ -colorable graphs with $4 - \epsilon$ colors:

COROLLARY 1.5. $PCSP(\mathbf{K}_{p/q}, \mathbf{K}_{p'/q'})$ is NP-hard for all $2 < p/q \le p'/q' < 4$.

A graph is said to be *square-free* if it does not contain the 4-cycle C_4 as a subgraph. This includes all graphs of girth at least 5 and thus graphs of arbitrarily high chromatic number. Again, it will be a simple observation that square-free graphs satisfy the condition of Theorem 1.4. Therefore, we confirm Conjecture 1.2 for square-free graphs G.

COROLLARY 1.6. $PCSP(\mathbf{H}, \mathbf{G})$ is NP-hard for all nonbipartite loopless graphs \mathbf{H}, \mathbf{G} such that $\mathbf{H} \to \mathbf{G}$ and \mathbf{G} is square-free.

Since the key assumption of Theorem 1.4 is topological, this raises a question of whether topology is in some sense necessary to settle Conjecture 1.2. Using adjointness, we argue that it is indeed the case, proving the following (see Theorem 4.32 for a formal statement).

THEOREM 1.7 (informal). For any graph \mathbf{G} , the property that $PCSP(\mathbf{H}, \mathbf{G})$ is NP-hard for all nonbipartite \mathbf{G} -colorable graphs \mathbf{H} depends only on the topology (and \mathbb{Z}_2 -action) of $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})|$.

Returning to the right half of the chain (1.1) (the classical coloring problem), we first show that, to prove NP-hardness of c-coloring k-colorable graphs for all constants $c \ge k \ge 3$, it is enough to prove it for any fixed k (and all $c \ge k$).

THEOREM 1.8. Suppose there is an integer k such that $PCSP(\mathbf{K}_k, \mathbf{K}_c)$ is NP-hard for all $c \geq k$. Then $PCSP(\mathbf{K}_3, \mathbf{K}_c)$ is NP-hard for all $c \geq 3$.

Following the reasoning in [40], the above theorem implies NP-hardness of all problems $PCSP(\mathbf{K}_k, \mathbf{K}_c)$ with $c \ge k \ge 3$ if the *d*-to-1 conjecture of Khot holds for any fixed $d \ge 2$. (The paper [40] used an earlier version of Theorem 1.8 with 4 in place of 3.)

Furthermore, we strengthen the best known asymptotic hardness: Huang [46] showed that $\text{PCSP}(\mathbf{K}_k, \mathbf{K}_c)$ is NP-hard for all sufficiently large k and $c = 2^{\Omega(k^{1/3})}$. We improve this in two ways, using Huang's result as a black box. First, we improve the asymptotics from subexponential $c = 2^{\Omega(k^{1/3})}$ to single-exponential $c = \binom{k}{\lfloor k/2 \rfloor} \in \Theta(2^k/\sqrt{k})$. Second, we show the claim holds for k starting as low as 4.

THEOREM 1.9. For all $k \ge 4$ and $c = \binom{k}{\lfloor k/2 \rfloor} - 1$, $PCSP(\mathbf{K}_k, \mathbf{K}_c)$ is NP-hard.

In comparison, the previous best result relevant for all integers k was obtained in [5] where NP-hardness of PCSP($\mathbf{K}_k, \mathbf{K}_{2k-1}$) is proved for all $k \geq 3$. For k = 3, 4 we obtain no new results and for k = 5 the two bounds coincide: $\binom{k}{\lfloor k/2 \rfloor} - 1 = 9 = 2k - 1$. However, already for k = 6 we improve the bound from 2k - 1 = 11 to $\binom{k}{\lfloor k/2 \rfloor} - 1 = 19$, and, for larger k, the improvement is even more dramatic.

The organization of the paper. Section 2 briefly describes the algebraic framework of [5]: minions (sets of polymorphisms of a PCSP), minion homomorphism (which provide log-space reductions between corresponding problems), and a condition on minions that guarantees NP-hardness. Section 3 details the topological method and its application, Theorem 1.4. The bulk of its content is devoted to expounding standard definitions with examples and then proving these definitions behave well when identifying variables of polymorphisms. Section 4 introduces adjunction in a wider context, in particular, relating it to gadget reductions and minion homomorphisms. Adjoint functions that give reductions for approximate graph coloring are presented in section 4.3. Finally section 4.4 uses another adjoint function to prove Theorem 1.7: that whether a graph \mathbf{G} satisfies the Brakensiek–Guruswami conjecture for all \mathbf{H} depends only on the topology of \mathbf{G} .

2. Preliminaries.

2.1. Promise graph homomorphism problems. The approximate graph coloring 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 digraphs. For general definitions, see, e.g., [5].

A digraph **H** is a pair $\mathbf{H} = (V(H), E(H))$, where V(H) is a set of vertices and $E(H) \subseteq \{(u, v) \mid u, v \in V(H)\}$ is a set of (directed) edges. Unless stated otherwise, our digraphs are finite and can have loops. We view undirected graphs as digraphs where each (nonloop) edge is present in both directions.

Copyright \bigodot by SIAM. Unauthorized reproduction of this article is prohibited.

DEFINITION 2.1. A homomorphism from a digraph $\mathbf{H} = (V(H), E(H))$ to another digraph $\mathbf{G} = (V(G), E(G))$ is a map $h: V(H) \to V(G)$ such that $(h(u), h(v)) \in E(G)$ for every $(u, v) \in E(H)$. In this case we write $h: \mathbf{H} \to \mathbf{G}$, and simply $\mathbf{H} \to \mathbf{G}$ to indicate that a homomorphism exists.

We now define formally the promise (di)graph homomorphism problem.

DEFINITION 2.2. Fix two digraphs **H** and **G** such that $\mathbf{H} \rightarrow \mathbf{G}$.

- The search variant of PCSP(H,G) is, given an input digraph I that maps homomorphically to H, find a homomorphism h: I→G.
- The decision variant of $PCSP(\mathbf{H}, \mathbf{G})$ requires, given an input digraph \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.

We remark that the (decision) problem $PCSP(\mathbf{H}, \mathbf{H})$ is nothing else but the constraint satisfaction problem $CSP(\mathbf{H})$, also known as **H**-coloring.

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 $PCSP(\mathbf{H}, \mathbf{G})$.

It is obvious that if at least one of \mathbf{H}, \mathbf{G} is undirected and bipartite, then the problem can be solved in polynomial time by using an algorithm for 2-coloring. If one of the graphs contains a loop, the problem is trivial. Recall that Brakensiek and Guruswami conjectured (see Conjecture 1.2) that, for undirected graphs, the problem is NP-hard in all the other cases.

All applications in this paper concern undirected graphs, even though some proofs use digraphs. We remark that, as shown in Theorem F.3 of the arXiv version of [16] (generalizing the corresponding result for CSPs [30]), a complexity classification of all problems $PCSP(\mathbf{H}, \mathbf{G})$ for digraphs is equivalent to such a classification for all PCSPs (for arbitrary relational structures).

Two (di)graphs \mathbf{H}_1 and \mathbf{H}_2 are called *homomorphically equivalent* if both $\mathbf{H}_1 \rightarrow \mathbf{H}_2$ and $\mathbf{H}_2 \rightarrow \mathbf{H}_1$. The binary relation $\mathbf{H}_1 \rightarrow \mathbf{H}_2$ defines a preorder on the class of all digraphs (or all graphs), called the *homomorphism preorder*. We will use this preorder in section 4.

We also define digraph powers, which are essential for the notion of polymorphisms.

DEFINITION 2.3. The nth direct (or tensor) power of a digraph \mathbf{H} is the digraph \mathbf{H}^n whose vertices are all n-tuples of vertices of \mathbf{H} (i.e., $V(H^n) = V(H)^n$), and whose edges are defined as follows: we have an edge from (u_1, \ldots, u_n) to (v_1, \ldots, v_n) in \mathbf{H}^n if and only if (u_i, v_i) is an edge of \mathbf{H} for all $i \in \{1, \ldots, n\}$.

2.2. Polymorphisms. We use the notions of polymorphisms [3, 16], minions, and minion homomorphisms [5, 23]. We introduce these notions in the special case of digraphs below. General definitions and more insights can be found in [5, 7].

DEFINITION 2.4. An n-ary polymorphism from a digraph **H** to a digraph **G** is a homomorphism from \mathbf{H}^n to **G**. To spell this out, it is a mapping $f: V(H)^n \to V(G)$ such that, for all tuples $(u_1, v_1), \ldots, (u_n, v_n)$ of edges of **H**, we have

$$(f(u_1,\ldots,u_n),f(v_1,\ldots,v_n)) \in E(G).$$

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

Example 2.5. The *n*-ary polymorphisms from a digraph **H** to the *k*-clique \mathbf{K}_k are the *k*-colorings of \mathbf{H}^n .

The set of all polymorphisms between any two digraphs has a certain algebraic structure, which we now describe. We denote by [n] the set $\{1, 2, ..., n\}$.

DEFINITION 2.6. An n-ary function $f: A^n \to B$ is called a minor of an m-ary function $g: A^m \to B$ if there is a map $\pi: [m] \to [n]$ such that

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

for all $x_1, \ldots, x_n \in A$. In this case, we write $f = g^{\pi}$.

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 2.7. For sets A, B, let $\mathcal{O}(A, B) = \{f : A^n \to B \mid n \geq 1\}$. A (function) minion \mathcal{M} on a pair of sets (A, B) is a nonempty 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} .

It is easy to see that $Pol(\mathbf{H}, \mathbf{G})$ is a minion whenever $\mathbf{H} \rightarrow \mathbf{G}$.

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

DEFINITION 2.8. A coordinate *i* of a function $f: A^n \to B$ is called essential if *f* depends on *it*, that is, if there exist a_1, \ldots, a_n and b_i in A such that

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

A coordinate of f that is not essential is called inessential.

DEFINITION 2.9. A minion \mathscr{M} is said to have essential arity at most k, if each function $f \in \mathscr{M}$ 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.

Example 2.10. It is well known (see, e.g., [37]), and not hard to check, that the minion $\text{Pol}(\mathbf{K}_3, \mathbf{K}_3)$ has essential arity at most 1. However for any odd k > 3, the minion $\text{Pol}(\mathbf{C}_k, \mathbf{K}_3)$ does not have bounded essential arity. Indeed, fix a homomorphism $h: \mathbf{C}_k \to \mathbf{K}_3$ 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 $f \in Pol(\mathbf{C}_k, \mathbf{K}_3)$. By using Definition 2.8 with $a_1 = \cdots = a_n = 1$ and $b_i = 0$, one can verify that every coordinate *i* of *f* is essential.

DEFINITION 2.11. Let \mathscr{M} and \mathscr{N} be two minions (not necessarily on the same pairs of sets). A mapping $\xi \colon \mathscr{M} \to \mathscr{N}$ 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 $\pi: \{1, \ldots, m\} \to \{1, \ldots, n\}$ and each $g \in \mathscr{M}^{(m)}$ we have $\xi(g)^{\pi} = \xi(g^{\pi})$, i.e.,

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

Copyright (c) by SIAM. Unauthorized reproduction of this article is prohibited.

We refer to [5, Example 2.22] for examples of minion homomorphisms.

Our proof of Theorem 1.4 is based on the following result. It is a special case of a result in [5] (that generalized [3, 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 2.12 ([5, Proposition 5.15]). Let \mathbf{H}, \mathbf{G} be digraphs such that $\mathbf{H} \to \mathbf{G}$. Assume that there exists a minion homomorphism $\xi \colon \operatorname{Pol}(\mathbf{H}, \mathbf{G}) \to \mathscr{M}$ for some minion \mathscr{M} on a pair of (possibly infinite) sets such that \mathscr{M} has bounded essential arity and does not contain a constant function (i.e., a function without essential variables). Then $\operatorname{PCSP}(\mathbf{H}, \mathbf{G})$ is NP-hard.

To prove Theorem 1.4, we will use Theorem 2.12 with the minion $\mathcal{M} = \mathscr{Z}_{\leq N}$ for some constant N > 0. The set $\mathscr{Z}_{\leq N}$ is defined to consist of all linear functions f on \mathbb{Z} of the form

$$f(x_1,\ldots,x_n) = \sum_{i=1}^n c_i x_i$$

for some $c_i \in \mathbb{Z}$ such that $\sum_{i=1}^n c_i$ is odd and $\sum_{i=1}^n |c_i| \leq N$. It easy to see that $\mathscr{Z}_{\leq N}$ is indeed a minion and that all functions in it have between 1 and N nonzero coefficients, meaning that it has bounded essential arity and contains no constant function.

3. Topology. All graphs in this section are assumed to be undirected and loop-less.

3.1. Simplicial complexes. An *(abstract) simplicial complex* is a family of nonempty sets K that is downwards closed, i.e., if $\sigma_1 \in \mathsf{K}$, $\sigma_2 \neq \emptyset$, and $\sigma_2 \subseteq \sigma_1$, then $\sigma_2 \in \mathsf{K}$. Each $\sigma \in \mathsf{K}$ is called a *face*. The elements in these sets are *vertices* of K. We denote the set of all vertices of K by $V(\mathsf{K})$, i.e., $V(\mathsf{K}) := \bigcup_{\sigma \in \mathsf{K}} \sigma$. A *simplicial map* between complexes K and K' is a function $f: V(\mathsf{K}) \to V(\mathsf{K}')$ that preserves faces, i.e., if $\sigma \in \mathsf{K}$, then $f(\sigma) := \{f(v) \mid v \in \sigma\} \in \mathsf{K}'$. Two simplicial complexes K and K' are *isomorphic* if there are simplicial maps $\alpha: \mathsf{K} \to \mathsf{K}'$ and $\beta: \mathsf{K}' \to \mathsf{K}$ such that both $\alpha\beta$ and $\beta\alpha$ are identity maps.

We will use the following notion of a product of simplicial complexes (see also [61, section 2.2] and [53, Definition 4.25]).

DEFINITION 3.1. Let K_1, \ldots, K_n be simplicial complexes. We define the product $K_1 \times \cdots \times K_n$ to be the simplicial complex with vertices

$$V(\mathsf{K}_1 \times \cdots \times \mathsf{K}_n) = V(\mathsf{K}_1) \times \cdots \times V(\mathsf{K}_n),$$

so that $\sigma \subseteq V(\mathsf{K}_1 \times \cdots \times \mathsf{K}_n)$ is a face if there are faces $\sigma_1 \in \mathsf{K}_1, \ldots, \sigma_n \in \mathsf{K}_n$ such that $\sigma \subseteq \sigma_1 \times \cdots \times \sigma_n$.

3.1.1. From graphs to simplicial complexes. As mentioned before, there are several ways to assign a simplicial complex to a graph. For our use, the most convenient is the homomorphism complex. Our definition of this complex is slightly different from that in [4, 53], but the difference is superficial (as we explain in Appendix A). The vertices of such a complex are homomorphisms, while faces are determined by multihomomorphisms defined below.

DEFINITION 3.2. A multihomomorphism from **K** to **G** is a mapping $f: V(K) \rightarrow 2^{V(G)}$ such that, for each $(u, v) \in E(K)$, we have $f(u) \times f(v) \subseteq E(G)$.

Copyright (c) by SIAM. Unauthorized reproduction of this article is prohibited.

DEFINITION 3.3. Let **K** and **G** be two graphs. We define a simplicial complex Hom(**K**, **G**) as follows. Its vertices are homomorphisms from **K** to **G**, and $\sigma = \{f_1, \ldots, f_\ell\}$ is a face if the mapping $u \mapsto \{f_1(u), \ldots, f_\ell(u)\}$ is a multihomomorphism from **K** to **G**.

We work almost exclusively with complexes $\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})$, where \mathbf{K}_2 is the twoelement clique. Such complexes (with our definition) appeared before, e.g., in [60], where they are called box complexes (which is not the traditional use of this name) and in [58], where these complexes appear under the name $\mathsf{B}_{\mathsf{edge}}(\mathbf{G})$. The complex $\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})$ can be also described in the following way. The vertices of $\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})$ are all (oriented) edges of \mathbf{G} . The faces are directed bipartite subgraphs that can be extended to a complete directed bipartite subgraph of \mathbf{G} (with all edges directed from one part to the other); more precisely, σ is a face if there are $U, V \subseteq V(G)$ such that $\sigma \subseteq U \times V \subseteq E(G)$. The complexes $\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})$ have an additional structure obtained from the automorphism of \mathbf{K}_2 that switches the two vertices. The group \mathbb{Z}_2 then acts on the vertices of $\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})$ by reversing the direction of edges, i.e., -(a, b) = (b, a).

Example 3.4. Let us consider the complex $\operatorname{Hom}(\mathbf{K}_2, \mathbf{C}_k)$. Its vertices are all oriented edges of the k-cycle which means pairs of the form (i, i + 1) and (i + 1, i) where the addition is considered modulo k. It is not hard to see that the only directed complete bipartite subgraphs of \mathbf{C}_k are either two outgoing edges from a single vertex, or two incoming edges to a single vertex. The only nontrivial faces of $\operatorname{Hom}(\mathbf{K}_2, \mathbf{C}_k)$ are therefore of the form $\{(i - 1, i), (i + 1, i)\}$ or $\{(i, i - 1), (i, i + 1)\}$. The resulting complex can be drawn as a graph (see Figure 3.1 for such a drawing of $\operatorname{Hom}(\mathbf{K}_2, \mathbf{C}_5)$). The exact structure depends on the parity of k. If k is odd, the complex is a single 2k-cycle, where (i, j) is opposite to (j, i). The \mathbb{Z}_2 -action acts as the central reflection. If k is even, the complex consists of two disjoint k-cycles such that one contains all edges of the form $(2i, 2i \pm 1)$ and the other all edges of the from $(2i \pm 1, 2i)$. The \mathbb{Z}_2 -action in this case switches the two parts.

Example 3.5. A slightly more complicated example is $\text{Hom}(\mathbf{K}_2, \mathbf{K}_4)$. See Figure 3.2 for graphical representations of this complex. There are two types of maximal directed complete bipartite subgraphs of \mathbf{K}_4 : either all three incoming/outgoing edges of a single vertex, or 4 directed edges from a two-element subset of \mathbf{K}_4 to its complement. These, and all their nonempty subsets, are the faces of $\text{Hom}(\mathbf{K}_2, \mathbf{K}_4)$. In the pictures, the incoming/outgoing edges correspond to the triangular faces, and the



FIG. 3.1. Representations of $Hom(\mathbf{K}_2, \mathbf{C}_5)$ and $Hom(\mathbf{K}_2, \mathbf{C}_6)$.



FIG. 3.2. Two representations of $Hom(K_2, K_4)$. The tetragons in both pictures represent tetrahedrons with the same vertices.

faces containing 4 edges correspond to tetrahedrons that are represented as tetragons. Naturally, all subsets of these tetragons are also faces, nevertheless they are omitted from the picture for better readability. Also note that the outer face of the left diagram forms such a tetrahedron (corresponding to the bipartite subgraph $\{1,3\} \times \{0,2\}$). The \mathbb{Z}_2 -symmetry of this complex is given by reversing edges; this corresponds to the antipodality on the spherical representation.

DEFINITION 3.6. $A \mathbb{Z}_2$ -(simplicial) complex is a simplicial complex K with a simplicial map $-: \mathsf{K} \to \mathsf{K}$ such that -(-v) = v for each $v \in V(\mathsf{K})$. We also write $-\sigma$ for the face $\{-v \mid v \in \sigma\}$. We say that a simplicial map f from one \mathbb{Z}_2 -complex K to another \mathbb{Z}_2 -complex K' is a \mathbb{Z}_2 -simplicial map, if f(-v) = -f(v) for each $v \in V(\mathsf{K})$ (note that the first - is taken in K, while the second is taken in K').

The map $v \mapsto -v$ can be also viewed as an action of the group \mathbb{Z}_2 on K by simplicial maps. We remark that a product $K_1 \times \cdots \times K_n$ of \mathbb{Z}_2 -complexes is also a \mathbb{Z}_2 -complex with the action defined componentwise, as $-(v_1, \ldots, v_n) = (-v_1, \ldots, -v_n)$. Every graph homomorphism $f: \mathbf{H} \to \mathbf{G}$ induces a \mathbb{Z}_2 -simplicial map $f': \operatorname{Hom}(\mathbf{K}_2, \mathbf{H}) \to$ $\operatorname{Hom}(\mathbf{K}_2, \mathbf{G})$ defined by f'((a, b)) = (f(a), f(b)).¹

3.2. Topological spaces. The spaces assigned to \mathbb{Z}_2 -complexes inherit the \mathbb{Z}_2 symmetry.

DEFINITION 3.7. A \mathbb{Z}_2 -space is a topological space \mathcal{X} with a distinguished continuous function $-: \mathcal{X} \to \mathcal{X}$ such that -(-x) = x for each x. A \mathbb{Z}_2 -map between two \mathbb{Z}_2 -spaces \mathcal{X} and \mathcal{Y} is a continuous function $f: \mathcal{X} \to \mathcal{Y}$ which preserves the action of \mathbb{Z}_2 , i.e., f(-x) = -f(x) for each $x \in \mathcal{X}$ (note that the first - is taken in \mathcal{X} , while the second is taken in \mathcal{Y}).

As is the case for \mathbb{Z}_2 -complexes, \mathbb{Z}_2 -spaces are topological spaces with an action of the group \mathbb{Z}_2 by continuous functions.

Example 3.8. Prime examples of \mathbb{Z}_2 -spaces are spheres: we define S^n as a subspace of \mathbb{R}^{n+1} consisting of all unit vectors, i.e., $S^n = \{(x_1, \ldots, x_{n+1}) \in \mathbb{R}^{n+1} \mid x_1^2 + \cdots + x_{n+1}^2 = 1\}$, with *antipodality* as the chosen \mathbb{Z}_2 -action, i.e., $-(x_1, \ldots, x_{n+1}) =$

¹This is an instance of a more general fact that f induces a simplical map $f': \operatorname{Hom}(\mathbf{K}, \mathbf{H}) \to \operatorname{Hom}(\mathbf{K}, \mathbf{G})$ for each \mathbf{K} . This map is defined by $f'(g): x \mapsto f(g(x))$.

 $(-x_1,\ldots,-x_{n+1})$. Other common \mathbb{Z}_2 -spaces are toruses. An *n*-torus \mathcal{T}^n is defined as the *n*th power $\mathcal{S}^1 \times \cdots \times \mathcal{S}^1$, and is therefore naturally equipped with a \mathbb{Z}_2 -action defined to act coordinatewise.

A \mathbb{Z}_2 -complex K is free if $-\sigma \neq \sigma$ for each $\sigma \in K$ (equivalently, $\{-v,v\} \notin K$ for all vertices v of K). Note that, for a loopless undirected graph **G**, the complex Hom(\mathbf{K}_2, \mathbf{G}) is always a free \mathbb{Z}_2 -complex. To ease a technical annoyance in the proofs below, we rephrase the definition of a geometric realization (see also [53, Definition 2.27]) of a free \mathbb{Z}_2 -complex.

DEFINITION 3.9. Let K be a free \mathbb{Z}_2 -simplicial complex. Let $v_1, -v_1, \ldots, v_n, -v_n$ be all vertices of K. We define $|\mathsf{K}|$, a geometric realization of K, as a subspace of \mathbb{R}^n . First, we identify the canonical unit vectors with v_1, \ldots, v_n , so that $v_1 = (1, 0, \ldots, 0)$, etc., and $-v_1, \ldots, -v_n$ with their opposites, so $-v_1 = (-1, 0, \ldots, 0)$, etc. Second, for each face $\sigma \subseteq V(\mathsf{K})$, we define $\Delta^{\sigma} \subseteq \mathbb{R}^n$ to be the convex hull of σ , i.e., $\Delta^{\sigma} = \{\sum_{v \in \sigma} \lambda_v v \mid \sum_{v \in \sigma} \lambda_v = 1, \lambda_v \geq 0\}$. Finally, we set

$$|\mathsf{K}| = \bigcup_{\sigma \in \mathsf{K}} \Delta^{\sigma} = \left\{ \sum_{v \in \sigma} \lambda_v v \mid \sigma \in \mathsf{K}, \sum_{v \in \sigma} \lambda_v = 1, \lambda_v \ge 0 \right\}$$

The action of \mathbb{Z}_2 on $|\mathsf{K}|$ maps a point $\sum_{v \in \sigma} \lambda_v v$ to the point $\sum_{-v \in -\sigma} \lambda_v(-v)$ which can be equivalently described as reversing the sign of a vector, i.e., as $-(x_1, \ldots, x_n) = (-x_1, \ldots, -x_n)$.

With the above definition, we can view $V(\mathsf{K})$ as a subset of $|\mathsf{K}|$ —this will be useful in the technical proofs below. Also note that -v has two meanings that result in the same object: either it is a \mathbb{Z}_2 -counterpart of $v \in V(\mathsf{K})$, or the opposite vector to $v \in |\mathsf{K}|$. Note that the geometric realization of a free \mathbb{Z}_2 -complex is a *free* \mathbb{Z}_2 -*space*, i.e., a \mathbb{Z}_2 -space \mathcal{X} such that $-x \neq x$ for all $x \in \mathcal{X}$.

To express abstractly what it means for two \mathbb{Z}_2 -spaces to be the same, we use the notion of \mathbb{Z}_2 -homeomorphism which is an analogue of the notion of homeomorphism. We remark that this is a strong notion of equivalence of topological spaces, akin to isomorphism, and that we will also use weaker notions of topological equivalence (see also Appendix A).

DEFINITION 3.10. Two \mathbb{Z}_2 -topological spaces \mathcal{X} and \mathcal{Y} are \mathbb{Z}_2 -homeomorphic if there are \mathbb{Z}_2 -maps $f: \mathcal{X} \to \mathcal{Y}$ and $g: \mathcal{Y} \to \mathcal{X}$ such that fg is the identity on \mathcal{Y} and gfis the identity on \mathcal{X} .

Example 3.11. It is not hard to see that the geometric representation of the homomorphism complex $\operatorname{Hom}(\mathbf{K}_2, \mathbf{C}_k)$ of an odd cycle \mathbf{C}_k is \mathbb{Z}_2 -homeomorphic to \mathcal{S}^1 (see Figure 3.3). Let us define one such \mathbb{Z}_2 -homeomorphism $f: |\operatorname{Hom}(\mathbf{K}_2, \mathbf{C}_k)| \to \mathcal{S}^1$. Choose k points on the circle in a regular pattern. Let us denote these vectors x_1, \ldots, x_k . Note that since k is odd, $-x_i \notin \{x_1, \ldots, x_k\}$ for all i. We first define a map $f_0: |\operatorname{Hom}(\mathbf{K}_2, \mathbf{C}_k)| \to \mathbb{R}^2$ as follows: $f_0(v) = x_b - x_a$ for $v \in V(\mathsf{K}), v = (a, b)$, and extend it linearly. Note that the image of $|\operatorname{Hom}(\mathbf{K}_2, \mathbf{C}_k)|$ forms a regular 2k-gon centered in the origin. We project the polygon onto \mathcal{S}^1 by putting $f(x) = f_0(x)/|f_0(x)|$. It is clear that f is continuous and f(-x) = -f(x), and therefore it is a \mathbb{Z}_2 -map. It is also not hard to see that it is 1-to-1 and therefore invertible, and that the inverse is a \mathbb{Z}_2 -map.

Remark 3.12. While there is always a continuous function between two topological spaces \mathcal{X} and \mathcal{Y} (simply map everything to a single point), there might not be a

Copyright (C) by SIAM. Unauthorized reproduction of this article is prohibited.



FIG. 3.3. A representation of $\text{Hom}(\mathbf{K}_2, \mathbf{C}_5)$ as S^1 and $\text{Hom}(\mathbf{K}_2, \mathbf{C}_5^2)$ as \mathcal{T}^2 .

 \mathbb{Z}_2 -map between two \mathbb{Z}_2 -spaces. In particular, the Borsuk–Ulam theorem [12] (see also [59]) states that there is no \mathbb{Z}_2 -map from a sphere \mathcal{S}^m to a sphere \mathcal{S}^n of smaller dimension (i.e., if m > n).

Every \mathbb{Z}_2 -simplicial map $f: \mathsf{K} \to \mathsf{K}'$ induces a \mathbb{Z}_2 -map $|f|: |\mathsf{K}| \to |\mathsf{K}'|$ defined as a piecewise linear extension of f:

$$|f|: \sum_{v \in \sigma} \lambda_v v \mapsto \sum_{v \in \sigma} \lambda_v f(v).$$

(Here, we use that $v \in V(\mathsf{K})$ is also a point in $|\mathsf{K}|$, and therefore a vector in \mathbb{R}^n .) Consequently, every graph homomorphism $\mathbf{H} \to \mathbf{G}$ induces a \mathbb{Z}_2 -map from $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{H})|$ to $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})|$.

3.2.1. The fundamental group. We briefly recall the definition of the fundamental group assigned to a topological space \mathcal{X} , denoted $\pi_1(\mathcal{X})$. For more details, see [43, Chapter 1]. The elements of the group are *homotopy classes* of maps $f: \mathcal{S}^1 \to \mathcal{X}$ defined as follows. Intuitively, two maps are *homotopic* if one can be continuously transformed into the other.

DEFINITION 3.13. We say that two continuous maps $f, g: \mathcal{X} \to \mathcal{Y}$ are homotopic if there is a continuous map $h: \mathcal{X} \times [0,1] \to \mathcal{Y}$ such that h(x,0) = f(x) and h(x,1) = g(x) for each $x \in \mathcal{X}$. Any such map h is called a homotopy. The homotopy class of $f: \mathcal{X} \to \mathcal{Y}$ is the set of all continuous maps $g: \mathcal{X} \to \mathcal{Y}$ that are homotopic to f. We denote such a class by $[f]^2$.

Formally, the fundamental group is defined relative to a point $x_0 \in \mathcal{X}$, but the choice of the point is irrelevant if the space \mathcal{X} is path connected (see [43, Proposition 1.5]), i.e., if any two points in \mathcal{X} are connected by a path. Fix one such choice $x_0 \in \mathcal{X}$. The elements of $\pi_1(\mathcal{X})$ are all homotopy classes of maps $\ell \colon \mathcal{S}^1 \to \mathcal{X}$ such that $\ell((1,0)) = x_0$. The group operation is given by so-called *loop composition*: seeing maps $\ell_1, \ell_2 \colon \mathcal{S}^1 \to \mathcal{X}$ as closed walks originating in x_0 , the product $\ell_1 \cdot \ell_2$ is the closed walk that follows first ℓ_1 and then ℓ_2 . While this product is not a group operation as is, it induces a group operation on the homotopy classes defined as $[\ell_1] \cdot [\ell_2] = [\ell_1 \cdot \ell_2]$ (see [43, Proposition 1.3] for a proof).

Any map $f: \mathcal{X} \to \mathcal{Y}$ induces a group homomorphism $f_*: \pi_1(\mathcal{X}) \to \pi_1(\mathcal{Y})$ defined by $f_*([\ell]) = [f\ell]$ for each $\ell: \mathcal{S}^1 \to \mathcal{X}$.

²We use the notation [*] both for sets $\{1, \ldots, n\}$ and for homotopy classes; the meaning will always be clear from the context.

The fundamental groups of many spaces are described in the literature. For example, we have the following.

LEMMA 3.14 ([43, Theorem 1.7]). $\pi_1(\mathcal{S}^1)$ is isomorphic to \mathbb{Z} .

We also define \mathbb{Z}_2 -homotopy which is a strengthening of homotopy, restricting it to \mathbb{Z}_2 -maps.

DEFINITION 3.15. Let $f, g: \mathcal{X} \to \mathcal{Y}$ be \mathbb{Z}_2 -maps. A homotopy h from f to g is a \mathbb{Z}_2 -homotopy if the map $h_t: x \mapsto h(x,t)$ is a \mathbb{Z}_2 -map for each $t \in [0,1]$. We say that f and g are \mathbb{Z}_2 -homotopic, if there is a \mathbb{Z}_2 -homotopy between them.

3.3. Polymorphisms of complexes, spaces, and groups. A polymorphism from one graph to another is defined as a homomorphism from a power. In the same way, we can define polymorphisms of any objects as long as we have a notion of a homomorphism and of a power.

DEFINITION 3.16.

1. Let K, K' be two \mathbb{Z}_2 -simplicial complexes. An n-ary polymorphism from K to K' is a \mathbb{Z}_2 -simplicial map from the nth power of K to K', i.e., $f: V(K)^n \to V(K')$ such that $f(-v_1, \ldots, -v_n) = -f(v_1, \ldots, v_n)$ for all $v_i \in V(K)$ and

$$f(\sigma_1 \times \cdots \times \sigma_n) = \{f(a_1, \dots, a_n) \colon a_i \in \sigma_i\}$$
 is in K'

for all $\sigma_1, \ldots, \sigma_n \in K$. We denote by Pol(K, K') the set of all polymorphisms from K to K'.

- 2. Let \mathcal{X}, \mathcal{Y} be two \mathbb{Z}_2 -spaces. An n-ary polymorphism from \mathcal{X} to \mathcal{Y} is a \mathbb{Z}_2 -map from \mathcal{X}^n to \mathcal{Y} , i.e., a continuous map $f: \mathcal{X}^n \to \mathcal{Y}$ such that $f(-x_1, \ldots, -x_n) = -f(x_1, \ldots, x_n)$ for all $x_i \in \mathcal{X}$. Again, $\operatorname{Pol}(\mathcal{X}, \mathcal{Y})$ denotes the set of all polymorphisms from \mathcal{X} to \mathcal{Y} .
- Let H, G be two groups. An n-ary polymorphism from H to G is a group homomorphism from Hⁿ to G, i.e., a mapping f: Hⁿ→G such that

 $f(g_1 \cdot h_1, \dots, g_n \cdot h_n) = f(g_1, \dots, g_n) \cdot f(h_1, \dots, h_n)$

for all $g_i, h_i \in H$. We denote the set of all polymorphisms from \mathbb{H} to \mathbb{G} by $\operatorname{Pol}(\mathbb{H}, \mathbb{G})$.

In all the cases above, it is easy to see that polymorphisms are closed under taking minors, and therefore Pol(-, -) is always a minion. This allows us to talk about minion homomorphisms between minions of polymorphisms of different objects (graphs, simplicial complexes, topological spaces, or groups).

Example 3.17. By definition, $\operatorname{Pol}(\mathbb{Z},\mathbb{Z})$ consists of all group homomorphisms from \mathbb{Z}^n to \mathbb{Z} for all n > 0. It is straightforward to check that such an *n*-ary polymorphism in $\operatorname{Pol}(\mathbb{Z},\mathbb{Z})$ is a linear function, i.e., of the form $(x_1,\ldots,x_n) \mapsto \sum_{i=1}^n c_i x_i$ for some $c_1,\ldots,c_n \in \mathbb{Z}$ and, conversely, any such function is a group homomorphism from \mathbb{Z}^n to \mathbb{Z} .

3.4. Proofs of Theorems 1.3 and 1.4. We recall the statement of Theorem 1.4.

THEOREM 1.4. Let \mathbf{H}, \mathbf{G} be nonbipartite loopless graphs such that $\mathbf{H} \to \mathbf{G}$, and there is a \mathbb{Z}_2 -map from $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})|$ to S^1 . Then $\mathrm{PCSP}(\mathbf{H}, \mathbf{G})$ is NP-hard. Theorem 1.3 is a direct corollary of the above and Example 3.11. In the rest of the section we prove Theorem 1.4 by using Theorem 2.12. We show that there is a minion homomorphism from $\text{Pol}(\mathbf{H}, \mathbf{G})$ to the minion $\mathscr{Z}_{\leq N}$ for some N. Recall that $\mathscr{Z}_{\leq N} \subset \text{Pol}(\mathbb{Z}, \mathbb{Z})$ is defined to consist of all linear functions f on \mathbb{Z} of the form

$$f(x_1,\ldots,x_n) = \sum_{i=1}^n c_i x_i$$

for some $c_i \in \mathbb{Z}$ such that $\sum_{i=1}^{N} c_i$ is odd and $\sum_{i=1}^{N} |c_i| \leq N$. This is achieved in two steps.

In the first step, we provide a minion homomorphism from $\text{Pol}(\mathbf{H}, \mathbf{G})$ to $\text{Pol}(\mathbb{Z}, \mathbb{Z})$. This is achieved by following the constructions described above, i.e., the transformations

$$\operatorname{graph} \overset{\mathsf{Hom}(\mathbf{K}_2,*)}{\longrightarrow} \mathbb{Z}_2 \operatorname{-complex} \overset{|*|}{\longrightarrow} \mathbb{Z}_2 \operatorname{-space} \overset{\pi_1(*)}{\longrightarrow} \operatorname{group},$$

and showing that pushing a polymorphism through this sequence of constructions preserves minors. This essentially follows from the well-known facts that these constructions behave well with respect to products. A detailed proof is presented in section 3.4.1.

When we push a polymorphism $f \in Pol(\mathbf{H}, \mathbf{G})$ through these constructions, we first obtain a \mathbb{Z}_2 -simplicial map $f' \in Pol(Hom(\mathbf{K}_2, \mathbf{H}), Hom(\mathbf{K}_2, \mathbf{G}))$, which in turn induces a \mathbb{Z}_2 -map $|f'| \in Pol(|Hom(\mathbf{K}_2, \mathbf{H})|, |Hom(\mathbf{K}_2, \mathbf{G})|)$. Then, by composing with the assumed \mathbb{Z}_2 -map $s : |Hom(\mathbf{K}_2, \mathbf{G})| \to S^1$ (and assuming without loss of generality that \mathbf{H} is an odd cycle), we obtain a polymorphism g of S^1 , and then finally a polymorphism g_* of the group $\pi_1(S^1) \simeq \mathbb{Z}$. As discussed before, g_* can be described more concretely as a linear function whose coefficients c_i are the winding numbers of maps $S^1 \to S^1$ defined by $t \mapsto s \circ g(x_0, \ldots, t, \ldots, x_0)$, where $x_0 \in S^1$ is an arbitrary (but fixed) point.

In the second step, we use the discrete structure of the graphs \mathbf{H} and \mathbf{G} , as well as the action of \mathbb{Z}_2 , to show that the image of $\text{Pol}(\mathbf{H}, \mathbf{G})$ under the constructed minion homomorphism is contained in $\mathscr{Z}_{\leq N}$ for some N. This is described in section 3.4.2.

We note that there are several ways to present the proof of Theorem 1.3. These presentations would look different on the surface, but in fact they use the same underlying topological concepts, just hidden to various extents. For example, the proof that appeared in the conference version [54] hides topology in a more direct combinatorial approach. Yet another version of the proof can be given in the language of *recolorings*: the required minion homomorphism would map two polymorphisms fand f' from Pol(\mathbf{H}, \mathbf{G}) to the same function if and only if f can be recolored to f' by changing one output value at a time, that is, if there is a sequence of polymorphisms $f_0, \ldots, f_n \in Pol(\mathbf{H}, \mathbf{G})$ such that $f_0 = f, f_n = f'$ and, for each $i \in \{0, \ldots, n-1\}$, there is a unique tuple \bar{t}_i with $f_i(\bar{t}_i) \neq f_{i+1}(\bar{t}_i)$. However, as shown in [72], recolorability is inherently a topological notion, so the resulting proof would have the same essence as the one presented here. We chose the current presentation because we believe that it reflects the 'true essence' of the proof.

3.4.1. A minion homomorphism. As mentioned above, the required minion homomorphism is obtained as a composition of three mappings. The first one is a minion homomorphism from polymorphisms of graphs to polymorphisms of simplicial complexes. We implicitly use that for any graphs \mathbf{H}_1 and \mathbf{H}_2 , there is a natural isomorphism of \mathbb{Z}_2 -simplicial complexes

3

$$\mathsf{Hom}(\mathbf{K}_2,\mathbf{H}_1)\times\mathsf{Hom}(\mathbf{K}_2,\mathbf{H}_2)\simeq\mathsf{Hom}(\mathbf{K}_2,\mathbf{H}_1\times\mathbf{H}_2)$$

given by the \mathbb{Z}_2 -simplicial map: $((a, b), (a', b')) \mapsto ((a, a'), (b, b'))$. See Figure 3.3 for an example.

LEMMA 3.18. For graphs H, G, the mapping

$$\mu_1 \colon \operatorname{Pol}(\mathbf{H}, \mathbf{G}) \to \operatorname{Pol}(\operatorname{\mathsf{Hom}}(\mathbf{K}_2, \mathbf{H}), \operatorname{\mathsf{Hom}}(\mathbf{K}_2, \mathbf{G}))$$

defined as

$$\mu_1(f)((u_1, v_1), \dots, (u_n, v_n)) := (f(u_1, \dots, u_n), f(v_1, \dots, v_n))$$

is a minion homomorphism.

Proof. Let us first check that $\mu_1(f)$ is indeed a simplicial map. Assume that $\sigma_1, \ldots, \sigma_n$ are faces in $\text{Hom}(\mathbf{K}_2, \mathbf{H})$, i.e., σ_i is a subset of edges of some complete directed bipartite subgraph of \mathbf{H} . We may assume without loss of generality that σ_i is the set of all edges of a complete directed bipartite subgraph of \mathbf{H} which gives $\sigma_i = U_i \times V_i$ for some $U_i, V_i \subseteq V(G)$, which form a bipartition of some complete bipartite subgraph of \mathbf{G} . By definition,

$$\mu_1(f)(\sigma_1 \times \dots \times \sigma_n) = \{(f(u_1, \dots, u_n), f(v_1, \dots, v_n)) : u_i \in U_i, v_i \in V_i\} = f(U) \times f(V),$$

where $U = U_1 \times \cdots \times U_n$ and $V = V_1 \times \cdots \times V_n$. By the definition of graph product, all edges between U and V are present in \mathbf{H}^n . Consequently, $f(U) \times f(V)$ is a complete directed bipartite subgraph of **G**, since f is a polymorphism. This implies that $\mu_1(f)(\sigma_1 \times \cdots \times \sigma_n)$ is a face of $\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})$.

It is straightforward that the mapping μ_1 preserves both the \mathbb{Z}_2 -action and taking minors.

The next step is from \mathbb{Z}_2 -simplicial complexes to \mathbb{Z}_2 -spaces. The map that we construct will not be a minion homomorphism, it will satisfy a weaker condition that will be sufficient later.

DEFINITION 3.19. Assume that \mathcal{X}, \mathcal{Y} are \mathbb{Z}_2 -spaces and let \mathscr{M} be a minion. We say that a mapping $\xi \colon \mathscr{M} \to \operatorname{Pol}(\mathcal{X}, \mathcal{Y})$ preserves minors up to \mathbb{Z}_2 -homotopy if for all $n, m > 0, f \in \mathscr{M}^{(n)}$ and $\pi \colon [n] \to [m]$, we have that $\xi(f^{\pi})$ is \mathbb{Z}_2 -homotopic to $\xi(f)^{\pi}$.

We recall that the points in the geometric representation of K can be viewed as convex combinations of vertices of K, more precisely $|\mathsf{K}| = \{\sum_{v \in \sigma} \lambda_v v \mid \sigma \in \mathsf{K}, \lambda_v \geq 0, \sum_{v \in \sigma} \lambda_v = 1\}$ and the representation of each point is unique (as we identified $v \in V(K)$ with affinely independent vectors in $\mathbb{R}^{V(K)}$, namely, the basis vectors). This is used in the following lemma.

LEMMA 3.20. Let K, K' be two \mathbb{Z}_2 -simplicial complexes. Let $\mu_2 \colon \text{Pol}(\mathsf{K},\mathsf{K}') \to \text{Pol}(|\mathsf{K}|,|\mathsf{K}'|)$ be the linear extension, i.e., μ_2 takes $f \in \text{Pol}(\mathsf{K},\mathsf{K}')$ to

$$\mu_2(f) \colon \left(\sum_{v \in \sigma_1} \lambda_{1,v} v, \dots, \sum_{v \in \sigma_n} \lambda_{n,v} v\right) \mapsto \sum_{v_1 \in \sigma_1, \dots, v_n \in \sigma_n} \lambda_{1,v_1} \cdots \lambda_{n,v_n} f(v_1, \dots, v_n)$$

for a point in $\Delta^{\sigma_1} \times \cdots \times \Delta^{\sigma_n} \subseteq |\mathsf{K}|^n$. Then μ_2 preserves minors up to \mathbb{Z}_2 -homotopy.

Copyright (C) by SIAM. Unauthorized reproduction of this article is prohibited.

Proof. Let $f: \mathsf{K}^n \to \mathsf{K}'$ be a simplicial map and pick any $\pi: [n] \to [m]$. Then

$$\mu_2(f^{\pi})\left(\sum_{v\in\sigma_1}\lambda_{1,v}v,\ldots,\sum_{v\in\sigma_m}\lambda_{m,v}v\right) = \sum_{v_1\in\sigma_1,\ldots,v_m\in\sigma_m}\lambda_{1,v_1}\cdots\lambda_{m,v_m}f^{\pi}(v_1,\ldots,v_m)$$
$$= \sum_{v_1\in\sigma_1,\ldots,v_m\in\sigma_m}\lambda_{1,v_1}\cdots\lambda_{m,v_m}f(v_{\pi(1)},\ldots,v_{\pi(n)})$$

On the other hand, if we take the induced map first and then the minor, we obtain

$$\mu_2(f)\left(\sum_{v\in\sigma_{\pi(1)}}\lambda_{\pi(1),v}v,\ldots,\sum_{v\in\sigma_{\pi(n)}}\lambda_{\pi(n),v}v\right)$$
$$=\sum_{v_1\in\sigma_{\pi(1)},\ldots,v_n\in\sigma_{\pi(n)}}\lambda_{\pi(1),v_1}\cdots\lambda_{\pi(n),v_n}f(v_1,\ldots,v_n).$$

Both points lie in $\Delta^{\sigma} \subseteq |\mathsf{K}'|$ for $\sigma = \{f(v_1, \ldots, v_n) \mid v_i \in \sigma_{\pi(i)}\} \in \mathsf{K}'$. We can thus continuously move from one to the other. Formally, we define a homotopy $h \colon |\mathsf{K}|^m \times [0,1] \to |\mathsf{K}'|$ by

$$h(x_1,\ldots,x_m,t) = t\mu_2(f)(x_{\pi(1)},\ldots,x_{\pi(n)}) + (1-t)\mu_2(f^{\pi})(x_1,\ldots,x_m).$$

It is clear that $h_t: \mathbf{x} \mapsto h(\mathbf{x}, t)$ is a well-defined \mathbb{Z}_2 -map, and therefore h is the required \mathbb{Z}_2 -homotopy. (As a side note, observe that the homotopy is constant on vertices: for any vertices v_1, \ldots, v_m of K and each $t \in [0, 1], h_t(v_1, \ldots, v_m)$ is equal to $f^{\pi}(v_1, \ldots, v_m) = f(v_{\pi(1)}, \ldots, v_{\pi(n)})$.)

LEMMA 3.21. Let **H** be a nonbipartite graph. Then there is a \mathbb{Z}_2 -map $r: S^1 \to |\text{Hom}(\mathbf{K}_2, \mathbf{H})|$.

Proof. Since the graph **H** is not bipartite, there is a homomorphism $h: \mathbf{C}_k \to \mathbf{H}$ for some odd k. This induces a \mathbb{Z}_2 -map $|h'|: |\mathsf{Hom}(\mathbf{K}_2, \mathbf{C}_k)| \to |\mathsf{Hom}(\mathbf{K}_2, \mathbf{H})|$, and since $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{C}_k)|$ is \mathbb{Z}_2 -homeomorphic to S^1 (see Example 3.11), the claim follows.

LEMMA 3.22. Let \mathbf{H}, \mathbf{G} be two graphs such that \mathbf{H} is nonbipartite, $\mathbf{H} \to \mathbf{G}$, and there is a \mathbb{Z}_2 -map $s: |\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})| \to S^1$. Then $\mu: \operatorname{Pol}(\mathbf{H}, \mathbf{G}) \to \operatorname{Pol}(S^1, S^1)$ defined as

$$\mu(f)(x_1,\ldots,x_n) := s(\mu_2\mu_1(f)(r(x_1),\ldots,r(x_n))),$$

where μ_1 , μ_2 , and r are from Lemmas 3.18, 3.20, and 3.21, respectively, preserves minors up to \mathbb{Z}_2 -homotopy.

Proof. Assume that f is a polymorphism from **H** to **G** of arity n and let $\pi: [n] \to [m]$. We want to prove that $\mu(f)^{\pi}$ is \mathbb{Z}_2 -homotopic to $\mu(f^{\pi})$. From Lemmas 3.18 and 3.20, we have that $\mu_2\mu_1(f)^{\pi}$ and $\mu_2\mu_1(f^{\pi})$ are \mathbb{Z}_2 -homotopic; let h' be a \mathbb{Z}_2 -homotopy that witnesses this fact. We define a \mathbb{Z}_2 -homotopy $h: \mathcal{T}^m \times [0,1] \to \mathcal{S}^1$ by

$$h(x_1, \dots, x_m, t) = sh'(r(x_1), \dots, r(x_m), t).$$

Indeed, for t = 0,

$$h(x_1, \dots, x_m, 0) = s\left(\mu_2 \mu_1(f)\left(r(x_{\pi(1)}), \dots, r(x_{\pi(n)})\right)\right) = \mu(f)^{\pi}(x_1, \dots, x_m)$$

Copyright (c) by SIAM. Unauthorized reproduction of this article is prohibited.

while for t = 1,

$$h(x_1, \dots, x_m, 1) = s(\mu_2 \mu_1(f^{\pi})(r(x_1), \dots, r(x_m))) = \mu(f^{\pi})(x_1, \dots, x_m).$$

This concludes the proof.

The final step is from \mathbb{Z}_2 -spaces to the fundamental groups. Recall that, for a continuous function $f: \mathcal{X} \to \mathcal{Y}$, we have a group homomorphism $f_*: \pi_1(\mathcal{X}) \to \pi_1(\mathcal{Y})$ defined as $f_*([\ell]) = [f\ell]$. We will also need a group homomorphism from $\pi_1(\mathcal{X})^n$ to $\pi_1(\mathcal{X}^n)$ that is guaranteed to exist for any path connected space \mathcal{X} by [43, Proposition 1.12]. One such homomorphism is the mapping $e_n: \pi_1(\mathcal{X})^n \to \pi_1(\mathcal{X}^n)$ defined as

$$e_n([\ell_1],\ldots,[\ell_n]) := [t \mapsto (\ell_1(t),\ldots,\ell_n(t))]$$

for $\ell_1, \ldots, \ell_n \colon \mathcal{S}^1 \to \mathcal{X}$.

LEMMA 3.23. Let \mathcal{X}, \mathcal{Y} be two path connected \mathbb{Z}_2 -spaces. Then the mapping

 $\nu(f) := f_* e_n$

is a minion homomorphism from $\operatorname{Pol}(\mathcal{X}, \mathcal{Y})$ to $\operatorname{Pol}(\pi_1(\mathcal{X}), \pi_1(\mathcal{Y}))$. Moreover, if f and g are homotopic then $\nu(f) = \nu(g)$.

Proof. Assume that $f \in \text{Pol}(\mathcal{X}, \mathcal{Y})$ is of arity $n, \pi \colon [n] \to [m]$, and $\ell_1, \ldots, \ell_m \colon \mathcal{S}^1 \to \mathcal{X}$. To simplify notation, let $\ell(t) = (\ell_1(t), \ldots, \ell_m(t))$ go from \mathcal{S}^1 to \mathcal{X}^n and $\ell_\pi(t) = (\ell_{\pi(1)}(t), \ldots, \ell_{\pi(n)}(t))$ go from \mathcal{S}^1 to \mathcal{X}^m . Using the definitions of f_* and e_n , we get

 $\nu(f^{\pi})(\ell_1,\ldots,\ell_m) = [f^{\pi}\ell] = [f\ell_{\pi}] = \nu(f)^{\pi}(\ell_1,\ldots,\ell_m).$

Finally, since $f_* = g_*$ if f and g are homotopic, we also get that $\nu(f) = \nu(g)$.

We recall that $\pi_1(S^1) \simeq \mathbb{Z}$ (Lemma 3.14). In the following statement we identify the two isomorphic groups to obtain a minion homomorphism to $\operatorname{Pol}(\mathbb{Z},\mathbb{Z})$, the minion of all linear functions on \mathbb{Z} .

COROLLARY 3.24. Let \mathbf{H}, \mathbf{G} be two graphs such that \mathbf{H} is nonbipartite, $\mathbf{H} \to \mathbf{G}$, and there is a \mathbb{Z}_2 -map s: $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})| \to S^1$. The mapping $\nu\mu$ is a minion homomorphism from $\mathsf{Pol}(\mathbf{H}, \mathbf{G})$ to $\mathsf{Pol}(\mathbb{Z}, \mathbb{Z})$ assuming μ is as in Lemma 3.22 and ν is as in Lemma 3.23.

Proof. Clearly, $\nu\mu$: Pol(**H**, **G**) \rightarrow Pol(\mathbb{Z}, \mathbb{Z}) is a well-defined mapping that preserves arities. We need to show that it also preserves minors. This follows from the facts that μ preserves minors up to \mathbb{Z}_2 -homotopy (Lemma 3.22) and that ν is a minion homomorphism that is constant on \mathbb{Z}_2 -homotopy classes (Lemma 3.23). More precisely, assume $f \in \text{Pol}(\mathbf{H}, \mathbf{G})$ is of arity n and $\pi \colon [n] \to [m]$. Then $\mu(f^{\pi})$ and $\mu(f)^{\pi}$ are \mathbb{Z}_2 -homotopic, and therefore

$$\nu(\mu(f^{\pi})) = \nu(\mu(f)^{\pi}) = \nu(\mu(f))^{\pi},$$

where the second equality follows from minor preservation by ν .

We remark that, for any (*n*-ary) function $f \in Pol(\mathbf{H}, \mathbf{G})$, the coefficients of the linear function $\nu \mu(f) = \sum_{i=1}^{n} c_i x_i$ can be naturally thought of as the degrees of f at the corresponding coordinates. Such degrees can be defined in a combinatorial way (see [54])—the intuitions in that approach are still topological, but the technical proofs become somewhat ad hoc.

3.4.2. Bounding essential arity. To finish the analysis of polymorphisms from **H** to **G** necessary for applying Theorem 2.12, we need to bound the essential arity of functions in the image of $\nu\mu$ (defined above) and show that none of these functions is a constant function. We achieve this by proving that the image of $\nu\mu$ is contained in the minion $\mathscr{Z}_{\leq N}$ for some N. Recall that this minion is defined to be the set of all functions $f: \mathbb{Z}^n \to \mathbb{Z}$ of the form $f(x_1, \ldots, x_n) = c_1 x_1 + \cdots + c_n x_n$ for some $c_1, \ldots, c_n \in \mathbb{Z}$ with $\sum_{i=1}^n |c_i| \leq N$ and $\sum_{i=1}^n c_i$ odd.

The oddness of the sum of coefficients follows from a well-known fact about \mathbb{Z}_2 maps on \mathcal{S}^1 . We recall that the degree of a map $f: \mathcal{S}^1 \to \mathcal{S}^1$ is the integer d_f such that the induced map f_* on $\pi_1(\mathcal{S}^1) = \mathbb{Z}$ is $x \mapsto d_f \cdot x$.

LEMMA 3.25 ([43, Proposition 2B.6]). The degree of any \mathbb{Z}_2 -map $f: S^1 \to S^1$ is odd.

LEMMA 3.26. Let \mathbf{H}, \mathbf{G} be two graphs such that \mathbf{H} is nonbipartite, $\mathbf{H} \to \mathbf{G}$, and there is a \mathbb{Z}_2 -map s: $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})| \to S^1$. Let μ be as in Lemma 3.22, ν as in Lemma 3.23, and let $\mathscr{H} = \mathrm{Pol}(\mathbf{H}, \mathbf{G})$. If $f \in \mathrm{Pol}(\mathbf{H}, \mathbf{G})$ and $\nu \mu(f) : (x_1, \ldots, x_n) \mapsto \sum_{i=1}^n c_i x_i$, then $\sum_{i=1}^n c_i$ is odd.

Proof. Consider the (unique) unary minor $h(x) := f(x, \ldots, x)$ of f. Since $\mu(h)$ is a \mathbb{Z}_2 -map, by Lemma 3.25 it has an odd degree, i.e., $\nu\mu(h) : x \mapsto d_h x$ for some odd d_h . Finally, since $\nu\mu$ preserves minors, we get that $d_h = \sum_{i \in [n]} c_i$ which we wanted to show to be odd.

The bound on the sum of absolute values of coefficients is given by the discrete structure of the involved graphs. The key here is that there are only finitely many polymorphisms of a fixed arity between two given finite graphs.

LEMMA 3.27. Let \mathscr{M} be a minion on finite sets A, B. Assume that $\xi \colon \mathscr{M} \to \operatorname{Pol}(\mathbb{Z},\mathbb{Z})$ is a minion homomorphism. Then there exists N such that for all $f \in \mathscr{M}$, if $\xi(f) \colon (x_1, \ldots, x_n) \mapsto \sum_{i=1}^n c_i x_i$, then $\sum_{i=1}^n |c_i| \leq N$.

Proof. We first consider binary functions. There are only finitely many functions $f \in \mathscr{M}^{(2)}$, so clearly the sum of the absolute values of the coefficients of $\xi(f)$ is bounded by some N. We argue that the same N provides a bound for all other arities as well. Let $f \in \mathscr{M}$ and $\xi(f)(x_1, \ldots, x_n) = c_1 x_1 + \cdots + c_n x_n$. Let $\sigma \colon [n] \to \{0, 1\}$ be defined as

$$\sigma(i) = \begin{cases} 0 & \text{if } c_i \le 0, \text{ and} \\ 1 & \text{if } c_i > 0. \end{cases}$$

That is, $\sigma^{-1}(1)$ is the set of all coordinates of $\xi(f)$ with positive coefficients, and $\sigma^{-1}(0)$ of those with negative or zero coefficients. Now, let $g = f^{\sigma}$, that is, g is the minor of f defined by

$$g(x_0, x_1) = f(x_{\sigma(1)}, \dots, x_{\sigma(n)}),$$

so g is obtained by identifying all variables of f that induce positive coefficients on $\xi(f)$, and also all those that induce negative coefficients. We let $\xi(g)(x_0, x_1) = c^+x_1 + c^-x_0$. Since ξ preserves minors, we get that $c^+ = \sum_{c_i>0} c_i$ and $c^- = \sum_{c_i<0} c_i$. Note that $c^+ \ge 0$ and $c^- \le 0$. Finally,

$$\sum_{i \in [n]} |c_i| = \sum_{c_i > 0} c_i - \sum_{c_i < 0} c_i = c^+ - c^- = |c^+| + |c^-| \le N,$$

where the last inequality follows from the definition of N.

We can now finish the proof of Theorem 1.4.

Proof of Theorem 1.4. We assume that \mathbf{H}, \mathbf{G} are two graphs such that \mathbf{H} is nonbipartite, $\mathbf{H} \to \mathbf{G}$, and there is a \mathbb{Z}_2 -map $s \colon |\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})| \to S^1$. Let μ be as in Lemma 3.22, ν as in Lemma 3.23, and N be the bound obtained from Lemma 3.27 for $\mathscr{M} = \operatorname{Pol}(\mathbf{H}, \mathbf{G})$ and $\xi = \nu \mu$. We claim that $\nu \mu(f) \in \mathscr{Z}_{\leq N}$ for each $f \in \operatorname{Pol}(\mathbf{H}, \mathbf{G})$. Assume $\nu\mu(f): (x_1, \ldots, x_n) \mapsto \sum_{i=1}^n c_i x_i$. We have that $\sum_{i=1}^n c_i$ is odd from Lemma 3.26, and $\sum_{i=1}^{n} |c_i| \leq N$ by the choice of N. This concludes that there is a minim homomorphism from $Pol(\mathbf{H}, \mathbf{G})$ to a minion of bounded essential arity, namely, $\mathscr{Z}_{< N}$, and thus Theorem 1.4 follows from Theorem 2.12.

3.5. Proofs of Corollaries 1.5 and 1.6. To show that Theorem 1.4 implies Corollaries 1.5 and 1.6, we need the following facts about the structure of $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})|$ for the relevant graphs **G**. These facts seem to be folklore, but we include proofs for completeness.

LEMMA 3.28. For any 2 < p/q < 4 and any square-free nonbipartite graph G, there exist \mathbb{Z}_2 -maps

1. s_1 : $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{K}_{p/q})| \to S^1$, and 2. s_2 : $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})| \to S^1$.

Proof. (1) We first define a \mathbb{Z}_2 -map $g: |\mathsf{Hom}(\mathbf{K}_2, \mathbf{K}_{p/q})| \to \mathbb{R}^2$. We will show that the origin 0 is not in the image of g, which then implies that the map $x \mapsto g(x)/|g(x)|$ is a \mathbb{Z}_2 -map to \mathcal{S}^1 .

First, we define g on vertices of the complex, i.e., oriented edges of $\mathbf{K}_{p/q}$: Place p points x_0, \ldots, x_{p-1} on S^1 in a regular *p*-gon pattern. Map the edge (a, b) to the point $x_b - x_a$. Then extend g linearly to the interior points of faces. Clearly, g is a \mathbb{Z}_2 -map. See Figure 3.4 for a visualization of g for $\mathbf{K}_{7/2}$.

Let σ be a face. That is, $\sigma \subseteq A \times B \subseteq E(K_{p/q})$ for some nonempty sets of vertices A, B. For an edge (a, b), the distance between x_a, x_b on the circle (the length of the shorter arc between them) is at least $2\pi \cdot q/p$. Since p/q < 4, this is greater than $\pi/2$. Hence there are no $a, a' \in A, b, b' \in B$ such that $x_a, x_b, x_{a'}, x_{b'}$ occur in this order on the circle, as the distances would add up to more than 2π . Therefore, there is a line in \mathbb{R}^2 that strictly separates $\{x_a \mid a \in A\}$ from $\{x_b \mid b \in B\}$ (indeed, scanning the circle clockwise, there is exactly one interval from A to B and exactly one from B to A, both of length at least $\pi/2$; any line crossing these intervals will do). This implies that the convex hull of vectors $x_a - x_b$ cannot contain 0 (since each such vector has a positive dot product with the normal vector of the line).

(2) This statement follows from observing that $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})|$ is essentially 1dimensional which loosely follows from the facts that there are no copies of the complete bipartite graph $\mathbf{K}_{2,2}$ and that every free \mathbb{Z}_2 -space of dimension 1 maps to \mathcal{S}^1 (see [59, Proposition 5.3.2(v)]). We present a compressed version of the two arguments.

Let $E^+ \cup E^- = V(\mathsf{Hom}(\mathbf{K}_2, \mathbf{G}))$ be an arbitrary partition into two sets that are swapped by reversing the edges, i.e., $-E^+ = E^-$ and $-E^- = E^+$. This means that we choose an orientation for each edge of G, and denote by E^+ the set of all edges of **G** oriented the chosen way, while E^- is the set of all edges oriented the opposite way. We define a mapping $h: |\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})| \to S^1$ on the vertices of $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})|$ by setting h(e) = (1,0) if $e \in E^+$ and h(e) = (-1,0) if $e \in E^-$.

We extend this mapping to inner points of faces. First, observe that for every face $\sigma \in Hom(\mathbf{K}_2, \mathbf{G})$ with at least two elements, there is a vertex $u \in V(G)$ such



FIG. 3.4. $\mathbf{K}_{7/2}$ with a complete bipartite subgraph (on the left) and the corresponding face of $|\mathsf{Hom}(\mathbf{K}_2,\mathbf{K}_{7/2})|$ after mapping to \mathbb{R}^2 (on the right).

that either $\sigma \subseteq \{(u,v) \in E(G)\}$ or $\sigma \subseteq \{(v,u) \in E(G)\}$, as otherwise we can find a copy of \mathbf{C}_4 in \mathbf{G} . We map $|\sigma|$ for faces of the first form to the arc connecting (1,0) and (-1,0) with positive y coordinates, and $|\sigma|$ for faces of the second form to the arc with negative y coordinates. More precisely, if $\sigma = \{(u,v_1), \ldots, (u,v_n)\}$, we let $a = \sum_{i=1}^n \lambda_i(u,v_i)$ be a point of $|\sigma|$. Put

$$x = \sum_{(u,v_i)\in E^+} \lambda_i - \sum_{(u,v_i)\in E^-} \lambda_i$$

and $y = \sqrt{1 - x^2}$ (note that $|x| \le 1$, so y is well-defined), and define h(a) = (x, y). Now to preserve the \mathbb{Z}_2 -action, we map the geometric representations of the faces of the second form to the arc with negative y coordinates analogously putting $y = -\sqrt{1 - x^2}$. Clearly, the mapping h defined this way is continuous and it is easy to check that indeed h(-a) = -h(a) for each $a \in |\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})|$.

4. Adjunction. In this section we will use both graphs and digraphs, which by default are allowed to have loops. We will work with certain (di)graph constructions that can be seen as functions from the set of all finite (di)graphs to itself. On one occasion in this section (section 4.2.1), we will allow the image of a finite digraph to be an infinite digraph; this will be specified. We denote the set of all finite graphs and digraphs by Gr and Dgr, respectively. The class of all (finite and infinite) digraphs is denoted by Dgr_{∞} .

In this section, we explain what *adjunction* is and how it can be used to obtain reductions between PCSPs. The notion of adjointness we present is a special case of the more general notion of adjoint functors in category theory. We restrict our attention to an order-theoretic version thereof (i.e., to posetal or thin categories), which only considers the existence of homomorphisms; this is also known as a (monotone) Galois connection. Generally, a monotone Galois connection between two preordered sets P_1 and P_2 is a pair of maps $\lambda: P_1 \to P_2$ and $\gamma: P_2 \to P_1$ such that, for all $a \in P_1$ and $b \in P_2$,

(4.1)
$$\lambda(a) \le b$$
 if and only if $a \le \gamma(b)$.

Copyright (c) by SIAM. Unauthorized reproduction of this article is prohibited.

For us, the preorder \leq will always be the homomorphism preorder \rightarrow , and the sets P_1 and P_2 will be either Dgr or Gr. In this case, Λ and Γ are adjoint if, for all (di)graphs **H** and **G**, we have

(4.2)
$$\Lambda \mathbf{H} \to \mathbf{G}$$
 if and only if $\mathbf{H} \to \Gamma \mathbf{G}$.

In this case Λ is a *left adjoint* and Γ is a *right adjoint*. If, for some Λ , there exists such a Γ we also say that Λ has (or admits) a right adjoint. Similarly, we say that Γ has a *left adjoint* if there exists such Λ .

Adjunction is an abstraction of a few concepts that are already present in the theory of (P)CSPs; notably, the Inv-Pol Galois correspondences of Geiger [36], and Bodnarchuk et al., [10, 11] and Pippenger [63] between sets of functions and sets of relations can also been seen as adjunctions where, in (4.1), the preorder on one side is the inclusion and the preorder on the other side is the inverse inclusion. We remark that many constructions described in [5, sections 3 and 4] (see e.g., [5, Lemma 4.4]) form pairs of adjoint functions. We also remark that condition (4.2) makes perfect sense when Λ and Γ are maps between the sets of relational structures of different signatures (say, between the set of all finite digraphs and the set of all finite 3-uniform hypergraphs), and all results in section 4.2 hold in this more general setting.

This section is organized as follows. In section 4.1, we show that the standard gadget reductions from the algebraic approach to the CSP can be seen as a special case of adjunction. In section 4.2, we give general results about adjunctions and reductions between PCSPs. In section 4.3, we apply specific cases of adjunction to prove our results about the hardness of approximate graph coloring and demonstrate that the reductions between PCSPs obtained there cannot be captured by the algebraic theory from [5, 23]. Finally, in section 4.4, we use another specific adjunction to prove that, in a precise technical sense, the complexity of the promise graph homomorphism problem depends only on the topological properties of graphs.

To emphasize that many of our proofs in this section do not assume computability of reductions, we will use the following definition.

DEFINITION 4.1. Let Λ be a function from Dgr to Dgr or from Gr to Gr. We say that Λ is

- a reduction from PCSP(H₁, G₁) to PCSP(H₂, G₂) if it preserves the yes- and no-answers of the two problems, i.e., for any I, I → H₁ implies ΛI → H₂ and I → G₁ implies ΛI → G₂. Preserving yes-answers is also called completeness and preserving no-answers soundness;
- log-space/polynomial-time computable if there is a log-space/polynomial-time algorithm that on input I outputs ΛI .

4.1. Adjunction in CSPs. The algebraic approach to the CSP studies certain constructions on templates of CSPs called *pp-powers*, and it asserts that if \mathbf{A} is a pp-power of \mathbf{B} then there is a log-space reduction from $\text{CSP}(\mathbf{A})$ to $\text{CSP}(\mathbf{B})$. This reduction is a function from instances of $\text{CSP}(\mathbf{A})$ to $\text{CSP}(\mathbf{B})$ computable in log-space that we call a *gadget replacement*. In fact, any such reduction is a left adjoint to the corresponding pp-power construction. We present the notions simplified for digraphs and refer to [7, section 3.1] for more background. We will use the constructions from Example 1.1 as a running example in this subsection.

Both functions are parameterized by a gadget or a primitive positive formula (a pp-formula), thus giving a reduction for each gadget. A digraph gadget formula of arity n is a logical formula $\phi(x_1, \ldots, x_n, y_1, \ldots, y_n)$ of the form

$$\exists z_1, \dots, z_m \, (u_1, v_1) \in E \land \dots \land (u_k, v_k) \in E \land (u_1' = v_1') \land \dots \land (u_\ell' = v_\ell')$$

where $u_i, v_i, u'_i, v'_i \in \{x_1, \ldots, x_n, y_1, \ldots, y_n, z_1, \ldots, z_m\}$. Such a formula can also be represented by a gadget digraph \mathbf{J}_{ϕ} , which is a digraph with distinguished vertices labeled $x_1, \ldots, x_n, y_1, \ldots, y_n$, obtained from vertices $\{x_1, \ldots, x_n, y_1, \ldots, y_n, z_1, \ldots, z_m\}$ and edges $\{(u_i, v_i) \mid i \in [k]\}$ by identifying some of the vertices (according to the equalities in ϕ).

If we want a gadget to transform an undirected graph to an undirected graph, we require that the gadget is *symmetric*, i.e., that the formula $\phi(x_1, \ldots, x_n, y_1, \ldots, y_n)$ is logically equivalent to $\phi(y_1, \ldots, y_n, x_1, \ldots, x_n)$ for all graphs, or that the gadget graph has an automorphism switching x_i and y_i for each i.

The subdivision from Example 1.1 is defined by using the following digraph gadget formula of arity 1:

$$\phi_k(x,y) = \exists z_1, \dots, z_{k-1} \, (x, z_1) \in E \land (z_1, z_2) \in E \land \dots \land (z_{k-1}, y) \in E$$

which is symmetric if we only consider undirected graphs. The corresponding gadget would be an (unoriented) path of length k connecting the two distinguished vertices x and y.

DEFINITION 4.2. The gadget replacement Λ_{ϕ} assigned to a digraph gadget ϕ is then defined by applying the following construction. Starting with a digraph **H**,

- 1. for each vertex $v \in V(H)$, introduce new vertices $v_1, \ldots, v_n \in V(\Lambda_{\phi}H)$;
- 2. for each edge $(u,v) \in E(H)$, introduce a fresh copy of the gadget digraph \mathbf{J}_{ϕ} while identifying x_1, \ldots, x_n with u_1, \ldots, u_n and y_1, \ldots, y_n with v_1, \ldots, v_n ; we denote the remaining vertices of this copy of \mathbf{J}_{ϕ} by $z_{u,v}$ for $z \in \{z_1, \ldots, z_m\}$.

Note that some of the vertices u_i, v_j above might get identified, which can also result in long chains of identifications. Nevertheless, Λ_{ϕ} is log-space computable.

The subdivision Λ_k of Example 1.1 is the same as the gadget replacement Λ_{ϕ_k} . We show an example of application in Figure 4.1.

DEFINITION 4.3. The pp-power $\Gamma_{\phi} \mathbf{G}$ of a digraph \mathbf{G} defined by ϕ is obtained by the following construction.

- 1. $V(\Gamma_{\phi}G) = V(G)^n$, and
- 2. $((u_1, \ldots, u_n), (v_1, \ldots, v_n)) \in E(\Gamma_{\phi}G)$ if $\phi(u_1, \ldots, u_n, v_1, \ldots, v_n)$ is true in **G**, in other words, there exists a homomorphism $e_{u,v}$ from the gadget digraph \mathbf{J}_{ϕ} to **G** such that $e_{u,v}(x_i) = u_i$ and $e_{u,v}(y_i) = v_i$ for all i.

Again, it is not hard to check that $\Gamma_k \colon \mathsf{Dgr} \to \mathsf{Dgr}$ from Example 1.1 is the same as the pp-power Γ_{ϕ_k} , i.e., it is the graph on the same vertex set where two vertices



FIG. 4.1. Example of gadget replacement and a pp-power.

are connected by an edge if and only if they are connected by a path of length k in the original graph. Again, see Figure 4.1 for an application.

The standard reductions used in the algebraic approach are of the form Λ_{ϕ} : it is well known (see [7, Theorem 13]) that Λ_{ϕ} is a reduction from $\text{CSP}(\Gamma_{\phi}\mathbf{G})$ to $\text{CSP}(\mathbf{G})$ for any digraph \mathbf{G} . For example, it is not hard to see that Λ_{ϕ_3} from our running example is really a reduction from $\text{CSP}(\mathbf{K}_3)$ to $\text{CSP}(\mathbf{C}_9)$ since \mathbf{K}_3 is homomorphically equivalent to $\Gamma_{\phi_3}\mathbf{C}_9$, and hence it has the same CSP. This observation also follows immediately from the fact that Λ_{ϕ} and Γ_{ϕ} are adjoint, which we show here directly.

Observation 4.4. Let ϕ be a pp-formula. Then Λ_{ϕ} and Γ_{ϕ} are adjoint.

Proof. To prove that indeed Λ_{ϕ} and Γ_{ϕ} are adjoint, first assume $h: \mathbf{H} \to \Gamma_{\phi} \mathbf{G}$ is a homomorphism. Such a homomorphism is a map $h: V(H) \to V(G)^n$ which can be seen as an *n*-tuple of mappings $h_1, \ldots, h_n: V(H) \to V(G)$. Further, since *h* preserves edges, we have that for each $(u, v) \in E(H)$, $\phi(h_1(u), \ldots, h_n(u), h_1(v), \ldots, h_n(v))$ is true in \mathbf{G} , which gives a homomorphism $e_{u,v}: \mathbf{J}_{\phi} \to \mathbf{G}$ such that $e_{u,v}(x_i) = h_i(u)$ and $e_{u,v}(y_i) = h_i(v)$. We use these $e_{u,v}$'s to define a homomorphism $h': \Lambda_{\phi} \mathbf{H} \to \mathbf{G}$:

1. put $h'(u_i) = h_i(u)$ for each $u \in \mathbf{H}$ and i;

2. extend h' to new vertices introduced by the second step of gadget replacement

of the edge $(u, v) \in E(H)$ by putting $h'(z) = e_{u,v}(z_{u,v})$ for all $z \in \{z_1, \ldots, z_m\}$. Clearly, h' is a homomorphism since each $e_{u,v}$ is and there are no edges in $\Lambda_{\phi} \mathbf{H}$ that are not included in some copy of \mathbf{J}_{ϕ} . For the other implication, assume $g \colon \Lambda_{\phi} \mathbf{H} \to \mathbf{G}$. We define $g' \colon \mathbf{H} \to \Gamma_{\phi} \mathbf{G}$ as $g'(u) = (g(u_1), \ldots, g(u_n))$ for each $u \in V(H)$. It is straightforward to check that g' is indeed a homomorphism. This concludes the proof.

One of the main strengths of the algebraic approach lies in a description of when such reductions apply, by means of polymorphisms and minion homomorphisms; see [7, Theorem 38] (originally appeared in [8]) and [5, Theorem 4.12] for the promise setting. We return to this later in this section (Example 4.9).

4.2. General results about adjunction for PCSPs. In the following lemma, we give a few basic and useful properties of adjoint functors that are well known in category theory. We provide proofs of these facts for completeness. We say that a function $\Lambda: \text{Dgr} \to \text{Dgr}$ is monotone if $\Lambda \mathbf{H} \to \Lambda \mathbf{G}$ for all \mathbf{H}, \mathbf{G} such that $\mathbf{H} \to \mathbf{G}$; and it *preserves disjoint unions* if $\Lambda(\mathbf{H}_1 + \mathbf{H}_2)$ and $\Lambda \mathbf{H}_1 + \Lambda \mathbf{H}_2$ are homomorphically equivalent for all digraphs $\mathbf{H}_1, \mathbf{H}_2$ (we denote disjoint union with +).

LEMMA 4.5. Let $\Lambda, \Gamma: \mathsf{Dgr} \to \mathsf{Dgr}$ be adjoint. Then

- 1. $\mathbf{G} \rightarrow \Gamma \Lambda \mathbf{G}$ for all digraphs \mathbf{G} ;
- 2. $\Lambda\Gamma\mathbf{H} \rightarrow \mathbf{H}$ for all digraphs \mathbf{H} ;
- 3. both Λ and Γ are monotone; and
- 4. Λ preserves disjoint unions.

Proof. We start by proving (1): $\Lambda \mathbf{G} \to \Lambda \mathbf{G}$ implies that $\mathbf{G} \to \Gamma \Lambda \mathbf{G}$ by adjunction. Similarly for (2), observe that $\Gamma \mathbf{H} \to \Gamma \mathbf{H}$ implies $\Lambda \Gamma \mathbf{H} \to \mathbf{H}$ by adjunction.

For (3) assume $\mathbf{H} \to \mathbf{G}$. Then by (1), we have $\mathbf{H} \to \mathbf{G} \to \Gamma \Lambda \mathbf{G}$, and therefore by adjunction $\Lambda \mathbf{H} \to \Lambda \mathbf{G}$. This concludes that Λ is monotone. Similarly from (2), we have $\Lambda \Gamma \mathbf{H} \to \mathbf{H} \to \mathbf{G}$ and hence by adjunction $\Gamma \mathbf{H} \to \Gamma \mathbf{G}$, so Γ is monotone.

For (4), consider the disjoint union of digraphs $\mathbf{H}_1, \mathbf{H}_2$. Note that $\mathbf{H}_i \to \mathbf{H}_1 + \mathbf{H}_2$ for i = 1, 2 implies $\Lambda \mathbf{H}_1 + \Lambda \mathbf{H}_2 \to \Lambda (\mathbf{H}_1 + \mathbf{H}_2)$ by monotonicity. To show the other direction, observe that $\mathbf{H}_i \to \Gamma \Lambda \mathbf{H}_i \to \Gamma (\Lambda \mathbf{H}_1 + \Lambda \mathbf{H}_2)$ for i = 1, 2 by (1) and

monotonicity of Γ , hence $\mathbf{H}_1 + \mathbf{H}_2 \rightarrow \Gamma(\Lambda \mathbf{H}_1 + \Lambda \mathbf{H}_2)$ and therefore $\Lambda(\mathbf{H}_1 + \mathbf{H}_2) \rightarrow \Lambda \mathbf{H}_1 + \Lambda \mathbf{H}_2$ by adjunction.

The next result is the main theorem of this subsection. It describes when an adjunction provides a reduction between two PCSPs. This theorem will be applied in the following two subsections to provide new reductions between promise digraph homomorphism problems of the sort that has not been described before.

THEOREM 4.6. Let $\Lambda, \Gamma: \mathsf{Dgr} \to \mathsf{Dgr}$ be adjoint. Let $\mathbf{H}_1, \mathbf{G}_1, \mathbf{H}_2, \mathbf{G}_2$ be digraphs such that $\mathbf{H}_i \to \mathbf{G}_i$ for i = 1, 2. Then Λ is a reduction from $\mathrm{PCSP}(\mathbf{H}_1, \mathbf{G}_1)$ to $\mathrm{PCSP}(\mathbf{H}_2, \mathbf{G}_2)$ if and only if $\mathbf{H}_1 \to \Gamma \mathbf{H}_2$ and $\Gamma \mathbf{G}_2 \to \mathbf{G}_1$.

Proof. Assume first that $\mathbf{H}_1 \to \Gamma \mathbf{H}_2$ and $\Gamma \mathbf{G}_2 \to \mathbf{G}_1$. Then Λ preserves yesinstances because $\mathbf{I} \to \mathbf{H}_1$ implies $\mathbf{I} \to \Gamma \mathbf{H}_2$ (since $\mathbf{H}_1 \to \Gamma \mathbf{H}_2$ by assumption) and then $\Lambda \mathbf{I} \to \mathbf{H}_2$ by adjointness. It also preserves no-instances because $\Lambda \mathbf{I} \to \mathbf{G}_2$ implies $\mathbf{I} \to \Gamma \mathbf{G}_2$ by adjointness and then $\mathbf{I} \to \mathbf{G}_1$ because $\Gamma \mathbf{G}_2 \to \mathbf{G}_1$ (by assumption). Hence Λ is a reduction, as claimed.

For the converse, preserving yes-instances means that for $\mathbf{I} \in \mathsf{Dgr}$, $\mathbf{I} \to \mathbf{H}_1$ implies $\Lambda \mathbf{I} \to \mathbf{H}_2$. Using this with $\mathbf{I} = \mathbf{H}_1$, we get that $\Lambda \mathbf{H}_1 \to \mathbf{H}_2$ and thus $\mathbf{H}_1 \to \Gamma \mathbf{H}_2$ by adjointness. Preserving no-instances means that $\Lambda \mathbf{I} \to \mathbf{G}_2$ implies $\mathbf{I} \to \mathbf{G}_1$. Take $\mathbf{I} = \Gamma \mathbf{G}_2$. Since $\Lambda \Gamma \mathbf{G}_2 \to \mathbf{G}_2$ by Lemma 4.5(2), we have $\Gamma \mathbf{G}_2 \to \mathbf{G}_1$.

Naturally, we use the above theorem in the case that Λ is log-space computable to obtain a log-space reduction between the two PCSPs. In the same way, it can also be applied if Λ is polynomial-time computable, if the goal is to get a polynomialtime reduction, etc. Note that, in such applications, Γ need not be computable to guarantee the correctness of the reduction.

Remark 4.7. We note (a well-known categorical fact) that any two right adjoints Γ_1 and Γ_2 of Λ are homomorphically equivalent in the following sense: for all \mathbf{G} , $\Gamma_1 \mathbf{G}$ and $\Gamma_2 \mathbf{G}$ are homomorphically equivalent. This follows, for example, from the above theorem: we have that Λ is a reduction from PCSP($\Gamma_1 \mathbf{G}, \Gamma_1 \mathbf{G}$) to PCSP(\mathbf{G}, \mathbf{G}) since Γ_1 is a right adjoint to Λ and, consequently, $\Gamma_1 \mathbf{G} \leftrightarrow \Gamma_2 \mathbf{G}$ since Γ_2 is a right adjoint.

COROLLARY 4.8. Let $\Lambda, \Gamma: \mathsf{Dgr} \to \mathsf{Dgr}$ be adjoint. Then

- 1. Λ is a reduction from PCSP($\mathbf{H}, \Gamma \mathbf{G}$) to PCSP($\Lambda \mathbf{H}, \mathbf{G}$) for all graphs \mathbf{H}, \mathbf{G} such that $\mathbf{H} \to \Gamma \mathbf{G}$ (or, equivalently, $\Lambda \mathbf{H} \to \mathbf{G}$);
- 2. A is a reduction from $PCSP(\Gamma \mathbf{H}, \Gamma \mathbf{G})$ to $PCSP(\mathbf{H}, \mathbf{G})$ for all graphs \mathbf{H}, \mathbf{G} such that $\mathbf{H} \to \mathbf{G}$.

Proof. For (1), the first condition of Theorem 4.6 is equivalent to $\mathbf{H} \to \Gamma \Lambda \mathbf{H}$, which holds by adjunction (see Lemma 4.5(1)); the second condition is trivial: $\Gamma \mathbf{G} \to \Gamma \mathbf{G}$. For (2), both conditions are trivial.

We remark that all reductions described in Theorem 4.6 can be deduced from the special case in Corollary 4.8(1) by composing it with trivial reductions (that map every instance to itself). Recall that there is a trivial reduction from PCSP(\mathbf{H}', \mathbf{G}') to PCSP(\mathbf{H}, \mathbf{G}) if (and only if) $\mathbf{H}' \to \mathbf{H}$ and $\mathbf{G} \to \mathbf{G}'$; this is referred to as a homomorphic relaxation [5, Definition 4.6]. If digraphs $\mathbf{H}_1, \mathbf{G}_1, \mathbf{H}_2, \mathbf{G}_2$ satisfy the conditions of Theorem 4.6, then we have the following sequence of reductions:

$$\operatorname{PCSP}(\mathbf{H}_1, \mathbf{G}_1) \xrightarrow{\operatorname{tr.}} \operatorname{PCSP}(\mathbf{H}_1, \Gamma \mathbf{G}_2) \xrightarrow{\operatorname{Cor.} 4.8(1)} \operatorname{PCSP}(\Lambda \mathbf{H}_1, \mathbf{G}_2) \xrightarrow{\operatorname{tr.}} \operatorname{PCSP}(\mathbf{H}_2, \mathbf{G}_2)$$

Copyright (c) by SIAM. Unauthorized reproduction of this article is prohibited.

Similarly, Corollary 4.8(2) implies all reductions in Theorem 4.6:

 $\operatorname{PCSP}(\mathbf{H}_1, \mathbf{G}_1) \xrightarrow{\operatorname{tr.}} \operatorname{PCSP}(\Gamma \mathbf{H}_2, \Gamma \mathbf{G}_2) \xrightarrow{\operatorname{Cor.} 4.8(2)} \operatorname{PCSP}(\mathbf{H}_2, \mathbf{G}_2).$

Example 4.9. What we described in section 4.1 in the context of CSPs, can be generalized to PCSPs as follows. The following are equivalent:

- 1. there is a pp-formula ϕ such that Λ_{ϕ} is a log-space reduction from PCSP $(\mathbf{H}_1, \mathbf{G}_1)$ to PCSP $(\mathbf{H}_2, \mathbf{G}_2)$ (i.e., there exists some gadget reduction between the two);
- 2. there is a pp-formula ϕ such that $\mathbf{H}_1 \to \Gamma_{\phi} \mathbf{H}_2$ and $\Gamma_{\phi} \mathbf{G}_2 \to \mathbf{G}_1$ (i.e., $(\mathbf{H}_1, \mathbf{G}_1)$ is a homomorphic relaxation of a pp-power of $(\mathbf{H}_2, \mathbf{G}_2)$);
- 3. there is a minion homomorphism $\operatorname{Pol}(\mathbf{H}_2, \mathbf{G}_2) \to \operatorname{Pol}(\mathbf{H}_1, \mathbf{G}_1)$.

The equivalence of (2) and (3) is by [5, Theorem 4.12]. The equivalence of the last two items and (1) is implicit in [5] (see, e.g., [5, Lemma 4.11]), but the equivalence of (1) and (2) follows directly from Theorem 4.6 above.

For example, all NP-hard (nonpromise) CSPs are reducible to one another in this way. The understanding that one can get simple reductions between CSPs by relating their sets of polymorphisms goes back at least as far as [47]. The use of pp-formulas and minion homomorphisms was initiated in [21] and [8], respectively.

4.2.1. Are all reductions given by adjunctions?. Theorem 4.6 raises a question of whether all reductions between PCSPs are given by adjunctions, in the sense that every reduction is a left adjoint from some adjoint pair. By Lemma 4.5(3-4) we have to restrict this question to reductions that are monotone and preserve disjoint unions. We will show that the answer to this question is positive, with a small technical caveat that the right adjoint might produce infinite digraphs on a finite input. Note that this caveat is not an issue, since the right adjoint does not need to be computable.

This suggests that looking at classes of adjoints that generalize the simple gadget constructions Λ_{ϕ} could lead to understanding an essential part of all reductions between PCSPs. In particular, we hope that the use of the PCP theorem in proving NP-hardness of PCSPs (see [5, section 5]) can be superseded this way. We remark that, e.g., the reduction in Dinur's proof of the PCP theorem [25] is not necessarily monotone: this is due to the fact that the number of repetitions of the gap amplification depends on the size of the input.

The core of the argument in the proof of the following theorem is a well-known categorical statement (the adjoint functor theorem). Again, we provide a full proof for completeness.

THEOREM 4.10. Let $\mathbf{H}_1, \mathbf{G}_1, \mathbf{H}_2, \mathbf{G}_2$ be finite digraphs such that $\mathbf{H}_i \to \mathbf{G}_i$ for i = 1, 2, and let $\Lambda : \mathsf{Dgr} \to \mathsf{Dgr}$ be a reduction from $\mathsf{PCSP}(\mathbf{H}_1, \mathbf{G}_1)$ to $\mathsf{PCSP}(\mathbf{H}_2, \mathbf{G}_2)$. Assume additionally that Λ is monotone and preserves disjoint unions. Then there is a function $\Gamma : \mathsf{Dgr} \to \mathsf{Dgr}_{\infty}$ with possibly infinite images such that, for all finite digraphs \mathbf{H} and \mathbf{G} , we have $\mathbf{H} \to \Gamma \mathbf{G}$ if and only if $\Lambda \mathbf{H} \to \mathbf{G}$. Moreover, we have $\mathbf{H}_1 \to \Gamma \mathbf{H}_2$ and $\Gamma \mathbf{G}_2 \to \mathbf{G}_1$ for any such Γ .

Proof. We define $\Gamma \mathbf{G}$ to be the disjoint union of all finite digraphs \mathbf{I} such that $\Lambda \mathbf{I} \to \mathbf{G}$. Assuming that \mathbf{H} and \mathbf{G} are finite digraphs, we immediately get that $\Lambda \mathbf{H} \to \mathbf{G}$ implies $\mathbf{H} \to \Gamma \mathbf{G}$. We first prove the other implication for connected \mathbf{H} : assuming that $\mathbf{H} \to \Gamma \mathbf{G}$, we get that \mathbf{H} maps to some connected component of $\Gamma \mathbf{G}$ and thus $\mathbf{H} \to \mathbf{I}$ for some finite \mathbf{I} such that $\Lambda \mathbf{I} \to \mathbf{G}$. This gives that $\Lambda \mathbf{H} \to \Lambda \mathbf{I} \to \mathbf{G}$

since Λ is monotone. For disconnected **H**, we use that Λ preserves disjoint unions, so we may repeat the above argument for each component separately.

The "moreover" claim is proved similarly to Theorem 4.6. In particular, the proof that if Λ is a reduction then $\mathbf{H}_1 \to \Gamma \mathbf{H}_2$ is identical to the one in Theorem 4.6. To prove that $\Gamma \mathbf{G}_2 \to \mathbf{G}_1$, we cannot simply use preservation of no-instances on the possibly infinite $\Gamma \mathbf{G}_2$. Instead, we get that for every \mathbf{I} finite, $\mathbf{I} \to \Gamma \mathbf{G}_2$ implies $\Lambda \mathbf{I} \to \mathbf{G}_2$, which implies $\mathbf{I} \to \mathbf{G}_1$ (because Λ is a reduction). A homomorphism from the possibly infinite $\Gamma \mathbf{G}_2$ to the finite \mathbf{G}_1 is then given by compactness.

While monotonicity is a key assumption in Theorem 4.10, preservation of disjoint unions can always be enforced on any reduction by first precomputing connected components of the input (which can be done in log-space due to [66]), and then applying the original reduction on each of the components separately.

We note that all the proofs in this section reduce between decision problems; they can be adapted for search problems. For that we need to additionally assume that there is an efficient way to find a homomorphism $\mathbf{I} \to \Gamma \mathbf{G}$ given a homomorphism $\Lambda \mathbf{I} \to \mathbf{G}$ on input (note that \mathbf{G} is fixed here). All the adjoint pairs that we use in the following subsections indeed have this property.

From now on, we return to considering only finite digraphs.

4.2.2. Reductions that have both a left and a right adjoint. In the two applications below, we use reductions that are a left adjoint from one adjoint pair and, at the same time, the right adjoint from another adjoint pair. (In fact, these reductions will be of the form of a pp-power Γ_{ϕ} , as described in section 4.1, for special gadgets ϕ). The property of being both left and right adjoint has the following consequence.

THEOREM 4.11. Let Γ be a log-space computable function that has a right adjoint Ω , and a log-space computable left adjoint Λ . Then PCSP($\Gamma \mathbf{H}, \mathbf{G}$) and PCSP($\mathbf{H}, \Omega \mathbf{G}$) are log-space equivalent for all digraphs \mathbf{H}, \mathbf{G} such that $\Gamma \mathbf{H} \to \mathbf{G}$.

Proof. Corollary 4.8(1) applied for Γ and Ω gives that Γ is a reduction from PCSP($\mathbf{H}, \Omega \mathbf{G}$) to PCSP($\Gamma \mathbf{H}, \mathbf{G}$). We claim that Λ is a reduction from PCSP($\Gamma \mathbf{H}, \mathbf{G}$) to PCSP($\mathbf{H}, \Omega \mathbf{G}$). This follows from Theorem 4.6 applied to Λ, Γ: we need to check that $\Gamma \mathbf{H} \to \Gamma \mathbf{H}$, which holds trivially, and that $\Gamma \Omega \mathbf{G} \to \mathbf{G}$, which follows by Lemma 4.5(2), since Γ and Ω are adjoint.

4.3. The arc digraph construction. Let **D** be a digraph. The *arc digraph* (or *line digraph*) of **D**, denoted δ **D**, is the digraph whose vertices are arcs (directed edges) of **D** and whose arcs are pairs of the form ((u, v), (v, w)). In other words, $\delta: \text{Dgr} \to \text{Dgr}$ is the pp-power Γ_{ϕ} corresponding to (n = 2 and) the following gadget digraph:



or to the pp-formula $\phi = (x_1, x_2) \in E \land (y_1, y_2) \in E \land x_2 = y_1$. It thus has a left adjoint $\delta_L = \Lambda_{\phi}$, though we will not need it. More surprisingly, δ has a right adjoint $\delta_R \colon \mathsf{Dgr} \to \mathsf{Dgr}$.

DEFINITION 4.12. For a digraph **D**, let $\delta_R \mathbf{D}$ be the digraph that has a vertex for each pair $S, T \subseteq V(D)$, where S or T can be empty, such that $S \times T \subseteq E(D)$, and an arc from (S,T) to (S',T') if and only if $T \cap S' \neq \emptyset$. We give a proof of the adjunction below for completeness. While δ will be the reduction we use, δ_R will be useful for understanding the best reduction we can get from δ .

LEMMA 4.13 ([33]). δ and δ_R are adjoint.

Proof. Let \mathbf{H}, \mathbf{G} be digraphs and let $h: \delta \mathbf{H} \to \mathbf{G}$ be a homomorphism. That is, h(u, v) is a vertex of \mathbf{G} for each arc (u, v) of \mathbf{H} , and for every pair of arcs (u, v), (v, w)in \mathbf{H} , there is an arc from h(u, v) to h(v, w) in \mathbf{G} . We can define a homomorphism $\mathbf{H} \to \delta_R \mathbf{G}$ as $v \mapsto (s(v), t(v))$, where $s(v) := \{h(u, v) \mid (u, v) \in E(H)\}$ and $t(v) := \{h(v, w) \mid (v, w) \in E(H)\}$. Then $s(v) \times t(v) \subseteq E(G)$, so (s(v), t(v)) is indeed a vertex of $\delta_R \mathbf{G}$. Moreover, for every arc (u, v) of $\mathbf{H}, t(u) \cap s(v)$ is nonempty, as it contains (u, v); hence (s, t) is a homomorphism to $\delta_R \mathbf{G}$.

Conversely, let (s,t) define a homomorphism $\mathbf{H} \to \delta_R \mathbf{G}$. That is, s(v), t(v) are subsets of V(G) such that $s(v) \times t(v) \subseteq E(G)$ and for every arc (u, v) of $\mathbf{H}, t(u) \cap s(v) \neq \emptyset$. We define a homomorphism $h: \delta \mathbf{H} \to \mathbf{G}$ as follows: choose h(u, v) to be an arbitrary vertex in $t(u) \cap s(v)$. For any two arcs (u, v), (v, w) in \mathbf{H} , we have that h(u, v) is a vertex in s(v) and h(v, w) is a vertex in t(v), hence (h(u, v), h(v, w)) is an arc of \mathbf{G} . Thus h is indeed a homomorphism $\delta \mathbf{H} \to \mathbf{G}$.

By Corollary 4.8(2), δ is a reduction from $\text{PCSP}(\delta_R \mathbf{H}, \delta_R \mathbf{G})$ to $\text{PCSP}(\mathbf{H}, \mathbf{G})$, for all digraphs \mathbf{H}, \mathbf{G} . Let us see what this gives for classical colorings, i.e., when \mathbf{H} and \mathbf{G} are cliques. Let us denote $b(n) := \binom{n}{\lfloor n/2 \rfloor}$.

Observation 4.14. For all $n \in \mathbb{N}$, there are homomorphisms

$$\mathbf{K}_{b(n)} \rightarrow \delta_R \mathbf{K}_n \rightarrow \mathbf{K}_{2^n}$$

Proof. Consider vertices of the form $(S, V(K_n) \setminus S)$ in $\delta_R \mathbf{K}_n$, for subsets S of $V(K_n)$ of size exactly $\lfloor n/2 \rfloor$. Clearly for any two such different S, S', the set S' intersects $V(K_n) \setminus S$, so these vertices from a clique of size b(n) in $\delta_R \mathbf{K}_n$. For the other bound, note that mapping a vertex (S,T) to $(S, V(K_n) \setminus S)$ gives a homomorphism from $\delta_R \mathbf{K}_n$ to its subgraph of size at most 2^n , and therefore to the clique \mathbf{K}_{2^n} .

In other words, if $\chi(\delta \mathbf{G}) \leq n$ (i.e., if $\delta \mathbf{G} \to \mathbf{K}_n$), then $\mathbf{G} \to \delta_R \mathbf{K}_n \to \mathbf{K}_{2^n}$, hence $\chi(\mathbf{G}) \leq 2^n$. Similarly, if $\chi(\mathbf{G}) \leq b(n)$, then $\chi(\delta \mathbf{G}) \leq n$. Therefore, δ has the remarkable property of decreasing the chromatic number roughly logarithmically (even though it is computable in log-space!). This was first proved by Harner and Entringer in [42].

Observation 4.14 can be made tight if we use another, somewhat trivial adjunction between digraphs and graphs: Let sym **D** be the symmetric closure of a digraph **D** and let sub**D** be the maximal symmetric subgraph of **D**; so sub $\mathbf{D} \to \mathbf{D} \to \text{ sym } \mathbf{D}$ by the identity maps. Observe that sym and sub are adjoint: sym $\mathbf{D} \to \mathbf{D}'$ if and only if $\mathbf{D} \to \text{sub } \mathbf{D}'$ for all digraphs \mathbf{D}, \mathbf{D}' .³ Composing the two adjunctions, we get that δ sym is adjoint to sub δ_R . Therefore, for any digraphs \mathbf{H}, \mathbf{G} with $\mathbf{H} \to \mathbf{G}, \delta$ sym is a reduction from PCSP(sub $\delta_R \mathbf{H}, \text{sub } \delta_R \mathbf{G})$ to PCSP(\mathbf{H}, \mathbf{G}) by Corollary 4.8(2). For cliques, Poljak and Rödl [65] showed the following.

LEMMA 4.15 ([65]). For all $n \in \mathbb{N}$, $\operatorname{sub} \delta_R \mathbf{K}_n$ is homomorphically equivalent to $\mathbf{K}_{b(n)}$.

³This is in fact the composition of two adjoint pairs: taking sym and sub as functions from digraphs to graphs and the inclusion function ι from graphs to digraphs, we have sym $\mathbf{D} \to \mathbf{G}$ if and only if $\mathbf{D} \to \iota \mathbf{G}$ and $\iota \mathbf{G} \to \mathbf{D}$ if and only if $\mathbf{G} \to \operatorname{sub} \mathbf{D}$ for all graphs \mathbf{G} and digraphs \mathbf{D} .

Proof. As before, mapping a vertex (S,T) of $\operatorname{sub} \delta_R \mathbf{K}_n$ to $(S, V(K_n) \setminus S)$ gives a homomorphism to the subgraph induced by vertices of the form $(S, V(K_n) \setminus S)$, so we can restrict our attention to it. A (bidirected) clique corresponds exactly to an antichain of sets S (in the subset lattice), so by Sperner's theorem (on the maximal size of such antichains) the largest clique has size b(n). Independent sets in this subgraph correspond exactly to chains of sets S, thus by Dilworth's theorem (on poset width) the subgraph can be covered with b(n) independent sets, giving a b(n)-coloring.

Lemma 4.15 is equivalent to the statement that for an undirected graph \mathbf{G} , $\delta \mathbf{G} = \delta \operatorname{sym} \mathbf{G} \to \mathbf{K}_n$ if and only if $\mathbf{G} \to \mathbf{K}_{b(n)}$. This, in particular, means that the chromatic number of $\delta \mathbf{G}$ is determined by $\chi(\mathbf{G})$, namely, $\chi(\delta \mathbf{G}) = \min\{n \mid \chi(\mathbf{G}) \leq b(n)\}$. This together with Corollary 4.8(2) implies that $\delta \operatorname{sym}$ gives the following reduction for approximate coloring:

LEMMA 4.16. $\text{PCSP}(\mathbf{K}_{b(k')}, \mathbf{K}_{c'})$ log-space reduces to $\text{PCSP}(\mathbf{K}_k, \mathbf{K}_c)$ for all k' > 1, $c' \ge b(k') > 1$, $c \ge k$ such that $b(c) \le c'$ and $k' \le k$.

Proof. By Theorem 4.6, δ sym is a reduction between the two problems if $\mathbf{K}_{b(k')} \rightarrow$ sub $\delta_R \mathbf{K}_k$ and sub $\delta_R \mathbf{K}_c \rightarrow \mathbf{K}_{c'}$. Lemma 4.15 then implies that the first condition is satisfied if $b(k') \leq b(k)$, which is implied by $k' \leq k$, and the second is satisfied if $b(c) \leq c'$.

Remark 4.17. We remark that the reduction in Lemma 4.16 cannot be obtained by using a standard gadget reduction captured by the algebraic approach [5] (see Example 4.9). In detail, since b(4) = 6, PCSP($\mathbf{K}_6, \mathbf{K}_{b(c)}$) log-space reduces to PCSP($\mathbf{K}_4, \mathbf{K}_c$) for all $c \ge 4$. This contrasts with [5, Proposition 10.3] which says that there exists a c such that Pol($\mathbf{K}_4, \mathbf{K}_c$) admits no minion homomorphism to any Pol($\mathbf{K}_{k'}, \mathbf{K}_{c'}$) for $c' \ge k' > 4$. Therefore, constructions like δ change the set of polymorphisms in an essential way and we believe that understanding the relation between Pol($\mathbf{H}, \delta_R \mathbf{G}$) and Pol($\delta \mathbf{H}, \mathbf{G}$) is an important question for future work.

4.3.1. Proof of Theorem 1.9. One consequence we derive from Lemma 4.16 is a strengthening of Huang's result.

THEOREM 4.18 (Huang [46]). For all sufficiently large k and $c = 2^{\Omega(k^{1/3})}$, PCSP ($\mathbf{K}_k, \mathbf{K}_c$) is NP-hard.

We improve the asymptotics from a subexponential c to a single exponential: from $c = \Theta(2^{k^{1/3}})$ to $c = b(k) - 1 = \Theta(2^k/\sqrt{k})$ while at the same time relaxing the condition from "sufficiently large k" to $k \ge 4$.

THEOREM 1.9. For all $k \ge 4$ and $c = \binom{k}{\lfloor k/2 \rfloor} - 1$, $PCSP(\mathbf{K}_k, \mathbf{K}_c)$ is NP-hard.

This theorem is proved by starting from Theorem 4.18 and repeatedly using the reduction δ sym. Roughly speaking, each step improves the asymptotics a little. After a few steps, this results in a single-exponential function, and with slightly more precision, this results in exactly b(k) - 1. Moreover, one can notice that the requirements on sufficiently large k get relaxed with every step. This allows us after sufficiently many steps to arrive at any $k \geq 4$.

We note it would *not* be sufficient to start from a quasi-polynomial $c = k^{\Theta(\log k)}$ in Khot's [50] earlier result in place of Huang's Theorem 4.18.

Proof of Theorem 1.9. We start with Theorem 4.18 which asserts a constant C > 0 such that

(4.3) $PCSP(\mathbf{K}_k, \mathbf{K}_{2|C \cdot k^{1/3}|})$ is NP-hard, for sufficiently large k.

After one reduction using Lemma 4.16, we obtain the following.

CLAIM 4.19. $PCSP(\mathbf{K}_k, \mathbf{K}_{\lfloor 2^{k/4} \rfloor})$ is NP-hard for sufficiently large k.

Proof. Substituting b(k) for k in (4.3), we get that $\text{PCSP}(\mathbf{K}_{b(k)}, \mathbf{K}_{2\lfloor C \cdot b(k)^{1/3} \rfloor})$ is NP-hard for sufficiently large k. We apply Lemma 4.16; to show that it implies the claim, we need $b(\lfloor 2^{k/4} \rfloor) \leq 2^{\lfloor C \cdot b(k)^{1/3} \rfloor}$. This follows since $b(m) \leq 2^m$ for all m and $2^{k/4} \leq (2^k/k)^{1/3} \leq |C \cdot b(k)^{1/3}|$ for sufficiently large k, and therefore

$$b(|2^{k/4}|) \le 2^{\lfloor 2^{k/4} \rfloor} \le 2^{\lfloor C \cdot b(k)^{1/3} \rfloor}$$

as we wanted to show.

The second reduction gives the following.

CLAIM 4.20. $PCSP(\mathbf{K}_k, \mathbf{K}_{\lfloor 2^k/4k \rfloor})$ is NP-hard for sufficiently large k.

Proof. Substitute b(k) for k in Claim 4.19 to get that $PCSP(\mathbf{K}_{b(k)}, \mathbf{K}_{\lfloor 2^{b(k)/4} \rfloor})$ is NP-hard. Observe that $2^k/4k \leq b(k)/4$ for sufficiently large k, and therefore

$$b(\lfloor 2^k/4k \rfloor) \leq b(\lfloor b(k)/4 \rfloor) \leq 2^{\lfloor b(k)/4 \rfloor} \leq \lfloor 2^{b(k)/4} \rfloor$$

Hence by Lemma 4.16, $PCSP(\mathbf{K}_{b(k)}, \mathbf{K}_{\lfloor 2^{b(k)/4} \rfloor})$ reduces to $PCSP(\mathbf{K}_k, \mathbf{K}_{\lfloor 2^k/4k \rfloor})$.

After the third reduction, we get the following.

CLAIM 4.21. $PCSP(\mathbf{K}_k, \mathbf{K}_{b(k-1)})$ is NP-hard for sufficiently large k.

Proof. Again, substitute b(k) for k in Claim 4.20 to obtain that $PCSP(\mathbf{K}_{b(k)}, \mathbf{K}_{\lfloor 2^{b(k)}/4b(k) \rfloor})$ is NP-hard. Observe that $b(k) \geq \frac{3}{2}b(k-1)$ for all $k \geq 1$, since

$$b(2k) = \binom{2k}{k} = \binom{2k-1}{k-1} \frac{2k}{k} = 2 \cdot b(2k-1) \ge \frac{3}{2}b(2k-1), \text{ and}$$
$$b(2k+1) = \binom{2k+1}{k} = \binom{2k}{k} \frac{2k+1}{k+1} = b(2k)\left(2 - \frac{1}{k+1}\right) \ge \frac{3}{2}b(2k).$$

Therefore,

$$b(b(k-1)) \le b\left(\frac{2}{3}b(k)\right) \le 2^{\frac{2}{3}b(k)} \le 2^{b(k)}/4b(k)$$

for sufficiently large k. Hence by Lemma 4.16, $PCSP(\mathbf{K}_{b(k)}, \mathbf{K}_{\lfloor 2^{b(k)}/4b(k) \rfloor})$ reduces to $PCSP(\mathbf{K}_k, \mathbf{K}_{b(k-1)})$.

Finally after the fourth reduction.

CLAIM 4.22. $PCSP(\mathbf{K}_k, \mathbf{K}_{b(k)-1})$ is NP-hard for sufficiently large k.

Proof. Substitute b(k) for k in Claim 4.21 to get that $PCSP(\mathbf{K}_{b(k)}, \mathbf{K}_{b(b(k)-1)})$ is NP-hard. By Lemma 4.16 this reduces to $PCSP(\mathbf{K}_k, \mathbf{K}_{b(k)-1})$.

This concludes the improvement in asymptotics. To relax the requirements for k, we repeat the reduction enough times. Each step is given by the following claim.

CLAIM 4.23. $PCSP(\mathbf{K}_{b(k)}, \mathbf{K}_{b(b(k))-1})$ log-space reduces to $PCSP(\mathbf{K}_k, \mathbf{K}_{b(k)-1})$ for all $k \ge 4$.

Proof. Lemma 4.16 gives the reduction since b(k) is strictly increasing and b(k) > 4 for $k \ge 4$, and hence $b(b(k) - 1) \le b(b(k)) - 1$.

To finish the proof, assume that $k' \ge 4$ and let k_0 be sufficiently large so Claim 4.22 is true for all $k \ge k_0$. Since $k' \ge 4$ there is n such that $k_0 \le b^{(n)}(k')$ (where $b^{(n)}$ denotes the n-fold composition of b). Applying Claim 4.23 n times gives a log-space reduction from PCSP($\mathbf{K}_{b^{(n)}(k')}, \mathbf{K}_{b^{(n+1)}(k')-1}$) to PCSP($\mathbf{K}_{k'}, \mathbf{K}_{b(k')-1}$), and since the first problem is NP-hard (Claim 4.22) this concludes that PCSP($\mathbf{K}_{k'}, \mathbf{K}_{b(k')-1}$) is NP-hard for all $k' \ge 4$.

4.3.2. Proof of Theorem 1.8. In this section we prove a slightly more general result than Theorem 1.8.

THEOREM 4.24. If, for some loopless digraph \mathbf{H} and all loopless digraphs \mathbf{G} such that $\mathbf{H} \to \mathbf{G}$, PCSP(\mathbf{H}, \mathbf{G}) is NP-hard, then PCSP(\mathbf{K}_3, \mathbf{G}) is NP-hard for all loopless digraphs \mathbf{G} such that $\mathbf{K}_3 \to \mathbf{G}$.

Indeed, it is easy to see that the assumption of Theorem 4.24 is slightly weaker than the assumption of Theorem 1.8, while the conclusions of the two theorems are equivalent.

We prove Theorem 4.24 by iterating the reduction given by δ in a similar way as in the last paragraph of the proof of Theorem 1.9; in fact, that part of the proof could be used with only minor changes to prove that $PCSP(\mathbf{K}_4, \mathbf{K}_c)$ is NP-hard for all $c \ge 4$ if $PCSP(\mathbf{K}_k, \mathbf{K}_c)$ is NP-hard for some k and all $c \ge k$. We get to \mathbf{K}_3 by omitting the intermediate use of sym, i.e., we keep the orientation of the edges. One step of the reduction is given by the following lemma.

LEMMA 4.25. Let **H** be a loopless digraph. If, for all loopless digraphs **G**' such that $\mathbf{H} \to \mathbf{G}'$, PCSP(\mathbf{H}, \mathbf{G}') is NP-hard, then PCSP($\delta \mathbf{H}, \mathbf{G}$) is NP-hard for all loopless digraphs **G** such that $\delta \mathbf{H} \to \mathbf{G}$.

Proof. Let **H** be a digraph that satisfies the premise, and let **G** be a loopless digraph such that $\delta \mathbf{H} \to \mathbf{G}$. We aim to prove that $\mathrm{PCSP}(\delta \mathbf{H}, \mathbf{G})$ is NP-hard. Corollary 4.8(1) gives a log-space reduction from $\mathrm{PCSP}(\mathbf{H}, \delta_R \mathbf{G})$ to $\mathrm{PCSP}(\delta \mathbf{H}, \mathbf{G})$. We claim that the digraph $\delta_R \mathbf{G}$ is loopless, which follows from the construction of δ_R (see Definition 4.12) and the assumption that **G** is loopless: Indeed, if a vertex (S,T) in $\delta_R \mathbf{G}$ has a loop, then $S \cap T \neq \emptyset$. Consequently, **G** has a loop on any vertex $v \in S \cap T$ since $S \times T \subseteq E(G)$. Clearly also $\mathbf{H} \to \delta_R \mathbf{G}$ since δ and δ_R are adjoint. Therefore, we have that $\mathrm{PCSP}(\delta \mathbf{H}, \mathbf{G})$ is NP-hard by the assumption of the lemma, and we conclude that $\mathrm{PCSP}(\delta \mathbf{H}, \mathbf{G})$ is NP-hard as well.

To finish the proof of Theorem 4.24, we will need the following two lemmas. The first one, which is a special case of the second, was independently discovered by Zhu [74], Poljak [64], and Schmerl (unpublished; see [68]). For the sake of completeness, we include the proof of Zhu [74].

LEMMA 4.26. There is a homomorphism $\delta(\delta \mathbf{K}_4) \rightarrow \mathbf{K}_3$.

Proof. The vertices of $\delta(\delta \mathbf{K}_4)$ are two consecutive pairs of arcs, i.e., they correspond to triples (i, j, k) in $\{0, 1, 2, 3\}$ such that $i \neq j$ and $j \neq k$. Two such triples (i, j, k) and (i', j', k') are adjacent if j = i' and k = j'. We define $h: \delta(\delta \mathbf{K}_4) \to \mathbf{K}_3$ so that $h: (i, j, k) \mapsto j$ if $j \in \{0, 1, 2\}$, and $h: (i, 3, k) \mapsto c$ for some $c \in \{0, 1, 2\} \setminus \{i, k\}$. It is straightforward to check that such an h is a valid coloring.

We note that $\delta(\operatorname{sym}\delta \mathbf{K}_4)$ is *not* 3-colorable, so it is important here to iterate δ rather than $\operatorname{sym}\delta$.

The next lemma essentially shows that iterating δ many times can bring a chromatic number of any finite loopless digraph down to 3. We use $\delta^{(i)}\mathbf{D}$ to denote the digraph obtained from \mathbf{D} by applying δi times. LEMMA 4.27. For every loopless digraph **D** there exists $i \ge 0$ s.t. $\delta^{(i)} \mathbf{D} \to \mathbf{K}_3$.

Proof. If **D** is a loopless digraph then $\mathbf{D} \to \mathbf{K}_{|V(D)|}$. Therefore, since δ is monotone, it is enough to prove the statement for $\mathbf{D} = \mathbf{K}_c$ and all c > 3 (the case $c \leq 3$ is trivial). Recall that $\mathbf{K}_{b(n)} \to \delta_R \mathbf{K}_n$ for all $n \geq 1$ (Observation 4.14) and, consequently, $\delta \mathbf{K}_{b(n)} \to \mathbf{K}_n$. Recall that b(n) is strictly increasing. Now since b(n) > 4 for $n \geq 4$, there is j such that $b^{(j)}(4) \geq c$. Since $\mathbf{K}_c \to \mathbf{K}_{b^{(j)}(4)}$, the monotonicity of δ implies $\delta \mathbf{K}_c \to \delta \mathbf{K}_{b^{(j)}(4)}$, which, with $\delta \mathbf{K}_{b^{(j)}(4)} \to \mathbf{K}_{b^{(j-1)}(4)}$, implies $\delta \mathbf{K}_c \to \mathbf{K}_{b^{(j-1)}(4)}$. Similarly, we get $\delta^{(2)} \mathbf{K}_c \to \delta \mathbf{K}_{b^{(j-1)}(4)} \to \mathbf{K}_{b^{(j-2)}(4)}$. Proceeding in the same way, we eventually get $\delta^{(j)} \mathbf{K}_c \to \mathbf{K}_4$. This together with Lemma 4.26 gives that $\delta^{(j+2)} \mathbf{K}_c \to \mathbf{K}_3$ which gives the claim for i = j + 2.

Proof of Theorem 4.24. Assume that **H** is a loopless digraph s.t. $PCSP(\mathbf{H}, \mathbf{G})$ is NP-hard for all loopless digraphs **G** with $\mathbf{H} \to \mathbf{G}$. Let *i* be such that $\delta^{(i)}\mathbf{H} \to \mathbf{K}_3$ which exists from Lemma 4.27. Now, iterating Lemma 4.25 *i* times gives that $PCSP(\delta^{(i)}\mathbf{H}, \mathbf{G})$ is NP-hard for all **G** such that $\delta^{(i)}\mathbf{H} \to \mathbf{G}$. Since $\delta^{(i)}\mathbf{H} \to \mathbf{K}_3$, $PCSP(\delta^{(i)}\mathbf{H}, \mathbf{G})$ trivially reduces to $PCSP(\mathbf{K}_3, \mathbf{G})$.

We remark that, if one iterates the reduction using δ further, one cannot improve Theorem 4.24 to imply the full extent of Conjecture 1.2. In particular, iterating δ cannot be used to show PCSP($\mathbf{C}_5, \mathbf{K}_c$) is NP-hard for all $c \geq 3$ given that PCSP($\mathbf{K}_3, \mathbf{K}_c$) is NP-hard for all $c \geq 3$. This is because $\delta^{(i)}\mathbf{K}_3$ contains a directed cycle of length three for all integers *i*.

Remark 4.28. Theorems 1.8 and 4.24 can be phrased in algebraic terms using so-called **H**-loop conditions, which recently gained popularity in universal algebra (see, e.g., [62]) and was featured in [5, section 6]. Fix a digraph **H** and let $(a_1, b_1), \ldots, (a_m, b_m)$ be the full list of arcs of **H**. The **H**-loop condition is the following identity (i.e., function equation) involving two minors of a function f:

(4.4)
$$f(x_{a_1}, \dots, x_{a_m}) = f(x_{b_1}, \dots, x_{b_m}).$$

(This definition is equivalent to the one used in [5, section 6].) One says that a minion \mathcal{M} satisfies the **H**-loop condition if it contains a function f satisfying (4.4).

The following two statements are equivalent for each loopless digraph H:

- PCSP(H, G) is NP-hard for all loopless digraphs G with $H \rightarrow G$; and
- PCSP(**A**, **B**) is NP-hard for all pairs of similar structures **A**, **B** such that Pol(**A**, **B**) does not satisfy the **H**-loop condition.

This claim can be proved in a similar way to [5, Theorems 6.9 and 6.12].

In this interpretation, Lemma 4.25 can be rephrased as follows: if, for all PCSPs, the failure to satisfy the **H**-loop condition implies NP-hardness, then so does the failure to satisfy the δ **H**-loop condition. Can this perspective be used to bring some ideas from algebra (e.g., [62]) to obtain better conditional hardness?

4.4. Only topology matters. All graphs in this subsection are undirected.

Recall Example 1.1. The functions Λ_k and Γ_k from this example are adjoint for all k. More surprisingly, for odd k, Γ_k is itself the *left* adjoint of a certain function $\Omega_k : \mathbf{Gr} \to \mathbf{Gr}$, i.e., for all graphs \mathbf{H} and \mathbf{G} , $\Gamma_k \mathbf{H} \to \mathbf{G}$ if and only if $\mathbf{H} \to \Omega_k \mathbf{G}$. The graph $\Omega_k \mathbf{G}$ for $k = 2\ell + 1$ is defined as follows; the vertices of $\Omega_k \mathbf{G}$ are tuples (A_0, \ldots, A_ℓ) of vertex subsets $A_i \subseteq V(G)$ such that A_0 contains exactly one vertex. Two such tuples (A_0, \ldots, A_ℓ) and (B_0, \ldots, B_ℓ) are adjacent if $A_i \subseteq B_{i+1}$, $B_i \subseteq A_{i+1}$ for all $i = 0 \ldots \ell - 1$, and $A_\ell \times B_\ell \subseteq E(G)$.

If there is a homomorphism $f: \Gamma_k \mathbf{H} \to \mathbf{G}$, then a homomorphism $\mathbf{H} \to \Omega_k \mathbf{G}$ is

obtained by mapping v to $(f(N^0(v)), \ldots, f(N^\ell(v)))$, where N^i denotes the set of vertices reachable via a walk of length exactly *i*. Conversely, if there is a homomorphism $f: \mathbf{H} \to \Omega_k \mathbf{G}$, then a homomorphism $\Gamma_k \mathbf{H} \to \mathbf{G}$ is obtained by mapping v to the unique vertex in the first set of $f(v) = (A_0, \ldots, A_\ell)$.

We note that Λ_k and Γ_k are log-space computable, for all odd k; however, Ω_k is not: $\Omega_k \mathbf{G}$ is exponentially larger than \mathbf{G} . Below, we will use the following observation that can be found together with more properties of the functions Λ_k, Γ_k , and Ω_k in [71, Lemma 2.3].

LEMMA 4.29. For any graph **G** and all odd k, $\Lambda_k \mathbf{G} \to \Omega_k \mathbf{G}$.

Proof. Since Γ_k and Ω_k are adjoint, the claim is equivalent to $\Gamma_k \Lambda_k \mathbf{G} \to \mathbf{G}$ (note that $\mathbf{G} \to \Gamma_k \Lambda_k \mathbf{G}$ since Λ_k and Γ_k are adjoint, so we will prove that the two graphs are homomorphically equivalent). We describe one such homomorphism $h \colon \Gamma_k \Lambda_k \mathbf{G} \to \mathbf{G}$. Note that $V(\Gamma_k \Lambda_k G) = V(\Lambda_k G) \supseteq V(G)$. We put h(v) = v for each $v \in V(G)$ and extend this to vertices introduced to $\Lambda_k \mathbf{G}$ by replacing an edge $(u, v) \in E(G)$ with a path of length k by alternatively mapping vertices on this path to u and v in such a way that h restricted to the path is a homomorphism from the odd path to (u, v). It is straightforward to check that h maps any two vertices of $\Lambda_k \mathbf{G}$ that are connected by a path of length k to an edge of \mathbf{G} , and therefore, it is a homomorphism from $\Gamma_k \Lambda_k \mathbf{G}$ to \mathbf{G} .

The next lemma gives the key reduction for the main result of this subsection, Theorem 4.32.

LEMMA 4.30. Let k be odd and **G** be a graph. If $PCSP(\mathbf{H}', \Omega_k \mathbf{G})$ is NP-hard for all nonbipartite graphs $\mathbf{H}' \to \Omega_k \mathbf{G}$, then $PCSP(\mathbf{H}, \mathbf{G})$ is NP-hard for all nonbipartite graphs $\mathbf{H} \to \mathbf{G}$.

Proof. Let \mathbf{H}, \mathbf{G} be nonbipartite with $\mathbf{H} \to \mathbf{G}$. By Corollary 4.8(2), Γ_k is a reduction from $\mathrm{PCSP}(\Omega_k \mathbf{H}, \Omega_k \mathbf{G})$ to $\mathrm{PCSP}(\mathbf{H}, \mathbf{G})$. To conclude that $\mathrm{PCSP}(\mathbf{H}, \mathbf{G})$ is NP-hard it remains to show that $\Omega_k \mathbf{H}$ is nonbipartite. Observe that since \mathbf{H} is nonbipartite and k is odd then also $\Lambda_k \mathbf{H}$ is nonbipartite. Furthermore, from the above lemma, $\Lambda_k \mathbf{H} \to \Omega_k \mathbf{H}$ which implies that $\Omega_k \mathbf{H}$ is also nonbipartite.

Thus (by the above lemma and homomorphic relaxation) if we know one graph \mathbf{G}' such that $\mathrm{PCSP}(\mathbf{H}', \mathbf{G}')$ is NP-hard for all nonbipartite \mathbf{H}' , then we can conclude the same for all \mathbf{G} such that $\Omega_k \mathbf{G} \to \mathbf{G}'$ for some odd k. When does such a k exists? The answer, given in [71], turns out to be topological. We remark that the results in [71] use the so-called box complex of \mathbf{G} instead of $\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})$. However, there exist \mathbb{Z}_2 -maps (in both directions) between the two complexes; see maps (M2), (M3), and (M7) in [58, Proposition 4] for explicit maps. This is enough for our purposes, but a stronger claim is true—the two complexes are \mathbb{Z}_2 -homotopy equivalent (as defined in Appendix A) [24, 77].

Intuitively, while the operation Γ_k gives a "thicker" graph, the operation Ω_k gives a "thinner" one. In fact, Ω_k behaves similarly to barycentric subdivision in topology: it preserves the topology of a graph (formally, $|\mathsf{Hom}(\mathbf{K}_2, \Omega_k \mathbf{G})|$ is \mathbb{Z}_2 -homotopy equivalent to $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})|$ [71]) but refines its geometry. With increasing k, this eventually allows one to model any continuous map with a graph homomorphism; in particular we have the following.

THEOREM 4.31 ([71]). There exists a \mathbb{Z}_2 -map $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{H})| \to_{\mathbb{Z}_2} |\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})|$ if and only if $\Omega_k \mathbf{H} \to \mathbf{G}$ for some odd k.

Copyright (C) by SIAM. Unauthorized reproduction of this article is prohibited.

We now conclude the proof of whether the Brakensiek–Guruswami conjecture holds for a graph **G** (and all relevant **H**) depends only on the topology of **G**—this was informally stated earlier in Theorem 1.7). In fact, it only matters which \mathbb{Z}_2 -maps $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})|$ admits.

THEOREM 4.32. Let \mathbf{G} , \mathbf{G}' be graphs such that $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})|$ admits a \mathbb{Z}_2 -map to $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{G}')|$ and suppose $\mathsf{PCSP}(\mathbf{H}, \mathbf{G}')$ is NP-hard for all nonbipartite graphs \mathbf{H} such that $\mathbf{H} \to \mathbf{G}'$. Then $\mathsf{PCSP}(\mathbf{H}, \mathbf{G})$ is NP-hard for all nonbipartite graphs \mathbf{H} such that $\mathbf{H} \to \mathbf{G}$.

Proof. By Theorem 4.31, $\Omega_k \mathbf{G} \to \mathbf{G}'$ for some odd k. Since $\text{PCSP}(\mathbf{H}, \mathbf{G}')$ is NP-hard for all nonbipartite graphs \mathbf{H} , we also have that $\text{PCSP}(\mathbf{H}, \Omega_k \mathbf{G})$ is NP-hard by a trivial reduction. Now, Lemma 4.30 gives the claim.

In particular, Theorem 4.32 implies that Theorems 1.3 and 1.4 are equivalent.

5. Conclusion. We presented two new methodologies, based on topology and adjunction, to analyze the complexity of PCSPs and provided some applications of these methodologies to considerably improve state-of-the-art in the complexity of approximate graph coloring and promise graph homomorphism problems.

As mentioned before, there are many ways in which topology can potentially be applied in the analysis of polymorphisms from **H** to **G**, for graphs or for general relational structures. With the approach that we used, we made a few choices for our analysis. Specifically, we used (a) the graph \mathbf{K}_2 to construct simplicial complexes $Hom(K_2, H)$ and $Hom(K_2, G)$, (b) \mathbb{Z}_2 -action on our complexes and topological spaces, and (c) fundamental groups of topological spaces. One can try to perform similar analyses, by (a) replacing \mathbf{K}_2 by any other graph \mathbf{K} (or, for general PCSPs, by another appropriate structure), (b) using any subgroup of the automorphism group of **K** to account for symmetry (called "equivariance" in the topological literature) of the complexes and topological spaces, and (c) replacing the fundamental group with a different topological invariant of spaces or continuous functions involved. Some examples of different choices, though not in the context of polymorphisms, can be found, e.g., in [53, 59]. These are the obvious first choices of adapting the approach. Naturally, it can be changed in a more fundamental way: the most prominent seems to be directly analyzing the topological structure of the simplicial complexes $\mathsf{Hom}(\mathbf{H}^n, \mathbf{G})$ (see [53, section 9.2.4 for related general suggestions). One advantage of this approach is that the analysis would depend only on the function minion $Pol(\mathbf{H}, \mathbf{G})$, rather than on the specific choice of **H** and **G**.

In this paper, we applied topology together with the algebraic theory from [5] to prove complexity results about promise graph homomorphism. However, our application can be seen as plugging the topological analysis into an algebraic result. Since topology appears to be naturally present in minions of polymorphisms, it would be interesting to further develop the algebraic theory from [5] to properly incorporate topology. Similarly, we used adjunction to obtain some reductions for approximate graph coloring problems that provably cannot be captured by the algebraic theory from [5]—it is natural to ask whether a more general theory can be constructed that incorporates both the current algebraic theory and adjunction.

It would be interesting to find further specific applications of our methodologies, for example, in approximate graph and hypergraph coloring and their variations, or possibly even beyond constraint satisfaction. Naturally, one would want to extend our methodologies to approximate graph coloring problems $PCSP(\mathbf{K}_3, \mathbf{K}_c)$ or promise graph homomorphism problems $PCSP(\mathbf{C}_k, \mathbf{K}_c)$ for $c \geq 4$. If one applies the same transformation of these graph problems into homomorphism complexes and topological spaces as we use in this paper, one would need to analyze the (\mathbb{Z}_2 -)polymorphisms from \mathcal{S}^1 to \mathcal{S}^m for $m \geq 2$. (Note that $\pi_1(\mathcal{S}^m)$ is trivial for $m \geq 2$, so the fundamental group is of no use in this case.) These polymorphisms are \mathbb{Z}_2 -maps from tori \mathcal{T}^n , $n \geq 1$, to \mathcal{S}^m . This is somewhat related to some well-known hard open questions from algebraic topology, such as classification of maps from one sphere \mathcal{S}^{m_1} to another \mathcal{S}^{m_2} . However, to the best of our knowledge, the equivariant version of such questions was never considered. Moreover, for our purposes, it would suffice to get any classification of Pol($\mathcal{S}^1, \mathcal{S}^m$) that can be connected with the algebraic theory, e.g., with Theorem 2.12 above, or with [5, Theorem 5.22] or [19, Corollary 4.2]. Of course, it is possible that some modification of our approach will need to be used. In any case, we believe that topology will play an important part in settling the complexity of approximate graph coloring and the Brakensiek–Guruswami conjecture.

Finally, we remark that the standard reductions from the algebraic approach, i.e., reductions of the form Λ_{ϕ} (see section 4.1), can be thought of as replacing individual constraints in an instance with gadgets (possibly consisting of many constraints). Similarly, certain reductions of the form Γ_{ϕ} , such as δ and Γ_k presented in sections 4.3 and 4.4, can be thought of as replacing gadgets (i.e., certain parts of input) with individual constraints. The latter turned out to be particularly useful when they themselves admit some right adjoint Ω (as was the case for δ and Γ_k); however, such reductions must have a rather restricted form [33, Theorem 2.5]. Thus, a natural extension would be to investigate reductions which replace gadgets with gadgets (i.e., introduce a copy of one gadget for each homomorphism from another gadget).

Appendix A. Equivalence of homomorphism complexes. There is a superficial distinction between the abstract simplicial complex Hom(H, G) as we defined it and the definition of the homomomorphism complex in [4, 53]. We will show that the two definitions give topological spaces that are equivalent in the following sense.

DEFINITION A.1. Two \mathbb{Z}_2 -spaces \mathcal{X}, \mathcal{Y} are \mathbb{Z}_2 -homotopy equivalent if there are \mathbb{Z}_2 -maps $\alpha: \mathcal{X} \to \mathcal{Y}$ and $\beta: \mathcal{Y} \to \mathcal{X}$ such that $\beta \alpha$ and $\alpha \beta$ are \mathbb{Z}_2 -homotopic (recall Definition 3.15) to the identity maps on \mathcal{X} and \mathcal{Y} , respectively.

This notion is coarser than \mathbb{Z}_2 -homeomorphism (which required fg and gf to be equal to identity maps); for example, $\mathbb{R}^2 \setminus \{(0,0)\}$ is \mathbb{Z}_2 -homotopy equivalent to S^1 but not \mathbb{Z}_2 -homeomorphic to it. Nevertheless, \mathbb{Z}_2 -homotopy equivalent spaces admit the same \mathbb{Z}_2 -maps and have isomorphic fundamental groups, for example, they are thus indeed equivalent for our purposes.

We remark that in [53] and other topological literature, very little attention is given to the distinction between abstract (simplicial) complexes and their geometric realizations. In particular, in [4, 53], the Hom complex of graphs \mathbf{H} , \mathbf{G} is defined as a topological space and not an abstract simplicial complex (it is in fact a so-called prodsimplicial complex; see [53, p. 28]). The following definition is an equivalent formulation of Definition 9.23 in [53] using the comments thereafter.

DEFINITION A.2. For a set V, we denote by Δ^V the standard simplex with vertices V that is defined as a subspace of \mathbb{R}^V , where the canonical unit vectors are identified with elements of V, obtained as the convex hull of V, i.e., $\Delta^V = \{\sum_{v \in V} \lambda_v v \mid \lambda_v \in [0,1] \text{ for each } v \in V, \text{ and } \sum_{v \in V} \lambda_v = 1\}.$ Hom^{prod}(\mathbf{K}_2, \mathbf{G}) is a subspace of $\Delta^{V(G)} \times \Delta^{V(G)}$. Thus a point of this space is de-

Hom^{prod}(**K**₂, **G**) is a subspace of $\Delta^{V(G)} \times \Delta^{V(G)}$. Thus a point of this space is described as a pair of formal sums ($\sum_{u \in V(G)} \lambda_u u, \sum_{v \in V(G)} \rho_v v$) such that $\sum_{u \in V(G)} \lambda_u u = \sum_{v \in V(G)} \rho_v v = 1$. Using this description, Hom^{prod}(**K**₂, **G**) is defined as the subspace



FIG. A.1. $\mathsf{Hom}^{prod}(\mathbf{K}_2, \mathbf{K}_4)$.

consisting of those points $(\sum_{u \in V(G)} \lambda_u u, \sum_{v \in V(G)} \rho_v v)$ such that $\{u \mid \lambda_u > 0\} \times \{v \mid \rho_v > 0\} \subseteq E(G)$ is a complete bipartite graph. The action of \mathbb{Z}_2 on this complex is given by switching the two coordinates, i.e., $-(\sum_{u \in V(G)} \lambda_u u, \sum_{v \in V(G)} \rho_v v) = (\sum_{v \in V(G)} \rho_v v, \sum_{u \in V(G)} \lambda_u u).$

On the other hand, the geometric realization of our $\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})$ can be described in the following way. It consists of convex combinations $\sum_{(u,v)\in E(G)} \lambda_{u,v}(u,v)$ (the points $(u,v)\in E(G)$ are identified with certain unit vectors in \mathbb{R}^n , where n = |E(G)|/2so that -(u,v) = (v,u)) such that $\{u \mid \lambda_{u,v} > 0\} \times \{v \mid \lambda_{u,v} > 0\} \subseteq E(G)$ is a complete bipartite subgraph.

Both complexes are therefore defined using complete bipartite subgraphs, and the similarity is apparent. As an example that highlights the small differences, let us note that $\mathsf{Hom}^{\mathrm{prod}}(\mathbf{K}_2, \mathbf{K}_4)$ is \mathbb{Z}_2 -homeomorphic to \mathcal{S}^2 ; the space is depicted in Figure A.1. This is since, unlike in $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{K}_4)|$, the tetragonal faces on the picture correspond to actual squares in $\mathsf{Hom}^{\mathrm{prod}}(\mathbf{K}_2, \mathbf{K}_4)$.

PROPOSITION A.3. Let **H**, **G** be graphs. Then the \mathbb{Z}_2 -spaces $|\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})|$ and $\mathsf{Hom}^{\mathrm{prod}}(\mathbf{K}_2, \mathbf{G})$ are \mathbb{Z}_2 -homotopy equivalent.

Proof. We define continuous maps α and β between the two spaces that witness the \mathbb{Z}_2 -homotopy equivalence.

$$\begin{array}{ccc} |\mathsf{Hom}(\mathbf{K}_{2},\mathbf{G})| & \mathsf{Hom}^{\mathrm{prod}}(\mathbf{K}_{2},\mathbf{G}) \\ \alpha \colon & \sum_{(u,v)\in E(G)}\lambda_{(u,v)}(u,v) & \mapsto & \left(\sum_{(u,v)\in E(G)}\lambda_{(u,v)}u,\sum_{(u,v)\in E(G)}\lambda_{(u,v)}v\right) \\ & \sum_{(u,v)\in E(G)}\lambda_{u}\rho_{v}(u,v) & \left(\sum_{u\in V(G)}\lambda_{u}u,\sum_{v\in V(G)}\rho_{v}v\right) & :\beta. \end{array}$$

It is straightforward to check that α and β are \mathbb{Z}_2 -maps, that $\alpha\beta = \mathrm{id}$, and that $\beta\alpha$ maps each point to a point in the same face. Thus, a homotopy from $\beta\alpha$ to id can be defined by linearly interpolating between the two: $(p,t) \mapsto (1-t)p + t\beta(\alpha(p))$ for $p \in |\mathsf{Hom}(\mathbf{K}_2, \mathbf{G})|$ and $t \in [0, 1]$.

We remark that the above proposition and its proof generalizes to arbitrary Hom complexes $\mathsf{Hom}(\mathbf{H}, \mathbf{G})$ with the action of $\operatorname{Aut}(\mathbf{H})$, more precisely, the complex $\mathsf{Hom}^{\operatorname{prod}}(\mathbf{H}, \mathbf{G})$ (as defined in [53, Definition 9.23]) is $\operatorname{Aut}(\mathbf{H})$ -homotopy equivalent to $|\mathsf{Hom}(\mathbf{H}, \mathbf{G})|$ for any two graphs \mathbf{H}, \mathbf{G} .

Acknowledgments. A. K. and J. O. would like to thank John Hunton for consultations on algebraic topology and Libor Barto and Antoine Mottet for inspiring discussions.

REFERENCES

- P. AUSTRIN, A. BHANGALE, AND A. POTUKUCHI, Simplified Inapproximability of Hypergraph Coloring via t-Agreeing Families, preprint, arXiv:1904.01163, 2019.
- [2] P. AUSTRIN, A. BHANGALE, AND A. POTUKUCHI, Improved inapproximability of rainbow coloring, in Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'20), SIAM, Philadelphia, 2020, pp. 1479–1495, https://doi.org/10.1137/ 1.9781611975994.90.
- [3] P. AUSTRIN, V. GURUSWAMI, AND J. HÅSTAD, (2+ε)-Sat is NP-hard, SIAM J. Comput., 46 (2017), pp. 1554–1573, https://doi.org/10.1137/15M1006507.
- [4] E. BABSON AND D. N. KOZLOV, Complexes of graph homomorphisms, Israel J. Math., 152 (2006), pp. 285–312, https://doi.org/10.1007/BF02771988.
- [5] L. BARTO, J. BULÍN, A. KROKHIN, AND J. OPRŠAL, Algebraic approach to promise constraint satisfaction, J. ACM, 68 (2021), 28, https://doi.org/10.1145/3457606.
- [6] L. BARTO AND M. KOZIK, Robustly solvable constraint satisfaction problems, SIAM J. Comput., 45 (2016), pp. 1646–1669, https://doi.org/10.1137/130915479.
- [7] L. BARTO, A. KROKHIN, AND R. WILLARD, Polymorphisms, and how to use them, in The Constraint Satisfaction Problem: Complexity and Approximability, Dagstuhl Follow-Ups 7, Schloss Dagstuhl, Dagstuhl, Germany, 2017, pp. 1–44, https://doi.org/10.4230/DFU.Vol7.15301.1.
- [8] L. BARTO, J. OPRŠAL, AND M. PINSKER, The wonderland of reflections, Israel J. Math., 223 (2018), pp. 363–398, https://doi.org/10.1007/s11856-017-1621-9.
- [9] A. BHANGALE, NP-hardness of coloring 2-colorable hypergraph with poly-logarithmically many colors, in Proceedings of the 45th International Colloquium on Automata, Languages, and Programming (ICALP'18), LIPIcs, Schloss Dagstuhl, Wadern, Germany, 2018, 15, https://doi.org/10.4230/LIPIcs.ICALP.2018.15.
- [10] V. G. BODNARCHUK, L. A. KALUZHNIN, V. N. KOTOV, AND B. A. ROMOV, Galois theory for post algebras. I, Cybernetics, 5 (1969), pp. 243–252.
- [11] V. G. BODNARCHUK, L. A. KALUZHNIN, V. N. KOTOV, AND B. A. ROMOV, Galois theory for post algebras. II, Cybernetics, 5 (1969), pp. 531–539.
- [12] K. BORSUK, Drei Sätze über die n-dimensionale euklidische Sphäre, Fund. Math., 20 (1933), pp. 177–190.
- [13] J. BRAKENSIEK AND V. GURUSWAMI, New hardness results for graph and hypergraph colorings, in Proceedings of the 31st Conference on Computational Complexity (CCC 2016), LIPIcs Leibniz Int. Proc. Inform. 50, Schloss Dagstuhl, Wadern, Germany, 2016, 14, https://doi.org/10.4230/LIPIcs.CCC.2016.14.
- [14] 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, SIAM, Philadelphia, 2019, pp. 436–455, https://doi.org/10.1137/1.9781611975482.28.
- [15] J. BRAKENSIEK AND V. GURUSWAMI, Symmetric polymorphisms and efficient decidability of promise CSPs, in Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, SIAM, Philadelphia, 2020, pp. 297–304, https://doi.org/10.1137/1.9781611975994.18.
- [16] J. BRAKENSIEK AND V. GURUSWAMI, Promise constraint satisfaction: Algebraic structure and a symmetric Boolean dichotomy, SIAM J. Comput., 50 (2021), pp. 1663–1700, https://doi.org/10.1137/19M128212X.
- [17] J. BRAKENSIEK AND V. GURUSWAMI, The quest for strong inapproximability results with perfect completeness, ACM Trans. Algorithms, 17 (2021), 17, https://doi.org/10.1145/3459668.
- [18] J. BRAKENSIEK, V. GURUSWAMI, M. WROCHNA, AND S. ŽIVNÝ, The power of the combined basic linear programming and affine relaxation for promise constraint satisfaction problems, SIAM J. Comput., 49 (2020), pp. 1232–1248, https://doi.org/10.1137/20M1312745.
- [19] A. BRANDTS, M. WROCHNA, AND S. ŽIVNÝ, The complexity of promise SAT on non-Boolean domains, ACM Trans. Comput. Theory, 13 (2021), 26, https://doi.org/10.1145/3470867.
- [20] R. BRICEÑO, A. A. BULATOV, V. DALMAU, AND B. LAROSE, Dismantlability, connectedness, and mixing in relational structures, J. Combin. Theory Ser. B, 147 (2021), pp. 37–70, https://doi.org/10.1016/j.jctb.2020.10.001.
- [21] A. BULATOV, P. JEAVONS, AND A. KROKHIN, Classifying the complexity of constraints using finite algebras, SIAM J. Comput., 34 (2005), pp. 720–742, https://doi.org/10.1137/S0097539700376676.
- [22] A. A. BULATOV, A dichotomy theorem for nonuniform CSPs, in Proceedings of the 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS), IEEE Computer Society, Los Alamitos, CA, 2017, pp. 319–330, https://doi.org/10.1109/FOCS.2017.37.

- [23] J. BULÍN, A. KROKHIN, AND J. OPRŠAL, Algebraic approach to promise constraint satisfaction, in Proceedings of the 51st Annual ACM SIGACT Symposium on the Theory of Computing (STOC '1409), ACM, New York, 2019, pp. 602–613, https://doi.org/10.1145/3313276.3316300.
- [24] P. CSORBA, On the simple Z₂-homotopy types of graph complexes and their simple Z₂universality, Canad. Math. Bull., 51 (2008), pp. 535–544, https://doi.org/10.4153/CMB-2008-053-9.
- [25] I. DINUR, The PCP theorem by gap amplification, J. ACM, 54 (2007), 12, https://doi.org/10.1145/1236457.1236459.
- [26] I. DINUR, S. KHOT, G. KINDLER, D. MINZER, AND M. SAFRA, Towards a proof of the 2-to-1 games conjecture?, in Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC'18), ACM, New York, 2018, pp. 376–389, https://doi.org/10.1145/3188745.3188804.
- [27] I. DINUR, E. MOSSEL, AND O. REGEV, Conditional hardness for approximate coloring, SIAM J. Comput., 39 (2009), pp. 843–873, https://doi.org/10.1137/07068062X.
- [28] I. DINUR, O. REGEV, AND C. SMYTH, The hardness of 3-uniform hypergraph coloring, Combinatorica, 25 (2005), pp. 519–535, https://doi.org/10.1007/s00493-005-0032-4.
- [29] I. DINUR AND I. SHINKAR, On the conditional hardness of coloring a 4-colorable graph with super-constant number of colors, in Proceedings of the 13th International Workshop on Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques and the 14th International Workshop on Randomization and Computation (APPROX-RANDOM'10), Lecture Notes in Comput. Sci. 6302, Springer, 2010, pp. 138– 151, https://link.springer.com/chapter/10.1007/978-3-642-15369-3_11.
- [30] 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., 28 (1998), pp. 57–104, https://doi.org/10.1137/S0097539794266766.
- [31] M. FICAK, M. KOZIK, M. OLŠÁK, AND S. STANKIEWICZ, Dichotomy for symmetric Boolean PCSPs, in Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP'19), LIPIcs Leibniz Int. Proc. Inform. 132, Schloss Dagstuhl, Wadern, Germany, 2019, 57, https://doi.org/10.4230/LIPIcs.ICALP.2019.57.
- [32] J. FONIOK AND C. TARDIF, Adjoint Functors in Graph Theory, preprint, arXiv:1304.2215, 2013.
- [33] J. FONIOK AND C. TARDIF, Digraph functors which admit both left and right adjoints, Discrete Math., 338 (2015), pp. 527–535, https://doi.org/10.1016/j.disc.2014.10.018.
- [34] J. FONIOK AND C. TARDIF, Hedetniemi's conjecture and adjoint functors in thin categories, Appl. Categ. Structures, 26 (2018), pp. 113–128, https://doi.org/10.1007/s10485-017-9484-0.
- [35] M. R. GAREY AND D. S. JOHNSON, The complexity of near-optimal graph coloring, J. ACM, 23 (1976), pp. 43–49, https://doi.org/10.1145/321921.321926.
- [36] D. GEIGER, Closed systems of functions and predicates, Pacific J. Math., 27 (1968), pp. 95–100.
- [37] D. L. GREENWELL AND L. LOVÁSZ, Applications of product colouring, Acta Math. Acad. Sci. Hungar., 25 (1974), pp. 335–340, https://doi.org/10.1007/BF01886093.
- [38] V. GURUSWAMI AND S. KHANNA, On the hardness of 4-coloring a 3-colorable graph, SIAM J. Discrete Math., 18 (2004), pp. 30–40, https://doi.org/10.1137/S0895480100376794.
- [39] V. GURUSWAMI AND E. LEE, Strong inapproximability results on balanced rainbow-colorable hypergraphs, Combinatorica, 38 (2018), pp. 547–599, https://doi.org/10.1007/s00493-016-3383-0.
- [40] V. GURUSWAMI AND S. SANDEEP, d-to-1 hardness of coloring 3-colorable graphs with O(1) colors, in Proceedings of the 47th International Colloquium on Automata, Languages, and Programming (ICALP 2020), Leibniz Int. ProcInform 168, Dagstuhl, Germany, 2020, Schloss Dagstuhl, Wadern, Germany, 2020, 62, https://doi.org/10.4230/LIPIcs.ICALP. 2020.62.
- [41] V. GURUSWAMI AND S. SANDEEP, Rainbow coloring hardness via low sensitivity polymorphisms, SIAM J. Discrete Math., 34 (2020), pp. 520–537, https://doi.org/10.1137/19M127731X.
- [42] C. C. HARNER AND R. C. ENTRINGER, Arc colorings of digraphs, J. Combin. Theory Ser. B, 13 (1972), pp. 219–225, https://doi.org/10.1016/0095-8956(72)90057-3.
- [43] A. HATCHER, Algebraic Topology, Cambridge University Press, Cambridge, 2001.
- [44] P. HELL AND J. NEŠETŘIL, On the complexity of H-coloring, J. Combin. Theory Ser. B, 48 (1990), pp. 92–110, https://doi.org/10.1016/0095-8956(90)90132-J.
- [45] P. HELL AND J. NEŠETŘIL, Graphs and Homomorphisms, Oxford University Press, Oxford, 2004.

- [46] 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, Springer, Berlin, 2013, pp. 233–243, https://doi.org/10.1007/ 978-3-642-40328-6_17.
- [47] P. JEAVONS, D. A. COHEN, AND M. GYSSENS, Closure properties of constraints, J. ACM, 44 (1997), pp. 527–548, https://doi.org/10.1145/263867.263489.
- [48] K. KAWARABAYASHI AND M. THORUP, Coloring 3-colorable graphs with less than n^{1/5} colors, J. ACM, 64 (2017), 4, https://doi.org/10.1145/3001582.
- [49] S. KHANNA, N. LINIAL, AND S. SAFRA, On the hardness of approximating the chromatic number, Combinatorica, 20 (2000), pp. 393–415, https://doi.org/10.1007/s004930070013.
- [50] S. KHOT, Improved inaproximability results for maxclique, chromatic number and approximate graph coloring, in Proceedings of the 42nd Annual Symposium on Foundations of Computer Science (FOCS'01), IEEE Computer Society, Los Alamitos, CA, 2001, pp. 600–609, https://doi.org/10.1109/SFCS.2001.959936.
- [51] S. KHOT, On the power of unique 2-prover 1-round games, in Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC'02), ACM, New York, 2002, pp. 767– 775, https://doi.org/10.1145/509907.510017.
- [52] S. KHOT, D. MINZER, AND M. SAFRA, Pseudorandom sets in Grassmann graph have nearperfect expansion, in Proceedings of the 59th IEEE Annual Symposium on Foundations of Computer Science (FOCS'18), IEEE Computer Society, Los Alamitos, CA, 2018, pp. 592–601, https://doi.org/10.1109/FOCS.2018.00062.
- [53] D. KOZLOV, Combinatorial Algebraic Topology, Algorithms Comput. Math. 21, Springer, Berlin, 2008, https://doi.org/10.1007/978-3-540-71962-5.
- [54] A. KROKHIN AND J. OPRŠAL, The complexity of 3-colouring H-colourable graphs, in Proceedings of the 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS'19), IEEE, Piscataway, NJ, 2019, pp. 1227–1239, https://doi.org/10.1109/FOCS.2019.00076.
- [55] A. KROKHIN AND S. ŽIVNÝ, The Constraint Satisfaction Problem: Complexity and Approximability, Dagstuhl Follow-Ups 7, Schloss Dagstuhl, Wadern, Germany 2017, https://doi.org/10.4230/DFU.Vol7.15301.
- [56] B. LAROSE, C. LOTEN, AND C. TARDIF, A characterisation of first-order constraint satisfaction problems, Log. Methods Comput. Sci., 3 (2007), https://doi.org/10.2168/LMCS-3(4:6)2007.
- [57] L. LOVÁSZ, Kneser's conjecture, chromatic number, and homotopy, J. Combin. Theory Ser. A, 25 (1978), pp. 319–324, https://doi.org/10.1016/0097-3165(78)90022-5.
- [58] J. MATOUŠEK AND G. ZIEGLER, Topological lower bounds for the chromatic number: A hierarchy, Jahresber. Dtsch. Math.-Ver., 106 (2004), pp. 71–90.
- [59] J. MATOUŠEK, Using the Borsuk-Ulam Theorem, Lectures on Topological Methods in Combinatorics and Geometry, 1st ed., Springer, Berlin, 2003, https://doi.org/10.1007/978-3-540-76649-0.
- [60] T. MATSUSHITA, Box complexes and homotopy theory of graphs, Homology Homotopy Appl., 19 (2017), pp. 175–197.
- [61] T. MATSUSHITA, Z₂-indices and Hedetniemi's conjecture, Discrete Comput. Geom., 62 (2019), pp. 662–673, https://doi.org/10.1007/s00454-019-00090-1.
- [62] M. OLŠÁK, Loop conditions for strongly connected digraphs, Internat. J. Algebra Comput., 30 (2020), pp. 467–499, https://doi.org/10.1142/S0218196720500083.
- [63] N. PIPPENGER, Galois theory for minors of finite functions, Discrete Math., 254 (2002), pp. 405–419, https://doi.org/10.1016/S0012-365X(01)00297-7.
- [64] S. POLJAK, Coloring digraphs by iterated antichains, Comment. Math. Univ. Carolin., 32 (1991), pp. 209–212, http://eudml.org/doc/247261.
- [65] S. POLJAK AND V. RÖDL, On the arc-chromatic number of a digraph, J. Combin. Theory Ser. B, 31 (1981), pp. 190–198, https://doi.org/10.1016/S0095-8956(81)80024-X.
- [66] O. REINGOLD, Undirected connectivity in log-space, J. ACM, 55 (2008), 17, https://doi.org/10.1145/1391289.1391291.
- [67] Y. SHITOV, Counterexamples to Hedetniemi's conjecture, Ann. Math. (2), 190 (2019), pp. 663– 667, https://doi.org/10.4007/annals.2019.190.2.6.
- [68] C. TARDIF, Hedetniemi's conjecture, 40 years later, Graph Theory Notes NY, 54 (2008), pp. 46–57, https://mast.queensu.ca/~ctardif/articles/gtn5406rp.pdf.
- [69] C. TARDIF AND M. WROCHNA, Hedetniemi's conjecture and strongly multiplicative graphs, SIAM J. Discrete Math., 33 (2019), pp. 2218–2250, https://doi.org/10.1137/19M1245013.
- [70] M. WROCHNA, Square-free graphs are multiplicative, J. Combin. Theory Ser. B, 122 (2017), pp. 479–507, https://doi.org/10.1016/j.jctb.2016.07.007.

- [71] M. WROCHNA, On inverse powers of graphs and topological implications of Hedetniemi's conjecture, J. Combin. Theory Ser. B, 139 (2019), pp. 267–295, https://doi.org/10.1016/j.jctb.2019.02.008.
- [72] M. WROCHNA, Homomorphism reconfiguration via homotopy, SIAM J. Discrete Math., 34 (2020), pp. 328–350, https://doi.org/10.1137/17M1122578.
- [73] M. WROCHNA AND S. ŽIVNÝ, Improved hardness for H-colourings of G-colourable graphs, in Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'20), SIAM, Philadelphia, 2020, pp. 1426–1435, https://doi.org/10.1137/1.9781611975994.86.
- [74] X. ZHU, A survey on Hedetniemi's conjecture, Taiwanese J. Math., 2 (1998), pp. 1–24, http://www.jstor.org/stable/43834380.
- [75] D. ZHUK, A proof of CSP dichotomy conjecture, in Proceedings of the 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS'17), IEEE Computer Society, Los Alamitos, CA, 2017, pp. 331–342, https://doi.org/10.1109/FOCS.2017.38.
- [76] D. ZHUK, A proof of the CSP dichotomy conjecture, J. ACM, 67 (2020), 30, https://doi.org/10.1145/3402029.
- [77] R. T. ŽIVALJEVIĆ, WI-posets, graph complexes and Z₂-equivalences, J. Combin. Theory Ser. A, 111 (2005), pp. 204–223, https://doi.org/10.1016/j.jcta.2004.12.002.