Improved hardness for H-colourings of G-colourable graphs*

Marcin Wrochna and Stanislav Živný†

Abstract

We present new results on approximate colourings of graphs and, more generally, approximate H-colourings and promise constraint satisfaction problems.

First, we show NP-hardness of colouring k-colourable graphs with $\binom{k}{\lfloor k/2 \rfloor} - 1$ colours for every $k \geq 4$. This improves the result of Bulín, Krokhin, and Opršal [STOC'19], who gave NP-hardness of colouring k-colourable graphs with $2k-1$ colours for $k \geq 3$, and the result of Huang [APPROX-RANDOM'13], who gave NP-hardness of colouring k-colourable graphs with $2^{\Omega(k^{1/3})}$ colours for sufficiently large k. Thus, for $k \geq 4$, we improve from known linear/sub-exponential gaps to exponential gaps.

Second, we show that the topology of the box complex of H alone determines whether H-colouring of G-colourable graphs is NP-hard for all (non-bipartite, H-colourable) G. This formalises the topological intuition behind the result of Krokhin and Opršal [FOCS'19] that 3-colouring of G-colourable graphs is NP-hard for all (3-colourable, non-bipartite) G. We use this technique to establish NP-hardness of H-colouring of G-colourable graphs for H that include but go beyond K_3, including square-free graphs and circular cliques (leaving K_4 and larger cliques open).

Underlying all of our proofs is a very general observation that adjoint functors give reductions between promise constraint satisfaction problems.

The full version [55] containing detailed proofs is available at https://arxiv.org/abs/1907.00872.

1 Introduction

Graph colouring is one of the most fundamental and studied problems in combinatorics and computer science. A graph G is called k-colourable if there is an assignment of colours $\{1, 2, \ldots, k\}$ to the vertices of G so that any two adjacent vertices are assigned different colours. The chromatic number of G, denoted by $\chi(G)$, is the smallest integer k for which G is k-colourable. Deciding whether $\chi(G) \leq k$ appeared on Karp's original list of 21 NP-complete problems [35], and is NP-hard for every $k \geq 3$. In particular, it is NP-hard to decide whether $\chi(G) \leq 3$ or $\chi(G) > 3$. Put differently (thanks to self-reducibility of graph colouring), it is NP-hard to find a 3-colouring of G even if G is promised to be 3-colourable.

In the *approximate graph colouring* problem, we are allowed to use more colours than needed. For instance, given a 3-colourable graph G on n vertices, can we find a colouring of G using significantly fewer than n colours? On the positive side, the currently best polynomial-time algorithm of Kawarabayashi and Thorup [36] finds a colouring using $O(n^{0.19996})$ colours. Their work continues a long line of research and is based on a semidefinite relaxation. On the negative side, it is believed that finding a c-colouring of a k-colourable graph is NP-hard for all constants $3 \leq k \leq c$. Already in this regime (let alone for non-constant c) our understanding remains rather limited, despite lots of work and the development of complex techniques, as we will survey in Section 1.1.

A natural and studied generalisation of graph colourings is that of graph homomorphisms and, more generally, constraint satisfaction problems [32].

Given two graphs G and H, a map $h : V(G) \to V(H)$ is a *homomorphism* from G to H if h preserves edges; that is, if $\{h(u), h(v)\} \in E(H)$ whenever $\{u, v\} \in E(G)$ [31]. A celebrated result of Hell and Nešetřil established a dichotomy for the homomorphism problem with a fixed target graph H, also known as the *H-colouring problem*: deciding whether an input graph G has a homomorphism to H is solvable in polynomial time if H is bipartite or if H has a loop; for all other H this problem is NP-hard [30]. Note that the H-colouring problem for $H = K_k$, the complete graph on k vertices, is precisely the graph colouring problem with k colours.

The constraint satisfaction problem (CSP) is a generalisation of the graph homomorphism problem from graphs to arbitrary relational structures. One type of CSP that has attracted a lot of attention is the one with a fixed target structure, also known as the *non-uniform* CSP; see, e.g., the work of Jeavons, Cohen, and Gyssens [34], Bulatov [14, 16], and Barto and Kozik [5, 6]. Following the above mentioned dichotomy of Hell and Nešetřil for the H-colouring [30] and a dichotomy result of Schaefer for Boolean CSPs [48], Feder and Vardi famously conjectured a dichotomy for all non-uniform CSPs [23]. The Feder-Vardi conjecture was

*Stanislav Živný was supported by a Royal Society University Research Fellowship. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 714532). 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.

†University of Oxford, UK. marcin.wrochna@cs.ox.ac.uk, standa.zivny@cs.ox.ac.uk

recently confirmed independently by Bulatov [17] and Zhuk [57]. In fact, both proofs establish the so-called "algebraic dichotomy", conjectured by Bulatov, Jeavons, and Krokhin [15], which delineates the tractability boundary in algebraic terms. A high-level idea of the tractability boundary is that of higher-order symmetries, called polymorphisms, which allow to combine several solutions to a CSP instance into a new solution. The lack of non-trivial[1] polymorphisms guarantees NP-hardness, as shown already in [15]. The work of Bulatov and Zhuk show that *any* non-trivial polymorphism guarantees tractability. We refer the reader to a recent accessible survey by Barto, Krokhin, and Willard on the algebraic approach to CSPs [7].

Given two graphs G an H such that G is H-colourable (i.e., there is a homomorphism from G to H), the *promise constraint satisfaction problem* parametrised by G and H, denoted by $\mathrm{PCSP}(G, H)$, is the following computational problem: given a G-colourable graph, find an H-colouring of this graph.[2] More generally, G and H do not have to be graphs but arbitrary relational structures. Note that if $G = H$ then we obtain the (search version of the) standard H-colouring and constraint satisfaction problem.

PCSPs have been studied as early as in the classic work of Garey and Johnson [26] on approximate graph colouring but a systematic study originated in the paper of Austrin, Guruswami, and Håstad [3], who studied a promise version of $(2k+1)$-SAT, called $(2+\epsilon)$-SAT. In a series of papers [10, 12, 13], Brakensiek and Guruswami linked PCSPs to the universal-algebraic methods developed for the study of non-uniform CSPs [7]. In particular, the notion of weak polymorphisms, identified in [3], allowed for some ideas developed for CSPs to be be used in the context of PCSPs. The algebraic theory of PCSPs was then lifted to an abstract level by Bulín, Krokhin, and Opršal in [18]. Consequently, this theory was used by Ficak, Kozik, Olšák, and Stankiewicz to obtain a dichotomy for symmetric Boolean PCSPs [24], thus improving on an earlier result from [12], which gave a dichotomy for symmetric Boolean PCSP with folding (negations allowed).

[1] We note that projections/dictators are not the only trivial polymorphims, cf. [7, Example 41].

[2] What we described is the "search version" of PCSPs. In the "decision version", the goal is to say YES if the input graph is G-colourable and NO if the input graph is not H-colourable. The decision PCSP reduces to the search PCSP but they are not known to be equivalent in general. However, as far as we know, all known positive results are for the search version, while all known negative results, including the new results from this paper, are for the decision version.

1.1 Prior and related work

While the NP-hardness of finding a 3-colouring of a 3-colourable graph was obtained by Karp [35] in 1972, the NP-hardness of finding a 4-colouring of a 3-colourable graph was only proved in 2000 by Khanna, Linial, and Safra [37] (see also the work of Guruswami and Khanna for a different proof [27]). This result implied NP-hardness of finding a $(k + 2\lfloor k/3\rfloor - 1)$-colouring of a k-colourable graph for $k \geq 3$ [37]. Early work of Garey and Johnson established NP-hardness of finding a $(2k - 5)$-colouring of a k-colourable graph for $k \geq 6$ [26]. In 2016, Brakensiek and Guruswami proved NP-hardness of a $(2k - 2)$-colouring of a k-colourable graph for $k \geq 3$ [10]. Only very recently, Bulín, Krokhin, and Opršal showed that finding a 5-colouring of a 3-colourable graph, and more generally, finding a $(2k - 1)$-colouring of a k-colourable graph for any $k \geq 3$, is NP-hard [18].

In 2001, Khot proved an asymptotic result: for sufficiently large k, finding a $k^{\frac{1}{25}(\log k)}$-colouring of a k-colourable graph is NP-hard [38]. In 2013, Huang improved the gap by showing the hardness of finding a $2^{\Omega(k^{1/3})}$-colouring of a k-colourable graph [33].

The NP-hardness of colouring (k-colourable graphs) with $(2k - 1)$ colours for $k \geq 3$ from [18] and with $2^{\Omega(k^{1/3})}$ colours for sufficiently large k from [33] constitute the currently strongest known NP-hardness results for approximate graph colouring.

Under stronger assumptions (Khot's 2-to-1 Conjecture [39] for $k \geq 4$ and its non-standard variant for $k = 3$), Dinur, Mossel, and Regev showed that finding a c-colouring of a k-colourable graph is NP-hard for all constants $3 \leq k \leq c$ [20] A variant of Khot's 2-to-1 Conjecture with imperfect completeness has recently been proved [19, 40], which implies hardness for approximate colouring variants where most but not all of the graph is guaranteed to be k-colourable.

Hypergraphs colourings, a special case of PCSPs, is another line of work intensively studied. A k-colouring of a hypergraph is an assignment of colours $\{1, 2, \ldots, k\}$ to its vertices that leaves no hyperedge monochromatic. Dinur, Regev, and Smyth showed that for any constants $2 \leq k \leq c$, it is NP-hard to find a c-colouring of given 3-uniform k-colourable hypergraph [21]. Other notions of colourings (such as different types of rainbow colourings) for hypergraphs were studied by Brakensiek and Guruswami [10, 11], Guruswami and Lee [28], and Austrin, Bhangale, and Potukuchi [2].

Some results are also known for colourings with a super-constant number of colours. For graphs, conditional hardness was obtained by Dinur and Shinkar [22]. For hypergraphs, NP-hardness results were obtained in recent work of Bhangale [8] and Austrin, Bhangale, and Potukuchi [1].

2 Results

For two graphs or digraphs G, H, we write $G \to H$ if there exists a homomorphism from G to H.[3] We are interested in the following computational problem.

Definition 2.1. *Fix two graphs G and H with $G \to H$. The (decision variant of the)* $\mathrm{PCSP}(G, H)$ *is, given an input graph I, output* YES *if $I \to G$, and* NO *if $I \not\to H$.*

To state our results it will be convenient to use the following definition.

Definition 2.2. *A graph H is* left-hard *if for every non-bipartite graph G with $G \to H$,* $\mathrm{PCSP}(G, H)$ *is NP-hard. A graph G is* right-hard *if for every loop-less graph H with $G \to H$,* $\mathrm{PCSP}(G, H)$ *is NP-hard.*

If $G \to G'$ and $H' \to H$, then $\mathrm{PCSP}(G, H)$ trivially reduces to $\mathrm{PCSP}(G', H')$ (this is called *homomorphic relaxation* [18]; intuitively, increasing the promise gap makes the problem easier). Therefore, if H is a left-hard graph, then all graphs left of H (that is, H' such that $H' \to H$) are trivially left-hard.[4] If G is right-hard, then all graphs right of G are right-hard.

For the same reason, since every non-bipartite graph admits a homomorphism from an odd cycle, to show that H is left-hard it suffices to show that $\mathrm{PCSP}(C_n, H)$ is NP-hard for arbitrarily large odd n, where C_n denotes the cycle on n vertices. Dually, since every loop-less graph admits a homomorphism to a clique, to show that G is right-hard it suffices to show that $\mathrm{PCSP}(G, K_k)$ is NP-hard for arbitrarily large k.

Brakensiek and Guruswami conjectured that all non-trivial PCSPs for (undirected) graphs are NP-hard, greatly extending Hell and Nešetřil's theorem:

Conjecture 2.1. ([12]) $\mathrm{PCSP}(G, H)$ *is NP-hard for every non-bipartite loop-less G, H. Equivalently, every loop-less graph is left-hard. Equivalently, every non-bipartite graph is right-hard.*

In addition to the results on classical colourings discussed above (the case where G and H are cliques), the following result was recently obtained in a novel application of topological ideas.

Theorem 2.1. (Krokhin and Opršal [41]) K_3 *is left-hard.*

[3] In this paper, we allow graphs to have loops: the existence of homomorphisms for such graphs is trivial, but this allows us to make statements about graph constructions that will work without exceptions.

[4] Note that by our definition, bipartite graphs are vacuously left-hard.

2.1 Improved hardness of classical colouring In Section 3, we focus on right-hardness. We use a simple construction called the *arc digraph* or *line digraph*, which decreases the chromatic number of a graph in a controlled way. The construction allows to conclude the following, in a surprisingly simple way:

Proposition 2.1. *There exists a right-hard graph if and only if K_4 is right-hard.*[5]

More concretely, we show in particular that $\mathrm{PCSP}(K_6, K_{2^k})$ log-space reduces to $\mathrm{PCSP}(K_4, K_k)$, for all $k \geq 4$. This contrasts with [4, Proposition 10.3],[6] where it is shown to be impossible to obtain such a reduction with *minion homomorphisms*: an algebraic reduction, described in the full version [55], central to the framework of [18, 4] (in particular, there exists a k such that $\mathrm{PCSP}(K_4, K_k)$ admits no minion homomorphism to any $\mathrm{PCSP}(K_{n'}, K_{k'})$ for $4 < n' \leq k'$).

Furthermore, we strengthen the best known asymptotic hardness: Huang [33] showed that for all sufficiently large n, $\mathrm{PCSP}(K_n, K_{2^{n^{1/3}}})$ is NP-hard. We improve this in two ways, using Huang's result as a black-box. First, we improve the asymptotics from sub-exponential $2^{n^{1/3}}$ to single-exponential $\binom{n}{\lfloor n/2 \rfloor} \sim \frac{2^n}{\sqrt{\pi n/2}}$. Second, we show the claim holds for n as low as 4.

Theorem 2.2. (**Main Result #1**) *For all $n \geq 4$,* $\mathrm{PCSP}(K_n, K_{\binom{n}{\lfloor n/2 \rfloor}-1})$ *is NP-hard.*

In comparison, the previous best result relevant for all integers n was proved in [18]: $\mathrm{PCSP}(K_n, K_{2n-1})$ is NP-hard for all $n \geq 3$. For $n = 3$ we are unable to obtain any results; for $n = 4$ the new bound $\binom{n}{\lfloor n/2 \rfloor} - 1 = 5$ is worse than $2n - 1 = 7$, while for $n = 5$ the two bounds coincide at 9. However, already for $n = 6$ we improve the bound from $2n - 1 = 11$ to $\binom{n}{\lfloor n/2 \rfloor} - 1 = 19$.

2.2 Left-hardness and topology In Section 4, we focus on left-hardness. The main idea behind Krokhin and Opršal's [41] proof that K_3 is left-hard is simple to state. To prove that $\mathrm{PCSP}(C_n, H)$ is NP-hard for all odd n, the algebraic framework of [18] shows that it is sufficient to establish certain properties of *polymorphisms*: homomorphisms $f\colon C_n^L \to H$ for $L \in \mathbb{N}$

[5] Jakub Opršal and Andrei Krokhin realised that in this Proposition, 4 can be improved to 3 by using the fact that $\delta(\delta(K_4))$ is 3-colourable, as proved by Rorabaugh, Tardif, Wehlau, and Zaguia [46]. Details will appear in a future journal version.

[6] [4] is a full version of [18]. Proposition 10.3 in [4] is Proposition 5.31 in the previous two versions of [4].

(where $G^L = G \times \cdots \times G$ is the L-fold tensor product[7]). For large n the graph C_n^L looks like an L-torus: an L-fold product of circles, so the pertinent information about f seems to be subsumed by its topological properties (such as *winding numbers*, when H is a cycle). We refer to [41] for further details, but this general principle applies to any H and in fact we prove (in Theorem 2.3 below) that whether H is left-hard or not depends *only* on its topology.

The topology we associate with a graph is its *box complex*. See the full version [55] for formal definitions and statements. Intuitively, the box complex $|\mathrm{Box}(H)|$ is a topological space built from H by taking the tensor product $H \times K_2$ and then gluing faces to each four-cycle and more generally, gluing higher-dimensional faces to complete bipartite subgraphs. The added faces ensure that the box complex of a product of graphs is the same as the product space of their box complexes: thanks to this, $\left|\mathrm{Box}(C_n^L)\right|$ is indeed equivalent to the L-torus. The product with K_2 equips the box complex with a symmetry that swaps the two sides of $H \times K_2$. This make the resulting space a $\mathbb{Z}_2$-space: a topological space together with a continuous involution from the space to itself, which we denote simply as $-$. A $\mathbb{Z}_2$-*map* between two $\mathbb{Z}_2$-spaces is a continuous function which preserves this symmetry: $f(-x) = -f(x)$. This allows to concisely state that a given map is "non-trivial" (in contrast, there is always *some* continuous function from one space to another: just map everything to a single point). The main use of the box complex is then the statement that every graph homomorphism $G \to H$ induces a $\mathbb{Z}_2$-map from $|\mathrm{Box}(G)|$ to $|\mathrm{Box}(H)|$. Graph homomorphisms can thus be studied with tools from algebraic topology.

The classical example of this is an application of the Borsuk-Ulam theorem: there is no $\mathbb{Z}_2$-map from $\mathcal{S}^n$ to $\mathcal{S}^m$ for $n > m$, where $\mathcal{S}^n$ denotes the n-dimensional sphere with antipodal symmetry. Hence if G and H are graphs such that $|\mathrm{Box}(G)|$ and $|\mathrm{Box}(H)|$ are equivalent to $\mathcal{S}^n$ and $\mathcal{S}^m$, respectively, then there can be no graph homomorphism $G \to H$. See Figure 1.

This is essentially the idea in Lovász' proof [42] of Kneser's conjecture that the chromatic number of Kneser graphs $KG(n,k)$ is $n-2k+2$. In the language of box complexes, the proof amounts to showing that the box complex of a clique K_c is equivalent to $\mathcal{S}^{c-2}$, while the box complex of a Kneser graph contains $\mathcal{S}^{n-2k}$. We refer to [43] for an in-depth, yet accessible reference.

[7]The *tensor* (or *categorical*) *product* $G \times H$ of graphs G, H has pairs $(g,h) \in V(G) \times V(H)$ as vertices and (g,h) is adjacent to (g',h') whenever g is adjacent to g' (in G) and h is adjacent to h' (in H).

We show that the left-hardness of a graph depends only on the topology of its box complex (in fact, it is only important what $\mathbb{Z}_2$-maps it admits, which is significantly coarser than $\mathbb{Z}_2$-homotopy equivalence):

THEOREM 2.3. (**Main Result #2**) *If H is left-hard and H' is a graph such that $|\mathrm{Box}(H')|$ admits a $\mathbb{Z}_2$-map to $|\mathrm{Box}(H)|$, then H' is left-hard.*

Using Krokhin and Opršal's result that K_3 is left-hard (Theorem 2.1), since $|\mathrm{Box}(K_3)|$ is the circle $\mathcal{S}^1$ (up to $\mathbb{Z}_2$-homotopy equivalence), we immediately obtain the following:

COROLLARY 2.1. *Every graph H for which $|\mathrm{Box}(H)|$ admits a $\mathbb{Z}_2$-map to $\mathcal{S}^1$ is left-hard.*

Two examples of such graphs (other than 3-colourable graphs) are loop-less square-free graphs and circular cliques $K_{p/q}$ with $2 < \frac{p}{q} < 4$ (see the full version [55] for proofs), which we introduce next. *Square-free graphs* are graphs with no cycle of length exactly 4. In particular, this includes all graphs of girth at least 5 and hence graphs of arbitrarily high chromatic number (but incomparable to K_4 and larger cliques, in terms of the homomorphism $\to$ relation). The *circular clique* $K_{p/q}$ (for $p, q \in \mathbb{N}, \frac{p}{q} > 2$) is the graph with vertex set $\mathbb{Z}_p$ and an edge from i to every integer at least q apart: $i+q, i+q+1, \ldots, i+p-q$. They generalise cliques $K_n = K_{n/1}$ and odd cycles $C_{2n+1} \simeq K_{(2k+1)/k}$. Their basic property is that $K_{p/q} \to K_{p'/q'}$ if and only if $\frac{p}{q} \le \frac{p'}{q'}$. Thus circular cliques refine the chain of cliques and odd cycles, corresponding to rational numbers between integers. For example:

$$\begin{aligned}\cdots \to C_7 \to C_5 \to C_3 = \\ = K_3 \to K_{7/2} \to K_4 \to K_{9/2} \to K_5 \to \ldots\end{aligned}$$

The *circular chromatic number* $\chi_c(G)$ is the infimum over $\frac{p}{q}$ such that $G \to K_{p/q}$. Therefore:

COROLLARY 2.2. *For every $2 < r \le r' < 4$, it is NP-hard to distinguish graphs G with $\chi_c(G) \le r$ from those with $\chi_c(G) > r'$.*

In this sense, we conclude that $K_{4-\varepsilon}$ is left-hard, thus extending the result for K_3. However, the closeness to K_4 is only deceptive and no conclusions on 4-colourings follow. For K_4, since the box complex is equivalent to the standard 2-dimensional sphere, we can at least conclude that to prove left-hardness of K_4 it would be enough to prove left-hardness of any other graph with the same topology: these include all non-bipartite quadrangulations of the projective plane, in

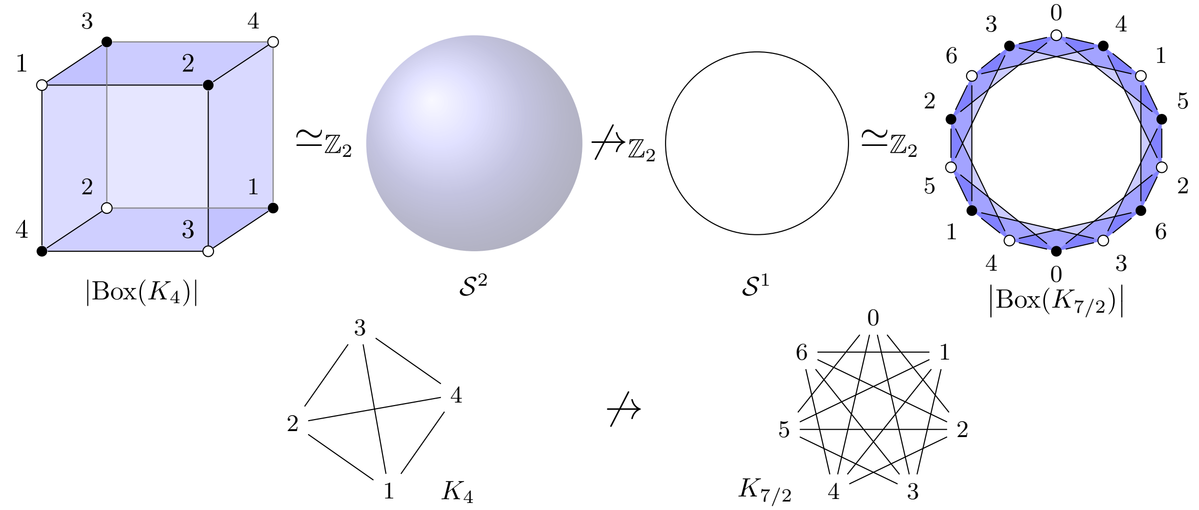

Figure 1: The box complex of K_4 is the hollow cube (informally speaking; the drawing skips some irrelevant faces). It is equivalent ($\mathbb{Z}_2$-homotopy equivalent) to the sphere. The box complex of the circular clique $K_{7/2}$ is equivalent to the circle. Thus there cannot be a homomorphism from K_4 to $K_{7/2}$ (of course in this case it is easier to show this directly).

particular the Grötzsch graph, 4-chromatic generalised Mycielskians, and 4-chromatic Schrijver graphs [43, 9]. In this sense, the exact geometry of K_4 is irrelevant. However, the fact that it is a finite graph, with only finitely many possible maps from C_n^L for any fixed n, L should still be relevant, as it is for K_3. It is also quite probable that any proof for a "spherical" graph would apply just as well to K_4, where the proof could be just notationally much simpler.

In the full version [55] we rephrase Krokhin and Opršal's [41] proof of Theorem 2.1 in terms of the box complex. In particular, left-hardness of K_3 follows from some general principles and the fact that $|\mathrm{Box}(K_3)|$ is a circle. The proof also extends to all graphs H such that $|\mathrm{Box}(H)|$ admits a $\mathbb{Z}_2$-map to $\mathcal{S}^1$, giving an independent, self-contained proof of Corollary 2.1 (and Theorem 2.1 in particular).

The general principle is that a homomorphism $C_n^L \to H$ induces a $\mathbb{Z}_2$-map $(\mathcal{S}^1)^L \to |\mathrm{Box}(H)|$, in a way that preserves minors (identifications within the L variables) and automorphisms. (In the language of category theory, the box complex is a functor from the category of graphs to that of $\mathbb{Z}_2$-spaces, and the functor preserves products). In turn, the $\mathbb{Z}_2$-map induces a group homomorphism between the fundamental group of $(\mathcal{S}^1)^L$, which is just $\mathbb{Z}^L$, and that of $|\mathrm{Box}(H)|$. This is essentially the map $\mathbb{Z}^L \to \mathbb{Z}$ obtained in [41]. While this rephrasing requires a bit more technical definitions, the main advantage is that it allows to replace a tedious combinatorial argument (about winding numbers preserving minors) with straightforward statements about preserving products.

2.3 Methodology – adjoint functors While the proof of the first main result is given elementarily in Section 3, it fits together with the second main result in a much more general pattern. The underlying principle is that pairs of graph constructions satisfying a simple duality condition give reductions between PCSPs. To introduce them, let us consider a concrete example. For a graph G and an odd integer k, $\Lambda_k G$ is the graph obtained by subdividing each edge into a path of k edges; $\Gamma_k G$ is the graph obtained by taking the k-th 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 of length exactly k in G. (For example $\Gamma_3 G$ has loops if G has triangles).

We say a graph construction Λ (a function from graphs to graphs) is a *thin (graph) functor* if $G \to H$ implies $\Lambda G \to \Lambda H$ (for all G, H). A pair of thin functors (Λ, Γ) is a *thin adjoint pair* if

$$\Lambda G \to H \text{ if and only if } G \to \Gamma H.$$

We call Λ the *left adjoint* of Γ and Γ the *right adjoint* of Λ.

For all odd k, (Λ_k, Γ_k) are a thin adjoint pair. For example, since $\Gamma_3 C_5 = K_5$, we have $G \to K_5$ if and only if $\Lambda_k G \to C_5$. This is a basic reduction that shows the NP-hardness of C_5-colouring; in fact adjointness of various graph construction is the principal tool behind the original proof of Hell and Nešetřil's theorem (characterising the complexity of H-colouring) [30].

In category theory, there is a stronger and more technical notion of (non-thin) *functors* and *adjoint pairs*. A thin graph functor is in fact a functor in the *thin category of graphs*, that is, the category whose objects are graphs, and with at most one morphism from one graph to another, indicating whether a homomorphism exists or not. In other words, we are only interested in the existence of homomorphisms, and not in their identity and how they compose. Equivalently, we look only at the preorder of graphs by the $G \to H$ relation (we can also make this a poset by considering graphs up to homomorphic equivalence). In order-theoretic language, thin functors are just order-preserving maps, while thin adjoint functors are known as Galois connections. We prefer the categorical language as most of the constructions we consider are in fact functors (in the non-thin category of graphs), which is important for connections to the algebraic framework of [18], as we discuss in the full version [55]. While unnecessary for our main results, we believe it may be important to understand these deeper connections to resolve the conjectures completely.

Thin adjoint functors give us a way to reduce one PCSP to another. We say that a graph functor Γ is log-space computable if, given a graph G, ΓG can be computed in logarithmic space in the size of G.

Observation 2.1. *Let Λ, Γ be thin adjoint graph functors and suppose Λ is log-space computable. Then* $\mathrm{PCSP}(G, \Gamma H)$ *reduces to* $\mathrm{PCSP}(\Lambda G, H)$ *in log-space, for all graphs G, H.*

Proof. Let F be an instance of $\mathrm{PCSP}(G, \Gamma H)$. Then ΛF is an appropriate instance of $\mathrm{PCSP}(\Lambda G, H)$. Indeed, if $F \to G$, then $\Lambda F \to \Lambda G$ (because Λ is a thin functor). If $\Lambda F \to H$, then $F \to \Gamma H$ by adjointness. ∎

In some cases, a thin functor Γ that is a thin right adjoint in a pair (Λ, Γ) is also a thin left adjoint in a pair (Γ, Ω). This allows to get a reduction in the opposite direction:

Observation 2.2. *Let (Λ, Γ) and (Γ, Ω) be thin adjoint pairs of functors. Then* $\mathrm{PCSP}(\Gamma G, H)$ *and* $\mathrm{PCSP}(G, \Omega H)$ *are log-space equivalent (assuming Λ and Γ are log-space computable).*

Proof. The previous observation gives a reduction from $\mathrm{PCSP}(G, \Omega H)$ to $\mathrm{PCSP}(\Gamma G, H)$. For the other direction, let F be an instance of $\mathrm{PCSP}(\Gamma G, H)$. Then ΛF is an appropriate instance of $\mathrm{PCSP}(G, \Omega H)$. Indeed, if $F \to \Gamma G$, then $\Lambda F \to G$. If $\Lambda F \to \Omega H$, then $F \to \Gamma\Omega H \to H$. The last arrow follows from the trivial $\Omega H \to \Omega H$. ∎

The proofs of Observations 2.1 and 2.2 of course extend to digraphs and general relational structures. Note that the above proofs reduce decision problems; they work just as well for search problems: all the thin adjoint pairs (Λ, Γ) we consider with Λ log-space computable also have the property that a homomorphism $\Lambda F \to H$ can be computed from a homomorphism $F \to \Gamma H$ and vice versa, in space logarithmic in the size of F.

As we discuss in Section 4, all of our results follow from reductions that are either trivial (homomorphic relaxations) or instantiations of Observation 2.1. While for the first main result we prefer to first give a direct proof that avoids this formalism (in Section 3), it will be significantly more convenient for the second main result (in Section 4), where we use a certain thin right adjoint Ω_k to the k-th power Γ_k.

2.4 Hedetniemi's conjecture Another leitmotif of this paper is the application of various tools developed in research around Hedetniemi's conjecture. A graph K is *multiplicative* if $G \times H \to K$ implies $G \to K$ or $H \to K$. The conjecture states that all cliques $K = K_n$ are multiplicative. Equivalently, $\chi(G \times H) = \min(\chi(G), \chi(H))$; see [56, 47, 51] for surveys. In a very recent breakthrough, Shitov [49] proved that the conjecture is false (for large n).

The arc digraph construction, which we will use in Section 3 to prove Theorem 2.2, was originally used by Poljak and Rödl [45] to show certain asymptotic bounds on chromatic numbers of products. The functors $\Lambda_k, \Gamma_k, \Omega_k$ were applied by Tardif [50] to show that colourings to circular cliques $K_{p/q}$ ($2 < \frac{p}{q} < 4$) satisfy the conjecture. Matsushita [44] used the box complex to show that Hedetniemi's conjecture would imply an analogous conjecture in topology. This was independently proved by the first author [54] using Ω_k functors, while the box complex was used to show that square-free graphs are multiplicative [53]. See [25] for a survey on applications of adjoint functors to the conjecture.

The refutation of Hedetniemi's conjecture and the fact that methods for proving the multiplicativity of K_3 extend to $K_{4-\varepsilon}$ and square-free graphs, but fail to extend to K_4, might suggest that the Conjecture 2.1 is doomed to the same fate. However, it now seems clear that proving multiplicativity requires more than just topology [52]: known methods do not even extend to all graphs H such that $|\mathrm{Box}(H)|$ is a circle. This contrasts with Theorem 2.3: topological tools work much more gracefully in the setting of PCSPs.

3 The arc digraph construction

Let D be a digraph. The *arc digraph* (or *line digraph*) of D, denoted δD , is the digraph whose vertices are arcs (directed edges) of D and whose arcs are pairs of the form $((u,v),(v,w))$. We think of undirected graphs as symmetric relations: digraphs in which for every arc (u,v) there is an arc (v,u). So for an undirected graph G, $\delta(G)$ has $2|E(G)|$ vertices and is a directed graph: the directions will not be important in this section, but will be in the full version [55]. The chromatic number of a digraph is the chromatic number of the underlying undirected graph (obtained by symmetrising each arc; so $\chi(D) \le n$ if and only if $D \to K_n$).

The crucial property of the arc digraph construction is that it decreases the chromatic number in a controlled way (even though it is computable in log-space!). We include a short proof for completeness. We denote by $[n]$ the set $\{1,2,\dots,n\}$.

LEMMA 3.1. (HARNER AND ENTRINGER [29]) *For any graph G:*

- *if $\chi(\delta(G)) \le n$, then $\chi(G) \le 2^n$,*
- *if $\chi(G) \le \binom{n}{\lfloor n/2 \rfloor}$, then $\chi(\delta(G)) \le n$.*

Proof. Suppose δG has an n-colouring. Recall that we think of G as a digraph with two arcs (u,v) and (v,u) for each edge $\{u,v\} \in E(G)$; thus δG contains two vertices (u,v) and (v,u), as well as (by definition of δ) two arcs from one pair to the other. In particular, an n-colouring of δG gives distinct colours to (u,v) and (v,u). Define a 2^n-colouring ϕ of G by assigning to each vertex v the set $\phi(v)$ of colours of incoming arcs. For any edge $\{u,v\}$ of G, $\phi(v)$ contains the colour c of the arc (u,v). Since every arc incoming to u gets a different colour from (u,v), the set $\phi(u)$ does not contain c. Hence $\phi(u) \neq \phi(v)$, so ϕ is a proper colouring.

Suppose G has a $\binom{n}{\lfloor n/2 \rfloor}$-colouring ϕ. We interpret colours $\phi(v)$ as $\lfloor n/2 \rfloor$-element subsets of $[n]$. Define an n-colouring of δG by assigning to each arc (u,v) an arbitrary colour in $\phi(u) \setminus \phi(v)$ (the minimum, say). Such a colour exists because $\phi(u) \neq \phi(v)$. For arcs (u,v), (v,w) clearly $\phi(u) \setminus \phi(v)$ is disjoint from $\phi(v) \setminus \phi(w)$, so this is a proper colouring of $\delta(G)$. ∎

The proofs in fact works for digraphs as well. For graphs, it is not much harder to show an exact correspondence (we note however that most conclusions only require the above approximate correspondence). Let us denote $b(n) := \binom{n}{\lfloor n/2 \rfloor}$.

LEMMA 3.2. (POLJAK AND RÖDL [45]) *For a (symmetric) graph G, $\chi(\delta(G)) = \min\{n \mid \chi(G) \le b(n)\}$. In other words, $\delta G \to K_n$ if and only if $G \to K_{b(n)}$.*

This immediately gives the following implication for approximate colouring:

LEMMA 3.3. $\mathrm{PCSP}(K_{b(n)}, K_{b(k)})$ *log-space reduces to* $\mathrm{PCSP}(K_n, K_k)$*, for all $n, k \in \mathbb{N}$.*

Proof. Let G be an instance of the first problem. Then δG is a suitable instance of $\mathrm{PCSP}(K_n, K_k)$: if $G \to K_{b(n)}$, then $\delta G \to K_n$. If $\delta G \to K_k$, then $G \to K_{b(k)}$. ∎

REMARK 3.1. *As a side note, adding a universal vertex gives the following obvious reduction:* $\mathrm{PCSP}(K_n, K_k)$ *log-space reduces to* $\mathrm{PCSP}(K_{n+1}, K_{k+1})$*, for $n, k \in \mathbb{N}$.*

Recall also that if $n \le n' \le k' \le k$, then $\mathrm{PCSP}(K_n, K_k)$ trivially reduces to $\mathrm{PCSP}(K_{n'}, K_{k'})$. One corollary of Lemma 3.3 is that if any clique of size at least 4 is right-hard, then all of them are:

PROPOSITION 3.1. *For all integers $n, n' \ge 4$, $\mathrm{PCSP}(K_n, K_k)$ is NP-hard for all $k \ge n$ if and only if $\mathrm{PCSP}(K_{n'}, K_{k'})$ is NP-hard for all $k' \ge n'$.*

Proof. Let $n \le n'$. For one direction, right-hardness of K_n trivially implies right-hardness of $K_{n'}$.

On the other hand, we claim that if $K_{b(n)}$ is right-hard, then so is K_n. Indeed, suppose $\mathrm{PCSP}(K_{b(n)}, K_k)$ is hard for all $k \ge b(n)$. In particular it is hard for all k of the form $k = b(k')$ for an integer $k' \ge n$. Hence by Lemma 3.3, $\mathrm{PCSP}(K_n, K_{k'})$ is hard for all $k' \ge n$.

Suppose K_n is not right-hard. Then $K_{b(n)}$ is not right-hard, $K_{b(b(n))}$ is not right-hard and so on. Since starting with $n \ge 4$, the sequence $b(b(\dots n \dots))$ grows to infinity, we conclude that $K_{n''}$ is not right-hard for some $n'' \ge n'$. Therefore, trivially $K_{n'}$ is not right-hard. ∎

In other words if any loop-less graph H is right-hard, then trivially some large enough clique $K_{\chi(H)}$ is right-hard; by the above, K_4 and all graphs right of it are right-hard. This proves Proposition 2.1. The proof fails to extend to K_3 because $b(3) = \binom{3}{\lfloor 3/2 \rfloor}$ is not strictly greater than 3.

The other consequence we derive from Lemma 3.3 is a strengthening of Huang's result:

THEOREM 3.1. (HUANG [33]) *For all sufficiently large n, $\mathrm{PCSP}(K_n, K_{2^{\Omega(n^{1/3})}})$ is NP-hard.*

THEOREM 3.2. (**Main Result #1**) *For all $n \ge 4$, $\mathrm{PCSP}(K_n, K_{\binom{n}{\lfloor n/2 \rfloor}-1})$ is NP-hard.*

We thus improve the asymptotics from sub-exponential $f(n) := 2^{n^{1/3}}$ to single-exponential $b(n) = \binom{n}{\lfloor n/2 \rfloor} \sim \frac{2^n}{\sqrt{\pi n/2}}$. The informal idea of the proof (delayed to the full version [55]) is as follows. Suppose n vs $f(n)$-colouring is hard (that is, $\mathrm{PCSP}(K_n, K_{f(n)})$ is NP-hard for all sufficently large n). Substituting $n = b(m)$, we conclude that $b(m)$ vs $f(b(m))$ is hard (for sufficiently large m). Trivially, we can then replace $f(b(m))$ with any smaller number of the form $b(k)$. So $b(m)$ vs $b(k)$ is hard for $k = b^{-1}(f(b(m)))$ where $b^{-1}(x) := \max\{k \colon b(k) \leq x\}$. By Lemma 3.3, we conclude that m vs k is hard. That is, m vs $b^{-1}(f(b(m)))$ is hard for all sufficiently large m.

Thus any $f(n)$ can be improved to $b^{-1}(f(b(n)))$. Since $b(n)$ is roughly exponential and $b^{-1}(n)$ is roughly logarithmic, starting from a function $f(n)$ of order $\exp^{(i+1)}(\alpha \cdot \log^{(i)}(n))$ with i-fold compositions and a constant $\alpha > 0$, such as $f(n) = 2^{n^{1/3}} = 2^{2^{\frac{1}{3}\log n}}$ from Huang's hardness, results in

$$\begin{aligned} b^{-1}(f(b(n))) &\approx \log\Big(\exp^{(i+1)}\big(\alpha \cdot \log^{(i)}(\exp(n))\big)\Big) \\ &= \exp^{(i)}(\alpha \cdot \log^{(i-1)}(n)), \end{aligned}$$

so a similar composition but with i decreased. In a constant number of steps, this results in a single-exponential function. In fact using one more step, but without approximating the function $b(n)$, this results in exactly $b(n) - 1$. We note it would not be sufficient to start from a quasi-polynomial $f(n)$, like $n^{\Theta(\log n)}$ in Khot's [38] result. To relax the requirement from "for sufficiently large n" to "for all $n \geq 4$", the argument is similar as in Proposition 3.1.

4 Adjoint functors and topology

Recall that Λ_k and Γ_k denote k-subdivision and the k-th power of a graph; they are thin adjoints:

$$\Lambda_k G \to H \text{ if and only if } G \to \Gamma_k H.$$

(for all odd k). More surprisingly, Γ_k is itself the thin *left* adjoint of a certain thin functor Ω_k:

$$\Gamma_k G \to H \text{ if and only if } G \to \Omega_k H.$$

This characterizes $\Omega_k G$ up to homomorphic equivalence; the exact definition will be irrelevant (we give it in the full version [55]). We note that Λ_k and Γ_k are log-space computable, for all odd k; however, Ω_k is not: $\Omega_k G$ is exponentially larger than G. See [54] for more about the thin functors $\Lambda_k, \Gamma_k, \Omega_k$ and their properties.

Observation 2.1 tells us that $\mathrm{PCSP}(G, \Omega_k H)$ log-space reduces to $\mathrm{PCSP}(\Gamma_k G, H)$. To give conclusions on left-hardness, we will need only two more facts about these functors. First, $\Omega_k G \to G$; second, $\Lambda_k G \to \Omega_k G$ for all G and odd k (see Lemma 2.3(iv) and (vi) in [54]).

Lemma 4.1. *For every odd k, $\Omega_k H$ is left-hard if and only if H is left-hard.*

Proof. If H is left-hard, then trivially so is $\Omega_k H$ because $\Omega_k H \to H$. For the other implication, suppose $\Omega_k H$ is left-hard, that is, $\mathrm{PCSP}(G, \Omega_k H)$ is hard for every non-bipartite G such that $G \to \Omega_k H$. By Observation 2.1, this implies $\mathrm{PCSP}(\Gamma_k G, H)$ is hard. Let G' be any non-bipartite graph such that $G' \to H$. We want to show that $\mathrm{PCSP}(G', H)$ is hard. Observe that $\Omega_k G'$ is non-bipartite, because $\Lambda_k G' \to \Omega_k G'$ and Λ_k subdivides each edge of G' an odd number of times. Since $\Omega_k G' \to \Omega_k H$, using $G := \Omega_k G'$ we conclude that $\mathrm{PCSP}(\Gamma_k \Omega_k G', H)$ is hard. Since $\Gamma_k \Omega_k G' \to G'$, this implies $\mathrm{PCSP}(G', H)$ is hard. ∎

As an example, consider the circular clique $K_{7/2}$ (we have $K_3 \to K_{7/2} \to K_4$). Knowing that K_3 is left-hard, one could check that $\Omega_3(K_{7/2})$ is 3-colorable and hence left-hard as well; the above lemma then allows to conclude that $K_{7/2}$ is left-hard. What other graphs could one use in place of $K_{7/2}$? The answer turns out to be topological. Intuitively, while the operation Γ_k gives a "thicker" graph, the operation Ω_k gives a "thinner" one. In fact, Ω_k behaves like barycentric subdivision in topology: it preserves the topology of a graph (formally: its box complex is $\mathbb{Z}_2$-homotopy equivalent to the original graph's box complex) but refines its geometry. With increasing k, this eventually allows to model any continuous map with a graph homomorphism; in particular:

Theorem 4.1. ([54]) *There exists a $\mathbb{Z}_2$-map $|\mathrm{Box}(G)| \to_{\mathbb{Z}_2} |\mathrm{Box}(H)|$ if and only if for some odd k, $\Omega_k G \to H$.*

This concludes our second main result:

Proof. [Proof of Theorem 2.3] Let H be left-hard and let H' be a graph s.t. $|\mathrm{Box}(H')| \to_{\mathbb{Z}_2} |\mathrm{Box}(H)|$. By Theorem 4.1, $\Omega_k H' \to H$ for some odd k. Hence $\Omega_k H'$ is left-hard. By Lemma 4.1, H' is left-hard. ∎

Conclusions In the full version [55], we consider other examples of thin adjoint functors. In particular we discuss how results of Section 3 follow from Observation 2.1 by considering a thin right adjoint δ_R of δ. We also introduce the algebraic framework of [18] and contemplate how Observation 2.1 and 2.2 could fit into it. A few of the thin adjoint functors we considered are in fact adjoint functors, in the category-theoretic sense. This does allow to deduce some instances of Obs. 2.1 and 2.2

from algebraic reductions known as minion homomorphisms. However, the most useful examples: (δ, δ_R) and (Γ_k, Ω_k), are only thin adjoint. Hence the question remains whether the notion of minion homomorphisms can be extended to account for them. It could also be interesting to consider how δ or δ_R affect the topology of a graph. Another direction could be to look at Huang's Theorem 3.1 not as a black-box: could constructions like δ be useful to say something directly about PCPs?

References

[1] P. Austrin, A. Bhangale, and A. Potukuchi. Simplified inpproximability of hypergraph coloring via t-agreeing families. Technical report, 2019, arXiv:1904.01163.

[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)*, 2020, arXiv:1810.02784. (to appear).

[3] P. Austrin, V. Guruswami, and J. Håstad. $(2+\epsilon)$-SAT Is NP-hard. *SIAM J. Comput.*, 46(5):1554–1573, 2017.

[4] L. Barto, J. Bulín, A. Krokhin, and J. Opršal. Algebraic approach to promise constraint satisfaction. Technical report, 2019, arXiv:1811.00970v3.

[5] L. Barto and M. Kozik. Constraint Satisfaction Problems solvable by local consistency methods. *Journal of the ACM*, 61(1), 2014.

[6] L. Barto and M. Kozik. Robustly solvable constraint satisfaction problems. *SIAM Journal on Computing*, 45(4):1646–1669, 2016, arXiv:1512.01157.

[7] L. Barto, A. Krokhin, and R. Willard. Polymorphisms, and how to use them. In A. Krokhin and S. Živný, editors, *Complexity and approximability of Constraint Satisfaction Problems*, volume 7 of *Dagstuhl Follow-Ups*, pages 1–44. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.

[8] 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)*, volume 107 of *LIPIcs*, pages 15:1–15:11. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.

[9] A. Björner and M. de Longueville. Neighborhood complexes of stable kneser graphs. *Combinatorica*, 23(1):23–34, 2003.

[10] J. Brakensiek and V. Guruswami. New Hardness Results for Graph and Hypergraph Colorings. In *Proceedings of the 31st Conference on Computational Complexity (CCC'16)*, volume 50 of *LIPIcs*, pages 14:1–14:27. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.

[11] J. Brakensiek and V. Guruswami. The quest for strong inapproximability results with perfect completeness. In *Proceedings of the 21st International Workshop on Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques and the 20th International Workshop on Randomization and Computation (APPROX-RANDOM'17)*, volume 81 of *LIPIcs*, pages 4:1–4:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.

[12] J. Brakensiek and V. Guruswami. Promise Constraint Satisfaction: Structure Theory and a Symmetric Boolean Dichotomy. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'18)*, pages 1782–1801. SIAM, 2018, arXiv:1704.01937.

[13] J. Brakensiek and V. Guruswami. An Algorithmic Blend of LPs and Ring Equations for Promise CSPs. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'19)*, pages 436–455. SIAM, 2019, arXiv:1807.05194.

[14] A. Bulatov. A dichotomy theorem for constraint satisfaction problems on a 3-element set. *Journal of the ACM*, 53(1):66–120, 2006.

[15] A. Bulatov, P. Jeavons, and A. Krokhin. Classifying the complexity of constraints using finite algebras. *SIAM Journal on Computing*, 34(3):720–742, 2005.

[16] A. A. Bulatov. Complexity of conservative constraint satisfaction problems. *ACM Transactions on Computational Logic*, 12(4), 2011. Article 24.

[17] A. A. Bulatov. A Dichotomy Theorem for Nonuniform CSPs. In *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science (FOCS'17)*, pages 319–330. IEEE, 2017, arXiv:1703.03021.

[18] J. Bulín, A. Krokhin, and J. Opršal. Algebraic approach to promise constraint satisfaction. In *Proceedings of the 51st Annual ACM Symposium on Theory of Computing (STOC'19)*, pages 602–613. ACM, 2019, arXiv:1811.00970.

[19] 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)*, pages 376–389, 2018.

[20] I. Dinur, E. Mossel, and O. Regev. Conditional hardness for approximate coloring. *SIAM J. Comput.*, 39(3):843–873, 2009, arXiv:cs/0504062.

[21] I. Dinur, O. Regev, and C. D. Smyth. The hardness of 3-uniform hypergraph coloring. *Combinatorica*, 25(5):519–535, 2005.

[22] 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)*, volume 6302 of *Lecture Notes in Computer Science*, pages 138–151. Springer, 2010.

[23] T. Feder and M. Y. Vardi. The Computational Structure of Monotone Monadic SNP and Constraint Satisfaction: A Study through Datalog and Group Theory. *SIAM Journal on Computing*, 28(1):57–104, 1998.

[24] 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. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019, arXiv:1904.12424. (to appear).

[25] J. Foniok and C. Tardif. Hedetniemi's conjecture and adjoint functors in thin categories. *Applied Categorical Structures*, 26:113–128, 2018, arXiv:1608.02918.

[26] M. R. Garey and D. S. Johnson. The complexity of near-optimal graph coloring. *J. ACM*, 23(1):43–49, 1976.

[27] V. Guruswami and S. Khanna. On the hardness of 4-coloring a 3-colorable graph. *SIAM J. Discrete Math.*, 18(1):30–40, 2004.

[28] V. Guruswami and E. Lee. Strong inapproximability results on balanced rainbow-colorable hypergraphs. *Combinatorica*, 38(3):547–599, 2018.

[29] C. Harner and R. Entringer. Arc colorings of digraphs. *J. Comb. Theory, Ser. B*, 13(3):219–225, 1972.

[30] P. Hell and J. Nešetřil. On the complexity of H-coloring. *Journal of Combinatorial Theory, Series B*, 48(1):92–110, 1990.

[31] P. Hell and J. Nešetřil. *Graphs and Homomorphisms*, volume 28 of *Oxford Lecture Series in Mathematics and Its Applications*. Oxford University Press, 2004.

[32] P. Hell and J. Nešetřil. Colouring, constraint satisfaction, and complexity. *Computer Science Review*, 2(3):143–163, 2008.

[33] S. Huang. Improved hardness of approximating chromatic number. In *Proceedings of the 16th International Workshop on Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques and the 17th International Workshop on Randomization and Computation (APPROX-RANDOM'13)*, pages 233–243. Springer, 2013, arXiv:1301.5216.

[34] P. Jeavons, D. A. Cohen, and M. Gyssens. Closure properties of constraints. *Journal of the ACM*, 44(4):527–548, 1997.

[35] R. M. Karp. Reducibility Among Combinatorial Problems. In *Proceedings of a Symposium on the Complexity of Computer Computations*, pages 85–103, 1972.

[36] K. Kawarabayashi and M. Thorup. Coloring 3-Colorable Graphs with Less than $n^{1/5}$ Colors. *J. ACM*, 64(1):4:1–4:23, 2017.

[37] S. Khanna, N. Linial, and S. Safra. On the hardness of approximating the chromatic number. *Combinatorica*, 20(3):393–415, 2000.

[38] 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)*, pages 600–609. IEEE Computer Society, 2001.

[39] 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)*, pages 767–775. ACM, 2002.

[40] S. Khot, D. Minzer, and M. Safra. Pseudorandom sets in grassmann graph have near-perfect expansion. In M. Thorup, editor, *Proceedings of the 59th IEEE Annual Symposium on Foundations of Computer Science (FOCS'18)*, pages 592–601. IEEE Computer Society, 2018.

[41] A. Krokhin and J. Opršal. The complexity of 3-colouring H-colourable graphs. In *Proceedings of the 60th Annual IEEE Symposium on Foundations of Computer Science (FOCS'19)*, 2019, arXiv:1904.03214. (to appear).

[42] L. Lovász. Kneser's conjecture, chromatic number, and homotopy. *J. Comb. Theory, Ser. A*, 25(3):319–324, 1978.

[43] J. Matoušek. *Using the Borsuk-Ulam theorem: lectures on topological methods in combinatorics and geometry.* Universitext. Springer, 2008.

[44] T. Matsushita. $\mathbb{Z}_2$-indices and Hedetniemi's conjecture. *Discrete & Computational Geometry*, 2019, arXiv:1710.05290.

[45] S. Poljak and V. Rödl. On the arc-chromatic number of a digraph. *J. Comb. Theory, Ser. B*, 31(2):190–198, 1981.

[46] D. Rorabaugh, C. Tardif, D. Wehlau, and I. Zaguia. Iterated arc graphs. Technical report, 2016, arXiv:1610.01259.

[47] N. Sauer. Hedetniemi's conjecture – a survey. *Discrete Math.*, 229(1-3):261–292, 2001.

[48] T. J. Schaefer. The Complexity of Satisfiability Problems. In *Proceedings of the 10th Annual ACM Symposium on Theory of Computing (STOC'78)*, pages 216–226. ACM, 1978.

[49] Y. Shitov. Counterexamples to Hedetniemi's conjecture. *Annals of Mathematics*, 190(2):663–667, 2019, arXiv:1905.02167.

[50] C. Tardif. Multiplicative graphs and semi-lattice endomorphisms in the category of graphs. *J. Comb. Theory, Ser. B*, 95(2):338–345, 2005.

[51] C. Tardif. Hedetniemi's conjecture, 40 years later. *Graph Theory Notes NY*, 54(46-57):2, 2008.

[52] C. Tardif and M. Wrochna. Hedetniemi's conjecture and strongly multiplicative graphs. *SIAM J. Discrete Math.*, 2019, arXiv:1808.04778. (to appear).

[53] M. Wrochna. Square-free graphs are multiplicative. *J. Comb. Theory, Ser. B*, 122:479–507, 2017, arXiv:1601.04551.

[54] M. Wrochna. On inverse powers of graphs and topological implications of Hedetniemi's conjecture. *J. Comb. Theory, Ser. B*, 2019, arXiv:1712.03196.

[55] M. Wrochna and S. Živný. Improved hardness for H-colourings of G-colourable graphs. Technical report, 2019, arXiv:1907.00872.

[56] X. Zhu. A survey on Hedetniemi's conjecture. *Taiwanese J. Math.*, 2(1):1–24, 1998.

[57] D. Zhuk. A Proof of CSP Dichotomy Conjecture. In *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science (FOCS'17)*, pages 331–342. IEEE, 2017, arXiv:1704.01914.