

Computation of Hadwiger Number and Related Contraction Problems: Tight Lower Bounds

FEDOR V. FOMIN, University of Bergen, Norway

DANIEL LOKSHTANOV, University of California, Santa Barbara, USA

IVAN MIHAJLIN, University of California, San Diego, USA

SAKET SAURABH, The Institute for Mathematical Sciences, HBNI, India and University of Bergen, Norway

MEIRAV ZEHAVI, Ben-Gurion University, Israel

We prove that the Hadwiger number of an n -vertex graph G (the maximum size of a clique minor in G) cannot be computed in time $n^{o(n)}$, unless the Exponential Time Hypothesis (ETH) fails. This resolves a well-known open question in the area of exact exponential algorithms. The technique developed for resolving the Hadwiger number problem has a wider applicability. We use it to rule out the existence of $n^{o(n)}$ -time algorithms (up to the ETH) for a large class of computational problems concerning edge contractions in graphs.

CCS Concepts: • **Theory of computation** → **Parameterized complexity and exact algorithms; Graph algorithms analysis;**



Additional Key Words and Phrases: Hadwiger number, exponential-time hypothesis, exact algorithms, edge contraction problems

ACM Reference format:

Fedor V. Fomin, Daniel Lokshantov, Ivan Mihajlin, Saket Saurabh, and Meirav Zehavi. 2021. Computation of Hadwiger Number and Related Contraction Problems: Tight Lower Bounds. *ACM Trans. Comput. Theory* 13, 2, Article 10 (March 2021), 25 pages.

<https://doi.org/10.1145/3448639>

A preliminary version of this article appeared in ICALP 2020 [10]. Fomin received funding from Research Council of Norway via the project MULTIVAL. Saurabh received funding from European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant no. 819416), and Swarnajayanti Fellowship grant

DST/SJF/MSA-01/2017-18.   Lokshantov and Zehavi received funding from the United States–Israel Binational Science Foundation Grant no. 2018302. Meirav Zehavi received funding from the Israel Science Foundation (ISF) under the Individual Research Grant program (Grant no. 1176/18).

Authors' addresses: F. V. Fomin, University of Bergen, Bergen, Thormohlensgate Street, Norway; email: fomin@ii.uib.no; D. Lokshantov, University of California, Santa Barbara, University Road Street, California, USA; email: daniello@ucsb.edu; I. Mihajlin, University of California, San Diego, Alcalá Part Street, California; email: imikhail@cs.ucsd.edu; S. Saurabh, The Institute for Mathematical Sciences, HBNI, Chennai, Cross Street, India, University of Bergen, Bergen, Thormohlensgate Street, Norway; email: imikhail@cs.ucsd.edu; M. Zehavi, Ben-Gurion University, Beersheba, Ben-Gurion Street, Israel; email: meiravze@bgu.ac.il.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

1942-3454/2021/03-ART10 \$15.00

<https://doi.org/10.1145/3448639>

1 INTRODUCTION

The *Hadwiger number* $h(G)$ of a graph G is the largest number h for which the complete graph K_h is a minor of G . Equivalently, $h(G)$ is the maximum size of the largest complete graph that can be obtained from G by contracting edges. It is named after Hugo Hadwiger, who conjectured in 1943 that the Hadwiger number of G is always at least as large as its chromatic number. According to Bollobás, Catlin, and Erdős, this conjecture remains “one of the deepest unsolved problems in graph theory” [5].

The Hadwiger number of an n -vertex graph G can be easily computed in time $n^{O(n)}$ by brute-forcing through all possible partitions of the vertex set of G into connected sets, contracting each set into one vertex and checking whether the resulting graph is a complete graph. The question whether the Hadwiger number of a graph can be computed in single-exponential $2^{O(n)}$ time was previously asked in [2, 7, 16]. Our main result provides a negative answer to this open question.

THEOREM 1.1. *Unless the Exponential Time Hypothesis (ETH) is false, there does not exist an algorithm computing the Hadwiger number of an n -vertex graph in time $n^{o(n)}$.*

The interest in the complexity of the Hadwiger number is naturally explained by the recent developments in the area of exact exponential algorithms, that is, algorithms solving intractable problems significantly faster than the trivial exhaustive search, though still in exponential time [9]. Within the last decade, significant progress on upper and lower bounds of exponential algorithms has been achieved. Drastic improvements over brute-force algorithms were obtained for a number of fundamental problems like GRAPH COLORING [4] and HAMILTONICITY [3]. On the other hand, by making use of the ETH, lower bounds could be obtained for 2-CSP [18] or for SUBGRAPH ISOMORPHISM and GRAPH HOMOMORPHISM [7].

GRAPH MINOR (deciding whether a graph G contains a graph H as a minor) is a fundamental problem in graph theory and graph algorithms. GRAPH MINOR could be seen as special case of a general graph embedding problem where one wants to embed a graph H into graph G . In what follows, we will use n to denote the number of vertices in G and h to denote the number of vertices in H . By the theorem of Robertson and Seymour [17], there exists a computable function f and an algorithm that, for given graphs G and H , checks in time $f(h) \cdot n^3$ whether H is a minor of G . Thus, the problem is fixed-parameter tractable (FPT) being parameterized by H . On the other hand, Cygan et al. [7] proved that unless the ETH fails, this problem cannot be solved in time $n^{o(n)}$ even in the case when $|V(G)| = |V(H)|$. Other interesting embedding problems that are strongly related to GRAPH MINOR include the following problems.

- SUBGRAPH ISOMORPHISM: Given two graphs G and H , decide whether G contains a subgraph isomorphic to H . This problem cannot be solved in time $n^{o(n)}$ when $|V(G)| = |V(H)|$, unless the ETH fails [7]. In the special case called CLIQUE, when H is a clique, a brute-force algorithm checking for every vertex subset of G whether it is a clique of size h solves the problem in time $n^{O(h)}$. The same algorithm also runs in single-exponential time $O(2^n n^2)$. It is also known that CLIQUE is $W[1]$ -hard parameterized by h and cannot be solved in time $f(h) \cdot n^{o(h)}$ for any function f unless the ETH fails [6, 8].
- GRAPH HOMOMORPHISM: Given two graphs G and H , decide whether there exists a homomorphism from H to G . (A *homomorphism* $H \rightarrow G$ from an undirected graph H to an undirected graph G is a mapping from the vertex set of H to that of G such that the image of every edge of H is an edge of G .) This problem is trivially solvable in time $n^{O(h)}$, and an algorithm of running time $n^{o(h)}$ for this problem would yield the failure of the ETH [7]. However, for the special case of G being a clique, GRAPH HOMOMORPHISM is equivalent to n -COLORING (deciding whether the chromatic number of H is at most n), and thus is solvable in single-exponential time $2^h \cdot h^{O(1)}$ [4, 15]. When the graph H is a complete

graph, the problem is equivalent to finding a clique of size h in G , and then is solvable in time $2^n \cdot n^{O(1)}$.

- TOPOLOGICAL GRAPH MINOR:** Given two graphs G and H , decide whether G contains H as a topological minor. (We say that a graph H is a subdivision of a graph G if G can be obtained from H by contracting only edges incident with at least one vertex of degree two. A graph H is called a *topological minor* of a graph G if a subdivision of H is isomorphic to a subgraph of G .) This problem is, perhaps, the closest “relative” of GRAPH MINOR. Grohe et al. [12] gave an algorithm of running time $f(h) \cdot n^3$ for this problem for some computable function f . Similar to GRAPH MINOR and SUBGRAPH ISOMORPHISM, this problem cannot be solved in time $n^{o(n)}$ when $|V(G)| = |V(H)|$, unless the ETH fails [7]. However, for the special case of the problem with H being a complete graph, Lingas and Wahlen [16] gave a single-exponential algorithm solving the problem in time $2^{O(n)}$.

Thus, all the above graph embedding “relatives” of GRAPH MINOR are solvable in single-exponential time when graph H is a clique. However, from the perspective of exact exponential algorithms, Theorem 1.1 implies that finding the largest clique minor is the most difficult problem out of them all. This is why we find the lower bound provided by Theorem 1.1 surprising. Moreover, from the perspective of parameterized complexity, finding a clique minor of size h , which is FPT, is actually easier than finding a clique (as a subgraph) of size h , which is W[1]-hard, as well as from finding an h -coloring of a graph, which is para-NP-hard.

Theorem 1.1 also answers another question of Cygan et al. [7], who asked whether deciding if a graph H can be obtained from a graph G only by edge contractions, could be resolved in single-exponential time. By Theorem 1.1, the existence of such an algorithm is highly unlikely even when the graph H is a complete graph. Moreover, the technique developed to prove Theorem 1.1, appears to be extremely useful to rule out the existence of $n^{o(n)}$ -time algorithms for various contraction problems. We formalize our results with the following \mathcal{F} -CONTRACTION problem. Let \mathcal{F} be a graph class. Given a graph G and $t \in \mathbb{N}$, the task is to decide whether there exists a subset $F \subseteq E(G)$ of size at most t such that $G/F \in \mathcal{F}$ (where G/F is the graph obtained from G by contracting the edges in F). We prove that in each of the cases of \mathcal{F} -CONTRACTION where \mathcal{F} is the family of chordal graphs, interval graphs, proper interval graphs, threshold graphs, trivially perfect graphs, split graphs, complete split graphs, and perfect graphs, unless the ETH fails, \mathcal{F} -CONTRACTION is not solvable in time $n^{o(n)}$.

Technical Details. A summary of the reductions presented in this article is given in Figure 1. To prove our lower bounds, we first revisit the proof of Cygan et al. [7] for the ETH-hardness of a problem called LIST SUBGRAPH ISOMORPHISM. Informally, in this problem we are given two graphs G and H on the same number of vertices, as well as a list of vertices in H for each vertex in G , and we need to find a copy of G in H so that each vertex u in G is mapped to a vertex v in H that belongs to its list (i.e., v belongs to the list of u). We prove that the instances produced by the reduction (after some modification) of [7] have a very useful property that we crucially exploit later. Specifically, we construct a proper coloring of G as well as a proper coloring of H , and show that every vertex v in H that belongs to the list of some vertex u is, in fact, of the same color as u .

Having proved the above, we turn to prove the ETH-hardness of a special case of CLIQUE CONTRACTION where the input graph is highly structured. To this end, we introduce an intermediate problem called CROSS MATCHING. Informally, in this problem we are given a graph L with a partition (A, B) of its vertex set, and need to find a perfect matching between A and B whose contraction gives a clique. To see the connection between this problem and LIST SUBGRAPH ISOMORPHISM, think of the subgraph of L induced by one side of the partition—say, A —as a representation of the

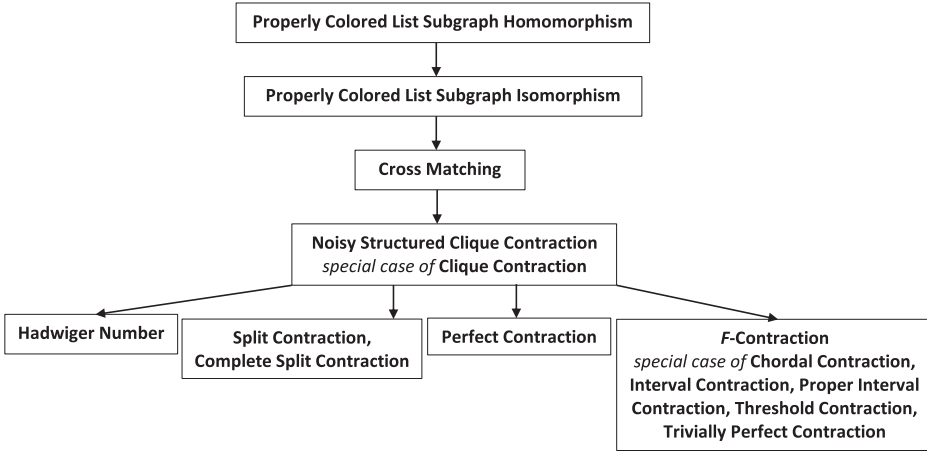


Fig. 1. A summary of the problems considered in this article, and the reductions between them.

complement of G , and the subgraph of L induced by the other side of the partition as a representation of H . Then, the edges that go across A and B in a perfect matching can be thought of as a mapping of the vertices of G to the vertices of H . The crossing edges of L are easily defined such that necessarily a vertex of G can only be matched to a vertex in its list. In particular, we would like to enforce that every “non-edge” of the complement of G (which corresponds to an edge of G) would have to be mapped to an edge of H in order to obtain a clique. However, the troublesome part is that non-edges of the complement of G may also be “filled” (to eventually get a clique) using crossing edges rather than only edges of H . To argue that this critical issue does not arise, we crucially rely on the proper colorings of G and H .

Now, for the connection between CROSS MATCHING and CLIQUE CONTRACTION, note that a solution to an instance of CROSS MATCHING is clearly a solution to the instance of CLIQUE CONTRACTION defined by the same graph, but the other direction is not true. By adding certain vertices and edges to the graph of an instance of CROSS MATCHING, we enforce all solutions to be perfect matchings between A and B . In particular, we construct the instances of CLIQUE CONTRACTION in a highly structured manner that allows us to derive not only the ETH-hardness of CLIQUE CONTRACTION itself, but to build upon them and further derive the ETH-hardness for a wide variety of other contraction problems. In particular, we show that the addition of “noise” (i.e., extra vertices and edges) to any structured instance of CLIQUE CONTRACTION has very limited effect. Roughly speaking, we show that the edges in the “noise” and the edges going across the “noise” and core of the graph (i.e., the original vertices corresponding to the structured instance of CLIQUE CONTRACTION) are not “helpful” when trying to create a clique on the core (i.e., it is not helpful to try to use these edges in order to fill non-edges between vertices in the core). Depending on the contraction problem at hand, the noise is slightly different, but the proof technique stays the same—first showing that the core must yield a clique, and then using the argument above (in fact, in all cases but that of perfect graphs, we are able to invoke the argument as a black box) to show that the noise is, in a sense, irrelevant.

2 PRELIMINARIES

For a vector R with L entries and $i \in \{1, \dots, L\}$, let $R[i]$ be the value of the i -th entry of R . Unless specified otherwise, bases of logarithms are assumed to be 2.

Given a graph G , let $V(G)$ and $E(G)$ denote its vertex set and edge set, respectively. Given a subset $U \subseteq V(G)$, let $G[U]$ denote the subgraph of G induced by U , that is, $V(G[U]) = U$ and $E(G[U]) = \{\{u, v\} \in E(G) : u, v \in U\}$. Given a subset $F \subseteq E(G)$, let $V(F)$ denote the set of vertices that are incident in G to at least one edge in F , and let $G[F] = G[V(F)]$. We say that G contains a graph H as an induced subgraph if there exists $U \subseteq V(G)$ such that $G[U]$ and H are identical up to relabeling vertices (more precisely, isomorphic). The set of neighbors of a vertex $u \in V(G)$ is denoted by $N_G(u)$, that is, $N_G(u) = \{v \in V(G) : \{u, v\} \in E(G)\}$. When G is clear from context, we drop it from subscripts of notations. A *matching* M in G is a subset of $E(G)$ such that no two edges in M share an endpoint. In case every vertex in $V(G)$ is an endpoint of an edge in M , that is, $|M| = |V(G)|/2$, it is said that M is *perfect*. A function $c : V(G) \rightarrow \mathbb{N}$ is a *proper coloring* of G if for every edge $\{u, v\} \in E(G)$, $c(u) \neq c(v)$. The *complement* of G , denoted by \bar{G} , is the graph with vertex set $V(G)$ and edge set $\{\{u, v\} \notin E(G) : u, v \in V(G), u \neq v\}$. The square G^2 of G is the graph on vertex set $V(G)$ and edge set $\{\{u, v\} : \{u, v\} \in E(G) \text{ or there exists } w \in V(G) \text{ with } \{u, w\}, \{v, w\} \in E(G)\}$.

Given an edge $e = \{u, v\} \in E(G)$, the *contraction* of e in G is the operation that replaces u and v by a new vertex that is adjacent to all vertices previously adjacent to u or v (or both), where the resulting graph is denoted by G/e . In other words, $V(G/e) = (V(G) \setminus \{u, v\}) \cup \{x\}$ for some new vertex x , and $E(G/e) = \{\{s, t\} \in E(G) : s, t \notin \{u, v\}\} \cup \{\{s, x\} : s \in N(u) \cup N(v)\}$. More generally, given a subset $F \subseteq E(G)$, the *contraction* of F in G is the operation that replaces each connected component C of $G[F]$ by a new vertex x_C that is adjacent to all vertices previously adjacent to at least one vertex in C , where the resulting graph is denoted by G/F . A graph H is said to be a *minor* of a graph G if H can be obtained from G by a series of vertex deletions, edge deletions, and edge contractions. For any $h \in \mathbb{N}$, the clique on h vertices is denoted by K_h , and the cycle on h vertices is denoted by C_h . The *Hadwiger number* of a graph G is the largest $h \in \mathbb{N}$ such that K_h is a minor of G .

To obtain (essentially) tight conditional lower bounds for the running times of algorithms, we rely on the *ETH* [13, 14]. To formalize its statement, we remind one that given a formula φ in conjunctive normal form (CNF) with n variables and m clauses, the task of CNF-SAT is to decide whether there is a truth assignment to the variables that satisfies φ . In the p -CNF-SAT problem, each clause is restricted to have at most p literals. Then, the ETH asserts that 3-CNF-SAT cannot be solved in time $2^{o(n)}$.

3 LOWER BOUND: PROP-COLORED LIST SUBGRAPH ISOMORPHISM

In this section, we build upon the work of Cygan et al. [7] and show a lower bound for a problem called Properly Colored List Subgraph Isomorphism (PROP-COL LSI). Intuitively, PROP-COL LSI is a variant of SPANNING SUBGRAPH ISOMORPHISM where given two graphs G and H , we ask whether G is isomorphic to some spanning subgraph of H . The input to the variant consists also of proper colorings of G and H and an additional labeling of vertices in G by subsets of vertices in H of the same color, so that each vertex in G can be mapped only to vertices in H contained in its list. Formally, it is defined as follows.

PROPERLY COLORED LIST SUBGRAPH ISOMORPHISM (PROP-COL LSI)

Input: Graphs G and H with proper colorings $c_G : V(G) \rightarrow \{1, \dots, k\}$ and $c_H : V(H) \rightarrow \{1, \dots, k\}$ for some $k \in \mathbb{N}$, respectively, and a function $\ell : V(G) \rightarrow 2^{V(H)}$ such that for every $u \in V(G)$ and $v \in \ell(u)$, $c_G(u) = c_H(v)$.

Question: Does there exist a bijective function $\varphi : V(G) \rightarrow V(H)$ such that (i) for every $\{u, v\} \in E(G)$, $\{\varphi(u), \varphi(v)\} \in E(H)$, and (ii) for every $u \in V(G)$, $\varphi(u) \in \ell(u)$?

Notice that as the function φ above is bijective rather than only injective, we seek a spanning subgraph. Our objective is to prove the following statement.

LEMMA 3.1. *Unless the ETH is false, there does not exist an algorithm that solves PROP-COL LSI in time $n^{o(n)}$ where $n = |V(G)|$.*

In [7], the authors considered the two problems defined below. Intuitively, the second problem is defined as PROP-COL LSI when no proper colorings of H and G are given (and hence the labeling of vertices in G is not restricted accordingly). Moreover, the first problem is defined as the second problem when we seek a homomorphism rather than an isomorphism (i.e., the sought function φ may not be injective) and also $|V(G)|$ may not be equal to $|V(H)|$ (thus φ may neither be onto).

List Subgraph Homomorphism (LSH)

Input: Graphs G and H , and a function $\ell : V(G) \rightarrow 2^{V(H)}$.

Question: Does there exist a function $\varphi : V(G) \rightarrow V(H)$ such that (i) for every $\{u, v\} \in E(G)$, $\{\varphi(u), \varphi(v)\} \in E(H)$, and (ii) for every $u \in V(G)$, $\varphi(u) \in \ell(u)$?

List Subgraph Isomorphism (LSI)

Input: Graphs G and H where $|V(G)| = |V(H)|$, and a function $\ell : V(G) \rightarrow 2^{V(H)}$.

Question: Does there exist a bijective function $\varphi : V(G) \rightarrow V(H)$ such that (i) for every $\{u, v\} \in E(G)$, $\{\varphi(u), \varphi(v)\} \in E(H)$, and (ii) for every $u \in V(G)$, $\varphi(u) \in \ell(u)$?

We remark that LSH may be termed LIST GRAPH HOMOMORPHISM since no subgraph is sought—homomorphism, unlike isomorphism, does not require one to “use” all edges, or even all vertices, of the target graph H ; however, we preferred LSH to emphasize this and, in particular, to be consistent with [7].

The proof of hardness of LSI consists of two parts:

- Showing the ETH-hardness of LSH.
- Giving a fine-grained reduction from LSH to LSI.

We cannot use the hardness of LSI as a black box because PROP-COL LSI is a special case of LSI. Nevertheless, we will prove that the instances generated by the reduction (with a minor crucial modification) of Cygan et al. [7] have the additional properties required to make them instances of our special case.

Lower Bound: PROPERLY COLORED SUBGRAPH HOMOMORPHISM. Adapting the scheme of Cygan et al. [7] to our purpose, we will first show that finding a homomorphism remains hard if it has to preserve a given proper coloring.

Properly Colored List Subgraph Homomorphism (Prop-Col LSH)

Input: Graphs G and H with proper colorings $c_G : V(G) \rightarrow \{1, \dots, k\}$ and $c_H : V(H) \rightarrow \{1, \dots, k\}$ for some $k \in \mathbb{N}$, respectively, and a function $\ell : V(G) \rightarrow 2^{V(H)}$ such that for every $u \in V(G)$ and $v \in \ell(u)$, $c_G(u) = c_H(v)$.

Question: Does there exist a function $\varphi : V(G) \rightarrow V(H)$ such that (i) for every $\{u, v\} \in E(G)$, $\{\varphi(u), \varphi(v)\} \in E(H)$, and (ii) for every $u \in V(G)$, $\varphi(u) \in \ell(u)$?

In [7], the authors gave a reduction from the 3-COLORING problem on n -vertex graphs of degree 4 (which is known not to be solvable in time $2^{o(n)}$ unless the ETH fails), which generates equivalent instances (G', H', ℓ) of LSH where both $|V(G')|$ and $|V(H')|$ are bounded by $O(\frac{n}{\log n})$. This proves that LSH is not solvable in time $n^{o(n)}$ where $n = \max\{|V(G)|, |V(H)|\}$ unless the ETH fails. For their reduction, Cygan et al. [7] considered the notion of a *grouping* (also known as *quotient graph*) \tilde{G} of

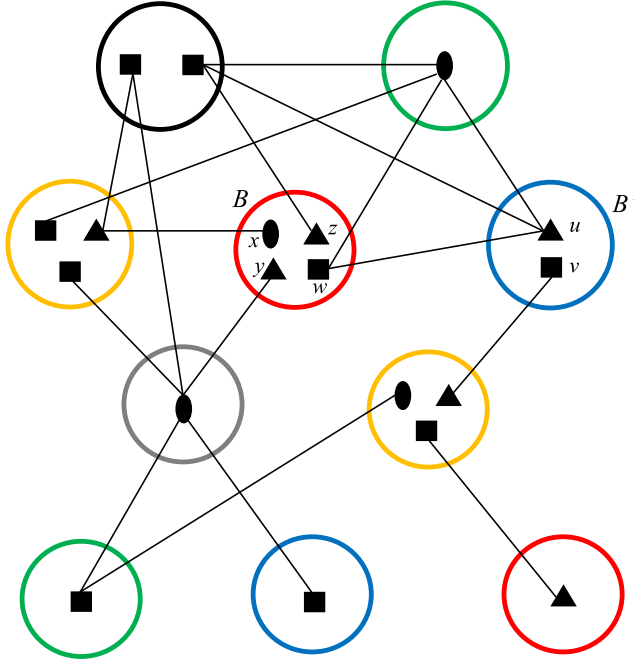


Fig. 2. The reduction in Definition 3.3. The vertices of G are depicted by black shapes, where each distinct shape represents a different color (say, square is 1, rectangle is 2, and oval is 3), and the vertices of \tilde{G} are depicted by circles enclosing the vertex sets identifies with them, where the color of a vertex is the color of its circle (say, black is 1, green is 2, yellow is 3, red is 4, blue is 5, and gray is 6). Edges (of both graphs) are depicted by black lines. (The graph \tilde{H} is not shown). Then, the function ϕ_B is defined as follows: $\phi_B(1) = z$, $\phi_B(2) = \phi_B(5) = w$, $\phi_B(3) = x$, $\phi_B(4) = 0$, and $\phi_B(6) = y$. Moreover, the function $\phi_{B'}$ is defined as follows: $\phi_{B'}(1) = \phi_{B'}(2) = \phi_{B'}(4) = u$, $\phi_{B'}(3) = v$, and $\phi_{B'}(5) = \phi_{B'}(6) = 0$. With respect to B and B' , the labeling ℓ is defined as follows: $\ell(B) = \{(R, 4) : R[1] \neq 0, R[2] = R[5] \neq 0, R[3] \neq 0, R[4] = 0, R[6] \neq 0\}$ and $\ell(B') = \{(R, 5) : R[1] = R[2] = R[4] \neq 0, R[3] \neq 0, R[5] = R[6] = 0\}$.

a graph G is a graph with vertex set $V(\tilde{G}) = \{B_1, B_2, \dots, B_t\}$ where (B_1, B_2, \dots, B_t) is a partition of $V(G)$ for some $t \in \mathbb{N}$ and for any distinct $i, j \in \{1, \dots, t\}$, the vertices B_i and B_j are adjacent in \tilde{G} if and only if there exist $u \in B_i$ and $v \in B_j$ that are adjacent in G . Specifically, they computed a grouping with a coloring having specific properties as stated in the following lemma (see also Figure 2).

LEMMA 3.2 (LEMMA 3.2 IN [7]). *For any constant $d \geq 1$, there exist positive integers $\lambda = \lambda(d)$, $n_0 = n_0(d)$ and a polynomial time algorithm that for a given graph G on $n \geq n_0$ vertices of maximum degree d and a positive integer $r \leq \sqrt{\frac{n}{2\lambda}}$, finds a grouping \tilde{G} of G and a coloring $\tilde{c} : V(\tilde{G}) \rightarrow [\lambda r]$ with the following properties:*

- (1) $|V(\tilde{G})| \leq |V(G)|/r$;
- (2) the coloring \tilde{c} is a proper coloring of \tilde{G}^2 ;
- (3) each vertex of \tilde{G} is an independent set in G ;
- (4) for any edge $\{B_i, B_j\} \in E(\tilde{G})$, there exists exactly one pair $(u, v) \in B_i \times B_j$ such that $\{u, v\} \in E(G)$.

Now, we describe the reduction of [7]. Here, without loss of generality, it is assumed that G has no isolated vertices, else they can be removed. An explanation of the intuition behind this somewhat technical definition is given below it.

Definition 3.3. For any instance G of 3-COLORING where G has degree d and a positive integer $r = o(\sqrt{|V(G)|})$, the instance $\text{reduce}(G) = (\tilde{G}, \tilde{H}, \ell)$ of LSH is defined as follows.

- **The graph \tilde{G} .** Let \tilde{G} and $\tilde{c} : V(\tilde{G}) \rightarrow \{1, 2, \dots, L\}$ be the grouping and coloring given by Lemma 3.2 where $L = \lambda(d)r$. Additionally, for each $B \in V(\tilde{G})$, define $\phi_B : \{1, 2, \dots, L\} \rightarrow B \cup \{0\}$ as follows: for any $i \in \{1, 2, \dots, L\}$, if there exists (u, v, B') such that $u \in B$ and $v \in B'$, $\{u, v\} \in E(G)$ and $\tilde{c}(B') = i$, then $\phi_B(i) = u$, and otherwise $\phi_B(i) = 0$.
- **The graph \tilde{H} .** Let $V(\tilde{H}) = \{(R, l) : R \in \{0, 1, 2, 3\}^L, l \in L\}$ and $E(\tilde{H}) = \{(R, l), (R', l')\} : R[l] \neq R'[l]\}$.
- **The labeling ℓ .** For any $B \in V(\tilde{G})$, let $\ell(B)$ contain all vertices $(R, l) \in V(\tilde{H})$ such that $\tilde{c}(B) = l$, and there exists $f : B \rightarrow \{1, 2, 3\}$ such that for all $i \in \{1, 2, \dots, L\}$, either $\phi_B(i) = R[i] = 0$ or both $\phi_B(i) \neq 0$ and $f(\phi_B(i)) = R[i]$.

We remark that the uniqueness of u (if it exists), and thus the validity of ϕ_B , follows from Properties 2 and 4 in Lemma 3.2.

Intuitively, for every vertex $B \in V(\tilde{G})$, the function ϕ_B can be interpreted as follows. It is the assignment, for every possible color $i \in \{1, \dots, L\}$, of the unique vertex u within the vertex set identified with B itself that is adjacent to some vertex in the vertex subset identified with some vertex $B' \in V(\tilde{G})$ colored i , if such a vertex u exists (else the assignment is of 0). In a sense, B thus stores the information on the identity of each vertex within it that is adjacent (in G) to some vertex outside of it, where each such internal vertex is uniquely accessed by specifying the color of the vertex in \tilde{G} whose identified vertex set contains the *neighbor*. With respect to the graph \tilde{H} and labeling ℓ , we interpret each vertex $(R, l) \in V(\tilde{H})$ as a “placeholder” (i.e., potential assignment of the sought function φ) for any vertex $B \in V(\tilde{G})$ that “complies with the pattern encoded by the pair (R, l) ” as follows. First and straightforwardly, B must be colored l . Here, we remind one that the colors of vertices in \tilde{G} belong to $\{1, \dots, L\}$, while vertices in G are colored 1, 2, or 3 only. Then, the second requirement is that we can recolor (by f) the vertices in B so that the color of each vertex in B that is adjacent (in G) to some vertex outside B is as encoded by the vector R —that is, for each color $i \in \{1, \dots, L\}$, if the vertex $\phi_B(i)$ is defined (i.e., $\phi_B(i) \neq 0$), then its color (which is 1, 2, or 3) must be equal to the i -th entry of R . (Further intuition is given in Figure 2.)

Now, we state the correctness of the reduction.

LEMMA 3.4 (LEMMA 3.3 IN [7]). *For any instance G of 3-COLORING where G is an n -vertex graph of degree d , and a positive integer $r = o(\sqrt{|V(G)|})$, the instance $\text{reduce}(G) = (\tilde{G}, \tilde{H}, \ell)$ is computable in time polynomial in the sizes of G , \tilde{G} , and \tilde{H} , and has the following properties.*

- G is a Yes-instance of 3-COLORING if and only if $(\tilde{G}, \tilde{H}, \ell)$ is a Yes-instance of LSH.
- $|V(\tilde{G})| \leq n/r$, and $|V(\tilde{H})| \leq \gamma(d)^r$ where γ is some computable function of d .

We next prove that we can add colorings to the instance $\text{reduce}(G) = (\tilde{G}, \tilde{H}, \ell)$ of LSH in order to cast it as an instance of PROP-COL LSH while making a minor mandatory modification to the graph \tilde{H} .

LEMMA 3.5. *Given an instance $\text{reduce}(G) = (\tilde{G}, \tilde{H}, \ell)$ of LSH, an equivalent instance $(\tilde{G}, \tilde{H}', c_{\tilde{G}}, c_{\tilde{H}'}, \ell)$ of PROP-COL LSH, where \tilde{H}' is a subgraph of \tilde{H} , is computable in polynomial time.*

PROOF. Define $c_{\tilde{G}} = \tilde{c}$ where \tilde{c} is the coloring of \tilde{G} in Definition 3.3. Additionally, let \tilde{H}' be the subgraph of \tilde{H} induced by the vertex set $\{(R, l) \in V(\tilde{H}) : \text{there exists } B \in V(\tilde{G}) \text{ such that } (R, l) \in \ell(B)\}$. Then, define $c_{\tilde{H}'} : V(\tilde{H}') \rightarrow \{1, 2, \dots, L\}$ as follows: for any $(R, l) \in V(\tilde{H}')$, define $c_{\tilde{H}'}((R, l)) = l$. Notice that, by the definition of $V(\tilde{H}')$, every set assigned by ℓ is a subset of $V(\tilde{H}')$.

First, we assert that $(\tilde{G}, \tilde{H}', c_{\tilde{G}}, c_{\tilde{H}'}, \ell)$ is an instance of PROP-COL LSH. To this end, we need to verify that the three following properties hold.

- (1) $c_{\tilde{G}}$ is a proper coloring of \tilde{G} .
- (2) $c_{\tilde{H}'}$ is a proper coloring of \tilde{H}' .
- (3) For every $B \in V(\tilde{G})$ and $(R, l) \in \ell(B)$, it holds that $c_{\tilde{G}}(B) = c_{\tilde{H}'}((R, l))$.

By the definition of $c_{\tilde{G}}$, it is a proper coloring of \tilde{G}^2 , which is a supergraph of \tilde{G} . Thus, $c_{\tilde{G}}$ is a proper coloring of \tilde{G} .

Now, we argue that $c_{\tilde{H}'}$ is a proper coloring of \tilde{H}' . To this end, consider some edge $\{(R, l), (R', l')\} \in E(\tilde{H}')$. We need to show that $c_{\tilde{H}'}((R, l)) \neq c_{\tilde{H}'}((R', l'))$. By the definition of $c_{\tilde{H}'}$, we have that $c_{\tilde{H}'}((R, l)) = l$ and $c_{\tilde{H}'}((R', l')) = l'$, and therefore it suffices to show that $l \neq l'$. By the definition of $E(\tilde{H})$ (which is a superset of $E(\tilde{H}')$), we have that $R[l'] \neq R'[l]$. Thus, necessarily at least one among $R[l']$ and $R'[l]$ is not 0, and so we suppose w.l.o.g. that $R[l']$ is not 0. Furthermore, since $(R, l) \in V(\tilde{H}')$, we have that there exists $B \in E(\tilde{G})$ such that $(R, l) \in \ell(B)$. Thus,

$$-\tilde{c}(B) = l.$$

– There exists $f : B \rightarrow \{1, 2, 3\}$ such that for all $i \in \{1, 2, \dots, L\}$, either $\phi_B(i) = R[i] = 0$ or both $\phi_B(i) \neq 0$ and $f(\phi_B(i)) = R[i]$.

From the second property, and because $R[l'] \neq 0$, we necessarily have that both $\phi_B(l') \neq 0$ and $f(\phi_B(l')) = R[l']$. In particular, by the definition of ϕ_B , having $\phi_B(l') \neq 0$ means that there exists (u, v, B') such that $u \in B, v \in B', \{u, v\} \in E(G)$, and $\tilde{c}(B') = l'$. By the definition of \tilde{G} as a grouping of G , having $u \in B, v \in B'$, and $\{u, v\} \in E(G)$ implies that $\{B, B'\} \in E(\tilde{G})$. Because \tilde{c} is a proper coloring of \tilde{G} , this means that $\tilde{c}(B) \neq \tilde{c}(B')$. Since $\tilde{c}(B) = l$ and $\tilde{c}(B') = l'$, we derive that $l \neq l'$. Hence, $c_{\tilde{H}'}$ is indeed a proper coloring of \tilde{H}' .

To conclude that $(\tilde{G}, \tilde{H}', c_{\tilde{G}}, c_{\tilde{H}'}, \ell)$ is indeed an instance of PROP-COL LSH, it remains to assert that for every $B \in V(\tilde{G})$ and $(R, l) \in \ell(B)$, it holds that $c_{\tilde{G}}(B) = c_{\tilde{H}'}((R, l))$. To this end, consider some $B \in V(\tilde{G})$ and $(R, l) \in \ell(B)$. By the definition of ℓ (recall Definition 3.3), $(R, l) \in \ell(B)$ implies that $\tilde{c}(B) = l$. As $c_{\tilde{G}} = \tilde{c}$, we have that $c_{\tilde{G}}(B) = l$. Moreover, the definition of $c_{\tilde{H}'}$ directly implies that $c_{\tilde{H}'}((R, l)) = l$. Thus, $c_{\tilde{G}}(B) = c_{\tilde{H}'}((R, l))$.

Finally, we argue that $(\tilde{G}, \tilde{H}, \ell)$ is a Yes-instance of LSH if and only if $(\tilde{G}, \tilde{H}', c_{\tilde{G}}, c_{\tilde{H}'}, \ell)$ is a Yes-instance of PROP-COL LSH. In one direction, because \tilde{H}' is a subgraph of \tilde{H} , it is immediate that if $(\tilde{G}, \tilde{H}', c_{\tilde{G}}, c_{\tilde{H}'}, \ell)$ is a Yes-instance of PROP-COL LSH, then so is $(\tilde{G}, \tilde{H}, \ell)$. For the other direction, suppose that $(\tilde{G}, \tilde{H}, \ell)$ is a Yes-instance of LSH. Thus, there exists a function $\varphi : V(\tilde{G}) \rightarrow V(\tilde{H})$ such that (i) for every $\{B, B'\} \in E(\tilde{G})$, $\{\varphi(B), \varphi(B')\} \in E(\tilde{H})$, and (ii) for every $B \in V(\tilde{G})$, $\varphi(B) \in \ell(B)$. In particular, directly by the definition of $V(\tilde{H}')$, the second condition implies that for every $B \in V(\tilde{G})$, it holds that $\varphi(B) \in V(\tilde{H}')$. Thus, because \tilde{H}' is an induced subgraph of \tilde{H} , it holds that for every $\{B, B'\} \in E(\tilde{G})$, $\{\varphi(B), \varphi(B')\} \in E(\tilde{H}')$. Therefore, φ witnesses that $(\tilde{G}, \tilde{H}', c_{\tilde{G}}, c_{\tilde{H}'}, \ell)$ is a Yes-instance of PROP-COL LSH. \square

We are now ready to assert the hardness of PROP-COL LSH. The proof is based on Lemmas 3.2, 3.4, and 3.5.

LEMMA 3.6. *Unless the ETH is false, there does not exist an algorithm that solves PROP-COL LSH in time $n^{o(n)}$ where $n = \max(|V(G)|, |V(H)|)$.*

PROOF. Targeting a contradiction, suppose that there exists an algorithm, denoted by LSHAlg, that solves PROP-COL LSH in time $n^{o(n)}$ where $n = \max(|V(G)|, |V(H)|)$ for input graphs G and H . We will show that this implies the existence of an algorithm, denoted by ColAlg, that solves 3-COLORING on graphs of maximum degree 4 in time $2^{o(n)}$ where n is the number of vertices of the input graph, which contradicts the ETH and hence completes the proof.

The execution of ColAlg is as follows. Given an instance G of 3-COLORING on graphs of maximum degree 4, ColAlg constructs the instance $\text{reduce}(G) = (\tilde{G}, \tilde{H}, \ell)$ of LSH in Definition 3.3 with $r = \lceil \log_{\gamma(4)}(n/\log n) \rceil$ where $n = |V(G)|$. Here, γ is the function stated in Lemma 3.4. By Lemma 3.4, $\text{reduce}(G) = (\tilde{G}, \tilde{H}, \ell)$ is computable in time polynomial in the sizes of G , \tilde{G} , and \tilde{H} , and has the following properties:

- G is a Yes-instance of 3-COLORING if and only if $(\tilde{G}, \tilde{H}, \ell)$ is a Yes-instance of LSH.
- $|V(\tilde{G})| \leq n/r = O(n/\log n)$ and $|V(\tilde{H})| \leq \gamma(4)^r = O(n/\log n)$.

Then, ColAlg calls the polynomial-time algorithm in Lemma 3.5 with $(\tilde{G}, \tilde{H}, \ell)$ to construct an equivalent instance $(\tilde{G}, \tilde{H}', c_{\tilde{G}}, c_{\tilde{H}'}, \ell)$ of PROP-COL LSH, where \tilde{H}' is a subgraph of \tilde{H} . Lastly, ColAlg calls LSHAlg with $(\tilde{G}, \tilde{H}', c_{\tilde{G}}, c_{\tilde{H}'}, \ell)$ as input, and returns its answer.

Since the instance G of 3-COLORING was argued above to be equivalent to the instance $(\tilde{G}, \tilde{H}', c_{\tilde{G}}, c_{\tilde{H}'}, \ell)$ of PROP-COL LSH, the correctness of ColAlg directly follows. For the running time, denote $M = \max(|V(\tilde{G})|, |V(\tilde{H}')|)$, and notice that $M \leq O(n/\log n)$. Thus, because LSHAlg runs in time $M^{o(M)} \leq (n/\log n)^{o(n/\log n)} \leq 2^{o(n)}$, it follows that ColAlg runs in time $2^{o(n)}$. This completes the proof. \square

From Graph Homomorphism to Subgraph Isomorphism. In this part, we observe that the reduction of [7] from LSH to LSI can be essentially used as is to serve as a reduction from PROP-COL LSH to PROP-COL LSI. For the sake of completeness, we give the full details. We begin by adapting the Turing reduction of [7] from LSH to LSI.

LEMMA 3.7. *There is an $2^{O(n)}$ -time algorithm that, given an instance (G, H, c_G, c_H, ℓ) of PROP-COL LSH, returns $2^{O(n)}$ instances of PROP-COL LSI having input graphs on at most n vertices for $n := \max(|V(G)|, |V(H)|)$, such that (G, H, c_G, c_H, ℓ) is a Yes-instance of PROP-COL LSH if and only if at least one of the returned instances is a Yes-instance of PROP-COL LSI.*

PROOF. Given an instance (G, H, c_G, c_H, ℓ) of PROP-COL LSH, the algorithm works as follows. Without loss of generality, suppose that $V(H) = \{1, 2, \dots, |V(H)|\}$. Let $\mathcal{P} = \{P \in \mathbb{N}_0^{|V(H)|} : \sum_{i=1}^{|V(H)|} P[i] = |V(G)|\}$. That is, \mathcal{P} contains every vector with $|V(H)|$ entries that are non-negative integers whose sum is $|V(G)|$. Then, for each $P \in \mathcal{P}$, the algorithm returns one instance $(G, H_P, c_G, c_{H_P}, \ell_P)$ of PROP-COL LSI that is constructed as follows.

- The graph H_P is constructed from H by replacing each vertex $v \in V(H)$ with $P[v]$ copies of it, denoted $v_1, v_2, \dots, v_{P[v]}$. (Note that $P[v]$ can be equal to 0.) Then, we connect two vertices v_i to u_j in H_P if and only if v is connected to u in H . That is, $V(H_P) = \{v_i : v \in V(H), i \in \{1, 2, \dots, P[v]\}\}$ and $E(H_P) = \{\{u_i, v_j\} : \{u, v\} \in E(H), u_i, v_j \in V(H_P)\}$.
- For every vertex $u_i \in V(H_P)$, let $c_{H_P}(u_i) = c_H(u)$.
- For every vertex $u \in V(G)$, let $\ell_P(u) = \{v_i \in V(H_P) : v \in \ell(u)\}$.

This completes the description of the algorithm.

First, we consider some $P \in \mathcal{P}$ and assert that $(G, H_P, c_G, c_{H_P}, \ell_P)$ is indeed an instance of PROP-COL LSI. By the construction of $V(H_P)$ and since $\sum_{i=1}^{|V(H)|} P[i] = |V(G)|$, we have that $|V(G)| = |V(H_P)|$. Clearly, as (G, H, c_G, c_H, ℓ) is an instance of PROP-COL LSH, we have that c_G is a proper coloring of G . Now, consider an edge $\{u_i, v_j\} \in E(H_P)$. Then, $\{u, v\} \in E(H)$, and since c_H is a proper coloring of H (as (G, H, c_G, c_H, ℓ) is an instance of PROP-COL LSH), this means that $c_H(u) \neq c_H(v)$. By definition, $c_{H_P}(u_i) = c_H(u)$ and $c_{H_P}(v_j) = c_H(v)$, and therefore $c_{H_P}(u_i) \neq c_{H_P}(v_j)$. Thus, c_{H_P} is a proper coloring of H_P . Lastly, consider some vertices $u \in V(G)$ and $v_i \in \ell_P(u)$. By the definition of ℓ_P , we have that $v \in \ell(P)$. Therefore, as (G, H, c_G, c_H, ℓ) is an instance of PROP-COL LSH, $c_G(u) = c_H(v)$. Thus, because $c_{H_P}(v_i) = c_H(v)$, we have that $c_G(u) = c_{H_P}(v_i)$.

Now, we consider the number of instances returned by the algorithm along with its running time. Towards this, first note that $|\mathcal{P}| = \binom{|V(G)| + |V(H)| - 1}{|V(H)| - 1} \leq 4^n$. As the number of returned instances equals $|\mathcal{P}|$, it is upper bounded by $2^{O(n)}$ as required. Because each instance is computed in polynomial time, we also get that the running time of the algorithm is bounded by $2^{O(n)}$.

Finally, we consider the correctness of the algorithm. In one direction, suppose that at least one of the returned instances is a Yes-instance of PROP-COL LSI. Then, there exists $P \in \mathcal{P}$ such that $(G, H_P, c_G, c_{H_P}, \ell_P)$ is a Yes-instance of PROP-COL LSI. Thus, there exists a bijective function $\varphi_P : V(G) \rightarrow V(H_P)$ such that (i) for every $\{u, v\} \in E(G)$, $\{\varphi_P(u), \varphi_P(v)\} \in E(H_P)$, and (ii) for every $u \in V(G)$, $\varphi_P(u) \in \ell_P(u)$. We define a function $\varphi : V(G) \rightarrow V(H)$ as follows: for every $u \in V(G)$, let $\varphi(u) = v$ where $v \in V(H)$ is the vertex for which there exists $i \in \{1, 2, \dots, P[v]\}$ such that $\varphi_P(u) = v_i$. We now verify that φ is a solution to the instance (G, H, c_G, c_H, ℓ) of PROP-COL LSH. Firstly, by item (i) above, for every $\{u, v\} \in E(G)$, we have that $\{x_i, y_i\} \in E(H_P)$ where $x_i = \varphi_P(u)$ and $y_i = \varphi_P(v)$; by the definition of H_P , this means that $\{x, y\} \in E(H)$, and as $x = \varphi(u)$ and $y = \varphi(v)$ (by the definition of φ), we get that $\{\varphi(u), \varphi(v)\} \in E(H)$. Secondly, by item (ii) above, for every $u \in V(G)$, $v_i \in \ell_P(u)$ where $v_i = \varphi_P(u)$; by the definition of ℓ_P , we have that $v \in \ell_P(u)$, and by the definition of φ , we have that $v = \varphi(u)$, therefore $\varphi(u) \in \ell(u)$. Thus, we conclude that (G, H, c_G, c_H, ℓ) is a Yes-instance of PROP-COL LSH.

In the other direction, suppose that (G, H, c_G, c_H, ℓ) is a Yes-instance of PROP-COL LSH. Then, there exists a function $\varphi : V(G) \rightarrow V(H)$ such that (i) for every $\{u, v\} \in E(G)$, $\{\varphi(u), \varphi(v)\} \in E(H)$, and (ii) for every $u \in V(G)$, $\varphi(u) \in \ell(u)$. Let P be the vector with $|V(H)|$ entries where for each $i \in \{1, 2, \dots, |V(H)|\}$, $P[i] = |\varphi^{-1}(i)|$. Then, $\sum_{i=1}^{|V(H)|} P[i] = \sum_{i=1}^{|V(H)|} |\varphi^{-1}(i)| = |V(G)|$, and therefore $P \in \mathcal{P}$. Choose some arbitrary order $<$ on $V(G)$. Now, we define a function $\varphi_P : V(G) \rightarrow V(H_P)$ as follows: for every $u \in V(G)$, let $\varphi_P(u) = v_i$ where $v = \varphi(u)$ and $i = |\{w \in V(G) : w \leq u, v = \varphi(w)\}|$. It should be clear that φ_P is a bijection. Moreover, analogously to the previous direction, we assert that (i) for every $\{u, v\} \in E(G)$, $\{\varphi_P(u), \varphi_P(v)\} \in E(H_P)$, and (ii) for every $u \in V(G)$, $\varphi_P(u) \in \ell_P(u)$. Thus, $(G, H_P, c_G, c_{H_P}, \ell_P)$ is a Yes-instance of PROP-COL LSI, which means that at least one of the returned instances is a Yes-instance of PROP-COL LSI. \square

We are ready to complete the proof of Lemma 3.1.

PROOF OF LEMMA 3.1. Targeting a contradiction, suppose that there exists an algorithm, denoted by LSIAlg, that solves PROP-COL LSI in time $n^{o(n)}$ where $n = \max(|V(G)|, |V(H)|)$ for input graphs G and H . We will show that this implies the existence of an algorithm, denoted by LSHAlg, that solves PROP-COL LSH in time $n^{o(n)}$ where $n = \max(|V(G)|, |V(H)|)$ for input graphs G and H , which contradicts Lemma 3.6 and hence completes the proof.

The execution of LSHAlg is as follows. Given an instance (G, H, c_H, c_G, ℓ) of PROP-COL LSH, LSHAlg calls the algorithm in Lemma 3.7 so that in time $2^{O(n)}$ it obtains $2^{O(n)}$ instances of PROP-COL LSI having input graphs on at most n vertices for $n := \max(|V(G)|, |V(H)|)$, such that (G, H, c_G, c_H, ℓ) is a Yes-instance of PROP-COL LSH if and only if at least one of the returned




Fig. 3. The construction of an instance of CROSS MATCHING in the proof of Lemma 4.1.

instances is a Yes-instance of PROP-COL LSI. Then, it calls LSIAlg on each of the returned instances, and returns Yes if and only if at least one of these calls returns Yes. It should be clear that LSHAlg runs in time $n^{o(n)}$ and that it is correct. \square

4 LOWER BOUND FOR THE CROSS MATCHING PROBLEM

In this section, toward the proof of a lower bound for CLIQUE CONTRACTION, we prove a lower bound for an intermediate problem called CROSS MATCHING that somewhat resembles CLIQUE CONTRACTION, and which is defined as follows.

CROSS MATCHING

Input: A graph G with a partition (A, B) of $V(G)$ where $|A| = |B|$.

Question: Does there exist a perfect matching M in G such that every edge in M has one endpoint in A and the other in B , and G/M is a clique?

Our objective is to prove the following statement.

LEMMA 4.1. *Unless the ETH is false, there does not exist an algorithm that solves CROSS MATCHING in time $n^{o(n)}$ where $n = |A|$.*

PROOF. Towards a contradiction, suppose that there exists an algorithm, denoted by MatchingAlg, that solves CROSS MATCHING in time $n^{o(n)}$ where n is the number of vertices in the set A in the input. We will show that this implies the existence of an algorithm, denoted by LSIAlg, that solves PROP-COL LSI in time $n^{o(n)}$ where n is the number of vertices in the input graph G , thereby contradicting Lemma 3.1 and hence completing the proof.

We define the execution of LSIAlg as follows. Given an instance (G, H, c_G, c_H, ℓ) of PROP-COL LSI, LSIAlg constructs an instance (L, A, B) of CROSS MATCHING as follows (see Figure 3):

- $V(L) = V(\overline{G}) \cup V(H)$.
- $E(L) = E(\overline{G}) \cup E(H) \cup \{\{u, v\} : u \in V(G), v \in L(u)\}$.
- $A = V(\overline{G})$ and $B = V(H)$.

Then, LSIAlg calls MatchingAlg with (L, A, B) as input, and returns the answer of this call.

Denote $n = |V(G)|$, and notice that $|A| = |B| = n$. Thus, because MatchingAlg runs in time $|A|^{o(|A|)} = n^{o(n)}$, so does LSIAlg.

For the correctness of the algorithm, first suppose that (G, H, c_G, c_H, ℓ) is a Yes-instance of PROP-COL LSI. This means that there exists a bijective function $\varphi : V(G) \rightarrow V(H)$ such that (i) for every $\{u, v\} \in E(G)$, $\{\varphi(u), \varphi(v)\} \in E(H)$, and (ii) for every $u \in V(G)$, $\varphi(u) \in L(u)$. Having φ at hand, we will show that (L, A, B) is a Yes-instance, which will imply that the call to MatchingAlg with (L, A, B) as input returns Yes, and hence LSIAlg returns Yes.

Based on φ , we define a subset $M \subseteq E(L)$ as follows: $M = \{\{u, \varphi(u)\} : u \in A\}$. Notice that the containment of M in $E(L)$ follows from the definition of $E(L)$ and Condition (ii) above. Moreover, by the definition of A, B and because φ is bijective, it further follows that M is a perfect matching in L such that every edge in M has one endpoint in A and the other in B . Thus, to conclude that (L, A, B) is a Yes-instance, it remains to argue that L/M is a clique. To this end, we consider two arbitrary vertices x and y of L/M , and prove that they are adjacent in L/M . Necessarily, x is a vertex that replaced two vertices $u \in A$ and $u' \in B$ such that $\{u, u'\} \in M$, and y is a vertex that replaced two vertices $v \in A \setminus \{u\}$ and $v' \in B \setminus \{u'\}$ such that $\{v, v'\} \in M$. By the definition of contraction, to show that x and y are adjacent in L/M , it suffices to show that u and v are adjacent in L or u' and v' are adjacent in L (or both). To this end, suppose that u and v are not adjacent in L , else we are done. By the definition of $E(L)$, this means that $\{u, v\} \notin E(\overline{G})$ and hence $\{u, v\} \in E(G)$. By Condition (i) above, we derive that $\{\varphi(u), \varphi(v)\} \in E(H)$. By the definition of M , we know that $u' = \varphi(u)$ and $v' = \varphi(v)$, therefore $\{u', v'\} \in E(H)$. In turn, by the definition of $E(L)$, we get that $\{u', v'\} \in E(L)$. Thus, the proof of the forward direction is complete.

Now, suppose that LSAAlg returns Yes, which means that the call to MatchingAlg with (L, A, B) returns Yes. Thus, (L, A, B) is a Yes-instance, which means that there exists a perfect matching M in G such that every edge in M has one endpoint in A and the other in B , and G/M is a clique. We define a function $\varphi : A \rightarrow B$ as follows. For every $u \in V(G)$, let $\varphi(u) = v$ where v is the unique vertex in B such that $\{u, v\} \in M$; the existence and uniqueness of v follows from the supposition that M is a perfect matching such that every edge in M has one endpoint in A and the other in B . Furthermore, by the definition of A, B and the edges in $E(L)$ with one endpoint in A and the other in B , it directly follows that φ is a bijective mapping between $V(G)$ and $V(H)$ such that for every $u \in V(G)$, it holds that $\varphi(u) \in L(u)$. Thus, it remains to argue that for every edge $\{u, v\} \in E(G)$, it holds that $\{\varphi(u), \varphi(v)\} \in E(H)$. To this end, consider some arbitrary edge $\{u, v\} \in E(G)$, and denote $u' = \varphi(u)$ and $v' = \varphi(v)$. Because L/M is a clique and M is a matching that, by the definition of φ , necessarily contains both $\{u, u'\}$ and $\{v, v'\}$, we derive that at least one of the following four conditions must be satisfied: (i) $\{u, v\} \in E(L)$; (ii) $\{u', v'\} \in E(L)$; (iii) $\{u, v'\} \in E(L)$; (iv) $\{v, u'\} \in E(L)$. Because $\{u, v\} \in E(G)$, we have that $\{u, v\} \notin E(\overline{G})$ and therefore $\{u, v\} \notin E(L)$. Thus, we are left with Conditions (ii), (iii), and (iv). Now, we will crucially rely on the proper colorings of G and H to rule out the satisfaction of Conditions (iii) and (iv). \square

CLAIM 1. *For any two edges $\{x, x'\}, \{y, y'\} \in E(L)$ such that $\{x, y\} \in E(G)$ and $x', y' \in V(H)$, it holds that neither $\{x, y'\}$ nor $\{y, x'\}$ belongs to $E(L)$.*

PROOF OF CLAIM 1. Because c_G is a proper coloring of G and $\{x, y\} \in E(G)$, it holds that $c_G(x) \neq c_G(y)$. Because $\{x, x'\}, \{y, y'\} \in E(L)$, $x, y \in V(G)$, and $x', y' \in V(H)$, and by the definition of $E(L)$, it holds that $x' \in L(x)$ and $y' \in L(y)$, and therefore $c_G(x) = c_H(x')$ and $c_G(y) = c_H(y')$. Thus, $c_G(x) \neq c_H(y')$ and $c_G(y) \neq c_H(x')$, implying that $y' \notin L(x)$ and $x' \notin L(y)$. In turn, by the definition of $E(L)$, this means that neither $\{x, y'\}$ nor $\{y, x'\}$ belongs to $E(L)$. This completes the proof of the claim. \diamond

We now return to the proof of the lemma. By Claim 1, we are only left with Condition (ii), that is, $\{u', v'\} \in E(L)$. However, by the definition of $E(L)$, this means that $\{u', v'\} \in E(H)$. As argued earlier, this completes the proof of the reverse direction.

5 LOWER BOUNDS: CLIQUE CONTRACTION AND HADWIGER NUMBER

In this section, we prove a lower bound for CLIQUE CONTRACTION and consequently for HADWIGER NUMBER, defined as follows.




Fig. 4. An instance of NOISY STRUCTURED CLIQUE CONTRACTION where dashed lines represent non-edges.

CLIQUE CONTRACTION

Input: A graph G and $t \in \mathbb{N}$.

Question: Is there a subset $F \subseteq E(G)$ of size at most t such that G/F is a clique?

HADWIGER NUMBER

Input: A graph G and $h \in \mathbb{N}$.

Question: Is the Hadwiger number of G at least as large as h ?

Our objective is to prove the following statement, where the analogous statement for HADWIGER NUMBER (called Theorem 1.1 in the Introduction) will follow as a corollary.

THEOREM 5.1. *Unless the ETH is false, there does not exist an algorithm that solves CLIQUE CONTRACTION in time $n^{o(n)}$ where $n = |V(G)|$.*

To make our approach adaptable to extract analogous statements for other contraction problems, we will first define a new problem called NOISY STRUCTURED CLIQUE CONTRACTION (which will arise in Section 6) along with a special case of it that is also a special case of CLIQUE CONTRACTION. Then, we will prove a crucial property of instances of NOISY STRUCTURED CLIQUE CONTRACTION, and afterwards we will use this property to prove Theorem 5.1 and its corollary. The definition of the new problem is as follows (see Figure 4).

NOISY STRUCTURED CLIQUE CONTRACTION

Input: A graph G on at least $6n$ vertices for some $n \in \mathbb{N}$, and a partition (A, B, C, D, N) of $V(G)$ such that $|A| = |B| = n$, $|C| = |D| = 2n$, no vertex in A is adjacent to any vertex in D , and no vertex in B is adjacent to any vertex in C .

Question: Does there exist a subset $F \subseteq E(G)$ of size at most n such that $G[A \cup B \cup C \cup D \cup X]/F$ is a clique,^a where $X = \{u \in N : \text{there exists a vertex } v \in A \cup B \cup C \cup D \text{ such that } u \text{ and } v \text{ belong to the same connected component of } G[F]\}$?

^aNote that F might contain edges outside $G[A \cup B \cup C \cup D \cup X]$. Then, we slightly abuse notation so that $G[A \cup B \cup C \cup D \cup X]/F$ refers to $G[A \cup B \cup C \cup D \cup X]/(F \cap E(G[A \cup B \cup C \cup D \cup X]))$.

Intuitively, the vertex set X consists of the noise (represented by N) that “interacts” with non-noise (represented by $V(G) \setminus N$) through contracted edges (in F), that is, the vertices in N that lie together with at least one vertex in $V(G) \setminus N$ in a component that will be contracted and thereby

replaced by a single vertex. We refer to the special case of NOISY STRUCTURED CLIQUE CONTRACTION where $N = \emptyset$ as STRUCTURED CLIQUE CONTRACTION. Note that STRUCTURED CLIQUE CONTRACTION is also a special case of CLIQUE CONTRACTION.

Solutions to instances of NOISY STRUCTURED CLIQUE CONTRACTION exhibit the following property, which will be crucial in the proof of Theorem 5.1 as well as results in Section 6.

LEMMA 5.2. *Let F be a solution to an instance (G, A, B, C, D, N, n) of NOISY STRUCTURED CLIQUE CONTRACTION. Then, F is a matching of size n in G such that each edge in F has one endpoint in A and the other in B . In particular, $X = \emptyset$.*

PROOF. We first argue that every vertex in $A \cup B$ is incident to at least one edge in F . Targeting a contradiction, suppose that there exists a vertex $u \in A \cup B$ that is not incident to any edge in F . Because $|A \cup B \cup C \cup D| = 6n$, $|F| \leq n$ and $G[A \cup B \cup C \cup D \cup X]/F$ is a clique (where the last two properties follow from the supposition that F is a solution), it holds that $G[A \cup B \cup C \cup D \cup X]/F$ is a clique on at least $5n + |X|$ vertices. Hence, the degree of every vertex in $G[A \cup B \cup C \cup D \cup X]/F$, and in particular of u , should be $5n - 1 + |X|$ in $G[A \cup B \cup C \cup D \cup X]/F$. However, because no vertex in A is adjacent to any vertex in D and no vertex in B is adjacent to any vertex in C , the degree of any vertex in $A \cup B$, and in particular of u , is at most $|A \cup B| - 1 + |C \cup D|/2 + |X| = 4n - 1 + |X|$ in $G[A \cup B \cup C \cup D \cup X]$. Because u is not incident to any edge in F , its degree in $G[A \cup B \cup C \cup D \cup X]/F$ is at most its degree in $G[A \cup B \cup C \cup D \cup X]$. This is a contradiction, thus we get that indeed every vertex in $A \cup B$ is incident to at least one edge in F . From this, because $|F| \leq n$ and $|A \cup B| = 2n$, we derive that F is a perfect matching in $G[A \cup B]$.

It remains to argue that every edge in F has one endpoint in A and the other in B . Targeting a contradiction, suppose that this is false. Because F is a perfect matching in $G[A \cup B]$, this means that there exist two vertices $a, a' \in A$ such that $\{a, a'\} \in F$. By the definition of NOISY STRUCTURED CLIQUE CONTRACTION, neither a nor a' is adjacent to any vertex in D . Moreover, note that $D \subseteq V(G[A \cup B \cup C \cup D \cup X]/F)$. In particular, the vertex of $G[A \cup B \cup C \cup D \cup X]/F$ yielded by the contraction of $\{a, a'\}$ is not adjacent to any vertex of D in $G[A \cup B \cup C \cup D \cup X]/F$. However, this is a contradiction because $G[A \cup B \cup C \cup D \cup X]/F$ is a clique. \square

We now prove a lower bound for STRUCTURED CLIQUE CONTRACTION. Because it is a special case of CLIQUE CONTRACTION, this will directly yield the correctness of Theorem 5.1.

LEMMA 5.3. *Unless the ETH is false, there does not exist an algorithm that solves STRUCTURED CLIQUE CONTRACTION in time $n^{o(n)}$ where $n = |V(G)|$.*

PROOF. Targeting a contradiction, suppose that there exists an algorithm, denoted by CliConAlg, that solves STRUCTURED CLIQUE CONTRACTION in time $n^{o(n)}$ where n is the number of vertices in the input graph. We will show that this implies the existence of an algorithm, denoted by MatchingAlg, that solves CROSS MATCHING in time $n^{o(n)}$ where n is the size of the set A in the input, thereby contradicting Lemma 4.1 and hence completing the proof.

We define the execution of MatchingAlg as follows. Given an instance (G, A, B) of CROSS MATCHING, MatchingAlg constructs an instance (H, A, B, C, D, n) of STRUCTURED CLIQUE CONTRACTION as follows (see Figure 5):

- Let $n = |A|$, and K be a clique on $4n$ new vertices. Let (C, D) be a partition of $V(K)$ such that $|C| = |D|$.
- $V(H) = V(G) \cup V(K)$.
- $E(H) = E(G) \cup E(K) \cup \{\{a, c\} : a \in A, c \in C\} \cup \{\{b, d\} : b \in B, d \in D\}$.

Then, MatchingAlg calls CliConAlg with (H, A, B, C, D, n) as input, and returns the answer.




Fig. 5. The construction of an instance of STRUCTURED CLIQUE CONTRACTION in the proof of Lemma 5.3 where dashed lines represent non-edges.

First, note that by construction, $|V(H)| = 6n$. Thus, because CliConAlg runs in time $|V(H)|^{o(|V(H)|)} \leq n^{o(n)}$, it follows that MatchingAlg runs in time $n^{o(n)}$.

For the correctness of the algorithm, first suppose that (G, A, B) is a Yes-instance of CROSS MATCHING. This means that there exists a perfect matching M in G such that every edge in M has one endpoint in A and the other in B , and G/M is a clique. By the definition of $E(H)$, $M \subseteq E(H)$. We will show that H/M is a clique. As $|M| = n$, this will mean that (H, A, B, C, D, n) is a Yes-instance of STRUCTURED CLIQUE CONTRACTION, which will mean, in turn, that the call to CliConAlg with (H, A, B, C, D, n) as input returns Yes, and hence MatchingAlg returns Yes.

Note that $V(H/M) = V(K) \cup V(G/M)$. To show that H/M is a clique, we consider two arbitrary vertices $u, v \in V(H/M)$, and show that they are adjacent in H/M . If $u, v \in V(K)$, then because K is a clique, it is clear that $\{u, v\} \in E(H/M)$. Moreover, if $u, v \in G/M$, then because G/M is a clique, it is clear that $\{u, v\} \in E(H/M)$. Thus, one of the vertices u and v belongs to $V(G/M)$ and the other belongs to $V(K)$. We suppose w.l.o.g. that $u \notin V(K)$. Because M is a perfect matching in G such that every edge in M has one endpoint in A and the other in B , it follows that u resulted from the contraction of the edge between some $a \in A$ and some $b \in B$. If $v \in C$, then $\{a, v\} \in E(H)$, and otherwise $v \in D$ and so $\{b, v\} \in E(H)$. Thus, by the definition of contraction, we conclude that $\{u, v\} \in E(H/M)$. This completes the proof of the forward direction.

Now, suppose that MatchingAlg returns Yes, which means that the call to CliConAlg with (H, A, B, C, D, n) returns Yes. Thus, (H, A, B, C, D, n) is a Yes-instance, which means that there exists a subset $F \subseteq E(H)$ of size at most n such that H/F is a clique. We will show that F is a perfect matching in G such that every edge in F has one endpoint in A and the other in B . Because H/F is a clique, this will imply that G/F is a clique and thus that (G, A, B) is a Yes-instance of CROSS MATCHING. To achieve this, notice that by Lemma 5.2, F is a matching of size n in H such that each edge in F has one endpoint in A and the other in B . Because $G = H[A \cup B]$, we have that F is a perfect matching in G . Thus, the proof of the reverse direction is complete. \square

COROLLARY 5.4. *Unless the ETH is false, there does not exist an algorithm that solves HADWIGER NUMBER in time $n^{o(n)}$ where $n = |V(G)|$.*

PROOF. Targeting a contradiction, suppose that there exists an algorithm, denoted by HadwigerAlg, that solves HADWIGER NUMBER in time $n^{o(n)}$ where n is the number of vertices in the input graph. We will show that this implies the existence of an algorithm, denoted by CliConAlg, that solves CLIQUE CONTRACTION in time $n^{o(n)}$ where n is the number of vertices in the input graph, thereby contradicting Theorem 5.1 and hence completing the proof.

We define the execution of CliConAlg as follows. Given an instance (G, t) of CLIQUE CONTRACTION, if G is not connected, then CliConAlg returns No, and otherwise it returns Yes if and only if




Fig. 6. A two-cliques graph (see Definition 6.2).

HadwigerAlg returns Yes when called with $(G, |V(G)| - t)$ as input. Because the call to HadwigerAlg with input $(G, |V(G)| - t)$ runs in time $n^{o(n)}$ where $n = |V(G)|$, we have that CliConAlg runs in time $n^{o(n)}$ as well.

For the correctness of the algorithm, first observe that if G is not connected, then no sequence of edge contractions can yield a clique, and hence it is correct to return No. Thus, now assume that G is connected. First, suppose that (G, t) is a Yes-instance of CLIQUE CONTRACTION. This means that there exists a sequence of at most t edge contractions that transforms G into a clique. In particular, this clique must have at least $|V(G)| - t$ vertices, and therefore the Hadwiger number of G is at least as large as $|V(G)| - t$. By the correctness of HadwigerAlg, its call with $(G, |V(G)| - t)$ returns Yes, and therefore CliConAlg returns Yes.

Now, suppose that CliConAlg returns Yes, which means that the call to HadwigerAlg with $(G, |V(G)| - t)$ returns Yes. By the correctness of HadwigerAlg, the clique K_h for $h = |V(G)| - t$ is a minor of G . This means that there is a sequence of vertex deletions, edge deletions, and edge contractions that transforms G into K_h . In particular, this sequence can contain at most t vertex deletions and edge contractions in total. Furthermore, by replacing each vertex deletion for a vertex v by an edge contraction for some edge e incident to v (which exists because G is connected) and dropping all edge deletions, we obtain another sequence that transforms G into K_h . Because this sequence contains only edge contractions, and at most t of them, we conclude that (G, t) is a Yes-instance of CLIQUE CONTRACTION. \square

6 LOWER BOUNDS FOR CONTRACTION TO GRAPH CLASSES PROBLEMS

In this section, we prove lower bounds for several cases of the \mathcal{F} -CONTRACTION problem, defined as follows. Here, \mathcal{F} is a (possibly infinite) family of graphs.

\mathcal{F} -CONTRACTION
Input: A graph G and $t \in \mathbb{N}$.
Question: Does there exist a subset $F \subseteq E(G)$ of size at most t such that $G/F \in \mathcal{F}$?

Notice that CLIQUE CONTRACTION is the case of \mathcal{F} -CONTRACTION where \mathcal{F} is the family of cliques. In this section, we consider the cases of \mathcal{F} -CONTRACTION where \mathcal{F} is the family of chordal graphs, interval graphs, proper interval graphs, threshold graphs, trivially perfect graphs, split graphs, complete split graphs, and perfect graphs, also called CHORDAL CONTRACTION, INTERVAL CONTRACTION, PROPER INTERVAL CONTRACTION, THRESHOLD CONTRACTION, TRIVIAALLY PERFECT CONTRACTION, SPLIT CONTRACTION, COMPLETE SPLIT CONTRACTION, and PERFECT CONTRACTION, respectively. Before we define these classes formally, it will be more enlightening to first define only the class of chordal graphs as well as somewhat artificial classes of graphs that will help us prove lower bounds for many of the classes above in a unified manner.

Definition 6.1 (Chordal Graphs). A graph is *chordal* if it does not contain C_ℓ for all $\ell \geq 4$ as an induced subgraph.

Our first class of graphs is defined as follows (see Figure 6).

Definition 6.2 (Two-Cliques Graphs). A two-cliques graph is a graph G such that there exist $A, B \subseteq V(G)$ such that $A \cup B = V(G)$, $G[A]$ and $G[B]$ are cliques, and there do not exist vertices $a \in A \setminus B$ and $b \in B \setminus A$ such that $\{a, b\} \in E(G)$. The two-cliques class is the class of all two-cliques graphs.

It should be clear that the two-cliques class is a subclass of the class of chordal graphs. Keeping this in mind, we further define a family of classes of graphs that includes the two-cliques class.

Definition 6.3 (Non-Trivial Chordal Class). We say that a class of graphs \mathcal{F} is non-trivial chordal if it is a subclass of the class of chordal graphs, and a superclass of the two-cliques class.

Clearly, the class of cliques is not a non-trivial chordal class, and the class of chordal graphs is a non-trivial chordal class. The rest of this section is divided as follows. First, in Section 6.1, we prove a lower bound for any non-trivial chordal class. Then, in Section 6.2, we prove a lower bound for some graph classes that are not non-trivial chordal.

6.1 Non-Trivial Chordal Graph Classes

The main objective of this subsection is to prove the following theorem. Afterwards, we will derive lower bounds for several known graph classes as corollaries.

THEOREM 6.4. *Let \mathcal{F} be any non-trivial chordal graph class. Unless the ETH is false, there does not exist an algorithm that solves \mathcal{F} -CONTRACTION in time $n^{o(n)}$ where $n = |V(G)|$.*

For the proof of this theorem, the following well-known property of chordal graphs will come in handy. This property is a direct consequence of the alternative characterization of the class of chordal graphs as the class of graphs that admit clique-tree decompositions (see [11]).

PROPOSITION 6.5. *Let G be a chordal graph, and let u and v be two non-adjacent vertices in G . Then, $G[N(u) \cap N(v)]$ is a clique.*

We are now ready to prove Theorem 6.4.

PROOF OF THEOREM 6.4. Targeting a contradiction, suppose that there exists an algorithm, denoted by `NonTrivChordAlg`, that solves \mathcal{F} -CONTRACTION in time $n^{o(n)}$ where n is the number of vertices in the input graph. We will show that this implies the existence of an algorithm, denoted by `CliConAlg`, that solves STRUCTURED CLIQUE CONTRACTION in time $n^{o(n)}$ where n is the number of vertices in the input graph, thereby contradicting Lemma 5.3 and hence completing the proof.

We define the execution of `CliConAlg` as follows. Given an instance (G, A, B, C, D, n) of STRUCTURED CLIQUE CONTRACTION, `CliConAlg` constructs an instance (H, n) of \mathcal{F} -CONTRACTION as follows (see Figure 7):

- Let $n = |A|$. Moreover, let K and K' be two cliques, each on $2n$ new vertices.
- $V(H) = V(G) \cup V(K) \cup V(K')$.
- $E(H) = E(G) \cup E(K) \cup E(K') \cup \{\{u, v\} : u \in V(G), v \in V(K) \cup V(K')\}$.

Then, `CliConAlg` calls `NonTrivChordAlg` with (H, n) as input, and returns the answer of this call.

First, note that by construction, $|V(H)| = 10n$. Thus, because `NonTrivChordAlg` runs in time $|V(H)|^{o(|V(H)|)} \leq n^{o(n)}$, it follows that `CliConAlg` runs in time $n^{o(n)}$.

For the correctness of the algorithm, first suppose that (G, A, B, C, D, n) is a Yes-instance of STRUCTURED CLIQUE CONTRACTION. This means that there exists a subset $F \subseteq E(G)$ of size at most n such that G/F is a clique. By the definition of H , we directly derive that H/F is a two-cliques graph, and therefore it belongs to \mathcal{F} . Thus, (H, n) is a Yes-instance of \mathcal{F} -CONTRACTION, which




Fig. 7. The construction of an instance of \mathcal{F} -CONTRACTION in the proof of Theorem 6.4 where dashed lines represent non-edges.

means that the call to `NonTrivChordAlg` with (H, n) as input returns Yes, and hence `CliConAlg` returns Yes.

Now, suppose that `CliConAlg` returns Yes, which means that the call to `NonTrivChordAlg` with (H, n) returns Yes. Thus, (H, n) is a Yes-instance of \mathcal{F} -CONTRACTION, which means that there exists a subset $F \subseteq E(H)$ of size at most n such that $H/F \in \mathcal{F}$. In particular, H/F is a chordal graph. Based on Proposition 6.5, we will first show that $H[A \cup B \cup C \cup D \cup X]/F$ is a clique, where $X = \{u \in V(K) \cup V(K') : \text{there exists a vertex } v \in A \cup B \cup C \cup D \text{ such that } u \text{ and } v \text{ belong to the same connected component of } H[F]\}$.

Targeting a contradiction, suppose that $H[A \cup B \cup C \cup D \cup X]/F$ is not a clique, and therefore there exist two non-adjacent vertices u and v in this graph. By the definition of X , $H[A \cup B \cup C \cup D \cup X]/F$ is equal to the subgraph of H/F induced by the set of vertices derived from connected components that contain at least one vertex from $A \cup B \cup C \cup D$. In particular, u and v are also non-adjacent vertices in H/F . By Proposition 6.5, this implies that $(H/F)[N_{H/F}(u) \cap N_{H/F}(v)]$ is a clique. Let \mathcal{C}_1 (\mathcal{C}_2 , respectively) be the set of connected components of $H[F]$ that contain at least one vertex from $V(K_1)$ ($V(K_2)$, respectively). Because $|F| \leq n$ and $|V(K_1)| = |V(K_2)| = 2n$, there exists at least one component $C_1 \in \mathcal{C}_1$ ($C_2 \in \mathcal{C}_2$, respectively) that does not contain any vertex from $A \cup B \cup C \cup D$. Let c_1 and c_2 be the vertices of H/F yielded by the replacement of C_1 and C_2 , respectively. As all vertices in $V(K_1) \cup V(K_2)$ are adjacent to all vertices in $A \cup B \cup C \cup D$, we have that $c_1, c_2 \in N_{H/F}(u) \cap N_{H/F}(v)$. However, there do not exist a vertex in $V(K_1)$ and a vertex in $V(K_2)$ that are adjacent in H , and for every vertex in $V(K_1) \cup V(K_2)$, its neighborhood outside this set is contained in $A \cup B \cup C \cup D$. Thus, c_1 and c_2 must be non-adjacent in H/F . However, this is a contradiction to the argument that $(H/F)[N_{H/F}(u) \cap N_{H/F}(v)]$ is a clique. From this, we derive that $H[A \cup B \cup C \cup D \cup X]/F$ is indeed a clique.

Now, notice that (H, A, B, C, D, N, n) where $N = V(K_1) \cup V(K_2)$ is an instance of NOISY STRUCTURED CLIQUE CONTRACTION. Furthermore, since $|F| \leq n$ and we have already shown that $H[A \cup B \cup C \cup D \cup X]/F$ is a clique, we have that F is a solution to this instance. Therefore, by Lemma 5.2, F is a matching of size n in H such that each edge in F has one endpoint in A and the other in B . In particular, $F \subseteq E(G)$ and hence $X = \emptyset$. Because $G = H[A \cup B \cup C \cup D]$, we thus derive that G/F is a clique. Thus, we conclude that (G, A, B, C, D, n) is a Yes-instance of STRUCTURED CLIQUE CONTRACTION. This completes the proof of the reverse direction. \square

Now, we give definitions for several classes of graphs for which lower bounds will follow from Theorem 6.5. First, a graph is an *interval graph* if there exists a set of intervals on the real line such

that the vertices of the graph are in bijection with these intervals, and there exists edge between two vertices if and only if their intervals intersect. A graph is a *proper interval graph* if, in the former definition, we also add the constraint that all intervals must have the same length. A graph is a *threshold graph* if it can be constructed from a one-vertex graph by repeated applications of the following two operations: addition of a single isolated vertex to the graph; addition of a single vertex that is connected to all other vertices. A graph is *trivially perfect* if in each of its induced subgraphs, the maximum size of an independent set equals the number of maximal cliques.

It is well-known that every graph that is a (proper) interval graph, or a threshold graph, or a trivially perfect graph, is also a chordal graph (see [11]). Moreover, it is immediate to verify that the two-cliques class is a subclass of the classes of (proper) interval graphs, threshold graphs, and trivially perfect graphs. Thus, these classes are non-trivial chordal graphs classes, and therefore Theorem 6.4 directly implies lower bounds for them as stated below.

COROLLARY 6.6. *Unless the ETH is false, none of the following problems admits an algorithm that solves it in time $n^{o(n)}$ where $n = |V(G)|$: CHORDAL CONTRACTION, INTERVAL CONTRACTION, PROPER INTERVAL CONTRACTION, THRESHOLD CONTRACTION, and TRIVIALY PERFECT CONTRACTION.*

6.2 Other Graph Classes

In Section 5, we have already proved a lower bound for a class of graphs that is not non-trivial chordal, namely, the class of cliques. In this section, we show that our approach can yield lower bounds for other classes of graphs that are not non-trivially chordal. For illustrative purposes, we consider the classes of SPLIT GRAPHS, COMPLETE SPLIT GRAPHS, and PERFECT GRAPHS.

A graph G is a *split graph* if there exists a partition (I, K) of $V(G)$ such that $G[I]$ is edgeless and $G[K]$ is a clique. In case $\{\{i, k\} : i \in I, k \in K\} \subseteq E(G)$, we further say that G is a *complete split graph*. Notice that the two-cliques class is not a subclass of the class of split graphs, and hence the class of (complete) split graphs is not non-trivially chordal.

For the class of (complete) split graphs, we prove the following statement.

THEOREM 6.7. *Unless the ETH is false, there does not exist an algorithm that solves SPLIT CONTRACTION (or COMPLETE SPLIT CONTRACTION) in time $n^{o(n)}$ where $n = |V(G)|$.*

PROOF. Targeting a contradiction, suppose that there exists an algorithm, denoted by SplitAlg, that solves SPLIT CONTRACTION (or COMPLETE SPLIT CONTRACTION) in time $n^{o(n)}$ where n is the number of vertices in the input graph. We will show that this implies the existence of an algorithm, denoted by CliConAlg, that solves STRUCTURED CLIQUE CONTRACTION in time $n^{o(n)}$ where n is the number of vertices in the input graph, thereby contradicting Lemma 5.3 and hence completing the proof.

We define the execution of CliConAlg as follows. Given an instance (G, A, B, C, D, n) of STRUCTURED CLIQUE CONTRACTION, CliConAlg constructs an instance (H, n) of SPLIT CONTRACTION (or COMPLETE SPLIT CONTRACTION) as follows (see Figure 8):

$$- V(H) = V(G) \cup S \text{ where } S \text{ is a set of } n + 2 \text{ new vertices.}$$

$$- E(H) = E(G) \cup \{\{u, v\} : u \in V(G), v \in S\}.$$

Then, CliConAlg calls SplitAlg with (H, n) as input, and returns the answer of this call.

First, note that by construction, $|V(H)| = 7n + 2$. Thus, because SplitAlg runs in time $|V(H)|^{o(|V(H)|)} \leq n^{o(n)}$, it follows that CliConAlg runs in time $n^{o(n)}$.

For the correctness of the algorithm, first suppose that (G, A, B, C, D, n) is a Yes-instance of STRUCTURED CLIQUE CONTRACTION. This means that there exists a subset $F \subseteq E(G)$ of size at most n such that G/F is a clique. By the definition of H , we derive that H/F is a complete split

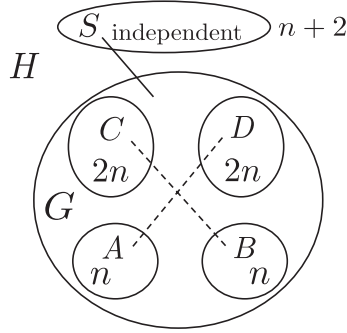


Fig. 8. The construction of an instance of SPLIT CONTRACTION in the proof of Theorem 6.7 where dashed lines represent non-edges.

graph: $(S, V(G/F))$ is a partition of $V(H/F)$ where S induces an independent set, $V(G/F)$ induces a clique, and every vertex in S is adjacent to every vertex in $V(G/F)$. Thus, (H, n) is a Yes-instance of COMPLETE SPLIT CONTRACTION (as well as of SPLIT CONTRACTION), which means that the call to SplitAlg with (H, n) returns Yes, and hence CliConAlg returns Yes.

Now, suppose that CliConAlg returns Yes, which means that the call to SplitAlg with (H, n) returns Yes. Thus, (H, n) is a Yes-instance of SPLIT CONTRACTION (even if SplitAlg solves COMPLETE SPLIT CONTRACTION), which means that there exists a subset $F \subseteq E(H)$ of size at most n such that H/F is a split graph. Let (I, K) be a partition of $V(H/F)$ into an independent set and a set of vertices that induce a clique. Because $|S| = n + 2$ and $H[S]$ is an independent set, there exist at least two vertices $s_1, s_2 \in S$ that are not incident to any edge in F . As these vertices are not adjacent to one another in H , and because they are adjacent to all vertices in $V(G)$ (and hence to all vertices in $V(H/F) \setminus S$), it follows that $s_1, s_2 \in I$ and $V(H/F) \setminus S \subseteq K$. In particular, $(H/F)[V(H/F) \setminus S]$ is a clique. Let $X = \{u \in S : \text{there exists a vertex } v \in V(G) \text{ such that } u \text{ and } v \text{ belong to the same connected component of } G[F]\}$. Then, we have that $H[V(G) \cup X]/F$ is a clique.

Now, notice that (H, A, B, C, D, S, n) is an instance of NOISY STRUCTURED CLIQUE CONTRACTION. Furthermore, since $|F| \leq n$ and we have already shown that $H[A \cup B \cup C \cup D \cup X]/F$ is a clique, we have that F is a solution to this instance. Therefore, by Lemma 5.2, F is a matching of size n in H such that each edge in F has one endpoint in A and the other in B . In particular, $F \subseteq E(G)$ and hence $X = \emptyset$. Because $G = H[A \cup B \cup C \cup D]$, we thus derive that G/F is a clique. Thus, we conclude that (G, A, B, C, D, n) is a Yes-instance of STRUCTURED CLIQUE CONTRACTION. This completes the proof of the reverse direction. \square

A graph G is a *perfect graph* if the chromatic number of every induced subgraph of G equals the size of the largest clique of that subgraph. Here, the chromatic number of a graph is the minimum number of colors required to color its vertices so that every pair of adjacent vertices are assigned different colors. For the class of perfect graphs, we prove the following statement.

THEOREM 6.8. *Unless the ETH is false, there does not exist an algorithm that solves PERFECT CONTRACTION in time $n^{o(n)}$ where $n = |V(G)|$.*

PROOF. Targeting a contradiction, suppose that there exists an algorithm, denoted by PerfectAlg, that solves PERFECTCONTRACTION in time $n^{o(n)}$ where n is the number of vertices in the input graph. We will show that this implies the existence of an algorithm, denoted by CliConAlg, that solves STRUCTURED CLIQUE CONTRACTION in time $n^{o(n)}$ where n is the number of vertices in the input graph, thereby contradicting Lemma 5.3 and hence completing the proof.




Fig. 9. The construction of an instance of PERFECT CONTRACTION in the proof of Theorem 6.8 where dashed lines represent non-edges.

We define the execution of CliConAlg as follows. Given an instance (G, A, B, C, D, n) of STRUCTURED CLIQUE CONTRACTION, CliConAlg constructs an instance (H, n) of PERFECT CONTRACTION as follows (see Figure 9):

- Let $K = \{u' : u \in V(G)\}$ where each element u' is a new vertex referred to as the *tagged copy* of u . Additionally, let I be a set of $n + 1$ new vertices.
- $V(H) = V(G) \cup K \cup I$.
- $E(H) = E(G) \cup \{\{u, u'\} : u \in V(G)\} \cup \{\{u', v'\} : u', v' \in K\} \cup \{\{u, i\} : u \in V(G), i \in I\}$.

Then, CliConAlg calls PerfectAlg with (H, n) as input, and returns the answer of this call.

First, note that by construction, $|V(H)| \leq 13n + 1$. Thus, because PerfectAlg runs in time $|V(H)|^{o(|V(H)|)} \leq n^{o(n)}$, it follows that CliConAlg runs in time $n^{o(n)}$.

In what follows, given a subset $U \subseteq V(G)$, we denote $U' = \{u' \in K : u \in U\}$. For the correctness of the algorithm, first suppose that (G, A, B, C, D, n) is a Yes-instance of STRUCTURED CLIQUE CONTRACTION. This means that there exists a subset $F \subseteq E(G)$ of size at most n such that G/F is a clique. Now, we will show that H/F is a perfect graph. To this end, consider some induced subgraph S of H/F .

In case the maximum size of a clique in S is 2, then S can contain at most two vertices from K , and either (i) at most one vertex from outside $K \cup I$, or (ii) two vertices from outside $K \cup I$ and no vertex from I (because $H[V(G)]/F$ is a clique and all of its vertices are adjacent to all vertices in I). In case (ii), S contains at most four vertices, and it is straightforward to verify that any graph on at most four vertices is perfect. In case (i), we color the vertex in $V(S) \setminus (K \cup I)$ (if it exists) in one color, and all vertices in $V(S) \cap I$ in the second color; among the (at most) two vertices in $V(S) \cap K$, at most one is adjacent to the vertex in $V(S) \setminus (K \cup I)$, and that one is colored by the same color as the vertices in $V(S) \cap I$ (while the other vertex is colored by the other color).

Thus, in what follows, suppose that the maximum size of a clique in S is at least 3. Observe that every clique of maximum size in S is of one of two types: it consists either (i) of no vertex from K (in which case it can contain at most one vertex from I), or (ii) only of vertices from K . First, suppose that there exists a clique \widehat{C} of maximum size in S that consists of no vertex from K . In this case, $|V(S) \cap K| \leq |\widehat{C}|$. Then, color each vertex in \widehat{C} by a distinct color. Clearly, because $|V(S) \cap K| \leq |\widehat{C}|$, all vertices in $V(S) \cap K$ can be colored using the same set of colors so that a vertex and its tagged copy are assigned distinct colors. Further, all the vertices left uncolored in $V(S)$ belong to I , and if there exists at least one such vertex left uncolored, then $|V(\widehat{C}) \cap I| = 1$, and all vertices

in I can receive the same color as the vertex from $V(\widehat{C}) \cap I$. Second, suppose that all cliques of maximum size in S consist only of vertices in K —then, S has exactly one clique of maximum size, which is $S[V(S) \cap K]$, and which we denote by \widehat{C} . In this case, $|V(S) \setminus (K \cup I)| \leq |V(\widehat{C})| - 1$. Then, color each vertex in \widehat{C} by a distinct color. Clearly, because $|V(S) \setminus (K \cup I)| \leq |V(\widehat{C})| - 1$, all vertices in $V(S) \setminus (K \cup I)$ can be colored using the same set of colors minus one color, so that a vertex and its tagged copy are assigned distinct colors. So, all vertices in I can be colored using the same color that is some color used by a vertex in \widehat{C} but by no vertex in $V(S) \setminus (K \cup I)$. In either case, we obtain that the chromatic number of S equals its maximum clique size. Thus, (H, n) is a Yes-instance of PERFECT CONTRACTION, which means that the call to PerfectAlg with (H, n) returns Yes, and hence CliConAlg returns Yes.

Now, suppose that CliConAlg returns Yes, which means that the call to PerfectAlg with (H, n) returns Yes. Thus, (H, n) is a Yes-instance of PERFECT CONTRACTION, which means that there exists a subset $F \subseteq E(H)$ of size at most n such that H/F is a perfect graph. We first argue that there does not exist a vertex $a \in A \cup B$ such that neither a nor a' is incident to at least one edge in F . Targeting a contradiction, suppose that there exists $a \in A \cup B$ such that neither a nor a' is incident to at least one edge in F . Assume that $a \in A$ as the other case is symmetric. Because $|F| \leq n$ and $|D| = 2n$, there either exists a vertex $d \in D$ such that neither d nor d' is incident to at least one edge in F , or F is a perfect matching in either $G[D]$ or $G[D']$, where in the latter case we let d denote some arbitrarily chosen vertex from D . Additionally, since I is an independent set of size $n + 1$, there exists a vertex $i \in I$ that is not incident to any edge in F . Now, consider the cycle $i - a - a' - d' - d - i$ (on five vertices) in H . This cycle is an induced cycle in H , because no vertex in A is adjacent to any vertex in D , and by the construction of H , i is not adjacent to a' and d' , a is not adjacent to d' , and a' is not adjacent to d . Furthermore, as i , a , and a' are not incident to any edge in F , and if any of d and d' is incident to an edge in F , then F is a perfect matching in either $G[D]$ or $G[D']$, we obtain that $i - a - a' - \widehat{d} - \widehat{d}' - i$ is an induced cycle (on five vertices) in H/F , where \widehat{d} and \widehat{d}' are the vertices yielded by the replacement of the connected components of $H[F]$ that contain d and d' , respectively, if such components exist (otherwise, $\widehat{d} = d$ and $\widehat{d}' = d'$). However, an induced cycle on five vertices has chromatic number 3 and maximum clique size 2, thus we derive a contradiction to the supposition that H/F is perfect.

So far, we derived that there does not exist a vertex $a \in A \cup B$ such that neither a nor a' is incident to at least one edge in F . As $|F| \leq n$ and $|A| = |A'| = |B| = |B'| = n$, this means that every edge in F has both endpoints in $A \cup A' \cup B \cup B'$ and that for each $u \in A \cup B$, exactly one vertex among u and u' is incident to an edge in F . Now, we will show that each vertex $a \in A \cup B$ is incident to at least one edge in F . Targeting a contradiction, suppose that there exists a vertex $a \in A \cup B$ that is not incident to any edge in F . Assume that $a \in A$, as the other case is symmetric. Denote i , d , d' , \widehat{d} , and \widehat{d}' as before, and again consider the induced cycle $i - a - a' - d' - d - i$ in H . Unlike before, now a' belongs to some connected component of $H[F]$, yet we know that this connected component consists only of a' and some vertex in B' . Let \widehat{a}' be the vertex yielded by the replacement of this component. As no vertex in B' is adjacent to any vertex in $I \cup D$, we again have that $i - a - \widehat{a}' - \widehat{d}' - \widehat{d} - i$ is an induced cycle in H/F , which gives rise to a contradiction. Thus, as $|F| \leq n$ and $|A| = |B| = n$, we know that F is a perfect matching in $G[A \cup B]$.

Next, we will show that G/F is a clique. This will imply that (G, A, B, C, D, n) is a Yes-instance of STRUCTURED CLIQUE CONTRACTION and thereby complete the proof. Targeting a contradiction, suppose that G/F is not a clique, and therefore there exist two non-adjacent vertices u and v in G/F . As F is a matching in $G[A \cup B]$, we can let x and y be two vertices in $A \cup B$ that belonged to the connected components of $H[F]$ that yielded u and v , respectively. Notice that the only vertex in

$A \cup B$ adjacent to x' is x , and the analogous claim holds for y' and y . As F is a matching in $G[A \cup B]$ that does not match x and y (since otherwise u and v would not be distinct vertices), we have that neither u is adjacent to y' in H/F nor v is adjacent to x' in H/F . From this, by the construction of H and since F is a matching in $G[A \cup B]$, we immediately derive that $i - u - x' - y' - v - i$ is an induced cycle in H/F where i is some arbitrarily chosen vertex from I . However, as before, the existence of such a cycle contradicts the supposition that H/F is a perfect graph. Thus, the proof of the reverse direction is complete. \square

7 CONCLUSION

In this article, we proved that the HADWIGER NUMBER problem cannot be solved in time $n^{o(n)}$ under the ETH. In particular, this stands in contrast to the related SUBGRAPH ISOMORPHISM, GRAPH HOMOMORPHISM, and TOPOLOGICAL GRAPH MINOR problems, which can be solved in time $2^{O(n)}$ when one of the input graphs is a clique. Our technique could also be applied to rule out (under the ETH) the existence of $n^{o(n)}$ -time algorithms for a large number of \mathcal{F} -CONTRACTION problems. We believe that our technique can be applied to additional contraction problems as well. However, some non-trivial contraction problems, such as PATH CONTRACTION [1], can be solved in time $2^{O(n)}$. So, for future research, we suggest to attain better understanding in terms of as general as possible dichotomy theorems that identify which contraction problems can, and which cannot, be solved in time $2^{O(n)}$.

REFERENCES

- [1] Akanksha Agrawal, Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, and Prafullkumar Tale. 2020. Path contraction faster than 2^n . *SIAM J. Discret. Math.* 34, 2 (2020), 1302–1325.
- [2] Akanksha Agrawal, Daniel Lokshtanov, Saket Saurabh, and Meirav Zehavi. 2017. Split contraction: The untold story. In *Proceedings of the 34th Symposium on Theoretical Aspects of Computer Science (STACS 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [3] Andreas Björklund. 2014. Determinant sums for undirected hamiltonicity. *SIAM J. Comput.* 43, 1 (2014), 280–299.
- [4] Andreas Björklund, Thore Husfeldt, and Mikko Koivisto. 2009. Set partitioning via inclusion–exclusion. *SIAM J. Comput.* 39, 2 (2009), 546–563.
- [5] B. Bollobás, P. A. Catlin, and P. Erdős. 1980. Hadwiger’s conjecture is true for almost every graph. *European J. Combin.* 1, 3 (1980), 195–199. DOI: [https://doi.org/10.1016/S0195-6698\(80\)80001-1](https://doi.org/10.1016/S0195-6698(80)80001-1)
- [6] Jianer Chen, Benny Chor, Michael R. Fellows, Xiuzhen Huang, David Juedes, Iyad A. Kanj, and Ge Xia. 2005. Tight lower bounds for certain parameterized NP-hard problems. *Inf. Comput.* 201, 2 (2005), 216–231. DOI: <https://doi.org/10.1016/j.ic.2005.05.001>
- [7] Marek Cygan, Fedor V. Fomin, Alexander Golovnev, Alexander S. Kulikov, Ivan Mihajlin, Jakub Pachocki, and Arkadiusz Socala. 2017. Tight lower bounds on graph embedding problems. *J. ACM* 64, 3 (2017), 18:1–18:22. DOI: <https://doi.org/10.1145/3051094>
- [8] Rodney G. Downey and Michael R. Fellows. 1999. *Parameterized Complexity*. Springer-Verlag, New York.
- [9] Fedor V. Fomin and Dieter Kratsch. 2010. *Exact Exponential Algorithms*. Springer. An EATCS Series: Texts in Theoretical Computer Science.
- [10] Fedor V. Fomin, Daniel Lokshtanov, Ivan Mihajlin, Saket Saurabh, and Meirav Zehavi. 2020. Computation of Hadwiger number and related contraction problems: Tight lower bounds. In *Proceedings of the 47th International Colloquium on Automata, Languages, and Programming (ICALP’20) (Virtual Conference)*, 49:1–49:18.
- [11] Martin Charles Golumbic. 2004. *Algorithmic Graph Theory and Perfect Graphs*. North-Holland Publishing Co., Amsterdam, The Netherlands.
- [12] Martin Grohe, Ken-ichi Kawarabayashi, Dániel Marx, and Paul Wollan. 2011. Finding topological subgraphs is fixed-parameter tractable. In *Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC’11)*, 479–488.
- [13] Russell Impagliazzo and Ramamohan Paturi. 1999. Complexity of k-SAT. In *Proceedings of the 14th Annual IEEE Conference on Computational Complexity*, 237–240.
- [14] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. 2001. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.* 63, 4 (2001), 512–530.
- [15] Eugene L. Lawler. 1976. A note on the complexity of the chromatic number problem. *Inform. Process. Lett.* 5, 3 (1976), 66–67.

- [16] Andrzej Lingas and Martin Wahlen. 2009. An exact algorithm for subgraph homeomorphism. *J. Discrete Algorithms* 7, 4 (2009), 464–468. DOI : <https://doi.org/10.1016/j.jda.2008.10.003>
- [17] Neil Robertson and Paul D. Seymour. 1995. Graph minors. XIII. The disjoint paths problem. *J. Comb. Theory Ser. B* 63, 1 (1995), 65–110.
- [18] Patrick Traxler. 2008. The time complexity of constraint satisfaction. In *Parameterized and Exact Computation*. Springer, 190–201.

Received April 2020; revised December 2020; accepted December 2020