Electronic Colloquium on Computational Complexity, Report No. 145 (2014)

On the AC^0 Complexity of Subgraph Isomorphism

Yuan Li* Alexander Razborov† Benjamin Rossman‡

October 15, 2014

Abstract

Let P be a fixed graph (hereafter called a "pattern"), and let $\textsc{Subgraph}(P)$ denote the problem of deciding whether a given graph G contains a subgraph isomorphic to P. We are interested in AC^0-complexity of this problem, determined by the smallest possible exponent $C(P)$ for which $\textsc{Subgraph}(P)$ possesses bounded-depth circuits of size $n^{C(P)+o(1)}$. Motivated by the previous research in the area, we also consider its "colorful" version $\textsc{Subgraph}_{\mathsf{col}}(P)$ in which the target graph G is $V(P)$-colored, and the average-case version $\textsc{Subgraph}_{\mathsf{ave}}(P)$ under the distribution $G(n, n^{-\theta(P)})$, where $\theta(P)$ is the threshold exponent of P. Defining $C_{\mathsf{col}}(P)$ and $C_{\mathsf{ave}}(P)$ analogously to $C(P)$, our main contributions can be summarized as follows.

- $C_{\mathsf{col}}(P)$ coincides with the tree-width of the pattern P within a logarithmic factor. This shows that the previously known upper bound by Alon, Yuster, Zwick [3] is almost tight.
- We give a characterization of $C_{\mathsf{ave}}(P)$ in purely combinatorial terms within a multiplicative factor of 2. This shows that the lower bound technique of Rossman [26] is essentially tight, for any pattern P whatsoever.
- We prove that if Q is a minor of P then $\textsc{Subgraph}_{\mathsf{col}}(Q)$ is reducible to $\textsc{Subgraph}_{\mathsf{col}}(P)$ via a linear-size monotone projection. At the same time, we show that there is *no* monotone projection whatsoever that reduces $\textsc{Subgraph}(M_3)$ to $\textsc{Subgraph}(P_3 + M_2)$ (P_3 is a path on 3 vertices, M_k is a matching with k edges, and "+" stands for the disjoint union). This result strongly suggests that the colorful version of the subgraph isomorphism problem is much better structured and well-behaved than the standard (worst-case, uncolored) one.

1 Introduction

The *subgraph isomorphism problem* takes as its input two graphs H and G and asks to determine whether or not G contains a subgraph (not necessarily induced) isomorphic to H. This is one of the most basic NP-complete problems that includes $\textsc{Clique}$ and $\textsc{Hamiltonian Cycle}$ as special cases, and little more can be said about its complexity in full generality.

A significant body of research, motivated both by the framework of parameterized complexity and practical applications, has been devoted to the case when the graph H is fixed and possesses

*University of Chicago, yuanli@cs.uchicago.edu.

†University of Chicago, razborov@cs.uchicago.edu. Part of this work was done while the author was at Steklov Mathematical Institute, supported by the Russian Foundation for Basic Research, and at Toyota Technological Institute, Chicago.

‡National Institute of Informatics, rossman@nii.ac.jp. Supported by JST ERATO Kawarabayashi Large Graph Project.

ISSN 1433-8092

some useful structure (see e.g. the sources [3, 9, 10, 20, 21, 23] related to the subject of our paper). To stress its nature in this situation, the graph H is traditionally called a *pattern* and designated by the letter P; we also follow this convention and denote by $\textsc{Subgraph}(P)$ the corresponding restriction of the general subgraph isomorphism problem.

The sources above (among many others!) provide quite non-trivial improvements on the obvious size bound $O(n^{|V(P)|})$ in many cases of interest. But for *unconditional* lower bounds we, given our current state of knowledge, have to resort to restricted models, and, indeed, a substantial amount of work has been done here in the context of both bounded-depth circuits and monotone circuits. In this paper we focus on the former model.

As for upper bounds, it was observed by Amano [4] that the color-coding algorithm by Alon, Yuster and Zwick [3] can be adapted to our context and gives AC^0 circuits for $\textsc{Subgraph}(P)$ of size $\widetilde{O}(n^{tw(P)+1})$, where $tw(P)$ is the treewidth of the pattern P. Our paper is motivated by the following natural question:

How tight is this bound?

Or, in other words,

Question 1. *Is it possible to give good general lower bounds on the* AC^0 *complexity of* $\textsc{Subgraph}(P)$ *in terms of the treewidth of* P *only?*

Prior to our work, Rossman [26] answered this question in affirmative for the case of a k-clique by proving a lower bound of $\Omega(n^{k/4})$ on the AC^0 complexity of $\textsc{Subgraph}(K_k)$. Generalizing Rossman's method, Amano [4] gave a general lower bound that holds for arbitrary patterns P. It in particular implied an $n^{\Omega(k)}$ lower bound (and, thus, an affirmative answer to Question 1) for the $k \times k$ grid $G_{k,k}$: this result is very interesting since $G_{k,k}$ is the "canonical" example of a sparse graph with large treewidth.

Before discussing our results, it will be convenient to introduce the following handy notation: given a pattern P, we let $C(P)$ be the minimal real number $c \geq 0$ for which $\textsc{Subgraph}(P)$ is solvable on n-vertex graphs by AC^0 circuits of size $n^{c+o(1)}$. In this notation, the previous results mentioned above can be stated as $C(P) \leq tw(P)+1$ ([3, 4], P any pattern), $C(K_k) \geq k/4$ [26] and $C(G_{k,k}) \geq \Omega(k)$ [4].

Our contributions.

We formulate explicitly and study two modifications that already played a great role in the previous research. The first of them is the *colorful P-subgraph isomorphism problem* $\textsc{Subgraph}_{\mathsf{col}}(P)$ in which the target graph G comes with a coloring $\chi : V(G) \to V(P)$ (that w.l.o.g. can and will be assumed to be a graph homomorphism), and we are looking only for properly colored P-subgraphs. Let $C_{\mathsf{col}}(P)$ be defined analogously to $C(P)$. Then the very first thing done by the algorithm of Alon, Yuster and Zwick is a simple reduction from $\textsc{Subgraph}(P)$ to $\textsc{Subgraph}_{\mathsf{col}}(P)$ thus establishing $C(P) \leq C_{\mathsf{col}}(P)$. After that they work exclusively with the colorful version that leads to

$$C(P) \leq C_{\mathsf{col}}(P) \leq tw(P) + 1.$$

We settle in the affirmative (up to a logarithmic factor) our motivating Question 1 for the colorful version by proving the following

Theorem 1.1. $C_{\mathsf{col}}(P) \geq \Omega(tw(P)/\log tw(P))$.

By previous work of Marx [20], it was known that $\text{SUBGRAPH}_{\mathsf{col}}(P)$ has no $n^{o(tw(P)/\log tw(P))}$ algorithm unless the Exponential Time Hypothesis fails. Theorem 1.1 establishes the same lower bound *unconditionally* for AC^0 circuits. (We say more about Marx's result and related work of Alon and Marx [1] in Section 6.)

We show that the colorful version is quite well-behaved by proving that it is minor-monotone: if Q is a minor of P, then $C_{\mathsf{col}}(Q) \le C_{\mathsf{col}}(P)$ (Theorem 5.1).[1] Whether a similar result holds for $C(P)$ is open, but we give a strong evidence (Theorem 5.6) that even if this is true, the proof will most likely require totally different techniques. One possible interpretation is that perhaps the colorful version is in fact a cleaner and more natural model to study than the standard (uncolored) version. We also observe that if the pattern P is a core (i.e., every homomorphism from P to P is an automorphism), then $C(P) = C_{\mathsf{col}}(P)$ and thus our lower bound from Theorem 1.1 transfers to the uncolored case. What happens to $C(P)$ at the opposite side of the spectrum, say, for bipartite patterns P, remains wide open.

All lower bounds surveyed above, including our proof of Theorem 1.1, were actually achieved in the context of average-case complexity. Prior to our work, the only distribution that was considered for this purpose is the *Erdős-Rényi* model $G(n, n^{-\theta(P)})$, where $\theta(P)$ is the uniquely defined *threshold exponent* for which the probability of containing a copy of P is bounded away from 0 and 1 (see [18] or Section 2.4 below). Accordingly, we define $C_{\mathsf{ave}}(P)$ analogously to $C(P)$, but only require that our circuit outputs the correct answer a.a.s. (asymptotically almost surely) when the input is drawn from $G(n, n^{-\theta(P)})$. Clearly, $C_{\mathsf{ave}}(P) \le C(P)$ so the whole picture now looks like

$$C_{\mathsf{ave}}(P) \le C(P) \le C_{\mathsf{col}}(P) \approx tw(P),$$

where $\approx$ means approximation within a logarithmic factor. Also, $C_{\mathsf{ave}}(K_k) \ge k/4$ [26] and $C_{\mathsf{ave}}(G_{k,k}) \ge \Omega(k)$ [4] where K_k is the complete graph on k vertices and $G_{k,k}$ is the k-by-k grid.

We explicitly define a combinatorial parameter $\kappa(P)$ and prove the following

Theorem 1.2. $\kappa(P) \le C_{\mathsf{ave}}(P) \le 2\kappa(P) + O(1)$.

In other words, we give lower and upper bounds on the average-case AC^0 complexity for an arbitrary pattern P, matching within a quadratic factor. The proof of Theorem 1.2 exploits a duality in the definition of $\kappa(P)$, which has equivalent min-max and max-min formulations (the former suited to upper bounds and the latter to lower bounds). The lower bound $C_{\mathsf{ave}}(P) \ge \kappa(P)$ generalizes the proof of $C_{\mathsf{ave}}(K_k) \ge k/4$ in Rossman [26] and improves a previous lower bound of Amano [4] for general patterns P. (A detailed comparison with previous work is given in Section 2.6 following the definition of $\kappa(P)$.)

Finally, let us say a few words about the proof of Theorem 1.1. Itself a worst-case lower bound, it is obtained as the maximum of a family of average-case lower bounds with respect to P-colored random graphs. These random graphs generalize Erdős-Rényi random graphs in the P-colored setting by allowing different edge probabilities according to the color classes of vertices, and we believe that this generalization may be of independent interest. Each P-colored random graph in this family is parameterized by a point in a certain convex polytope, denoted $\theta_{\mathsf{col}}(P)$. We rely on results of [11, 20] that characterize the treewidth of P in terms of the existence of a certain concurrent flow on P, which we convert to a suitable point in $\theta_{\mathsf{col}}(P)$.

[1] It is worth observing that this fact, along with the recent result [8] by Chekura and Chuzhoy and Amano's bound $C_{\mathsf{col}}(G_{k,k}) \ge \Omega(k)$ [4] already implies the weaker bound $C_{\mathsf{col}}(P) \ge tw(P)^{\Omega(1)}$. But the exponent given by this approach will be disappointingly small.

The paper is organized as follows. In Section 2 we give the necessary definitions and preliminaries; in particular, in Section 2.5 we present the parameters $\kappa(P)$ and $\kappa_{\mathsf{col}}(P)$ that are our main technical tools in this paper. Section 3 is devoted to the proof of Theorem 1.2, and it also paves way to the proof of Theorem 1.1 that, up to a certain point, goes in parallel to the former. The proof of Theorem 1.1 is completed in Section 4. Section 5 contains our structural results about the behavior of $\textsc{Subgraph}(P)$ and $\textsc{Subgraph}_{\mathsf{col}}(P)$ with respect to minors and subgraphs. The paper is concluded with a brief discussion and a list of open problems in Section 6.

2 Definitions and Preliminaries

Let $[k] := \{1, \ldots, k\}$.

2.1 Graphs

We start off with terminology and notation for graphs. Throughout this paper, *graphs* are finite simple graphs $G = (V(G), E(G))$ where $E(G)$ is a subset of $\binom{V(G)}{2}$. We often write $v(G)$ for $|V(G)|$ and $e(G)$ for $|E(G)|$.

A graph H is a *subgraph* of G, denoted $H \subseteq G$, if $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$. For arbitrary G and H, $G+H$ and $G \times H$ respectively denote the *disjoint union* and *Cartesian product* of graphs G and H (where $E(G \times H) := \{\{(v, v'), (w, w')\} : \{v, w\} \in E(G) \text{ and } \{v', w'\} \in E(H)\}$).

A *homomorphism* from G to H is a function $\varphi : V(G) \to V(H)$ such that $\{\varphi(v), \varphi(w)\} \in E(H)$ for all $\{v, w\} \in E(G)$. A graph G is a *core* if every homomorphism from G to G is an automorphism.

The *treewidth* of G is denoted by $tw(G)$ (for the definition and background, see e.g. [6]). Other relevant facts about treewidth will be stated where needed.

K_k is a *clique* on k vertices, and $G_{k,k}$ is a $k \times k$ *grid*. These graphs have treewidth $tw(K_k) = k-1$ and $tw(G_{k,k}) = k$.

2.2 Monotone Projections

Definition 2.1. Let I, J be arbitrary sets.

(i) For a function $p : J \to I \cup \{0,1\}$ and $x \in \{0,1\}^I$, we write $p^*(x)$ for the unique $y \in \{0,1\}^J$ such that $y_j = x_{p(j)}$ if $p(j) \in I$, and $y_j = p(j)$ if $p(j) \in \{0,1\}$.

(ii) For boolean functions $f : \{0,1\}^I \to \{0,1\}$ and $g : \{0,1\}^J \to \{0,1\}$, we say that f is *reducible via a monotone projection* to g, denoted $f \le_{\mathrm{mp}} g$, if there exists $p : J \to I \cup \{0,1\}$ such that $f(x) = g(p^*(x))$ for all $x \in \{0,1\}^I$. (Note that $\le_{\mathrm{mp}}$ is transitive.)

Any decision problem L can be represented as a sequence of Boolean functions $\{L^n\}$ in n variables. We say that L_1 is *reducible via a monotone projection* to another decision problem L_2 if for any n there exists[2] $m(n)$ such that $L_1^n \le_{\mathrm{mp}} L_2^{m(n)}$. If in addition $m(n) \le O(n)$, we call this projection *linear*.

[2]Uniformity issues do not play any role in this paper.

2.3 Subgraph Isomorphism Problems

Throughout this paper, the letters P, Q represent arbitrary fixed graphs that should be intuitively thought of as "patterns". G stands for a (large) "input" graph for the P-subgraph isomorphism problem. Subgraphs of G (not necessarily induced) which are isomorphic to P will be called *P-subgraphs*.

We also consider *P-colored graphs*, defined as pairs (G, χ) where G is a graph and $\chi : V(G) \to V(P)$ is a homomorphism. We usually suppress χ and simply refer to G as *P-colored graph*. In this setting, given a sub-pattern $Q \subseteq P$ (not necessarily induced), a *Q-subgraph* of G is a subgraph of G (again, not necessarily induced) that is isomorphic to Q under χ. In particular, its vertices are mapped bijectively to $V(Q)$ via χ.

We consider two versions ("uncolored" and "colored") of the P-subgraph isomorphism problem:

- $\textsc{Subgraph}(P)$ is the problem, given a graph G, of determining whether or not G contains a P-subgraph.
- $\textsc{Subgraph}_{\mathsf{col}}(P)$ is the problem, given a P-colored graph (G, χ), of determining whether of not G contains a (properly colored) P-subgraph.

This problem is also known in the literature as the "partitioned" or "colorful" variant, and in this paper we mostly adopt the latter term.

It will be convenient to introduce a notation for the AC^0 complexity of these problems. (Recall that AC^0 is the class of problems solvable by polynomial-size constant-depth boolean circuits with unbounded fan-in.)

Definition 2.2. Let $C(P)$ (resp. $C_{\mathsf{col}}(P)$) denote the minimum real number $c > 0$ such that $\textsc{Subgraph}(P)$ (resp. $\textsc{Subgraph}_{\mathsf{col}}(P)$) is solvable (in the worst-case) on n-vertex graphs by AC^0 circuits of size[3] $O(n^{c+\varepsilon})$ for every $\varepsilon > 0$.

Note that if $\textsc{Subgraph}(P)$ is reducible to $\textsc{Subgraph}(Q)$ via a linear monotone projection then $C(P) \le C(Q)$, and this remains true if we add the subscript col to both sides.

Lemma 2.3.

1. *$C(P) \le C_{\mathsf{col}}(P) \le tw(P) + 1$.*
2. *If P is a core, then $C(P) = C_{\mathsf{col}}(P)$.*

Proof. (1): The second inequality $C_{\mathsf{col}}(P) \le tw(P) + 1$ is by the color-coding algorithm of Alon, Yuster and Zwick [3] (adapted to the P-colored setting), which can be implemented in AC^0 as observed by Amano [4]. The first inequality $C(P) \le C_{\mathsf{col}}(P)$ is also implicitly proved there by reducing $\textsc{Subgraph}(P)$ to $\textsc{Subgraph}_{\mathsf{col}}(P)$: the reduction searches through logarithmically many different colorings $\chi_1, \chi_2, \dots : V(G) \to V(P)$ of the same target graph G, picked at random. An easy counting argument shows that a.a.s. every P-subgraph of G will be properly colored with respect to at least one of the colorings χ_i.

(2): This observation goes back at least to Grohe [12]. If P is a core, then $(G, \chi) \mapsto G$ is a reduction from $\textsc{Subgraph}_{\mathsf{col}}(P)$ to $\textsc{Subgraph}(P)$. To see why, it suffices to show that every P-subgraph of G is properly colored with respect to every homomorphism $\chi : G \to P$. Suppose

[3] In this paper, the size of all constant-depth circuits is measured by the number of *gates*.

H is a P-subgraph of G. Then $H = \varphi(P)$ for some one-to-one homomorphism $\varphi : P \to G$. Since P is a core, the homomorphism $\chi \circ \varphi : P \to P$ is an automorphism of P. It follows that the homomorphism $\chi|_{V(H)} : H \to P$ is one-to-one. Since $|E(H)| = |E(P)|$, it must be an isomorphism, that is H is properly colored with respect to χ. □

2.4 The Average Case

We now define the random graphs which appear in our average-case lower bounds for $\textsc{Subgraph}(P)$ and $\textsc{Subgraph}_{\mathsf{col}}(P)$. In the uncolored setting, we consider the Erdős-Rényi random graph $G(n, p(n))$ for an appropriately chosen threshold function $p(n)$. Also, in what follows we assume that P is non-empty, that is contains at least one edge.

Definition 2.4.

(i) The *threshold exponent* of P is defined by $\theta(P) := \min_{Q \subseteq P} v(Q)/e(Q)$.

(ii) P is *balanced* if $v(P)/e(P) = \theta(P)$.

(iii) P is *strictly balanced* if $v(Q)/e(Q) > \theta(P)$ for every nonempty proper subgraph $Q \subset P$.

(iv) Let $\mathsf{Bal}(P) := \bigcup\{Q \subseteq P : v(Q)/e(Q) = \theta(P)\}$.

Lemma 2.5.

1. *P is balanced if and only if $P = \mathsf{Bal}(P)$.*
2. *For every P, $\mathsf{Bal}(P)$ is balanced and $\theta(\mathsf{Bal}(P)) = \theta(P)$.*

Proof. It suffices to show that $\mathcal{B} := \{Q \subseteq P : v(Q)/e(Q) = \theta(P)\}$ is closed under unions (in fact, it is closed under intersections as well). For all $Q_1, Q_2 \in \mathcal{B}$, we have

$$\begin{aligned} v(Q_1 \cup Q_2) + v(Q_1 \cap Q_2) &= v(Q_1) + v(Q_2) \\ &= \theta(P)e(Q_1) + \theta(P)e(Q_2) \\ &= \theta(P)e(Q_1 \cup Q_2) + \theta(P)e(Q_1 \cap Q_2). \end{aligned} \tag{1}$$

By definition of $\theta(P)$,

$$v(Q_1 \cup Q_2) \geq \theta(P)e(Q_1 \cup Q_2) \quad \text{and} \quad v(Q_1 \cap Q_2) \geq \theta(P)e(Q_1 \cap Q_2). \tag{2}$$

Together (1) and (2) imply that equality holds in (2), that is, $Q_1 \cup Q_2$ and $Q_1 \cap Q_2$ are both in $\mathcal{B}$. □

Recall that $G(n, p)$ is the Erdős-Rényi random graph with vertex set $[n]$, in which each $e \in \binom{[n]}{2}$ occurs as an edge independently with probability p. The next lemma states that $p = n^{-\theta(P)}$ is a threshold function for $\textsc{Subgraph}(P)$ and that detecting P-subgraphs on $G(n, n^{-\theta(P)})$ is equivalent to detecting $\mathsf{Bal}(P)$-subgraphs. (Lemma 2.6(1) is a standard fact about random graphs (see [18]); Lemma 2.6(2) was proved in [7].)

Lemma 2.6.

1. $\Pr[\,G(n, n^{-\theta(P)})$ has a P-subgraph$\,]$ *is bounded away from* 0 *and* 1.

2. *Asymptotically almost surely, if $G(n, n^{-\theta(P)})$ contains a $\mathsf{Bal}(P)$-subgraph, then it contains a P-subgraph.*

With slight abuse of notation, we denote by $\textsc{Subgraph}_{\mathsf{ave}}(P)$ the algorithmic problem of solving $\textsc{Subgraph}(P)$ on $G(n, n^{-\theta(P)})$ correctly a.a.s, that is with probability that tends to 1 as n tends to ∞. (We remark that our results are unchanged if $n^{-\theta(P)}$ is replaced by any other threshold function $p(n) \in \Theta(n^{-\theta(P)})$.) Similarly to Definition 2.2, let $C_{\mathsf{ave}}(P)$ be the smallest $c > 0$ for which this problem can be solved by AC^0-circuits of size $n^{c+o(1)}$.

Remark 2.7. Obviously, $C_{\mathsf{ave}}(P) \leq C(P)$, but the gap between them can be arbitrarily large. Assume e.g. that $P = K_4 + G_{k,k}$ where $k \to \infty$. Then $\mathsf{Bal}(P) = K_4$ and thus Lemma 2.6(2) implies that $C_{\mathsf{ave}}(P) = C_{\mathsf{ave}}(K_4) \leq 4$. On the other hand, $\textsc{Subgraph}(G_{k,k})$ is reduced to $\textsc{Subgraph}(P)$ via an obvious linear monotone projection that takes G to K_4+G. This proves $C(P) \geq C(G_{k,k}) \geq \Omega(k)$ by the result from [4].

One might argue that this example is not "fair" since it heavily exploits the fact that the pattern P is highly unbalanced. It is, however, possible to give nearly the same separation (albeit, more complicated) with a strictly balanced pattern P. Say, let $d > 0$ be a sufficiently large constant, and $V(P) = [k]$, where $k \gg d$. We start building $E(P)$ with the clique on the set $[d]$, and then for every $i \in \{d+1, \ldots, k\}$ pick at random d different vertices $j_1, \ldots, j_d < i$ and add all d edges $\{j_\nu, i\}$. Then P will be strictly balanced, and randomness in selecting the edges will imply that a.a.s. $tw(P) \geq \Omega(k)$ and that P is a core. Given these facts, the bounds $C_{\mathsf{ave}}(P) \leq O(d)$ and $C(P) \geq \Omega(k/\log k)$ readily follow from the main results of our paper.

We now move onto the notion of average case complexity for $\textsc{Subgraph}_{\mathsf{col}}(P)$. In contrast to the uncolored setting, there is no single most obvious distribution on P-colored random graphs. Instead, we consider a family of P-colored random graphs, denoted $G_{\alpha,\beta}(n)$, which are parameterized by certain pairs of functions $\alpha : V(P) \to [0,1]$ and $\beta : E(P) \to [0,2]$ called "threshold pairs". (Note: Unlike $G(n,p)$, the vertex set of $G_{\alpha,\beta}(n)$ is not $[n]$, but rather consists of $|V(P)|$ disjoint parts of different sizes.)

Definition 2.8. (P-colored random graph $G_{\alpha,\beta}(n)$)

(i) A *threshold pair* for P is a pair (α, β) of functions $\alpha : V(P) \to [0,1]$ and $\beta : E(P) \to [0,2]$ such that

- $\alpha(P) = \beta(P)$,
- $\alpha(Q) \geq \beta(Q)$ for all $Q \subseteq P$,

where $\alpha(Q) := \sum_{v \in V(Q)} \alpha(v)$ and $\beta(Q) := \sum_{e \in E(Q)} \beta(e)$.

(ii) $\theta_{\mathsf{col}}(P)$ denotes the set of threshold pairs for P. Note that $\theta_{\mathsf{col}}(P)$ is a polytope in $\mathbb{R}^{V(P) \cup E(P)}$ and its section $\{\beta : (1, \beta) \in \theta_{\mathsf{col}}(P)\}$ is a polytope in $\mathbb{R}^{E(P)}$. We view elements of $\theta_{\mathsf{col}}(P)$ as the "P-colored" analogue of $\theta(P)$.

(iii) We say that $(\alpha, \beta) \in \theta_{\mathsf{col}}(P)$ is *nontrivial* if α and β are not identically zero.

(iv) We say that $(\alpha, \beta) \in \theta_{\mathsf{col}}(P)$ is *strictly balanced* if $\alpha(Q) > \beta(Q)$ for every nonempty proper subgraph $Q \subset P$.

(v) For all $(\alpha, \beta) \in \theta_{\mathsf{col}}(P)$, let $G_{\alpha,\beta}(n)$ denote the random graph with vertex set $V_\alpha(n) := \{(v,i) : v \in V(P),\ 1 \le i \le \lfloor n^{\alpha(v)} \rfloor\}$ where each $\{(v,i),(w,j)\}$ with $\{v,w\} \in E(P)$ is an edge, independently, with probability $n^{-\beta(\{v,w\})}$. The P-coloring of $G_{\alpha,\beta}(n)$ is the obvious one: $(v,i) \mapsto v$.

Remark 2.9. Note that if P is a balanced pattern, then the pair of constant functions ($\alpha \equiv 1,\ \beta \equiv \theta(P)$) is a threshold pair for P; moreover, P is strictly balanced if and only if this (α, β) is strictly balanced. Thus, Definition 2.8 is indeed a generalization of threshold exponent for balanced patterns. The following lemma makes the analogy even more clear (and justifies the terminology "threshold pair").

Lemma 2.10. *For every pattern P and nontrivial threshold pair $(\alpha, \beta) \in \theta_{\mathsf{col}}(P)$,*

1. $\liminf_{n\to\infty} [\, G_{\alpha,\beta}(n)$ *contains no P-subgraph*$\,] \ge \dfrac{1}{e}$,

2. $\liminf_{n\to\infty} [\, G_{\alpha,\beta}(n)$ *contains exactly one P-subgraph*$\,] \ge \dfrac{1}{e^{|E(P)|}}$.

The proof is included in Appendix A. With a bit of work, it is possible to completely characterize the asymptotic distribution of the number of P-subgraphs in $G_{\alpha,\beta}(n)$; this distribution is a function of independent Poisson random variables (in the uncolored setting, see [7] for a characterization of the asymptotic number of P-subgraphs in $G(n, n^{-\theta(P)})$).

In the context of $\textsc{Subgraph}_{\mathsf{col}}(P)$, we speak of the *average-case complexity with respect to* $G_{\alpha,\beta}(P)$, meaning the size of an AC^0 circuit which solves $\textsc{Subgraph}_{\mathsf{col}}(P)$ on $G_{\alpha,\beta}(P)$ with probability that tends to 1 as n tends to ∞. We do not introduce any special notation like $C_{\alpha,\beta}(P)$ as this concept is intended to be auxiliary.

2.5 Parameters $\kappa(P)$ and $\kappa_{\mathsf{col}}(P)$

We now introduce the parameters $\kappa(P)$ and $\kappa_{\mathsf{col}}(P)$ which figure in our lower bounds. The definitions, which might appear unmotivated at first glance, are derived from the lower bound technique of [26], which we explain in the next section.

Definition 2.11. (Union sequences and hitting sets) A *union sequence* for P is a sequence $Q_1, \dots, Q_t$ of subgraphs of P such that $Q_t = P$ and for all $1 \le k \le t$, either Q_k is a single vertex or a single edge or $Q_k = Q_i \cup Q_j$ for some $1 \le i < j < k$. A *hitting set for union sequences* (or *hitting set* for short) is a set $\mathcal{H}$ of subgraphs of P such that $\mathcal{H}$ contains at least one element from every union sequence.

Definition 2.12. (Parameters $\kappa(P)$, $\kappa_{\alpha,\beta}(P)$ and $\kappa_{\mathsf{col}}(P)$)

(i) If P is balanced, then $\kappa(P)$ is defined by

$$\kappa(P) := \min_{\text{union seq. } Q_1,\dots,Q_t} \max_{i \in [t]} v(Q_i) - \theta(P) e(Q_i).$$

For P which is not balanced, we define $\kappa(P) := \kappa(\mathsf{Bal}(P))$.

(ii) For $(\alpha, \beta) \in \theta_{\mathsf{col}}(P)$, let

$$\kappa_{\alpha,\beta}(P) := \min_{\text{union seq. } Q_1,\dots,Q_t} \max_{i \in [t]} \alpha(Q_i) - \beta(Q_i).$$

(iii) Let $\kappa_{\mathsf{col}}(P) := \max_{(\alpha,\beta)\in\theta_{\mathsf{col}}(P)} \kappa_{\alpha,\beta}(P)$.

Remark 2.13. Later on we will see that in this definition we could restrict ourselves to threshold pairs with $\alpha \equiv 1$ (Corollary 4.2). But since arbitrary threshold pairs appear quite naturally in our lower bound proofs in Section 4.2, we prefer to give this more general definition at once.

The next lemma is key to linking our upper and lower bounds on the average-case AC^0 complexity of SUBGRAPH(P).

Lemma 2.14. (Minimax principle for $\kappa(P)$ and $\kappa_{\alpha,\beta}(P)$)

1. *If P is balanced, then*
$$\kappa(P) = \max_{\mathcal{H}} \min_{Q\in\mathcal{H}} v(Q) - \theta(P)e(Q),$$
where $\mathcal{H}$ ranges over hitting sets for P.

2. *Similarly, $\kappa_{\alpha,\beta}(P) = \max_{\mathcal{H}} \min_{Q\in\mathcal{H}} \alpha(Q) - \beta(Q)$ for all $(\alpha,\beta) \in \theta_{\mathsf{col}}(P)$.*

Proof. The argument is the same for (1) and (2). Let $f(Q) := v(Q) - \theta(P)e(Q)$ (the proof works for any real-valued objective function). First, we will prove that $\max_{\mathcal{H}} \min_{Q\in\mathcal{H}} f(Q) \leq \kappa(P)$. Since $\mathcal{H}$ is a hitting set, for any union sequence $\{Q_i\}$, there exists some $Q_i \in \mathcal{H}$. It follows that $\min_{Q\in\mathcal{H}} f(Q) \leq \max_i f(Q_i)$, and thus $\min_{Q\in\mathcal{H}} f(Q) \leq \kappa(P)$ as $\{Q_i\}$ is taken arbitrarily.

On the other hand, let us prove $\kappa(P) \leq \max_{\mathcal{H}} \min_{Q\in\mathcal{H}} f(Q)$. Enumerate all union sequences $\{Q_i^{(j)}\}$, $j = 1, 2, \ldots$ (each $\{Q_i^{(j)}\}$ is a finite sequence). For each j, take the subgraph $S^{(j)}$ in $\{Q_i^{(j)}\}$ with maximal $f(Q_i^{(j)})$. Let $\mathcal{S} = \{S^{(1)}, S^{(2)}, \ldots\}$. It is easily seen that $\mathcal{S}$ is a hitting set, as every union sequence has some element in it. By definition,
$$\max_{\mathcal{H}} \min_{Q\in\mathcal{H}} f(Q) \geq \min_{S^{(j)}\in\mathcal{S}} f(S^{(j)}) = \min_j \max_i f(P_i^{(j)}) = \kappa(P),$$
which completes the proof. ☐

2.6 Comparison with previous work

The dual (max-min) expression for $\kappa(P)$ given by Lemma 2.14(1) is naturally suited to *lower bounds*. It is this dual version of $\kappa(P)$ which we use to prove $C_{\mathsf{ave}}(P) \geq \kappa(P)$ in the next section. This dual expression—which maximizes over hitting sets for a pattern P—generalizes Rossman's proof of $C_{\mathsf{ave}}(K_k) \geq k/4$ in [26], which considers a specific hitting set for K_k.[4]

Previous work of Amano [4] also generalizes the technique of [26] to obtain a lower bound $C_{\mathsf{ave}}(P) \geq \ell(P)$ for general patterns P. The function $\ell(P)$ defined by Amano (which is denoted $Z_P^\star$ in [4]) is similar to the dual expression for $\kappa(P)$, except it restricts attention to hitting sets of a particular form:
$$\ell(P) := \max_{s\,:\,2\leq s\leq v(P)} \; \min_{Q\subseteq P\,:\,s/2<v(Q)\leq s} v(Q) - \theta(P)e(Q).$$
Clearly, $\ell(P) \leq \kappa(P)$ for all patterns P. In some cases of interest, such as grid $G_{k,k}$, Amano shows that $\ell(G_{k,k}) = \Omega(k)$. However, $\ell(P)$ is slack in general (for example, $\ell(K_k) = 2k/9 + O(1)$ while

[4]This hitting set consists of the "medium" subgraphs of K_k, defined as subgraphs Q such that (i) Q has $> k/2$ vertices and (ii) Q is a union of two subgraphs with $\leq k/2$ vertices.

$\kappa(K_k) = k/4 + O(1)$). A key insight of the present paper is that the stronger parameter $\kappa(P)$ – in its primal form, which minimizes over union sequences—leads to *upper bounds* on $C_{\mathsf{ave}}(P)$ which are tight within a multiplicative constant.

Another result of Amano [4] is a construction of nearly optimal AC^0 circuits for the average-case k-clique problem, which match the lower bound of [26] by showing $C_{\mathsf{ave}}(K_k) \leq k/4 + O(1)$. Nakagawa and Watanabe [22] observed that Amano's construction generalizes to an upper bound $C_{\mathsf{ave}}(P) \leq u(P) + O(1)$ where $u(P)$ is defined by

$$u(P) := \min_{\text{linear orderings } v_1 < \cdots < v_k \text{ of } V(P)} \max_{j \in [k]} j - \theta(P) e(\{v_1, \ldots, v_j\})$$

and $e(\{v_1, \ldots, v_i\})$ is the number of edges in P among vertices $v_1, \ldots, v_i$. This parameter $u(P)$ is similar to the (primal) definition of $\kappa(P)$, except that $u(P)$ is restricted to union sequences $Q_1, \ldots, Q_t$ where $|V(Q_{i+1}) \setminus V(Q_i)| \leq 1$. Thus, $u(P) \geq \kappa(P)$. However, in contrast to $\kappa(P)$, Nakagawa and Watanabe showed that $u(P)$ is not bounded by any function of $C_{\mathsf{ave}}(P)$: there is a sequence of patterns $P_1, P_2, \ldots$ with $C_{\mathsf{ave}}(P_i) = O(1)$ while $\lim_i u(P_i) = \infty$.

In summary, our bounds $\kappa(P) \leq C_{\mathsf{ave}}(P) \leq 2\kappa(P) + O(1)$ (Theorem 1.2) both achieve a tighter generalization of [26] and close the (arbitrarily large) gap between the previous bounds $\ell(P) \leq C_{\mathsf{ave}}(P) \leq u(P) + O(1)$ of [4, 22]. Our results on $C_{\mathsf{col}}(P)$, including the definitions of $\theta_{\mathsf{col}}(P)$ and $\kappa_{\mathsf{col}}(P)$, are completely new to this paper (the colored setting was not considered in [4, 22, 26]).

3 Average-Case AC^0 Complexity

In this section, we prove Theorem 1.2 ($\kappa(P) \leq C_{\mathsf{ave}}(P) \leq 2\kappa(P) + O(1)$), which gives a combinatorial characterization of the AC^0-complexity of $\textsc{Subgraph}_{\mathsf{ave}}(P)$ up to a quadratic factor. More generally, we prove a family of average-case lower and upper bounds for the average-case colorful P-subgraph isomorphism problem:

Theorem 3.1. *For every pattern P and $(\alpha, \beta) \in \theta_{\mathsf{col}}(P)$, the average-case AC^0-complexity of $\textsc{Subgraph}_{\mathsf{col}}(P)$ on the P-colored random graph $G_{\alpha,\beta}(n)$ is between $n^{\kappa_{\alpha,\beta}(P)-o(1)}$ and $n^{2\kappa_{\alpha,\beta}(P)+O(1)}$.*

Rather than proving Theorem 1.2 and Theorem 3.1 separately, to avoid redundancy we present a proof of the latter only. For balanced P the proof of Theorem 1.2 looks exactly like the proof of Theorem 3.1 in the special case where $\alpha \equiv 1$ and $\beta \equiv \theta(P)$ (see Remark 2.9). The general case is reduced to the balanced one since for an arbitrary pattern P we have $\kappa(P) = \kappa(\mathsf{Bal}(P))$ (by definition of $\kappa(P)$) and $C_{\mathsf{ave}}(P) = C_{\mathsf{ave}}(\mathsf{Bal}(P))$ (by Lemma 2.6).

Theorem 3.1 also plays a key role in our other main result, Theorem 1.1 (the worst-case lower bound $C_{\mathsf{col}}(P) \geq \Omega(tw(P)/\log tw(P))$). Since the worst-case AC^0-complexity of $\textsc{Subgraph}_{\mathsf{col}}(P)$ is lower-bounded by the average-case AC^0-complexity of $\textsc{Subgraph}_{\mathsf{col}}(P)$ on $G_{\alpha,\beta}(n)$ for every $(\alpha, \beta) \in \theta_{\mathsf{col}}(P)$, Theorem 3.1 directly implies:

Corollary 3.2. *$C_{\mathsf{col}}(P) \geq \kappa_{\mathsf{col}}(P)$.*

In Section 4, we will show that $\kappa_{\mathsf{col}}(P) \geq \Omega(tw(P)/\log tw(P))$; together with Corollary 3.2, this proves Theorem 1.1. The remainder of this section contains the proof of Theorem 3.1. The $n^{2\kappa_{\alpha,\beta}(P)+O(1)}$ upper bound is proved in Section 3.1, followed by the $n^{\kappa_{\alpha,\beta}(P)-o(1)}$ lower bound in Section 3.2.

3.1 Proof of Theorem 3.1 (Upper Bound)

Fix a pattern P and a threshold pair $(\alpha, \beta) \in \theta_{\mathsf{col}}(P)$. For a P-colored graph G and $Q \subseteq P$, let $\mathrm{sub}(Q, G)$ denote the number of (colored) Q-subgraphs of G. We write $\mathbf{G}$ for the P-colored random graph $G_{\alpha,\beta}(n)$. Note that $\mathrm{E}[\,\mathrm{sub}(Q, \mathbf{G})\,] \leq n^{\alpha(Q)-\beta(Q)}$.

Let $\mathcal{G}_{\alpha,\beta}(n)$ denote the support of $\mathbf{G}$, that is, the class of P-colored graphs with vertex set $V_\alpha(n) := \{(v, i) : v \in V(P),\ 1 \leq i \leq \lfloor n^{\alpha(v)} \rfloor\}$ and the vertex-coloring $(v, i) \mapsto v$. Let also

$$\mathcal{G}'_{\alpha,\beta}(n) := \{G \in \mathcal{G}_{\alpha,\beta}(n) : \mathrm{sub}(Q, G) \leq n^{\alpha(Q)-\beta(Q)+1} \text{ for all } Q \subseteq P\}.$$

The next lemma says that $\mathbf{G}$ is extremely unlikely to contain significantly more than $n^{\alpha(Q)-\beta(Q)}$ Q-subgraphs for any $Q \subseteq P$. It is proved by a straightforward application of Markov's inequality.

Lemma 3.3. $\Pr[\,\mathbf{G} \notin \mathcal{G}'_{\alpha,\beta}(n)\,] = o(1)$.

We wish to construct a *deterministic* AC^0-circuit C which solves $\textsc{Subgraph}_{\mathsf{col}}(P)$ correctly on $\mathbf{G}$ with probability $1 - o(1)$. We will invert the role of randomness and instead construct a *random* AC^0-circuit $\mathbf{C}$ which solves $\textsc{Subgraph}_{\mathsf{col}}(P)$ correctly with probability $1 - o(1)$ on *every* $G \in \mathcal{G}'_{\alpha,\beta}(n)$. That is, we will show

Lemma 3.4. *There exists a random* AC^0 *circuit* $\mathbf{C}$ *of size* $n^{2\kappa_{\alpha,\beta}(P)+O(1)}$ *and depth*[5] $O(e(P))$ *such that for every* $G \in \mathcal{G}'_{\alpha,\beta}(n)$,

$$\Pr[\,\mathbf{C}(G) = 1 \Leftrightarrow \mathrm{sub}(P, G) \geq 1\,] = 1 - o(1).$$

The upper bound of Theorem 3.1 follows as a corollary of Lemmas 3.3 and 3.4.

Proposition 3.5. *There exists a* AC^0 *circuit* C *of size* $n^{2\kappa_{\alpha,\beta}(P)+O(1)}$ *such that*

$$\Pr[\,\mathsf{C}(\mathbf{G}) = 1 \Leftrightarrow \mathrm{sub}(P, \mathbf{G}) \geq 1\,] = 1 - o(1).$$

Proof. Lemmas 3.3 and 3.4 imply that $\Pr[\,\mathbf{C}(\mathbf{G}) = 1 \Leftrightarrow \mathrm{sub}(P, \mathbf{G}) \geq 1\,] = 1 - o(1)$. Now Proposition 3.5 follows by a straightforward application of Yao's Principle [28]. □

The random circuit C.

It remains to define the randomized AC^0-algorithm solving $\textsc{Subgraph}_{\mathsf{col}}(P)$ with high probability on every $G \in \mathcal{G}'_{\alpha,\beta}(n)$. We first describe the algorithm informally. We then check that this algorithm can be implemented by circuits of size $n^{2\kappa_{\alpha,\beta}(P)+O(1)}$ and depth $O(e(P))$.

By definition of $\kappa_{\alpha,\beta}(P)$, there exists a union sequence $Q_1, \ldots, Q_t$ with $Q_t = P$ such that $\kappa_{\alpha,\beta}(P) = \max_{i \in [t]} \alpha(Q_i) - \beta(Q_i)$. The idea behind the algorithm is simple: given a graph $G \in \mathcal{G}'_{\alpha,\beta}(n)$ (the input), we will compute a sequence $L_1, \ldots, L_t$ of lists, where L_k contains all of the Q_k-subgraphs of G (with high probability). The list L_k will contain $n^{\alpha(Q_k)-\beta(Q_k)+O(1)}$ entries (enough to accommodate all of the Q_k-subgraphs in G). Many entries in L_k will be *blank* (signified by $\emptyset$); by construction, every non-blank entry of L_k will contain the description of a Q_k-subgraph of G

[5] In fact, the depth is linear in the *height* of the optimal union sequence, where the height is defined as the length of the longest path from the root to a leaf in the directed acyclic graph induced by a union sequence.

(as a string of length $\alpha(Q_k)\log n$). Note that blank and non-blank entries in L_k will in general be interleaved (as AC^0 is not powerful enough to sort them).

Some notation: we write ℓ_k for the number of entries in the list L_k. For $a \in [\ell_k]$, we write $L_k(a)$ for the contents of the ath entry in L_k (either $\emptyset$ or a Q_k-subgraph of G). We say that L_k is *good* (with respect to G and the randomness of the algorithm) if L_k contains all Q_k-subgraphs of G exactly once.

Lists $L_1, \dots, L_t$ are computed, in order, as follows. For $k \in [t]$, assume that $L_1, \dots, L_{k-1}$ have been computed and are good. In the case that Q_k is a single edge of P, let L_k have $\ell_k := n^{\alpha(Q_k)}$ entries, indexed by the potential Q_k-subgraphs of G. For $a \in [\ell_k]$, the ath entry $L_k(a)$ will contain the ath potential Q_k-subgraph iff it is a Q_k subgraph of G; otherwise $L_k(a)$ is blank. Clearly L_k is good.

If Q_k is not a single edge, then by the definition of union sequence, $Q_k = Q_i \cup Q_j$ for some $1 \le i < j < k$. We compute L_k in three steps as follows.

Step 1: Let M_k be the $\ell_i \times \ell_j$ array where, for $a \in [\ell_i]$ and $b \in [\ell_j]$,

$$M_k(a,b) := \begin{cases} L_i(a) \cup L_j(b) & \text{if } L_i(a) \text{ and } L_j(b) \text{ are consistent on } V(Q_i) \cap V(Q_j), \\ \emptyset & \text{otherwise.} \end{cases}$$

(Note that, since L_i and L_j are good, M_k contains each Q_k-subgraph of G exactly once. That is, M_k satisfies the "good" condition that we want for L_k.)

Step 2: We hash M_k down to a smaller number of entries to obtain the list L_k. Let $\mathrm{Supp}(M_k) \subseteq [\ell_i] \times [\ell_j]$ denote the set of nonempty entries of M_k. Let $m_k := n^{\alpha(Q_k)-\beta(Q_k)+1}$ and note that $m_k \ge \#\{Q_k\text{-subgraphs of } G\} = |\mathrm{Supp}(M_k)|$. Let $\mathbf{h}_k$ be a uniform random function

$$\mathbf{h}_k : [\ell_i] \times [\ell_j] \to [m_k].$$

(Restricted to the $\le m_k$ nonempty entries of M_k, this gives a uniform random packing of $\le m_k$ balls into m_k bins.)

Step 3: Let $\ell_k := m_k \ln m_k$. Indexing entries of L_k by pairs $(p,q) \in [m_k] \times [\ln m_k]$ (rather than elements of $[\ell_k]$), let

$$L_k(p,q) := \begin{cases} \text{the } q\text{th element of } \mathbf{h}_k^{-1}(p) \cap \mathrm{Supp}(M_k) & \text{if } |\mathbf{h}_k^{-1}(p) \cap \mathrm{Supp}(M_k)| \ge q, \\ \emptyset & \text{otherwise.} \end{cases}$$

Note that L_k is good if, and only if,

$$\bigwedge_{p \in [m_k]} |\mathbf{h}_k^{-1}(p) \cap \mathrm{Supp}(M_k)| \le \ln m_k.$$

After computing the final list L_t, the algorithm outputs 1 iff L_t has non-blank entries. Note that the output of the algorithm will be correct provided L_t is good.

To analyze the success probability of the algorithm, note the following elementary fact about balls-into-bins, established by a simple union bound.[6]

$$(3) \qquad \begin{cases} \textit{For any } \widetilde{m} \le m, \textit{ the maximum load of a random function of } \widetilde{m} \textit{ balls to } m \textit{ bins is} \\ \le \ln m \text{ with probability } \ge 1 - 1/m. \end{cases}$$

[6] A tighter analysis than we require shows that the maximum load is $\le 3\ln m/\ln\ln m$ with probability $\ge 1 - 1/m$.

From this fact, we have

$$\Pr_{\mathbf{h}_k}[\, L_k \text{ is not good} \mid L_1,\dots,L_{k-1} \text{ are good}] \le \Pr_{\mathbf{h}_k}[\, \bigvee_{p\in[m_k]} |\mathbf{h}_k^{-1}(p) \cap \mathrm{Supp}(M_k)| > \ln m_k \,] \le \frac{1}{m_k} \;\le\; \frac{1}{n}.$$

It follows that

$$\begin{aligned}\Pr_{\mathbf{h}_1,\dots,\mathbf{h}_t}[\,\text{erroneous output}\,] &= \sum_{k\in[t]} \Pr_{\mathbf{h}_1,\dots,\mathbf{h}_k}[\, L_k \text{ is not good}, L_1,\dots,L_{k-1} \text{ are good}\,]\\ &\le \sum_{k\in[t]} \Pr_{\mathbf{h}_1,\dots,\mathbf{h}_k}[\, L_k \text{ is not good} \mid L_1,\dots,L_{k-1} \text{ are good}]\\ &\le tn^{-1} \;\le\; o(1).\end{aligned}$$

Therefore, the algorithm correctly solves $\textsc{Subgraph}_{\mathsf{col}}(P)$ with high probability for every $G \in \mathcal{G}'_{\alpha,\beta}(P)$.

It remains to show that this algorithm can be implemented by a random circuit $\mathbf{C}$ of size $n^{2\kappa_{\alpha,\beta}(P)+O(1)}$ and depth $O(e(P))$. We will make an additional assumption about the random functions $\mathbf{h}_1,\dots,\mathbf{h}_t$:

$$(4) \qquad |\mathbf{h}_k^{-1}(p)| \le \frac{2\ell_i\ell_j}{m_k} \quad \text{for all } k \in [t] \text{ and } p \in [m_k].$$

That is, $|\mathbf{h}_k^{-1}(p)|$ is at most twice its expectation for all k and p. By Chernoff and union bounds, (4) holds with probability $1 - \exp(-n^{\Omega(1)})$. So even with this assumption, the error probability of the circuits we describe remains $o(1)$.

Let us now *fix*[7] any particular hash functions $h_1,\dots,h_t$ such that (3) and (4) hold for any $k \in [t]$. We will design constant-depth circuits of size $n^{2\kappa_{\alpha,\beta}(P)+O(1)}$ computing the lists that correspond to our particular choice of $h_1,\dots,h_t$.

We describe the sub-circuit which computes the list L_k given lists $L_1,\dots,L_{k-1}$. In the case that Q_k is a single edge, the list L_k is clearly computable by a depth-2 circuit of size $\widetilde{O}(n^{\alpha(Q_k)})$ (the $\widetilde{O}()$ coming from the fact that it takes $\alpha(Q_k)\log n$ gates to encode each entry of L_k). In the case that $Q_k = Q_i \cup Q_j$, first note that we can compute the array M_k by a circuit of size $\widetilde{O}(n\ell_i\ell_j)$ and depth $O(1)$ (sitting on top of the sub-circuits which compute lists L_i and L_j); this is because checking that $L_i(a)$ and $L_j(b)$ agree on all vertices of $V(Q_i) \cap V(Q_j)$ requires only $O(n)$ size and depth 2. Having computed M_k, computing the entries $L_k(p,q)$ requires sorting $\le \ln m_k = O(\log n)$ elements from the list $h_k^{-1}(p)$ of size $O(\frac{\ell_i\ell_j}{m_k})$. By [15], this can be done by a constant-depth circuit of size $n^{o(1)}\frac{\ell_i\ell_j}{m_k}$. As there are altogether $\ell_k \le \widetilde{O}(m_k)$ pairs (p,q), we get a constant-depth circuit which computes L_k (given L_i and L_j) with total size $n^{o(1)}\ell_i\ell_j \le n^{\alpha(Q_i)-\beta(Q_i)+\alpha(Q_j)-\beta(Q_j)+O(1)} \le n^{2\kappa_{\alpha,\beta}(P)+O(1)}$.

After computing all lists $L_1,\dots,L_t$, we have a circuit of size $n^{2\kappa_{\alpha,\beta}(P)+O(1)}$. Finally, note that the depth of this circuit will be $O(d)$ where d is the height of the poset where $i,j \prec k$ iff $Q_k = Q_i \cup Q_j$ for $i,j < k$. Clearly, $d \le e(P)$ as long as all graphs in the sequence are pairwise distinct.

[7] It is important for our argument that the random functions $\mathbf{h}_1,\dots,\mathbf{h}_t$ are generated in advance and that we do not attempt to make our construction uniform in $\mathbf{h}_1,\dots,\mathbf{h}_t$. Putting it less dramatically, if we are not careful about this point then the multiplicative factor 2 will get increased to 3.

3.2 Proof of Theorem 3.1 (Lower Bound)

This subsection gives the proof of the lower bound in Theorem 3.1 (the average-case AC^0-complexity of $\textsc{Subgraph}_{\mathsf{col}}(P)$ on $G_{\alpha,\beta}(n)$ is at least $n^{\kappa_{\alpha,\beta}(P)-o(1)}$). The argument closely follows the technique of [26, 27].

It will be convenient to work with an alternative characterization of AC^0 as boolean circuits with fan-in 2. We distinguish between "type-I" and "type-II" AC^0 circuits as follows.

type-I: polynomial-size constant-depth $\{\mathrm{AND}_\infty, \mathrm{OR}_\infty, \mathrm{NOT}\}$-circuits with unbounded fan-in (this is the standard definition of AC^0),

type-II: polynomial-size $\{\mathrm{AND}_2, \mathrm{OR}_2, \mathrm{NOT}\}$-circuits with fan-in 2 and arbitrary depth, but $O(1)$ alternations between AND and OR gates (where w.l.o.g. NOT gates are on the bottom level).

The conversion from type-I to type-II replaces each AND_∞ (resp. OR_∞) gate with an almost balanced binary tree of AND_2 (resp. OR_2) gates. Note that this converts a type-I circuit with g gates (and hence $w \leq O(g^2)$ wires) into a type-II circuit with $O(w)$ gates. Since we measure size by the number of gates, a lower bound of S on the size of type-II circuits implies a lower bound of $\Omega(\sqrt{S})$ on the size of type-I circuits.

We will prove an $n^{\kappa_{\alpha,\beta}(P)-o(1)}$ lower bound on the size of type-II circuits solving $\textsc{Subgraph}_{\mathsf{col}}(P)$ in the average-case on $G_{\alpha,\beta}(n)$. This implies a weaker $n^{\kappa_{\alpha,\beta}(P)/2-o(1)}$ lower bound for type-I circuits. The stronger $n^{\kappa_{\alpha,\beta}(P)-o(1)}$ lower bound for type-I circuits can be shown by an additional argument, which we omit (see Section 3.4 of [27]).

Now comes a key definition.

Definition 3.6. Let f be a boolean function on [labeled, n-vertex] graphs, and let H be any graph. The *f-sensitive subgraph* of H, denoted $\mathsf{Sens}(f, H)$, is defined as the unique minimal subgraph $S \subseteq H$ such that $f(H') = f(H' \cap S)$ for every $H' \subseteq H$. We say that f is *sensitive over* H if $\mathsf{Sens}(f, H) = H$.

For all f and H, observe that

(5) f is sensitive over $\mathsf{Sens}(f, H)$ (i.e. $\mathsf{Sens}(f, \mathsf{Sens}(f, H)) = \mathsf{Sens}(f, H)$),

(6) if f is the AND or OR of functions f_1 and f_2, then $\mathsf{Sens}(f, H) \subseteq \mathsf{Sens}(f_1, H) \cup \mathsf{Sens}(f_2, H)$.

We say that a single-output boolean circuit whose variables encode potential edges in a graph is *sensitive over* H if its output function is so.

Lemma 3.7. *Let C be an arbitrary boolean circuit with fan-in 2 that is sensitive over some nonempty graph H. Then there exists a union sequence $H_1, \ldots, H_t = H$ and a sequence $\mathsf{C}_1, \ldots, \mathsf{C}_t$ of sub-circuits of C such that C_i is sensitive over H_i for all $i \in \{1, \ldots, t\}$.*

Proof. We argue by induction on boolean circuits with fan-in 2. In the base case, C is a variable (corresponding to a possible edge). The assumption that C is sensitive over H implies that H is a single edge. Therefore, H itself is a union sequence of length 1 which satisfies the condition of the lemma.

For the induction step, note that if $\mathsf{C} = \mathrm{NOT}(\mathsf{C}')$, then C' is sensitive over H; therefore, the lemma holds by the induction hypothesis for C'. Finally, suppose C is the AND or OR of subcircuits C_1 and C_2. If C_1 or C_2 is sensitive over H, then appealing to the induction hypothesis, we are done. So we will assume that neither C_1 nor C_2 are sensitive over H. Let $H_i := \mathsf{Sens}(\mathsf{C}_i, H)$ for $i = 1, 2$. Then C_i is sensitive over H_i by observation (5). By observation (6),

$$H \;=\; \mathsf{Sens}(\mathsf{C}, H) \;\subseteq\; \mathsf{Sens}(\mathsf{C}_1, H) \cup \mathsf{Sens}(\mathsf{C}_2, H) \;=\; H_1 \cup H_2.$$

Hence $H = H_1 \cup H_2$. By the induction hypothesis, there exist union sequence $S_1, \ldots, S_s = H_1$ and $T_1, \ldots, T_t = H_2$ which satisfy the condition in the lemma with respect to C_1, H_1 and C_2, H_2 respectively. Then $S_1, \ldots, S_s, T_1, \ldots, T_t, H$ is a union sequence which satisfies the condition in the lemma with respect to C and H. $\square$

Definition 3.8. If f is a boolean functions on (labelled, n-vertex) graphs and G is any graph, then let $f^{\cup G}$ denote the function $f^{\cup G}(H) := f(G \cup H)$.

Note that if a boolean circuit C computes a function f on graphs, then the circuit $\mathsf{C}^{\cup G}$ that substitutes 1 for variables corresponding to edges in G computes $f^{\cup G}$.

We now fix a pattern P and a threshold pair $(\alpha, \beta) \in \theta_{\mathsf{col}}(P)$. Without loss of generality, we assume that $\beta(e) > 0$ for all $e \in E(P)$ (otherwise we replace P with the subgraph with edge set $\{e \in E(P) : \beta(e) > 0\}$). We continue to write $\mathbf{G}$ for the P-colored random graph $\mathcal{G}_{\alpha,\beta}(n)$. Let $\mathbf{P}$ be a uniform random "planted" P-subgraph (independent of $\mathbf{G}$). (That is, $\mathbf{P}$ is the P-subgraph with vertex set $\{(v, i_v) : v \in V(P)\}$ where i_v is uniform random in $\{1, \ldots, \lfloor n^{\alpha(v)} \rfloor\}$.) For a subgraph $Q \subseteq P$, let $\mathbf{Q}$ denote the corresponding subgraph of $\mathbf{P}$.

We next state two technical lemmas.

Lemma 3.9. *Suppose $f : \mathcal{G}_{\alpha,\beta}(n) \to \{0,1\}$ solves* $\textsc{Subgraph}_{\mathsf{col}}(P)$ *in the average-case on* $\mathbf{G}$*, that is,*

$$\Pr_{\mathbf{G}}[\, f(\mathbf{G}) = 1 \Leftrightarrow \mathbf{G} \textit{ has a } P\textit{-subgraph}\,] = 1 - o(1). \tag{7}$$

Then

$$\liminf_{n \to \infty} \Pr_{\mathbf{G},\mathbf{P}}[\, f^{\cup \mathbf{G}} \textit{ is sensitive over } \mathbf{P}\,] > 0.$$

Specifically, the proof of Lemma 3.9 shows the following: the event "$E(\mathbf{P}) \cap E(\mathbf{G}) = \emptyset$ and $\mathbf{P}$ is the unique P-subgraph in $\mathbf{G} \cup \mathbf{P}$" holds with constant probability. After that, the assumption (7) rather straightforwardly implies that, conditioned on this event, $f^{\cup \mathbf{G}}$ is almost surely the AND function over edges of $\mathbf{P}$ (i.e. for all $Q \subseteq \mathbf{P}$, $f(\mathbf{G} \cup Q) = 1$ iff $Q = \mathbf{P}$).

The proof of Lemma 3.9 requires a delicate analysis of the random graph $G_{\alpha,\beta}(n)$ and may be considered separately from the central technique in our lower bound. For this reason, we defer the proof to Appendix B. (We remark that Lemma 3.9 may be avoided entirely, or at least greatly simplified, if we only care about worst-case lower bounds, rather than average-case lower bounds at the distribution $G_{\alpha,\beta}(n)$.)

The second technical lemma relies on Håstad's Switching Lemma [14] and its proof closely follows Proposition 3.11 of [27]. (The reader who wishes to skip the technical details is encouraged to jump ahead to Theorem 3.11.)

Lemma 3.10. *Suppose $f : \mathcal{G}_{\alpha,\beta}(n) \to \{0,1\}$ is* AC^0*-computable. Then for every Q,*

$$\Pr_{\mathbf{G},\mathbf{Q}}[\, f^{\cup\mathbf{G}} \text{ is sensitive over } \mathbf{Q}\,] \le n^{-\alpha(Q)+\beta(Q)+o(1)}.$$

Proof. Let $\mathcal{E}$ be the set of potential edges of graphs in $\mathcal{G}_{\alpha,\beta}(n)$. We view elements of $\mathcal{E}$ as the variables of f (i.e. we view f as a boolean function on $\{0,1\}^{\mathcal{E}}$). For $e \in \mathcal{E}$, let $\widehat{e} \in E(P)$ be the corresponding edge of P (under the $V(P)$-coloring of the common vertex set of graphs in $\mathcal{G}_{\alpha,\beta}(n)$).

For sufficiently small constant $\delta > 0$ (which depends on P, but not n), we generate a random restriction $\rho : \mathcal{E} \to \{0,1,\star\}$ where, independently for all $e \in \mathcal{E}$,

$$\Pr_{\rho}[\,\rho(e) = \star\,] = n^{-\beta(\widehat{e})-\delta}, \qquad \Pr_{\rho}[\,\rho(e) = 1 \mid \rho(e) \neq \star\,] = n^{-\beta(\widehat{e})}. \tag{8}$$

Let H_ρ denote the P-colored graph with edge set $E(H_\rho) = \rho^{-1}(\star)$. Note that H_ρ has distribution $\mathcal{G}_{\alpha,\beta+\delta}$. In particular, $\mathrm{E}_\rho[\,\mathrm{sub}(Q,H_\rho)\,] = n^{\alpha(Q)-\beta(Q)-\delta|E(Q)|}$. We assume δ is sufficiently small so that $\alpha(Q)-\beta(Q)-\delta|E(Q)| > 0$ (here we assume $\alpha(Q)-\beta(Q) > 0$ since otherwise the lemma is trivial). Using the lower-tail version of Janson's Inequality [18], it can be shown that

$$\Pr_{\rho}[\,\mathrm{sub}(Q,H_\rho) < \tfrac{1}{2} n^{\alpha(Q)-\beta(Q)-\delta|E(Q)|}\,] = n^{-\omega(1)}. \tag{9}$$

That is, with very high probability, H_ρ contains at least half the expected number of Q-subgraphs. (Since P and $\delta > 0$ are fixed, $n^{-\omega(1)}$ is $O(n^{-c})$ for every constant $c = c(P,\delta)$ which may depend on P and δ.)

Let L_ρ denote the subgraph of H_ρ with edge set

$$E(L_\rho) = \{e \in \rho^{-1}(\star) : \text{restricted function } f{\restriction_\rho} : \{0,1\}^{\rho^{-1}(\star)} \to \{0,1\} \text{ depends}^{8} \text{ on coordinate } e\}.$$

Note that in the language of graphs predominantly used in this proof, we have $E(L_\rho) = \mathsf{Sens}(f{\restriction_\rho}, \rho^{-1}(\{1,\star\}))$.

We now use the fact that f is computed by an AC^0-circuit (in particular, a type-I AC^0-circuit). A bottom-up depth-reduction argument using Håstad's Switching Lemma [14] shows that

$$\Pr_{\rho}[\,|L_\rho| > n^{\delta}\,] = n^{-\omega(1)}. \tag{10}$$

Details of the argument (which is nearly identical to Lemma 3.7 of [27]) can be found in Appendix C.

Let $\mathcal{A} = \mathcal{A}(\rho)$ denote the event that $\mathrm{sub}(Q,H_\rho) \ge \frac{1}{2} n^{\alpha(Q)-\beta(Q)-\delta|E(Q)|}$ and $|L_\rho| \le n^{\delta}$. Note that $\Pr[\,\neg\mathcal{A}\,] = n^{-\omega(1)}$ (by (9) and (10)) and

$$\mathcal{A} \implies \frac{\mathrm{sub}(Q,L_\rho)}{\mathrm{sub}(Q,H_\rho)} \le 2n^{-\alpha(Q)+\beta(Q)+\delta|E(Q)|+\delta|V(Q)|}.$$

We now generate a pair $(\mathbf{G}',\mathbf{Q}')$ of random variables by the following two-step process.

- independently, for all $e \in \mathcal{E}$, let

$$\Pr[\,e \in E(\mathbf{G}') \mid \rho\,] = \begin{cases} \rho(e) & \text{if } \rho(e) \in \{0,1\}, \\ n^{-\beta(\widehat{e})} & \text{if } \rho(e) = \star. \end{cases} \tag{11}$$

[8] A function $g : \{0,1\}^I \to \{0,1\}$ *depends* on a coordinate $i \in I$ if there exists $x \in \{0,1\}^I$ such that $g(x) \neq g(x^{(i)})$ where $x^{(i)}$ is x with its ith coordinate flipped.

- If $\mathcal{A}(\rho)$ holds (in particular, $\mathrm{sub}(Q, H_\rho) \neq \emptyset$), we let $\mathbf{Q}'$ be a uniform random Q-subgraph of H_ρ. Otherwise, we let $\mathbf{Q}' := \perp$ ($\perp$ stands for "undefined").

We claim that $(\mathbf{G}', \mathbf{Q}')$ under the condition $\mathbf{Q}' \neq \perp$ is distributed identically with $(\mathbf{G}, \mathbf{Q})$.

Indeed, by inspecting the definitions (8) and (11), we see that $\Pr_{\mathbf{G}'}[\, e \in \mathbf{G}' \,] = n^{-\beta(\widehat{e})}$. This implies that, firstly, $\mathbf{G}' \sim G_{\alpha,\beta}(n)$ and, secondly, $\mathbf{G}'$ and H_ρ are independent. As $\mathbf{Q}'$ is a function of H_ρ, it is independent of $\mathbf{G}'$ as well. Finally, $\mathbf{Q}'$ conditioned by the event $\mathbf{Q}' \neq \perp$ is distributed identically with $\mathbf{Q}$ simply by symmetry.

This observation, along with the crucial fact $\Pr[\, \neg\mathcal{A} \,] = n^{-\omega(1)}$ mentioned above, implies that we can rephrase the inequality we are proving as

$$\Pr_{\rho, \mathbf{G}', \mathbf{Q}'}[\, f^{\cup \mathbf{G}'} \text{ is sensitive over } \mathbf{Q}' \mid \mathcal{A}(\rho) \,] \leq n^{-\alpha(Q)+\beta(Q)+o(1)}.$$

We now fix an arbitrary ρ such that $\mathcal{A}$ holds.

Since $\mathbf{Q}' \subseteq H_\rho = \{e : \rho(e) = \star\}$, it follows from definitions that if $f^{\cup \mathbf{G}'}$ is sensitive over $\mathbf{Q}'$, then $\mathbf{Q}' \subseteq L_\rho$. We now have

$$\Pr_{\mathbf{G}', \mathbf{Q}'}[\, f^{\cup \mathbf{G}'} \text{ is sensitive over } \mathbf{Q}' \,] \leq 2n^{-\alpha(Q)+\beta(Q)+\delta|E(Q)|+\delta|V(Q)|}.$$

The lemma follows, since we can choose $\delta > 0$ arbitrarily small relative to $|E(P)|$. □

Finally, the main result of this subsection (the lower bound of Theorem 3.1):

Theorem 3.11. *Suppose* C *is a type-II circuit which solves* $\textsc{Subgraph}_{\mathsf{col}}(P)$ *in the average-case on* $G_{\alpha,\beta}(P)$*. Then* C *has size* $n^{\kappa_{\alpha,\beta}(P)-o(1)}$.

Proof. For contradiction, assume C has size $\leq n^{\kappa_{\alpha,\beta}(P)-\varepsilon}$ and d alternations for constants $\varepsilon, d > 0$ (which may depend on P but not n). By Lemma 2.14, there exists a hitting set $\mathcal{H}$ for P such that $\kappa_{\alpha,\beta}(P) = \min_{Q \in \mathcal{H}} \alpha(Q) - \beta(Q)$. Note that every sub-circuit C' of C is computable by a type-I AC^0 circuit of depth d (by combining all adjacent AND and OR gates in C'). Therefore, by Lemma 3.10,

$$\begin{aligned}
\Pr_{\mathbf{G}}[\bigvee_{Q \in \mathcal{H}} \mathsf{C}'^{\cup \mathbf{G}} \text{ is sensitive over } \mathbf{Q}] &\leq \sum_{Q \in \mathcal{H}} n^{-\alpha(Q)+\beta(Q)+o(1)} \\
&\leq |\mathcal{H}| \max_{Q \in \mathcal{H}} n^{-\alpha(Q)+\beta(Q)+o(1)} \\
&= n^{-\kappa_{\alpha,\beta}(P)+o(1)} \quad (\text{since } |\mathcal{H}| \leq 2^{|E(P)|} = n^{o(1)}).
\end{aligned}$$

Taking a union bound over the $\leq n^{\kappa_{\alpha,\beta}(P)-\varepsilon}$ sub-circuits of C, we have

$$\Pr_{\mathbf{G}}[\bigvee_{\text{sub-circuits } \mathsf{C}'} \bigvee_{Q \in \mathcal{H}} \mathsf{C}'^{\cup \mathbf{G}} \text{ is sensitive over } \mathbf{Q}] \leq n^{-\varepsilon+o(1)} = o(1). \tag{12}$$

We now derive a contradiction to (12). By Lemma 3.7 (with respect to the circuit $\mathsf{C}^{\cup \mathbf{G}}$), if $\mathsf{C}^{\cup \mathbf{G}}$ is sensitive over $\mathbf{P}$, then there exist $Q \in \mathcal{H}$ and a sub-circuit C' of C such that $\mathsf{C}'^{\cup \mathbf{G}}$ is sensitive over $\mathbf{Q}$. It follows that

$$\begin{aligned}
\Pr_{\mathbf{G}}[\bigvee_{\text{sub-circuits } \mathsf{C}'} \bigvee_{Q \in \mathcal{H}} \mathsf{C}'^{\cup \mathbf{G}} \text{ is sensitive over } \mathbf{Q}] &\geq \Pr_{\mathbf{G}, \mathbf{P}}[\, \mathsf{C}^{\cup \mathbf{G}} \text{ is sensitive over } \mathbf{P} \,] \\
&= \Omega(1) \quad (\text{by Lemma 3.9}). \qquad (13)
\end{aligned}$$

Inequalities (12) and (13) give a contradiction, which completes the proof. □

4 Bounds on $\kappa_{\mathsf{col}}(P)$

In the previous section, we proved that $C_{\mathsf{col}}(P) \geq \kappa_{\mathsf{col}}(P)$, that is, $n^{\kappa_{\mathsf{col}}(P)-o(1)}$ is a lower bound on the AC^0 complexity of $\textsc{Subgraph}_{\mathsf{col}}(P)$. In Section 4.2 below we will complete the proof of Theorem 1.1 by showing that $\kappa_{\mathsf{col}}(P) \geq \Omega(tw(P)/\log tw(P))$. But, as a warm-up, let us do a simple combinatorial *upper* bound on $\kappa_{\mathsf{col}}(P)$ (and also present a useful construction in Lemma 4.1 that we will need for lower bound proofs).

4.1 Upper Bound on $\kappa_{\mathsf{col}}(P)$

We have already established that $\kappa_{\mathsf{col}}(P) \leq C_{\mathsf{col}}(P)$ (Corollary 3.2) and $C_{\mathsf{col}}(P) \leq tw(P)+1$ (Lemma 2.3(1)). By these lower and upper bounds in circuit complexity, it follows that $\kappa_{\mathsf{col}}(P) \leq tw(P)+1$. In this subsection, we give a direct proof that $\kappa_{\mathsf{col}}(P) \leq tw(P)+1$.

We need the following fact, which shows that the max in the definition $\kappa_{\mathsf{col}}(P) := \max_{(\alpha,\beta)\in\theta_{\mathsf{col}}(P)} \kappa_{\alpha,\beta}(P)$ is always achieved by some $(\alpha,\beta) \in \theta_{\mathsf{col}}(P)$ with $\alpha \equiv 1$.

Lemma 4.1. *Assume that $(\alpha,\beta) \in \theta_{\mathsf{col}}(P)$ and define $\beta' : E(P) \to [0,2]$ by the formula*

$$\beta'(\{v,w\}) := \beta(\{v,w\}) + \frac{1-\alpha(v)}{d_P(v)} + \frac{1-\alpha(w)}{d_P(w)},$$

where $d_P(v)$ is the degree of the vertex v. Then $(1,\beta') \in \theta_{\mathsf{col}}(P)$ and $|V(Q)| - \beta'(Q) \geq \alpha(Q) - \beta(Q)$ for any $Q \subseteq P$.

Proof. For all $Q \subseteq P$, we have

$$\begin{aligned}
|V(Q)| - \beta'(Q) &= |V(Q)| - \sum_{\{v,w\}\in E(Q)} \Big(\beta(\{v,w\}) + \frac{1-\alpha(v)}{d_P(v)} + \frac{1-\alpha(w)}{d_P(w)}\Big) \\
&= \sum_{v\in V(Q)} \underbrace{\Big(1 - \frac{d_Q(v)}{d_P(v)}(1-\alpha(v))\Big)}_{\geq \alpha(v)} - \sum_{\{v,w\}\in E(Q)} \beta(\{v,w\}) \\
&\geq \alpha(Q) - \beta(Q)
\end{aligned}$$

with equality when $Q = P$. □

Corollary 4.2. *For all P, there exists $\beta : E(P) \to [0,2]$ such that $(1,\beta) \in \theta_{\mathsf{col}}(P)$ and $\kappa_{\mathsf{col}}(P) = \kappa_{1,\beta}(P)$.*

Proof. Let $(\alpha,\beta) \in \theta_{\mathsf{col}}(P)$ be such that $\kappa_{\mathsf{col}}(P) = \kappa_{\alpha,\beta}(P)$. Then the element $(1,\beta') \in \theta_{\mathsf{col}}(P)$ constructed from (α,β) as in Lemma 4.1, has the desired property. □

Proposition 4.3. $\kappa_{\mathsf{col}}(P) \leq tw(P)+1$.

Proof. In fact, we will prove $\kappa_{\mathsf{col}}(P) \leq bw(P)$ where $bw(P)$ is the branch-width of P (it is known that $bw(P) \leq tw(P)+1$ by [25]). Recall that a *branch decomposition* of P is a pair (T,b) where T is a ternary tree and b is bijection from Leaves(T) to E. Each edge in T determines a partition of Leaves(T) (and hence of E) into two sets. The *width* of (T,b) is the maximum of $|V(E_1) \cap V(E_2)|$ over partitions $E = E_1 \uplus E_2$ determined by the edges of T, and the *branch-width* $bw(P)$ is the minimum possible width of a branch decomposition of P.

Suppose $bw(P) = k$. This means that there exists a branch decomposition (T, b) of width k. Fix an arbitrary root of T and let $x_1, \dots, x_t$ be a post-order traversal of nodes in T (in particular, x_t is the root). For every $i \in [t]$, let Q_i and $\overline{Q_i}$ be the subgraphs of P with $E(Q_i) := \{b(y) : y \text{ is a leaf of } T \text{ lying below } x_i\}$ and $E(\overline{Q_i}) := E(P) \setminus E(Q_i)$ and $V(Q_i) := \bigcup_{e \in E(Q_i)} e$ and $V(\overline{Q_i}) := \bigcup_{e \in E(\overline{Q_i})} e$. Note that $Q_1, \dots, Q_t$ is a union sequence of P. Since (T, b) has width $\leq k$, we have $|V(Q_i) \cap V(\overline{Q_i})| \leq k$ for all $i \in [t]$.

By Corollary 4.2, there exists $\beta : E(P) \to [0, 2]$ such that $(1, \beta) \in \theta_{\mathsf{col}}(P)$ and $\kappa_{\mathsf{col}}(P) = \kappa_{1,\beta}(P)$. For all $i \in [t]$, we have

$$\begin{aligned}
|V(Q_i)| - \beta(Q_i) &\leq |V(Q_i)| - \beta(Q_i) + |V(\overline{Q_i})| - \beta(\overline{Q_i}) && (\text{since } |V(\overline{Q_i})| \geq \beta(\overline{Q_i})) \\
&= |V(Q_i)| + |V(\overline{Q_i})| - |V(P)| && (\text{since } \beta(Q_i) + \beta(\overline{Q_i}) = \beta(P) = |V(P)|) \\
&= |V(Q_i) \cap V(\overline{Q_i})| \ \leq \ k.
\end{aligned}$$

Therefore, $\kappa_{\mathsf{col}}(P) = \kappa_{1,\beta}(P) \leq \max_{i \in [t]} |V(Q_i)| - \beta(Q_i) \leq k \leq bw(P) \leq tw(P) + 1$. □

4.2 Lower Bounds on $\kappa_{\mathsf{col}}(P)$

We give two lower bounds on $\kappa_{\mathsf{col}}(P)$. The first applies to all patterns P.

Theorem 4.4. *$\kappa_{\mathsf{col}}(P) \geq \Omega(tw(P)/\log tw(P))$.*

Together with the fact that $C_{\mathsf{col}}(P) \geq \kappa_{\mathsf{col}}(P)$ (Corollary 3.2), this completes the proof of our main theorem (Theorem 1.1).

Our proof of Theorem 4.4 uses a characterization of treewidth from Marx [20] (based on results of Feige et al [11]): for every P with $tw(P) = k$, there is a subset $W \subseteq V(P)$ of size $|W| = \Omega(k)$ and a concurrent flow on P which routes $\Omega(1/k \log k)$ flow between every pair of distinct vertices in W (Lemma 4.8). Given such a concurrent flow on P, we construct a corresponding threshold pair $(\alpha, \beta) \in \theta_{\mathsf{col}}(P)$ and show that $\kappa_{\alpha,\beta}(P)$ gives the desired bound.

We also include a lower bound on $\kappa_{\mathsf{col}}(P)$ in terms of the expansion of P (Theorem 4.9), which improves Theorem 4.4 in the case where P is a constant-degree expander.

Definition 4.5.

(i) Let $\mathsf{Paths}(P)$ denote the set of *paths* in P (i.e. subgraphs of P isomorphic to an (undirected, simple) path of length ≥ 1).

(ii) Let $\mathsf{Flows}(P)$ denote the set of *concurrent flows* on P with node-capacity 1, that is, functions $f : \mathsf{Paths}(P) \to [0, 1]$ such that for all $v \in V(P)$, $\sum_{\pi \in \mathsf{Paths}(P) \,:\, v \in V(\pi)} f(\pi) \leq 1$.

(iii) For $f \in \mathsf{Flows}(P)$ and disjoint $S, T \subseteq V(P)$, let $f(S, T)$ denote the total flow that f sends between S and T, that is,

$$f(S,T) := \sum_{\pi \in \mathsf{Paths}(P) \,:\, \pi \text{ has endpoints in } S \text{ and } T} f(\pi).$$

For two distinct vertices v, w, we let $f(v, w) := f(\{v\}, \{w\})$.

(iv) For $\pi \in \mathsf{Paths}(P)$, define $\alpha_\pi : V(P) \to [0,1]$ and $\beta_\pi : E(P) \to [0,2]$ by

$$\alpha_\pi(v) := \begin{cases} 1/2 & \text{if } v \text{ is an endpoint of } \pi, \\ 1 & \text{if } v \text{ is an interior vertex of } \pi, \\ 0 & \text{if } v \notin V(\pi), \end{cases} \qquad \beta_\pi(e) := \begin{cases} 1 & \text{if } e \in E(\pi), \\ 0 & \text{if } e \notin E(\pi). \end{cases}$$

(v) For $f \in \mathsf{Flows}(P)$, define $\alpha_f : V(P) \to [0,1]$ and $\beta_f : E(P) \to [0,2]$ by

$$\alpha_f(v) := \sum_{\pi \in \mathsf{Paths}(P)} f(\pi) \cdot \alpha_\pi(v), \qquad \beta_f(e) := \sum_{\pi \in \mathsf{Paths}(P)} f(\pi) \cdot \beta_\pi(e).$$

Lemma 4.6. *$(\alpha_f, \beta_f) \in \theta_{\mathsf{col}}(P)$ for all $f \in \mathsf{Flows}(P)$.*

Proof. Clearly, $\alpha_\pi(P) = \beta_\pi(P)$ $(= |E(\pi)|)$ and $\alpha_\pi(Q) \geq \beta_\pi(Q)$ for all $Q \subseteq P$ and $\pi \in \mathsf{Paths}(P)$. $(\alpha_f, \beta_f) \in \theta_{\mathsf{col}}(P)$ follows by convexity. $\square$

Lemma 4.7. *For all $Q \subseteq P$ and $f \in \mathsf{Flows}(P)$,*

$$\alpha_f(Q) - \beta_f(Q) \geq \frac{1}{2} f(V(Q), \overline{V(Q)}).$$

Proof. Note that $f(S,T) = \sum_{\pi \in \mathsf{Paths}(P)} f(\pi) \cdot \pi(S,T)$ where

$$\pi(S,T) := \begin{cases} 1 & \text{if } \pi \text{ has one endpoint in } S \text{ and another in } T, \\ 0 & \text{otherwise.} \end{cases}$$

Therefore, it suffices to show, for all $\pi \in \mathsf{Paths}(P)$, that

$$\alpha_\pi(Q) - \beta_\pi(Q) \geq \frac{1}{2} \pi(V(Q), \overline{V(Q)}). \tag{14}$$

If both endpoints of π belong to the same set among $V(Q)$, $\overline{V(Q)}$, then $\frac{1}{2}\pi(V(Q), \overline{V(Q)}) = 0$ while $\alpha_\pi(Q) - \beta_\pi(Q) \geq 0$ by Lemma 4.6 (since $(\alpha_\pi, \beta_\pi) \in \theta_{\mathsf{col}}(P)$); so (14) holds. On the other hand, if π has one endpoint in $V(Q)$ and another in $\overline{V(Q)}$, then $\frac{1}{2}\pi(V(Q), \overline{V(Q)}) = \frac{1}{2}$, while

$$\alpha_\pi(Q) - \beta_\pi(Q) = \frac{1}{2} |\{\text{edges of } \pi \text{ that cross between } V(Q) \text{ and } \overline{V(Q)}\}| \geq \frac{1}{2},$$

so again (14) holds. $\square$

Our lower bound on $\kappa_{\mathsf{col}}(P)$ relies on a characterization of treewidth in terms of concurrent flows:

Lemma 4.8 ([11, 20])**.** *If P has treewidth k, then there exist $W \subseteq V(P)$ with $|W| \geq 2k/3$ and $f \in \mathsf{Flows}(P)$ such that $f(v,w) \geq 1/ck \log k$ for all distinct $v, w \in W$ where $c > 0$ is a universal constant.*

Proof of Theorem 4.4. Suppose $tw(P) = k$ and fix $W \subseteq V(P)$ and $f \in \mathsf{Flows}(P)$ as in the Lemma 4.8. Let $\mathcal{H}$ be the set of subgraphs $Q \subseteq P$ such that $2k/9 \leq |W \cap V(Q)| \leq 4k/9$. Clearly $\mathcal{H}$ is a hitting set for P (i.e. every union sequence for P contains a graph in this set). For every $Q \in \mathcal{H}$, we have

$$\begin{aligned}\alpha_f(Q) - \beta_f(Q) &\geq \frac{f(W \cap V(Q), W \setminus V(Q))}{2} \\ &\geq \frac{|W \cap V(Q)| \cdot |W \setminus V(Q)|}{2ck \log k} \geq \frac{4k}{81c \log k} = \Omega\left(\frac{k}{\log k}\right).\end{aligned}$$

Therefore, $\kappa_{\mathsf{col}}(P) \geq \kappa_{\alpha_f,\beta_f}(P) = \Omega(k/\log k)$. □

Tight lower bound for expanders.

We conclude this section by giving a second lower bound on $\kappa_{\mathsf{col}}(P)$ in terms of edge expansion; this gives the optimal $\Omega(tw(P))$ lower bound in the case that P is an expander such as K_k or $G_{k,k}$. Let $\Delta(P)$ denote the maximum degree of P. For $S \subseteq V(P)$, let $e_P(S, \overline{S}) := |\{\{v,w\} \in E(P) : v \in S \text{ and } w \in V(P) \setminus S\}|$. Recall that the *edge expansion* of P is defined by

$$h(P) := \min_{S \,:\, \emptyset \subset S \subset V(P)} \frac{e_P(S, \overline{S})}{\min\{|S|, |\overline{S}|\}}.$$

Theorem 4.9. $\kappa_{\mathsf{col}}(P) \geq \dfrac{h(P)|V(P)|}{3\Delta(P)}$.

Proof. Let us apply the construction from Lemma 4.1 to the pair $(0,0)$. This gives us the function $\beta : E(P) \to [0,2]$ defined by

$$\beta(\{v,w\}) := \frac{1}{d_P(v)} + \frac{1}{d_P(w)}$$

such that $(1, \beta) \in \theta_{\mathsf{col}}(P)$.

Consider the hitting set $\mathcal{H}$ consisting of subgraphs $Q \subseteq P$ such that $\frac{1}{3}|V(P)| \leq |V(Q)| \leq \frac{2}{3}|V(P)|$. For every $Q \in \mathcal{H}$ the calculation in the proof of Lemma 4.1 gives us

$$\begin{aligned}|V(Q)| - \beta(Q) = \sum_{v \in V(Q)} \left(1 - \frac{d_Q(v)}{d_P(v)}\right) &\geq \frac{1}{\Delta(P)} \sum_{v \in V(Q)} (d_P(v) - d_Q(v)) \\ &= \frac{e_P(V(Q), \overline{V(Q)})}{\Delta(P)} \\ &\geq \frac{h(P) \min\{|V(Q)|, |V(P)| - |V(Q)|\}}{\Delta(P)} \\ &\geq \frac{h(P)|V(P)|}{3\Delta(P)}.\end{aligned}$$

Completing the proof,

$$\kappa_{\mathsf{col}}(P) \geq \kappa_{1,\beta}(P) \geq \min_{Q \in \mathcal{H}} (|V(Q)| - \beta(Q)) \geq \frac{h(P)|V(P)|}{3\Delta(P)}.$$ □

5 Minor-Monotonicity and Monotone Projections

In this section, we prove that $\kappa_{\mathsf{col}}(P)$ and $C_{\mathsf{col}}(P)$ are minor-monotone graph parameters. First, a few definitions.

Recall that a *minor* of G is any graph that can be obtained from G by a sequence of vertex deletions, edge deletions, and edge contractions. A real-valued graph parameter f is *minor-monotone* if $f(G) \le f(G')$ whenever G is a minor of G'.

Theorem 5.1. *$\kappa_{\mathsf{col}}(P)$ and $C_{\mathsf{col}}(P)$ are minor-monotone.*

The algorithmic problem $\textsc{Subgraph}_{\mathsf{col}}(P)$ was defined in Section 2.3 in such a way that the coloring $\chi : G \to P$ is a part of its input. We first observe that the parameter $C_{\mathsf{col}}(P)$ does not change if we consider instead its more structured version $\textsc{Subgraph}_{\mathsf{col},n}(P)$ in which (cf. Definition 2.8) we demand that the target graph G has the vertex set $V(P) \times [n]$, and χ is the projection onto the first coordinate. An easy AC^0-reduction from $\textsc{Subgraph}_{\mathsf{col}}(P)$ to $\textsc{Subgraph}_{\mathsf{col},n}(P)$ works as follows. Assume that we are given an input (G, χ) to the problem $\textsc{Subgraph}_{\mathsf{col}}(P)$, and assume w.l.o.g. that $V(G) = [n]$. We map it to the pair (G', χ'), where $V(G') := V(P) \times [n]$, χ' is the projection onto the first coordinate, and $E(G')$ is defined as follows: $E(G') := \{\{(v, \chi(v)), (w, \chi(w))\} : \{v, w\} \in E(G)\}$. (Thus, all vertices (v, i) with $i \neq \chi(v)$ remain isolated.) Hence Theorem 5.1 readily follows from the following lemma.

Lemma 5.2. *Suppose P is a minor of P'. Then*

1. *for every $(\alpha, \beta) \in \theta_{\mathsf{col}}(P)$, there exists $(\alpha', \beta') \in \theta_{\mathsf{col}}(P')$ such that $\kappa_{\alpha,\beta}(P) \le \kappa_{\alpha',\beta'}(P')$,*
2. *$\textsc{Subgraph}_{\mathsf{col},n}(P) \ \le_{\mathrm{mp}}\ \textsc{Subgraph}_{\mathsf{col},n}(P')$.*

Proof. It suffices to show that the lemma holds in the two cases where P is a subgraph of P', and where P is obtained from P' by contracting a single edge $\{x, y\}$ where x, y have no common neighbor (Otherwise, we perform necessary edge deletions before contraction).

Subgraph Case. Suppose P is a subgraph of P'.

For (1): Consider any $(\alpha, \beta) \in \theta_{\mathsf{col}}(P)$. Define $\alpha' : V(P') \to [0, 1]$ and $\beta' : E(P') \to [0, 2]$ by

$$\alpha'(v) := \begin{cases} \alpha(v) & \text{if } v \in V(P), \\ 0 & \text{otherwise,} \end{cases} \qquad \beta'(e) := \begin{cases} \beta(e) & \text{if } e \in E(P), \\ 0 & \text{otherwise.} \end{cases}$$

It is easily seen that $(\alpha', \beta') \in \theta_{\mathsf{col}}(P')$ and $\kappa_{\alpha,\beta}(P) = \kappa_{\alpha',\beta'}(P')$.

For (2): The monotone projection p is defined as follows:

$$p(\{(v, i), (w, j)\}) := \begin{cases} \{(v, i), (w, j)\} & \text{if } \{v, w\} \in E(P), \\ 1 & \text{if } \{v, w\} \in E(P') \setminus E(P). \end{cases}$$

Thus, p^* takes an input G to the problem $\textsc{Subgraph}_{\mathsf{col},n}(P)$ and converts it into an input G' to $\textsc{Subgraph}_{\mathsf{col},n}(P')$ by filling in complete bipartite graphs between $\{v\} \times [n]$ and $\{w\} \times [n]$ for all new edges $\{v, w\} \in E(P') \setminus E(P)$.

Contraction Case. Now suppose P is obtained from P' by contracting a single edge $\{x, y\}$ where x, y have no common neighbor. Let z label the contracted vertex in P, so that $V(P) \setminus V(P') = \{z\}$

and $V(P') \setminus V(P) = \{x,y\}$. Let $\rho : V(P') \to V(P)$ be the function $x, y \mapsto z$ and $v \mapsto v$ for all $v \in V(P') \setminus \{x,y\}$. For $e = \{v,w\} \in E(P') \setminus \{x,y\}$, let $\rho(e) := \{\rho(v), \rho(w)\} \in E(P)$ ($\rho(\{x,y\})$ is undefined).

For (1): Consider any $(\alpha, \beta) \in \theta_{\mathsf{col}}(P)$. Define $\alpha' : V(P') \to [0,1]$ and $\beta' : E(P') \to [0,2]$ by

$$\alpha'(v) := \alpha(\rho(v)), \qquad \beta'(e) := \begin{cases} \alpha(z) & \text{if } e = \{x,y\}, \\ \beta(\rho(e)) & \text{otherwise.} \end{cases}$$

We now check that $(\alpha', \beta') \in \theta_{\mathsf{col}}(P')$ and $\kappa_{\alpha',\beta'}(P') \geq \kappa_{\alpha,\beta}(P)$. For that consider the mapping $\widehat{\rho} : Q' \to \rho(Q' \setminus \{\{x,y\}\})$ that takes subgraphs of P' to subgraphs of P. It is easy to see that $\alpha'(Q') - \beta'(Q') \geq \alpha(\widehat{\rho}(Q')) - \beta(\widehat{\rho}(Q'))$, and that this is tight for $Q' = P'$: in the only non-trivial case $\{x,y\} \in E(Q')$ we have $\alpha(\widehat{\rho}(Q')) = \alpha'(Q') - \alpha(z)$ and $\beta(\widehat{\rho}(Q')) = \beta'(Q') - \alpha(z)$. This proves the first claim $(\alpha', \beta') \in \theta_{\mathsf{col}}(P')$. To see that $\kappa_{\alpha,\beta}(P) \leq \kappa_{\alpha',\beta'}(P')$, it suffices to observe that $\widehat{\rho}$ takes union sequences for P' into union sequences for P and thus $\widehat{\rho}^{-1}$ maps hitting sets for P into hitting sets for P'.

For (2): This time the monotone projection p is defined by

$$p(\{(v,i),(w,j)\}) := \begin{cases} 1 & \text{if } \{v,w\} = \{x,y\} \text{ and } i = j, \\ 0 & \text{if } \{v,w\} = \{x,y\} \text{ and } i \neq j, \\ \{(\rho(v),i),(\rho(w),j)\} & \text{otherwise.} \end{cases}$$

(That is, p^* duplicates $\{z\} \times [n]$ into two sets $\{x\} \times [n]$, $\{y\} \times [n]$ and then plants a perfect matching between twins.) This p is clearly a monotone projection from $\textsc{Subgraph}_{\mathsf{col},n}(P)$ to $\textsc{Subgraph}_{\mathsf{col},n}(P')$. □

5.1 Negative results in the uncolored setting

In the colored setting, we have seen that $\textsc{Subgraph}_{\mathsf{col}}(P)$ is minor-monotone via linear-size monotone projections. Highlighting a difference between the uncolored and colored settings, we now show that there is *no* monotone projection whatsoever that reduces $\textsc{Subgraph}(M_3)$ to $\textsc{Subgraph}(P_3 + M_2)$ (where P_3 is a path on 3 vertices and M_k is a matching with k edges). While it remains an open problem whether $C(P)$ is (even approximately) minor-monotone under general AC^0 reductions, this result strongly suggests that the colorful version of the subgraph isomorphism problem is much better structured and well-behaved than the standard (uncolored) one.

We begin with some properties of $P_3 + M_2$-free graphs.

Lemma 5.3. *Every $P_3 + M_2$-free graph G satisfies one of the following conditions:*

(i) *G has $\leq C$ edges for some absolute constant C,*

(ii) *G is a matching,*

(iii) *G contains a vertex of degree ≥ 6.*

Note: Lemma 5.3 is true with any integer replacing 6 in (iii), for an appropriate constant C in (i). The choice of 6 is simply convenient in the proof of Theorem 5.6 later on.

Proof. Assume G is $P_3 + M_2$-free, not a matching, and has maximum degree ≤ 5. We will show that G has $O(1)$ edges. Since G is not a matching, it contains a vertex u of degree ≥ 2. Since G has maximum degree ≤ 5, there is a constant C' such that if G has $> C'$ non-isolated vertices, then it contains vertices v, w such that any two of u, v, w have distance ≥ 3; in that case, we would have $P_3 + M_2$-subgraph of G by taking any two edges containing u plus any two edges containing v and w respectively. Therefore, G has $\leq C'$ non-isolated vertices. It follows that G has $\leq 5C'/2$ edges. ☐

Lemma 5.4. *If G is $P_3 + M_2$-free and contains an M_4-subgraph, then G is a matching.*

Proof. Suppose G contains an M_4-subgraph H, but G is not a matching. We will show that G contains a $P_3 + M_2$-subgraph. Since G is not a matching, it contains a P_3-subgraph K. If K is vertex-disjoint from H, then K plus any two edges from H is a $P_3 + M_2$-subgraph of G. Now assume that K intersects a vertex in H. Then there is a P_3-subgraph K' which contains an edge in H. This K' is vertex-disjoint from at least two edges in H; then K' plus these two edges is a $P_3 + M_2$-subgraph of G. ☐

Lemma 5.5. *Suppose G contains a $P_3 + M_2$-subgraph and a vertex u of degree ≥ 6. Then G contains a $P_3 + M_2$-subgraph in which u is the degree-2 vertex.*

Proof. Let H be any $P_3 + M_2$-subgraph of G. H contains an M_2-subgraph H' which does not include the vertex u. Since u has degree ≥ 6, it has two distinct neighbors v and w such that $\{u, v, w\} \cap V(H') = \emptyset$. Then H' plus edges $\{u, v\}$ and $\{u, w\}$ is a $P_3 + M_2$-subgraph of G in which u is the degree-2 vertex. ☐

Now the main result of this subsection:

Theorem 5.6. $\textsc{Subgraph}(M_3)$ *is not a monotone projection of* $\textsc{Subgraph}(P_3 + M_2)$.

Proof. Toward a contradiction, assume there exists a monotone projection $p : \binom{[N]}{2} \to \binom{[n]}{2} \cup \{0, 1\}$ from $\textsc{Subgraph}(M_3)$ on n-vertex graphs to $\textsc{Subgraph}(P_3 + M_2)$ on N-vertex graphs for some $n, N \in \mathbb{N}$ where $n \geq C + 2$ with C the constant from Lemma 5.3. That is, for every graph G with vertex set $[n]$, we have

$$G \text{ contains an } M_3\text{-subgraph} \Leftrightarrow p^*(G) \text{ contains a } P_3 + M_2\text{-subgraph}$$

where $p^*(G)$ is the graph with edge set $p^{-1}(E(G) \cup \{1\})$.

We proceed by showing a sequence of claims. The first claim establishes that the monotone projection p^* depends on all variables (i.e. on all edges in $\binom{[n]}{2}$).

Claim 5.7. $p^{-1}(e)$ is non-empty for all $e \in \binom{[n]}{2}$.

▶ Let G be an n-vertex graph consisting of three disjoint edges, one of which is e. Then G contains an M_3-subgraph, so $p^*(G)$ contains a $P_3 + M_2$-subgraph. However, $G - e$ is M_3-free, so $p^*(G - e)$ is $P_3 + M_2$-free. We conclude that $p^{-1}(e)$ is nonempty, since $p^*(G) = p^*(G - e) \cup p^{-1}(e)$. $\square_{\text{claim}}$

For $a \in [n]$, let S_a denote the n-vertex star centered at a (with edge set $\{e \in \binom{[n]}{2} : a \in e\}$). Let $F_a := p^{-1}(S_a)$ (so that $p^*(S_a)$ is the disjoint union of F_a and $p^{-1}(1)$). Over the next few claims, we will show that F_a are stars of degree ≥ 6 with distinct centers. Claim 5.7 implies that F_a contains at least $n - 1$ ($> C$) edges.

Since S_a is M_3-free, $p^*(S_a)$ is $P_3 + M_2$-free, hence F_a is $P_3 + M_2$-free. By Lemma 5.3, it follows that either F_a is a matching or F_a contains a vertex of degree ≥ 6. The next claim eliminates the first possibility.

Claim 5.8. For every $a \in [n]$, F_a is not a matching.

► For contradiction, assume F_a is a matching for some $a \in [n]$. Consider any $b \in [n]$. Note that $S_a \cup S_b$ is M_3-free, hence $p^*(S_a \cup S_b)$ is $P_3 + M_2$-free. Since $p^*(S_a \cup S_b) \supseteq F_a \cup F_b$ and F_a contains a M_4-subgraph, Lemma 5.4 implies that $F_a \cup F_b$ is a matching. This argument shows that $F_a \cup F_b$ is a matching for *all* $a, b \in [n]$. Therefore, the entire pre-image $p^{-1}(K_n)$ is a matching, where K_n is the complete graph on vertices $[n]$.

Since K_n contains an M_3-subgraph, $p^*(K_n)$ $(= p^{-1}(K_n) \cup p^{-1}(1))$ contains a $P_3 + M_2$-subgraph. It follows that either $p^{-1}(1)$ contains a path of length 2, or $p^{-1}(1)$ contains an edge with an endpoint in $V(p^{-1}(K_n))$. In both cases we get a contradiction, as it follows that $p^*(S_c)$ contains a $P_3 + M_2$-subgraph for some $c \in [n]$, even though S_c is M_3-free. (If $p^{-1}(1)$ contains a P_3-subgraph, then any $c \in [n]$ will do; if $p^{-1}(1)$ contains an edge with an endpoint $v \in V(p^{-1}(K_n))$, then any $c \in [n]$ with $v \in V(F_c)$ will do.) $\square_{\text{claim}}$

For all $a \in [n]$, we have established that F_a is $P_3 + M_2$-free, has $> C$ edges and is not a matching. By Lemma 5.3, we conclude that F_a contains at least one vertex of degree ≥ 6. Let us now fix a function $z : [n] \to [N]$ such that $z(a)$ is a vertex of degree ≥ 6 in F_a for all $a \in [n]$.

Claim 5.9. z is (≤ 2)-to-1.

► For contradiction, assume there exist distinct $a, b, c \in [n]$ such that $v := z(a) = z(b) = z(c)$. By Lemma 5.5, $p^*(S_a \cup S_b \cup S_c)$ contains a $P_3 + M_2$-subgraph in which v is the degree-2 vertex. Let $e, f \in \binom{[N]}{2}$ be the two edges in this subgraph which are not adjacent to v. Without loss of generality, $\{e, f\} \subseteq p^*(S_a \cup S_b)$. Since v has degree ≥ 6 in $p^*(S_a \cup S_b)$, we can find a different path of length 2 through v which is vertex-disjoint from edges e and f. Therefore, $p^*(S_a \cup S_b)$ contains a $P_3 + M_2$-subgraph. Since $S_a \cup S_b$ is M_3-free, this contradicts our assumption about p. $\square_{\text{claim}}$

Claim 5.10. F_a is a star with center $z(a)$ for all $a \in [n]$.

► For contradiction, assume F_a is not a star with center $z(a)$. Then F_a contains an edge e with $z(a) \notin e$. Since z is (≤ 2)-to-1, there exists $b \in [n]$ such that $z(b) \notin e$. We may find a $P_3 + M_2$-subgraph within $F_a \cup F_b$ by taking e together with a disjoint path of length 2 through $z(a)$ and a disjoint edge containing $z(b)$. This contradicts the fact that $p^*(S_a \cup S_b)$ is $P_3 + M_2$-free. $\square_{\text{claim}}$

Claim 5.11. z is 1-to-1.

► For contradiction, assume $v := z(a) = z(b)$ for some $a \neq b$. Let $c \in [n] \setminus \{a, b\}$. Then $z(c) \neq v$ and $p^*(S_a \cup S_b \cup S_c) = F_a \cup F_b \cup F_c \cup p^{-1}(1)$ contains a $P_3 + M_2$-subgraph H. We may assume that H contains edges $\{v, u\} \in E(F_a) \setminus E(F_b)$ and $\{v, w\} \in E(F_b) \setminus E(F_a)$ since otherwise H would be a subgraph of either $p^*(S_a \cup S_c)$ or $p^*(S_b \cup S_c)$ contradicting $P_3 + M_2$-freeness of these graphs. Note that $u \neq w$. Since v has degree ≥ 6 in F_a, we can find an edge $\{v, w'\} \in E(F_a)$ such that $w \notin V(H)$. Let H' be the graph obtained by substituting the edge $\{v, w'\}$ for $\{v, w\}$. Then H' is a $P_3 + M_2$-subgraph of $p^*(S_a \cup S_c)$, which is again a contradiction. $\square_{\text{claim}}$

At this point, we have established that graphs F_a $(a \in [n])$ are stars of degree ≥ 6 with distinct centers.

Claim 5.12. $|p^{-1}(e)| = 1$ for all $e \in \binom{[n]}{2}$.

▶ Suppose $e = \{a, b\}$. Since F_a and F_b are stars with different centers and $p^{-1}(e) \subseteq F_a \cap F_b$, there can be only one possibility for $p^{-1}(e)$. Since $p^{-1}(e)$ is nonempty by Claim 5.7, it follows that $|p^{-1}(e)| = 1$. $\square_{\text{claim}}$

Claim 5.13. $p^{-1}(1)$ is nonempty.

▶ Let G be any copy of M_3 (i.e. any three disjoint edges) among n-vertices. Then $p^{-1}(G)$ has only three edges by Claim 5.12. Since $p^*(G) = p^{-1}(G) \cup p^{-1}(1)$ has $P_3 + M_2$-subgraph, it contains at least 4 edges. Therefore, $p^{-1}(1)$ is nonempty. $\square_{\text{claim}}$

Fix any edge e in $p^{-1}(1)$ and any $a \neq b \in [n]$ such that $z(a), z(b) \notin e$. Then $p^*(S_a \cup S_b)$ contains a $P_3 + M_2$-subgraph, even though $S_a \cup S_b$ is M_3-free. This, finally, is the contradiction which completes the proof of Theorem 1.2. $\square$

6 Conclusion

With the results of this paper, the state of knowledge on the average/worst-case AC^0 complexity of the uncolored/colorful P-subgraph isomorphism problem now stands as

$$\begin{array}{c}\Omega(\frac{tw(P)}{\log tw(P)}) \leq \kappa_{\mathsf{col}}(P) \leq C_{\mathsf{col}}(P) \leq tw(P) + 1 \\ \veeeq \\ C(P) \\ \veeeq \\ \kappa(P) \leq C_{\mathsf{ave}}(P) \leq 2\kappa(P) + O(1).\end{array}$$

We have examples showing that the gap between $C_{\mathsf{ave}}(P)$ and $C(P)$ (i.e. the average-case vs. worst-case AC^0 complexity of $\textsc{Subgraph}(P)$) can be arbitrarily large (see Remark 2.7). We do not know of any gap between $C(P)$ and $C_{\mathsf{col}}(P)$. Equivalently, we can ask whether $C(P)$ is bounded from below by any function of $tw(P)$. Restating Question 1 from the introduction:

Question 1. *Is it possible to give general lower bounds on the worst-case* AC^0 *complexity of* $\textsc{Subgraph}(P)$ *(uncolored P-subgraph isomorphism) in terms of the treewidth of P only?*

When P is a core, we know that $C(P) = C_{\mathsf{col}}(P) = \widetilde{\Theta}(tw(P))$. At the opposite end of the spectrum, Question 1 is wide open for bipartite patterns P.

The next two questions seek to improve the parameters in our main results.

Question 2. *Can the upper bound $C_{\mathsf{ave}}(P) \leq 2\kappa(P) + O(1)$ of Theorem 1.2 be improved to $\kappa(P) + O(1)$?*

Question 3. *Can the* $\log tw(P)$ *factor be eliminated from our lower bounds on $\kappa_{\mathsf{col}}(P)$ (Theorem 1.1) or at least $C_{\mathsf{col}}(P)$?*

We are able to answer Question 3 affirmatively in the special case where P is a constant-degree expander (Theorem 4.9).

Another question raised by this work is whether the AC^0 complexity of $\textsc{Subgraph}(P)$ is monotone with respect to minors or subgraphs. In contrast to the colorful setting, we showed that monotone projections (the simplest form of reduction) fail to give any reduction whatsoever from $\textsc{Subgraph}(Q)$ to $\textsc{Subgraph}(P)$, even when Q is only a subgraph of P.

Question 4. *Is $C(P)$ minor-monotone or monotone under subgraphs?*

More modestly, if Q is a minor (or subgraph) of P, is there a reduction from $\textsc{Subgraph}(Q)$ to $\textsc{Subgraph}(P)$ by AC^0-circuits of size $O(n^c)$ for a constant c independent of P and Q? That would imply $C(Q) \le O(C(P))$; currently we do not know if $C(Q)$ can be bounded by *any* function in $C(P)$.

Finally, it would be interesting to investigate the relationship between $\kappa_{\mathsf{col}}(P)$ and the complexity of $\textsc{Subgraph}_{\mathsf{col}}(P)$ beyond AC^0. In particular, we recall the result of Marx [20] that $\textsc{Subgraph}_{\mathsf{col}}(P)$ has no $n^{o(tw(P)/\log tw(P))}$-time algorithm unless the Exponential Time Hypothesis (ETH) fails. Follow-up work of Alon and Marx [1] looked at the question of removing the $\log tw(P)$ factor loss in the exponent of this result (toward the goal of showing that $n^{\Theta(tw(P))}$ is the true complexity of $\textsc{Subgraph}_{\mathsf{col}}(P)$, at least assuming the ETH). Alon and Marx specifically identified constant-degree expanders as a case where "substantially different methods" are needed to eliminate the $\log tw(P)$ factor loss incurred by the reduction of [20]. In light of our lower bounds $C_{\mathsf{col}}(P) \ge \kappa_{\mathsf{col}}(P) = \Omega(|V(P)|)$ when P is a constant-degree expander, it becomes interesting to ask:

Question 5. *Can it be shown that* $\textsc{Subgraph}_{\mathsf{col}}(P)$ *has no $n^{o(\kappa_{\mathsf{col}}(P))}$-time algorithm unless the ETH fails?*

Acknowledgements

We thank the anonymous reviewers of the conference version of this paper for many helpful comments.

References

[1] Noga Alon and Dániel Marx. Sparse balanced partitions and the complexity of subgraph problems. *SIAM Journal on Discrete Mathematics* 25 (2), 631-644, 2011.

[2] Noga Alon and Joel H. Spencer. *The Probabilistic Method.* John Wiley & Sons, Inc., 2004.

[3] Noga Alon, Raphael Yuster and Uri Zwick. Color-coding. *J. ACM,* 42(4), 844-856, 1995.

[4] Kazuyuki Amano. k-Subgraph isomorphism on AC^0 circuits. *Computational Complexity,* 19(2):183-210, 2010.

[5] Andrew D. Barbour, Lars Holst and Svante Janson. *Poisson Approximation.* Oxford University Press, Oxford, UK, 1992.

[6] Hans L. Bodleander. Discovering treewidth. In *Proceedings of the 31st International Conference on Current Trends in theory and Practice of Computer Science*, volume 3381 of *Lecture Notes in Computer Science*, pages 1–16. Springer, 2005.

[7] Béla Bollobás and John C. Wierman. Subgraph counts and containment probabilities of balanced and unbalanced subgraphs in a large random graph. *Annals of the New York Academy of Sciences,* 576: 63-70, 1989.

[8] Chandra Chekuri and Julia Chuzhoy. Polynomial bounds for the grid-minor theorem. In Proceedings of the 46th Annual Symposium on the Theory of Computing (STOC), pages 60-69, 2014.

[9] David Eppstein. Subgraph isomorphism in planar graphs and related problems. *Journal of Graph Algorithms and Applications*, 3(3):1–27, 1999.

[10] David Eppstein. Diameter and treewidth in minor-closed graph families. *Algorithmica*, 27:275–291, 2000.

[11] Uriel Feige, Mohammadtaghi Hajiaghayi and James R. Lee. Improved approximation algorithms for minimum weight vertex separators. *SIAM Journal on Computing*, 38(2), 629-657, 2008.

[12] Martin Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *Journal of the ACM*, 54(1):1-24, 2007.

[13] Martin Grohe and Dániel Marx. On tree width, bramble size, and expansion. *J. Comb. Theory, Ser. B,* 99(1):218-228, 2009.

[14] Johan Håstad. *Computational limitations of small-depth circuits.* MIT press, 1987.

[15] Johan Håstad, Ingo Wegener, Norbert Wurm and Sang-Zin Yi. Optimal depth, very small size circuit for symmetric functions in AC^0. *Information and Computation,* 108(2):200-211, 1994.

[16] Svante Janson, Poisson approximation for large deviations. *Random Structures and Algorithms,* 1(2):221-230, 1990.

[17] Svante Janson, Coupling and Poisson approximation. *Acta Appl. Math.* 34, 7-15, 1994.

[18] Svante Janson, Tomasz Luczak and Andrzej Rucinski. *Random Graphs.* Wiley-Interscience, 2000.

[19] Stasys Jukna. *Boolean Function Complexity: Advances and Frontiers.* Vol. 27. Springer-verlag Berlin Heidelberg, 2012.

[20] Dániel Marx. Can you beat treewidth?. In *Proc. 48th IEEE Symposium on Foundations of Computer Science*, 169–179, 2007.

[21] Dániel Marx and Michał Pilipczuk. Everything you always wanted to know about the parameterized complexity of Subgraph Isomorphism (but were afraid to ask). In *Proc. 31st International Symposium on Theoretical Aspects of Computer Science*, 542-553, 2014.

[22] Kotaro Nakagawa and Osamu Watanabe. Gap between two combinatorial measures for constant depth circuit complexity of subgraph isomorphism. Technical Report, Tokyo Institute of Technology, 2011.

[23] Jaroslav Nešetřil and Patrice Ossona de Mendez. Linear time low tree-width partitions and algorithmic consequences. In *Proc.* 38*th ACM Symposium on the Theory of Computing*, 391–400, 2006.

[24] Jürgen Plehn and Bernd Voigt. Finding minimally weighted subgraphs. *Graph-Theoretic Concepts in Computer Science*. Springer Berlin Heidelberg, 1991.

[25] Neil Robertson and Paul D. Seymour. Graph minors X. Obstructions to tree-decomposition. *Journal of Combinatorial Theory* 52(2): 153-190, 1991.

[26] Benjamin Rossman. On the constant-depth complexity of k-clique. *Proc. 40th ACM Symposium on Theory of Computing,* 721-730, 2008.

[27] Benjamin Rossman. Average-case complexity of detecting cliques. Ph.D. thesis, MIT, 2010.

[28] Andrew Yao. Probabilistic computations: Toward a unified measure of complexity. *Proc. 18th IEEE Symposium on Foundations of Computer Science,* 222-227, 1977.

A Proof of Lemma 2.10

Fix a pattern P and a nontrivial threshold pair $(\alpha,\beta) \in \theta_{\mathsf{col}}(P)$. We can assume w.l.o.g. that $\beta(e) > 0$ for all $e \in E(G)$ (as the edges with $\beta(e) = 0$ can be removed). Following the approach of Bollobás and Wierman [7], we fix a chain of (necessarily induced) subgraphs

$$\emptyset = Q_0 \subset Q_1 \subset \cdots \subset Q_{t-1} \subset Q_t = P$$

satisfying

- $\alpha(Q_i) = \beta(Q_i)$ for all $0 \le i \le t$, and
- $\alpha(R) > \beta(R)$ for all $1 \le i \le t$ and $Q_{i-1} \subset R \subset Q_i$.

Call such a sequence $(Q_0,\dots,Q_t)$ an (α,β)-*grading* of G. Clearly, at least one (α,β)-grading exists. Note that $1 \le t \le |E(P)|$, since (α,β) is nontrivial. (It is known that t is the same for (α,β)-gradings; however, we will not use this fact.)

Let $\mathbf{G} := G_{\alpha,\beta}(P)$. For $0 \le i \le t$, define random variable $\mathbf{X}_i$ as the number of Q_i-subgraphs in $\mathbf{G}$. Obviously, $\mathbf{X}_0 = 1$ (with probability 1). For $1 \le i \le t$, let $\mathcal{L}(\mathbf{X}_i \mid \mathbf{X}_{i-1} = 1)$ denote the distribution of $\mathbf{X}_i$ conditioned on the event $\mathbf{X}_{i-1} = 1$. We prove Lemma 2.10 by showing the following

Lemma A.1. *For all* $1 \le i \le t$*,* $\mathcal{L}(\mathbf{X}_i \mid \mathbf{X}_{i-1} = 1)$ *is asymptotically the Poisson distribution* $\mathrm{Po}(1)$*. In particular,*

$$\lim_{n\to\infty} \Pr[\,\mathbf{X}_i = 0 \mid \mathbf{X}_{i-1} = 1\,] = \lim_{n\to\infty} \Pr[\,\mathbf{X}_i = 1 \mid \mathbf{X}_{i-1} = 1\,] = \frac{1}{e}.$$

The first inequality of Lemma 2.10 follows immediately, as we have

$$\begin{aligned}\liminf_{n\to\infty} \Pr[\,\mathbf{G} \text{ has no } P\text{-subgraph}\,] &\ge \liminf_{n\to\infty} \Pr[\,\mathbf{G} \text{ has no } Q_1\text{-subgraph}\,]\\ &= \liminf_{n\to\infty} \Pr[\,\mathbf{X}_1 = 0 \mid \mathbf{X}_0 = 1\,]\\ &= \frac{1}{e}.\end{aligned}$$

For the second inequality, we have

$$\begin{aligned}\liminf_{n\to\infty} \Pr[\,\mathbf{G} \text{ has a unique } P\text{-subgraph}\,] &= \liminf_{n\to\infty} \Pr[\,\mathbf{X}_t = 1\,]\\ &\geq \liminf_{n\to\infty} \Pr[\,\mathbf{X}_0 = \cdots = \mathbf{X}_t = 1\,]\\ &= \liminf_{n\to\infty} \prod_{1\le i\le t} \Pr[\,\mathbf{X}_i = 1 \mid \mathbf{X}_{i-1} = 1\,]\\ &= \frac{1}{e^t} \;\geq\; \frac{1}{e^{|E(P)|}}.\end{aligned}$$

In the remainder of this appendix we give the proof of Lemma A.1. We will use the following result on Poisson approximation. Before stating it, recall that the *total variation distance* $d_{TV}(\mathbf{X}, \mathbf{Y})$ between two random variables $\mathbf{X}$ and $\mathbf{Y}$ with values in the same set (in particular, real-valued variables) is given by

$$d_{TV}(\mathbf{X}, \mathbf{Y}) := \sup_A |\Pr[\,\mathbf{X} \in A\,] - \Pr[\,\mathbf{Y} \in A\,]|.$$

$\{0,1\}$-valued random variables $\mathbf{I}_1, \ldots, \mathbf{I}_m$ are *positively related* if for any given $i \in [m]$ one can find $\{0,1\}$-valued random variables $\mathbf{J}_{ji}$ $(j \neq i)$ such that $\mathbf{J}_{ji} \geq \mathbf{I}_j$ and this tuple is distributed identically with the tuple $\mathbf{I}_j$ $(j \neq i)$ conditioned by the event $\mathbf{I}_i = 1$.

Lemma A.2 (Theorem 6.24 in [18])**.** *Suppose $\mathbf{I}_1, \ldots, \mathbf{I}_m$ are positively related $\{0,1\}$-valued random variables, and let $\mathbf{k} := \sum_i \mathbf{I}_i$. Then*

$$d_{TV}(\mathcal{L}(\mathbf{k}), \mathrm{Po}(\mathrm{E}[\mathbf{k}])) \leq \frac{\mathrm{Var}[\mathbf{k}]}{\mathrm{E}[\mathbf{k}]} - 1 + 2 \max_i \mathrm{E}[\mathbf{I}_i].$$

Proof of Lemma A.1. Fix $i \in \{1, \ldots, t\}$ and let $Q := Q_i$ and $Q' := Q_{i-1}$ and $\mathbf{X} := \mathbf{X}_i$ and $\mathbf{X}' := \mathbf{X}_{i-1}$. To show that $\mathcal{L}(\mathbf{X} \mid \mathbf{X}' = 1)$ is asymptotically $\mathrm{Po}(1)$, we would like to sample $\mathbf{G}$ conditioned on $\mathbf{X}' = 1$ (i.e. the event that $\mathbf{G}$ contains a unique Q'-subgraph). However, it will be convenient to condition on the entire $V(Q')$-colored part of $\mathbf{G}$ (i.e. the induced subgraph of G on the vertices which map to $V(Q')$ under the vertex-coloring of $\mathbf{G}$). We shall therefore fix an arbitrary $V(Q')$-colored graph G' such that

- G' equals the Q'-colored part of G for some G in the support of $\mathbf{G}$, and
- G' contains a unique Q'-subgraph, which we will denote by H'.

We denote by $\mathbf{G}|G'$ the random graph $\mathbf{G}$ conditioned on the event that the Q'-colored part of $\mathbf{G}$ equals G'. Note that $\mathbf{G}|G'$ is a product distribution on the unrestricted edges.

Let $\mathcal{Q}$ $(= \mathcal{Q}(H'))$ be the set of potential Q-subgraphs which extend H'. For each $H \in \mathcal{Q}$, let $\mathbf{I}_H$ be the indicator variable for the event that $\mathbf{G}|G'$ contains H. These random variables are positively related: just let $\mathbf{J}_{H'H}$ be the characteristic function of the event that $\mathbf{G}$ contains $E(H') \setminus E(H)$. Let $\mathbf{k} := \sum_{H \in \mathcal{Q}} \mathbf{I}_H$. We will show that $\mathcal{L}(\mathbf{k})$ is asymptotically $\mathrm{Po}(1)$ using Lemma A.2. Since the event $\{\mathbf{X}' = 1\}$ is the disjoint union of events $\{G' \text{ is the } Q'\text{-colored part of } \mathbf{G}\}$ over all G', it follows that $\mathcal{L}(\mathbf{X} \mid \mathbf{X}' = 1)$ is asymptotically $\mathrm{Po}(1)$ by the convexity of d_{TV}.

We will now calculate the expectation of $\mathbf{k}$. First, we have

$$|\mathcal{Q}| = \prod_{v \in V(Q) \setminus V(Q')} \lfloor n^{\alpha(v)} \rfloor = (1 - o(1)) n^{\alpha(Q) - \alpha(Q')}.$$

For each $H \in \mathcal{Q}$, we have

$$\mathrm{E}[\mathbf{I}_H] = n^{-\beta(Q)+\beta(Q')}.$$

(Note for the record that this is $o(1)$ since $\beta(Q') < \beta(Q)$ by the fact that $Q_0, \dots, Q_t$ is an (α, β)-grading.) Therefore,

$$\mathrm{E}[\mathbf{k}] = (1 - o(1))n^{\alpha(Q)-\alpha(Q')}n^{-\beta(Q)+\beta(Q')} = 1 - o(1),$$

using the fact that $\alpha(Q) = \beta(Q)$ and $\alpha(Q') = \beta(Q')$.

We next calculate $\mathrm{Var}[\mathbf{k}]$. For $H, K \in \mathcal{Q}$, let $U := \chi(V(H) \cap V(K))$ be the set of P-colors of vertices in the intersection of $V(H)$ and $V(K)$. Note that

$$V(Q') \subseteq U \subseteq V(Q).$$

Thus,

$$\mathrm{E}[\,\mathbf{I}_H \mathbf{I}_K\,] = n^{-2\beta(Q)+\beta(Q')+\beta(U)}$$

where $\beta(U) := \sum_{e \in E(P) \cap \binom{U}{2}} \beta(e)$. For all $V(Q') \subseteq U \subseteq V(Q)$, we have

$$\#\{(H, K) \in \mathcal{Q} \times \mathcal{Q} : \chi(V(H) \cap V(K)) = U\} = (1 - o(1))n^{2\alpha(Q)-\alpha(Q')-\alpha(U)}.$$

Therefore,

$$\begin{aligned}
\mathrm{Var}[\mathbf{k}] &= \sum_{H,K \in \mathcal{Q}\,:\,H \neq K} \mathrm{E}[\,\mathbf{I}_H \mathbf{I}_K\,] \\
&= \sum_{U\,:\,V(Q') \subseteq U \subset V(Q)} (1 - o(1))n^{2\alpha(Q)-\alpha(Q')-\alpha(U)}n^{-2\beta(Q)+\beta(Q')+\beta(U)} \\
&= \sum_{U\,:\,V(Q') \subseteq U \subset V(Q)} (1 - o(1))n^{\beta(U)-\alpha(U)}.
\end{aligned}$$

Note that $\beta(U) < \alpha(U)$ for all $V(Q') \subset U \subset V(Q)$ (otherwise, letting R denote the induced subgraph of Q on U, we would have $\alpha(R) = \beta(R)$, contradicting the fact that $Q_0, \dots, Q_t$ is an (α, β)-grading). It follows that

$$\mathrm{Var}[\mathbf{k}] = 1 \pm o(1).$$

Plugging the bounds $\mathrm{E}[\mathbf{k}] = 1 - o(1)$ and $\mathrm{Var}[\mathbf{k}] = 1 \pm o(1)$ and $\mathrm{E}[\mathbf{I}_H] = o(1)$ into Lemma A.2, we have

$$d_{TV}\big(\mathcal{L}(\mathbf{k}), \mathrm{Po}(\mu)\big) \le \frac{\mathrm{Var}[\mathbf{k}]}{\mathrm{E}[\mathbf{k}]} - 1 + 2 \max_{H \in \mathcal{Q}} \mathrm{E}[\mathbf{I}_H] = o(1).$$

Finally, since $d_{TV}\big(\mathrm{Po}(1), \mathrm{Po}(1 - o(1))\big) = o(1)$, we conclude that $\mathbf{k}$ is asymptotically $\mathrm{Po}(1)$, which completes the proof. $\square$

B Proof of Lemma 3.9

In this section we also assume that $\beta(e) > 0$ for all $e \in P$ (see the paragraph before the statement of Lemma 3.9). This assumption in particular implies that $E(\mathbf{G}) \cap E(\mathbf{P}) = \emptyset$ almost surely. Thus we only have to prove that with constant probability $\mathbf{G} \cup \mathbf{P}$ does not contain any P-subgraphs other than $\mathbf{P}$ itself (a formal argument is included at the end of this section).

Lemma B.1.

1. *For every P-colored graph G in the support of* $\mathbf{G}$ *and every subgraph $Q \subseteq P$,*

$$\frac{\Pr[\,\mathbf{G}\cup\mathbf{Q}=G\,]}{\Pr[\,\mathbf{G}=G\,]} = (1+o(1))\frac{\mathrm{sub}(Q,G)}{n^{\alpha(Q)-\beta(Q)}}.$$

2. *If $\mathcal{A}$ is a property of P-colored graphs which holds a.a.s. for* $\mathbf{G}$*, then $\mathcal{A}$ holds a.a.s. for* $\mathbf{G}\cup\mathbf{Q}$ *for every $Q \subseteq P$.*

3. $\liminf_{n\to\infty} \Pr[\,\mathrm{sub}(P,\mathbf{G}\cup\mathbf{P})=1\,] > 0$*. That is,* $\mathbf{G}\cup\mathbf{P}$ *has a unique P-subgraph (namely* $\mathbf{P}$*) with probability bounded away from* 0*.*

Proof. (1): Noting that the number of possible Q-subgraphs in $\mathbf{G}$ is $\prod_{v\in V(Q)}\lfloor n^{\alpha(v)}\rfloor = (1-o(1))n^{\alpha(Q)}$, we have

$$\Pr[\,\mathbf{G}\cup\mathbf{Q}=G\,] = (1+o(1))n^{-\alpha(Q)} \sum_{K\in\mathrm{Sub}(Q,G)} \ \sum_{H\,:\,G\setminus K\subseteq H\subseteq G} \Pr[\,\mathbf{G}=H\,].$$

For every $K\in\mathrm{Sub}(Q,G)$ and H such that $G\setminus K\subseteq H\subseteq G$, we have

$$\begin{aligned}\Pr[\,\mathbf{G}=H\,] &= \Pr[\,\mathbf{G}=G\,]\cdot \prod_{e\in E(K\setminus H)} \frac{1-n^{-\beta(\widehat{e})}}{n^{-\beta(\widehat{e})}}\\ &= (1-o(1))\Pr[\,\mathbf{G}=G\,]\cdot\begin{cases} n^{\beta(Q)} & \text{if } H=G\setminus K,\\ n^{\beta(Q)-\Omega(1)} & \text{otherwise (since } \beta \text{ positive).}\end{cases}\end{aligned}$$

(Above, $\widehat{e}$ is the edge in P corresponding to e under the vertex-coloring of G.) Since $n^{-\Omega(1)}$ dominates $2^{|E(Q)|}-1$ (i.e. the number of summands where $H\neq G\setminus K$), statement (1) follows.

(2): Suppose $\mathcal{A}$ holds a.a.s. with respect to $\mathbf{G}$ (i.e. $\liminf_{n\to\infty}\Pr[\,\mathbf{G}\notin\mathcal{A}\,]=0$) and let $Q\subseteq P$. Let $c>0$ be an arbitrary (large) constant. We split up the event $\{\mathbf{G}\cup\mathbf{Q}\notin\mathcal{A}\}$ as follows:

$$\begin{aligned}\Pr[\,\mathbf{G}\cup\mathbf{Q}\notin\mathcal{A}\,] &\le \Pr[\,\mathrm{sub}(Q,\mathbf{G}\cup\mathbf{Q})\ge cn^{\alpha(Q)-\beta(Q)}\,]\\ &\qquad + \Pr[\,\mathbf{G}\cup\mathbf{Q}\notin\mathcal{A} \text{ and } \mathrm{sub}(Q,\mathbf{G}\cup\mathbf{Q})\le cn^{\alpha(Q)-\beta(Q)}\,].\end{aligned}$$

We bound each of the righthand terms separately.

First, note that

$$\begin{aligned}\mathrm{E}[\,\mathrm{sub}(Q,\mathbf{G}\cup\mathbf{Q})\,] &= \sum_{R\subseteq Q}\mathrm{E}[\,|\{H\in\mathrm{Sub}(Q,\mathbf{G}\cup\mathbf{Q}) : \chi(H\cap\mathbf{Q})=R\}|\,]\\ &\le \sum_{R\subseteq Q} n^{\alpha(Q)-\alpha(R)}\cdot n^{-\beta(Q\setminus R)}\\ &\le 2^{|E(Q)|}\cdot n^{\alpha(Q)-\beta(Q)}\end{aligned}$$

(the latter inequality holds since $\alpha(R)\ge\beta(R)$ for all R). So by Markov's inequality,

$$\Pr[\,\mathrm{sub}(Q,\mathbf{G}\cup\mathbf{Q})\ge cn^{\alpha(Q)-\beta(Q)}\,] \le \frac{2^{|E(Q)|}}{c}.$$

Second, we have

$$
\begin{aligned}
\Pr[\,\mathbf{G}\cup\mathbf{Q} \notin \mathcal{A} \text{ and } \mathrm{sub}(Q,\mathbf{G}\cup\mathbf{Q}) \le cn^{\alpha(Q)-\beta(Q)}\,]
&= \sum_{G\,:\,G\notin\mathcal{A} \text{ and } \mathrm{sub}(Q,G)\le cn^{\alpha(Q)-\beta(Q)}} \Pr[\,\mathbf{G}\cup\mathbf{Q} = G\,]\\
&= \sum_{G\,:\,G\notin\mathcal{A} \text{ and } \mathrm{sub}(Q,G)\le cn^{\alpha(Q)-\beta(Q)}} (1+o(1))\Pr[\,\mathbf{G} = G\,]\frac{\mathrm{sub}(Q,G)}{n^{\alpha(Q)-\beta(Q)}} \text{ (by (1))}\\
&\le \sum_{G\,:\,G\notin\mathcal{A} \text{ and } \mathrm{sub}(Q,G)\le cn^{\alpha(Q)-\beta(Q)}} (1+o(1))c\Pr[\,\mathbf{G} = G\,]\\
&\le (1+o(1))c\Pr[\,\mathbf{G}\notin\mathcal{A}\,].
\end{aligned}
$$

Since $\liminf_{n\to\infty} \Pr[\,\mathbf{G}\notin\mathcal{A}\,] = 0$, it follows that

$$\liminf_{n\to\infty} \Pr[\,\mathbf{G}\cup\mathbf{Q}\notin\mathcal{A}\,] \le \frac{2^{|E(Q)|}}{c}.$$

Since c may be chosen arbitrarily large, we conclude that $\mathcal{A}$ holds a.a.s. with respect to $\mathbf{G}\cup\mathbf{Q}$.

(3): Note that for $Q = P$ and $\mathrm{sub}(P,G) = 1$ Lemma B.1(1) simplifies to $\Pr[\,\mathbf{G}\cup\mathbf{P} = G\,] = (1+o(1))\Pr[\,\mathbf{G} = G\,]$. Thus, we have

$$
\begin{aligned}
\liminf_{n\to\infty} \Pr[\,\mathrm{sub}(P,\mathbf{G}\cup\mathbf{P}) = 1\,] &= \liminf_{n\to\infty} \Pr[\,\mathrm{sub}(P,\mathbf{G}) = 1\,]\\
&> 0 \text{ (by Lemma 2.10).}
\end{aligned}
$$

□

Remark B.2. In the uncolored setting (where P is balanced and $\mathbf{G} = G(n, n^{-\theta(P)})$), in Lemma B.1(1) the denominator $n^{\alpha(Q)-\beta(Q)}$ should be replaced by $\mathrm{E}[\,\mathrm{sub}(Q,\mathbf{G})\,]$.

Proof of Lemma 3.9. Let $h : \mathcal{G}_{\alpha,\beta}(n) \to \{0,1\}$ denote the $\textsc{Subgraph}_{\mathsf{col}}(P)$ function, that is, $h(G) = 1 \Leftrightarrow G$ contains a P-subgraph. Assume $f : \mathcal{G}_{\alpha,\beta}(n) \to \{0,1\}$ solves $\textsc{Subgraph}_{\mathsf{col}}(P)$ in the average-case on $\mathbf{G}$, that is,

$$\Pr[\,f(\mathbf{G}) = h(\mathbf{G})\,] = 1 - o(1).$$

By Lemma B.1(2),

$$\Pr[\,f(\mathbf{G}\cup\mathbf{Q}) = h(\mathbf{G}\cup\mathbf{Q}) \text{ for all } Q \subseteq P\,] = 1 - o(1).$$

Since the event "$f^{\cup\mathbf{G}}$ is sensitive over $\mathbf{P}$" depends only on the values of $f(\mathbf{G}\cup\mathbf{Q})$ for $Q \subseteq P$, we have

$$\Pr[\,f^{\cup\mathbf{G}} \text{ is sensitive over } \mathbf{P} \Leftrightarrow h^{\cup\mathbf{G}} \text{ is sensitive over } \mathbf{P}\,] = 1 - o(1). \tag{15}$$

As we already indicated above,

$$h^{\cup\mathbf{G}} \text{ is sensitive over } \mathbf{P} \Leftrightarrow E(\mathbf{G}) \cap E(\mathbf{P}) = \emptyset \text{ and } \mathrm{sub}(P,\mathbf{G}\cup\mathbf{P}) = 1 \tag{16}$$

(with probability 1). To see why, first assume $E(\mathbf{G}) \cap E(\mathbf{P}) = \emptyset$ and $\mathrm{sub}(P,\mathbf{G}\cup\mathbf{P}) = 1$. It follows that $h^{\cup\mathbf{G}}(\mathbf{P} - \{e\}) = 0$ for all $e \in E(\mathbf{P})$. Since $h^{\cup\mathbf{G}}(\mathbf{P}) = 1$, this shows that $h^{\cup\mathbf{G}}$ is sensitive over

$\mathbf{P}$. For the opposite direction, consider the case that there exists $e \in E(\mathbf{G}) \cap E(\mathbf{P})$. Note that e does appear in $\mathsf{Sens}(h^{\cup G}, \mathbf{P})$. Therefore, $h^{\cup \mathbf{G}}$ is not sensitive over $\mathbf{P}$. Finally, consider the case that $\mathrm{sub}(P, \mathbf{G} \cup \mathbf{P}) > 1$. Then $\mathbf{G} \cup \mathbf{P}$ contains a P-subgraph other than $\mathbf{P}$; this P-subgraph necessarily does not include some edge $e \in E(P)$. Note that $\mathrm{sub}(P, \mathbf{G} \cup (\mathbf{P} - \{e\})) \geq 1$, which means that e does appear in $\mathsf{Sens}(h^{\cup G}, \mathbf{P})$. So again $h^{\cup \mathbf{G}}$ is not sensitive over $\mathbf{P}$.

From (15) and (16), we have

$$\begin{aligned}\Pr[\, f^{\cup \mathbf{G}} \text{ is sensitive over } \mathbf{P}\,] &\geq \Pr[\, E(\mathbf{G}) \cap E(\mathbf{P}) = \emptyset \text{ and } \mathrm{sub}(P, \mathbf{G} \cup \mathbf{P}) = 1\,] - o(1)\\ &\geq \Pr[\,\mathrm{sub}(P, \mathbf{G} \cup \mathbf{P}) = 1\,] - \Pr[\, E(\mathbf{G}) \cap E(\mathbf{P}) \neq \emptyset\,] - o(1).\end{aligned}$$

Since β is positive,

$$\Pr[\, E(\mathbf{G}) \cap E(\mathbf{P}) \neq \emptyset\,] \leq \sum_{e \in E(P)} n^{-\beta(e)} = o(1).$$

Completing the proof, we have

$$\liminf_{n\to\infty} \Pr[\, f^{\cup \mathbf{G}} \text{ is sensitive over } \mathbf{P}\,] \geq \liminf_{n\to\infty} \Pr[\,\mathrm{sub}(P, \mathbf{G} \cup \mathbf{P}) = 1\,] > 0$$

by Lemma B.1(3). □

C The switching lemma proof of (10)

Let C be a type-I (constant-depth, unbounded fan-in) AC^0 circuit of size $S = S(n) = n^{O(1)}$ and depth $d = d(n) = O(1)$ over the set of variables $\mathcal{E}$ (corresponding to potential edges of graphs in $\mathcal{G}_{\alpha,\beta}(n)$). Recall that $\rho : \mathcal{E} \to \{0, 1, *\}$ is the (coordinate-wise independent) random restriction with

$$\Pr[\,\rho(e) = \star\,] = n^{-\beta(\widehat{e})-\delta} \ (=: q_e), \qquad \Pr[\,\rho(e) = 1 \mid \rho(e) \neq \star\,] = n^{-\beta(\widehat{e})} \ (=: p_e),$$

where $\delta > 0$ is a fixed constant. Restating (10), we must show that

$$\Pr[\,\mathsf{C}{\restriction_\rho} \text{ depends on } > n^{\delta} \text{ variables}\,] = n^{-\omega(1)}.$$

The proof uses nothing more than Håstad's Switching Lemma [14]. We generate ρ as a composition of $d+1$ random restrictions $\rho_0, \ldots, \rho_d$ where

- $\rho_0 : \mathcal{E} \to \{0, 1, \star\}$ is the *unbalanced* random restriction with

$$\Pr[\,\rho_0(e) = \star\,] = p_e, \qquad \Pr[\,\rho_0(e) = 1 \mid \rho_0(e) \neq \star\,] = \frac{\frac{1}{2}(p_e + q_e) - p_e q_e}{1 - p_e} \ (=: r_e)$$

  (note that $r_e = \frac{1}{2}p_e + o(p_e)$ is a well-defined value in $(0, 1)$ since $q_e \ll p_e \ll 1$),

- for $i \in \{1, \ldots, d\}$, $\rho_i : \rho_{i-1}^{-1}(\star) \to \{0, 1, \star\}$ is the *balanced* random restriction with

$$\Pr[\,\rho_i(e) = \star\,] = n^{-\delta/d}, \qquad \Pr[\,\rho_i(e) = 1 \mid \rho_i(e) \neq \star\,] = 1/2.$$

It is easy to check that $\rho_d \circ \cdots \circ \rho_0$ is indeed the same distribution as ρ.

Let $D(\cdot)$ denote *decision tree depth* of a boolean function. Let g be the function computed by a gate of height $i \in \{1, \ldots, d\}$ in C. Let $g_1, \ldots, g_m$ be the children $g_1, \ldots, g_m$. (That is, g is the AND or the OR of $g_1, \ldots, g_m$ and has distance $d - i$ to the output of C.) Håstad's Switching Lemma tells us

$$\Pr\left[D(g{\restriction}_{\rho_i \circ \cdots \circ \rho_0}) > \delta \log n \;\middle|\; \bigwedge_{j \in [m]} D(g_j{\restriction}_{\rho_{i-1} \circ \cdots \circ \rho_0}) \leq \delta \log n \right] \leq (5n^{-\delta/d} \delta \log n)^{\delta \log n}.$$

By a union bound, it follows that

$$\Pr[\, \mathsf{C}{\restriction}_\rho \text{ depends on } > n^\delta \text{ variables} \,] \leq \Pr[\, D(\mathsf{C}{\restriction}_\rho) > \delta \log n \,] \leq S \cdot (5n^{-\delta/d} \delta \log n)^{\delta \log n}.$$

Since $S = n^{O(1)}$ and $d = O(1)$ and $\delta = \Omega(1)$ (with these constants only depending on the fixed pattern P), we have

$$S \cdot (5n^{-\delta/d} \delta \log n)^{\delta \log n} = n^{-\omega(1)}.$$

Note that this argument (as indeed all results in this paper) tolerates super-constant depth $d(n)$ up to $o(\log n / \log \log n)$.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de