2014 IEEE Annual Symposium on Foundations of Computer Science

On the AC^0 Complexity of Subgraph Isomorphism

Yuan Li
University of Chicago
Chicago, USA
yuanli@cs.uchicago.edu

Alexander Razborov
University of Chicago
Chicago, USA
razborov@cs.uchicago.edu

Benjamin Rossman
National Institute of Informatics
Tokyo, Japan
rossman@nii.ac.jp

***Abstract*—Let P be a fixed graph (hereafter called a "pattern"), and let 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 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 SUBGRAPH$_{\mathsf{col}}(P)$ in which the target graph G is $V(P)$-colored, and the average-case version 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 [21] is essentially tight, for any pattern P whatsoever.**
- **We prove that if Q is a minor of P then SUBGRAPH$_{\mathsf{col}}(Q)$ is reducible to 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 SUBGRAPH(M_3) to 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.**

I. 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 CLIQUE and 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 some useful structure (see e.g. the sources [3], [8], [16], [17] 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 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 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* SUBGRAPH(P) *in terms of the treewidth of* P *only?*

Prior to our work, Rossman [21] 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 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

0272-5428/14 $31.00 © 2014 IEEE
DOI 10.1109/FOCS.2014.44

to introduce the following handy notation: given a pattern P, we let $C(P)$ be the minimal real number $c \geq 0$ for which 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), $C(K_k) \geq k/4$ [21] and $C(G_{k,k}) \geq \Omega(k)$ [4].

A. *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*, 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 SUBGRAPH(P) to 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. $C_{\mathsf{col}}(P) \geq \Omega(tw(P)/\log tw(P))$.

By previous work of Marx [16], it was known that SUBGRAPH$_{\mathsf{col}}(P)$ has no $n^{o(tw(P)/\log tw(P))}$ algorithm unless the Exponential Time Hypothesis fails. Theorem 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 VI.)

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 $C_{\mathsf{col}}(Q) \leq C_{\mathsf{col}}(P)$ (Theorem 6).[1] Whether a similar result holds for $C(P)$ is open, but we give a strong evidence (Theorem 7) 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 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, 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 [14] or Section II-B 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, $C_{\mathsf{ave}}(P) \leq C(P)$ so the whole picture now looks like

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

where $\approx$ means approximation within a logarithmic factor. Also, $C_{\mathsf{ave}}(K_k) \geq k/4$ [21] and $C_{\mathsf{ave}}(G_{k,k}) \geq \Omega(k)$ [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 2. $\kappa(P) \leq C_{\mathsf{ave}}(P) \leq 2\kappa(P) + O(1)$.

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 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 $C_{\mathsf{ave}}(P) \geq \kappa(P)$ generalizes the proof of $C_{\mathsf{ave}}(K_k) \geq k/4$ in Rossman [21] and improves a previous lower bound of Amano [4] for general patterns P. (A detailed comparison with previous work is given in Section II-D following the definition of $\kappa(P)$.)

Finally, let us say a few words about the proof of Theorem 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 [9], [16] that characterize the treewidth of P in terms of

[1] It is worth observing that this fact, along with the recent result [7] by Chekura and Chuzhoy and Amano's bound $C_{\mathsf{col}}(G_{k,k}) \geq \Omega(k)$ [4] already implies the weaker bound $C_{\mathsf{col}}(P) \geq tw(P)^{\Omega(1)}$. But the exponent given by this approach will be disappointingly small.

the existence of a certain concurrent flow on P, which we convert to a suitable point in $\theta_{\mathsf{col}}(P)$.

The paper is organized as follows. In Section II we give the necessary definitions and preliminaries; in particular, in Section II-C we present the parameters $\kappa(P)$ and $\kappa_{\mathsf{col}}(P)$ that are our main technical tools in this paper. Section III is devoted to the proof of Theorem 2, and it also paves way to the proof of Theorem 1 that, up to a certain point, goes in parallel to the former. The proof of Theorem 1 is completed in Section IV. Section V contains structural results about the behavior of SUBGRAPH(P) and SUBGRAPH$_{\mathsf{col}}(P)$ with respect to minors and subgraphs. We conclude with a brief discussion and list of open problems in Section VI.

II. Definitions and Preliminaries

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

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 $E(G \times H) := \{\{(v, v'), (w, w')\} : \{v, w\} \in E(G)$ 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. [5]). 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 $tw(G_{k,k}) = k$.

A. Subgraph Isomorphism Problems

Throughout this paper, the letters P 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:

- SUBGRAPH(P) is the problem, given a graph G, of determining whether or not G contains a P-subgraph.
- 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.[2]

Definition 1. Let $C(P)$ (resp. $C_{\mathsf{col}}(P)$) denote the minimum real number $c > 0$ such that SUBGRAPH(P) (resp. 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 SUBGRAPH(P) is reducible to SUBGRAPH(Q) via a linear monotone projection then $C(P) \leq C(Q)$, and this remains true if we add the subscript col to both sides.

Lemma 1.

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

Proof: (1): The second inequality $C_{\mathsf{col}}(P) \leq tw(P) + 1$ 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 $C(P) \leq C_{\mathsf{col}}(P)$ is also implicitly proved there by reducing SUBGRAPH(P) to SUBGRAPH$_{\mathsf{col}}(P)$: the reduction searches through logarithmically many different colorings $\chi_1, \chi_2, \cdots : V(G) \to V(P)$ 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 [10]. If P is a core, then $(G, \chi) \mapsto G$ is a reduction from SUBGRAPH$_{\mathsf{col}}(P)$ to SUBGRAPH(P). To see why, it suffices to show that every P-subgraph of G is properly colored with respect to every homomorphism

[2]Recall that AC^0 is the class of problems solvable by polynomial-size constant-depth boolean circuits with unbounded fan-in. Uniformity issues do not play any role in this paper.

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

$\chi : G \to P$. Suppose 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. ∎

B. 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.

- (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.

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

The elementary proof is included in the full paper.

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 3.

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.

Lemma 3(1) is a standard fact about random graphs (see [14]); Lemma 3(2) was proved in [6].

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 1, 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 1. Obviously, $C_{\mathsf{ave}}(P) \le C(P)$, 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 3(2) implies that $C_{\mathsf{ave}}(P) = C_{\mathsf{ave}}(K_4) \le 4$. 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 $C(P) \ge C(G_{k,k}) \ge \Omega(k)$ 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. $tw(P) \ge \Omega(k)$ and that P is a core. Given these facts, the bounds $C_{\mathsf{ave}}(P) \le O(d)$ and $C(P) \ge \Omega(k/\log k)$ 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 3. **(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
  - $\alpha(P) = \beta(P)$,
  - $\alpha(Q) \ge \beta(Q)$ 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 *strictly balanced* if $\alpha(Q) > \beta(Q)$ for every nonempty proper subgraph $Q \subset P$.

(iv) 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. 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. For general (not necessary balanced) P, we consider functions $\alpha(v) := 1_{v \in V(\mathsf{Bal}(P))}$ and $\beta(e) := \theta(P) \cdot 1_{e \in E(\mathsf{Bal}(P))}$. This threshold pair corresponds to the (uncolored, average-case) problem $\text{SUBGRAPH}_{\mathsf{ave}}(P)$ (i.e. the average-case behavior of $\text{SUBGRAPH}(P)$ on $G(n, n^{-\theta(P)})$ parallels $\text{SUBGRAPH}_{\mathsf{col}}(P)$ on $G_{\alpha,\beta}(n)$). Thus, Definition 3 is indeed a generalization of the threshold exponent, and the following lemma makes the analogy even more clear.

Lemma 4. *For every pattern P and threshold pair $(\alpha, \beta) \in \theta_{\mathsf{col}}(P)$,*

$$\lim_{n\to\infty} \Pr[\, G_{\alpha,\beta}(n) \textit{ contains a } P\textit{-subgraph} \,] \in (0, 1).$$

See the full paper for a detailed analysis of the asymptotic distribution of P-subgraphs in $G_{\alpha,\beta}(n)$, including a proof of Lemma 4. (In the case of strictly balanced (α, β), it is shown that the number of P-subgraphs is asymptotically distributed according to the Poisson distribution $\mathrm{Po}(1)$.)

In the context of $\text{SUBGRAPH}_{\mathsf{col}}(P)$, we speak of the *average-case complexity with respect to $G_{\alpha,\beta}(P)$*, meaning the size of an AC^0 circuit which solves $\text{SUBGRAPH}_{\mathsf{col}}(P)$ 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.

C. 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 [21], which we explain in the next section.

Definition 4. (Union sequences and hitting sets) A *union sequence* for P is a sequence $Q_1, \ldots, Q_t$ 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 5. (Parameters $\kappa(P)$, $\kappa_{\alpha,\beta}(P)$, $\kappa_{\mathsf{col}}(P)$)

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

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

If P is not balanced, then $\kappa(P) := \kappa(\mathsf{Bal}(P))$.

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

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

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

Remark 3. Later on we will see that in this definition we could restrict ourselves to threshold pairs with $\alpha \equiv 1$, But since arbitrary threshold pairs appear quite naturally in our lower bound proofs, 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 $\text{SUBGRAPH}(P)$.

Lemma 5. (Dual form of $\kappa(P)$ 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, for all $(\alpha, \beta) \in \theta_{\mathsf{col}}(P)$,

$$\kappa_{\alpha,\beta}(P) = \max_{\mathcal{H}} \min_{Q\in\mathcal{H}} \alpha(Q) - \beta(Q).$$

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 $\max_{\mathcal{H}} \min_{Q\in\mathcal{H}} f(Q) \le \kappa(P)$. Since $\mathcal{H}$ is a hitting set, for any union sequence $\{Q_i\}$, there exists some $Q_i \in \mathcal{H}$. It follows that $\min_{Q\in\mathcal{H}} f(Q) \le \max_i f(Q_i)$, and thus $\min_{Q\in\mathcal{H}} f(Q) \le \kappa(P)$ as $\{Q_i\}$ is taken arbitrarily.

On the other hand, let us prove $\kappa(P) \le \max_{\mathcal{H}} \min_{Q\in\mathcal{H}} f(Q)$. 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,

$$\begin{aligned}\max_{\mathcal{H}} \min_{Q\in\mathcal{H}} f(Q) &\ge \min_{S^{(j)}\in\mathcal{S}} f(S^{(j)}) \\ &= \min_j \max_i f(P_i^{(j)}) = \kappa(P),\end{aligned}$$

which completes the proof. ∎

D. Comparison with previous work

The dual (max-min) expression for $\kappa(P)$ given by Lemma 5(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, comes from generalizing Rossman's proof of $C_{\mathsf{ave}}(K_k) \geq k/4$ [21].

Previous work of Amano [4] also generalizes the technique of [21] 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 $\kappa(K_k) = k/4 + O(1)$). A key insight of the present paper is that the stronger parameter $\kappa(P)$ leads to *upper bounds* on $C_{\mathsf{ave}}(P)$ which are tight within a multiplicative constant. (The primal (min-max) expression for $\kappa(P)$ given by Definition 5 is naturally suited to upper bounds.)

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 [21] by showing $C_{\mathsf{ave}}(K_k) \leq k/4 + O(1)$. Nakagawa and Watanabe [18] 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

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

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 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 2) both achieve a tighter generalization of [21] and close the (arbitrarily large) gap between the previous bounds $\ell(P) \leq C_{\mathsf{ave}}(P) \leq u(P) + O(1)$ of [4], [18]. 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], [18], [21]).

III. Average-Case AC^0 Complexity

In this section, we prove Theorem 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 $\mathrm{SUBGRAPH}_{\mathsf{ave}}(P)$ 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. *For every pattern P and $(\alpha, \beta) \in \theta_{\mathsf{col}}(P)$, the average-case AC^0-complexity of $\mathrm{SUBGRAPH}_{\mathsf{col}}(P)$ 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 2 and Theorem 3 separately, to avoid redundancy we present a proof of the latter only. For balanced P the proof of Theorem 2 looks exactly like the proof of Theorem 3 in the special case where $\alpha \equiv 1$ and $\beta \equiv \theta(P)$ (see Remark 2). 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 3).

Theorem 3 plays a key role in our other main result, Theorem 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 $\mathrm{SUBGRAPH}_{\mathsf{col}}(P)$ is lower-bounded by the average-case AC^0-complexity of $\mathrm{SUBGRAPH}_{\mathsf{col}}(P)$ on $G_{\alpha,\beta}(n)$ for every $(\alpha, \beta) \in \theta_{\mathsf{col}}(P)$, Theorem 3 directly implies:

Corollary 1. $C_{\mathsf{col}}(P) \geq \kappa_{\mathsf{col}}(P)$.

In Section IV, we will show that $\kappa_{\mathsf{col}}(P) \geq \Omega(tw(P)/\log tw(P))$; together with Corollary 1, this proves Theorem 1. The remainder of this section contains the proof of Theorem 3. The $n^{2\kappa_{\alpha,\beta}(P)+O(1)}$ upper bound is proved in Section III-A, below. See the full paper for the proof (generalizing [21]) of the matching $n^{\kappa_{\alpha,\beta}(P)-o(1)}$ lower bound.

A. Proof of Theorem 3 (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 (properly 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 $\mathcal{G}'_{\alpha,\beta}(n)$ denote the set of $G \in \mathcal{G}_{\alpha,\beta}(n)$ such that $\mathrm{sub}(Q, G) \leq n^{\alpha(Q)-\beta(Q)+1}$ 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 6. $\Pr[\,\mathbf{G} \notin \mathcal{G}'_{\alpha,\beta}(n)\,] = o(1)$.

We wish to construct a *deterministic* AC^0-circuit C which solves $\mathrm{SUBGRAPH}_{\mathsf{col}}(P)$ 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 $\mathrm{SUBGRAPH}_{\mathsf{col}}(P)$ correctly with probability $1 - o(1)$ on *every* $G \in \mathcal{G}'_{\alpha,\beta}(n)$. That is, we will show

Lemma 7. *There exists a random AC^0 circuit $\mathbf{C}$ of size $n^{2\kappa_{\alpha,\beta}(P)+O(1)}$ and depth $O(e(P))$ 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 follows as a corollary of Lemmas 6 and 7.

Proposition 1. *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 6 and 7 imply that $\Pr[\,\mathbf{C}(\mathbf{G}) = 1 \Leftrightarrow \mathrm{sub}(P,\mathbf{G}) \geq 1\,] = 1 - o(1)$. Now Proposition 1 follows by a straightforward application of Yao's Principle [23]. ∎

The random circuit $\mathbf{C}$*:* It remains to define the randomized AC^0-algorithm solving $\mathrm{SUBGRAPH}_{\mathsf{col}}(P)$ 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 $Q_1, \dots, Q_t$ 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 $L_1, \dots, L_t$ 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 (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 $1 \leq i < j < k$. We compute L_k in three steps as follows.

1: Let M_k be the $\ell_i \times \ell_j$ array where, for $a \in [\ell_i]$ and $b \in [\ell_j]$, the entry $M_k(a,b)$ contains the graph $L_i(a) \cup L_j(b)$ if this is a valid Q_k-subgraph (i.e. $L_i(a)$ and $L_j(b)$ are consistent on $V(Q_i) \cap V(Q_j)$); otherwise, $M_k(a,b)$ is empty. (Note that, since L_i and L_j are good, M_k contains each Q_k-subgraph of G exactly once.)

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 $m_k \geq \#\{Q_k\text{-subgraphs of } G\} = |\mathrm{Supp}(M_k)|$. Let $\mathbf{h}_k$ be a uniform random function $\mathbf{h}_k : [\ell_i] \times [\ell_j] \to [m_k]$. (Restricted to the $\leq m_k$ nonempty entries of M_k, this gives a uniform random packing of $\leq m_k$ balls into m_k bins.)

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 $L_k(p,q)$ contain the qth element of $\mathbf{h}_k^{-1}(p) \cap \mathrm{Supp}(M_k)$ if $|\mathbf{h}_k^{-1}(p) \cap \mathrm{Supp}(M_k)| \geq q$; otherwise, let $L_k(p,q)$ be blank. (Note that L_k is good if, and only if, $\bigwedge_{p\in[m_k]} |\mathbf{h}_k^{-1}(p) \cap \mathrm{Supp}(M_k)| \leq \ln m_k$.)

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.

Due to page limitations, we omit the analysis of the success probability of this algorithm and its implementation on AC^0 circuits of size $n^{2\kappa_{\alpha,\beta}(P)+O(1)}$. (See the full paper for details.)

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

In this section, we give upper and lower bounds on the parameter $\kappa_{\mathsf{col}}(P)$ for arbitrary patterns P.

Proposition 2. $\kappa_{\mathsf{col}}(P) \leq tw(P) + 1$.

The proof (included in the full paper) actually shows that $\kappa_{\mathsf{col}}(P) \leq bw(P)$ (the *branch-width* of P); it is well-known that $bw(P) \leq tw(P) + 1$ (see [20]).

Theorem 4. $\kappa_{\mathsf{col}}(P) \geq \Omega(tw(P)/\log tw(P))$.

Theorem 4, in conjunction with Corollary 1 ($C_{\mathsf{col}}(P) \geq \kappa_{\mathsf{col}}(P)$), directly implies one of our main results: $C_{\mathsf{col}}(P) \geq \Omega(tw(P)/\log tw(P))$ (Theorem 1). Our second lower bound on $\kappa_{\mathsf{col}}(P)$ (Theorem 5, stated in Section IV-A) eliminates the log-factor loss in Theorem 4 in the case that P is a constant-degree expander.

The proof of Theorem 4 uses a characterization of treewidth from Marx [16] (based on results of Feige et al [9]): 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 10). 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.

Definition 6.

(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 two distinct vertices v, w, we let $f(v,w) := f(\{v\},\{w\})$. For disjoint $S, T \subseteq V(P)$, let $f(S,T) := \sum_{v \in V,\, w \in T} f(v,w)$.

(iv) For $\pi \in \mathsf{Paths}(P)$, define $(\alpha_\pi, \beta_\pi) \in \theta_{\mathsf{col}}(P)$ by

$$\alpha_\pi(v) := \begin{cases} \frac{1}{2} & \text{if } v \text{ is an endpoint of } \pi, \\ 1 & \text{if } v \text{ is an interior vertex of } \pi, \\ 0 & \text{if } v \notin V(\pi), \end{cases}$$

$$\beta_\pi(e) := \begin{cases} 1 & \text{if } e \in E(\pi), \\ 0 & \text{if } e \notin E(\pi). \end{cases}$$

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

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

Lemma 8. $(\alpha_f, \beta_f) \in \theta_{\mathsf{col}}(P)$ *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. ■

Lemma 9. *For all* $Q \subseteq P$ *and* $f \in \mathsf{Flows}(P)$,

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

Proof: We have $f(S,T) = \sum_{\pi \in \mathsf{Paths}(P)} f(\pi) \cdot \pi(S,T)$ where

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

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

$$\text{(1)} \qquad \alpha_\pi(Q) - \beta_\pi(Q) \geq \tfrac{1}{2} \pi(V(Q), \overline{V(Q)}).$$

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 8 (since $(\alpha_\pi, \beta_\pi) \in \theta_{\mathsf{col}}(P)$); so (1) 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 $\alpha_\pi(Q) - \beta_\pi(Q)$ equals $\frac{1}{2}$ times the number of edges of π that cross between $V(Q)$ and $\overline{V(Q)}$; this number is at least 1, so again (1) holds. ■

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

Lemma 10 ([9], [16])**.** *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: Suppose $tw(P) = k$ and fix $W \subseteq V(P)$ and $f \in \mathsf{Flows}(P)$ as in the Lemma 10. 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

$$\begin{aligned} \alpha_f(Q) - \beta_f(Q) &\geq \tfrac{1}{2} f(W \cap V(Q), W \setminus V(Q)) \\ &\geq \tfrac{1}{2ck\log k} |W \cap V(Q)| \cdot |W \setminus V(Q)| \\ &\geq \tfrac{4k}{81c\log k}. \end{aligned}$$

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

A. Tight lower bound for expanders

We give a second lower bound on $\kappa_{\mathsf{col}}(P)$ in terms of edge expansion, which eliminates the log-factor loss in Theorem 4 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 5. $\kappa_{\mathsf{col}}(P) \geq \frac{h(P)|V(P)|}{3\Delta(P)}$.

Proof: Define $\beta : E(P) \to [0, 2]$ by

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

and note 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 $Q \in \mathcal{H}$, we have

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

Completing the proof,

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

■

V. Minor-Monotonicity and Monotone Projections

See the full paper for proofs of the two results in this section.

Definition 7. 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 $f \leq_{\mathrm{mp}} g$, if there exists $p : J \to I \cup \{0, 1\}$ such that $f(x) = g(p^*(x))$ for all $x \in \{0, 1\}^I$.

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 $m(n)$ such that $L_1^n \leq_{\mathrm{mp}} L_2^{m(n)}$. If in addition $m(n) \leq O(n)$, we call this projection *linear*.

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

Theorem 6. *$\kappa_{\mathsf{col}}(P)$ and $C_{\mathsf{col}}(P)$ are minor-monotone. Moreover, $C_{\mathsf{col}}(P)$ is minor-monotone via linear monotone projections.*

In contrast to the colorful setting, we also show (in the uncolored setting) 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 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.

Theorem 7. $\textsc{Subgraph}(M_3)$ *is not a monotone projection of* $\textsc{Subgraph}(P_3 + M_2)$.

VI. 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:

$$\begin{array}{ccccccc} \Omega(\frac{tw(P)}{\log tw(P)}) \leq \kappa_{\mathsf{col}}(P) & \leq & C_{\mathsf{col}}(P) & \leq & tw(P) + 1 \\ & & \rotatebox{90}{$\leq$} & & \\ & & C(P) & & \\ & & \rotatebox{90}{$\leq$} & & \\ \kappa(P) & \leq & C_{\mathsf{ave}}(P) & \leq & 2\kappa(P) + O(1). \end{array}$$

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 1). We do not know of any gap between $C(P)$ and $C_{\mathsf{col}}(P)$. More broadly, we can ask whether $C(P)$ is bounded 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 2 be improved to $\kappa(P) + O(1)$?*

Question 3. *Can the* $\log tw(P)$ *factor can be eliminated from our lower bounds on $\kappa_{\mathsf{col}}(P)$ (Theorem 1) or $C_{\mathsf{col}}(P)$?*

We are able to answer Question 3 affirmatively in the special case where P is a constant-degree expander (Theorem 5).

Another question raised by this work is whether the AC^0 complexity of 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 SUBGRAPH(Q) to SUBGRAPH(P), even when Q is only a subgraph of P.

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

In particular, if Q is a minor (or subgraph) of P, is there a reduction from SUBGRAPH(Q) to SUBGRAPH(P) by AC^0-circuits of size $O(n^c)$ for a constant c independent of P and Q? That would imply $C(Q) \leq O(C(P))$; 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 SUBGRAPH$_{\mathsf{col}}(P)$ beyond AC^0. In particular, we recall the result of Marx [16] that 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 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 [16]. In light of our lower bounds $C_{\mathsf{col}}(P) \geq \kappa_{\mathsf{col}}(P) = \Omega(|V(P)|)$ when P is a constant-degree expander, it becomes interesting to ask:

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

Acknowledgements: We thank the anonymous reviewers for many helpful comments. Part of this work was done while the second author was at Steklov Mathematical Institute, supported by the Russian Foundation for Basic Research, and at Toyota Technological Institute, Chicago. The third author is supported by the JST ERATO Kawarabayashi Large Graph Project.

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] Hans L. Bodleander. Discovering treewidth. In *Proc. 31st International Conference on Current Trends in Theory and Practice of Computer Science*, LNCS volume 3381, pages 1–16. Springer, 2005.

[6] 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.

[7] Chandra Chekuri and Julia Chuzhoy. Polynomial bounds for the grid-minor theorem. *CoRR, abs/1305.6577,* 2013.

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

[9] 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.

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

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

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

[13] 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.

[14] Svante Janson, Tomasz Luczak and Andrzej Rucinski. *Random Graphs.* Wiley-Interscience, 2011.

[15] Stasys Jukna. *Boolean Function Complexity: Advances and Frontiers*. Springer-Verlag, 2012.

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

[17] 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.

[18] 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.

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

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

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

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

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