

Beating Treewidth for Average-Case Subgraph Isomorphism

Gregory Rosenthal*
University of Toronto

Abstract

For any fixed graph G , the subgraph isomorphism problem asks whether an n -vertex input graph has a subgraph isomorphic to G . A well-known algorithm of Alon, Yuster and Zwick (1995) efficiently reduces this to the “colored” version of the problem, denoted G -SUB, and then solves G -SUB in time $O(n^{tw(G)+1})$ where $tw(G)$ is the treewidth of G . Marx (2010) conjectured that G -SUB requires time $\Omega(n^{\text{const}\cdot tw(G)})$ and, assuming the Exponential Time Hypothesis, proved a lower bound of $\Omega(n^{\text{const}\cdot emb(G)})$ for a certain graph parameter $emb(G) = \Omega(tw(G)/\log tw(G))$. With respect to the size of AC^0 circuits solving G -SUB, Li, Razborov and Rossman (2017) proved an unconditional average-case lower bound of $\Omega(n^{\kappa(G)})$ for a different graph parameter $\kappa(G) = \Omega(tw(G)/\log tw(G))$.

Our contributions are as follows. First, we show that $emb(G)$ is at most $O(\kappa(G))$ for all graphs G . Next, we show that $\kappa(G)$ can be asymptotically less than $tw(G)$; for example, if G is a hypercube then $\kappa(G)$ is $\Theta\left(tw(G)/\sqrt{\log tw(G)}\right)$. Finally, we construct AC^0 circuits of size $O(n^{\kappa(G)+\text{const}})$ that solve G -SUB in the average case, on a variety of product distributions. This improves an $O(n^{2\kappa(G)+\text{const}})$ upper bound of Li et al., and shows that the average-case complexity of G -SUB is $n^{o(tw(G))}$ for certain families of graphs G such as hypercubes.

1 Introduction

The subgraph isomorphism problem asks, given graphs X and G , whether X has a subgraph isomorphic to G . In the “colored” or “partitioned” version of the problem, each vertex of the larger graph X comes with a “color” from the vertex set of G , and we ask whether X has a subgraph that is isomorphic to G with respect to this coloring. We denote the uncolored and colored subgraph isomorphism problems by G -SUB_{uncol}(X) and G -SUB(X) respectively.

Subgraph isomorphism is NP-complete (e.g. if G is a clique or Hamiltonian cycle), so research has focused on algorithms for a variety of special cases in the context of parameterized complexity, surveyed in [MP14]. If G is a fixed graph on k vertices then G -SUB_{uncol} is solvable in time $O(n^k)$ by brute force, where (here and throughout this section) n is the order of the input graph. The color-coding algorithm of Alon, Yuster and Zwick [AYZ95] improves on this by efficiently reducing G -SUB_{uncol} to G -SUB and solving the latter in time $O(n^{tw(G)+1})$, where $tw(G)$ is the treewidth of the fixed graph G .

The exponent $tw(G) + 1$ can sometimes be improved using fast matrix multiplication [NP85; EG04], but no significantly faster algorithm is known for either the colored or uncolored subgraph isomorphism problem. Marx [Mar10] conjectured the following:

*Email: rosenthal@cs.toronto.edu. Supported by NSERC (PGS D). Part of this work was done while the author was visiting the Simons Institute for the Theory of Computing.

Conjecture 1.1. *There is no class \mathcal{G} of graphs with unbounded treewidth, no algorithm \mathbb{A} that on inputs G and X solves $G\text{-SUB}(X)$, and no function f such that if G is in \mathcal{G} then \mathbb{A} runs in time $f(G)n^{o(tw(G))}$.*

Marx [Mar10] came close to proving Conjecture 1.1 assuming the Exponential Time Hypothesis (ETH) [IPZ01], which is the hypothesis that solving 3SAT on n variables requires $2^{\Omega(n)}$ time. We state his result in terms of a parameter $emb(G)$ (short for “embedding”) which we will define in Section 4:

Theorem 1.2 ([Mar10]). *If there is a class \mathcal{G} of graphs with unbounded treewidth, an algorithm \mathbb{A} that on inputs G and X solves $G\text{-SUB}(X)$, and a function f such that if G is in \mathcal{G} then \mathbb{A} runs in time $f(G)n^{o(emb(G))}$, then ETH is false.*

Marx [Mar10] proved that $emb(G)$ is $\Omega(tw(G)/\log tw(G))$, so Theorem 1.2 comes within a logarithmic factor in the exponent of proving Conjecture 1.1. Our main result is a counterexample to an average-case analogue of Conjecture 1.1, in a sense that will be made precise in Section 3. Moreover, our result holds on circuits of depth depending only on G .

Li, Razborov and Rossman [LRR17] proved that for fixed G , the average-case AC^0 complexity of $G\text{-SUB}$ is between $n^{\kappa(G)-o(1)}$ and $n^{2\kappa(G)+c}$, where $\kappa(G)$ is a graph property defined in Section 3 and c is an absolute constant.¹ We tighten this gap, answering an open problem posed in [LRR17]:

Theorem 1.3. *There is a constant $c > 0$ such that for any fixed graph G , the average-case AC^0 complexity of $G\text{-SUB}$ is at most $n^{\kappa(G)+c}$.*

We observe that a similar result holds easily on Turing machines, using as a subroutine the *sort-merge join* algorithm from relational algebra. This involves sorting, which cannot be done in AC^0 [Häs86], so our circuit instead uses hashing that relies on concentration of measure for subgraphs of random graphs.

Li et al. [LRR17] also proved that $\kappa(G)$ is between $\Omega(tw(G)/\log tw(G))$ and $tw(G) + 1$, from which it follows that the *worst-case* complexity of $G\text{-SUB}$ on bounded-depth circuits is at least $n^{\Omega(tw(G)/\log tw(G))}$. Li et al. posed the question of whether $\kappa(G)$ is $\Theta(tw(G))$; an affirmative answer would have implied that Conjecture 1.1 holds on bounded-depth circuits.

Our main result is a separation of κ from treewidth. The Hamming graph K_q^d has vertex set $\{1, \dots, q\}^d$ and edges between every two vertices that differ in exactly one coordinate. It is already known that K_q^d has treewidth $\Theta(q^d/\sqrt{d})$ [CK06]. We prove the following:

Theorem 1.4. *$\kappa(K_q^d)$ is $\Theta(q^d/d)$.*

Thus, if G is the hypercube graph K_2^d for example, then $\kappa(G)$ is $\Theta(tw(G)/\sqrt{\log tw(G)})$. It follows that an average-case analogue of Conjecture 1.1 is false if \mathcal{G} is taken to be the set of all hypercubes. We also prove the following (for arbitrary graphs G):

Theorem 1.5. *$emb(G)$ is $O(\kappa(G))$.*

Because of Theorem 1.5, even if our upper bound generalizes to the worst case, it is still consistent with current knowledge (in particular Theorem 1.2) that ETH is true. Another consequence of Theorem 1.5 is that the lower bound from Theorem 1.2 holds unconditionally in AC^0 .

¹In [LRR17], the parameter $\kappa(G)$ was called $\kappa_{\text{col}}(G)$.

It follows from Theorems 1.4 and 1.5 that if G is a hypercube then $emb(G) \leq O(\kappa(G)) = o(tw(G))$, so proving that Conjecture 1.1 holds under ETH cannot be done by proving that $emb(G)$ is $\Theta(tw(G))$. In fact, this conclusion was already known: Alon and Marx [AM11] proved that if G is a 3-regular expander then $emb(G)$ is $\Theta(tw(G)/\log tw(G))$. Li et al. [LRR17] proved that if G is a 3-regular expander then $\kappa(G)$ is $\Theta(tw(G))$, which makes our separation of κ from treewidth more surprising. On the other hand, we will see that Theorem 1.5 is asymptotically tight in the case of Hamming graphs.

We can make a similar statement regarding AC^0 . Amano [Ama10] observed that the coloring algorithm for G -SUB can be implemented by AC^0 circuits of size $O(n^{tw(G)+1})$ for fixed G . Our separation of κ from treewidth implies that if Conjecture 1.1 holds in AC^0 , then this cannot be proved using average-case complexity as defined here and in [LRR17].

The paper is organized as follows. In Section 2 we introduce some notation and definitions. In Section 3 we define the average-case problem and $\kappa(G)$, and give an $\tilde{O}(n^{\kappa(G)})$ -time algorithm for the average-case problem. In Section 4 we define $emb(G)$ and prove that $emb(G)$ is $O(\kappa(G))$. In Section 5 we prove that $\kappa(K_q^d)$ is $\Theta(q^d/d)$, and obtain as a corollary that $emb(K_q^d)$ is $\Theta(q^d/d)$ as well. We also summarize the proof of Chandran and Kavitha [CK06] that $tw(K_q^d)$ is $\Theta(q^d/\sqrt{d})$. In Section 6 we prove our AC^0 upper bound.

2 Preliminaries

It will be convenient to define $\tilde{O}(f(n)) = f(n) \log^{O(1)} n$. (This differs from the standard notation when $f(n) = n^{o(1)}$.) We will often fix a graph G , in which case the constants hidden in asymptotic notation are allowed to depend on G .

We use **boldface** to denote random variables. The indicator variable $\mathbb{1}\{E\}$ equals 1 if the event E occurs and 0 otherwise. Expected value is denoted $\mathbb{E}[\cdot]$. An event occurs *asymptotically almost surely (a.a.s.)* if it occurs with probability $1 - o(1)$ as n goes to infinity.

Let $[k] = \{1, \dots, k\}$ for $k \in \mathbb{N}$. If a positive real number x is used in a context where a natural number is expected (for example $[x]$), it's because x can be rounded arbitrarily to $\lceil x \rceil$ or $\lfloor x \rfloor$ without affecting the asymptotic behavior of whatever is being considered.

2.1 Graphs

All graphs we consider are simple and undirected, and may have isolated vertices. If G is a graph then let $V(G)$ and $E(G)$ denote its vertex and edge sets, with respective cardinalities $v(G)$ and $e(G)$. If u and v are adjacent vertices then we denote the edge connecting them by uv or vu . A graph H is a subgraph of G , denoted $H \subseteq G$, if $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$.

Definition 2.1 (Colored subgraph isomorphism problem). For graphs G and X , where X comes with a coloring $\chi : V(X) \rightarrow V(G)$, the problem G -SUB(X) asks whether X has a subgraph G' such that χ (restricted to $V(G')$) is an isomorphism from G' to G . (Note that G' is not required to be an induced subgraph of X .)

For $U \subseteq V(G)$ let $G[U]$ be the induced subgraph of G on U , and more generally let $G[U_1, \dots, U_k] = G[U_1 \cup \dots \cup U_k]$. Let $G - U = G[V(G) - U]$, and for $H \subseteq G$ let $G - H = G - V(H)$.

When the parent graph G is clear in context, let $\deg(u)$ be the degree of a vertex u , and for disjoint $S, T \subseteq V(G)$ let $e(S, T)$ be the number of edges between S and T . Similarly, for vertex-disjoint graphs A and B let $e(A, B) = e(V(A), V(B))$.

Let $G \cap H$ be the graph with vertex set $V(G) \cap V(H)$ and edge set $E(G) \cap E(H)$, and define $G \cup H$ similarly. Note that $G \cap H$ may have isolated vertices even if G and H do not. If $A \subseteq B$ are graphs then let $[A, B] = \{H \mid A \subseteq H \subseteq B\}$, and let $(A, B]$ be the same interval without A , etc.

The Cartesian product of graphs G and H , denoted $G \square H$, has vertex set $V(G) \times V(H)$ and edges $(u, v_1)(u, v_2)$ for all $u \in V(G)$ and $v_1 v_2 \in E(H)$, and $(u_1, v)(u_2, v)$ for all $u_1 u_2 \in E(G)$ and $v \in V(H)$. Let G^d be the Cartesian product of d copies of G .

We denote by K_k the complete graph on k vertices, also called the k -clique. It follows that K_q^d has vertex set $[q]^d$, and two vertices are adjacent if and only if they differ in exactly one coordinate. Such graphs are called *Hamming graphs*. A special case is the d -dimensional hypercube $Q_d = K_2^d$; we will use $\{0, 1\}^d$ for its vertex set.

Definition 2.2 (Graph minor). A graph H is a *minor* of a graph G if there exists a *minor mapping* ϕ assigning a connected component of G to each vertex of H , such that $\phi(u)$ and $\phi(v)$ are vertex-disjoint for all $u \neq v$, and if $uv \in E(H)$ then there exists an edge in G with endpoints in $\phi(u)$ and $\phi(v)$.

In particular, any subgraph of G is also a minor of G (e.g. let ϕ be the identity).

Definition 2.3 (Treewidth). A *tree decomposition* of a graph G is a tree T whose vertices are subsets of $V(G)$ (called “bags”), such that each vertex and edge of G is contained in at least one of the bags, and for any $u \in V(G)$, the induced subgraph of T on the bags that contain u is a connected subtree of T . The *width* of T is one less than the size of the smallest bag, and the *treewidth* of G , denoted $tw(G)$, is the minimum width over all tree decompositions.

Roughly speaking, a graph has small treewidth if and only if it’s “similar to a tree”. See e.g. [Bod98; BK08] for further background, and [HW17] for a survey of parameters that are polynomially tied to treewidth.

The edge expansion of a graph G is defined as follows:

$$h(G) = \min_{\emptyset \subset U \subset V(G)} \frac{e(U, V(G) - U)}{\min(|U|, |V(G) - U|)}.$$

A bounded-degree expander is a graph with edge expansion $\Omega(1)$ and maximum degree $O(1)$ (see [HLW06] for a survey). Let $\lambda_i(G)$ be the i ’th largest eigenvalue of the adjacency matrix of G . We will use the following half of Cheeger’s Inequality:

Fact 2.4 ([AM85]). *If G is a d -regular graph then $h(G) \geq (d - \lambda_2(G))/2$.*

Finally, let $\mathbf{ER}(n, p)$ be the Erdős-Rényi graph on n vertices in which each possible edge exists independently with probability p .

3 The Average-Case Problem and the Parameter $\kappa(G)$

3.1 Threshold Random Graphs

First we will define *threshold weightings*, which assign weights to the vertices and edges of a graph subject to certain constraints. Then we will define a family of random graphs for each threshold weighting. The content in this subsection is essentially all from [LRR17].

Definition 3.1. A threshold weighting on a graph G is a pair $(\alpha, \beta) \in [0, 1]^{V(G)} \times [0, 2]^{E(G)}$ with the following property. For $H \subseteq G$ let $\alpha(H) = \sum_{u \in V(H)} \alpha(u)$ and $\beta(H) = \sum_{e \in E(H)} \beta(e)$, and let $\Delta(H) = \alpha(H) - \beta(H)$. Then, $\Delta(H) \geq 0$ for all $H \subseteq G$, and $\Delta(G) = 0$. Let $\theta(G)$ be the set of threshold weightings on G .

We will often denote $\Delta = (\alpha, \beta)$ in a slight abuse of notation. (Since $\Delta(u) = \alpha(u)$ if u is a single vertex, the pair (α, β) is uniquely determined by Δ .) The requirement that α be nonnegative is redundant because it's a special case of the requirement that Δ be nonnegative. The requirement that $\beta \leq 2$ is also redundant because for every edge uv ,

$$0 \leq \Delta(uv) = \alpha(u) + \alpha(v) - \beta(uv) \leq 2 - \beta(uv).$$

It will sometimes be convenient to define $\beta(e) = 0$ for $e \notin E(G)$, e.g. for disjoint sets $S, T \subseteq V(G)$ let $\beta(S, T) = \sum_{u \in S, v \in T} \beta(uv)$, and for vertex-disjoint $A, B \subseteq G$ let $\beta(A, B) = \beta(V(A), V(B))$.

A trivial example is $(\alpha, \beta) = (0, 0)$, i.e. all vertices and edges have a weight of zero. The following example is more general:

Example 3.2 (Markov Chains). Let $M \in \mathbb{R}_{\geq 0}^{V(G) \times V(G)}$ be a column stochastic matrix (meaning each column sums to 1) such that if $M_{u,v} \neq 0$ then either $u = v$ or $uv \in E(G)$. Let $\alpha(u) = 1 - M_{u,u}$ for all u , and $\beta(uv) = M_{u,v} + M_{v,u}$ for all $u \neq v$. Then for all $H \subseteq G$,

$$\Delta(H) = \sum_{\substack{v \in V(H) \\ uv \in E(G) - E(H)}} M_{u,v} \geq 0, \tag{1}$$

with equality if $H = G$. In fact, we prove that every threshold weighting is equivalent to at least one Markov Chain (Appendix A).

The following threshold weighting will be especially important, and can be thought of as representing a uniform random walk on G :

Definition 3.3. If G lacks isolated vertices then let $\Delta_o = (1, \beta_o) \in \theta(G)$ be the threshold weighting generated in Example 3.2 when $M_{u,v} = \mathbb{1}\{uv \in E(G)\} / \deg(v)$. That is, $\Delta_o = (\alpha, \beta)$, where $\alpha(u) = 1$ for all u and $\beta(uv) = 1/\deg(u) + 1/\deg(v)$ for all $u \neq v$. If G is d -regular then this simplifies to $\Delta_o = (1, \beta_o) = (1, 2/d)$.

Now we define threshold random graphs:

Definition 3.4. For $\Delta = (\alpha, \beta) \in \theta(G)$ let $\mathbf{X}_{\Delta, n}$ be the graph with vertices u_i for $u \in V(G)$ and $i \in [n^{\alpha(u)}]$, and for $uv \in E(G)$, each edge $u_i v_j$ independently with probability $n^{-\beta(uv)}$. The graph $\mathbf{X}_{\Delta, n}$ comes with the coloring to G defined by $u_i \mapsto u$.

For $H \subseteq G$ and X in the support of $\mathbf{X}_{\Delta, n}$, let $\text{Sub}_X(H)$ be the set of subgraphs $H' \subseteq X$ such that the aforementioned coloring (restricted to $V(H')$) is an isomorphism from H' to H . We say that such a graph H' is “ H -colored”. Note that $\text{Sub}_X(H)$ can be identified with a subset of $\prod_{u \in V(H)} [n^{\alpha(u)}]$.

Lemma 3.5. If $\Delta \in \theta(G)$ and $H \subseteq G$ then $\mathbb{E}[|\text{Sub}_{\mathbf{X}_{\Delta, n}}(H)|] = n^{\Delta(H)}(1 \pm o(1))$.

Proof. Let $(\alpha, \beta) = \Delta$. The set $\text{Sub}_{\mathbf{X}_{\Delta, n}}(H)$ contains each of its $n^{\alpha(H)}$ possible elements with probability $n^{-\beta(H)}$, so the result follows from linearity of expectation. (The $1 \pm o(1)$ accounts for having to round $n^{\alpha(\cdot)}$ to an integer.) \square

Lemma 3.5 motivates the requirements that Δ be nonnegative everywhere and that $\Delta(G) = 0$. Recall that the problem $G\text{-SUB}(X)$ asks whether $\text{Sub}_X(G)$ is the empty set. Since $\Delta(G)$ is required to be zero, it follows that $\text{Sub}_{\mathbf{X}_{\Delta,n}}(G)$ has (approximately) one element on average, and the probability that $\text{Sub}_{\mathbf{X}_{\Delta,n}}(G)$ is empty is known to be bounded away from 0 and 1 as n goes to infinity [LRR17].

3.2 The Parameter $\kappa(G)$ and an Algorithm for the Average Case

We now define $\kappa(G)$:

Definition 3.6 ([LRR17]). Let G be a graph with no isolated vertices. Let $\text{Seq}(G)$ be the set of *union sequences*, meaning sequences (H_1, \dots, H_k) of distinct subgraphs of G such that $H_k = G$ and each H_i is either an edge or the union of two previous graphs in the sequence. For $\Delta \in \theta(G)$ let $\kappa_\Delta(G) = \min_{S \in \text{Seq}(G)} \max_{H \in S} \Delta(H)$. Finally, let $\kappa(G) = \max_{\Delta \in \theta(G)} \kappa_\Delta(G)$.

To simplify the exposition, whenever we refer to $\kappa(G)$, the graph G is implicitly assumed to lack isolated vertices. Li et al. [LRR17] proved that for any fixed G , constant-depth circuits solving $G\text{-SUB}(\mathbf{X}_{\Delta,n})$ a.a.s. require size at least $n^{\kappa_\Delta(G)-o(1)}$ and at most $n^{2\kappa_\Delta(G)+c}$ (where c is an absolute constant). The results about average-case complexity described in Section 1 are with respect to a Δ such that $\kappa_\Delta(G) = \kappa(G)$.

Theorem 3.7. *The problem $G\text{-SUB}(\mathbf{X}_{\Delta,n})$ can be solved in time $\tilde{O}(n^{\kappa_\Delta(G)}) \leq \tilde{O}(n^{\kappa(G)})$ a.a.s. for any fixed G .*

Proof. First we prove a weaker upper bound of $\tilde{O}(n^{2\kappa_\Delta(G)})$, in a manner analogous to the circuit from [LRR17], and then we describe a modification (on Turing machines) that removes the factor of 2 from the exponent. Later on we will remove the factor of 2 in AC^0 using a different approach, summarized at the beginning of Section 6.

Let S be a union sequence such that $\kappa_\Delta(G) = \max_{H \in S} \Delta(H)$. For any $H \in S$, by Lemma 3.5 and Markov's Inequality, $P(|\text{Sub}_{\mathbf{X}_{\Delta,n}}(H)| > n^{\Delta(H)} \log n) \leq 1/\log n$. (We will obtain a tighter bound of $P(|\text{Sub}_{\mathbf{X}_{\Delta,n}}(H)| > \tilde{O}(n^{\Delta(H)})) \leq n^{-\omega(1)}$ in Section 6.1.) By a union bound it follows that if $X \sim \mathbf{X}_{\Delta,n}$ then $\max_{H \in S} |\text{Sub}_X(H)| \leq \tilde{O}(n^{\kappa_\Delta(G)})$ a.a.s. Assume this condition holds for X . For each successive H in S , compute $\text{Sub}_X(H)$ as follows. If H is a single edge then this is trivial. Otherwise $H = A \cup B$ for some previous $A, B \in S$, in which case $\text{Sub}_X(H)$ is the set of $\mathcal{A} \cup \mathcal{B}$ such that $\mathcal{A} \in \text{Sub}_X(A), \mathcal{B} \in \text{Sub}_X(B)$ and the projections of \mathcal{A} and \mathcal{B} onto $[n]^{V(A \cap B)}$ are equal. Therefore $\text{Sub}_X(H)$ can be computed by brute force in time $\tilde{O}(|\text{Sub}_X(A)| \cdot |\text{Sub}_X(B)|) \leq \tilde{O}(n^{2\kappa_\Delta(G)})$. Finally, check whether $\text{Sub}_X(G)$ is empty.

We can save a quadratic factor by computing $\text{Sub}_X(H)$ from $\text{Sub}_X(A)$ and $\text{Sub}_X(B)$ as follows. (This is a case of the *sort-merge join* algorithm for computing the *natural join* of two relations, as defined in database theory [SKS11].) Fix an efficiently computable total order on $[n]^{V(A \cap B)}$, e.g. interpret elements of $[n]^{V(A \cap B)}$ as $v(A \cap B)$ -digit base- n numbers in increasing order, and then define a partial order on $[n]^{V(A)} \cup [n]^{V(B)}$ by first projecting onto $[n]^{V(A \cap B)}$. Sort $\text{Sub}_X(A)$ and $\text{Sub}_X(B)$ in nondecreasing order, and for convenience add the symbol \perp to the end of both sorted lists. Let \mathcal{A} and \mathcal{B} be the first elements of $\text{Sub}_X(A)$ and $\text{Sub}_X(B)$ respectively, and initialize an empty accumulator (which will ultimately equal $\text{Sub}_X(H)$). While $\mathcal{A} \neq \perp$ and $\mathcal{B} \neq \perp$, do the following. If $\mathcal{A} < \mathcal{B}$ then let \mathcal{A} be the next element of $\text{Sub}_X(A)$. If $\mathcal{B} < \mathcal{A}$ then let \mathcal{B} be the next element of $\text{Sub}_X(B)$. Otherwise, let $\mathcal{B}' = \mathcal{B}$, and while $\mathcal{B}' \neq \perp$ and the projections of \mathcal{A} and \mathcal{B}' onto

$[n]^{V(A \cap B)}$ are equal, add $\mathcal{A} \cup \mathcal{B}'$ to the accumulator and let \mathcal{B}' be the next element of $\text{Sub}_X(B)$. Then (once the procedure involving \mathcal{B}' has finished) let \mathcal{A} be the next element of $\text{Sub}_X(A)$.

Sorting $\text{Sub}_X(A)$ and $\text{Sub}_X(B)$ takes $\tilde{O}(|\text{Sub}_X(A)| + |\text{Sub}_X(B)|)$ comparisons, and then computing $\text{Sub}_X(H)$ takes $\tilde{O}(|\text{Sub}_X(A)| + |\text{Sub}_X(B)| + |\text{Sub}_X(H)|) \leq \tilde{O}(n^{\kappa_\Delta(G)})$ time. \square

We will use the following graph-theoretic properties of $\kappa(G)$:

Theorem 3.8 ([LRR17]²). *Let G be a graph with no isolated vertices.*

- (i) *There exists $\Delta = (1, \beta) \in \theta(G)$ (meaning $\Delta(u) = 1$ for all vertices u) such that $\kappa(G) = \kappa_\Delta(G)$.*
- (ii) *$\kappa(G) \geq v(G)h(G)/(3 \max_{u \in V(G)} \deg(u))$, where $h(G)$ is the edge expansion of G .*
- (iii) *If G is a minor of some graph H then $\kappa(G) \leq \kappa(H)$.*

Li et al. [LRR17] observed the following as well:

Corollary 3.9. *If G is a bounded-degree expander then $\kappa(G)$ is $\Theta(v(G))$.*

Proof. Theorem 3.8(ii) implies that $\kappa(G)$ is $\Omega(v(G))$. Recall from Section 1 that $\kappa(G) \leq tw(G) + 1$ [LRR17], and it is well known that $tw(G) + 1 \leq v(G)$. \square

We note the following immediate consequence of Theorem 3.8(ii) and Fact 2.4:

Corollary 3.10. *If G is a d -regular graph then $\kappa(G) \geq v(G)(1 - \lambda_2(G)/d)/6$.*

4 The Parameter $emb(G)$ and Proof that $emb(G)$ is $O(\kappa(G))$

Recall that $emb(G)$ is significant because of its role in Marx's ETH-hardness result for G -SUB, namely Theorem 1.2.

Definition 4.1 ($emb(G)$). Let $G^{(q)}$ be the graph formed by replacing each vertex of G with a q -clique, i.e. it has vertices u_i for all $u \in V(G)$ and $i \in [q]$, and edges $u_i v_j$ for all $u_i \neq v_j$ such that either $u = v$ or $uv \in E(G)$. Let $emb(G)$ be the supremum of all $r > 0$ for which there exists $m_0 = m_0(G, r)$ such that if H is any graph with $m \geq m_0$ edges and no isolated vertices, then H is a minor of $G^{\lceil m/r \rceil}$, and furthermore a minor mapping from H to $G^{\lceil m/r \rceil}$ can be computed in time $f(G)m^{O(1)}$ for some function f .

Although the requirement that such a minor mapping be efficiently computable is crucial in Theorem 1.2, none of the other results about $emb(G)$ that we reference or derive depend on this requirement, so we may safely ignore it going forward. The following example illustrates Definition 4.1:

Example 4.2 ($emb(K_k)$ [Mar10]). Since $K_k^{\lceil m/r \rceil} = K_{k \lceil m/r \rceil}$, any graph H with m edges is a minor of $K_k^{\lceil m/r \rceil}$ if and only if $v(H) \leq k \lceil m/r \rceil$. If H has no isolated vertices then H could have up to $2m$ vertices, so $2m \leq k \lceil m/r \rceil$. Therefore $emb(K_k) = k/2$: it is sufficient for $2m$ to be at most km/r (i.e. $r \leq k/2$), and no $r > k/2$ satisfies $2m \leq k \lceil m/r \rceil$ for arbitrarily large m .

²Specifically, Corollary 4.2, Theorem 4.9, and Theorem 5.1 of [LRR17] correspond to Theorems 3.8(i) to 3.8(iii) respectively.

Remark. The name $emb(G)$ comes from the fact that Marx [Mar10] called a minor mapping from H to $G^{(q)}$ an “embedding of depth q ” from H into G . Marx used the notation $G^{(q)}$, but the parameter $emb(G)$ is new in the current paper, all results about $emb(G)$ in [Mar10; AM11] having been stated in terms of embeddings of some depth.

The following is used in our proof that $emb(G)$ is $O(\kappa(G))$:

Lemma 4.3. $\kappa(G^{(q)}) \leq q \max(\kappa(G), 2)$.

Proof. Let $\Delta = (\alpha, \beta) \in \theta(G^{(q)})$ such that $\kappa(G^{(q)}) = \kappa_\Delta(G^{(q)})$. Define a threshold weighting $\Delta' = (\alpha', \beta') \in \theta(G)$ as follows: For all $u \in V(G)$ and $uv \in E(G)$,

$$\alpha'(u) = \frac{\Delta(u^{(q)})}{q} = \frac{1}{q} \left(\sum_{i=1}^q \alpha(u_i) - \sum_{1 \leq i < j \leq q} \beta(u_i u_j) \right),$$

$$\beta'(uv) = \frac{1}{q} \sum_{i=1}^q \sum_{j=1}^q \beta(u_i v_j).$$

This is a threshold weighting because if $H \subseteq G$ then $\Delta'(H) = \Delta(H^{(q)})/q \geq 0$, with equality if $H = G$. It’s also normalized to $\alpha' \leq 1$.

Let S' be an optimal union sequence for G with respect to Δ' . Construct a union sequence S for $G^{(q)}$ as follows:

1. For each $e \in E(G)$ append an arbitrary union sequence for $e^{(q)}$.
2. For each $H \in S'$ (in order) append $H^{(q)}$.

If $H \subseteq e^{(q)}$ then $\Delta(H) \leq \alpha(e^{(q)}) \leq 2q$, and we’ve already seen that $\Delta(H^{(q)}) = q\Delta'(H)$ for all $H \in S'$. Therefore,

$$\kappa(G^{(q)}) = \kappa_\Delta(G^{(q)}) \leq \max_{H \in S} \Delta(H) \leq \max \left(2q, \max_{H \in S'} q\Delta'(H) \right) = q \max(\kappa_{\Delta'}(G), 2) \leq q \max(\kappa(G), 2). \quad \square$$

Now we prove that $emb(G)$ is $O(\kappa(G))$ (Theorem 1.5), using an argument similar to the proof by Marx [Mar10] that $emb(G)$ is $O(tw(G))$:

Proof. Let $r > 0$, and assume there exists an arbitrarily large 3-regular expander H that’s a minor of $G^{\lceil e(H)/r \rceil}$. Then by Corollary 3.9, Theorem 3.8(iii), and Lemma 4.3,

$$e(H) = \Theta(v(H)) = \Theta(\kappa(H)) \leq O \left(\kappa \left(G^{\lceil e(H)/r \rceil} \right) \right) \leq O(\kappa(G)e(H)/r),$$

so r must be $O(\kappa(G))$. □

Li et al. [LRR17] posed the question of whether Theorem 1.2 holds with $\kappa(G)$ in place of $emb(G)$. By Theorem 1.5 this would be a stronger bound, which makes the question even more interesting. This problem is open even in the case of 3-regular expanders: recall from Section 1 that if G is a 3-regular expander then $emb(G)$ is $\Theta(tw(G)/\log tw(G))$ and $\kappa(G)$ is $\Theta(tw(G))$ [AM11; LRR17].

The fact that $\kappa(G)$ is $\Omega(emb(G))$ gives an alternate proof, besides the one in [LRR17], that $\kappa(G)$ is $\Omega(tw(G)/\log tw(G))$.

5 Separating κ from Treewidth

In Section 5.1 we prove that $\kappa(K_k) = k/4 + O(1)$, which is a special case of the more general result that $\kappa(K_q^d) = \Theta(q^d/d)$. We obtain tighter multiplicative constants in the case $d = 1$, and it provides an opportunity to illustrate the main ideas of our proof in a simpler setting, but it may be skipped without penalty. In Section 5.2 we prove that $\kappa(K_q^d)$ is $O(q^d/d)$ when q is even, which is sufficient to separate κ from treewidth. Again, this case is cleaner than the general case and conveys most of the intuition behind it. In Appendix B we prove that $\kappa(K_q^d)$ is $O(q^d/d)$ for all q . In Section 5.3 we prove that $\kappa(K_q^d)$ is $\Omega(q^d/d)$ in two different ways, completing the proof that $\kappa(K_q^d)$ is $\Theta(q^d/d)$ (Theorem 1.4), and we obtain as a corollary that $emb(K_q^d)$ is $\Theta(q^d/d)$ as well. In Section 5.4 we summarize the proof of Chandran and Kavitha [CK06] that $tw(K_q^d)$ is $\Theta(q^d/\sqrt{d})$.

5.1 Proof that $\kappa(K_k) = k/4 + O(1)$

Remark. It was already observed in [LRR17] that $\kappa(K_k)$ is $\Theta(k)$.

Rosman [Ros08] proved that $\kappa_{\Delta_o}(K_k) \geq k/4$, so it suffices to prove the upper bound. By Theorem 3.8(i) it suffices to prove that $\kappa_{\Delta}(K_k) \leq k/4 + O(1)$ for an arbitrary $\Delta = (1, \beta) \in \theta(G)$. First we construct a sequence $U_1 \subseteq \dots \subseteq U_k = V(K_k)$ such that U_i is an i -element subset of $V(K_k)$, and $\beta(K_k[U_i]) \geq \beta_o(K_k[U_i])$ for all i . The set $U_k = V(K_k)$ satisfies this requirement because $\beta(K_k)$ and $\beta_o(K_k)$ are both equal to k . Given U_i , let \mathbf{U}_{i-1} be an $(i-1)$ -element subset of U_i chosen uniformly at random. Each pair of elements in U_i is included in \mathbf{U}_{i-1} with the same probability $p_i (= 1 - 2/i)$, so it follows from linearity of expectation that

$$\mathbb{E}[\beta(K_k[\mathbf{U}_{i-1}])] = \sum_{e \in E(K_k[U_i])} \beta(e)p_i = p_i \beta(K_k[U_i]) \geq p_i \beta_o(K_k[U_i]) = \mathbb{E}[\beta_o(K_k[\mathbf{U}_{i-1}])].$$

Therefore there exists a fixed U_{i-1} such that $\beta(K_k[U_{i-1}]) \geq \beta_o(K_k[U_{i-1}])$.

We construct a union sequence S for K_k as follows. Start by enumerating the edges, and then for i from 1 to $k-1$, append $(K_k[U_i] \cup e_1, K_k[U_i] \cup e_1 \cup e_2, \dots, K_k[U_{i+1}])$, where e_1, e_2, \dots are the edges between U_i and $U_{i+1} - U_i$. Then,

$$\max_{H \in S} \Delta(H) \leq \max_i \Delta(K_k[U_i]) + 1 \leq \max_i \Delta_o(K_k[U_i]) + 1.$$

As observed in [Ros08], it follows from Eq. (1) that $\Delta_o(K_k[U_i]) = i(k-i)/k$, which is at most $k/4$ (when $i = k/2$). Therefore $\kappa_{\Delta}(K_k) \leq k/4 + 1$.

5.2 Proof that $\kappa(K_q^d)$ is $O(q^d/d)$ if q is Even

First we reduce this to the case $q = 2$. The graph K_q^d is a subgraph of $Q_d^{((q/2)^d)}$ (recall Definition 4.1), as evidenced by the following argument. Let $\phi_L : [q] \rightarrow \{0, 1\}$ and $\phi_R : [q] \rightarrow [q/2]$ such that $\phi_L \times \phi_R$ is a bijection from $[q]$ to $\{0, 1\} \times [q/2]$, and let $\psi : [q/2]^d \rightarrow [(q/2)^d]$ be another arbitrary bijection. Then the following map is an injective homomorphism from K_q^d to $Q_d^{((q/2)^d)}$:

$$(x_1, \dots, x_d) \mapsto ((\phi_L(x_1), \dots, \phi_L(x_d)), \psi(\phi_R(x_1), \dots, \phi_R(x_d))).$$

By Theorem 3.8(iii) and Lemma 4.3, if $\kappa(Q_d)$ is $O(2^d/d)$ then

$$\kappa\left(K_q^d\right) \leq \kappa\left(Q_d^{\left(\left(\frac{q}{2}\right)^d\right)}\right) \leq O\left(\left(\frac{q}{2}\right)^d \kappa\left(Q_d\right)\right) \leq O\left(\left(\frac{q}{2}\right)^d \frac{2^d}{d}\right) = O\left(q^d/d\right).$$

Now we prove that $\kappa(Q_d)$ is $O(2^d/d)$, following some brief definitions and a high-level overview of the argument. Fix d . We identify each $u \in \{0,1\}^d$ with $\sum_{i=0}^{d-1} u_i 2^i$. For $0 \leq a \leq 2^d$ let $G(a) = Q_d[0, \dots, a-1]$. Recall that $\Delta_o = (1, \beta_o) = (1, 2/d)$ is a threshold weighting on Q_d (Definition 3.3). Let $\mu = \max_{0 \leq a \leq 2^d} \Delta_o(G(a))$.

Remark. The intuition behind μ is as follows. The reader may note that $\kappa_{\Delta_o}(Q_d) \leq \mu + 1$, by reasoning analogous to that in Section 5.1. That is, for each vertex u of Q_d in increasing lexicographic order, add to an accumulator all edges uv for which $v < u$.

There is another union sequence captured by μ as well. If a subgraph $B \subseteq Q_d$ isomorphic to Q_k for some k , then since Q_k is isomorphic to $G(2^k)$ (and β_o is uniform) it follows that $\Delta_o(B) \leq \mu$. Consider a depth- d binary tree in which each node at depth k is a subgraph of Q_d isomorphic to Q_{d-k} (in particular, the root is Q_d and the leaves are vertices), and each interior node is the union of its two children along with some additional edges corresponding to a coordinate cut. This tree describes a union sequence S for Q_d : recursively obtain the graphs L and R corresponding to the children of Q_d , and then take $L \cup R$ and add the missing edges. Note that $\max_{H \in S} \Delta_o(H) = 2 \max_{0 \leq k \leq d} \Delta_o(G(2^k)) \leq 2\mu$.

Analogous to Section 5.1, the upper bound is obtained by comparing $\kappa_{\Delta}(Q_d)$ to μ for each Δ , and bounding μ . For this purpose we will consider the two union sequences mentioned above, as well as hybrids of them.

The proof is as follows:

$$\begin{aligned} \kappa(Q_d) &= \max_{\beta} \kappa_{(1,\beta)}(Q_d) && \text{Theorem 3.8(i)} \\ &\leq 2\mu && \text{Lemma 5.1 (below) with } a = 0 \text{ and } k = d \\ &< 4/3 \cdot 2^d/d. && \text{Lemma 5.2 (below)} \end{aligned}$$

For each threshold weighting $\Delta \in \theta(Q_d)$, it will be convenient in the following to generalize κ_{Δ} to subgraphs $H \subseteq Q_d$ by $\kappa_{\Delta}(H) = \min_{S \in \text{Seq}(H)} \max_{F \in S} \Delta(F)$. (This is a nontrivial generalization of the definition of κ_{Δ} , because if $\Delta(H) > 0$ then the restriction of Δ to subgraphs of H is not a threshold weighting on H .) Also if v is a single vertex then let $\kappa_{\Delta}(v) = 0$.

Lemma 5.1. *Let $0 \leq a \leq 2^d$ and $0 \leq k \leq d$ such that 2^k divides a . Let $\Delta = (1, \beta) \in \theta(Q_d)$ such that $\beta(G(a)) \geq \beta_o(G(a))$ and $\beta(G(a + 2^k)) \geq \beta_o(G(a + 2^k))$, and $\kappa_{\Delta}(G(a)) \leq 2\mu$. Then $\kappa_{\Delta}(G(a + 2^k)) \leq 2\mu$.*

Proof. The proof is by induction on k . The inductive hypothesis will actually be (slightly) stronger in the following way: given a labeling of the vertices of Q_d with the elements of $\{0,1\}^d$, the labels can be rearranged according to any of the $2^d d!$ isomorphisms of Q_d , and the inductive hypothesis is required to hold with respect to any such labeling. (The value of μ doesn't depend on the labeling used in its definition because of the symmetry of β_o .)

Let $B = G(a + 2^k) - G(a)$. Since 2^k divides a , it follows that B is isomorphic to Q_k . In the inductive step we handle separately the cases where $\beta(B) \geq \beta_o(B)$ and $\beta(B) < \beta_o(B)$. The base

case is a special case of the former because if B is a single vertex then $\beta(B)$ and $\beta_o(B)$ are both zero.

Case 1: $\beta(B) \geq \beta_o(B)$. If $k = 0$ then $\kappa_\Delta(B) = 0 \leq 2\mu$; we now obtain the same result in the case where $k > 0$. For $0 \leq i < k$ and $b \in \{0, 1\}$ let $B(i, b) = B[v \in V(B) \mid v_i = b]$. Choose $\mathbf{i} \in \{0, \dots, k-1\}$ and $\mathbf{b} \in \{0, 1\}$ independently and uniformly at random. By symmetry, each $e \in E(B)$ is in $B(\mathbf{i}, \mathbf{b})$ with the same probability p . (Specifically, $p = (k-1)/2k$: For any edge uv , there is a unique index i in which u and v differ. If $\mathbf{i} = i$ then exactly one of u and v is in $B(\mathbf{i}, \mathbf{b})$; otherwise uv is in $B(\mathbf{i}, \mathbf{b})$ with probability $1/2$ depending on \mathbf{b} .) By linearity of expectation,

$$\mathbb{E}[\beta(B(\mathbf{i}, \mathbf{b}))] = \sum_{e \in E(B)} P(e \in B(\mathbf{i}, \mathbf{b}))\beta(e) = p\beta(B).$$

Similarly, $\mathbb{E}[\beta_o(B(\mathbf{i}, \mathbf{b}))] = p\beta_o(B)$. By our assumption that $\beta(B) \geq \beta_o(B)$,

$$\mathbb{E}[\beta(B(\mathbf{i}, \mathbf{b}))] = p\beta(B) \geq p\beta_o(B) = \mathbb{E}[\beta_o(B(\mathbf{i}, \mathbf{b}))].$$

Therefore there exist fixed i and b such that $\beta(B(i, b)) \geq \beta_o(B(i, b))$.

Now our claim that $\kappa_\Delta(B) \leq 2\mu$ follows from two applications of the inductive hypothesis. Since we required the inductive hypothesis to hold for all labelings of Q_d , we can assume without loss of generality that $i = k-1$ and $b = 0$. Ignoring $G(a)$, an application of the inductive hypothesis with $a' = 0$ and $k' = k-1$ reveals that $\kappa_\Delta(B(i, b)) \leq 2\mu$, and then a second application of the inductive hypothesis with $a'' = 2^{k-1}$ and $k'' = k-1$ reveals that $\kappa_\Delta(B) \leq 2\mu$.

Let S be an optimal (with respect to Δ) union sequence for $G(a)$, followed by an optimal union sequence for B , followed by $G(a) \cup B, G(a) \cup B \cup e_1, \dots, G(a) \cup B \cup \{e_j\}$, where the $\{e_j\}$ are the edges between $G(a)$ and B in Q_d . (If B is a single vertex then omit $G(a) \cup B$ from the sequence.) Then,

$$\max_{H \in S} \Delta(H) \leq \max(\kappa_\Delta(G(a)), \kappa_\Delta(B), \Delta(G(a)) + \Delta(B)).$$

We proceed to bound each of these three terms by 2μ , completing the proof. We have assumed that $\kappa_\Delta(G(a)) \leq 2\mu$, and proved that $\kappa_\Delta(B) \leq 2\mu$. We have also assumed that $\beta(G(a)) \geq \beta_o(G(a))$, and since Δ and Δ_o both evaluate to 1 on all vertices, it follows that $\Delta(G(a)) \leq \Delta_o(G(a)) \leq \mu$ (with the last step following from the definition of μ). Similarly, since B is isomorphic to $G(2^k)$ it follows that $\Delta(B) \leq \Delta_o(B) \leq \mu$. Therefore $\Delta(G(a)) + \Delta(B) \leq 2\mu$.

Case 2: $\beta(B) < \beta_o(B)$. For $i < k$ and $b \in \{0, 1\}$ let $H(i, b) = Q_d[0, \dots, a-1, V(B(i, b))]$ (where $B(i, b)$ is defined as above). Choose $\mathbf{i} < k$ and $\mathbf{b} \in \{0, 1\}$ independently and uniformly at random. Note that $\beta(G(a + 2^k)) = \beta(G(a)) + \beta(G(a), B) + \beta(B)$. By reasoning similar to that in the previous case (and applying our various assumptions),

$$\begin{aligned} \mathbb{E}[\beta(H(\mathbf{i}, \mathbf{b}))] &= \beta(G(a)) + \frac{1}{2}\beta(G(a), B) + \frac{k-1}{2k}\beta(B) \\ &= \frac{1}{2}\beta(G(a)) + \frac{1}{2}\beta(G(a + 2^k)) - \frac{1}{2k}\beta(B) \\ &> \frac{1}{2}\beta_o(G(a)) + \frac{1}{2}\beta_o(G(a + 2^k)) - \frac{1}{2k}\beta_o(B) \\ &= \mathbb{E}[\beta_o(H(\mathbf{i}, \mathbf{b}))]. \end{aligned}$$

Therefore $\beta(H(i, b)) > \beta_o(H(i, b))$ for some fixed i and b .

Assume without loss of generality that $i = k - 1$ and $b = 0$; then $H(i, b) = G(a + 2^{k-1})$. Applying the inductive hypothesis with $a' = a$ and $k' = k - 1$ reveals that $\kappa_\Delta(G(a + 2^{k-1})) \leq 2\mu$, and then applying the inductive hypothesis with $a'' = a + 2^{k-1}$ and $k'' = k - 1$ reveals that $\kappa_\Delta(G(a + 2^k)) \leq 2\mu$. \square

Lemma 5.2. $\mu < 2/3 \cdot 2^d/d$.

Proof. For any $0 \leq a \leq 2^d$, it follows from Eq. (1) that $\Delta_o(G(a)) = e(G(a), Q_d - G(a))/d$, so it suffices to prove that $e(G(a), Q_d - G(a)) < 2^{d+1}/3$ for all a . Let $G(a, b) = Q_d[a, \dots, b - 1]$. Since

$$e(G(0, a), G(a, 2^d)) = e(G(0, 2^d - a), G(2^d - a, 2^d))$$

(as can be seen by applying the automorphism $(x_1, \dots, x_d) \mapsto (1 - x_1, \dots, 1 - x_d)$ to Q_d), we can restrict our search to $a \in [0, 2^{d-1}]$. In that case,

$$\begin{aligned} e(G(0, a), G(a, 2^d)) &= e(G(0, a), G(a, 2^{d-1})) + e(G(0, a), G(2^{d-1}, 2^d)) \\ &= e(G(0, a), G(a, 2^{d-1})) + v(G(0, a)) \\ &= e(G(0, a), G(a, 2^{d-1})) + a. \end{aligned}$$

By the same reasoning,

$$e(G(0, 2^{d-1} - a), G(2^{d-1} - a, 2^d)) = e(G(0, 2^{d-1} - a), G(2^{d-1} - a, 2^{d-1})) + 2^{d-1} - a.$$

Since $e(G(0, a), G(a, 2^{d-1})) = e(G(0, 2^{d-1} - a), G(2^{d-1} - a, 2^{d-1}))$ (consider a similar automorphism), it follows that if $a < 2^{d-2}$ then

$$e(G(0, 2^{d-1} - a), G(2^{d-1} - a, 2^d)) - e(G(0, a), G(a, 2^d)) = 2^{d-1} - 2a > 0.$$

Therefore we can restrict our search to $a \in [2^{d-2}, 2^{d-1}]$, in which case

$$\begin{aligned} e(G(0, a), G(a, 2^d)) &= e(G(0, a), G(a, 2^{d-1})) + a \\ &= e(G(2^{d-2}, a), G(a, 2^{d-1})) + e(G(0, 2^{d-2}), G(a, 2^{d-1})) + a \\ &= e(G(2^{d-2}, a), G(a, 2^{d-1})) + 2^{d-1}. \end{aligned}$$

By induction it follows that $\mu = 2^{d-1} + 2^{d-3} + 2^{d-5} + \dots + (2 \text{ or } 1) < 2^{d+1}/3$. \square

Remark. Harper [Har04] proved that out of all subgraphs of Q_d with a vertices, $G(a) = G(0, a)$ has the fewest outgoing edges [Fil15].

5.3 Proof that $\kappa(K_q^d)$ is $\Omega(q^d/d)$ and $\text{emb}(K_q^d)$ is $\Theta(q^d/d)$

Alon and Marx [AM11, Theorem 4.3] proved that $\text{emb}(K_q^d)$ is $\Omega(q^d/d)$, and it follows from Theorem 1.5 that $\text{emb}(K_q^d) \leq O(\kappa(K_q^d)) \leq O(q^d/d)$. Therefore $\text{emb}(K_q^d)$ is $\Theta(q^d/d)$.

It is implicit in the above argument that $\kappa(K_q^d) \geq \Omega(\text{emb}(K_q^d)) \geq \Omega(q^d/d)$; we now present an alternate proof that $\kappa(K_q^d)$ is $\Omega(q^d/d)$ based on edge expansion. Since K_q^d is $d(q-1)$ -regular, by Corollary 3.10 it suffices to prove that $1 - \lambda_2(K_q^d)/d(q-1)$ is $\Omega(1/d)$. We use the following well-known fact, where graphs are identified with their adjacency matrices:

Fact 5.3. *The eigenvalues of $G \square H$ are $\lambda_i(G) + \lambda_j(H)$ for $i \in [v(G)], j \in [v(H)]$.*

Proof. Observe that $G \square H = G \otimes I + I \otimes H$, where the symbols \otimes and I denote the tensor product and the identity matrix respectively. Let u_i (resp. w_i) be the i 'th eigenvector of G (resp. H); clearly $u_i \otimes w_j$ is an eigenvector of $G \square H$ with eigenvalue $\lambda_i(G) + \lambda_j(H)$. Since a real symmetric matrix (in particular G or H) has an orthogonal eigenbasis, it follows that the $u_i \otimes w_j$ are also orthogonal. Since $v(G \square H) = v(G)v(H)$, there are no other eigenvalues of $G \square H$. \square

Since $\lambda_i(K_q)$ equals $q - 1$ if $i = 1$ and -1 otherwise, repeated application of Fact 5.3 reveals that $\lambda_2(K_q^d) = (q - 1)(d - 1) - 1 = d(q - 1) - q$, so $1 - \lambda_2(K_q^d) / d(q - 1) = q / (q - 1)d$, as desired.

Remark. Fact 2.4 is an equality in the case of hypercubes (see e.g. [HLW06]): let $i \in [d]$ and define a cut by partitioning the vertices according to the values of their i 'th coordinates. So for hypercubes, all slack in the application of Corollary 3.10 comes from Theorem 3.8(ii).

5.4 Proof that $tw(K_q^d)$ is $\Theta(q^d / \sqrt{d})$, Summarized

(See [CK06] for the full proof.) The proof that $tw(K_q^d)$ is $O(q^d / \sqrt{d})$ reduces to the case $q = 2$ by reasoning analogous to that in the beginning of Section 5.2. For $k \in [d]$ let U_k be the set of vertices of Q_d with exactly k or $k - 1$ ones. The path (U_1, \dots, U_d) is a tree decomposition of Q_d with width approximately $2\binom{d}{d/2}$, and by Stirling's approximation this is $\Theta(2^d / \sqrt{d})$.³

For a graph G let $\phi(G)$ be the minimum over all $U \subseteq V(G), v(G)/4 \leq |U| \leq v(G)/2$ of the number of vertices in $V(G) - U$ with at least one neighbor in U . From a result of Robertson and Seymour [RS86] it follows that $tw(G) \geq \phi(G) - 1$, and from a result of Harper [Har99] it follows that $\phi(K_q^d)$ is $\Omega(q^d / \sqrt{d})$. (Also note the parallels between $tw(G) \geq \phi(G) - 1$ and Theorem 3.8(ii); interestingly, we've seen that both are tight to within a constant factor in the case of K_q^d .)

6 AC⁰ Upper Bound

An AC⁰ circuit is a constant-depth circuit with polynomially many unbounded-fanin AND and OR gates and NOT gates. Fix a graph G and threshold weighting $\Delta \in \theta(G)$ for the remainder of this section. We prove the following, which is a more precise statement of Theorem 1.3:

Theorem 6.1. *There exists a constant-depth circuit with $n^{\kappa_\Delta(G)+c}$ wires that solves $G\text{-SUB}(\mathbf{X}_{\Delta,n})$ with probability $1 - n^{-\omega(1)}$, where $c > 0$ is an absolute constant.*

Since in any circuit the number of gates is at most one plus the number of wires, the circuit from Theorem 6.1 has size $n^{\kappa_\Delta(G)+O(1)} \leq n^{\kappa(G)+O(1)}$. (In this discussion, all $\pm O(1)$ terms in an exponent are independent of G .) For comparison, it was proved in [LRR17] (building on a line of previous work [Ros08; Ama10; Ros10; NW11]) that the average-case AC⁰ complexity of $G\text{-SUB}(\mathbf{X}_{\Delta,n})$ is between $n^{\kappa_\Delta(G)-o(1)}$ and $n^{2\kappa_\Delta(G)+O(1)}$. Another related result, regarding the uncolored k -clique problem, is that the average-case AC⁰ complexity of $K_k\text{-SUB}_{\text{uncol}}(\mathbf{ER}(n, n^{-2/(k-1)}))$ is at most $n^{k/4+O(1)}$ [Ama10; Ros14] ($= n^{\kappa(K_k)\pm O(1)}$ by Section 5.1). See [Ros18] for a survey of the average-case circuit complexity of subgraph isomorphism more generally.

³Compared to the tree decomposition from [CK06], this one is a simpler variant whose width is larger by up to a constant factor.

One challenge to implementing the algorithm behind Theorem 3.7 in AC^0 is that sorting cannot be done in AC^0 [Hås86]. The $n^{2\kappa_\Delta(G)+O(1)}$ -size circuit from [LRR17] computes $\text{Sub}_X(A \cup B)$ by finding the relevant pairs in $\text{Sub}_X(A) \times \text{Sub}_X(B)$ by brute force with $\tilde{O}(|\text{Sub}_X(A)| \cdot |\text{Sub}_X(B)|)$ gates. Our circuit differs in that we represent $\text{Sub}_X(H)$ as a depth- $v(H)$ tree, where the non-root vertices are assigned labels in $[n]$, and the (sequences of labels along the) paths from the root to the leaves correspond to the elements of $\text{Sub}_X(H)$. This will allow us to compute $\text{Sub}_X(A \cup B)$ with high probability given $\text{Sub}_X(A)$ and $\text{Sub}_X(B)$, on a constant-depth circuit of size nearly linear in $|\text{Sub}_X(A)| + |\text{Sub}_X(B)|$. A key fact in our construction is that constant-depth circuits can (with high probability) convert between representations of $\text{Sub}_X(H)$ corresponding to different orderings of $V(H)$.

Our construction requires fairly precise estimates for how many children to assign each node. Luckily this number is highly concentrated around its mean if the input graph is $\mathbf{X}_{\Delta,n}$. This result will follow from the concentration inequality below, whose statement requires several definitions:

Definition 6.2. Let X be in the support of $\mathbf{X}_{\Delta,n}$, and let $U \subseteq G$ be an arbitrary graph (which we think of as a “universe”). Let $\text{Sub}_n(U)$ be the set of all possible elements of $\text{Sub}_{\mathbf{X}_{\Delta,n}}(U)$; note that this can be identified with $\prod_{v \in V(U)} [n^{\alpha(v)}]$. If $A \subseteq U$ and $\mathcal{A} \in \text{Sub}_n(A)$ then let \mathcal{A} *extend to U in X* if there exists a graph $\mathcal{U} \in \text{Sub}_X(U)$ (called a *U -extension of \mathcal{A}*) such that $\mathcal{A} \subseteq \mathcal{U}$. (In context, X or \mathbf{X} will be implicit.) Equivalently, \mathcal{A} could be required to be in $\text{Sub}_X(A)$ rather than $\text{Sub}_n(A)$ in the latter definition.

Let $\Delta_U^*(A) = \min_{A \subseteq H \subseteq U} \Delta(H)$. Let X be *good* if for all graphs $U \subseteq G$ and $A \subseteq U$, and for all $\mathcal{A} \in \text{Sub}_n(A)$ and vertices $v \in V(U) - V(A)$, there are $\tilde{O}(n^{\Delta_U^*(A \cup v) - \Delta_U^*(A)})$ values of $i \in [n^{\alpha(v)}]$ such that $\mathcal{A} \cup v_i$ extends to U . (Recall our unconventional definition of $\tilde{O}(\cdot)$ from Section 2, e.g. $\tilde{O}(1)$ denotes $\log^{O(1)} n$.) Finally, let an event occur *with high probability (w.h.p.)* if it occurs with probability $1 - n^{-\omega(1)}$.

We prove the following:

Theorem 6.3. *The graph $\mathbf{X}_{\Delta,n}$ is good w.h.p.*

Observe that this is a substantially stronger concentration bound than the application of Markov’s Inequality in the proof of Theorem 3.7. In Section 6.1 we prove Theorem 6.3, and then in Section 6.2 we use this result to prove Theorem 6.1. Both proofs use the following concentration inequality, which is proved by a Chernoff bound:

Lemma 6.4. *If $\mathbf{S} = \mathbf{S}(n)$ is a sum of independent Bernoulli random variables then w.h.p. $\mathbf{S} \leq \max(\mathbb{E}[\mathbf{S}], 1) \cdot \tilde{O}(1)$.*

Proof. Let $\mathbf{S} = \sum_i \mathbf{B}_i$ be a decomposition of \mathbf{S} as a sum of independent Bernoulli random variables. Let $p_i = \mathbb{E}[\mathbf{B}_i]$ and $\mu = \mathbb{E}[\mathbf{S}] = \sum_i p_i$. Then for $r, t \geq 0$,

$$\begin{aligned} P(\mathbf{S} \geq r) &= P(\exp(t\mathbf{S}) \geq \exp(tr)) \leq \exp(-tr) \prod_i \mathbb{E}[\exp(t\mathbf{B}_i)] \\ &= \exp(-tr) \prod_i (1 - p_i + p_i e^t) \leq \exp(-tr) \prod_i \exp(p_i e^t) \\ &= \exp(-tr + \mu e^t). \end{aligned}$$

Letting $t = \log(r/\mu)$ gives $P(\mathbf{S} \geq r) \leq (e\mu/r)^r$ assuming $t \geq 0$, and then (for example) letting $r = \max(\mu, 1) \log^2 n$ gives, for sufficiently large n ,

$$P(\mathbf{S} \geq r) \leq (e/\log^2 n)^{\log^2 n} \leq (1/e)^{\log^2 n} = n^{-\log n}. \quad \square$$

6.1 Proof of Theorem 6.3

First we derive some algebraic properties of the threshold weighting Δ .

Lemma 6.5. *If $A, B \subseteq G$ then $\Delta(A) + \Delta(B) = \Delta(A \cap B) + \Delta(A \cup B)$.*

Proof. Each vertex or edge in one (resp. two) of A and B is also in one (resp. two) of $A \cap B$ and $A \cup B$. \square

Definition 6.6. For $A \subseteq U \subseteq G$ let $\Gamma_U(A) = \bigcap \{H \in [A, U] \mid \Delta(H) = \Delta_U^*(A)\}$, and let A be a U -base if $\Delta(A) = \Delta_U^*(A)$.

Throughout this subsection, U will be an arbitrary subgraph of G unless additional structure is imposed on it, and missing subscripts on Δ^* and Γ default to U .

Lemma 6.7. *If $A \subseteq U$ then $\Delta(\Gamma(A)) = \Delta^*(A)$ and $A \subseteq \Gamma(A)$.*

Proof. It suffices to show that the set $S = \{H \in [A, U] \mid \Delta(H) = \Delta^*(A)\}$ is closed under intersection. Let $B, C \in S$. By the definition of S , Lemma 6.5, and the fact that $A \subseteq B \cup C$,

$$2\Delta^*(A) = \Delta(B) + \Delta(C) = \Delta(B \cap C) + \Delta(B \cup C) \geq \Delta(B \cap C) + \Delta^*(A),$$

so $\Delta(B \cap C) \leq \Delta^*(A)$. On the other hand, $\Delta(B \cap C) \geq \Delta^*(A)$ because $A \subseteq B \cap C$. Therefore $\Delta(B \cap C) = \Delta^*(A)$, so $B \cap C \in S$. \square

Lemma 6.8. *If $A \subseteq \Gamma(A) \subseteq U' \subseteq U$ then $\Gamma(A)$ is a U' -base.*

Proof. Since the interval $[A, U]$ includes the interval $[\Gamma(A), U']$, it follows from Lemma 6.7 that $\Delta(\Gamma(A)) = \Delta_U^*(A) \leq \Delta_{U'}^*(\Gamma(A)) \leq \Delta(\Gamma(A))$. Therefore $\Delta(\Gamma(A)) = \Delta_{U'}^*(\Gamma(A))$. \square

Lemma 6.9. *If $A \subseteq B \subseteq U$ then $\Gamma(A) \subseteq \Gamma(B)$.*

Proof. Since $B \subseteq \Gamma(B) \subseteq \Gamma(B) \cup \Gamma(A)$ it follows that $\Delta^*(B) \leq \Delta(\Gamma(B) \cup \Gamma(A))$, so by Lemmas 6.5 and 6.7,

$$\Delta(\Gamma(A) \cap \Gamma(B)) + \Delta^*(B) \leq \Delta(\Gamma(A) \cap \Gamma(B)) + \Delta(\Gamma(A) \cup \Gamma(B)) = \Delta(\Gamma(A)) + \Delta(\Gamma(B)) = \Delta^*(A) + \Delta^*(B).$$

Therefore $\Delta^*(A) \geq \Delta(\Gamma(A) \cap \Gamma(B))$. On the other hand, since $A \subseteq \Gamma(A)$ and $A \subseteq B \subseteq \Gamma(B)$ it follows that $A \subseteq \Gamma(A) \cap \Gamma(B)$, so $\Delta^*(A) \leq \Delta(\Gamma(A) \cap \Gamma(B))$. Therefore $\Delta^*(A) = \Delta(\Gamma(A) \cap \Gamma(B))$, so it follows from the definition of $\Gamma(A)$ that $\Gamma(A) \subseteq \Gamma(A) \cap \Gamma(B) \subseteq \Gamma(B)$. \square

We now analyze the concentration of $\mathbf{X}_{\Delta, n}$, making liberal use of the fact that if $n^{O(1)}$ events occur with uniformly high probability then their conjunction also occurs w.h.p. by a union bound. For the rest of this subsection, “extensions” are with respect to an implicit $\mathbf{X} \equiv \mathbf{X}_{\Delta, n}$.

Lemma 6.10. *If $A \subseteq U$ and $\Gamma_U(A) = U$ (i.e. $\Delta(H) > \Delta(U)$ for all $H \in [A, U)$) then the number of U -extensions of any $\mathcal{A} \in \text{Sub}_n(A)$ is $\tilde{O}(1)$ w.h.p.*

(The above conditions are equivalent because, by the definition of $\Gamma(A)$, we have $\Gamma(A) = U$ if and only if U is the unique $H \in [A, U]$ that minimizes $\Delta(H)$.)

Proof. The result is trivial for $A = U$; assume it's true for all $B \in (A, U]$ and that $A \neq U$. (Since $\Delta(H) > \Delta(U)$ for all $H \in [A, U]$, for any $B \in (A, U]$ it is the case that $\Delta(H) > \Delta(U)$ for all $H \in [B, U]$.) Assume without loss of generality that $A = U[V(A)]$, since all U -extensions of \mathcal{A} are also U -extensions of \mathcal{A} 's unique possible $U[V(A)]$ -extension. Also condition on $\mathcal{A} \subseteq \mathbf{X}$, since otherwise \mathcal{A} trivially has zero U -extensions.

There are $n^{\alpha(U)-\alpha(A)}$ possible U -extensions of \mathcal{A} , so there are at most $n^{(\alpha(U)-\alpha(A)) \log n}$ sets of $\log n$ possible U -extensions of \mathcal{A} whose projections onto $\text{Sub}_n(U - A)$ are pairwise vertex-disjoint. (This is true even if we omit the condition about vertex-disjointness.) For each of these sets, all of its elements are subgraphs of \mathbf{X} with probability $n^{(-\beta(U)+\beta(A)) \log n}$, so this occurs for at least one such set with probability at most $n^{(\Delta(U)-\Delta(A)) \log n}$ (by a union bound). By assumption, $\Delta(U) - \Delta(A) < 0$, so w.h.p. any set of U -extensions of \mathcal{A} whose projections onto $\text{Sub}_n(U - A)$ are pairwise vertex-disjoint has $\tilde{O}(1)$ elements.

Let S be one such set, such that S is maximal. It follows that every U -extension of \mathcal{A} agrees with some element of S on some vertex in $V(U) - V(A)$. Therefore \mathcal{A} has at most $\sum_{\mathcal{U} \in S} \sum_{H \in (A, U]} E(\mathcal{U}, H)$ U -extensions, where $E(\mathcal{U}, H)$ is the number of U -extensions of \mathcal{A} that agree with \mathcal{U} on precisely H . By the inductive hypothesis, $E(\mathcal{U}, H)$ is $\tilde{O}(1)$ w.h.p. for all \mathcal{U} and H (independent of S), so \mathcal{A} has $\tilde{O}(1)$ U -extensions w.h.p. by a union bound. \square

Lemma 6.11. *If A is a U -base then any $\mathcal{A} \in \text{Sub}_n(A)$ has $\tilde{O}(n^{\Delta(U)-\Delta(A)})$ U -extensions w.h.p.*

Proof. Again, we assume that A is an induced subgraph of U and condition on $\mathcal{A} \subseteq \mathbf{X}$. Also assume without loss of generality that β is strictly positive on $E(U)$. The proof is by induction on $v(U) - v(A)$, for all $U \subseteq G$. The base case $A = U$ is trivial. Fix an arbitrary vertex $v \in V(U) - V(A)$. First we consider the case where $\Gamma(A \cup v) \neq U$. The number of U -extensions of \mathcal{A} equals the sum over all $\gamma \in \{\Gamma(A \cup v)\text{-extensions of } \mathcal{A}\}$ of the number of U -extensions of γ . Clearly A is a $\Gamma(A \cup v)$ -base, and Lemma 6.8 implies that $\Gamma(A \cup v)$ is a U -base. It follows from our assumptions that $v(A) < v(\Gamma(A \cup v)) < v(U)$, so we can apply the inductive hypothesis twice: w.h.p. \mathcal{A} has $\tilde{O}(n^{\Delta(\Gamma(A \cup v)) - \Delta(A)})$ extensions to $\Gamma(A \cup v)$, each of which has $\tilde{O}(n^{\Delta(U) - \Delta(\Gamma(A \cup v))})$ extensions to U , and the result follows.

Now assume that $\Gamma(A \cup v) = U$. Lemma 6.10 implies that $\mathcal{A} \cup v_i$ has $\tilde{O}(1)$ U -extensions w.h.p. for any i , so it suffices to show that w.h.p. there are $\tilde{O}(n^{\Delta(U) - \Delta(A)})$ values of i such that $\mathcal{A} \cup v_i$ extends to U . Let $\mathbf{W} = \mathbf{X}[u_i \mid u \in V(U) - v, i \in [n^{\alpha(u)}]]$, and if \mathcal{A} has $\tilde{O}(n^{\Delta(U-v) - \Delta(A)})$ extensions to $U - v$ when $\mathbf{W} = W$ then let W be “okay”. Since A is a $(U - v)$ -base, \mathbf{W} is okay w.h.p. by the inductive hypothesis. Let $\mathbf{Z}_i = \mathbb{1}\{\mathcal{A} \cup v_i \text{ extends to } U\}$, and let E be the event that $\sum_i \mathbf{Z}_i > \tilde{O}(n^{\Delta(U) - \Delta(A)})$. Then,

$$\begin{aligned} P(E) &= P(E \mid \mathbf{W} \text{ is okay})P(\mathbf{W} \text{ is okay}) + P(E \mid \mathbf{W} \text{ isn't okay})P(\mathbf{W} \text{ isn't okay}) \\ &\leq P(E \mid \mathbf{W} \text{ is okay}) + P(\mathbf{W} \text{ isn't okay}) \\ &\leq \max_{\text{okay } W} P(E \mid \mathbf{W} = W) + n^{-\omega(1)}, \end{aligned}$$

so it suffices to prove that $P(E \mid \mathbf{W} = W) \leq n^{-\omega(1)}$ for a worst-case W subject to W being okay.

The \mathbf{Z}_i are independent Bernoulli random variables (given W). By a union bound, $\mathbb{E}[\mathbf{Z}_i]$ is at most the number of $(U - v)$ -extensions of \mathcal{A} times the probability that the requisite edges between any one of them and v_i are in \mathbf{X} , i.e.

$$\mathbb{E}[\mathbf{Z}_i] \leq \tilde{O}\left(n^{\Delta(U-v) - \Delta(A)}\right) n^{\beta(U-v) - \beta(U)} = \tilde{O}\left(n^{\Delta(U) - \Delta(A) - \alpha(v)}\right).$$

Since A is a U -base, $\Delta(U) - \Delta(A) \geq 0$, so it follows from Lemma 6.4 that $\sum_{i=1}^{n^{\alpha(v)}} \mathbf{Z}_i$ is $\tilde{O}(n^{\Delta(U) - \Delta(A)})$ w.h.p. \square

Remark. It follows from Lemma C.1 that Lemma 6.11 is essentially tight.

Now we prove that $\mathbf{X}_{\Delta,n}$ is good w.h.p.:

Proof of Theorem 6.3. Let $A \subseteq U$, $\mathcal{A} \in \text{Sub}_n(A)$ and $v \in V(U) - V(A)$. By a union bound it suffices to prove that w.h.p. there are $\tilde{O}(n^{\Delta^*(A \cup v) - \Delta^*(A)})$ values of i such that $\mathcal{A} \cup v_i$ extends to U . The number of such i is at most the number of i such that $\mathcal{A} \cup v_i$ extends to $\Gamma(A \cup v)$, which is at most the number of $\Gamma(A \cup v)$ -extensions of \mathcal{A} . Since $\Gamma(A) \subseteq \Gamma(A \cup v)$ (Lemma 6.9), this equals the sum over all $\gamma \in \{\Gamma(A)\text{-extensions of } \mathcal{A}\}$ of the number \mathbf{E}_γ of $\Gamma(A \cup v)$ -extensions of γ .

It follows from Lemma 6.10 that \mathcal{A} has $\tilde{O}(1)$ extensions to $\Gamma(A)$ w.h.p. (To see this, note that if $A \subseteq H \subset \Gamma(A)$ then $\Delta(H) \geq \Delta^*(A) = \Delta(\Gamma(A))$ (Lemma 6.7), and if $\Delta(H) = \Delta^*(A)$ then it follows from the definition of $\Gamma(A)$ that $\Gamma(A) \subseteq H$, a contradiction.) Since $\Gamma(A)$ is a $\Gamma(A \cup v)$ -base (Lemma 6.8), it follows from Lemma 6.11 that any \mathbf{E}_γ is $\tilde{O}(n^{\Delta(\Gamma(A \cup v)) - \Delta(\Gamma(A))})$ w.h.p. ($= \tilde{O}(n^{\Delta^*(A \cup v) - \Delta^*(A)})$ by Lemma 6.7). \square

6.2 The Circuit

If D is a data structure then let $|D|$ denote the number of bits used to represent it according to whatever schema we describe. If A is a bit array and b is a bit then let $(A \vee b)_i = A_i \vee b$ and $(A \wedge b)_i = A_i \wedge b$ for all $i \in [|A|]$. When there is a null element we represent it by the all-zeros string.

We now prove that there exists a constant-depth circuit with $\tilde{O}(n^{\kappa \Delta(G)+3})$ wires that solves $G\text{-SUB}(\mathbf{X}_{\Delta,n})$ w.h.p.

Proof of Theorem 6.1. Since $\mathbf{X}_{\Delta,n}$ is good w.h.p. (Theorem 6.3) it suffices to prove the existence of a (small, constant-depth) circuit \mathbf{C} such that $P_{X \sim \mathbf{X}_{\Delta,n}}(\mathbf{C}(X) = G\text{-SUB}(X) \mid X \text{ is good}) = 1 - n^{-\omega(1)}$. (This is explained below.) By Yao's Principle [Yao77] it suffices to prove the existence of a (small, constant-depth) random circuit \mathbf{C} such that $P(\mathbf{C}(X) = G\text{-SUB}(X)) = 1 - n^{-\omega(1)}$ for any fixed good X . More precisely,

$$\begin{aligned} & \max_{\mathbf{C}: P(\mathbf{C}=\mathbf{C}) > 0} P_{X \sim \mathbf{X}_{\Delta,n}}(\mathbf{C}(X) = G\text{-SUB}(X)) \geq \\ & \max_{\mathbf{C}: P(\mathbf{C}=\mathbf{C}) > 0} P_{X \sim \mathbf{X}_{\Delta,n}}(\mathbf{C}(X) = G\text{-SUB}(X) \mid X \text{ is good}) P(\mathbf{X}_{\Delta,n} \text{ is good}) \geq \\ & P_{X \sim \mathbf{X}_{\Delta,n}}(\mathbf{C}(X) = G\text{-SUB}(X) \mid X \text{ is good}) \cdot \left(1 - n^{-\omega(1)}\right) \geq \\ & \left(1 - n^{-\omega(1)}\right) \min_{\text{good } X} P(\mathbf{C}(X) = G\text{-SUB}(X)). \end{aligned}$$

The following result is essentially implicit in [LRR17] (as is the argument above) and helps keep the random circuit small:

Lemma 6.12 (Random Hashing). *Let S be a set containing a null element, and assume all elements of S are represented using the same number of bits. Let $l = l(n) \leq n^{O(1)}$ and $m = m(n)$ be functions of n . Then there exists a random, constant-depth circuit $\mathbf{C} : S^l \rightarrow S^{\tilde{O}(m)}$ such that if A is an array of l values in S , of which all but at most m are null, then \mathbf{C} has at most $|A|n^{O(1)}$ gates and $|A|\tilde{O}(l/m)$ wires, and w.h.p. the multiset of non-null elements of $\mathbf{C}(A)$ is the same as that of A .*

We remark that Lemma 6.12 will only be called with $l \leq \tilde{O}(n)$.

Proof. The result is trivial if $l \leq m$ (simply return A) so assume otherwise. Let $\mathbf{h} : [l] \rightarrow [m]$ be a uniform random function. Let \mathbf{B} be an $m \times \tilde{O}(1)$ array of values in S , where $\mathbf{B}[p, q]$ is the q 'th non-null element of $\mathbf{A}^{(p)} := A[\mathbf{h}^{-1}(p)]$ if this set has at least q elements, and $\mathbf{B}[p, q]$ is null otherwise. Each of the at most m non-null elements of A is independently in $\mathbf{A}^{(p)}$ with probability $1/m$, so for any particular p , the sub-array $\mathbf{B}[p, :]$ is large enough to store the non-null elements of $\mathbf{A}^{(p)}$ w.h.p. (Lemma 6.4). It follows from a union bound that \mathbf{B} has the same non-null elements as A w.h.p. Also assume that $|\mathbf{h}^{-1}(p)|$ is $\tilde{O}(l/m)$ for all p ; this occurs w.h.p. by Lemma 6.4. Under these conditions it suffices to compute \mathbf{B} , and this can be done as follows.

For $x \in \{0, 1\}^N$ let $T_k^N(x) = \mathbb{1}\{x \text{ has at least } k \text{ ones}\}$. Then $\mathbf{J}^{(p)}[i] := \mathbb{1}\{\mathbf{A}^{(p)}[i] \neq \text{null}\}$ can be computed by applying a single OR gate to all elements of $\mathbf{A}^{(p)}[i]$, and

$$\mathbf{B}[p, q] = \bigvee_{i \in [|\mathbf{h}^{-1}(p)|]} \left(T_q^i \left(\mathbf{J}^{(p)}[1 : i] \right) \wedge \neg T_q^{i-1} \left(\mathbf{J}^{(p)}[1 : i-1] \right) \wedge \mathbf{A}^{(p)}[i] \right).$$

Fact 6.13 ([Hås+94, Theorem 6]). *If $k = \lfloor \log^\gamma N \rfloor$ for constant γ , then T_k^N can be computed for $m = \lfloor \gamma \rfloor + 1$ by monotone unbounded fan-in circuits of depth $m + 2$ with $2^{O(\log^{\gamma/m} N \log \log N)}$ gates, where $\gamma/m < 1$, and $O(N \log^{2\gamma+2} N)$ wires.*

Let $N = \tilde{O}(l/m) = n^{O(1)}$, and let γ be a constant such that the dimensions of \mathbf{B} are at most $m \times k$ where $k = \lfloor \log^\gamma N \rfloor$. Let \mathbf{T} be the $N^{O(1)}$ -size (hence $n^{O(1)}$ -size) circuit from Fact 6.13 that computes T_k^N . Observe that $T_q^i(x) = \mathbf{T}(x, y)$ where $y \in \{0, 1\}^{N-i}$ is an arbitrary fixed string with exactly $k - q$ ones that can be hard-coded in. Therefore $\mathbf{B}[p, q]$ can be computed by a constant-depth circuit of size $\sum_{i \in [|\mathbf{h}^{-1}(p)|]} (n^{O(1)} + |\mathbf{A}^{(p)}[i]|) \leq |\mathbf{A}^{(p)}| n^{O(1)}$. Summing over p and q , the total number of gates is $|A| n^{O(1)}$. To count wires instead of gates, replace $n^{O(1)}$ with $\tilde{O}(N) = \tilde{O}(l/m)$. \square

Given $H \subseteq G$ and an ordering $\pi = (\pi^1, \dots, \pi^{v(H)})$ of $V(H)$, we can represent $\text{Sub}_X(H)$ as a tree in the following way. Start with a rooted, depth- $v(H)$ tree (meaning the root has depth 0 and the leaves have depth $v(H)$) in which each interior node has n unordered children labeled $1, \dots, n$. Then take the induced subtree of this tree on the union of all root-to-leaf paths (root, l_1, \dots, l_d) such that⁴ $\pi_{l_1}^1, \dots, \pi_{l_v}^{v(H)}$ are the vertices of an H -colored subgraph of X .

With respect to an implicit H and π , let $\delta_i = \Delta_H^*(\pi^1 \cup \dots \cup \pi^i)$ for $0 \leq i \leq v(H)$, and let $\phi_i = \delta_{i+1} - \delta_i$ for $0 \leq i < v(H)$.

Lemma 6.14. $0 \leq \phi_i \leq 1$ for all i .

Proof. Clearly $\delta_i \leq \delta_{i+1}$. Let $A \subseteq H$ such that $\pi^1, \dots, \pi^i \in V(A)$ and $\Delta(A) = \delta_i$. Then $\delta_{i+1} \leq \Delta(A \cup \pi^{i+1}) \leq \Delta(A) + \alpha(\pi^{i+1}) \leq \delta_i + 1$. \square

Let $T = T(H, \pi)$ be a depth- $v(H)$ tree in which each node at depth $i < v(H)$ has $n^{\phi_i} \log^{c_i} n$ children, where c_i is a sufficiently large constant. Each non-root node N has a label $\mathcal{L}(N) \in \{\text{null}\} \cup [n]$, and the root is labeled “root”. It is required that if we ignore the null vertices of T , then T is isomorphic to the tree representation of $\text{Sub}_X(H)$ described above.

If the underlying tree structure of T (that is, everything except the labels) is implicit, then we can represent T by an array of values in $\{\text{null}\} \cup [n]$, indexed by the nodes of T . Each of these

⁴Recall that $(\pi^j)_{l_j}$ is a π^j -colored vertex in X .

values can be associated with a bit string in a natural way. We will consider circuits that compute T according to this representation.

Let S be an *immediate subtree* of T (resp. of a node N), denoted $S \in T$ (resp. $S \in N$), if S 's root is a child of T 's root (resp. of N). Any subtree is considered to have the same label as its root.

Lemma 6.15. $|T|$ is $\tilde{O}(n^{\Delta(H)})$.

Proof. $\delta_0 = \Delta(\emptyset) = 0$ and $\delta_{v(H)} = \Delta_H^*(V(H)) = \Delta(H)$. It takes $\tilde{O}(1)$ bits to store an element of $[n]^{V(H)}$, and each ϕ_i is nonnegative (Lemma 6.14), so

$$|T| = \tilde{O} \left(\prod_{i=0}^{v(H)-1} n^{\phi_i} \right) = \tilde{O} \left(n^{\sum_{i=0}^{v(H)-1} \phi_i} \right) = \tilde{O} \left(n^{\delta_{v(H)} - \delta_0} \right) = \tilde{O} \left(n^{\Delta(H)} \right). \quad \square$$

Lemma 6.16. For all $H \subseteq G$ there exists a random, constant-depth circuit with $\tilde{O}(n^{\Delta(H)+2})$ wires, independent of X , that computes $T(H, \pi')$ from $T(H, \pi)$ w.h.p.

Proof. Assume that π and π' differ only in positions d and $d+1$. (The general case can be reduced to at most $\binom{v(H)}{2}$ copies of this circuit in succession.) Define δ'_i and ϕ'_i analogously to δ_i and ϕ_i , but with respect to π' rather than π . Clearly $\delta_i = \delta'_i$ for $i \neq d$, so $\phi_i = \phi'_i$ for $i \notin \{d-1, d\}$.

For each depth- $(d-1)$ node N of $T(H, \pi)$, in parallel, do the following. For $\sigma \in N, j \in [n]$ let

$$\tau'_{\sigma j} = \bigvee_{\tau \in \sigma} \left((\mathcal{L}(\tau) = j) \wedge \tau^{(\mathcal{L}(\sigma))} \right),$$

where $\tau^{(\mathcal{L}(\sigma))}$ is formed from τ by replacing its (root's) label with $\mathcal{L}(\sigma)$. Let σ'_j be the tree whose immediate subtrees are $\tau'_{\sigma j}$ for $\sigma \in N$, and whose label is $(\bigvee_{\sigma \in N} \bigvee_{\tau \in \sigma} (\mathcal{L}(\tau) = j)) \wedge \bar{j}$ where \bar{j} is the bit-string representation of j . Hash the number of immediate subtrees of σ'_j down to $\tilde{O}(n^{\phi'_d})$ for each j in parallel, and hash the number of σ'_j down to $\tilde{O}(n^{\phi'_d-1})$. (The hashing uses Lemma 6.12 and succeeds w.h.p. because X is good; also note that $\phi_d + \phi_{d-1} = \delta_{d+1} - \delta_{d-1} = \phi'_d + \phi'_{d-1}$.) Finally, the new children of N are the remaining σ'_j .

Computing $\tau'_{\sigma j}$ takes $\tilde{O}(\sum_{\tau \in \sigma} |\tau|) = \tilde{O}(|\sigma|)$ wires, so computing σ'_j takes $\tilde{O}(\sum_{\sigma \in N} |\sigma|) = \tilde{O}(|N|)$ wires, and doing this for all N and j takes $\tilde{O}(n|T|) = \tilde{O}(n^{\Delta(H)+1})$ wires (Lemma 6.15). The hashing increases the number of wires by a factor of $\tilde{O}(n)$. \square

For $uv \in E(G)$ we can construct $T(uv)$ as follows. Suppose we're given the adjacency matrix $A \in \{0, 1\}^{n^{\alpha(u)} \times n^{\alpha(v)}}$ such that $A_{ij} = \mathbb{1}\{u_i v_j \in E(X)\}$. Let $\tau'_{ij} = A_{ij} \wedge \bar{i}$. Let σ'_j be the tree with children τ'_{ij} for $i \in [n^{\alpha(u)}]$, and label $(\bigvee_i A_{ij}) \wedge \bar{j}$. This setup is equivalent to the situation immediately before the hashing in the proof of Lemma 6.16, and the rest of the construction is the same. This takes $\tilde{O}(n^3)$ wires, including the hashing.

Lemma 6.17. For all $H, H' \subseteq G$ there exists a random, constant-depth circuit, independent of X , with $\tilde{O}(n^{\max(\Delta(H), \Delta(H'))+2})$ wires, that computes $T(H \cup H', \hat{\pi})$ from $T(H, \pi)$ and $T(H', \pi')$ w.h.p. for some $\hat{\pi}$.

Proof. Let $T = T(H, \pi)$ and $T' = T(H', \pi')$. By Lemma 6.16 we can assume without loss of generality that $\{\pi^1, \dots, \pi^{v(H \cap H')}\} = \{\pi'^1, \dots, \pi'^{v(H \cap H')}\} = V(H \cap H') = V(H) \cap V(H')$, and that

$\pi^k = \pi'^k = \hat{\pi}^k$ for $k \in [v(H \cap H')]$. Define ϕ' and $\hat{\phi}$ with respect to (H', π') and $(H \cup H', \hat{\pi})$ respectively.

Let $\psi_i = \min(\phi_i, \phi'_i)$. For $0 \leq d \leq v(H \cap H')$ let S_d be a depth- d tree in which each node at depth $i < d$ (including $i = 0$) has $\tilde{O}(n^{\psi_i})$ children. Again, each non-root node N of S_d has a label $\mathcal{L}(N) \in \{\text{null}\} \cup [n]$, and the root is labeled “root”. It is required that if we ignore null vertices, then S_d is isomorphic to the intersection of the depth- d truncations of T and T' . Furthermore, each leaf ℓ of S_d is associated with the pair (σ, σ') of subtrees of T and T' such that the root(S)-to- ℓ path in S , the root(T)-to-root(σ) path in T , and the root(T')-to-root(σ') path in T' are all the same sequence of labels.

The tree S_0 is the single node (T, T') , and we can compute S_{d+1} from S_d by doing the following for each leaf (σ, σ') of S_d in parallel. Assume without loss of generality that $\psi_d = \phi_d$. (If $\psi_d = \phi'_d$, reverse the roles of σ and σ' in the following construction.) For $\tau \in \sigma$ let ρ_τ be the immediate subtree of σ' with the same label as τ (if this exists), i.e. $\rho_\tau = \bigvee_{\tau' \in \sigma'} ((\mathcal{L}(\tau) = \mathcal{L}(\tau')) \wedge \tau')$. Replace (σ, σ') with a new node with children $(\rho_\tau \neq \text{null}) \wedge (\tau, \rho_\tau)$ for all $\tau \in \sigma$. Assign the node replacing (σ, σ') the same label as (σ, σ') , and assign (τ, ρ_τ) the same label as τ and ρ_τ .

Computing ρ_τ takes $\tilde{O}(\sum_{\tau' \in \sigma'} |\tau'|) = \tilde{O}(|\sigma'|)$ wires, and there are at most n values of τ (Lemma 6.14), so computing ρ takes $\tilde{O}(n|\sigma'|)$ wires. Given ρ , computing the leaves of the replacement for (σ, σ') takes $O(\sum_{\tau \in \sigma} (|\tau| + |\rho_\tau|)) = O(|\sigma| + |\sigma'|)$ wires. Since the roles of σ and σ' might be reversed above, all of this takes at most $\tilde{O}(n|\sigma| + n|\sigma'|)$ wires. Since S_d has $\tilde{O}(n^{\sum_{i < d} \psi_i})$ leaves, the number of wires is at most

$$\tilde{O}\left(n^{\sum_{i < d} \psi_i} (n|\sigma| + n|\sigma'|)\right) \leq \tilde{O}\left(n^{1 + \sum_{i < d} \phi_i} |\sigma| + n^{1 + \sum_{i < d} \phi'_i} |\sigma'|\right) = \tilde{O}(n|T| + n|T'|).$$

Let $S = S_{v(H \cap H')}$. For d from $v(H \cap H') - 1$ down to 0, for each depth- d node N in S , hash (Lemma 6.12) the number of immediate subtrees of N down from $\tilde{O}(n^{\psi_d})$ to $\tilde{O}(n^{\hat{\phi}_d})$, and if all of N 's children are null and $d > 0$ then set N to null. (Note that $\hat{\phi}_d \leq \psi_d$ by Lemma C.3.) This takes $\tilde{O}(|S|n) \leq \tilde{O}((|T| + |T'|)n) = \tilde{O}(n^{\max(\Delta(H), \Delta(H'))+1})$ wires (Lemma 6.15). By induction on d , a node retains its label if and only if it should retain its label in $T(H \cup H', \hat{\pi})$, so the hashing succeeds w.h.p. because X is good.

Finally, for each leaf (τ, τ') of S , append a copy of τ' to each leaf of τ , and put this in place of (τ, τ') in S . This operation is purely semantic and requires no wires. The resulting tree does in fact have the proper dimensions to be $T(H \cup H', (\pi^1, \dots, \pi^{v(H)}, \pi^{v(H \cap H')+1}, \dots, \pi^{v(H')}))$ by Lemma C.4, but without this knowledge we could instead use hashing on τ and τ' as above, without knowing whether or not it succeeds vacuously. \square

For each successive H in an optimal union sequence, compute $T(H)$ as described above, and then apply a single OR gate to all leaves of $T(G)$. \square

Acknowledgments

Thanks to Benjamin Rossman for introducing me to this topic, and for having many helpful discussions about the research and about drafts of this paper. Thanks to Henry Yuen and the anonymous reviewers for their feedback as well.

A Equivalence of Threshold Weightings and Markov Chains

Theorem A.1. *For any threshold weighting $(\alpha, \beta) \in \theta(G)$ there exists a function $M : V(G) \times V(G) \rightarrow \mathbb{R}_{\geq 0}$ such that*

1. $M(u, u) = 0$ for all u ,
2. $M(u, v) + M(v, u) = \beta(uv)$ for all $u \neq v$, and
3. $\sum_{v \in V(G)} M(u, v) = \alpha(u)$ for all u .

Proof. Let $\Delta = (\alpha, \beta)$. The proof is by induction on $v(G)$. If G is a single vertex u then $\theta(G)$ consists only of $\alpha = 0$, so setting $M(u, u) = 0$ satisfies the requirements. Now assume $v(G) > 1$. For $A, B \subseteq G$ let $M(A, B) = \sum_{u \in V(A), v \in V(B)} M(u, v)$ (once $M(u, v)$ is specified). Assume without loss of generality that G is a clique, since we can assign $\beta = 0$ on nonexistent edges.

Let $H = \operatorname{argmin}_{F \subseteq G, 0 < v(F) < v(G)} \Delta(F)$, where ties are broken arbitrarily subject to H being an induced subgraph of G . Since $\Delta(G) = 0$,

$$\beta(H, G - H) = \Delta(G) + \beta(H, G - H) = \Delta(H) + \Delta(G - H) \geq \Delta(H),$$

so for $u \in V(H), v \in V(G - H)$ we can define $M(u, v) \in [0, \beta(uv)]$ such that $M(H, G - H) = \Delta(H)$. For $u \in V(H)$ let $\alpha_H(u) = \alpha(u) - M(u, G - H)$, and let Δ_H be the restriction of $\alpha_H - \beta$ to subgraphs of H . For any $\emptyset \subset F \subseteq H$,

$$\Delta_H(F) = \Delta(F) - M(F, G - H) \geq \Delta(F) - M(H, G - H) \geq \Delta(H) - M(H, G - H) = 0,$$

with equality if $F = H$. Therefore Δ_H is a threshold weighting on H . Recursively define a restriction of M to $V(H) \times V(H)$ such that this restriction is a Markov Chain on H that is equivalent to Δ_H .

For $u \in V(G - H), v \in V(H)$ let $M(u, v) = \beta(uv) - M(v, u)$. For $u \in V(G - H)$ let $\alpha_{G-H}(u) = \alpha(u) - M(u, H)$, and let Δ_{G-H} be the restriction of $\alpha_{G-H} - \beta$ to subgraphs of $G - H$. Then,

$$\begin{aligned} \Delta_{G-H}(G - H) &= \Delta(G - H) - M(G - H, H) = \Delta(G - H) - \beta(G - H, H) + M(H, G - H) \\ &= \Delta(G - H) - \beta(G - H, H) + \Delta(H) = \Delta(G) = 0. \end{aligned}$$

For any $\emptyset \subset F \subset G - H$, if $v(F) < v(G - H)$ then

$$\Delta_{G-H}(F) = \Delta(F) - M(F, H) \geq \Delta(F) - \beta(F, H) \geq \Delta(G[V(H), V(F)]) - \Delta(H) \geq 0,$$

and if $v(F) = v(G - H)$ then $\Delta_{G-H}(F) \geq \Delta_{G-H}(G - H) = 0$. Therefore Δ_{G-H} is a threshold weighting on $G - H$. Recursively define a restriction of M to $V(G - H) \times V(G - H)$ such that this restriction is a Markov Chain on $G - H$ that is equivalent to Δ_{G-H} .

We now verify that $M(u, G) = \alpha(u)$ for all u ; the other requirements follow easily by induction. If $u \in V(H)$ then $M(u, H) = \alpha_H(u)$ by induction, and $M(u, G - H) = \alpha(u) - \alpha_H(u)$ by the definition of α_H . Similarly, if $u \in V(G - H)$ then $M(u, G - H) = \alpha_{G-H}(u)$ and $M(u, H) = \alpha(u) - \alpha_{G-H}(u)$. Therefore $M(u, G) = M(u, H) + M(u, G - H) = \alpha(u)$ for all u . \square

B Proof that $\kappa(K_q^d)$ is $O(q^d/d)$ for all q

The proof below is self-contained; however in places with clear analogues in Section 5.2 we will give less detailed explanations of the intermediate steps and intuition.

Proof. Fix q and d . Let a *query tree* be a binary tree in which each node is labeled with some $U_1 \times \cdots \times U_d$ where each $U_i \subseteq [q]$. The root is labeled with $[q]^d$, each leaf is labeled with a singleton set, and for any interior node N labeled with $U_1 \times \cdots \times U_d$ there exist $i \in [d]$ and $k \in U_i$ such that the left and right children of N are labeled with $U_1 \times \cdots \times U_{i-1} \times (U_i - k) \times U_{i+1} \times \cdots \times U_d$ and $U_1 \times \cdots \times U_{i-1} \times \{k\} \times U_{i+1} \times \cdots \times U_d$ respectively. (In the latter case, U_i necessarily has at least two elements.)

With respect to an implicit query tree T , let $\ell_0, \dots, \ell_{q^d-1}$ be the leaves in increasing order from left to right, and for $0 \leq a \leq q^d$ let $G(a) = K_q^d[\ell_0, \dots, \ell_{a-1}]$. Let $\mu_T = \max_a \Delta_o(G(a))$ and let μ be the maximum of μ_T over all query trees T . For a threshold weighting $\Delta \in \theta(K_q^d)$ and $H \subseteq K_q^d$ let $\kappa_\Delta(H) = \min_{S \in \text{Seq}(H)} \max_{F \in S} \Delta(F)$, and for a vertex v let $\kappa_\Delta(v) = 0$. By Theorem 3.8(i) it suffices to prove that $\kappa_\Delta(K_q^d)$ is $O(q^d/d)$ for all threshold weightings $\Delta = (1, \beta)$.

Lemma B.1. *Fix a query tree T . Let $0 \leq a < b \leq q^d$ such that $\ell_a, \dots, \ell_{b-1}$ are exactly the leaves descended from some node of T . Let $\Delta = (1, \beta) \in \theta(K_q^d)$ such that $\beta(G(a)) \geq \beta_o(G(a))$ and $\beta(G(b)) \geq \beta_o(G(b))$, and $\kappa_\Delta(G(a)) \leq 2\mu$. Then $\kappa_\Delta(G(b)) \leq 2\mu$.*

Proof. Let N be the node of T such that the leaves descended from N are exactly $\ell_a, \dots, \ell_{b-1}$. Let $U_1 \times \cdots \times U_d$ be the label of N . Let $B = G(b) - G(a)$; it follows that $B = K_q^d[U_1 \times \cdots \times U_d] = K_q^d[\ell_a, \dots, \ell_{b-1}]$, where ℓ_i is identified with the unique element in its label.

The proof is by induction on $\sum_i (|U_i| - 1)$, for all query trees T . (It follows from the definitions that $|U_i| \geq 1$, with equality for all i if and only if N is a leaf.) In the inductive step we handle separately the cases where $\beta(B) \geq \beta_o(B)$ and $\beta(B) < \beta_o(B)$. The base case is a special case of the former because if B is a single vertex then $\beta(B)$ and $\beta_o(B)$ are both zero.

Case 1: $\beta(B) \geq \beta_o(B)$. If B is a single vertex then $\kappa_\Delta(B) = 0 \leq 2\mu$; we now obtain the same result in the case where B is not a single vertex. Let $\mathcal{I} = \{i \in [d] \mid |U_i| \geq 2\}$, and note that \mathcal{I} is nonempty. For $i \in \mathcal{I}$ and $k \in U_i$ let $B(i, k) = B[v \in V(B) \mid v_i \neq k]$. Choose a pair (\mathbf{i}, \mathbf{k}) uniformly at random out of all pairs (i, k) such that $i \in \mathcal{I}$ and $k \in U_i$. Each edge in B is also in $B(\mathbf{i}, \mathbf{k})$ with the same probability $p = 1 - (|\mathcal{I}| + 1) / \sum_{i \in \mathcal{I}} |U_i|$ (since adjacent vertices differ in a unique coordinate), so by linearity of expectation,

$$\mathbb{E}[\beta(B(\mathbf{i}, \mathbf{k}))] = p\beta(B) \geq p\beta_o(B) = \mathbb{E}[\beta_o(B(\mathbf{i}, \mathbf{k}))].$$

Therefore $\beta(B(i, k)) \geq \beta_o(B(i, k))$ for some fixed i and k .

Now our claim that $\kappa_\Delta(B) \leq 2\mu$ follows from two applications of the inductive hypothesis. Let T' be any query tree in which the sequence of labels along the path from the root to the leftmost leaf includes $U_1 \times \cdots \times U_d$ followed by $U_1 \times \cdots \times (U_i - k) \times \cdots \times U_d$. With respect to T' , an application of the inductive hypothesis with $a' = 0$ and $b' = (1 - 1/|U_i|) \prod_j |U_j|$ reveals that $\kappa_\Delta(B(i, k)) \leq 2\mu$, and then an application of the inductive hypothesis with $a'' = (1 - 1/|U_i|) \prod_j |U_j|$ and $b'' = \prod_j |U_j|$ ($= b - a$) reveals that $\kappa_\Delta(B) \leq 2\mu$.

The rest of the proof is essentially identical to the case $q = 2$. Let S be an optimal (with respect to Δ) union sequence for $G(a)$, followed by an optimal union sequence for B , followed by

$G(a) \cup B, G(a) \cup B \cup e_1, \dots, G(a) \cup B \cup \{e_j\}$, where the $\{e_j\}$ are the edges between $G(a)$ and B in K_q^d . (If B is a single vertex then omit $G(a) \cup B$ from the sequence.) Then,

$$\max_{H \in \mathcal{S}} \Delta(H) \leq \max(\kappa_\Delta(G(a)), \kappa_\Delta(B), \Delta(G(a)) + \Delta(B)).$$

We proceed to bound each of these three terms by 2μ , completing the proof. We have assumed that $\kappa_\Delta(G(a)) \leq 2\mu$, and proved that $\kappa_\Delta(B) \leq 2\mu$. We have also assumed that $\beta(G(a)) \geq \beta_o(G(a))$, and since Δ and Δ_o both evaluate to 1 on all vertices, it follows that $\Delta(G(a)) \leq \Delta_o(G(a)) \leq \mu$ (with the last step following from the definition of μ). Similarly, $\Delta(B) \leq \Delta_o(B) \leq \mu$, and it follows that $\Delta(G(a)) + \Delta(B) \leq 2\mu$.

Case 2: $\beta(B) < \beta_o(B)$. For $i \in \mathcal{I}$ and $k \in U_i$ let $H(i, k) = K_q^d[\ell_0, \dots, \ell_{a-1}, V(B(i, k))]$ (where \mathcal{I} and $B(i, k)$ are defined as above). Choose a pair (\mathbf{i}, \mathbf{k}) uniformly at random out of all pairs (i, k) such that $i \in \mathcal{I}$ and $k \in U_i$. Then there exist $p_0 > p_1 > p_2 \geq 0$ (specifically, $p_0 = 1$, $p_1 = 1 - |\mathcal{I}| / \sum_{i \in \mathcal{I}} |U_i|$, and $p_2 = 1 - (|\mathcal{I}| + 1) / \sum_{i \in \mathcal{I}} |U_i|$) such that

$$\begin{aligned} \mathbb{E}[\beta(H(\mathbf{i}, \mathbf{k}))] &= p_0\beta(G(a)) + p_1\beta(G(a), B) + p_2\beta(B) \\ &= (p_0 - p_1)\beta(G(a)) + p_1\beta(G(b)) + (p_2 - p_1)\beta(B) \\ &> (p_0 - p_1)\beta_o(G(a)) + p_1\beta_o(G(b)) + (p_2 - p_1)\beta_o(B) \\ &= \mathbb{E}[\beta_o(H(\mathbf{i}, \mathbf{k}))]. \end{aligned}$$

Therefore $\beta(H(i, k)) > \beta_o(H(i, k))$ for some fixed i and k .

Preparing to apply the inductive hypothesis, let T'' be any query tree structured and labeled exactly like T on all ancestors of ℓ_j for all $j < a$, and on all ancestors of N , but now the left child of N is labeled with $U_1 \times \dots \times (U_i - k) \times \dots \times U_d$. With respect to T'' , an application of the inductive hypothesis with $a' = a$ and $b' = a + (1 - 1/|U_i|) \prod_j |U_j|$ reveals that $\kappa_\Delta(G(a + (1 - 1/|U_i|) \prod_j |U_j|)) \leq 2\mu$, and a second application of the inductive hypothesis with $a'' = a + (1 - 1/|U_i|) \prod_j |U_j|$ and $b'' = b (= a + \prod_j |U_j|)$ reveals that $\kappa_\Delta(G(b)) \leq 2\mu$. \square

Lemma B.2. μ is $O(q^d/d)$.

Proof. We use a cruder bound here than in the case $q = 2$. Let T be an arbitrary query tree and $a \in [q^d]$. Let N_0 be the nearest common ancestor of ℓ_0 and ℓ_{a-1} . If N_0 is a leaf then clearly $\Delta_o(G(a))$ is $O(q^d/d)$, so assume otherwise. Let N_L and N_R be the left and right children of N_0 , and note that ℓ_{a-1} is a descendant of N_R . By Eq. (1), since K_q^d is $(q-1)d$ -regular it suffices to prove that $e(G(a), K_q^d - G(a))$ is $O(q^{d+1})$. Suppose N_0 is labeled with $U_1 \times \dots \times U_d$ and N_L is labeled with $U_1 \times \dots \times (U_i - k) \times \dots \times U_d$. Since each vertex in $G(a)$ is a descendant of N_0 , all edges between $G(a)$ and $K_q^d - G(a)$ are in one of the following classes:

1. *Edges (in K_q^d) between a leaf descended from N_0 and a leaf not descended from N_0 .* Each leaf descended from N_0 is adjacent to $\sum_{j=1}^d (q - |U_j|)$ leaves not descended from N_0 , so this amounts to $(dq - \sum_j |U_j|) \prod_j |U_j|$ edges in total. By the AM-GM inequality, this is at most

$$\left(\frac{(dq - \sum_j |U_j|) + \sum_j |U_j|}{d+1} \right)^{d+1} = \left(\frac{dq}{d+1} \right)^{d+1} = q^{d+1} \left(1 - \frac{1}{d+1} \right)^{d+1} < q^{d+1}/e.$$

2. *Edges (in K_q^d) between a leaf descended from N_L and a leaf descended from N_R .* Each leaf descended from N_L is adjacent to one leaf descended from N_R , so this amounts to at most $\prod_j |U_j| \leq q^d$ edges in total.

3. Edges (in K_q^d) between a leaf descended from N_R that's in $G(a)$, and a leaf descended from N_R that's in $K_q^d - G(a)$. This is at most what the value of μ would be if d were $d - 1$ instead. (Eliminate coordinate i , and replace U_j with $[q]$ for all $j \neq i$.)

The total number of edges in all classes is therefore $O(q^{d+1} + q^d + \dots) = O(q^{d+1})$. \square

Finally, it follows from Lemmas B.1 and B.2 that $\kappa(K_q^d) \leq 2\mu \leq O(q^d/d)$. \square

Remark. The above argument holds even if we relax the definition of threshold weightings to allow Δ to take on negative values (where all definitions in terms of threshold weightings are with respect to this revised definition).

C Properties of Threshold Weightings and Threshold Random Graphs

Lemma C.1. *If $A \subseteq U \subseteq G$ are graphs, $\Delta \in \theta(G)$, and $\Delta(A) < \Delta(H)$ for all $H \in (A, U]$, then conditional on $\mathcal{A} \in \text{Sub}_{\mathbf{X}_{\Delta, n}}(A)$, there are at least $n^{\Delta(U) - \Delta(A)}(1 - o(1))$ U -extensions of \mathcal{A} a.a.s.*

Li et al. [LRR17] stated without proof that a similar result can be obtained using Janson's Inequality [Jan90]:

Fact C.2 (Janson's Inequality). *Let $\mathbf{B}_1, \dots, \mathbf{B}_\ell$ be independent Bernoulli random variables, let $W_1, \dots, W_k \subseteq [\ell]$, and for $i \in [k]$ let $\mathbf{I}_i = \prod_{j \in W_i} \mathbf{B}_j$. Also for $i, j \in [k], i \neq j$ let $i \sim j$ if $W_i \cap W_j \neq \emptyset$. Let $\mathbf{S} = \sum_i \mathbf{I}_i$ and $\mu = \mathbb{E}[\mathbf{S}]$. Then for any $0 \leq \epsilon \leq 1$,*

$$P(\mathbf{S} \leq (1 - \epsilon)\mu) \leq \exp\left(-\frac{\epsilon^2}{2} \cdot \frac{\mu^2}{\mu + \sum_{i \sim j} \mathbb{E}[\mathbf{I}_i \mathbf{I}_j]}\right).$$

Proof of Lemma C.1. Let $\mathcal{U}_1, \dots, \mathcal{U}_k$ be the possible U -extensions of \mathcal{A} , and let $\mathbf{I}_i = \mathbb{1}\{\mathcal{U}_i \subseteq \mathbf{X}\}$. Define μ as in Fact C.2; clearly $\mu = n^{\Delta(U) - \Delta(A)}$, by reasoning similar to the proof of Lemma 3.5. If $i \sim j$ then the projection of $\mathcal{U}_i \cap \mathcal{U}_j$ onto U must be some $H \in (A, U)$ that includes at least one edge. Summing over such H ,

$$\sum_{i \sim j} \mathbb{E}[\mathbf{I}_i \mathbf{I}_j] \leq \sum_H \mu n^{\Delta(U) - \Delta(H)} = \mu^2 \sum_H n^{\Delta(A) - \Delta(H)} = o(\mu^2).$$

Since μ is also $o(\mu^2)$, it follows that $\mu^2 / \left(\mu + \sum_{i \sim j} \mathbb{E}[\mathbf{I}_i \mathbf{I}_j]\right) \geq \mu^2 / o(\mu^2) = \omega(1)$, and the result follows from Fact C.2. \square

Lemma C.3. *For all $A \subseteq B \subseteq F \subseteq H \subseteq G$ and $\Delta \in \theta(G)$,*

$$\Delta_H^*(B) - \Delta_H^*(A) \leq \Delta_F^*(B) - \Delta_F^*(A).$$

Proof. Since $B \subseteq \Gamma_F(B) \subseteq \Gamma_H(A) \cup \Gamma_F(B)$ it follows that $\Delta_H^*(B) \leq \Delta(\Gamma_H(A) \cup \Gamma_F(B))$, and since $A \subseteq \Gamma_H(A)$ and $A \subseteq B \subseteq \Gamma_F(B)$ it follows that $\Delta_F^*(A) \leq \Delta(\Gamma_H(A) \cap \Gamma_F(B))$. So by Lemmas 6.5 and 6.7,

$$\begin{aligned} \Delta_H^*(B) + \Delta_F^*(A) &\leq \Delta(\Gamma_H(A) \cup \Gamma_F(B)) + \Delta(\Gamma_H(A) \cap \Gamma_F(B)) \\ &= \Delta(\Gamma_H(A)) + \Delta(\Gamma_F(B)) \\ &= \Delta_H^*(A) + \Delta_F^*(B). \end{aligned} \quad \square$$

Lemma C.4. *Let $\Delta \in \theta(G)$ and assume $L \cap R \subseteq A \subseteq B \subseteq L \subseteq G$ and $L \cap R \subseteq C \subseteq R \subseteq G$. Then, $\Delta_{L \cup R}^*(B \cup C) - \Delta_{L \cup R}^*(A \cup C) = \Delta_L^*(B) - \Delta_L^*(A)$.*

Proof. For any $F \in [A, L]$ and $H \in [C, R]$,

$$\begin{aligned} F \cap H &\subseteq L \cap R && (F \subseteq L \text{ and } H \subseteq R) \\ &\subseteq A \cap C && (\text{by assumption}) \\ &\subseteq F \cap H, && (A \subseteq F \text{ and } C \subseteq H), \end{aligned}$$

so by Lemma 6.5,

$$\begin{aligned} \Delta_{L \cup R}^*(A \cup C) &= \min_{\substack{A \subseteq F \subseteq L \\ C \subseteq H \subseteq R}} \Delta(F \cup H) \\ &= \min_{\substack{A \subseteq F \subseteq L \\ C \subseteq H \subseteq R}} (\Delta(F) + \Delta(H) - \Delta(F \cap H)) \\ &= \min_{A \subseteq F \subseteq L} \Delta(F) + \min_{C \subseteq H \subseteq R} \Delta(H) - \Delta(L \cap R) \\ &= \Delta_L^*(A) + \Delta_R^*(C) - \Delta(L \cap R). \end{aligned}$$

The same reasoning applies with B in place of A , so

$$\Delta_{L \cup R}^*(A \cup C) - \Delta_L^*(A) = \Delta_R^*(C) - \Delta(L \cap R) = \Delta_{L \cup R}^*(B \cup C) - \Delta_L^*(B). \quad \square$$

References

- [AM11] Noga Alon and Dániel Marx. “Sparse balanced partitions and the complexity of subgraph problems”. In: *SIAM J. Discrete Math.* 25.2 (2011), pp. 631–644. DOI: [10.1137/100812653](https://doi.org/10.1137/100812653) (pp. 3, 7, 8, 12).
- [AM85] Noga Alon and Vitali D. Milman. “ λ_1 , isoperimetric inequalities for graphs, and superconcentrators”. In: *J. Combin. Theory Ser. B* 38.1 (1985), pp. 73–88. DOI: [10.1016/0095-8956\(85\)90092-9](https://doi.org/10.1016/0095-8956(85)90092-9) (p. 4).
- [Ama10] Kazuyuki Amano. “ k -subgraph isomorphism on AC^0 circuits”. In: *Comput. Complexity* 19.2 (2010), pp. 183–210. DOI: [10.1007/s00037-010-0288-y](https://doi.org/10.1007/s00037-010-0288-y) (pp. 3, 13).
- [AYZ95] Noga Alon, Raphael Yuster, and Uri Zwick. “Color-coding”. In: *J. ACM* 42.4 (1995), pp. 844–856. DOI: [10.1145/210332.210337](https://doi.org/10.1145/210332.210337) (p. 1).
- [BK08] Hans L. Bodlaender and Arie M. C. A. Koster. “Combinatorial optimization on graphs of bounded treewidth”. In: *Comput. J.* 51.3 (2008), pp. 255–269. DOI: [10.1093/comjnl/bxm037](https://doi.org/10.1093/comjnl/bxm037) (p. 4).
- [Bod98] Hans L. Bodlaender. “A partial k -arboretum of graphs with bounded treewidth”. In: *Theoret. Comput. Sci.* 209.1-2 (1998), pp. 1–45. DOI: [10.1016/S0304-3975\(97\)00228-4](https://doi.org/10.1016/S0304-3975(97)00228-4) (p. 4).
- [CK06] L. Sunil Chandran and Telikepalli Kavitha. “The treewidth and pathwidth of hypercubes”. In: *Discrete Math.* 306.3 (2006), pp. 359–365. DOI: [10.1016/j.disc.2005.12.011](https://doi.org/10.1016/j.disc.2005.12.011) (pp. 2, 3, 9, 13).

- [EG04] Friedrich Eisenbrand and Fabrizio Grandoni. “On the complexity of fixed parameter clique and dominating set”. In: *Theoret. Comput. Sci.* 326.1-3 (2004), pp. 57–67. DOI: [10.1016/j.tcs.2004.05.009](https://doi.org/10.1016/j.tcs.2004.05.009) (p. 1).
- [Fil15] Yuval Filmus. “Harper’s isoperimetric inequality”. 2015. URL: <http://www.cs.toronto.edu/~yuvalf/Harper.pdf> (p. 12).
- [Har04] L. H. Harper. *Global methods for combinatorial isoperimetric problems*. Vol. 90. Cambridge Studies in Advanced Mathematics. Cambridge University Press, Cambridge, 2004. Chap. 1.2.3 and 3.3.5. DOI: [10.1017/CB09780511616679](https://doi.org/10.1017/CB09780511616679) (p. 12).
- [Har99] L. H. Harper. “On an isoperimetric problem for Hamming graphs”. In: *Discrete Appl. Math.* 95.1-3 (1999), pp. 285–309. DOI: [10.1016/S0166-218X\(99\)00082-7](https://doi.org/10.1016/S0166-218X(99)00082-7) (p. 13).
- [Hås+94] Johan Håstad, Ingo Wegener, Norbert Wurm, and Sang-Zin Yi. “Optimal depth, very small size circuits for symmetric functions in AC^0 ”. In: *Inform. and Comput.* 108.2 (1994), pp. 200–211. DOI: [10.1006/inco.1994.1008](https://doi.org/10.1006/inco.1994.1008) (p. 18).
- [Hås86] Johan Håstad. “Almost optimal lower bounds for small depth circuits”. In: *Proc. 18th Ann. ACM Symp. on Theory of Computing (STOC)*. 1986, pp. 6–20. DOI: [10.1145/12130.12132](https://doi.org/10.1145/12130.12132) (pp. 2, 14).
- [HLW06] Shlomo Hoory, Nathan Linial, and Avi Wigderson. “Expander graphs and their applications”. In: *Bull. Amer. Math. Soc. (N.S.)* 43.4 (2006), pp. 439–561. DOI: [10.1090/S0273-0979-06-01126-8](https://doi.org/10.1090/S0273-0979-06-01126-8) (pp. 4, 13).
- [HW17] Daniel J. Harvey and David R. Wood. “Parameters tied to treewidth”. In: *J. Graph Theory* 84.4 (2017), pp. 364–385. DOI: [10.1002/jgt.22030](https://doi.org/10.1002/jgt.22030) (p. 4).
- [IPZ01] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. “Which problems have strongly exponential complexity?” In: *J. Comput. System Sci.* 63.4 (2001), pp. 512–530. DOI: [10.1006/jcss.2001.1774](https://doi.org/10.1006/jcss.2001.1774) (p. 2).
- [Jan90] Svante Janson. “Poisson approximation for large deviations”. In: *Random Structures Algorithms* 1.2 (1990), pp. 221–229. DOI: [10.1002/rsa.3240010209](https://doi.org/10.1002/rsa.3240010209) (p. 24).
- [LRR17] Yuan Li, Alexander Razborov, and Benjamin Rossman. “On the AC^0 complexity of subgraph isomorphism”. In: *SIAM J. Comput.* 46.3 (2017), pp. 936–971. DOI: [10.1137/14099721X](https://doi.org/10.1137/14099721X) (pp. 2–4, 6–9, 13, 14, 17, 24).
- [Mar10] Dániel Marx. “Can you beat treewidth?” In: *Theory Comput.* 6.1 (2010), pp. 85–112. DOI: [10.4086/toc.2010.v006a005](https://doi.org/10.4086/toc.2010.v006a005) (pp. 1, 2, 7, 8).
- [MP14] 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 Int. Symp. on Theoretical Aspects of Computer Science (STACS)*. Vol. 25. 2014, pp. 542–553. DOI: [10.4230/LIPIcs.STACS.2014.542](https://doi.org/10.4230/LIPIcs.STACS.2014.542) (p. 1).
- [NP85] Jaroslav Nešetřil and Svatopluk Poljak. “On the complexity of the subgraph problem”. In: *Comment. Math. Univ. Carolin.* 26.2 (1985), pp. 415–419 (p. 1).
- [NW11] K. Nakagawa and O. Watanabe. *Gap Between Two Combinatorial Measures for Constant Depth Circuit Complexity of Subgraph Isomorphism*. Technical report. Tokyo Institute of Technology, 2011 (p. 13).

- [Ros08] Benjamin Rossman. “On the constant-depth complexity of k -clique”. In: *Proc. 40th Ann. ACM Symp. on Theory of Computing (STOC)*. 2008, pp. 721–730. DOI: [10.1145/1374376.1374480](https://doi.org/10.1145/1374376.1374480) (pp. 9, 13).
- [Ros10] Benjamin Rossman. “Average-Case Complexity of Detecting Cliques”. Ph.D. MIT, 2010 (p. 13).
- [Ros14] Benjamin Rossman. “The monotone complexity of k -clique on random graphs”. In: *SIAM J. Comput.* 43.1 (2014), pp. 256–279. DOI: [10.1137/110839059](https://doi.org/10.1137/110839059) (p. 13).
- [Ros18] Benjamin Rossman. “Lower bounds for subgraph isomorphism”. In: *Proc. Int. Congress of Mathematicians (ICM)*. Vol. 3. 2018, pp. 3409–3430. DOI: [10.1142/9789813272880_0187](https://doi.org/10.1142/9789813272880_0187). URL: <https://eta.impa.br/dl/051.pdf> (p. 13).
- [RS86] Neil Robertson and P. D. Seymour. “Graph minors. II. Algorithmic aspects of tree-width”. In: *J. Algorithms* 7.3 (1986), pp. 309–322. DOI: [10.1016/0196-6774\(86\)90023-4](https://doi.org/10.1016/0196-6774(86)90023-4) (p. 13).
- [SKS11] Abraham Silberschatz, Henry F. Korth, and S. Sudarshan. *Database System Concepts*. 6th ed. McGraw-Hill Book Company, 2011. Chap. 12.5.4.1 (p. 6).
- [Yao77] Andrew Chi-Chih Yao. “Probabilistic computations: toward a unified measure of complexity”. In: *Proc. 18th Ann. IEEE Symp. on Foundations of Computer Science*. 1977, pp. 222–227. DOI: [10.1109/SFCS.1977.24](https://doi.org/10.1109/SFCS.1977.24) (p. 17).