



Tree-depth and the Formula Complexity of Subgraph Isomorphism

Deepanshu Kush
University of Toronto

Benjamin Rossman
Duke University

April 28, 2020

Abstract

For a fixed “pattern” graph G , the *colored G -subgraph isomorphism problem* (denoted $\text{SUB}(G)$) asks, given an n -vertex graph H and a coloring $V(H) \rightarrow V(G)$, whether H contains a properly colored copy of G . The complexity of this problem is tied to parameterized versions of $P =? NP$ and $L =? NL$, among other questions. An overarching goal is to understand the complexity of $\text{SUB}(G)$, under different computational models, in terms of natural invariants of the pattern graph G .

In this paper, we establish a close relationship between the *formula complexity* of $\text{SUB}(G)$ and an invariant known as *tree-depth* (denoted $\text{td}(G)$). $\text{SUB}(G)$ is known to be solvable by monotone AC^0 formulas of size $O(n^{\text{td}(G)})$. Our main result is an $n^{\Omega(\text{td}(G)^{1/3})}$ lower bound for formulas that are monotone or have sub-logarithmic depth. This complements a lower bound of Li, Razborov and Rossman [8] relating tree-width and AC^0 circuit size. As a corollary, it implies a stronger homomorphism preservation theorem for first-order logic on finite structures [14].

The technical core of this result is an $n^{\Omega(k)}$ lower bound in the special case where G is a complete binary tree of height k , which we establish using the *pathset framework* introduced in [15]. (The lower bound for general patterns follows via a recent excluded-minor characterization of tree-depth [4, 6].) Additional results of this paper extend the pathset framework and improve upon both, the best known upper and lower bounds on the average-case formula size of $\text{SUB}(G)$ when G is a path.

Contents

1	Introduction	2
1.1	Minor-monotonicity	3
1.2	Corollary in finite model theory	4
1.3	Improved bounds for average-case $\text{SUB}(P_k)$	5
1.3.1	Upper bound	5
1.3.2	Lower bound	7
1.4	Related work	7
2	Preliminaries	7
2.1	Graphs	7
2.2	Threshold weightings	7
2.3	Join-trees and parameters $\kappa(G)$ and $\tau(G)$	9
2.4	Observations about Φ_θ	11
2.5	Lower bounds on Φ_θ	11
3	Lower bound $\tau(T_k) = \Omega(k)$	12
3.1	Proof of Theorem 3.9	16
3.2	Proofs of Lemmas 3.10, 3.12, 3.13	18
4	A better potential function	21

5	The pathset framework	22
6	Lower bound $\tau(P_k) \geq \log_{\sqrt{5+5}}(k)$	26
7	Randomized AC^0 formulas computing the product of k permutations	37
7.1	Upper bounds on $\chi(\vec{a}) + \ \vec{a}\ $	41
7.1.1	Recursive doubling	42
7.1.2	Maximally overlapping joins	42
7.1.3	Fibonacci overlapping joins	43
7.2	Tightness of upper bounds	44
8	Open problems	44
A	Appendix: Lower bound $\tau(P_k) \geq \frac{1}{2} \log_{\sqrt{13+1}}(k)$ from [15]	45

1 Introduction

Let G be a fixed “pattern” graph. In the COLORED G -SUBGRAPH ISOMORPHISM PROBLEM, denoted $\text{SUB}(G)$, we are given an n -vertex “host” graph H and a vertex-coloring $c : V(H) \rightarrow V(G)$ as input and required to determine whether or not H contains a properly colored copy of G (i.e., a subgraph $G' \subseteq H$ such that the restriction of c to $V(G')$ constitutes an isomorphism from G' to G). This general problem includes, as special cases, several important problems in parameterized complexity. In particular, $\text{SUB}(G)$ is equivalent (up to AC^0 reductions) to the k -CLIQUE and DISTANCE- k CONNECTIVITY problems when G is a clique or path of order k .

For any fixed pattern graph G , the problem $\text{SUB}(G)$ is solvable by brute-force search in polynomial time $O(n^{|V(G)|})$. Understanding the fine-grained complexity of $\text{SUB}(G)$ — in this context, we mean the exponent of n in the complexity of $\text{SUB}(G)$ under various computational models — for general patterns G is an important challenge that is tied to major open questions including $P =? NP$, $L =? NL$, $NC^1 =? L$, and their parameterized versions ($FPT =? W[1]$, etc.) An overarching goal is to bound the fine-grained complexity of $\text{SUB}(G)$ in terms of natural invariants of the graph G .

Two key invariants arising in this connection are tree-width (tw) and tree-depth (td). The *tree-depth* of G is the minimum height of a rooted forest whose ancestor-descendant closure contains G as a subgraph. This invariant has a number of equivalent characterizations and plays a major role in structural graph theory and parameterized complexity [10]. *Tree-width* is even more widely studied in graph theory and parameterized complexity [5]. It is defined in terms of a different notion of tree decomposition and provides a lower bound on tree-depth ($\text{tw} + 1 \leq \text{td}$).

These two invariants provide well-known upper bounds on the *circuit size* and *formula size* of $\text{SUB}(G)$. To state this precisely, we regard $\text{SUB}(G)$ as a sequence of boolean functions $\{0, 1\}^{|E(G)| \cdot n^2} \rightarrow \{0, 1\}$ where the input encodes a host graph H with vertex set $V(G) \times \{1, \dots, n\}$ under the vertex-coloring that maps (v, i) to v . (Restricting attention to this class of inputs is without loss of generality.) Throughout this paper, we consider circuits and formulas in the unbounded fan-in basis $\{\text{AND}_\infty, \text{OR}_\infty, \text{NOT}\}$; we measure *size* of both circuits and formulas by the number of gates. A circuit or formula is *monotone* if it contains no NOT gates. We use AC^0 as an adjective that means “depth $O(1)$ ” in reference to upper bounds and “depth $o(\log n)$ ” in reference to lower bounds on formula size.¹

Theorem 1.1 (Folklore upper bounds). *For all pattern graphs G , $\text{SUB}(G)$ is solvable by monotone AC^0 circuits (respectively, formulas) of size $O(n^{\text{tw}(G)+1})$ (respectively, $O(n^{\text{td}(G)})$).*

It is conjectured that $\text{SUB}(G)$ requires circuit size $n^{\Omega(\text{tw}(G))}$ for all graphs G ; if true this would imply $FPT \neq W[1]$ and $P \neq NP$ in a very strong way. As evidence for this conjecture, Marx [9] proved a

¹Here and elsewhere, asymptotic notation hides constants that may depend on G . In other contexts, e.g. $\Omega(\text{td}(G))$, hidden constants are absolute.

conditional $n^{\Omega(\text{tw}(G)/\log \text{tw}(G))}$ lower bound assuming the Exponential Time Hypothesis. Providing further evidence, Li, Razborov and Rossman [8] established an unconditional $n^{\Omega(\text{tw}(G)/\log \text{tw}(G))}$ lower bound for AC^0 circuits, via a technique that extends to (unbounded depth) monotone circuits. This result is best stated in terms of a certain graph invariant $\kappa(G)$ introduced in [8]:

Theorem 1.2 (Lower bound on the restricted circuit size of $\text{SUB}(G)$ [8]). *For all pattern graphs G , the circuit size of $\text{SUB}(G)$ — in both the AC^0 and monotone settings — is at least $n^{\kappa(G)-o(1)}$ where $\kappa(G)$ is a graph invariant satisfying $\Omega(\text{tw}(G)/\log \text{tw}(G)) \leq \kappa(G) \leq \text{tw}(G) + 1$.²*

Shifting our focus from circuits to formulas, it is natural to conjecture that $\text{SUB}(G)$ requires formula size $n^{\Omega(\text{td}(G))}$. This statement generalizes the prominent conjecture that $\text{DISTANCE-}k$ CONNECTIVITY requires formula size $n^{\Omega(\log k)}$, which as a consequence implies $NC^1 \neq NL$. (There is also an average-case version of this conjecture which implies $NC^1 \neq L$, as we explain shortly.)

In this paper, we carry out the final step in the proof of an analogous result to Theorem 1.2 that lower bounds the restricted formula size of $\text{SUB}(G)$ in terms of an invariant $\tau(G)$ that is polynomially related to tree-depth:

Theorem 1.3 (Lower bound on the restricted formula size of $\text{SUB}(G)$). *For all patterns graphs G , the formula size of $\text{SUB}(G)$ — in both the AC^0 and monotone settings — is at least $n^{\tau(G)-o(1)}$ where $\tau(G)$ is a graph invariant satisfying $\tilde{\Omega}(\text{td}(G)^{1/3}) \leq \tau(G) \leq \text{td}(G)$.*

The invariant $\tau(G)$ was introduced in [16], where it was also shown that $n^{\tau(G)-o(1)}$ is a lower bound on the formula size of $\text{SUB}(G)$ in the AC^0 and monotone settings. The results of [16] generalized lower bounds for $\text{SUB}(P_k)$ from papers [13, 15], which showed that $\tau(P_k) = \Omega(\log k)$ (where P_k is the path graph of length k). As we will explain shortly, this lower bound for $\tau(P_k)$ implies that $\tau(G) = \Omega(\log \text{td}(G))$ for all graphs G . The contribution of the present paper lies in improving this logarithmic lower bound to a polynomial one by showing $\tau(G) = \tilde{\Omega}(\text{td}(G)^{1/3})$.

Remark 1.4. It is helpful to keep in mind the related inequalities:

$$\text{circuit size} \leq \text{formula size}, \quad \text{tw} + 1 \leq \text{td}, \quad \kappa \leq \tau.$$

It is further known that $\text{td}(G) \leq (\text{tw}(G) + 1) \log |V(G)|$ [10]. A nearly maximal separation between invariants td and tw is witness by *bounded-degree trees* T , which have tree-width 1 but tree-depth $\Omega(\log |V(T)|)$. This class includes paths and complete binary trees, the two families of pattern graphs studied in this paper.

For trees T , we point out that $\text{SUB}(T)$ is computable by monotone AC^0 circuits of size $c(T) \cdot n^2$ for a constant $c(T)$ depending on T . (This follows from Theorem 1.1, since all trees have tree-width 1.) Although formulas are a weaker model than circuits, establishing formula lower bounds for $\text{SUB}(T)$ of the form $n^{\Omega(\log |V(T)|)}$, as we do in this paper, is a subtle task which requires techniques that distinguish formulas from circuits. Accordingly, Theorem 1.3 involves greater machinery than Theorem 1.2. The invariant $\tau(G)$ is also significantly harder to define and analyze compared to $\kappa(G)$.

1.1 Minor-monotonicity

Recall that a graph F is a *minor* of G if F can be obtained from G by a sequence of edge deletions and contractions (i.e., remove an edge and identify its two endpoint). A graph invariant p is said to be *minor-monotone* if $p(F) \leq p(G)$ whenever F is a minor of G . As observed in [8], the complexity of $\text{SUB}(G)$ (under any reasonable class of circuits) is minor-monotone in the following sense:

²It is actually shown that $\kappa(G)$ is at most the *branch-width* of G , an invariant related to tree-width by $\frac{2}{3}(\text{tw}+1) \leq \text{bw} \leq \text{tw}+1$. The relationship between $\kappa(G)$ and $\text{tw}(G)$ was further investigated by Rosenthal [11], who identified the separating example $\kappa(Q) = \Theta(\text{tw}(Q)/\sqrt{\log \text{tw}(Q)})$ for hypercubes Q .

Lemma 1.5. *If F is a minor of G , then there is a reduction from $\text{SUB}(F)$ to $\text{SUB}(G)$ via a monotone projection.³*

In the quest to characterize the complexity of $\text{SUB}(G)$ in terms of invariants of G , it makes sense to focus on minor-monotone ones. Indeed, invariants tw , td , κ , τ are all minor-monotone. This feature is useful in bounding the complexity of $\text{SUB}(G)$. For example, we can combine the result of [2] that every graph with tree-width at least k^9 polylog k contains a $(k \times k)$ -grid minor, with the lower bound $\kappa((k \times k)\text{-grid graph}) = \Omega(k)$ from [1], in order to conclude that $\kappa(G) = \tilde{\Omega}(\text{tw}(G)^{1/9})$ for all graphs G . (Notation $\tilde{O}(\cdot)$ and $\tilde{\Omega}(\cdot)$ suppresses poly-logarithmic factors.) The stronger $\kappa(G) = \Omega(\text{tw}(G)/\log \text{tw}(G))$ bound of Theorem 1.2 is obtained by a more nuanced analysis of the invariant κ .

In a similar manner, we can combine the fact that every graph G contains a path of length $\text{td}(G)$ [10], with the lower bound $\tau(P_k) = \Omega(\log k)$ [15], in order to conclude that $\tau(G) = \Omega(\log \text{td}(G))$ for all graphs G . With the goal of improving this lower bound to $\Omega(\text{poly td}(G))$ (that is, $\Omega(\text{td}(G)^\varepsilon)$ for some constant $\varepsilon > 0$), Kawarabayashi and Rossman [6] established a polynomial excluded-minor characterization of tree-depth, which was subsequently sharpened by Czerwiński, Nadara and Pilipczuk [4].

Theorem 1.6 (Excluded-minor characterization of tree-depth [6, 4]). *Every graph G with tree-depth $\Omega(k^3)$ satisfies at least one of the following:*

- (i) G has tree-width $\geq k$,
- (ii) G contains a path of length 2^k ,
- (iii) G contains a T_k -minor, where T_k is the complete binary tree of height k .

Theorem 1.6 reduces the task of proving $\tau(G) = \Omega(\text{poly td}(G))$ to the task of proving $\tau(T_k) = \Omega(\text{poly } k)$. It is this final step that we tackle in this paper.⁴

Theorem 1.7 (Main result of this paper). $\tau(T_k) = \Omega(k)$.

This lower bound is proved in Section 3 using a certain potential function (described in Section 2), which further reduces our task to a combinatorial problem concerning *join-trees* over T_k , that is, rooted binary trees whose leaves are labeled by edges of T_k . This is the same combinatorial framework as the $\tau(P_k) = \Omega(\log k)$ lower bound of [15]; however, the task of analyzing join-trees over T_k turned out to be significantly harder compared with P_k .

Theorems 1.6 and 1.7 combine to prove Theorem 1.3 (the bound $\tau(G) = \tilde{\Omega}(\text{td}(G)^{1/3})$) as follows. For a graph G with tree-depth $\Omega(k^3)$, we can see that $\tau(G) = \Omega(k/\log k)$ by considering the three cases given by Theorem 1.6:

- (i) If G has tree-width $\geq k$, then $\tau(G) \geq \kappa(G) = \Omega(k/\log k)$ by Theorem 1.2.
- (ii) If G contains a path of length 2^k , then $\tau(G) \geq \tau(P_{2^k}) = \Omega(k)$ by the lower bound of [15].
- (iii) If G contains a T_k -minor, then $\tau(G) \geq \tau(T_k) = \Omega(k)$ by Theorem 1.7.

1.2 Corollary in finite model theory

Theorem 1.3 has a striking consequence in finite model theory, observed in the paper [14].

³That is, for every n , there is a reduction from $\text{SUB}(F)$ to $\text{SUB}(G)$, viewed as boolean functions $\{0, 1\}^{|E(F)| \cdot n} \rightarrow \{0, 1\}$ and $\{0, 1\}^{|E(G)| \cdot n} \rightarrow \{0, 1\}$, via a monotone projection that maps each variable of $\text{SUB}(G)$ to a variable of $\text{SUB}(F)$ or a constant 0 or 1.

⁴Theorem 1.7 delivers on a promise in papers [6, 14, 16], which cite $\tau(T_k) = \Omega(\text{poly } k)$ as an unpublished result of upcoming work. Let us mention that, after finding many devils in the details of an earlier sketch of an $\Omega(\sqrt{k})$ bound by the second author, we worked out an entirely different approach in this paper, which moreover gives a *linear* lower bound (which is tight up to a constant since $\tau(T_k) \leq \text{td}(T_k) = k$).

Corollary 1.8 (Polynomial-rank homomorphism preservation theorem over finite structures). *Every first-order sentence of quantifier-rank r that is preserved under homomorphisms of finite structures is logically equivalent on finite structures to an existential-positive first-order sentence of quantifier-rank $\tilde{O}(r^3)$.*

The polynomial upper bound of Corollary 1.8 improves an earlier *non-elementary* upper bound of [12]. This surprising connection between circuit complexity and finite model theory was in fact the original motivation behind Theorems 1.3 and 1.6, as well as the present paper.

1.3 Improved bounds for average-case $\text{SUB}(P_k)$

Additional results of this paper improve both the average-case upper and lower bounds for $\text{SUB}(P_k)$ [13]. Here *average-case* refers to the p -biased product distribution on $\{0, 1\}^{kn^2}$ where $p = n^{-(k+1)/k}$. This input distribution corresponds to a random graph \mathbf{X} , comprised of $k + 1$ layers of n vertices, where every pair of vertices in adjacent layers is connected by an edge independently with probability p . For this choice of p , the probability that \mathbf{X} contains a path of length k containing one vertex from each layer is bounded away from 0 and 1.

Theorem 1.9 ([15]). *$\text{SUB}(P_k)$ is solvable on \mathbf{X} with probability $1 - o(1)$ by monotone AC^0 formulas of size $n^{\frac{1}{2} \lceil \log_2(k) \rceil + o(1)}$. On the other hand, AC^0 formulas solving $\text{SUB}(P_k)$ on \mathbf{X} with probability ≥ 0.9 require size $n^{\tau(P_k) - o(1)}$ where $\tau(P_k) \geq \frac{1}{2} \log_{\sqrt{13}+1}(k) (\geq 0.22 \log_2(k))$.*

A similar average-case lower bound for (unbounded depth) monotone formulas was subsequently shown in [13]. Precisely speaking, that paper gives an $n^{\frac{1}{2} \tau(P_k) - o(1)}$ lower bound under \mathbf{X} , as well as an $n^{\tau(P_k) - o(1)}$ lower bound under the distribution that, half of the time, is \mathbf{X} and, the other half, is a uniform random path of length k with no additional edges.

1.3.1 Upper bound

The *average-case* upper bound of Theorem 1.9 can be recast, in stronger terms, as a *worst-case randomized* upper bound for the problem of multiplying k ($n \times n$)-permutation matrices Q_1, \dots, Q_k . This problem is solvable by deterministic (non-randomized) AC^0 formulas of size $n^{\log_2(k) + O(1)}$ via the classic “recursive doubling” procedure: recursively compute matrix products $L := Q_1 \cdots Q_{\lceil k/2 \rceil}$ and $R := Q_{\lceil k/2 \rceil + 1} \cdots Q_k$ and then obtain $Q_1 \cdots Q_k = LR$ by a single matrix multiplication.

Randomization lets us achieve quadratically smaller formula size $n^{\frac{1}{2} \log_2(k) + O(1)}$. The idea is as follows. Generate $m := \tilde{O}(\sqrt{n})$ independent random sets $\mathbf{I}_1, \dots, \mathbf{I}_m \subseteq [n]$, each of size \sqrt{n} . Rather than compute all entries of the permutation matrix L using n^2 subformulas, we will encode the information in L more efficiently using $(2 \log n + 1)m^2 = \tilde{O}(n)$ subformulas (note that $\log(n!) = O(n \log n)$ bits are required to encode a permutation matrix). For each $(r, s) \in [m]^2$, we recursively construct

- one subformula that indicates⁵ whether or not there exists a *unique* pair $(a, b) \in \mathbf{I}_r \times \mathbf{I}_s$ such that $L_{a,b} = 1$, and
- $2 \log n$ additional subformulas that give the binary representation of a and b whenever such (a, b) uniquely exists.

Similarly, with respect to the permutation matrix R , for each $(s, t) \in [m]^2$, we have $2 \log n + 1$ recursively constructed subformulas that indicate whether there exists a *unique* pair $(b, c) \in \mathbf{I}_s \times \mathbf{I}_t$ such that $R_{b,c} = 1$, and if so, give the binary representation of b and c . Using these subformulas for subproblems L and R , we construct the corresponding formulas for $Q_1 \cdots Q_k$ which, for each $(r, t) \in [m]^2$, indicate whether there exists a *unique* pair $(a, c) \in \mathbf{I}_r \times \mathbf{I}_t$ such that $R_{a,c} = 1$, and if so, give the binary representation of a and c . These formulas check, for each $s \in [m]$, whether the (r, s) - and (s, t) -subformulas of the L - and R -subproblems output (a, b) and (b', c) , respectively such that $b = b'$. These formulas are therefore larger than

⁵When describing the behavior of randomized formulas in this subsection (using verbs like “indicate”, “output”, etc.), we leave implicit that the description holds *correctly with high probability* for any input.

the subformulas for subproblems L and R by a factor $\tilde{O}(m)$. This implies an upper bound $\tilde{O}(m)^{\lceil \log_2(k) \rceil} = n^{\frac{1}{2} \log_2(k) + O(1)}$ on size and $O(\log k)$ on depth of the resulting randomized AC^0 formulas.

A similar construction solves $\text{SUB}(P_k)$ in the average-case. This yields an upper bound on $\frac{1}{2} \log k + O(1)$ on the parameter $\tau(P_k)$, which we initially guessed might be optimal. However, in the course of trying prove a matching lower bound, we were surprised to discover a better upper bound!

Theorem 1.10. *There exist randomized AC^0 formulas of size $n^{\frac{1}{3} \log_\varphi(k) + O(1)}$ ($\leq n^{0.49 \log_2(k) + O(1)}$), where $\varphi = (\sqrt{5} + 1)/2$ is the golden ratio, which compute the product of k permutation matrices.*

The algorithm generalizes the randomized “recursively doubling” method outlined above. Here we give a brief sketch (full details are given in Section 7). Let $k = \text{Fib}(\ell)$ where $\ell \geq 3$ (i.e., the ℓ^{th} Fibonacci number, which satisfies $\text{Fib}(\ell) = \text{Fib}(\ell - 1) + \text{Fib}(\ell - 2)$). We will represent information about the product $Q_1 \cdots Q_k$ by constructing formulas that enumerate all triples $(a, c, d) \in [n]^3$ such that

$$(1) \quad (Q_1 \cdots Q_{\text{Fib}(\ell-1)})_{a,c} = (Q_{\text{Fib}(\ell-1)+1} \cdots Q_k)_{c,d} = 1.$$

This is accomplished by generating $m := \tilde{O}(n^{1/3})$ independent random sets $\mathbf{I}_1, \dots, \mathbf{I}_m$, each of size $n^{2/3}$, and recording the *unique* triples $(a, c, d) \in \mathbf{I}_r \times \mathbf{I}_t \times \mathbf{I}_u$ for which (1) holds.

The recursive construction breaks into a “left” subproblem on $(Q_1, \dots, Q_{\text{Fib}(\ell-1)})$ and a “right” subproblem on $(Q_{\text{Fib}(\ell-2)+1}, \dots, Q_k)$.⁶ (In contrast to the “recursive doubling” method, here the “left” and “right” subproblems involve overlapping subsequences of permutation matrices.) In the “left” subproblem: for each $(r, s, t) \in [m]^3$, we have

- $3 \log n + 1$ subformulas that indicate whether there exists a *unique* triple $(a, b, c) \in \mathbf{I}_r \times \mathbf{I}_s \times \mathbf{I}_t$ such that $(Q_1 \cdots Q_{\text{Fib}(\ell-2)})_{a,b} = (Q_{\text{Fib}(\ell-2)+1} \cdots Q_{\text{Fib}(\ell-1)})_{b,c} = 1$, and if so, give the binary representation of a, b, c .

In the “right” subproblem: for each $(r, s, t) \in [m]^3$, we have

- $3 \log n + 1$ subformulas that indicate whether there exists a *unique* triple $(b, c, d) \in \mathbf{I}_s \times \mathbf{I}_t \times \mathbf{I}_u$ such that $(Q_{\text{Fib}(\ell-2)+1} \cdots Q_{\text{Fib}(\ell-1)})_{b,c} = (Q_{\text{Fib}(\ell-1)+1} \cdots Q_k)_{c,d} = 1$, and if so, give the binary representation of b, c, d .

The subformulas in the “left” and “right” subproblems may be combined to produce the analogous (left-handed) formulas for the original input (Q_1, \dots, Q_k) : for each $(r, t, u) \in [m]^3$, we construct

- $3 \log n + 1$ subformulas that indicate whether there exists a *unique* triple $(a, c, d) \in \mathbf{I}_r \times \mathbf{I}_t \times \mathbf{I}_u$ such that (1) holds, and if so, give the binary representation of a, c, d .

These formulas check, for each $s \in [m]$, whether the (r, s, t) - and (s, t, u) -subformulas in the “left” and “right” subproblems output triples (a, b, c) and (b', c', d) , respectively, such that $b = b'$ and $c = c'$. These formulas are therefore larger than the subformulas in the “left” and “right” subproblems by a factor $\tilde{O}(m)$. Taking $k = \text{Fib}(3) = 2$ as our base case with formula size $n^{O(1)}$, this gives an upper bound $\tilde{O}(m)^{\ell-3} \cdot n^{O(1)} = n^{\frac{1}{3} \log_\varphi(k) + O(1)}$ for all $k = \text{Fib}(\ell)$ (which extends as well to non-Fibonacci numbers k).

In Section 7, we introduce a broad class of randomized algorithms (based on a simplification of the pathset complexity measure) that generalize both the “recursive doubling” and “Fibonacci overlapping” algorithms outlined above. We also discuss reasons, including experimental data, which suggest that $n^{\frac{1}{3} \log_\varphi(k) + O(1)}$ might in fact be the asymptotically *tight* bound on the randomized formula size of multiplying k permutations.

⁶The “right” subproblem on $(Q_{\text{Fib}(\ell-2)+1}, \dots, Q_k)$ can also be viewed as a “left” subproblem on $(P_1, \dots, P_{\text{Fib}(\ell-1)})$ where P_i is the transpose of Q_{k-i+1} .

1.3.2 Lower bound

The final result of this paper improves the $\tau(P_k) \geq \frac{1}{2} \log_{\sqrt{13}+1}(k)$ ($\geq 0.22 \log_2(k)$) lower bound of Theorem 1.9.

Theorem 1.11. $\tau(P_k) \geq \log_{\sqrt{5}+5}(k) - 1$ ($\geq 0.35 \log_2(k) - 1$)

More significant than the quantitative improvement we obtain in Theorem 1.11 is the fact that our proof further develops *pathset framework* by introducing a new potential function that gives stronger lower bounds on $\tau(G)$. This development and the proof of Theorem 1.11 are presented in detail in Sections 4, 5 and 6.

Since $\frac{1}{3} \log_\varphi(k) = \log_{\sqrt{5}+2}(k)$, our upper and lower bounds are off by exactly 3 in the base of the logarithm. It would be very interesting to completely close this gap.

1.4 Related work

There have been several papers, including [3, 7, 9], which give conditional lower bounds (under ETH and other assumptions) on the *circuit size* of $\text{SUB}(G)$ and its uncolored variant. We are not aware of any conditional hardness results for the *formula size* of $\text{SUB}(G)$. It would be interesting to show that $\text{SUB}(G)$ requires (unrestricted) formula size $n^{\Omega(\text{td}(G))}$ under a natural assumption.

2 Preliminaries

For a natural number n , $[n]$ denotes the set $\{1, \dots, n\}$. For simplicity of presentation, we occasionally omit floors and ceilings, e.g., treating quantities like \sqrt{n} as natural numbers). This is always without loss of parameters in our results. When no base is indicated, $\log(\cdot)$ denotes the base-2 logarithm.

2.1 Graphs

In this paper, *graphs* are simple graphs, i.e., pairs $G = (V(G), E(G))$ where $V(G)$ is a set and $E(G)$ is a subset of $\binom{V(G)}{2}$ (the set of unordered pairs $\{v, w\}$ where v, w are distinct elements of $V(G)$). Unless explicitly stated otherwise, graphs are assumed to be locally finite (i.e., every vertex has finite degree) and without isolated vertices (i.e., $V(G) = \bigcup_{e \in E(G)} e$). For a vertex $v \in V(G)$, $\deg_G(v)$ or simply $\deg(v)$ denotes the degree of v in G .

We regard G as a fixed (possibly infinite) “pattern” graph. F shall consistently denote a finite subgraph of G . We write \subseteq for the subgraph relation and \subset (or sometimes \subsetneq) for the proper subgraph relation. If F is a subgraph of G , then $G \setminus F$ denotes the graph with edge set $E(G) \setminus E(F)$ (and no isolated vertices).

Two important graphs in this paper are paths and complete binary trees. P_k denotes the path graph of length k (with $k + 1$ vertices and k edges). T_k denotes the complete binary tree of height k (with $2^{k+1} - 1$ vertices and $2^{k+1} - 2$ edges). We also consider infinite versions of these graphs. P_∞ is the path graph with vertex set \mathbb{Z} and edge set $\{(i, i + 1) : i \in \mathbb{Z}\}$. T_∞ is the union $\bigcup_{k=1}^\infty T_k$ under the nesting $T_1 \subset T_2 \subset T_3 \subset \dots$ where $\text{Leaves}(T_1) \subset \text{Leaves}(T_2) \subset \text{Leaves}(T_3) \subset \dots$. Thus, T_∞ is an infinite, rootless, layered binary tree, with leaves in layer 0, their parents in level 1, etc.

We use terms *graph invariant* and *graph parameter* interchangeably in reference to real-valued functions on graphs that are invariant under isomorphism.

2.2 Threshold weightings

We describe a family of edge-weightings on graphs G , which in the case of finite graphs correspond to product distributions that are balanced with respect to the problem $\text{SUB}(G)$. (Definitions in this section are adapted from [8].)

Definition 2.1. For any graph G and function $\theta : E(G) \rightarrow \mathbb{R}$, we denote by $\Delta_\theta : \{\text{finite subgraphs of } G\} \rightarrow \mathbb{R}$ the function

$$\Delta_\theta(F) := |V(F)| - \sum_{e \in E(F)} \theta(e).$$

Definition 2.2. A *threshold weighting* for a graph G is a function $\theta : E(G) \rightarrow [0, 2]$ such that $\Delta_\theta(F) \geq 0$ for all finite subgraphs $F \subseteq G$; if G is finite, we additionally require that $\Delta_\theta(G) = 0$.

We refer to the pair (G, θ) as a *threshold-weighted graph*. When θ is fixed, we will at times simply write $\Delta(\cdot)$ instead of $\Delta_\theta(\cdot)$.

Definition 2.3. A *Markov chain* on a graph G is a matrix $[0, 1]^{V(G) \times V(G)}$ that satisfies

- $\sum_{w \in V(G)} M_{v,w} = 1$ for all $v \in V(G)$ and
- $M_{v,w} > 0 \implies \{v, w\} \in E(G)$ for all $v, w \in V(G)$.

Lemma 2.4. *Every Markov chain M on G induces a threshold weighting θ on G defined by*

$$\theta(\{v, w\}) := M_{v,w} + M_{w,v}.$$

This threshold weighting satisfies

$$\Delta_\theta(F) = \sum_{v \in V(F)} \sum_{w \in V(G) : \{v,w\} \notin E(F)} M_{v,w}.$$

We remark that this lemma has a converse (shown in [11]): *Every threshold weighting on G is induced by a (not necessarily unique) Markov chain on G .* Lemma 2.4 also gives us a way to define *threshold weightings* when G is an infinite graph; this will be useful later on.

Example 2.5. Let M be the transition matrix of the uniform random walk on T_k where $k \geq 2$. That is,

$$M_{v,w} := \begin{cases} 1/\deg(v) & \text{if } \{v, w\} \in E(T_k), \\ 0 & \text{otherwise.} \end{cases}$$

For the associated threshold weighting $\theta : E(T_k) \rightarrow [0, 2]$, we have

$$\theta(e) = \begin{cases} 4/3 & \text{if } e \text{ contains a leaf,} \\ 5/6 & \text{if } e \text{ contains the root,} \\ 2/3 & \text{otherwise.} \end{cases}$$

A key property of this θ that we will use later on (Lemma 3.8) is that

$$\Delta_\theta(F) \geq \frac{|V(F) \cap V(T_k \setminus F)|}{3}$$

(that is, $\Delta_\theta(F)$ is at least one-third the size of the boundary of F) for all graphs $F \subseteq T_k$. This is a straightforward consequence of Lemma 2.4, which is also true in the infinite tree T_∞ .

Example 2.6. Let P_k be the path of length k (with $k + 1$ vertices and k edges). The constant function $\theta \equiv 1 + \frac{1}{k}$ is a threshold weighting for P_k . (This is different from the threshold function induced by the uniform random walk on P_k , which has value $1/2$ on the two outer edges of P_k and value 1 on the inner edges.)

This example again makes sense for $k = \infty$. The constant function $E(P_\infty) \mapsto \{1\}$ is a threshold weighting for P_∞ . This threshold function has the nice property that

$$\Delta(F) = |V(F)| - |E(F)| = \#\{\text{connected components of } F\}$$

for all finite subgraphs $F \subset P_\infty$.

Definition 2.7. Let G be a finite graph, let θ be a threshold weighting on G , and let n be a positive integer. We denote by $\mathbf{X}_{\theta,n}$ be the random $V(G)$ -colored graph (i.e., input distribution to $\text{SUB}(G)$) with vertex set $V(G) \times [n]$, vertex-coloring $(v, i) \mapsto v$, and random edge relation given by

$$\mathbb{P}[\{(v, i), (w, j)\} \text{ is an edge of } \mathbf{X}_{\theta,n}] = \begin{cases} 1/n^{\theta(\{v,w\})} & \text{if } \{v, w\} \in E(G), \\ 0 & \text{otherwise,} \end{cases}$$

independently for all $\{(v, i), (w, j)\} \in (V(G) \times [n])^2$.

Lemma 2.8 ([8]). *The probability that $\mathbf{X}_{\theta,n}$ is a YES-instance of $\text{SUB}(G)$ is bounded away from 0 and 1.*

The lower bounds of Theorem 1.2 and 1.3 are in fact average-case lower bounds for $\text{SUB}(G)$ under $\mathbf{X}_{\theta,n}$ for arbitrary threshold weightings θ . Parameters $\kappa(G)$ and $\tau(G)$ are obtained by taking the optimal choice of threshold weighting θ , as we describe in the next subsection.

2.3 Join-trees and parameters $\kappa(G)$ and $\tau(G)$

Parameters $\kappa(G)$ and $\tau(G)$ are defined in terms of a notion called *join-trees* for subgraphs of G . A join-tree is simply a “formula” computing a subgraph of G , starting from individual edges, where union (\cup) is the only operation.

Definition 2.9. A *join-tree over G* is a finite rooted binary tree A together with a labeling $\text{label}_A : \text{Leaves}(A) \rightarrow E(G) \cup \{\perp\}$ (which may also be viewed as a partial function $\text{Leaves} \rightarrow E(G)$). We reserve symbols A, B, C, D, E for join-trees. (F will always denote a subgraph of G .)

The *graph* of A , denoted G_A , is the subgraph of G with edge set $E(G) \cap \text{Range}(\text{label}_A)$. (Note that G_A is always finite.) As a matter of notation, we write $E(A)$ for $E(G_A)$ and $V(A)$ for $V(G_A)$. We also write $\Delta_\theta(A)$ for $\Delta_\theta(G_A)$ where θ is a threshold weighting on G .

We write $\langle \rangle$ for the single-node join-tree labeled by \perp . For $e \in E(G)$, we write $\langle e \rangle$ for the single-node join-tree labeled by e . For join-trees B and C , we write $\langle B, C \rangle$ for the join-tree consisting a root with B and C as children (with the inherited labels, i.e., $\text{label}_{\langle B, C \rangle} = \text{label}_B \cup \text{label}_C$). Note that $G_{\langle B, C \rangle} = G_B \cup G_C$.

Every join-tree A is clearly either $\langle \rangle$, or $\langle e \rangle$ where $e \in E(G)$, or $\langle B, C \rangle$ where B, C are join-trees. In the first two cases, we say that A is *atomic*; in the third case, we say that A is *non-atomic*.

We say that B is a *child* of A if $A \in \{\langle B, C \rangle, \langle C, B \rangle\}$ for some C . We say that D is a *sub-join-tree* of A (denoted $D \preceq A$) if $D = \langle \rangle$ or $D = A$ or D is a sub-join-tree of a child of A . We say that D is a *proper sub-join-tree* (denoted $D \prec A$) if $D \preceq A$ and $D \neq A$.

We are now able to define the invariant $\kappa(G)$ in Theorem 1.2, which lower bounds the restricted circuit size of $\text{SUB}(G)$. (In fact, $\kappa(G)$ also provides a nearly tight upper bound on the average-case AC^0 circuit size of $\text{SUB}(G)$ [11].)

Definition 2.10 (The invariant $\kappa(G)$ of Theorem 1.2). For finite graphs G , let

$$\kappa(G) := \max_{\text{threshold weightings } \theta \text{ for } G} \min_{\text{join-trees } A \text{ with graph } G} \max_{B \preceq A} \Delta_\theta(B).$$

The invariant of $\tau(G)$ of Theorem 1.3 is significantly more complicated to define. We postpone the definition to Section 5 and, in the meantime, focus on a simpler “potential function” on join-trees, denoted $\Phi_\theta(A)$, which we use to lower bound $\tau(G)$. In order to state the definition of $\Phi_\theta(A)$, we require the following operation \ominus (“restriction away from”) on graphs and join-trees.

Definition 2.11 (The operation \ominus on graphs and join-trees). For $F \subseteq G$ and a subset $S \subseteq V(G)$, we denote by $F \ominus S$ the graph consisting of the connected components of F that are vertex-disjoint from S .

For a join-tree A , we denote $A \ominus S$ the join-tree with the same rooted tree structure as A and leaf labeling function

$$\text{label}_{A \ominus S}(l) := \begin{cases} \text{label}_A(l) & \text{if } \text{label}_A(l) \in E(G_A \ominus S), \\ \perp & \text{otherwise.} \end{cases}$$

That is, $A \ominus S$ deletes all labels except for edges in $G_A \ominus S$. Note that $G_{A \ominus S} = G_A \ominus S$.

As a matter of notation, if B is another join-tree, we write $A \ominus B$ for $A \ominus V(B)$ and $A \ominus (S \cup B)$ for $A \ominus (S \cup V(B))$.




Figure 1: An example where A is a join-tree whose graph G_A consists of two paths of length 2 with edges e, f, g, h . S is the set containing just the external endpoint of h . The join-tree $A \ominus S$ is depicted to the right.

Definition 2.12 (The potential function Φ_θ on join-trees). Fix a threshold weighting θ on a graph G . The potential function $\Phi_\theta : \{\text{join-trees over } G\} \rightarrow \mathbb{R}_{\geq 0}$ is the unique pointwise minimum function satisfying the following inequalities for all join-trees A, B, C, D :

- (†) $\Phi_\theta(A) \geq \Phi_\theta(D) + \Delta_\theta(C \ominus D) + \Delta_\theta(A \ominus (C \cup D))$ if $A \in \{\langle B, C \rangle, \langle C, B \rangle\}$ and $D \preceq B$,
- (‡) $\Phi_\theta(A) \geq \frac{1}{2} \left(\Phi_\theta(D) + \Phi_\theta(E \ominus D) + \Delta_\theta(A) + \Delta_\theta(A \ominus (D \cup E)) \right)$ if $D, E \prec A$.

Alternatively, $\Phi_\theta(A)$ has the following recursive characterization:

- If A is an atomic join-tree, then

$$\Phi_\theta(A) := \Delta_\theta(A) = \begin{cases} 0 & \text{if } A = \langle \rangle, \\ 2 - \theta(e) & \text{if } A = \langle e \rangle \text{ where } e \in E(G). \end{cases}$$

(Obs: In the case $A = \langle e \rangle$, the constraint $\Phi_\theta(A) \geq \Delta_\theta(A)$ is forced by (‡) where $B = C = \langle \rangle$.)

- If $A = \langle B, C \rangle$, then

$$\Phi_\theta(A) := \max \left\{ \begin{array}{l} \max_{D \preceq B} \Phi_\theta(D) + \Delta_\theta(C \ominus D) + \Delta_\theta(A \ominus (C \cup D)), \\ \max_{D \preceq C} \Phi_\theta(D) + \Delta_\theta(B \ominus D) + \Delta_\theta(A \ominus (B \cup D)), \\ \max_{D, E \prec A} \frac{1}{2} \left(\Phi_\theta(D) + \Phi_\theta(E \ominus D) + \Delta_\theta(A) + \Delta_\theta(A \ominus (D \cup E)) \right) \end{array} \right\}.$$

That is, at least one of inequalities (†) or (‡) is tight for each join-tree A .

This definition, although opaque at first, will be clarified later (in Sections 4 and 5). The key property of $\Phi_\theta(A)$ is that it provides a lower bound the invariant $\tau(G)$, which in turn provides a lower bound on the restricted formula complexity of $\text{SUB}(G)$.

Theorem 2.13 ([16]). *The invariant $\tau(G)$ of Theorem 1.3 satisfies*

$$\tau(G) \geq \max_{\text{threshold weightings } \theta \text{ for } G} \min_{\text{join-trees } A \text{ with graph } G} \Phi_\theta(A).$$

The definition of $\tau(G)$ and proof of Theorem 2.13 are postponed to Section 5. First, in Section 3, we will present our combinatorial main lemma, which gives a lower bound on $\Phi_\theta(A)$ for all join-trees with graph T_k under the threshold weighting θ of Example 2.5.

2.4 Observations about Φ_θ

Note that inequality (†) implies $\Phi_\theta(A) \geq \Phi_\theta(D)$ for all $D \preceq A$ (since $\Delta_\theta(\cdot)$ is nonnegative). Also note that inequality (‡) implies $\Phi_\theta(A) \geq \Delta_\theta(A)$ in the special case $B = C = \langle \rangle$ (since $\Phi_\theta(\langle \rangle) = 0$ and $A \ominus$ (the empty graph) = A). Combining these observations, we see that $\Phi_\theta(A) \geq \Delta_\theta(D)$ for all $D \preceq A$. It follows that $\tau(G) \geq \kappa(G)$ for all graphs G , which makes sense in light of the fact that $\kappa(G)$ bounds circuit size and $\tau(G)$ bounds formula size.

Next, observe that $\Phi_\theta(A)$ always equals either $\Phi_\theta(D) + (\text{some } \Delta_\theta(\cdot)\text{-terms})$ or $\frac{1}{2}(\Phi_\theta(D) + \Phi_\theta(E \ominus D)) + (\text{some } \Delta_\theta(\cdot)\text{-terms})$ where D and E are proper sub-join-trees of A . This can be expanded out until we get a nonnegative linear combination of $\Delta_\theta(\cdot)$ -terms. Looking closely, we see that

$$\Phi_\theta(A) = \sum_{F \subseteq G} c_F \cdot \Delta_\theta(F)$$

where coefficients c_F (which depend on both θ and A) are nonnegative dyadic rational numbers coming from the tight instances of inequalities (†) and (‡). We may further observe, for any $v \in V(G)$, that

$$\sum_{F \subseteq G: v \in V(F)} c_F \leq 1.$$

This is easily shown by induction using the fact that graphs F_1 and $F_2 \ominus F_1$ and $F_3 \ominus (F_1 \cup F_2)$ are pairwise disjoint for any $F_1, F_2, F_3 \subseteq G$.

One consequence of this observation is the following lemma, which we will use in Sections 3 and 6.

Lemma 2.14. *Suppose (G, θ) and (G^*, θ^*) are threshold-weighted graphs such that $G \subseteq G^*$ and $\theta^*(e) \leq \theta(e)$ for all $e \in E(G)$. Then for any join-tree A with graph G , we have*

$$\Phi_\theta(A) \geq \Phi_{\theta^*}(A) - \sum_{e \in E(G)} (\theta(e) - \theta^*(e)).$$

Proof. Let $\{c_F\}_{F \subseteq G}$ be nonnegative dyadic rationals — arising from the tight instances of inequalities (†) and (‡) in the recursive definition of $\Phi_{\theta^*}(A)$ — such that $\Phi_{\theta^*}(A) = \sum_{F \subseteq G} c_F \cdot \Delta_{\theta^*}(F)$. We may apply inequalities (†) and (‡) in the exact same way to get the bound $\Phi_\theta(A) \geq \sum_{F \subseteq G} c_F \cdot \Delta_\theta(F)$. We now have

$$\begin{aligned} \Phi_{\theta^*}(A) - \Phi_\theta(A) &\leq \sum_{F \subseteq G} c_F (\Delta_{\theta^*}(F) - \Delta_\theta(F)) \\ &= \sum_{F \subseteq G} c_F \sum_{e \in E(F)} (\theta(e) - \theta^*(e)) \\ &= \sum_{e \in E(G)} (\theta(e) - \theta^*(e)) \sum_{F \subseteq G: e \in E(F)} c_F \\ &\leq \sum_{e \in E(G)} (\theta(e) - \theta^*(e)), \end{aligned}$$

using the fact that $\theta(e) - \theta^*(e) \geq 0$ and $\sum_{F \subseteq G: v \in V(F)} c_F \leq 1$ for all $v \in V(G)$. □

2.5 Lower bounds on Φ_θ

Having introduced the potential function Φ_θ and described its connection to τ in Theorem 2.13, we conclude this section by briefly explaining how it is used to derive lower bounds $\tau(P_k)$ and $\tau(T_k)$. The main combinatorial lemma behind the lower bound of Theorem 1.9 is the following:

Lemma 2.15 ([15]). *Let θ be the constant $1 + \frac{1}{k}$ threshold weighting on P_k . For every join-tree A with graph P_k , we have $\Phi_\theta(A) \geq \frac{1}{2} \log_{\sqrt{13}+1}(k)$. (Therefore, $\tau(G) \geq \frac{1}{2} \log_{\sqrt{13}+1}(k)$.)*

The proof is included in Appendix A, for the sake of comparison with our two lower bounds below. We remark that this proof makes crucial use of both (\dagger) and (\ddagger) ; it was shown in [15] that no lower bound better than $\Phi_\theta(A) = \Omega(1)$ is provable using (\dagger) alone or (\ddagger) alone.

Our lower bound $\tau(T_k) = \Omega(k)$ (Theorem 1.7) is an immediate consequence of the following:

Lemma 2.16. *Let θ be the threshold weighting arising from the uniform random walk on T_k (Example 2.5). For every join-tree A with graph T_k , we have $\Phi_\theta(A) \geq k/30 - 1/5$.*

Our proof, given in the next section, is purely graph-theoretic. Interestingly, the argument essentially uses only inequality (\ddagger) ; we do not require (\dagger) , other than in the weak form $\Phi_\theta(A) \geq \Phi_\theta(D)$ for all $D \prec A$.

It is worth mentioning that the choice of threshold weighting is important in Lemma 2.16. A different, perhaps more obvious, threshold weighting is the constant function with value $\frac{|V(T_k)|}{|E(T_k)|} (= \frac{2^{k+1}-1}{2^{k+1}-2})$. With respect to this threshold weighting, no lower bound better than $\Omega(1)$ is possible.

Finally, our improved lower bound $\tau(P_k) \geq \log_{\sqrt{5+5}}(k)$ (Theorem 1.11) is obtained via the following lemma. This result involves a 2-parameter extension of $\Phi_\theta(A)$ denoted $\Phi_\theta(A|S)$ (where $S \subseteq V(G)$), which we introduce in Section 4.

Lemma 2.17. *Let θ be the constant $1 + \frac{1}{k}$ threshold weighting on P_k . For every join-tree A with graph P_k , we have $\Phi_\theta(A|\emptyset) \geq \log_{\sqrt{5+5}}(k) - 1$.*

This lemma is proved in Section 6, after we show how $\Phi_\theta(\cdot|\cdot)$ provides a lower bound on $\tau(\cdot)$ in Section 5.

3 Lower bound $\tau(T_k) = \Omega(k)$

We fix the infinite pattern graph T_∞ with the threshold weighting θ induced by the uniform random walk. Recall that $T_\infty = \bigcup_{k=1}^\infty T_k$ under a nesting $T_1 \subset T_2 \subset T_3 \subset \dots$ with $\text{Leaves}(T_1) \subset \text{Leaves}(T_2) \subset \text{Leaves}(T_3) \subset \dots$. F, G, H will represent finite subgraphs of T_∞ , and A, B, C, D, E will be join-trees over T_∞ . (In particular, note that G no longer denotes the ambient pattern graph.)

We next recall the definition of θ from Example 2.5. Let $M \in [0, 1]^{V(T_\infty) \times V(T_\infty)}$ be the transition matrix of the uniform random walk on T_∞ , that is,

$$M_{v,w} = \begin{cases} 1 & \text{if } \{v, w\} \in E(T_\infty) \text{ and } v \text{ is a leaf,} \\ 1/3 & \text{if } \{v, w\} \in E(T_\infty) \text{ and } v \text{ is a non-leaf,} \\ 0 & \text{if } \{v, w\} \notin E(T_\infty). \end{cases}$$

This induces the threshold weighting $\theta : E(T_\infty) \rightarrow [0, 2]$ given by

$$\theta(\{v, w\}) := M_{v,w} + M_{w,v} = \begin{cases} 2/3 & \text{if } v \text{ or } w \text{ is a leaf of } T_\infty, \\ 4/3 & \text{otherwise.} \end{cases}$$

Since θ is fixed, we will suppress it when writing $\Delta(F)$ and $\Phi(A)$.

Definition 3.1. For all $k \geq 0$, let

$$V_k := \{v \in V(T_k) : v \text{ has distance } k \text{ from a leaf}\}.$$

Thus, V_0 is the set of leaves in T_∞ , V_1 is the set of parents of leaves, etc. Note that $V(T_\infty) = \bigcup_{k=0}^\infty V_k$. We shall refer to the V_k as the various *levels* of T_∞ .

For $k \geq 1$ and $x \in V_k$, let $T_x \subset T_\infty$ be the complete binary tree of height k rooted at x (in the case $k = 0$, we regard T_x as a single isolated vertex). We denote by T_x^+ the graph obtained from T_x by including an extra edge between x and its parent. Note that

$$|V(T_x)| = 2^{k+1} - 1, \quad |V(T_x^+)| = 2^{k+1}, \quad |E(T_x)| = 2^{k+1} - 2, \quad |V(T_x^+)| = 2^{k+1} - 1.$$

For $j \in \{0, \dots, k\}$, let $V_j(T_x) := V_j \cap V(T_x)$.

Observation 3.2. If $x \in V_k$, then for $j \in \{0, \dots, k\}$, $|V_j(T_x)| = 2^{k-j}$.

We next define two useful parameters of finite subgraphs of T_∞ .

Definition 3.3 (Max-complete height). For a finite subgraph F of T_∞ , define the *max-complete height* $\lambda(F)$ to be the maximum $k \in \mathbb{N}$ for which there exists $x \in V_k$ with $T_x \subseteq F$; $\lambda(F)$ is defined to be zero when no such x exists (in particular, this happens when $V(F) \cap V_0 = \emptyset$).

Observation 3.4. For any $x \in V_k$, $\lambda(T_x) = \lambda(T_x^+) = k$.

Definition 3.5 (Boundary size). Let $\partial(F)$ denote the size of *boundary* of F in T_∞ :

$$\partial(F) := |V(F) \cap V(T_\infty \setminus F)|.$$

Observation 3.6. For any $x \in V_k$, we have $\partial(T_x) = \partial(T_x^+) = 1$, as the boundaries in the respective graphs are simply the singletons $\{x\}$ and $\{\text{parent}(x)\}$. Another example is as follows: if $x \in V_k$ for some $k \geq 2$ and F is the subgraph of T_x induced by the set of vertices $V(T_x) \setminus V_0$, then $\partial(F) = 2^{k-1} + 1$ as all vertices in $V_1(T_x)$ (along with x) lie in the boundary of F .

Definition 3.7 (Grounded and ungrounded subgraphs of T_∞). Let F, H be finite subgraphs of T_∞ . We say that F is *grounded* if it is connected and $V(F) \cap V_0 \neq \emptyset$ (that is, F is a tree, at least one of whose leaves is also a leaf of T_∞). We say that H is *ungrounded* if it is non-empty and connected and $V(H) \cap V_0 = \emptyset$ (that is, H is a non-empty tree, none of whose leaves is a leaf of T_∞).




Figure 2: Example of a grounded graph $F \subset T_\infty$. The dotted lines indicate the various levels that $V(F)$ intersects and dashed lines indicate (some of the) edges in $T_\infty \setminus F$. The nodes colored white lie in the boundary of F , and therefore $\partial(F) = 7$ and $\Delta(F) = 11/3$. The max-complete height $\lambda(F) = 2$, since $T_x \subset F$ for $x \in V_2$ and x is the highest such node. The subgraph $H = F \cap T_y$ is ungrounded with $\lambda(H) = 0$.

We shall think of the function $\partial(F)$ as essentially a proxy for $\Delta(F)$, as it has the advantage of having a simple combinatorial definition. This is justified by the following:

Lemma 3.8. For every $F \subset T_\infty$, we have $\Delta(F) \geq \partial(F)/3$.

It also holds that $\Delta(F) \leq 2\partial(F)/3$ for F without isolated vertices (or $\Delta(F) \leq \partial(F)$ if we allow isolated vertices), but we will not need this upper bound.

Proof. Since θ is the threshold weighting induced by M , Lemma 2.4 tells us

$$\Delta(F) = \sum_{v \in V(F)} \sum_{w \in V(T_\infty) : \{v, w\} \in E(T_\infty \setminus F)} M_{v, w}.$$

Note that $v \in V(F)$ contributes to this sum if, and only if, it belongs to the boundary of F (i.e., $v \in V(F) \cap V(T_\infty \setminus F)$). Since $M_{v, w} \geq 1/3$ whenever $\{v, w\} \in E(F)$, the claim follows. \square

We are now ready to state the main theorem of the section.

Theorem 3.9. *Let $\varepsilon = 1/30$ and $\delta = 2/5$. Then for every join-tree A ,*

$$\Phi(A) \geq \varepsilon\lambda(A) + \delta\Delta(A).$$

Theorem 3.9 directly implies Lemma 2.16, which in turn yields the lower bound $\tau(T_k) = \Omega(k)$ of Theorem 1.7. To see why, let θ' be the threshold weighting on T_k coming from the uniform random walk (Example 2.5). Viewing T_k a subgraph of T_∞ , note that $\sum_{e \in E(T_k)} (\theta'(e) - \theta(e)) = 2(\frac{5}{6} - \frac{2}{3}) = \frac{1}{3}$. For any join-tree A with graph T_k , Lemma 2.14 and Theorem 3.9 now imply

$$\Phi_{\theta'}(A) \geq \Phi(A) - \frac{1}{3} \geq \frac{1}{30}\lambda(T_k) + \frac{2}{5}\Delta(T_k) - \frac{1}{3} = \frac{1}{30}k + \frac{2}{5} \cdot \frac{1}{3} - \frac{1}{3} = \frac{1}{30}k - \frac{1}{5}.$$

We proceed with a few definitions and lemmas needed for the proof of Theorem 3.9 in Section 3.1. In order to present the main arguments first, proofs of three auxiliary lemmas (3.10, 3.12, 3.13) are postponed to Section 3.2.

Lemma 3.10. *Let H be a non-empty finite subgraph of T_∞ , all of whose components are ungrounded. Then $\partial(H) \geq \frac{1}{2}(|E(H)| + 3)$.*

We make a note of its following corollary here.

Corollary 3.11. *Suppose H be a finite connected subgraph of T_∞ and $y \in V(H)$ such that $E(H) \cap E(T_y) \neq \emptyset$ and H does not contain any path from y to a leaf of T_y . Then $\partial(H) \geq \frac{1}{2}(|E(H) \cap E(T_y)| + 1)$.*

Proof. Let F be the graph with edge set $E(F) := E(H) \cap E(T_y)$. Note that F is non-empty, connected and ungrounded. Observe that $\partial(H) \geq \partial(F) - 1$ because all vertices in the boundary of H , with the only possible exception of y , also lie in the boundary of F . Hence by Lemma 3.10, $\partial(H) \geq \frac{1}{2}(|E(H) \cap E(T_y)| + 1)$. \square

The second auxiliary lemma gives a useful inequality relating $\partial(G)$, $\lambda(G)$ and $|E(G)|$.

Lemma 3.12. *For every finite subgraph G of T_∞ , we have $\lambda(G) + \partial(G) \geq \log(|E(G)| + 1)$.*

(This is tight when $G = T_x^+$ for some $x \in V_k$, in which case $\lambda(G) = k$ and $\partial(G) = 1$ and $|E(G)| = 2^{k+1} - 1$.) The third auxiliary lemma shows that subgraphs of T_k that contain at most half the edges of T_k and have boundary size j ($\leq k/2$) have empty intersection with a large complete subtree of T_k of height $k - j$.

Lemma 3.13. *Let $x \in V_k$ and suppose $G \subseteq T_x$ such that $|E(G \cap T_x)| \leq 2^k - 1$ and $\partial(G) \leq k/6$. Then there exists a vertex $z \in V_{k-\partial(G)}(T_x)$ such that $E(G) \cap E(T_z^+) = \emptyset$.*

We now state and prove the main lemma used in the proof of Theorem 3.9.

Lemma 3.14. *For any integers $1 \leq t \leq \ell$, let $z \in V_\ell$ and suppose A is a join-tree such that $T_z \subseteq G_A$. Then one the following conditions holds:*

- (i) *There exists $D \preceq A$ such that $\partial(D) \geq t$ and $\lambda(D) + \partial(D) \geq \ell$.*
- (ii) *There exists $C \prec A$ with $\lambda((C \cap T_z) \ominus \{z\}) + \partial((C \cap T_z) \ominus \{z\}) \geq \ell - t$.*
- (iii) *There exists $E \prec A$ such that $\partial(E) \geq \ell - t$.*

Proof. Descend in the join-tree A until reaching a $B \preceq A$ such that B contains a path P from z to a leaf of T_z , but no $B' \prec B$ contains a path from z to a leaf of T_z .

Let $j \in \{1, \dots, \ell\}$ be maximal such that there exists $y \in V_j(T_z) \cap P$ such that $T_y \subseteq G_B$. We claim that $\partial(B) \geq \ell - j$. To see this, note that for every vertex $v \neq y$ on the path from z to y (a subpath of P), if $c(v)$ denotes the child of v that is not on the path P , then G does not contain $T_{c(v)}^+$ (because if it did then $\lambda(G) > k$). As a result, it must be the case that for every vertex $v \neq y$ on the path from z to y , some vertex in $V(G_B) \cap V(T_{c(v)}^+)$ lies in the boundary of G and so, $\partial(G) \geq \ell - j$.

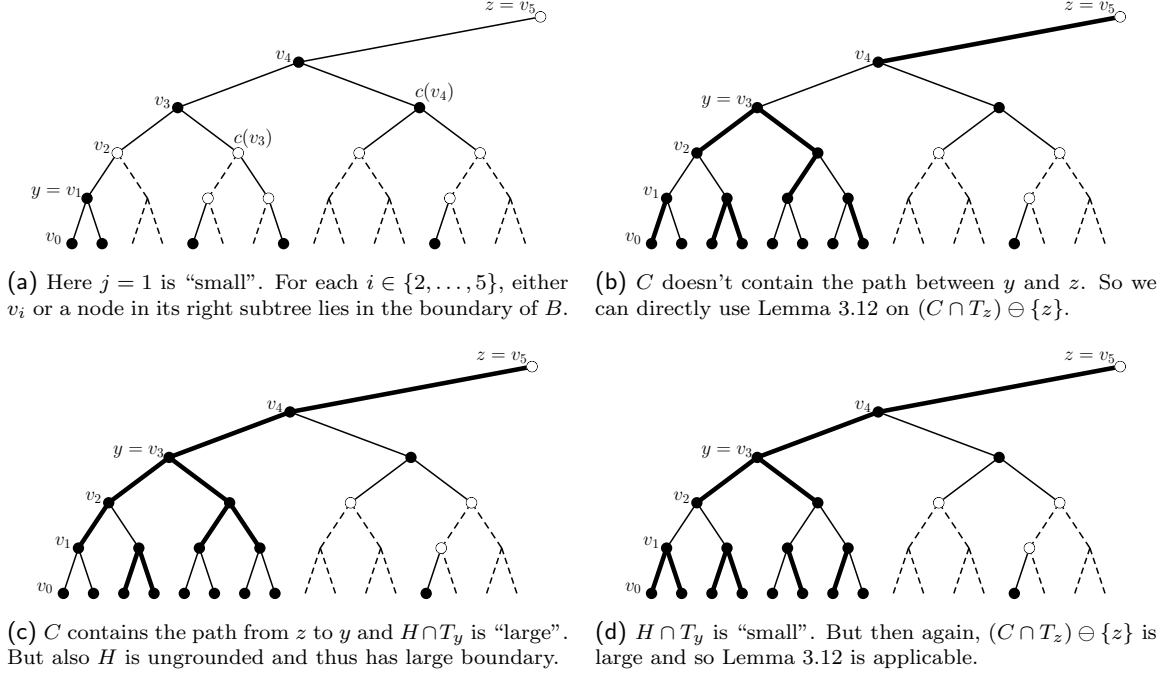


Figure 3: An illustration with $\ell = 5$, $t = 3$. In each example, solid lines denote the edges of B , thickly shaded lines denote the subgraph C , and dashed lines denote edges of $T_z \setminus B$. Nodes shaded white lie in the boundary of B (thought not necessarily the boundary of C). P is the path $v_0v_1v_2v_3v_4v_5$.

As an illustration of this argument, consider the graph G_B in Figure 3a. Then for every vertex $v \in \{v_2, \dots, v_5\}$, either v itself is in the boundary (like v_2 and v_5) or some vertex in $T_{c(v)}$ is in the boundary (for v_3 , it is $c(v_3)$ for example and for v_4 , it is either child of $c(v_4)$).

Consider the case that $j \leq \ell - t$ (again see Figure 3a). Letting $D := B$, we have $\partial(D) \geq \ell - j \geq t$ and $\lambda(D) \geq j$, so condition (i) is satisfied. We shall therefore proceed under the assumption that

$$j \geq \ell - t + 1.$$

Since $T_y \subseteq G_B$, at least one child C of B satisfies

$$|E(C) \cap E(T_y)| \geq \frac{1}{2}|E(T_y)| = 2^j - 1.$$

Fix one such C .

Consider the case that C does not contain the path between z and y (see Figure 3b). Then $C \cap T_y \subseteq C \ominus \{z\}$, so by Lemma 3.12,

$$\lambda((C \cap T_z) \ominus \{z\}) + \partial((C \cap T_z) \ominus \{z\}) \geq \log(|E(C \cap T_y)| + 1) \geq \log(2^j) \geq \ell - t + 1.$$

In this case, we satisfy condition (ii). We shall therefore proceed under the assumption that C contains the path between z and y .

Note that C does not contain a path from y to any leaf of T_y (since otherwise C would contain a path from z to a leaf of T_z , contradicting the way we choose $B \preceq A$). Let H be the connected component of G_C that contains y (and hence also contains the path between z and y).

We now consider two final cases, depending on the size of $|E(H) \cap E(T_y)|$. First, assume $|E(H) \cap E(T_y)| \geq 2(\ell - t)$ (see Figure 3c). In this case, Corollary 3.11 implies

$$\partial(C) \geq \partial(H) \geq \frac{1}{2}(|E(H) \cap E(T_y)| + 1) \geq \ell - t.$$

We satisfy condition (iii) setting $E := C$.

Finally, assume $|E(H) \cap E(T_y)| \leq 2(\ell - t) - 1$ (see Figure 3d). We have

$$\begin{aligned} |E((C \cap T_z) \ominus \{z\})| &\geq |E(C) \cap E(T_y)| - |E(H) \cap E(T_y)| \\ &\geq (2^j - 1) - (2(\ell - t) - 1) \\ &\geq 2^{\ell-t+1} - 2(\ell - t). \end{aligned}$$

Lemma 3.12 now implies

$$\begin{aligned} \lambda((C \cap T_z) \ominus \{z\}) + \partial((C \cap T_z) \ominus \{z\}) &\geq \log(|E((C \cap T_z) \ominus \{z\})| + 1) \\ &\geq \log(2^{\ell-t+1} - 2(\ell - t) + 1) \\ &> \ell - t \end{aligned}$$

since $2^{x+1} - 2x + 1 > 2^x$ for all $x \geq 0$. Therefore, condition (ii) is again satisfied in this final case. \square

3.1 Proof of Theorem 3.9

We now prove Theorem 3.9: the lower bound $\Phi(A) \geq \varepsilon\lambda(A) + \delta\Delta(A)$ where $\varepsilon = 1/30$ and $\delta = 2/5$.

We argue by a structural induction on join-trees A . First, consider the case that G_A is empty, then $\Phi(A) = 0$ and $\varepsilon\lambda(A) + \delta\Delta(A) = 0$. We shall assume that G_A is non-empty.

Next consider the base case where A is the atomic join-tree $\langle e \rangle$ for an edge $e \in E(T_\infty)$. In this case, we have $\lambda(A) \leq 1$ and

$$\Phi(A) = \Delta(A) = \begin{cases} 2/3 & \text{if } e \text{ contains a leaf,} \\ 4/3 & \text{otherwise.} \end{cases}$$

Therefore, $\varepsilon\lambda(A) + \delta\Delta(A) \leq (1/30) + (2/5)(4/3) = 17/30$. We are done, since $\Phi(A) \geq 2/3 > \varepsilon\lambda(A) + \delta\Delta(A)$.

From now on, let A be a non-atomic join-tree with whose graph is non-empty. Let

$$k := \lambda(A).$$

Our goal is thus to prove $\Phi(A) \geq \varepsilon k + \delta\Delta(A)$, which we do by analyzing numerous cases.

Consider first the case that $k = 0$. In this case, we clearly have $\Phi(A) \geq \varepsilon k + \delta\Delta(A)$ (since $\Phi \geq \Delta \geq 0$ and $\delta < 1$). So shall proceed on the assumption that $k \geq 1$.

Since $\Phi \geq \Delta$, we are done if $\Delta(A) \geq \varepsilon k + \delta\Delta(A)$. So we shall proceed on the additional assumption that

$$(2) \quad \Delta(A) \leq \frac{\varepsilon}{1-\delta} k = \frac{1}{18} k.$$

Fixing $x \in T_k$ with $T_k \subseteq G_A$: By definition of $\lambda(A)$, there exists a vertex $x \in T_k$ such that $T_k \subseteq G_A$. Let us fix any such x .

Fixing $B \preceq A$ with $2^{k-1} \leq |E(B) \cap E(T_x)| \leq 2^k - 1$: We next fix a sub-join-tree $B \preceq A$ satisfying $2^{k-1} \leq |E(B) \cap E(T_x)| \leq 2^k - 1$. To see that such B exists, first note that $|E(A) \cap E(T_x)| = |E(T_x)| = 2^{k+1} - 2$. Consider a walk down A which at each step descends to a child C which maximizes $|E(C) \cap E(T_x)|$. This quantity shrinks by a factor $\geq 1/2$ at each step, eventually reaching size 1. Therefore, at some stage, we reach a sub-join-tree B such that the intersection size $|E(B) \cap E(T_x)|$ is between 2^{k-1} and $2^k - 1$.

Observe that $\Phi(A) \geq \Phi(B) \geq \Delta(B)$ (by \dagger) and the fact that $\Phi \geq \Delta$ for all join-trees). Therefore, we are done if $\Delta(B) \geq \varepsilon k + \delta\Delta(A)$. So we shall proceed under the additional assumption that

$$(3) \quad \Delta(B) \leq \varepsilon k + \delta\Delta(A) = \frac{1}{30} k + \frac{2}{5} \Delta(A) \leq \left(\frac{1}{30} + \frac{2}{5} \cdot \frac{1}{18} \right) k = \frac{1}{18} k.$$

Since $|E(B)| \geq |E(B) \cap E(T_x)| \geq 2^{k-1}$, Lemma 3.12 tells us that $\lambda(B) + \partial(B) \geq k - 1$. We make note of the fact that this implies

$$\begin{aligned}
(4) \quad \Phi(B) &\geq \varepsilon\lambda(B) + \delta\Delta(B) && \text{(induction hypothesis)} \\
&\geq \varepsilon\left(k - \partial(B) - 1\right) + \delta\Delta(B) \\
&\geq \varepsilon k + (\delta - 3\varepsilon)\Delta(B) - \varepsilon && \text{(using } -\partial \geq -3\Delta\text{)}.
\end{aligned}$$

Fixing $z \in V_{k-\partial(B)}(T_x)$ **with** $E(B) \cap E(T_z^+) = \emptyset$: Note that $\partial(B) \leq 3\Delta(B) \leq k/6$ by (3). Since $|E(B) \cap E(T_x)| \leq 2^k - 1$, the hypotheses of Lemma 3.13 are satisfied with respect to the vertex x and the graph $G_B \cap T_x$. Therefore, we may fix a vertex $z \in V_{k-\partial(B)}(T_x)$ such that $E(B) \cap E(T_z^+) = \emptyset$.

We next introduce a parameter

$$t := 6\Delta(A) + 3\Delta(B).$$

Note that t is an integer, since 3Δ is integral. Our choice of parameters moreover ensure that $1 \leq t \leq k/2$, since $\Delta(A), \Delta(B) \leq k/18$ by (2),(3) and $\Delta(A), \Delta(B) \geq 1/3$ for all nonempty graphs. Since $k - \partial(B) \geq k - 3\Delta(B) \geq 5k/6$, it follows that $t < k - \partial(B)$.

Since $z \in V_{k-\partial(B)}$ and $T_z \subseteq T_x \subseteq G_A$, Lemma 3.14 (with respect to z and $1 \leq t < \ell := k - \partial(B)$) tells us that one of the following conditions holds:

- (i) There exists $D \preceq A$ such that $\partial(D) \geq t$ and $\lambda(D) + \partial(D) \geq k - \partial(B)$.
- (ii) There exists $C \prec A$ with $\lambda((C \cap T_z) \ominus \{z\}) + \partial((C \cap T_z) \ominus \{z\}) \geq k - \partial(B) - t$.
- (iii) There exists $E \prec A$ with $\partial(E) \geq k - \partial(B) - t$.

We will show that $\Phi(A) \geq \varepsilon k + \delta\Delta(A)$ in each of these three cases.

Case (i): Suppose $D = A$. Then it follows that $\partial(A) \geq t > 6\Delta(A)$, but this contradicts Lemma 3.8 by which $\Delta(A) \geq \partial(A)/3$, as boundary of a non-empty graph is always non-empty. Thus, $D \prec A$ and we have

$$\begin{aligned}
\Phi(A) &\stackrel{(\dagger)}{\geq} \Phi(D) \geq \varepsilon\lambda(D) + \delta\Delta(D) && \text{(induction hypothesis)} \\
&\geq \varepsilon\lambda(D) + \frac{\delta}{3}\partial(D) && \text{(since } \Delta \geq \partial/3\text{)} \\
&\geq \varepsilon\left(k - \partial(B)\right) + \left(\frac{\delta}{3} - \varepsilon\right)\partial(D) && \text{(since } \lambda(D) + \partial(D) \geq k - \partial(B) \text{ by Case (i))} \\
&\geq \varepsilon\left(k - \partial(B)\right) + \left(\frac{\delta}{3} - \varepsilon\right)t && \text{(since } \delta/3 > \varepsilon \text{ and } \partial(D) \geq t \text{ by Case (i))} \\
&= \varepsilon\left(k - \partial(B)\right) + \left(\frac{\delta}{3} - \varepsilon\right)\left(6\Delta(A) + 3\Delta(B)\right) \\
&\geq \varepsilon\left(k - 6\Delta(A) - 6\Delta(B)\right) + \delta\left(2\Delta(A) + \Delta(B)\right) && \text{(since } -\partial \geq -3\Delta\text{)} \\
&= \varepsilon k + \delta\Delta(A) + \left(\Delta(A) + \Delta(B)\right)(\delta - 6\varepsilon) \\
&\geq \varepsilon k + \delta\Delta(A) && \text{(since } \delta > 6\varepsilon\text{)}.
\end{aligned}$$

Case (ii): Recall that $z \in V_{k-\partial(B)}(T_x)$ was chosen such that $E(B) \cap E(T_z^+) = \emptyset$. It follows that the graph of $(C \cap T_z) \ominus \{z\}$ is a union of connected components of $C \ominus S$. Therefore, $\lambda(C \ominus B) \geq \lambda((C \cap T_z) \ominus \{z\})$ and $\Delta(C \ominus B) \geq \Delta((C \cap T_z) \ominus \{z\})$. Case (ii) now implies

$$\lambda(C \ominus B) + \partial(C \ominus B) \geq k - \partial(B) - t$$

It follows that

$$\begin{aligned} \Phi(C \ominus B) &\geq \varepsilon\lambda(C \ominus B) + \delta\Delta(C \ominus B) && \text{(induction hypothesis)} \\ &\geq \varepsilon(k - \partial(B) - \partial(C \ominus B) - t) + \delta\Delta(C \ominus B) && \text{(by the above inequality)} \\ &\geq \varepsilon(k - 6\Delta(A) - 6\Delta(B) - 3\Delta(C \ominus B)) + \delta\Delta(C \ominus B) && \text{(using } -\partial \geq -3\Delta) \\ &= \varepsilon k - 6\varepsilon(\Delta(A) + \Delta(B)) + (\delta - 3\varepsilon)\Delta(C \ominus B) \\ &\geq \varepsilon k - 6\varepsilon(\Delta(A) + \Delta(B)) && \text{(since } \delta > 3\varepsilon). \end{aligned}$$

By inequality (†) and the induction hypothesis, we have

$$\begin{aligned} \Phi(A) &\stackrel{(\ddagger)}{\geq} \frac{1}{2}\Phi(B) + \frac{1}{2}\Phi(C \ominus B) + \frac{1}{2}\Delta(A) \\ &\geq \frac{1}{2}(\varepsilon k + (\delta - 3\varepsilon)\Delta(B) - \varepsilon) + \frac{1}{2}(\varepsilon k - 6\varepsilon(\Delta(A) + \Delta(B))) + \frac{1}{2}\Delta(A) \quad \text{(by (4) and the above)} \\ &= \varepsilon k + \frac{1-6\varepsilon}{2}\Delta(A) + \frac{\delta-9\varepsilon}{2}\Delta(B) - \frac{\varepsilon}{2} \\ &= \varepsilon k + \delta\Delta(A) + \frac{1}{20}\Delta(B) - \frac{1}{60} \\ &\geq \varepsilon k + \delta\Delta(A). \end{aligned}$$

The final step above uses $\Delta(B) \geq \partial(B)/3 \geq 1/3$ since G_B is a nonempty subgraph of T_∞ .

Case (iii): We have

$$\begin{aligned} \Phi(A) &\stackrel{(\ddagger)}{\geq} \Phi(E) \geq \Delta(E) \geq \frac{1}{3}\partial(E) \geq \frac{1}{3}(k - \partial(B) - t) && \text{(by the inequality of Case (iii))} \\ &\geq \frac{1}{3}(k - 6\Delta(A) - 6\Delta(B)) && \text{(using } -\partial \geq -3\Delta) \\ &= \varepsilon k + \delta\Delta(A) + \left(\frac{1}{3} - \varepsilon\right)k - (2 + \delta)\Delta(A) - 2\Delta(B) \\ &= \varepsilon k + \delta\Delta(A) + \left(\frac{3}{10}k - \frac{12}{5}\Delta(A) - 2\Delta(B)\right). \end{aligned}$$

Recalling that $\Delta(A), \Delta(B) \leq k/18$ by (2), (3), we have

$$\frac{3}{10}k - \frac{12}{5}\Delta(A) - 2\Delta(B) \geq \left(\frac{3}{10} - \frac{12}{5} \cdot \frac{1}{18} - 2 \cdot \frac{1}{18}\right)k = \frac{5}{90}k > 0.$$

This establishes that $\Phi(A) \geq \varepsilon k + \delta\Delta(A)$ in the final case, which concludes the proof of the theorem. \square

3.2 Proofs of Lemmas 3.10, 3.12, 3.13

Proof of Lemma 3.10. Let H be a non-empty finite subgraph of T_∞ , all of whose components are ungrounded. Then $\partial(H) \geq \frac{1}{2}(|E(H)| + 3)$.

As $\partial(\cdot)$ and $|E(\cdot)|$ are additive over disjoint components, it suffices to prove the lemma in the case where H is connected. Let y be the unique highest vertex in H (i.e., belonging to V_k for the maximal k), which we view as the “root” of H .

We now argue by induction on the size of $E(H)$. If $|E(H)| = 1$, then both y and one of its children lie in the boundary of H and hence, we are done. If $|E(H)| = 2$, then either H is the graph induced by y and its two children, or H is a path of length 2 emanating from y . In either case, as H is ungrounded, all three vertices are in the boundary of H and the claim follows.

So assume that $|E(H)| > 2$. We consider two cases:

Case 1: H contains a vertex v such that both its children v_1 and v_2 are leaves of H . Let H' be the subgraph of H induced by $V(H) \setminus \{v_1, v_2\}$. By the induction hypothesis, $\partial(H') \geq \frac{1}{2}(|E(H')| + 3)$. But $\partial(H) = \partial(H') + 1$ because v_1, v_2 are in the boundary of H while v is not. Therefore,

$$\partial(H) = \partial(H') + 1 \geq \frac{1}{2}(|E(H')| + 5) = \frac{1}{2}(|E(H)| + 3).$$

Case 2: There exists a leaf u of H such that $v := \text{parent}(u) \in V(H)$ and $w := \text{parent}(v) \in V(H)$ and $\deg_H(v) = 2$. Define H' to be the subgraph of H induced by $V(H) \setminus \{u, v\}$. By the induction hypothesis, $\partial(H') \geq \frac{1}{2}(|E(H')| + 3)$. But $\partial(H) \geq \partial(H') + 1$ because u, v are in the boundary of H while w may or may not be. In either case, we again have,

$$\partial(H) \geq \partial(H') + 1 \geq \frac{1}{2}(|E(H')| + 5) = \frac{1}{2}(|E(H)| + 3),$$

which completes the proof. \square

Proof of Lemma 3.12. For all finite subgraphs G of T_∞ , $\lambda(G) + \partial(G) \geq \log(|E(G)| + 1)$.

We will prove the following equivalent inequality

$$(5) \quad |E(G)| + 1 \leq 2^{\partial(G) + \lambda(G)}.$$

We claim that it suffices to establish (5) for connected graphs G . To see why, assume (5) holds for two vertex-disjoint graphs G and H . Then we have

$$\begin{aligned} |E(G \cup H)| + 1 &= (|E(G)| + 1) + (|E(H)| + 1) - 1 \leq 2^{\partial(G) + \lambda(G)} + 2^{\partial(H) + \lambda(H)} - 1 \\ &\leq 2^{\max\{\lambda(G), \lambda(H)\}} \cdot (2^{\partial(G)} + 2^{\partial(H)}) - 1 \\ &\leq 2^{\max\{\lambda(G), \lambda(H)\} + \partial(G) + \partial(H)} - 1 \\ &= 2^{\partial(G \cup H) + \lambda(G \cup H)} - 1 < 2^{\partial(G \cup H) + \lambda(G \cup H)} \end{aligned}$$

proving (5) for the graph $G \cup H$.

We now prove (5) assuming G is connected. If G is ungrounded, then we have $\lambda(G) = 0$ and $\partial(G) \geq \frac{1}{2}(|E(G)| + 3)$ by Lemma 3.10, and so

$$|E(G)| + 1 \leq 2\partial(G) - 2 < 2^{\partial(G)} = 2^{\partial(G) + \lambda(G)}.$$

So assume now that G is grounded. Let $y \in V_m$ be the unique vertex in G of maximum height. Let $k := \lambda(G)$ (note that $0 \leq k \leq m$) and fix a choice of $x \in V_k$ such that $T_x \subseteq G$. If $y = x$, then $G = T_x$ and therefore $\lambda(G) = k$ and $\partial(G) = 1$ and $|E(G)| = 2^{k+1} - 2$, so the inequality follows. So assume that $y \neq x$.

As G is connected, it contains the path P from y to x . Consider the case when G contains only one child y' of y . (See Figure 4a for an example.) Then $G \subseteq T_{y'}^+$ and therefore $|E(G)| \leq |E(T_{y'}^+)| = 2^m - 1$. Further, we claim that $\partial(G) \geq m - k$. To see this, note that for every vertex $v \neq x$ on the path P , if $c(v)$ denotes the child of v that is not on the path P , then G does not contain $T_{c(v)}^+$ (because otherwise, $\lambda(G) > k$). As a result, it

must be the case that for every $v \neq x$ on the path P , some vertex in $V(G) \cap V(T_{c(v)}^+)$ lies in the boundary of G and so, $\partial(G) \geq m - k$. Therefore, the desired inequality follows as $|E(G)| + 1 \leq 2^m = 2^{(m-k)+k} \leq 2^{\partial(G)+\lambda(G)}$.

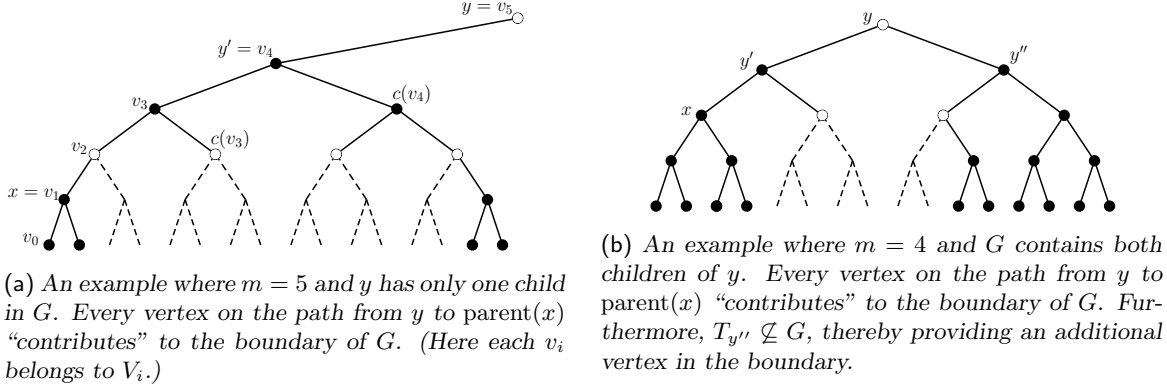


Figure 4: Examples of two cases that arise when the graph G is grounded.

Next, suppose that G contains both children of y . Let y' be the child of y on the path to x , and let y'' be its sibling. (See Figure 4b for an example.) Note that G cannot contain the complete binary tree $T_{y''}$ (since otherwise $\lambda(G) > k$ contradicting our choice of x). Therefore, there is at least one vertex in $T_{y''}$ that lies in the boundary of G . As a result, $\partial(G) \geq m - k + 1$ (the $m - k$ vertices identified by the previous argument, along with the additional boundary element in $V(G) \cap V(T_{y''})$). Also as $G \subseteq T_y$, we have $|E(G)| \leq 2^{m+1} - 2$ and thus, $|E(G)| + 1 < 2^{m+1} = 2^{(m-k+1)+k} \leq 2^{\partial(G)+\lambda(G)}$. \square

Proof of Lemma 3.13. Let $x \in V_k$ and suppose $G \subseteq T_x$ such that $|E(G)| \leq 2^k - 1$ and $\partial(G) \leq k/6$. Then there exists a vertex $z \in V_{k-\partial(G)}(T_x)$ such that $E(G) \cap E(T_z^+) = \emptyset$.

The claim is easy to establish when G is empty or of the form T_y or T_y^+ where $y \in V(T_\infty)$ (that is, in the cases where $\partial(G) \leq 1$). So we shall assume that $\partial(G) \geq 2$.

Note that T_x contains $2^{\partial(G)}$ vertices with height $k - \partial(G)$, that is, $|V_{k-\partial(G)}(T_x)| = 2^{\partial(G)}$ (Observation 3.2). Let

$$Y := \{y \in V_{k-\partial(G)}(T_x) : T_y^+ \subseteq G\},$$

$$Z := \{z \in V_{k-\partial(G)}(T_x) : \emptyset \subsetneq G \cap T_z^+ \subsetneq G\}.$$

It suffices to show that $|Y| + |Z| < 2^{\partial(G)}$.

Since $\bigsqcup_{y \in Y} E(T_y^+) \subseteq E(G)$, it follows that

$$|Y| \cdot (2^{k-\partial(G)+1} - 1) \leq |E(G)| \leq 2^k - 1.$$

We next observe that $\partial(G) \geq |Z| + 1$. This is because for each $z \in Z$, G has at least one boundary element in the set $V(G \cap T_z)$; and we get an additional boundary element by consider any element $w \in V(G)$ of maximum height, noting that w cannot lie in $V(T_z)$ for any $z \in Z$. Therefore,

$$|Y| + |Z| \leq \frac{2^k - 1}{2^{k-\partial(G)+1} - 1} + \partial(G) - 1.$$

Letting $b := \partial(G)$, it suffices to show that

$$\frac{2^k - 1}{2^{k-b+1} - 1} + b - 1 < 2^b, \text{ or equivalently, } 2^b + (b - 1)(2^{k-b+1} - 1) < 2^k + 1.$$

This numerical inequality is simple to verify for all $2 \leq b \leq k/6$, so we are clearly done by the assumption that $2 \leq \partial(G) \leq k/6$.

4 A better potential function

We again fix a graph G and a threshold weighting θ . In this section, we define a potential function $\Phi_\theta(A|S)$ with two parameters: a join-tree A and a set $S \subseteq V(G)$. This improved potential function serves the same purpose of lower-bounding $\tau(G)$. In Section 6 we use $\Phi_\theta(A|S)$ to obtain a better lower bound on $\tau(P_k)$.

Let us begin by recalling the defining inequalities for $\Phi_\theta(A)$:

$$\begin{aligned} (\dagger) \quad \Phi_\theta(A) &\geq \Phi_\theta(D) + \Delta_\theta(C \ominus D) + \Delta_\theta(A \ominus (C \cup D)) && \text{if } A \in \{\langle B, C \rangle, \langle C, B \rangle\} \text{ and } D \preceq B, \\ (\ddagger) \quad \Phi_\theta(A) &\geq \frac{1}{2} \left(\Phi_\theta(D) + \Phi_\theta(E \ominus D) + \Delta_\theta(A) + \Delta_\theta(A \ominus (D \cup E)) \right) && \text{if } D, E \prec A. \end{aligned}$$

A first observation toward improving $\Phi_\theta(A)$ is that we could have included an additional inequality in the definition of $\Phi_\theta(A)$ while maintaining Theorem 2.13 (the lower bound on $\tau(G)$ in terms of $\Phi_\theta(A)$). We call this inequality (\downarrow) , since it is a variant of (\dagger) :

$$(\downarrow) \quad \Phi_\theta(A) \geq \Phi_\theta(D \ominus C) + \Delta_\theta(C) + \Delta_\theta(A \ominus (C \cup D)) \quad \text{if } A \in \{\langle B, C \rangle, \langle C, B \rangle\} \text{ and } D \preceq B.$$

A second observation is that, in the recursive view of $\Phi_\theta(A)$, we “shrink” more than necessary by passing to $\Phi_\theta(D \ominus C)$ in (\downarrow) and $\Phi_\theta(E \ominus D)$ in (\ddagger) . Recall that $D \ominus C$ is a join-tree with graph $G_D \ominus V(C)$ formed by the connected components of G_D that are vertex-disjoint from $V(C)$. Rather than recursing on $D \ominus C$, we can instead simply recurse on D while treating “ $\ominus C$ ” as an extra parameter. These two observations lead to the definition of $\Phi_\theta(A|S)$ below.

Notation 4.1. The following alternative notation for Δ_θ will be convenient in what follows. For a graph $F \subseteq G$ and a set $S \subseteq V(G)$, we write $\Delta_\theta(F|S)$ for $\Delta_\theta(F \ominus S)$. Similarly, for a join-tree A , we write $\Delta_\theta(A|S)$ for $\Delta_\theta(G_A \ominus S)$.

Definition 4.2 (The potential function $\Phi_\theta(A|S)$). Let $\Phi_\theta : \{\text{join-trees for subgraphs of } G\} \times \{\text{subsets of } V(G)\} \rightarrow \mathbb{R}_{\geq 0}$ be the unique pointwise minimum function — written as $\Phi_\theta(A|S)$ rather than $\Phi_\theta(A, S)$ — satisfying the following inequalities for all sets $S \subseteq V(G)$ and join-trees A, B, C, D :

$$\begin{aligned} (\dagger) \quad \Phi_\theta(A|S) &\geq \Phi_\theta(B|S) + \Delta_\theta(C|S \cup B) && \text{if } A \in \{\langle B, C \rangle, \langle C, B \rangle\}, \\ (\downarrow) \quad \Phi_\theta(A|S) &\geq \Delta_\theta(B|S) + \Phi_\theta(C|S \cup B) && \text{if } A \in \{\langle B, C \rangle, \langle C, B \rangle\}, \\ (\ddagger) \quad \Phi_\theta(A|S) &\geq \frac{1}{2} \left(\Phi_\theta(D|S) + \Phi_\theta(A|S \cup D) + \Delta_\theta(A|S) \right) && \text{if } D \prec A. \end{aligned}$$

Alternatively, $\Phi_\theta(A|S)$ has the following recursive characterization:

- If A is an atomic join-tree, then $\Phi_\theta(A|S) := \Delta_\theta(A|S)$.
- If $A = \langle B, C \rangle$, then

$$\Phi_\theta(A|S) := \max \left\{ \begin{array}{l} \Phi_\theta(B|S) + \Delta_\theta(C|S \cup B), \quad \Phi_\theta(C|S) + \Delta_\theta(B|S \cup C), \\ \Delta_\theta(B|S) + \Phi_\theta(C|S \cup B), \quad \Delta_\theta(C|S) + \Phi_\theta(B|S \cup C), \\ \max_{D \prec A} \frac{1}{2} \left(\Phi_\theta(D|S) + \Phi_\theta(A|S \cup D) + \Delta_\theta(A|S) \right) \end{array} \right\}.$$

(To avoid circularity, we take this $\max_{D \prec A}$ over proper sub-join-trees $D \prec A$ such that $V(D) \not\subseteq S$.)

That is, at least one among inequalities (\dagger) , (\downarrow) , (\ddagger) is tight for each join-tree A .

Remark 4.3. We could have defined $\Phi_\theta(A|S)$ in a stronger manner that ensures $\Phi_\theta(A \ominus S) \leq \Phi_\theta(A|S)$ for all A and S (in order to claim that $\Phi_\theta(A|\emptyset)$ improves $\Phi_\theta(A)$) by using the following more general versions of (†), (‡), (‡′):

$$\begin{aligned} (\dagger') \quad & \Phi_\theta(A|S) \geq \Phi_\theta(D|S) + \Delta_\theta(C|S \cup D) + \Delta_\theta(A|S \cup C \cup D) && \text{if } A \in \{\langle B, C \rangle, \langle C, B \rangle\} \text{ and } D \preceq B, \\ (\ddagger') \quad & \Phi_\theta(A|S) \geq \Phi_\theta(D|S \cup C) + \Delta_\theta(C|S) + \Delta_\theta(A|S \cup C \cup D) && \text{if } A \in \{\langle B, C \rangle, \langle C, B \rangle\} \text{ and } D \preceq B, \\ (\ddot{\ddagger}') \quad & \Phi_\theta(A|S) \geq \frac{1}{2} \left(\Phi_\theta(D|S) + \Phi_\theta(E|S \cup D) + \Delta_\theta(A|S) + \Delta_\theta(A|S \cup D \cup E) \right) && \text{if } D, E \prec A. \end{aligned}$$

We chose the simpler Definition 4.2 since it is sufficient for our lower bound on $\tau(P_k)$ in Section 6. Definition 4.2 also leads to a mildly simpler proof of Theorem 5.10 in Section 5, compared with inequalities (†′), (‡′), (‡′′) above.

The following theorem shows that $\Phi_\theta(A|\emptyset)$ serves the same purpose as $\Phi_\theta(A)$ of lower-bounding the invariant $\tau(G)$.

Theorem 4.4. *The invariant $\tau(G)$ in Theorem 1.3 satisfies*

$$\tau(G) \geq \max_{\text{threshold weightings } \theta \text{ for } G} \min_{\text{join-trees } A \text{ with graph } G} \Phi_\theta(A|\emptyset).$$

In the next section, we will finally state the definition of $\tau(G)$ and prove Theorem 4.4. We will then use this theorem to prove a lower bound on $\tau(P_k)$ using $\Phi_\theta(A|S)$ in Section 5. (The argument in Section 5 is purely graph-theoretic and does not require the material in Section 5 if Theorem 4.4 is taken for granted.)

Remark 4.5. The authors first proved the lower bound for $\tau(T_k)$ using $\Phi_\theta(A)$ before considering the improved potential function $\Phi_\theta(A|S)$. It is conceivable that the use of $\Phi_\theta(A|S)$ would simplify or improve the constant in our $k/30$ lower bound. On the other hand, a suitable induction hypothesis would have to account for the additional parameter S , so it is unclear whether a dramatic simplification can be achieved.

5 The pathset framework

In this section we present the *pathset framework*, state the definition of $\tau(G)$, and prove Theorem 4.4 (which bounds $\tau(G)$ in terms of the potential function $\Phi_\theta(A|S)$). All definitions and results in this section are from papers [13, 15, 16] (with minor modifications); a few straightforward lemmas are stated without proof. The reader is referred to those papers for much more context, illustrative examples, and an explanation of how $\tau(G)$ provides a lower bound on the AC^0 and monotone formulas size of $\text{SUB}(G)$.

Throughout this section, we fix a threshold-weighted graph (G, θ) and an arbitrary positive integer n . Let F range over subgraphs of G , let S, T range over subsets of $V(G)$, and let A, B, C, D range over join-trees for subgraphs of G .

Definition 5.1 (Relations, density, join, projection, restriction). Let V, W, T be arbitrary finite sets.

- For a tuple $x \in [n]^V$ and subset $U \subset V$, let $x_U \in [n]^U$ denote the restriction of x to coordinates in U .
- For a relation $\mathcal{A} \subseteq [n]^V$, the *density* of \mathcal{A} is denoted

$$\mu(\mathcal{A}) := |\mathcal{A}|/n^{|V|}.$$

- For relations $\mathcal{A} \subseteq [n]^V$ and $\mathcal{B} \subseteq [n]^W$, the join $\mathcal{A} \bowtie \mathcal{B} \subseteq [n]^{V \cup W}$ is the relation defined by

$$\mathcal{A} \bowtie \mathcal{B} := \{z \in [n]^{V \cup W} : x_V \in \mathcal{A} \text{ and } z_W \in \mathcal{B}\}.$$

- For $\mathcal{A} \subseteq [n]^V$ and $U \subseteq V$, the *projection* $\text{proj}_U(\mathcal{A}) \subseteq [n]^U$ is defined by

$$\text{proj}_U(\mathcal{A}) := \{x \in [n]^U : \exists y \in \mathcal{A} \text{ s.t. } y_U = x\}.$$

- For $\mathcal{A} \subseteq [n]^V$ and $z \in [n]^T$, the *restriction* $\text{rest}_{V \setminus T}(\mathcal{A}|z) \subseteq [n]^{V \setminus T}$ is defined by

$$\text{rest}_{V \setminus T}(\mathcal{A}|z) := \{x \in [n]^{V \setminus T} : \exists y \in \mathcal{A} \text{ s.t. } y_{V \setminus T} = x_{V \setminus T} \text{ and } y_{V \cap T} = z_{V \cap T}\}.$$

The next two lemmas bound the density of relations in terms of projections and restrictions.

Lemma 5.2. *For every relation $\mathcal{A} \subseteq [n]^V$ and $U \subseteq V$,*

$$\mu(\mathcal{A}) \leq \mu(\text{proj}_U(\mathcal{A})) \cdot \max_{z \in [n]^U} \mu(\text{rest}_{V \setminus U}(\mathcal{A}|z)).$$

Lemma 5.3. *For all relations $\mathcal{A} \subseteq [n]^V$ and $\mathcal{B} \subseteq [n]^W$,*

$$\mu(\mathcal{A} \bowtie \mathcal{B}) \leq \mu(\mathcal{A}) \cdot \max_{z \in [n]^W} \mu(\text{rest}_W(\mathcal{B}|z)).$$

For subgraph $F \subseteq G$ and sets $S \subseteq V(G)$, we will be interested in relations $\mathcal{A} \subseteq [n]^{V(G) \setminus S}$ called $G|S$ -*pathsets* that satisfy certain density constraints. These density constraints are related to subgraph counts in the random graph $\mathbf{X}_{\theta, n}$ (see [16]).

Definition 5.4 (Pathsets). Let $F \subseteq G$ and $S \subseteq V(G)$.

- We write $[n]^{F|S}$ for $[n]^{V(F) \setminus S}$.
- We write $\text{Relation}_{F|S}$ for the set of relations $\mathcal{A} \subseteq [n]^{F|S}$. (That is, $\text{Relation}_{F|S}$ is the power set of $[n]^{F|S}$.)
- For $\mathcal{A} \in \text{Relation}_{F|S}$ and $F' \subseteq F$, let $\text{proj}_{F'|S}(\mathcal{A}) := \text{proj}_{V(F') \setminus S}(\mathcal{A})$.
- For $\mathcal{A} \in \text{Relation}_{F|S}$ and $T \supseteq S$ and $z \in [n]^T$, let $\text{rest}_{F|T}(\mathcal{A}|z) := \text{rest}_{V(F) \setminus T}(\mathcal{A}|z)$.
- A relation $\mathcal{A} \in \text{Relation}_{F|S}$ is a $F|S$ -*pathset* if it satisfies

$$\mu(\text{rest}_{F|T}(\mathcal{A}|z)) \leq (1/n)^{\Delta(F|T)}$$

for all $T \supseteq S$ and $z \in [n]^T$. The set of all $F|S$ -pathsets is denoted by $\text{Pathset}_{F|S}$.

The next lemma is immediate from the definition of $\text{Pathset}_{F|S}$.

Lemma 5.5 (Pathsets are closed under restriction). *For all $\mathcal{A} \in \text{Pathset}_{F|S}$ and $T \supseteq S$ and $z \in [n]^T$, we have $\text{rest}_{F|T}(\mathcal{A}|z) \in \text{Pathset}_{F|T}$.*

We next introduce, for each join-tree A and set S , a complexity measure $\chi_{A|S}$ on relations $\mathcal{A} \subseteq [n]^{V(A) \setminus S}$. Very roughly speaking, $\chi_{A|S}$ measures the cost of “constructing” \mathcal{A} via operations \cup and \bowtie , where all intermediate relations are subject to pathset constraints and the pattern of joins is given by A .

Definition 5.6 (Pathset complexity $\chi_{A|S}(\mathcal{A})$).

- For a join-tree A , let $\text{Relation}_{A|S} := \text{Relation}_{G_A|S}$, $\text{Pathset}_{A|S} := \text{Pathset}_{G_A|S}$, etcetera.
- For an atomic join-tree A and relation $\mathcal{A} \in \text{Relation}_{A|S}$, the $A|S$ -*pathset complexity* of \mathcal{A} , denoted $\chi_{A|S}(\mathcal{A})$, is the minimum number m of pathsets $\mathcal{A}_1, \dots, \mathcal{A}_m \in \text{Pathset}_{A|S}$ such that $\mathcal{A} \subseteq \bigcup_{i=1}^m \mathcal{A}_i$.
- For a non-atomic join-tree $A = \langle B, C \rangle$ and relation $\mathcal{A} \in \text{Relation}_{A|S}$, the $A|S$ -*pathset complexity* of \mathcal{A} , denoted $\chi_{A|S}(\mathcal{A})$, is the minimum value of $\sum_{i=1}^m \max\{\chi_{B|S}(\mathcal{B}_i), \chi_{C|S}(\mathcal{C}_i)\}$ over families $\{(\mathcal{A}_i, \mathcal{B}_i, \mathcal{C}_i)\}_{i \in [m]}$ satisfying
 - $(\mathcal{A}_i, \mathcal{B}_i, \mathcal{C}_i) \in \text{Pathset}_{A|S} \times \text{Pathset}_{B|S} \times \text{Pathset}_{C|S}$ for all i ,

- $\mathcal{A}_i \subseteq \mathcal{B}_i \bowtie \mathcal{C}_i$ for all i , and
- $\mathcal{A} \subseteq \bigcup_{i=1}^m \mathcal{A}_i$.

We express this concisely as:

$$\chi_{A|S}(\mathcal{A}) := \min_{\{(\mathcal{A}_i, \mathcal{B}_i, \mathcal{C}_i)\}_i} \sum_i \max\{\chi_{B|S}(\mathcal{B}_i), \chi_{C|S}(\mathcal{C}_i)\}.$$

We refer to any family $\{(\mathcal{A}_i, \mathcal{B}_i, \mathcal{C}_i)\}_i$ achieving this minimum as a *witnessing family* for $\chi_{A|S}(\mathcal{A})$.

The next lemma lists three properties of $\chi_{A|S}$, which are easily deduced from the definition.

Lemma 5.7 (Properties of $\chi_{A|S}$).

- $\chi_{A|S}$ is subadditive: $\chi_{A|S}(\mathcal{A}_1 \cup \mathcal{A}_2) \leq \chi_{A|S}(\mathcal{A}_1) + \chi_{A|S}(\mathcal{A}_2)$,
- $\chi_{A|S}$ is monotone: $\mathcal{A}_1 \subseteq \mathcal{A}_2 \implies \chi_{A|S}(\mathcal{A}_1) \leq \chi_{A|S}(\mathcal{A}_2)$,
- if $A = \langle B, C \rangle$ and $\mathcal{B} \in \text{Pathset}_{B|S}$ and $\mathcal{C} \in \text{Pathset}_{C|S}$, then $\chi_{A|S}(\mathcal{B} \bowtie \mathcal{C}) \leq \max\{\chi_{B|S}(\mathcal{B}), \chi_{C|S}(\mathcal{C})\}$.

The next two lemmas show that pathset complexity is non-increasing under restrictions, as well as under projections to sub-join-trees.

Lemma 5.8 (Projection Lemma). *For all $\mathcal{A} \in \text{Relation}_{A|S}$ and $B \preceq A$, we have*

$$\chi_{B|S}(\text{proj}_{B|S}(\mathcal{A})) \leq \chi_{A|S}(\mathcal{A}).$$

Proof. It suffices to prove the lemma in the case where $A = \langle B, C \rangle$ (since \prec is the transitive closure of “ B is a child of A ”). Fix a witnessing family $\{(\mathcal{A}_i, \mathcal{B}_i, \mathcal{C}_i)\}_i$ for $\chi_{A|S}(\mathcal{A})$. Note that $\text{proj}_{B|S}(\mathcal{A}) \subseteq \bigcup_i \mathcal{B}_i$, since $\mathcal{A} \subseteq \bigcup_i \mathcal{A}_i$ and $\mathcal{A}_i \subseteq \mathcal{B}_i \bowtie \mathcal{C}_i$. It follows that

$$\begin{aligned} \chi_{B|S}(\text{proj}_{B|S}(\mathcal{A})) &\leq \chi_{B|S}(\bigcup_i \mathcal{B}_i) \\ &\leq \sum_i \chi_{B|S}(\mathcal{B}_i) \\ &\leq \sum_i \max\{\chi_{B|S}(\mathcal{B}_i), \chi_{C|S}(\mathcal{C}_i)\} \\ &= \chi_{A|S}(\mathcal{A}). \end{aligned} \quad \square$$

Lemma 5.9 (Restriction Lemma). *For all $\mathcal{A} \in \text{Relation}_{A|S}$ and $T \supseteq S$ and $z \in [n]^T$, we have*

$$\chi_{A|T}(\text{rest}_{A|T}(\mathcal{A}|z)) \leq \chi_{A|S}(\mathcal{A}).$$

Proof. We argue by induction on join-trees A . The lemma is trivial in the base case where A is atomic. So assume $A = \langle B, C \rangle$. Fix a witnessing family $\{(\mathcal{A}_i, \mathcal{B}_i, \mathcal{C}_i)\}_i$ for $\chi_{A|S}(\mathcal{A})$. Observe that the family of restricted triples $\{(\text{rest}_{A|T}(\mathcal{A}_i|z), \text{rest}_{B|T}(\mathcal{B}_i|z), \text{rest}_{C|T}(\mathcal{C}_i|z))\}_i$ satisfies:

- $(\text{rest}_{A|T}(\mathcal{A}_i|z), \text{rest}_{B|T}(\mathcal{B}_i|z), \text{rest}_{C|T}(\mathcal{C}_i|z)) \in \text{Pathset}_{A|T} \times \text{Pathset}_{B|T} \times \text{Pathset}_{C|T}$ for all i ,
- $\text{rest}_{A|T}(\mathcal{A}_i|z) \subseteq \text{rest}_{B|T}(\mathcal{B}_i|z) \bowtie \text{rest}_{C|T}(\mathcal{C}_i|z)$ for all i ,
- $\text{rest}_{A|T}(\mathcal{A}|z) \subseteq \bigcup_i \text{rest}_{A|T}(\mathcal{A}_i|z)$.

Therefore,

$$\begin{aligned} \chi_{A|T}(\text{rest}_{A|T}(\mathcal{A}|z)) &\leq \sum_i \max\{\chi_{B|T}(\text{rest}_{B|T}(\mathcal{B}_i|z)), \chi_{C|T}(\text{rest}_{C|T}(\mathcal{C}_i|z))\} \\ &\leq \sum_i \max\{\chi_{B|S}(\mathcal{B}_i), \chi_{C|S}(\mathcal{C}_i)\} \quad (\text{induction hypothesis}) \\ &= \chi_{A|S}(\mathcal{A}). \end{aligned} \quad \square$$

We now prove the key theorem bounding pathset complexity $\chi_{A|S}(\mathcal{A})$ in terms of the density of \mathcal{A} and the potential function $\Phi(A|S)$.

Theorem 5.10. *For every join-tree A and set S and relation $\mathcal{A} \in \text{Relation}_{A|S}$, we have*

$$\mu(\mathcal{A}) \leq (1/n)^{\Phi(A|S)} \cdot \chi_{A|S}(\mathcal{A}).$$

Proof. We argue by induction on join-trees A . The base case where A is atomic is straightforward. Suppose $\chi_{A|S}(\mathcal{A}) = m$ where $\mathcal{A}_1, \dots, \mathcal{A}_m \in \text{Pathset}_{A|S}$ with $\mathcal{A} \subseteq \bigcup_i \mathcal{A}_i$. We have $\mu(\mathcal{A}) \leq \sum_i \mu(\mathcal{A}_i)$. For each i , we have $\mu(\mathcal{A}_i) \leq (1/n)^{\Phi(A|S)}$ (by definition of $\mathcal{A}_i \in \text{Pathset}_{A|S}$). Therefore, $\mu(\mathcal{A}) \leq (1/n)^{\Phi(A|S)} \cdot m = (1/n)^{\Phi(A|S)} \cdot \chi_{A|S}(\mathcal{A})$ as required.

So we assume $A = \langle B, C \rangle$ is non-atomic. Fix a witnessing family $\{(\mathcal{A}_i, \mathcal{B}_i, \mathcal{C}_i)\}_i$ for $\chi_A(\mathcal{A})$. Since $\mu(\mathcal{A}) \leq \sum_i \mu(\mathcal{A}_i)$, it suffices to show, for each i , that

$$(6) \quad \mu(\mathcal{A}_i) \leq (1/n)^{\Phi(A|S)} \cdot \max\{\chi_{B|S}(\mathcal{B}_i), \chi_{C|S}(\mathcal{C}_i)\}.$$

We establish (6) by considering the three different cases for $\Phi(A|S)$, according to which of the inequalities (†), (‡) (‡) is tight.

Case (†): $\Phi(A|S) = \Phi(B|S) + \Delta(C|S \cup B)$ (or symmetrically $\Phi(A|S) = \Phi(C|S) + \Delta(B|S \cup C)$)

Since $\mathcal{A}_i \subseteq \mathcal{B}_i \bowtie \mathcal{C}_i$ and $\mathcal{B}_i \in \text{Pathset}_{B|S}$, we have $\text{proj}_{B|S}(\mathcal{A}_i) \subseteq \mathcal{B}_i$ and $\mu(\text{rest}_{C|S \cup B}(\mathcal{A}_i|z)) \subseteq \mathcal{C}_i$ for any $z \in [n]^{B|S}$. We now have (6) as follows:

$$\begin{aligned} \mu(\mathcal{A}_i) &\leq \mu(\text{proj}_{B|S}(\mathcal{A}_i)) \max_{z \in [n]^{B|S}} \mu(\text{rest}_{C|S \cup B}(\mathcal{A}_i|z)) && \text{(Lemma 5.3)} \\ &\leq \mu(\mathcal{B}_i) \max_{z \in [n]^{B|S}} \mu(\text{rest}_{C|S \cup B}(\mathcal{C}_i|z)) && \text{(by above observations)} \\ &\leq (1/n)^{\Delta(C|S \cup B)} \cdot \mu(\mathcal{B}_i) && \text{(since } \mathcal{C}_i \in \text{Pathset}_{C|S}) \\ &\leq (1/n)^{\Delta(C|S \cup B)} \cdot (1/n)^{\Phi(B|S)} \cdot \chi_{B|S}(\mathcal{B}_i) && \text{(induction hypothesis)} \\ &\leq (1/n)^{\Phi(A|S)} \cdot \max\{\chi_{B|S}(\mathcal{B}_i), \chi_{C|S}(\mathcal{C}_i)\}. \end{aligned}$$

Case (‡): $\Phi(A|S) = \Delta(B|S) + \Phi(C|S \cup B)$ (or symmetrically $\Phi(A|S) = \Delta(C|S) + \Phi(B|S \cup C)$)

We show (6) as follows:

$$\begin{aligned} \mu(\mathcal{A}_i) &\leq \mu(\mathcal{B}_i) \max_{z \in [n]^{B|S}} \mu(\text{rest}_{C|S \cup B}(\mathcal{C}_i|z)) && \text{(as in the previous case)} \\ &\leq (1/n)^{\Delta(B|S)} \max_{z \in [n]^{B|S}} \mu(\text{rest}_{C|S \cup B}(\mathcal{C}_i|z)) && \text{(since } \mathcal{B}_i \in \text{Pathset}_{B|S}) \\ &\leq (1/n)^{\Delta(B|S)} \max_{z \in [n]^{B|S}} (1/n)^{\Phi(C|S \cup B)} \cdot \chi_{C|S \cup B}(\text{rest}_{C|S \cup B}(\mathcal{C}_i|z)) && \text{(induction hypothesis)} \\ &\leq (1/n)^{\Delta(B|S)} \cdot (1/n)^{\Phi(C|S \cup B)} \cdot \chi_{C|S}(\mathcal{C}_i) && \text{(Restriction Lemma 5.9)} \\ &\leq (1/n)^{\Phi(A|S)} \cdot \max\{\chi_{B|S}(\mathcal{B}_i), \chi_{C|S}(\mathcal{C}_i)\}. \end{aligned}$$

Case (‡): $\Phi(A|S) = \frac{1}{2}(\Phi(D|S) + \Phi(A|S \cup D) + \Delta(A|S))$ for some $D \prec A$

We have

$$\begin{aligned} \mu(\mathcal{A}_i) &\leq \mu(\text{proj}_{D|S}(\mathcal{A}_i)) \max_{z \in [n]^{D|S}} \mu(\text{rest}_{A|S \cup D}(\mathcal{A}_i|z)) && \text{(Lemma 5.2)} \\ &\leq (1/n)^{\Phi(D|S) + \Phi(A|S \cup D)} \cdot \chi_{D|S}(\text{proj}_{D|S}(\mathcal{A}_i)) \max_{z \in [n]^{D|S}} \chi_{A|S \cup D}(\text{rest}_{A|S \cup D}(\mathcal{A}_i|z)) && \text{(induction hyp.)} \\ &\leq (1/n)^{\Phi(D|S) + \Phi(A|S \cup D)} \cdot \chi_{A|S}(\mathcal{A}_i) \cdot \chi_{A|S \cup D}(\text{rest}_{A|S \cup D}(\mathcal{A}_i|z)) && \text{(Proj. Lemma 5.8)} \\ &\leq (1/n)^{\Phi(D|S) + \Phi(A|S \cup D)} \cdot (\chi_{A|S}(\mathcal{A}_i))^2 && \text{(Rest. Lemma 5.9)}. \end{aligned}$$

Since $\mathcal{A}_i \in \text{Pathset}_{A|S}$, we also have

$$\mu(\mathcal{A}_i) \leq (1/n)^{\Delta(A|S)}.$$

Taking the product of the square roots of these two bounds on $\mu(\mathcal{A}_i)$, we conclude

$$\begin{aligned} \mu(\mathcal{A}_i) &\leq (1/n)^{\frac{1}{2}(\Phi(D|S)+\Phi(A|S \cup D)+\Delta(A|S))} \cdot \chi_{A|S}(\mathcal{A}_i) \\ &\leq (1/n)^{\Phi(A|S)} \cdot \chi_{A|S}(\mathcal{B}_i \bowtie \mathcal{C}_i) && \text{(since } \mathcal{A}_i \subseteq \mathcal{B}_i \bowtie \mathcal{C}_i \text{)} \\ &\leq (1/n)^{\Phi(A|S)} \cdot \max\{\chi_{B|S}(\mathcal{B}_i), \chi_{C|S}(\mathcal{C}_i)\} && \text{(Lemma 5.7).} \end{aligned}$$

Having established (6) in all three cases, the proof is complete. \square

Finally, we define the graph invariant $\tau(G)$. Since we no longer fix a particular threshold weighting θ and positive integer n , we include these as subscripts to the pathset complexity function by writing $\chi_{\theta,n,A|S}(\mathcal{A})$ for relations $\mathcal{A} \subseteq [n]^{V(G) \setminus S}$.

Definition 5.11 (The parameter $\tau(G)$ of Theorem 1.3). For a graph G , let $\tau(G) \in \mathbb{R}_{\geq 0}$ be the minimum real number such that

$$\chi_{\theta,n,A|S}(\mathcal{A}) \geq n^{\tau(G)} \cdot \mu(\mathcal{A})$$

for every threshold weighting θ on G , join-tree A with graph G , positive integer n , and relation $\mathcal{A} \subseteq [n]^{V(G)}$.

It is evident from this definition that Theorem 4.4 is an immediate corollary of Theorem 5.10.

6 Lower bound $\tau(P_k) \geq \log_{\sqrt{5}+5}(k)$

Throughout this section, we fix the infinite pattern graph P_∞ and threshold weighting $\theta : E(P_\infty) \rightarrow \{1\}$. Let F range over finite subgraphs of P_∞ , and let S range over finite subsets of $V(P_\infty)$ ($= \mathbb{Z}$). We suppress θ , writing $\Delta(F|S)$ instead of $\Delta_\theta(F|S)$.

For integers $i < j$, let $P_{i,j} \subseteq P_\infty$ be the path from i to j (with edges $\{i, i+1\}, \{i+1, i+2\}, \dots, \{j-1, j\}$). For $k \geq 0$, let $P_k := P_{0,k}$.

Definition 6.1 (Open/half-open/closed components of $F|S$). For integers $i < j$ such that $P_{i,j} \subseteq F$, we say that

- (i, j) is an *open component* of $F|S$ if $V(P_{i,j}) \cap S = \{i, j\}$,
- $[i, j)$ is a *half-open component* of $F|S$ if $V(P_{i,j}) \cap S = \{i\}$ and $\{j, j+1\} \notin E(G)$,
- $(i, j]$ is a *half-open component* of $F|S$ if $V(P_{i,j}) \cap S = \{j\}$ and $\{i-1, i\} \notin E(F)$,
- $[i, j]$ is a *closed component* of $F|S$ if $V(P_{i,j}) \cap S = \emptyset$ and $\{i, i-1\} \notin E(F)$ and $\{j, j+1\} \notin E(F)$.

We shall use the term ‘interval’ and component interchangeably. In each of the four cases above, we define the length of the interval to be $j - i$ and refer to i and j as the left and right ‘end-points’ of that interval, respectively. We shall also refer to the length of an interval I by $|I|$. We treat a join-tree A as its graph G_A when speaking of the *open/half-open/closed components* of $A|S$.

As a matter of notation, let $V((i, j)) = \{i+1, \dots, j-1\}$ and $V([i, j)) = \{i+1, \dots, j\}$ and $V((i, j]) = \{i+1, \dots, j\}$ and $V([i, j]) = \{i, \dots, j\}$.




Figure 5: An example where F is the graph $P_{0,4} \cup P_{6,8} \cup P_{9,14} \cup P_{17,21}$ and $S = \{0, 6, 8, 21\}$. Then $F|S$ consists of four components: $(0, 4]$, $(6, 8)$, $[9, 14]$, and $[17, 21]$. We follow the same format for the rest of the figures in this section: ‘(’, ‘)’ indicate that the corresponding (left or right, respectively) end-point of that component is in S , while ‘|’ indicates that it is not. We shall skip marking the internal vertices of an interval.

Lemma 6.2. $\Delta(F|S)$ equals the number of closed components of $F|S$.

Proof. We have $\Delta(F|S) = \Delta(F \ominus S) = |V(F \ominus S)| - |E(F \ominus S)| = \#\{\text{connected components of } F \ominus S\} = \#\{\text{closed components of } F|S\}$. \square

Recall that for any any graph F and set T , we denote by $F[T]$ the subgraph induced by the vertices of $V(F) \cap T$. We have the following observation that follows immediately from Lemma 6.2:

Observation 6.3. For a graph F , set S and a set T such that for every open/half-open/closed component K of $F|S$, either $T \cap V(K) = \emptyset$ or $V(K) \subseteq T$,

$$\Delta(F|S) = \Delta(F[T]|S) + \Delta(F|S \cup T)$$

where note that $\Delta(F[T]|S)$ equals the number of closed components of $F|S$ that T contains.

The following lemma will allow us to “zoom” into any component of $A|S$ when bounding $\Phi(A|S)$ and gain 1 for each closed component that we discard.

Lemma 6.4. Let A be a join-tree and let S, T be subsets of $V(A)$ such that, for every open/half-open/closed component K of $A|S$, either $T \cap V(K) = \emptyset$ or $V(K) \subseteq T$. Then

$$\Phi(A|S) \geq \Phi(A|S \cup T) + \Delta(A[T]|S).$$

Proof. We argue by induction on join-trees A and by a backward induction on the set S . The lemma is trivial when A is atomic, or when $S = V(A)$. So assume $A = \langle B, C \rangle$. Note that the given condition means that $A[T]|S$ is a union of components of $A|S$ and since $B|S, C|S$ are subgraphs, $B[T]|S$ is a union of components of $B|S$ (and similarly for $C|S$ and $C|S \cup B$). We start with the observation that

$$(7) \quad \Delta(B[T]|S) + \Delta(C[T]|S \cup B) \geq \Delta(A[T]|S).$$

To see this, note that as the graph $A[T]|S$ is a union of the graphs $B[T]|S$ and $C[T]|S$, each closed component of $A[T]|S$ either contains at least one closed component of $B[T]|S$, or it does not. In the latter case, it is then clear that it must also be a component of $C[T]|S \cup B$. See Figure 6 for an illustration.




Figure 6: An example where the set T and the graphs $A|S, B|S, C|S$ are as depicted. Then $B[T]|S = B_3 \cup B_4$ and $C[T]|S \cup B = G_2 \cup G_3 \cup G_4$ with $\Delta(A|S) = \Delta(A[T]|S) = 2$, $\Delta(B[T]|S) = 2$ and $\Delta(C[T]|S \cup B) = 1$.

We will now consider three cases according to whether (\dagger) , (\ddagger) or $(\ddot{\dagger})$ is tight for $\Phi(A|S \cup T)$.

First, assume $\Phi(A|S \cup T) = \Phi(B|S \cup T) + \Delta(C|S \cup T \cup B)$. We have

$$\begin{aligned}
\Phi(A|S) &\stackrel{(\dagger)}{\geq} \Phi(B|S) + \Delta(C|S \cup B) \\
&\geq \Phi(B|S \cup T) + \Delta(B[T]|S) + \Delta(C|S \cup B) && \text{(by induction hypothesis)} \\
&= \Phi(B|S \cup T) + \Delta(B[T]|S) + \Delta(C[T]|S \cup B) + \Delta(C|S \cup B \cup T) && \text{(by observation 6.3 for } C|S \cup B) \\
&\geq \Phi(A|S \cup T) + \Delta(A[T]|S) && \text{(by assumption and by eq. (7))}
\end{aligned}$$

Next, assume $\Phi(A|S \cup T) = \Delta(B|S \cup T) + \Phi(C|S \cup T \cup B)$. We have

$$\begin{aligned}
\Phi(A|S) &\stackrel{(\dagger)}{\geq} \Delta(B|S) + \Phi(C|S \cup B) \\
&= \Delta(B|S \cup T) + \Delta(B[T]|S) + \Phi(C|S \cup B) && \text{(by observation 6.3 for } B|S) \\
&\geq \Delta(B|S \cup T) + \Delta(B[T]|S) + \Phi(C|S \cup B \cup T) + \Delta(C[T]|S \cup B) && \text{(by induction hypothesis)} \\
&\geq \Phi(A|S \cup T) + \Delta(A[T]|S) && \text{(by assumption and by eq. (7))}
\end{aligned}$$

Finally, assume $\Phi(A|S \cup T) = \frac{1}{2}(\Phi(D|S \cup T) + \Phi(A|S \cup T \cup D) + \Delta(A|S \cup T))$ for some $D \prec A$. We have

$$\begin{aligned}
\Phi(A|S) &\stackrel{(\dagger)}{\geq} \frac{1}{2} \left(\Phi(D|S) + \Phi(A|S \cup D) + \Delta(A|S) \right) \\
&\geq \frac{1}{2} \left(\Phi(D|S \cup T) + \Phi(A|S \cup D \cup T) + \Delta(A|S \cup T) + \Delta(D[T]|S) + \Delta(A[T]|S \cup D) + \Delta(A[T]|S) \right)
\end{aligned}$$

where the last inequality follows from the induction hypothesis on $D|S$ and $A|S \cup D$ (this is where we use the backward induction for sets). It thus suffices to check that

$$\Delta(D[T]|S) + \Delta(A[T]|S \cup D) \geq \Delta(A[T]|S).$$

But this is straightforward as each closed component of $A[T]|S$ either contains at least one closed component of $D[T]|S$, or it does not. In the latter case, it is then clear that it must also be a component of $A[T]|S \cup D$. \square

Theorem 6.5. *For every join-tree A and set S , and a component K of $A|S$ of length k ,*

$$\begin{aligned}
\Phi(A|S) &\geq \log_c(\varepsilon \delta k) + \Delta(A|S \cup K) && \text{if } K \text{ is open,} \\
\Phi(A|S) &\geq \log_c(\delta k) + \Delta(A|S \cup K) && \text{if } K \text{ is half-open,} \\
\Phi(A|S) &\geq \log_c(k) + \Delta(A|S \cup K) && \text{if } K \text{ is closed.}
\end{aligned}$$

where $c = \sqrt{5} + 5$, $\delta = (c - 3)/c$ and $\varepsilon = 1/2$.

Proof. We prove this by a structural induction on join-trees as well as a backward induction on the set S . Note that we may assume without loss of generality that $S \subseteq V(A)$. Assume $A|S$ is non-empty as the statement follows immediately, otherwise. If A is atomic, then as $k = 1$ in each case, the statement is trivial. So let us assume that A is non-atomic with $A = \langle B, C \rangle$ and that the theorem statement holds for any proper sub-join-tree $D \prec A$ and set $T \subseteq \mathbb{Z}$ with the given parameter settings of c, δ, ε . Moreover, we shall also assume so for the given join-tree A and for every S' such that $S \subsetneq S' \subseteq V_A$.

Fix a component K of $A|S$. Let T be the union of the vertex sets of all components of $A|S$ excluding K . Then note that $\Delta(A[T]|S) = \Delta(A|S \cup K)$. Therefore, by Lemma 6.4, it suffices to show that

$$\begin{aligned}
\Phi(A|S \cup T) &\geq \log_c(\varepsilon \delta k) && \text{if } K \text{ is open,} \\
\Phi(A|S \cup T) &\geq \log_c(\delta k) && \text{if } K \text{ is half-open,} \\
\Phi(A|S \cup T) &\geq \log_c(k) && \text{if } K \text{ is closed,}
\end{aligned}$$

where we know that the graph of $A|S \cup T$ is simply K . Hence, we may now assume without loss of generality that $A|S$ is connected and has length k .

Henceforth, we shall think of ε, δ , and c as indeterminates, imposing constraints on them as we move along the proof. We will eventually verify that the parameter settings specified in the theorem statement indeed satisfy these constraints. We shall proceed by considering the three cases one by one: that $A|S$ is (i) open, (ii) half-open, or (iii) closed. The strategy in each case is to suitably apply one of the three rules, namely (\dagger) , (\downarrow) , or (\ddagger) in order to obtain a lower bound on $\Phi(A|S)$. We remark that cases (i) and (ii) are similar in terms of the ideas involved to arrive at the lower bound: in particular, neither uses the (\ddagger) rule, which is solely used in case (iii). Roughly speaking, we shall see that case (i) determines the value of ε , while case (ii) that of δ and finally, case (iii) that of c .

Case (i): $A|S$ is open. Suppose $A|S = (0, k)$. Let $A = \langle B, C \rangle$ and suppose $B|S = B_1 \sqcup \cdots \sqcup B_{n(B)}$ (respectively $C|S = C_1 \sqcup \cdots \sqcup C_{n(C)}$) where B_i 's (respectively C_i 's) are the components of $B|S$ (respectively $C|S$), sorted in the increasing order of left end-points. Then with the possible exception of B_1 and $B_{n(B)}$, every other B_i is a closed interval (similarly for C). Also note that for any i and j , $B_i \neq C_j$. We define

$$I(B) = \{B_i : \nexists j \text{ such that } C_j \supseteq B_i\} \text{ and } I(C) = \{C_i : \nexists j \text{ such that } B_j \supseteq C_i\}$$

It follows that $I(B) \cup I(C)$ is a covering of the interval $(0, k)$ and moreover, each of the end-points 0 and k lie in a *unique* (and half-open) interval among the members of $I(B) \cup I(C)$, which we denote by I_0 and I_k respectively. Observe that if we arrange the intervals in $I(B) \cup I(C)$ in the increasing order of left end-points (starting with I_0 obviously), then they alternate membership between $I(B)$ and $I(C)$. Finally, we denote by $\mathcal{C}(A) := I(B) \cup I(C) \setminus \{I_0, I_k\}$, the set of all *closed* intervals in the covering $I(B) \cup I(C)$ and define $t(A) := |\mathcal{C}(A)|$. See Figure 7 for an example that illustrates these notions.




Figure 7: An example where the graphs $A|S$, $B|S$ and $C|S$ are as depicted above. Here $I(B) = \{B_1, B_3, B_4\}$, $I(C) = \{C_2, C_4\}$, $I_0 = B_1$ and $I_k = B_4$. Thus, $\mathcal{C}(A) = \{C_2, B_3, C_4\}$ and $t(A) = 3$. Also, notice that $I(B) \cup I(C) = \{B_1, C_2, B_3, C_4, B_4\}$ and thus its members alternate membership between $I(B)$ and $I(C)$ when arranged in increasing order of left end-points.

We now consider four sub-cases according to whether $t(A) = 0$, $t(A) = 1$, $t(A) \geq 2$ and $t(A)$ is even, or $t(A) \geq 3$ and $t(A)$ is odd. They all involve very similar ideas except for the $t(A) = 1$ sub-case, which is slightly more subtle and requires the application of (\downarrow) unlike the other sub-cases. The reader should bear in mind that the variables a, b, x, y are taken to hold a distinct meaning in each of these sub-cases.

Open sub-case $t(A) = 0$: Either B or C then has a half-open component of size at least $k/2$ and without loss of generality, suppose that is B . See Figure 8 for an example.




Figure 8: An example for $t(A) = 0$ where $I_0 \in I(B)$ and $I_k \in I(C)$, with $|I_0| \geq k/2$. Note that both $B|S$ and $C|S$ each may have additional components, such as the unlabelled intervals depicted in this example.

We have

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \Phi(B|S) + \Delta(C|S \cup B) \geq \Phi(B|S) \geq \log_c(\delta k/2) \geq \log_c(\varepsilon \delta k)$$

as long as $\varepsilon \leq 1/2$. Note that this sub-case is tight for our setting of ε but works for any $c > 1, 0 < \delta \leq 1$.

Open sub-case $t(A) = 1$: Both I_0 and I_k belong to the same collection i.e., either $I(B)$ or $I(C)$. Without loss of generality, assume that it is the former. Then we know that $|I(C)| = 1$, so let $I = [i, j]$ be that closed interval in $I(C)$, for some $0 < i < j < k$. Note that $\max\{i, k - j\} \geq (k - |I|)/2$ so assume without loss of generality that $i \geq (k - |I|)/2$ (in words, that there is more ‘space’ to the left of I than to its right). Next, let J_1, \dots, J_s be the closed intervals in $C|S$ (labelled in the increasing order of left end-points) that appear ‘before’ I i.e., their right end-points are (strictly) less than i (of course it may be the case that $s = 0$, when there is no such interval) and let $y := \frac{1}{k}(\max_{J \in \{I, J_1, \dots, J_s\}} |J|)$. Let J_0 be the component (if it exists) in $C|S$ containing 0, and let $b := |J_0|/k$ (define $b = 0$ if it does not exist).

Let $G_1 \sqcup \dots \sqcup G_{s+1} = (0, i] \setminus (J_0 \cup \dots \cup J_s \cup S) \subseteq B|(S \cup C)$ be the ‘gap’ intervals (labelled in the increasing order of left end-points) as shown in Figure 9. Now with the possible exception of G_1 which may be half-open or open (depending on whether J_0 is empty or not), all other G_i are open. Let gk be the length of the longest gap interval. Then it follows that $(s + 1)g + b + sy \geq \frac{i}{k} \geq \frac{1-y}{2}$.




Figure 9: An example for $t(A) = 1$ where $s = 2$. The gap intervals G_i are shown in dashed lines. Both $B|S$ and $C|S$ each may have additional components, such as the unlabelled intervals depicted in this example.

Now if either of the following occurs, we are immediately done.

- If $a \geq \varepsilon$, then simply use the induction hypothesis on B . We have

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \Phi(B|S) + \Delta(C|S \cup B) \geq \Phi(B|S) \geq \log_c(\delta ak) \geq \log_c(\varepsilon \delta k).$$

- If $b \geq \varepsilon/c^{s+1}$, then as $C|S$ contains at least $s + 1$ other closed components, we have

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \Phi(C|S) + \Delta(B|S \cup C) \geq \Phi(C|S) \geq \log_c(\delta bk) + (s + 1) \geq \log_c(\varepsilon \delta k).$$

- If $y \geq (\varepsilon \delta)/c^s$, then we have

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \Phi(C|S) + \Delta(B|S \cup C) \geq \Phi(C|S) \geq \log_c(yk) + s \geq \log_c(\varepsilon \delta k).$$

So we may assume that $a < \varepsilon$, $b < \varepsilon/c^{s+1}$, and $y < (\varepsilon \delta)/c^s$, which together imply that

$$g > \frac{1}{s + 1} \left(\frac{1}{2} - \frac{\varepsilon}{c^{s+1}} - \left(s + \frac{1}{2} \right) \cdot \frac{\varepsilon \delta}{c^s} \right).$$

We have (this time by (\dagger))

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \Delta(C|S) + \Phi(B|S \cup C) \geq s + 1 + \log_c \left(\frac{\varepsilon \delta k}{s + 1} \left(\frac{1}{2} - \frac{\varepsilon}{c^{s+1}} - \left(s + \frac{1}{2} \right) \cdot \frac{\varepsilon \delta}{c^s} \right) \right)$$

which is at least $\log_c(\varepsilon\delta k)$ as long as for all $s \geq 0$,

$$\frac{c^{s+1} - 2\varepsilon - 2(s + \frac{1}{2}) \cdot \varepsilon\delta c}{2(s+1)} \geq 1$$

which upon plugging in $\varepsilon \leq 1/2$, reduces to showing for all $s \geq 0$ that

$$c^{s+1} \geq 2s + 3 + \left(s + \frac{1}{2}\right) \cdot \delta c$$

which is indeed true for $c > 6$, $0 < \delta \leq 1$. We note that this sub-case is not tight for our parameter settings $c = \sqrt{5} + 5$, $\delta = \frac{c-3}{c}$.

Open sub-case $t(A) \geq 2$ and $t(A)$ is even: If $s = t(A)/2$, then both $B|S$ and $C|S$ contain (exactly) one half-open interval among I_0, I_k and $s \geq 1$ closed intervals each. Similar to the previous case, we define $a := \frac{1}{k}(\max\{|I_0|, |I_k|\})$ and $x := \frac{1}{k}(\max_{I \in \mathcal{C}(A)} |I|)$. It follows that $2a + 2sx \geq 1$. See Figure 10 for an example.




Figure 10: An example where $t(A) = 4$. Here, $I(B) = \{I_0, B_2, B_3\}$ and $I(C) = \{C_2, C_3, I_k\}$. a is therefore defined to be $|I_k|/k$ and x is defined to be $|B_3|/k$.

If $a \geq \varepsilon/c^s$, then as one of $B|S$ or $C|S$ has a half-open interval of length ak along with at least s closed intervals, we have

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \max\{\Phi(B|S), \Phi(C|S)\} \geq \log_c(\delta ak) + s \geq \log_c(\varepsilon\delta k)$$

and we are done. So assume that $a < \varepsilon/c^s$, implying $x > \frac{1}{s}(\frac{1}{2} - \frac{\varepsilon}{c^s})$. Therefore, as one of $B|S$ or $C|S$ has a closed interval of length xk with at least $s - 1$ other closed intervals, we have

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \max\{\Phi(B|S), \Phi(C|S)\} \geq \log_c(xk) + (s - 1) > \log_c\left(\frac{c^{s-1}k}{s} \left(\frac{1}{2} - \frac{\varepsilon}{c^s}\right)\right)$$

and it is enough to check that this last expression is at least $\log_c(\varepsilon\delta k)$, which reduces to showing that for all $s \geq 1$,

$$c^s - 2\varepsilon \geq 2c\varepsilon\delta s$$

which is clearly true when $c > 4$, $0 < \delta \leq (c-1)/c$ and $\varepsilon \leq 1/2$. Therefore, this sub-case is also not tight for our parameter settings $c = \sqrt{5} + 5$, $\delta = \frac{c-3}{c}$.

Open sub-case $t(A) \geq 3$ and $t(A)$ is odd: Both I_0 and I_k must then lie in $I(B)$, without loss of generality. Then if $s = (t(A) + 1)/2$, then $I(C)$ has $s \geq 2$ closed components while $I(B)$ has $s - 1$. Define $a := \frac{1}{k}(\max\{|I_0|, |I_k|\})$, $x := \frac{1}{k}(\max_{I \in I(B) \setminus \{I_0, I_k\}} |I|)$, and $y := \frac{1}{k}(\max_{I \in I(C)} |I|)$. It follows that $2a + (s - 1)x + sy \geq 1$.

Refer back to the example described in Figure 7. For those particular graphs $A|S$, $B|S$, and $C|S$, we would have $a = |B_1|/k$, $x = |B_3|/k$, and $y = |C_2|/k$.

Now if either of the following occurs, we are immediately done.

- If $a \geq \varepsilon/c^{s-1}$, simply use the induction hypothesis on B : one of I_0 or I_k is a half-open component of length at least ak and $B|S$ has $s-1$ closed components. We have

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \Phi(B|S) + \Delta(C|S \cup B) \geq \Phi(B|S) \geq \log_c(\delta ak) + (s-1) \geq \log_c(\varepsilon \delta k).$$

- Similarly if $x \geq (\varepsilon \delta)/c^{s-2}$, then there is closed component of length at least xk in $B|S$ along with $s-2 \geq 0$ other closed components. We have

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \Phi(B|S) + \Delta(C|S \cup B) \geq \Phi(B|S) \geq \log_c(xk) + (s-2) \geq \log_c(\varepsilon \delta k).$$

Hence, we may assume that $a < \varepsilon/c^{s-1}$ and $x < (\varepsilon \delta)/c^{s-2}$, which together imply that

$$y > \frac{1}{s} \left(1 - \frac{2\varepsilon}{c^{s-1}} - \frac{\varepsilon \delta (s-1)}{c^{s-2}} \right).$$

Now $C|S$ has a closed component of length at least yk along with $s-1$ other components. We have

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \Phi(C|S) \geq \log_c(yk) + (s-1) > \log_c \left(\frac{c^{s-1}k}{s} \left(1 - \frac{2\varepsilon}{c^{s-1}} - \frac{\varepsilon \delta (s-1)}{c^{s-2}} \right) \right)$$

and it is enough to show that this last expression is at least $\log_c(\varepsilon \delta k)$, which in turn reduces to checking that for all $s \geq 2$,

$$\frac{c^{s-1}}{\varepsilon \delta s} - \frac{2}{\delta s} - \frac{c(s-1)}{s} \geq 1$$

which is straightforward to verify for any $4 < c, 0 < \delta \leq 1$, and $\varepsilon \leq 1/2$. In particular, this sub-case is also not tight for our parameter settings $c = \sqrt{5} + 5$, $\delta = \frac{c-3}{c}$.

Remark: The above inequality is not true for $s = 1$ irrespective of the choice of c , which is precisely why we have a separate argument for the sub-case $t(A) = 1$.

Case (ii): $A|S$ is half-open. Suppose without loss of generality that $A|S = [0, k)$. Again, let $A = \langle B, C \rangle$ and suppose $B|S = B_1 \sqcup \dots \sqcup B_{n(B)}$ (respectively $C|S = C_1 \sqcup \dots \sqcup C_{n(C)}$) where B_i s (respectively C_i s) are the connected components of $B|S$ (respectively $C|S$). Then with the possible exception of $B_{n(B)}$, every other B_i is a closed interval (similarly for C). Also note that for any i and j , $B_i \neq C_j$. We define

$$I(B) = \{B_i : \nexists j \text{ such that } C_j \supseteq B_i\} \text{ and } I(C) = \{C_i : \nexists j \text{ such that } B_j \supseteq C_i\}$$

It follows that $I(B) \cup I(C)$ is a covering of the interval $[0, k)$ and moreover, each of the end-points 0 and k lies in a *unique* (closed and half-open, respectively) interval among the members of $I(B) \cup I(C)$, which we denote by I_0 and I_k respectively. Again, observe that if we arrange the intervals in $I(B) \cup I(C)$ in the increasing order of left end-points, then they alternate membership between $I(B)$ and $I(C)$. Finally, we denote by $\mathcal{C}(A) := I(B) \cup I(C) \setminus \{I_k\}$, the set of all *closed* intervals in the covering $I(B) \cup I(C)$ and define $t(A) := |\mathcal{C}(A)|$. See Figure 11 for an example that illustrates these notions.




Figure 11: An example where the graphs $A|S$, $B|S$ and $C|S$ are as depicted above. Here $I(B) = \{B_1, B_2, B_3\}$, $I(C) = \{C_1, C_2, C_5\}$, $I_0 = C_1$ and $I_k = C_5$. Thus, $\mathcal{C}(A) = \{C_1, B_1, C_2, B_3\}$ and $t(A) = 4$. Also, notice that $I(B) \cup I(C) = \{C_1, B_1, C_2, B_3, C_5\}$ and thus its members alternate membership between $I(B)$ and $I(C)$ when arranged in increasing order of left end-points.

We again consider four sub-cases according to whether $t(A) = 1$, $t(A) = 2$, $t(A) \geq 3$ and $t(A)$ is odd, or $t(A) \geq 4$ and $t(A)$ is even. The first two sub-cases are slightly more subtle and require the application of (\dagger) unlike the rest. Again, the variables a, b, x, y are taken to hold a distinct meaning in each of these sub-cases.

Half-open sub-case $t(A) = 1$: This means that I_0 is the unique interval in $\mathcal{C}(A)$, also implying that I_k and I_0 overlap. So suppose without loss of generality that $I(B) = \{I_k\}$, $I(C) = \{I_0\}$, and let J_1, \dots, J_s be the closed intervals that appear in $C|S$ ‘after’ I_0 . Further, let J be the half-open interval in $C|S$ that contains k and let $b := |J|/k$ where $b := 0$ if such an interval does not exist. We define $y := \frac{1}{k}(\max_{I \in \{I_0, J_1, \dots, J_s\}} |I|)$. See Figure 12 for an example.




Figure 12: An example for $t(A) = 1$ where $s = 2$. The gap intervals G_i are shown in dashed lines. $B|S$ may have additional components, such as the unlabelled interval depicted in this example.

Now if $y \geq \delta/c^s$, then

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \Phi(C|S) + \Delta(B|S \cup C) \geq \Phi(C|S) \geq \log_c(yk) + s \geq \log_c(\delta k)$$

and we are done. So assume that $y < \delta/c^s$. Further, if $b \geq 1/c^{s+1}$, then

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \Phi(C|S) + \Delta(B|S \cup C) \geq \Phi(C|S) \geq \log_c(bk) + s + 1 \geq \log_c(\delta k)$$

and so we assume that $b < 1/c^{s+1}$. Next, suppose $G_1 \sqcup \dots \sqcup G_{s+1} = I_k \setminus (I_0 \cup J_1 \cup \dots \cup J_s \cup S) \subseteq B|(C \cup S)$ are the ‘gap’ intervals as shown in Figure 12. Let gk be the length of the longest gap interval. Then it follows that $(s+1)g + sy + b \geq 1 - y$. We have (this time by (\dagger))

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \Delta(C|S) + \Phi(B|S \cup C) \geq s + 1 + \log_c(\varepsilon \delta g k) \geq s + 1 + \log_c \left(\frac{\varepsilon \delta k}{s+1} \left(1 - \frac{\delta(s+1)}{c^s} - \frac{1}{c^{s+1}} \right) \right).$$

The task of showing that this expression is at least $\log_c(\delta k)$ reduces to showing that for all $s \geq 0$,

$$\varepsilon \left(\frac{c^{s+1} - 1}{s+1} - \delta c \right) \geq 1$$

which is clearly true when $c > 4$, $\delta \leq (c-3)/c$, and $\varepsilon = 1/2$. We note that this sub-case is indeed tight for our parameter settings $\delta = \frac{c-3}{c}$ and $\varepsilon = 1/2$, however it works for any $c > 4$.

Half-Open sub-case $t(A) = 2$: Let $I(C) = \{I_0, I_k\}$, $I(B) = \{J\}$ and $a := |I_k|/k$. If $a \geq 1/c$, then we have (\dagger)

$$\Phi(A|S) \geq \Phi(C|S) \geq \log_c(\delta a k) + 1 \geq \log_c(\delta k)$$

and we are done. So we may assume that $a < 1/c$. Let J_1, \dots, J_s (s may be zero) be the closed intervals in $C|S$ that appear ‘before’ J as shown in Figure 13 and suppose $y := \frac{1}{k}(\max_{I \in \{J, J_1, \dots, J_s\}} |I|)$. If $y \geq \delta/c^s$, then

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \Phi(B|S) \geq \log_c(yk) + s \geq \log_c(\delta k)$$

and we are done. So assume that $y < \delta/c^s$. Now let i be the left end-point of J and suppose $G_1 \sqcup \dots \sqcup G_{s+1} = (0, i] \mid (J \cup J_1 \cup \dots \cup J_s \cup S) \subseteq C \mid (S \cup B)$ are the ‘gap’ intervals as shown in the Figure 13. Let gk be the length of the longest gap interval. Then it follows that $(s+1)g + sy \geq \frac{i}{k} \geq 1 - a - y \geq 1 - \frac{\delta}{c^s} - \frac{1}{c}$.




Figure 13: An example for $t(A) = 2$ where $s = 2$. The gap intervals G_i are shown in dashed lines. Both $B|S$ and $C|S$ each may have additional components, such as the unlabelled intervals depicted in this example.

We have (by (\dagger))

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \Delta(B|S) + \Phi(C|S \cup B) \geq s + 1 + \log_c(\varepsilon \delta g k) \geq s + 1 + \log_c \left(\frac{\varepsilon \delta k}{s+1} \left(1 - \frac{\delta(s+1)}{c^s} - \frac{1}{c} \right) \right)$$

The task of showing that this expression is at least $\log_c(\delta k)$ reduces to showing that for all $s \geq 0$,

$$\varepsilon \left(\frac{c^{s+1} - c^s}{s+1} - \delta c \right) \geq 1$$

which is clearly true when $c > 4$, $\delta \leq (c-3)/c$, and $\varepsilon = 1/2$. This sub-case is also tight for our parameter settings $\delta = \frac{c-3}{c}$ and $\varepsilon = 1/2$, however it works for any $c > 4$.

Half-open sub-case $t(A) \geq 3$ and $t(A)$ is odd: Suppose I_k is in $I(B)$, without loss of generality. Then if $s = (t(A) + 1)/2$, then $I(C)$ has $s \geq 2$ closed components while $I(B)$ has $s - 1$. Let $a := |I_k|/k$, $x := \frac{1}{k}(\max_{I \in I(B) \setminus \{I_k\}} |I|)$, and $y := \frac{1}{k}(\max_{I \in I(C)} |I|)$ (see Figure 14 for an example). It follows that $a + (s-1)x + sy \geq 1$.




Figure 14: An example for $t(A) = 5$. Here, we would have $x = |B_3|/k$ and $y = |C_2|/k$.

Now if either of the following occurs, we are immediately done.

- If $a \geq 1/c^{s-1}$, then

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \Phi(B|S) \geq \log_c(\delta a k) + (s-1) \geq \log_c(\delta k).$$

- Similarly if $x \geq \delta/c^{s-2}$, then there is closed component of length at least xk in $B|S$ along with $s-2 \geq 0$ other closed components. We have

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \Phi(B|S) \geq \log_c(xk) + (s-2) \geq \log_c(\delta k).$$

Hence, we may assume that $a < 1/c^{s-1}$ and $x < \delta/c^{s-2}$, which together imply that

$$y > \frac{1}{s} \left(1 - \frac{1}{c^{s-1}} - \frac{\delta(s-1)}{c^{s-2}} \right).$$

Now $C|S$ has a closed component of length at least yk along with $s-1$ other components. Thus,

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \Phi(C|S) \geq \log_c(yk) + (s-1) > \log_c \left(\frac{c^{s-1}k}{s} \left(1 - \frac{1}{c^{s-1}} - \frac{\delta(s-1)}{c^{s-2}} \right) \right)$$

and it is enough to show that this last expression is at least $\log_c(\delta k)$, which in turn reduces to checking that for all $s \geq 2$,

$$\frac{c^{s-1}}{\delta s} - \frac{1}{\delta s} - \frac{c(s-1)}{s} \geq 1$$

which is straightforward to verify for $4 < c$ and $0 < \delta \leq (c-1)/(c+2)$. In particular, this is true for our parameter settings $c = \sqrt{5} + 5$, $\delta = \frac{c-3}{c}$ and it follows that this sub-case is not tight.

Half-open sub-case $t(A) \geq 4$ and $t(A)$ is even: Suppose $a := |I_k|/k$ without loss of generality, $I_k \in I(C)$. Then if $s = t(A)/2$, then both $B|S$ and $C|S$ contain at least $s \geq 2$ closed intervals each. We define $x := \frac{1}{k}(\max_{I \in \mathcal{C}(A)} |I|)$. It follows that $a + 2sx \geq 1$.

Refer back to the example depicted in Figure 11. For that particular example, we would then have $a = |C_5|/k$ and $x = |B_3|/k$, as B_3 is the longest interval in $\mathcal{C}(A)$.

If $a \geq 1/c^s$, then

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \Phi(C|S) \geq \log_c(\delta ak) + s \geq \log_c(\delta k)$$

and we are done. So assume that $a < 1/c^s$, implying $x > \frac{1}{2s}(1 - \frac{1}{c^s})$. Therefore,

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \max\{\Phi(B|S), \Phi(C|S)\} \geq \log_c(xk) + (s-1) > \log_c \left(\frac{c^{s-1}k}{2s} \left(1 - \frac{1}{c^s} \right) \right)$$

and it is enough to check that this last expression is at least $\log_c(\delta k)$, which reduces to showing that for all $s \geq 2$,

$$c^s - 1 \geq 2c\delta s$$

which is clearly true when $c > 4$ and $0 < \delta \leq 1$. This sub-case is also not tight for our parameter settings $c = \sqrt{5} + 5$, $\delta = \frac{c-3}{c}$.

Case (iii): $A|S$ is closed. Finally, suppose $A|S = [0, k]$. First, consider the sub-case that there exists $D \prec A$ such that $D = [0, j]$ or $D = [j, k]$ with

$$\frac{1}{2}(1 - \sqrt{x})k \leq j \leq \frac{1}{2}(1 + \sqrt{x})k$$

for $x = \frac{1}{(\sqrt{5}+2)^2}$, as shown in Figure 15. Then $j(k-j) \geq \frac{1-x}{4}k^2$. We have (this time by (\ddagger))

$$\begin{aligned} \Phi(A|S) &\stackrel{(\ddagger)}{\geq} \frac{1}{2} \left(\Phi(D|S) + \Phi(A|S \cup D) + \Delta(A|S) \right) \\ &\geq \frac{1}{2} \left(\log_c(k-j) + \log_c(\delta j) + 1 \right) \quad (\text{induction hypothesis on } A|S \cup D) \\ &= \frac{1}{2} \log_c(c\delta j(k-j)) \geq \frac{1}{2} \log_c \left(\frac{c\delta(1-x)}{4} k^2 \right) = \log_c(k) + \frac{1}{2} \log_c \left(\frac{c\delta(1-x)}{4} \right). \end{aligned}$$

Note that for our choice of x , and the given parameter settings $c = \sqrt{5} + 5$ and $\delta = \frac{c-3}{c}$, $\frac{c\delta(1-x)}{4} = 1$, thereby establishing the claim if such a $D \prec A$ exists. We also note that our parameter settings are indeed tight in this case.




Figure 15: An example where $A|S$ is closed and there exists a connected $D \prec A$ containing one of the end-points of $A|S$ and having ‘intermediate’ length. We apply the (\ddagger) rule in this situation.

Therefore, now assume that no such D exists.

Furthermore, we may assume that if there exists $D \prec A$ such that D contains a component of length at least $\frac{1}{c}k$, then D is connected (call this assumption $(*)$). Because otherwise, we have the following:

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \Phi(D|S) \geq 1 + \log_c(k/c) = \log_c(k)$$

Note that $\frac{1}{c} < \frac{1}{2}(1 - \sqrt{x})$.

It follows from the application of (\ddagger) above and $(*)$ that the only case left now is when there exists a (connected) $D \prec A$ and a child E of D such that $D = [0, j]$ and the component $J_0 = [0, i]$ of E containing 0 satisfy

$$i \leq \frac{1}{2}(1 - \sqrt{x})k \quad \text{and} \quad \frac{1}{2}(1 + \sqrt{x})k \leq j.$$

Let E have s other components J_1, \dots, J_s apart from J_0 . We now consider sub-cases according to whether s is zero or not.

Closed sub-case I: $s = 0$. It follows that E is connected. Then note that $D|S \cup E = G = (i, j)$ is as shown in Figure 16, and we have

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \Phi(D|S) \stackrel{(\ddagger)}{\geq} \Delta(E|S) + \Phi(D|S \cup E) \geq 1 + \log_c(\delta(j-i)) \geq 1 + \log_c(\delta\sqrt{x}k) = \log_c(k) + \log_c(c\delta\sqrt{x}).$$

So it only remains to check that $x \geq 1/(c\delta)^2$, which is indeed true for our parameter settings. We note that this sub-case is tight for our parameter settings.




Figure 16: An example where $A|S$ is closed and $E \prec D$ is connected.

Closed sub-case II: $s \geq 1$. We still apply (\ddagger) as in the previous sub-case, but the analysis is different. As E has $s+1$ components in all, we may assume that each J_i has length at most $\frac{1}{c^s}k$ as otherwise, we immediately obtain $\Phi(A|S) \geq \Phi(D|S) \geq \log_c(k)$ from the (\dagger) rule. As a consequence, if $G_1 \sqcup \dots \sqcup G_{s+1} = D|S \cup E$ are

the ‘gap’ intervals as shown in Figure 17, some G_i must be an open (or half-open) interval of length at least $\frac{1}{s+1}(j - \frac{s+1}{c^s})$.




Figure 17: An example where $A|S$ is closed and $s = 2$.

We have

$$\Phi(A|S) \stackrel{(\dagger)}{\geq} \Phi(D|S) \stackrel{(\ddagger)}{\geq} \Delta(E|S) + \Phi(D|S \cup E) \geq s + 1 + \log_c \left(\frac{\varepsilon \delta}{s+1} \left(j - \frac{s+1}{c^s} k \right) \right)$$

which is at least $\log_c(k)$ as long as for all $s \geq 1$,

$$\frac{\varepsilon \delta c^{s+1}}{s+1} \left(\frac{j}{k} - \frac{s+1}{c^s} \right) - 1 \geq 0$$

By assumption, $\frac{j}{k} \geq \frac{1}{2}(1 + \sqrt{x}) = \frac{\sqrt{5}-1}{2}$. Thus, we plug in $\varepsilon = \frac{1}{2}$ and $c\delta = \sqrt{5} + 2$ and see that it is enough to verify that for all $s \geq 1$,

$$\frac{\sqrt{5} + 2}{2(s+1)} \left(\frac{c^s(\sqrt{5}-1)}{2} - (s+1) \right) - 1 \geq 0$$

which is easily seen to be true for $c = \sqrt{5} + 5$. □

Lemma 2.17 and hence Theorem 1.11 ($\tau(P_k) \geq \log_{\sqrt{5}+5}(k) - 1$) now follow as a corollary to Theorem 6.5 and Lemma 2.14 (in the latter, we simply plug in $G^* = P_\infty$, $G = P_k$, θ and θ^* to be the constant $1 + \frac{1}{k}$ and 1 threshold weightings respectively).

7 Randomized AC^0 formulas computing the product of k permutations

In this section we define a broad class of randomized AC^0 formulas for computing the product of k permutations. The size of these formulas corresponds to a complexity measure related to pathset complexity, but much simpler and easier to analyze.

Definition 7.1. For integers $k \geq 1$, let $\mathcal{P}(k)$ be the set of sequences $\vec{a} = (a_0, \dots, a_k) \in [0, 1]^{\{0, \dots, k\}}$ such that $\|\vec{a}\| := a_0 + \dots + a_k \geq 1$. We denote by $\vec{a} \leq \vec{b}$ that each $a_i \leq b_i$.

We define a complexity measure $\chi : \bigcup_{k \geq 1} \mathcal{P}(k) \rightarrow \mathbb{R}_{\geq 0}$ by the following induction:

- In the base case $k = 1$, let $\chi(a_0, a_1) := 0$.
- For $k \geq 2$, let

$$\chi(a_0, \dots, a_k) := \min_{\substack{0 < i \leq j < k, \\ \vec{b} \in \mathcal{P}(k): \\ \vec{a} \leq \vec{b}, \\ (b_0, \dots, b_j) \in \mathcal{P}(j), \\ (b_i, \dots, b_k) \in \mathcal{P}(k-i)}} \|\vec{b} - \vec{a}\| + \max\{\chi(b_0, \dots, b_j), \chi(b_i, \dots, b_k)\}.$$

Note the following properties of χ :

1. For all $\vec{a} \in \mathcal{P}(k)$ and $1 < i \leq j < k$, if $(a_0, \dots, a_j) \in \mathcal{P}(j)$ and $(a_i, \dots, a_k) \in \mathcal{P}(k-i)$, then $\chi(\vec{a}) \leq \max\{\chi(a_0, \dots, a_j), \chi(a_i, \dots, a_k)\}$.
2. For all $\vec{a}, \vec{b} \in \mathcal{P}(k)$, if $\vec{a} \leq \vec{b}$, then $\chi(\vec{a}) \leq \|\vec{b} - \vec{a}\| + \chi(\vec{b})$.
3. For all $k \geq 1$, we have $\chi(\underbrace{\frac{1}{2}, \dots, \frac{1}{2}}_{k+1 \text{ times}}) = 0$.

The complexity measure $\chi(\vec{a})$ is a simplified version of pathset complexity $\chi_A(\mathcal{A})$. In fact, $\chi(\vec{a})$ provides an upper bound on pathset complexity!

Remark 7.2. Consider the infinite pattern graph P_∞ under the constant 1 threshold weighting. For join-trees A over P_∞ , we will write Pathset_A for $\text{Pathset}_{A|\emptyset}$ and $\chi_A(\cdot)$ for $\chi_{A|\emptyset}(\cdot)$.

Each $\vec{a} \in \mathcal{P}(k)$ corresponds to a P_k -pathset

$$\mathcal{A}_{\vec{a}} := \{x \in [n]^{\{0, \dots, k\}} : x_h \in S_h \text{ for all } h \in \{0, \dots, k\}\}$$

where S_0, \dots, S_k are arbitrary subsets of $[n]$ of size $|S_h| \leq n^{1-a_h}$. Then there exists a join-tree A with graph P_k such that

$$\chi_A(\mathcal{A}_{\vec{a}}) \leq n^{\chi(\vec{a})+o(1)}.$$

This join-tree arises from the optimal $1 < i \leq j < k$ and \vec{b} in the definition of $\chi(\vec{a})$: namely, $A = \langle B, C \rangle$ where B is the join-tree for $P_{0,j}$ associated with (b_0, \dots, b_j) and C is the join-tree for $P_{i,k}$ associated with (b_i, \dots, b_k) . (Note that A has the property that G_D is a path for each $D \preceq A$; so, not all join-trees with graph P_k arise in this way.)

The above bound on $\chi_A(\mathcal{A})$ (now dropping the subscript as \vec{a} is fixed) is justified as follows. Letting $m := n^{\|\vec{b}-\vec{a}\|+o(1)}$, there exist sets $T_{\ell,h} \subseteq S_h$ of size $|T_{\ell,h}| \leq n^{1-b_h}$, indexed over $\ell \in [m]$ and $h \in \{0, \dots, k\}$, such that $\bigcup_{\ell \in [m]} (T_{\ell,0} \times \dots \times T_{\ell,k}) = S_0 \times \dots \times S_k$. We then have $\mathcal{A} = \bigcup_{\ell \in [m]} \mathcal{B}_\ell \bowtie \mathcal{C}_\ell$ where

$$\begin{aligned} \mathcal{B}_\ell &:= \{y \in [n]^{\{0, \dots, j\}} : y_h \in T_{\ell,h} \text{ for all } 0 \leq h \leq j\}, \\ \mathcal{C}_\ell &:= \{z \in [n]^{\{i, \dots, k\}} : z_h \in T_{\ell,h} \text{ for all } i \leq h \leq k\}. \end{aligned}$$

Arguing by induction on proper subsequences (b_0, \dots, b_j) and (b_i, \dots, b_k) (note that the base case $k=1$ is trivial as $\chi(\vec{a})=0$ and $\chi_A(\mathcal{A})=1$ as \mathcal{A} itself is a pathset), it follows that

$$\begin{aligned} \chi_A(\mathcal{A}) &\leq \sum_{\ell \in [m]} \max\{\chi_B(\mathcal{B}_\ell), \chi_C(\mathcal{C}_\ell)\} \\ &\leq \sum_{\ell \in [m]} \max\{n^{\chi(b_0, \dots, b_j)+o(1)}, n^{\chi(b_i, \dots, b_k)+o(1)}\} \leq m \cdot n^{\chi(\vec{a})-\|\vec{b}-\vec{a}\|+o(1)} = n^{\chi(\vec{a})+o(1)}. \end{aligned}$$

As a consequence of these observations, we see that $\chi(\vec{a})$ is lower-bounded by $\log_n(\chi_A(\mathcal{A}))$. Since $\chi_A(\mathcal{A}) \leq n^{\Phi(A)} \cdot \mu(\mathcal{A}) \leq n^{\Phi(A)} \cdot n^{-\|\vec{a}\|+o(1)}$ (by Theorem 5.10), it follows that

$$\chi(\vec{a}) \geq \max_{\substack{\text{join-trees } A \text{ with graph } P_k \text{ s.t.} \\ G_D \text{ is connected for all } D \preceq A}} \Phi(A) - \|\vec{a}\|.$$

In particular, our lower bound of Section 6 implies that $\chi(\vec{a}) \geq \log_{\sqrt{5+5}}(k) - \|\vec{a}\| - 1$ for all $\vec{a} \in \mathcal{P}(k)$.

Finally, note that by covering the complete relation $[n]^{\{0, \dots, k\}}$ by $n^{\|\vec{a}\|+o(1)}$ shifted copies of rectangles $S_0 \times \dots \times S_k$, we get an upper bound

$$\chi_A([n]^{\{0, \dots, k\}}) \leq n^{\chi(\vec{a})+\|\vec{a}\|+o(1)}.$$

By a similar construction, we will show that $n^{\chi(\vec{a})+\|\vec{a}\|+o(1)}$ is an upper bound on the randomized AC^0 formula size of computing the product of k permutations.

Definition 7.3. Let $\vec{\pi} = (\pi_1, \dots, \pi_k)$ be a sequence of permutations $[n] \xrightarrow{\cong} [n]$. For a sequence $\vec{x} = (x_0, \dots, x_k) \in [n]^{\{0, \dots, k\}}$, we say that \vec{x} is a $\vec{\pi}$ -path if $\pi_h(x_{h-1}) = x_h$ for all $h \in \{1, \dots, k\}$.

If \vec{x} is a $\vec{\pi}$ -path and $\vec{S} = (S_0, \dots, S_k)$ is a sequence of sets $S_0, \dots, S_k \subseteq [n]$, we will say that \vec{S} isolates \vec{x} if $\vec{x} \in S_0 \times \dots \times S_k$ and \vec{x} is the *only* $\vec{\pi}$ -path in $S_0 \times \dots \times S_k$.

Definition 7.4. For a set U and $p \in [0, 1]$, notation $\mathcal{S} \subseteq_p U$ denotes that \mathcal{S} is a random subset of U that contains each element independently with probability p .

Given $\vec{a} = (a_0, \dots, a_k) \in \mathcal{P}(k)$, we will denote by $\vec{\mathcal{S}} = (\mathcal{S}_0, \dots, \mathcal{S}_k)$ the sequence of independent random sets $\mathcal{S}_h \subseteq_{n^{-a_h}} [n]$.

We now state the key lemma for our construction.

Lemma 7.5. For every $\vec{a} \in \mathcal{P}(k)$ and sequence $\vec{S} = (S_0, \dots, S_k)$ of sets $S_h \subseteq [n]$, there exist randomized AC^0 formulas

$$\mathbf{f}_{\vec{a}, \vec{S}} \text{ and } \vec{\mathbf{g}}_{\vec{a}, \vec{S}} = \{\mathbf{g}_{\vec{a}, \vec{S}}^{(h,t)}\}_{r \in \{0, \dots, k\}, t \in \{1, \dots, \lceil \log(n+1) \rceil\}}$$

each of depth $O(k)$ and size $n^{\chi(\vec{a})+o(1)}$ and taking as input a sequence $\vec{\pi} = (\pi_1, \dots, \pi_k)$ of permutations $[n] \xrightarrow{\cong} [n]$, such that on every input $\vec{\pi}$ then with probability $1 - n^{-\omega(1)}$ (with respect to both \vec{S} and the randomness of $\mathbf{f}_{\vec{a}, \vec{S}}$ and $\vec{\mathbf{g}}_{\vec{a}, \vec{S}}$):

1. $\mathbf{f}_{\vec{a}, \vec{S}}(\vec{\pi})$ outputs 1 if, and only if, \vec{S} isolates some $\vec{\pi}$ -path.
2. If \vec{S} isolates a (necessarily unique) $\vec{\pi}$ -path $\vec{x} = (x_0, \dots, x_k)$, then formulas $\vec{\mathbf{g}}_{\vec{a}, \vec{S}}(\vec{\pi})$ output the binary representation of integers $x_0, \dots, x_k \in [n]$.

Proof. The construction mimics the pathset complexity upper bound in Remark 7.2. In the base case $k = 1$, we have sets $S_0, S_1 \subseteq [n]$ and need to determine if a permutation $\pi : [n] \xrightarrow{\cong} [n]$ satisfies $\pi(x) = y$ for a unique pair $(x, y) \in S_0 \times S_1$. This is accomplished by the following AC^0 formula (writing $1_{\pi(x)=y}$ for the input variable that is 1 if and only if $\pi(x) = y$):

$$\begin{aligned} \mathbf{f}_{\vec{a}, \vec{S}}(\pi) := & \bigvee_{(x,y) \in S_0 \times S_1} 1_{\pi(x)=y} \\ & \wedge \bigwedge_{t \in \{1, \dots, \lceil \log(n+1) \rceil\}} \neg \left(\left(\bigvee_{(x,y) \in S_0 \times S_1 : \text{the } t^{\text{th}} \text{ bit of } x \text{ is } 0} 1_{\pi(x)=y} \right) \wedge \left(\bigvee_{(x,y) \in S_0 \times S_1 : \text{the } t^{\text{th}} \text{ bit of } x \text{ is } 1} 1_{\pi(x)=y} \right) \right) \\ & \wedge \bigwedge_{t \in \{1, \dots, \lceil \log(n+1) \rceil\}} \neg \left(\left(\bigvee_{(x,y) \in S_0 \times S_1 : \text{the } t^{\text{th}} \text{ bit of } y \text{ is } 0} 1_{\pi(x)=y} \right) \wedge \left(\bigvee_{(x,y) \in S_0 \times S_1 : \text{the } t^{\text{th}} \text{ bit of } y \text{ is } 1} 1_{\pi(x)=y} \right) \right). \end{aligned}$$

This formula has depth $O(1)$ and size $O(\log n)$ (as measured by number of gates). Since $\chi(\vec{a}) = 0$, this size bound is $n^{\chi(\vec{a})+o(1)}$ as required. Formulas $\vec{\mathbf{g}}_{\vec{a}, \vec{S}}$ giving the binary representation of x and y (whenever (x, y) uniquely exists) have just a single OR gate:

$$\begin{aligned} \mathbf{g}_{\vec{a}, \vec{S}}^{(0,t)}(\pi) := & \bigvee_{(x,y) \in S_0 \times S_1 : \text{the } t^{\text{th}} \text{ bit of } x \text{ is } 1} 1_{\pi(x)=y}, \\ \mathbf{g}_{\vec{a}, \vec{S}}^{(1,t)}(\pi) := & \bigvee_{(x,y) \in S_0 \times S_1 : \text{the } t^{\text{th}} \text{ bit of } y \text{ is } 1} 1_{\pi(x)=y}. \end{aligned}$$

Onto the induction step where $k \geq 2$. Fix $0 < i \leq j < k$ and $\vec{b} \in \mathcal{P}(k)$ with $\vec{a} \leq \vec{b}$ and

$$\begin{aligned}\vec{b}' &:= (b_0, \dots, b_j) \in \mathcal{P}(j), \\ \vec{b}'' &:= (b_i, \dots, b_k) \in \mathcal{P}(k-i), \\ \chi(\vec{a}) &= \|\vec{b} - \vec{a}\| + \max\{\chi(\vec{b}'), \chi(\vec{b}'')\}.\end{aligned}$$

Letting $m := n^{\|\vec{b} - \vec{a}\| + o(1)}$, we sample independent random sequences of sets $\vec{T}_1, \dots, \vec{T}_m$ where for each $\ell \in [m]$, we have $\vec{T}_\ell = (\mathbf{T}_{\ell,0}, \dots, \mathbf{T}_{\ell,k})$ with $\mathbf{T}_{\ell,h} \subseteq_{n^{-b_h + a_h}} S_h$.⁷

Writing \vec{T}'_ℓ for $(\mathbf{T}_{\ell,0}, \dots, \mathbf{T}_{\ell,j})$ and \vec{T}''_ℓ for $(\mathbf{T}_{\ell,i}, \dots, \mathbf{T}_{\ell,k})$ and $\vec{\pi}'$ for (π_1, \dots, π_j) and $\vec{\pi}''$ for $(\pi_{i+1}, \dots, \pi_k)$, we now introduce auxiliary randomized formulas $\mathbf{join}_1, \dots, \mathbf{join}_m$ defined by

$$\mathbf{join}_\ell(\vec{\pi}) := \mathbf{f}_{\vec{b}', \vec{T}'_\ell}(\vec{\pi}') \wedge \mathbf{f}_{\vec{b}'', \vec{T}''_\ell}(\vec{\pi}'') \wedge \bigwedge_{\substack{h \in \{i, \dots, j\} \\ t \in \{1, \dots, \lceil \log(n+1) \rceil\}}} \left(\mathbf{g}_{\vec{b}', \vec{T}'_\ell}^{(h,t)}(\vec{\pi}') \leftrightarrow \mathbf{g}_{\vec{b}'', \vec{T}''_\ell}^{(h-i,t)}(\vec{\pi}'') \right).$$

(Here $P \leftrightarrow Q$ abbreviates the formula $(P \wedge Q) \vee (\neg P \wedge \neg Q)$.) If we consider the random sequence \vec{S} (in place of the arbitrary fixed sequence \vec{S}), then for every input $\vec{\pi}$, with high probability, the formula $\mathbf{join}_\ell(\vec{\pi})$ outputs 1 if, and only if, there exists a $\vec{\pi}$ -path (x_0, \dots, x_k) such that \vec{T}'_ℓ isolates the $\vec{\pi}'$ -path (x_0, \dots, x_j) and \vec{T}''_ℓ isolates the $\vec{\pi}''$ -path (x_i, \dots, x_k) .

Note that the number of $\vec{\pi}$ -paths in \vec{S} has expectation $n^{1 - \|\vec{a}\|} (\leq 1)$; it is easily shown that this number is at most $n^{o(1)}$ with high probability. For each $\vec{\pi}$ -path \vec{x} and $\ell \in [m]$, we have (by independence)

$$\mathbb{P} \left[\vec{x} \in \mathbf{T}_{\ell,0} \times \dots \times \mathbf{T}_{\ell,k} \mid \vec{x} \in \mathbf{S}_0 \times \dots \times \mathbf{S}_k \right] = n^{-\|\vec{b} - \vec{a}\|}.$$

A further argument⁸ shows that

$$\mathbb{P} \left[\vec{T}'_\ell \text{ isolates } (x_0, \dots, x_j) \text{ and } \vec{T}''_\ell \text{ isolates } (x_i, \dots, x_k) \mid \vec{x} \in \mathbf{T}_{\ell,0} \times \dots \times \mathbf{T}_{\ell,k} \right] = 1 - o(1).$$

By independence of $\vec{T}_1, \dots, \vec{T}_m$, we next have

$$\begin{aligned}\mathbb{P} \left[\bigvee_{\ell \in [m]} \left(\vec{T}'_\ell \text{ isolates } (x_0, \dots, x_j) \text{ and } \vec{T}''_\ell \text{ isolates } (x_i, \dots, x_k) \right) \mid \vec{x} \in \mathbf{S}_0 \times \dots \times \mathbf{S}_k \right] &\leq 1 - \left(1 - \Omega(n^{-\|\vec{b} - \vec{a}\|}) \right)^m \\ &\leq 1 - \exp(-\Omega(n^{-\|\vec{b} - \vec{a}\|} m)).\end{aligned}$$

Recalling that $m = n^{\|\vec{b} - \vec{a}\| + o(1)}$, the above bound will be $1 - n^{-\omega(1)}$ (i.e., “with high probability”) for a suitable choice of $o(1)$ in the exponent of m (for instance, if we set $m = n^{\|\vec{b} - \vec{a}\|} (\log n)^c$ for any constant $c > 1$).

We have shown that, with high probability, for every $\vec{\pi}$ -path (x_0, \dots, x_k) in $\mathbf{S}_0 \times \dots \times \mathbf{S}_k$, there exists $\ell \in [m]$ such that $\mathbf{join}_\ell(\vec{\pi})$ outputs 1. This justifies defining

$$\begin{aligned}\mathbf{f}_{\vec{a}, \vec{S}}(\vec{\pi}) &:= \bigvee_{\ell \in [m]} \mathbf{join}_\ell(\vec{\pi}) \\ &\wedge \bigwedge_{\substack{h \in \{0, \dots, j\} \\ t \in \{1, \dots, \lceil \log(n+1) \rceil\}}} \neg \left(\left(\bigvee_{\ell \in [m]} \mathbf{join}_\ell(\vec{\pi}) \wedge \mathbf{g}_{\vec{b}', \vec{T}'_\ell}^{(h,t)}(\vec{\pi}') \right) \wedge \left(\bigvee_{\ell \in [m]} \mathbf{join}_\ell(\vec{\pi}) \wedge \neg \mathbf{g}_{\vec{b}', \vec{T}'_\ell}^{(h,t)}(\vec{\pi}') \right) \right) \\ &\wedge \bigwedge_{\substack{h \in \{j+1, \dots, k\} \\ t \in \{1, \dots, \lceil \log(n+1) \rceil\}}} \neg \left(\left(\bigvee_{\ell \in [m]} \mathbf{join}_\ell(\vec{\pi}) \wedge \mathbf{g}_{\vec{b}'', \vec{T}''_\ell}^{(h-i,t)}(\vec{\pi}'') \right) \wedge \left(\bigvee_{\ell \in [m]} \mathbf{join}_\ell(\vec{\pi}) \wedge \neg \mathbf{g}_{\vec{b}'', \vec{T}''_\ell}^{(h-i,t)}(\vec{\pi}'') \right) \right).\end{aligned}$$

⁷A minor technicality arises when $a_h = b_h$: in this case, we should instead sample $\mathbf{T}_{\ell,h} \subseteq_{1/2} S_h$. This case can be also avoided by approximating $\chi(\vec{a})$ to an arbitrary additive constant $\varepsilon > 0$: if we instead consider \vec{c} defined by $c_h := b_h + (\varepsilon/k)$, then we have $\max\{\chi(c_0, \dots, c_j), \chi(c_i, \dots, c_k)\} \leq \max\{\chi(b_0, \dots, b_j) + (j+1)\varepsilon, \chi(b_i, \dots, b_k) + (k-i+1)\varepsilon\} \leq \chi(\vec{a}) + \varepsilon$.

⁸This is a straightforward union bound. Here is where we use the assumption that $b_h > a_h$ for all $h \in \{0, \dots, k\}$.

In light of the above discussion, for random \vec{S} , the subformula $\bigvee_{\ell \in [m]} \mathbf{join}_\ell(\vec{\pi})$ with high probability asserts the existence of a $\vec{\pi}$ -path in \vec{S} , while the remainder of $\mathbf{f}_{\vec{a}, \vec{S}}(\vec{\pi})$ asserts uniqueness. Formulas $\vec{g}_{\vec{a}, \vec{S}}$ are defined by

$$\mathbf{g}_{\vec{a}, \vec{S}}^{(h,t)}(\vec{\pi}) := \begin{cases} \bigvee_{\ell \in [m]} \mathbf{join}_\ell(\vec{\pi}) \wedge \mathbf{g}_{\vec{b}', \vec{T}'_\ell}^{(h,t)}(\vec{\pi}') & \text{if } h \in \{0, \dots, j\}, \\ \bigvee_{\ell \in [m]} \mathbf{join}_\ell(\vec{\pi}) \wedge \mathbf{g}_{\vec{b}'', \vec{T}''_\ell}^{(h-i,t)}(\vec{\pi}'') & \text{if } h \in \{j+1, \dots, k\}. \end{cases}$$

If formulas $f_{\vec{b}', \vec{T}'_\ell}$ and $\vec{g}_{\vec{b}', \vec{T}'_\ell}$ (respectively, $f_{\vec{b}'', \vec{T}''_\ell}$) and $\vec{g}_{\vec{b}'', \vec{T}''_\ell}$ have depth at most d' and size at most z' (respectively, d'' and z''), then it is readily seen that formulas $\mathbf{f}_{\vec{a}, \vec{S}}$ and $\vec{g}_{\vec{a}, \vec{S}}$ have depth d and size z where

$$d = \max\{d', d''\} + O(1), \quad z = O(m \cdot (k \log n)^2 \cdot (z' + z'')) = n^{\|\vec{b}-\vec{a}\|+o(1)} \cdot (z' + z'').$$

This recurrence justifies the bounds $d = O(k)$ and $z = n^{\chi(\vec{a})+o(1)}$.

As for the error probability of Properties (1) and (2), it should be clear that every usage of “with high probability” in this argument can be made to be $1 - n^{-\omega(1)}$ by setting $m = n^{\|\vec{b}-\vec{a}\|} (\log n)^{c_k}$ for suitable constants $c_k > 1$. \square

Lemma 7.5 has the following corollary, which for each $\vec{a} \in \mathcal{P}(k)$, gives a collection of randomized AC^0 formulas of aggregate size $n^{\chi(\vec{a})+\|\vec{a}\|+o(1)}$ that compute the product of k permutations. Moreover, these formulas, on input $\vec{\pi}$, produce a list of all paths $\vec{x} = (x_0, \dots, x_k) \in [n]^{\{0, \dots, k\}}$ such that $\pi_h(x_{h-1}) = x_h$ for all $h \in \{1, \dots, k\}$.

Corollary 7.6. *For every $\vec{a} \in \mathcal{P}(k)$, there exists a matrix of randomized AC^0 formulas*

$$\vec{h}_{\vec{a}} = \{\mathbf{h}_{\vec{a}}^{(\ell,t)}\}_{\ell \in \{1, \dots, n^{\|\vec{a}\|+o(1)}\}, t \in \{1, \dots, (k+1) \log[(n+1)]\}}$$

each of depth $O(k)$ and size $n^{\chi(\vec{a})+o(1)}$ and taking a sequence $\vec{\pi} = (\pi_1, \dots, \pi_k)$ of permutations $[n] \xrightarrow{\cong} [n]$ as input, such that the following properties hold with probability $1 - n^{-\omega(1)}$:

1. Each row in $\vec{h}_{\vec{a}}(\vec{\pi})$ is either the all-0 string or contains the binary representation of integers x_0, \dots, x_k for some $\vec{\pi}$ -path $\vec{x} \in [n]^{\{0, \dots, k\}}$.
2. For every $\vec{\pi}$ -path $\vec{x} \in [n]^{\{0, \dots, k\}}$, the binary representation of integers $x_0, \dots, x_k \in [n]$ is given by at least one row of $\vec{h}_{\vec{a}}(\vec{\pi})$.

Similar to the bound $\chi_A([n]^{\{0, \dots, k\}}) \leq n^{\chi(\vec{a})+\|\vec{a}\|+o(1)}$ in Remark 7.2, Corollary 7.6 is obtained from Lemma 7.5 by covering $[n]^{\{0, \dots, k\}}$ with $m := n^{\|\vec{a}\|+o(1)}$ random rectangles $\mathcal{S}_{\ell,0} \times \dots \times \mathcal{S}_{\ell,k}$ where each $\vec{\mathcal{S}}_\ell = (\mathcal{S}_{\ell,0}, \dots, \mathcal{S}_{\ell,k})$ has the same distribution as \vec{S} (i.e., $\mathcal{S}_{\ell,h} \subseteq_{n^{-a_h}} [n]$). The rows of $\vec{h}_{\vec{a}}(\vec{\pi})$ are then given by the conjunction of $\mathbf{f}_{\vec{a}, \vec{\mathcal{S}}_\ell}(\vec{\pi})$ with formulas $\vec{g}_{\vec{a}, \vec{\mathcal{S}}_\ell}(\vec{\pi})$. Property (1) is immediate from Lemma 7.5, while Property (2) follows by noting that, with high probability, every $\vec{\pi}$ -path in $[n]^{\{0, \dots, k\}}$ is isolated by a rectangle $\mathcal{S}_{\ell,0} \times \dots \times \mathcal{S}_{\ell,k}$ for some $\ell \in [m]$.

7.1 Upper bounds on $\chi(\vec{a}) + \|\vec{a}\|$

We describe a few different constructions giving upper bounds on $\chi(\vec{a}) + \|\vec{a}\|$ for sequences $\vec{a} \in \mathcal{P}(k)$. Thanks to Corollary 7.6, each of these constructions corresponds to randomized AC^0 formulas of size $n^{\chi(\vec{a})+\|\vec{a}\|+o(1)}$ computing the product of k permutations. Our best bound, $\frac{1}{3} \log_{(\sqrt{5}+1)/2}(k) + O(1)$, is obtained via a construction we call “Fibonacci overlapping”.

7.1.1 Recursive doubling

For $k \geq 2$, let

$$\vec{a}_k := \left(\underbrace{\frac{1}{2}, 0, \dots, 0}_{\lceil k/2 \rceil - 1}, \frac{1}{2}, \underbrace{0, \dots, 0}_{\lfloor k/2 \rfloor - 1}, \frac{1}{2} \right)$$

Then $\chi(\vec{a}_2) = \chi(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}) = \max\{\chi(\frac{1}{2}, \frac{1}{2}), \chi(\frac{1}{2}, \frac{1}{2})\} = 0$ and for $k \geq 3$,

$$\chi(\vec{a}_k) \leq \max\{\chi(\frac{1}{2}, \underbrace{0, \dots, 0}_{\lceil k/2 \rceil - 1}, \frac{1}{2}), \chi(\frac{1}{2}, \underbrace{0, \dots, 0}_{\lfloor k/2 \rfloor - 1}, \frac{1}{2})\} \leq \max\{\frac{1}{2} + \chi(\vec{a}_{\lceil k/2 \rceil}), \frac{1}{2} + \chi(\vec{a}_{\lfloor k/2 \rfloor})\} \leq \frac{1}{2} \lceil \log_2(k) \rceil.$$

This construction achieves

$$\chi(\vec{a}_k) + \|\vec{a}_k\| \leq \frac{1}{2} \lceil \log_2(k) \rceil + 1.$$

7.1.2 Maximally overlapping joins

For $k \geq 2$, let

$$\vec{a}_k := \left(\underbrace{\frac{1}{k}, \dots, \frac{1}{k}}_{k+1} \right) \in \mathcal{P}(k).$$

Then $\chi(\vec{a}_2) = \chi(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}) = 0$ and for $k \geq 3$,

$$\chi(\vec{a}_k) \leq \max\{\chi(\underbrace{\frac{1}{k}, \dots, \frac{1}{k}}_k), \chi(\underbrace{\frac{1}{k}, \dots, \frac{1}{k}}_k)\} \leq \chi(\vec{a}_{k-1}) + k\left(\frac{1}{k-1} - \frac{1}{k}\right) = \chi(\vec{a}_{k-1}) + \frac{1}{k-1} \leq \frac{1}{2} + \dots + \frac{1}{k-1}.$$

This construction achieves

$$\chi(\vec{a}_k) + \|\vec{a}_k\| \leq 1 + \frac{1}{2} + \dots + \frac{1}{k-1} + \frac{1}{k} = \ln(k) + O(1).$$

Since $\ln(k) \approx 0.69 \log_2(k)$, this upper bound is worse than the one from recursive doubling.

It turns out that a $\frac{1}{2} \log_2(k) + O(1)$ upper bound is achievable via a different construction using the maximally overlapping join-tree. If $k = 2^\ell + t$ where $\ell \geq 0$ and $t \in \{0, \dots, 2^\ell - 1\}$, we instead define

$$\vec{a}_k := \left(\underbrace{\frac{1}{2^{\ell+1}}, \dots, \frac{1}{2^{\ell+1}}}_t, \underbrace{\frac{1}{2^\ell}, \dots, \frac{1}{2^\ell}}_{2^\ell - t + 1}, \underbrace{\frac{1}{2^{\ell+1}}, \dots, \frac{1}{2^{\ell+1}}}_t \right).$$

(Note that $\vec{a} \in \mathcal{P}(k)$ since $\|\vec{a}\| = (2^\ell - t + 1)\frac{1}{2^\ell} + 2t\frac{1}{2^{\ell+1}} = 1 + \frac{1}{2^\ell} > 1$.) In the base case $k = 1$ (i.e., $\ell = t = 0$), we have $\chi(\vec{a}_2) = \chi(1, 1) = 0$. When $\ell \geq 1$ and $t \geq 1$, we have

$$\begin{aligned} \chi(\vec{a}_k) &\leq \max\left\{ \chi\left(\underbrace{\frac{1}{2^{\ell+1}}, \dots, \frac{1}{2^{\ell+1}}}_t, \underbrace{\frac{1}{2^\ell}, \dots, \frac{1}{2^\ell}}_{2^\ell - t + 1}, \underbrace{\frac{1}{2^{\ell+1}}, \dots, \frac{1}{2^{\ell+1}}}_{t-1}\right), \chi\left(\underbrace{\frac{1}{2^{\ell+1}}, \dots, \frac{1}{2^{\ell+1}}}_{t-1}, \underbrace{\frac{1}{2^\ell}, \dots, \frac{1}{2^\ell}}_{2^\ell - t + 1}, \underbrace{\frac{1}{2^{\ell+1}}, \dots, \frac{1}{2^{\ell+1}}}_t\right) \right\} \\ &= \chi\left(\underbrace{\frac{1}{2^{\ell+1}}, \dots, \frac{1}{2^{\ell+1}}}_t, \underbrace{\frac{1}{2^\ell}, \dots, \frac{1}{2^\ell}}_{2^\ell - t + 1}, \underbrace{\frac{1}{2^{\ell+1}}, \dots, \frac{1}{2^{\ell+1}}}_{t-1}\right) \quad (\text{by symmetry}) \\ &= \chi\left(\underbrace{\frac{1}{2^{\ell+1}}, \dots, \frac{1}{2^{\ell+1}}}_{t-1}, \underbrace{\frac{1}{2^\ell}, \dots, \frac{1}{2^\ell}}_{2^\ell - t + 2}, \underbrace{\frac{1}{2^{\ell+1}}, \dots, \frac{1}{2^{\ell+1}}}_{t-1}\right) + \frac{1}{2^\ell} - \frac{1}{2^{\ell+1}} \quad (t^{\text{th}} \text{ coordinate increases from } \frac{1}{2^{\ell+1}} \text{ to } \frac{1}{2^\ell}) \\ &= \chi(\vec{a}_{k-1}) + \frac{1}{2^{\ell+1}}. \end{aligned}$$

When $\ell \geq 1$ and $t = 0$ (i.e., $k = 2^\ell$), we have

$$\begin{aligned} \chi(\vec{a}_k) &= \chi\left(\underbrace{\frac{1}{2^\ell}, \dots, \frac{1}{2^\ell}}_{2^\ell + 1}\right) \leq \chi\left(\underbrace{\frac{1}{2^\ell}, \dots, \frac{1}{2^\ell}}_{2^\ell}\right) \leq \chi\left(\underbrace{\frac{1}{2^\ell}, \dots, \frac{1}{2^\ell}}_{2^{\ell-1}-1}, \frac{1}{2^{\ell-1}}, \frac{1}{2^{\ell-1}}, \underbrace{\frac{1}{2^\ell}, \dots, \frac{1}{2^\ell}}_{2^{\ell-1}-1}\right) + 2\left(\frac{1}{2^{\ell-1}} - \frac{1}{2^\ell}\right) \\ &= \chi(\vec{a}_{k-1}) + \frac{1}{2^{\ell-1}}. \end{aligned}$$

For any $k = 2^\ell + t$, this recurrence shows

$$\begin{aligned}
\chi(\vec{a}_k) &\leq t \cdot \frac{1}{2^{\ell+1}} + \chi(\vec{a}_{2^\ell}) \\
&\leq t \cdot \frac{1}{2^{\ell+1}} + \frac{1}{2^{\ell-1}} + \chi(\vec{a}_{2^{\ell-1}}) \\
&\leq t \cdot \frac{1}{2^{\ell+1}} + \left(\frac{1}{2^{\ell-1}} + (2^{\ell-1} - 1) \cdot \frac{1}{2^\ell}\right) + \chi(\vec{a}_{2^{\ell-1}}) \\
&= t \cdot \frac{1}{2^{\ell+1}} + \left(\frac{1}{2} + \frac{1}{2^\ell}\right) + \chi(\vec{a}_{2^{\ell-1}}) \\
&\leq t \cdot \frac{1}{2^{\ell+1}} + \sum_{j=1}^{\ell} \left(\frac{1}{2} + \frac{1}{2^j}\right) \\
&= \frac{1}{2}(\ell + \frac{t}{2^\ell}) + 1 - \frac{1}{2^\ell}.
\end{aligned}$$

This second construction thus achieves

$$\chi(\vec{a}) + \|\vec{a}\| \leq \frac{1}{2}(\ell + \frac{t}{2^\ell}) + 2 = \frac{1}{2} \log_2(k) + O(1).$$

7.1.3 Fibonacci overlapping joins

Let $\text{Fib}(1) = \text{Fib}(2) = 1$ and for $\ell \geq 3$, let $\text{Fib}(\ell) := \text{Fib}(\ell - 1) + \text{Fib}(\ell - 2)$. For $\ell \geq 4$, let

$$\chi(\vec{a}_{\text{Fib}(\ell)}) := \left(\frac{1}{3}, \underbrace{0, \dots, 0}_{\text{Fib}(\ell-2)-1}, \frac{1}{3}, \underbrace{0, \dots, 0}_{\text{Fib}(\ell-3)-1}, \frac{1}{3}, \underbrace{0, \dots, 0}_{\text{Fib}(\ell-2)-1}, \frac{1}{3}\right) \in \mathcal{P}(\text{Fib}(\ell))$$

We have $\text{Fib}(4) = 3$ and

$$\chi(\vec{a}_{\text{Fib}(4)}) = \chi\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right) \leq \chi\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right) \leq \frac{1}{3} + \max\left\{\chi\left(\frac{1}{3}, \frac{2}{3}\right), \chi\left(\frac{2}{3}, \frac{1}{3}\right)\right\} = \frac{1}{3}.$$

For $\ell \geq 5$, we have

$$\begin{aligned}
\chi(\vec{a}_{\text{Fib}(\ell)}) &\leq \max\left\{\chi\left(\frac{1}{3}, \underbrace{0, \dots, 0}_{\text{Fib}(\ell-2)-1}, \frac{1}{3}, \underbrace{0, \dots, 0}_{\text{Fib}(\ell-3)-1}, \frac{1}{3}\right), \chi\left(\frac{1}{3}, \underbrace{0, \dots, 0}_{\text{Fib}(\ell-3)-1}, \frac{1}{3}, \underbrace{0, \dots, 0}_{\text{Fib}(\ell-2)-1}, \frac{1}{3}\right)\right\} \\
&= \chi\left(\frac{1}{3}, \underbrace{0, \dots, 0}_{\text{Fib}(\ell-2)-1}, \frac{1}{3}, \underbrace{0, \dots, 0}_{\text{Fib}(\ell-3)-1}, \frac{1}{3}\right) \quad (\text{by symmetry}) \\
&\leq \frac{1}{3} + \chi\left(\frac{1}{3}, \underbrace{0, \dots, 0}_{\text{Fib}(\ell-3)-1}, \frac{1}{3}, \underbrace{0, \dots, 0}_{\text{Fib}(\ell-4)-1}, \frac{1}{3}, \underbrace{0, \dots, 0}_{\text{Fib}(\ell-3)-1}, \frac{1}{3}\right) = \frac{1}{3} + \chi(\vec{a}_{\text{Fib}(\ell-1)}) = \frac{1}{3}\ell - 1.
\end{aligned}$$

For $k = \text{Fib}(\ell)$ with $\ell \geq 4$, this construction gives $\vec{a}_k \in \mathcal{P}(k)$ with

$$\chi(\vec{a}_k) + \|\vec{a}_k\| \leq \frac{1}{3}(\ell + 1).$$

For $\text{Fib}(\ell - 1) < k \leq \text{Fib}(\ell)$, the bound $\chi(\vec{a}_k) + \|\vec{a}_k\| \leq \frac{1}{3}(\ell + 1)$ extends to all $\vec{a}_k \in \mathcal{P}(k)$ of the form

$$\vec{a}_k := \left(\frac{1}{3}, \underbrace{0, \dots, 0}_{\leq \text{Fib}(\ell-2)-1}, \frac{1}{3}, \underbrace{0, \dots, 0}_{\leq \text{Fib}(\ell-3)-1}, \frac{1}{3}, \underbrace{0, \dots, 0}_{\leq \text{Fib}(\ell-2)-1}, \frac{1}{3}\right).$$

This construction proves the following

Theorem 7.7. *For all $k \geq 1$, there exists $\vec{a} \in \mathcal{P}(k)$ with*

$$\chi(\vec{a}) + \|\vec{a}\| = \frac{1}{3} \log_\varphi(k) + O(1)$$

where $\varphi = (\sqrt{5} + 1)/2$ is the golden ratio.

Since $\frac{1}{3} \log_\varphi(k) = \log_{\sqrt{5}+2}(k) \leq 0.49 \log_2(k)$, Theorem 7.7 improves the $\frac{1}{2} \log_2(k) + O(1)$ upper bounds from the recursive doubling and maximally overlapping join-trees described above. As a corollary of Corollary 7.6 and Theorem 7.7, we have

Corollary 7.8. *There exist randomized AC⁰ formulas of size $n^{\frac{1}{3} \log_\varphi(k) + O(1)}$ that compute the product of k ($n \times n$)-permutation matrices.*

7.2 Tightness of upper bounds

We say that a join-tree A (over P_∞) is *connected* if G_D is connected for all $D \preceq A$. For every connected join-tree A with graph P_k , we can consider the constrained complexity measure $\chi_A : \mathcal{P}(k) \rightarrow \mathbb{R}_{\geq 0}$ where parameters $0 < i \leq j < k$ in the definition of $\chi(\vec{a})$ are fixed according to A . As described in Remark 7.2, the potential function $\Phi(A)$ implies a lower bound $\chi_A(\vec{a}) \geq \Phi(A) - \|\vec{a}\|$.

Let RD_k , MO_k and FO_k be the “recursive doubling”, “maximally overlapping” and “Fibonacci overlapping” connected join-trees recursively defined by $\text{RD}_1 = \text{MO}_1 = \text{FO}_1 = \langle \{0, 1\} \rangle$ and

$$\begin{aligned} \text{RD}_k &= \text{RD}_{0,k} := \langle \text{RD}_{0, \lceil k/2 \rceil}, \text{RD}_{\lceil k/2 \rceil, k} \rangle && \text{for } k \geq 2, \\ \text{MO}_k &= \text{MO}_{0,k} := \langle \text{MO}_{0, k-1}, \text{MO}_{1, k} \rangle && \text{for } k \geq 2, \\ \text{FO}_k &= \text{FO}_{0,k} := \langle \text{FO}_{0, \text{Fib}(\ell-1)}, \text{FO}_{\text{Fib}(\ell-1), k} \rangle && \text{for } k = \text{Fib}(\ell), \ell \geq 3. \end{aligned}$$

The upper bounds on $\chi(\cdot)$ in the previous subsection respectively apply to constrained complexity measures $\chi_{\text{RD}_k}(\cdot)$, $\chi_{\text{MO}_k}(\cdot)$ and $\chi_{\text{FO}_k}(\cdot)$.

With respect to these particular join-trees, the upper bounds of previous subsection are in fact tight! By (†) we have

$$\begin{aligned} \Phi(\text{RD}_k) &\geq \Phi(\text{RD}_{\lceil \lceil k/2 \rceil / 2 \rceil}) + 1 && \text{for } k \geq 4, \\ \Phi(\text{FO}_{\text{Fib}(\ell)}) &\geq \Phi(\text{FO}_{\text{Fib}(\ell-3)}) + 1 && \text{for } k = \text{Fib}(\ell), \ell \geq 5. \end{aligned}$$

It follows that $\Phi(\text{RD}_k) \geq \frac{1}{2} \log_2(k)$ and $\Phi(\text{FO}_k) \geq \frac{1}{3} \log_\varphi(k) - O(1)$. We get a lower bound on $\Phi(\text{MO}_k) \geq \frac{1}{2} \log_2(k) - O(1)$ via (‡):

$$\Phi(\text{MO}_k) \geq \frac{1}{2} (\Phi(\text{MO}_{0, \lfloor (k-1)/2 \rfloor}) + \Phi(\text{MO}_{\lfloor (k-1)/2 \rfloor, k}) + 1) \quad \text{for } k \geq 3.$$

Therefore, for all $\vec{a} \in \mathcal{P}(k)$, we have

$$\begin{aligned} \chi_{\text{RD}_k}(\vec{a}) + \|\vec{a}\| &\geq \frac{1}{2} \log_2(k), \\ \chi_{\text{MO}_k}(\vec{a}) + \|\vec{a}\| &\geq \frac{1}{2} \log_2(k) - O(1), \\ \chi_{\text{FO}_k}(\vec{a}) + \|\vec{a}\| &\geq \frac{1}{3} \log_\varphi(k) - O(1). \end{aligned}$$

This establishes the tightness of our upper bounds for these specific join-trees. It is open whether a different connected join-tree achieves a better bound. (The best lower bound on $\chi(\vec{a}) + \|\vec{a}\|$ that we could determine is $\log_{\sqrt{5}+4}(k) - 1$ via a strengthening of the argument in Section 6 in the case of connected join-trees.)

Experimental results. For any connected join-tree A , the value $\min_{\vec{a} \in \mathcal{P}(k)} \chi_A(\vec{a})$ is computable by a linear program with $O(\sum_{D \preceq A} |V(D)|)$ variables and $O(\sum_{D \preceq A} |V(D)|)$ constraints. In fact, our second upper bound for the maximally overlapping pattern (achieving $\chi_{\text{MO}_k}(\vec{a}_k) + \|\vec{a}_k\| = \frac{1}{2} \log_2(k) + O(1)$) was found with the help of this linear programs!

We experimentally searched for connected join-trees A that beat the $\frac{1}{3} \log_\varphi(k) + O(1)$ upper bound via Fibonacci overlapping by evaluating $\min_{\vec{a} \in \mathcal{P}(k)} \chi_A(\vec{a}) + \|\vec{a}\|$ on various examples, both structured and randomly generated. We could not find any better upper bound. (In particular, join-trees FO_k appears optimal among a broad class of “recursively overlapping” join-trees.) It is tempting to conjecture that FO_k in fact gives the optimal bound, that is, $\chi_A(\vec{a}) + \|\vec{a}\| \leq \frac{1}{3} \log_\varphi(k) + O(1)$ for all $\vec{a} \in \mathcal{P}(k)$. We leave this as an intriguing open question.

8 Open problems

We conclude by mentioning some open questions raised by this work.

Problem 1. Prove that $\tau(G) = \Omega(\text{td}(G))$ for all graphs G . (An $\tilde{\Omega}(\text{td}(G))$ would also be interesting.)

Problem 1 is unlikely to follow from any excluded-minor approximation of tree-depth along the lines of Theorem 1.6. A first step to resolving this problem is to identify, for each graph G , a particular threshold weighting θ such that $\mathbf{X}_{\theta,n}$ is a “hard” input distribution with respect to the average-case AC^0 formula size of $\text{SUB}(G)$. (The paper [8] does precisely this with respect to AC^0 circuit size.)

Problem 1 should be easier to tackle in the special case of trees. (We remark that our lower bounds for P_k and T_k , combined with results in [4, 6], imply that $\tau(T) = \Omega(\sqrt{\text{td}(T)})$ for all trees T .)

Problem 2. Prove that $\tau(T) = \Omega(\text{td}(T))$ for all trees T .

A solution to Problem 2 could perhaps be shown by a common generalization of our lower bounds for P_k and T_k .

A third open problem is to nail down the exact average-case AC^0 formula size of $\text{SUB}(P_k)$ (or the related problem of multiplying k permutations).

Problem 3. Prove that $\tau(P_k) = \frac{1}{3} \log_{\varphi}(k)$ or find an upper bound improving Theorem 1.10.

Finally and most ambitiously:

Problem 4. Prove that $n^{\tau(G)-o(1)}$ is a lower bound the *unrestricted* formula size $\text{SUB}(G)$.

This of course would imply $NC^1 \neq NL$. An $n^{\Omega(\log k)}$ lower bound for $\text{SUB}(P_k)$ in the average-case (or for the problem of multiplying k permutations) would moreover imply $NC^1 \neq L$. Although Problem 3 lies beyond current techniques, the applicability of the pathset framework in establishing $n^{\tau(G)-o(1)}$ lower bounds in the disparate AC^0 and monotone settings is possibly reason for optimism.

A Appendix: Lower bound $\tau(P_k) \geq \frac{1}{2} \log_{\sqrt{13}+1}(k)$ from [15]

This appendix gives the proof of Lemma 2.15 from [15]. As in Section 6, we consider infinite pattern graph P_{∞} with the constant threshold weighting $E(P_{\infty}) \rightarrow \{1\}$.

Definition A.1. For a join-tree A , let $\lambda(A)$ denote the length of the longest path in A (i.e., the number of edges in the largest connected component of A).

We omit the proof of the following lemma, which is similar to Lemma 6.4.

Lemma A.2. For every join-tree A and set S , if S intersects t distinct connected components of G_A , then

$$\Phi(A) \geq \Phi(A \ominus S) + t.$$

We now present the result from [15] that implies Lemma 2.15. (The precise value of c in Lemma A.3, below, is thanks to an optimization suggested by an anonymous referee of the journal paper [15].)

Lemma A.3. For every join-tree A , $\Phi(A) \geq \frac{1}{c} \log(\lambda(A)) + \Delta(A)$ where $c = 2 \log(\sqrt{13} + 1)$.

Proof. Here c is chosen such that $\frac{1}{2} - \frac{1}{2^{c/2}} - \frac{1}{2^{c-1}} = \frac{1}{2^{c-2}}$.

We argue by induction on join-trees. The base case where A is atomic is trivial. For the induction step, let A be a non-atomic join-tree and assume the lemma holds for all smaller join-trees. We will consider a sequence of cases, which will be summarized at the end of the proof.

First, consider the case that G_A is disconnected. Let $t = \Delta(A) (\geq 2)$. Let S be the set of all vertices of G_A , except those in the largest connected component of G_A . We have

$$\begin{aligned} \Phi(A) &\geq \Phi(A \upharpoonright S) + t - 1 && \text{(Lemma A.2)} \\ &\geq \frac{1}{c} \log(\lambda(A \ominus S)) + \Delta(A \upharpoonright S) + t - 1 && \text{(induction hypothesis)} \\ &= \frac{1}{c} \log(\lambda(A)) + \Delta(A). \end{aligned}$$

This proves the lemma in the case where G_A is disconnected.

Therefore, we proceed under the assumption that G_A is connected (i.e. $\Delta(A) = 1$). Without loss of generality, we assume that $G_A = P_k$ (i.e. $\lambda(A) = k$). Our goal is to show that

$$\Phi(A) \geq \frac{1}{c} \log(k) + 1.$$

Consider the case that there exists a sub-join-tree $A' \preceq A$ such that $|E_{A'}| \geq \frac{1}{2^{c-1}}k$ and $\Delta(A') \geq 2$. Note that $\lambda(A') \geq |E_{A'}|/\Delta(A')$ (i.e. the number of edges in the largest component of $G_{A'}$ is at least the number of edges in $G_{A'}$ divided by the number of components in $G_{A'}$). We have

$$\begin{aligned} \Phi(A) &\geq \Phi(A') \\ &\geq \frac{1}{c} \log(\lambda(A')) + \Delta(A') && \text{(induction hypothesis)} \\ &\geq \frac{1}{c} \log(k) - \frac{c-1}{c} - \frac{1}{c} \log(\Delta(A')) + \Delta(A') && (\lambda(A') \geq |E_{A'}|/\Delta(A') \geq \frac{1}{2^{c-1}}k\Delta(A')) \\ &\geq \frac{1}{c} \log(k) - \frac{c-1}{c} - \frac{1}{c} \log(2) + 2 && (\Delta(A') \geq 2 \text{ and } x - \frac{1}{c} \log x \text{ increasing}) \\ &= \frac{1}{c} \log(k) + 1. \end{aligned}$$

This proves the lemma in this case.

Therefore, we proceed under the following assumption:

$$(\otimes) \quad \text{for all } A' \preceq A, \text{ if } |E_{A'}| \geq \frac{1}{2^{c-1}}k \text{ then } \Delta(A') = 1.$$

Going forward, the following notation will be convenient: for a proper sub-join-tree $B \prec A$, let B^\dagger denote the parent of B in A , and let B^\sim denote the sibling of B in A . Note that $B^\dagger = \{B, B^\sim\} \preceq A$.

By walking down the join-tree A , we can proper sub-join-trees $B, Z \prec A$ such that

$$v_0 \in V_B, \quad v_k \in V_Z, \quad |E_B|, |E_Z| < \frac{1}{2^{c/2}}k, \quad |E_{B^\dagger}|, |E_{Z^\dagger}| \geq \frac{1}{2^{c/2}}k.$$

Fix any choice of such B and Z . Note that G_{B^\dagger} and G_{Z^\dagger} are connected by (\otimes) . In particular, G_{B^\dagger} is a path of length $|E_{B^\dagger}|$ with initial endpoint v_0 , and G_{Z^\dagger} is a path of length $|E_{Z^\dagger}|$ with final endpoint v_k .

Consider the case that B^\dagger and Z^\dagger are vertex-disjoint. Note that $\frac{1}{2^{c/2}}k \geq \frac{1}{2^{c-1}}k$, so the assumption (\otimes) implies that B^\dagger and Z^\dagger are connected and $\lambda(B^\dagger), \lambda(Z^\dagger) \geq \frac{1}{2^{c/2}}k$. Let Y denote the least common ancestor of B^\dagger and Z^\dagger in A . We have

$$\begin{aligned} \Phi(A) &\geq \Phi(Y) \\ &\geq \frac{1}{2}(\Phi(B^\dagger) + \Phi(Z^\dagger \ominus B^\dagger) + \Delta(Y) + \Delta(Y \ominus \{B^\dagger, Z^\dagger\})) && \text{(by } (\ddagger)\text{)} \\ &\geq \frac{1}{2}(\Phi(B^\dagger) + \Phi(Z^\dagger)) + \frac{1}{2} && (\Delta(Y) \geq 1 \text{ and } Z^\dagger \ominus B^\dagger = Z^\dagger) \\ &\geq \frac{1}{2}(\frac{1}{c} \log(\lambda(B^\dagger)) + \Delta(B^\dagger) + \frac{1}{c} \log(\lambda(Z^\dagger)) + \Delta(Z^\dagger)) + \frac{1}{2} && \text{(induction hypothesis)} \\ &\geq \frac{1}{2}(\frac{1}{c} \log(\frac{1}{2^{c/2}}k) + 1 + \frac{1}{c} \log(\frac{1}{2^{c/2}}k) + 1) + \frac{1}{2} \\ &= \frac{1}{c} \log(k) + 1. \end{aligned}$$

Therefore, we proceed under the assumption that B^\dagger and Z^\dagger are not vertex-disjoint. It follows that $\lambda(B^\dagger) \geq k/2$ or $\lambda(Z^\dagger) \geq k/2$. Without loss of generality, we assume that $\lambda(B^\dagger) \geq k/2$. (We now forget about Z and Z^\dagger .)

Before continuing, let's take stock of the assumptions we have made so far:

$$G_A = P_k, \quad (\otimes), \quad B \preceq A, \quad v_0 \in V_B, \quad |E_B| < \frac{1}{2^{c/2}}k, \quad |E_{B^\dagger}| = \lambda(B^\dagger) \geq k/2.$$

Going forward, we will define vertices v_r, v_s, v_t where $0 < r < s < t \leq k$. The following illustration will be helpful for what follows:




Figure 18: Schematic of sub-join-trees in the argument.

We first define $v_r \in B$ and $v_t \in B^\sim$ as follows: Let $\{v_0, \dots, v_r\}$ be the component of G_B containing v_0 . (That is, the component of v_0 in G_B is a path whose initial vertex is v_0 ; let v_r be the final vertex in this path.) Let v_t be the vertex in V_{B^\sim} with maximal index t (i.e. farthest away from v_0).

Note that E_B contains edges $\text{edge}(v_i, v_{i+1})$ for all $i \in \{0, \dots, r-1\} \cup \{t, \dots, \lceil k/2 \rceil - 1\}$. (In the event that $t < k/2$, since $G_{B^\uparrow} = G_B \cup G_{B^\sim}$ is a path of length $\geq k/2$ and G_{B^\sim} does not contain vertices $v_{t+1}, \dots, v_{\lceil k/2 \rceil}$, it follows that G_B contains all edges between v_t and $v_{\lceil k/2 \rceil}$.) Therefore, $r + (k/2) - t \leq |E_B| < \frac{1}{2^{c/2}}k$. It follows that

$$t - r > \left(\frac{1}{2} - \frac{1}{2^{c/2}}\right)k.$$

Next, note that $|E_{B^\sim}| \geq |E_{B^\uparrow}| - |E_B| \geq \left(\frac{1}{2} - \frac{1}{2^{c/2}}\right)k > \frac{1}{2^{c-1}}k$. We now walk down B^\sim to find a proper sub-join-tree $C \prec B^\sim$ such that

$$v_t \in V_C, \quad |E_C| < \frac{1}{2^{c-1}}k, \quad |E_{C^\uparrow}| \geq \frac{1}{2^{c-1}}k.$$

Fix any choice of such C . Note that G_{C^\uparrow} is connected by (\otimes) .

Consider the case that $|E_{C^\uparrow}| < \left(\frac{1}{2} - \frac{1}{2^{c/2}}\right)k$. Since G_{C^\uparrow} is connected and $v_t \in V_{C^\uparrow}$ and $t - r > \left(\frac{1}{2} - \frac{1}{2^{c/2}}\right)k$, it follows that $V_{C^\uparrow} \cap \{v_0, \dots, v_r\} = \emptyset$ and hence $\Delta(B \ominus C^\uparrow) \geq 1$. We have

$$\begin{aligned} \Phi(A) &\geq \Phi(B^\uparrow) \geq \Phi(C^\uparrow) + \Delta(B \ominus C^\uparrow) + \Delta(B^\uparrow \ominus \{B, C^\uparrow\}) && \text{(by (†))} \\ &\geq \Phi(C^\uparrow) + 1 \\ &\geq \frac{1}{c} \log(\lambda(C^\uparrow)) + \Delta(C^\uparrow) + 1 && \text{(induction hypothesis)} \\ &\geq \frac{1}{c} \log\left(\frac{1}{2^{c-1}}k\right) + 2 \\ &> \frac{1}{c} \log(k) + 1. \end{aligned}$$

Therefore, we proceed under the assumption that $|E_{C^\uparrow}| \geq \left(\frac{1}{2} - \frac{1}{2^{c/2}}\right)k$. Since $E_{C^\uparrow} = E_C \cup E_{C^\sim}$, we have

$$|E_{C^\sim}| \geq |E_{C^\uparrow}| - |E_C| > \left(\frac{1}{2} - \frac{1}{2^{c/2}} - \frac{1}{2^{c-1}}\right)k = \frac{1}{2^{c-2}}k.$$

We now define vertex $v_s \in V_C$. Since v_t is the vertex of G_{B^\sim} with maximal index, it follows that $\text{edge}(v_t, v_{t+1}) \notin E_{B^\sim}$ and hence $\text{edge}(v_t, v_{t+1}) \notin E_C$ (since $C \prec B^\sim$). Therefore, the component of G_C

containing v_t is a path with final vertex v_t ; let v_s be the initial vertex in this path. That is, $\{v_s, \dots, v_t\}$ is the component of G_C which contains v_t . Recall that $t - r > (\frac{1}{2} - \frac{1}{2^{c/2}})k$ and note that $t - s \leq |E_C| < \frac{1}{2^{c-1}}k$. Therefore,

$$s - r = (t - r) - (t - s) > (\frac{1}{2} - \frac{1}{2^{c/2}} - \frac{1}{2^{c-1}})k = \frac{1}{2^{c-2}}k.$$

We now claim that there exists a proper sub-join-tree $D \prec C^\sim$ such that

$$\frac{1}{2^{c-1}}k \leq |E_D| < \frac{1}{2^{c-2}}k.$$

To see this, note that there exists a chain of sub-join-trees $C^\sim = D_0 \succ D_1 \succ \dots \succ D_j$ such that D_j is atomic and $D_{i-1} = D_i^\uparrow$ and $|E_{D_i}| \geq |E_{D_i^\sim}|$ for all $i \in \{1, \dots, j\}$. Since $|E_{D_0}| > \frac{1}{2^{c-2}}k$ and $|E_{D_j}| = 1$ and $|E_{D_{i-1}}| \leq |E_{D_i}| + |E_{D_i^\sim}| \leq 2|E_{D_i}|$, it must be the case that there exists $i \in \{1, \dots, j\}$ such that $\frac{1}{2^{c-1}}k \leq |E_{D_i}| < \frac{1}{2^{c-2}}k$.

Since $|E_D| \geq \frac{1}{2^{c-1}}k$, (\otimes) implies that G_D is connected. Since $|E_D| < \frac{1}{2^{c-2}}k$ and $s - r > \frac{1}{2^{c-2}}k$, it follows that V_D cannot contain both v_r and v_s . We are now down to our final two cases: either $v_r \notin V_D$ or $v_s \notin V_D$.

First, suppose that $v_r \notin V_D$. We have $\Delta(B \ominus D) \geq 1$ and hence

$$\begin{aligned} \Phi(A) &\geq \Phi(B^\uparrow) \geq \Phi(D) + \Delta(B \ominus D) + \Delta(B^\uparrow \ominus \{B, D\}) && \text{(by } (\dagger) \text{)} \\ &\geq \Phi(D) + 1 \\ &\geq \frac{1}{c} \log(\lambda(D)) + \Delta(D) + 1 && \text{(induction hypothesis)} \\ &\geq \frac{1}{c} \log(\frac{1}{2^{c-1}}k) + 2 \\ &> \frac{1}{c} \log(k) + 1. \end{aligned}$$

Finally, we are left with the alternative that $v_s \notin V_D$. In this case $\Delta(C \ominus D) \geq 1$ and hence (substituting C for B in the above), we have

$$\Phi(A) \geq \Phi(C^\uparrow) \geq \Phi(D) + \Delta(C \ominus D) + \Delta(C^\uparrow \ominus \{C, D\}) \geq \Phi(D) + 1 > \frac{1}{c} \log(k) + 1.$$

We have now covered all cases. In summary, we considered cases in the following sequence:

1. $\Delta(A) \geq 2$ otherwise assume $G_A = P_k$ w.l.o.g.,
2. $\exists A' \prec A$ with $\Delta(A') \geq 2$ and $\lambda(A') \geq \frac{1}{2^{c-1}}k$ otherwise assume (\otimes) ,
3. B^\uparrow and Z^\uparrow are vertex-disjoint otherwise assume $|E_{B^\uparrow}| \geq k/2$ w.l.o.g.,
4. $|E_{C^\uparrow}| < (\frac{1}{2} - \frac{1}{2^{c/2}})k$ otherwise assume $|E_{C^\uparrow}| \geq (\frac{1}{2} - \frac{1}{2^{c/2}})k$,
5. $v_r \notin V_D$ or $v_s \notin V_D$.

Since $\Phi(A) \geq \frac{1}{c} \log(\lambda(A)) + \Delta(A)$ in each case, the proof is complete. □

References

- [1] Kazuyuki Amano. k -Subgraph isomorphism on AC^0 circuits. *Computational Complexity*, 19(2):183–210, 2010.
- [2] Julia Chuzhoy and Zihan Tan. Towards tight(er) bounds for the excluded grid theorem. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1445–1464. SIAM, 2019.
- [3] Marek Cygan, Fedor V Fomin, Alexander Golovnev, Alexander S Kulikov, Ivan Mihajlin, Jakub Pachocki, and Arkadiusz Socała. Tight bounds for graph homomorphism and subgraph isomorphism. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 1643–1649. SIAM, 2016.

- [4] Wojciech Czerwiński, Wojciech Nadara, and Marcin Pilipczuk. Improved bounds for the excluded-minor approximation of treedepth. *arXiv preprint arXiv:1904.13077*, 2019.
- [5] Reinhard Diestel. Graph theory 3rd ed. *Graduate texts in mathematics*, 173, 2005.
- [6] Ken-ichi Kawarabayashi and Benjamin Rossman. A polynomial excluded-minor characterization of treedepth. In *29th ACM-SIAM Symposium on Discrete Algorithms*, pages 234–246, 2018.
- [7] Robert Krauthgamer and Ohad Trabelsi. Conditional lower bound for subgraph isomorphism with a tree pattern. *arXiv preprint arXiv:1708.07591*, 2017.
- [8] Yuan Li, Alexander Razborov, and Benjamin Rossman. On the AC^0 complexity of subgraph isomorphism. *SIAM Journal on Computing*, 46(3):936–971, 2017.
- [9] Dániel Marx. Can you beat treewidth? *Theory of Computing*, 6:85–112, 2010.
- [10] Jaroslav Nešetřil and Patrice Ossona De Mendez. Tree-depth, subgraph coloring and homomorphism bounds. *European Journal of Combinatorics*, 27(6):1022–1041, 2006.
- [11] Gregory Rosenthal. Beating treewidth for average-case subgraph isomorphism. In *14th International Symposium on Parameterized and Exact Computation (IPEC 2019)*, volume 148 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 24:1–24:14, 2019.
- [12] Benjamin Rossman. Homomorphism preservation theorems. *Journal of the ACM*, 55(3):15, 2008.
- [13] Benjamin Rossman. Correlation bounds against monotone NC^1 . In *30th Conference on Computational Complexity*, volume 33 of *LIPIcs*, pages 392–411, 2015.
- [14] Benjamin Rossman. An improved homomorphism preservation theorem from lower bounds in circuit complexity. In *8th Innovations in Theoretical Computer Science*, volume 67 of *LIPIcs*, pages 27:1–17, 2017.
- [15] Benjamin Rossman. Formulas versus circuits for small distance connectivity. *SIAM Journal on Computing*, 47(5):1986–2028, 2018.
- [16] Benjamin Rossman. Lower bounds for subgraph isomorphism. In *Proc. International Congress of Mathematicians (ICM 2018, Rio de Janeiro)*, volume 3, pages 3409–3430, 2018.