

Formulas vs. Circuits for Small Distance Connectivity

Benjamin Rossman*

November 28, 2013

Abstract

We give the first super-polynomial separation in the power of bounded-depth boolean formulas vs. circuits. Specifically, we consider the problem $\text{DISTANCE } k(n) \text{ CONNECTIVITY}$, which asks whether two specified nodes in a graph of size n are connected by a path of length at most $k(n)$. This problem is solvable (by the recursive doubling technique) on **circuits** of depth $O(\log k)$ and size $O(kn^3)$. In contrast, we show that solving this problem on **formulas** of depth $\log n / (\log \log n)^{O(1)}$ requires size $n^{\Omega(\log k)}$ for all $k(n) \leq \log \log n$. As corollaries:

- (i) It follows that polynomial-size circuits for $\text{DISTANCE } k(n) \text{ CONNECTIVITY}$ require depth $\Omega(\log k)$ for all $k(n) \leq \log \log n$. This matches the upper bound from recursive doubling and improves a previous $\Omega(\log \log k)$ lower bound of Beame, Pitassi and Impagliazzo [BIP98].
- (ii) We get a tight lower bound of $s^{\Omega(d)}$ on the size required to simulate size- s depth- d circuits by depth- d formulas for all $s(n) = n^{O(1)}$ and $d(n) \leq \log \log \log n$. No lower bound better than $s^{\Omega(1)}$ was previously known for any $d(n) \not\leq O(1)$.

Our proof technique is centered on a new notion of *pathset complexity*, which roughly speaking measures the minimum cost of constructing a set of (partial) paths in a universe of size n via the operations of union and relational join, subject to certain density constraints. Half of our proof shows that bounded-depth formulas solving $\text{DISTANCE } k(n) \text{ CONNECTIVITY}$ imply upper bounds on pathset complexity. The other half is a combinatorial lower bound on pathset complexity.

*National Institute of Informatics, 2-1-2 Hitotsubashi, Chiyoda-ku, Tokyo 101-8430, Japan

Contents

1	Introduction	3
2	Our Results	6
3	Proof Overview	7
4	Preliminaries	9
5	Pathset Complexity	12
6	From Formulas to Pathset Complexity	16
6.1	Inputs of F	17
6.2	Gates of F	18
6.3	Output of F	18
6.4	Reduction to Pathset Complexity	19
7	Small Pathsets from Random Restrictions	21
7.1	Proof of Lemma 7.3	25
7.2	Proof of Lemma 7.4	26
8	Patterns	28
9	From χ to $\bar{\chi}$	30
10	Projection and Restriction	31
10.1	$\bar{\chi}$ Decreases Under Projection	32
10.2	$\bar{\chi}$ Decreases Under Restriction	32
10.3	The Operation $A \ominus B$	33
11	Lower Bound for $\bar{\chi}$	34
11.1	Definition of Φ_A	34
11.2	Showing $\bar{\chi}_A(\mathcal{A}) \geq \tilde{n}^{\Phi_A} \delta(\mathcal{A})$	35
11.3	Showing $\Phi_A \geq \frac{1}{6} \log(\ell_A) + \Delta_A$	37
12	Conclusion	41
A	Key Examples	44
B	Lower Bound for $\bar{\chi}$: Special Cases	45
C	Rectangular Pathsets	48

1 Introduction

Understanding the relative power of formulas vs. circuits is a central challenge in complexity theory, especially in the important boolean setting. Whereas boolean circuits are the most general non-uniform model of computation (in particular, boolean circuits can efficiently simulate Turing machines), there is a strong intuition that boolean formulas (= tree-like circuits with fan-out 1) are a very weak model of computation. Many natural problems solvable by small circuits, such as st-connectivity, are believed to require large formulas. However, no super-polynomial gap between the formula complexity and circuit complexity of any problem has ever been established. The existence of such a gap is a major open question.

Question 1.1. *Are polynomial-size boolean circuits strictly more powerful than polynomial-size boolean formulas?*

There are two versions of Question 1.1 for the **uniform** and **non-uniform** settings.¹ In terms of complexity classes, this is equivalent to asking whether uniform-NC^1 (resp. NC^1) is a proper subclass of P (resp. P/poly).² Both the uniform and non-uniform versions of this question are wide open:

- An obvious prerequisite of the separation of uniform-NC^1 from P is a super-polynomial lower bound on the formula complexity of *any* explicit boolean function. However, despite the fact that *almost all* boolean functions have formula complexity $\Omega(2^n/\log n)$ by a classic theorem of Riordan and Shannon [RS42], the best lower bound for any explicit function, due to Håstad [Hås98], is only $\Omega(n^{3-o(1)})$. Unfortunately, n^3 is known to be the limit of existing techniques, and it appears that any improvement will require a major breakthrough.
- The situation is no better in the non-uniform setting. By a striking theorem of Savický and Woods [SW98], for every constant $k > 1$, almost all boolean functions with formula complexity $\leq n^k$ have circuit complexity $\geq n^k/k$. This shows that NC^1 cannot be separated from P/poly by a straightforward counting argument (in contrast with results like the Circuit Size Hierarchy Theorem, see [Juk12]). Other than by counting arguments, it is not clear how to take advantage of non-uniformity.

In short, it appears that we are a long way from answering Question 1.1. In the meantime, we can hope to gain insight by studying the question of formulas vs. circuits in *restricted settings* where strong lower bounds are available. In particular, Question 1.1 has natural analogues in both the **monotone** setting and the **bounded-depth** (boolean) setting, where exponential lower bounds have been around for decades. However, as we will explain, while question of monotone formulas vs. circuits has been settled for 25 years, essentially nothing was known in bounded-depth setting prior to the results of this paper.

¹Whenever we speak of a circuit (or formula), this is understood to mean a sequence $(C_n)_{n=1}^\infty$ of circuits, one for each input size n . In the *uniform* setting, there is an underlying algorithm which, given 1^n as input, outputs a description of the circuit C_n . In the *non-uniform* setting, C_n are arbitrary. All bounds mentioned in this paper may be interpreted in the stronger sense: uniform upper bounds and non-uniform lower bounds.

²By Spira's Theorem [Spi71], NC^1 is equivalent to the class of languages recognized by polynomial-size boolean formulas (of unbounded depth).

Monotone Formulas vs. Circuits: The separation of monotone formulas from monotone circuits was shown by Karchmer and Wigderson [KW90] via a lower bound for directed st-connectivity (*STCONN*).

Theorem 1.2. *Monotone formulas solving STCONN require size $n^{\Omega(\log n)}$.*

As it was already known that *STCONN* has polynomial-size monotone circuits, Theorem 1.2 implies the separation of monotone classes mNC^1 and mP (in fact, it shows $\text{mNC}^1 \neq \text{mAC}^1$). (In a notable recent development, Potechin [Pot10] showed that monotone switching networks for *STCONN* require size $n^{\Omega(\log n)}$. This result strengthens Theorem 1.2 and implies the sharper separation $\text{mL} \neq \text{mNL}$.)

Bounded-Depth Formulas vs. Circuits: The *bounded-depth* setting refers to the class of unbounded fan-in boolean circuits and formulas of depth $\leq d(n)$ for some (not necessarily constant) function $d : \mathbb{N} \rightarrow \mathbb{N}$. Unlike the monotone setting, the question of bounded-depth formulas vs. questions gives a natural approach to Question 1.1: by comparing the power of depth- d formulas vs. depth- d circuits, we can hope to get a separation for as large a depth $d(n)$ as possible, noting that a super-polynomial separation for any $d(n) = \log n$ would imply $\text{NC}^1 \neq \text{AC}^1$ (answering Question 1.1).

We write $\text{Circuit}(s, d)$ (resp. $\text{Formula}(s, d)$) for the class of languages computable by unbounded fan-in boolean circuits (resp. formulas) of size $\leq s(n)$ and depth $\leq d(n)$. Consider the elementary fact that $\text{Circuit}(s, d) \subseteq \text{Formula}(s^d, d)$, that is, every depth- d circuit of size s is equivalent to a depth- d formula of size $\leq s^d$. In the naive simulation of circuits by formulas, we simply replace overlapping sub-circuits with non-overlapping copies until the circuit becomes a tree. Note that this give a slightly better upper bound of fan-in ^{d} . It is natural to ask: is this naive simulation of depth- d circuits by depth- d formulas asymptotically optimal? To make this question meaningful, we focus on the case where $s(n)$ is any $n^{O(1)}$ and $d(n) \leq \log n$. Thus, $\text{Circuit}(n^{O(1)}, d) \subseteq \text{Formula}(n^{O(d)}, d)$ and we can ask whether $n^{O(d)}$ can be improved to $n^{o(d)}$.

Question 1.3. *For which functions $d(n) \leq \log n$ do we have*

$$(*) \quad \text{Circuit}(n^{O(1)}, d) \not\subseteq \text{Formula}(n^{o(d)}, d)?$$

On the basis of problems like *STCONN*, we conjecture that $(*)$ holds for all $d(n) \leq \log n$. Of course, since this (more than) implies $\text{NC}^1 \neq \text{AC}^1$, we should not expect to prove $(*)$ all the way to depth $\log n$ anytime soon. On the other hand, more modest depths like $O(\log \log n)$ are well within the range of techniques like switching lemmas (after all, the super-polynomial lower bounds for parity extend to depth $o(\log n / \log \log n)$ [Hås87]). For this reason, it might seem that $(*)$ is the kind of statement that ought to be known (or follow from known results) for modest but super-constant $d(n)$. (Note that $(*)$ is trivial for constant $d(n) \leq O(1)$.) However, it turns out that the status of $(*)$ was entirely unknown for all $d(n) \not\leq O(1)$. Even the weakest possible separation $\text{Circuit}(n^{O(1)}, d) \not\subseteq \text{Formula}(n^{O(1)}, d)$ (i.e. $\text{Formula}(n^{O(1)}, d) \subsetneq \text{Circuit}(n^{O(1)}, d)$) was not known to hold for any $d(n) \not\leq O(1)$. In this paper, we improve this state of affairs by showing that $(*)$ holds for all $d(n) \leq \log \log \log n$ (Corollary 2.3).

At this point, we should ask: why do the previous techniques (in particular, switching lemmas [Hås87] and approximation by low-degree polynomials [Raz87, Smo87]) fail to distinguish formulas from circuits? In other words, why don't these techniques imply stronger lower bounds for depth- d

formula complexity vis-à-vis depth- d circuit complexity of a given boolean function? We suggest that this is the consequence of a certain kind of *bottom-up* depth-reduction argument.³ At the same time, *top-down* lower bound techniques (in particular, Karchmer-Wigderson games [KW90]) have never been successfully used in the boolean setting, in contrast to the monotone setting.⁴ Our lower bound technique gets around the limitations of previous techniques by a novel combination of bottom-up and top-down arguments. In particular, the part of our proof which distinguishes formulas from circuits is a new top-down argument (Lemma 6.7).

Distance $k(n)$ Connectivity: As with the separation of monotone formulas vs. circuits in [KW90], our separation of bounded-depth formulas vs. circuits comes by way of a lower bound for (a parameterized version of) st-connectivity. As Wigderson wrote in his excellent survey on graph connectivity [Wig92], “Of all computational problems, graph connectivity is the one that has been studied on the largest variety of computational models, such as Turing machines, PRAMs, Boolean circuits, decision trees and communication complexity. It has proven a fertile test case for comparing basic resources such as time vs. space, nondeterminism vs. randomness vs. determinism, and sequential vs. parallel computation.” There has been some significant progress in the 20 years since [Wig92]. One notable result is Reingold’s theorem [Rei08] that $USTCONN$ (undirected st-connectivity) $\in DSPACE(\log n)$. However, many questions remain open. Chief among these is the space complexity of $STCONN$. Savitch’s theorem [Sav70] that $STCONN \in DSPACE(\log^2 n)$ is still the best known upper bound.

As for lower bounds for $STCONN$, in addition to various results in monotone models of computation [KW90, Pot10, RW89, SS79, TT94], there are results on structured models of computations whose basic operations manipulate pebbings on graphs. One result of this type, due to Edmonds, Poon and Achlioptas [EPA99], gives a tight space lower bound of $\Omega(\log^2 n)$ on the NNJAG model. Another interesting result, in the unusual restricted model of arithmetic circuits with \times gates of odd fan-in, is a tight lower bound of $n^{\Omega(\log n)}$ for $STCONN$ (or more accurately its algebraic cousin, iterated matrix multiplication) was shown by Nisan and Wigderson [NW96] using the method of partial derivatives.

In this paper, we consider a version of $STCONN$ parameterized by distance. For a function $k : \mathbb{N} \rightarrow \mathbb{N}$ with $k(n) \leq n$, *distance $k(n)$ connectivity*, denoted $STCONN(k(n))$, is the following problem: given a directed graph with n vertices and specified vertices s and t , determine whether or not there is a path of length at most $k(n)$ from s to t . Unlike $STCONN$ and $USTCONN$, the

³ This style of lower bound has the following elements:

- For some notion of “simple” functions and some notion of “approximation”, there is a lemma of the form: if g is the AND or OR of simple functions $f_1, \dots, f_{\text{poly}(n)}$, then g is approximated by a (slightly less) simple function. (For example: if f_i are small decision trees, then after a random restriction, g simplifies to a small decision tree with high probability; if f_i are low-degree polynomials, then g agrees with a low-degree polynomial up to small error.)
- Given a polynomial-size depth- d circuit C sitting on top of a layer of input variables (themselves simple functions), the approximation lemma is applied *independently* to all gates at the bottom level (directly above the inputs). C is thus transformed into a circuit C' of depth $d - 1$ sitting on top of a layer of simple functions. This depth-reduction step is repeated d times. Finally, we have a function which is sufficiently simple to imply a lower bound for whatever (non-simple) function one has in mind.

It is precisely because the approximation lemma is applied *independently* to all bottom-level gates that the distinct between circuits and formulas is lost.

⁴One exception is a top-down lower bound for depth-3 circuits due to Jukna, Pudlák and Håstad [HJP95], who pose the problem of proving *any* super-polynomial lower bound at depth 4 by a purely top-down argument.

directed and undirected versions of distance $k(n)$ connectivity are essentially equivalent.⁵ It is easy to show that $STCONN(k(n))$ has circuits (moreover, semi-unbounded monotone circuits) of size $O(kn^3)$ and depth $2 \log k$ using the recursive doubling (a.k.a. repeated squared) method of Savitch [Sav70]. (At the expense of larger depth, one can get smaller circuits of size $O(kn^{2.38})$ using fast matrix multiplication.)

An important relationship between $STCONN$ and its parameterized version $STCONN(k(n))$ is the fact every algorithm for $STCONN(k(n))$ “scales up” to an algorithm for $STCONN$ by recursive k th powering. (Conversely, every lower bound for $STCONN$ “scales down” to a lower bound for $STCONN(k(n))$. In particular, Theorem 1.2 implies that monotone formulas solving $STCONN(k(n))$ require size $n^{\Omega(\log k)}$.) For circuits, we have the implication:

$$STCONN(k(n)) \in \text{Circuit}(s, d) \implies STCONN \in \text{Circuit}(n^{O(1)} \cdot s, \frac{\log n}{\log k} \cdot d).$$

As noted in [Wig92], if $STCONN(k(n))$ has polynomial-size circuits of depth $o(\log k)$, then $STCONN$ has polynomial-size circuits of depth $o(\log n)$ and hence $STCONN \in \text{DSPACE}(o(\log^2 n))$. This observation strongly motivates the following:

Question 1.4. *What is the minimum depth of polynomial-size circuits solving $STCONN(k(n))$?*

Furst, Saxe and Sipser [FSS84] showed that $STCONN \notin \text{AC}^0$ via the reduction from parity to $STCONN$. Via the same reduction, it follows from the parity lower bound of Håstad [Hås87] that $STCONN(k(n)) \notin \text{AC}^0$ for all $k(n) \not\leq \log^{O(1)} n$. However, this says nothing when $k(n) \leq \log^{O(1)} n$.

Ajtai [Ajt89] proved the first lower bound for small distances $k(n)$, showing that $STCONN(k(n)) \notin \text{AC}^0$ for all $k(n) \not\leq O(1)$. Via an explicit version of Ajtai’s originally non-constructive proof, Bellantoni, Pitassi and Urquhart [BPU92] proved a lower bound of $\Omega(\log^* k)$ on the depth of polynomial-size circuits solving $STCONN(k(n))$. This was subsequently improved to $\Omega(\log \log k)$ for all $k(n) \leq \log^{O(1)} n$ by Beame, Impagliazzo and Pitassi [BIP98], using a special-purpose “connectivity switching lemma” tailored to $STCONN(k(n))$. It was left as an open problem to further narrow the gap between the $O(\log k)$ and $\Omega(\log \log k)$ upper and lower bounds. In this paper, we completely close this gap by proving a lower bound of $\Omega(\log k)$ for all $k(n) \leq \log \log n$ (Corollary 2.2). (While our current proof is restricted to $k(n) \leq \log \log n$, we believe this can be extended $k(n) \leq \log^{O(1)} \log n$ as in [BIP98].) The significance of this result is that, for small but super-constant $k(n)$, we rule out the possibility of showing that $STCONN \in \text{DSPACE}(o(\log^2 n))$ by constructing polynomial-size circuits for $STCONN(k(n))$ of depth $o(\log k)$.

2 Our Results

Our main theorem is a tight lower bound on the size of bounded-depth formulas solving distance $k(n)$ connectivity.

Theorem 2.1 (Main Result). *Formulas of depth $\log n / (\log \log n)^{O(1)}$ solving $STCONN(k(n))$ have size $n^{\Omega(\log k)}$ for all $k(n) \leq \log \log n$.*

⁵The reduction from $STCONN(k(n))$ to $USTCONN(k'(n'))$ converts a directed graph on n vertices into a layered undirected graph on $n' = (k+1)n$ vertices where $k'(n') = k(n)$.

To be precise, we get a lower bound of $n^{(1/6)\log k - O(1)}$ for formulas of depth $c \log n / \max\{k^3 \log^2 k, k \log k \log \log n\}$ where $c > 0$ is an absolute constant. Moreover, this lower bound is not only worst-case: it applies to formulas solving $STCONN(k(n))$ in the natural *average-case* sense (see §12).

The following two corollaries of Theorem 2.1 were already mentioned in the introduction. As discussed, these corollaries answer Questions 1.3 and 1.4 for a limited range of $d(n)$ and $k(n)$.

Corollary 2.2. *Polynomial-size circuits solving $STCONN(k(n))$ require depth $\Omega(\log k)$ for all $k(n) \leq \log \log n$.*

Proof. For contradiction, assume C is a circuit of size $s(n) = n^{O(1)}$ and depth $d(n) = o(\log k)$ solving $STCONN(k(n))$ for some $k(n) \leq \log \log n$. By the naive simulation of circuits by formulas, C is equivalent to a depth- d formula of size at most $s^d = n^{o(\log k)}$. But since $d(n) = o(\log \log \log n) \ll \log n / (\log \log n)^{O(1)}$, we get contradiction with Theorem 2.1. \square

Corollary 2.3. *It is impossible to simulate polynomial-size depth- d circuits by depth- d formulas of size $n^{o(d)}$ (that is, we get the optimal separation $\text{Circuit}(n^{O(1)}, d) \not\subseteq \text{Formula}(n^{o(d)}, d)$) for all $s(n) = n^{O(1)}$ and $d(n) \leq \log \log \log n$.*

Proof. The separating language is $STCONN(k(n))$ where $k(n) = 2^{d(n)/2} (\leq \log \log n)$. We have $STCONN(k(n)) \in \text{Circuit}(n^{O(1)}, d)$ by the circuits (of depth $2 \log k = d$) which implement recursive doubling. The lower bound $STCONN(k(n)) \notin \text{Formula}(n^{o(d)}, d)$ is by Theorem 2.1, noting that $d(n) \leq \log \log \log n \ll \log n / (\log \log n)^{O(1)}$. \square

3 Proof Overview

Our proof technique is centered on a new notion of *pathset complexity*. Informally, a *pathset* is a subset $\mathcal{A} \subseteq [n]^{k+1}$ whose elements represent potential paths of length k in a graph of size n . The *pathset complexity* of \mathcal{A} , denoted $\chi(\mathcal{A})$, measures of the minimum number of operations required to construct \mathcal{A} via unions (\cup) and relational join (\bowtie), subject to certain density constraints. (The formal definition of $\chi(\mathcal{A})$, given in §5, is not important for this overview.)

The proof of Theorem 2.1 has two parts. Part 1 shows that every bounded-depth formula F solving $STCONN(k(n))$ implies an upper bound on the pathset complexity of a certain (random) pathset \mathcal{A}^Γ . Part 2 is a general lower bound on $\chi(\mathcal{A})$ for arbitrary pathsets \mathcal{A} . Combining these two parts, we get the desired $n^{\Omega(\log k)}$ lower bound on the size of F .

Before explaining Parts 1 and 2 in more detail, we state the key property of $STCONN(k(n))$ which our proof exploits. Instances for $STCONN(k(n))$ are directed graphs with vertex set $[n]$ and distinguished vertices s and t (without loss of generality $s = 1$ and $t = 2$). An *st-path* is a sequence $(x_0, \dots, x_k) \in [n]^{k+1}$ such that $x_0 = s$ and $x_k = t$ and $x_i \neq x_j$ for all $i \neq j$.

Denote by Γ the random directed graph with edge probability $1/n$. (Note that $1/n$ is below the threshold for $STCONN(k(n))$, that is, almost surely Γ contains no *st-path* of length k .) Define \mathcal{A}^Γ as the set of *st-paths* $(x_0, \dots, x_k) \in [n]^{k+1}$ such that

- $(x_0, x_1), \dots, (x_{k-1}, x_k)$ are non-edges of Γ ,
- $\Gamma \cup \{(x_0, x_1), \dots, (x_{k-1}, x_k)\}$ contains a unique *st-path* of length k (namely, (x_0, \dots, x_k)).

Then the (average-case) property of $STCONN(k(n))$ that our proof exploits is:

Key Property (§6.3): *Almost surely, \mathcal{A}^Γ contains 99% of st-paths of length k .*

We now state Parts 1 and 2 of the proof of Theorem 2.1 in more detail.

Part 1 (§6–7): *Suppose F is a formula of depth $\log n / (\log \log n)^{O(1)}$ solving $\text{STCONN}(k(n))$. Then, almost surely (with respect to Γ),*

$$(1) \quad \text{size}(F) \geq 2^{-O(k^2)} \cdot n^{-O(1)} \cdot \chi(\mathcal{A}^\Gamma).$$

Part 2 (§8–11): *For all pathsets $\mathcal{A} \subseteq [n]^{k+1}$, writing $\delta(\mathcal{A}) := |\mathcal{A}|/n^{k+1}$ for the density of \mathcal{A} ,*

$$(2) \quad \chi(\mathcal{A}) \geq 2^{-O(2^k)} \cdot n^{\Omega(\log k)} \cdot \delta(\mathcal{A}).$$

Combining (1) and (2) with $\delta(\mathcal{A}^\Gamma) \geq .99n^{-2}$ (by the key property), we get the lower bound $\text{size}(F) \geq 2^{-O(2^k)} \cdot n^{\Omega(\log k)}$. Since $2^{-O(2^k)}$ is $n^{-O(1)}$ for $k(n) \leq \log \log n$, Theorem 2.1 is proved.

Part 1 builds on the technique of [Ros08, Ros10]. An essential new ingredient, which distinguishes formulas from circuits, is a top-down argument (Lemma 6.7) relating formula size to pathset complexity.

For Part 2, we develop a combinatorial framework for studying pathset complexity. This involves analyzing the *pattern* of joins which predominates the construction of a given pathset \mathcal{A} . In §8 we define an auxiliary notion of *pathset complexity with respect to a pattern*, denoted $\bar{\chi}(\mathcal{A})$. Part 2 then consists of 2a and 2b:

Part 2a (§9): *For every pathset \mathcal{A} , there exists $\mathcal{A}' \subseteq \mathcal{A}$ such that $\chi(\mathcal{A}) \geq \bar{\chi}(\mathcal{A}')$ and $\delta(\mathcal{A}') \geq 2^{-O(2^k)} \cdot \delta(\mathcal{A})$.*

Part 2b (§11): *For all pathsets \mathcal{A}' , $\bar{\chi}(\mathcal{A}') \geq n^{\Omega(\log k)} \cdot \delta(\mathcal{A}')$.*

Part 2a is relatively straightforward. This move from χ to $\bar{\chi}$ is precisely where we lose the factor of $2^{O(2^k)}$, which is the reason that our main theorem is limited to $k(n) \leq \log \log n$. (If this factor can be removed, which I believe is possible (with a lot more work) within the current framework, then Theorem 2.1 and Corollary 2.2 would hold up to $k(n) \leq \log^{1/3} n$ and Corollary 2.3 would hold up to $d(n) \leq \log \log n$.)

Part 2b is the true combinatorial lower bound at the heart the paper. The proof involves an intricate induction on patterns.

Organization of the Paper: Section 4 sets out the basic terminology and notation for the paper. Section 5 introduces the key notion of *pathset complexity*. Sections 6–7 contain Part 1 of the proof of Theorem 2.1. Sections 8–11 contain Part 2 of the proof. We state some conclusions and discuss future directions in Section 12. Three appendices (Sections A–C) contain supplementary material including key examples and relatively easier special cases of our main lower bound.

4 Preliminaries

Let n be an arbitrary positive integer (which we view as growing to infinity). Let $[n] := \{1, \dots, n\}$. We note that, for all purposes in this paper, $[n]$ may be regarded as an arbitrary fixed set of size n . Let $k = k(n)$ and $d = d(n)$ be arbitrary functions of n . As parameters, k represents *distance* and d represents *depth*. No bound on k or d is assumed throughout the paper; assumptions like $k(n) \leq \log \log n$ are explicitly stated where needed. All constants in asymptotic notation ($O(\cdot)$, etc.) are universal (with no dependence on n, k, d).

Circuits and Formulas: The *circuits* and *formulas* considered in this paper are unbounded fan-in boolean circuits and formulas with a single output node and NOT gates at the bottom level. Formally, a *circuit* is a finite acyclic directed graph with a unique output (node of out-degree 0) where each input (node of in-degree 0) is labeled by a literal (i.e. X_i or \overline{X}_i) and each gate (node of in-degree ≥ 1) is labeled by AND or OR. A *formula* is a tree-like circuit in which every node other than the output has out-degree 1. The *size* of a circuit is the number of gates, while the *size* of a formula is the number of leaves. (For a formula F , the circuit-size of F equals the formula-size of F minus 1.)

Graphs: All *graphs* in this paper are directed graph $G = (V_G, E_G)$ where V_G is a (possibly empty) set and $E_G \subseteq V_G \times V_G$. The edge from v to w is written simply as vw to cut down on unnecessary parentheses.

Two important graphs in this paper are P_k (the directed path of length k) and $P_{k,n}$ (the “complete k -layered graph” with $k + 1$ layers of n vertices and kn^2 edges). Formally, let

$$P_k = (V_k, E_k) \text{ where } V_k = \{v_0, \dots, v_k\} \text{ and } E_k = \{v_i v_{i+1} : 0 \leq i < k\}$$

where v_0, \dots, v_k are fixed abstract vertices. We will usually omit subscripts writing simply v and vw for arbitrary elements of V_k and E_k . To define $P_{k,n}$, we create $(k + 1)n$ fresh vertices denoted v^i for each $v \in V_k$ and $i \in [n]$. Then

$$P_{k,n} = (V_{k,n}, E_{k,n}) \text{ where } V_{k,n} = \{v^i : v \in V_k, i \in [n]\} \text{ and } E_{k,n} = \{v^i w^j : vw \in E_k, i, j \in [n]\}.$$

We refer to subgraphs $\Gamma \subseteq P_{k,n}$ with $V_\Gamma = V_{k,n}$ as *k-layered graphs*. Throughout the paper, Γ consistently represents a (random) k -layered graph, while G, H, K are reserved for subgraphs of P_k . We sometimes view Γ as the input to a circuit or formula; in this case, we identify the set of layered graphs with $\{0, 1\}^N$ where N is a set of kn^2 variables indexed by elements of $E_{k,n}$.

Layered Distance $k(n)$ Connectivity: As with previous lower bounds for distance $k(n)$ connectivity [Ajt89, BIP98], we consider a variant of the problem on k -layered graphs. Let s, t denote vertices v_0^1, v_k^1 respectively. *Layered distance $k(n)$ connectivity* is the problem of determining whether a layered graph $\Gamma \in \{0, 1\}^N$ contains a path from s to t . Following [BIP98], we denote this problem by $DISTCONN(k, n)$. The layered and unlayered versions of distance $k(n)$ connectivity are essentially equivalent.⁶ This allows us to restate Theorem 2.1 as a lower bound on $DISTCONN(k, n)$:

⁶Since k -layered graphs are graphs with $(k + 1)n$ vertices, there is a trivial reduction from $DISTCONN(k, n)$ to $STCONN(k'(n'))$ where $n' = (k + 1)n$ and $k'(n') = k$. In the opposite direction, there is a simple reduction from $STCONN(k(n))$ to $DISTCONN(k, n)$ which converts graphs to k -layered graphs.

Theorem 2.1. (restated) *Solving DISTCONN(k, n) on formulas of depth $\frac{\log n}{k^3 \log \log n}$ requires size $n^{(1/6) \log k - O(1)}$ for all $k(n) \leq \log \log n$.*

This restatement includes explicit expressions $n^{(1/6) \log k - O(1)}$ for $n^{\Omega(\log k)}$ and $\frac{\log n}{k^3 \log \log n}$ for $\log n / (\log \log n)^{O(1)}$.⁷

Boolean Functions and Restrictions: Let $f : \{0, 1\}^I \rightarrow \{0, 1\}$ be a boolean function where I is an arbitrary finite set (of “variables”). We say that a variable $i \in I$ is *live* with respect to f if there exists $x \in \{0, 1\}^I$ such that $f(x) \neq f(x')$ where x' equals x with its i th coordinate flipped. Let $\text{Live}(f) := \{i \in I : i \text{ is live w.r.t. } f\}$.

A *restriction* on I is any function $\theta : I \rightarrow \{0, 1, *\}$. We denote by $f[\theta : \{0, 1\}^{\theta^{-1}(*)} \rightarrow \{0, 1\}$ the function (over the “unrestricted” variables i such that $\theta(i) = *$) obtained from f by applying the restriction θ .

Probabilistic Notation: For a finite set I and $0 \leq p, q \leq 1$, we write:

- $x \in \{0, 1\}_p^I$ for the random tuple $x \in \{0, 1\}^I$ where $\mathbb{P}[x_i = 1] = p$ independently for all $i \in I$ (in particular, we will consider the random layered graph $\Gamma \in \{0, 1\}_{1/n}^N$),
- $R \subseteq_p I$ for the random subset R of I where $i \in R$ independently with probability p for all $i \in I$,
- $\theta \in \mathcal{R}(p, q)$ for the random restriction $\theta : I \rightarrow \{0, 1, *\}$ where $\mathbb{P}[\theta(i) = *] = q$ and $\mathbb{P}[\theta(i) = 1] = (1 - q)p$ for all $i \in I$.

Whenever we say *almost surely*, this is understood to mean *asymptotically almost surely* as $n \rightarrow \infty$ (i.e. with probability that goes to 1 as $n \rightarrow \infty$).

Tuples and Relations: The following notation pertains to “ V -ary” tuples $x \in [n]^V$ and relations $\mathcal{A} \subseteq [n]^V$ where V is an arbitrary finite set.

Definition 4.1 (V -tuples). For $x \in [n]^V$ and $S \subseteq V$, we denote by $x_S \in [n]^S$ the restriction of x to coordinates in S . For $x \in [n]^V$ and $y \in [n]^W$ where $V \cap W = \emptyset$, let $xy \in [n]^{V \cup W}$ denote the unique $z \in [n]^{V \cup W}$ such that $z_i = x_i$ for all $i \in V$ and $z_j = y_j$ for all $j \in W$; here $xy = yx$, as there is no intrinsic linear order on $V \cup W$. We adopt the convention $[n]^\emptyset = \{()\}$ where $()$ denotes the unique \emptyset -tuple.

Definition 4.2 (Join). For finite sets V and W and $\mathcal{A} \subseteq [n]^V$ and $\mathcal{B} \subseteq [n]^W$, the *join* of \mathcal{A} and \mathcal{B} is the set

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

The join operation \bowtie is a hybrid of intersection \cap and cartesian product \times : if $V = W$ then $\mathcal{A} \bowtie \mathcal{B} = \mathcal{A} \cap \mathcal{B}$, and if $V \cap W = \emptyset$ then $\mathcal{A} \bowtie \mathcal{B}$ is the product $\mathcal{A} \times \mathcal{B}$. Note that $\mathcal{A} \bowtie \emptyset = \emptyset$ and $\mathcal{A} \bowtie \{()\} = \mathcal{A}$.

⁷As mentioned earlier, our proof actually extends to depth $O(\log n / \max\{k^3 \log^2 k, k \log k \log \log n\})$; in particular, this is $O(\frac{\log n}{k \log k \log \log n})$ for very small $k(n) \leq \log^{1/3} \log n$. We state Theorem 2.1 with depth $\frac{\log n}{k^3 \log \log n}$ for the sake of simplicity.

Definition 4.3 (Density, Projection, Restriction). Let $\mathcal{A} \subseteq [n]^V$.

(i) The *density* of \mathcal{A} is defined by $\delta(\mathcal{A}) := |\mathcal{A}| / n^{|V|}$.

(ii) For $S \subseteq V$, the *S-projection* and *S-projection density* of \mathcal{A} are defined by

$$\text{proj}_S(\mathcal{A}) := \{x_S : x \in \mathcal{A}\}, \quad \pi_S(\mathcal{A}) := \delta(\text{proj}_S(\mathcal{A})).$$

That is, $\pi_S(\mathcal{A}) = |\text{proj}_S(\mathcal{A})| / n^{|S|}$, as δ here refers to the density of the S -ary relation $\text{proj}_S(\mathcal{A}) \subseteq [n]^S$.

(iii) For $S \subseteq V$ and $z \in [n]^{V \setminus S}$, the *S-restriction* of \mathcal{A} at z and *maximum S-restriction density* of \mathcal{A} are defined by

$$\mathcal{A}|_S^z := \{y \in [n]^S : yz \in \mathcal{A}\}, \quad \mu_S(\mathcal{A}) := \max_{z \in [n]^{V \setminus S}} \delta(\mathcal{A}|_S^z).$$

It will be convenient (later on in §10) to extend this notation as follows: for any sets S and \bar{S} such that $S \cap \bar{S} = \emptyset$ and $V \subseteq S \cup \bar{S}$ and any $z \in [n]^{\bar{S}}$, let $\mathcal{A}|_S^z$ be understood as $\mathcal{A}|_{V \cap S}^{z'}$ where $z' = z_{V \cap \bar{S}}$.

We conclude this section with a lemma which gives some basic inequalities relating the densities of projections, restrictions and joins. In particular, inequality (c), bounding the density of a join, will play a crucial role later on.

Lemma 4.4. For all $\mathcal{A} \subseteq [n]^V$ and $\mathcal{B} \subseteq [n]^W$ and $S^- \subseteq S \subseteq S^+ \subseteq V$ and $T \subseteq W$,

- (a) $\mu_{S^+}(\mathcal{A}) \leq \mu_S(\mathcal{A}) \leq \mu_S(\text{proj}_{S^+}(\mathcal{A})) \leq \pi_S(\mathcal{A}) \leq \pi_{S^-}(\mathcal{A})$,
- (b) $\delta(\mathcal{A}) \leq \pi_S(\mathcal{A}) \cdot \mu_{V \setminus S}(\mathcal{A})$,
- (c) $\delta(\mathcal{A} \bowtie \mathcal{B}) \leq \pi_S(\mathcal{A}) \cdot \mu_{T \setminus S}(\text{proj}_T(\mathcal{B})) \cdot \mu_{(V \cup W) \setminus (S \cup T)}(\mathcal{A} \bowtie \mathcal{B})$.

Proof. Inequalities (a) and (b) are obvious once the notation is understood. Inequality (c) is mainly derived by two applications of inequality (b). We first project $\mathcal{A} \bowtie \mathcal{B}$ to $S \cup T$:

$$\delta(\mathcal{A} \bowtie \mathcal{B}) \leq \pi_{S \cup T}(\mathcal{A} \bowtie \mathcal{B}) \cdot \mu_{(V \cup W) \setminus (S \cup T)}(\mathcal{A} \bowtie \mathcal{B}).$$

We then project $\text{proj}_{S \cup T}(\mathcal{A} \bowtie \mathcal{B})$ to S :

$$\pi_{S \cup T}(\mathcal{A} \bowtie \mathcal{B}) \leq \pi_S(\mathcal{A} \bowtie \mathcal{B}) \cdot \mu_{T \setminus S}(\text{proj}_{S \cup T}(\mathcal{A} \bowtie \mathcal{B})).$$

Finally, we have $\pi_S(\mathcal{A} \bowtie \mathcal{B}) \leq \pi_S(\mathcal{A})$ and $\mu_{T \setminus S}(\text{proj}_{S \cup T}(\mathcal{A} \bowtie \mathcal{B})) \leq \mu_{T \setminus S}(\text{proj}_T(\mathcal{A} \bowtie \mathcal{B})) \leq \mu_{T \setminus S}(\text{proj}_T(\mathcal{B}))$. Combined, these inequalities give (c). \square

5 Pathset Complexity

In this section, we define the key notion of pathset complexity, state our lower bound for pathset complexity (Theorem 5.8, to be proved in §8–11), and present a matching upper bound (Proposition 5.11).

Definition 5.1 (Pattern Graph). Recall that $P_k = (V_k, E_k)$ is the directed path of length k where $V_k = \{v_i : 0 \leq i \leq k\}$ and $E_k = \{v_i v_{i+1} : 0 \leq i < k\}$. A *pattern graph* is a subgraph of P_k with no isolated vertices. That is, $G = (V_G, E_G)$ is a pattern graph if, and only if, $E_G \subseteq E_k$ and $V_G = \bigcup_{vw \in E_G} \{v, w\}$. We write \wp_k for the set of pattern graphs. (We appropriate the power set notation, since pattern graphs are in 1-1 correspondence with subsets of E_k .)

Note that every pattern graph is a (possibly empty) disjoint union of directed paths of length ≥ 1 . We refer to maximal connected subsets of V_G simply as *components* of G . Two important parameters of pattern graphs are the number of components (= the number of maximal paths) and the length of the longest path (= the number of edges in the largest component). These are denoted by

$$\begin{aligned} \Delta_G &:= \# \text{ of components in } G (= |V_G| - |E_G|), \\ \ell_G &:= \text{length of the longest path in } G. \end{aligned}$$

Definition 5.2 (Pathset). For a pattern graph G , let \mathcal{P}_G denote the power set of $[n]^{V_G}$. We refer to elements of \mathcal{P}_G as *G-pathsets* (or just *pathsets* if G is clear from context).

The intuition for pathsets is as follows. For a pattern graph G , we view each $x \in [n]^{V_G}$ as corresponding to a “lifting” of G inside the complete layered graph $P_{k,n}$, namely isomorphic copy of G with vertex set $\{v^i \in V_{k,n} : i = x_v\}$ and edge set $\{v^i w^j \in E_{k,n} : i = x_v \text{ and } j = x_w\}$. In this view, a pathset $\mathcal{A} \subseteq [n]^{V_G}$ corresponds to a set of liftings of G . I have chosen to define *pathset* as a relation (a subset of $[n]^{V_G}$) rather than a set of liftings (which better matches intuition) in order to more naturally apply operations like \bowtie and proj_S and μ_S , etc.

Definition 5.3 (G -small).

- (i) Let $\varepsilon := 1/\log k$ and $\tilde{n} := n^{1-\varepsilon}$.
- (ii) A pathset $\mathcal{A} \in \mathcal{P}_G$ is *G-small* (we simply say *small* when G is understood from context) if, for all $1 \leq t \leq \Delta_G$ and $S \subseteq V_G$ such that S is the union of t components of G , \mathcal{A} satisfies the density constraint

$$\mu_S(\mathcal{A}) \leq \tilde{n}^{-t}, \quad \text{that is,} \quad \frac{|\{x \in \mathcal{A} : x_{V_G \setminus S} = y\}|}{n^{|S|}} \leq \tilde{n}^{-t} \text{ for all } y \in [n]^{V_G \setminus S}.$$

- (iii) The set of G -small pathsets is denoted $\mathcal{P}_G^{\text{small}}$.

A few quick remarks:

- As the terminology suggests, G -smallness is a monotone decreasing property (i.e. if \mathcal{A} is G -small, then so is every $\mathcal{A}' \subseteq \mathcal{A}$).

- G -smallness consists of $2^{\Delta_G} - 1$ density constraints on \mathcal{A} , corresponding to the nonempty unions of the Δ_G components of G . Note that for $t = \Delta_G$ and $S = V_G$, the constraint $\mu_S(\mathcal{A}) \leq \tilde{n}^{-t}$ is equivalent to $\delta(\mathcal{A}) \leq \tilde{n}^{-\Delta_G}$. In the special case that G is connected (i.e. $\Delta_G = 1$), \mathcal{A} is G -small $\iff \delta(\mathcal{A}) \leq \tilde{n}^{-1}$.
- The precise value of ε is not important: any ε between $1/k$ and $1/2$ would suit our purposes, modulo a slight weakening in the parameters of our main theorem.⁸

Example 5.4. Let G be the pattern graph with components $U = \{v_1, v_2, v_3\}$ and $U' = \{v_5, v_6\}$ (i.e. $V_G = \{v_1, v_2, v_3, v_5, v_6\}$ and $E_G = \{v_1v_2, v_2v_3, v_5v_6\}$). A pattern $\mathcal{A} \in \mathcal{P}_G$ is G -small if, and only if,

$$\delta(\mathcal{A}) \leq \tilde{n}^{-2}, \quad \mu_U(\mathcal{A}) \leq \tilde{n}^{-1}, \quad \mu_{U'}(\mathcal{A}) \leq \tilde{n}^{-1}.$$

For example, the pathset $\mathcal{A}_1 := \{x : x_1 = x_5 = 1\}$ is G -small (here x ranges over $[n]^{V_G}$ and we write x_i for x_{v_i}) since $\delta(\mathcal{A}_1) = n^{-2} < \tilde{n}^{-2}$ and $\mu_U(\mathcal{A}_1) = \mu_{U'}(\mathcal{A}_1) = n^{-1} < \tilde{n}^{-1}$. The pathset $\mathcal{A}_2 := \{x : x_1 = x_5 \text{ and } x_2 = x_6\}$ is G -small as well since $\delta(\mathcal{A}_2) = \mu_U(\mathcal{A}_2) = \mu_{U'}(\mathcal{A}_2) = n^{-2}$. However, pathsets

$$\mathcal{A}_3 := \{x : x_1 = x_2 = 1\}, \quad \mathcal{A}_4 := \{x : x_1 = x_5\}$$

are not G -small since $\mu_{U'}(\mathcal{A}_3) = 1 > \tilde{n}^{-1}$ and $\delta(\mathcal{A}_4) = n^{-1} > \tilde{n}^{-2}$.

The next lemma shows that smallness is preserved under joins. (Note to the reader: Although it natural to state Lemma 5.5 now, we will not use this lemma until §11.)

Lemma 5.5. *If \mathcal{A} is a small G -pathset and \mathcal{B} is a small H -pathset, then $\mathcal{A} \bowtie \mathcal{B}$ is a small $G \cup H$ -pathset.*

Proof. Assume \mathcal{A} is a small G -pathset and \mathcal{B} is a small H -pathset. To show that $\mathcal{A} \bowtie \mathcal{B}$ is a small $G \cup H$ -pathset, consider any $1 \leq t \leq \Delta_{G \cup H}$ and $S \subseteq V_G \cup V_H$ such that S contains t distinct components U_1, \dots, U_t of $G \cup H$. We must show that $\mu_S(\mathcal{A} \bowtie \mathcal{B}) \leq \tilde{n}^{-t}$.

Without loss of generality, assume U_1, \dots, U_t are ordered such that, for some $t' \leq t$, we have $U_i \cap V_G \neq \emptyset$ for all $1 \leq i \leq t'$ and $U_j \cap V_G = \emptyset$ for all $t' < j \leq t$. Let $S' = S \cap V_G$ and $S'' = U_{t'+1} \cup \dots \cup U_t$. Then S' contains $\geq t'$ components of G , since $U_i \cap V_G$ contains ≥ 1 component of G for all $1 \leq i \leq t'$. Next note that U_j is a component of H for all $t' < j \leq t$, hence S'' is a union of $t - t'$ components of H . By G -smallness of \mathcal{A} and H -smallness of \mathcal{B} , it follows that

$$\mu_{S'}(\mathcal{A}) \leq \tilde{n}^{-t'} \quad \text{and} \quad \mu_{S''}(\mathcal{B}) \leq \tilde{n}^{t'-t}.$$

Now fix $z \in [n]^{(V_G \cup V_H) \setminus S}$ which maximizes $\delta((\mathcal{A} \bowtie \mathcal{B})|_S^z)$. Using the basic properties of restrictions and joins (Lemma 4.4(a,b)), we have

$$\begin{aligned} \mu_S(\mathcal{A} \bowtie \mathcal{B}) &= \delta((\mathcal{A} \bowtie \mathcal{B})|_S^z) = \delta((\mathcal{A}|_{S'}^z) \bowtie (\mathcal{B}|_{S''}^z)) \\ &\leq \delta(\mathcal{A}|_{S'}^z) \cdot \mu_{S \setminus V_G}(\mathcal{B}|_{S''}^z) \\ &\leq \mu_{S'}(\mathcal{A}) \cdot \mu_{S''}(\mathcal{B}). \end{aligned}$$

It follows that $\mu_S(\mathcal{A} \bowtie \mathcal{B}) \leq \tilde{n}^{-t}$, which completes the proof. \square

⁸We choose $\varepsilon = 1/\log k$ so that $\tilde{n}^{\Omega(\log k)} = n^{\Omega(\log k)}$ with the same constant in the $\Omega(\log k)$. The proof of Lemma 7.5 is the only place where ε really shows up. Outside this lemma, the difference between \tilde{n} and n may be ignored (in particular, all statements in §8–11 are valid if $\tilde{n} = n$).

Definition 5.6 (Pathset Complexity). For every pattern graph G and pathset $\mathcal{A} \in \mathcal{P}_G$, the *pathset complexity* $\chi_G(\mathcal{A})$ of \mathcal{A} with respect to G is defined by the following induction:

- (i) If G is the empty graph, then $\chi_G(\mathcal{A}) := 0$.
- (ii) If G consists of a single edge, then $\chi_G(\mathcal{A}) := |\mathcal{A}|$.
- (iii) If G has ≥ 2 edges, then

$$\chi_G(\mathcal{A}) := \min_{(H_i, K_i, \mathcal{B}_i, \mathcal{C}_i)_i} \sum_i \chi_{H_i}(\mathcal{B}_i) + \chi_{K_i}(\mathcal{C}_i)$$

where $(H_i, K_i, \mathcal{B}_i, \mathcal{C}_i)_i$ ranges over sequences⁹ where

$$H_i, K_i \subset G, \quad H_i \cup K_i = G, \quad \mathcal{B}_i \in \mathcal{P}_{H_i}^{\text{small}}, \quad \mathcal{C}_i \in \mathcal{P}_{K_i}^{\text{small}} \quad \text{and} \quad \mathcal{A} \subseteq \bigcup_i \mathcal{B}_i \bowtie \mathcal{C}_i.$$

In plain language, we consider *coverings* of \mathcal{A} by *joins* of *small* pathsets over *proper* subgraphs of G . The pathset complexity $\chi_G(\mathcal{A})$ is the minimum possible value—over all such coverings—of the sum of pathset complexities of the constituent small pathsets.

Note that pathset complexity satisfies the following inequalities:

(base case)	$\chi_\emptyset(\{\emptyset\}) \leq 0$ and $\chi_G(\mathcal{A}) \leq 1$	if $ E_G = \mathcal{A} = 1$,
(monotonicity)	$\chi_G(\mathcal{A}') \leq \chi_G(\mathcal{A})$	if $\mathcal{A}' \subseteq \mathcal{A}$,
(sub-additivity)	$\chi_G(\mathcal{A}_1 \cup \mathcal{A}_2) \leq \chi_G(\mathcal{A}_1) + \chi_G(\mathcal{A}_2)$	for all $\mathcal{A}_1, \mathcal{A}_2$,
(join rule)	$\chi_{G \cup H}(\mathcal{A} \bowtie \mathcal{B}) \leq \chi_G(\mathcal{A}) + \chi_H(\mathcal{B})$	if $\mathcal{A} \in \mathcal{P}_G^{\text{small}}, \mathcal{B} \in \mathcal{P}_H^{\text{small}}$.

We will refer to these inequalities repeatedly throughout the paper.

Remark 5.7. Pathset complexity has a *dual characterization* as the unique pointwise maximal function from pairs (G, \mathcal{A}) to \mathbb{R} which satisfies (base case), (monotonicity), (sub-additivity) and (join rule). We will expand on this observation later in Remark 8.4.

We now state our lower bound on pathset complexity (to be proved in §8–11).

Theorem 5.8 (Pathset Complexity Lower Bound). *For all $\mathcal{A} \in \mathcal{P}_{P_k}$,*

$$\chi_{P_k}(\mathcal{A}) \geq \frac{n^{(1/6) \log k}}{2^{O(2^k)}} \cdot \delta(\mathcal{A}).$$

In particular, for $k \leq \log \log n$ and non-negligible $\delta(\mathcal{A}) = n^{-O(1)}$, Theorem 5.8 implies $\chi_{P_k}(\mathcal{A}) \geq n^{(1/6) \log k - O(1)}$. In a moment, we will give an upper bound (Proposition 5.11) which shows that Theorem 5.8 is tight in the regime of $k \leq \log \log n$ and non-negligible $\delta(\mathcal{A})$. First, a couple of remarks which give a different perspective on the definition of $\chi_G(\mathcal{A})$:

⁹Without loss of generality, i ranges over \mathbb{N} since $H_i = K_i = G$ and $\mathcal{B}_i = \mathcal{C}_i = \emptyset$ can occur infinitely often.

Remark 5.9 (Pathset Complexity as Construction Cost). Pathset complexity can be seen as a minimum construction cost. In this view, the goal is to construct a pathset $\mathcal{A} \in \mathcal{P}_G$ out of the fewest possible “atomic” pathsets (i.e., individual edges). The rules of construction are as follows:

- (a) A single “atomic” pathset of the form $\mathcal{A} \in \mathcal{P}_G$ where $|E_G| = |\mathcal{A}| = 1$ may be bought for unit cost.
- (b) Once a pathset \mathcal{A} has been constructed, we may freely discard elements from \mathcal{A} (i.e. replace \mathcal{A} with any smaller $\mathcal{A}' \subseteq \mathcal{A}$).
- (c) Having constructed two G -pathsets \mathcal{A} and \mathcal{A}' , we may merge \mathcal{A} and \mathcal{A}' into a single G -pathset $\mathcal{A} \cup \mathcal{A}'$ (i.e. replace \mathcal{A} and \mathcal{A}' with $\mathcal{A} \cup \mathcal{A}'$) at no additional cost.
- (d) Having constructed a G -pathset \mathcal{A} and a H -pathset \mathcal{B} , provided both \mathcal{A} and \mathcal{B} are *small*, we may join \mathcal{A} and \mathcal{B} into a single $G \cup H$ -pathset $\mathcal{A} \bowtie \mathcal{B}$ at no additional cost.

For a pathset $\mathcal{A} \in \mathcal{P}_G$, $\chi_G(\mathcal{A})$ is equal to the minimum cost of constructing \mathcal{A} according to these rules. Construction rules (a), (b), (c), (d) respectively correspond to inequalities (base case), (monotonicity), (sub-additivity), (join rule). Only applications of rule (a) increase cost (so minimum construction cost = fewest application of rule (a)). Rule (b) can be used to convert a non-small pathset into a small pathset (in order to use rule (d), for example). Note that only rule (c) can increase the density of pathsets.

Remark 5.10 (The Role of Smallness). Suppose we modify construction rule (d) by dropping the smallness constraint on \mathcal{A} and \mathcal{B} (this is equivalent to substituting \mathcal{P}_{H_i} and \mathcal{P}_{K_i} for $\mathcal{P}_{H_i}^{\text{small}}$ and $\mathcal{P}_{K_i}^{\text{small}}$ in Definition 5.6(iii)). We could then construct the complete P_k -pathset $[n]^{V_k}$ at a total cost of kn^2 simply by joining pathsets $[n]^{\{v_i, v_{i+1}\}}$ for $0 \leq i < k$. This shows that the smallness constraint on joins is essential to Theorem 5.8. Intuitively, smallness is responsible for bottlenecks which drive up the cost of constructing sufficiently dense pathsets. However, smallness is not necessarily an obstacle for very sparse pathsets like $[\sqrt{n}]^{P_k}$: since $[\sqrt{n}]^{\{v_i, v_{i+1}\}}$ are small, we can take joins showing $\chi_{P_k}([\sqrt{n}]^{P_k}) \leq kn$.

We conclude this section with an upper bound.

Proposition 5.11 (Pathset Complexity Upper Bound). *For all $\mathcal{A} \in \mathcal{P}_{P_k}$,*

$$\chi_{P_k}(\mathcal{A}) \leq O(kn^{(1/2)\lceil \log k \rceil + 2}).$$

For $k \leq \log \log n$ and $\mathcal{A} \in \mathcal{P}_{P_k}$ with $\delta(\mathcal{A}) = n^{-O(1)}$, our lower and upper bounds show that $\chi_{P_k}(\mathcal{A}) = n^{\Theta(\log k)}$ where the constant in $\Theta(\log k)$ is between $\frac{1}{6}$ and $\frac{1}{2}$.

Notation 5.12. For a pattern graph G and an integer s , we denote by $G^{\triangleright s}$ the s -shifted pattern graph with vertex set $\{v_{i+s} : v_i \in V_G\}$ and edge set $\{v_{i+s}v_{i+s+1} : v_iv_{i+1} \in E_G\}$. For a pathset $\mathcal{A} \in \mathcal{P}_G$, we denote by $\mathcal{A}^{\triangleright s} \in \mathcal{P}_{G^{\triangleright s}}$ the corresponding s -shifted pathset. Note that pathset complexity is invariant under shifts (i.e. $\chi_G(\mathcal{A}) = \chi_{G^{\triangleright s}}(\mathcal{A}^{\triangleright s})$).

Proof of Proposition 5.11. For simplicity we assume \sqrt{n} is an integer. For all $k \geq 1$, define $\mathcal{A}_k \in \mathcal{P}_{P_k}^{\text{small}}$ by

$$\mathcal{A}_k := \{x \in [n]^{\{0, \dots, k\}} : x_0, x_k \leq \sqrt{n}\}.$$

(Note that $\delta(\mathcal{A}_k) = 1/n < 1/\tilde{n}$, so \mathcal{A}_k is indeed P_k -small.)

Letting $j = \lceil k/2 \rceil$, we have

$$\mathcal{A}_j \bowtie \mathcal{A}_{k-j}^{\bowtie j} = \{x \in [n]^{\{0, \dots, k\}} : x_0, x_j, x_k \leq \sqrt{n}\}.$$

Note that \mathcal{A}_k is covered by \sqrt{n} “copies” of $\mathcal{A}_j \bowtie \mathcal{A}_{k-j}^{\bowtie j}$ where, for $1 \leq t \leq \sqrt{n}$,

$$\text{Copy}_t(\mathcal{A}_j \bowtie \mathcal{A}_{k-j}^{\bowtie j}) := \{x \in [n]^{\{0, \dots, k\}} : x_0, x_k \leq \sqrt{n} \text{ and } (t-1)\sqrt{n} < x_j \leq t\sqrt{n}\}.$$

Note that pathset complexity is invariant under “copies” in this sense (i.e. χ_G is invariant under the action of coordinate-wise permutations of $[n]$ on \mathcal{P}_G):

$$\begin{aligned} \chi_{P_k}(\text{Copy}_t(\mathcal{A}_j \bowtie \mathcal{A}_{k-j}^{\bowtie j})) &= \chi_{P_k}(\mathcal{A}_j \bowtie \mathcal{A}_{k-j}^{\bowtie j}) && \text{(invariance under “copies”)} \\ &= \chi_{P_j}(\mathcal{A}_j) + \chi_{P_{k-j}^{\bowtie j}}(\mathcal{A}_{k-j}^{\bowtie j}) && \text{(join rule)} \\ &= \chi_{P_j}(\mathcal{A}_j) + \chi_{P_{k-j}}(\mathcal{A}_{k-j}) && \text{(invariance under shifts)}. \end{aligned}$$

Since $\mathcal{A}_k \subseteq \bigcup_{1 \leq t \leq \sqrt{n}} \text{Copy}_t(\mathcal{A}_j \bowtie \mathcal{A}_{k-j}^{\bowtie j})$, sub-additivity of χ implies

$$\chi_{P_k}(\mathcal{A}_k) \leq \sum_{1 \leq t \leq \sqrt{n}} \chi_{P_k}(\text{Copy}_t(\mathcal{A}_j \bowtie \mathcal{A}_{k-j}^{\bowtie j})) = \sqrt{n} \cdot (\chi_{P_j}(\mathcal{A}_j) + \chi_{P_{k-j}}(\mathcal{A}_{k-j})).$$

This recurrence implies

$$\chi_{P_k}(\mathcal{A}_k) \leq (2\sqrt{n})^{\lceil \log k \rceil} \cdot \chi_{P_1}(\mathcal{A}_1) = O(kn^{(1/2)\lceil \log k \rceil + 1}).$$

Now note that the complete P_k -pathset $[n]^{V_k}$ is covered by n “copies” of P_k . Therefore, by a similar argument,

$$\chi_{P_k}([n]^{V_k}) \leq n \cdot \chi_{P_k}(\mathcal{A}) = O(kn^{(1/2)\lceil \log k \rceil + 2}).$$

Finally, monotonicity of χ implies that $\chi_{P_k}(\mathcal{A}) \leq O(kn^{(1/2)\lceil \log k \rceil + 2})$ for all $\mathcal{A} \in \mathcal{P}_{P_k}$. \square

6 From Formulas to Pathset Complexity

In this section we derive our main result (Theorem 2.1) from our lower bound on pathset complexity (Theorem 5.8). Let F_0 be a formula of depth $d(n)$ which solves $\text{DISTCONN}(k, n)$ where $k(n) \leq \log \log n$ and $d(n) \leq \log n / k^3 \log \log n$. We must show that F_0 has size $n^{\Omega(\log k)}$.

As a first preliminary step: without loss of generality, we assume that F_0 has minimal size among all depth $d(n)$ formulas solving $\text{DISTCONN}(k, n)$. In particular, we have $\text{size}(F_0) \leq kn^{k-1}$ since $\text{DISTCONN}(k, n)$ has DNFs of this size.

As a second preliminary step, we convert F_0 into a fan-in 2 formula F by replacing each unbounded fan-in AND/OR gate by a balanced binary tree of fan-in 2 AND/OR gates. We have

$$\text{size}(F) = \text{size}(F_0) \leq n^k \quad \text{and} \quad \text{depth}(F) \leq \text{depth}(F_0) \cdot \log(\text{size}(F_0)) \leq \log^2 n.$$

We write F_{in} for the set of inputs (i.e. leaves) in F , and F_{gate} for the set of gates in F , and f_{out} for the output gate in F . Note that each $f \in F$ is computed by an (unbounded fan-in) formula of size $\leq n^k$ and depth $\leq d(n)$ (by collapsing all adjacent AND/OR gates below f).

In order to lower bound $\text{size}(F)$ in terms of pathset complexity, we define a family of pathsets $\mathcal{A}_{f,G}^\Gamma$ associated with each $f \in F$ and $G \in \wp_k$ and $\Gamma \in \{0, 1\}^N$. Recall that we identify $\{0, 1\}^N$ with the set of k -layered graphs where $N = E_{k,n} = \{v^i w^j : vw \in E_k, i, j \in [n]\}$.

Definition 6.1 (Pathsets $\mathcal{A}_{f,G}^\Gamma$). For all $G \in \wp_k$ and $x \in [n]^{V_G}$ and $\Gamma \in \{0,1\}^N$ and $f \in F$:

- (i) Let $N_{G,x} := \{v^i w^j \in N : i = x_v \text{ and } j = x_w\}$ ($= \{v^{x_v} w^{x_w} : vw \in E_G\}$).
- (ii) Let $\rho_{G,x}^\Gamma : N \rightarrow \{0,1,*\}$ be the restriction which equals $*$ over $N_{G,x}$ and agrees with Γ over $N \setminus N_{G,x}$. In particular, applying $\rho_{G,x}^\Gamma$ to f , we get a function $f[\rho_{G,x}^\Gamma] : \{0,1\}^{N_{G,x}} \rightarrow \{0,1\}$ (whose variables correspond to edges of G via the bijection $N_{G,x} \cong E_G$).
- (iii) Let $\mathcal{A}_{f,G}^\Gamma$ be the G -pathset defined by

$$\mathcal{A}_{f,G}^\Gamma := \{x \in [n]^{V_G} : \text{Live}(f[\rho_{G,x}^\Gamma]) = N_{G,x}\}.$$

That is, $\mathcal{A}_{f,G}^\Gamma$ is the set of $x \in [n]^{V_G}$ such that the restricted function $f[\rho_{G,x}^\Gamma]$ depends on all $|N_{G,x}|$ ($= |E_G|$) of its variables.

In the next three subsections, we prove a sequence of claims about pathsets $\mathcal{A}_{f,G}^\Gamma$ in three cases where $f \in F_{\text{in}}$ and $f \in F_{\text{gate}}$ and $f = f_{\text{out}}$.

Remark 6.2. Claims 6.3, 6.4, 6.5 rely on few assumptions about F . In particular, these claims do not depend on the assumption that F_0 has bounded depth (i.e. F has bounded alternations), nor even that F is a formula as opposed to a circuit. In fact, these claims are valid if F is any B_2 -circuit computing $\text{DISTCONN}(k, n)$ where B_2 is the full binary basis.

Of course, we will eventually use both assumptions that (I) F_0 has bounded depth (i.e. F has bounded alternations), and (II) F is a formula as opposed to a circuit. Our main technical lemma (Lemma 6.6) relies on (I) but not (II) (not surprisingly, since the proof uses the Switching Lemma, which does not distinguish between circuits and formulas). A second key lemma (Lemma 6.7) relies on (II) but not (I) (using a novel top-down argument which only works for formulas).

6.1 Inputs of F

Suppose f is an input in F labeled by a literal (i.e. a variable or its negation) corresponding to some $v^i w^j \in N$. Then we have the following explicit description of $\mathcal{A}_{f,G}^\Gamma$:

- if G is the empty graph, then $\mathcal{A}_{f,G}^\Gamma = \{()\}$ (i.e. the singleton containing the 0-tuple),
- if $E_G = \{vw\}$, then $\mathcal{A}_{f,G}^\Gamma = \{x\}$ for the unique $x \in [n]^{\{v,w\}}$ with $x_v = i$ and $x_w = j$,
- otherwise (i.e. if $|E_G| \geq 2$), $\mathcal{A}_{f,G}^\Gamma = \emptyset$.

By the base case conditions (i) and (ii) in Definition 5.6 of pathset complexity, we have $\chi_\emptyset(\mathcal{A}) = 0$ and $\chi_G(\mathcal{A}) = |\mathcal{A}|$ if G has a single edge. The upshot of these observations is the following claim.

Claim 6.3 (Inputs of F). For all $f \in F_{\text{in}}$, $\sum_{G \in \wp_k} \chi_G(\mathcal{A}_{G,f}^\Gamma) = 1$.

6.2 Gates of F

Suppose f is an AND or OR gate in F with children f_1 and f_2 . Consider any $G \in \wp_k$ and $x \in \mathcal{A}_{f,G}^\Gamma$ (assuming $\mathcal{A}_{f,G}^\Gamma$ is nonempty). By definition of $\mathcal{A}_{f,G}^\Gamma$, the function $f[\rho_{G,x}^\Gamma : \{0,1\}^{N_{G,x}} \rightarrow \{0,1\}$ depends on all variables in $N_{G,x}$. Since $f[\rho_{G,x}^\Gamma$ is the AND or OR of functions $f_1[\rho_{G,x}^\Gamma$ and $f_2[\rho_{G,x}^\Gamma$, each variable in $N_{G,x}$ is a live variable for one or both $f_1[\rho_{G,x}^\Gamma$ and $f_2[\rho_{G,x}^\Gamma$.

Define sub-pattern graph $G_1 \subseteq G$ as follows: for each $vw \in E_G$, let vw be an edge in G_1 if and only if $v^{x_v}w^{x_w} \in N_{G,x}$ is a live variable for the function $f_1[\rho_{G,x}^\Gamma$. Define $G_2 \subseteq G$ in the same way with respect to f_2 . Since

$$\{v^{x_v}w^{x_w} : vw \in E_G\} = N_{G,x} = \text{Live}(f[\rho_{G,x}^\Gamma) = \text{Live}(f_1[\rho_{G,x}^\Gamma) \cup \text{Live}(f_2[\rho_{G,x}^\Gamma),$$

it follows that $G_1 \cup G_2 = G$.

Let $y = x_{V_{G_1}}$ be the restriction of x ($\in [n]^{V_G}$) to coordinates in V_{G_1} . By definition of G_1 , we have

- $v^{y_v}w^{y_w} = v^{x_v}w^{x_w} \in \text{Live}(f_1[\rho_{G,x}^\Gamma)$ for all $vw \in E_{G_1}$, and
- $v^{x_v}w^{x_w} \notin \text{Live}(f_1[\rho_{G,x}^\Gamma)$ for all $vw \in E_G \setminus E_{G_1}$.

It follows that $\text{Live}(f_1[\rho_{G_1,y}^\Gamma) = \text{Live}(f_1[\rho_{G,x}^\Gamma) = N_{G_1,y}$, hence $y \in \mathcal{A}_{f_1,G_1}^\Gamma$. Similarly, for $z = x_{V_{G_2}}$, we have $z \in \mathcal{A}_{f_2,G_2}^\Gamma$. This shows that $x \in \mathcal{A}_{f_1,G_1}^\Gamma \bowtie \mathcal{A}_{f_2,G_2}^\Gamma$.

The observation may be succinctly expressed as

$$\mathcal{A}_{f,G}^\Gamma \subseteq \bigcup_{G_1, G_2 \subseteq G : G_1 \cup G_2 = G} \mathcal{A}_{f_1,G_1}^\Gamma \bowtie \mathcal{A}_{f_2,G_2}^\Gamma.$$

Splitting this union into the cases that $G_1 = G$ or $G_2 = G$ or $G_1, G_2 \subset G$, we have proved:

Claim 6.4 (Gates of F). *For every $f \in F_{\text{gates}}$ with children f_1, f_2 and every $G \in \wp_k$,*

$$\mathcal{A}_{f,G}^\Gamma \subseteq \mathcal{A}_{f_1,G}^\Gamma \cup \mathcal{A}_{f_2,G}^\Gamma \cup \bigcup_{G_1, G_2 \subset G : G_1 \cup G_2 = G} \mathcal{A}_{f_1,G_1}^\Gamma \bowtie \mathcal{A}_{f_2,G_2}^\Gamma.$$

6.3 Output of F

We now use the fact that F computes $\text{DISTCONN}(k, n)$. Our previous Claims 6.3 and 6.4 applied to arbitrary $\Gamma \in \{0,1\}^N$. We now shift perspective and consider *random* $\Gamma \in \{0,1\}_{1/n}^N$. That is, Γ is the random k -layered graph (i.e. subgraph of $P_{k,n}$) with edge probability $1/n$. Recall that $V_{k,n} = \{v^i : v \in V_k \text{ and } i \in [n]\}$ and s, t are the vertices v_0^1, v_k^1 . Each $x \in [n]^{V_k}$ corresponds to a path of length k in $P_{k,n}$, where x is an st -path if and only if $x_0 = x_k = 0$ (writing x_i instead of x_{v_i} for the coordinates of x).

Almost surely, Γ satisfies the following properties:

- Γ contains no st -path, and
- all vertices in Γ have total degree (in-degree plus out-degree) $\leq \log^2 n$.

Both (i) and (ii) follow from simple union bounds. For (i), the number of st -paths is n^{k-1} , and each st -path only has probability n^{-k} of being in Γ . For (ii), the number of vertices is kn^2 , and the probability of any given vertex having total degree $\geq \log^2 n$ is $\leq \binom{2n}{\log^2 n} n^{-\log^2 n} \leq \left(\frac{2e}{\log^2 n}\right)^{\log^2 n} \leq n^{-\omega(1)}$.

For an st -path x , we will say that x is Γ -independent if Γ contains no path from x_i to x_j for all $0 \leq i < j \leq k$. We claim that, if Γ satisfies (i) and (ii), then 99% of st -paths are Γ -independent. To see this, consider the following greedy procedure for constructing a Γ -independent st -path. Sequentially, for $i = 1, \dots, k-1$, choose any x_i in the i th layer of $V_{k,n}$ such that Γ contains no path from s to x_i (this kills $\leq \log^{2i} n$ choices for x_i), nor a path from x_i to t (this kills $\leq \log^{2(k-i)} n$ choices), nor a path from $x_{i'}$ to x_i for any $1 \leq i' < i$ (this kills $\leq \sum_{i'=1}^{i-1} \log^{2(i-i')} n$ choices). Setting $x_0 = s$ and $x_k = t$, note that x is Γ -independent. In total we get $\geq (n - k^2 \log^{2k} n)^{k-1} \geq .99n^{k-1}$ distinct Γ -independent st -paths.

Suppose x is a Γ -independent st -path and let e_1, \dots, e_k be the k edges in x . We claim that $\Gamma \cup \{e_1, \dots, e_{i-1}, e_{i+1}, \dots, e_k\}$ contains no st -path for all $1 \leq i \leq k$. To see this, assume for the sake of contradiction that x' is an st -path in $\Gamma \cup \{e_1, \dots, e_{i-1}, e_{i+1}, \dots, e_k\}$. Let e'_1, \dots, e'_k be the edges of x' . Since e_i is a non-edge of Γ , we have $e_i \neq e'_i$. Starting at the endpoint of e'_i , we can follow the path x' forwards until reaching a vertex in x ; we can also follow x' backwards from the initial vertex of e'_i until reaching a vertex in x . This segment of x' is a path in Γ between two vertices of x , contradiction Γ -independence of x .

Since f_{out} computes $\text{DISTCONN}(k, n)$, it follows that

$$f_{\text{out}}(\Gamma \cup \{e_1, \dots, e_k\}) = 1 \quad \text{and} \quad f_{\text{out}}(\Gamma \cup \{e_1, \dots, e_{i-1}, e_{i+1}, \dots, e_k\}) = 0 \quad \text{for all } 1 \leq i \leq k.$$

This shows that the restricted function $f_{\text{out}}[\rho_{P_k, x}^\Gamma]$ depends on all k unrestricted variables (corresponding to the edges of x); in fact, $f_{\text{out}}[\rho_{P_k, x}^\Gamma]$ is the AND function. Therefore, $x \in \mathcal{A}_{f_{\text{out}}, P_k}^\Gamma$ for every Γ -independent st -path x .

By this argument, we have proved:

Claim 6.5 (Output of F). $\lim_{n \rightarrow \infty} \mathbb{P}_{\Gamma \in \{0,1\}_{1/n}^N} [\delta(\mathcal{A}_{f_{\text{out}}, P_k}^\Gamma) \geq .99n^{-2}] = 1$.

6.4 Reduction to Pathset Complexity

We now present the two main lemmas in the reduction from formula size to pathset complexity. Lemma 6.6, below, is the main technical lemma (the proof, which relies in part on the switching lemma, is given in §7). This lemma is the only place in the overall proof of Theorem 5.8 which depends on the assumption that F has bounded depth (though not on the fact that F is a formula as opposed to a circuit).

Lemma 6.6 (Pathsets $\mathcal{A}_{f,G}^\Gamma$ are Small). *Suppose $f : \{0, 1\}^N \rightarrow \{0, 1\}$ is computed by a circuit of depth $\leq \frac{\log n}{k^3 \log \log n}$ and size $\leq n^k$. Then, for all $G \in \wp_k$,*

$$\mathbb{P}_{\Gamma \in \{0,1\}_{1/n}^N} [\mathcal{A}_{f,G}^\Gamma \text{ is not } G\text{-small}] \leq O(n^{-2k}).$$

Lemma 6.7, below, is the nexus between formula size and pathset complexity. The proof involves a novel top-down argument, which is key to distinguishing formulas and circuits. (Though we will apply Lemma 6.7 to the formula F which we have been considering so far, Lemma 6.7 is stated in general terms for arbitrary boolean functions with fan-in 2.)

Lemma 6.7 (“Top-Down Lemma”). *Let F be any fan-in 2 formula and let $\Gamma \in \{0,1\}^N$. If $\mathcal{A}_{f,G}^\Gamma \in \mathcal{P}_G^{\text{small}}$ for all $f \in F$ and $G \in \wp_k$, then*

$$\chi_{P_k}(\mathcal{A}_{f_{\text{out}},P_k}^\Gamma) \leq 2^{O(k^2)} \cdot \text{depth}(F)^k \cdot \text{size}(F).$$

Proof. Assume $\mathcal{A}_{f,G}^\Gamma \in \mathcal{P}_G^{\text{small}}$ for all $f \in F$ and $G \in \wp_k$. Consider any $f \in F_{\text{gates}}$ with children f_1 and f_2 . By Claim 6.4, together with the key properties (monotonicity), (sub-additivity) and (join rule) of pathset complexity, we have

$$\begin{aligned} \chi_G(\mathcal{A}_{f,G}^\Gamma) &\leq \chi_G\left(\mathcal{A}_{f_1,G}^\Gamma \cup \mathcal{A}_{f_2,G}^\Gamma \cup \bigcup_{G_1, G_2 \subset G: G_1 \cup G_2 = G} \mathcal{A}_{f_1, G_1}^\Gamma \bowtie \mathcal{A}_{f_2, G_2}^\Gamma\right) \\ &\leq \chi_G(\mathcal{A}_{f_1,G}^\Gamma) + \chi_G(\mathcal{A}_{f_2,G}^\Gamma) + \sum_{G_1, G_2 \subset G: G_1 \cup G_2 = G} \left(\chi_{G_1}(\mathcal{A}_{f_1, G_1}^\Gamma) + \chi_{G_2}(\mathcal{A}_{f_2, G_2}^\Gamma)\right) \\ &\leq \left(\chi_G(\mathcal{A}_{f_1,G}^\Gamma) + 2^k \sum_{H \subset G} \chi_H(\mathcal{A}_{f_1, H}^\Gamma)\right) + \left(\chi_G(\mathcal{A}_{f_2,G}^\Gamma) + 2^k \sum_{H \subset G} \chi_H(\mathcal{A}_{f_2, H}^\Gamma)\right). \end{aligned}$$

If we start from $\chi_{P_k}(\mathcal{A}_{f_{\text{out}},P_k}^\Gamma)$ and repeatedly apply the above inequality until reaching the inputs of F , we get a bound of the form

$$\chi_{P_k}(\mathcal{A}_{f_{\text{out}},P_k}^\Gamma) \leq \sum_{f \in F_{\text{in}}, G \in \wp_k} c_{f,G} \cdot \chi_G(\mathcal{A}_{f,G}^\Gamma)$$

for some $c_{f,G} \in \mathbb{Z}_{\geq 0}$. We claim that

$$c_{f,G} \leq \sum_{i, H_0, \dots, H_i: P_k = H_0 \supset \dots \supset H_i = G} 2^{ik} \cdot \binom{\text{depth of } f \text{ in } F}{i} \leq 2^{O(k^2)} \cdot \text{depth}(F)^k.$$

To see this, consider any $f \in F_{\text{in}}$ and $G \in \wp_k$ and let $f_{\text{out}} = f_0, \dots, f_d = f$ be the branch in F from the output gate down to f . Then in the expansion of $\chi_{P_k}(\mathcal{A}_{f_{\text{out}},P_k}^\Gamma)$, we get a contribution of 2^{ik} ($\leq 2^{k^2}$) from each sequence $(i, t_0, H_0, t_1, H_1, \dots, t_i, H_i)$ where $0 = t_0 < \dots < t_i = d$ and $P_k = H_0 \supset \dots \supset H_i = G$; here t_i is the location where the expansion of $\chi_{P_k}(\mathcal{A}_{f_{\text{out}},P_k}^\Gamma)$ branches as we move from $\chi_{H_{i-1}}(\mathcal{A}_{f_{i-1}, H_{i-1}}^\Gamma)$ to $2^k \chi_{H_i}(\mathcal{A}_{f_i, H_i}^\Gamma)$. Finally, we bound the number of (t_0, \dots, t_i) by $\binom{d}{i}$ ($\leq \text{depth}(F)^k$) and the number of (H_0, \dots, H_i) by 2^{ik} ($\leq 2^{k^2}$). Summing over i adds only a factor of k , so in total we get $c_{f,G} \leq 2^{O(k^2)} \cdot \text{depth}(F)^k$.

We now use the fact that $\sum_{G \in \wp_k} \chi_G(\mathcal{A}_{f,G}^\Gamma) = 1$ for all $f \in F_{\text{in}}$ (Claim 6.3) and $\text{size}(F) = |F_{\text{in}}|$ (since F is a formula!). Concluding the proof, we have

$$\chi_{P_k}(\mathcal{A}_{f_{\text{out}},P_k}^\Gamma) \leq 2^{O(k^2)} \cdot \text{depth}(F)^k \cdot \text{size}(F). \quad \square$$

We conclude this section by giving the proof of Theorem 2.1 assuming our pathset complexity lower bound (Theorem 5.8) and main technical lemma (Lemma 6.6).

Reduction 6.8. Theorem 5.8 and Lemma 6.6 \implies Theorem 2.1.

Proof. Assuming Theorem 5.8 and Lemma 6.6, we must show that $\text{size}(F) \geq n^{\Omega(\log k)}$. By Claim 6.5 and Lemma 6.6, there exists $\Gamma \in \{0,1\}^N$ such that $\delta(\mathcal{A}_{f_{\text{out}},P_k}^\Gamma) \geq .99n^{-2}$ and $\mathcal{A}_{f,G}^\Gamma \in \mathcal{P}_G^{\text{small}}$ for

all $f \in F$ and $G \in \wp_k$. Fix any such Γ . We now have

$$\begin{aligned} \text{size}(F) &\geq \frac{1}{2^{O(k^2)} \cdot \text{depth}(f)^k} \cdot \chi_{P_k}(\mathcal{A}_{f_{\text{out}}, P_k}^\Gamma) && \text{(Lemma 6.7)} \\ &\geq \frac{1}{2^{O(k^2)} \cdot \text{depth}(f)^k} \cdot \frac{n^{(1/6) \log k}}{2^{O(2^k)}} \cdot \delta(\mathcal{A}_{f_{\text{out}}, P_k}^\Gamma) && \text{(Theorem 5.8)}. \end{aligned}$$

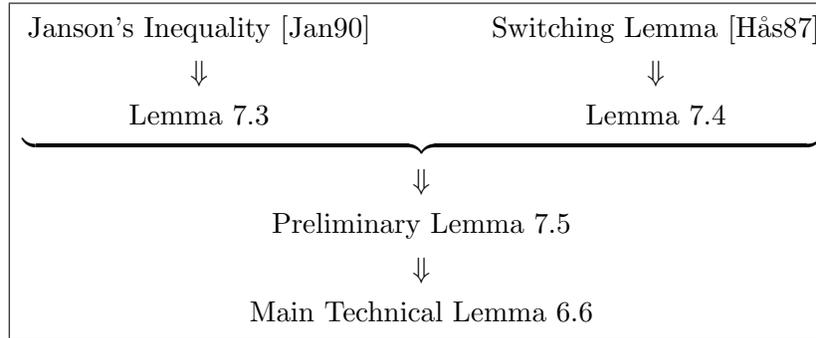
Using inequalities

$$\text{depth}(F) \leq \log^2 n, \quad \delta(\mathcal{A}_{f_{\text{out}}, P_k}^\Gamma) \geq .99n^{-2}, \quad k \leq \log \log n,$$

we get the desired bound $\text{size}(F) \geq n^{(1/6) \log k - O(1)}$. \square

7 Small Pathsets from Random Restrictions

In this section, we prove Lemma 6.6 showing that, with high probability over random $\Gamma \in \{0, 1\}_{1/n}^N$, pathsets $\mathcal{A}_{f, G}^\Gamma$ are small for all $f \in F$ and $G \in \wp_k$. The proof has the following scheme:



The central argument is contained in the proof of Preliminary Lemma 7.5 (from which Lemma 6.6 essentially follows as a corollary). In the interest of presenting this central argument first, the proofs of Lemmas 7.3 and 7.4 are given afterwards in §7.1 and §7.2.

Remark 7.1. Lemma 6.6 is similar to the main technical lemma in the k -clique lower bound of [Ros08, Ros10]. One important difference is that here we require a concentration of measure inequality in a place where a mere expectation bound sufficed for the k -clique result.¹⁰ This makes the proof of Lemma 6.6 somewhat more complicated.

Recall that G -smallness consists of $2^{\Delta_G} - 1$ density constraints corresponding to the nonempty unions of components of G . We say that non-small pathset \mathcal{A} is G -critical if it violates only the “top” constraint $\delta(\mathcal{A}) > \tilde{n}^{-\Delta_G}$. Formally:

¹⁰Consider the fact that functions $f : \{0, 1\}^n \rightarrow \{0, 1\}$ in AC^0 have low average sensitivity. This is a statement of the form $\mathbb{P}_{x \in \{0, 1\}^n, i \in [n]}[i \in S(f, x)] \leq \varepsilon$ where $S(f, x)$ is the set $\{i \in [n] : f(x) \neq f(x \oplus i)\}$. The main technical lemma of [Ros08] is an inequality of the same form; the difference is that $S(f, x)$, rather than the set of sensitive coordinates, is instead the set of *sensitive G -shaped sets of coordinates* where G is a pattern graph (which, in the context of k -clique, means a subgraph of K_k). By contrast, Lemma 6.6 is a concentration of measure inequality analogous to showing $\mathbb{P}_x[\mathbb{P}_i[i \in S(f, x)] > \varepsilon] \leq \delta$.

Definition 7.2 (Critical Pathsets). For a pattern graph G and pathset $\mathcal{A} \in \mathcal{P}_G$, we say that \mathcal{A} is G -critical if $\delta(\mathcal{A}) > \tilde{n}^{-\Delta_G}$ and $\mu_{V_G \setminus S}(\mathcal{A}) \leq \tilde{n}^{s-\Delta_G}$ for all $1 \leq s < \Delta_G$ and $S \subseteq V_G$ such that S intersects exactly s components of G .

Our first lemma, proved in §7.1, gives a concentration of measure inequality for critical pathsets. (Recall that $\varepsilon = 1/\log k$ from Def. 5.3(i).)

Lemma 7.3. *Let G be a pattern graph and suppose \mathcal{A} is a G -critical pathset. Let $q = (1/n)^{1+(\varepsilon/2k)}$. Then*

$$\mathbb{P}_{R \subseteq_q N} \left[\#\{x \in \mathcal{A} : N_{G,x} \subseteq R\} \leq \frac{n^{\varepsilon/2}}{2} \right] \leq \exp \left(-\Omega \left(\frac{n^{\varepsilon/2}}{2^k} \right) \right).$$

Our second lemma, proved in §7.2, is a straightforward corollary of Håstad's Switching Lemma [Hås87].

Lemma 7.4. *Suppose f is a boolean function¹¹ computed by a circuit of size s and depth d . For all $0 < q \leq p \leq 1/2$ and $\tau, r > 0$,*

$$d \leq \frac{\log(p/q)}{r^{-1} \log(s/\tau) + \log(5r)} \implies \mathbb{P}_{\theta \in \mathcal{R}(p,q)} [|\text{Live}(f[\theta])| > 2^r] \leq \tau.$$

We now give the core argument in the proof of Lemma 6.6.

Lemma 7.5. *Suppose $f : \{0,1\}^N \rightarrow \{0,1\}$ is computed by a circuit of size $\leq n^k$ and depth $\leq \frac{\log n}{k^3 \log \log n}$. Let G, H be pattern graphs with $V_G \cap V_H = \emptyset$ and let $y \in [n]^{V_H}$. For $\Gamma \in \{0,1\}^N$, define G -pathset \mathcal{A}^Γ by*

$$\mathcal{A}^\Gamma := \{x \in [n]^{V_G} : N_{G,x} \subseteq \text{Live}(f[\rho_{G \cup H,xy}^\Gamma])\}.$$

Then $\mathbb{P}_{\Gamma \in \{0,1\}_{1/n}^N} [\mathcal{A}^\Gamma \text{ is } G\text{-critical}] \leq O(n^{-10k})$.

Note that Lemma 7.5 does not rely on the assumption that $k \leq \log \log n$. It holds up to $k \leq \log^{1/3} n$ (at which point the statement is trivial). We also remark that, while the bound $O(n^{-10k})$ is sufficient for our purposes, we could easily get a stronger bound like $O(n^{-k \log k})$.

Proof. Define $\mathcal{I} \subseteq \{0,1\}^N$ by

$$\mathcal{I} := \{I \in \{0,1\}^N : I_\nu = 0 \text{ for all } \nu \in N \setminus N_{H,y}\}.$$

Note that $|\mathcal{I}| = 2^{|N_{H,y}|} = 2^{|E_H|} \leq 2^k$.

For $I \in \mathcal{I}$, define $f_I : \{0,1\}^N \rightarrow \{0,1\}$ by $f_I(\Gamma) = f(\Gamma \oplus I)$. For all $x \in [n]^{V_G}$ and $\Gamma \in \{0,1\}^N$, we have

$$(3) \quad \text{Live}(f[\rho_{G \cup H,xy}^\Gamma]) = \bigcup_{I \in \mathcal{I}} \text{Live}(f_I[\rho_{G,x}^\Gamma]).$$

¹¹This statement is independent of the number of variables of f .

For $R \subseteq N$, let $\theta_R^\Gamma : N \rightarrow \{0, 1, *\}$ be the restriction taking value $*$ over R and equal to Γ over $N \setminus R$. Define pathsets \mathcal{B}_R and \mathcal{C}_R^Γ by

$$\begin{aligned}\mathcal{B}_R &:= \{x \in [n]^{V_G} : N_{G,x} \subseteq R\}, \\ \mathcal{C}_R^\Gamma &:= \{x \in [n]^{V_G} : N_{G,x} \subseteq \bigcup_{I \in \mathcal{I}} \text{Live}(f_I \upharpoonright \theta_R^\Gamma)\}.\end{aligned}$$

It follows from (3) that $\mathcal{A}^\Gamma \cap \mathcal{B}_R \subseteq \mathcal{C}_R^\Gamma$. Also, since $|N_{G,x}| = |E_G| \leq k$ and $|\mathcal{I}| \leq 2^k$,

$$(4) \quad |\mathcal{C}_R^\Gamma|^{1/k} \leq \left| \bigcup_{I \in \mathcal{I}} \text{Live}(f_I \upharpoonright \theta_R^\Gamma) \right| \leq 2^k \cdot \max_{I \in \mathcal{I}} |\text{Live}(f_I \upharpoonright \theta_R^\Gamma)|.$$

We now consider independent random $\Gamma \in \{0, 1\}_{1/n}^N$ and random $R \subseteq_q N$ where $q = (1/n)^{1+(\varepsilon/2k)}$. Note that θ_R^Γ has distribution $\mathcal{R}(1/n, q)$. Also, note that \mathcal{A}^Γ and \mathcal{B}_R are independent, as \mathcal{A}^Γ depends only on Γ and \mathcal{B}_R depends only on R .

We may assume that $k \leq \log^{1/3} n$, since otherwise the lemma is trivial. In particular, $2^k = o(n^{\varepsilon/2})$ (recall that $\varepsilon = 1/\log k$) and hence $\exp(-\Omega(n^{\varepsilon/2}/2^k)) = o(1)$. We have

$$\begin{aligned}\mathbb{P}_\Gamma [\mathcal{A}^\Gamma \text{ is } G\text{-critical}] &\leq \mathbb{P}_\Gamma \left[\mathbb{P}_R \left[|\mathcal{A}^\Gamma \cap \mathcal{B}_R| \leq \frac{n^{\varepsilon/2}}{2} \right] \leq \exp \left(-\Omega \left(\frac{n^{\varepsilon/2}}{2^k} \right) \right) \right] && \text{(Lemma 7.3)} \\ &= \mathbb{P}_\Gamma \left[\mathbb{P}_R \left[|\mathcal{A}^\Gamma \cap \mathcal{B}_R| > \frac{n^{\varepsilon/2}}{2} \right] \geq 1 - o(1) \right] \\ &\leq (1 + o(1)) \mathbb{P}_{\Gamma, R} \left[|\mathcal{A}^\Gamma \cap \mathcal{B}_R| > \frac{n^{\varepsilon/2}}{2} \right] && \text{(Markov ineq.)} \\ &\leq (1 + o(1)) \mathbb{P}_{\Gamma, R} \left[|\mathcal{C}_R^\Gamma| > \frac{n^{\varepsilon/2}}{2} \right] && (\mathcal{A}^\Gamma \cap \mathcal{B}_R \subseteq \mathcal{C}_R^\Gamma) \\ &\leq (1 + o(1)) \mathbb{P}_{\Gamma, R} \left[\max_{I \in \mathcal{I}} |\text{Live}(f_I \upharpoonright \theta_R^\Gamma)| > \frac{n^{\varepsilon/2k}}{2^{k+1}} \right] && \text{(by (4))} \\ &\leq (1 + o(1)) \sum_{I \in \mathcal{I}} \mathbb{P}_{\Gamma, R} \left[|\text{Live}(f_I \upharpoonright \theta_R^\Gamma)| > \frac{n^{\varepsilon/2k}}{2^{k+1}} \right].\end{aligned}$$

Using $|\mathcal{I}| \leq 2^k$ and the fact that θ_R^Γ has distribution $\mathcal{R}(1/n, q)$, it suffices to show that for every $I \in \mathcal{I}$ and sufficiently large n ,

$$(5) \quad \mathbb{P}_{\theta \in \mathcal{R}(1/n, q)} \left[|\text{Live}(f_I \upharpoonright \theta)| > \frac{n^{\varepsilon/2k}}{2^{k+1}} \right] \leq n^{-11k}.$$

In order to apply Lemma 7.4, let

$$p = n^{-1}, \quad s = n^k, \quad d = \frac{\log n}{k^3 \log \log n}, \quad \tau = n^{-11k}, \quad r = \log \left(\frac{n^{\varepsilon/2k}}{2^{k+1}} \right).$$

Recall that $q = (1/n)^{1+(\varepsilon/2k)}$ and $k \leq \log^{1/3} n$. We have $\log(p/q) = (\varepsilon/2k) \log n$ and $\log(5r) = O(\log \log n)$ and $r^{-1} \log(s/\tau) = ((\varepsilon/2k) \log n - k - 1)^{-1} 12k \log n = O(k^2/\varepsilon)$. Since $\varepsilon = 1/\log k$, we have

$$\frac{\log(p/q)}{r^{-1} \log(s/\tau) + \log(5r)} \geq \Omega \left(\frac{\log n}{(k^3/\varepsilon^2) + (k/\varepsilon) \log \log n} \right) \geq \omega(d).$$

Note that f_I has circuits of the same size ($\leq s$) and depth ($\leq d$) as f , since the operation $\Gamma \mapsto \Gamma \oplus I$ simply exchanges the positive and negative literals for variables corresponding to coordinates of I with value 1. Therefore, for sufficiently large n , Lemma 7.4 implies (5). This completes the proof. \square

Finally, we derive Lemma 6.6 from Lemma 7.5.

Proof of Lemma 6.6. Suppose $f : \{0, 1\}^N \rightarrow \{0, 1\}$ is computed by circuits of size $\leq n^k$ and depth $\leq \log n/k^3 \log \log n$. Fix a pattern graph G . We must show

$$\mathbb{P}_{\Gamma \in \{0,1\}_{1/n}^N} [\mathcal{A}_{f,G}^\Gamma \text{ is not } G\text{-small}] \leq O(n^{-2k}).$$

Suppose $\Gamma \in \{0, 1\}^N$ is any layered graph such that $\mathcal{A}_{f,G}^\Gamma$ is not G -small. We claim that there exist $S \subseteq V_G$ and $z \in [n]^{V_G \setminus S}$ such that S is a nonempty union of components of G and the pathset $\mathcal{B}_{S,z}^\Gamma$ is $G|_S$ -critical where $G|_S$ is the induced subgraph of G on S and

$$\mathcal{B}_{S,z}^\Gamma := \{y \in [n]^S : N_{G|_S,y} \subseteq \text{Live}(f[\rho_{G,yz}^\Gamma])\}.$$

We first note that it suffices to prove this claim. Since there are $2^{\Delta_G} - 1$ ($\leq 2^k$) choices for S and $\leq n^k$ choices for z , assuming the claim we have

$$\begin{aligned} \mathbb{P}_{\Gamma \in \{0,1\}_{1/n}^N} [\mathcal{A}_{f,G}^\Gamma \text{ is not } G\text{-small}] &\leq \mathbb{P}_{\Gamma \in \{0,1\}_{1/n}^N} \left[\bigvee_{S,z} \mathcal{B}_{S,z}^\Gamma \text{ is } G|_S\text{-critical} \right] \\ &\leq 2^k n^k O(n^{-10k}) \quad (\text{by Lemma 7.5}) \\ &\leq O(n^{-2k}). \end{aligned}$$

To see why the claim holds, assume that $\mathcal{A}_{f,G}^\Gamma$ is not G -small and consider the following procedure. Initially set $S \leftarrow V_G$ and $z \leftarrow ()$ (the empty tuple). If $\mathcal{B}_{S,z}^\Gamma$ is $G|_S$ -critical, then we are done. Otherwise, since $\mathcal{B}_{S,z}^\Gamma$ is neither $G|_S$ -small nor $G|_S$ -critical, there is a proper subset $T \subset S$ such that T is a union of $t \geq 1$ components of $V_{G|_S}$ and $\mu_T(\mathcal{B}_{S,z}^\Gamma) > \tilde{n}^{-t}$. By definition of μ_T , there exists $y \in [n]^{S \setminus T}$ such that $\delta(\mathcal{B}_{S,z}^\Gamma|_T^y) > \tilde{n}^{-t}$. Note that T is a union of components of G and $yz \in [n]^{V_G \setminus T}$. Also, for all $u \in [n]^T$, we have $N_{G|_T,u} \subseteq N_{G|_S,uy}$ and hence

$$\begin{aligned} u \in \mathcal{B}_{S,z}^\Gamma|_T^y &\implies uy \in \mathcal{B}_{S,z}^\Gamma \\ &\implies N_{G|_S,uy} \subseteq \text{Live}(f[\rho_{G,uyz}^\Gamma]) \\ &\implies N_{G|_T,u} \subseteq \text{Live}(f[\rho_{G,uyz}^\Gamma]) \\ &\implies u \in \mathcal{B}_{T,yz}^\Gamma. \end{aligned}$$

Therefore, $\mathcal{B}_{S,z}^\Gamma|_T^y \subseteq \mathcal{B}_{T,yz}^\Gamma$. It follows that $\delta(\mathcal{B}_{T,yz}^\Gamma) > \tilde{n}^{-t}$ and, hence, $\mathcal{B}_{T,yz}^\Gamma$ is not $G|_T$ -small. We now update $S \leftarrow T$ and $z \leftarrow yz$. Since $\mathcal{B}_{S,z}^\Gamma$ is not $G|_S$ -small, we may repeat this process so long as $\mathcal{B}_{S,z}^\Gamma$ is not $G|_S$ -critical. Since S shrinks with every step, eventually this process will terminate, at which point $\mathcal{B}_{S,z}^\Gamma$ is $G|_S$ -critical (and S is nonempty by definition of $G|_S$ -criticality). Thus, the claim holds and the lemma is proved. \square

7.1 Proof of Lemma 7.3

For the proof of Lemma 7.3 we use a concentration of measure inequality due to Janson [Jan90].

Lemma 7.6 (Janson's Inequality [Jan90]). *Let Ω be a finite universal set and let R be a random subset of Ω given by $\mathbb{P}[r \in R] = p_r$, these events mutually independent over $r \in \Omega$. Let $\{S_i\}_{i \in I}$ be an indexed family of subsets of Ω . Define λ and Υ by*

$$\lambda := \sum_{i \in I} \mathbb{P}[S_i \subseteq R], \quad \Upsilon := \sum_{(i,j) \in I^2 : i \neq j, S_i \cap S_j \neq \emptyset} \mathbb{P}[S_i \cup S_j \subseteq R].$$

Then, for all $0 \leq t \leq \lambda$, $\mathbb{P}\left[\#\{i \in I : S_i \subseteq R\} \leq \lambda - t\right] \leq \exp\left(-\frac{t^2}{2(\lambda + \Upsilon)}\right)$.

Proof of Lemma 7.3. Let G be a nonempty pattern graph, let \mathcal{A} be a G -critical pathset, and let $q = (1/n)^{1+(\varepsilon/2k)}$. We must show

$$\mathbb{P}_{R \subseteq qN} \left[\#\{x \in \mathcal{A} : N_{G,x} \subseteq R\} \leq \frac{n^{\varepsilon/2}}{2} \right] \leq \exp\left(-\Omega\left(\frac{n^{\varepsilon/2}}{2^k}\right)\right).$$

As in Janson's Inequality, define λ and Υ by

$$\lambda := \sum_{x \in \mathcal{A}} \mathbb{P}[N_{G,x} \subseteq R], \quad \Upsilon := \sum_{(x,y) \in \mathcal{A}^2 : x \neq y, N_{G,x} \cap N_{G,y} \neq \emptyset} \mathbb{P}[N_{G,x} \cup N_{G,y} \subseteq R].$$

Taking $t = \lambda/2$ in Lemma 7.6, we get

$$(6) \quad \mathbb{P}\left[\#\{x \in \mathcal{A} : N_{G,x} \subseteq R\} \leq \frac{\lambda}{2}\right] \leq \exp\left(-\frac{1}{16} \min\left\{\lambda, \frac{\lambda^2}{\Upsilon}\right\}\right).$$

Recall that $\Delta_G = |V_G| - |E_G|$ and $\tilde{n} = n^{1-\varepsilon}$. By G -criticality of \mathcal{A} ,

$$|\mathcal{A}| = n^{|V_G|} \delta(\mathcal{A}) > n^{|V_G|} \tilde{n}^{-\Delta_G} = n^{|E_G| + \varepsilon \Delta_G}.$$

Note that $\mathbb{P}[N_{G,x} \subseteq R] = q^{|E_G|}$ for all $x \in \mathcal{A}$. Since $|E_G| \leq k$ and $\Delta_G \geq 1$, it follows that

$$(7) \quad \lambda = |\mathcal{A}| \cdot q^{|E_G|} > n^{\varepsilon(\Delta_G - |E_G|/2k)} \geq n^{\varepsilon(\Delta_G - (1/2))} \geq n^{\varepsilon/2}.$$

To complete the proof, it suffices to show that $\frac{\lambda^2}{\Upsilon} \geq \frac{n^{\varepsilon/2}}{2^k}$.

For all $(x, y) \in \mathcal{A}^2$, let

$$T_{x,y} := \{vw \in E_G : x_v = y_v \text{ and } x_w = y_w\}.$$

Note that $x = y$ iff $T_{x,y} = E_G$, and $N_{G,x} \cap N_{G,y} \neq \emptyset$ iff $T_{x,y} \neq \emptyset$, and $|N_{G,x} \cup N_{G,y}| = 2|E_G| - |T_{x,y}|$. Next, note that $\Upsilon = \sum_{T: \emptyset \subset T \subset E_G} \Upsilon_T$ where

$$\Upsilon_T := \sum_{(x,y) \in \mathcal{A}^2 : T_{x,y} = T} \mathbb{P}[N_{G,x} \cup N_{G,y} \subseteq R] = \#\{(x,y) \in \mathcal{A}^2 : T_{x,y} = T\} \cdot q^{2|E_G| - |T|}.$$

Now consider any fixed $\emptyset \subset T \subset E_G$. Let $S = \bigcup_{vw \in T} \{v, w\}$ and let s be the number of components of G which S intersects. Note that $1 \leq s \leq |S| - |T|$, since $|S| - |T|$ equals the number of components in the induced subgraph $G|_S$. We have

$$\begin{aligned}
\#\{(x, y) \in \mathcal{A}^2 : T_{x,y} = T\} &= \sum_{z \in [n]^S} |\mathcal{A}|_{V_G \setminus S}^z|^2 \\
&\leq \left(\sum_{z \in [n]^S} |\mathcal{A}|_{V_G \setminus S}^z \right) \left(\max_{z \in [n]^S} |\mathcal{A}|_{V_G \setminus S}^z \right) \\
&= |\mathcal{A}| \cdot n^{|V_G| - |S|} \cdot \mu_{V_G \setminus S}(\mathcal{A}) \\
&\leq |\mathcal{A}| \cdot n^{|V_G| - |S|} \cdot \tilde{n}^{s - \Delta_G} \quad (\text{by } G\text{-criticality of } \mathcal{A}) \\
&= |\mathcal{A}| \cdot n^{|E_G| - |S| + s + \varepsilon(\Delta_G - s)}.
\end{aligned}$$

It follows that

$$\begin{aligned}
\Upsilon_T &= \#\{(x, y) \in \mathcal{A}^2 : T_{x,y} = T\} \cdot q^{2|E_G| - |T|} \\
&\leq |\mathcal{A}| \cdot n^{|E_G| - |S| + s + \varepsilon(\Delta_G - s)} \cdot q^{2|E_G| - |T|} \\
&\leq \lambda \cdot n^{|T| - |S| + s + \varepsilon(\Delta_G - s)} \quad (\text{using } \lambda = |\mathcal{A}| \cdot q^{|E_G|} \text{ and } q \leq n^{-1}) \\
&\leq \lambda \cdot n^{\varepsilon(\Delta_G - 1)} \quad (\text{using } 1 \leq s \leq |S| - |T|).
\end{aligned}$$

We now have

$$\begin{aligned}
(8) \quad \frac{\lambda^2}{\Upsilon} &= \frac{\lambda^2}{\sum_{T: \emptyset \subset T \subset E_G} \Upsilon_T} \geq \frac{\lambda \cdot n^{-\varepsilon(\Delta_G - 1)}}{2^k} \quad (\text{by the above}) \\
&\geq \frac{n^{\varepsilon/2}}{2^k} \quad (\text{since } \lambda \geq n^{\varepsilon(\Delta_G - (1/2))} \text{ by (7)}).
\end{aligned}$$

Plugging (7) and (8) into (6) completes the proof. \square

7.2 Proof of Lemma 7.4

For a boolean function f , we write $D(f)$ for the *decision-tree depth* of f (i.e. the minimum depth of a decision tree computing f). Note that $|Live(f)| \leq 2^{D(f)}$.

The following lemma is a special case of the original Switching Lemma of Håstad [Hås87]. (For simplicity, we consider depth- r decision trees as opposed to r -DNFs and s -CNFs.)

Lemma 7.7 (Switching Lemma [Hås87]). *Suppose f is a boolean function which is an AND or OR of (arbitrary many) depth- r decision trees. Then for all $q \in [0, 1/2]$,*

$$\mathbb{P}_{\theta \in \mathcal{R}(1/2, q)} [D(f[\theta]) > r] \leq (5qr)^r.$$

Lemma 7.4 follows directly from the following lemma via the fact that $|Live(f)| \leq 2^{D(f)}$. (This lemma originally appeared in the author's Ph.D. thesis [Ros10].)

Lemma 7.8. *Suppose f is a boolean function computed by a circuit of size s and depth d . For all $0 < q \leq p \leq 1/2$ and $\tau, r > 0$,*

$$d \leq \frac{\log(p/q)}{r^{-1} \log(s/\tau) + \log(5r)} \implies \mathbb{P}_{\theta \in \mathcal{R}(p, q)} [D(f[\theta]) > r] \leq \tau.$$

This lemma originally appeared in the author's Ph.D. thesis [Ros10]. The proof is included here for completeness.

Proof. We generate a sequence $\theta_0, \dots, \theta_d$ of random restrictions as follows:

- let $\theta_0 \in \mathcal{R}(p_0, p)$ where $p_0 = \frac{1}{1-p} \left(\frac{p+q}{2} - pq \right)$ (note that $0 < p_0 \leq 1$),
- for $i \in \{1, \dots, d\}$, let $\theta_i \in \mathcal{R}(1/2, (q/p)^{1/d})$ applied to the variables left unrestricted by $\theta_0, \dots, \theta_{i-1}$.

For $i \in \{0, \dots, d\}$, let Θ_i denote the composition of restrictions $\theta_0, \dots, \theta_i$. Note that Θ_d has distribution $\mathcal{R}(p, q)$:

$$\begin{aligned} \mathbb{P}[\Theta_d = *] &= \mathbb{P}[\theta_0 = \dots = \theta_d = *] = p \left(\left(\frac{q}{p} \right)^{1/d} \right)^d = q, \\ \mathbb{P}[\Theta_d = 1] &= \mathbb{P}[\theta_0 = 1] + \sum_{i=1}^d \mathbb{P}[\theta_0 = \dots = \theta_{i-1} = * \text{ and } \theta_i = 1] \\ &= p_0(1-p) + \sum_{i=1}^d \frac{p}{2} \left(\frac{q}{p} \right)^{(i-1)/d} \left(1 - \left(\frac{q}{p} \right)^{1/d} \right) \\ &= p_0(1-p) + \frac{p-q}{2} \\ &= (1-q)p. \end{aligned}$$

Therefore, Θ_d has distribution $\mathcal{R}(p, q)$.

Let C be the circuit of size s and depth d computing f . For each input/gate $g \in C$ at height i from the bottom (where inputs have height 0 and the output gate g_{out} has height d), let X_g denote the event that $D(g[\Theta_i]) \leq r$. Let $C_{<g}$ denote the set of $g' \in C$ such that g' lies below g .

If g has height 0, then $\mathbb{P}[X_g] = 1$. If g has height $i \geq 1$, then

$$\begin{aligned} \mathbb{P}[\neg X_g \mid \bigwedge_{g' \in C_{<g}} X_{g'}] &= \mathbb{P}[D((g[\Theta_{i-1}])[\theta_i]) > r \mid \bigwedge_{g' \in C_{<g}} D(g'[\Theta_{i-1}]) \leq r] \\ &\leq \left(5r \left(\frac{q}{p} \right)^{1/d} \right)^r && \text{(Lemma 7.7)} \\ &\leq \frac{\tau}{s} && \left(d \leq \frac{\log(p/q)}{\log(5r(s/\tau)^{1/r})} \right). \end{aligned}$$

Completing the proof, we have

$$\mathbb{P}_{\theta \in \mathcal{R}(p,q)} [D(f[\theta]) > r] = \mathbb{P}[\neg X_{g_{\text{out}}}] \leq \mathbb{P}[\bigvee_{g \in C} \neg X_g] \leq \sum_{g \in C} \mathbb{P}[\neg X_g \mid \bigwedge_{g' \in C_{<g}} X_{g'}] \leq \tau. \quad \square$$

8 Patterns

At this point in the paper, it only remains to prove Theorem 5.8, our lower bound on pathset complexity χ . As the first step in the proof, we introduce the notion of *patterns* and *pathset complexity w.r.t. patterns*, denoted $\bar{\chi}$. Intuitively, a pattern is a blueprint for constructing a pattern graph via pairwise unions starting from individual edges. This leads to a more constrained notion of pathset complexity where the allowable joins are prescribed by a given pattern.

Fixing the pattern of allowable joins can only increase the cost of constructing a pathset, hence $\chi \leq \bar{\chi}$. Counterintuitively, the lower bound on χ is derived from a lower bound on $\bar{\chi}$. This lower bound on $\bar{\chi}$ is the true combinatorial lower bound in this paper. (Unfortunately, in the shift from χ to $\bar{\chi}$ we lose a factor of $2^{O(2^k)}$, which is the reason that Theorem 5.8 only holds up to $k(n) \leq \log \log n$.)

In this section, we present the definition of $\bar{\chi}$ and state our lower bound for $\bar{\chi}$ (Theorem 8.3). The reduction from Theorem 8.3 to Theorem 5.8 is given in §9. In §10 we prove some preliminary lemmas (on properties of $\bar{\chi}$ with respect to projection and restriction). Finally we prove Theorem 8.3 in §11.

Definition 8.1 (Patterns). A *pattern* is a (rooted, unordered) binary tree whose leaves are labeled by edges of P_k (i.e. elements of $E_k = \{v_i v_{i+1} : 0 \leq i < k\}$). Every pattern A is associated with a pattern graph denoted $G_A = (V_A, E_A)$ where E_A is the set of edges of P_k which label leaves in A .

The *empty pattern* (of size 0) is denoted \emptyset . Patterns of size 1 (corresponding to elements of E_k) are said to be *atomic*. Patterns of size ≥ 2 are *non-atomic*. Throughout, A and B represent non-empty patterns. Let $\{A, B\}$ ($= \{B, A\}$) denote the pattern with children A and B . Note that every non-atomic pattern has the form $\{A, B\}$ for some A and B ; also, $G_{\{A, B\}} = G_A \cup G_B$.

For a pattern A , *sub-patterns* of A are sub-trees of A consisting a node in A and all nodes below that node with the inherited labeling of leaves. The sub-pattern and strict sub-pattern relations are denoted by \preceq and \prec respectively.

To simplify notation, for a pattern A we write \mathcal{P}_A for \mathcal{P}_{G_A} and proj_A for proj_{V_A} and ℓ_A for ℓ_{G_A} , etc. We consistently write $\mathcal{A}, \mathcal{B}, \mathcal{C}$ for pathsets with underlying patterns A, B, C respectively.

Definition 8.2 (Pathset Complexity w.r.t. Patterns). For every pattern A and pathset $\mathcal{A} \in \mathcal{P}_A$, the *pathset complexity* of \mathcal{A} with respect to A , denoted $\bar{\chi}_A(\mathcal{A})$, is defined by the following induction:

- (i) $\bar{\chi}_\emptyset(\{()\}) := 0$, that is, the pathset complexity of $\{()\}$ w.r.t. the empty pattern \emptyset is 0.
- (ii) If A is atomic and $|\mathcal{A}| = 1$, then $\bar{\chi}_A(\mathcal{A}) := 1$.
- (iii) For non-atomic $A = \{B, C\}$,

$$\bar{\chi}_A(\mathcal{A}) := \min_{(\mathcal{B}_i, \mathcal{C}_i)_i} \sum_i \bar{\chi}_B(\mathcal{B}_i) + \bar{\chi}_C(\mathcal{C}_i)$$

where $(\mathcal{B}_i, \mathcal{C}_i)_i$ ranges over sequences such that $\mathcal{B}_i \in \mathcal{P}_B^{\text{small}}$, $\mathcal{C}_i \in \mathcal{P}_C^{\text{small}}$ and $\mathcal{A} \subseteq \bigcup_i \mathcal{B}_i \bowtie \mathcal{C}_i$.

In Appendices A–C we present some key examples of patterns and prove upper and lower bounds for $\bar{\chi}$ with respect to some special classes of patterns. The material in these appendices is not directly needed for our main results. However, these appendices serve as a warm-up and motivation for the lower bound that follows.

The following inequalities (analogous to the inequalities following Definition 5.6 of χ) are essentially built into Definition 8.2 of $\bar{\chi}$:

$$\begin{array}{lll}
(\text{base case}) & \bar{\chi}_\emptyset(\{\emptyset\}) \leq 0 \text{ and } \bar{\chi}_A(\mathcal{A}) \leq 1 & \text{if } A \text{ is atomic and } |\mathcal{A}| = 1, \\
(\text{monotonicity}) & \bar{\chi}_A(\mathcal{A}') \leq \bar{\chi}_A(\mathcal{A}) & \text{if } \mathcal{A}' \subseteq \mathcal{A}, \\
(\text{sub-additivity}) & \bar{\chi}_A(\mathcal{A}_1 \cup \mathcal{A}_2) \leq \bar{\chi}_A(\mathcal{A}_1) + \bar{\chi}_A(\mathcal{A}_2) & \text{for all } \mathcal{A}_1, \mathcal{A}_2, \\
(\text{join rule}) & \bar{\chi}_{\{A,B\}}(\mathcal{A} \bowtie \mathcal{B}) \leq \bar{\chi}_A(\mathcal{A}) + \bar{\chi}_B(\mathcal{B}) & \text{if } \mathcal{A} \in \mathcal{P}_A^{\text{small}}, \mathcal{B} \in \mathcal{P}_B^{\text{small}}.
\end{array}$$

The essential difference between χ and $\bar{\chi}$ is that χ allows arbitrary joins, while $\bar{\chi}$ only allows joins as prescribed by the given pattern. Viewed as a minimum construction cost (see Remark 5.9), this means that $\bar{\chi}$ has more highly constrained rules of construction compared with χ . Consequently, $\chi_{G_A}(\mathcal{A}) \leq \bar{\chi}_A(\mathcal{A})$ for every pattern A and $\mathcal{A} \in \mathcal{P}_G$. Note that this inequality goes in the wrong direction for the purpose of proving a lower bound on χ . In §9 we give a different inequality between χ and $\bar{\chi}$ in the right direction.

Theorem 8.3 (Lower Bound for $\bar{\chi}$). *For every pattern A and pathset $\mathcal{A} \in \mathcal{P}_A$,*

$$\bar{\chi}_A(\mathcal{A}) \geq \tilde{n}^{(1/6) \log(\ell_A) + \Delta_A} \cdot \delta(\mathcal{A}).$$

The game plan for the rest of the paper is as follows: in §9 we derive our lower bound for χ (Theorem 5.8) from Theorem 8.3. In §10 we establish some important properties of $\bar{\chi}$. Finally, in §11 we give the proof of Theorem 8.3.

Remark 8.4 (Dual Characterization of $\bar{\chi}$). Similar to the dual characterization of χ mentioned in Remark 5.7, $\bar{\chi}$ has a *dual characterization* as the unique pointwise maximal function from $\{(A, \mathcal{A}) : A \text{ is a pattern and } \mathcal{A} \in \mathcal{P}_A\}$ to \mathbb{R} which satisfies inequalities (base case), (monotonicity), (sub-additivity) and (join rule). This fact is established by a straightforward induction on patterns (omitted here since we don't actually use this dual characterization in our lower bound).

This dual characterization suggests an obvious “direct method” for proving a lower bound on $\bar{\chi}$: find an explicit function from pairs (A, \mathcal{A}) to \mathbb{R} and show that this function satisfies inequalities (base case), (monotonicity), (sub-additivity) and (join rule). This is analogous to the “direct method” of proving a formula size lower bound via a *complexity measure*, defined as a function M from $\{\text{boolean functions on } n \text{ variables}\}$ to \mathbb{R} satisfying inequalities $M(f \wedge g) \leq M(f) + M(g)$ and $M(f \vee g) \leq M(f) + M(g)$ in addition to base case inequalities $M(f) \leq 0$ if f is constant and $M(f) \leq 1$ if f is a coordinate function.

Using the direct method, we were only able to prove lower bounds on $\bar{\chi}$ for a few restricted classes patterns (see Appendix B). For general patterns, we could not prove a lower bound along the lines of Theorem 8.3 using the direct method. We still do not know of any explicit function which satisfies (base case), (monotonicity), (sub-additivity) and (join rule) and maps $(A, [n]^{P_k})$ to $n^{\Omega(\log k)}$ for all patterns A with graph P_k . A priori, it is not even clear whether any such nice explicit function exists.¹²

¹²A natural approach is to consider functions of the form $n^{c_A} \cdot \nu(\mathcal{A})$ where c_A is a constant depending only on A and $\nu : \mathcal{P}_A \rightarrow \mathbb{R}$ is a monotone sub-additive function, such as δ or μ_S or π_S or any norm on $\mathbb{R}^{[n]^{V_A}}$ (viewing $\mathcal{P}_A \cong \{0, 1\}^{[n]^{V_A}}$ as a subset of $\mathbb{R}^{[n]^{V_A}}$). For such functions, one only needs to show (join rule); (base case) can be handled by appropriate scaling.

The proof of Theorem 8.3 which we present in §11 does not proceed via the direct method. In particular, neither the function $\tilde{n}^{(1/6)\log(\ell_{\mathcal{A}})+\Delta_{\mathcal{A}}}\cdot\delta(\mathcal{A})$ nor $\tilde{n}^{\Phi_{\mathcal{A}}}\cdot\delta(\mathcal{A})$ (defined in §11.1) satisfies inequality (join rule). Rather, our proof involves a more subtle induction on patterns.

9 From χ to $\bar{\chi}$

In this section, we prove:

Reduction 9.1. Theorem 8.3 (lower bound on $\bar{\chi}$) \implies Theorem 5.8 (lower bound on χ).

The following definition of *strict pattern* is only needed in this section. Rather than A, B, C , we write α, β, γ for this special class of patterns.

Definition 9.2. A pattern α is *strict* if $G_{\alpha''} \subset G_{\alpha'}$ for all $\alpha'' \prec \alpha' \preceq \alpha$. For a pattern graph G , let $Strict(G)$ denote the set of strict patterns α with graph G .

It is important that the number of strict patterns with a given pattern graph is bounded (though doubly exponential in $|E_G|$).

Lemma 9.3. *For every pattern graph G with r edges, there are only $2^{O(2^r)}$ strict patterns with graph G .*

Proof. Denote by $s(r)$ the number of strict patterns supported on any fixed set of r edges. (Note that $|Strict(G)|$ depends only on $|E_G|$.) Then we have $s(1) = 1$ and $s(r) \leq (r \cdot s(r-1))^2$ for all $r \geq 2$. Therefore,

$$s(r) \leq r^2(r-1)^4(r-2)^8 \dots 3^{2^{r-2}} 2^{2^{r-1}} 1^{2^r} = 2^{O(2^r)}. \quad \square$$

We now give the main lemma needed for Reduction 9.1.

Lemma 9.4. *For every pattern graph G and pathset $\mathcal{A} \in \mathcal{P}_G$, there is an indexed family $\{\mathcal{A}^{(\alpha)}\}_{\alpha \in Strict(G)}$ of sub-pathsets $\mathcal{A}^{(\alpha)} \subseteq \mathcal{A}$ such that*

$$\mathcal{A} = \bigcup_{\alpha \in Strict(G)} \mathcal{A}^{(\alpha)} \quad \text{and} \quad \forall \alpha \in Strict(G), \quad \bar{\chi}_{\alpha}(\mathcal{A}^{(\alpha)}) \leq \chi_G(\mathcal{A}).$$

Proof. By induction on $|E_G|$. The lemma is trivial if $|E_G| \leq 1$ (since in this case $|Strict(G)| = 1$). For the induction step, suppose G is a pattern graph with ≥ 2 edges. By Definition 5.6 of χ , there exists a sequence $(H_i, K_i, \mathcal{B}_i, \mathcal{C}_i)_i$ with

$$H_i, K_i \subset G, \quad H_i \cup K_i = G, \quad \mathcal{B}_i \in \mathcal{P}_{H_i}^{\text{small}}, \quad \mathcal{C}_i \in \mathcal{P}_{K_i}^{\text{small}}$$

such that

$$\mathcal{A} \subseteq \bigcup_i \mathcal{B}_i \bowtie \mathcal{C}_i \quad \text{and} \quad \chi_G(\mathcal{A}) = \sum_i \chi_{H_i}(\mathcal{B}_i) + \chi_{K_i}(\mathcal{C}_i).$$

For each $\alpha = \{\beta, \gamma\} \in Strict(G)$, define $\mathcal{A}^{(\alpha)}$ inductively by

$$\mathcal{A}^{(\alpha)} := \mathcal{A} \cap \bigcup_{i: H_i=G_{\beta}, K_i=G_{\gamma}} \mathcal{B}_i^{(\beta)} \bowtie \mathcal{C}_i^{(\gamma)}.$$

First, we show that $\mathcal{A} = \bigcup_{\alpha \in \text{Strict}(G)} \mathcal{A}^{(\alpha)}$. The inclusion \supseteq is obvious. For the inclusion \subseteq , consider any $x \in \mathcal{A}$. Then x belongs to $\mathcal{B}_i \bowtie \mathcal{C}_i$ for some i . This means that $x_{V_{H_i}} \in \mathcal{B}_i$ and $x_{V_{K_i}} \in \mathcal{C}_i$. By the induction hypothesis, there exist $\beta \in \text{Strict}(H_i)$ and $\gamma \in \text{Strict}(K_i)$ such that $x_{V_{H_i}} \in \mathcal{B}_i^{(\beta)}$ and $x_{V_{K_i}} \in \mathcal{C}_i^{(\gamma)}$. Let $\alpha = \{\beta, \gamma\}$ and note that $\alpha \in \text{Strict}(G)$. Since $x \in \mathcal{B}_i^{(\beta)} \bowtie \mathcal{C}_i^{(\gamma)}$, it follows that $x \in \mathcal{A}^{(\alpha)}$, proving the inclusion \subseteq .

Finally, for all $\alpha \in \text{Strict}(G)$, we show $\bar{\chi}_\alpha(\mathcal{A}) \leq \chi_G(\mathcal{A})$ as follows:

$$\begin{aligned} \bar{\chi}_\alpha(\mathcal{A}) &\leq \bar{\chi}_\alpha\left(\bigcup_{i: H_i=G_\beta, K_i=G_\gamma} \mathcal{B}_i^{(\beta)} \bowtie \mathcal{C}_i^{(\gamma)}\right) && \text{(monotonicity)} \\ &\leq \sum_{i: H_i=G_\beta, K_i=G_\gamma} \bar{\chi}_\alpha(\mathcal{B}_i^{(\beta)} \bowtie \mathcal{C}_i^{(\gamma)}) && \text{(sub-additivity)} \end{aligned}$$

Noting that $\mathcal{B}_i^{(\beta)}$ and $\mathcal{C}_i^{(\gamma)}$ are small (since $\mathcal{B}_i^{(\beta)} \subseteq \mathcal{B}_i \in \mathcal{P}_{H_i}^{\text{small}}$ and $\mathcal{C}_i^{(\gamma)} \subseteq \mathcal{C}_i \in \mathcal{P}_{K_i}^{\text{small}}$), we continue:

$$\begin{aligned} &\leq \sum_{i: H_i=G_\beta, K_i=G_\gamma} \bar{\chi}_\beta(\mathcal{B}_i^{(\beta)}) + \bar{\chi}_\gamma(\mathcal{C}_i^{(\gamma)}) && \text{(join rule)} \\ &\leq \sum_{i: H_i=G_\beta, K_i=G_\gamma} \chi_{H_i}(\mathcal{B}_i) + \chi_{K_i}(\mathcal{C}_i) && \text{(ind. hyp.)} \\ &\leq \chi_G(\mathcal{A}). \end{aligned} \quad \square$$

The next corollary follows directly from Lemma 9.4.

Corollary 9.5. *For every pattern graph G and pathset $\mathcal{A} \in \mathcal{P}_G$, there is a strict pattern $\alpha \in \text{Strict}(G)$ and a sub-pathset $\mathcal{A}' \subseteq \mathcal{A}$ such that $\bar{\chi}_\alpha(\mathcal{A}') \leq \chi_G(\mathcal{A})$ and $\delta(\mathcal{A}) \leq |\text{Strict}(G)| \cdot \delta(\mathcal{A}')$. \square*

We conclude this section with the proof of Reduction 9.1.

Proof of Reduction 9.1. Assume Theorem 8.3 and consider arbitrary $\mathcal{A} \in \mathcal{P}_{P_k}$. By Corollary 9.5, there exist $\alpha \in \text{Strict}(P_k)$ and $\mathcal{A}' \subseteq \mathcal{A}$ such that $\chi_\alpha(\mathcal{A}') \leq \chi_{P_k}(\mathcal{A})$ and $\delta(\mathcal{A}) \leq |\text{Strict}(P_k)| \cdot \delta(\mathcal{A}') \leq 2^{O(2^k)} \cdot \delta(\mathcal{A}')$ (Lemma 9.3). We now have

$$\begin{aligned} \chi_{P_k}(\mathcal{A}) &\geq \chi_\alpha(\mathcal{A}') \geq \tilde{n}^{(1/6)\log(\ell_\alpha) + \Delta_\alpha} \cdot \delta(\mathcal{A}') && \text{(Theorem 8.3)} \\ &\geq \frac{\tilde{n}^{(1/6)\log(k)+1}}{2^{O(2^k)}} \cdot \delta(\mathcal{A}) \\ &\geq \frac{n^{(1/6)\log k}}{2^{O(2^k)}} \cdot \delta(\mathcal{A}) && \text{(as } \tilde{n} = n^{1-\varepsilon} = n^{1-(1/\log k)}\text{)}. \end{aligned}$$

This shows that Theorem 5.8 holds, which completes the proof of the reduction. \square

10 Projection and Restriction

In this section we establish two key properties of $\bar{\chi}$: it is monotone decreasing with respect to projection to sub-patterns (Lemma 10.2) and restriction to unions of components (Lemma 10.4). We also introduce an operation on patterns $A \ominus B$ (Definition 10.7), read as “ A restricted away from B ”. This notation will be extremely convenient in §11.

10.1 $\bar{\chi}$ Decreases Under Projection

Claim 10.1. *For every non-atomic pattern $\{A, B\}$ and pathset $\mathcal{C} \in \mathcal{P}_{\{A, B\}}$, we have $\bar{\chi}_A(\text{proj}_A(\mathcal{C})) \leq \bar{\chi}_{\{A, B\}}(\mathcal{C})$.*

Proof. By Definition 8.2(iii) of $\bar{\chi}_{\{A, B\}}(\mathcal{C})$, there is a sequence $(\mathcal{A}_i, \mathcal{B}_i)_i$ such that

$$\mathcal{A}_i \in \mathcal{P}_A^{\text{small}}, \quad \mathcal{B}_i \in \mathcal{P}_B^{\text{small}}, \quad \mathcal{C} \subseteq \bigcup_i \mathcal{A}_i \bowtie \mathcal{B}_i \quad \text{and} \quad \bar{\chi}_{\{A, B\}}(\mathcal{C}) = \sum_i \bar{\chi}_A(\mathcal{A}_i) + \bar{\chi}_B(\mathcal{B}_i).$$

Note that $\text{proj}_A(\mathcal{C}) \subseteq \text{proj}_A(\bigcup_i \mathcal{A}_i \bowtie \mathcal{B}_i) \subseteq \bigcup_i \mathcal{A}_i$. By monotonicity and sub-additivity of $\bar{\chi}_A$, it follows that

$$\bar{\chi}_A(\text{proj}_A(\mathcal{C})) \leq \bar{\chi}_A(\bigcup_i \mathcal{A}_i) \leq \sum_i \bar{\chi}_A(\mathcal{A}_i) \leq \bar{\chi}_{\{A, B\}}(\mathcal{C}). \quad \square$$

Lemma 10.2 ($\bar{\chi}$ decreases under projections). *For every pattern A and pathset $\mathcal{A} \in \mathcal{P}_A$ and sub-pattern $A' \preceq A$, $\bar{\chi}_{A'}(\text{proj}_{A'}(\mathcal{A})) \leq \bar{\chi}_A(\mathcal{A})$.*

Proof. Induction using Claim 10.1. □

10.2 $\bar{\chi}$ Decreases Under Restriction

For a pattern A and a pathset $\mathcal{A} \in \mathcal{P}_A$, Lemma 10.2 concerns projections of \mathcal{A} of the form $\text{proj}_{A'}(\mathcal{A})$ where A' is a sub-pattern of A . The restrictions of \mathcal{A} that we consider next are not restrictions of the form $\mathcal{A}|_{A'}^z$, where $z \in [n]^{V_A \setminus V_{A'}}$. Note that $\mathcal{A}|_{A'}^z \subseteq \text{proj}_{A'}(\mathcal{A})$, so we already have $\bar{\chi}_{A'}(\mathcal{A}|_{A'}^z) \leq \bar{\chi}_A(\mathcal{A})$ by Lemma 10.2 and monotonicity of $\bar{\chi}_{A'}$.

Rather than restrictions over sub-patterns, we instead consider restrictions of the form $\mathcal{A}|_S^z$ where $z \in [n]^{V_A \setminus S}$ and $S \subseteq V_A$ is a union of components of G_A . We define an operation of restriction on patterns; the restriction $A|_S$ is a pattern with $V_{A|_S} = S$. Even though $A|_S$ is not necessarily a sub-pattern of A , we will show that $\bar{\chi}_{A|_S}(\mathcal{A}|_S^z) \leq \bar{\chi}_A(\mathcal{A})$.

Definition 10.3 (Restriction of Patterns).

(i) For all $S \subseteq V_k$, let \bar{S} denote the complement $V_k \setminus S$ of S in V_k .

(ii) For a pattern A , we say that S is *A-respecting* if $V_A \cap S$ is a union of components of G_A .

Note that S is *A-respecting* $\iff \bar{S}$ is *A-respecting* \iff every leaf in A is labeled by an edge $v_i v_{i+1} \in E_k$ such that $\{v_i, v_{i+1}\} \subseteq S$ or $\{v_i, v_{i+1}\} \subseteq \bar{S}$. Also note that if S is $\{A, B\}$ -respecting, then it is both *A-respecting* and *B-respecting* and $\{A, B\}|_S = \{A|_S, B|_S\}$.

(iii) If S is *A-respecting*, we denote by $A|_S$ the pattern obtained from A by pruning all leaves labeled by elements of $\bar{S} \times \bar{S}$.

For example, if A is the pattern $\{\{v_1 v_2, v_5 v_6\}, \{v_2 v_3, v_6 v_7\}\}$ and S is the *A-respecting* set $\{v_1, v_2, v_3\}$, then $A|_S = \{v_1 v_2, v_2 v_3\}$. Note that $A|_S = \{v_1 v_2, v_2 v_3\}$ also when S is the *A-respecting* set $\{v_1, v_2, v_3, v_4\}$; in general, $A|_S = A|(V_A \cap S)$. Also note that $A|_S$ is not a sub-pattern of A in this example.

Before stating the main lemma of this subsection, recall our convention concerning notation $\mathcal{A}|_S^z$ (see Definition 4.3): for every pattern A and pathset $\mathcal{A} \in \mathcal{P}_A$ and $S \subseteq V_k$ and $z \in [n]^{\bar{S}}$, the pathset $\mathcal{A}|_S^z \in \mathcal{P}_{A|_S}$ is defined by $\mathcal{A}|_S^z := \mathcal{A}|_{V_A \cap S}^{z'} = \{y \in V_A \cap S : yz' \in \mathcal{A}\}$ where $z' = z_{V_A \setminus S}$.

Lemma 10.4 ($\bar{\chi}$ decreases under restrictions). *For every pattern A and pathset $\mathcal{A} \in \mathcal{P}_A$ and A -respecting $S \subseteq V_k$ and $z \in [n]^{\bar{S}}$, we have $\bar{\chi}_{A \upharpoonright S}(\mathcal{A}|_S^z) \leq \bar{\chi}_A(\mathcal{A})$.*

Proof. By induction on patterns. The lemma is trivial for empty and atomic patterns. For the induction step, consider a non-atomic pattern $\{A, B\}$ and assume the lemma holds for A and B . Let $\mathcal{C} \in \mathcal{P}_{\{A, B\}}$, let S be a $\{A, B\}$ -respecting subset of V_k , and let $z \in [n]^{\bar{S}}$. By Def. 8.2(iii) of $\bar{\chi}_{\{A, B\}}(\mathcal{C})$, there is a sequence $(\mathcal{A}_i, \mathcal{B}_i)_i$ such that

$$\mathcal{A}_i \in \mathcal{P}_A^{\text{small}}, \quad \mathcal{B}_i \in \mathcal{P}_B^{\text{small}}, \quad \mathcal{C} \subseteq \bigcup_i \mathcal{A}_i \bowtie \mathcal{B}_i \quad \text{and} \quad \bar{\chi}_{\{A, B\}}(\mathcal{C}) = \sum_i \bar{\chi}_A(\mathcal{A}_i) + \bar{\chi}_B(\mathcal{B}_i).$$

By Lemma 10.5, $\mathcal{A}_i|_S^z \in \mathcal{P}_{A \upharpoonright S}^{\text{small}}$ and $\mathcal{B}_i|_S^z \in \mathcal{P}_{B \upharpoonright S}^{\text{small}}$. We now have

$$\begin{aligned} \bar{\chi}_{\{A, B\} \upharpoonright S}(\mathcal{C}|_S^z) &\leq \bar{\chi}_{\{A, B\} \upharpoonright S} \left(\bigcup_i (\mathcal{A}_i \bowtie \mathcal{B}_i)|_S^z \right) && \text{(monotonicity)} \\ &\leq \sum_i \bar{\chi}_{\{A, B\} \upharpoonright S}((\mathcal{A}_i \bowtie \mathcal{B}_i)|_S^z) && \text{(sub-additivity)} \\ &= \sum_i \bar{\chi}_{\{A \upharpoonright S, B \upharpoonright S\}}(\mathcal{A}_i|_S^z \bowtie \mathcal{B}_i|_S^z) \\ &\leq \sum_i \bar{\chi}_{A \upharpoonright S}(\mathcal{A}_i|_S^z) + \bar{\chi}_{B \upharpoonright S}(\mathcal{B}_i|_S^z) && \text{(join rule)} \\ &\leq \sum_i \bar{\chi}_A(\mathcal{A}_i) + \bar{\chi}_B(\mathcal{B}_i) && \text{(ind. hyp.)} \\ &= \bar{\chi}_{\{A, B\}}(\mathcal{C}). \end{aligned} \quad \square$$

Lemma 10.5 (Smallness is preserved under restriction). *For every pattern A and small pathset $\mathcal{A} \in \mathcal{P}_A^{\text{small}}$ and A -respecting $S \subseteq V_k$ and $z \in [n]^{\bar{S}}$, we have $\mathcal{A}|_S^z \in \mathcal{P}_{A \upharpoonright S}^{\text{small}}$.*

Proof. Immediate from Definition 5.3 of small pathsets. □

Remark 10.6. Smallness is preserved under joins (Lemma 5.5) and restrictions to union of components (Lemma 10.5). However, smallness is not preserved under projection to unions of components. A counterexample is the pattern $A = \{v_1 v_2, v_3 v_4\}$ and pathset $\mathcal{A} = \{x \in [n]^{V_A} : x_1 = x_3 \text{ and } x_2 = x_4\} \in \mathcal{P}_A^{\text{small}}$. Letting A' be the atomic sub-pattern $v_1 v_2$ of A , we have $\pi_{A'}(\mathcal{A}) = 1$, hence $\text{proj}_{A'}(\mathcal{A}) \notin \mathcal{P}_{A'}^{\text{small}}$.

10.3 The Operation $A \ominus B$

As a convenient notation, we introduce an operation $A \ominus B$ on patterns, read as A restricted away from B .

Definition 10.7. For patterns A and B , we write $A \ominus B$ for the pattern $A \upharpoonright S$ where $S \subseteq V_A$ consists of the components of G_A which do not intersect V_B . For example, if $A = \{\{v_1 v_2, v_4 v_5\}, \{v_2 v_3, v_5 v_6\}\}$ (so G_A is the union of paths $v_1 v_2 v_3$ and $v_4 v_5 v_6$) and $B = \{v_6 v_7\}$, then $A \ominus B = \{v_1 v_2, v_2 v_3\}$.

We conclude this section with two basic lemmas about this operation.

Lemma 10.8. For all patterns $C = \{A, B\}$ and $A' \preceq A$ and $B' \preceq B$,

$$\Delta_C \leq \Delta_{A'} + \Delta_{B' \ominus A'} + \Delta_{C \ominus \{A', B'\}}.$$

Proof. Each component of G_C contains at least one component from at least one of the three vertex-disjoint graphs $G_{A'}$, $G_{B' \ominus A'}$ and $G_{C \ominus \{A', B'\}}$. \square

Lemma 10.9. For all patterns $C = \{A, B\}$ and $A' \preceq A$ and $B' \preceq B$ and pathsets $\mathcal{A} \in \mathcal{P}_A$ and $\mathcal{B} \in \mathcal{P}_B$,

$$\delta(\mathcal{A} \bowtie \mathcal{B}) \leq \pi_{A'}(\mathcal{A}) \cdot \mu_{B' \ominus A'}(\text{proj}_{B'}(\mathcal{B})) \cdot \mu_{C \ominus \{A', B'\}}(\mathcal{A} \bowtie \mathcal{B}).$$

Proof. By Lemma 4.4(c),

$$\delta(\mathcal{A} \bowtie \mathcal{B}) \leq \pi_{A'}(\mathcal{A}) \cdot \mu_{V_{B'} \setminus V_{A'}}(\text{proj}_{B'}(\mathcal{B})) \cdot \mu_{V_C \setminus (V_{A'} \cup V_{B'})}(\mathcal{A} \bowtie \mathcal{B}).$$

Since $V_{B' \ominus A'} \subseteq V_{B'} \setminus V_{A'}$ and $V_{C \ominus \{A', B'\}} \subseteq V_C \setminus (V_{A'} \cup V_{B'})$, by Lemma 4.4(a),

$$\begin{aligned} \mu_{V_{B'} \setminus V_{A'}}(\text{proj}_{B'}(\mathcal{B})) &\leq \mu_{B' \ominus A'}(\text{proj}_{B'}(\mathcal{B})), \\ \mu_{V_C \setminus (V_{A'} \cup V_{B'})}(\mathcal{A} \bowtie \mathcal{B}) &\leq \mu_{C \ominus \{A', B'\}}(\mathcal{A} \bowtie \mathcal{B}). \end{aligned}$$

Combining these inequalities finishes the proof. \square

11 Lower Bound for $\bar{\chi}$

In this section we prove Theorem 8.3, our lower bound for $\bar{\chi}$. Recall that ℓ_A denote the length of the longest path in G_A , i.e., the number of edges in the largest component of G_A .

Theorem 8.3. (restated) For every pattern A and pathset $\mathcal{A} \in \mathcal{P}_A$,

$$\bar{\chi}_A(\mathcal{A}) \geq \tilde{n}^{(1/6) \log(\ell_A) + \Delta_A} \cdot \delta(\mathcal{A}).$$

To prove Theorem 8.3, first we define an auxiliary function $\Phi : \{\text{patterns}\} \rightarrow \mathbb{R}$. We then prove two lemmas: $\bar{\chi}_A(\mathcal{A}) \geq \tilde{n}^{\Phi_A} \delta(\mathcal{A})$ (Lemma 11.2) and $\Phi_A \geq \frac{1}{6} \log(\ell_A) + \Delta_A$ (Lemma 11.4).

11.1 Definition of Φ_A

Definition 11.1. Let $\Phi : \{\text{patterns}\} \rightarrow \mathbb{R}$ be the unique minimal function such that the following hold:

- $\Phi_A = 0$ if A is empty, and $\Phi_A = 2$ if A is atomic,
- for every non-atomic pattern $C = \{A, B\}$ and sub-patterns $A' \preceq A$ and $B' \preceq B$,

$$\begin{aligned} (\dagger)_{A', B}^C \quad \Phi_C &\geq \Phi_{A'} + \Delta_{B \ominus A'} + \Delta_{C \ominus \{A', B\}}, \\ (\ddagger)_{A', B'}^C \quad \Phi_C &\geq \frac{\Phi_{A'} + \Phi_{B' \ominus A'} + \Delta_C + \Delta_{C \ominus \{A', B'\}}}{2}. \end{aligned}$$

We refer to (\dagger) and (\ddagger) as the “one-sided” and “balanced” inequalities. Note that since $\{A, B\}$ and $\{B, A\}$ are considered to be the same pattern, we also have the reverse inequalities $(\dagger)_{B', A}^C$ and $(\ddagger)_{B', A'}^C$. For better readability, we write Φ_A inline and $\Phi^{(A)}$ in superscripts.

Some brief remarks on this definition:

- Minimality of Φ among functions satisfying these inequalities means that for every non-atomic pattern $C = \{A, B\}$, at least one of the four inequalities $(\dagger)_{A',B}^C$, $(\dagger)_{B',A}^C$, $(\ddagger)_{A',B'}^C$, $(\ddagger)_{B',A'}^C$ is tight (i.e. holds with equality) for some $A' \preceq A$ and $B' \preceq B$.
- Note that Φ is monotone decreasing with respect to sub-patterns, that is, $\Phi_{A'} \leq \Phi_A$ for all $A' \preceq A$ (by inequalities (\dagger)).
- Φ increases by means of the contribution of Δ 's: if we remove the Δ 's from $(\dagger)_{A',B}^C$ and $(\ddagger)_{A',B'}^C$ (replacing these inequalities by $\Phi_C \geq \Phi_{A'}$ and $\Phi_C \geq \frac{1}{2}(\Phi_{A'} + \Phi_{B'})$ respectively), then we would have $\Phi_A = 2$ for every nonempty pattern A . Intuitively, in the attempt to lower bound Φ_A , the objective of the game is to pick up as many Δ 's as possible.
- For the patterns A_k and B_k defined in Appendix A, we have $\Phi_{A_k} \geq \Phi_{A_{\lceil k/4 \rceil}} + 1$ by (\dagger) and $\Phi_{B_k} \geq \Phi_{B_{\lfloor (k-1)/2 \rfloor}} + \frac{1}{2}$ by (\ddagger) for all $k \geq 4$. It follows that $\Phi_{A_k} \geq \frac{1}{2} \log k - O(1)$ and $\Phi_{B_k} \geq \frac{1}{2} \log k - O(1)$.

11.2 Showing $\bar{\chi}_A(\mathcal{A}) \geq \tilde{n}^{\Phi_A} \delta(\mathcal{A})$

We now prove the most important lemma in the overall proof of Theorem 8.3. Lemma 11.2 accounts for the definition of Φ_A (essentially Φ_A is the maximum function for which the argument of Lemma 11.2 is valid). The two cases (\dagger) and (\ddagger) in the proof are inspired by the special cases proved in Appendix B.

Lemma 11.2. *For every pattern A and pathset $\mathcal{A} \in \mathcal{P}_A$, $\bar{\chi}_A(\mathcal{A}) \geq \tilde{n}^{\Phi(A)} \delta(\mathcal{A})$.*

Proof. We argue by induction on patterns. The base case where A is empty or atomic is trivial. For the induction step, consider a non-atomic pattern $C = \{A, B\}$ and assume the lemma holds for all smaller patterns.

We claim that it suffices to show that

$$(9) \quad \tilde{n}^{\Phi(C)} \delta(\mathcal{A} \bowtie \mathcal{B}) \leq \bar{\chi}_A(\mathcal{A}) + \bar{\chi}_B(\mathcal{B})$$

for all $\mathcal{A} \in \mathcal{P}_A^{\text{small}}$, $\mathcal{B} \in \mathcal{P}_B^{\text{small}}$. To see that this suffices, consider any $\mathcal{C} \in \mathcal{P}_C$. By the dual characterization of pathset complexity, there exists a covering $\mathcal{C} \subseteq \bigcup_i \mathcal{A}_i \bowtie \mathcal{B}_i$ by joins of small pathsets \mathcal{A}_i and \mathcal{B}_i such that $\bar{\chi}_C(\mathcal{C}) = \sum_i \bar{\chi}_A(\mathcal{A}_i) + \bar{\chi}_B(\mathcal{B}_i)$. Note that

$$\delta(\mathcal{C}) \leq \delta\left(\bigcup_i \mathcal{A}_i \bowtie \mathcal{B}_i\right) \leq \sum_i \delta(\mathcal{A}_i \bowtie \mathcal{B}_i).$$

Assuming (9) holds for all \mathcal{A}_i and \mathcal{B}_i , we have

$$\tilde{n}^{\Phi(C)} \delta(\mathcal{C}) \leq \sum_i \tilde{n}^{\Phi(C)} \delta(\mathcal{A}_i \bowtie \mathcal{B}_i) \leq \sum_i \bar{\chi}_A(\mathcal{A}_i) + \bar{\chi}_B(\mathcal{B}_i) = \bar{\chi}_C(\mathcal{C}).$$

We now turn to proving inequality (9). Fix small pathsets $\mathcal{A} \in \mathcal{P}_A^{\text{small}}$ and $\mathcal{B} \in \mathcal{P}_B^{\text{small}}$. Note that $\mathcal{A} \bowtie \mathcal{B} \in \mathcal{P}_C^{\text{small}}$ by Lemma 5.5. Recall that at least one of the four inequalities $(\dagger)_{A',B}^C$, $(\ddagger)_{A',B'}^C$, $(\dagger)_{B',A}^C$, $(\ddagger)_{B',A'}^C$ is tight for some $A' \preceq A$ and $B' \preceq B$. By symmetry of the argument, we consider only the first two possibilities without loss of generality.

Case (†) (one-sided induction case): Assume that there exists $A' \preceq A$ such that $(\dagger)_{A',B}^C$ is tight, that is,

$$(10) \quad \Phi_C = \Phi_{A'} + \Delta_{B \ominus A'} + \Delta_{C \ominus \{A', B\}}.$$

By Lemma 10.9, we have

$$\delta(\mathcal{A} \bowtie \mathcal{B}) \leq \pi_{A'}(\mathcal{A}) \cdot \mu_{B \ominus A'}(\mathcal{B}) \cdot \mu_{C \ominus \{A', B\}}(\mathcal{A} \bowtie \mathcal{B}).$$

Since \mathcal{B} is B -small and $\mathcal{A} \bowtie \mathcal{B}$ is C -small, we have

$$\mu_{B \ominus A'}(\mathcal{B}) \leq \tilde{n}^{-\Delta(B \ominus A')} \quad \text{and} \quad \mu_{C \ominus \{A', B\}}(\mathcal{A} \bowtie \mathcal{B}) \leq \tilde{n}^{-\Delta(C \ominus \{A', B\})}.$$

Combining these inequalities (and substituting $\delta(\text{proj}_{A'}(\mathcal{A}))$ for $\pi_{A'}(\mathcal{A})$), we have

$$(11) \quad \delta(\mathcal{A} \bowtie \mathcal{B}) \leq \tilde{n}^{-\Delta(B \ominus A') - \Delta(C \ominus \{A', B\})} \delta(\text{proj}_{A'}(\mathcal{A})).$$

Using the fact that $\bar{\chi}$ decreases under projections, together with the induction hypothesis, we have

$$\begin{aligned} \tilde{n}^{\Phi(C)} \delta(\mathcal{A} \bowtie \mathcal{B}) &= \tilde{n}^{\Phi(A') + \Delta(B \ominus A') + \Delta(C \ominus \{A', B\})} \delta(\mathcal{A} \bowtie \mathcal{B}) && \text{(by (10))} \\ &\leq \tilde{n}^{\Phi(A')} \delta(\text{proj}_{A'}(\mathcal{A})) && \text{(by (11))} \\ &\leq \bar{\chi}_{A'}(\text{proj}_{A'}(\mathcal{A})) && \text{(ind. hyp.)} \\ &\leq \bar{\chi}_A(\mathcal{A}) && \text{(Lemma 10.2).} \end{aligned}$$

Therefore, (9) holds in this case.

Case (‡) (balanced induction case): Assume that there exist $A' \preceq A$ and $B' \preceq B$ such that $(\ddagger)_{A', B'}^C$ is tight, that is,

$$(12) \quad \Phi_C = \frac{\Phi_{A'} + \Phi_{B' \ominus A'} + \Delta_C + \Delta_{C \ominus \{A', B'\}}}{2}.$$

By Lemma 10.9, we have

$$\delta(\mathcal{A} \bowtie \mathcal{B}) \leq \pi_{A'}(\mathcal{A}) \cdot \mu_{B' \ominus A'}(\text{proj}_{B'}(\mathcal{B})) \cdot \mu_{C \ominus \{A', B'\}}(\mathcal{A} \bowtie \mathcal{B}).$$

By definition of $\mu_{B' \ominus A'}$, there exists $z \in [n]^{V_{B'} \setminus V_{B' \ominus A'}}$ such that

$$\mu_{B' \ominus A'}(\text{proj}_{B'}(\mathcal{B})) = \delta(\text{proj}_{B'}(\mathcal{B})|_{B' \ominus A'}^z).$$

C -smallness of $\mathcal{A} \bowtie \mathcal{B}$ implies both

$$\delta(\mathcal{A} \bowtie \mathcal{B}) \leq \tilde{n}^{-\Delta(C)} \quad \text{and} \quad \mu_{C \ominus \{A', B'\}}(\mathcal{A} \bowtie \mathcal{B}) \leq \tilde{n}^{-\Delta(C \ominus \{A', B'\})}.$$

Taking square roots and combining these inequalities, we have

$$(13) \quad \delta(\mathcal{A} \bowtie \mathcal{B}) \leq \sqrt{\tilde{n}^{-\Delta(C) - \Delta(C \ominus \{A', B'\})} \cdot \pi_{A'}(\mathcal{A}) \cdot \delta(\text{proj}_{B'}(\mathcal{B})|_{B' \ominus A'}^z)}.$$

Using the fact that $\bar{\chi}$ decreases under projections and restrictions (Lemmas 10.2 and 10.4), together with the induction hypothesis, we have

$$(14) \quad \begin{aligned} \tilde{n}^{\Phi(A')} \pi_{A'}(\mathcal{A}) &= \tilde{n}^{\Phi(A')} \delta(\text{proj}_{A'}(\mathcal{A})) \leq \bar{\chi}_{A'}(\text{proj}_{A'}(\mathcal{A})) && \text{(ind. hyp.)} \\ &\leq \bar{\chi}_A(\mathcal{A}) && \text{(Lemma 10.2)} \end{aligned}$$

and also

$$(15) \quad \begin{aligned} \tilde{n}^{\Phi(B' \ominus A')} \delta(\text{proj}_{B'}(\mathcal{B})|_{B' \ominus A'}^z) &\leq \bar{\chi}_{B' \ominus A'}(\text{proj}_{B'}(\mathcal{B})|_{B' \ominus A'}^z) && \text{(ind. hyp.)} \\ &\leq \bar{\chi}_{B'}(\text{proj}_{B'}(\mathcal{B})) && \text{(Lemma 10.4)} \\ &\leq \bar{\chi}_B(\mathcal{B}) && \text{(Lemma 10.2)}. \end{aligned}$$

We now finish the proof using the inequality of arithmetic and geometric means (AM-GM inequality):

$$\begin{aligned} \tilde{n}^{\Phi(C)} \delta(\mathcal{A} \bowtie \mathcal{B}) &= \sqrt{\tilde{n}^{\Phi(A') + \Phi(B' \ominus A') + \Delta(C) + \Delta(C \ominus \{A', B'\})} \cdot \delta(\mathcal{A} \bowtie \mathcal{B})} && \text{(by (12))} \\ &\leq \sqrt{\tilde{n}^{\Phi(A') + \Phi(B' \ominus A')} \cdot \pi_{A'}(\mathcal{A}) \cdot \delta(\text{proj}_{B'}(\mathcal{B})|_{B' \ominus A'}^z)} && \text{(by (13))} \\ &\leq \frac{1}{2} \left(\tilde{n}^{\Phi(A')} \pi_{A'}(\mathcal{A}) + \tilde{n}^{\Phi(B' \ominus A')} \delta(\text{proj}_{B'}(\mathcal{B})|_{B' \ominus A'}^z) \right) && \text{(AM-GM ineq.)} \\ &\leq \frac{1}{2} \left(\bar{\chi}_A(\mathcal{A}) + \bar{\chi}_B(\mathcal{B}) \right) && \text{(by (14), (15)).} \end{aligned}$$

Therefore, (9) holds in this case also, which concludes the proof. \square

11.3 Showing $\Phi_A \geq \frac{1}{6} \log(\ell_A) + \Delta_A$

We now complete the proof of Theorem 8.3 by proving Lemma 11.4 ($\Phi_A \geq \frac{1}{6} \log(\ell_A) + \Delta_A$ for all patterns A). We require one preliminary lemma.

Lemma 11.3. *For every pattern A and A -respecting $S \subseteq V_k$, we have $\Phi_A \geq \Phi_{A \upharpoonright S} + \Delta_{A \upharpoonright \bar{S}}$.*

Proof. We argue by induction on patterns. The lemma is trivial when A is empty or atomic. For the induction step, consider any non-atomic pattern $C = \{A, B\}$ and assume the lemma holds for all smaller patterns. Let S be any C -respecting subset of V_k .

Noting that $C \upharpoonright S = \{A \upharpoonright S, B \upharpoonright S\}$ and every sub-pattern of $A \upharpoonright S$ has the form $A' \upharpoonright S$ where $A' \preceq A$ (and similarly for $B \upharpoonright S$), it follows that that at least one the four inequalities

$$(\dagger)_{A' \upharpoonright S, B \upharpoonright S}^{C \upharpoonright S}, \quad (\ddagger)_{A' \upharpoonright S, B' \upharpoonright S}^{C \upharpoonright S}, \quad (\dagger)_{B' \upharpoonright S, A \upharpoonright S}^{C \upharpoonright S}, \quad (\ddagger)_{B' \upharpoonright S, A' \upharpoonright S}^{C \upharpoonright S}$$

is tight for some $A' \preceq A$ and $B' \preceq B$. Once again, without loss of generality, we consider just the first two possibilities.

First, consider the case that there exists $A' \preceq A$ for which $(\dagger)_{A' \upharpoonright S, B \upharpoonright S}^{C \upharpoonright S}$ is tight, that is,

$$(16) \quad \Phi_{C \upharpoonright S} = \Phi_{A' \upharpoonright S} + \Delta_{(B \ominus A') \upharpoonright S} + \Delta_{(C \ominus \{A', B\}) \upharpoonright S}.$$

In this case, we have

$$\begin{aligned}
\Phi_C &\geq \Phi_{A'} + \Delta_{B \ominus A'} + \Delta_{C \ominus \{A', B\}} && \text{(by } (\dagger)_{A', B}^C \text{)} \\
&\geq \Phi_{A'} + \Delta_{B \ominus A'} + \Delta_{C \ominus \{A', B\}} \\
&\quad + \Delta_{C \uparrow \bar{S}} - \Delta_{A' \uparrow \bar{S}} - \Delta_{(B \ominus A') \uparrow \bar{S}} - \Delta_{(C \ominus \{A', B\}) \uparrow \bar{S}} && \text{(Lemma 10.8)} \\
&= \Phi_{A'} - \Delta_{A' \uparrow \bar{S}} + \Delta_{B \ominus A' \uparrow S} + \Delta_{C \ominus \{A', B\} \uparrow S} + \Delta_{C \uparrow \bar{S}} \\
&\geq \Phi_{A' \uparrow S} + \Delta_{B \ominus A' \uparrow S} + \Delta_{C \ominus \{A', B\} \uparrow S} + \Delta_{C \uparrow \bar{S}} && \text{(ind. hyp.)} \\
&\geq \Phi_{C \uparrow S} + \Delta_{C \uparrow \bar{S}} && \text{(by (16)).}
\end{aligned}$$

Finally, consider the alternative that there exist $A' \preceq A$ and $B' \preceq B$ for which $(\ddagger)_{A' \uparrow S, B' \uparrow S}^{C \uparrow S}$ is tight, that is,

$$(17) \quad \Phi_{C \uparrow S} = \frac{\Phi_{A' \uparrow S} + \Phi_{(B' \ominus A') \uparrow S} + \Delta_{C \uparrow S} + \Delta_{(C \ominus \{A', B'\}) \uparrow S}}{2}.$$

In this case, we have

$$\begin{aligned}
\Phi_C &\geq \frac{1}{2}(\Phi_{A'} + \Phi_{B' \ominus A'} + \Delta_C + \Delta_{C \ominus \{A', B'\}}) && \text{(by } (\ddagger)_{A', B'}^C \text{)} \\
&\geq \frac{1}{2}(\Phi_{A'} + \Phi_{B' \ominus A'} + (\Delta_{C \uparrow S} + \Delta_{C \uparrow \bar{S}}) + \Delta_{C \ominus \{A', B'\}}) + \\
&\quad \frac{1}{2}(\Delta_{C \uparrow \bar{S}} - \Delta_{A' \uparrow \bar{S}} - \Delta_{(B' \ominus A') \uparrow \bar{S}} - \Delta_{(C \ominus \{A', B'\}) \uparrow \bar{S}}) && \text{(Lemma 10.8)} \\
&= \frac{1}{2}(\Phi_{A'} - \Delta_{A' \uparrow \bar{S}} + \Phi_{B' \ominus A'} - \Delta_{(B' \ominus A') \uparrow \bar{S}} + \Delta_{C \uparrow S} + \Delta_{(C \ominus \{A', B'\}) \uparrow S}) + \Delta_{C \uparrow \bar{S}} \\
&\geq \frac{1}{2}(\Phi_{A' \uparrow S} + \Phi_{(B' \ominus A') \uparrow S} + \Delta_{C \uparrow S} + \Delta_{(C \ominus \{A', B'\}) \uparrow S}) + \Delta_{C \uparrow \bar{S}} && \text{(ind. hyp.)} \\
&\geq \Phi_{C \uparrow S} + \Delta_{C \uparrow \bar{S}} && \text{(by (17)).}
\end{aligned}$$

Having shown $\Phi_C \geq \Phi_{C \uparrow S} + \Delta_{C \uparrow \bar{S}}$ in both cases, we are done. \square

Lemma 11.4. *For every pattern A , $\Phi_A \geq \frac{1}{6} \log(\ell_A) + \Delta_A$.*

Proof. We argue by induction on patterns. The base case where A is empty or atomic is trivial. For the induction step, let A be a non-atomic pattern and assume the lemma holds for all smaller patterns. We will consider a sequence of cases. In each case, after showing that $\Phi_A \geq \frac{1}{6} \log(\ell_A) + \Delta_A$ under a given hypothesis, we will proceed assuming the negation of that hypothesis. The sequences of cases is summarized at the end of the proof.

First, consider the case that G_A is disconnected (i.e. $\Delta_A \geq 2$). Let S be the largest component of G_A . We have

$$\begin{aligned}
\Phi_A &\geq \Phi_{A \uparrow S} + \Delta_{A \uparrow \bar{S}} && \text{(Lemma 11.3)} \\
&\geq \frac{1}{6} \log(\ell_{A \uparrow S}) + \Delta_{A \uparrow S} + \Delta_{A \uparrow \bar{S}} && \text{(ind. hyp.)} \\
&= \frac{1}{6} \log(\ell_A) + \Delta_A.
\end{aligned}$$

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. $\ell_A = k$). Our goal is to show that

$$\Phi_A \geq \frac{1}{6} \log(k) + 1.$$

Since $\Phi_A \geq 2$ for all nonempty patterns, we may further assume that $k > 2^6$. (Below, we will only need the assumption that $k > 8$.)

Consider the case that there exists a sub-pattern $A' \preceq A$ such that $|E_{A'}| \geq k/8$ and $\Delta_{A'} \geq 2$. Note that $\ell_{A'} \geq |E_{A'}|/\Delta_{A'}$ (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

$$\begin{aligned}
\Phi_A \geq \Phi_{A'} &\geq \frac{1}{6} \log(\ell_{A'}) + \Delta_{A'} && (\text{ind. hyp.}) \\
&\geq \frac{1}{6} \log(k) - \frac{1}{2} - \frac{1}{6} \log(\Delta_{A'}) + \Delta_{A'} && (\ell_{A'} \geq |E_{A'}|/\Delta_{A'} \geq k/8\Delta_{A'}) \\
&\geq \frac{1}{6} \log(k) - \frac{1}{2} - \frac{1}{6} \log(2) + 2 && (\Delta_{A'} \geq 2) \\
&= \frac{1}{6} \log(k) + \frac{4}{3} \\
&> \frac{1}{6} \log(k) + 1.
\end{aligned}$$

This proves the lemma in this case.

Therefore, we proceed under the following assumption:

$$(*) \quad \text{for all } A' \preceq A, \text{ if } |E_{A'}| \geq k/8 \text{ then } \Delta_{A'} = 1.$$

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

It is easy to see that there exist proper sub-patterns $B, Z \prec A$ such that

$$v_0 \in V_B, \quad v_k \in V_Z, \quad |E_B|, |E_Z| < k/8, \quad |E_{B^\dagger}|, |E_{Z^\dagger}| \geq k/8.$$

Fix any choice of such B and Z . Note that G_{B^\dagger} and G_{Z^\dagger} are connected by $(*)$. 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 $\ell_{B^\dagger} < k/2$ and $\ell_{Z^\dagger} < k/2$. Note that V_{B^\dagger} and V_{Z^\dagger} are disjoint and hence $Z^\dagger \ominus B^\dagger = Z^\dagger$. Let Y denote the least common ancestor of B^\dagger and Z^\dagger in A . We have

$$\begin{aligned}
\Phi_A \geq \Phi_Y &\geq \frac{1}{2}(\Phi_{B^\dagger} + \Phi_{Z^\dagger \ominus B^\dagger} + \Delta_Y + \Delta_{Y \ominus \{B^\dagger, Z^\dagger\}}) && (\text{by } (\ddagger)_{B^\dagger, Z^\dagger}^Y) \\
&= \frac{1}{2}(\Phi_{B^\dagger} + \Phi_{Z^\dagger}) + \frac{1}{2} && (\Delta_Y \geq 1) \\
&\geq \frac{1}{2}(\frac{1}{6} \log(\ell_{B^\dagger}) + \Delta_{B^\dagger} + \frac{1}{6} \log(\ell_{Z^\dagger}) + \Delta_{Z^\dagger}) + \frac{1}{2} && (\text{ind. hyp.}) \\
&\geq \frac{1}{6} \log(k/8) + \frac{3}{2} \\
&= \frac{1}{6} \log(k) + 1.
\end{aligned}$$

(We remark that this is the only place in the proof where the inequality (\ddagger) is used and the only tight case which forces $1/6$.)

Therefore, we proceed under the assumption that $\ell_{B^\dagger} \geq k/2$ or $\ell_{Z^\dagger} \geq k/2$. Without loss of generality, we assume that $\ell_{B^\dagger} \geq k/2$. (We now forget about Z and Z^\dagger .)

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

$$G_A = P_k, \quad (*), \quad B \preceq A, \quad v_0 \in V_B, \quad |E_B| < k/8, \quad |E_{B^\dagger}| = \ell_{B^\dagger} \geq k/2.$$

Going forward, we will define vertices v_r, v_s, v_t where $0 < r < s < t \leq k$.

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. furthest away from v_0).

Note that E_B contains edges $v_i v_{i+1}$ for all $i \in \{0, \dots, r-1\} \cup \{t, \dots, \lceil k/2 \rceil - 1\}$. (In the event that $t < k/2$, since $G_{B^\dagger} = G_B \cup G_{B^\sim}$ 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, $r + (k/2) - t \leq |E_B| < k/8$. It follows that

$$t - r > 3k/8.$$

Next, note that $|E_{B^\sim}| \geq |E_{B^\dagger}| - |E_B| \geq (k/2) - (k/8) > k/8$. It follows that there exists a proper sub-pattern $C \prec B^\sim$ such that

$$v_t \in V_C, \quad |E_C| < k/8, \quad |E_{C^\dagger}| \geq k/8.$$

Fix any choice of such C .

Consider the case that $\ell_{C^\dagger} < 3k/8$. Since G_{C^\dagger} is connected (by $(*)$) and $v_t \in V_{C^\dagger}$ and $t - r > 3k/8$, it follows that $V_{C^\dagger} \cap \{v_0, \dots, v_r\} = \emptyset$ and hence $\Delta_{B \ominus C^\dagger} \geq 1$. We have

$$\begin{aligned} \Phi_A &\geq \Phi_{B^\dagger} \geq \Phi_{C^\dagger} + \Delta_{B \ominus C^\dagger} + \Delta_{B^\dagger \ominus \{B, C^\dagger\}} && \text{(by } (\dagger)_{C^\dagger, B}^{B^\dagger} \text{)} \\ &\geq \Phi_{C^\dagger} + 1 \\ &\geq \frac{1}{6} \log(\ell_{C^\dagger}) + \Delta_{C^\dagger} + 1 && \text{(ind. hyp.)} \\ &\geq \frac{1}{6} \log(k/8) + 2 \\ &> \frac{1}{6} \log(k) + 1. \end{aligned}$$

Therefore, we proceed under the assumption that $\ell_{C^\dagger} \geq 3k/8$. Since $E_{C^\dagger} = E_C \cup E_{C^\sim}$, we have

$$|E_{C^\sim}| \geq |E_{C^\dagger}| - |E_C| > (3k/8) - (k/8) = k/4.$$

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 $v_t v_{t+1} \notin E_{B^\sim}$ and hence $v_t v_{t+1} \notin E_C$ (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 $t - r > 3k/8$ and note that $t - s \leq |E_C| < k/8$. Therefore,

$$s - r = (t - r) - (t - s) > (3k/8) - (k/8) = k/4.$$

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

$$k/8 \leq |E_D| < k/4.$$

To see this, note that there exists a chain of sub-patterns $C^\sim = D_0 \succ D_1 \succ \dots \succ D_j$ such that D_j is atomic and $D_i = D_{i-1}^\dagger$ and $|E_{D_i}| \geq |E_{D_i^\sim}|$ for all $i \in \{1, \dots, j\}$. Since $|E_{D_0}| > k/4$ and $|E_{D_j}| = 1$ and $|E_{D_{i-1}}| = |E_{D_i}| + |E_{D_i^\sim}| \leq 2|E_{D_i}|$, it must be the case that there exists $i \in \{1, \dots, j\}$ such that $k/8 \leq |E_{D_i}| < k/4$.

Since $|E_D| \geq k/8$, $(*)$ implies that G_D is connected. Since $|E_D| < k/4$ and $s - r > k/4$, 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

$$\begin{aligned}
\Phi_A &\geq \Phi_{B^\dagger} \geq \Phi_D + \Delta_{B \ominus D} + \Delta_{B^\dagger \ominus \{B, D\}} && \text{(by } (\dagger)_{D, B}^{B^\dagger} \text{)} \\
&\geq \Phi_D + 1 \\
&\geq \frac{1}{6} \log(\ell_D) + \Delta_D + 1 && \text{(ind. hyp.)} \\
&\geq \frac{1}{6} \log(k/8) + 2 \\
&> \frac{1}{6} \log(k) + 1.
\end{aligned}$$

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

$$\Phi_A \geq \Phi_{C^\dagger} \geq \Phi_D + \Delta_{C \ominus D} + \Delta_{C^\dagger \ominus \{C, D\}} \geq \Phi_D + 1 > \frac{1}{6} \log(k) + 1.$$

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

1. $\Delta_A \geq 2$ else assume $w \log G_A = P_k$,
2. $\exists A' \prec A$ with $\Delta_{A'} \geq 2$ and $\ell_{A'} \geq k/8$ else assume $(*)$,
3. $|E_{B^\dagger}| < k/2$ and $|E_{Z^\dagger}| < k/2$ else assume $w \log |E_{B^\dagger}| \geq k/2$,
4. $|E_{C^\dagger}| < 3k/8$ else assume $|E_{C^\dagger}| \geq 3k/8$,
5. $v_r \notin E_D$ or $v_s \notin E_D$.

Since $\Phi_A \geq \frac{1}{6} \log(\ell_A) + \Delta_A$ in each case, the proof is complete. □

As we have now proved Lemmas 11.2 and 11.4, this completes the proof of Theorem 8.3 and hence also of Theorem 2.1.

12 Conclusion

We proved the first super-polynomial separation in the power of bounded-depth boolean formulas vs. circuits via technique based on the notion of pathset complexity. The most obvious question for future research is whether pathset complexity can be used to derive lower bounds for distance $k(n)$ connectivity in other models of computation.

We conclude with a comment extending our results to the *average-case* setting. Let $p(n) = \Theta(n^{-\frac{k+1}{k}})$ be the exact threshold function such that

$$\mathbb{P}_{G=G(n,p)} [G \in STCONN(k(n))] = \frac{1}{2}$$

where $G(n, p)$ is the Erdős-Rényi random graph with edge probability $p(n)$. Our proof of Theorem 2.1 is easily adapted to give the same $n^{(1/6) \log k - O(1)}$ lower bound for bounded-depth formulas F which satisfy

$$\mathbb{P}_{G=G(n,p)} [F(G) = 1 \iff G \in STCONN(k(n))] \geq \frac{1}{2} + \varepsilon$$

for any constant $\varepsilon > 0$. Using the idea behind Proposition 5.11, we can construct formulas F of size $n^{(1/2)\log k + O(1)}$ (the best worst-case upper bound I know of is size $n^{\log k + O(1)}$) and depth $O(\log k)$ which solve $STCONN(k(n))$ in a strong average-case sense:

$$\mathbb{P}_{G=G(n,p)} [F(G) = 1 \iff G \in STCONN(k(n))] \geq 1 - \exp(-n^{\Omega(1)}).$$

It would be interesting to close the gap between $\frac{1}{6} \log k$ and $\frac{1}{2} \log k$ in these bounds.

Acknowledgements

I want to thank Osamu Watanabe and Rahul Santhanam for many helpful discussions, Stasys Jukna for his interest in this work and stimulating discussions at a time when I was badly stuck on the proof, and Igor Carboni Oliveira for valuable feedback on an earlier draft of this paper.

References

- [Ajt89] Miklós Ajtai. First-order definability on finite structures. *Annals of Pure and Applied Logic*, 45(3):211–225, 1989.
- [BIP98] Paul Beame, Russell Impagliazzo, and Toniann Pitassi. Improved depth lower bounds for small distance connectivity. *Computational Complexity*, 7(4):325–345, 1998.
- [BPU92] Stephen Bellantoni, Toniann Pitassi, and Alasdair Urquhart. Approximation and small-depth frege proofs. *SIAM Journal on Computing*, 21(6):1161–1179, 1992.
- [EPA99] Jeff Edmonds, Chung Keung Poon, and Dimitris Achlioptas. Tight lower bounds for st-connectivity on the NNJAG model. *SIAM Journal on Computing*, 28(6):2257–2284, 1999.
- [FSS84] Merrick L. Furst, James B. Saxe, and Michael Sipser. Parity, circuits, and the polynomial-time hierarchy. *Mathematical Systems Theory*, 17:13–27, 1984.
- [Hås87] Johan Håstad. *Computational limitations of small-depth circuits*. MIT press, 1987.
- [Hås98] Johan Håstad. The shrinkage exponent of de Morgan formulas is 2. *SIAM Journal on Computing*, 27(1):48–64, 1998.
- [HJP95] Johan Håstad, Stasys Jukna, and Pavel Pudlák. Top-down lower bounds for depth-three circuits. *Computational Complexity*, 5(2):99–112, 1995.
- [Jan90] Svante Janson. Poisson approximation for large deviations. *Random Structures & Algorithms*, 1(2):221–229, 1990.
- [Juk12] Stasys Jukna. *Boolean Function Complexity: Advances and Frontiers*, volume 27. Springer-Verlag Berlin Heidelberg, 2012.
- [KW90] Mauricio Karchmer and Avi Wigderson. Monotone circuits for connectivity require super-logarithmic depth. *SIAM Journal on Discrete Mathematics*, 3(2):255–265, 1990.

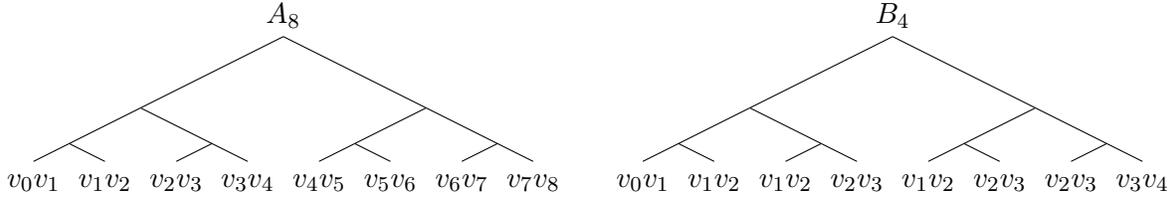
- [NW96] Noam Nisan and Avi Wigderson. Lower bounds on arithmetic circuits via partial derivatives. *Computational Complexity*, 6(3):217–234, 1996.
- [Pot10] Aaron Potechin. Bounds on monotone switching networks for directed connectivity. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 553–562. IEEE, 2010.
- [Raz87] Alexander A. Razborov. Lower bounds on the size of bounded depth circuits over a complete basis with logical addition. *Math. Notes*, 41:333–338, 1987.
- [Rei08] Omer Reingold. Undirected connectivity in log-space. *Journal of the ACM (JACM)*, 55(4):17, 2008.
- [Ros08] Benjamin Rossman. On the constant-depth complexity of k -clique. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, pages 721–730. ACM, 2008.
- [Ros10] Benjamin Rossman. *Average-case complexity of detecting cliques*. PhD thesis, Massachusetts Institute of Technology, 2010.
- [RR97] Alexander A. Razborov and Steven Rudich. Natural proofs. *J. Comput. Syst. Sci.*, 55(1):24–35, 1997.
- [RS42] John Riordan and Claude E Shannon. The number of two-terminal series-parallel networks. *J. Math. Phys*, 21(2):83–93, 1942.
- [RW89] Ran Raz and Avi Wigderson. Probabilistic communication complexity of boolean relations. In *Foundations of Computer Science, 1989., 30th Annual Symposium on*, pages 562–567. IEEE, 1989.
- [Sav70] Walter J Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of computer and system sciences*, 4(2):177–192, 1970.
- [Smo87] Roman Smolensky. Algebraic methods in the theory of lower bounds for boolean circuit complexity. In *STOC '87: Proceedings of the 19th Annual ACM Symposium on Theory of Computing*, pages 77–82, 1987.
- [Spi71] PM Spira. On time-hardware complexity tradeoffs for boolean functions. In *Proceedings of the 4th Hawaii Symposium on System Sciences*, pages 525–527, 1971.
- [SS79] Eli Shamir and Marc Snir. On the depth complexity of formulas. *Mathematical Systems Theory*, 13(1):301–322, 1979.
- [SW98] Petr Savický and Alan R Woods. The number of boolean functions computed by formulas of a given size. *Random Structures & Algorithms*, 13(3-4):349–382, 1998.
- [TT94] Prasoon Tiwari and Martin Tompa. A direct version of Shamir and Snir’s lower bounds on monotone circuit depth. *Information Processing Letters*, 49(5):243–248, 1994.
- [Wig92] Avi Wigderson. The complexity of graph connectivity. In *Proceedings of the 17th International Symposium on Mathematical Foundations of Computer Science*, pages 112–132. Springer-Verlag, 1992.

A Key Examples

We introduce two key examples of patterns, denoted A_k and B_k , and present upper bounds for $\bar{\chi}$ with respect to these patterns. In the next section, we prove lower bounds for two classes of patterns which generalize A_k and B_k . The arguments in these special cases show up in the two cases (†) and (‡) of our main lower bound (Theorem 8.3).

Notation A.1. Recall Notation 5.12 for s -shifted pattern graphs $G^{\triangleright s}$ and pathsets $\mathcal{A}^{\triangleright s}$. For a pattern A and integer s , we define the s -shifted pattern $A^{\triangleright s}$ analogously by replacing each label $v_i v_{i+1}$ with the label $v_{i+s} v_{i+s+1}$.

Definition A.2 (Patterns A_k and B_k). We define patterns A_k and B_k for all $k \geq 1$ by the following induction. Let $A_1 = B_1 :=$ the atomic pattern labeled by $v_0 v_1$. For $k \geq 2$, let $A_k := \{A_j, A_{k-j}^{\triangleright j}\}$ where $j = \lceil k/2 \rceil$, and let $B_k := \{B_{k-1}, B_{k-1}^{\triangleright 1}\}$. For example, the explicit pictures of A_8 and B_4 are:



Intuitively, the pattern A_k corresponds to the recursive doubling algorithm for $DISTCONN(k, n)$. Note that we have essentially already encountered this pattern in the proof of Proposition 5.11 (our upper bound for χ_{P_k}). In fact, this proof shows:

Corollary A.3. For all $\mathcal{A} \in \mathcal{P}_{A_k}$, $\bar{\chi}_{A_k}(\mathcal{A}) \leq O(kn^{(1/2)\lceil \log k \rceil + 2})$. \square

The pattern B_k has a different nature than A_k . Whereas sub-patterns A_j and $A_{k-j}^{\triangleright j}$ of A_k overlap at only a single vertex v_j , sub-patterns B_{k-1} and $B_{k-1}^{\triangleright 1}$ of B_k overlap to the maximum possible extent. Despite this difference, it turns out that there is also a reasonable upper bound for $\bar{\chi}_{B_k}$.

Proposition A.4. For all $\mathcal{B} \in \mathcal{P}_{B_k}$, $\bar{\chi}_{B_k}(\mathcal{B}) \leq 2^k n^{\ln(k+1)+1}$.

Proof. We present a similar argument to the proof of Proposition 5.11. For all $k \geq 1$, define $\mathcal{B}_k \in \mathcal{P}_{B_k}$ by

$$\mathcal{B}_k := \{x \in [n]^{V_k} : x_0, \dots, x_k \leq n^{1-1/(k+1)}\}.$$

We have $\mathcal{B}_{k-1} \bowtie \mathcal{B}_{k-1}^{\triangleright 1} = \{x \in [n]^{V_k} : x_0, \dots, x_k \leq n^{1-1/k}\}$. For all $1 \leq t_0, \dots, t_k \leq n^{1/k(k+1)}$, let

$$\text{Copy}_{t_0, \dots, t_k}(\mathcal{B}_{k-1} \bowtie \mathcal{B}_{k-1}^{\triangleright 1}) := \{x \in [n]^{V_k} : t_i - 1 < \frac{x_i}{n^{1-1/k}} \leq t_i \text{ for all } 0 \leq i \leq k\}.$$

Note that

$$\mathcal{B}_k = \bigcup_{1 \leq t_0, \dots, t_k \leq n^{1/k(k+1)}} \text{Copy}_{t_0, \dots, t_k}(\mathcal{B}_{k-1} \bowtie \mathcal{B}_{k-1}^{\triangleright 1}).$$

Using (sub-additivity) and (join rule), together with the invariance of $\bar{\chi}$ under coordinate-wise permutations of $[n]$ and under shifts, we have

$$\bar{\chi}_{B_k}(\mathcal{B}_k) \leq 2n^{1/k} \bar{\chi}_{B_{k-1}}(\mathcal{B}_{k-1}).$$

This recurrence, together with the base case $\bar{\chi}_{B_1}(\mathcal{B}_1) = n$, implies

$$\bar{\chi}_{B_k}(\mathcal{B}_k) \leq 2^{k-1} n^{1+(1/2)+\dots+(1/k)} \leq 2^k n^{\ln(k+1)}.$$

Noting that $[n]^{V_k}$ is covered by n copies of \mathcal{B}_k , we have $\bar{\chi}_{B_k}([n]^{V_k}) \leq 2^k n^{\ln(k+1)+1}$. The proposition then follows using (monotonicity). \square

In Appendix B we prove matching lower bounds for $\bar{\chi}_{A_k}$ and $\bar{\chi}_{B_k}$. In fact, these lower bounds apply to two classes of patterns which include A_k and B_k . While the upper bounds for $\bar{\chi}_{A_k}$ and $\bar{\chi}_{B_k}$ are quite similar, our lower bound arguments are significantly different. The arguments in these two special cases—a “one-sided” induction for $\bar{\chi}_{A_k}$ and a “balanced” induction using the AM-GM inequality for $\bar{\chi}_{B_k}$ —show up in the two cases (†) and (‡) of our general lower bound (Theorem 8.3). For this reason, the reader might find the results in Appendix B to be a helpful warm-up.

Remark A.5. The pathsets \mathcal{A}_k and \mathcal{B}_k which show up in the proofs of our upper bounds are of a particularly simple form: they are *rectangular* subsets of $[n]^{V_k}$. In Appendix C we discuss a notion of *rectangular pathset complexity* $\bar{\chi}^{\text{rect}}$. Proving lower bounds for $\bar{\chi}^{\text{rect}}$ turns out to be much easier than for $\bar{\chi}$. We present an example (the “palindrome pathset”) which illustrates the difficulty in attempting to generalize this easier lower bound to the non-rectangular setting.

B Lower Bound for $\bar{\chi}$: Special Cases

We prove easier special cases of our lower bound for $\bar{\chi}$ with respect to two classes of patterns which include the key examples A_k and B_k introduced in §A. Although the results of this appendix are not used in the main body of the paper, the arguments in the proof show up in the two cases (†) and (‡) of our general lower bound.

Definition B.1.

- (i) For a pattern A ,
 - let $\text{Ends}(A) \subseteq V_A$ denote the set of *endpoints* in G_A (i.e. vertices of in-degree or out-degree zero), and let $\text{Interior}(A) := V_A \setminus \text{Ends}(A)$ denote the set of *interior vertices* in G_A ,
 - let $\mathcal{I}(A)$ denote the set of *intervals* in G_A (i.e. subsets of V_A which are connected in G_A).

Note that $\ell_A = \max_{I \in \mathcal{I}(A)} |I| - 1$ and $\Delta_A = |\text{Ends}(A)| / 2$.

- (ii) The classes of *end-joining* and *fully connected* patterns are defined as follows:

- A is *end-joining* if no edge of P_k labels more than one leaf of A (equivalently, $E_{A_1} \cap E_{A_2} = \emptyset$ for all non-atomic sub-patterns $\{A_1, A_2\} \preceq A$),
- A is *fully connected* if $G_{A'}$ is connected (i.e. $\Delta_{A'} = 1$) for all sub-patterns $A' \preceq A$.

Note that patterns A_k and B_k are both fully connected, while only A_k is end-joining (for $k \geq 3$).

- (iii) Functions $\psi_A, \xi_A : \mathcal{P}_A \rightarrow \mathbb{R}$ are defined as follows:

— for end-joining patterns A ,

$$\psi_A(\mathcal{A}) := \tilde{n}^{\frac{1}{2}(\log(\ell_A) + \Delta_A)} \sqrt{\mathbb{E}_{z \in [n]^{\text{Ends}(A)}} [\delta(\mathcal{A}|_{\text{Interior}(A)}^z)^2]},$$

— for fully connected patterns A ,

$$\xi_A(\mathcal{A}) := \max_{I \in \mathcal{I}(A)} \tilde{n}^{\frac{1}{4}(\log(|I|+1) + |I \cap \text{Ends}(A)|)} \cdot \pi_I(\mathcal{A}).$$

For non-end-joining patterns A , we set $\psi_A(\mathcal{A}) := 0$, and for non-fully connected patterns A , we set $\xi_A(\mathcal{A}) := 0$.

Proposition B.2. *Both ψ and ξ are lower bounds on pathset complexity. That is, for every pattern A and pathset $\mathcal{A} \in \mathcal{P}_A$, we have $\psi_A(\mathcal{A}) \leq \bar{\chi}_A(\mathcal{A})$ and $\xi_A(\mathcal{A}) \leq \bar{\chi}_A(\mathcal{A})$. In particular,*

$$\begin{aligned} \bar{\chi}_{A_k}([n]^{V_k}) &\geq \psi_{A_k}([n]^{V_k}) \geq \tilde{n}^{\frac{1}{2}(\log(k)+1)} \geq n^{\frac{1}{2} \log k}, \\ \bar{\chi}_{B_k}([n]^{V_k}) &\geq \xi_{B_k}([n]^{V_k}) \geq \tilde{n}^{\frac{1}{4}(\log(k+1)+2)} \geq n^{\frac{1}{4} \log k}. \end{aligned}$$

Recall from Remark 8.4 the dual characterization of $\bar{\chi}$ as the unique coordinate-wise maximal function from pairs (A, \mathcal{A}) to \mathbb{R} which satisfies inequalities (base case), (monotone), (sub-additive) and (join rule). It is easy to see that ψ and ξ satisfy the first three of these inequalities. To prove Proposition B.2, it suffices to show that ψ and ξ also satisfy inequality (join rule). We begin with ψ .

Lemma B.3. *For every non-atomic end-joining pattern $C = \{A, B\}$ and small pathsets $\mathcal{A} \in \mathcal{P}_A^{\text{small}}$ and $\mathcal{B} \in \mathcal{P}_B^{\text{small}}$,*

$$\psi_C(\mathcal{A} \bowtie \mathcal{B}) \leq \max\{\psi_A(\mathcal{A}), \psi_B(\mathcal{B})\}.$$

Proof. Without loss of generality, assume that $\ell_A \geq \ell_B$. After making three observations, will show that $\psi_C(\mathcal{A} \bowtie \mathcal{B}) \leq \psi_A(\mathcal{A})$.

First, note that each connected component of $G_C (= G_A \cup G_B)$ is the union of at most $\Delta_A + \Delta_B - \Delta_C + 1$ components of G_A and G_B . It follows that $\ell_C \leq (\Delta_A + \Delta_B - \Delta_C + 1) \cdot \ell_A$.

Since C is end-joining, $\text{Ends}(C)$ is the symmetric difference of $\text{Ends}(A)$ and $\text{Ends}(B)$. By the Cauchy-Schwartz inequality,

$$\begin{aligned} &\mathbb{E}_{c \in [n]^{\text{Ends}(C)}} [\delta((\mathcal{A} \bowtie \mathcal{B})|_{\text{Interior}(C)}^c)^2] \\ &= \mathbb{E}_{\substack{x \in [n]^{\text{Ends}(A) \setminus \text{Ends}(B)} \\ y \in [n]^{\text{Ends}(B) \setminus \text{Ends}(A)}}} \left[\left(\mathbb{E}_{z \in [n]^{\text{Ends}(A) \cap \text{Ends}(B)}} [\delta(\mathcal{A}|_{\text{Interior}(A)}^{xz}) \cdot \delta(\mathcal{B}|_{\text{Interior}(B)}^{yz})] \right)^2 \right] \\ &\leq \mathbb{E}_{\substack{x \in [n]^{\text{Ends}(A) \setminus \text{Ends}(B)} \\ z \in [n]^{\text{Ends}(A) \cap \text{Ends}(B)}}} [\delta(\mathcal{A}|_{\text{Interior}(A)}^{xz})^2] \mathbb{E}_{\substack{y \in [n]^{\text{Ends}(B) \setminus \text{Ends}(A)} \\ z \in [n]^{\text{Ends}(A) \cap \text{Ends}(B)}}} [\delta(\mathcal{B}|_{\text{Interior}(B)}^{yz})^2] \\ &= \mathbb{E}_{a \in [n]^{\text{Ends}(A)}} [\delta(\mathcal{A}|_{\text{Interior}(A)}^a)^2] \mathbb{E}_{b \in [n]^{\text{Ends}(B)}} [\delta(\mathcal{B}|_{\text{Interior}(B)}^b)^2]. \end{aligned}$$

We next note that B -smallness of \mathcal{B} implies

$$\mathbb{E}_{b \in [n]^{\text{Ends}(B)}} [\delta(\mathcal{B}|_{\text{Interior}(B)}^b)^2] \leq \mathbb{E}_{b \in [n]^{\text{Ends}(B)}} [\delta(\mathcal{B}|_{\text{Interior}(B)}^b)] = \delta(\mathcal{B}) \leq \tilde{n}^{-\Delta_B}.$$

Putting these inequalities together, we have

$$\begin{aligned}
\psi_C(\mathcal{A} \bowtie \mathcal{B}) &= \tilde{n}^{\frac{1}{2}} \left(\log(\ell_C) + \Delta_C \right) \sqrt{\mathbb{E}_{c \in [n]^{\text{Ends}(C)}} \left[\delta(\mathcal{A} \bowtie \mathcal{B})_{\text{Interior}(C)}^c \right]^2} \\
&\leq \tilde{n}^{\frac{1}{2}} \left(\log(\ell_A) + \log(\Delta_A + \Delta_B - \Delta_C + 1) + \Delta_C - \Delta_B \right) \sqrt{\mathbb{E}_{a \in [n]^{\text{Ends}(A)}} \left[\delta(\mathcal{A})_{\text{Interior}(A)}^a \right]^2} \\
&= \tilde{n}^{\frac{1}{2}} \left(\log(\Delta_A + \Delta_B - \Delta_C + 1) + \Delta_C - \Delta_B - \Delta_A \right) \cdot \psi_A(\mathcal{A}) \\
&\leq \psi_A(\mathcal{A})
\end{aligned}$$

using the fact that $\log(s+1) \leq s$ for every integer $s \geq 0$. We get $\psi_C(\mathcal{A} \bowtie \mathcal{B}) \leq \max\{\psi_A(\mathcal{A}), \psi_B(\mathcal{B})\}$ as required. \square

We next show that ξ satisfies inequality (join rule).

Lemma B.4. *For every non-atomic fully connected pattern $C = \{A, B\}$ and small pathsets $\mathcal{A} \in \mathcal{P}_A^{\text{small}}$ and $\mathcal{B} \in \mathcal{P}_B^{\text{small}}$,*

$$\xi_C(\mathcal{A} \bowtie \mathcal{B}) \leq \frac{\xi_A(\mathcal{A}) + \xi_B(\mathcal{B})}{2}.$$

Proof. Fix $I \in \mathcal{I}(A)$ such that

$$\xi_C(\mathcal{A} \bowtie \mathcal{B}) = \tilde{n}^{\frac{1}{4}} \left(\log(|I|+1) + |I \cap \text{Ends}(C)| \right) \pi_I(\mathcal{A} \bowtie \mathcal{B}).$$

We consider various cases depending on $|I \cap \text{Ends}(C)| \in \{0, 1, 2\}$. The most important case is where $|I \cap \text{Ends}(C)| = 2$ (i.e. I contains both endpoints of G_C). Because G_C is connected, this means that $I = V_C (= V_A \cup V_B)$ and hence $\pi_I(\mathcal{A} \bowtie \mathcal{B}) = \delta(\mathcal{A} \bowtie \mathcal{B})$.

Within this case, the most important sub-case is where $|E_A|, |E_B| \geq \frac{1}{2}|E_C|$. In this sub-case, we argue as follows. Without loss of generality, $V_C = \{v_0, \dots, v_k\}$ (i.e. G_C is the path P_k) and $v_0 \in V_A$ and $v_k \in V_B$. Let $j = \lfloor \frac{k-1}{2} \rfloor$ and $J = \{v_0, \dots, v_j\}$ and $K = \{v_{k-j}, \dots, v_k\}$ and note that $J \in \mathcal{I}(A)$ and $K \in \mathcal{I}(B)$. Since $v_0 \in J \cap \text{Ends}(A)$ and $v_k \in K \cap \text{Ends}(B)$, we have

$$(18) \quad \log(k+2) \leq \log(|J|+1) + |J \cap \text{Ends}(A)| \quad \text{and} \quad \log(k+2) \leq \log(|K|+1) + |K \cap \text{Ends}(B)|.$$

Next, observe that $\delta(\mathcal{A} \bowtie \mathcal{B}) \leq \tilde{n}^{-1}$ by C -smallness of $\mathcal{A} \bowtie \mathcal{B}$. We also have the bound $\delta(\mathcal{A} \bowtie \mathcal{B}) \leq \pi_J(\mathcal{A}) \cdot \pi_K(\mathcal{B})$ (since $J \cap K = \emptyset$). Taking the geometric mean of these two inequalities, we have

$$\delta(\mathcal{A} \bowtie \mathcal{B}) \leq \tilde{n}^{-1/2} \sqrt{\pi_J(\mathcal{A}) \cdot \pi_K(\mathcal{B})}.$$

Putting these pieces together, we have

$$\begin{aligned}
\xi_C(\mathcal{A} \bowtie \mathcal{B}) &= \tilde{n}^{\frac{1}{4}} \left(\log(k+2) + 2 \right) \delta(\mathcal{A} \bowtie \mathcal{B}) \\
&\leq \tilde{n}^{\frac{1}{4} \log(k+2)} \sqrt{\pi_J(\mathcal{A}) \cdot \pi_K(\mathcal{B})} \\
&\leq \frac{1}{2} \left(\tilde{n}^{\frac{1}{4} \log(k+2)} \pi_J(\mathcal{A}) + \tilde{n}^{\frac{1}{4} \log(k+2)} \pi_K(\mathcal{B}) \right) && \text{(AM-GM ineq.)} \\
&\leq \frac{1}{2} \left(\tilde{n}^{\frac{1}{4} \left(\log(|J|+1) + |J \cap \text{Ends}(A)| \right)} \pi_J(\mathcal{A}) + \tilde{n}^{\frac{1}{4} \left(\log(|K|+1) + |K \cap \text{Ends}(B)| \right)} \pi_K(\mathcal{B}) \right) && \text{(by (18))} \\
&\leq \frac{1}{2} \left(\xi_A(\mathcal{A}) + \xi_B(\mathcal{B}) \right) && \text{(ind. hyp.)}
\end{aligned}$$

In all other cases (i.e. when $|I \cap \text{Ends}(C)| < 2$ or $\min\{|E_A|, |E_B|\} < |E_C|/2$), the inequality is proved by finding $J \in \mathcal{I}(A)$ or $K \in \mathcal{I}(B)$ such that $|I \cap \text{Ends}(C)| < |J \cap \text{Ends}(A)|$ or $|K \in \text{Ends}(B)|$. We omit the analysis of these cases, since the arguments are not relevant to our main pathset complexity lower bound. \square

Having shown that ψ and ξ both satisfying (join rule), the proof of Proposition B.2 is complete. Combining our upper and lower bounds for $\bar{\chi}_{A_k}$ and $\bar{\chi}_{B_k}$ (Corollary A.3 and Propositions A.4 and B.2), we have

Corollary B.5. *With respect to patterns A_k and B_k , the pathset complexity of the complete P_k -pathset $[n]^{V_k}$ has the following bounds:*

$$\begin{aligned} n^{\frac{1}{2} \log k - O(1)} &\leq \bar{\chi}_{A_k}([n]^{V_k}) \leq kn^{\frac{1}{2} \log k + O(1)}, \\ n^{\frac{1}{4} \log k - O(1)} &\leq \bar{\chi}_{B_k}([n]^{V_k}) \leq 2^k n^{\ln k + O(1)}. \end{aligned}$$

Since $\Phi_{B_k} = \frac{1}{2} \log k - O(1)$ (as noted in §11.1), Theorem 8.3 gives the stronger lower bound $\bar{\chi}_{B_k}([n]^{V_k}) \geq \tilde{n}^{(1/2) \log k - O(1)} = n^{(1/2) \log k - O(1)}$. Even after extensively studying this special case, we were unable to narrow the gap between $\frac{1}{2} \log k$ and $\ln k$ ($\approx 0.69 \log k$) in the exponent of n in $\bar{\chi}_{B_k}([n]^{V_k})$.

C Rectangular Pathsets

A set $X \subseteq [n]^V$ is *rectangular* if there exist sets $S_i \subseteq [n]$, $i \in V$, such that $X = \{x \in [n]^V : x_i \in S_i \text{ for all } i \in V\}$. For a pattern graph G , let $\mathcal{R}_G = \{\mathcal{A} \in \mathcal{P}_G : \mathcal{A} \text{ is rectangular}\}$ and $\mathcal{R}_G^{\text{small}} = \mathcal{R}_G \cap \mathcal{P}_G^{\text{small}}$. For $\mathcal{A} \in \mathcal{R}_G$, we define *rectangular pathset complexity* $\bar{\chi}_G^{\text{rect}}(\mathcal{A})$ exactly like pathset complexity $\bar{\chi}_G(\mathcal{A})$ (Definition 5.6) except with \mathcal{R}_G and $\mathcal{R}_G^{\text{small}}$ replacing \mathcal{P}_G and $\mathcal{P}_G^{\text{small}}$. Analogously, we define $\bar{\chi}_A^{\text{rect}}(\mathcal{A})$ for patterns A . Note that $\bar{\chi}_A(\mathcal{A}) \leq \bar{\chi}_A^{\text{rect}}(\mathcal{A})$ for all $\mathcal{A} \in \mathcal{R}_A$.

Remark C.1. I venture to guess that $\bar{\chi}_A(\mathcal{A}) = \bar{\chi}_A^{\text{rect}}(\mathcal{A})$ for all $\mathcal{A} \in \mathcal{R}_A$, but have no idea how to prove this.

We have remarked that our upper bounds on $\bar{\chi}_{A_k}$ and $\bar{\chi}_{B_k}$ (Corollary A.3 and Proposition A.4) involved only rectangular pathsets. It follows that the same upper bounds apply to $\bar{\chi}_{A_k}^{\text{rect}}$ and $\bar{\chi}_{B_k}^{\text{rect}}$.

As for lower bounds on $\bar{\chi}^{\text{rect}}$, this turns out to be significantly easier than our lower bound for $\bar{\chi}$. Similar to our lower bound for fully connected patterns in Appendix B, we can lower bound $\bar{\chi}_G^{\text{rect}}(\mathcal{A})$ for all $\mathcal{A} \in \mathcal{R}_G$ in terms of the projection densities $\pi_S(\mathcal{A})$ where $S \in \mathcal{I}(G)$ via a function similar to ψ .

A key difference when it comes rectangular pathsets is that $\pi_S = \mu_S$ (projection density = maximum restriction density) and hence smallness of rectangular pathsets is preserved under projections to a union of components (cp. Remark 10.6). This fact turns out to greatly simplify the task of proving a lower bound for $\bar{\chi}^{\text{rect}}$.

The next example shows that projections of non-rectangular pathsets can be tricky. This illustrates the difficulty in generalizing the lower bound for $\bar{\chi}^{\text{rect}}$ to the non-rectangular setting.

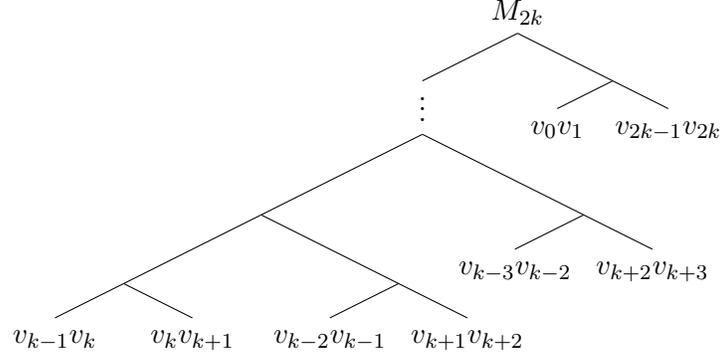
Example C.2. For $k \geq 1$, let $\mathcal{Pal}_{2k} \in \mathcal{P}_{P_{2k}}$ be the “palindrome pathset”

$$\mathcal{Pal}_{2k} = \{x \in [n]^{0, \dots, 2k} : x_{k-i} = x_{k+i} \text{ for all } 0 \leq i \leq k\}.$$

The palindrome pathset \mathcal{Pal}_{2k} has low density, while having the maximum projection over vertices v_0, \dots, v_k :

$$\delta(\mathcal{Pal}_{2k}) = n^{-k} \quad \text{and} \quad \pi_{\{v_0, \dots, v_k\}}(\mathcal{Pal}_{2k}) = 1.$$

It turns out that \mathcal{Pal}_{2k} is inexpensive to construct, given the right pattern. Let M_{2k} be the pattern



It is easy to show that

$$\bar{\chi}_{M_{2k}}(\mathcal{Pal}_{2k}) \leq O(kn^2).$$

On the other hand, for any fully connected pattern C with graph P_{2k} (such as A_{2k} or B_{2k}), the lower bound of Appendix B implies

$$\xi_C(\mathcal{Pal}_{2k}) \geq \xi_C(\mathcal{Pal}_{2k}) \geq \tilde{n}^{\frac{1}{4}} \left(\log(|\{v_0, \dots, v_k\}| + 1) + |\{v_0, \dots, v_k\} \cap \text{Ends}(A)| \right) \cdot \pi_{\{v_0, \dots, v_k\}}(\mathcal{A}) = n^{\Omega(\log k)}.$$