


Depth-3 Circuit Lower Bounds for k -OV

Tameem Choudhury 

Department of Computer Science and Engineering, IIT Hyderabad, India

Karteek Sreenivasaiah 

Department of Computer Science and Engineering, IIT Hyderabad, India

Abstract

The 2-Orthogonal Vectors (2-OV) problem is the following: given two tuples A and B of n Boolean vectors, each of dimension d , decide if there exist vectors $u \in A$, and $v \in B$, such that u and v are orthogonal. This problem, and its generalization k -OV defined analogously for k tuples, are central problems in the area of fine-grained complexity. One of the major conjectures in fine-grained complexity is that k -OV cannot be solved by a randomised algorithm in $n^{k-\epsilon} \text{poly}(d)$ time for any constant $\epsilon > 0$.

In this paper, we are interested in unconditional lower bounds against k -OV, but for weaker models of computation than the general Turing Machine. In particular, we are interested in circuit lower bounds to computing k -OV by Boolean circuit families of depth 3 of the form OR-AND-OR, or equivalently, a *disjunction of CNFs*.

We show that for all $k \leq d$, any disjunction of t -CNFs computing k -OV requires size $\Omega((n/t)^k)$. In particular, when k is a constant, any disjunction of k -CNFs computing k -OV needs to use $\Omega(n^k)$ CNFs. This matches the brute-force construction, and for each fixed $k > 2$, this is the first unconditional $\Omega(n^k)$ lower bound against k -OV for a computation model that can compute it in size $O(n^k)$. Our results partially resolve a conjecture by Kane and Williams [17] (page 12, conjecture 10) about depth-3 AC^0 circuits computing 2-OV.

As a secondary result, we show an exponential lower bound on the size of $\text{AND} \circ \text{OR} \circ \text{AND}$ circuits computing 2-OV when d is very large. Since 2-OV reduces to k -OV by projections trivially, this lower bound works against k -OV as well.

2012 ACM Subject Classification Theory of computation \rightarrow Circuit complexity; Theory of computation \rightarrow Problems, reductions and completeness

Keywords and phrases fine grained complexity, k -OV, circuit lower bounds, depth-3 circuits

Digital Object Identifier 10.4230/LIPIcs.STACS.2024.25

Related Version *Previous Version:* <https://ecc.weizmann.ac.il/report/2023/014/#revision3>

Acknowledgements The authors would like to thank the anonymous referees for their valuable comments that helped improve the presentation of this paper in several respects.

1 Introduction

The area of fine-grained complexity is a branch of computational complexity that studies the complexity of functions with a finer lens than the usual approach that makes a coarse distinction between polynomial time and super-polynomial time. The area has been focused on functions in P that find uses in a variety of contexts. In the seminal paper by Vassilevska Williams and Williams [26], they show eight problems that are subcubic time equivalent to one another. Hence a truly subcubic time algorithm for any one of these problems will also imply a subcubic algorithm for the others.

The holy grail of computation complexity is to show *unconditional* lower bounds to resources used in computing an *explicit* function. Unfortunately, the state of affairs in terms of unconditional lower bounds for computation, in its full generality, is rather bleak. The best known unconditional lower bounds for the running time of computing an explicit function are



© Tameem Choudhury and Karteek Sreenivasaiah;
licensed under Creative Commons License CC-BY 4.0

41st International Symposium on Theoretical Aspects of Computer Science (STACS 2024).

Editors: Olaf Beyersdorff, Mamadou Moustapha Kanté, Orna Kupferman, and Daniel Lokshantov;

Article No. 25; pp. 25:1–25:17



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



merely linear. Even for functions such as SAT that do not have any polynomial time running algorithms till date, we do not know how to show super-linear lower bounds. We do know from the time hierarchy theorem¹ that there are languages in $\text{DTIME}(n^2)$ that are not in $\text{DTIME}(n^c)$ for any $c < 2$. However the languages constructed in a proof of the time hierarchy are not natural, and not as explicit as we would like. Results such as [26] and [7] that show equivalences among several important functions help in identifying candidate functions that might witness the time hierarchy theorem for their time class. One such candidate function for quadratic time² is the *2-Orthogonal Vectors problem*.

The 2-Orthogonal Vectors problem $2\text{-OV}_{n,d}$ is defined as follows: Given as input two tuples $A \subseteq \{0,1\}^d$ and $B \subseteq \{0,1\}^d$ of n vectors each, decide if there is a vector $a \in A$ and a vector $b \in B$ such that a and b are orthogonal. To define a generalization of this problem, we think of each vector from $\{0,1\}^d$ as a characteristic vector of a subset from $[d]$. Then a natural generalization of $2\text{-OV}_{n,d}$ is the problem $k\text{-OV}_{n,d}$ that takes as input k tuples $A_1, A_2, \dots, A_k \subseteq \{0,1\}^d$ of n vectors each, and the task is to decide if there exists vectors $a_1 \in A_1, a_2 \in A_2, \dots, a_k \in A_k$ such that $a_1 \cap a_2 \cap \dots \cap a_k = \emptyset$. The problems 2-OV and $k\text{-OV}$ have emerged as central problems in fine-grained complexity. An important hypothesis is that no deterministic, or randomized, algorithm computing $2\text{-OV}_{n,d}$ can run in time $O(n^{2-\epsilon} \text{poly}(d))$ for any $\epsilon > 0$. This is essentially saying that the brute force algorithm is also the best. Interestingly, Ryan Williams in [24], shows that under the *strong exponential time hypothesis* (SETH)³, 2-OV (3-OV) requires $n^{2-o(1)}$ time ($n^{3-o(1)}$ time respectively).

In the absence of techniques that can show unconditional lower bounds, two natural directions of research emerge: (i) Conditional lower bounds to help us understand connections between various such problems, and “bottlenecks” to better algorithms. (ii) Unconditional lower bounds for weaker models of computation.

The first line of research has seen a tremendous body of results. There are numerous fine-grained reductions, and lower bounds, conditioned on SETH, and the hardness of functions such as $2\text{-OV}_{n,d}$, and $k\text{-OV}_{n,d}$. In the 2018 survey [25], Vassilevska Williams aptly describes it as “*an explosion of hardness results based on OV*”, and lists nineteen problems whose complexity is connected to that of $k\text{-OV}$. The fact that better algorithms for so many problems would imply better algorithms for $k\text{-OV}$, is perhaps not surprising. Intuitively, the $k\text{-OV}$ function looks “canonical” in a certain sense, and has managed to hide itself inside several other problems that look quite different at the surface. These include seemingly unrelated problems such as Longest Common Subsequence [1], Edit Distance [2], Fréchet distance [4, 5], Regular Expressions Matching [3], to name a few. Their survey [25] is an excellent source for those looking for a thorough treatment of fine-grained complexity, and in particular, this line of research.

The second direction, of showing lower bounds against weaker models of computation, seems to be lacking the same attention. To the best of our knowledge, the only paper that addresses this line is that of Kane and Williams [17]. In their paper they show tight lower bounds for formulas and branching programs computing 2-OV . We do not know any non-trivial lower bounds for computing 2-OV by models stronger than branching programs.

Note that if a uniform circuit family of bounded fan-in, and size $O(s(n, d))$ computes $k\text{-OV}_{n,d}$, then an algorithm that simply evaluates the circuit, computes $k\text{-OV}_{n,d}$ in time $\tilde{O}(s(n, d))$. So if the $k\text{-OV}$ hypothesis is true, then we can expect any uniform circuit family computing $k\text{-OV}_{n,d}$ to have size $\Omega(n^k)$. This begs the question:

¹ Such hierarchy theorems go through for the unit cost RAM model as well.

² We are being imprecise here so as to remain informal. The input length of $2\text{-OV}_{n,d}$ is actually nd . So “quadratic in n ” is not the same as $\text{DTIME}(n^2)$

³ [15],[6] For every $\epsilon > 0$, $\exists k$ such that $k\text{-SAT}$ problem on n variables cannot be solved in $O(2^{(1-\epsilon)n})$ time

What is the largest class of circuits for which we can show $\Omega(n^k \text{ poly}(d))$ size lower bounds against computing $k\text{-OV}_{n,d}$?

One class of Boolean circuits that has been extensively studied in terms of lower bounds is AC^0 (gates from $\{\wedge, \vee, \neg\}$, unbounded fan-in, $O(1)$ -depth). In fact we know exponential lower bounds against this class of circuits. So a good target would be to show that $k\text{-OV}_{n,d}$ requires AC^0 circuits of size $\Omega(n^k \text{ poly}(d))$. We note that $k\text{-OV}_{n,d}$ can indeed be computed by depth-3 AC^0 circuits of size $n^k d$, as shown later in equation 2. Can we show matching lower bounds?

The best known lower bound against depth-3 AC^0 circuits is $2^{\Omega(\sqrt{n})}$ for computing majority. This bound can be obtained by several classic techniques from the 80s including the switching lemma by Håstad [13], the polynomial method by Razborov [21] and Smolensky [22], and finite-limit vectors by [14]. One of the most important problems in circuit complexity is to prove $2^{\omega(n/\log \log n)}$ lower bounds to the size of depth-3 AC^0 circuits computing an explicit function. This would imply superlinear lower bounds against $O(\log n)$ depth circuits (of bounded fan-in) due to the depth reduction procedure described by Valiant [23] (alternatively, see Chapter 11 of Jukna [16]). With the aim of making progress on this front, Goldreich and Wigderson proposed a new framework in [11] where they define a new model of arithmetic circuits that use *multilinear gates*, as opposed to allowing gates computing sum or product alone, and a new complexity measure on this model. The main motivation being that lower bounds to their complexity measure implies lower bounds to a specific class of Boolean depth-3 circuits that they call *D-canonical*. The best lower bounds obtained for this class of depth-3 Boolean circuits, using their framework, is $\Omega(2^{n^{3/5}})$ by Goldreich and Tal [10]. In fact, the brute force depth-3 AC^0 circuits computing the negation of $k\text{-OV}$, described later in equation 3, bears close resemblance to D-canonical circuits since it is a product of set-multilinear functions, but over the Boolean algebra, as opposed to $\text{GF}(2)$.

More recently, the status of depth-3 $\text{AC}^0[\oplus]$ circuits (gates computing xor are allowed in addition to the usual \wedge, \vee, \neg) got an update. The lower bound for computing majority using depth-3 $\text{AC}^0[\oplus]$ circuits was improved from $2^{\Omega(n^{1/4})}$ to $2^{\Omega(\sqrt{n})}$ by Oliveira, Santhanam and Srinivasan [20]. This closed the gap between upper and lower bounds up to a logarithmic factor in the exponent.

While techniques such as the switching lemma and the polynomial method work in a “bottom-up” fashion, the techniques in [14] is a “top-down” approach specifically for depth-3 AC^0 circuits. To the best of our knowledge, the only top-down strategies for circuit lower bounds are the *Karchmer-Wigderson game* by Karchmer and Wigderson [18], the *discriminator lemma* for depth-2 threshold circuits by Hajnal, Masse, Pudlák, Szegedy, Turán [12], and *finite-limits* by Håstad, Jukna, Pudlak [14]. Our results in this paper can be seen as a non-trivial application of the techniques of Håstad, Jukna, Pudlak [14].

Kane and Williams [17] conjecture that any depth-3 AC^0 circuit computing $2\text{-OV}_{n,d}$ requires $\Omega(n^2)$ wires (see page 12, conjecture 10 in [17]). Observe that $2\text{-OV}_{n,d}$ (and $k\text{-OV}_{n,d}$) can be computed by $\text{OR} \circ \text{AND} \circ \text{OR}$ circuits with $2n^2 d$ wires (and $kn^k d$ wires respectively):

$$2\text{-OV}_{n,d}(A, B) = \bigvee_{i_1, i_2 \in [n]} \bigwedge_{j \in [d]} (\neg a_{i_1}[j] \vee \neg b_{i_2}[j]) \quad (1)$$

$$k\text{-OV}_{n,d}(A_1, \dots, A_k) = \bigvee_{i_1, \dots, i_k \in [n]} \bigwedge_{j \in [d]} (\neg a_{i_1}[j] \vee \dots \vee \neg a_{i_k}[j]) \quad (2)$$

Hence, informally, their conjecture for $2\text{-OV}_{n,d}$, and by extension $k\text{-OV}_{n,d}$, is that the brute-force circuit is also the best.

A second important question in [17] is about generalizing lower bounds from 2-OV to k -OV. As they have noted, generalizing their lower bounds to $k > 2$ would beat the state of the art in branching program lower bounds. Our results for depth-3 AC^0 circuits generalize to $k > 2$, and scale well when the bottom fan-in is bounded.

Our Results

In this paper, we show lower bounds against the size of certain classes of depth-3 AC^0 circuit families computing k -OV $_{n,d}$. Our main result shows lower bounds against a restricted class of $OR \circ AND \circ OR$ circuits computing k -OV $_{n,d}$, while our secondary result deals with $AND \circ OR \circ AND$ circuits computing a special case of k -OV $_{n,d}$. Our main result is the following:

► **Theorem 1.** *For all $k \leq d$, any $OR \circ AND \circ OR$ circuit with bottom fan-in t computing k -OV $_{n,d}$ requires top fan-in $\Omega\left(\left(\frac{n}{t}\right)^k\right)$.*

Circuit families of the type $OR \circ AND \circ OR$ can also be understood as a *disjunction of CNFs*. Therefore Theorem 1 is equivalent to the following statement:

“Any disjunction of t -CNFs computing k -OV $_{n,d}$ requires size $\Omega\left((n/t)^k\right)$.”

(Here, by ‘ t -CNF’, we mean a CNF whose clauses have at most t literals, and by ‘size’ we mean the number of CNFs being used.)

The brute-force circuit described earlier in equation 2 for k -OV $_{n,d}$, is a disjunction of n^k many k -CNFs, and the lower bound from Theorem 1 for this model is $\Omega\left((n/k)^k\right)$. Hence for all constant $k > 1$, the complexity of computing k -OV $_{n,d}$ as a disjunction of k -CNFs is $\Theta(n^k)$.

The proof technique used for Theorem 1 actually goes through for a more general class of depth-3 circuits where the bottom gates can have arbitrary fan-in as long as the number of negated literals among their inputs is at most t . We describe this in the next subsection. The more general theorem is the following. Let \mathcal{C}_t^- be the set of all unate functions (see Definition 7) that are negative unate on at most t variables.

► **Theorem 2.** *For all $k \leq d$, any $OR \circ AND \circ \mathcal{C}_t^-$ circuit computing k -OV $_{n,d}$ requires top fan-in $\Omega\left(\left(\frac{n}{t}\right)^k\right)$.*

It is important to note that the usual trick of using random restrictions to get rid of the bottom fan-in restriction in Theorem 1 is very unlikely to work as it is known that 2-OV becomes easy to compute by AC^0 circuits with high probability under random restrictions [17] (section 3).

As a secondary result, we show an exponential lower bound on the size of $AND \circ OR \circ AND$ circuits computing 2-OV $_{n,d}$ when d is very large:

► **Theorem 3.** *For all $\ell \leq d$, any $AND \circ OR \circ AND$ circuit computing 2-OV $_{n,d}$ requires size $s \in \Omega\left(\min\left\{2^\ell, \left(\frac{d}{n\ell}\right)^n\right\}\right)$. In particular, for $\ell = d/2n$ and $d \in \Omega(n^2)$, $s \in \Omega(2^n)$.*

Since 2-OV $_{n,d}$ reduces to k -OV $_{n,d}$ by projections trivially, the above theorem holds for k -OV $_{n,d}$ as well. It must also be noted that the input size for 2-OV $_{n,d}$ is $2nd$ which, for the choice of d in Theorem 3, is $2n^3$. Hence with respect to an input size of n , the lower bound for 2-OV $_{n,d}$ from Theorem 3 is actually $2^{\Omega(n^{1/3})}$.

An important fact to be noted about depth-3 AC^0 circuits is that, in general, the computational power of $OR \circ AND \circ OR$ circuits and $AND \circ OR \circ AND$ circuits are incomparable. As demonstrated by [14], the *iterated intersection* function on $2n$ variables (see Definition 26)

is computable by $\text{AND} \circ \text{OR} \circ \text{AND}$ circuits of linear size, but any $\text{OR} \circ \text{AND} \circ \text{OR}$ circuit family computing it requires size $2^{\Omega(\sqrt{n})}$. A more thorough discussion on this topic can be found in Chapter 11.5 of Jukna [16]. This is true in the context of $\text{k-OV}_{n,d}$ and our bounds also: an idea used by Kane and Williams (Proposition 4 in [17]) can be used to show $\text{AND} \circ \text{OR} \circ \text{AND}$ circuits computing $\text{k-OV}_{n,d}$ that are smaller than our lower bounds in Theorem 1 for $\text{OR} \circ \text{AND} \circ \text{OR}$ circuits when $d \in O(1)$.

We show in Section 5 a general construction for $\text{k-OV}_{n,d}$ that achieves a trade-off between top fan-in and bottom fan-in. This shows that in general, for circuits with bottom fan-in t our lower bound against the top fan-in of $\text{OR} \circ \text{AND} \circ \text{OR}$ circuits computing $\text{k-OV}_{n,d}$ is at least a factor of t^{k-1}/k away from the corresponding upper bound.

Techniques

We note that throughout this paper, we work with the function $\text{k-Int}_{n,d}$ defined as the negation of $\text{k-OV}_{n,d}$. We do this because $\text{k-Int}_{n,d}$ is a monotone function, and hence allows us several conveniences with regard to notation. Thus our lower bounds to $\text{AND} \circ \text{OR} \circ \text{AND}$ circuits computing $\text{k-Int}_{n,d}$ transfer directly to $\text{OR} \circ \text{AND} \circ \text{OR}$ circuits computing $\text{k-OV}_{n,d}$. More formally, $\text{k-Int}_{n,d}$ is defined as

$$\text{k-Int}_{n,d}(A_1, \dots, A_k) = \bigwedge_{i_1, \dots, i_k \in [n]} \bigvee_{j \in [d]} (a_{i_1}[j] \wedge \dots \wedge a_{i_k}[j]) \quad (3)$$

Main result. For our main result, the strategy we use is that of *finite limit vectors*. This is a top-down strategy that was used by Håstad, Jukna, and Pudlák in [14] for proving depth-3 AC^0 circuit lower bounds. We briefly describe the approach.

Assume an $\text{AND} \circ \text{OR} \circ \text{AND}$ circuit $C = C_1 \wedge \dots \wedge C_{s(n)}$ computes a function f . Then for any $\mathcal{N} \subseteq f^{-1}(0)$, by an averaging argument, there is a C_i that correctly outputs 0 on at least $1/s$ fraction of inputs in \mathcal{N} . Hence showing an upper bound to $|C_i^{-1}(0) \cap \mathcal{N}|$ implies a lower bound to $s(n)$ as $s \geq |\mathcal{N}|/|C_i^{-1}(0) \cap \mathcal{N}|$.

The technique of *finite limits* by [14] is used to show that C_i cannot be correct on many inputs in \mathcal{N} . The idea is to show that if $C_i^{-1}(0) \cap \mathcal{N}$ is large, then we can construct a 1-input y such that for any set of t input positions, it looks identical to *some string* in $C_i^{-1}(0) \cap \mathcal{N}$. Such a string y is called a *t-limit* for the set $C_i^{-1}(0) \cap \mathcal{N}$. Then if the bottom gates in C_i can each see only t bits of the input, the string y fools all of them into evaluating to 0 *simultaneously*, and hence C_i will output 0 on y . This is a contradiction since $y \in C^{-1}(1)$ by construction, but $C_i(y) = 0$ implies $C(y) = 0$. It is not hard to see that if the *t-limit* string y has the additional property that $y \geq x$ for all $x \in C_i^{-1}(0) \cap \mathcal{N}$, and each bottom gate in C_i has at most t positive literals among its inputs, the same argument goes through. We call such a y an *upper t-limit* for the set $C_i^{-1}(0) \cap \mathcal{N}$ (as opposed to the term ‘*lower t-limit*’ used in [14] for the case when $y \leq x$). We shall also use the term “bottom positive fan-in” to indicate how many of the input literals are allowed to be positive for each bottom gate.

The key idea behind our construction of a *t-limit* is to first model any subset of *maxterms* of $\text{k-Int}_{n,d}$ as a k -partite hypergraph such that the maxterms in the subset and the hyperedges are in bijection. We call this hypergraph as a “Support Graph”, and construct it in Section 3.2. Then we construct a *t-limit* for the case of $2\text{-Int}_{n,d}$ by using König’s theorem on this graph. To deal with the general case of $\text{k-Int}_{n,d}$, we first show a sunflower lemma on this support graph, and then use the sunflower structure to construct a *t-limit*. We show a version of the sunflower lemma on our hypergraph that is *very slightly* less demanding than the standard sunflower lemma [9]. We note that this does not improve the asymptotic complexity of our final bound.

We remark here that all t -limit strings that we construct in this paper are also upper t -limit strings. Hence all our lower bounds for k -Int $_{n,d}$ go through for the circuit class $\text{AND} \circ \text{OR} \circ \mathcal{C}_t^+$ where \mathcal{C}_t^+ is the set of all unate functions that are positive unate on at most t variables. Informally, this means that the bottom gates can compute any unate functions, have unbounded fan-in, but at most t of the inputs can be positive literals. (The dual statement for k -OV $_{n,d}$ is Theorem 2 stated in the previous section.) As an example, lower bounds using this technique will also work against depth-3 circuits where the top and middle layers are AND and OR respectively, and the bottom layer consists of homogeneous linear threshold functions, each of which is defined by a vector of weights that has at most t positive weights.

An important observation about the technique described above is that it is impervious to the fan-in of the middle OR gates. So we could use a suitable DNF for each bottom gate and convert an $\text{AND} \circ \text{OR} \circ \mathcal{C}_t^+$ circuit to an $\text{AND} \circ \text{OR} \circ \text{AND}$ circuit with bottom positive fan-in at most t and a possibly larger middle fan-in. Since the technique gives lower bounds to top fan-in regardless of middle fan-in, all lower bounds that we can derive against $\text{AND} \circ \text{OR} \circ \text{AND}$ circuits with bottom positive fan-in t using this technique, transfer to $\text{AND} \circ \text{OR} \circ \mathcal{C}_t^+$ without any change. Hence throughout this paper, we focus our attention to $\text{AND} \circ \text{OR} \circ \text{AND}$ circuits.

Secondary result. The exponential lower bound of [14] for $\text{OR} \circ \text{AND} \circ \text{OR}$ circuits computing the iterated intersection function $S_{n,d}$ for $d \in \sqrt{n}$ is of particular interest to us. The function $S_{n,d}$ bears a close resemblance to 2 -Int $_{n,d}$. While $S_{n,d}$ is the *iterated* intersection, 2 -Int $_{n,d}$ can be seen as “all-pairs” intersection.

We show a reduction (via projections) from $S_{n,d/n}$ to 2 -Int $_{n,d}$. The blow-up in the dimension of vectors is rather large, and we can conclude non-trivial lower bounds only for $d \in \omega(n)$.

2 Preliminaries

We often interpret a d -dimensional vector $u \in \{0,1\}^d$ as the characteristic vector of a subset of $[d]$.

► **Definition 4** (k -OV $_{n,d}$). For tuples $A_1, A_2, \dots, A_k \subseteq \{0,1\}^d$ where $\forall i \in [k], |A_i| = n$.

$$k\text{-OV}_{n,d}(A_1, A_2, \dots, A_k) = 1 \iff \exists a_1 \in A_1, \exists a_2 \in A_2, \dots, \exists a_k \in A_k, \text{ such that} \\ a_1 \cap a_2 \cap \dots \cap a_k = \emptyset$$

For notational convenience, we work with the negation of k -OV $_{n,d}$ throughout the paper. We use k -Int $_{n,d}$ to denote the negation of k -OV $_{n,d}$, and is defined as follows:

► **Definition 5** (k -Int $_{n,d}$). For tuples $A_1, A_2, \dots, A_k \subseteq \{0,1\}^d$ where $\forall i \in [k], |A_i| = n$.

$$k\text{-Int}_{n,d}(A_1, A_2, \dots, A_k) = 1 \iff \forall a_1 \in A_1, \forall a_2 \in A_2, \dots, \forall a_k \in A_k, \text{ we have} \\ a_1 \cap a_2 \cap \dots \cap a_k \neq \emptyset$$

An input to the function k -Int $_{n,d}$ has nk vectors, each of dimension d . Hence $nk d$ many input bits in total.

For any $x, y \in \{0,1\}^d$, we write $x \leq y$ if $\forall i, x_i \leq y_i$. Similarly, we write $x \oplus y$ to denote the string obtained by a point-wise xor between x and y .

► **Definition 6** (Monotone function). We say that a Boolean function f is monotone if $\forall x, y \in \{0, 1\}^d$ such that $x \leq y$, we have $f(x) \leq f(y)$.

The notion of monotone can be generalized to the notion of being *unate*:

► **Definition 7** (Unate function). A Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is unate if there exists a monotone Boolean function $g : \{0, 1\}^n \rightarrow \{0, 1\}$ and a string $s \in \{0, 1\}^n$ such that for all inputs x , we have $f(x) = g(x \oplus s)$.

Further, a unate function is *positive unate* (negative unate) on a variable x_i if $s_i = 0$ ($s_i = 1$ respectively).

For monotone functions such as $\text{k-Int}_{n,d}$, we can define *maximal 0-inputs*:

► **Definition 8** (Maximal 0-input). Let f be a monotone Boolean function. An input x is a maximal 0-input for f if $f(x) = 0$ and for all strings y such that $x < y$, $f(y) = 1$.

Throughout this article, we will use the term “*maxterm*” and “maximal 0-inputs” interchangeably. This deviates from the standard definition of *maxterm*, but is very convenient in our context.

For a vector $u \in \{0, 1\}^d$, and a set of indices $S \subseteq [d]$, we denote the restriction of u to the indices in S as $u|_S$.

► **Definition 9** (t -limit). A vector $y \in \{0, 1\}^m$ is said to be a t -limit for a set $B \subseteq \{0, 1\}^m$ if and only if $\forall S \subseteq [m]$ with $|S| = t$, $\exists x \in B$ such that $y \neq x$ but $y|_S = x|_S$. Further, $y \in \{0, 1\}^m$ is said to be an upper t -limit if $y \geq x$.

Note that if a string y is a t -limit for a set B and $B \subseteq B'$, then y is also a t -limit for B' .

We will be using König’s theorem in our proofs, which is stated as follows:

► **Proposition 10** ([19], [8]). The maximum cardinality of a matching in a bipartite graph \mathcal{G} is equal to the minimum cardinality of a vertex cover of its edges.

We will always assume that the depth-3 circuits we consider are layered. i.e., inputs are read directly by only the gates at the bottom layer, and every layer reads outputs from the layer below it. This assumption does not affect asymptotic complexity. We say a depth-3 circuit C has *bottom positive fan-in* (bottom negation fan-in) t if for every gate in the bottom layer, at most t of its inputs are positive literals (negated literals respectively).

We denote the permutation group on k distinct elements with S_k . Let $\mathcal{P} = (P_1, \dots, P_k)$ be an ordered partition of $[d]$ into k parts. For any permutation $\sigma \in S_k$, we use \mathcal{P}_σ to denote the ordered partition obtained by permuting the parts of \mathcal{P} using σ . i.e., $\mathcal{P}_\sigma \triangleq (P_{\sigma(1)}, \dots, P_{\sigma(k)})$

3 AND ◦ OR ◦ AND circuits

To describe the lower bound for $\text{k-Int}_{n,d}$ against AND ◦ OR ◦ AND circuits, we first identify a special set of maxterms (maximal 0-inputs) of $\text{k-Int}_{n,d}$. We do this by explicitly constructing such inputs.

3.1 Maxterms of $\text{k-Int}_{n,d}$

Fix any choice of integers $k, d \in \mathbb{N}$ such that $1 < k \leq d$. For any choice of $n_1, \dots, n_k \in [n]$, and any ordered partition $\mathcal{P} = (P_1, \dots, P_k)$ of $[d]$ into k parts, we will construct an input $N = (A_1, \dots, A_k)$ where $A_i \subseteq \{0, 1\}^d$ with $|A_i| = n$ such that N is a maxterm for $\text{k-Int}_{n,d}$. Throughout, we will denote the j ’th vector in A_i by a_i^j .

The input $N = (A_1, \dots, A_k) \in \{0, 1\}^{nkd}$ is constructed as follows:

- Set every vector other than $a_1^{n_1}, \dots, a_k^{n_k}$ to all 1s.
- In each $a_i^{n_i}$, set the indices contained in P_i to 0s. Set every other position to 1. Formally, for all $i \in [k]$, set $a_i^{n_i}|_{P_i} \leftarrow \vec{0}$ and $a_i^{n_i}|_{[d] \setminus P_i} \leftarrow \vec{1}$.

We shall call $((n_1, \dots, n_k), \mathcal{P})$ the *support* of N , and denote it by $\text{sup}(N)$.

To see that N is indeed a maxterm of $\mathbf{k}\text{-Int}_{n,d}$, observe that since \mathcal{P} is a partition of $[d]$, for every position $\ell \in [d]$, there is a *unique* $i \in [k]$ such that $\ell \in P_i$. Therefore, by construction of N , $a_i^{n_i}[\ell] = 0$. So for every position ℓ , there is some vector among $a_1^{n_1}, \dots, a_k^{n_k}$ that is 0 in position ℓ , and hence $a_1^{n_1} \cap \dots \cap a_k^{n_k} = \emptyset$. Moreover, due to i being unique for each such ℓ , we also have $a_j^{n_j}[\ell] = 1$ for all $j \neq i$. So changing $a_i^{n_i}[\ell]$ from 0 to 1 results in the vectors intersecting at ℓ . Combining this with the fact that every vector in N other than $a_1^{n_1}, \dots, a_k^{n_k}$ is the all-1s vector, we conclude that N is indeed a maximal 0-input.

We will be particularly interested in a subset of such maxterms of $\mathbf{k}\text{-Int}_{n,d}$ that are formed by the permutations of the parts of some fixed partition into non-empty parts. We define this formally as follows.

► **Definition 11** (Permutation-maxterms). *Fix an ordered partition $\mathcal{P} = (P_1, \dots, P_k)$ of $[d]$ into k non-empty parts. A permutation-maxterm with respect to \mathcal{P} is any maxterm N constructed as above that has $\text{sup}(N) = ((n_1, \dots, n_k), \mathcal{P}_\sigma)$ for some $n_1, \dots, n_k \in [n]$ and $\sigma \in \mathbf{S}_k$.*

We shall use $\mathcal{N}_{\mathcal{P}}^{n,k,d}$ to denote the set of all permutation-maxterms of $\mathbf{k}\text{-Int}_{n,d}$ with respect to some ordered partition \mathcal{P} of $[d]$ into k non-empty parts. We drop the subscript, and superscripts if it is clear from context.

Note that for any partition \mathcal{P} as in the definition above, $|\mathcal{N}_{\mathcal{P}}^{n,k,d}| = n^k k!$ as there are n^k many k -tuples (n_1, \dots, n_k) and $k!$ many permutations in \mathbf{S}_k .

► **Remark 12.** The proofs in this paper do not depend on the exact partition chosen. Any arbitrary ordered partition of $[d]$ into k non-empty parts will work. For a further simplification, one could assume $k = d$, and fix the permutation $\mathcal{P} = (P_1, \dots, P_k)$ to be $P_i = \{i\}$ for all $i \in [d]$.

3.2 Support Graph

We define a k -partite hypergraph to encode, and reason about, the relationship between permutation-maxterms of $\mathbf{k}\text{-Int}_{n,d}$. Here, by k -partite hypergraph we mean that every hyperedge must contain exactly one vertex from each part.

Fix $k \geq 2$ and $d \geq k$, and any ordered partition \mathcal{P} of $[d]$ into k non-empty parts. For any subset $S \subseteq \mathcal{N}_{\mathcal{P}}^{n,k,d}$ of permutation-maxterms of $\mathbf{k}\text{-Int}_{n,d}$, we define the *support graph* of S as a k -partite hypergraph $\mathcal{G}_S = (V_1 \cup \dots \cup V_k, E)$ such that each maxterm in S corresponds to a hyperedge in the graph. Recall that an input to $\mathbf{k}\text{-Int}_{n,d}$ consists of k tuples, each having n vectors of dimension d . In each V_i , we include a total of nk vertices as follows: for each $j \in [n]$, we include k vertices in V_i indexed as $v_i^{j,1}, \dots, v_i^{j,k}$. So for all $i \in [k]$, we have $|V_i| = nk$ and hence the graph \mathcal{G}_S is defined on nk^2 many vertices.

We define the set E of hyperedges as follows:

$$\left(v_1^{n_1, b_1}, \dots, v_k^{n_k, b_k} \right) \in E \iff \exists \text{ maxterm } N \in S \text{ such that} \\ \text{sup}(N) = ((n_1, \dots, n_k), \mathcal{P}_\sigma) \text{ and } b_i = \sigma(i) \forall i \in [k]$$

► **Remark 13.** Note that the set of maxterms $S \subseteq \mathcal{N}_{\mathcal{P}}$ and the set of hyperedges in \mathcal{G}_S are in bijection. More precisely, a maxterm N with $\text{sup}(N) = ((n_1, \dots, n_k), \mathcal{P}_\sigma)$ corresponds to the hyperedge $\left(v_1^{n_1, \sigma(1)}, \dots, v_k^{n_k, \sigma(k)} \right)$ and vice-versa.

► **Definition 14** (Co-disjoint). We call two vectors $u \in \{0, 1\}^d$ and $v \in \{0, 1\}^d$ as co-disjoint if and only if $\bar{u} \cap \bar{v} = \emptyset$. i.e., the set of positions where u is 0, and the set where v is 0 are disjoint.

For two tuples of vectors $A = (a_1, \dots, a_n)$ and $B = (b_1, \dots, b_n)$ where $a_i, b_i \in \{0, 1\}^d$, we say A and B are co-disjoint if for all $i \in [n]$, a_i and b_i are co-disjoint.

Maxterms $M = (M_1, \dots, M_k)$ and $N = (N_1, \dots, N_k)$, both from $\mathcal{N}_{\mathcal{P}}^{n,k,d}$, are said to be co-disjoint if and only if for all $i \in [k]$, M_i and N_i are co-disjoint.

Intuitively, the graph \mathcal{G}_S records where the 0s in each of the maxterms in S appear. This gives us the following close connection between co-disjointness of vectors across maxterms, and disjointness of their hyperedges.

► **Lemma 15.** Let $S \subseteq \mathcal{N}_{\mathcal{P}}^{n,k,d}$, and let $\mathcal{G}_S = (V_1 \cup \dots \cup V_k, E)$ be its support graph. Let $M = (M_1, \dots, M_k)$ and $N = (N_1, \dots, N_k)$ be two maxterms from S and let E_M , and E_N respectively, denote their corresponding hyperedges in \mathcal{G}_S . Then for each $i \in [k]$, we have the following two properties:

1. If E_M and E_N share a vertex in V_i , then $M_i = N_i$.
2. If E_M and E_N contain different vertices from V_i , then M_i and N_i are co-disjoint.

Proof. Let $\text{sup}(M) = (a_1, \dots, a_k, \mathcal{P}_\sigma)$ and $\text{sup}(N) = (b_1, \dots, b_k, \mathcal{P}_\pi)$.

Proof of (1). If E_M and E_N share a vertex in V_i for some $i \in [k]$, then $v_i^{a_i, \sigma(i)} = v_i^{b_i, \pi(i)}$ and so we have $a_i = b_i$ and $\sigma(i) = \pi(i)$. Let $\ell = a_i = b_i$, and let $q = \sigma(i) = \pi(i)$. Then by construction of the maxterms M and N , all vectors in M_i other than m_i^ℓ are all 1s, and similarly all vectors in N_i other than n_i^ℓ are all 1s. The vector m_i^ℓ and n_i^ℓ both have 0s in indices from the part P_q , and 1s elsewhere. So $m_i^\ell = n_i^\ell$. Hence the tuple M_i and N_i are identical.

Proof of (2). If E_M and E_N have different vertices from V_i , then $v_i^{a_i, \sigma(i)} \neq v_i^{b_i, \pi(i)}$. So either $a_i \neq b_i$ or $\sigma(i) \neq \pi(i)$ (or both). The claim holds in both cases:

- If $a_i \neq b_i$, then recall that by construction, the only vector that has 0s in M_i is the vector $m_i^{a_i}$. Every other vector in M_i , and in particular $m_i^{b_i}$ is the all 1s vector by construction. So the tuples of vectors M_i and N_i cannot both be 0 in any vector in any position.
- Else $a_i = b_i$ and $\sigma(i) \neq \pi(i)$. By our construction of maxterms, the 0s in the vectors $m_i^{a_i}$ and $n_i^{b_i=a_i}$ are in the indices given by $P_{\sigma(i)}$ and $P_{\pi(i)}$ respectively. Since \mathcal{P} is a partition, and $\sigma(i) \neq \pi(i)$, $P_{\sigma(i)} \cap P_{\pi(i)} = \emptyset$. Therefore there cannot be an index where both $m_i^{a_i}$ and $n_i^{b_i}$ are both 0. ◀

The following lemma follows directly from Lemma 15:

► **Lemma 16.** Let $S \subseteq \mathcal{N}_{\mathcal{P}}^{n,k,d}$ be a set of maxterms such that all hyperedges in \mathcal{G}_S are pairwise vertex-disjoint. Then the maxterms in S are pairwise co-disjoint. (i.e., for all positions $\ell \in [nkd]$, there is at most one maxterm in S that has 0 in the ℓ 'th position.)

Proof. Let $M, N \in S$ be any two maxterms, and let the vertex set of \mathcal{G}_S be $V = V_1 \cup \dots \cup V_k$. The hyperedges E_M and E_N , corresponding to M , and N respectively, are vertex-disjoint from the premise. So for each $i \in [k]$, E_M and E_N contain different vertices from V_i . Applying Lemma 15 to \mathcal{G}_S , we obtain that M_i and N_i are co-disjoint for all $i \in [k]$. Hence there is no position where both M and N are 0 by definition of co-disjoint. ◀

3.3 Warm-up: 2-Int $_{n,d}$

We give a self-contained proof of our lower bound for the case of 2-Int $_{n,d}$ that demonstrates the strategy behind the proof for the general case.

► **Theorem 17.** *For all $d > 1$, any AND \circ OR \circ AND circuit with bottom fan-in t computing 2-Int $_{n,d}$ requires top fan-in at least $2n^2/t^2$.*

Proof. Let $C = C_1 \wedge C_2 \wedge \cdots \wedge C_s$ be an AND \circ OR \circ AND circuit with bottom fan-in t computing 2-Int $_{n,d}$. Let $\mathcal{P} = (P_1, P_2)$ be any ordered partition of $[d]$ into two non-empty parts. Consider the permutation-maxterms $\mathcal{N} = \mathcal{N}_{\mathcal{P}}^{n,2,d}$ of 2-Int $_{n,d}$ as described in definition 11. Since \mathcal{N} is a subset of the 0-inputs of 2-Int $_{n,d}$, the circuit C outputs 0 on every input in \mathcal{N} . By an averaging argument, there exists $i \in [s]$ such that C_i correctly outputs 0 on at least $1/s$ fraction of inputs in \mathcal{N} . We will show that $|C_i^{-1}(0) \cap \mathcal{N}| \leq t^2$. Then the theorem follows as:

$$\frac{2n^2}{s} = \frac{1}{s} |\mathcal{N}| \leq |C_i^{-1}(0) \cap \mathcal{N}| \leq t^2.$$

Let $S = C_i^{-1}(0) \cap \mathcal{N}$. Suppose, for the sake of contradiction, $|S| > t^2$. We will show later in the proof that in such a case, there is a t -limit $y \in C^{-1}(1)$ for S . This leads to a contradiction as follows: let $C_i = g_1 \vee g_2 \vee \cdots \vee g_\ell$ with each g_j having fan-in at most t . Then, by definition of t -limit, for all $T \subseteq [nkd]$ with $|T| = t$, there exists $x \in S$ such that $x|_T = y|_T$. Now each of the gates g_j is a function of at most t variables, and we know that for all inputs $x \in S$, we have $g_j(x) = 0$ for all $j \in [\ell]$. Since y looks identical to some string in S when restricted to these t positions, all the g_j will output 0 on y too. This forces $C_i(y) = 0$, but this cannot happen since $y \in C^{-1}(1)$. Hence it could not have been the case that $|S| > t^2$.

We now construct a t -limit for any $S \subseteq \mathcal{N}$ when $|S| > t^2$. Let $S \subseteq \mathcal{N}$ be any set with size $|S| > t^2$ and let \mathcal{G}_S be its support graph. Note that since $k = 2$, \mathcal{G}_S is a bipartite graph with simple edges rather than hyperedges, and every maxterm in S corresponds to an edge in \mathcal{G}_S and vice versa. We claim at least one of the following is true for \mathcal{G}_S :

- (i) There exists a matching of size $t + 1$ in \mathcal{G}_S .
- (ii) There exists a vertex of degree at least $t + 1$ in \mathcal{G}_S .

Indeed this is a consequence of Kőnig's Theorem (stated in Proposition 10): suppose the size of a maximum matching is at most t , then the minimum vertex-cover has size at most t . Since there are $|S|$ many edges in \mathcal{G}_S , there must be a vertex v in the vertex cover with degree at least $\frac{|S|}{t}$. Since $|S| > t^2$, it must be that $\deg(v) > t$ which satisfies (ii). In both the above cases, we construct a string $y \in C^{-1}(1)$ that is a t -limit for S .

- Case (i): Consider the set S' of maxterms corresponding to the edges in a maximum matching of \mathcal{G}_S . Then S' is a set of at least $t + 1$ pairwise co-disjoint maxterms. Then $y \triangleq \vec{1}$ is a t -limit for S' . To see why, consider any set of t positions. By Lemma 16, at each of these positions, at most one of the maxterms in S' can be 0. Since there are $t + 1$ such maxterms and only t positions, there must be a maxterm where the value at all the given positions is 1, thus looking identical to y .
- Case (ii): Let the vertex set of \mathcal{G}_S be $V = V_1 \cup V_2$. Without loss of generality, let the vertex v with $\deg(v) > t$ be in V_1 . Let E be the edges that have v as one endpoint, and let $M_E \subseteq S$ be the maxterms corresponding to the edges in E . Then by property (1) of Lemma 15, the first tuple of vectors in all these maxterms is the same. Let A_1 be the first tuple of vectors. We construct the input $y = (Y_1, Y_2)$ as follows: set $Y_1 \leftarrow A_1$, and set $Y_2 \leftarrow \vec{1}$.

Since the string y was obtained by taking first tuple of a *maxterm*, and setting every vector in the 2nd tuple to 1, it must be a 1-input.

To see that y is a t -limit, take any subset of indices $T \subseteq [2nd]$ with $|T| = t$. We will show that one of the maxterms in M_E looks identical to y in these t positions. For every position from $[nd]$ (the 1st tuple of vectors), *every* maxterm in M_E is identical to y since $Y_1 = A_1$. So assume that all indices in T are from the range $\{nd + 1, \dots, 2nd\}$. By construction, y is all-1s in this range of indices. Since edges in E have distinct endpoints in V_2 , property (2) of Lemma 15 tells us that the second tuple of vectors in the maxterms in T are pairwise co-disjoint. This is similar to case (i): we have $|M_E| \geq t + 1$ many maxterms such that for any position in T , at most one of them is 0, and there are only t positions in T . So by the pigeon-hole principle, there must be a maxterm in M_E that has 1 in all positions from T , thus looking identical to y in these positions. ◀

Since $2\text{-OV}_{n,d}$ is the negation of $2\text{-Int}_{n,d}$, the following is an immediate corollary of Theorem 17.

► **Corollary 18.** *For all $d > 1$, any $\text{OR} \circ \text{AND} \circ \text{OR}$ circuit with bottom fan-in t computing $2\text{-OV}_{n,d}$ requires top fan-in at least $2n^2/t^2$.*

► **Remark 19.** It is easy to see that the t -limit string y constructed in the proof of Theorem 17 is in fact an *upper* t -limit. Therefore the lower bound shown for $2\text{-Int}_{n,d}$ works against a slightly more general class of circuits – $\text{AND} \circ \text{OR} \circ \text{AND}$ circuits that have each bottom AND-gate seeing at most t positive literals. Analogously the lower bound for $2\text{-OV}_{n,d}$ works against $\text{OR} \circ \text{AND} \circ \text{OR}$ circuits where each bottom gate has at most t negated inputs.

3.4 General case: $k\text{-Int}_{n,d}$

We will need the following lemma on k -partite hypergraphs:

► **Lemma 20.** *Let G be a k -partite hypergraph with m many hyperedges. Then for all $t > 0$ at least one of the following holds:*

- (i) *There are more than t vertex-disjoint hyperedges in G .*
- (ii) *There is a vertex u such that $\deg(u) > \lfloor \frac{m}{kt} \rfloor$.*

Proof. Let G be a k -partite hypergraph with m hyperedges. Let S be a largest set of vertex-disjoint hyperedges in G . If $|S| > t$, then the lemma is true. Suppose $|S| \leq t$. Let V_S be the set of vertices participating in the hyperedges in S . Since each hyperedge contains exactly k many vertices, $|V_S| \leq kt$. Also, since S is a largest such set, each of the remaining hyperedges must contain at least one vertex from V_S . Therefore, by an averaging argument, there is a vertex $u \in V_S$ that is part of at least $\frac{m - |S|}{|V_S|}$ many hyperedges outside S , and 1 hyperedge in S . Therefore, we have:

$$\deg(u) \geq \frac{m - |S|}{|V_S|} + 1 \geq \frac{m - t}{kt} + 1 = \frac{m}{kt} - \frac{1}{k} + 1 > \left\lfloor \frac{m}{kt} \right\rfloor \quad \blacktriangleleft$$

We use Lemma 20 to show that if we start with enough hyperedges, then there is a subset of them such that in each part, either all of them coincide, or they are all distinct.

► **Lemma 21.** *Let $k \geq 2$, and let $G = (V_1 \cup \dots \cup V_k, E)$ be a k -partite hypergraph with $|E| > \frac{kt^k}{2}$. Then there exists $S \subseteq E$ with $|S| > t$ such that for each $i \in [k]$, exactly one of the following holds:*

1. *There exists a vertex $u \in V_i$ such that all hyperedges in S share the vertex u .*
2. *No two hyperedges in S share the same vertex in V_i .*

Proof. Induction on k . Base case $k = 2$ is a consequence of König's theorem (Proposition 10): Since $k = 2$, G is just a bipartite graph. If there is a matching in G of size more than t , then let S be the edges in such a matching. Clearly the edges in S are vertex-disjoint and statement (2) holds. Else the maximum matching has size $\leq t$. Then Proposition 10 implies that the minimum vertex cover has size at most t . By an averaging argument, there must exist a vertex u such that $\deg(u) > |E|/t > \frac{k!t^k}{2t} = \frac{2t^2}{2t} = t$. Define S to be the set of edges that share u . Without loss of generality, let $u \in V_1$. Then all edges in S must have distinct vertices in V_2 . Therefore in V_1 , they all coincide, and in V_2 they are all distinct.

Case $k > 2$. Apply Lemma 20 to G . If (i) holds, then we have a set S of more than t vertex-disjoint hyperedges. This means for all $i \in [k]$, statement (2) holds and we are done.

Suppose (ii) holds, then there is a vertex u such that $\deg(u) > \lfloor m/kt \rfloor = \frac{(k-1)!t^{k-1}}{2}$. Let S be the set of all hyperedges that contain vertex u . Then $|S| = \deg(u)$. Let $z \in [k]$ be such that $u \in V_z$.

We construct a $(k-1)$ -partite hypergraph $G' = (V', E')$ by removing V_z , and the z 'th coordinate from each edge. More formally:

$$V' \triangleq V_1 \cup \dots \cup V_{z-1} \cup V_{z+1}, \dots \cup V_k$$

$$E' \triangleq \{(v_1, \dots, v_{z-1}, v_{z+1}, \dots, v_k) \mid (v_1, \dots, v_{z-1}, u, v_{z+1}, v_k) \in S\}$$

(Informally, an edge $e' \in E'$ is just an edge $e \in S$ with its z 'th coordinate removed.)

We define $V'_i = V_i$ for all $i \neq z$ as the $k-1$ parts of V' . Note that $|E'| = |S|$. This is because $\forall e_1, e_2 \in S$ such that $e_1 \neq e_2$, the edges e_1 and e_2 share the vertex u in V_z . So there must exist $j \neq z$ such that e_1 and e_2 use different vertices in V_j . Hence $e'_1 \neq e'_2$. Further, observe that for any $i \neq z$, $e'_1, e'_2 \in E'$ share a vertex in V'_i if and only if e_1 and e_2 share the same vertex in V_i .

Now G' is a $(k-1)$ -partite hypergraph with $|E'| = |S| > \frac{(k-1)!t^{k-1}}{2}$ many hyperedges. By induction on G' , there must exist a set $S' \subseteq E'$ with $|S'| > t$ such that for each $i \neq z$, either all hyperedges in S' share a vertex in V'_i , or they use distinct vertices in V'_i . We already know that since $S' \subseteq S$, all edges in S' share the same vertex in V_z , namely u . Hence for all $i \in [k]$, the edges in S' satisfy (1) or (2). ◀

► **Remark 22.** The statement of Lemma 21 can be seen as a sunflower lemma. Take any vertex u in the graph G that participates in at least one hyperedge from S . Then exactly one of the following holds: (i) The vertex u participates in exactly one hyperedge in S , or (ii) The vertex u participates in all hyperedges in S . The standard sunflower lemma would require more than $k!t^k$ hyperedges, while our statement needs half of that.

We now describe how to construct an upper t -limit in the general case.

► **Lemma 23.** *Let $\mathcal{M} \subseteq \mathcal{N}_{\mathcal{P}}^{n,k,d}$ be any set of permutation-maxterms of k -Int $_{n,d}$ for any $k \geq 2$ and $d \geq k$. If $|\mathcal{M}| > \frac{k!t^k}{2}$, then there is a string $y \in k$ -Int $_{n,d}^{-1}(1)$ that is an upper t -limit for \mathcal{M} .*

Proof. Let $G_{\mathcal{M}} = (V, E)$ be the k -partite support graph of \mathcal{M} (defined in section 3.2), and let $V = V_1 \cup \dots \cup V_k$. By Lemma 21, there exists a set of hyperedges $S \subseteq E$ with $|S| \geq t+1$ such that for each $i \in [k]$, either all edges in S share the same vertex in V_i , or no two edges share a vertex of V_i . Let M_S be the set of maxterms corresponding to S .

Let $B \subseteq [k]$ be the set of all indices $i \in [k]$ such that all edges in S share the same vertex in V_i . Then \overline{B} contains indices of parts where the edges in S use distinct vertices. (Observe that \overline{B} is non-empty because otherwise all maxterms would share all vertices, and hence would be one and the same. But we know that $|S| \geq t+1 > 1$, so this cannot happen.) By

property (1) of Lemma 15, this implies that for each $i \in B$, the i 'th tuple of vectors in the maxterms in M_S are identical. For each $i \in B$, denote the i 'th tuple of vectors in all these maxterms as A_i .

We construct the string $y = (Y_1, \dots, Y_k)$ as follows:

$$\forall i \in B, \text{ set } Y_i \leftarrow A_i$$

$$\forall j \in \overline{B}, \text{ set } Y_j \leftarrow \vec{1}$$

y is a 1-input of $\mathbf{k}\text{-Int}_{n,d}$

Observe that y can also be obtained by starting with any maxterm $N = (N_1, \dots, N_k)$ from S , and setting to 1s all vectors in N_j for all $j \in \overline{B}$. Since N is a maxterm (maximal 0-input), the string y must be a 1-input. This also means that the string y is point-wise greater than or equal to any maxterm in S .

y is a t -limit

Let $T \subseteq [nkd]$ with $|T| = t$ be a set of any t positions. For all $i \in B$, the string y is identical to *every* maxterm in M_S . So assume that T only has positions that fall into tuples indexed by \overline{B} . By property (2) of Lemma 15, the maxterms in M_S are pairwise co-disjoint on all such positions. i.e., for any position $\ell \in T$, at most one maxterm in M_S can be 0. So we have t positions, and $|M_S| = |S| \geq t + 1$ maxterms. By pigeon-hole principle, there exists a maxterm in M_S that is 1 on all these t positions, thus looking identical to y .

Since y is point-wise greater or equal to every maxterm in S , we conclude that indeed y is an upper t -limit for \mathcal{M} . \blacktriangleleft

► **Lemma 24.** *Let C be any OR \circ AND circuit with bottom positive fan-in t computing a function f on n variables. Let y be any string that is an upper t -limit for $f^{-1}(0)$. Then $C(y) = 0$.*

Proof. Let g be any bottom AND-gate of C . Let $P \subseteq [n]$ ($Q \subseteq [n]$) be the variables whose positive literals (negated literals resp.) are input to g . Then $|P| \leq t$ by assumption.

Since y is an upper t -limit for $g^{-1}(0)$, it must be that for every set T of t positions there exists a string $x^{(T)} \in g^{-1}(0)$ such that $y|_T = x^{(T)}|_T$. In particular, this holds for the set P . So in all positions from P , the gate g sees no difference between y and $x^{(T)}$.

The gate g sees negative literals of all variables from Q . Since y is an *upper* t -limit, we have $x^{(T)}|_Q \leq y|_Q$. Hence for all $i \in Q$ such that $\neg x_i = 0$, we also have $\neg y_i = 0$. Hence $g(y) \leq g(x^{(T)}) = 0$ as $x^{(T)} \in g^{-1}(0)$. \blacktriangleleft

► **Theorem 25.** *For all k, d such that $k \leq d$, any AND \circ OR \circ AND circuit with bottom positive fan-in t computing $\mathbf{k}\text{-Int}_{n,d}$ requires top fan-in $\Omega\left(\left(\frac{n}{t}\right)^k\right)$.*

Proof. Let $C = C_1 \wedge \dots \wedge C_s$ be an AND \circ OR \circ AND circuit with bottom positive fan-in t , computing $\mathbf{k}\text{-Int}_{n,d}$. Consider the set $\mathcal{N} = \mathcal{N}_{\mathcal{P}}^{n,k,d}$ of all permutation-maxterms of $\mathbf{k}\text{-Int}_{n,d}$ with respect to any ordered permutation \mathcal{P} of $[d]$ into k non-empty parts (see Definition 11, and Remark 12). Since C outputs 0 on all inputs from \mathcal{N} , there must be some OR \circ AND $_t$ subcircuit C_i that correctly outputs 0 on at least $1/s$ fraction of inputs in \mathcal{N} . We will show that $|C_i^{-1}(0) \cap \mathcal{N}| \leq k! t^k / 2$, and the theorem follows since:

$$\frac{k! n^k}{s} = \frac{1}{s} |\mathcal{N}| \leq |C_i^{-1}(0) \cap \mathcal{N}| \leq \frac{k! t^k}{2}$$

25:14 Depth-3 Circuit Lower Bounds for k -OV

Let $\mathcal{M} = C_i^{-1}(0) \cap \mathcal{N}$. Suppose, for the sake of contradiction, $|\mathcal{M}| > k!t^k/2$. Since $\mathcal{M} \subseteq \mathcal{N}$, we apply Lemma 23 to conclude that there exists a string $y \in \text{k-Int}_{n,d}^{-1}(1)$ that is an upper t -limit y for \mathcal{M} . Then by Lemma 24, it must be that $C(y) = 0$. But this is a contradiction since $y \in \text{k-Int}_{n,d}^{-1}(1)$. ◀

Since $\text{k-OV}_{n,d}$ is the negation of $\text{k-Int}_{n,d}$, the following is an immediate corollary of Theorem 25.

► **Theorem 1.** *For all $k \leq d$, any $\text{OR} \circ \text{AND} \circ \text{OR}$ circuit with bottom fan-in t computing $\text{k-OV}_{n,d}$ requires top fan-in $\Omega\left(\left(\frac{n}{t}\right)^k\right)$.*

4 OR ◦ AND ◦ OR circuits

In this section, we show that any $\text{OR} \circ \text{AND} \circ \text{OR}$ circuit requires exponential size to compute $2\text{-Int}_{n,d}$ for any $d \in \Omega(n^2)$. This result is a consequence of a known lower bound for the iterated intersection function defined as follows:

► **Definition 26** (Iterated Intersection). *Let $A = (a_1, a_2, \dots, a_n)$ and $B = (b_1, b_2, \dots, b_n)$ be tuples of vectors from $\{0, 1\}^d$,*

$$S_{n,d}(A, B) = 1 \iff \forall i \in [n] \text{ we have } a_i \cap b_i \neq \emptyset$$

Observe that $S_{n,d}(A, B)$ differs from $2\text{-Int}_{n,d}(A, B)$ in that the intersection between two vectors a_i and b_j when $i \neq j$ does not affect the value of $S_{n,d}$ at all. Recall the definition of $2\text{-Int}_{n,d}(A, B)$:

$$2\text{-Int}_{n,d}(A, B) = 1 \iff \forall i, j \in [n] \text{ we have } a_i \cap b_j \neq \emptyset$$

The function $S_{n,d}$ can also be defined using an $\text{AND} \circ \text{OR} \circ \text{AND}_2$ circuit of size nd :

$$S_{n,d}(A, B) = \bigwedge_{i=1}^n \bigvee_{j=1}^d a_i[j] \wedge b_i[j].$$

The result by Håstad, Jukna, Pudlák in [14] shows the following lower bound for computing $S_{n,d}$ by $\text{OR} \circ \text{AND} \circ \text{OR}$ circuits:

► **Proposition 27** ([14]). *For all $\ell \leq nd$, any $\text{OR} \circ \text{AND} \circ \text{OR}$ circuit computing $S_{n,d}$ requires size $\min\{2^\ell, (d/\ell)^n\}$.*

In particular, Proposition 27 shows that $S_{\sqrt{n}, \sqrt{n}}$ requires $2^{\Omega(\sqrt{n})}$ size $\text{OR} \circ \text{AND} \circ \text{OR}$ circuits. This can be used to show lower bounds for $2\text{-Int}_{n,d}$:

► **Theorem 28.** *Let C be an $\text{OR} \circ \text{AND} \circ \text{OR}$ circuit computing $2\text{-Int}_{n,d}$. Then for all $\ell \leq d$, size of C is at least $\min\{2^\ell, \left(\frac{d}{n\ell}\right)^n\}$.*

Proof. We show this by reducing $S_{n, \lfloor d/n \rfloor}$ to $2\text{-Int}_{n,d}$ via projections. Let $d' = \lfloor d/n \rfloor$. Take any instance $A = (a_1, \dots, a_n)$ and $B = (b_1, \dots, b_n)$ with $a_i, b_i \in \{0, 1\}^{d'}$ of $S_{n,d'}$. We create two tuples of d -dimensional vectors $A' = (a'_1, \dots, a'_n)$ and $B' = (b'_1, \dots, b'_n)$ that serve as an instance of $2\text{-Int}_{n,d}$ as follows – for all $i \in [n]$, define $a'_i = 1^{(i-1)d'} a_i 1^{(n-i)d'}$ and $b'_i = 0^{(i-1)d'} b_i 0^{(n-i)d'}$. Note that the dimension of each a'_i and b'_i is $nd' \leq d$.

Observe that a_i and b_i are disjoint if and only if a'_i and b'_i are disjoint. So if (A, B) was a 0-instance of $S_{n,d'}$, then (A', B') is a 0-instance of $2\text{-Int}_{n,d}$.

Further, if $b_j \neq \vec{0}$ for some $j \in [n]$, then for all $i \neq j$, we have $a'_i \cap b'_j \neq \emptyset$. To see this, observe that if $b_j \neq \vec{0}$, then there is some position $p \in [(j-1)d' + 1, jd']$ such that $b'_j[p] = 1$. But by construction, the vector a'_i is 1 everywhere outside the interval $[(i-1)d' + 1, id']$. Since $i \neq j$, the vector a'_i must be 1 at position p .

If (A, B) was a 1-instance of $S_{n,d'}$, then all a_i intersect b_i . This means all b_i are non-zero vectors. Thus for all $i, j \in [n]$, $a'_i \cap b'_j \neq \emptyset$.

The above reduction shows that C can be used to compute $S_{n, \lfloor d/n \rfloor}$. Applying Proposition 27 to C tells us that C must have size at least $\min\{2^\ell, \left(\frac{d}{n\ell}\right)^n\}$ for all $\ell \leq d$. ◀

Our reduction in proof of Theorem 28 inflates the dimension of vectors by a factor of n making the obtained bound trivial when $d \in O(n)$. However, we can still conclude an exponential lower bound by substituting $\ell = d/2n$ that gives us a lower bound of $\min\{2^{d/2n}, 2^n\} \in 2^{\Omega(n)}$ when $d \in \Omega(n^2)$.

Since $2\text{-OV}_{n,d}$ is the negation of $2\text{-Int}_{n,d}$, the following is an immediate corollary.

▶ **Theorem 3.** *For all $\ell \leq d$, any $\text{AND} \circ \text{OR} \circ \text{AND}$ circuit computing $2\text{-OV}_{n,d}$ requires size $s \in \Omega(\min\{2^\ell, \left(\frac{d}{n\ell}\right)^n\})$. In particular, for $\ell = d/2n$ and $d \in \Omega(n^2)$, $s \in \Omega(2^n)$.*

5 A General Upper Bound

In this section, we describe a more general construction of a depth-3 circuit to compute $k\text{-Int}_{n,d}$ that allows a trade-off between the top fan-in and bottom fan-in. We recall the construction given by equation 3 here:

$$k\text{-Int}_{n,d}(A_1, \dots, A_k) = \bigwedge_{i_1, \dots, i_k \in [n]} \bigvee_{j \in [d]} (a_{i_1}[j] \wedge \dots \wedge a_{i_k}[j]) \quad (3)$$

We generalize this construction to obtain the following trade-off between top and bottom fan-in:

▶ **Proposition 29.** *For any integer $1 \leq t \leq n^k$, the function $k\text{-Int}_{n,d}$ can be computed by a monotone depth-3 $\text{AND} \circ \text{OR} \circ \text{AND}$ circuit with top fan-in $\lceil \frac{n^k}{t} \rceil$, middle fan-in d^t , and bottom fan-in at most kt .*

Proof. Let C be the circuit described in equation 3. Observe that each $\text{OR} \circ \text{AND}$ subcircuit of C is checking whether a particular choice $a_{i_1} \in A_1, a_{i_2} \in A_2, \dots, a_{i_k} \in A_k$ of vectors are intersecting or not. Since there are n^k many such choices, the top fan-in in C is n^k . Checking if a particular choice of k vectors intersects at some fixed coordinate uses an AND of fan-in k , and hence the bottom fan-in in C is k .

To obtain the trade-off in the theorem statement, the idea is to construct an $\text{AND} \circ \text{OR} \circ \text{AND}$ circuit where each $\text{OR} \circ \text{AND}$ subcircuit checks whether t many such choices of vectors intersect. Each choice can be written as a k -tuple of vectors $(a_{i_1}, \dots, a_{i_k})$. For convenience, let's assume that t divides n^k . Let $T = \{T_1, T_2, \dots, T_{n^k/t}\}$ be a partition of the set of n^k possible k -tuples of vectors into n^k/t parts with each T_l containing exactly t many k -tuples. For the vectors in any particular k -tuple in T_l to have non-empty intersection, there must exist a position $i \in [d]$ where all the k vectors in the k -tuple are 1. Hence to check if each of the k -tuples of vectors in T_l have non-zero intersection, it suffices to check if there exist t positions $i_1, i_2, \dots, i_t \in [d]$ such that the j 'th k -tuple of vectors intersect in i_j .

Let $A_l^j[i]$ be the AND of the bits in the i^{th} position of the vectors in the j^{th} tuple in T_l . This is an AND gate with fan-in k because there are k many vectors in each tuple. We construct the following circuit where the ℓ^{th} OR \circ AND subcircuit checks if each k -tuple of vectors in T_ℓ have non-zero intersection:

$$G_t = \bigwedge_{l \in \{1, \dots, \frac{n^k}{t}\}} \bigvee_{i_1, i_2, \dots, i_t \in [d]} (A_l^1[i_1] \wedge A_l^2[i_2] \wedge \dots \wedge A_l^t[i_t])$$

Observe that G_t has top fan-in as n^k/t , middle fan-in as d^t , and bottom fan-in kt as desired. \blacktriangleleft

References

- 1 Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. Tight hardness results for LCS and other sequence similarity measures. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 59–78. IEEE Computer Society, 2015. doi:10.1109/FOCS.2015.14.
- 2 Arturs Backurs and Piotr Indyk. Edit distance cannot be computed in strongly subquadratic time (unless SETH is false). In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 51–58. ACM, 2015. doi:10.1145/2746539.2746612.
- 3 Arturs Backurs and Piotr Indyk. Which regular expression patterns are hard to match? In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*, pages 457–466. IEEE Computer Society, 2016. doi:10.1109/FOCS.2016.56.
- 4 Karl Bringmann. Why walking the dog takes time: Fréchet distance has no strongly subquadratic algorithms unless SETH fails. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 661–670. IEEE Computer Society, 2014. doi:10.1109/FOCS.2014.76.
- 5 Karl Bringmann and Wolfgang Mulzer. Approximability of the discrete fréchet distance. *J. Comput. Geom.*, 7(2):46–76, 2016. doi:10.20382/jocg.v7i2a4.
- 6 Chris Calabro, Russell Impagliazzo, and Ramamohan Paturi. On the exact complexity of evaluating quantified k -cnf. In *Parameterized and Exact Computation - 5th International Symposium, IPEC 2010, Chennai, India, December 13-15, 2010. Proceedings*, volume 6478 of *Lecture Notes in Computer Science*, pages 50–59. Springer, 2010. doi:10.1007/978-3-642-17493-3_7.
- 7 Lijie Chen and Ryan Williams. An equivalence class for orthogonal vectors. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 21–40. SIAM, 2019. doi:10.1137/1.9781611975482.2.
- 8 Reinhard Diestel. *Graph Theory, 4th Edition*, volume 173 of *Graduate texts in mathematics*. Springer, 2012.
- 9 Paul Erdős and Richard Rado. Intersection theorems for systems of sets. *Journal of the London Mathematical Society*, 1(1):85–90, 1960.
- 10 Oded Goldreich and Avishay Tal. On constant-depth canonical boolean circuits for computing multilinear functions. In *Computational Complexity and Property Testing - On the Interplay Between Randomness and Computation*, volume 12050 of *Lecture Notes in Computer Science*, pages 306–325. Springer, 2020. doi:10.1007/978-3-030-43662-9_17.
- 11 Oded Goldreich and Avi Wigderson. On the size of depth-three boolean circuits for computing multilinear functions. In *Computational Complexity and Property Testing - On the Interplay Between Randomness and Computation*, volume 12050 of *Lecture Notes in Computer Science*, pages 41–86. Springer, 2020. doi:10.1007/978-3-030-43662-9_6.

- 12 András Hajnal, Wolfgang Maass, Pavel Pudlák, Mario Szegedy, and György Turán. Threshold circuits of bounded depth. *J. Comput. Syst. Sci.*, 46(2):129–154, 1993. doi:10.1016/0022-0000(93)90001-D.
- 13 Johan Håstad. Almost optimal lower bounds for small depth circuits. In *Proceedings of the 18th Annual ACM Symposium on Theory of Computing, May 28-30, 1986, Berkeley, California, USA*, pages 6–20. ACM, 1986. doi:10.1145/12130.12132.
- 14 Johan Håstad, Stasys Jukna, and Pavel Pudlák. Top-down lower bounds for depth-three circuits. *Computational Complexity*, 5(2):99–112, 1995. doi:10.1007/BF01268140.
- 15 Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-sat. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001. doi:10.1006/jcss.2000.1727.
- 16 Stasys Jukna. *Boolean Function Complexity - Advances and Frontiers*, volume 27 of *Algorithms and combinatorics*. Springer, 2012. doi:10.1007/978-3-642-24508-4.
- 17 Daniel M. Kane and Richard Ryan Williams. The orthogonal vectors conjecture for branching programs and formulas. In *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, volume 124 of *LIPICs*, pages 48:1–48:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.ITCS.2019.48.
- 18 Mauricio Karchmer and Avi Wigderson. Monotone circuits for connectivity require super-logarithmic depth. *SIAM J. Discret. Math.*, 3(2):255–265, 1990. doi:10.1137/0403021.
- 19 Dénes König. Graphen und matrizen. *Mat. Fiz. Lapok*, 38(10), 1931.
- 20 Igor Carboni Oliveira, Rahul Santhanam, and Srikanth Srinivasan. Parity helps to compute majority. In *34th Computational Complexity Conference, CCC 2019, July 18-20, 2019, New Brunswick, NJ, USA*, volume 137 of *LIPICs*, pages 23:1–23:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.CCC.2019.23.
- 21 Alexander A. Razborov. On the method of approximations. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing, May 14-17, 1989, Seattle, Washington, USA*, pages 167–176. ACM, 1989. doi:10.1145/73007.73023.
- 22 Roman Smolensky. Algebraic methods in the theory of lower bounds for boolean circuit complexity. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing, 1987, New York, New York, USA*, pages 77–82. ACM, 1987. doi:10.1145/28395.28404.
- 23 Leslie G. Valiant. Exponential lower bounds for restricted monotone circuits. In *Proceedings of the 15th Annual ACM Symposium on Theory of Computing, 25-27 April, 1983, Boston, Massachusetts, USA*, pages 110–117. ACM, 1983. doi:10.1145/800061.808739.
- 24 Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theoretical Computer Science*, 348(2-3):357–365, 2005. doi:10.1016/j.tcs.2005.09.023.
- 25 Virginia Vassilevska Williams. On some fine-grained questions in algorithms and complexity. In *Proceedings of the International Congress of Mathematicians: Rio de Janeiro 2018*, pages 3447–3487. World Scientific, 2018.
- 26 Virginia Vassilevska Williams and Ryan Williams. Subcubic equivalences between path, matrix and triangle problems. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 645–654. IEEE Computer Society, 2010. doi:10.1109/FOCS.2010.67.