

On monotone circuits with local oracles and clique lower bounds

Jan Krajíček Igor C. Oliveira

Received April 24, 2017; Revised December 18, 2017; Published March 25, 2018

Abstract: We investigate monotone circuits with local oracles [Krajíček, 2016], i.e., circuits containing additional inputs $y_i = y_i(\vec{x})$ that can perform unstructured computations on the input string \vec{x} . Let $\mu \in [0, 1]$ be the locality of the circuit, a parameter that bounds the combined strength of the oracle functions $y_i(\vec{x})$, and $U_{n,k}, V_{n,k} \subseteq \{0, 1\}^m$ be the set of k -cliques and the set of complete $(k-1)$ -partite graphs, respectively (similarly to [Razborov, 1985]). Our results can be informally stated as follows.

- (i) For an appropriate extension of depth-2 monotone circuits with local oracles, we show that the size of the smallest circuits separating $U_{n,3}$ (triangles) and $V_{n,3}$ (complete bipartite graphs) undergoes two phase transitions according to μ .
- (ii) For $5 \leq k(n) \leq n^{1/4}$, arbitrary depth, and $\mu \leq 1/50$, we prove that the monotone circuit size complexity of separating the sets $U_{n,k}$ and $V_{n,k}$ is $n^{\Theta(\sqrt{k})}$, under a certain restrictive assumption on the local oracle gates.

The second result, which concerns monotone circuits with restricted oracles, extends and provides a matching upper bound for the exponential lower bounds on the monotone circuit size complexity of k -clique obtained in [Alon and Boppana, 1987].

Key words and phrases: monotone circuits, proof complexity, k -clique problem

1 Introduction and motivation

We establish initial lower bounds on the power of monotone circuits with local oracles (monotone CLOs), an extension of monotone circuits introduced in [10] motivated by problems in proof complexity. Interestingly, while the model has been conceived as part

of an approach to establish new length-of-proofs lower bounds, our results indicate that investigating such circuits can benefit our understanding of classical results obtained in the usual setting of monotone circuit complexity, where no oracle gates are present (see the discussion on the Alon-Boppana exponential lower bounds for k -clique [1] presented later in this section).

Before describing the circuit model and our contributions in more detail, which require no background in proof complexity, we explain the main motivation that triggered our investigations.

Relation to proof complexity. A major open problem in proof complexity is to obtain lower bounds on proof length in $F_d[\oplus]$, depth- d Frege systems extended with parity connectives (cf. [8]). It is known that strong enough lower bounds for $F_3[\oplus]$, the depth-3 version of this system, imply related lower bounds for each system $F_d[\oplus]$, where $d \in \mathbb{N}$ is arbitrary [4]. A natural restriction of $F_3[\oplus]$ for which proving general lower bounds is still open is the proof system $R(\text{Lin}/\mathbb{F}_2)$ (cf. [5], [10]). It corresponds to an extension of Resolution where clauses involve linear functions over \mathbb{F}_2 .¹

In order to attack this and other related problems, [10] proposed a generalization of the feasible interpolation method to randomized feasible interpolation. Among other results, [10] established that lower bounds on the size of monotone circuits with local oracles separating the sets $U_{n,k}$ and $V_{n,k}$ (defined below) imply lower bounds on the size of general (dag-like) $R(\text{Lin}/\mathbb{F}_2)$ proofs. In addition, it was shown that strong lower bounds in the new circuit model would provide a unifying approach to important length-of-proofs lower bounds established via feasible interpolation (cf. [10, Section 6], [11]).

Motivated by these connections and by the important role of feasible interpolation in proof complexity, we start in this work a more in-depth investigation of the power and limitations of monotone circuits with local oracles. We focus on the complexity of the k -clique problem over the classical sets of negative and positive instances considered in monotone circuit complexity [12, 1]. While the monotone complexity of k -clique has been investigated over other input distributions of interest (cf. [14]), we remark that the structure of these instances is particularly useful in proof complexity (cf. [9, 11, 2]). The corresponding tautologies have appeared in several other works.

We provide next a brief introduction to the circuit model and to the set of instances of k -clique that are relevant to our results.

An extension of monotone circuits. A monotone circuit with local oracles $C(\vec{x}, \vec{y})$ is a monotone boolean circuit containing extra inputs y_j (local oracles) that compute an arbitrary *monotone* function of \vec{x} . In order to limit the power of these oracles, there is a locality parameter $\mu \in [0, 1]$ that controls the sets of positive and negative inputs on which the inputs y_i can be helpful. In more detail, we consider circuits computing a monotone function $f: \{0, 1\}^m \rightarrow \{0, 1\}$, and associate to each input y_i a rectangle $U_i \times V_i$, with $U_i \subseteq f^{-1}(1)$ and $V_i \subseteq f^{-1}(0)$. We restrict attention to sets of rectangles whose union have measure at most μ according to an appropriate distribution \mathcal{D} that depends on f . We are guaranteed that $y_i(U_i) = 1$ and $y_i(V_i) = 0$ but, crucially, the

¹Lower bounds for tree-like $R(\text{Lin}/\mathbb{F}_2)$ -proofs were established in [5].

computation of $C(\vec{x}, \vec{y})$ must be correct no matter the interpretation of each y_i outside its designated sets U_i and V_i .

The k -clique function and the sets $U_{n,k}$ and $V_{n,k}$. We focus on the monotone boolean function $f: \{0, 1\}^m \rightarrow \{0, 1\}$ that outputs 1 on an n -vertex graph $G \in \{0, 1\}^m$ if and only if it contains a clique of size k , where $m = \binom{n}{k}$. More specifically, we investigate its complexity as a partial boolean function over $U_{n,k} \cup V_{n,k}$, where $U_{n,k}$ is the set of inputs corresponding to k -cliques over the set $[n]$ of vertices, and $V_{n,k}$ is the set of complete ζ -partite graphs over $[n]$, where $\zeta = k - 1$. Roughly speaking, for this choice of f , we measure the size of a subset $\mathcal{B} \subseteq U_{n,k} \times V_{n,k}$ using the product distribution obtained from the uniform distribution over the k -cliques in $U_{n,k}$, and the distribution supported over $V_{n,k}$ obtained by sampling a random coloring $\chi: [n] \rightarrow [k - 1]$ of $[n]$ using exactly $\zeta = k - 1$ colors, and considering the associated complete ζ -partite graph $G(\chi)$.²

A more rigorous treatment of the circuit model and of the problem investigated in our work appears in Section 2.

1.1 Our results

We observe a phase transition for an extension of depth-2 monotone circuits with local oracles that separate triangles from complete bipartite graphs.

Theorem 1.1 (Phase transitions in depth-2). *Let $s = s(n, \mu)$ be the minimum size of a depth-2 monotone circuit (DNF) on inputs \vec{x} , $y_i(\vec{x})$, and $g_j(\vec{y})$ that separates $U_{n,3}$ and $V_{n,3}$, where the y -inputs have locality $\leq \mu$, and each g_j is an arbitrary monotone function on \vec{y} . Then, for every $\varepsilon > 0$,*

$$s = \begin{cases} 1 & \text{if } \mu = 1, \\ \Theta_\varepsilon(n^2) & \text{if } 1/2 + \varepsilon \leq \mu \leq 1 - \varepsilon, \\ \Theta_\varepsilon(n^3) & \text{if } 0 \leq \mu \leq 1/2 - \varepsilon. \end{cases}$$

Furthermore, the upper bounds on $s(n, \mu)$ do not require the extra inputs $g_j(\vec{y})$.

Observe that the lower bounds remain valid in the presence of the functions $g_j(\vec{y})$. In other words, in the restricted setting of depth-2 circuits, a small locality parameter does not help, even if arbitrary monotone computations that depend on the output of the local oracle gates are allowed in the circuit. (As explained in Section 3, the monotone functions $g_j(\vec{y})$ can be handled in a generic way, and add no power to the model.)

The proof of Theorem 1.1 is presented in Section 3. The argument considers different bottlenecks in the computation based on the value of μ . In our opinion, the main conceptual message of Theorem 1.1 is that an interesting complexity-theoretic behavior appears already at depth two. Indeed, the oracle gates can interact with the standard input variables in unexpected ways, and the main difficulty when analyzing

²Some authors consider as negative instances the larger set of complete ζ -partite graphs where ζ ranges from 1 to $k - 1$. For technical reasons, we work with exactly $(k - 1)$ -partite graphs (cf. Claim 2.1). In most lower bound contexts this is inessential, as a random coloring $\chi: [n] \rightarrow [k - 1]$ under a bounded $k(n)$ contains non-empty color classes except with an exponentially small probability.

general monotone CLOs is the arbitrary nature of these gates, which are limited only by the locality parameter.³

We obtain stronger results for larger $k = k(n)$ and with respect to unrestricted monotone circuits (i.e., arbitrary depth), but our approach requires an extra condition on the set of rectangles that appear in the definition of the oracle gates. Our assumption, denoted by \mathcal{A}_d , says that if each oracle variable y_i is associated to the rectangle $U_i \times V_i$, then the intersection of every collection of $d + 1$ sets U_i is empty.

Theorem 1.2 (Upper and lower bounds for monotone circuits with restricted oracles). *For every $k = k(n)$ satisfying $5 \leq k \leq n^{1/4}$, the following holds.*

1. *If $D(\vec{x}, \vec{y})$ is a monotone circuit with local oracles that separates $U_{n,k}$ and $V_{n,k}$ and its y -variables have locality $\mu \leq 1/16$ and satisfy condition \mathcal{A}_d , then $\text{size}(D) = n^{\Omega(\sqrt{k}/d)}$.*
2. *For every $\varepsilon > 0$, there exists a monotone circuit with local oracles $C(\vec{x}, \vec{y})$ of size $n^{O_\varepsilon(\sqrt{k})}$ separating $U_{n,k}$ and $V_{n,k}$ whose y -variables have locality $\mu \leq \varepsilon$ and satisfy condition \mathcal{A}_1 .*

The proof of Theorem 1.2 appears in Section 4. The lower bound extends results on the monotone circuit size complexity of k -clique for large $k = k(n)$ obtained in [1].⁴ Indeed, our argument relies on their analysis of Razborov’s approximation method [12], with extra work required to handle the oracle gates. The upper bound is achieved by an explicit description of a monotone CLO generalizing the construction from Theorem 1.1. The following corollary, stated for reference, is immediate from Theorem 1.2.

Corollary 1.3. *Let $5 \leq k(n) \leq n^{1/4}$, $\mu = 1/50$, and assume rectangles are mapped to local oracle gates in a way that no k -clique is associated to more than a constant number of rectangles. Then the monotone circuit size complexity of separating the sets $U_{n,k}$ and $V_{n,k}$ is $n^{\Theta(\sqrt{k})}$.*

(We note that the constant $1/50$ appearing in this statement is not particularly important, and that any small enough constant locality parameter μ suffices.) To our knowledge, Corollary 1.3 provides the first explanation for the tightness of the Alon-Boppana [1] exponential lower bounds for k -clique. In particular, in order to prove monotone circuit lower bounds for this problem stronger than $n^{\sqrt{k}}$ in the regime where $k(n) \gg \text{poly}(\log n)$, one has to consider either a different set of instances, or employ a technique that does not apply to circuits with local oracles of constant locality.⁵

We discuss some directions for future investigations in Section 5, where we also say a few more words on the connection to proof complexity.⁶

³It is plausible that the analysis behind the proof of Theorem 1.1 extends to larger k , but we have not pursued this direction in the context of depth-2 circuits. See also the related discussion on Section 5.

⁴For $k \leq \log n$, near-optimal results were proved in [12].

⁵We remark that much tighter monotone lower bounds of the form $n^k/\text{poly}(\log n)$ are known in the regime where k is constant or slightly super-constant [12, 1]. Interestingly, these results do not generalize to circuits with local oracles due to the different choice of parameters employed in the corresponding legitimate lattices.

⁶We have made no attempt to optimize the constants and the asymptotic notation appearing in Theorems 1.1 and 1.2.

2 Notation and basic facts

Let $[e]$ denote the set $\{1, 2, \dots, e\}$, $e \in \mathbb{N}$. For a set B , we use $\binom{B}{\ell}$ to denote the family of subsets of B of size exactly ℓ . The function $\log(\cdot)$ refers to logarithm in base 2. For a set V , we use $v \sim V$ to denote a uniformly distributed element from V . We are interested in the computation of partial boolean functions over $\{0, 1\}^m$. For $A \subseteq \{0, 1\}^m$, a function $f: A \rightarrow \{0, 1\}$ is monotone if $x, y \in A$ and $x \preceq y$ (i.e. $x_i \leq y_i$ for all $i \in [m]$) imply $f(x) \leq f(y)$.

Monotone CLOs. A monotone boolean circuit $C(x_1, \dots, x_n, y_1, \dots, y_e)$ on n variables and e local oracles (monotone CLO for short) is a (non-empty) directed acyclic graph containing $\leq n + e + 2$ sources and one sink (the output node). The non-source nodes have in-degree 2. Source nodes are labeled by elements in $\{x_1, \dots, x_n\} \cup \{y_1, \dots, y_e\} \cup \{0, 1\}$, and each non-source node is labeled by a gate symbol in $\{\wedge, \vee\}$. We say that C has size s if the total number of nodes in the underlying graph is s , including source nodes. The computation of C on an input string $(a, b) \in \{0, 1\}^n \times \{0, 1\}^e$ is defined in the natural way.

The formulation above is consistent with the statement of Theorem 1.2. In Theorem 1.1, which concerns bounded-depth circuits, we allow the internal $\{\wedge, \vee\}$ -nodes to have unbounded fan-in.

We consider the computation of $C(\vec{x}, \vec{y})$ on input pairs where each bit in the second input \vec{y} is a function of \vec{x} . Furthermore, we will restrict our analysis to monotone computations over a set $A \subseteq \{0, 1\}^n$ of interest. For this reason, to specify the computation of C on a string $x \in A$, we will associate to each local oracle variable y_i a corresponding monotone function $f_i: A \rightarrow \{0, 1\}$.

In order to obtain a non-trivial notion of circuit complexity in this model, we use a real-valued parameter $\mu \in [0, 1]$ to control the family of admissible functions f_i . Each function f_i separates a particular pair of sets $U_i \subseteq f^{-1}(1) \subseteq A$ and $V_i \subseteq f^{-1}(0) \subseteq A$, but C must be correct no matter the choice of the functions f_i separating these sets. The parameter μ captures the measure of $\bigcup_i U_i \times V_i$. This is formalized by the definitions introduced next.

Correctness and locality. Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$, $U \subseteq f^{-1}(1)$, $V \subseteq f^{-1}(0)$, $W = (U, V)$, and $A = U \cup V$. Moreover, let $U_1, \dots, U_e \subseteq U$ and $V_1, \dots, V_e \subseteq V$ be sets of inputs, and for convenience, let $\mathcal{W} = (W_i)_{i \in [e]}$ denote the sequence of pairs $W_i = (U_i, V_i)$. Finally, let \mathcal{D} be a probability distribution supported over $U \times V$. We say that \mathcal{W} has locality μ with respect to \mathcal{D} if, for $\mathcal{B} = \bigcup_{i \in [e]} U_i \times V_i$,

$$\Pr_{(u,v) \sim \mathcal{D}} [(u,v) \in \mathcal{B}] \leq \mu.$$

We say that a pair $W' = (U', V')$ is included in the pair $W = (U, V)$ if $U' \subseteq U$ and $V' \subseteq V$, and that a sequence $\mathcal{W} = (U_i, V_i)_{i \in [e]}$ of pairs is included in W if each member $W_i = (U_i, V_i)$ of \mathcal{W} is included in W . Let $g: A \rightarrow \{0, 1\}$ be an arbitrary monotone boolean function over $A = U \cup V$. We say that g separates a pair (U', V') if $g(U') = 1$ and $g(V') = 0$. Let $\mathcal{F} = (f_1, \dots, f_e)$ be a sequence of functions, where each $f_i \in A \rightarrow \{0, 1\}$ is monotone. We say that \mathcal{F} separates \mathcal{W} if each f_i separates

(U_i, V_i) . For convenience, we also say in this case that \mathcal{F} is a \mathcal{W} -separating sequence of functions.

Given a monotone CLO pair (C, \mathcal{W}) as above, and a \mathcal{W} -separating sequence \mathcal{F} of monotone functions, let

$$C(\vec{x}, \mathcal{F}) \stackrel{\text{def}}{=} C(x_1, \dots, x_n, f_1(\vec{x}), \dots, f_e(\vec{x}))$$

denote the function in $A \rightarrow \{0, 1\}$ that agrees with the output of C when each oracle input y_i is set to $f_i(x)$. Observe that $C(x, \mathcal{F})$ is a *monotone* function over $A = U \cup V$, since C is a monotone circuit and each f_i is a monotone function over A . We will sometimes abuse notation and view $C(x, \mathcal{F})$ as a circuit. We say that the pair (C, \mathcal{W}) computes the function $f: A \rightarrow \{0, 1\}$ if for every \mathcal{W} -separating sequence \mathcal{F} of monotone functions, we have $C(a, \mathcal{F}) = f(a)$ for all $a \in A$. (We stress that the monotone CLO pair must be correct on every input string, and on every \mathcal{W} -separating sequence.)

Finally, let $f \in \{0, 1\}^n \rightarrow \{0, 1\}$ be a monotone function, $A = U \cup V$ for sets $U \subseteq f^{-1}(1)$ and $V \subseteq f^{-1}(0)$, and $W = (U, V)$. We say that f can be computed over $A \subseteq \{0, 1\}^n$ by a monotone circuit with local oracles of size s and locality μ (with respect to a distribution \mathcal{D}) if there exists a monotone circuit $C(\vec{x}, \vec{y})$ of size $\leq s$ and a sequence $\mathcal{W} = (U_i, V_i)_{i \in [e]}$ of length $e \leq s$ that is included in W and has locality $\leq \mu$ such that the monotone CLO pair (C, \mathcal{W}) computes f over A .

For convenience of notation, we will sometimes write $y_i = y[U_i, V_i]$ to indicate a local oracle over the pair $W = (U_i, V_i)$.

Defining $U_{n,k}$, $V_{n,k}$, and $\mathcal{D}_{n,k}$. Let $m = \binom{n}{k}$, where $n \geq 4$, and let $k \in \mathbb{N}$ be an integer satisfying $3 \leq k < n$. We view $[n]$ as a set of vertices, and $[m]$ as its associated set of (undirected) edges. For $B \subseteq [n]$, we use $K_B \in \{0, 1\}^m$ to denote the graph (also viewed as a string) corresponding to a clique over B . Let

$$\begin{aligned} U_{n,k} &\stackrel{\text{def}}{=} \left\{ K_B \in \{0, 1\}^m \mid B \in \binom{[n]}{k} \right\}, \text{ and} \\ V_{n,k} &\stackrel{\text{def}}{=} \left\{ H \in \{0, 1\}^m \mid H \text{ is a non-trivial complete } \zeta\text{-partite graph, where } \zeta = k - 1 \right\}, \\ A_{n,k} &\stackrel{\text{def}}{=} U_{n,k} \cup V_{n,k}. \end{aligned}$$

Clearly, $U_{n,k} \cap V_{n,k} = \emptyset$. It is convenient to associate to each coloring $\chi: [n] \rightarrow [k-1]$ a corresponding graph $G(\chi)$, where $e = \{v_1, v_2\} \in E(G(\chi))$ if and only if $\chi(v_1) \neq \chi(v_2)$. Let

$$V_{n,k}^\chi \stackrel{\text{def}}{=} \{ \chi \mid \chi: [n] \rightarrow [k-1] \}$$

be the family of all possible colorings of $[n]$ using at most $k-1$ colors. Under our definitions, for a given coloring $\chi \in V_{n,k}^\chi$ we have $G(\chi) \in V_{n,k}$ if and only if $|\chi([n])| = k-1$. We measure the locality of monotone CLO pairs (C, \mathcal{W}) separating $U_{n,k}$ and $V_{n,k}$ with respect to a product distribution $\mathcal{D}_{n,k} \stackrel{\text{def}}{=} \mathcal{D}_{n,k}^U \times \mathcal{D}_{n,k}^V$, whose components are defined as follows. $\mathcal{D}_{n,k}^U$ is simply the uniform distribution over the k -cliques in $U_{n,k}$, while $\mathcal{D}_{n,k}^V$ assigns to each fixed graph $H \in V_{n,k}$ probability mass $\mathcal{D}_{n,k}^V(H) \stackrel{\text{def}}{=} \Pr_{\chi \sim V_{n,k}^\chi} [G(\chi) = H \mid G(\chi) \in V_{n,k}]$.⁷ (This is simply the uniform distribu-

⁷Note that the probability that a random coloring $\chi: [n] \rightarrow [k-1]$ contains less than $k-1$ non-trivial color classes is exponentially small in n for the values of $k(n)$ investigated in Theorems 1.1 and 1.2.

tion over $V_{n,k}$, but this is not the most convenient point of view in some estimates.)

The sequence \mathcal{F}^* . The definition introduced above agrees with the formulation of monotone circuits with oracles from [10]. We stress that a source of difficulty when computing a function $f: A \rightarrow \{0, 1\}$ using a monotone circuit $C(\vec{x}, \vec{y})$ and a sequence $\mathcal{W} = (W_i)$ of pairs included in $W = (f^{-1}(0), f^{-1}(1))$ is that $C(x, \mathcal{F})$ must be correct for every \mathcal{W} -separating sequence $\mathcal{F} = (f_i)$ of monotone functions. In order to prove lower bounds against a monotone CLO pair (C, \mathcal{W}) , we will consider a particular instantiation of the monotone functions $f_i: A \rightarrow \{0, 1\}$, discussed next.

Let $y_i = y[U_i, V_i]$ be a local oracle variable associated with the pair $W_i = (U_i, V_i)$. We define the function $f_{W_i}^*: A \rightarrow \{0, 1\}$ as follows:

$$f_{W_i}^*(x) = \begin{cases} 1 & \text{if } x \in U_i \cup (V \setminus V_i), \\ 0 & \text{otherwise.} \end{cases}$$

Observe that $f_{W_i}^*(U_i) = 1$ and $f_{W_i}^*(V_i) = 0$. In particular, $f_i^* \stackrel{\text{def}}{=} f_{W_i}^*$ separates the pair W_i . We use $\mathcal{F}^* \stackrel{\text{def}}{=} (f_i^*)$ to denote the corresponding sequence of functions for a given choice of $\mathcal{W} = (W_i)$.

For an arbitrary monotone function $f: A \rightarrow \{0, 1\}$, $U_i \subseteq U \subseteq f^{-1}(1)$, and $V_i \subseteq V \subseteq f^{-1}(0)$, f_i^* is not necessarily monotone. However, for the problem investigated in our work f_i^* is always monotone, as stated next.

Claim 2.1. *Let $3 \leq k < n$. For every pair $W_i = (U_i, V_i)$ with $U_i \subseteq U_{n,k}$ and $V_i \subseteq V_{n,k}$, the function $f_i^*: A_{n,k} \rightarrow \{0, 1\}$ is monotone.*

Proof. It is enough to observe that, under these assumptions, there are no distinct strings $a_1, a_2 \in A_{n,k}$ satisfying $a_1 \preceq a_2$. Here we crucially used that the $(k-1)$ -partite graphs in $V_{n,k}$ have exactly $k-1$ non-empty parts. \square

The use of \mathcal{F}^* to prove lower bounds against monotone CLO pairs (C, \mathcal{W}) computing a monotone function $f: A \rightarrow \{0, 1\}$ is justified by the following observation, which describes an extremal property of \mathcal{F}^* .

Claim 2.2. *Let $\mathcal{F} = (f_i)$ be an arbitrary \mathcal{W} -separating sequence of monotone functions $f_i: A \rightarrow \{0, 1\}$. If $C(x, \mathcal{F})$ is incorrect on an input $a \in A$, then $C(x, \mathcal{F}^*)$ is also incorrect on a .*

Proof. Assume that $a \in U$. Consequently, $f(a) = 1$, and the assumption that $C(x, \mathcal{F})$ is incorrect means that $C(x, \mathcal{F}) = 0$. Using that each f_i separates $W_i = (U_i, V_i)$ and the definition of f_i^* , we get $f_i^*(a) \leq f_i(a)$. By the monotonicity of the circuit C , it follows that $C(a, \mathcal{F}^*) \leq C(a, \mathcal{F})$. Thus $C(a, \mathcal{F}^*)$ is incorrect on input a as well. The case where $a \in V$ is analogous. \square

Therefore, \mathcal{F}^* is the hardest separating-sequence, meaning that any circuit that computes f under \mathcal{F}^* computes f under any separating-sequence.

Remark 2.3 (Simulating negated inputs). It is possible to simulate negated input variables in C using oracle gates. For instance, if $x_{\{1,2\}}$ corresponds to the input edge $\{1, 2\}$, we define an oracle gate $y[U', V']$ with $U' = \{K_B \in U_{n,k} \mid \neg x_{\{1,2\}}(K_B) = 1\}$ and

$V' = \{H \in V_{n,k} \mid \neg x_{\{1,2\}}(H) = 0\}$. It is well-known that $U_{n,k}$ and $V_{n,k}$ can be separated by counting input edges and using a single negation gate. However, it is easy to see that, by combining the latter construction with the trick above, we get monotone circuits with oracles of huge locality.

Indeed, for the problem investigated here, monotone circuits with local oracles can be seen as an intermediary model between monotone and non-monotone circuits, where the locality parameter μ restricts the computation of the extra input variables y_i .

In order to be precise, we rephrase the hypothesis \mathcal{A}_d employed in Theorem 1.2 using the notation introduced in this section.

The assumption \mathcal{A}_d . Let $d \in \mathbb{N}$, and (C, \mathcal{W}) be a monotone CLO pair with $\mathcal{W} = (W_i)_{i \in I}$, $W_i = (U_i, V_i)$, $U_i \subseteq U$ and $V_i \subseteq V$. We say that (C, \mathcal{W}) satisfies \mathcal{A}_d if there exists no $u \in U$ and $I' \subseteq I$, $|I'| > d$ such that $u \in \bigcap_{i' \in I'} U_{i'}$.

3 Phase transitions in depth-2: Proof of Theorem 1.1

Our argument relies on Claims 2.1 and 2.2 described in Section 2. We start with a straightforward adaptation of a lemma from [10].

Lemma 3.1. *Let $C(\vec{x}, \vec{y})$ be a monotone circuit, $A = U \cup V$ be a disjoint union, $W = (U, V)$, and $\mathcal{W} = (W_i)_{i \in [e]}$ be a sequence of pairs included in W , where each $W_i = (U_i, V_i)$. Then,*

1. *Over inputs $a \in A$, for every $i, j \in [e]$, the following holds:*

$$\begin{aligned} f_{(U_i, V_i)}^* \vee f_{(U_j, V_j)}^* &= f_{(U_i \cup U_j, V_i \cap V_j)}^* \\ f_{(U_i, V_i)}^* \wedge f_{(U_j, V_j)}^* &= f_{(U_i \cap U_j, V_i \cup V_j)}^* \end{aligned}$$

2. *Let $\mathcal{B} \stackrel{\text{def}}{=} \bigcup_{i \in [e]} U_i \times V_i \subseteq U \times V$, and $i, j \in [e]$. Then $(U_i \cap U_j) \times (V_i \cup V_j) \subseteq \mathcal{B}$ and $(U_i \cup U_j) \times (V_i \cap V_j) \subseteq \mathcal{B}$.*

Proof. Immediate from the definitions. □

First, we prove a weaker version of Theorem 1.1 that forbids the extra inputs $g_j(\vec{y})$. Then we use Lemma 3.1 to observe that our argument extends to the more general class of circuits.

Let $\varepsilon > 0$ be a fixed constant, and n be sufficiently large.

Case 1: $\mu = 1$. Obviously, there is a trivial monotone CLO pair (C, \mathcal{W}) with locality $\mu = 1$ that separates $U_{n,3}$ and $V_{n,3}$: C contains a single node y_1 , and $W_1 = (U_{n,3}, V_{n,3})$.

Case 2: $1/2 + \varepsilon \leq \mu \leq 1 - \varepsilon$. We start with the upper bound. In other words, we construct a monotone CLO of size $O(n^2)$ and locality $\leq 1/2 + o(1)$.⁸ Let $x_{\{i,j\}}$ for $i \neq j \in [n]$ denote the input variable corresponding to edge $\{i, j\} \in \binom{[n]}{2}$. Consider the following monotone circuit:

⁸This construction is inspired by discussions in [13].

$$C(\vec{x}, \vec{y}) \stackrel{\text{def}}{=} \bigvee_{i < j} (x_{\{i,j\}} \wedge y_{\{i,j\}}).$$

We associate to each $y_{\{i,j\}} = y[U_{\{i,j\}}, V_{\{i,j\}}]$ the sets

$$U_{\{i,j\}} \stackrel{\text{def}}{=} \{K_B \in U_{n,3} \mid \{i,j\} \subseteq B \text{ and these are the smallest elements in } B\}, \text{ and}$$

$$V_{\{i,j\}} \stackrel{\text{def}}{=} \{H \in V_{n,3} \mid \text{vertices } i \text{ and } j \text{ are in different parts of } H\}.$$

Observe that C has size $O(n^2)$.

First, we argue that this monotone CLO is correct. If the input graph is a triangle $K_B \in \{0,1\}^m$ with $B = \{i,j,k\}$, where $i < j < k$, then $x_{\{i,j\}}(K_B) = 1$. Moreover, for any monotone function $f_{\{i,j\}}$ that separates $(U_{\{i,j\}}, V_{\{i,j\}})$, we must have $f_{\{i,j\}}(K_B) = 1$, since $K_B \in U_{\{i,j\}}$ by construction. Thus $C(K_B, \mathcal{F})$ must accept K_B for all separating sequences $\mathcal{F} = (f_{\{i,j\}})$. Now let $H \in V_{n,3}$ be a complete bipartite graph over $[n]$ with non-empty parts V_1^H and V_2^H partitioning $[n]$. We show that for $i < j$ it holds that $x_{\{i,j\}}(H) \wedge y_{\{i,j\}}(H) = 0$. If for some $x_{\{i,j\}}$ we have $x_{\{i,j\}}(H) = 1$, then i, j are in different parts of H . By construction, any $f_{\{i,j\}}$ separating the pair $(U_{\{i,j\}}, V_{\{i,j\}})$ must output 0 on H . Consequently, the output of the circuit on H is 0, under any sequence \mathcal{F} of separating functions.

Next, we upper bound the locality of the y -variables. Let $\mathcal{B} = \bigcup_{i < j} U_{\{i,j\}} \times V_{\{i,j\}} \subseteq U_{n,3} \times V_{n,3}$. Let (K_B, H) be a fixed input pair in $U_{n,3} \times V_{n,3}$. Observe that this pair is in \mathcal{B} if and only if there exist $i, j \in [n]$ with $i < j$ such that:

- (1) $\{i, j\} \in B$,
- (2) these are the smallest elements in B , and
- (3) the vertices i and j belong to different parts of H .

Therefore, the locality μ of the monotone CLO defined above is upper bounded by

$$\begin{aligned} \Pr_{(K_B, H) \sim \mathcal{D}_{n,3}} [\exists i < j \text{ satisfying (1), (2), (3)}] &\leq \sum_{i < j} \Pr[(i, j) \text{ satisfies (1), (2), (3)}] \\ \text{(using independence)} &= \sum_{i < j} \Pr_{H \sim \mathcal{D}_{n,3}^V} [(i, j) \text{ satisfies (3)}] \cdot \Pr_{K_B \sim \mathcal{D}_{n,3}^U} [(i, j) \text{ satisfies (1), (2)}] \\ &= \Pr_{\chi \sim V_{n,3}^\chi} [\chi(1) \neq \chi(2) \mid \chi([n]) = \{1, 2\}] \cdot \sum_{i < j} \frac{n-j}{\binom{n}{3}} \\ &= (1/2 + o(1)) \cdot 1 \leq 1/2 + \varepsilon. \end{aligned}$$

We argue next the lower bound on circuit size for this range of μ . In other words, we prove that if $\mu \leq 1 - \varepsilon$ then the circuit size is $\Omega_\varepsilon(n^2)$. Let (C, \mathcal{W}) be a monotone CLO pair, where $C(\vec{x}, \vec{y})$ is a monotone DNF with $t \leq s$ terms, $\mathcal{W} = (W_i)_{i \in [e]}$, $e \leq s$, $W_i = (U_i, V_i)$, and each W_i is included in the pair $(U_{n,3}, V_{n,3})$. Further, let $\mathcal{B} = \bigcup_i U_i \times V_i$. Assume the pair (C, \mathcal{W}) computes 3-clique over $A_{n,3}$. In order to establish a lower bound, we consider the sequence \mathcal{F}^* , as defined in Section 2. Then, using Lemma 3.1, we can write this circuit in an equivalent way as follows:

$$C(\vec{x}, \mathcal{F}^*) = \bigvee_{j \in [t]} \left(\bigwedge_{e \in S_j} x_e \wedge f_{(U_j^*, V_j^*)}^*(\vec{x}) \right), \quad (3.1)$$

where $S_j \subseteq \binom{[n]}{2}$ and $U'_j \times V'_j \subseteq \mathcal{B}$, for each $j \in [t]$. This is without loss of generality, since terms that did not originally include a y -variable can be represented using $f_{(U_{n,3}, \emptyset)}^*$, which is equivalent to the constant 1 function over inputs in $A_{n,3}$.

Next, observe that if $|S_j| > 3$ for some $j \in [t]$ then the corresponding term cannot accept an input from $U_{n,3}$. Thus we can assume without loss of generality that $0 \leq |S_j| \leq 3$. Partition the terms of $C(\vec{x}, \mathcal{F}^*)$ into sets T_ℓ , $0 \leq \ell \leq 3$, with T_ℓ containing all terms for which $|S_j| = \ell$.

Every triangle K_B accepted by a term from T_0 forces a measure $\geq 1/\binom{n}{3}$ in \mathcal{B} , since the corresponding functions $f_{(U'_j, V'_j)}^*$ must satisfy $V'_j = V_{n,3}$ in order for the term not to accept a complete bipartite graph $H \in V_{n,3}$. Consequently, using that $\mu \leq 1 - \varepsilon$, a total number of at most $r = (1 - \varepsilon)\binom{n}{3}$ triangles can be accepted by terms in T_0 .

Now each term in T_2 or in T_3 accepts at most one triangle, and each term in T_1 accepts at most n triangles. Therefore, using the preceding paragraph, in order for the circuit to accept all $\binom{n}{3}$ triangles in $U_{n,3}$, we must have:

$$|T_1| \cdot n + |T_2| + |T_3| \geq \binom{n}{3} - r = \Omega(n^3).$$

This implies that at least one of $|T_1|$, $|T_2|$, and $|T_3|$ must be $\Omega(n^2)$. In particular, the original circuit must have size at least $\Omega(n^2)$.

Case 3: $0 \leq \mu \leq 1/2 - \varepsilon$. The $O(n^3)$ size upper bound at $\mu = 0$ is achieved by the trivial monotone circuit for 3-clique. For the lower bound, we adapt the argument presented above. Using the same notation, we assume there is a correct circuit as described in (3.1). By the same reasoning, $|S_j| \leq 3$ for each $j \in [t]$. Furthermore, we can assume that the edges corresponding to each S_j are contained in some triangle from $U_{n,3}$.

Rewrite $C(\vec{x}, \mathcal{F}^*)$ as an equivalent circuit C' :

$$C'(\vec{x}, \mathcal{F}^*) \stackrel{\text{def}}{=} \bigvee_{\ell \in I_{\leq 2}} \left(\bigwedge_{e \in S_\ell} x_e \wedge f_{(U_\ell, V_\ell)}^*(\vec{x}) \right) \vee \bigvee_{i \in I_3} \left(\bigwedge_{e \in S_i} x_e \wedge f_{(U_i, V_i)}^*(\vec{x}) \right), \quad (3.2)$$

where $I_{\leq 2}$ contains the indexes of the original sets S_j such that the edges obtained from S_j touch at most 2 vertices, and I_3 contains the indexes corresponding to sets S_j whose edges span exactly 3 vertices.

First, suppose there exists $\ell \in I_{\leq 2}$ such that $\mathcal{D}_{n,3}^V(V_\ell) \leq 1/2 - \varepsilon/4$. This implies that f_ℓ^* rejects a subset of $V_{n,3}$ of measure at most $1/2 - \varepsilon/4$. Moreover, using that $\ell \in I_{\leq 2}$, $\bigwedge_{e \in S_\ell} x_e$ rejects a subset of $V_{n,3}$ of measure at most $1/2 + \varepsilon/8$. Consequently, the ℓ -th term of the original circuit $C(\vec{x}, \mathcal{F}^*)$ must accept some negative input from $V_{n,3}$. This violates the assumption that the initial monotone CLO pair computes 3-clique over $A_{n,3}$.

We get from the previous argument that for every $\ell \in I_{\leq 2}$, $\mathcal{D}_{n,3}^V(V_\ell) \geq 1/2 - \varepsilon/4$. Consider now the quantity $\eta = |\bigcup_{\ell \in I_{\leq 2}} U_\ell|/|U_{n,3}|$, and observe that $\mu \geq \eta \cdot (1/2 - \varepsilon/4)$ by the previous density lower bound. Since we are in the case where $\mu \leq 1/2 - \varepsilon$, we obtain $\eta \leq 1 - \Omega_\varepsilon(1)$.

In turn, using the definition of η and of \mathcal{F}^* , it follows that the left-hand side of $C'(\vec{x}, \mathcal{F}^*)$ in (3.2) accepts at most a η -fraction of $U_{n,3}$. By the correctness of $C(x, \mathcal{F}^*)$, the right-hand side of the equivalent circuit $C'(\vec{x}, \mathcal{F}^*)$ must accept at least a $\Omega_\varepsilon(1)$ -fraction of the triangles in $U_{n,3}$. Now observe that for each $i \in I_3$, the corresponding

term $\bigwedge_{e \in S_i} x_e$ accepts exactly one triangle. Therefore, we must have $|I_3| \geq \Omega_\varepsilon\left(\binom{n}{3}\right)$. This completes the proof that $t = \Omega(n^3)$.

In order to prove lower bounds in the presence of $g_j(\vec{y})$ input variables, observe that the following holds. First, all lower bounds were obtained using \mathcal{F}^* . Due to Lemma 3.1, each $g_j(\vec{y})$ is equivalent over $A_{n,3}$ to $f_{(U'_j, V'_j)}^*$, for an appropriate pair (U'_j, V'_j) satisfying $U'_j \times V'_j \subseteq \mathcal{B}$. Finally, in addition to the locality bound, the inclusion in \mathcal{B} is the only information about the y -variables that was employed in the proofs. In other words, each $g_j(\vec{y})$ can be treated as a new y -variable in the arguments above, without affecting the locality bounds.

This extends the lower bound to the desired class of circuits, and completes the proof of Theorem 1.1.

4 Circuits with restricted oracles: Proof of Theorem 1.2

We start with the upper bound.

Lemma 4.1. *Let $3 \leq k \leq n^{1/4}$ and $2 \leq \ell < k$. There exists a monotone circuit with local oracles $E(\vec{x}, \vec{y})$ of size $O\left(\binom{n}{\ell} \cdot \binom{\ell}{2}\right)$ and locality $\mu \leq \exp(-\Omega(\ell^2/k))$ that computes k -clique over $A_{n,k}$. Furthermore, the local oracles associated to E satisfy condition \mathcal{A}_1 .*

Proof. We generalize a construction in the proof of Theorem 1.1. For every set $B \in \binom{[n]}{k}$, let $F(B) \in \binom{[B]}{\ell}$ be the lexicographic first ℓ -sized subset of B . Consider the following monotone circuit with local oracles:

$$E(\vec{x}, \vec{y}) \stackrel{\text{def}}{=} \bigvee_{D \in \binom{[n]}{\ell}} \left(\bigwedge_{e \in \binom{D}{2}} x_e \wedge y_D \right),$$

where to each y_D we associate a pair (U_D, V_D) with $U_D \times V_D \subseteq U_{n,k} \times V_{n,k}$, defined as follows:

$$U_D \stackrel{\text{def}}{=} \{K_B \in U_{n,k} \mid F(B) = D\} \quad \text{and} \quad V_D \stackrel{\text{def}}{=} \{H \in V_{n,k} \mid K_D \subseteq H\}.$$

By construction, $U_D \cap U_{D'} = \emptyset$ for distinct $D, D' \in \binom{[n]}{\ell}$. In other words, assumption \mathcal{A}_1 is satisfied. Further, the size of E is $O\left(\binom{n}{\ell} \cdot \binom{\ell}{2}\right)$. The correctness of this monotone CLO can be established by a straightforward generalization of the argument from Section 3. It remains to estimate its locality parameter μ .

Fix a set $D \in \binom{[n]}{\ell}$, and let $\gamma_D \stackrel{\text{def}}{=} \mathcal{D}_{n,k}^V(V_D)$. By symmetry, $\gamma_D = \gamma_{D'}$ for every $D' \in \binom{[n]}{\ell}$. Since distinct sets U_D are pairwise disjoint and locality is measured with respect to the product distribution $\mathcal{D}_{n,k} = \mathcal{D}_{n,k}^U \times \mathcal{D}_{n,k}^V$, the locality of the oracle rectangles

associated with E is at most γ_D . This value can be upper bounded as follows:

$$\begin{aligned}
 \gamma_D &= \Pr_{H \sim \mathcal{D}_{n,k}^V} [K_D \subseteq H] = \Pr_{\chi \sim V_{n,k}^\chi} [K_D \subseteq G(\chi) \mid G(\chi) \in V_{n,k}] \\
 &= \frac{\Pr_\chi [K_D \subseteq G(\chi) \wedge G(\chi) \in V_{n,k}]}{\Pr_\chi [G(\chi) \in V_{n,k}]} \\
 &\leq \frac{\Pr_\chi [K_D \subseteq G(\chi)]}{\Pr_\chi [|\chi([n])| = k-1]} \\
 &\text{(using } 3 \leq k \leq n^{1/4} \text{ and } n \rightarrow \infty) \leq (1 + o(1)) \cdot \frac{(k-1)(k-2) \dots (k-\ell)}{(k-1)^\ell} \\
 &\leq (1 + o(1)) \cdot \frac{(k - \lfloor \ell/2 \rfloor)^{\ell/2}}{(k-1)^{\ell/2}} \\
 &= (1 + o(1)) \cdot \left(1 - \frac{\lfloor \ell/2 \rfloor - 1}{k-1}\right)^{\ell/2} \\
 &\text{(using } (1-x) \leq e^{-x} \text{ and } 0 \leq x \leq 1) \leq \exp(-\Omega(\ell^2/k)).
 \end{aligned}$$

This completes the proof of Lemma 4.1. \square

The upper bound in Theorem 1.2 follows immediately from Lemma 4.1, by taking a large enough $\ell = O(\sqrt{k})$. Observe that, more generally, one can get a trade-off between circuit size and locality.

We move on now to the lower bound part, which relies on a sequence of lemmas. For a set $X \subseteq [n]$, we let $[X] \stackrel{\text{def}}{=} \bigwedge_{\{i,j\} \in \binom{X}{2}} x_{\{i,j\}}$ be the corresponding clique indicator circuit. For convenience, we define $[X] \stackrel{\text{def}}{=} 1$ if X is a singleton or the empty set. Also, note that $[X] = [X] \wedge f_{U_{n,k}, \emptyset}^*$ over $A_{n,k}$. Under this notation, we don't need to consider standalone terms in the lemma below, which adapts to our setting a result from [10].

Lemma 4.2. *Let $\mathcal{W} = (W_i)$ with $W_i = (U_i, V_i)$ be a sequence of pairs included in $(U_{n,k}, V_{n,k})$. Let $C(\vec{x}, \vec{y})$ be a monotone circuit with local oracles of the form*

$$C(x, y) = \bigvee_{i \in [t]} \left([X_i] \wedge y[U_i, V_i] \right),$$

where t is arbitrary, $|X_i| \leq \lfloor \sqrt{k} \rfloor$, $k(n) \geq 5$, and all rectangles $U_i \times V_i \subseteq \mathcal{B}$, for some set $\mathcal{B} \subseteq U_{n,k} \times V_{n,k}$ of locality $\mu \leq 1/16$. Then, for large enough n , the following holds.

1. $C(x, \mathcal{F}^*)$ accepts a subset of $V_{n,k}$ of measure at least $1/10$, or
2. $C(x, \mathcal{F}^*)$ rejects a subset of $U_{n,k}$ of measure at least $1/10$.

Proof. If $t = 0$ the circuit computes a constant function, and consequently one of the items above must hold. Otherwise, for each $i \in [t]$, since $U_i \times V_i \subseteq \mathcal{B}$ and $\mathcal{D}_{n,k} = \mathcal{D}_{n,k}^U \times \mathcal{D}_{n,k}^V$, we have that either $\mathcal{D}_{n,k}^U(U_i) \leq \mu^{1/2}$ or $\mathcal{D}_{n,k}^V(V_i) \leq \mu^{1/2}$. We consider two cases.

First, assume there is $i \in [t]$ such that $\mathcal{D}_{n,k}^V(V_i) \leq \mu^{1/2} \leq 1/4$. Then,

$$\Pr_{H \sim \mathcal{D}_{n,k}^V} \left[([X_i] \wedge f_i^*)(H) = 1 \right] \geq 1 - \Pr[X_i(H) = 0] - \Pr[H \in V_i] \geq 3/4 - \Pr[X_i(H) = 0].$$

The latter probability is 0 if $|X_i| \leq 1$. Otherwise, it can be upper bounded by

$$\begin{aligned} \Pr_{\chi \sim V_{n,k}^\chi} [|\chi(X_i)| < |X_i| \mid G(\chi) \in V_{n,k}] &\leq (1 + o(1)) \cdot \sum_{\{a,b\} \in \binom{X_i}{2}} \Pr_{\chi \sim V_{n,k}^\chi} [\chi(a) = \chi(b)] \\ &\text{(since } |X_i| \leq \lfloor \sqrt{k} \rfloor \text{)} \leq (1 + o(1)) \cdot \binom{\lfloor \sqrt{k} \rfloor}{2} \cdot \frac{k-1}{(k-1)^2}. \end{aligned}$$

This shows that item 1 above holds, using $k \geq 5$ and the previous estimate.

If there is no $i \in [t]$ satisfying $\mathcal{D}_{n,k}^V(V_i) \leq \mu^{1/2}$, by the observation in the first paragraph of this proof we get that $\mathcal{D}_{n,k}^U(U_i) \leq \mu^{1/2}$ and $\mathcal{D}_{n,k}^V(V_i) > \mu^{1/2}$ for all $i \in [t]$. Recall that the measure of \mathcal{B} is at most $\mu \leq 1/16$. Therefore, it must be the case that $|\bigcup_i U_i|/|U_{n,k}| \leq \mu^{1/2}$, as each K_B in this union contributes at least $\mu^{1/2}$ to the measure of \mathcal{B} . Due to our choice of \mathcal{F}^* and the structure of C , $C(\vec{x}, \mathcal{F}^*)$ will accept at most a $(1/4)$ -fraction of $U_{n,k}$, and item 2 holds. \square

Crucially, Lemma 4.2 requires no upper bound on the number of terms appearing in C , and this will play a fundamental role in the argument below.

For the rest of the proof, let $D(\vec{x}, \vec{y})$ be a monotone CLO of size s that computes k -clique over $A_{n,k}$, and $W_i = (V_i, U_i)$ for $i \leq e$ be its associated pairs, where $e \leq s$. As usual, we set $\mathcal{B} = \bigcup_i U_i \times V_i$. Recall the extra condition on the local oracle gates.

Assumption \mathcal{A}_d : If $J \subseteq [e]$ and $|J| > d$, then $\bigcap_{j \in J} U_j = \emptyset$.

We can assume without loss of generality that different oracle variables appearing in the description of the circuit are associated to distinct subsets of $U_{n,k}$. Indeed, due to monotonicity (cf. Claim 2.2), we can always take a larger subset of $V_{n,k}$ if different oracle variables are associated to the same subset of $U_{n,k}$. A bit more precisely, if $y_i = y_i[U', V_i]$ and $y_j = y_j[U', V_j]$, we can redefine these local oracles to use the pair $(U', V_i \cup V_j)$. This does not increase the overall locality, and does not change the correctness of the computation. Note that this transformation produces oracle variables associated to the same pair of subsets, but since we use boolean circuits instead of boolean formulas, oracle variables don't need to be repeated in the description of the circuit.

For $J \subseteq [e]$, we use $D_J(\vec{x})$ to denote the circuit with y_j substituted by 1 if $j \in J$, and by 0 otherwise. In particular, each D_J is a monotone circuit in the usual sense, i.e., it does not contain local oracle gates. Moreover, $\text{size}(D_J) \leq \text{size}(D)$.

Lemma 4.3. *Under Assumption \mathcal{A}_d , for every input graph $G \in A_{n,k}$,*

$$D(G, \mathcal{F}^*) = \bigvee_{J \in \binom{[e]}{\leq d}} D_J(G) \wedge f_{(U_J, V_J)}^*(G),$$

where $U_J \stackrel{\text{def}}{=} \bigcap_{j \in J} U_j$ and $V_J \stackrel{\text{def}}{=} \bigcup_{j \in J} V_j$ (here an empty intersection is $U_{n,k}$ and an empty union is \emptyset , corresponding to the case where $J = \emptyset$).

Proof. First, observe that for inputs in $A_{n,k}$,

$$D(\vec{x}, \mathcal{F}^*) \equiv \bigvee_{J \subseteq [e]} \left(D_J(\vec{x}) \wedge \bigwedge_{j \in J} f_j^*(\vec{x}) \wedge \bigwedge_{j \notin J} \neg f_j^*(\vec{x}) \right),$$

using our definition of $D_J(\vec{x})$. As we explain below, this circuit is further equivalent to a circuit where we drop the negated part:

$$D(\vec{x}, \mathcal{F}^*) \equiv \bigvee_{J \subseteq [e]} \left(D_J(\vec{x}) \wedge \bigwedge_{j \in J} f_j^*(\vec{x}) \right).$$

Clearly, by eliminating some “literals” we can only accept more inputs. However, by monotonicity the latter is not going to happen. Indeed, if we have a term and a negative input $H \in V_{n,k}$ such that $D_J(H) \wedge \bigwedge_{j \in J} f_j^*(H) = 1$ but $\bigwedge_{j \notin J} \neg f_j^*(H) = 0$, then there is a set J' with $J \subseteq J' \subseteq [e]$ such that $D_{J'}(H) \wedge \bigwedge_{j \in J'} f_j^*(H) \wedge \bigwedge_{j \notin J'} \neg f_j^*(H) = 1$, where we have used the monotonicity of $D(\vec{x}, \vec{y})$ in order to claim that $D_{J'}(H) \geq D_J(H)$. This is impossible, since by assumption $D(\vec{x}, \mathcal{F}^*)$ separates $U_{n,k}$ and $V_{n,k}$.

Using Lemma 3.1, we know that $\bigwedge_{j \in J} f_j^* = f_{(U_J, V_J)}^*$, for U_J and V_J as in the statement of the lemma. Under assumption \mathcal{A}_d , whenever $|J| > d$ we get $U_J = \emptyset$. Therefore,

$$D(\vec{x}, \mathcal{F}^*) \equiv \bigvee_{J \in \binom{[e]}{\leq d}} \left(D_J(\vec{x}) \wedge f_{(U_J, V_J)}^*(\vec{x}) \right) \vee \bigvee_{J \in \binom{[e]}{> d}} \left(D_J(\vec{x}) \wedge f_{(\emptyset, V_J)}^*(\vec{x}) \right). \quad (4.1)$$

Using the equivalences established above and the correctness of the original circuit, the circuit in (4.1) accepts every input in $U_{n,k}$, and rejects every input in $V_{n,k}$. Now observe that the right-hand terms of the circuit cannot accept an input in $V_{n,k}$, due to the presence of the functions $f_{(\emptyset, V_J)}^*$. Thus such terms can be discarded, and the circuit obtained after this simplification still accepts $U_{n,k}$ and rejects $V_{n,k}$. This completes the proof of the lemma. \square

Observe that $U_J \times V_J \subseteq \mathcal{B}$ for every $J \subseteq [e]$, due to Lemma 3.1. In particular, the simplification above is well-behaved with respect to the new oracle rectangles introduced in the transformation.

The next steps of our argument rely on results from Alon and Boppana [1] related to the approximation method [12]. We follow the terminology of the exposition in Boppana and Sipser [3, Section 4.2]. For the rest of the proof, we let $\ell \stackrel{\text{def}}{=} \lfloor \sqrt{k} \rfloor$, $p \stackrel{\text{def}}{=} \lceil 10\sqrt{k} \log n \rceil$, and $m \stackrel{\text{def}}{=} (p-1)^\ell \cdot \ell!$. (Recall that ℓ is the size of each indicator set $[X_i]$, m is the maximum number of indicators in each approximator, and p is an auxiliary parameter.⁹)

Approximate each individual circuit $D_J(\vec{x})$ as in Boppana-Sipser, obtaining a corresponding depth-2 approximator $\tilde{D}_J(\vec{x})$. Since each $D_J(\vec{x})$ is a monotone circuit of size at most s , our choice of $U_{n,k}$ and $V_{n,k}$ and the argument in [3] provide the following bounds.

Lemma 4.4. [3, Lemma 4.3]. *For each $J \subseteq [e]$, the number of positive test graphs $G \in U_{n,k}$ for which $D_J(G) \leq \tilde{D}_J(G)$ does not hold is at most $E^+ \stackrel{\text{def}}{=} s \cdot m^2 \cdot \binom{n-\ell-1}{k-\ell-1}$.*

Lemma 4.5. [3, Lemma 4.4]. *For each $J \subseteq [e]$, the number of negative test graphs (colorings) $\chi \in V_{n,k}^\chi$ for which $D_J(G(\chi)) \geq \tilde{D}_J(G(\chi))$ does not hold is at most $E^- \stackrel{\text{def}}{=} s \cdot m^2 \cdot \left[\binom{l}{2} / (k-1) \right]^p \cdot (k-1)^n$.*

⁹Do not confuse this definition of m with the number of edges in the input graph, which will not be needed in the rest of the proof.

Now define using D and the individual approximators \tilde{D}_J a corresponding monotone circuit $\tilde{D}(\vec{x}, \vec{y})$ with access to the functions $f_{(U_J, V_J)}^*$:

$$\tilde{D}(\vec{x}, \mathcal{F}^*) \stackrel{\text{def}}{=} \bigvee_{J \in \binom{[e]}{\leq d}} (\tilde{D}_J(\vec{x}) \wedge f_{(U_J, V_J)}^*(\vec{x})). \quad (4.2)$$

Clearly, $D(G, \mathcal{F}^*) \neq \tilde{D}(G, \mathcal{F}^*)$ on an input $G \in A_{n,k}$ only if for some approximator \tilde{D}_J we have $\tilde{D}_J(G) \neq D_J(G)$. Furthermore, at most $\sum_{j=0}^d \binom{e}{j} \leq (e+1)^d \leq (s+1)^d$ distinct circuits D_J are approximated. Combining this with Lemmas 4.4 and 4.5, a union bound, and the fact that the original circuit is correct on every input graph in $A_{n,k}$, we get:

$$\Pr_{G \sim \mathcal{D}_{n,k}^U} [\tilde{D}(G, \mathcal{F}^*) = 1] \geq 1 - (s+1)^d \cdot \frac{E^+}{\binom{n}{k}},$$

and similarly,

$$\begin{aligned} \Pr_{H \sim \mathcal{D}_{n,k}^V} [\tilde{D}(H, \mathcal{F}^*) = 0] &\geq (1 - o(1)) \cdot \Pr_{\chi \sim V_{n,k}^\chi} [\tilde{D}(G(\chi), \mathcal{F}^*) = 0 \wedge G(\chi) \in V_{n,k}] \\ &\geq (1 - o(1)) \cdot (1 - \Pr_{\chi} [\tilde{D}(G(\chi), \mathcal{F}^*) = 1] - o(1)) \\ &\geq (1 - o(1)) \cdot \left(1 - (s+1)^d \cdot \frac{E^-}{(k-1)^n}\right). \end{aligned}$$

We can assume each one of these probabilities $\rightarrow 1$ as $n \rightarrow \infty$, since otherwise we get that $s \geq n^{\Omega(\sqrt{k}/d)}$ using the values of E^- , E^+ , p , ℓ , and m , completing the proof of Theorem 1.2. In more detail, let $\delta > 0$ be an arbitrary small constant, and suppose that:

$$(s+1)^d \cdot \frac{s \cdot m^2 \cdot \binom{n-\ell-1}{k-\ell-1}}{\binom{n}{k}} \geq \delta \quad \text{or} \quad (s+1)^d \cdot \frac{s \cdot m^2 \cdot \binom{\ell}{2} / (k-1)^p \cdot (k-1)^n}{(k-1)^n} \geq \delta.$$

Due to the upper bound on k in the statement of Theorem 1.2, using estimates entirely analogous to the ones employed in [3] (which are routine and left to the reader), it follows in each case that:

$$(s+1)^{d+1} \geq n^{\Omega(\sqrt{k})}.$$

This justifies the claim made above on the convergence of the probabilities.

Now expand each term $\tilde{D}_J(\vec{x}) \wedge f_{(U_J, V_J)}^*(\vec{x})$ in $\tilde{D}(\vec{x}, \mathcal{F}^*)$ (Equation 4.2), using that (see [3]) each circuit $\tilde{D}_J(\vec{x})$ is either a union of clique indicators of bounded size:

$$\tilde{D}_J(\vec{x}) \equiv \bigvee_{i \in [m_J]} [X_i^J]$$

for $m_J \leq m$ and an appropriate choice of sets $X_i^J \subseteq [n]$ satisfying $0 \leq |X_i^J| \leq \ell$, or $\tilde{D}_J \equiv 0$. This produces a circuit equivalent to $\tilde{D}(\vec{x}, \mathcal{F}^*)$ over inputs in $A_{n,k}$, and it can be written in the following form:

$$\tilde{D}(\vec{x}, \mathcal{F}^*) \equiv \bigvee_{i \in [t]} \left([X_i] \wedge f_{(U_i', V_i')}^*(\vec{x}) \right) \quad (4.3)$$

Here t can be arbitrarily large, but observe that $U_i' \times V_i' \subseteq \mathcal{B}$ for every $i \in [t]$ (due to Lemmas 3.1 and 4.3). We don't assume that $(U_i', V_i') \neq (U_{i'}', V_{i'}')$ when $i \neq i'$, and similarly for X_i and $X_{i'}$.

Finally, we know that the circuit in Equation 4.3 accepts a subset of $U_{n,k}$ of measure $1 - o(1)$, and that it rejects a subset of $V_{n,k}$ of measure $1 - o(1)$. By construction, each clique indicator in the description of \tilde{D} has size at most $\ell \leq \lfloor \sqrt{k} \rfloor$. Together with $U_i' \times V_i' \subseteq \mathcal{B}$ for every $i \in [t]$ and the upper bound on the locality of \mathcal{B} , we get a contradiction to Lemma 4.2.

The proof of Theorem 1.2 is complete. Observe that, under the same assumptions, it is possible to obtain a slightly stronger trade-off of the form: $e^d \cdot s \geq n^{\Omega(\sqrt{k})}$.

5 Concluding remarks

We discuss below some questions and directions motivated by our results, and elaborate a bit more on the connection to proof complexity.

Monotone circuit complexity. The main open problem in the context of circuit complexity is to understand the size of monotone circuits of small locality separating the sets $U_{n,k}$ and $V_{n,k}$, under no further assumption on the y -variables. It is not clear if the hypothesis \mathcal{A}_d in Theorem 1.2 is an artifact of our proof. As far as we know, it is conceivable that smaller circuits can be designed by increasing the overlap between the sets U_i .¹⁰

However, if one is more inclined to lower bounds, we mention that the fusion approach described in [7] can be easily adapted to monotone circuit with local oracles, and that this point of view might be helpful in future investigations of unrestricted monotone CLOs.

Another question of combinatorial interest is whether the phase transitions observed in Theorem 1.1 extend to more expressive classes of monotone circuits beyond depth two. More broadly, are the phase transitions observed here particular to k -clique, or an instance of a more general phenomenon connected to computations using monotone circuits extended with oracle gates?

Corollary 1.3 suggests the following problem. Is it possible to refine the approach from [1], and to prove that the monotone circuit size complexity of k -clique is $n^{\Omega(k)}$ for a larger range of k ? In a related direction, it would be interesting to understand if monotone CLOs can shed light into the difficulties in proving stronger monotone circuit size lower bounds for other boolean functions of interest, such as the matching problem on graphs (see e.g. [1, Section 5] and [6, Section 9.11]).

Proof complexity. Back to the original motivation from proof complexity, we have been unable so far to transform proofs in $\mathbf{R}(\text{Lin}/\mathbb{F}_2)$ into monotone CLOs satisfying \mathcal{A}_d , for $d \leq k^{1/2-\epsilon}$, or certain variations of \mathcal{A}_d under which Theorem 1.2 still holds.

¹⁰We notice that non-monotone polynomial size circuits containing oracles of small locality can compute any boolean function (see [10, Section 3]). A similar phenomenon appears in the adaptation of real-valued monotone circuits to general real-valued circuits [11, Section 7], but in that case strong lower bounds are known against monotone real-valued circuits.

Observe that, using the connections established in [10], this would be sufficient for exponential lower bounds on proof size.

The reduction from randomized feasible interpolation actually provides a distribution on monotone CLOs C_r with a common bound on their sizes such that each is correct and they satisfy:

$$\Pr_r[(u, v) \in \mathcal{B}_r] \leq \mu \text{ for every fixed pair } (u, v) \in U \times V,$$

where \mathcal{B}_r is the union of the oracle rectangles in C_r . An averaging argument then yields a fixed monotone CLO whose locality is bounded by μ . One might lose some information useful for a lower bound in this last step depending on the choice of the distribution \mathcal{D} supported over $U \times V$.

Even though our initial attempts at establishing new length-of-proofs lower bounds have been unsuccessful, we feel that in order to prove limitations for $\text{R}(\text{Lin}/\mathbb{F}_2)$ and for other proof systems via randomized feasible interpolation it should be sufficient to establish lower bounds against monotone CLOs under an appropriate assumption on the oracle gates. (In particular, the existence of monotone CLOs of small size and small locality separating $U_{n,k}$ and $V_{n,k}$ does not imply that the approach presented in [10] is fruitless.) For instance, while \mathcal{A}_d is a semantic condition on the (unstructured) sets U_i and V_i , one can try to explore the syntactic information obtained on these sets from a given proof, such as upper bounds on the circuit complexity of separating each pair U_i and V_i , or other related structural information.

Acknowledgements. We would like to thank Pavel Pudlák for discussions on monotone circuits with local oracles and proof complexity. The second author is grateful to Michal Garlík for several related conversations. Finally, we thank the anonymous reviewers for helpful comments regarding the presentation.

References

- [1] NOGA ALON AND RAVI B. BOPANA: The monotone circuit complexity of boolean functions. *Combinatorica*, 7(1):1–22, 1987. [2](#), [4](#), [14](#), [16](#)
- [2] MARIA LUISA BONET, TONIANN PITASSI, AND RAN RAZ: Lower bounds for cutting planes proofs with small coefficients. *J. Symbolic Logic*, 62(3):708–728, 1997. [2](#)
- [3] RAVI B. BOPANA AND MICHAEL SIPSER: The complexity of finite functions. In *Handbook of Theoretical Computer Science, Volume A: Algorithms and Complexity*, pp. 757–804. 1990. [14](#), [15](#)
- [4] SAMUEL BUSS, LESZEK KOŁODZIEJCZYK, AND KONRAD ZDANOWSKI: Collapsing modular counting in bounded arithmetic and constant depth propositional proofs. *Transactions of the American Mathematical Society*, 367(11):7517–7563, 2015. [2](#)
- [5] DMITRY ITSYKSON AND DMITRY SOKOLOV: Lower bounds for splittings by linear combinations. In *Mathematical Foundations of Computer Science (MFCS)*, pp. 372–383, 2014. [2](#)

- [6] STASYS JUKNA: *Boolean Function Complexity - Advances and Frontiers*. Springer, 2012. [16](#)
- [7] MAURICIO KARCHMER: On proving lower bounds for circuit size. In *Structure in Complexity Theory Conference (CCC)*, pp. 112–118, 1993. [16](#)
- [8] JAN KRAJÍČEK: *Bounded Arithmetic, Propositional Logic, and Complexity Theory*. Cambridge University Press, 1995. [2](#)
- [9] JAN KRAJÍČEK: Interpolation theorems, lower bounds for proof systems, and independence results for bounded arithmetic. *J. Symbolic Logic*, 62(2):457–486, 1997. [2](#)
- [10] JAN KRAJÍČEK: Randomized feasible interpolation and monotone circuits with a local oracle. Available at arXiv:1611.08680, 2016. [1](#), [2](#), [7](#), [8](#), [12](#), [16](#), [17](#)
- [11] PAVEL PUDLÁK: Lower bounds for resolution and cutting plane proofs and monotone computations. *J. Symbolic Logic*, 62(3):981–998, 1997. [2](#), [16](#)
- [12] ALEXANDER A. RAZBOROV: Lower bounds on the monotone complexity of some boolean functions. *Soviet Math. Doklady*, 31:354–357, 1985. [2](#), [4](#), [14](#)
- [13] ROBERT ROBERE: Average-case lower bounds for monotone switching networks, 2013. (Masters thesis, University of Toronto). [8](#)
- [14] BENJAMIN ROSSMAN: The monotone complexity of k -clique on random graphs. *SIAM J. Comput.*, 43(1):256–279, 2014. [2](#)

AUTHORS

Jan Krajíček
Professor
Faculty of Mathematics and Physics, Charles University, Czech Republic
krajicek@karlin.mff.cuni.cz
<https://www.karlin.mff.cuni.cz/~krajicek/>

Igor C. Oliveira
Researcher
Department of Computer Science, University of Oxford, United Kingdom
igor.carboni.oliveira@cs.ox.ac.uk
<https://www.cs.ox.ac.uk/people/igor.carbonioliveira/>