

Monotone Circuit Lower Bounds from Resolution*

Ankit Garg
Microsoft Research
USA

Pritish Kamath
MIT
USA

Mika Göös
Harvard University
USA

Dmitry Sokolov
KTH Royal Institute of Technology
Sweden

ABSTRACT

For any unsatisfiable CNF formula F that is hard to refute in the *Resolution* proof system, we show that a gadget-composed version of F is hard to refute in any proof system whose lines are computed by efficient communication protocols—or, equivalently, that a monotone function associated with F has large monotone circuit complexity. Our result extends to monotone *real* circuits, which yields new lower bounds for the *Cutting Planes* proof system.

CCS CONCEPTS

• **Theory of computation** → **Computational complexity and cryptography**;

KEYWORDS

Circuit complexity, proof complexity

ACM Reference Format:

Ankit Garg, Mika Göös, Pritish Kamath, and Dmitry Sokolov. 2018. Monotone Circuit Lower Bounds from Resolution. In *Proceedings of 50th Annual ACM SIGACT Symposium on the Theory of Computing (STOC'18)*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3188745.3188838>

1 APPETIZER

Dag-like communication protocols [38, 44, 50], generalizing the usual notion of *tree-like* communication protocols [30, 34, 39], provide a useful abstraction to study two kinds of objects in complexity theory:

- **Monotone circuits.** Let f be a monotone boolean function. The *monotone circuit complexity* of f can be characterized in the language of dag-like protocols. Namely, it equals the least size of a dag-like protocol that solves the *monotone Karchmer–Wigderson (mKW)* search problem associated with f .
- **Propositional proofs.** Let F be a CNF contradiction (an unsatisfiable CNF formula). Lower bounds for the *Resolution refutation size (aka length) complexity* of F —or indeed

lower bounds for any propositional proof system whose lines are computed by efficient communication protocols—can be proved via dag-like protocols. Namely, a lower bound is given by the least size of a dag-like protocol that solves a certain CNF search problem associated with F .

In this paper, we prove a **query-to-communication lifting theorem** that escalates lower bounds for a dag-like query model (essentially Resolution) to lower bounds for dag-like communication protocols. In particular, this yields a new technique to prove size lower bounds for monotone circuits and several types of proof systems (including Cutting Planes).

The result can be interpreted as a **converse to monotone feasible interpolation** [10, 32], which is a popular method to prove refutation size lower bounds for proof systems (such as Resolution and Cutting Planes) by reductions to monotone circuit lower bounds. A theorem of this type was conjectured by Beame, Pitassi, and Huynh [5, §6]. We also note that lifting theory for deterministic *tree-like* protocols—with applications to monotone *formula* size, *tree-like* refutation size, and size–space tradeoffs—has been developed in quite some detail [11, 13, 19, 20, 27, 40, 52]. We import techniques from this line of work into the dag-like setting.

We formalize our result in Section 3 after we have carefully defined our dag-like models in Section 2.

2 DAG-LIKE MODELS

We define all computational models as solving *search problems*, defined by a relation $S \subseteq \mathcal{I} \times \mathcal{O}$ for some finite input and output sets \mathcal{I} and \mathcal{O} . On input $x \in \mathcal{I}$ the search problem is to find some output in $S(x) := \{o \in \mathcal{O} : (x, o) \in S\}$. We always assume S is *total* so that $S(x) \neq \emptyset$ for all $x \in \mathcal{I}$. We also define $S^{-1}(o) := \{x \in \mathcal{I} : (x, o) \in S\}$. For applications, the two most important examples of search problems, one associated with a monotone function $f: \{0, 1\}^n \rightarrow \{0, 1\}$, another with an n -variable CNF contradiction $F = \bigwedge_i D_i$ (where D_i are disjunctions of literals), are as follows.

*The full version of this work is available online at [16]

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

STOC'18, June 25–29, 2018, Los Angeles, CA, USA

© 2018 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-5559-9/18/06...\$15.00

<https://doi.org/10.1145/3188745.3188838>

mKW search problem S_f :

input: a pair $(x, y) \in f^{-1}(1) \times f^{-1}(0)$

output: a coordinate $i \in [n]$ such that $x_i > y_i$

CNF search problem S_F :

input: an n -variable truth assignment $z \in \{0, 1\}^n$

output: clause D of F unsatisfied by z , i.e., $D(z) = 0$

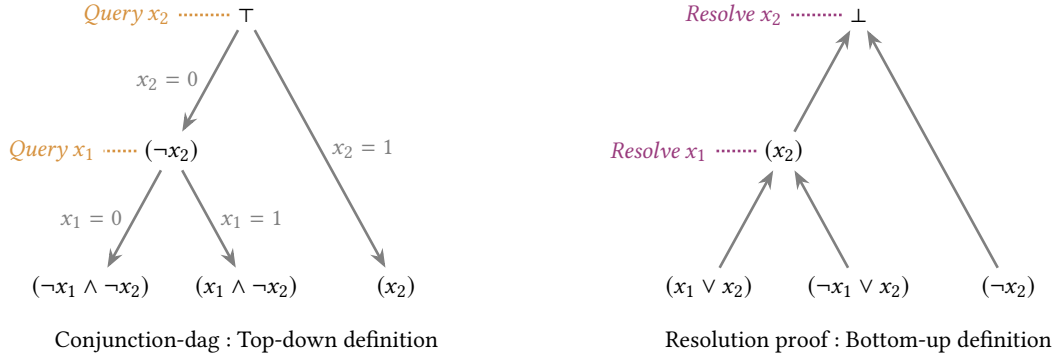


Figure 1: Two equivalent ways to view a Resolution refutation, illustrated in the tree-like case (see [30, §18.2] for more discussion of the tree-like case).

2.1 Abstract Dags

We work with a *top-down* definition of dag-like models. A version of the following definition (with a specialized \mathcal{F}) was introduced by [44] and subsequently simplified in [38, 50].

Top-down definition. Let \mathcal{F} be a family of functions $\mathcal{I} \rightarrow \{0, 1\}$. An \mathcal{F} -dag solving $S \subseteq \mathcal{I} \times \mathcal{O}$ is a directed acyclic graph of fan-out ≤ 2 where each node v is associated with a function $f_v \in \mathcal{F}$ (we call $f_v^{-1}(1)$ the *feasible set* for v) satisfying the following:

- (1) *Root:* There is a distinguished root node r (fan-in 0), and $f_r \equiv 1$ is the constant 1 function.
- (2) *Non-leaves:* For each non-leaf node v with children u, u' (perhaps $u = u'$), we have $f_v^{-1}(1) \subseteq f_u^{-1}(1) \cup f_{u'}^{-1}(1)$.
- (3) *Leaves:* Each leaf node v is labeled with an output $o_v \in \mathcal{O}$ such that $f_v^{-1}(1) \subseteq S^{-1}(o_v)$.

The *size* of an \mathcal{F} -dag is its number of nodes. If we specialize S to be a CNF search problem S_F , the above specializes to the familiar definition of refutations in a proof system whose lines are *negations* of functions in \mathcal{F} . Here is that dual definition, specialized to $S = S_F$.

Bottom-up definition. Let \mathcal{G} be a family of functions $\{0, 1\}^n \rightarrow \{0, 1\}$. (To match up with the top-down definition, one should take $\mathcal{G} := \{\neg f : f \in \mathcal{F}\}$.) A (semantic) \mathcal{G} -refutation of an n -variable CNF contradiction F is a directed acyclic graph of fan-out ≤ 2 where each node (or *line*) v is associated with a function $g_v \in \mathcal{G}$ satisfying the following:

- (1) *Root:* There is a distinguished root node r (fan-in 0), and $g_r \equiv 0$ is the constant 0 function.
- (2) *Non-leaves:* For each non-leaf node v with children u, u' (perhaps $u = u'$), we have $g_v^{-1}(1) \supseteq g_u^{-1}(1) \cap g_{u'}^{-1}(1)$.
- (3) *Leaves:* Each leaf node v is labeled with a clause D of F such that $g_v^{-1}(1) \supseteq D^{-1}(1)$.

2.2 Concrete Dags

We now instantiate the abstract model for the purposes of communication and query complexity.

Rectangle-dags (dag-like protocols). Consider a bipartite input domain $\mathcal{I} := \mathcal{X} \times \mathcal{Y}$ so that Alice holds $x \in \mathcal{X}$, Bob holds $y \in \mathcal{Y}$, and let \mathcal{F} be the set of all indicator functions of (*combinatorial*) *rectangles over $\mathcal{X} \times \mathcal{Y}$* (sets of the form $X \times Y$ with $X \subseteq \mathcal{X}$, $Y \subseteq \mathcal{Y}$). Call such \mathcal{F} -dags simply *rectangle-dags*. For a search problem $S \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{O}$ we define its *rectangle-dag complexity* by

$$\text{rect-dag}(S) := \text{least size of a rectangle-dag that solves } S.$$

In circuit complexity, a straightforward generalization of the Karchmer–Wigderson depth characterization [31] shows that the monotone circuit complexity of any monotone function f equals $\text{rect-dag}(S_f)$; see [38, 50].

In proof complexity, a useful-to-study semantic proof system is captured by \mathcal{F}_c -dags solving CNF search problems S_F where \mathcal{F}_c is the family of all functions $\mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$ (where $\mathcal{X} \times \mathcal{Y} = \{0, 1\}^n$ corresponds to a bipartition of the n input variables of S_F) that can be computed by tree-like protocols of communication cost c , say for $c = \text{polylog}(n)$. Such a proof system can simulate other systems (such as Resolution and Cutting Planes with bounded coefficients), and hence lower bounds against \mathcal{F}_c -dags imply lower bounds for other concrete proof systems. Moreover, any \mathcal{F}_c -dag can be simulated by a rectangle-dag with at most a factor 2^c blow-up in size, and hence we do not lose much generality by studying only rectangle-dags.

Conjunction-dags (essentially Resolution). Consider the n -bit input domain $\mathcal{I} := \{0, 1\}^n$ and let \mathcal{F} be the set of all *conjunctions of literals over the n input variables*. Call such \mathcal{F} -dags simply *conjunction-dags*. We define the *width* of a conjunction-dag Π as the maximum width of a conjunction associated with a node of Π . For a search problem $S \subseteq \{0, 1\}^n \times \mathcal{O}$ we define

$$\begin{aligned} \text{conj-dag}(S) &:= \text{least size of a conjunction-dag that solves } S, \\ w(S) &:= \text{least width of a conjunction-dag that solves } S. \end{aligned}$$

In the context of CNF search problems $S = S_F$, conjunction-dags are equivalent to Resolution refutations; see also Figure 1. Indeed, $\text{conj-dag}(S_F)$ is just the Resolution refutation size complexity of F , and $w(S_F)$ is the Resolution width complexity of F [8].

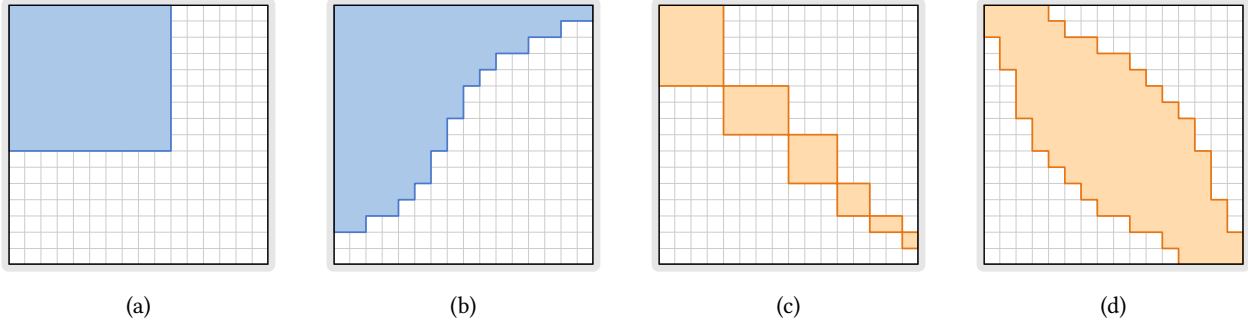


Figure 2: We show lifting theorems for dags whose feasible sets are (a) *rectangles* or (b) *triangles*. It remains open (see Section 9) to prove any lower bounds for explicit mKW/CNF search problems when the feasible sets are (c) *block-diagonal*, which a special case of (d) *intersections of 2 triangles*.

The complexity measures introduced so far are related as follows; here S' is any two-party version of S obtained by choosing some bipartition $\mathcal{X} \times \mathcal{Y} = \{0, 1\}^n$ of the input domain of S :

$$\text{rect-dag}(S') \leq \text{conj-dag}(S) \leq n^{O(w(S))}. \quad (1)$$

The first inequality holds because each conjunction can be simulated by a rectangle. The second inequality holds since there are at most $n^{O(w)}$ many distinct width- w conjunctions, and we may assume wlog that any $f \in \mathcal{F}$ is associated with at most one node in an \mathcal{F} -dag (any incoming edge to a node v can be rewired to the *lowest* node u , in topological order, such that $f_v = f_u$).

3 OUR RESULTS

Our first theorem is a characterization of the rectangle-dag complexity for *composed* search problems of the form $S \circ g^n$. Here $S \subseteq \{0, 1\}^n \times \mathcal{O}$ is an arbitrary n -bit search problem, and $g: \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$ is some carefully chosen two-party *gadget* that helps to distribute each input bit of S between the two parties. More precisely, $S \circ g^n \subseteq \mathcal{X}^n \times \mathcal{Y}^n \times \mathcal{O}$ is the search problem where Alice holds $x \in \mathcal{X}^n$, Bob holds $y \in \mathcal{Y}^n$, and their goal is to find some $o \in S(z)$ for $z := g^n(x, y) = (g(x_1, y_1), \dots, g(x_n, y_n))$.

Our concrete choice for a gadget is the usual m -bit *index* function $\text{IND}_m: [m] \times \{0, 1\}^m \rightarrow \{0, 1\}$ mapping $(x, y) \mapsto y_x$. For large enough m , we show that the bounds (1) are tight.

THEOREM 1. *Let $m = m(n) := n^\delta$ for a large enough constant δ . For any $S \subseteq \{0, 1\}^n \times \mathcal{O}$,*

$$\text{rect-dag}(S \circ \text{IND}_m^n) = n^{\Theta(w(S))}.$$

Implications. The primary advantage of such a lifting theorem is that we obtain, in a generic fashion, a large class of hard (explicit) monotone functions and CNF contradictions. Indeed, let us see an example of how to apply our theorem. We can start with any n -variable k -CNF contradiction F of Resolution width w , and conclude from **Theorem 1** that the composed problem $S' := S_F \circ \text{IND}_m^n$ has rectangle-dag complexity $n^{\Theta(w)}$. Then we can use known reductions to translate S' back to a mKW/CNF search problem. We recall such reductions in **Section 8**, but the upshot will be that:

- S' reduces to $S_{f'}$, where f' is some N -bit monotone function with $N := n^{O(k)}$.

- S' reduces to $S_{F'}$ where F' is some $n^{O(1)}$ -variable $2k$ -CNF contradiction.

A disadvantage, stemming from the large gadget size $m = \text{poly}(n)$, is that we get at best (using $w = \Theta(n)$) a monotone circuit lower bound of $\exp(N^\epsilon)$ for a small constant $\epsilon > 0$. This falls especially short of the current best record of $\exp(N^{1/3-o(1)})$ shown for an explicit monotone function by Harnik and Raz [24]. For this reason (and others), it is an important open problem to develop a lifting theory for gadgets of size $m = O(1)$. In particular, an optimal $2^{\Omega(N)}$ lower bound would follow from an appropriate constant-size-gadget version of **Theorem 1**; see **Section 8** for details.

Techniques. We use tools developed in the context of tree-like lifting theorems, specifically from [18, 21]. These tools allow us to relate large rectangles in the input domain of $S \circ \text{IND}_m^n$ with large subcubes in the input domain of S ; see **Section 4**. Given these tools, the proof of **Theorem 1** is relatively short (two pages). The proof is extremely direct: from any rectangle-dag of size n^d solving $S \circ \text{IND}_m^n$ we extract a width- $O(d)$ conjunction-dag solving S .

Classical works on monotone circuit lower bounds have typically focused on specific monotone functions [1, 3, 22, 42, 48] and more generally on studying the power of the underlying proof methods [2, 9, 43, 45, 49, 51]. A notable exception is Jukna's criterion [29], recently applied in [14, 26], which is a general sufficient condition for a monotone function to require large monotone circuit complexity. Our perspective is seemingly even more abstract, as our result is phrased for arbitrary search problems (not just of mKW/CNF type). However, it remains unclear exactly how the power of our methods compare with the classical techniques; for example, can our result be rephrased in the language of Razborov's method of approximations?

3.1 Extension: Monotone Real Circuits

Triangle-dags. Consider a bipartite input domain $\mathcal{I} := \mathcal{X} \times \mathcal{Y}$ and let \mathcal{F} be the set of all indicator functions of (*combinatorial*) *triangles over $\mathcal{X} \times \mathcal{Y}$* ; here a *triangle* $T \subseteq \mathcal{X} \times \mathcal{Y}$ is a set that can be written as $T = \{(x, y) \in \mathcal{X} \times \mathcal{Y} : a_T(x) < b_T(y)\}$ for some labeling of the rows $a_T: \mathcal{X} \rightarrow \mathbb{R}$ and columns $b_T: \mathcal{Y} \rightarrow \mathbb{R}$ by real numbers; see **Figure 2b**. In particular, every rectangle is a triangle. Call such \mathcal{F} -dags simply *triangle-dags*. For a search problem $S \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{O}$

we define

$\text{tri-dag}(S) :=$ least size of a triangle-dag that solves S .

Hrubeš and Pudlák [25] showed recently that the *monotone real circuit complexity* of an f equals $\text{tri-dag}(S_f)$. Monotone real circuits [23, 36] generalize monotone circuits by allowing the wires to carry arbitrary real numbers and the binary gates to compute arbitrary monotone functions $\mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$. The original motivation to study such circuits, and what interests us here, is that lower bounds for monotone real circuits imply lower bounds for the *Cutting Planes* proof system [12]. In our language, semantic Cutting Planes refutations are equivalent to \mathcal{L} -dags solving CNF search problems, where \mathcal{L} is the family of linear threshold functions (each $f \in \mathcal{L}$ is defined by some $(n+1)$ -tuple $a \in \mathbb{R}^{n+1}$ so that $f(x) = 1$ iff $\sum_{i \in [n]} a_i x_i > a_{n+1}$).

Our second theorem states that [Theorem 1](#) holds more generally with rectangle-dags replaced with triangle-dags. The proof is however more involved than the proof for [Theorem 1](#).

THEOREM 2. *Let $m = m(n) := n^\delta$ for a large enough constant δ . For any $S \subseteq \{0, 1\}^n \times \mathcal{O}$,*

$$\text{tri-dag}(S \circ \text{IND}_m^n) = n^{\Theta(w(S))}.$$

A pithy corollary is that if we start with any CNF contradiction F that is hard for Resolution and compose F with a gadget (as described in [Section 8](#)), the formula becomes hard for Cutting Planes. Previously, only few examples of hard contradictions were known for Cutting Planes, all proved via feasible interpolation [14, 23, 26, 36]. A widely-asked question has been to improve this state-of-the-art by developing alternative lower bound methods; see the surveys [6, §4] and [47, §5]. In particular, Jukna [30, Research Problem 19.17] asked to find a more intuitive “combinatorial” proof method “explicitly showing what properties of [contradictions] force long derivations.” It is unclear how “combinatorial” our method is, but at least it does afford a simple intuition: the hardness is simply borrowed from the realm of Resolution (where we understand very well what makes formulas hard).

4 SUBCUBES FROM RECTANGLES

In this section, as preparation, we recall some technical notions from [18, 21] concerning the index gadget $g := \text{IND}_m$. Namely, writing $G := g^n: [m]^n \times \{0, 1\}^{mn} \rightarrow \{0, 1\}^n$ for n copies of g , we explain how large rectangles in G 's domain are related with large subcubes in G 's codomain.

4.1 Structured Rectangles

For a partial assignment $\rho \in \{0, 1, *\}^n$ we let $\text{free } \rho := \rho^{-1}(*)$ denote its *free* coordinates, and $\text{fix } \rho := [n] \setminus \text{free } \rho$ denote its *fixed* coordinates. The number of fixed coordinates $|\text{fix } \rho|$ is the *width* of ρ . Width- d partial assignments are naturally in 1-to-1 correspondence with width- d conjunctions: for any ρ we define $C_\rho: \{0, 1\}^n \rightarrow \{0, 1\}$ as the width- $|\text{fix } \rho|$ conjunction that accepts an $x \in \{0, 1\}^n$ iff x is consistent with ρ . Thus $C_\rho^{-1}(1) = \{x \in \{0, 1\}^n : x_i = \rho_i \text{ for all } i \in \text{fix } \rho\}$ is a subcube. We say that $R \subseteq [m]^n \times \{0, 1\}^{mn}$ is ρ -like if the image of R under G is precisely the subcube of n -bit

strings consistent with ρ , that is, in short,

$$R \text{ is } \rho\text{-like} \iff G(R) = C_\rho^{-1}(1).$$

For a random variable \mathbf{x} we let $\mathbf{H}_\infty(\mathbf{x}) := \min_x \log(1/\Pr[\mathbf{x} = x])$ denote the usual *min-entropy* of \mathbf{x} . When $\mathbf{x} \in [m]^J$ for some index set J , we write $\mathbf{x}_I \in [m]^I$ for the marginal distribution of \mathbf{x} on a subset $I \subseteq J$ of coordinates. For a set X we use the boldface X to denote a random variable uniformly distributed over X .

Definition 1 ([18, 21]). A rectangle $R := X \times Y \subseteq [m]^n \times \{0, 1\}^{mn}$ is ρ -structured if

- (1) $X_{\text{fix } \rho}$ is fixed, and every $z \in G(R)$ is consistent with ρ , that is, $G(R) \subseteq C_\rho^{-1}(1)$.
- (2) $X_{\text{free } \rho}$ is 0.9-dense: for every nonempty $I \subseteq \text{free } \rho$, X_I has *min-entropy rate* ≥ 0.9 , that is, $\mathbf{H}_\infty(X_I) \geq 0.9 \cdot |I| \log m$.
- (3) Y is large enough: $\mathbf{H}_\infty(Y) \geq mn - n^3$.

Lemma 3 ([17, 21]). *Every ρ -structured rectangle is ρ -like.*

In this work we need a slight strengthening of [Lemma 3](#): for a ρ -structured R , there is a *single* row of R that is already ρ -like. The proof is given in the full version [16].

Lemma 4. *Let $X \times Y$ be ρ -structured. There exists an $x \in X$ such that $\{x\} \times Y$ is ρ -like.*

4.2 Rectangle Partition Scheme

We claim that, given any rectangle $R := X \times Y \subseteq [m]^n \times \{0, 1\}^{mn}$, we can partition most of $X \times Y$ into ρ -structured subrectangles with $|\text{fix } \rho|$ bounded in terms of the size of $X \times Y$. Indeed, we describe a simple 2-round partitioning scheme from [21] below; see also [Figure 3](#). In the 1st round of the algorithm, we partition the rows as $X = \bigsqcup_i X^i$ where each X^i will be fixed on some blocks $I_i \subseteq [n]$ and 0.95-dense on the remaining blocks $[n] \setminus I_i$. In the 2nd round, each $X^i \times Y$ is further partitioned along columns so as to fix the outputs of the gadgets on coordinates I_i .

Rectangle Scheme

Input: $R = X \times Y \subseteq [m]^n \times \{0, 1\}^{mn}$.

Output: A partition of R into subrectangles.

- 1: **1st round:** Iterate the following for $i = 1, 2, \dots$, until X becomes empty:
 - (i) Let $I_i \subseteq [n]$ be a *maximal* subset (possibly $I_i = \emptyset$) such that X_{I_i} has min-entropy rate < 0.95 , and let $\alpha_i \in [m]^{I_i}$ be an outcome witnessing this: $\Pr[X_{I_i} = \alpha_i] > m^{-0.95|I_i|}$
 - (ii) Define $X^i := \{x \in X : x_{I_i} = \alpha_i\}$
 - (iii) Update $X \leftarrow X \setminus X^i$
 - 2: **2nd round:** For each part X^i and $\gamma \in \{0, 1\}^{I_i}$, define $Y^{i,\gamma} := \{y \in Y : g^{I_i}(\alpha_i, y_{I_i}) = \gamma\}$
 - 3: **return** $\{R^{i,\gamma} := X^i \times Y^{i,\gamma} : Y^{i,\gamma} \neq \emptyset\}$
-

All the properties of [Rectangle Scheme](#) that we will subsequently need are formalized below; see also [Figure 3](#). For terminology, given a subset $A' \subseteq A$ we define its *density* (inside A) as $|A'|/|A|$. The proof of the following lemma is postponed to [Section 7](#).




Figure 3: (a) Rectangle Scheme partitions $R = X \times Y$ first along rows, then along columns. **(b) Rectangle Lemma** illustrated: most subrectangles are ρ -structured for low-width ρ , except some error parts (highlighted in figure) that are contained in few error rows/columns $X_{\text{err}}, Y_{\text{err}}$.

Rectangle Lemma. Fix any parameter $k \leq n \log n$. Given a rectangle $R \subseteq [m]^n \times \{0, 1\}^{mn}$, let $R = \sqcup_i R^i$ be the output of *Rectangle Scheme*. Then there exist “error” sets $X_{\text{err}} \subseteq [m]^n$ and $Y_{\text{err}} \subseteq \{0, 1\}^{mn}$, both of density $\leq 2^{-k}$ (inside their respective sets), such that for each i , one of the following holds:

- **Structured case:** R^i is ρ^i -structured for some ρ^i of width at most $O(k/\log n)$.
- **Error case:** R^i is covered by error rows/columns, i.e., $R^i \subseteq X_{\text{err}} \times \{0, 1\}^{mn} \cup [m]^n \times Y_{\text{err}}$.

Finally, a *query alignment* property holds: for every $x \in [m]^n \setminus X_{\text{err}}$, there exists a subset $I_x \subseteq [n]$ with $|I_x| \leq O(k/\log n)$ such that every “structured” R^i intersecting $\{x\} \times \{0, 1\}^{mn}$ has $\text{fix } \rho^i \subseteq I_x$.

5 LIFTING FOR RECTANGLE-DAGS

In this section we prove the nontrivial direction of **Theorem 1**: Let Π be a rectangle-dag solving $S \circ G$ of size n^d for some d . Our goal is to show that $w(S) \leq O(d)$.

5.1 Game Semantics for Dags

For convenience (and fun), we use the language of two-player competitive games, introduced in [4, 37], which provide an alternative way of thinking about conjunction-dags solving $S \subseteq \{0, 1\}^n \times \mathcal{O}$. The game involves two competing players, *Explorer* and *Adversary*, and proceeds in rounds. The state of the game in each round is modeled as a partial assignment $\rho \in \{0, 1, *\}^n$. At the start of the game, $\rho := *$. In each round, Explorer makes one of two moves:

- *Query a bit:* Explorer specifies an $i \in \text{free } \rho$, and Adversary responds with a bit $b \in \{0, 1\}$. The state ρ is updated by $\rho_i \leftarrow b$.
- *Forget a bit:* Explorer specifies an $i \in \text{fix } \rho$, and the state is updated by $\rho_i \leftarrow *$.

An important detail is that Adversary is allowed to choose $b \in \{0, 1\}$ freely even if the i -th bit was queried (with response different from b) and subsequently forgotten during past play. The game ends when a solution to S can be inferred from ρ , that is, when $C_\rho^{-1}(1) \subseteq S^{-1}(o)$ for some $o \in \mathcal{O}$.

Explorer’s goal is to end the game while keeping the width of the game state ρ as small as possible. Indeed, Atserias and Dalmau [4] prove that $w(S)$ is characterized (up to an additive ± 1) as the least w such that the Explorer has a strategy for ending the game that keeps the width of the game state at most w throughout the game. (A similar characterization exists for dag size [37].) Hence our goal becomes to describe a Explorer-strategy for S such that the width of the game state never exceeds $O(d)$ regardless of how the Adversary plays.

5.2 Simplified Proof

To explain the basic idea, we first give a simplified version of the proof: We assume that all rectangles R involved in Π —call them the *original* rectangles—can be partitioned *errorlessly* into ρ -structured subrectangles for ρ of width $O(d)$. That is, invoking *Rectangle Scheme* for each original R , we assume that

- (*) *Assumption:* All subrectangles in the partition $R = \sqcup_i R^i$ output by *Rectangle Scheme* satisfy the “structured” case of *Rectangle Lemma* for $k := 2d \log n$.


In **Section 5.3** we remove this assumption by explaining how the proof can be modified to work with some error rows/columns.

Overview. We extract a width- $O(d)$ Explorer-strategy for S by walking down the rectangle-dag Π , starting at the root. For each original rectangle R that is reached in the walk, we maintain a ρ -structured subrectangle $R' \subseteq R$ chosen from the partition of R . Note that ρ will have width $O(d)$ by our choice of k . The intention is that ρ will record the current state of the game. There are three issues to address: (1) Why is the starting condition of the game met? (2) How do we take a step from a node of Π to one of its children? (3) Why are we done once we reach a leaf?

(1) *Root case.* At start, the root of Π is associated with the original rectangle $R = [m]^n \times \{0, 1\}^{mn}$ comprising the whole domain. The partition of R computed by *Rectangle Scheme* is trivial: it contains a single part, the $*$ -structured R itself. Hence we simply maintain the $*$ -structured $R \subseteq R$, which meets the starting condition for the game.

(2) *Internal step.* This is the crux of the argument: Supposing the game has reached state $\rho_{R'}$ and we are maintaining some $\rho_{R'}$ -structured subrectangle $R' \subseteq R$ associated with an internal node v , we want to move to some $\rho_{L'}$ -structured subrectangle $L' \subseteq L$ associated with a child of v . Moreover, we must keep the width of the game state at most $O(d)$ during this move.

Since $R' =: X' \times Y'$ is $\rho_{R'}$ -structured, we have from [Lemma 4](#) that there exists some $x^* \in X'$ such that $\{x^*\} \times Y'$ is $\rho_{R'}$ -like. Let the two original rectangles associated with the children of v be L_0 and L_1 . Let $\sqcup_i L_b^i$ be the partition of L_b output by [Rectangle Scheme](#). By query alignment in [Rectangle Lemma](#), there is some $I_b^* \subseteq [n]$, $|I_b^*| \leq O(d)$, such that all L_b^i that intersect the x^* -th row are ρ^i -structured with $\text{fix } \rho^i \subseteq I_b^*$. As Explorer, we now query the input bits in coordinates $J := (I_0^* \cup I_1^*) \setminus \text{fix } \rho_{R'}$ (in any order) obtaining some response string $z_J \in \{0, 1\}^J$ from the Adversary. As a result, the state of the game becomes the extension of $\rho_{R'}$ by z_J , call it ρ^* , which has width $|\text{fix } \rho^*| = |\text{fix } \rho_{R'} \cup J| \leq O(d)$.



Note that there is some $y^* \in Y'$ (and hence $(x^*, y^*) \in R' \subseteq L_0 \cup L_1$) such that $G(x^*, y^*)$ is consistent with ρ^* ; indeed, the whole row $\{x^*\} \times Y'$ is $\rho_{R'}$ -like and ρ^* extends $\rho_{R'}$. Suppose $(x^*, y^*) \in L_0$; the case of L_1 is analogous. In the partition of L_0 , let L' be the unique part such that $(x^*, y^*) \in L'$. Note that L' is $\rho_{L'}$ -like for some $\rho_{L'}$ that is consistent with $G(x^*, y^*)$ and $\text{fix } \rho_{L'} \subseteq I_0^*$ (by query alignment). Hence ρ^* extends $\rho_{L'}$. As Explorer, we now forget all queried bits in ρ^* except those queried in $\rho_{L'}$.

We have recovered our invariant: the game state is $\rho_{L'}$ and we maintain a $\rho_{L'}$ -structured subrectangle L' of an original rectangle L_0 . Moreover, the width of the game state remained $O(d)$.

(3) *Leaf case.* Suppose the game state is ρ and we are maintaining an associated ρ -structured subrectangle $R' \subseteq R$ corresponding to a leaf node. The leaf node is labeled with some solution $o \in \mathcal{O}$ satisfying $R' \subseteq (S \circ G)^{-1}(o)$, that is, $G(R') \subseteq S^{-1}(o)$. But $G(R') = C_\rho^{-1}(1)$ by [Lemma 3](#) so that $C_\rho^{-1}(1) \subseteq S^{-1}(o)$. Therefore the game ends. This concludes the (simplified) proof.

5.3 Accounting for Error

Next, we explain how to get rid of the assumption (*) by accounting for the rows and columns that are classified as error in [Rectangle Lemma](#) for $k := 2d \log n$. The partitioning of Π 's rectangles is done more carefully: We sort all original rectangles in *reverse topological order* R_1, R_2, \dots, R_{n^d} from leaves to root, that is, if R_i is a descendant of R_j then R_i comes before R_j in the order. Then we process the rectangles in this order:

Initialize cumulative error sets $X_{\text{err}}^* = Y_{\text{err}}^* := \emptyset$. Iterate for $i = 1, 2, \dots, n^d$ rounds:

- (1) Remove from R_i the rows/columns $X_{\text{err}}^*, Y_{\text{err}}^*$. That is, update

$$R_i \leftarrow R_i \setminus (X_{\text{err}}^* \times \{0, 1\}^{mn} \cup [m]^n \times Y_{\text{err}}^*).$$

- (2) Apply the [Rectangle Scheme](#) for R_i . Output all resulting subrectangles that satisfy the “structured” case of [Rectangle Lemma](#) for $k := 2d \log n$. (All non-structured subrectangles are omitted). Call the resulting error rows/columns X_{err} and Y_{err} .

- (3) Update $X_{\text{err}}^* \leftarrow X_{\text{err}}^* \cup X_{\text{err}}$ and $Y_{\text{err}}^* \leftarrow Y_{\text{err}}^* \cup Y_{\text{err}}$.

In words, an original rectangle R_i is processed only after all of its descendants are partitioned. Each descendant may contribute some error rows/columns, accumulated into sets $X_{\text{err}}^*, Y_{\text{err}}^*$, which are deleted from R_i before it is partitioned. The partitioning of R_i will in turn contribute its error rows/columns to its ancestors.

We may now repeat the proof of [Section 5.2](#) verbatim using only the structured subrectangles output by the above process. That is, we still maintain the same invariant: when the game state is ρ , we maintain a ρ -structured R' (output by the above process) of an original R . We highlight only the key points below.

(1) *Root case.* The cumulative error at the end of the process is tiny: $X_{\text{err}}^*, Y_{\text{err}}^*$ have density at most $n^d \cdot n^{-2d} \leq 1/4$ by a union bound over all rounds. In particular, the root rectangle R_{n^d} (with errors removed) still has density $\geq 1/2$ inside $[m]^n \times \{0, 1\}^{mn}$, and so the partition output by [Rectangle Scheme](#) is trivial, containing only the ρ^* -structured R_{n^d} itself. This meets the starting condition for the game.

(2) *Internal step.* By construction, the cumulative error sets *shrink* when we take a step from a node to one of its children. This means that our error handling does not interfere with the internal step: each structured subrectangle R' of an original rectangle R is wholly covered by the structured subrectangles of R 's children.

- (3) *Leaf case.* This case is unchanged.

6 LIFTING FOR TRIANGLE-DAGS

In this section we prove the nontrivial direction of [Theorem 2](#): Let Π be a triangle-dag solving $S \circ G$ of size n^d for some d . Our goal is to show that $w(S) \leq O(d)$.

The proof is conceptually the same as for rectangle-dags. The only difference is that we need to replace [Rectangle Scheme](#) (and the associated [Rectangle Lemma](#)) with an algorithm that partitions a given triangle $T \subseteq [m]^n \times \{0, 1\}^{mn}$ into subtriangles that behave like conjunctions.

6.1 Triangle Partition Scheme

We introduce a triangle partitioning algorithm, [Triangle Scheme](#). Its definition is given in the full version [16]. For now, we only need its high-level description: On input a triangle T , [Triangle Scheme](#) outputs a disjoint cover $\sqcup_i R^i \supseteq T$ where R^i are rectangles. This induces a partition of T into subtriangles $T \cap R^i$. Each (non-error) rectangle R^i is ρ^i -structured (for low-width ρ^i) and is associated with a ρ^i -structured “inner” subrectangle $L^i \subseteq R^i$ satisfying $L^i \subseteq$




Figure 4: Structured case of Triangle Lemma: The subtriangle $T \cap R^i$ is sandwiched between two ρ^i -structured rectangles L^i and R^i .

$T \cap R^i \subseteq R^i$; see Figure 4. Hence $T \cap R^i$ is ρ^i -like, as it is sandwiched between two ρ^i -like rectangles.

More formally, all the properties of Triangle Scheme that we will subsequently need are formalized below (note the similarity with Rectangle Lemma); see the full version [16] for the proof.

Triangle Lemma. Fix any parameter $k \leq n \log n$. Given a triangle $T \subseteq [m]^n \times \{0, 1\}^{mn}$, let $\sqcup_i R^i$ be the output of Triangle Scheme. Then there exist “error” sets $X_{\text{err}} \subseteq [m]^n$ and $Y_{\text{err}} \subseteq \{0, 1\}^{mn}$, both of density $\leq 2^{-k}$ (inside their respective sets), such that for each i , one of the following holds:

- **Structured case:** R^i is ρ^i -structured for some ρ^i of width at most $O(k/\log n)$. Moreover, there exists an “inner” rectangle $L^i \subseteq T \cap R^i$ such that L^i is also ρ^i -structured.
- **Error case:** R^i is covered by error rows/columns, i.e., $R^i \subseteq X_{\text{err}} \times \{0, 1\}^{mn} \cup [m]^n \times Y_{\text{err}}$.

Finally, a **query alignment** property holds: for every $x \in [m]^n \setminus X_{\text{err}}$, there exists a subset $I_x \subseteq [n]$ with $|I_x| \leq O(k/\log n)$ such that every “structured” R^i intersecting $\{x\} \times \{0, 1\}^{mn}$ has $\text{fix } \rho^i \subseteq I_x$.

6.2 Simplified Proof

As in the rectangle case, we give a simplified proof assuming no errors. That is, invoking Triangle Scheme for each triangle T involved in Π , we assume that

- (†) *Assumption:* All rectangles in the cover $\sqcup_i R^i \supseteq T$ output by Triangle Scheme satisfy the “structured” case of Triangle Lemma for $k := 2d \log n$.

The argument for getting rid of the assumption (†) is the same as in the rectangle case, and hence we omit that step—one only needs to observe that removing cumulative error rows/columns from a triangle still leaves us with a triangle.

Overview. As before, we extract a width- $O(d)$ Explorer-strategy for S by walking down the triangle-dag Π , starting at the root. For each triangle T of Π that is reached in the walk, we maintain a ρ -structured inner rectangle $L \subseteq T$. Here ρ (of width $O(d)$) by the choice of k will record the current state of the game. There are the three steps (1)–(3) to address, of which (1) and (3) remain exactly the same as in the rectangle case. So we only explain step (2), which

requires us to replace the use of Rectangle Lemma with the new Triangle Lemma.

(2) *Internal step.* Supposing the game has reached state ρ_L and we are maintaining some ρ_L -structured inner rectangle $L \subseteq T$ associated with an internal node v , we want to move to some $\rho_{\tilde{L}}$ -structured inner rectangle $\tilde{L} \subseteq T$ associated with a child of v . Moreover, we must keep the width of the game state at most $O(d)$ during this move.

Since $L =: X' \times Y'$ is ρ_L -structured, we have from Lemma 4 that there exists some $x^* \in X'$ such that $\{x^*\} \times Y'$ is ρ_L -like. Let the two triangles associated with the children of v be T_0 and T_1 , so that $L \subseteq T_0 \cup T_1$.

Let $\sqcup_b R_b^i$ be the rectangle cover of T_b output by Triangle Scheme. By query alignment in Triangle Lemma, there is some $I_b^* \subseteq [n]$, $|I_b^*| \leq O(d)$, such that all R_b^i that intersect the x^* -th row are ρ^i -structured with $\text{fix } \rho^i \subseteq I_b^*$. As Explorer, we now query the input bits in coordinates $J := (I_0^* \cup I_1^*) \setminus \text{fix } \rho_L$ (in any order) obtaining some response string $z_J \in \{0, 1\}^J$ from the Adversary. As a result, the state of the game becomes the extension of ρ_L by z_J , call it ρ^* , which has width $|\text{fix } \rho^*| = |\text{fix } \rho_L \cup J| \leq O(d)$.

Note that there is some $y^* \in Y'$ (and hence $(x^*, y^*) \in L \subseteq T_0 \cup T_1$) such that $G(x^*, y^*)$ is consistent with ρ^* ; indeed, the whole row $\{x^*\} \times Y'$ is ρ_L -like and ρ^* extends ρ_L . Suppose $(x^*, y^*) \in T_0$; the case of T_1 is analogous. In the rectangle covering of T_0 , let R be the unique part such that $(x^*, y^*) \in R$. Note that R is ρ_R -like for some ρ_R that is consistent with $G(x^*, y^*)$ and $\text{fix } \rho_R \subseteq I_0^*$ (by query alignment). Hence ρ^* extends ρ_R . As Explorer, we now forget all queried bits in ρ^* except those queried in ρ_R . Also we move to the inner rectangle $\tilde{L} \subseteq R$ promised by Triangle Lemma that satisfies $\tilde{L} \subseteq T_0$ and is $\rho_{\tilde{L}} = \rho_R$ structured.

We have recovered our invariant: the game state is $\rho_{\tilde{L}}$ and we maintain a $\rho_{\tilde{L}}$ -structured subrectangle \tilde{L} of a triangle T_0 . Moreover, the width of the game state remained $O(d)$.

7 PARTITIONING RECTANGLES

In this section, we prove Rectangle Lemma. We use repeatedly the following simple fact about min-entropy.

Fact 5. Let X be a random variable and E an event. Then $\mathbf{H}_{\infty}(X \mid E) \geq \mathbf{H}_{\infty}(X) - \log 1/\Pr[E]$.

The proof is more-or-less implicit in [18, 21]. We start by recording a key property of the 1st round of Rectangle Scheme.

Claim 6. Each part X^i obtained in 1st round of Rectangle Scheme satisfies:

- Blockwise-density: $X_{[n] \setminus I_i}^i$ is 0.95-dense.
- Relative size: $|X^{\geq i}| \leq m^{n-0.05|I_i|}$ where $X^{\geq i} := \bigcup_{j \geq i} X^j$.

PROOF. By definition, $X^i = (X^{\geq i} \mid X_{I_i}^{\geq i} = \alpha_i)$. Suppose for contradiction that $X_{[n] \setminus I_i}^i$ is not 0.95-dense. Then there is some nonempty subset $K \subseteq [n] \setminus I_i$ and an outcome $\beta \in [m]^K$ violating the min-entropy condition, namely $\Pr[X_K^i = \beta] > m^{-0.95|K|}$. But this contradicts the maximality of I_i since the larger set $I_i \cup K$ now

violates the min-entropy condition for $X^{\geq i}$:

$$\begin{aligned} \Pr[X_{I_i \cup K}^{\geq i} = \alpha_i \beta] &= \Pr[X_{I_i}^{\geq i} = \alpha_i] \cdot \Pr[X_K^i = \beta] \\ &> m^{-0.95|I_i|} \cdot m^{-0.95|K|} = m^{-0.95(|I_i \cup K|)}. \end{aligned}$$

This shows the first property. For the second property, apply **Fact 5** for $X^i = (X^{\geq i} | X_{I_i}^{\geq i} = \alpha_i)$ to find that $H_\infty(X^i) \geq H_\infty(X^{\geq i}) - 0.95|I_i| \log m$. On the other hand, since X^i is fixed on I_i , we have $H_\infty(X^i) \leq (n - |I_i|) \log m$. Combining these two inequalities we get $H_\infty(X^{\geq i}) \leq (n - 0.05|I_i|) \log m$, which yields the second property. \square

Proof of Rectangle Lemma. Identifying $Y_{\text{err}}, X_{\text{err}}$. We define $Y_{\text{err}} := \bigcup_{i, \gamma} Y^{i, \gamma}$ subject to $|Y^{i, \gamma}| < 2^{mn-n^2}$. To bound the size of Y_{err} , we claim that there are at most $(4m)^n$ possible choices of i, γ . Indeed, each X^i is associated with a unique pair $(I_i \subseteq [n], \alpha_i \in [m]^{|I_i|})$, and there are at most 2^n choices of I_i and at most m^n choices of corresponding α_i . Also, for each X^i , there are at most 2^n possible assignments to $\gamma \in \{0, 1\}^{I_i}$. For each i, γ , we add at most 2^{mn-n^2} columns to Y_{err} . Thus, Y_{err} has density at most $(4m)^n \cdot 2^{-n^2} < 2^{-k}$ inside $\{0, 1\}^{mn}$.

We define $X_{\text{err}} := \bigsqcup_i X^i$ subject to $|I_i| > 20k/\log m$. Let i be the least index with $|I_i| > 20k/\log m$ so that $X_{\text{err}} \subseteq X^{\geq i}$. By **Claim 6**, $|X^{\geq i}| \leq m^{n-0.05|I_i|} < m^n \cdot 2^{-k}$ since $|I_i| > 20k/\log m$. In other words, $X^{\geq i}$, and hence X_{err} , has density at most 2^{-k} inside $[m]^n$.

Structured vs. error. Let $R^{i, \gamma} := X^i \times Y^{i, \gamma}$, where X^i is associated with (I_i, α_i) , be a rectangle *not* contained in the error rows/columns. By definition of $X_{\text{err}}, Y_{\text{err}}$, this means $|Y^{i, \gamma}| \geq 2^{mn-n^2}$ (so that $H_\infty(Y^{i, \gamma}) \geq mn - n^2$) and $|I_i| \leq 20k/\log m$. We have from **Claim 6** that $X_{[n] \setminus I_i}^i$ is 0.95-dense. Hence, $R^{i, \gamma}$ is ρ^i -structured where ρ^i equals γ on I_i and consists of stars otherwise.

Query alignment. For each $x \in [m]^n \setminus X_{\text{err}}$, we define $I_x = I_i$ where X^i is the unique part that contains x . It follows that any ρ -structured rectangle that intersects the x -th row is of the form $X^i \times Y^{i, \gamma}$ and hence has $\text{fix } \rho = I_i$. Since $X^i \not\subseteq X_{\text{err}}$, we have $|I_i| \leq O(k/\log n)$. \square

8 TRANSLATING BETWEEN mKW/CNF

In this section, for exposition, we recall some known reductions between mKW and CNF search problems. These reductions can be combined with our main theorems to yield applications in proof and monotone circuit complexity (as outlined in **Section 3**).

Certificates. The key property of an n -bit search problem $S \subseteq \{0, 1\}^n \times \mathcal{O}$ that facilitates an efficient reduction to a mKW/CNF search problem is having a low *certificate* (aka nondeterministic) complexity. A *certificate for* $(x, o) \in S$ is a partial assignment $\rho \in \{0, 1, *\}^n$ such that x is consistent with ρ and o is a valid output for every input consistent with ρ ; in short, $x \in C_\rho^{-1}(1) \subseteq S^{-1}(o)$. A *certificate for* x is a certificate for $(x, o) \in S$ for some $o \in S(x)$. The *certificate complexity of* x is the least width of a certificate for x . The *certificate complexity of* S is the maximum over all $x \in \{0, 1\}^n$ of the certificate complexity of x .

For any search problem S one can associate a ‘‘certification’’ search problem S_{cert} : on input x to S , output a certificate for x in S . Algorithmically speaking, such an S_{cert} is clearly at least as

hard as S : if we solve S_{cert} by finding a certificate for $(x, o) \in S$, we can solve S by outputting o .

CNF search \Leftrightarrow low certificate complexity. For any k -CNF contradiction F , the associated CNF search problem S_F has certificate complexity at most k . Conversely [35], for any total search problem $S \subseteq \{0, 1\}^n \times \mathcal{O}$, we can construct a k -CNF contradiction F , where k is the certificate complexity of S , such that S_F is a type of certification problem for S (and hence at least as hard as S). Namely, we can pick a collection C of width- k certificates, one for each $x \in \{0, 1\}^n$. The k -CNF formula F is then defined as $\bigwedge_{\rho \in C} \neg C_\rho$.

Gadget composition. For the purposes of query complexity, there are two ways to represent the first argument $x \in [m]$ to the index function $\text{IND}_m: [m] \times \{0, 1\}^m$ as a binary string. The simplest is to write x as a $\log m$ -bit string. Under this convention, IND_m has certificate complexity $\log m + 1$. If $S \subseteq \{0, 1\}^n \times \mathcal{O}$ has certificate complexity k , the composed problem $S \circ \text{IND}_m^n$ has certificate complexity $k(\log m + 1)$ (by composing certificates). For applications, this means that if we start with a k -CNF contradiction F , we may reduce $S_F \circ \text{IND}_m^n$ to solving $S_{F'}$ where F' is a $k(\log m + 1)$ -CNF contradiction over $O(mn)$ variables.

A better representation [5, 13], which does not blow up the certificate complexity (or CNF width), is to write x as an m -bit string of Hamming weight 1 (the index of the unique 1-entry encodes $x \in [m]$). Under this convention, $\text{IND}_m^n: \{0, 1\}^m \times \{0, 1\}^m \rightarrow \{0, 1\}$ becomes a *partial* function of certificate complexity 2. Hence, if S has certificate complexity k , the *partial* composed problem $S' := S \circ \text{IND}_m^n$ has certificate complexity $2k$.

Moreover, the partial problem S' can be extended into a *total* problem S_{tot} without making it any easier to solve for rectangle-dags. Indeed, we introduce new variables/certificates allowing us to say that an input (x, y) to S' is trivially solved with output $\perp \notin \mathcal{O}$, if for some $i \in [n]$, $x_i \in \{0, 1\}^m$ is *not* of Hamming weight 1. Specifically, Alice will receive new input bits $x' \in (\{0, 1\}^m)^n$ (in addition to the original $x \in (\{0, 1\}^m)^n$) and we say that an Alice input xx' is *good* if for each $i \in [n]$, the string $x'_i \in \{0, 1\}^m$ describes a non-decreasing sequence

$$0 = x'_{i,1} \leq x'_{i,2} \leq \dots \leq x'_{i,m} \leq x'_{i,m+1} := 1$$

(the last value being hardcoded by convention), and moreover $x_{i,j} = 1$ iff $x'_{i,j} < x'_{i,j+1}$. Note that if xx' is *not* good, there is a width-3 certificate witnessing this. Our total search problem $S_{\text{tot}} \subseteq \{0, 1\}^{2mn} \times \{0, 1\}^m \times (\mathcal{O} \cup \{\perp\})$ is defined by all these width-3 certificates (for output \perp) together with all the original certificates of S' . To see that S_{tot} is at least as hard as S' for rectangle-dags, we note that for any input (x, y) to S' , Alice can compute a unique x' so that xx' is *good*. Now any output $o \in S_{\text{tot}}(xx', y)$ is also such that $o \in S'(x, y)$.

In summary, we can reduce (in the context of rectangle-dags) $S_F \circ \text{IND}_m^n$ to solving $S_{F'}$ where F' is a $2k$ -CNF contradiction over $O(mn)$ variables.

mKW problems. A rectangle $R \subseteq \mathcal{X} \times \mathcal{Y}$ is *monochromatic* for a search problem $S \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{O}$ if $R \subseteq S^{-1}(o)$ for some $o \in \mathcal{O}$. The nondeterministic communication complexity of S is the logarithm of the least number of monochromatic rectangles that cover the whole input domain $\mathcal{X} \times \mathcal{Y}$. If S has nondeterministic

communication complexity $\log N$, then by a standard reduction (e.g., [15, Lemma 2.3]) S reduces to S_f for some monotone $f: \{0, 1\}^N \rightarrow \{0, 1\}$.

Consider a composed search problem $S_f \circ g^n$ obtained from a k -CNF contradiction with ℓ clauses. Its nondeterministic communication complexity is at most $\log \ell + k \cdot (\log m + 1)$; intuitively, it takes $\log \ell$ bits to specify an unsatisfied clause C , and $\log m + 1$ bits to verify the output of a single gadget, and there are k gadgets relevant to C . Suppose for a moment that a version of [Theorem 1](#), proving a $2^{\Omega(w)}$ lower bound, held for a gadget of constant size $m = O(1)$. Then we could lift any of the known CNF contradictions with parameters $k = O(1)$, $\ell = O(n)$, $w = \Omega(n)$, to obtain an explicit monotone function on $N = \Theta(n)$ variables, with essentially maximal monotone circuit complexity $2^{\Omega(N)}$. This gives some motivation to further develop lifting tools for small gadgets.

9 OPEN PROBLEMS

If the long line of work on *tree-like* lifting theory is of any indication, there should be much to explore also in the *dag-like* setting. We propose a few concrete directions.

Can our methods be extended to prove lower bounds for dags whose feasible sets are *intersections of k triangles* for $k \geq 2$? See [Figure 2](#). This would imply lower bounds for proofs systems such as width- k Resolution over Cutting Planes [33] and Resolution over linear equations [28, 41].

Question 1. *Prove a lifting theorem for \mathcal{F} -dags where $\mathcal{F} := \{\text{intersections of } k \text{ triangles}\}$.*

One of the most important open problems (e.g., [47, §5]) regarding semi-algebraic proof systems that manipulate low-degree polynomials—where \mathcal{F} is, say, degree- d polynomial threshold functions—is to prove lower bounds on their *dag-like* refutation length (*tree-like* lower bounds are known [7, 19]). Since degree- d polynomials can be efficiently evaluated by $(d + 1)$ -party number-on-forehead (NOF) protocols, one might hope to prove a dag-like NOF lifting theorem. However, we currently lack a good understanding of NOF lifting even in the tree-like case. We believe the first necessary step should be to settle the following (a two-party analogue of which was proved in [18]).

Question 2. *Prove a nondeterministic lifting theorem for NOF protocols.*

The proof of [Theorem 1](#), which extracts a width- $O(d)$ conjunction-dag from a size- n^d rectangle-dag, has the additional property of preserving the dag *depth* (up to an $O(d)$ factor). This raises the question of whether one could investigate size–depth tradeoffs for monotone circuits via lifting.

Question 3. *Does there exist, for any $d \geq 1$, an $f: \{0, 1\}^n \rightarrow \{0, 1\}$ computable with monotone circuits of size n^d such that any subexponential-size monotone circuit computing f has depth $n^{\Omega(d)}$?*

Razborov [46] has recently obtained related results for Resolution, but the parameters in his construction seem not to be good enough for a direct application of [Theorem 1](#).

ACKNOWLEDGMENTS

We thank Jakob Nordström for extensive feedback on an early draft of this work. We also thank Toniann Pitassi, Thomas Watson, and anonymous STOC reviewers for comments.

PK was funded in part by NSF grants CCF-1650733 and IIS-1741137.

REFERENCES

- [1] Noga Alon and Ravi Boppana. 1987. The monotone circuit complexity of Boolean functions. *Combinatorica* 7, 1 (1987), 1–22. <https://doi.org/10.1007/BF02579196>
- [2] Kazuyuki Amano and Akira Maruoka. 2004. The Potential of the Approximation Method. *SIAM J. Comput.* 33, 2 (2004), 433–447. <https://doi.org/10.1137/S009753970138445X>
- [3] Alexander Andreev. 1985. On a method for obtaining lower bounds for the complexity of individual monotone functions. *Doklady Akademii Nauk USSR* 281, 2 (1985), 1033–1037.
- [4] Albert Atserias and Víctor Dalmau. 2008. A combinatorial characterization of resolution width. *J. Comput. System Sci.* 74, 3 (2008), 323–334. <https://doi.org/10.1016/j.jcss.2007.06.025>
- [5] Paul Beame, Trinh Huynh, and Toniann Pitassi. 2010. Hardness Amplification in Proof Complexity. In *Proceedings of the 42nd Symposium on Theory of Computing (STOC)*. ACM, 87–96. <https://doi.org/10.1145/1806689.1806703>
- [6] Paul Beame and Toniann Pitassi. 2001. Propositional Proof Complexity: Past, Present, and Future. In *Current Trends in Theoretical Computer Science: Entering the 21st Century*. World Scientific, 42–70. https://doi.org/10.1142/9789812810403_0001
- [7] Paul Beame, Toniann Pitassi, and Nathan Segerlind. 2007. Lower Bounds for Lovász–Schrijver Systems and Beyond Follow from Multiparty Communication Complexity. *SIAM J. Comput.* 37, 3 (2007), 845–869. <https://doi.org/10.1137/060654645>
- [8] Eli Ben-Sasson and Avi Wigderson. 2001. Short Proofs Are Narrow—Resolution Made Simple. *J. ACM* 48, 2 (2001), 149–169. <https://doi.org/10.1145/375827.375835>
- [9] Christer Berg and Staffan Ullberg. 1999. Symmetric Approximation Arguments for Monotone Lower Bounds Without Sunflowers. *Computational Complexity* 8, 1 (1999), 1–20. <https://doi.org/10.1007/s000370050017>
- [10] Maria Bonet, Toniann Pitassi, and Ran Raz. 1997. Lower Bounds for Cutting Planes Proofs with Small Coefficients. *The Journal of Symbolic Logic* 62, 3 (1997), 708–728. <https://doi.org/10.2307/2275569>
- [11] Arkadev Chattopadhyay, Michal Koucký, Bruno Loff, and Sagnik Mukhopadhyay. 2017. *Simulation Theorems via Pseudorandom Properties*. Technical Report. arXiv:1704.06807
- [12] William Cook, Collette Coullard, and György Turán. 1987. On the complexity of cutting-plane proofs. *Discrete Applied Mathematics* 18, 1 (1987), 25–38. [https://doi.org/10.1016/0166-218X\(87\)90039-4](https://doi.org/10.1016/0166-218X(87)90039-4)
- [13] Susanna de Rezende, Jakob Nordström, and Marc Vinyals. 2016. How Limited Interaction Hinders Real Communication (and What It Means for Proof and Circuit Complexity). In *Proceedings of the 57th Symposium on Foundations of Computer Science (FOCS)*. IEEE, 295–304. <https://doi.org/10.1109/FOCS.2016.40>
- [14] Noah Fleming, Denis Pankratov, Toniann Pitassi, and Robert Robere. 2017. Random CNFs are Hard for Cutting Planes, In *Proceedings of the 58th Symposium on Foundations of Computer Science (FOCS)*. *Proceedings of the 58th Symposium on Foundations of Computer Science (FOCS)*. <https://doi.org/10.2307/2275569>
- [15] Anna Gál. 2001. A Characterization of Span Program Size and Improved Lower Bounds for Monotone Span Programs. *Computational Complexity* 10, 4 (2001), 277–296. <https://doi.org/10.1007/s000370100001>
- [16] Ankit Garg, Mika Göös, Pritish Kamath, and Dmitri Sokolov. 2017. *Monotone Circuit Lower Bounds from Resolution*. Technical Report TR17-175. Electronic Colloquium on Computational Complexity (ECCC). <https://eccc.weizmann.ac.il/report/2017/175/>
- [17] Mika Göös, Pritish Kamath, Toniann Pitassi, and Thomas Watson. 2017. Query-to-Communication Lifting for P^{NP}. In *Proceedings of the 32nd Computational Complexity Conference (CCC)*. Schloss Dagstuhl, 12:1–12:16. <https://doi.org/10.4230/LIPIcs.CCC.2017.12>
- [18] Mika Göös, Shachar Lovett, Raghu Meka, Thomas Watson, and David Zuckerman. 2016. Rectangles Are Nonnegative Juntas. *SIAM J. Comput.* 45, 5 (2016), 1835–1869. <https://doi.org/10.1137/15M103145X>
- [19] Mika Göös and Toniann Pitassi. 2014. Communication Lower Bounds via Critical Block Sensitivity. In *Proceedings of the 46th Symposium on Theory of Computing (STOC)*. ACM, 847–856. <https://doi.org/10.1145/2591796.2591838>
- [20] Mika Göös, Toniann Pitassi, and Thomas Watson. 2015. Deterministic Communication vs. Partition Number. In *Proceedings of the 56th Symposium on Foundations*

- of Computer Science (FOCS). IEEE, 1077–1088. <https://doi.org/10.1109/FOCS.2015.70>
- [21] Mika Göös, Toniann Pitassi, and Thomas Watson. 2017. Query-to-Communication Lifting for BPP. In *Proceedings of the 58th Symposium on Foundations of Computer Science (FOCS)*. IEEE. To appear.
- [22] Armin Haken. 1995. Counting Bottlenecks to Show Monotone $P \neq NP$. In *Proceedings of the 36th Symposium on Foundations of Computer Science (FOCS)*. 36–40. <https://doi.org/10.1109/SFCS.1995.492460>
- [23] Armin Haken and Stephen Cook. 1999. An Exponential Lower Bound for the Size of Monotone Real Circuits. *J. Comput. System Sci.* 58, 2 (1999), 326–335. <https://doi.org/10.1006/jcss.1998.1617>
- [24] Danny Harnik and Ran Raz. 2000. Higher Lower Bounds on Monotone Size. In *Proceedings of the 32nd Symposium on Theory of Computing (STOC)*. ACM, 378–387. <https://doi.org/10.1145/335305.335349>
- [25] Pavel Hruševš and Pavel Pudlák. 2017. *A note on monotone real circuits*. Technical Report TR17-048. Electronic Colloquium on Computational Complexity (ECCC). <https://eccc.weizmann.ac.il/report/2017/048/>
- [26] Pavel Hruševš and Pavel Pudlák. 2017. Random formulas, monotone circuits, and interpolation. In *Proceedings of the 58th Symposium on Foundations of Computer Science (FOCS)*. To appear.
- [27] Trinh Huynh and Jakob Nordström. 2012. On the Virtue of Succinct Proofs: Amplifying Communication Complexity Hardness to Time-Space Trade-Offs in Proof Complexity. In *Proceedings of the 44th Symposium on Theory of Computing (STOC)*. ACM, 233–248. <https://doi.org/10.1145/2213977.2214000>
- [28] Dmitry Itsykson and Dmitry Sokolov. 2014. Lower Bounds for Splittings by Linear Combinations. In *Proceedings of the 39th Mathematical Foundations of Computer Science (MFCS)*. Springer, 372–383. https://doi.org/10.1007/978-3-662-44465-8_32
- [29] Stasys Jukna. 1997. Finite Limits and Monotone Computations: The Lower Bounds Criterion. In *Proceedings of the 12th Computational Complexity Conference (CCC)*. 302–313. <https://doi.org/10.1109/CCC.1997.612325>
- [30] Stasys Jukna. 2012. *Boolean Function Complexity: Advances and Frontiers*. Algorithms and Combinatorics, Vol. 27. Springer.
- [31] Mauricio Karchmer and Avi Wigderson. 1988. Monotone circuits for connectivity require super-logarithmic depth. In *Proceedings of the 20th Symposium on Theory of Computing (STOC)*. ACM, 539–550. <https://doi.org/10.1145/62212.62265>
- [32] Jan Krajíček. 1997. Interpolation Theorems, Lower Bounds for Proof Systems, and Independence Results for Bounded Arithmetic. *Journal of Symbolic Logic* 62, 2 (1997), 457–486. <https://doi.org/10.2307/2275541>
- [33] Jan Krajíček. 1998. Discretely ordered modules as a first-order extension of the cutting planes proof system. *Journal of Symbolic Logic* 63, 4 (1998), 1582–1596. <https://doi.org/10.2307/2586668>
- [34] Eyal Kushilevitz and Noam Nisan. 1997. *Communication Complexity*. Cambridge University Press.
- [35] László Lovász, Moni Naor, Ilan Newman, and Avi Wigderson. 1995. Search Problems in the Decision Tree Model. *SIAM Journal on Discrete Mathematics* 8, 1 (1995), 119–132. <https://doi.org/10.1137/S0895480192233867>
- [36] Pavel Pudlák. 1997. Lower Bounds for Resolution and Cutting Plane Proofs and Monotone Computations. *The Journal of Symbolic Logic* 62, 3 (1997), 981–998. <https://doi.org/10.2307/2275583>
- [37] Pavel Pudlák. 2000. Proofs as Games. *The American Mathematical Monthly* 107, 6 (2000), 541–550. <https://doi.org/10.2307/2589349>
- [38] Pavel Pudlák. 2010. On extracting computations from propositional proofs (a survey). In *Proceedings of the 30th Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, Vol. 8. Schloss Dagstuhl, 30–41. <https://doi.org/10.4230/LIPLcs.FSTTCS.2010.30>
- [39] Anup Rao and Amir Yehudayoff. 2017. *Communication Complexity*. In preparation.
- [40] Ran Raz and Pierre McKenzie. 1999. Separation of the Monotone NC Hierarchy. *Combinatorica* 19, 3 (1999), 403–435. <https://doi.org/10.1007/s004930050062>
- [41] Ran Raz and Iddo Zameret. 2008. Resolution over linear equations and multilinear proofs. *Annals of Pure and Applied Logic* 155, 3 (2008), 194–224. <https://doi.org/10.1016/j.apal.2008.04.001>
- [42] Alexander Razborov. 1985. Lower bounds on the monotone complexity of some Boolean functions. *Doklady Akademii Nauk USSR* 285 (1985), 798–801.
- [43] Alexander Razborov. 1989. On the Method of Approximations. In *Proceedings of the 21st Symposium on Theory of Computing (STOC)*. 167–176. <https://doi.org/10.1145/73007.73023>
- [44] Alexander Razborov. 1995. Unprovability of lower bounds on circuit size in certain fragments of bounded arithmetic. *Izvestiya of the RAN* (1995), 201–224. Issue 1.
- [45] Alexander Razborov. 1997. On Small Size Approximation Models. In *The Mathematics of Paul Erdős I*. Springer, 385–392. https://doi.org/10.1007/978-3-642-60408-9_28
- [46] Alexander Razborov. 2016. A New Kind of Tradeoffs in Propositional Proof Complexity. *J. ACM* 63, 2 (2016), 16:1–16:14. <https://doi.org/10.1145/2858790>
- [47] Alexander Razborov. 2016. Proof Complexity and Beyond. *SIGACT News* 47, 2 (2016), 66–86. <https://doi.org/10.1145/2951860.2951875>
- [48] Benjamin Rossman. 2014. The Monotone Complexity of k -Clique on Random Graphs. *SIAM J. Comput.* 43, 1 (2014), 256–279. <https://doi.org/10.1137/110839059>
- [49] Janos Simon and Shi-Chun Tsai. 1997. A Note on the Bottleneck Counting Argument. In *Proceedings of the 12th Computational Complexity Conference (CCC)*. 297–301. <https://doi.org/10.1109/CCC.1997.612324>
- [50] Dmitry Sokolov. 2017. Dag-Like Communication and Its Applications. In *Proceedings of the 12th Computer Science Symposium in Russia (CSR)*. Springer, 294–307. https://doi.org/10.1007/978-3-319-58747-9_26
- [51] Avi Wigderson. 1993. The Fusion Method for Lower Bounds in Circuit Complexity. In *Combinatorics, Paul Erdős is Eighty*. János Bolyai Mathematical Society, 453–468.
- [52] Xiaodi Wu, Penghui Yao, and Henry Yuen. 2017. *Raz-McKenzie Simulation with the Inner Product Gadget*. Technical Report TR17-010. Electronic Colloquium on Computational Complexity (ECCC). <https://eccc.weizmann.ac.il/report/2017/010/>