





Multiparty Communication Complexity of Collision-Finding and Cutting Planes Proofs of Concise Pigeonhole Principles

Paul Beame   

University of Washington, Seattle, WA, USA

Michael Whitmeyer   

University of Washington, Seattle, WA, USA

Abstract

We prove several results concerning the communication complexity of a collision-finding problem, each of which has applications to the complexity of cutting-plane proofs, which make inferences based on integer linear inequalities.

In particular, we prove an $\Omega(n^{1-1/k} \log k / 2^k)$ lower bound on the k -party number-in-hand communication complexity of collision-finding. This implies a $2^{n^{1-o(1)}}$ lower bound on the size of tree-like cutting-planes refutations of the bit pigeonhole principle CNFs, which are compact and natural propositional encodings of the negation of the pigeonhole principle, improving on the best previous lower bound of $2^{\Omega(\sqrt{n})}$. Using the method of density-restoring partitions, we also extend that previous lower bound to the full range of pigeonhole parameters.

Finally, using a refinement of a bottleneck-counting framework of Haken and Cook and Sokolov for DAG-like communication protocols, we give a $2^{\Omega(n^{1/4})}$ lower bound on the size of fully general (not necessarily tree-like) cutting planes refutations of the same bit pigeonhole principle formulas, improving on the best previous lower bound of $2^{\Omega(n^{1/8})}$.

2012 ACM Subject Classification Mathematics of computing

Keywords and phrases Proof Complexity, Communication Complexity

Digital Object Identifier 10.4230/LIPIcs.ICALP.2025.21

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version:* <https://eccc.weizmann.ac.il/report/2025/057/> [3]

Funding Research supported by NSF grants CCF-2006359 and CCF-2422205.

1 Introduction

The pigeonhole principle, which asserts that there is no injective function $f : [m] \rightarrow [n]$ for $m > n$, is a cornerstone problem in the study of proof complexity. It is typically encoded as unsatisfiable conjunctive normal form formula (CNF), henceforth denoted PHP_n^m , on the variables $y_{i,j}$, each of which is an indicator that “pigeon” i is mapped to “hole” j .

It is well known that any refutation of PHP_n^{n+1} using resolution proofs requires size $2^{\Omega(n)}$ [12] and the same asymptotic bound holds for all m that are $O(n)$ [6]. On the other hand, if we allow our proof system to reason about linear inequalities (for example using cutting-planes proofs), then it is easy to see that refuting PHP_n^{n+1} becomes easy – indeed, there exist polynomial size refutations of PHP_n^{n+1} .

Despite the pigeonhole principle having short cutting-planes refutations, the related clique-coloring formulas, which state that a graph cannot have both k -cliques and $k - 1$ -colorings, requires exponential-size cutting-planes refutation [24].¹ The clique-coloring formula can

¹ Lower bounds for restricted cutting-planes refutations of these formulas were earlier shown in [17, 4]



© Paul Beame and Michael Whitmeyer;

licensed under Creative Commons License CC-BY 4.0

52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025).

Editors: Keren Censor-Hillel, Fabrizio Grandoni, Joël Ouaknine, and Gabriele Puppis

Article No. 21; pp. 21:1–21:20



Leibniz International Proceedings in Informatics

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



be viewed as a kind of indirect pigeonhole principle: The k nodes of the clique correspond to the pigeons and $k - 1$ colors correspond to the holes, but the representation of possible mappings is quite indirect.

It is natural to wonder about the extent to which indirection is required for the pigeonhole principle to be hard for cutting-planes reasoning. As part of studying techniques for cutting-planes proofs, Hrubeš and Pudlák [16] considered a very natural compact and direct way of expressing the pigeonhole principle, known as the *bit* or *binary* pigeonhole principle². The bit pigeonhole principle analog of PHP_n^m (henceforth denoted BPHP_n^m) has $m \log n$ variables $x_{i,j}$ for $i \in [m]$, $j \in [\log n]$ and the principle asserts that, when we organize these variables as an $m \times [\log n]$ matrix, the rows of the matrix all have distinct values. BPHP_n^m is the following CNF formula: for each $i \neq j \in [m]$, include the clauses of a CNF encoding that $x_i \neq x_j$. One can achieve this by including a clause for each $\alpha \in \{0, 1\}^{\log n}$ expressing that $x_i \neq \alpha \vee x_j \neq \alpha$. The end result is a CNF with $\binom{m}{2}n$ clauses of size $2 \log n$.

Using techniques related to those of [24], Hrubeš and Pudlák [16] showed that BPHP_n^m requires cutting-planes refutations of size $2^{\Omega(n^{1/8})}$ for any $m > n$, proving that even a very direct representation of the pigeonhole principle is hard for cutting-planes proofs. Their arguments, like those of Pudlák, also apply to any proof system that has proof lines consisting of integer linear inequalities with two antecedents per inference that are sound with respect to 01-valued variables; such proofs are known alternatively as *semantic* cutting-planes proofs or $\text{Th}(1)$ proofs [2].

Recently, Dantchev, Galesi, Ghani, and Martin [8] exhibited a $2^{\Omega(n/\log n)}$ lower bound on the size of any general resolution refutation of BPHP_n^m for all $m > n$. In fact, they showed that BPHP_n^m requires proofs of size $2^{\Omega(n^{1-\epsilon})}$ for a more powerful class of proof systems that extend resolution by operating on k -DNFs (known as $\text{Res}(k)$ proofs) for $k \leq \log^{1/2-\epsilon'} n$. (Note that any sound proof system operating on DNFs requires size at least $2^{n^{\Omega(1)}}$ to refute PHP_n^{n+1} [23, 20, 14].) In addition, [8] showed that BPHP_n^m has no refutations in the Sherali-Adams proof system [27] of size smaller than $2^{\Omega(n/\log^2 n)}$. Finally, just as PHP_n^m has polynomial-size Sum-of-Squares refutations [11], Dantchev et al. showed that BPHP_n^m has polynomial-sized Sum-of-Squares refutations.

Given the large lower bounds for resolution, $\text{Res}(k)$, and Sherali-Adams refutations of BPHP_n^m , it is natural to ask the extent to which the sub-exponential lower bounds can be improved for cutting-planes proofs; how close to a $2^{\Omega(n)}$ lower bound is possible?

1.1 Tree-like Proofs and Multiparty Communication

In prior work there has been progress towards this question for the restricted class of *tree-like* refutations. Tree-like proofs require that any time an inequality is used, it must be re-derived (i.e., the underlying graph of deductions is a tree); the polynomial-size cutting-planes refutations of PHP_n^{n+1} can be made tree-like. In contrast, Itsykson and Riazanov [18] showed that BPHP_n^m requires tree-like cutting-planes refutations of size $2^{\Omega(\sqrt{n})}$ when $m \leq n + \sqrt{n}$.

Our first result pushes this bound almost to its limit. Specifically, we prove that any tree-like semantic cutting-planes refutation of BPHP_n^m requires size $2^{n^{1-o(1)}}$ whenever $m \leq n + 2^2 \sqrt{\log n - 2}$.

In order to show this, we utilize a well-known connection between tree-like refutations and communication complexity. While the results of [18] for cutting planes rely on two-party communication complexity (and number-on-forehead multiparty communication for other

² This encoding of the pigeonhole principle was introduced in [1].

results that we mention below), our stronger results are based on multiparty number-in-hand communication. In particular they are based on a similar natural collision-finding communication problem $\text{Coll}_{m,\ell}^k$, in which each player $p \in [k]$ in the number-in-hand model receives an input in $x^{(p)} \in [\ell]^m$, and their goal is to communicate and find a pair $i \neq j \in [m]$ such that $x_i^{(p)} = x_j^{(p)}$ for all players $p \in [k]$. Such a communication problem is well-defined (in the sense that such a pair i, j always exists) when $m > \ell^k$.

This collision-finding problem is intimately related to the unsatisfiable BPHP_n^m formula via the following natural search problem associated with any unsatisfiable CNF formula: Given unsatisfiable CNF φ , the associated search problem Search_φ takes as input a truth assignment α to the variables of φ and requires the output of the index of a clause of φ that is falsified by α . In particular the connection follows by considering a natural k -party number-in-hand communication game that we denote by Search_φ^k wherein the assignment α to the variables of φ is evenly distributed among the k players, and the players must communicate to find an answer for $\text{Search}_\varphi(\alpha)$. It is not hard to see that if we have a communication protocol solving $\text{Search}_{\text{BPHP}_n^m}^k(\alpha)$ then such a protocol also solves $\text{Coll}_{m,n^{1/k}}^k$ on input α .

Our first result is a lower bound on $\text{Coll}_{m,n^{1/k}}^k$ that holds even when we allow randomized protocols.

► **Theorem 1.1.** *The randomized number-in-hand communication complexity of $\text{Coll}_{m,n^{1/k}}^k$ is $\Omega(n^{1-1/k} \log k / 2^k)$ whenever $n + 1 \leq m \leq n + 2^{k-2}n^{1/k}$.*

We pause here to note that this bound is nearly tight. There is a deterministic protocol wherein the first player sends a subset of coordinates of size $\lceil m/n^{1/k} \rceil$ in which their inputs are all equal. This requires $\log \binom{m}{m/n^{1/k}} \lesssim (m/n^{1/k}) \log(m/n^{1/k})$ bits; when $m \approx n$, this is $O(n^{1-1/k} \log(n^{1/k})) = O(n^{1-1/k} \log n / k)$ bits of communication. Player two then announces a subset of these coordinates on which they are equal of size $\lceil m/n^{2/k} \rceil$. The players can continue in this manner until they have found a collision (which is guaranteed by the pigeonhole principle). Note that the amount of communication is bounded by a geometric series, and is dominated by the first term, which results in communication $O(n^{1-1/k} \log n / k)$. This shows that up to logarithmic factors and a factor of 2^k , Theorem 1.1 is tight.

We state here a simplified corollary³ of Theorem 1.1 which formalizes our lower bounds for cutting-planes refutations of BPHP_n^m .

► **Theorem 1.2.** *Any tree-like semantic cutting-planes refutation of BPHP_n^m requires size $2^{n^{1-2/\sqrt{\log n - o(1/\sqrt{\log n})}}}$ when $m \leq n + 2^2\sqrt{\log n - 2}$.*

We remark that Itsykson and Riazanov [18] utilized the same connection between communication and proof complexity to achieve their results. They were also interested in a k -party number-on-forehead version of $\text{Coll}_{m,\ell}^k$ (in particular, in their version, the matrices are added rather than concatenated), which leads to weaker lower bounds in stronger proof systems $\text{Th}(k-1)$ that manipulate degree $k-1$ polynomial inequalities.

Itsykson and Riazanov also left as an open problem whether their bounds for $\text{Search}_{\text{BPHP}_n^m}$ could be extended to the regime of the “weak” pigeonhole principle when $m = n + \Omega(n)$. Göös and Jain [10] first answered this in the affirmative, giving an $\Omega(n^{1/12})$ lower bound on the randomized communication complexity of $\text{Coll}_{2n,n^{1/2}}^2$. Yang and Zhang [32] subsequently

³ When m is somewhat larger, we can obtain somewhat weaker lower bounds.

improved this to an $\Omega(n^{1/4})$ bound, which is tight for randomized computation. Because of the connection between communication protocols and tree-like communication, Yang and Zhang's bound showed that any tree-like cutting planes refutation of BPHP_n^m requires size $2^{\Omega(n^{1/4})}$. We extend the lower bounds of Yang and Zhang for all $k \geq 2$ as follows:

► **Theorem 1.3.** *For all $m > n$ and $2 \leq k \leq \log n / 4$, the randomized number-in-hand communication complexity of $\text{Coll}_{m, n^{1/k}}^k$ is $\Omega(n^{1/2-1/(2k)}/k)$.*

This bound, whose proof combines techniques of Yang, Zhang, and Wang [31, 32], is much weaker than that of Theorem 1.1 when m is close to n but it is tight, up to an $O(\log n)$ factor, for randomized protocols when m is $n + \Omega(n)$. For $k = \Omega(\log n)$, it also implies the following proof complexity bound which asymptotically matches the bound of [18], but applies for all $m > n$.

► **Corollary 1.4.** *For all $m > n$, any tree-like semantic cutting-planes refutation of BPHP_n^m requires size $2^{\tilde{\Omega}(\sqrt{n})}$.*

1.2 General Proofs and DAG-like Protocols

Our second result improves on the $2^{\Omega(n^{1/8})}$ lower bound on the size of cutting-planes refutations of BPHP_n^m (for all $m > n$) of Hrubeš and Pudlák [16], and subsumes the $2^{\Omega(n^{1/4})}$ lower bound for tree-like proofs of Yang and Zhang [32]. Hrubeš and Pudlák used the method of interpolation involving a version of the method of approximation due to Jukna [19]. On the other hand, we use a bottleneck-counting method inspired by recent work of Sokolov [29] that refines a method introduced by Haken and Cook [13].

► **Theorem 1.5.** *Any semantic cutting planes refutation of BPHP_n^m requires size at least $2^{n^{1/4}/\sqrt{2}-2}$, for any $m > n$.*

We prove Theorem 1.5 via a connection between cutting-planes refutations and certain types of DAG-like communication protocols [15, 28]. Specifically, we prove a lower bound on the size of any two party *triangle-DAG* protocol computing $\text{Search}_{\text{BPHP}_n^m}$. We define such protocols formally in Section 2.3.

For now, we state the connection as the following proposition, originally due to Sokolov and Hrubeš and Pudlák [28, 15]. Its proof can be found in Section 2.3 for completeness.

► **Proposition 1.6** ([28, 15]). *Given a semantic cutting planes refutation for CNF formula φ of size s and any partition of the variables of φ into two sets, there is a size s triangle-DAG computing Search_φ .*

In our setting, for BPHP_n^m , there is a natural partition of the variables such that each clause has exactly half of its variables coming from V_x and half of its variables coming from V_y . This can be achieved, for example, by letting V_x be the set of variables corresponding to the first $\frac{1}{2} \log n$ columns of the matrix associated with BPHP_n^m . In light of Proposition 1.6, we derive Theorem 1.5 by proving the following.

► **Theorem 1.7.** *For the natural partition of variables $V_x \sqcup V_y = [m \log n]$, any triangle-DAG computing $\text{Search}_{\text{BPHP}_n^m}$ requires size $2^{n^{1/4}/\sqrt{2}-2}$.*

2 Preliminaries

2.1 Proof Complexity

Proof complexity studies how the size required for refutations of unsatisfiable formulas depends on the size of the formulas themselves. The size of a refutation in general depends on the allowable structure (lines, derivation rules, etc.) of a proof. In general, a proof *system* corresponds to a polynomial-time verifier that can check proofs of a certain format. It generally suffices to derive refutations of unsatisfiable CNF formulas; so it is usual to focus solely on inputs of this form.

For most proof systems, a sequence of deductions can be thought of as a directed graph, where two (or possibly more) lines (whether given or derived) are combined soundly to create a new line. The underlying graph then has edges pointing from the derived inequality to its antecedents.⁴ We say that a proof is *tree-like* if every inequality is used as an antecedent at most once in the proof – that is, if we want to use an inequality twice, we must derive it twice.

For example, in the *resolution* proof system the lines are clauses, and we have the derivation rule $(A \vee x) \wedge (B \vee \neg x) \vdash A \vee B$.

A more powerful proof system is the cutting planes proof system⁵, denoted CP, where lines are linear inequalities. Clauses can be trivially converted into linear inequalities⁶ but, more generally, cutting planes can start with any system of inequalities $Ax \leq b$ where A is a matrix with integer entries. We say that the system is unsatisfiable iff it is unsatisfiable for any $x \in \{0, 1\}^n$. The most basic form of the cutting planes proof system consists of three rules: addition of inequalities, multiplication of inequalities by positive integers, and most crucially, the rounded division rule, also known as the Gomory-Chvátal rule. The rounded division rule is the simple but powerful observation that if a_1, \dots, a_n are all integers with a common factor c , and we have the inequality $a^T x \leq b$, then we can derive that $\frac{1}{c}a^T x \leq \lfloor \frac{b}{c} \rfloor$.

In general, there are many more sound derivation rules for integer/linear inequalities than just rounded division (such as saturation [9] for example), and even more generally, one may allow *any* sound derivation rule for linear inequalities, which yields what is known as *semantic* cutting planes or Th(1) – we use the two names interchangeably.

There is a well-known connection between communication complexity and tree-like proofs, which we will now detail. Given any unsatisfiable formula φ , an assignment α to the variables can be distributed among k players, who must then communicate in order to find a violated clause in φ . This is a search problem, and is denoted $\text{Search}_\varphi(\alpha)$. Short tree-like proofs of the unsatisfiability of φ can often be converted into short protocols for Search_φ using standard techniques.

For example, a short tree-like proof of unsatisfiability of φ using the resolution rule naturally corresponds to a decision tree for finding a violated clause of φ in the following way. For every derivation of the form $(A \vee x) \wedge (B \vee \neg x) \vdash (A \vee B)$ where $A \vee B$ is known to be false, we query x in order to see whether $A \vee x$ or $B \vee \neg x$ is necessarily false. We can continue in this manner, from the root of the tree-like refutation, until we hit an unsatisfied clause in the original formula.

⁴ The direction of arrows in this digraph may seem counterintuitive, but it is convenient when thinking of the graph as a search problem for a violated axiom. In this case, we can follow a path in the graph to find such a violated axiom on one of the leaves/sources.

⁵ Cutting planes proofs can simulate resolution proofs, see e.g. [19, 25].

⁶ For example, $x \vee y \vee z$ could be converted to the inequality $x + y + z \geq 1$.

On the other hand, tree-like semantic cutting planes refutations naturally correspond to *threshold* decision trees, which we now define.

► **Definition 2.1** (Threshold Decision Tree). *A threshold decision tree is a tree whose vertices are labeled with inequalities of the form $a_1x_1 + \dots + a_nx_n \leq b$ where a_1, \dots, a_n, b are integers. Edges are labelled with 0 or 1, and leaves are axioms of a system of inequalities $Ax \leq b$.*

We traverse a threshold decision tree by computing the threshold function at the root, following the corresponding edge, and continuing in this manner until we hit a leaf. We say that a threshold decision tree computes the search problem for a formula φ if this process leads to a leaf corresponding to a violated clause in φ .

First, we have the following well-known lemma, which states that one can derive a low-depth threshold decision tree from a small $\text{Th}(1)$ refutation.⁷

► **Proposition 2.2.** *Given a size S tree-like $\text{Th}(1)$ refutation of an unsatisfiable system $Ax \leq b$, there is a depth $O(\log S)$ threshold decision tree finding a violated axiom.*

Proposition 2.2 goes back to the work of Impagliazzo, Pitassi, and Urquhart [17], and can also be found for instance in [19]. We omit the proof, but the idea is a common one: find a node in the tree with roughly half (between $1/3$ and $2/3$) of the leaves as its descendants, make that the root of the threshold decision tree, and recurse.

2.2 Communication Complexity

We mainly focus on k -party number-in-hand communication, wherein each player p receives an input $x^{(p)}$, and the players' goal is to communicate as little as possible in order to compute a known function or relation involving their collective inputs. In general, players may have access to shared randomness, and we allow incorrect answers with probability $1/3$.

An important function for us is the number-in-hand disjointness problem with k players and input size n , henceforth denoted DISJ_n^k . This is the communication problem wherein each player has an input in $\{0, 1\}^n$, and they must decide if there exists a coordinate i for which they all have a 1 in that coordinate. Disjointness in general is an extremely well-studied problem [7, 25], and for the specific case of the NIH model, we have the following lower bound due to Braverman and Oshman.

► **Theorem 2.3** ([5]). *The randomized communication complexity of DISJ_n^k is $\Omega(n \log k)$.*

Our results rely on the following connection between threshold decision trees for finding violated clauses and k -party NIH communication. It is closely related to previous work.

► **Lemma 2.4.** *For $x \in \{0, 1\}^n$, if an unsatisfiable system $Ax \leq b$ has a threshold decision tree of depth $d \leq n$ finding a violated axiom, then for any partition of the input variables into k parts there is a randomized protocol for $\text{Search}_{\text{NIH}}^k(Ax \leq b)$ using $O(dk \log k \log n)$ bits of communication.*

Lemma 2.4 is similar for example to Lemma 5 in [17] or Lemma 19.11 in [19], but is slightly stronger so we give a brief proof. Using a well-known theorem of Muroga [21] we can assume, without loss of generality, that each of the threshold functions in the decision tree has integer coefficients of absolute value at most $2^{-n}(n+1)^{(n+1)/2}$.

The players evaluate these threshold functions using the following efficient randomized protocol given by Viola [30] which tightens and refines a protocol first given by Nisan [22].

⁷ As noted in [19], there is no meaningful converse to this statement, since if there are m inequalities in our unsatisfiable system, there exists a trivial depth m threshold decision tree finding one that is violated.

► **Proposition 2.5** ([30]). *Suppose that each player $p \in [k]$ receives an input $x^{(p)} \in [-2^n, 2^n]$. Then there is a randomized number-in-hand protocol with error at most ε that determines whether $\sum_p x^{(p)} > s$ and communicates $O(k \log k \log(n/\varepsilon))$ bits.*

► **Corollary 2.6.** *Suppose that each player $p \in [k]$ receives an input $x^{(p)} \in [2^t]$. Then they can execute a randomized number-in-hand protocol to determine whether $\sum_p a_p x^{(p)} \leq b$, where each $|a_p| \leq 2^w$ with error at most ε using at most $O(k \log k \log((w+t)/\varepsilon))$ bits of communication.*

Proof of Lemma 2.4. Given a threshold decision tree of depth d , which has coefficients of magnitude at most $2^{O(n \log n)}$ without loss of generality [21], we simply traverse it from the root, evaluating the threshold functions until a root is reached. By Corollary 2.6, the players can communicate $O(k \log((n \log n)/\varepsilon) \log k)$ bits to compute the threshold function with error probability ε . Setting $\varepsilon = \Theta(1/d)$ and continuing in this manner, by a union bound the threshold functions are all computed correctly with constant probability. Using the assumption that $d \leq n$, this yields a protocol communicating $O(dk \log k \log n)$ bits. ◀

Given a system of unsatisfiable inequalities $Ax \leq b$, and a partition of an assignment $\alpha \in \{0, 1\}^n$ between k players, there is a natural (number-in-hand) communication game wherein players must communicate to find an axiom violated by α . Lemma 2.4 implies the following result connecting communication complexity and proof complexity.

► **Lemma 2.7.** *For any partition of n variables into k parts, if $\text{Search}_{NIH}^k(Ax \leq b)$ requires t bits of communication, then any tree-like $\text{Th}(1)$ refutation of $Ax \leq b$ requires size $2^{\Omega(t/(k \log k \log n))}$.*

Proof. By Proposition 2.2, given a size S tree-like $\text{Th}(1)$ proof, we get a depth $d = O(\log S)$ threshold decision tree finding a violated axiom. By Lemma 2.4, there exists a communication protocol finding a violated axiom using $O(\log S k \log k \log n)$ bits of communication. This implies that $\log(S)k \log k \log n \geq ct$ for constant c , which in turn implies S is $2^{\Omega(t/(k \log k \log n))}$, as desired. ◀

2.3 Triangle-DAGs

Given a bipartite domain $X \times Y$, a *triangle* $T \subseteq X \times Y$ is any set that can be written as $T = \{(x, y) : a_T(x) < b_T(y)\}$ for some labelling $a_T : X \rightarrow \mathbb{R}$ and $b_T : Y \rightarrow \mathbb{R}$.

A triangle-DAG computing a search problem $\text{Search} \subseteq (X \times Y \times \mathcal{O})$ is a directed acyclic graph D of fan-out at most 2 where each node $u \in D$ is associated with a triangle $T_u \subseteq X \times Y$ satisfying the following:

- there is a distinguished root node r with fan-in zero, and $T_r = X \times Y$, and
- for each non-sink node u with children v, v' we have $T_u \subseteq T_v \cup T_{v'}$, and
- each sink node u is labeled with an output o such that $T_u \subseteq \text{Search}^{-1}(o)$.

Given these definitions, we restate Proposition 1.6 for convenience. Proofs may be found in [28, 15] or in our full paper [3].

► **Proposition 1.6** ([28, 15]). *Given a semantic cutting planes refutation for CNF formula φ of size s and any partition of the variables of φ into two sets, there is a size s triangle-DAG computing Search_φ .*

3 Multiparty Communication Lower Bounds for $m \leq n + \sqrt{n}$

In this section we prove Theorem 1.1, which we now recall.

► **Theorem 1.1.** *The randomized number-in-hand communication complexity of $\text{Coll}_{m, n^{1/k}}^k$ is $\Omega(n^{1-1/k} \log k / 2^k)$ whenever $n + 1 \leq m \leq n + 2^{k-2} n^{1/k}$.*

The idea for the proof is to exhibit a random reduction from the decision problem $\text{DISJ}_{m^{k-1}}^k$ to the collision problem. This is analogous to the approach of Itsykson and Riazanov [18] for number-on-forehead communication complexity which used lower bounds for disjointness in that model and a randomized decision-to-search reduction paradigm introduced by Raz and Wigderson [26] to prove lower bounds on the monotone depth complexity of matching. The details and parameters of our reduction are necessarily quite different.

In our setting, we embed k players' inputs to a disjointness problem into k matrices such that, when these matrices are concatenated, the resulting matrix has distinct rows if and only if the players' inputs were disjoint. We can then add a few “fake” rows to this matrix and run our algorithm for $\text{Coll}_{m, n^{1/k}}^k$, and see if the collisions it finds involve the fake rows or not. If so, we conclude that the inputs were disjoint and if not, we know that they were not disjoint.

The following key combinatorial lemma allows us to carry out the first step of this process.

► **Lemma 3.1.** *For all integers $k \geq 1$ there exist matrices $M_k^1 \in \{0, 1\}^{2^k \times k}$ and $M_k^0 \in \{0, 1\}^{2^k \times k}$ such that*

1. M_k^1 has 2^{k-1} unique pairs of identical rows.
2. For any string $b \in \{0, 1\}^k$, define the matrix $M_k(b_1, \dots, b_k)$ as the matrix formed by making its i -th column equal to the i -th column of M_k^0 if $b_i = 0$, and equal to the i -th column of M_k^1 if $b_i = 1$. Then $M_k(b)$ has unique rows for all $b \neq \vec{1}$.

We defer the proof of Lemma 3.1 to Section 3.1.

Proof of Theorem 1.1 using Lemma 3.1. As alluded to, we will reduce the NIH disjointness problem to our bit-pigeonhole problem. Namely, we will reduce $\text{DISJ}_{m^{k-1}}^k$ to a $\text{Coll}_{m, \tilde{\ell}}^k$ communication game, where $\tilde{m} = m^k 2^k + 2^{k-1} m$, and $\tilde{\ell} = 2m$.

The players get $x^{(1)}, \dots, x^{(k)} \subseteq [m^{k-1}]$ (viewed as bit strings of length m^{k-1}), and need to determine whether $x^{(1)} \cap \dots \cap x^{(k)} = \emptyset$. First, we define $B_{m^k}(j) := [\text{bin}_{m^k}(j), \dots, \text{bin}_{m^k}(j)]^T \in \{0, 1\}^{2^k \times k \log m}$, where the same row is repeated 2^k times.

For each $i \in [m^{k-1}]$, consider the matrix $M_k(x_i^{(1)}, \dots, x_i^{(k)})$ from Lemma 3.1. Note that each player p knows the p -th column of $M_k(x_i^{(1)}, \dots, x_i^{(k)})$ for all i , but without communication the other players do not know this column. Then we can define

$$\tilde{M} := \begin{bmatrix} M_k(x_1^{(1)}, \dots, x_1^{(k)}) & B_{m^k}(0) \\ M_k(x_1^{(1)}, \dots, x_1^{(k)}) & B_{m^k}(1) \\ \vdots & \vdots \\ M_k(x_1^{(1)}, \dots, x_1^{(k)}) & B_{m^k}(m-1) \\ M_k(x_2^{(1)}, \dots, x_2^{(k)}) & B_{m^k}(m) \\ \vdots & \vdots \\ M_k(x_2^{(1)}, \dots, x_2^{(k)}) & B_{m^k}(2m-1) \\ \vdots & \vdots \\ M_k(x_{m^{k-1}}^{(1)}, \dots, x_{m^{k-1}}^{(k)}) & B_{m^k}(m^k-1) \end{bmatrix} \in \{0, 1\}^{m^k \cdot 2^k \times (k \log m + k)}.$$

Observe that each player p can construct their “part” of this matrix without communicating by constructing the p -th column of every $M_k^{S_i}$ matrix (which only depends on $x^{(p)}$), and then taking the p -th part of each of the $B_{m^k}(j)$ matrices.

Lemma 3.1 lets us connect the distinctness property of this matrix with the disjointness property of the players’ inputs.

▷ **Claim 3.2.** If $(x^{(1)}, \dots, x^{(k)})$ are disjoint, then \widetilde{M} has distinct rows.

Proof. The only possible collisions happen in every group of 2^k rows, since $B_{m^k}(j)$ has every row different from $B_{m^k}(i)$ for all $i \neq j$. Within these groups, by Lemma 3.1, if the inputs are disjoint then there are no collisions. ◁

▷ **Claim 3.3.** If X is not disjoint, then there are at least $2^{k-1}m$ pairs of colliding rows in \widetilde{M} .

Proof. Any coordinate i for which $x_i^{(p)} = 1$ for all p generates $S_i = \emptyset$, which by Lemma 3.1 generates $2^{k-1}m$ such pairs, since input i was repeated m times. ◁

We cannot run any collision protocol for \widetilde{M} yet, as there are not guaranteed collisions. To address this, the players use shared randomness to put an additional $2^{k-1}m$ rows at the bottom of \widetilde{M} . These rows will be chosen randomly with the following two properties:

1. Each fake row will be distinct.
2. Each player’s “part” of the matrix (which consists of $\log m + 1$ columns) when restricted to these rows will repeat the $2m$ unique possible bit strings an addition $2^{k-1}m/(2m) = 2^{k-2}$ times.⁸

Denote this new matrix M . M now has “fake” collisions which involve any of the last $2^{k-1}m$ rows.

Let \mathcal{A} denote the randomized protocol solving the $\text{Coll}_{(2m)^k + 2^{k-1}m, 2m}^k$ problem.

Observe that if the inputs are disjoint, then the only collisions in M involve fake rows. The players would like to feed their parts of this matrix into \mathcal{A} and conclude that their inputs are disjoint if the output involves a fake row, and conclude that they were not disjoint if the output involves two non-fake rows. However, this is problematic, as we have no guarantees over how \mathcal{A} behaves, and it could always find a collision involving one of the last $2m$ rows (which it knows are fake), regardless of if there are other collisions. This necessitates the following random shuffling.

1. Each player applies (the same) random permutation $\pi : [2^k m^k + 2^{k-1}m] \rightarrow [2^k m^k + 2^{k-1}m]$ which shuffles the rows of M .
2. Each player applies an individual random permutation $\pi^{(p)} : [2m] \rightarrow [2m]$ to each of their rows. Note that this preserves collisions/distinctness in the concatenation.

Denote $\vec{\pi} := (\pi, \pi^{(1)}, \dots, \pi^{(k)})$, and call this final matrix $M_{\vec{\pi}}$.

Algorithm. The algorithm for disjointness is as follows: the players use their inputs and shared randomness to compute (without communication) their respective parts of 5 independent copies of an $M_{\vec{\pi}}$ constructed in the above manner, and run \mathcal{A} using these as inputs. The players then examine the outputs $(i_1, j_1), \dots, (i_5, j_5)$ of the algorithm on these five inputs. They then exchange an additional $O(k \log m)$ bits to determine if each claimed collision actually was a collision. Finally, if any of the claimed collisions were actually collisions on rows that were not fake (under the appropriate permutation), then the players can conclude with certainty that their inputs were not disjoint. Otherwise, if \mathcal{A} only ever finds colliding pairs that involve a row the players know is fake (or otherwise fail to find any collisions), then players guess that $(x^{(1)}, \dots, x^{(k)})$ were disjoint.

⁸ This is important because if we bias and have a certain fake row appear more often in the input for player p , then \mathcal{A} could potentially detect and use this to its advantage.

Analysis. We analyze one iteration of the algorithm. Suppose \mathcal{A} has error probability at most $1/3$ – that is, with probability at least $2/3$, it outputs two rows i, j that are equal.

Suppose $(x^{(1)}, \dots, x^{(k)})$ are disjoint. Then by assumption, \mathcal{A} finds a collision with probability at least $2/3$, and we know this collision will always involve a fake row by Claim 3.2. Therefore the players will correctly output that their inputs were distinct with probability at least $2/3$, and this is only improved by the five-fold repetition.

Otherwise, suppose that the players' inputs were not disjoint. Suppose further that \mathcal{A} successfully finds a collision – this happens with probability at least $2/3$. Recall that by Claim 3.3 $M_{\vec{\pi}}$ will have at least $2^{k-1}m$ distinct pairs of real collisions. Adding the $2^{k-1}m$ fake rows produced additional “fake” collisions. These fake rows could have created up to $2^{k-1}m$ additional unique pairs of fake collisions, or could have “joined” the real collisions, creating up to $2^{k-1}m$ groups of 3 equal rows in $M_{\vec{\pi}}$.

If \mathcal{A} outputs a collision from one of the groups of three, then because we applied random permutations to the rows, it is equally likely to have chosen any of the 3 possible pairs. Therefore, with probability at least $1/3$, it outputs a real collision, and the players successfully discover that they are not disjoint. Otherwise, if \mathcal{A} outputs one of the unique collision pairs, then (again because we have applied random permutations to the rows), any such unique collision is equally likely to be output. If t of the fake rows formed a group of three with real collisions, then that leaves at most $2^{k-1}m - t$ fake rows to collide with a different unique row. It also leaves $2^{k-1}m - t$ untouched real collisions, so \mathcal{A} outputs a real collision with probability at least $1/2$. Either way, the probability that \mathcal{A} outputs a real collision is at least $2/3 \cdot 1/3 = 2/9$.

Therefore, after repeating this 5 times independently, the probability of seeing at least one real collision is at least $1 - (7/9)^5 > 2/3$.

Let $n := (2m)^k$. We have shown that if $\text{Coll}_{m,\ell}^k \sim$ with input size $\tilde{m} = 2^k m^k + 2^{k-1}m = n + 2^{k-2}n^{1/k}$ can be solved with $o(n^{1-1/k} \log k/2^k)$ communication, then we can solve the decision disjointness problem with $o(n^{1-1/k} \log k/2^k) + O(k \log m)$ which is at most $o(m^{k-1} \log k)$, contradicting the $\Omega(m^{k-1} \log k)$ lower bound from Theorem 2.3. ◀

3.1 Proof of Lemma 3.1

We first recall Lemma 3.1.

► **Lemma 3.1.** *For all integers $k \geq 1$ there exist matrices $M_k^1 \in \{0, 1\}^{2^k \times k}$ and $M_k^0 \in \{0, 1\}^{2^k \times k}$ such that*

1. M_k^1 has 2^{k-1} unique pairs of identical rows.
2. For any string $b \in \{0, 1\}^k$, define the matrix $M_k(b_1, \dots, b_k)$ as the matrix formed by making its i -th column equal to the i -th column of M_k^0 if $b_i = 0$, and equal to the i -th column of M_k^1 if $b_i = 1$. Then $M_k(b)$ has unique rows for all $b \neq \vec{1}$.

Proof. Let $\mathcal{E}_k \subseteq \{0, \dots, k-1\}$ be the set of integers with an even number of 1s in their binary representation.

Define

$$M_k^1 = \begin{bmatrix} \text{bin}_k(y_1) \\ \text{bin}_k(y_1) \\ \text{bin}_k(y_2) \\ \text{bin}_k(y_2) \\ \vdots \\ \text{bin}_k(y_{2^{k-1}}) \\ \text{bin}_k(y_{2^{k-1}}) \end{bmatrix} \in \{0, 1\}^{2^k \times k},$$

where y_i is the i -th smallest number in \mathcal{E}_k .

We pause here to note that each of the columns of M_k^1 are actually each the truth table of a linear function. Let $f_1 : \{0, 1\}^k \rightarrow \{0, 1\}$ be the linear function $f_1(x) = \langle x, e_1 \rangle$, where e_1 is the first standard basis vector. Then we can describe the first column of M_k^1 as the truth table of f_1 . More generally, we have that for $i < k$ the i -th column of M_k^1 is the truth table of $f_i(x) = \langle x, e_i \rangle$, and the last column is the truth table of $f_k(x) = \langle x, e_1 + \dots + e_{k-1} \rangle$.

If we define F_k^1 to be the matrix whose column i is the vector whose inner product we are taking with x in f_i , and $B_k \in \mathbb{F}_2^{2^k \times k}$ whose rows are the binary strings written in order, then we have that

$$M_k^1 = \begin{bmatrix} \text{bin}_k(0) \\ \text{bin}_k(1) \\ \vdots \\ \text{bin}_k(2^k - 1) \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & \cdots & 1 \\ 0 & 1 & 0 & \cdots & 1 \\ 0 & 0 & 1 & \ddots & 1 \\ \vdots & & & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \\ 0 & 0 & 0 & \cdots & 0 \end{bmatrix} =: B_k F_k^1$$

where all operations are over \mathbb{F}_2 . M_k^1 has repeated rows precisely because F_k^1 has linearly dependent columns.

With this perspective in mind, if we can find a $k \times k$ matrix F_k^0 over \mathbb{F}_2 such that replacing any (nonzero) number of columns of F_k^1 with corresponding columns in F_k^0 produces a matrix with linearly independent columns, then we are done, as we can let $M_k^0 := B_k F_k^0$.

We define F_k^0 to be the following lower triangular matrix:

$$F_k^0 := \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 \\ 1 & 1 & 0 & \cdots & 0 \\ 1 & 1 & 1 & \ddots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & 1 & \cdots & 1 \end{bmatrix},$$

Clearly F_k^0 is full rank.

We claim that replacing any nonempty set S of columns of F_k^1 with the corresponding columns in F_k^0 produces a matrix F_k^S with linearly independent columns. Consider arbitrary k , and arbitrary nonempty $S \subseteq [k]$.

Case 1. Suppose $k \in S$, that is, the last column of F_k^S is e_k . Then our matrix is lower triangular with 1s on the diagonal, and so has linearly independent columns.

Case 2. Suppose $k \notin S$, that is, the last column of F_k^S is $e_1 + \dots + e_{k-1}$. In this case, the first $k-1$ columns are lower triangular with 1s on the diagonal, and therefore their span is equal to $\text{span}\{e_1, \dots, e_{k-1}\}$. It suffices then to show that also e_k is in the column span. Towards this goal, take the minimal $0 < i < k$ such that $i \in S$, and observe that summing the first i columns of F_k^S equals $\vec{1}$, the all 1s vector. Adding this to the last column produces e_k , so we are done. ◀

4 Tree-like Proof Complexity Lower Bounds

In this section, we prove the following more detailed version of Theorem 1.2.

► **Theorem 4.1.** *When $m \leq n + 2^{k-2}n^{1/k}$, any tree-like semantic cuttings planes refutation of BPHP_n^m must have size at least $2^{\Omega(n^{1-1/k}2^{-k}/(k \log n))}$.*

Proof. The theorem follows quite readily from Lemma 2.7 and the fact that $\text{Search}_{\text{BPHP}_n^m}^k$ reduces to $\text{Coll}_{m, n^{1/k}}^k$. If we translate BPHP_n^m to a system of inequalities, then there are $O(nm^2) = O(n^3)$ inequalities on $m \log n$ variables.

Lemma 2.7 then says that any tree-like semantic cutting planes proof of the unsatisfiability of BPHP_n^m must have size at least $2^{\Omega(t/(k \log k \log n))}$. By Theorem 1.1, $t = \Omega(n^{1-1/k} 2^{-k} \log k)$. Plugging this in yields that the tree-like size must be at least $2^{\Omega(n^{1-1/k} 2^{-k}/(k \log n))}$. ◀

From Theorem 4.1 we achieve Theorem 1.2, which we restate now.

► **Corollary 4.2.** *Any tree-like semantic cutting planes refutation of BPHP_n^m requires size $2^{n^{1-2/\sqrt{\log n} - o(1/\sqrt{\log n})}}$ when $m \leq n + 2^2 \sqrt{\log n} - 2$.*

Proof. Let $k = \sqrt{\log n}$. Plugging this into the bound from Theorem 4.1, we get that $m \leq n + 2^2 \sqrt{\log n} - 2$. In this regime, Theorem 4.1 gives the size lower bound

$$2^{\Omega(n^{1-1/\sqrt{\log n}} n^{-1/\sqrt{\log n}} / (\log^{3/2} n))} = 2^{cn^{1-2/\sqrt{\log n} - 1.5 \log \log n / \log n}},$$

for appropriate constant c . Using the fact that $c = n^{\log c / \log n}$, the bound becomes

$$2^{n^{1-2/\sqrt{\log n} - o(1/\sqrt{\log n})}}.$$

5 Triangle-DAG Lower Bounds

In this section, we prove Theorem 1.7, which we restate here for convenience.

► **Theorem 1.7.** *For the natural partition of variables $V_x \sqcup V_y = [m \log n]$, any triangle-DAG computing $\text{Search}_{\text{BPHP}_n^m}$ requires size $2^{n^{1/4}/\sqrt{2}-2}$.*

Before proving Theorem 1.7, we need the following notion of triangle slices. For a triangle $T \subseteq X \times Y$ and a given $x \in X$, we define the slice T^x to be $T \cap (\{x\} \times Y)$, and define the slice T^y analogously. Note that for any triangle T defined by functions a and b there is a natural total pre-order on X where $x \preceq x'$ iff $a(x') \leq a(x)$ and hence $T^x \subseteq T^{x'}$; similarly, there is a total pre-order on Y where $y \preceq y'$ iff $b(y) \leq b(y')$ and hence $T^y \subseteq T^{y'}$. This immediately implies the following.

► **Proposition 5.1.** *Given a triangle T , if $y \in Y$ is such that $|T^y|$ is maximized, then for all $(x, y') \in T$ we have $(x, y) \in T$.*

We will have the following proposition.

► **Proposition 5.2.** *For any triangle T and rectangle R , $T \cap R$ is a triangle.*

Proof. By definition, $T = \{(x, y) : a(x) < b(y)\}$ for some a, b . If $R = X' \times Y'$, simply modify $a(x) = \infty$ for all $x \notin X'$ and $b(y) = -\infty$ for all $y \notin Y'$. ◀

The proof of Theorem 1.7 uses the *bottleneck counting* method of Haken and Cook for cutting planes proofs [13] as adapted for direct use on triangle-DAGs by Sokolov [29].

Let $V(D)$ denote the nodes of a triangle-DAG D computing $\text{Search}_{\text{BPHP}_n^m}$ under the natural partition of variables. At a high level, to prove Theorem 1.7, we will construct a partial function $\mu : X \cup Y \rightarrow V(D)$ such that

- (a) (Lemma 5.5) most $x \in X$ and $y \in Y$ are mapped by μ to a node of D , and
- (b) (Lemma 5.6) there is no node of D with too many x or y assigned to it by μ .

■ **Algorithm 1** Defining the partial map $\mu : X \cup Y \rightarrow V(D)$.

Input: triangle-DAG D on input set $X \times Y$ with node set $V(D)$
Output: a partial map $\mu : X \cup Y \rightarrow V(D)$

```

1 Let  $X' := X$  and  $Y' := Y$  be the remaining coordinates we have not yet assigned by  $\mu$ 
2 for  $u$  in some topological ordering of  $V(D)$ , starting from the sinks do
3    $\tilde{T}_u := T_u \cap (X' \times Y')$ 
4   for  $x \in X'$  do
5     if  $w(\tilde{T}_u, x) \geq k$  then
6       Let  $\mu(x) := u$  and delete  $x$  from  $X'$ .
7    $\tilde{T}_u := T_u \cap (X' \times Y')$ 
8   for  $y \in Y'$  do
9     if  $w(\tilde{T}_u, y) \geq k$  then
10      Let  $\mu(y) := u$  and delete  $y$  from  $Y'$ .
11    $\tilde{T}_u := T_u \cap (X' \times Y')$ 
12 return  $\mu$ 

```

It is then easy to combine these two properties and conclude that $|V(D)|$ must be large⁹.

For any $i \neq j \in [m]$, there is an associated rectangle $R_{i,j} \subseteq X \times Y$ of inputs that violate at least one of the n clauses enforcing an inequality between coordinates i and j given by

$$R_{i,j} = \{(x, y) \in X \times Y \mid x_i = x_j, y_i = y_j\}.$$

Observe that the set of all rectangles $\{R_{i,j} : i \neq j \in [m]\}$ covers the entire space $X \times Y$ as long as $m > n$. (That is, $X \times Y = \bigcup_{i \neq j \in [m]} R_{i,j}$.) Note that in order for a triangle-DAG to correctly compute $\text{Search}_{\text{BPHP}_n^m}$, it must, at the very least, have each sink be contained in a single $R_{i,j}$.

► **Definition 5.3.** For a node $u \in V(D)$, let T_u denote the triangle associated with that node. Define the width $w(T, z)$ of z in triangle T to be the minimal size of any subset $S \subseteq \{R_{i,j} : i \neq j \in [m]\}$ such that $T^z \subseteq \bigcup_{R \in S} R$.

With this definition in hand, we can define our $\mu : X \cup Y \rightarrow V(D)$ using Algorithm 1 which depends on a parameter k that we will eventually set to $\sqrt{n}/4$.

Intuitively, in using Algorithm 1 we are going backwards through our DAG D , and whenever the protocol had too many options for what clause to output when the protocol was at a node u for a fixed $z \in X \cup Y$ (which corresponds to large width), we assign z to u . Arguing that $|\mu^{-1}(u)|$ is not too large is thus arguing that a protocol cannot make huge strides in learning about which clause to output for too many $z \in X \cup Y$ suddenly in one single node u .

⁹ The name bottleneck counting came from its original version due to Haken [12] in which a class of *full* input assignments was mapped by a function μ to nodes in a proof DAG and each input assignment could be viewed as “flowing” from the root to a particular sink node (defined by the assignment). In that case, μ identified a node on such a path, a bottleneck node, that did not permit many assignments to flow through it. Since there were many assignments, only few of which could pass through any bottleneck node identified by μ , there must be many bottleneck nodes in the DAG. In the modified form for cutting planes/triangle-DAGs, while the full assignment is flowing through the triangle-DAG, the mapping μ is based on half of the input assignment, either the x or y portion, and the sink node is not unique; not all assignments are assigned in this way, but many are.

We now make some basic observations about our definition of μ . First, note that no $z \in X \cup Y$ is assigned to any sink node of D , since every sink node $s \in V(D)$ is wholly contained in an $R_{i,j}$, and therefore we have that $w(T_s, z) = 1$ for all z . The following observation will be of critical importance in our analysis.

► **Claim 5.4.** During the execution of Algorithm 1, for every u in triangle-DAG D and every $z \in X' \cup Y'$, $w(\tilde{T}_u, z) \leq 2k$.

Proof. By definition in a triangle-DAG, if v and v' are children of node u , then we have that $T_u \subseteq T_v \cup T_{v'}$. In particular, we know that if $z \in X \cup Y$ was not assigned to nodes v or v' , then $w(\tilde{T}_u, z) \leq 2k$, since we can take the union of the coverings for \tilde{T}_v^z and $\tilde{T}_{v'}^z$. ◀

Finally, we note that at the end of Algorithm 1 at the root r , for all $z \in X' \cup Y'$, we have that $w(\tilde{T}_r, z) < k$.

► **Lemma 5.5.** Let $k \leq \sqrt{n}/2$. For μ constructed using Algorithm 1, at least $|X|/2$ elements of X are assigned by μ or at least $|Y|/2$ elements of Y are assigned by μ .

Proof. At the end of Algorithm 1 after processing the root r of D , since $T_r = X \times Y$, we have $\tilde{T}_r = X' \times Y'$ which is the rectangle of unassigned inputs. If $|X'| \leq |X|/2$ then we are done; otherwise, let $x \in X'$. Since $w(\tilde{T}_r, x) < k$, there exists a set S , $|S| < k$, such that for all $y \in Y'$, $y_i = y_j$ for at least one $(i, j) \in S$.

So, we obtain

$$\Pr_{y \sim Y}[y \in Y'] \leq \Pr[\exists (i, j) \in S \text{ s.t. } y_i = y_j] \leq |S| \max_{(i, j) \in S} \Pr[y_i = y_j] < k/\sqrt{n} \leq 1/2.$$

Hence, $|Y'| < |Y|/2$, as desired. ◀

This shows that a large number of elements of $X \cup Y$ are assigned by μ . It remains to show that each node in the protocol D only has a bounded number of elements of $X \cup Y$ assigned by μ .

► **Lemma 5.6.** Let μ be defined by Algorithm 1 with $k = \sqrt{n}/4$. For all $u \in V(D)$, μ maps at most $|X| \cdot 2^{-n^{1/4}/\sqrt{2}+1}$ elements of X to u ; the analogous bound also holds for Y .

Proof. We fix any $u \in V(D)$ and focus on the number of elements $x \in X$ that μ maps to u ; the bound for the number of such $y \in Y$ is identical. In particular, consider the values of X' and Y' and \tilde{T}_u in Algorithm 1 immediately before processing node u ; any $x \in X$ that μ maps to u must be in X' .

Our strategy will be to construct a set of *potential* coverings for the slices $\{\tilde{T}_u^x : x \in X'\}$ by monochromatic rectangles. We will then aim to show that most $x \in X'$ have width $< k$ even when restricted to this set of potential coverings, thereby limiting how many $x \in X'$ could be assigned by μ to u .

We now describe this potential covering in detail. Intuitively, our strategy will be to construct a tree \mathcal{T} , each of whose nodes is labelled by a triangle of inputs that are not yet covered, where each edge of the tree is labelled with a pair $i \neq j \in [m]$. We say that $x \in X'$ is *consistent* with a node $t \in \mathcal{T}$ iff for all (i, j) labels on the edges of the unique path to t in \mathcal{T} , we have that $x_i = x_j$. Furthermore, \mathcal{T} will be constructed so that if $x \in X'$ is consistent with a leaf t of \mathcal{T} , then \tilde{T}_u^x is covered by $\{R_{i,j} : (i, j) \text{ is on the path to } t\}$. We produce the tree \mathcal{T} of potential coverings of \tilde{T}_u^x for all $x \in X'$ using Algorithm 2 applied with $\tilde{T} = \tilde{T}_u$, $X = X'$, and $Y = Y'$ which satisfy the preconditions by Claim 5.4.

■ **Algorithm 2** Defining tree \mathcal{T} , which encodes a set of potential coverings of \tilde{T}^x for all $x \in X$.

```

Input: A triangle  $\tilde{T} \subseteq X \times Y$  such that  $w(\tilde{T}, y) \leq 2k$  for all  $y \in Y$ 
Output: A tree  $\mathcal{T}$  of potential coverings of  $\tilde{T}^x$  for all  $x \in X$ .
1 Initialize  $\mathcal{T}$  as a tree with a single node labelled  $\tilde{T}$ 
2 Let  $\text{Leaves} := \{\tilde{T}\}$ 
3 while  $\text{Leaves} \neq \emptyset$  do
4   Choose some  $T$  from  $\text{Leaves}$ , and remove it from  $\text{Leaves}$ 
5   Choose a  $y \in Y$  such that  $|T^y|$  is maximized
      /*  $\forall x \in X$  if  $T^x \neq \emptyset$  then  $(x, y) \in T$  by Proposition 5.1. */
6   Let  $\{R_{i_1, j_1}, \dots, R_{i_\ell, j_\ell}\} \subseteq \{R_{i, j} : i \neq j \in [m]\}$  be a minimal covering for  $T^y$ .
      /* Note that  $\ell \leq 2k$ ; see Figure 1 for a visualization of this step. */
7   for  $m = 1, \dots, \ell$  do
8     Let  $R := R_{i_m, j_m} = X_R \times Y_R$ 
9     Let  $T_R := T \cap (X_R \times (Y \setminus Y_R))$ 
      /* This is still a triangle by Proposition 5.2. */
10    Add edge to  $\mathcal{T}$  labelled by  $(i_m, j_m)$  from the leaf of  $\mathcal{T}$  labelled  $T$  to new node
        labelled by  $T_R$ 
11    if  $T_R \neq \emptyset$  then
12      Add  $T_R$  to  $\text{Leaves}$ .
13 return  $\mathcal{T}$ 


```

First, observe that any $x \in X'$ such that $\tilde{T}^x \neq \emptyset$ is consistent with at least one leaf in \mathcal{T} , by the choice of y in Algorithm 2 along with Proposition 5.1. We define the set of equality constraints of a leaf $t \in \mathcal{T}$ to be the set $E(t) := \{i \neq j \in [m] : (i, j) \text{ is on the unique path to } t \text{ in } \mathcal{T}\}$. Observe that we have the desired property that if x is consistent with a leaf t of \mathcal{T} , then \tilde{T}_u^x is covered by $\{R_{i, j} : (i, j) \in E(t)\}$, thereby ensuring that $w(\tilde{T}_u, x) \leq \text{depth}_{\mathcal{T}}(t)$.

Therefore, if μ maps $x \in X'$ to u then x must be consistent with a leaf of \mathcal{T} of depth at least k . Let \mathcal{T}' denote tree \mathcal{T} truncated at depth k . Thus, the total number of $x \in X'$ that μ maps to u is at most

$$\sum_{\substack{t \in \text{Leaves}(\mathcal{T}') \\ \text{depth}_{\mathcal{T}'}(t) = k}} |\{x \in X : x \text{ is consistent with } t\}|. \quad (1)$$

► **Remark 5.7.** Before explaining the details of our analysis, it is useful to discuss a naive bound on (1) that does not suffice to yield our bounds: Naively, there could be $(2k)^k$ leaves in \mathcal{T}' at depth k , and each such leaf t could have the property that $E(t)$ forms a clique on $\ell = O(\sqrt{k})$ coordinates when viewed as a graph. In this case, the fraction of $x \in X$ that could be consistent with the intersection $R_{i_1, j_1} \cap \dots \cap R_{i_\ell, j_\ell}$ is $(1/\sqrt{n})^\ell$. If we simply take a union bound over all $(2k)^k$ paths, the upper bound we would get on (1) is $(2k)^k \cdot (1/\sqrt{n})^\ell$ which, since ℓ is only $O(\sqrt{k})$, is larger than 1 unless $k^{\sqrt{k}}$ is at most polynomial in n , which requires that k is $o(\log^2 n)$. ┘



■ **Figure 1** A single iteration of Algorithm 2.

In light of Remark 5.7, we need to be more careful about how we bound (1). For $t \in V(\mathcal{T}')$, we define its *equality graph* $G(t)$ to be a graph with vertex set $[m]$ and edge set $E(t)$. We use this name because if x is consistent with t , then $x_i = x_j$ for all $(i, j) \in E(t)$. We say that an edge (i, j) denoting $x_i = x_j$ is *implied* by equality graph $G(t)$ if adding it completes a cycle in $G(t)$.

The key observation is the following.

▷ **Claim 5.8.** Let t be a node in \mathcal{T}' , and $G(t)$ its associated equality graph. If t has out-degree strictly larger than 1 in \mathcal{T}' , then none of the outgoing edges can be labelled by an (i, j) that is implied by $G(t)$.

Proof. We prove the contrapositive. Consider any x consistent with the path to t in \mathcal{T}' (meaning $x_i = x_j$ for all $(i, j) \in E(t)$). Let T be the triangle labelling node t and consider the y that we choose when processing T in Algorithm 2. If there is some outgoing edge from t labelled (i, j) that is implied by $G(t)$, then $R_{i,j} \cap T^y \neq \emptyset$ and hence $y_i = y_j$. Since we know that $x_i = x_j$ is already implied in T , the single rectangle $R_{i,j}$ covers all of T^y , so t will have out-degree 1 in \mathcal{T} . ◁

Claim 5.8 implies that every path in \mathcal{T}' goes through a node that branches at least $\sqrt{2k}$ times, since k edges must imply equality constraints on at least $\sqrt{2k}$ coordinates (the worst case is a clique). We simplify \mathcal{T}' further by collapsing nodes of \mathcal{T}' with out-degree 1. Then we are left with \mathcal{T}'' , which is a tree with every path of length between $\sqrt{2k}$ and k . Since the input distribution on X is uniform and none of the out-degree 1 edges of \mathcal{T}' involve new constraints by Claim 5.8, the probability that x is consistent with any leaf $t \in \mathcal{T}''$ is precisely $(1/\sqrt{n})^{\text{depth}_{\mathcal{T}''}(t)}$. Then we can bound (1) as follows:

$$\begin{aligned}
 & \sum_{\substack{t \in \mathcal{T}' \\ \text{depth}_{\mathcal{T}'}(t)=k}} |\{x \in X : x \text{ is consistent with } t\}| \\
 &= \sum_{t \in \text{Leaves}(\mathcal{T}'')} |\{x \in X : x \text{ is consistent with } t\}| \\
 &= \sum_{i=\sqrt{2k}}^k \sum_{\substack{t \in \text{Leaves}(\mathcal{T}'') \\ \text{depth}_{\mathcal{T}''}(t)=i}} |\{x \in X : x \text{ is consistent with } t\}|
 \end{aligned}$$

$$\begin{aligned}
&\leq \sum_{i=\sqrt{2k}}^k (2k)^i (1/\sqrt{n})^i && \text{since } \mathcal{T} \text{ and hence } \mathcal{T}'' \text{ has out-degree at most } 2k \\
&\leq |X| \cdot \sum_{i=\sqrt{2k}}^k (1/2)^i && \text{since } k = \sqrt{n}/4 \\
&\leq 2^{-\sqrt{2k}+1} |X|.
\end{aligned}$$

This implies that the number of $x \in X$ that get assigned to u by μ is at most $|X| \cdot 2^{-\sqrt{2k}+1} \leq |X| 2^{-n^{1/4}/\sqrt{2}+1}$ as required. The analogous bound holds for the number of y assigned to u by μ . ◀

We now have all we need to prove our main theorem.

Proof of Theorem 1.5. Let $k = \sqrt{n}/4$. By Lemma 5.5, μ given by Algorithm 1 maps at least $1/2$ of the elements of X or elements of Y to vertices of D . Choose the applicable X or Y . By Lemma 5.6, at most a $2^{-n^{1/4}/\sqrt{2}+1}$ fraction of these can map to a single vertex of D . Therefore D has at least $2^{n^{1/4}/\sqrt{2}-2}$ vertices. ◀

► **Remark 5.9.** Hrubeš and Pudlák [16] proved their lower bound using the interpolation method combined with a modified form of a general monotone switching lemma of Jukna [19]. The construction of Algorithm 2 has some flavor of the monotone switching lemma arguments, but the overall argument here seems fairly different. It would not be surprising that if one could modify this method to yield a similar lower bound to the one we give here, but exactly how to do that is not clear. In any case, we find the form of the bottleneck counting argument appealing.

6 Multipart Communication Lower Bounds for all $m > n$

► **Theorem 1.3.** *For all $m > n$ and $2 \leq k \leq \log n/4$, the randomized number-in-hand communication complexity of $\text{Coll}_{m,n^{1/k}}^k$ is $\Omega(n^{1/2-1/(2k)}/k)$.*

We note that this is essentially optimal in that a simple randomized protocol nearly matches our lower bound.

► **Proposition 6.1.** *For $m \geq n + n^{1-1/k}$ and $k \leq \log n/4$, there is a randomized k -party number-in-hand protocol for $\text{Coll}_{m,n^{1/k}}^k$ with complexity $O(n^{1/2-1/(2k)} \log n)$.*

Details of both proofs are in our full paper [3]. We sketch the proof of Theorem 1.3, which simplifies and generalizes ideas of Yang, Zhang, and Wang [32, 31] to the $k > 2$ -party setting for $\text{Coll}_{m,n^{1/k}}^k$. It is based on a method for decomposing the input space into *structured* rectangles on which a protocol's output is constant.

► **Definition 6.2.** *For a random variable X supported on Ω^J , we say it is $(1 - \delta)$ -dense if the min-entropy of all its projections is nearly as large as possible; that is, for all $I \subseteq J$, $H_\infty(X(I)) \geq (1 - \delta) |I| \log |\Omega|$, where $X(I)$ denotes X marginalized to the coordinates in I .*

For $\text{Coll}_{m,n^{1/k}}^k$ we have $\text{supp}(X_i) = \Omega^m$ for $\Omega = [n^{1/k}]$ and we will choose $\delta = \frac{1}{\log n}$.

► **Definition 6.3 (Structured rectangle).** *Given a rectangle $R = X_1 \times \dots \times X_k$, where each $X_i \subseteq \Omega^m$, we say that R is structured iff there exist $J_1, \dots, J_k \subseteq [m]$ and τ_1, \dots, τ_k where each $\tau_i : [m] \setminus J_i \rightarrow \Omega$ such that*

- X_i is $(1 - \delta)$ -dense on the coordinates in J_i and
- X_i is fixed to τ_i on the coordinates in $\overline{J_i} := [m] \setminus J_i$.

Using ideas from [31] we obtain:

► **Lemma 6.4.** *Let Π be a deterministic k -party number-in-hand protocol solving $\text{Coll}_{m,n^{1/k}}^k$. There is a partition of the inputs into structured rectangles, each of which yields a fixed output of Π , such that the probability distribution $\mathcal{R}^{\text{leaf}}$ on these structured rectangles induced by uniform inputs satisfies $\mathbf{E}_{R \sim \mathcal{R}^{\text{leaf}}} \sum_i |\overline{J_i}| \leq 2k \cdot CC(\Pi)$ where J_1, \dots, J_k are the sets associated with the structured property of R .*

Our proof considers rectangles in the support of $\mathcal{R}^{\text{leaf}}$ of three kinds: (1) Those where $\sum_i |\overline{J_i}|$ is larger than its expected value of $2k \cdot CC(\Pi)$, (2) those where it is not much larger and every element of the rectangle contains some fixed collision, and (3) the rest. The fraction of inputs in rectangles of type (1) is small by Markov's inequality and the fraction in rectangles of type (3) on which any fixed answer of Π can be correct, is small because they are structured. Analyzing the total measure of inputs in rectangles of type (2) relies on a subtle charging scheme that allocates the total number of such inputs to rectangles of all types using a per-rectangle fraction proportional to $\min(\sum_i \binom{|\overline{J_i}|}{2}, k^2 \cdot CC^2(\Pi)) / n^{1-1/k}$.

7 Discussion

We end by discussing some related problems and directions for future study. In particular, we highlight three possible directions.

- **Multiparty DAG-like communication lower bounds.** In general, it would be interesting to see if the $2^{\Omega(n^{1/4})}$ we give can be further improved, or if any nontrivial (size $2^{o(n)}$) cutting planes refutation of BPHP_n^{n+1} exists. One possible way to improve the lower bound could be to study a multiparty ($k > 2$) version of DAG-like protocols. Concretely, could one can prove DAG-like communication lower bounds for the k -player analog $\text{Search}_{\text{BPHP}_n^m}^k$? We have shown that generalizing to k players helps in the tree-like case, and perhaps this holds true in the DAG-like setting as well.
- **Stronger bounds for tree-like cutting-planes proof complexity of the weak bit pigeonhole principle** When $m = n + \Omega(n)$, we only obtain the weaker $2^{\Omega(\sqrt{n})}$ size lower bound for BPHP_n^m , rather than the $2^{n-o(1)}$ size lower bound we obtain for values of m closer to n . Can we extend the range of our strong lower bound to all values of $m > n$? To do so we must use a technique that does not use randomized communication complexity.
- Finally, we highlight that for m close to n , the loss of 2^k in the denominator of Theorem 1.1 could potentially be improved when $m \leq n + n^{1/k}$. Indeed, we do not suspect that it should be present at all, and we conjecture that $\text{Search}_{\text{BPHP}_n^m}^k$ should remain hard for k all the way up to $\log n$. However, it seems unlikely that any reduction from k -party disjointness would be able to achieve this since an additional input bit per player seems essential in implementing the reduction.

References

- 1 Albert Atserias, Moritz Müller, and Sergi Oliva. Lower bounds for DNF-refutations of a relativized weak pigeonhole principle. *J. Symb. Log.*, 80(2):450–476, 2015. doi:10.1017/JSL.2014.56.
- 2 Paul Beame, Toniann Pitassi, and Nathan Segerlind. Lower bounds for Lovász–Schrijver systems and beyond follow from multiparty communication complexity. *SIAM J. Comput.*, 37(3):845–869, 2007. doi:10.1137/060654645.

- 3 Paul Beame and Michael Whitmeyer. Multiparty communication complexity of collision-finding and cutting planes proofs of concise pigeonhole principles. *Electron. Colloquium Comput. Complex.*, TR25-057, 2025. URL: <https://eccc.weizmann.ac.il/report/2025/057>.
- 4 Maria Luisa Bonet, Toniann Pitassi, and Ran Raz. Lower bounds for cutting planes proofs with small coefficients. In *Proceedings of the Twenty-Seventh Annual ACM Symposium on Theory of Computing (STOC 1995)*, pages 575–584, Las Vegas, Nevada, 1995. ACM. doi:10.1145/225058.225275.
- 5 Mark Braverman and Rotem Oshman. The communication complexity of number-in-hand set disjointness with no promise. *Electron. Colloquium Comput. Complex.*, TR15-002, 2015. URL: <https://eccc.weizmann.ac.il/report/2015/002>.
- 6 Samuel R. Buss and György Turán. Resolution proofs of generalized pigeonhole principles. *Theor. Comput. Sci.*, 62(3):311–317, 1988. doi:10.1016/0304-3975(88)90072-2.
- 7 Arkadev Chattopadhyay and Toniann Pitassi. The story of set disjointness. *SIGACT News*, 41(3):59–85, 2010. doi:10.1145/1855118.1855133.
- 8 Stefan S. Dantchev, Nicola Galesi, Abdul Ghani, and Barnaby Martin. Proof complexity and the binary encoding of combinatorial principles. *SIAM J. Comput.*, 53(3):764–802, 2024. doi:10.1137/20M134784X.
- 9 Stephan Gocht, Jakob Nordström, and Amir Yehudayoff. On division versus saturation in pseudo-boolean solving. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019*, pages 1711–1718. ijcai.org, August 2019. doi:10.24963/IJCAI.2019/237.
- 10 Mika Göös and Siddhartha Jain. Communication complexity of collision. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2022*, volume 245 of *LIPICs*, pages 19:1–19:9. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICs.APPROX/RANDOM.2022.19.
- 11 Dima Grigoriev, Edward A. Hirsch, and Dmitrii V. Pasechnik. Complexity of semi-algebraic proofs. In *STACS 2002, 19th Annual Symposium on Theoretical Aspects of Computer Science, Proceedings*, volume 2285 of *Lecture Notes in Computer Science*, pages 419–430. Springer, 2002. doi:10.1007/3-540-45841-7_34.
- 12 Armin Haken. The intractability of resolution. *Theor. Comput. Sci.*, 39:297–308, 1985. doi:10.1016/0304-3975(85)90144-6.
- 13 Armin Haken and Stephen A. Cook. An exponential lower bound for the size of monotone real circuits. *J. Comput. Syst. Sci.*, 58(2):326–335, 1999. doi:10.1006/JCSS.1998.1617.
- 14 Johan Håstad. On small-depth Frege proofs for PHP. In *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023*, pages 37–49. IEEE, November 2023. doi:10.1109/FOCS57990.2023.00010.
- 15 Pavel Hrubes and Pavel Pudlák. A note on monotone real circuits. *Electron. Colloquium Comput. Complex.*, TR17-048, 2017. URL: <https://eccc.weizmann.ac.il/report/2017/048>.
- 16 Pavel Hrubes and Pavel Pudlák. Random formulas, monotone circuits, and interpolation. *Electron. Colloquium Comput. Complex.*, TR17-042, 2017. URL: <https://eccc.weizmann.ac.il/report/2017/042>.
- 17 Russell Impagliazzo, Toniann Pitassi, and Alasdair Urquhart. Upper and lower bounds for tree-like cutting planes proofs. In *Proceedings of the Ninth Annual Symposium on Logic in Computer Science (LICS '94)*, pages 220–228, Paris, France, 1994. IEEE Computer Society. doi:10.1109/LICS.1994.316069.
- 18 Dmitry Itsykson and Artur Riazanov. Proof complexity of natural formulas via communication arguments. In *Proceedings of the 36th Computational Complexity Conference, CCC 2021*, volume 200 of *LIPICs*, pages 3:1–3:34. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.CCC.2021.3.
- 19 Stasys Jukna. *Boolean Function Complexity - Advances and Frontiers*, volume 27 of *Algorithms and combinatorics*. Springer, 2012. doi:10.1007/978-3-642-24508-4.

- 20 Jan Krajíček, Pavel Pudlák, and Alan R. Woods. An exponential lower bound to the size of bounded depth Frege proofs of the pigeonhole principle. *Random Struct. Algorithms*, 7(1):15–40, 1995. doi:10.1002/RSA.3240070103.
- 21 Saburo Muroga, Iwao Toda, and Satoru Takasu. Theory of majority decision elements. *Journal of the Franklin Institute*, 271(5):376–418, 1961.
- 22 Noam Nisan. The communication complexity of threshold gates. *Combinatorics, Paul Erdos is Eighty*, 1(301-315):6, 1993.
- 23 Toniann Pitassi, Paul Beame, and Russell Impagliazzo. Exponential lower bounds for the pigeonhole principle. *Comput. Complex.*, 3:97–140, 1993. doi:10.1007/BF01200117.
- 24 Pavel Pudlák. Lower bounds for resolution and cutting plane proofs and monotone computations. *J. Symb. Log.*, 62(3):981–998, 1997. doi:10.2307/2275583.
- 25 Anup Rao and Amir Yehudayoff. *Communication Complexity: and Applications*. Cambridge University Press, 2020. doi:10.1017/9781108671644.
- 26 Ran Raz and Avi Wigderson. Monotone circuits for matching require linear depth. *J. ACM*, 39(3):736–744, 1992. doi:10.1145/146637.146684.
- 27 Hanif D. Sherali and Warren P. Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM J. Discret. Math.*, 3(3):411–430, 1990. doi:10.1137/0403036.
- 28 Dmitry Sokolov. Dag-like communication and its applications. In *Computer Science - Theory and Applications - 12th International Computer Science Symposium in Russia, CSR 2017, Proceedings*, volume 10304 of *Lecture Notes in Computer Science*, pages 294–307. Springer, 2017. doi:10.1007/978-3-319-58747-9_26.
- 29 Dmitry Sokolov. Random $(\log n)$ -CNF are hard for cutting planes (again). In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, Vancouver, BC, Canada, June 24-28, 2024*, pages 2008–2015. ACM, 2024. doi:10.1145/3618260.3649636.
- 30 Emanuele Viola. The communication complexity of addition. *Combinatorica*, 35(6):703–747, 2015. doi:10.1007/S00493-014-3078-3.
- 31 Shuo Wang, Guangxu Yang, and Jiapeng Zhang. Communication complexity of set-intersection problems and its applications. *Electron. Colloquium Comput. Complex.*, TR23-164, 2023. URL: <https://eccc.weizmann.ac.il/report/2023/164>.
- 32 Guangxu Yang and Jiapeng Zhang. Communication lower bounds for collision problems via density increment arguments. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024*, pages 630–639. ACM, June 2024. doi:10.1145/3618260.3649607.