



Extension Complexity of Independent Set Polytopes

Mika Göös¹ Rahul Jain² Thomas Watson³

¹*Department of Computer Science, Harvard University*

²*Centre for Quantum Technologies and Department of Computer Science,
National University of Singapore and MajuLab, UMI 3654, Singapore*

³*Department of Computer Science, University of Memphis*

December 7, 2017

Abstract

We exhibit an n -node graph whose independent set polytope requires extended formulations of size exponential in $\Omega(n/\log n)$. Previously, no explicit examples of n -dimensional 0/1-polytopes were known with extension complexity larger than exponential in $\Theta(\sqrt{n})$. Our construction is inspired by a relatively little-known connection between extended formulations and (monotone) circuit depth.

An extended abstract of this work was published as [GJW16].

Contents

1	Introduction	1
1.1	Our result	1
1.2	Our approach	2
1.3	Background	2
2	KW/EF Connection	3
2.1	Definitions	3
2.2	The connection	4
2.3	Example: Matchings	4
2.4	Minterms and maxterms	5
3	Tseitin Problem	5
3.1	Query version	5
3.2	Communication version	6
3.3	Statement of result	6
4	Reductions	6
4.1	Definition: Monotone CSP-SAT	6
4.2	From Tseitin to CSP-SAT	7
4.3	From CSP-SAT to independent sets	7
4.4	Proof of Theorem 1	8
5	Our Gadget	8
5.1	Flips and windows	9
5.2	Checking the existence of flips	10
6	Communication Lower Bound	11
6.1	High-level intuition	11
6.2	Preliminaries	12
6.3	Proof outline	13
6.4	Deriving the contradiction	14
6.5	Roadmap for the rest of the proof	16
6.6	Proof of the 1-vs-3 Lemma	17
6.7	Proof of the 3-vs-7 Lemma	18
6.8	Proof of the Homogeneity Lemma	20
7	Query Lower Bound	23
7.1	Conical juntas	23
7.2	A linear lower bound	23
7.3	Limitations of existing simulation theorems	25
	References	26

1 Introduction

A polytope $P \subseteq \mathbb{R}^n$ with many facets can sometimes admit a concise description as the projection of a higher dimensional polytope $E \subseteq \mathbb{R}^e$ with few facets. This phenomenon is studied in the theory of “extended formulations”. The *extension complexity* $\text{xc}(P)$ of a polytope P is defined as the minimum number of facets in any polytope E (called an *extended formulation* for P) such that

$$P = \{x \in \mathbb{R}^n : (x, y) \in E \text{ for some } y\}.$$

Extended formulations are useful for solving combinatorial optimization problems: instead of optimizing a linear function over P , we can optimize it over E —this may be more efficient since the runtime of LP solvers often depends on the number of facets.

Fiorini et al. [FMP⁺15] were the first to show (using methods from communication complexity [KN97, Juk12]) exponential extension complexity lower bounds for many explicit polytopes of relevance to combinatorial optimization, thereby solving an old challenge set by Yannakakis [Yan91]. For example, their results include a $2^{\Omega(m)}$ lower bound for the $\binom{m}{2}$ -dimensional *correlation/cut polytope*. In another breakthrough, Rothvoß [Rot14] proved a much-conjectured $2^{\Omega(m)}$ lower bound for the $\binom{m}{2}$ -dimensional *matching polytope*. By now, many accessible introductions to extended formulations are available; e.g., Roughgarden [Rou15, §5], Kaibel [Kai11], Conforti et al. [CCZ10] or their textbook [CCZ14, §4.10].

\sqrt{n} -frontier. Both of the results quoted above—while optimal for their respective polytopes—seem to get “stuck” at being exponential in the square root of their dimension. (For those graph problems, the dimension is the number of possible edges, but the lower bound is exponential in the number of nodes.) In fact, no explicit n -dimensional 0/1-polytope (convex hull of a subset of $\{0, 1\}^n$) was known with extension complexity asymptotically larger than $2^{\Theta(\sqrt{n})}$. The problem of proving higher lower bounds was explicitly posed in, e.g., [Wel16, Problem 7]. In comparison, Rothvoß [Rot12] showed via a counting argument that almost all n -dimensional 0/1-polytopes have extension complexity $2^{\Omega(n)}$.

1.1 Our result

Our main result is to construct an explicit 0/1-polytope in \mathbb{R}^n of near-maximal extension complexity $2^{\Omega(n/\log n)}$. Moreover, the polytope can be taken to be the *independent set polytope* P_G of an n -node graph G , i.e., the convex hull of (the indicator vectors of) the independent sets of G . Previously, a lower bound of $2^{\Omega(\sqrt{n})}$ was known for independent set polytopes [FMP⁺15].

Theorem 1. *There is an (explicit) family of n -node graphs G with $\text{xc}(P_G) \geq 2^{\Omega(n/\log n)}$.*

In fact, our graph family has bounded degree. Hence, using known reductions, we get as a corollary quantitative improvements—from $2^{\Omega(\sqrt{n})}$ to $2^{\Omega(n/\log n)}$ —for the extension complexity of, for instance, *3SAT* and *knapsack polytopes*; see [AT14, PV13] for details.

We strongly conjecture that our graph family actually satisfies $\text{xc}(P_G) \geq 2^{\Omega(n)}$, i.e., that the $\log n$ factor in the exponent is an artifact of our proof technique. We give concrete evidence for this by proving an optimal bound for a certain *query complexity* analogue of **Theorem 1**. In particular, the conjectured bound $\text{xc}(P_G) \geq 2^{\Omega(n)}$ would follow from quantitative improvements to the known communication-to-query simulation theorems ([GLM⁺16] in particular). Incidentally, this also answers a question of Lovász, Naor, Newman, and Wigderson [LNNW95]: we obtain a maximal $\Omega(n)$ lower bound on the randomized query complexity of a search problem with constant certificate complexity.

1.2 Our approach

Curiously enough, an analogous \sqrt{n} -frontier existed in the seemingly unrelated field of *monotone circuits*: Raz and Wigderson [RW92] proved an $\Omega(m)$ lower bound for the depth of any monotone circuit computing the *matching function* on $\binom{m}{2}$ input bits. This remained the largest monotone depth bound for an explicit function until the recent work of Göös and Pitassi [GP14], who exhibited an n -bit function with monotone depth $\Omega(n/\log n)$. In short, our idea is to prove an extension complexity analogue of this latter result.

The conceptual inspiration for our construction is a relatively little-known connection between Karchmer–Wigderson games [KW88] (which characterize circuit depth) and extended formulations. This “KW/EF connection” (see Section 2 for details) was pointed out by Hrubeš [Hru12] as a nonnegative analogue of a classic rank-based method of Razborov [Raz90]. In this work, we focus only on the monotone setting. For any monotone $f: \{0, 1\}^n \rightarrow \{0, 1\}$ we can study the convex hull of its 1-inputs, namely, the polytope

$$F := \text{conv } f^{-1}(1).$$

The upshot of the KW/EF connection is that extension complexity lower bounds for F follow from a certain type of *strengthening* of monotone depth lower bounds for f . For example, using this connection, it turns out that Rothvoß’s result [Rot14] implies the result of Raz and Wigderson [RW92] in a simple black-box fashion (Section 2.3).

Our main technical result is to strengthen the existing monotone depth lower bound from [GP14] into a lower bound for the associated polytope (though we employ substantially different techniques than were used in that paper). The key communication search problem studied in [GP14] is a communication version of the well-known *Tseitin* problem (see Section 3 for definitions), which has especially deep roots in proof complexity (e.g., [Juk12, §18.7]) and has also been studied in query complexity [LNNW95]. We use information complexity techniques to prove the required $\Omega(n/\log n)$ communication lower bound for the relevant variant of the Tseitin problem; information theoretic tools have been used in extension complexity several times [BM13, BP16, BP15]. One relevant work is Huynh and Nordström [HN12] (predecessor to [GP14]), whose information complexity arguments we extend in this work.

Instead of using information complexity, an alternative seemingly promising approach would be to “lift” a strong enough query complexity lower bound for Tseitin into communication complexity. Unfortunately, this approach runs into problems due to limitations in existing communication-to-query simulation theorems; we discuss this in Section 7.

Theorem 1 follows by reductions from the result for the Tseitin problem (Section 4). Indeed, it was known that the Tseitin problem reduces to the monotone KW game associated with an $f: \{0, 1\}^{O(n)} \rightarrow \{0, 1\}$ that encodes (in a monotone fashion) a certain CSP satisfiability problem. This gives us an extension complexity lower bound for the (explicit) polytope $F := \text{conv } f^{-1}(1)$. As a final step, we give a reduction from F to an independent set polytope.

1.3 Background

Let M be a nonnegative matrix. The *nonnegative rank* of M , denoted $\text{rk}^+(M)$, is the minimum r such that M can be decomposed as a sum $\sum_{i \in [r]} R_i$ where each R_i is a rank-1 nonnegative matrix.

Randomized protocols. Faenza et al. [FFGT14] observed that a nonnegative rank decomposition can be naturally interpreted as a type of randomized protocol that computes the matrix M “in expectation”. We phrase this connection precisely as follows: $\log \text{rk}^+(M) + \Theta(1)$ is the minimum communication cost of a private-coin protocol Π whose acceptance probability on each input (x, y)

satisfies $\mathbb{P}[\Pi(x, y) \text{ accepts}] = \alpha \cdot M_{x,y}$ where $\alpha > 0$ is an absolute constant of proportionality (depending on Π but not on x, y). All communication protocols in this paper are private-coin.

Slack matrices. The extension complexity of a polytope $P = \{x \in \mathbb{R}^n : Ax \geq b\}$ can be characterized in terms of the nonnegative rank of the *slack matrix* $M = M(P)$ associated with P . (We always assume P has positive dimension.) The entries of M are indexed by (v, i) where $v \in P$ is a vertex of P and i refers to the i -th facet-defining inequality $A_i x \geq b_i$ for P . We define $M_{v,i} := A_i v - b_i \geq 0$ as the distance (*slack*) of the i -th inequality from being tight for vertex v . Yannakakis [Yan91] showed that $\text{xc}(P) = \text{rk}^+(M(P))$.

A convenient fact for proving lower bounds on $\text{rk}^+(M)$ is that the nonnegative rank is unaffected by the addition of columns to M that each record the slack between vertices of P and some valid (but not necessarily facet-defining) inequality for P . For notation, let $P \subseteq Q$ be a polyhedral pair (P is a polytope; Q is a polytope or unbounded polyhedron). We define $M(P; Q)$ as the slack matrix whose rows correspond to vertices of P and columns correspond to the facets of Q (hence $M(P; P) = M(P)$). We have $\text{rk}^+(M(P)) \geq \text{rk}^+(M(P) \cup M(P; Q)) - 1 \geq \text{rk}^+(M(P; Q)) - 1$ where “ \cup ” denotes concatenation of columns.¹ We summarize all the above in the following.

Fact 2. *For all polyhedral pairs $P \subseteq Q$, we have $\text{xc}(P) = \text{rk}^+(M(P)) \geq \text{rk}^+(M(P; Q)) - 1$.*

2 KW/EF Connection

We now describe the connection showing that EF lower bounds follow from a certain type of strengthening of lower bounds for monotone KW games (and similarly, lower bounds for monotone KW games follow from certain strong enough EF lower bounds). This is not directly used in the proof of [Theorem 1](#), but it serves as inspiration by suggesting the approach we use in the proof.

2.1 Definitions

Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$ be a monotone function. We define $\text{KW}^+(f)$ as the deterministic communication complexity of the following *monotone KW game* associated with f .

KW⁺-game

Input: Alice gets $x \in f^{-1}(1)$, and Bob gets $y \in f^{-1}(0)$.

Output: An index $i \in [n]$ such that $x_i = 1$ and $y_i = 0$.

We often think of x and y as subsets of $[n]$. In this language, a feasible solution for the KW^+ -game is an $i \in x \cap \bar{y}$ where $\bar{y} := [n] \setminus y$. Given a monotone f , we denote by $F := \text{conv } f^{-1}(1)$ the associated polytope. We can express the fact that any pair $(x, y) \in f^{-1}(1) \times f^{-1}(0)$ admits at least one witness $i \in x \cap \bar{y}$ via the following system of linear inequalities

$$\forall y \in f^{-1}(0) : \sum_{i: y_i=0} x_i \geq 1. \tag{1}$$

Since (1) is valid for all the vertices $x \in F$, it is valid for the whole polytope F . Define $F_{\text{KW}} \supseteq F$ as the polyhedron whose facets are determined by the inequalities (1), as indexed by 0-inputs y . The

¹Specifically, Farkas’s Lemma implies that the slack of any valid inequality for P can be written as a nonnegative linear combination of the slacks of the facet-defining inequalities for P , plus a nonnegative constant [Zie95, Proposition 1.9]. Thus if we take $M(P) \cup M(P; Q)$ and subtract (possibly different) nonnegative constants from each of the “new” columns $M(P; Q)$, we get a matrix each of whose columns is a nonnegative linear combination of the “original” columns $M(P)$ and hence has the same nonnegative rank as $M(P)$. Since we subtracted a nonnegative rank-1 matrix, we find that $\text{rk}^+(M(P) \cup M(P; Q)) \leq \text{rk}^+(M(P)) + 1$.

(x, y) -th entry in the slack matrix $M(F; F_{\text{KW}})$ is then $\sum_{i: y_i=0} x_i - 1$. In words, this quantity counts the number of witnesses in the KW^+ -game on input (x, y) minus one.

More generally, let $S \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{Q}$ be *any* communication search problem (not necessarily a KW^+ -game, even though any S can be reformulated as such [Gál01, Lemma 2.3]). That is, \mathcal{Q} is some set of solutions/witnesses, and on input $x \in \mathcal{X}$ to Alice and $y \in \mathcal{Y}$ to Bob, the task is to find a solution $q \in S(x, y) := \{q \in \mathcal{Q} : (x, y, q) \in S\}$ (we assume $S(x, y) \neq \emptyset$ for all (x, y)). We associate with S the following natural “number of witnesses minus one” communication game.

($\#\exists-1$)-game

Input: Alice gets $x \in \mathcal{X}$, and Bob gets $y \in \mathcal{Y}$.

Output: Accept with probability proportional to $|S(x, y)| - 1$.

The communication complexity of this game is simply $\log \text{rk}^+(M^S) + \Theta(1)$ where $M_{x,y}^S := |S(x, y)| - 1$.

2.2 The connection

What Hrubeš [Hru12, Proposition 4] observed was that an efficient protocol for a search problem S implies an efficient protocol for the associated $(\#\exists-1)$ -game. In particular, for KW^+ -games,

$$\log \text{rk}^+(M(F; F_{\text{KW}})) \leq O(\text{KW}^+(f)). \quad (\text{KW/EF})$$

The private-coin protocol for $M(F; F_{\text{KW}})$ computes as follows. On input $(x, y) \in f^{-1}(1) \times f^{-1}(0)$ we first run the optimal deterministic protocol for the KW^+ -game for f to find a particular $i \in [n]$ witnessing $x_i = 1$ and $y_i = 0$. Then, Alice uses her private coins to sample a $j \in [n] \setminus \{i\}$ uniformly at random, and sends this j to Bob. Finally, the two players check whether $x_j = 1$ and $y_j = 0$ accepting iff this is the case. The acceptance probability of this protocol is proportional to the number of witnesses minus one, and the protocol has cost $\text{KW}^+(f) + \log n + O(1) \leq O(\text{KW}^+(f))$ (where we assume w.l.o.g. that f depends on all of its input bits so that $\text{KW}^+(f) \geq \log n$).

2.3 Example: Matchings

Rothvoß vs. Raz–Wigderson. Consider the monotone function $f: \{0, 1\}^{\binom{m}{2}} \rightarrow \{0, 1\}$ that outputs 1 iff the input, interpreted as a graph on m nodes (m even), contains a perfect matching. Then $F := \text{conv } f^{-1}(1)$ is the convex hull of all graphs that contain a perfect matching; in particular, the perfect matching polytope is a face of F . The inequalities (1) for f happen to include the so-called “odd set” inequalities, which were exploited by Rothvoß [Rot14] in showing that $\log \text{rk}^+(M(F; F_{\text{KW}})) \geq \Omega(m)$. Applying the (KW/EF) connection to Rothvoß’s lower bound implies in a black-box fashion that $\text{KW}^+(f) \geq \Omega(m)$, which is the result of Raz and Wigderson [RW92].

Converse to (KW/EF)? It is interesting to compare the above with the case of *bipartite* perfect matchings. Consider a monotone $f: \{0, 1\}^{m \times m} \rightarrow \{0, 1\}$ that takes a bipartite graph as input and outputs 1 iff the graph contains a perfect matching. It is well-known that $F := \text{conv } f^{-1}(1)$ admits a polynomial-size extended formulation [Sch03, Theorem 18.1]. By contrast, the lower bound $\text{KW}^+(f) \geq \Omega(m)$ from [RW92] continues to hold even in the bipartite case. This example shows that the converse inequality to (KW/EF) does not hold in general. Hence, a lower bound for the $(\#\exists-1)$ -game can be a strictly stronger result than a similar lower bound for the KW^+ -game.

2.4 Minterms and maxterms

A *minterm* $x \in f^{-1}(1)$ is a minimal 1-input in the sense that flipping any 1-entry of x into a 0 will result in a 0-input. Analogously, a *maxterm* $y \in f^{-1}(0)$ is a maximal 0-input. It is a basic fact that solving the KW^+ -game for minterms/maxterms is enough to solve the search problem on any input: Say that Alice's input x is not a minterm. Then Alice can replace x with any minterm $x' \subseteq x$ and run the protocol on x' . A witness $i \in [n]$ for (x', y) works also for (x, y) . A similar fact holds for the $(\#\exists-1)$ -game: we claim that the nonnegative rank does not change by much when restricted to minterms/maxterms. Say that Alice's input x is not a minterm. Then Alice can write $x = x' \cup x''$ (disjoint union) where x' is a minterm. Then $|x \cap \bar{y}| - 1 = (|x' \cap \bar{y}| - 1) + |x'' \cap \bar{y}|$ where the first term is the $(\#\exists-1)$ -game for (x', y) and the second term has nonnegative rank at most n (consider the protocol that samples a uniformly random $i \in [n]$ and accepts iff $i \in x'' \cap \bar{y}$). A similar argument works if Bob does not have a maxterm.

3 Tseitin Problem

3.1 Query version

Fix a connected node-labeled graph $G = (V, E, \ell)$ where the node-labeling $\ell \in \mathbb{Z}_2^V$ has *odd weight*, i.e., $\sum_{v \in V} \ell(v) = 1$ where the addition is modulo 2. For any edge-labeling $z \in \mathbb{Z}_2^E$ and a node $v \in V$ we write concisely $z(v) := \sum_{e \ni v} z(e)$ for the mod-2 sum of the edge-labels adjacent to v .

Tseitin problem: TSE_G

Input: Labeling $z \in \mathbb{Z}_2^E$ of the edges.
Output: A node $v \in V$ containing a *parity violation* $z(v) \neq \ell(v)$.

As a sanity check, we note that on each input z there must exist at least one node with a parity violation. This follows from the fact that, since each edge has two endpoints, the sum $\sum_v z(v)$ is even, whereas we assumed that the sum $\sum_v \ell(v)$ is odd.

Basic properties. The above argument implies more generally that the set of violations $\text{viol}(z) := \{v \in V : z(v) \neq \ell(v)\}$ is always of odd size. Conversely, for any odd-size set $S \subseteq V$ we can design an input z such that $\text{viol}(z) = S$. To see this, it is useful to understand what happens when we *flip a path* in an input z . Formally, suppose $p \in \mathbb{Z}_2^E$ is (an indicator vector of) a path. Define z^p as z with bits on the path p flipped (note that $z^p = z + p \in \mathbb{Z}_2^E$; however, the notation z^p will be more convenient later). Flipping p has the effect of flipping whether each endpoint of p is a violation. More precisely, the violated nodes in z^p are related to those in z as follows: (i) if both endpoints of p are violated in z then the flip causes that pair of violations to disappear; (ii) if neither endpoint of p is violated in z , then the flip introduces a pair of new violations; (iii) if precisely one endpoint of p was violated in z , then the flip moves a violation from one endpoint of p to the other. By applying (i)–(iii) repeatedly in a connected graph G , we can design an input z where $\text{viol}(z)$ equals any prescribed odd-size set S .

If z and z' have the same set of violations, $\text{viol}(z) = \text{viol}(z')$, then their difference $q := z - z' \in \mathbb{Z}_2^E$ satisfies $q(v) = 0$ for all $v \in V$. That is, q is an *eulerian* subgraph of G (for us, q need not be connected). On the other hand, for any eulerian graph q , the inputs z and z^q have the same violations. Consequently, to generate a random input with the same set of violations as some fixed

z , we need only pick a random eulerian graph q and output z^q . (Eulerian graphs form a subspace of \mathbb{Z}_2^E , sometimes called the *cycle space* of G .)

3.2 Communication version

The communication version of the Tseitin problem is obtained by composing (or *lifting*) TSE_G with a constant-size two-party gadget $g: \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$. In the lifted problem $\text{TSE}_G \circ g^n$, where $n := |E|$, Alice gets $x \in \mathcal{X}^n$ as input, Bob gets $y \in \mathcal{Y}^n$ as input, and their goal is to find a node $v \in V$ that is violated for

$$z := g^n(x, y) = (g(x_1, y_1), \dots, g(x_n, y_n)).$$

We define our gadget precisely in [Section 5](#). For now—in particular, for the reductions presented in the next section—the only important property of our gadget is that $|\mathcal{X}|, |\mathcal{Y}| \leq O(1)$.

3.3 Statement of result

We prove that there is a family of bounded-degree graphs G such that the $(\#\exists-1)$ -game associated with $\text{TSE}_G \circ g^n$ requires $\Omega(n/\log n)$ bits of communication. We prove our lower bound assuming only that $G = (V, E)$ is well-connected enough as captured by the following definition (also used in [\[GP14\]](#)). A graph G is *k-routable* iff there is a set of $2k + 1$ nodes $T \subseteq V$ called *terminals* such that for any *pairing* $\mathcal{P} := \{\{s_i, t_i\} : i \in [\kappa]\}$ (set of pairwise disjoint pairs) of 2κ terminals ($\kappa \leq k$), there exist κ edge-disjoint paths (called *canonical paths* for \mathcal{P}) such that the i -th path connects s_i to t_i . Furthermore, we tacitly equip G with an arbitrary odd-weight node-labeling.

Theorem 3. *There is a constant-size g such that for every k -routable graph G with n edges, the $(\#\exists-1)$ -game for $\text{TSE}_G \circ g^n$ requires $\Omega(k)$ bits of communication.*

If we choose G to be a sufficiently strong expander graph, we may take $k = \Theta(n/\log n)$ as shown by Frieze et al. [\[FZ00, Fri01\]](#). Alternative explicit constructions with $k = \Theta(n/\log n)$ exist based on bounded-degree “butterfly” graphs; see [\[Nor15, §5\]](#) for an exposition.

Corollary 4. *There is a constant-size g and an explicit bounded-degree graph G with n edges such that the $(\#\exists-1)$ -game for $\text{TSE}_G \circ g^n$ requires $\Omega(n/\log n)$ bits of communication.*

As a bonus, we also prove that the *query complexity* of the $(\#\exists-1)$ -game for TSE_G is $\Omega(n)$ on any expander G (see [Section 7](#)).

4 Reductions

The goal of this section is to show, via reductions, that a lower bound on the $(\#\exists-1)$ -game for $\text{TSE}_G \circ g^n$ (where $G = (V, E)$ is of bounded degree and $n := |E|$) translates directly into a lower bound on the extension complexity of P_K for an $O(n)$ -node bounded-degree graph K .

4.1 Definition: Monotone CSP-SAT

We start by describing a way of representing constraint satisfaction problems (CSP) as a monotone function; this was introduced in [\[GP14\]](#) and further studied by Oliveira [\[Oli15, Chapter 3\]](#). The function is defined relative to some finite alphabet Σ and a fixed constraint topology determined by a bipartite graph $H := (L \cup R, E)$. The left nodes L are thought of as *variables* (taking values in Σ)

and the right nodes R correspond to *constraints*. For a constraint $c \in R$, let $\text{var}(c) \subseteq L$ denote the variables involved in c (i.e., neighbors of c in H). Let d denote the maximum degree of a node in R . The function $\text{SAT} = \text{SAT}_{\Sigma, H}: \{0, 1\}^m \rightarrow \{0, 1\}$, where $m := \sum_{c \in R} |\Sigma^{\text{var}(c)}| \leq |R| \cdot |\Sigma|^d$, is now defined as follows. An input $x \in \{0, 1\}^m$ defines a CSP instance by specifying, for each $c \in R$, a truth table $\Sigma^{\text{var}(c)} \rightarrow \{0, 1\}$ that records which assignments to the variables $\text{var}(c)$ satisfy c . In short, x is simply a concatenation of truth tables. Then $\text{SAT}(x) := 1$ iff there is some global assignment $L \rightarrow \Sigma$ that satisfies all the constraints as specified by x . This is monotone: if we flip any 0 into a 1 in the truth table of a constraint, we are only making the constraint easier to satisfy.

4.2 From Tseitin to CSP-SAT

For completeness, we present the reduction (due to [GP14, §5.1]) from the search problem $\text{TSE}_G \circ g^n$ to the KW^+ -game for $\text{SAT} = \text{SAT}_{\mathcal{X}, H}: \{0, 1\}^m \rightarrow \{0, 1\}$. Here the alphabet is \mathcal{X} and the bipartite graph H is defined on $E(G) \cup V(G)$ such that there is an edge $(e, v) \in E(H)$ iff $v \in e$. Note that $m \leq O(n)$ provided that $|\mathcal{X}| \leq O(1)$ and that G is of bounded degree.

On input (x, y) to $\text{TSE}_G \circ g^n$ the two players proceed as follows:

- Alice maps her $x \in \mathcal{X}^{E(G)}$ into a CSP whose sole satisfying assignment is x . Namely, for each constraint $v \in V(G)$, the truth table $\mathcal{X}^{\text{var}(v)} \rightarrow \{0, 1\}$ is all-0 except for a unique 1 in position $x|_{\text{var}(v)}$ (restriction of x to coordinates in $\text{var}(v)$).
- Bob maps his $y \in \mathcal{Y}^{E(G)}$ into an unsatisfiable CSP. Namely, for each constraint $v \in V(G)$, the truth table $t_v: \mathcal{X}^{\text{var}(v)} \rightarrow \{0, 1\}$ is given by $t_v(\hat{x}) := 1$ iff $(g(\hat{x}_e, y_e))_{e \in \text{var}(v)} \in \{0, 1\}^{\text{var}(v)}$ is a partial edge-labeling of G that does *not* create a parity violation on v .

Let us explain why Bob really produces a 0-input of SAT. Suppose for contradiction that there is an $\hat{x} \in \mathcal{X}^{E(G)}$ that satisfies all of Bob's constraints: $t_v(\hat{x}|_{\text{var}(v)}) = 1$ for all v . By definition, this means that $z := g^n(\hat{x}, y)$ is an input to TSE_G without any violated nodes—a contradiction.

This reduction is parsimonious: it maps witnesses to witnesses in 1-to-1 fashion (so the number of Tseitin violations equals the number of bit positions where Alice's CSP has a 1 and Bob's CSP has a 0). Indeed, a node v is violated for $\text{TSE}_G \circ g^n$ if and only if Alice's truth table for v has its unique 1 in a coordinate where Bob has a 0. In conclusion, the $(\#\exists-1)$ -game associated with (the KW^+ -game for) SAT is at least as hard as the $(\#\exists-1)$ -game for $\text{TSE}_G \circ g^n$.

4.3 From CSP-SAT to independent sets

As a final step, we start with $\text{SAT} = \text{SAT}_{\Sigma, H}: \{0, 1\}^m \rightarrow \{0, 1\}$ and construct an m -node “conflict” graph K such that a slack matrix of the independent set polytope P_K embeds the $(\#\exists-1)$ -game for SAT (restricted to minterms). Let $H := (L \cup R, E)$ (as above) and define $n := |R|$ (above we had $n = |L|$, but in our case $|L| = \Theta(|R|)$ anyway).

The conflict graph K is defined as follows (this is reminiscent of a reduction from [FGL⁺96]).

- The m nodes of K are in 1-to-1 correspondence with the input bits of SAT. That is, for each constraint $c \in R$ we have $|\Sigma^{\text{var}(c)}|$ many nodes in K labeled with assignments $\text{var}(c) \rightarrow \Sigma$.
- There is an edge between any two nodes whose assignments are *inconsistent* with one another. (Here $\phi_i: \text{var}(c_i) \rightarrow \Sigma$, $i \in \{1, 2\}$, are inconsistent iff there is some $e \in \text{var}(c_1) \cap \text{var}(c_2)$ such that $\phi_1(e) \neq \phi_2(e)$.) In particular, the truth table of each constraint becomes a clique by definition.

(It can be seen that K has bounded degree if H has bounded left- and right-degree, which it does after our reduction from Tseitin for a bounded-degree G .)

The key property of this construction is the following:

The minterms of SAT are precisely the (indicator vectors of) maximal independent sets of K .

Indeed, the minterms $x \in \text{SAT}^{-1}(1)$ correspond to CSPs with a unique satisfying assignment $\phi: L \rightarrow \Sigma$; there is a single 1-entry in each of the n truth tables (so that $|x| = n$) consistent with ϕ . Such an x , interpreted as a subset of nodes, is independent in K as it only contains nodes whose labels are consistent with ϕ . Conversely, because every independent set $x \subseteq V(K)$ can only contain pairwise consistently labeled nodes, x naturally defines a partial assignment $L' \rightarrow \Sigma$ for some $L' \subseteq L$. A maximal independent set x corresponds to picking a node from each of the n constraint cliques consistent with some total assignment $\phi: L \rightarrow \Sigma$. Hence x is a 1-input to SAT with unique satisfying assignment ϕ .

Our goal is now to exhibit a set of valid inequalities for the independent set polytope P_K whose associated slack matrix embeds the $(\#\exists-1)$ -game for SAT. Let $x \subseteq V(K)$ be an independent set and $y \in \text{SAT}^{-1}(0)$. We claim that the following inequalities (indexed by y) are valid:

$$|x \cap y| = \sum_{i: y_i=1} x_i \leq n - 1. \quad (2)$$

Clearly (2) holds whenever $|x| \leq n - 1$. Since it is impossible to have $|x| \geq n + 1$, assume that x is maximal: $|x| = n$. As argued above, x is a minterm of SAT. Hence (x, y) is a valid pair of inputs to the KW^+ -game, and so they admit a witness: $|x \cap \bar{y}| \geq 1$. Therefore $|x \cap y| = n - |x \cap \bar{y}| \leq n - 1$. This shows that (2) is valid. The slack matrix associated with inequalities (2) has entries

$$n - 1 - |x \cap y| = |x \cap \bar{y}| - 1,$$

for any minterm x and any $y \in \text{SAT}^{-1}(0)$. But this is just the $(\#\exists-1)$ -game for SAT with Alice's input restricted to minterms.

4.4 Proof of Theorem 1

Here we simply string the above reductions together. By Corollary 4 there is a constant-size g and a bounded-degree G with n edges such that the $(\#\exists-1)$ -game for $\text{TSE}_G \circ g^n$ requires $\Omega(n/\log n)$ bits of communication. By the reduction of Section 4.2 this implies an $\Omega(n/\log n)$ lower bound for the $(\#\exists-1)$ -game associated with (the KW^+ -game for) a monotone function $\text{SAT}: \{0, 1\}^{O(n)} \rightarrow \{0, 1\}$. As discussed in Section 2.4, the complexity of the $(\#\exists-1)$ -game for SAT is affected only by $\pm \log n$ when restricted to minterms. Thus the minterm-restricted $(\#\exists-1)$ -game for SAT still has complexity $\Omega(n/\log n)$. (Alternatively, one can note that the reduction from Tseitin to CSP-SAT produced only minterms.) Hence the nonnegative rank of the matrix for that game is $2^{\Omega(n/\log n)}$. By the reduction of Section 4.3 there is a bounded-degree $O(n)$ -node graph K and a system of valid inequalities (2) for the independent set polytope P_K such that the slack matrix $M(P_K; Q)$, where Q is the polyhedron with facets determined by (2), embeds the matrix for the minterm-restricted $(\#\exists-1)$ -game for SAT. Thus $\log \text{rk}^+(M(P_K; Q)) \geq \Omega(n/\log n)$. By Fact 2 we have $\log \text{xc}(P_K) = \log \text{rk}^+(M(P_K)) \geq \log(\text{rk}^+(M(P_K; Q)) - 1) \geq \Omega(n/\log n)$.

5 Our Gadget

We define our two-party gadget $g: \{0, 1\}^3 \times \{0, 1\}^3 \rightarrow \{0, 1\}$ as follows; see Figure 1:

$$g(x, y) := x_1 + y_1 + x_2 y_2 + x_3 y_3 \pmod{2}.$$




Figure 1: Three ways to view our gadget $g: \{0, 1\}^3 \times \{0, 1\}^3 \rightarrow \{0, 1\}$ by permuting rows and columns. The white and gray cells represent 0- and 1-inputs, respectively. For (a), the rows and columns are in lexicographic order of $x_3x_2x_1$ and $y_3y_2y_1$.


We note that the smaller gadget $x_1 + y_1 + x_2y_2 \pmod{2}$ was considered in [She11, GP14].

5.1 Flips and windows

The most basic property of g is that it admits *Alice/Bob-flips*:


- (1) *Alice-flips*: There is a row permutation $\pi_A: \{0, 1\}^3 \rightarrow \{0, 1\}^3$ that flips the output of the gadget: $g(\pi_A(x), y) = \neg g(x, y)$ for all x, y . Namely, Alice just flips the value of x_1 .
- (2) *Bob-flips*: There is a column permutation $\pi_B: \{0, 1\}^3 \rightarrow \{0, 1\}^3$ that flips the output of the gadget: $g(x, \pi_B(y)) = \neg g(x, y)$ for all x, y . Namely, Bob just flips the value of y_1 .

A more interesting feature of our gadget (which we prove later, and which $x_1 + y_1 + x_2y_2$ does not possess) is that g embeds—in an especially uniform manner—certain 2×4 and 4×2 submatrices which we call “stretched AND” and “stretched NAND”. For terminology, we define a z -window where $z \in \{0, 1\}$ as a z -monochromatic rectangle of size 2 in the domain of g , i.e., an all- z submatrix of either horizontal shape 1×2 or vertical shape 2×1 . Here is an illustration of *horizontally* stretched AND/NAND, which are composed of four horizontally shaped windows (for *vertical* stretch, the illustration should be transposed):




A stretched AND/NAND embedding may use any rows and any columns of g , so e.g., the window corresponding to the stretched $(0, 0)$ -input need not appear “upper left” in the embedding.

The key property is that each z -window w is embedded as the stretched $(1, 1)$ -input to a *unique* embedding of stretched AND (if $z = 1$) or NAND (if $z = 0$) inside g . That is, for each w we can find the following unique submatrix (illustrated again for horizontal shapes), where we denote by w^\leftarrow , w^\wedge , and w^\uparrow the $(1 - z)$ -windows corresponding to the stretched $(1, 0)$ -, $(0, 0)$ -, and $(0, 1)$ -inputs to the stretched AND/NAND.



if w is a 1-window



if w is a 0-window


This defines three maps (“directed flips”) $w \mapsto w^\leftarrow$, $w \mapsto w^\nwarrow$, $w \mapsto w^\uparrow$, which turn out to be shape-maintaining *bijections* between the set of z -windows and the set of $(1 - z)$ -windows. In particular, if w is a uniformly random z -window of g , then each of w^\leftarrow , w^\nwarrow , w^\uparrow is a uniformly random $(1 - z)$ -window.

5.2 Checking the existence of flips

The properties of g claimed above can be verified by directly inspecting the gadget (by hand). Luckily, this task can be eased by exploiting symmetries.

- (3) *Transitive symmetry*: The gadget admits a group of symmetries (permutations of its rows and columns leaving g invariant) which splits the domain of g into two orbits, $g^{-1}(1)$ and $g^{-1}(0)$. Specifically, there is a group $\mathcal{S} \subseteq \mathfrak{S}_8 \times \mathfrak{S}_8$ (here \mathfrak{S}_8 is the symmetric group on 8 elements) such that when $(\pi_1, \pi_2) \in \mathcal{S}$ acts on g , the output remains invariant: $g(\pi_1(x), \pi_2(y)) = g(x, y)$ for all x, y ; and moreover, \mathcal{S} is transitive in the sense that for any two 1-inputs $(x, y), (x', y') \in g^{-1}(1)$ (or 0-inputs) there is a symmetry $(\pi_1, \pi_2) \in \mathcal{S}$ such that $(\pi_1(x), \pi_2(y)) = (x', y')$.

To see that g really does have property (3), we visualize g as constructed from $\text{XOR}(x_1, x_2) := x_1 + x_2 \pmod{2}$ by applying the following “ \rightsquigarrow ” transformation twice:



It is easy to see that XOR has the properties (1)–(3). We argue that if M is a boolean matrix with the properties (1)–(3) and $M \rightsquigarrow M'$, then M' has the properties (1)–(3). Suppose the entries of M are indexed by (x, y) ; we use (xa, yb) to index the entries of M' where $a, b \in \{0, 1\}$ are bits. If π_A, π_B are the Alice/Bob-flips for M , then Alice/Bob-flips for M' are

$$\begin{aligned} xa &\mapsto \pi_A(x)a, \\ yb &\mapsto \pi_B(y)b. \end{aligned}$$

Suppose \mathcal{S} is the transitive symmetry group for M . Then the transitive symmetry group for M' is generated by the following symmetries (here $\pi_A^0(x) := x$ and $\pi_A^1(x) := \pi_A(x)$ and similarly for π_B^b):

$$\begin{aligned} \forall (\pi_1, \pi_2) \in \mathcal{S}: \quad & (xa, yb) \mapsto (\pi_1(x)a, \pi_2(y)b), \\ & (xa, yb) \mapsto (\pi_A^a(x)a, y(1 - b)), \\ & (xa, yb) \mapsto (x(1 - a), \pi_B^b(y)b). \end{aligned}$$

The first family of symmetries makes each quadrant of M' transitive, whereas the last two symmetries map entries between quadrants. In the second-to-last symmetry, Bob swaps the left and right halves

while Alice applies her flip to the bottom half. In the last symmetry, Alice swaps the top and bottom halves while Bob applies his flip to the right half. This shows that g satisfies (1)–(3).

Rather than checking that each z -window w appears as the stretched $(1, 1)$ -input to a unique embedding of stretched AND/NAND and that the directed flips are bijections, it is equivalent to check that for all $\ell \in \{(0, 0), (0, 1), (1, 0), (1, 1)\}$ each w appears as the stretched ℓ -input to a unique embedding of stretched AND/NAND in g . Let us check this assuming w is a 0-window of shape 1×2 (the other possibilities can be checked similarly). By transitive symmetry, we may assume that w is picked among the four 0's of the first row of [Figure 1\(c\)](#) (so $\binom{4}{2}$ choices for w). The key observation is that the four columns corresponding to these 0's define a submatrix of g (left half of (c)) that contains each even Hamming weight row once, and that the other four columns (right half of (c)) also contain each even Hamming weight row once. We consider the four cases for ℓ .

- $\ell = (0, 0)$: To see that w is the stretched $(0, 0)$ -input to a unique embedding of stretched AND, find the unique other row that has 0's in the same columns as w . The other two columns in the left half of (c) have 0's in the top row and 1's in the other row.
- $\ell = (0, 1)$: To see that w is the stretched $(0, 1)$ -input to a unique embedding of stretched AND, find the unique other row that has 1's in the same columns as w and 0's in the other two columns of the left half of (c). These other two columns have 0's in the top row.
- $\ell = (1, 0)$: To see that w is the stretched $(1, 0)$ -input to a unique embedding of stretched AND, find the unique other row that has 0's in the same columns as w , then find the unique pair of columns in the right half of (c) that has 0's in that other row. This pair of columns has 1's in the first row.
- $\ell = (1, 1)$: To see that w is the stretched $(1, 1)$ -input to a unique embedding of stretched NAND, find the unique other row that has 1's in the same columns as w and 0's in the other two columns of the left half of (c), then find the unique pair of columns in the right half of (c) that has 1's in that other row. This pair of columns has 1's in the first row.

6 Communication Lower Bound

In this section we prove [Theorem 3](#), where g is the gadget from [Section 5](#).

6.1 High-level intuition

The high-level reason for why the $(\#\exists-1)$ -game for Tseitin (or really for any sufficiently unstructured search problem) is hard is the same as for the $(\#\exists-1)$ -game for matching [[Rot14](#)]: A correct protocol Π dare not accept its input before it has found at least two witnesses, lest it risk accepting with positive probability an input with a unique witness (which would contradict correctness). However, in an input with i witnesses, there are $\binom{i}{2}$ pairs of witnesses for the protocol to find. Hence one expects the acceptance probability of Π (that communicates too few bits and never errs when $i = 1$) to grow at least *quadratically* with i rather than linearly as required by the $(\#\exists-1)$ -game.

Formalizing this quadratic increase in acceptance probability for protocols takes some technical work given the current tools available in communication complexity. However, the quadratic increase phenomenon for Tseitin is easier to formalize in the query complexity setting, which we do in [Section 7](#). The reader may want to have a look at that simpler proof first, even though the query proof is somewhat incomparable to our approach for protocols (which revolves around k -routability).

6.2 Preliminaries

Probability and information theory. We use standard notions from information theory: $\mathbb{H}(X)$ is Shannon entropy; $\mathbb{H}(X|Y) := \mathbb{E}_{y \sim Y} \mathbb{H}(X|Y=y)$ is conditional entropy; $\mathbb{I}(X;Y) := \mathbb{H}(X) - \mathbb{H}(X|Y) = \mathbb{H}(Y) - \mathbb{H}(Y|X)$ is mutual information; $\Delta(X,Y) := \max_E |\mathbb{P}[X \in E] - \mathbb{P}[Y \in E]|$ is statistical (total variation) distance. We use upper-case letters for random variables and corresponding lower-case letters for particular outcomes. Throughout the whole proof, all random choices are assumed to be uniform in their respective domains unless otherwise stated.

Inputs and transcripts. Let XY be random inputs to a private-coin protocol Π . We denote by $\Pi = \Pi(X,Y)$ the transcript of the protocol on input XY , and we let $|\Pi|$ be the maximum length of a transcript (i.e., the communication cost of Π). Note that the transcript Π depends on both XY and the private coins of the players. We let $\Pi^{\text{acc}} := (\Pi | \Pi \text{ accepts})$ denote the transcript conditioned on the protocol accepting. For each input $z \in \mathbb{Z}_2^n$ to the query problem TSE_G we can associate in a natural way a pair of random inputs XY to the communication problem $\text{TSE}_G \circ g^n$ that are *consistent with* z in the sense that $g^n(X,Y) = z$; namely, we let XY be uniformly distributed on

$$(g^n)^{-1}(z) = g^{-1}(z_1) \times \cdots \times g^{-1}(z_n).$$

We write $\Pi|z$ as a shorthand for $\Pi(X,Y)$ where XY are drawn at random from the above set. Note that if the protocol is correct then conditioning on acceptance does not change the distribution of such XY , since all inputs consistent with z have the same acceptance probability. Thus $\Pi^{\text{acc}}|z$ can be equivalently defined by sampling a uniformly random XY from $(g^n)^{-1}(z)$ and then sampling a transcript conditioned on acceptance.

Windows. As is often the case with information complexity arguments, we need to introduce a conditioning variable W whose purpose is to make X and Y conditionally independent. To this end, we employ windows (Section 5.1): we call a rectangle $w := w_1 \times \cdots \times w_n \subseteq (g^n)^{-1}(z)$ a (multi-gadget) *window of* z iff each w_i is a z_i -window in g (so $w_i \subseteq g^{-1}(z_i)$). Now, to generate XY as above, we first pick W uniformly at random among all the windows of z , and then, conditioned on an outcome $W = w$, we pick $XY \in w$ uniformly at random. In conclusion, XY is uniform on $(g^n)^{-1}(z)$ (since each row and column of g is balanced) and X and Y are conditionally independent given W . We write $\Pi|w := (\Pi(X,Y) | W = w)$ for short.

Alice-flips. Let (x,y) be an input consistent with $z := g^n(x,y)$ and let $B \subseteq [n]$ be any subset of coordinates of z . (B stands for “block” by analogy with the concept of block sensitivity from query complexity.) We denote by (x^B,y) the input obtained from (x,y) by letting Alice flip the outputs of all gadgets corresponding to coordinates in B , i.e., for every $i \in B$ Alice replaces her input x_i with $\pi_A(x_i)$ where π_A is the row permutation from Section 5.1. Hence (x^B,y) is an input consistent with z^B . We can also have Alice flip whole windows: $w^B := \{(x^B,y) : (x,y) \in w\}$. We henceforth refer to such Alice-flips as just “flips”. (We could equally well have Bob be the flipper throughout the whole proof, but we needed to make an arbitrary choice between the players.)

Smooth protocols. Recall that if z is an input to TSE_G and $B \subseteq E(G)$ is an eulerian graph, then z and z^B have the same set of violations. Consequently, any protocol Π for the $(\#\exists-1)$ -game must accept inputs (x,y) and (x^B,y) with the same probability. We note that we may assume w.l.o.g. that the transcript distribution of Π is not sensitive to flipping eulerian graphs: if w is a window and B an eulerian graph, then $\Pi|w$ and $\Pi|w^B$ have the same distribution. Indeed, if Π does

not satisfy this, then we may replace it by a new “smoothed” protocol Π' that computes as follows on input (x, y) : Alice uses her private coins to choose a uniformly random eulerian graph B and then the players run Π on input (x^B, y) . The fact that we may assume Π is smooth is critically used later in the proof.

6.3 Proof outline

Recall the setting from [Section 3.3](#): We have a k -routable graph with terminal nodes T . Let us assume for the sake of contradiction that Π is a private-coin protocol of cost $|\Pi| \leq o(k)$ that accepts each input (x, y) with probability $\alpha \cdot (|\text{viol}(z)| - 1)$ where $\alpha > 0$ is a constant (independent of (x, y)) and $z := g^n(x, y)$. We call an input z (and any (x, y) consistent with z) an i -violation input if $|\text{viol}(z)| = i$ and all violations occur at the terminals T . We analyze the behavior of Π on i -violation inputs with $i \in \{1, 3, 7\}$ and show a contradiction via the following implication:

- (*) *If protocol Π accepts all 1-violation (resp. 3-violation) inputs with probability 0 (resp. 2α), then Π must mess up by accepting some 7-violation input with probability $> 6\alpha$.*

Henceforth, we use $o(1)$ to denote anonymous quantities that tend to 0 as $|\Pi|/k$ tends to 0.

The implication (*) can be derived cleanly from two types of limitations of our too-good-to-be-true Π . The first limitation concerns the situation where we start with a 1-violation input z , and consider 3-violation inputs z^{B_1} and z^{B_2} that are obtained from z by flipping either a typical canonical path B_1 or another typical canonical path B_2 that is edge-disjoint from B_1 (the endpoints of B_i are terminals). The protocol should accept both z^{B_1} and z^{B_2} (more precisely, any (x, y) consistent with them) with probability 2α , but it better not accept both inputs while generating the same transcript—otherwise we could exploit the rectangular nature of transcripts and *cut-and-paste* the inputs z^{B_1} and z^{B_2} together and fool Π into accepting z , which would contradict correctness. What we actually get is that the accepting transcripts for z^{B_1} and z^{B_2} should be near-disjoint:

1-vs-3 Lemma. *Let z be any 1-violation input and let \mathcal{P} be any pairing of the non-violated terminals with canonical edge-disjoint paths B_1, \dots, B_k . Let w be a random window of z , and choose distinct $i, j \in [k]$ at random. Then, with probability $\geq 1 - o(1)$,*

$$\Delta(\Pi^{\text{acc}}|w^{B_i}, \Pi^{\text{acc}}|w^{B_j}) \geq 1 - o(1).$$

The second limitation concerns the situation where we start with a 3-violation input z and flip a typical canonical path B to obtain a 5-violation input z^B . Consider a typical accepting transcript τ in $\Pi|z$. It is unlikely that the execution τ catches us making the tiny local change $z \mapsto z^B$ in the input, and one expects that τ continues to appear in $\Pi|z^B$ with noticeable probability. (This is the usual *corruption* property of large rectangles.) Formally, for windows w_1 and w_2 , we set² $p_\tau^i := \mathbb{P}[\Pi|w_i = \tau]$ and define

$$\Pi|w_1 \text{ overflows onto } \Pi|w_2 \quad \text{iff} \quad \sum_\tau \max(p_\tau^1 - p_\tau^2, 0) \leq o(\alpha), \quad (3)$$

where the sum is over accepting transcripts τ . Intuitively, this means that for a typical accepting transcript τ , we have $p_\tau^2 \geq (1 - o(1)) \cdot p_\tau^1$. (The definition of overflow makes sense for any distributions over transcripts; we will also apply it to $\Pi|z$.) For technical reasons (which will become apparent shortly), we shall flip two canonical paths instead of one in order to pass from 3-violation inputs to 7-violation inputs.

²Note that the event in $\mathbb{P}[\Pi|w_i = \tau]$ is to be parsed as “a sample from the distribution $(\Pi|w_i)$ yields τ ”.

3-vs-7 Lemma. Let z be any 3-violation input and let \mathcal{P} be any pairing of the non-violated terminals with canonical edge-disjoint paths B_1, \dots, B_{k-1} . Let w be a random window of z , and choose distinct $i, j \in [k-1]$ at random. Then, with probability $\geq 1 - o(1)$,

$$\Pi|w \text{ overflows onto } \Pi|w^{B_i \cup B_j}.$$


6.4 Deriving the contradiction

We now prove $(*)$ by applying the **1-vs-3 Lemma** and the **3-vs-7 Lemma** in a black-box fashion to find some 7-violation input that Π accepts with too high a probability $> 6\alpha$.

Define $F := ([7], E)$ as the *Fano plane* hypergraph on 7 nodes. See the figure on the right. This hypergraph has 7 hyperedges, each of which is incident to 3 nodes, and the hyperedges are pairwise uniquely intersecting. For each hyperedge $e \in E$ choose some arbitrary but fixed pairing \mathcal{P}^e of the remaining nodes in $[7] \setminus e$.

Probability space. Choose the following at random:

1. An injection of $[7]$ into T . Denote the result by $v_1, \dots, v_7 \in T$.
2. A pairing \mathcal{P} of the remaining terminals $T \setminus \{v_1, \dots, v_7\}$.
3. A 7-violation input z_7 with $\text{viol}(z_7) = \{v_1, \dots, v_7\}$.
4. A window w_7 of z_7 .



We do not make a distinction between the nodes of F and their embedding $\{v_1, \dots, v_7\}$ in T . In particular, we think of the hyperedges $e \in E$ as triples of terminals, and the \mathcal{P}^e as pairings of terminals. Associated with the pairing $\mathcal{P}^e \cup \mathcal{P}$ there is a canonical collection of edge-disjoint paths; let $\{B_1^e, B_2^e\}$ denote the two paths that connect \mathcal{P}^e in this collection.

Based on the above, we define seven 3-violation windows, indexed by $e \in E$:

$$\text{window } w_e := w_7^{B_1^e \cup B_2^e} \text{ of } z_e := z_7^{B_1^e \cup B_2^e} \quad (\text{note: } \text{viol}(z_e) = e).$$

The following claim (proved at the end of this subsection) follows directly from the **1-vs-3 Lemma** and the **3-vs-7 Lemma** as soon as we view our probability space from the right perspective.

Claim 5. In the following list of 28 events, each occurs with probability $\geq 1 - o(1)$:

- *Overflow for $e \in E$:* $\Pi|w_e$ overflows onto $\Pi|w_7$.
- *Near-disjointness for $\{e, e'\} \subseteq E$:* $\Delta(\Pi^{\text{acc}}|w_e, \Pi^{\text{acc}}|w_{e'}) \geq 1 - o(1)$.

By a union bound over all the 28 events in the above list, we can fix our random choices 1–4 to obtain a fixed 7-violation window w_7 and fixed 3-violation windows w_e such that

$$\text{Overflow: } \forall e \in E: \quad \sum_{\tau} \max(p_{\tau}^e - p_{\tau}^7, 0) \leq o(\alpha), \quad (4)$$

$$\text{Near-disjointness: } \forall \{e, e'\} \subseteq E: \quad \sum_{\tau} \min(p_{\tau}^e, p_{\tau}^{e'}) \leq o(\alpha). \quad (5)$$

Here $p_{\tau}^7 := \mathbb{P}[\Pi|w_7 = \tau]$, $p_{\tau}^e := \mathbb{P}[\Pi|w_e = \tau]$, and the sums are over accepting transcripts; we have also rephrased the near-disjointness property using the fact that $\mathbb{P}[\Pi|w_e \text{ accepts}] = 2\alpha$.

These two properties state that typical accepting transcripts for $\Pi|w_e$ contribute to the acceptance probability of $\Pi|w_7$, and these contributions are pairwise near-disjoint. Hence, roughly speaking, one expects $\mathbb{P}[\Pi|w_7 \text{ accepts}]$ to be at least $\sum_{e \in E} \mathbb{P}[\Pi|w_e \text{ accepts}] = 7 \cdot 2\alpha = 14\alpha > 6\alpha$. But then some 7-violation input in w_7 would be accepted with probability $> 6\alpha$, which completes the proof of $(*)$ (and hence **Theorem 3**). Indeed, we perform this calculation carefully as follows. We first partition

the set of accepting transcripts as $\bigcup_{e \in E} S_e$ where S_e consists of those τ 's for which $p_\tau^e = \max_{e'} p_\tau^{e'}$ (breaking ties arbitrarily). Then

$$\begin{aligned}
\mathbb{P}[\Pi|w_7 \text{ accepts}] &= \sum_{\tau} p_\tau^7 \\
&\geq \sum_{e \in E, \tau \in S_e} \min(p_\tau^7, p_\tau^e) \\
&= \sum_{e \in E, \tau \in S_e} (p_\tau^e - \max(p_\tau^e - p_\tau^7, 0)) \\
&\geq \sum_{e \in E, \tau \in S_e} p_\tau^e - \sum_{e \in E, \tau} \max(p_\tau^e - p_\tau^7, 0) \\
&\geq \sum_{e \in E, \tau \in S_e} p_\tau^e - 7 \cdot o(\alpha) && \text{(via (4))} \\
&= \sum_{e \in E, \tau} p_\tau^e - \sum_{e \in E, e' \in E \setminus \{e\}, \tau \in S_{e'}} p_\tau^e - o(\alpha) \\
&= \sum_{e \in E, \tau} p_\tau^e - \sum_{e \in E, e' \in E \setminus \{e\}, \tau \in S_{e'}} \min(p_\tau^e, p_\tau^{e'}) - o(\alpha) \\
&\geq \sum_{e \in E, \tau} p_\tau^e - \sum_{e \in E, e' \in E \setminus \{e\}, \tau} \min(p_\tau^e, p_\tau^{e'}) - o(\alpha) \\
&\geq \sum_{e \in E, \tau} p_\tau^e - 7 \cdot 6 \cdot o(\alpha) - o(\alpha) && \text{(via (5))} \\
&= \sum_{e \in E} \mathbb{P}[\Pi|w_e \text{ accepts}] - o(\alpha) \\
&= 7 \cdot 2\alpha - o(\alpha) \\
&= (14 - o(1)) \cdot \alpha \\
&> 6\alpha.
\end{aligned}$$

Proof of Claim 5. Overflow. For notational convenience, suppose $e = \{v_1, v_2, v_3\}$ and $\mathcal{P}^e = \{\{v_4, v_7\}, \{v_5, v_6\}\}$. An alternative way to generate a sample from our probability space is (in steps 1 and 6, we are really picking random injections):

1. Random $\{v_1, v_2, v_3\} \subseteq T$.
2. Random 3-violation input z_e subject to $\text{viol}(z_e) = \{v_1, v_2, v_3\}$.
3. Random pairing $\mathcal{P}' = \{P_1, \dots, P_{k-1}\}$ of $T \setminus \{v_1, v_2, v_3\}$ with canonical paths B_1, \dots, B_{k-1} .
4. Random window w_e of z_e .
5. Random distinct $i, j \in [k-1]$.
6. Random $\{v_4, v_7\} = P_i$ and $\{v_5, v_6\} = P_j$.
7. Deterministically, define $z_7 := z_e^{B_i \cup B_j}$ and $w_7 := w_e^{B_i \cup B_j}$ and $\mathcal{P} := \mathcal{P}' \setminus \{P_i, P_j\}$.

The choices made in steps 1–3 match the data that is quantified universally in the **3-vs-7 Lemma**, whereas steps 4 and 5 make random choices as in the **3-vs-7 Lemma**; hence the lemma applies.

Near-disjointness. For notational convenience, suppose $e = \{v_1, v_2, v_3\}$, $e' = \{v_3, v_4, v_5\}$, $\mathcal{P}^e = \{\{v_4, v_7\}, \{v_5, v_6\}\}$, and $\mathcal{P}^{e'} = \{\{v_1, v_7\}, \{v_2, v_6\}\}$ (it does not matter for the proof how \mathcal{P}^e and $\mathcal{P}^{e'}$ were chosen). An alternative way to generate a sample from our probability space is (see **Figure 2**):

1. Random $v_3 \in T$.
2. Random 1-violation input z_1 subject to $\text{viol}(z_1) = \{v_3\}$.
3. Random pairing $\mathcal{P}' = \{P'_1, \dots, P'_k\}$ of $T \setminus \{v_3\}$ with canonical paths B'_1, \dots, B'_k .
4. Random window w_1 of z_1 .
5. Random distinct $i, j, l \in [k]$.
6. Random $\{v_1, v_2\} = P'_i$ and $\{v_4, v_5\} = P'_j$ and $\{v_6, v_7\} = P'_l$.
7. Deterministically, define

- $z_e := z_1^{B'_i}$ and $w_e := w_1^{B'_i}$,
- $\hat{z}_{e'} := z_1^{B'_j}$ and $\hat{w}_{e'} := w_1^{B'_j}$,
- $\mathcal{P} := \mathcal{P}' \setminus \{P'_i, P'_j, P'_l\}$,

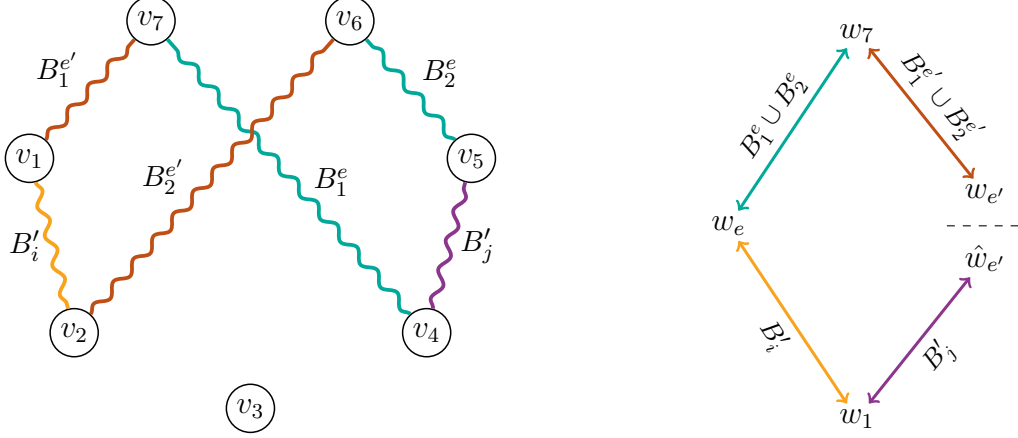


Figure 2: Illustration for the proof of Claim 5. *Left:* Paths flipped between terminals. *Right:* Relationships between windows.

- $\{B_1^e, B_2^e\}$ according to the canonical paths for $\mathcal{P}^e \cup \mathcal{P}$,
- $\{B_1^{e'}, B_2^{e'}\}$ according to the canonical paths for $\mathcal{P}^{e'} \cup \mathcal{P}$,
- $z_7 := z_e^{B_1^e \cup B_2^e}$ and $w_7 := w_e^{B_1^e \cup B_2^e}$,
- $z_{e'} := z_7^{B_1^{e'} \cup B_2^{e'}}$ and $w_{e'} := w_7^{B_1^{e'} \cup B_2^{e'}}$.

The choices made in steps 1–3 match the data that is quantified universally in the 1-vs-3 Lemma, whereas steps 4 and 5 (excluding variable l) make random choices as in the 1-vs-3 Lemma. Hence that lemma applies and shows that $\Pi^{\text{acc}}|w_e$ and $\Pi^{\text{acc}}|\hat{w}_{e'}$ are near-disjoint with high probability. Finally, we note that $\hat{w}_{e'}$ and $w_{e'}$ differ by the flipping of an eulerian graph, namely $B_j' \oplus B_i' \oplus B_1^e \oplus B_2^e \oplus B_1^{e'} \oplus B_2^{e'}$ (where \oplus means symmetric difference), so $\Pi|w_{e'}$ and $\Pi|\hat{w}_{e'}$ have the same distribution assuming w.l.o.g. that Π is smooth (as discussed in Section 6.2). Thus $\Pi^{\text{acc}}|w_e$ and $\Pi^{\text{acc}}|w_{e'}$ are also near-disjoint with high probability. \square

6.5 Roadmap for the rest of the proof

We prove the 1-vs-3 Lemma in Section 6.6 and the 3-vs-7 Lemma in Section 6.7. Both proofs rely on another technical lemma, the Homogeneity Lemma (stated below, proved in Section 6.8), which generalizes a lemma from (the full version of) [HN12, §5]. In fact, we prove the Homogeneity Lemma for any gadget g that is *regular* (as defined in Section 6.8), which our gadget is.

Homogeneity Lemma. *Fix an arbitrary $z \in \{0, 1\}^m$ for some m . Let W be a random window of z in g^m , let XY be a random input in W , and let R be an arbitrary random variable that is conditionally independent of W given XY . If $\mathbb{I}(R; XY | W) \leq o(1)$ then at least a $1 - o(1)$ fraction of windows w of z are such that $\Delta(R|w, R|z) \leq o(1)$.*

In the statement, $R|w$ is shorthand for $R|(W = w)$, and $R|z$ denotes the marginal distribution of R in the whole probability space, which is over uniformly random $XY \in (g^m)^{-1}(z)$. Furthermore, we mention that our proof shows that at least a $1 - o(1)$ fraction of $xy \in (g^m)^{-1}(z)$ are such that $\Delta(R|xy, R|z) \leq o(1)$, but for the 1-vs-3 Lemma and the 3-vs-7 Lemma we only require the property for windows.

In Section 5 we defined the directed flips $w^{\leftarrow}, w^{\leftarrow B}, w^{\uparrow}$ for a single-gadget window. We now also define directed flips for multi-gadget windows w : if B is a subset of coordinates then $w^{\leftarrow B}, w^{\leftarrow B}, w^{\uparrow B}$

are defined by applying the corresponding directed flips to the coordinates in B . Then we have the following key property of our gadget.

Fact 6. *If w is a uniformly random window of z , then each of $w^{\leftarrow B}, w^{\frown B}, w^{\uparrow B}$ is marginally a uniformly random window of z^B .*

This concept is used in the proofs of the **1-vs-3 Lemma** and the **3-vs-7 Lemma**. It turns out that the **3-vs-7 Lemma** can be proved (with a small modification to our proof) even for the simpler gadget that was used in [She11, GP14] (as can the **Homogeneity Lemma** since that gadget is regular), but our proof of the **1-vs-3 Lemma** crucially uses **Fact 6**, which does not hold for that simpler gadget.

6.6 Proof of the 1-vs-3 Lemma

Consider a probability space with the following random variables: $I \in [k]$, $J \in [k] \setminus \{I\}$, W is a random window of z^{B_I} , XY is a random input in W , and Π^{acc} is the random transcript of Π on input XY conditioned on acceptance. For convenience, denote $B := B_1 \cup \dots \cup B_k$ and $B_{-i} := B \setminus B_i$. We have

$$\mathbb{I}(\Pi^{\text{acc}}; (XY)_{B_{-I}} \mid IW) \leq \mathbb{H}(\Pi^{\text{acc}} \mid IW) \leq |\Pi| \leq o(k)$$

so by the standard direct sum property [BJKS04],

$$\begin{aligned} \mathbb{I}(\Pi^{\text{acc}}; (XY)_{B_J} \mid IJW) &= \frac{1}{k-1} \cdot \mathbb{E}_{i \sim I} \sum_{j \in [k] \setminus \{i\}} \mathbb{I}(\Pi^{\text{acc}}; (XY)_{B_j} \mid W, I = i) \\ &\leq \frac{1}{k-1} \cdot \mathbb{I}(\Pi^{\text{acc}}; (XY)_{B_{-I}} \mid IW) \\ &\leq o(1). \end{aligned}$$

Define $H := \{I, J\}$, and abbreviate $B_I \cup B_J$ as B_H and $W_{[n] \setminus (B_I \cup B_J)}$ as W_{-B_H} . By Markov's inequality, with probability $\geq 1 - o(1)$ over $h \sim H$ and $w_{-B_h} \sim W_{-B_h}$, we have

$$\mathbb{I}(\Pi^{\text{acc}}; (XY)_{B_J} \mid IJW_{B_h}, H = h, W_{-B_h} = w_{-B_h}) \leq o(1).$$

Fixing such h and w_{-B_h} (henceforth), say $h = \{1, 2\}$, it suffices to show that with probability $\geq 1 - o(1)$ over a random window w_{B_h} of z_{B_h} , we have $\Delta(\Pi^{\text{acc}}|w^{B_1}, \Pi^{\text{acc}}|w^{B_2}) \geq 1 - o(1)$ (where w is the combination of w_{B_h} and w_{-B_h}).

We rephrase the situation as follows. Consider a protocol Π_* that interprets its input as $(xy)_{B_h}$, uses private coins to sample random $(xy)_{-B_h}$ from w_{-B_h} , and runs Π on the input xy (the combination of $(xy)_{B_h}$ and $(xy)_{-B_h}$). Henceforth recycling notation by letting $z \in \{0, 1\}^{|B_h|}$ refer to z_{B_h} , and letting (I, J) be random in $\{(1, 2), (2, 1)\}$, W be a random window of (the new) z^{B_I} , and XY be a random input to Π_* in W , the situation is:

Assumption: $\mathbb{I}(\Pi_*^{\text{acc}}; (XY)_{B_J} \mid IJW) \leq o(1)$.

Want to show: For $\geq 1 - o(1)$ fraction of windows w of z , $\Delta(\Pi_*^{\text{acc}}|w^{B_1}, \Pi_*^{\text{acc}}|w^{B_2}) \geq 1 - o(1)$.

The assumption holds (with factor 2 loss in the $o(1)$) conditioned on either outcome of (I, J) ; let us tacitly condition on the outcome $(1, 2)$. Then $\mathbb{I}(\Pi_*^{\text{acc}}; (XY)_{B_2} \mid W) \leq o(1)$ where W is a random window of z^{B_1} . By Markov's inequality, with probability $\geq 1 - o(1)$ over $w_{B_1} \sim W_{B_1}$ we have $\mathbb{I}(\Pi_*^{\text{acc}}; (XY)_{B_2} \mid W_{B_2}, W_{B_1} = w_{B_1}) \leq o(1)$; call such a w_{B_1} *good*. Hence for a good w_{B_1} , we can apply the **Homogeneity Lemma** with $m := |B_2|$ and $R := \Pi_*^{\text{acc}}|_{(W_{B_1} = w_{B_1})}$ (note that $R|(xy)_{B_2}$ is the distribution of Π_*^{acc} on input $(XY)_{B_1}(xy)_{B_2}$ where $(XY)_{B_1}$ is random in w_{B_1}). This tells us that for a good w_{B_1} , with probability $\geq 1 - o(1)$ over $w_{B_2} \sim W_{B_2}$ we have $\Delta(\Pi_*^{\text{acc}}|w_{B_1}w_{B_2}, \Pi_*^{\text{acc}}|w_{B_1}z_{B_2}) \leq o(1)$, where the distribution $\Pi_*^{\text{acc}}|w_{B_1}z_{B_2}$ is over random $(XY)_{B_1} \in w_{B_1}$ and $(XY)_{B_2} \in (g^m)^{-1}(z_{B_2})$. We summarize the above with the following claim.

Claim 7. For $\geq 1 - o(1)$ fraction of windows w of z^{B_1} , we have $\Delta(\Pi_*^{\text{acc}}|w, \Pi_*^{\text{acc}}|w_{B_1}z_{B_2}) \leq o(1)$.

Conditioning on the other outcome $(I, J) = (2, 1)$ yields the symmetric property.

Claim 8. For $\geq 1 - o(1)$ fraction of windows w of z^{B_2} , we have $\Delta(\Pi_*^{\text{acc}}|w, \Pi_*^{\text{acc}}|z_{B_1}w_{B_2}) \leq o(1)$.

Now pick a random window w of z^{B_h} . Using [Fact 6](#), w^{B_2} and $w^{\wedge B_2}$ are both uniformly random (albeit correlated) windows of z^{B_1} , and w^{B_1} and $w^{\wedge B_1}$ are both uniformly random (albeit correlated) windows of z^{B_2} . Hence by [Claim 7](#), [Claim 8](#), and a union bound, with probability $\geq 1 - o(1)$ over the choice of w , the following four distances are simultaneously $\leq o(1)$: $\Delta(\Pi_*^{\text{acc}}|w^{B_2}, \Pi_*^{\text{acc}}|w_{B_1}z_{B_2})$, $\Delta(\Pi_*^{\text{acc}}|w^{\wedge B_2}, \Pi_*^{\text{acc}}|w_{B_1}z_{B_2})$, $\Delta(\Pi_*^{\text{acc}}|w^{B_1}, \Pi_*^{\text{acc}}|z_{B_1}w_{B_2})$, $\Delta(\Pi_*^{\text{acc}}|w^{\wedge B_1}, \Pi_*^{\text{acc}}|z_{B_1}w_{B_2})$.

We argue shortly that $\Delta(\Pi_*^{\text{acc}}|w^{\wedge B_1}, \Pi_*^{\text{acc}}|w^{\wedge B_2}) = 1$ with probability 1; putting everything together then shows that $\Delta(\Pi_*^{\text{acc}}|w^{B_1}, \Pi_*^{\text{acc}}|w^{B_2}) \geq 1 - o(1)$, as illustrated below. (This is equivalent to what we want to show, since sampling a window w of z^{B_h} and taking w^{B_1}, w^{B_2} is equivalent to sampling a window w of z and taking w^{B_2}, w^{B_1} .)

$$\begin{array}{ccccc}
\Pi_*^{\text{acc}}|w^{B_1} & \xleftrightarrow{\Delta \leq o(1)} & \Pi_*^{\text{acc}}|z_{B_1}w_{B_2} & \xleftrightarrow{\Delta \leq o(1)} & \Pi_*^{\text{acc}}|w^{\wedge B_1} \\
\uparrow \Delta \geq 1 - o(1) & & & & \downarrow \Delta = 1 \\
\Pi_*^{\text{acc}}|w^{B_2} & \xleftrightarrow{\Delta \leq o(1)} & \Pi_*^{\text{acc}}|w_{B_1}z_{B_2} & \xleftrightarrow{\Delta \leq o(1)} & \Pi_*^{\text{acc}}|w^{\wedge B_2}
\end{array}$$

To finish the proof, suppose for contradiction that some accepting transcript has positive probability under both $\Pi_*^{\text{acc}}|xy$ and $\Pi_*^{\text{acc}}|x'y'$ for some $xy \in w^{\wedge B_1}$ and $x'y' \in w^{\wedge B_2}$. Then Π_* would also accept xy' with positive probability, by the rectangular nature of transcripts. We claim that $g^{|B_h|}(xy') = z$. To see this, consider any coordinate c of z ; suppose $c \in B_1$ (the case $c \in B_2$ is similar). There is an embedding of stretched AND (if $z_c = 0$) or NAND (if $z_c = 1$) such that $w_c^{\wedge B_1}$ is the image of $(0, 0)$ (hence is z_c -monochromatic) and $w_c^{\wedge B_2} = w_c$ is the image of $(1, 1)$ (hence is $(1 - z_c)$ -monochromatic). Since $(xy)_c \in w_c^{\wedge B_1}$ and $(x'y')_c \in w_c$, it follows that $(xy')_c$ is in the image of $(0, 1)$, which is z_c -monochromatic. So $g((xy')_c) = z_c$ and the claim is proved.

Since Π_* accepts some input in $(g^{|B_h|})^{-1}(z)$ with positive probability (for the new z), it follows that Π accepts some input in $(g^n)^{-1}(z)$ with positive probability, for the original z , which is a contradiction since the original z has only one violation.

6.7 Proof of the 3-vs-7 Lemma

Assume for convenience that $k - 1$ is even. Note that sampling distinct $i, j \in [k - 1]$ is equivalent to sampling a permutation σ of $[k - 1]$ and an $h \in [\frac{k-1}{2}]$ and setting $i = \sigma(2h - 1)$, $j = \sigma(2h)$.

Thus we have a probability space with random variables Σ, H, I, J corresponding to the above, as well as the following: W is a random window of z , XY is a random input in W , and Π^{acc} is the random transcript of Π on input XY conditioned on acceptance. For convenience, denote $B := B_1 \cup \dots \cup B_{k-1}$ and $B_{ij} := B_i \cup B_j$. We have

$$\mathbb{I}(\Pi^{\text{acc}}; (XY)_B | W) \leq \mathbb{H}(\Pi^{\text{acc}} | W) \leq |\Pi| \leq o(k)$$

so by the standard direct sum property [BJKS04],

$$\begin{aligned}
\mathbb{I}(\Pi^{\text{acc}}; (XY)_{B_{IJ}} \mid W_{IJ}) &= \mathbb{I}(\Pi^{\text{acc}}; (XY)_{B_{IJ}} \mid W\Sigma H) \\
&= \frac{2}{k-1} \cdot \sum_{h \in [(k-1)/2]} \mathbb{I}(\Pi^{\text{acc}}; (XY)_{B_{IJ}} \mid W\Sigma, H = h) \\
&\leq \frac{2}{k-1} \cdot \mathbb{I}(\Pi^{\text{acc}}; (XY)_B \mid W\Sigma) \\
&= \frac{2}{k-1} \cdot \mathbb{I}(\Pi^{\text{acc}}; (XY)_B \mid W) \\
&\leq o(1).
\end{aligned}$$

Abbreviate $W_{[n] \setminus B_{ij}}$ as $W_{-B_{ij}}$. By Markov's inequality, with probability $\geq 1 - o(1)$ over $ij \sim IJ$ and $w_{-B_{ij}} \sim W_{-B_{ij}}$, we have $\mathbb{I}(\Pi^{\text{acc}}; (XY)_{B_{ij}} \mid W_{B_{ij}}, W_{-B_{ij}} = w_{-B_{ij}}) \leq o(1)$. Fixing such ij and $w_{-B_{ij}}$ (henceforth), it suffices to show that with probability $\geq 1 - o(1)$ over $w_{B_{ij}} \sim W_{B_{ij}}$, $\Pi|w$ overflows onto $\Pi|w^{B_{ij}}$ (where w is the combination of $w_{B_{ij}}$ and $w_{-B_{ij}}$).

We rephrase the situation as follows. Consider a protocol Π_* that interprets its input as $(xy)_{B_{ij}}$, uses private coins to sample random $(xy)_{-B_{ij}}$ from $w_{-B_{ij}}$, and runs Π on the input xy (the combination of $(xy)_{B_{ij}}$ and $(xy)_{-B_{ij}}$). Henceforth recycling notation by letting $z \in \{0, 1\}^{|B_{ij}|}$ refer to $z_{B_{ij}}$, letting B refer to B_{ij} , and letting W be a random window of (the new) z and XY be a random input to Π_* in W , the situation is:

Assumption: $\mathbb{I}(\Pi_*^{\text{acc}}; XY \mid W) \leq o(1)$.

Want to show: For $\geq 1 - o(1)$ fraction of windows w of z , $\Pi_*|w$ overflows onto $\Pi_*|w^B$.

Claim 9. For $\geq 1 - o(1)$ fraction of windows w of z^B , $\Pi_*|z$ overflows onto $\Pi_*|w$.

We prove Claim 9 shortly, but first we finish the proof of the 3-vs-7 Lemma assuming it. By the Homogeneity Lemma (with $m := |B|$ and $R := \Pi_*^{\text{acc}}$), Claim 9, and a union bound, at least a $1 - o(1)$ fraction of windows w of z are such that both $\Delta(\Pi_*^{\text{acc}}|w, \Pi_*^{\text{acc}}|z) \leq o(1)$ and $\Pi_*|z$ overflows onto $\Pi_*|w^B$ (since w^B is a uniform window of z^B if w is a uniform window of z). We show that this implies that $\Pi_*|w$ overflows onto $\Pi_*|w^B$ as follows (letting $p_\tau^z, p_\tau^w, p_\tau^{w^B}$ denote the probability of a transcript τ under the distributions $\Pi_*|z, \Pi_*|w, \Pi_*|w^B$ respectively, and summing only over accepting τ 's):

$$\sum_\tau \max(p_\tau^w - p_\tau^{w^B}, 0) \leq \sum_\tau \max(p_\tau^z - p_\tau^{w^B}, 0) + \sum_\tau |p_\tau^w - p_\tau^z| \leq o(\alpha) + o(\alpha) = o(\alpha).$$

Proof of Claim 9. By Fact 6, if w is a random window of z^B , then $w^{\leftarrow B}, w^{\wedge B}, w^{\uparrow B}$ are each marginally uniformly random windows of z . Thus by the Homogeneity Lemma (with $m := |B|$ and $R := \Pi_*^{\text{acc}}$) and a union bound, with probability $\geq 1 - o(1)$ over the choice of w , the following three distances are simultaneously $\leq o(1)$: $\Delta(\Pi_*^{\text{acc}}|w^{\leftarrow B}, \Pi_*^{\text{acc}}|z)$, $\Delta(\Pi_*^{\text{acc}}|w^{\wedge B}, \Pi_*^{\text{acc}}|z)$, $\Delta(\Pi_*^{\text{acc}}|w^{\uparrow B}, \Pi_*^{\text{acc}}|z)$. Now assuming this good event occurs for some particular w , we just need to show that $\Pi_*|z$ overflows onto $\Pi_*|w$.

(See Figure 3 for a proof-by-picture.) Let $p_\tau, p_\tau^{11}, p_\tau^{10}, p_\tau^{00}, p_\tau^{01}$ denote the probabilities of a transcript τ under $\Pi_*|z, \Pi_*|w, \Pi_*|w^{\leftarrow B}, \Pi_*|w^{\wedge B}, \Pi_*|w^{\uparrow B}$ respectively. Let $\gamma_\tau^{00} := |p_\tau - p_\tau^{00}|$, and for $ab \in \{01, 10\}$ let $\gamma_\tau^{ab} := |p_\tau^{00} - p_\tau^{ab}|$. We claim that for all τ , $p_\tau - p_\tau^{11} \leq \gamma_\tau^{00} + \gamma_\tau^{01} + \gamma_\tau^{10}$; this will finish the proof since then (summing only over accepting τ 's)

$$\sum_\tau \max(p_\tau - p_\tau^{11}, 0) \leq \sum_\tau (\gamma_\tau^{00} + \gamma_\tau^{01} + \gamma_\tau^{10}) \leq o(\alpha) + o(\alpha) + o(\alpha) = o(\alpha)$$

where the second inequality is because $\sum_\tau \gamma_\tau^{00}, \sum_\tau \gamma_\tau^{01}, \sum_\tau \gamma_\tau^{10} \leq o(\alpha)$ follow from (respectively) $\Delta(\Pi_*^{\text{acc}}|z, \Pi_*^{\text{acc}}|w^{\wedge B}), \Delta(\Pi_*^{\text{acc}}|w^{\wedge B}, \Pi_*^{\text{acc}}|w^{\uparrow B}), \Delta(\Pi_*^{\text{acc}}|w^{\wedge B}, \Pi_*^{\text{acc}}|w^{\leftarrow B}) \leq o(1)$.




Figure 3: Proof of [Claim 9](#) illustrated. The four windows w , $w^{\leftarrow B}$, $w^{\nwarrow B}$, $w^{\nearrow B}$ are rectangles of (x, y) 's. Each (x, y) can be further subdivided according to the private coins (r_A, r_B) of the players. The protocol Π_* partitions the extended input space of (xr_A, yr_B) 's into *transcript rectangles*—above, we have only drawn *accepting* transcript rectangles (in various colors). For a window w' , the probability $\mathbb{P}[\Pi_*|w' = \tau]$ is simply the *area* (appropriately scaled) of the transcript rectangle of τ inside w' . In the proof of [Claim 9](#), the relevant case is when all of $\Pi_*^{\text{acc}}|w^{\leftarrow B}$, $\Pi_*^{\text{acc}}|w^{\nwarrow B}$, $\Pi_*^{\text{acc}}|w^{\nearrow B}$ have roughly the same distribution, say, D (in fact, $D := \Pi_*^{\text{acc}}|z$). By the rectangular property of transcripts, this forces $\Pi_*|z$ to *overflow onto* $\Pi_*|w$. (Note that $\Pi_*^{\text{acc}}|w$ may contain additional transcripts to those in D , since the acceptance probability is higher.)

To verify the subclaim, it suffices to show that

$$p_\tau^{01} \cdot p_\tau^{10} \geq (p_\tau^{00})^2 - p_\tau^{00} \gamma_\tau^{01} - p_\tau^{00} \gamma_\tau^{10} \quad (6)$$

since by the rectangular nature of transcripts, we have $p_\tau^{00} \cdot p_\tau^{11} = p_\tau^{01} \cdot p_\tau^{10}$, and thus if $p_\tau^{00} > 0$ then

$$p_\tau^{11} = \frac{p_\tau^{01} \cdot p_\tau^{10}}{p_\tau^{00}} \geq p_\tau^{00} - \gamma_\tau^{01} - \gamma_\tau^{10} \geq p_\tau - \gamma_\tau^{00} - \gamma_\tau^{01} - \gamma_\tau^{10}$$

and if $p_\tau^{00} = 0$ then of course $p_\tau^{11} \geq p_\tau^{00} = p_\tau - \gamma_\tau^{00}$. To see (6), note that for some signs $\sigma_\tau^{01}, \sigma_\tau^{10} \in \{1, -1\}$, the left side of (6) equals $(p_\tau^{00} + \sigma_\tau^{01} \gamma_\tau^{01}) \cdot (p_\tau^{00} + \sigma_\tau^{10} \gamma_\tau^{10})$, which expands to

$$(p_\tau^{00})^2 + \sigma_\tau^{01} p_\tau^{00} \gamma_\tau^{01} + \sigma_\tau^{10} p_\tau^{00} \gamma_\tau^{10} + \sigma_\tau^{01} \sigma_\tau^{10} \gamma_\tau^{01} \gamma_\tau^{10}. \quad (7)$$

If $\sigma_\tau^{01} = \sigma_\tau^{10}$ then (7) is at least the right side of (6) since the last term of (7) is nonnegative. If $\sigma_\tau^{01} \neq \sigma_\tau^{10}$, say $\sigma_\tau^{01} = -1$ and $\sigma_\tau^{10} = 1$, then (7) is at least the right side of (6) since the sum of the last two terms in (7) is $p_\tau^{00} \gamma_\tau^{10} - \gamma_\tau^{01} \gamma_\tau^{10} = p_\tau^{01} \gamma_\tau^{10} \geq 0$. \square

6.8 Proof of the Homogeneity Lemma

Definition 10. For a gadget $g: \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$ and $b \in \{0, 1\}$, define the digraph \mathcal{G}^b as follows: the nodes are the b -inputs of g , and there is an edge from xy to $x'y'$ iff $x = x'$ or $y = y'$. (That is, each node has a self-loop, and all b -inputs in a given row or column have all possible edges between them.)

Definition 11. We say a gadget $g: \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$ is *regular* iff (i) $|\mathcal{X}| = |\mathcal{Y}|$ is even, (ii) each row and each column is balanced (number of 0's equals number of 1's), and (iii) \mathcal{G}^0 and \mathcal{G}^1 are both strongly connected.

Our gadget g is indeed regular, but we prove the lemma for any regular constant-size g .

The first part of the proof is inspired by a similar approach that was used in [HN12]. We augment the probability space with the following random variables: let $X'Y'$ be a random input in W that is conditionally independent of XY given W (note that in each of the m coordinates independently, XY and $X'Y'$ are equal with probability $1/2$), and let $E \in ((g^m)^{-1}(z))^2$ be chosen randomly from $\{(XY, X'Y'), (X'Y', XY)\}$. We have $\mathbb{H}(R|E) = \mathbb{H}(R|WE) \leq \mathbb{H}(R|W)$ since R is conditionally independent of W given E , and conditioning decreases entropy. We also have $\mathbb{H}(R|XYE) = \mathbb{H}(R|XY) = \mathbb{H}(R|XYW)$ since R is conditionally independent of WE given XY . Putting these together, we get

$$\mathbb{I}(R; XY|E) = \mathbb{H}(R|E) - \mathbb{H}(R|XYE) \leq \mathbb{H}(R|W) - \mathbb{H}(R|XYW) = \mathbb{I}(R; XY|W) \leq o(1).$$

By Markov's inequality, with probability $\geq 1 - o(1)$ over $e \sim E$, we have $\mathbb{I}(R; XY|E = e) \leq o(1)$, in which case if $e = (x^{(0)}y^{(0)}, x^{(1)}y^{(1)})$ then by Pinsker's inequality³, $\Delta(R|x^{(0)}y^{(0)}, R|x^{(1)}y^{(1)}) \leq o(1)$; let us use $\epsilon > 0$ for the latter $o(1)$ quantity. We describe what the above means in graph theoretic terms.


Define the digraph \mathcal{G}^z as follows: the nodes are the inputs in $(g^m)^{-1}(z)$, and there is an edge from one input to another iff there exists a window of z containing both inputs; this includes a self-loop at each node. Note that \mathcal{G}^z is the tensor product $\mathcal{G}^{z_1} \otimes \dots \otimes \mathcal{G}^{z_m}$, i.e., each node of \mathcal{G}^z corresponds to an m -tuple of nodes from those digraphs, and each edge of \mathcal{G}^z corresponds to an m -tuple of edges. For convenience, we make the dependence of the random variable E on z explicit using the notation E^z ; thus E^z is distributed over the edges of \mathcal{G}^z . By regularity, for $b \in \{0, 1\}$ the distribution of E^b over the edges of \mathcal{G}^b puts half its mass uniformly over the self-loops, and half its mass uniformly over the non-self-loops. Note that the distribution of E^z is the product of the distributions of E^{z_1}, \dots, E^{z_m} , i.e., E^z can be sampled by taking samples $(x^{(0,i)}y^{(0,i)}, x^{(1,i)}y^{(1,i)})$ from E^{z_i} (independent over $i \in [m]$) and forming the edge $(x^{(0,1)}y^{(0,1)} \dots x^{(0,m)}y^{(0,m)}, x^{(1,1)}y^{(1,1)} \dots x^{(1,m)}y^{(1,m)})$ in \mathcal{G}^z .

We say an edge $(x^{(0)}y^{(0)}, x^{(1)}y^{(1)})$ of \mathcal{G}^z is *great* iff $\Delta(R|x^{(0)}y^{(0)}, R|x^{(1)}y^{(1)}) \leq \epsilon$. Thus the great edges have at least $1 - o(1)$ probability mass under E^z .

Let L be the number of non-self-loop edges in \mathcal{G}^b (which is the same for $b = 0$ and $b = 1$).

Claim 12. *There exists a distribution over length- $2L$ walks on \mathcal{G}^z such that (i) the first and last nodes are independent and each marginally uniform, and (ii) each of the $2L$ edges on the walk is marginally distributed according to E^z .*

Proof. By the product structure of \mathcal{G}^z and E^z , it suffices to prove this claim for a bit b instead of z (as the claim for z follows by sampling m independent such walks on the \mathcal{G}^{z_i} 's and running them "in parallel"). By regularity, if we ignore the self-loops, there exists an eulerian tour in \mathcal{G}^b



Example of \mathcal{G}^1 for the regular gadget $x_1 + y_1 + x_2y_2$

³Specifically, if RB are jointly distributed random variables where $B \in \{0, 1\}$ is a uniformly random bit, and R_b denotes the distribution of $R|(B = b)$, then $\mathbb{I}(R; B) = \mathbb{D}(R_0 \| R)/2 + \mathbb{D}(R_1 \| R)/2 \geq \Omega(\Delta(R_0, R)^2/2 + \Delta(R_1, R)^2/2) \geq \Omega((\Delta(R_0, R)/2 + \Delta(R_1, R)/2)^2) \geq \Omega(\Delta(R_0, R_1)^2)$, where \mathbb{D} denotes KL-divergence, and the first inequality is Pinsker's, the second is by convexity of the square function, and the third is by the triangle inequality.

that uses all the non-self-loop edges exactly once, and pays an equal number of visits to each node. Let $v_0, v_1, \dots, v_{L-1}, v_0$ denote the sequence of nodes visited (with repeats) on a fixed such tour. We explicitly describe the distribution of walks $v_{i_0}, \dots, v_{i_{2L}}$ on \mathcal{G}^b , using mod- L arithmetic:

1. Independently sample i_0 and ℓ uniformly from $\{0, \dots, L-1\}$.
2. For $j = 1, \dots, \ell$, execute one of the following with probability $1/2$ each:
 - 2a. Use the self-loop then move forward (i.e., $i_{2j-1} = i_{2j-2}$ and $i_{2j} = i_{2j-1} + 1$).
 - 2b. Move forward then use the self-loop (i.e., $i_{2j-1} = i_{2j-2} + 1$ and $i_{2j} = i_{2j-1}$).
3. For $j = \ell + 1, \dots, L$, execute one of the following with probability $1/2$ each:
 - 3a. Use the self-loop twice (i.e., $i_{2j} = i_{2j-1} = i_{2j-2}$).
 - 3b. Move forward then backward (i.e., $i_{2j-1} = i_{2j-2} + 1$ and $i_{2j} = i_{2j-1} - 1$).

This procedure has L phases, each taking 2 steps of the walk. Each of the first ℓ phases has the effect of moving forward one node on the tour, and each of the last $L - \ell$ phases has the effect of ending up at the same node the phase started at. Thus $i_{2L} = i_0 + \ell$ and is hence independent of i_0 and uniform over $\{0, \dots, L-1\}$ (since ℓ is independent of i_0 and uniform); hence also v_{i_0} and $v_{i_{2L}}$ are independent and uniform (since the tour visits each node equally often) and so (i) is verified. Property (ii) holds even conditioned on any ℓ , and can be verified by a little case analysis; e.g., if $\ell > 1$ then the first edge is (v_{i_0}, v_{i_0}) with probability $1/2$, and is (v_{i_0}, v_{i_0+1}) with probability $1/2$ (this is a sample from E^b since v_{i_0} is a uniform node and (v_{i_0}, v_{i_0+1}) is a uniform non-self-loop edge). \square

If we sample a walk $x^{(0)}y^{(0)}, \dots, x^{(2L)}y^{(2L)}$ in \mathcal{G}^z as in [Claim 12](#), then by property (ii) and a union bound, with probability $\geq 1 - 2L \cdot o(1) = 1 - o(1)$, each of the edges on the walk is great, in which case by the triangle inequality, $\Delta(R|x^{(0)}y^{(0)}, R|x^{(2L)}y^{(2L)}) \leq 2L\epsilon$. In summary, by property (i), a $1 - o(1)$ fraction of pairs of inputs in $(g^m)^{-1}(z)$ are *good* in the sense that their conditional distributions of R are within statistical distance $2L\epsilon = o(1)$. Thus a $1 - o(1)$ fraction of inputs $xy \in (g^m)^{-1}(z)$ are such that (xy, \overline{xy}) is good for a $1 - o(1)$ fraction of $\overline{xy} \in (g^m)^{-1}(z)$, in which case (letting \overline{xy} be random in $(g^m)^{-1}(z)$ in the following)

$$\begin{aligned}
\Delta(R|xy, R) &= \Delta(R|xy, \mathbb{E}_{\overline{xy}} R|\overline{xy}) \\
&\leq \mathbb{E}_{\overline{xy}} \Delta(R|xy, R|\overline{xy}) \\
&\leq \mathbb{P}_{\overline{xy}}[(xy, \overline{xy}) \text{ is good}] \cdot o(1) + \mathbb{P}_{\overline{xy}}[(xy, \overline{xy}) \text{ is not good}] \cdot 1 \\
&\leq 1 \cdot o(1) + o(1) \cdot 1 \\
&= o(1)
\end{aligned}$$

where the second line is a basic general fact about statistical distance. Say xy is *typical* if $\Delta(R|xy, R) \leq o(1)$ as above. Note that in the original probability space, XY is marginally uniform over $(g^m)^{-1}(z)$ and thus with probability at least $1 - o(1)$ over sampling $w \sim W$ and $xy \sim XY \in w$, xy is typical. It follows that for at least $1 - o(1)$ fraction of w , at least $1 - o(1)$ fraction of $xy \in w$ are typical, in which case

$$\begin{aligned}
\Delta(R|w, R) &= \Delta(\mathbb{E}_{xy \in w} R|xy, R) \\
&\leq \mathbb{E}_{xy \in w} \Delta(R|xy, R) \\
&\leq \mathbb{P}_{xy \in w}[xy \text{ is typical}] \cdot o(1) + \mathbb{P}_{xy \in w}[xy \text{ is not typical}] \cdot 1 \\
&\leq 1 \cdot o(1) + o(1) \cdot 1 \\
&= o(1).
\end{aligned}$$

7 Query Lower Bound

An alternative approach for proving a lower bound for the $(\#\exists-1)$ -game for $\text{TSE}_G \circ g^n$ is:

Step 1: Prove an appropriate *query complexity* lower bound for TSE_G .

Step 2: Use a communication-to-query simulation theorem like [CLRS16, GLM⁺16, LRS15].

In this section, we carry out the first step by proving an optimal $\Omega(n)$ query complexity lower bound (in particular, this answers a question from [LNNW95]). The query lower bound is a lot simpler to prove than our $\Omega(n/\log n)$ communication lower bound in Section 6. Unfortunately, as we discuss in Section 7.3, it is not known how to perform the second step for small-enough gadgets g .

The result of this section can be interpreted as evidence that the right bound in Theorem 1 is $2^{\Omega(n)}$ and the right bound in Corollary 4 is $\Omega(n)$, and also as motivation for further work to improve parameters for simulation theorems.

7.1 Conical juntas

The query complexity analogue of nonnegative rank decompositions (nonnegative combinations of nonnegative rank-1 matrices) are *conical juntas*: nonnegative combinations of conjunctions of literals (input bits or their negations). We write a conical junta as $h = \sum_C w_C C$ where $w_C \geq 0$ and C ranges over all conjunctions $C: \{0, 1\}^n \rightarrow \{0, 1\}$. The *degree* of h is the maximum number of literals in a conjunction C with $w_C > 0$. Each conical junta naturally computes a nonnegative function $h: \{0, 1\}^n \rightarrow \mathbb{R}_{\geq 0}$.

A randomized decision tree \mathcal{T} of height d (probability distribution over deterministic decision trees of height d) can be converted into a degree- d conical junta h that on input $z \in \{0, 1\}^n$ outputs the acceptance probability of \mathcal{T} on input z ; the conversion is the same as for multilinear polynomials [BdW02, Theorem 15]. Conversely, given a degree- d conical junta $h = \sum_C w_C C$, consider the randomized height- d decision tree \mathcal{T} that on input $z \in \{0, 1\}^n$ samples a random conjunction C with probability proportional to w_C and accepts iff $C(z) = 1$. The resulting randomized decision tree has acceptance probability *proportional* to $h(z)$. Hence, for computational tasks (such as $(\#\exists-1)$ -games) where we care only about the acceptance probability *up to scaling*, randomized decision trees and conical juntas are equivalent.

We will study $(\#\exists-1)$ -games in query complexity: the query complexity of the $(\#\exists-1)$ -game for TSE_G is the least degree of a conical junta h that on input z outputs $h(z) = |\text{viol}(z)| - 1$.

7.2 A linear lower bound

Theorem 13. *There is a family of n -node bounded-degree graphs G such that the $(\#\exists-1)$ -game for TSE_G requires query complexity $\Omega(n)$.*

Relation to [LNNW95]. An analogue of the (KW/EF) connection holds for query complexity: if there is a deterministic decision tree of height d that solves the search problem TSE_G , we can convert this into a degree- $(d + O(1))$ conical junta for the associated $(\#\exists-1)$ -game. Moreover, if we only have a *randomized* ϵ -error decision tree for the search problem, then the connection gives us a conical junta h that *approximately* solves the $(\#\exists-1)$ -game: $h(z) \in (|\text{viol}(z)| - 1) \cdot (1 \pm \epsilon)$ for all z .

Our proof below is robust enough that the $\Omega(n)$ bound holds even for conical juntas that merely approximately solve the $(\#\exists-1)$ -game. Hence we get a randomized $\Omega(n)$ lower bound for TSE_G , which was conjectured by [LNNW95, p. 125]; note however that the earlier work [GP14] already got a near-optimal $\Omega(n/\log n)$ bound. In any case, to our knowledge, this is the first $O(1)$ -vs- $\Omega(n)$

separation between nondeterministic query complexity (which for TSE_G equals the maximum degree of G) and randomized query complexity for search problems.

The proof. Fix an n -node bounded-degree expander $G = (V, E)$. That is, for any subset $U \subseteq V$ of size $|U| \leq n/2$, the number of edges leaving U is $\Theta(|U|)$. We tacitly equip G with an arbitrary odd-weight node-labeling. Assume for the sake of contradiction that there is a conical junta $h = \sum w_C C$ of degree $o(n)$ for the $(\#\exists-1)$ -game for TSE_G . Let C be a conjunction with $w_C > 0$. Denote by $S \subseteq E$ the set of edges that C reads; hence $|S| \leq o(n)$. Below, we write $G \setminus S$ for the graph induced on the edges $E \setminus S$ (deleting nodes that become isolated).

Claim 14. *We may assume w.l.o.g. that $G \setminus S$ is connected.*

Proof. If $G \setminus S$ is not connected, we may replace C with a conjunction (actually, a sum of them) that reads more input variables; namely, we let C read a larger set of edges $S' \supseteq S$ including all edges from connected components of $G \setminus S$ of “small” size $\leq n/2$. When adding some small component $K \subseteq E$ to S' we note that, because G is expanding, the size of K is big- O of the size of the edge boundary of K (which is contained in S). On the other hand, every edge in S lies on the boundary of at most two components. It follows that $|S'| = O(|S|)$, i.e., we increased the degree of h only by a constant factor. Now in $G \setminus S'$ we have only components of size $> n/2$, but there can only be one such component. \square

Claim 15. *We may assume w.l.o.g. that C witnesses at least two fixed nodes with a parity violation (i.e., C reads all the edge labels incident to the two nodes).*

Proof. Suppose for contradiction that C witnesses at most one violation. Then we may fool C into accepting an input (and hence h into outputting a positive value on that input) where the number of violations is 1, which is a contradiction to the definition of the $(\#\exists-1)$ -game. Indeed, let z be some input accepted by C . Then we may modify z freely on the connected graph $G \setminus S$ (by Claim 14) without affecting C 's acceptance: we may eliminate pairs of violations from z by flipping paths (as in Section 3) until only one remains. (This is possible since by definition, all the non-witnessed violations of z remain in $G \setminus S$.) \square

Let μ_i (i odd) denote the distribution on inputs that have i violations at a random set of i nodes, and are otherwise random with this property. We may generate an input from μ_i as follows:

1. Choose an i -set $T_i \subseteq V$ of nodes at random.
2. Let $z \in \mathbb{Z}_2^E$ be any fixed input with $\text{viol}(z) = T_i$.
3. Let $q \in \mathbb{Z}_2^E$ be a random eulerian graph.
4. Output $z + q$.

Theorem 13 follows from the following lemma. Here we identify C with the set (subcube) of inputs it accepts, and define $\mu_i(C) := \mathbb{P}_{z \sim \mu_i}[C(z) = 1]$.

Lemma 16. $\mu_5(C) \geq (10/3 - o(1)) \cdot \mu_3(C)$.

Indeed, consider the expected output value $\mathbb{E}_{z_i \sim \mu_i}[h(z_i)]$. This should be 2 for $i = 3$, and 4 for $i = 5$, i.e., a factor 2 increase. However, the above lemma implies that the output value gets multiplied by more than a factor 3, which is the final contradiction.

Proof of Lemma 16. By Claim 15 let $\{v_1, v_2\}$ be a pair of nodes where C witnesses two violations. For $i = 3, 5$, let $z_i \sim \mu_i$ and denote by T_i the i -set of its violations. Then

$$\begin{aligned}\mu_3(C) &= \mathbb{P}[C(z_3) = 1] \\ &= \mathbb{P}[C(z_3) = 1 \text{ and } T_3 \supseteq \{v_1, v_2\}] \\ &= \binom{n-2}{1} / \binom{n}{3} \cdot \mathbb{P}[C(y_3) = 1], \quad (\text{for } y_3 := (z_3 \mid T_3 \supseteq \{v_1, v_2\}))\end{aligned}$$

$$\begin{aligned}\mu_5(C) &= \mathbb{P}[C(z_5) = 1] \\ &= \mathbb{P}[C(z_5) = 1 \text{ and } T_5 \supseteq \{v_1, v_2\}] \\ &= \binom{n-2}{3} / \binom{n}{5} \cdot \mathbb{P}[C(y_5) = 1]. \quad (\text{for } y_5 := (z_5 \mid T_5 \supseteq \{v_1, v_2\}))\end{aligned}$$

So their ratio is

$$\frac{\mu_5(C)}{\mu_3(C)} = \frac{10}{3} \cdot \frac{\mathbb{P}[C(y_5) = 1]}{\mathbb{P}[C(y_3) = 1]}.$$

Hence the following claim concludes the proof of Lemma 16. □

Claim 17. $\mathbb{P}[C(y_5) = 1] / \mathbb{P}[C(y_3) = 1] \geq 1 - o(1)$.

Proof. We can generate y_3 and y_5 jointly as follows:

- y_3 :** Choose $v_3 \in V \setminus \{v_1, v_2\}$ uniformly random and let x_3 be some input with $\text{viol}(x_3) = \{v_1, v_2, v_3\}$. Output $y_3 := x_3 + q$ where q is a random eulerian graph.
- y_5 :** Continuing from the above, choose $\{v_4, v_5\} \subseteq V \setminus \{v_1, v_2, v_3\}$ at random. If possible, let p be a path in $G \setminus S$ joining $\{v_4, v_5\}$ (a “good” event), otherwise let p be any path joining $\{v_4, v_5\}$. Output $y_5 := x_3 + p + q$.

It suffices to prove the claim conditioned on any particular v_3 (and hence also on x_3). By Claim 14 we have $\mathbb{P}[\text{“good”} \mid v_3] = \mathbb{P}[v_4, v_5 \in G \setminus S \mid v_3] \geq 1 - o(1)$ since $|S| \leq o(n)$. If the “good” event occurs, then C cannot distinguish between $y_3 = x_3 + q$ and $y_5 = x_3 + p + q$ so that $\mathbb{P}[C(y_3) = 1 \mid v_3] = \mathbb{P}[C(y_5) = 1 \mid \text{“good”}, v_3]$. The claim follows as

$$\begin{aligned}\mathbb{P}[C(y_5) = 1 \mid v_3] &\geq \mathbb{P}[C(y_5) = 1 \text{ and “good”} \mid v_3] \\ &= \mathbb{P}[C(y_5) = 1 \mid \text{“good”}, v_3] \cdot \mathbb{P}[\text{“good”} \mid v_3] \\ &= \mathbb{P}[C(y_3) = 1 \mid v_3] \cdot \mathbb{P}[\text{“good”} \mid v_3] \\ &\geq \mathbb{P}[C(y_3) = 1 \mid v_3] \cdot (1 - o(1)). \quad \square\end{aligned}$$

7.3 Limitations of existing simulation theorems

The main result of [GLM⁺16] is a simulation of randomized protocols (or nonnegative rank decompositions) by conical juntas: a cost- d protocol for a lifted problem $F \circ g^n$ can be simulated by a degree- $O(d)$ conical junta (approximately) computing F . While F here is arbitrary, the result unfortunately assumes that $g := \text{IP}_b$ is a logarithmic-size, $b := \Theta(\log n)$, inner-product function $\text{IP}_b: \{0, 1\}^b \times \{0, 1\}^b \rightarrow \{0, 1\}$ given by $\text{IP}_b(x, y) := \langle x, y \rangle \bmod 2$.

Plugging b -bit gadgets into the reductions of Section 4 would blow up the number of input bits of CSP-SAT exponentially in b . This is not only an artifact of our particular reduction! Consider more generally any reduction from a communication search problem $S \circ g^n$ to a KW^+ -game for a monotone $f: \{0, 1\}^m \rightarrow \{0, 1\}$. Since the KW^+ -game has *nondeterministic* communication complexity $\log m$ (number of bits the players must nondeterministically guess to find a witness), the reduction would

imply $c \leq \log m$ where c is the nondeterministic communication complexity of $S \circ g^n$. If merely computing g requires b bits of nondeterministic communication, then clearly $c \geq b$ so that $m \geq 2^b$. For $b = \Theta(\log n)$, as required by [GLM⁺16], the number of input bits m becomes polynomially larger than the bound $\Omega(n)$ we proved above, and hence the lower bound is weak as a function of m .

A major open problem is to develop simulation theorems for gadget size $b = O(1)$.

Acknowledgements

Thanks to Denis Pankratov, Toniann Pitassi, and Robert Robere for discussions. We also thank Samuel Fiorini and Raghu Meka for e-mail correspondence. Finally, many thanks to anonymous referees (of FOCS'16 and SICOMP) for thoughtful comments. M.G. admits to having a wonderful time at IBM while learning about extended formulations with T.S. Jayram and Jan Vondrak.

Part of this research was done while M.G. and R.J. were attending the *Semidefinite and Matrix Methods for Optimization and Communication* program at the Institute for Mathematical Sciences, National University of Singapore in 2016. This research was supported in part by NSERC, in part by NSF grant CCF-1657377, and in part by the Singapore Ministry of Education and the National Research Foundation, also through the Tier 3 Grant *Random numbers from quantum processes* MOE2012-T3-1-009. M.G. was partially supported by the Simons Award for Graduate Students in TCS.

References

- [AT14] David Avis and Hans Raj Tiwary. On the extension complexity of combinatorial polytopes. *Mathematical Programming*, 153(1):95–115, 2014. doi:10.1007/s10107-014-0764-2.
- [BdW02] Harry Buhrman and Ronald de Wolf. Complexity measures and decision tree complexity: A survey. *Theoretical Computer Science*, 288(1):21–43, 2002. doi:10.1016/S0304-3975(01)00144-X.
- [BJKS04] Ziv Bar-Yossef, T.S. Jayram, Ravi Kumar, and D. Sivakumar. An information statistics approach to data stream and communication complexity. *Journal of Computer and System Sciences*, 68(4):702–732, 2004. doi:10.1016/j.jcss.2003.11.006.
- [BM13] Mark Braverman and Ankur Moitra. An information complexity approach to extended formulations. In *Proceedings of the 45th Symposium on Theory of Computing (STOC)*, pages 161–170. ACM, 2013. doi:10.1145/2488608.2488629.
- [BP15] Gábor Braun and Sebastian Pokutta. The matching polytope does not admit fully-polynomial size relaxation schemes. In *Proceedings of the 26th Symposium on Discrete Algorithms (SODA)*, pages 837–846. ACM–SIAM, 2015. doi:10.1137/1.9781611973730.57.
- [BP16] Gábor Braun and Sebastian Pokutta. Common information and unique disjointness. *Algorithmica*, 76(3):597–629, 2016. doi:10.1007/s00453-016-0132-0.
- [CCZ10] Michele Conforti, Gérard Cornuéjols, and Giacomo Zambelli. Extended formulations in combinatorial optimization. *4OR*, 8(1):1–48, 2010. doi:10.1007/s10288-010-0122-z.
- [CCZ14] Michele Conforti, Gérard Cornuéjols, and Giacomo Zambelli. *Integer Programming*. Springer, 2014. doi:10.1007/978-3-319-11008-0.

- [CLRS16] Siu On Chan, James Lee, Prasad Raghavendra, and David Steurer. Approximate constraint satisfaction requires large LP relaxations. *Journal of the ACM*, 63(4):34:1–34:22, 2016. doi:10.1145/2811255.
- [FFGT14] Yuri Faenza, Samuel Fiorini, Roland Grappe, and Hans Raj Tiwary. Extended formulations, nonnegative factorizations, and randomized communication protocols. *Mathematical Programming*, 153(1):75–94, 2014. doi:10.1007/s10107-014-0755-3.
- [FGL⁺96] Uriel Feige, Shafi Goldwasser, László Lovász, Shmuel Safra, and Mario Szegedy. Interactive proofs and the hardness of approximating cliques. *Journal of the ACM*, 43(2):268–292, 1996. doi:10.1145/226643.226652.
- [FMP⁺15] Samuel Fiorini, Serge Massar, Sebastian Pokutta, Hans Raj Tiwary, and Ronald de Wolf. Exponential lower bounds for polytopes in combinatorial optimization. *Journal of the ACM*, 62(2):17:1–17:23, 2015. doi:10.1145/2716307.
- [Fri01] Alan Frieze. Edge-disjoint paths in expander graphs. *SIAM Journal on Computing*, 30(6):1790–1801, 2001. doi:10.1137/S0097539700366103.
- [FZ00] Alan Frieze and Lei Zhao. Optimal construction of edge-disjoint paths in random regular graphs. *Combinatorics, Probability and Computing*, 9(3):241–263, 2000. doi:10.1017/S0963548300004284.
- [Gál01] Anna Gál. A characterization of span program size and improved lower bounds for monotone span programs. *Computational Complexity*, 10(4):277–296, 2001. doi:10.1007/s000370100001.
- [GJW16] Mika Göös, Rahul Jain, and Thomas Watson. Extension complexity of independent set polytopes. In *Proceedings of the 57th Symposium on Foundations of Computer Science (FOCS)*, pages 565–572. IEEE, 2016. doi:10.1109/FOCS.2016.67.
- [GLM⁺16] Mika Göös, Shachar Lovett, Raghu Meka, Thomas Watson, and David Zuckerman. Rectangles are nonnegative juntas. *SIAM Journal on Computing*, 45(5):1835–1869, 2016. doi:10.1137/15M103145X.
- [GP14] Mika Göös and Toniann Pitassi. Communication lower bounds via critical block sensitivity. In *Proceedings of the 46th Symposium on Theory of Computing (STOC)*, pages 847–856. ACM, 2014. doi:10.1145/2591796.2591838.
- [HN12] Trinh Huynh and Jakob Nordström. 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)*, pages 233–248. ACM, 2012. doi:10.1145/2213977.2214000.
- [Hru12] Pavel Hrubeš. On the nonnegative rank of distance matrices. *Information Processing Letters*, 112(11):457–461, 2012. doi:10.1016/j.ipl.2012.02.009.
- [Juk12] Stasys Jukna. *Boolean Function Complexity: Advances and Frontiers*, volume 27 of *Algorithms and Combinatorics*. Springer, 2012.
- [Kai11] Volker Kaibel. Extended formulations in combinatorial optimization. Technical report, arXiv, 2011. arXiv:1104.1023.

- [KN97] Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, 1997.
- [KW88] Mauricio Karchmer and Avi Wigderson. Monotone circuits for connectivity require super-logarithmic depth. In *Proceedings of the 20th Symposium on Theory of Computing (STOC)*, pages 539–550. ACM, 1988. doi:10.1145/62212.62265.
- [LNNW95] László Lovász, Moni Naor, Ilan Newman, and Avi Wigderson. Search problems in the decision tree model. *SIAM Journal on Discrete Mathematics*, 8(1):119–132, 1995. doi:10.1137/S0895480192233867.
- [LRS15] James Lee, Prasad Raghavendra, and David Steurer. Lower bounds on the size of semidefinite programming relaxations. In *Proceedings of the 47th Symposium on Theory of Computing (STOC)*, pages 567–576. ACM, 2015. doi:10.1145/2746539.2746599.
- [Nor15] Jakob Nordström. New wine into old wineskins: A survey of some pebbling classics with supplemental results. Technical report, KTH Royal Institute of Technology, 2015.
- [Oli15] Igor Oliveira. *Unconditional Lower Bounds in Complexity Theory*. PhD thesis, Columbia University, 2015. doi:10.7916/D8ZP45KT.
- [PV13] Sebastian Pokutta and Mathieu Van Vyve. A note on the extension complexity of the knapsack polytope. *Operations Research Letters*, 41(4):347–350, 2013. doi:10.1016/j.orl.2013.03.010.
- [Raz90] Alexander Razborov. Applications of matrix methods to the theory of lower bounds in computational complexity. *Combinatorica*, 10(1):81–93, 1990. doi:10.1007/BF02122698.
- [Rot12] Thomas Rothvoß. Some 0/1 polytopes need exponential size extended formulations. *Mathematical Programming*, 142(1):255–268, 2012. doi:10.1007/s10107-012-0574-3.
- [Rot14] Thomas Rothvoß. The matching polytope has exponential extension complexity. In *Proceedings of the 46th Symposium on Theory of Computing (STOC)*, pages 263–272. ACM, 2014. doi:10.1145/2591796.2591834.
- [Rou15] Tim Roughgarden. Communication complexity (for algorithm designers). Technical report, arXiv, 2015. arXiv:1509.06257.
- [RW92] Ran Raz and Avi Wigderson. Monotone circuits for matching require linear depth. *Journal of the ACM*, 39(3):736–744, 1992. doi:10.1145/146637.146684.
- [Sch03] Alexander Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*, volume 24 of *Algorithms and Combinatorics*. Springer, 2003.
- [She11] Alexander Sherstov. The pattern matrix method. *SIAM Journal on Computing*, 40(6):1969–2000, 2011. doi:10.1137/080733644.
- [Wel16] Stefan Weltge. *Sizes of Linear Descriptions in Combinatorial Optimization*. PhD thesis, Otto-von-Guericke-Universität Magdeburg, 2016.
- [Yan91] Mihalis Yannakakis. Expressing combinatorial optimization problems by linear programs. *Journal of Computer and System Sciences*, 43(3):441–466, 1991. doi:10.1016/0022-0000(91)90024-Y.

[Zie95] Günter Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer, 1995.