

Toward Better Depth Lower Bounds: A KRW-like theorem for Strong Composition*

Or Meir[†]

February 12, 2025

Abstract

One of the major open problems in complexity theory is proving super-logarithmic lower bounds on the depth of circuits (i.e., $\mathbf{P} \not\subseteq \mathbf{NC}^1$). Karchmer, Raz, and Wigderson (Computational Complexity 5(3/4), 1995) suggested approaching this problem by proving that the depth complexity of a composition of functions $f \diamond g$ is roughly the sum of the depth complexities of f and g . They showed that the validity of this conjecture would imply that $\mathbf{P} \not\subseteq \mathbf{NC}^1$.

The intuition that underlies the KRW conjecture is that the composition $f \diamond g$ should behave like a “direct-sum problem”, in a certain sense, and therefore the depth complexity of $f \diamond g$ should be the sum of the individual depth complexities. Nevertheless, there are two obstacles toward turning this intuition into a proof: first, we do not know how to prove that $f \diamond g$ must behave like a direct-sum problem; second, we do not know how to prove that the complexity of the latter direct-sum problem is indeed the sum of the individual complexities.

In this work, we focus on the second obstacle. To this end, we study a notion called “strong composition”, which is the same as $f \diamond g$ except that it is forced to behave like a direct-sum problem. We prove a variant of the KRW conjecture for strong composition, thus overcoming the above second obstacle. This result demonstrates that the first obstacle above is the crucial barrier toward resolving the KRW conjecture. Along the way, we develop some general techniques that might be of independent interest.

*A preliminary version of this work appeared in FOCS 2023.

[†]Department of Computer Science, University of Haifa, Haifa 3498838, Israel. ormeir@cs.haifa.ac.il. Partially supported by the Israel Science Foundation (grant No. 716/20).

Contents

1	Introduction	3
1.1	Our result	4
1.2	Our techniques	6
1.2.1	Composition of multiplexors	7
1.2.2	Our structure theorem	7
1.2.3	Lower bounds using graph coloring	8
1.2.4	Prefix-thick sets	9
1.2.5	The proof of the structure theorem	10
1.2.6	Comparison with previous works	12
1.3	The organization of the paper	15
2	Preliminaries	16
2.1	Depth complexity and formula complexity	17
2.2	Communication complexity	18
2.2.1	Non-deterministic protocols	19
2.3	Karchmer-Wigderson and multiplexor relations	20
2.3.1	Fortification	20
2.3.2	Multiplexor relations	20
2.4	Half-duplex protocols	21
2.4.1	Partially half-duplex protocol and multiplexors	23
2.5	Linear codes and the Varshamov bound	24
3	Main theorem	24
3.1	Reduction to multiplexor lower bounds	25
3.2	The structure theorem	26
3.3	Proof of main theorem from structure theorem	27
4	Multiplexor lower bounds via graph coloring	29
4.1	The graph equality problem	30
4.2	Proof of Lemmas 4.2 and 4.4	31
5	Prefix-thick sets	35
6	Proof of the structure theorem	38
6.1	The construction of \mathcal{G}'	39
6.2	The independence number of \mathcal{G}'	43
7	A barrier to improving γ	45

1 Introduction

A major frontier of the research on circuit complexity is proving super-logarithmic lower bounds on the depth complexity of an explicit function, i.e., proving that $\mathbf{P} \not\subseteq \mathbf{NC}^1$. This question is an important milestone toward proving lower bounds on general circuits, and also captures the natural question of whether there are tractable computational tasks that cannot be parallelized. The state of the art is the work of Håstad [Hås93], who proved a lower bound of $(3 - o(1)) \cdot \log n$, following a long line of work [Sub61, Khr72, And87, PZ93, IN93]. This lower bound has not been improved for three decades except for the lower order terms [Tal14], and it is an important problem to break this barrier.

Karchmer, Raz, and Wigderson [KRW91] proposed approaching this problem by studying the composition of Boolean functions, defined as follows: if $f : \{0, 1\}^m \rightarrow \{0, 1\}$ and $g : \{0, 1\}^n \rightarrow \{0, 1\}$ are Boolean functions, then their composition $f \diamond g$ takes inputs in $(\{0, 1\}^n)^m$ and is defined by

$$f \diamond g(x_1, \dots, x_m) = f(g(x_1), \dots, g(x_m)). \quad (1)$$

Let us denote by $D(f)$ the minimal depth of a circuit with fan-in 2 that computes f . The circuit that computes $f \diamond g$ using Equation (1) has depth $D(f) + D(g)$. Karchmer et al. [KRW91] conjectured that this upper bound is roughly optimal:

Conjecture 1.1 (The KRW conjecture). *Let $f : \{0, 1\}^m \rightarrow \{0, 1\}$ and $g : \{0, 1\}^n \rightarrow \{0, 1\}$ be non-constant functions. Then*

$$D(f \diamond g) \approx D(f) + D(g). \quad (2)$$

Karchmer et al. observed that their conjecture, if proved, would imply that $\mathbf{P} \not\subseteq \mathbf{NC}^1$. The meaning of “approximate equality” in Equation (2) is intentionally left vague, since there are many variants that would imply the separation. In particular, the conjecture would imply that $\mathbf{P} \not\subseteq \mathbf{NC}^1$ even if it only holds for *some hard* function g rather than for *every* function g :

Conjecture 1.2 (weak KRW conjecture). *For every non-constant function $f : \{0, 1\}^m \rightarrow \{0, 1\}$ and every $n \in \mathbb{N}$ there exists a function $g : \{0, 1\}^n \rightarrow \{0, 1\}$ such that*

$$D(f \diamond g) \geq D(f) + n - O(\log(m \cdot n)). \quad (3)$$

Proposition 1.3 (folklore, see, e.g., [Mei20, Prop. 2.10]). *The weak KRW conjecture implies that $\mathbf{P} \not\subseteq \mathbf{NC}^1$.*

The KRW conjecture has been studied extensively in [KRW91, EIRS91, Hås93, HW93, GMWW14, DM16, KM18, dRMN⁺20, Mei20, FMT21, MS21]. These works succeeded in proving several special cases of the conjecture, and in identifying some variants of particular interest. And yet, the full KRW conjecture, or even just Conjecture 1.2, seem beyond our reach.

The communication complexity approach. It is useful to study the KRW conjecture through the lens of communication complexity, and specifically, using the framework of *Karchmer-Wigderson relations*. Let us denote the (deterministic) communication complexity of a problem R by $CC(R)$. The *Karchmer-Wigderson relation* of a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, denoted KW_f , is the communication problem in which the inputs of Alice and Bob are $x \in f^{-1}(1)$ and $y \in f^{-1}(0)$ respectively, and their goal is to find a coordinate i such that $x_i \neq y_i$. Karchmer and Wigderson [KW88] observed that $D(f) = CC(KW_f)$. This connection allows us to study the depth complexity of functions using techniques from communication complexity.

Now, let $f : \{0, 1\}^m \rightarrow \{0, 1\}$ and $g : \{0, 1\}^n \rightarrow \{0, 1\}$ be non-constant functions. For convenience, we define the notation $KW_f \diamond KW_g$ to denote the KW relation $KW_{f \diamond g}$ of the composed function. In this relation, Alice and Bob get $X \in (f \diamond g)^{-1}(1)$ and $Y \in (f \diamond g)^{-1}(0)$, viewed as $m \times n$ matrices, and their goal is to find an entry (i, j) such that $X_{i,j} \neq Y_{i,j}$. The KRW conjecture can be restated as:

$$\text{CC}(KW_f \diamond KW_g) \approx \text{CC}(KW_f) + \text{CC}(KW_g).$$

It is worth noting the obvious protocol for solving $KW_f \diamond KW_g$: Let $a = g(X)$ and $b = g(Y)$ be the column vectors that are obtained by applying g to the rows of X and Y respectively. Observe that they constitute an instance of KW_f , since by definition $f(g(X)) = 1$ and $f(g(Y)) = 0$. The players begin by solving KW_f on a and b , thus obtaining a coordinate $i \in [m]$ such that $a_i \neq b_i$. This implies that $g(X_i) \neq g(Y_i)$, and therefore the rows X_i and Y_i constitute an instance of KW_g . The players now solve KW_g on the rows X_i and Y_i , thus obtaining a coordinate $j \in [n]$ where $X_{i,j} \neq Y_{i,j}$. The communication complexity of this protocol is $\text{CC}(KW_f) + \text{CC}(KW_g)$, and the KRW conjecture says that this obvious protocol is roughly optimal.

Remark 1.4. The text of the introduction up to this point borrows heavily from [dRMN⁺20], with the permission of the authors.

1.1 Our result

In this work, we make progress toward the weak KRW conjecture by decoupling the two major obstacles toward proving it, and tackling one of them separately. Specifically, we prove the conjecture for a simpler notion of composition, called “strong composition”, which captures one of the barriers to proving the conjecture,

In order to motivate the notion of strong composition, recall that the KRW conjecture says that the above “obvious protocol” for $KW_f \diamond KW_g$ is essentially optimal. There is a very good intuition for why this should be the case: First, recall that in the problem $KW_f \diamond KW_g$, the players are looking for an entry (i, j) such that $X_{i,j} \neq Y_{i,j}$. It seems obvious that the players have to look for a solution (i, j) in a row where $a_i \neq b_i$, since in any other row there is no guarantee that a solution even exists. Now, observe that finding such a row i is equivalent to solving KW_f on a and b . Moreover, given such a row i , finding a solution (i, j) is equivalent to solving KW_g on X_i and Y_i . Hence, it seems that in order to solve $KW_f \diamond KW_g$, the players have to solve both KW_f and KW_g , and therefore have to transmit roughly $\text{CC}(KW_f) + \text{CC}(KW_g)$ bits.

Nevertheless, when attempting to turn the above intuition into a formal proof, one immediately encounters two significant obstacles:

- While it seems obvious that the players should look for a solution (i, j) in a row i where $a_i \neq b_i$, it is difficult to prove that they *must* do so.
- Even if we could prove that the players have to solve both KW_f and KW_g , it would still not necessarily imply that they must transmit at least $\text{CC}(KW_f) + \text{CC}(KW_g)$ bits. It could be the case that the players solve both relations using less bits by employing some clever strategy. For example, perhaps they could somehow “recycle” the bits that were used to solve KW_f in order to solve KW_g . In fact, showing that this cannot be done is an instance of the well-known “direct-sum question” in communication complexity.

Indeed, proofs of special cases of the KRW conjecture employ highly non-trivial arguments to deal with those obstacles. In particular, these arguments are tailored for specific cases at hand, and do not seem applicable to the general KRW conjecture or even its weak variant. In this work, we focus

on tackling the second obstacle above, while avoiding the first obstacle. To this end, we consider the following strong composition operation, which requires Alice and Bob to find a solution (i, j) in rows i where $a_i \neq b_i$.

Definition 1.5. Let $f : \{0, 1\}^m \rightarrow \{0, 1\}$ and $g : \{0, 1\}^n \rightarrow \{0, 1\}$ be non-constant functions. The *strong composition* of KW_f and KW_g , denoted $KW_f \circledast KW_g$, is the following communication problem: Alice and Bob take as inputs $X \in (f \diamond g)^{-1}(1)$ and $Y \in (f \diamond g)^{-1}(0)$. Let $a = g(X)$ and $b = g(Y)$. The goal of Alice and Bob is to find an entry (i, j) such that both $a_i \neq b_i$ and $X_{i,j} \neq Y_{i,j}$.

Observe that this notion of composition is indeed stronger than the standard notion, in the sense that every protocol that solves $KW_f \circledast KW_g$ also solves $KW_f \diamond KW_g$. This, in turn, means that *proving the KRW conjecture for strong composition is necessary for proving the original KRW conjecture* (and the same goes for the weak variant). Informally, our main result is the following analogue of the weak KRW conjecture for strong composition (see [Theorem 3.1](#) for the formal statement).

Theorem 1.6 (Main theorem, informal). *There exists a constant $\gamma > 0.04$ such that the following holds: for every non-constant function $f : \{0, 1\}^m \rightarrow \{0, 1\}$ and for every $n \in \mathbb{N}$ there exists a function $g : \{0, 1\}^n \rightarrow \{0, 1\}$ such that*

$$\text{CC}(KW_f \circledast KW_g) \geq \text{CC}(KW_f) + n - (1 - \gamma) \cdot m - O(\log(m \cdot n)). \quad (4)$$

A few remarks are in order:

- The strong composition $KW_f \circledast KW_g$ is not a KW relation, and hence our result does not imply depth lower bounds.
- The loss of $(1 - \gamma) \cdot m$ means that Theorem 1.6 is considerably weaker quantitatively compared to the weak KRW conjecture. Yet, the lower bound we get is far from trivial: if this theorem is proved for standard composition rather than strong composition, it will imply a new depth lower bound of $3.04 \cdot \log n$. Such a result would be the first significant improvement in depth lower bounds in three decades, and hence very exciting. Nevertheless, such a result would not be strong enough to imply $\mathbf{P} \not\subseteq \mathbf{NC}^1$.
- We believe that Theorem 1.6 can be strengthened to hold even when g is a random function (in which case the lower bound holds with high probability). Nevertheless, we did not verify it.
- The notion of strong composition $KW_f \circledast KW_g$ was defined in passing, without a name, by de Rezende et al. [dRMN⁺20]. Such a definition was considered even earlier in private discussions among researchers: for example, Sajin Koroth suggested this notion to us in 2019 (private communication).
- In the setting of monotone circuits, strong composition is equivalent to standard composition [KRW91, dRMN⁺20] (essentially, this holds because the min-terms and max-terms of $f \diamond g$ force the condition of strong composition; see [dRMN⁺24, Sec. 3.1.1] for a proof). Since our proof of Theorem 1.6 can be adapted to monotone circuits in a straightforward manner, this means that our main result actually holds *for standard composition* in the monotone setting. Again, we believe this is true even when g is a random (slice) function. Such a result was not known prior to our work. Unfortunately, the lower bounds that are implied by this result are weaker than the state-of-the-art lower bounds in the monotone setting.

- The only difference between the informal statement of our main theorem and the formal one is that the formal statement refers to the logarithm of the formula complexity of f rather than to $\text{CC}(KW_f)$.
- We did not attempt to optimize the value of the constant γ in Theorem 1.6. Nevertheless, improving γ to a value of 0.64 or larger would require significant new ideas (see Section 7 for details).
- We chose the name “strong composition” for the relation $KW_f \otimes KW_g$ since we view both $KW_f \otimes KW_g$ and $KW_f \diamond KW_g$ as ways to compose the *KW relations* KW_f and KW_g , where the former is stronger than the latter since it requires the output to satisfy a stronger requirement. Note, however, that $KW_f \diamond KW_g$ also corresponds to the composition of *the functions* f and g , whereas $KW_f \otimes KW_g$ does not correspond to a composition of *the functions* in any obvious way.

Previous work. In their original paper, Karchmer, Raz, and Wigderson [KRW91] defined the universal relation U , a simplification of KW relations, and proposed to prove their conjecture for the composition $U \diamond U$ as a step toward the full conjecture. This challenge was met by Edmonds et al. [EIRS91], and an alternative proof was given later by Håstad and Wigderson [HW93]. These results were extended to compositions of the form $KW_f \diamond U$ in the work Gavinsky et al. [GMWW14], and their result was improved quantitatively by Korothe and Meir [KM18].

Håstad [Hås93] proved the KRW conjecture implicitly for the composition $KW_f \diamond KW_{\text{parity}}$ for every function f using the method of random restrictions. Dinur and Meir [DM16] provided an alternative proof of that result using the communication complexity approach. Filmus, Meir, and Tal [FMT21] generalized the result of [Hås93] to compositions $KW_f \diamond KW_g$ where f is an any function, and g is any function for which the soft quantum adversary bound is tight. Unfortunately, the result of [FMT21] does not imply the weak KRW conjecture since the only functions g that satisfy the latter condition are relatively easy to compute ($D(g) \leq 2 \log n$).

De Rezende et al. [dRMN⁺20] proved the KRW conjecture in the *monotone* setting for every outer function f and for a large family of inner functions g : namely, those functions g for which there is a reduction from a lifted problem to the monotone KW relation mKW_g . They also introduced a new notion of a “semi-monotone composition”, and proved a semi-monotone KRW conjecture for $U \diamond mKW_g$ for a similar family of functions g .

Most relevant to the current paper is the work of Mihajlin and Smal [MS21], who proved the weak KRW conjecture for compositions of the form $U \diamond KW_g$. Our main result is incomparable to theirs: on the one hand, our result applies to a KW relation KW_f rather than the universal relation U ; on the other hand, their result applies to standard composition whereas our result applies only to strong composition.

1.2 Our techniques

In order to prove our main result, we develop several new techniques that we believe to be no less important than the result itself. In particular, we believe that these techniques will be useful for attacking the weak KRW conjecture (with standard composition). In what follows, we provide an overview of the the proof of our main result, including a detailed discussion of those new techniques.

1.2.1 Composition of multiplexors

The bulk of our proof is a lower bound for the “multiplexor composition” $KW_f \otimes MUX_n$: this communication problem is defined similarly to $KW_f \otimes KW_g$, but now the function g is given to the players as part of the input. More specifically, in the problem $KW_f \otimes MUX_n$, Alice takes as input a function $g : \{0, 1\}^n \rightarrow \{0, 1\}$ and a matrix $X \in (f \diamond g)^{-1}(1)$, and Bob takes *the same function* g and a matrix $Y \in (f \diamond g)^{-1}(0)$. As before, we let $a = g(X)$ and $b = g(Y)$. The goal of Alice and Bob is to find an entry (i, j) such that both $a_i \neq b_i$ and $X_{i,j} \neq Y_{i,j}$.

Multiple works have observed that lower bounds for such multiplexor composition problems can be used to derive depth lower bounds (see, e.g., [EIRS91, HW93, Mei20]). Most relevant to this paper is the following observation of Mihajlin and Smal [MS21]: for every function $f : \{0, 1\}^m \rightarrow \{0, 1\}$ and every $n \in \mathbb{N}$, there exists a function $g : \{0, 1\}^n \rightarrow \{0, 1\}$ such that

$$\text{CC}(KW_f \diamond KW_g) \geq \text{CC}^{\text{phd}}(KW_f \diamond MUX_n) - O(\log(m \cdot n)),$$

where CC^{phd} denotes communication complexity in a model called “partially half-duplex protocols”. Such a result holds just as well for strong composition. Therefore, to prove our main result, it suffices to prove that for every $f : \{0, 1\}^m \rightarrow \{0, 1\}$ and $n \in \mathbb{N}$ it holds that

$$\text{CC}^{\text{phd}}(KW_f \otimes MUX_n) \geq \text{CC}(KW_f) + n - (1 - \gamma) \cdot m - O(\log(m \cdot n)).$$

In the rest of this section, we describe the proof of the latter lower bound. In order to streamline the presentation, we ignore the differences between standard and partially half-duplex protocols in this overview.

1.2.2 Our structure theorem

Recall that the obvious protocol for $KW_f \diamond KW_g$ works in two stages: the players first solve KW_f on a, b , thus obtaining a row $i \in [m]$, and then solve KW_g on the rows X_i, Y_i . Our goal is to prove that this protocol is roughly optimal. Earlier works [EIRS91, DM16, KM18, dRMN⁺20] employed the following general strategy: Given any protocol Π , we break it into two parts, which roughly correspond to the two stages of the obvious protocol. We then show that Π transmits roughly $\text{CC}(KW_f)$ bits in the first part, and roughly $\text{CC}(KW_g)$ bits in the second part. This implies that Π must transmit roughly $\text{CC}(KW_f) + \text{CC}(KW_g)$ bits overall, as required.

We follow this general strategy in our proof as well. Specifically, we define a notion of “live transcripts”, which correspond to the first stage of the obvious protocol, and prove two claims:

- There exists a live transcript of length $\text{CC}(KW_f) - (1 - \gamma) \cdot m$.
- **The structure theorem:** after transmitting a live transcript, the protocol must transmit at least $n - O(\log(m \cdot n))$ more bits.

Together, the two claims imply our lower bound. The first claim is relatively easy to prove, and the bulk of this paper is devoted to proving the structure theorem.

Intuitively, a live transcript is a partial transcript π_1 in which the protocol Π has not finished solving KW_f on a and b yet, and has not revealed too much information about g , X , and Y . We now describe this notion in detail. To this end, we begin with setting up some notation. Fix a protocol Π and a partial transcript π_1 of Π . As usual in communication complexity, we associate with π_1 a set \mathcal{X}_{π_1} that consists of all the inputs (g, X) of Alice that are consistent with π_1 .

Similarly, we denote by \mathcal{Y}_{π_1} the set of all inputs (g, Y) of Bob that are consistent with π_1 . For every function $g : \{0, 1\}^n \rightarrow \{0, 1\}$ and strings $a \in f^{-1}(1)$ and $b \in f^{-1}(0)$, we denote

$$\begin{aligned} \mathcal{X}_{\pi_1}(g) &= \{X \in \{0, 1\}^{m \times n} : (g, X) \in \mathcal{X}_{\pi_1}\} & \mathcal{Y}_{\pi_1}(g) &= \{Y \in \{0, 1\}^{m \times n} : (g, Y) \in \mathcal{Y}_{\pi_1}\} \\ \mathcal{X}_{\pi_1}(g, a) &= \{X \in \mathcal{X}_{\pi_1}(g) : g(X) = a\} & \mathcal{Y}_{\pi_1}(g, b) &= \{Y \in \mathcal{Y}_{\pi_1}(g) : g(Y) = b\} \\ &= \mathcal{X}_{\pi_1}(g) \cap g^{-1}(a) & &= \mathcal{Y}_{\pi_1}(g) \cap g^{-1}(b). \end{aligned}$$

We also denote by $\mathcal{A}_{\pi_1}(g)$ and $\mathcal{B}_{\pi_1}(g)$ the sets of all strings $a \in f^{-1}(1)$ and $b \in f^{-1}(0)$ such that $\mathcal{X}_{\pi_1}(g, a)$ and $\mathcal{Y}_{\pi_1}(g, b)$ are non-empty, respectively. Finally, we denote by \mathcal{V}_{π_1} the set of functions $g : \{0, 1\}^n \rightarrow \{0, 1\}$ such that both $\mathcal{A}_{\pi_1}(g)$ and $\mathcal{B}_{\pi_1}(g)$ are non-empty. For convenience, we focus only on *balanced* functions g (i.e., functions that take the value 1 on exactly half of the inputs). Informally, we say that π_1 is *alive* if there exists a set $\mathcal{V} \subseteq \mathcal{V}_{\pi_1}$ of balanced functions that satisfies the following conditions:

- The set \mathcal{V} is large, i.e., it consists of at least $2^{-O(m)}$ fraction of all balanced functions. Intuitively, the transcript π_1 does not reveal too much information about g .
- For every $g \in \mathcal{V}$, it is still quite hard to solve $\text{CC}(KW_f)$ on inputs from the set $\mathcal{A}_{\pi_1}(g) \times \mathcal{B}_{\pi_1}(g)$. Specifically, we require that any protocol that solves KW_f on $\mathcal{A}_{\pi_1}(g) \times \mathcal{B}_{\pi_1}(g)$ has to transmit at least $(1 - \gamma) \cdot m$ bits.
- For every $g \in \mathcal{V}$ and strings $a \in \mathcal{A}_{\pi_1}(g)$ and $b \in \mathcal{B}_{\pi_1}(g)$, the sets $\mathcal{X}_{\pi_1}(g, a)$ and $\mathcal{Y}_{\pi_1}(g, b)$ are large. Specifically, the sets should consist of at least $2^{-\gamma \cdot m}$ fraction of the sets $g^{-1}(a)$ and $g^{-1}(b)$. Intuitively, the transcript π_1 does not reveal too much information about X, Y even conditioned on g, a , and b .

For the rest of this overview, we focus on proof of the above structure theorem.

1.2.3 Lower bounds using graph coloring

Our starting point for proving the structure theorem is the following result.

Lemma 1.7 (implicit in [MS21]). *Let Π be a protocol that solves $KW_f \diamond \text{MUX}_n$, and let π_1 be a transcript of Π . Suppose there is a subset of functions $\mathcal{V}_1 \subseteq \mathcal{V}_{\pi_1}$ such that, for every two distinct functions $g_A, g_B \in \mathcal{V}_1$, the following intersection property holds:*

- **Intersection Property:** *either $\mathcal{X}_{\pi_1}(g_A) \cap \mathcal{Y}_{\pi_1}(g_B) \neq \emptyset$ or $\mathcal{X}_{\pi_1}(g_B) \cap \mathcal{Y}_{\pi_1}(g_A) \neq \emptyset$.*

Then, after transmitting π_1 , the protocol Π must transmit at least

$$\log \log |\mathcal{V}_1| - \log \log \log |\mathcal{V}_1| - O(1)$$

more bits.

Essentially, [MS21] prove their lower bound for $U \diamond \text{MUX}$ by constructing a sufficiently long transcript π_1 with such a set \mathcal{V}_1 of size roughly 2^{2^n} . Lemma 1.7 then implies a lower bound of roughly n on the number of bits that Π must communicate after π_1 .

Lemma 1.7 is a powerful tool for proving lower bounds, but the requirement that the intersection property holds for *every* two functions $g_A, g_B \in \mathcal{V}_1$ is quite difficult to satisfy. In this work, we show that the latter requirement can be weakened substantially. To this end, we take a graph-theoretic perspective on Lemma 1.7. We define a graph \mathcal{G}_{π_1} over \mathcal{V}_{π_1} in which two distinct functions $g_A, g_B \in \mathcal{V}_{\pi_1}$ are neighbors if and only if they satisfy the intersection property. Taking this view, Lemma 1.7

gives a lower bound of roughly $\log \log(\omega(\mathcal{G}_{\pi_1}))$, where $\omega(\mathcal{G}_{\pi_1})$ is the *clique number* of \mathcal{G}_{π_1} (i.e., the size of the largest clique in \mathcal{G}_{π_1}).

Our first main technical contribution is the following strengthening of Lemma 1.7. Let $\chi(\mathcal{G}_{\pi_1})$ denote the *chromatic number* of \mathcal{G}_{π_1} , i.e., the minimum number of colors required to color the vertices of \mathcal{G}_{π_1} such that no two adjacent vertices share the same color, and recall that $\chi(\mathcal{G}_{\pi_1}) \geq \omega(\mathcal{G}_{\pi_1})$.

Lemma 1.8 (Lemma 4.2, informal). *Let Π be a protocol that solves $KW_f \diamond MUX$, and let π_1 be a transcript of Π . Then, after transmitting π_1 , the protocol Π must transmit at least*

$$\log \log \chi(\mathcal{G}_{\pi_1}) - \log \log \log \chi(\mathcal{G}_{\pi_1}) - O(1)$$

more bits.

Moving from $\omega(\mathcal{G}_{\pi_1})$ to $\chi(\mathcal{G}_{\pi_1})$ gives us a lot more flexibility. In particular, it means that instead of *lower bounding* the size of the largest *clique*, we can *upper bound* the size of the largest *independent set*. In fact, it suffices to upper bound the size of independent sets in a large sub-graph of \mathcal{G}_{π_1} of our choice. This task turns out to be more manageable. Indeed, the crux of our proof amounts to finding an appropriate sub-graph and upper bounding the size of its independent sets.

We stress that our Lemma 1.8 is proved for *standard* composition, so it could potentially be used to attack the weak KRW conjecture in its original form. Nevertheless, both Lemmas 1.7 and 1.8 can be adapted to strong composition as well. The only difference is that the above intersection property is replaced with the following, easier to satisfy, intersection property.

- **Weak Intersection Property:** there exist matrices $X \in \mathcal{X}_{\pi_1}(g_A)$ and $Y \in \mathcal{Y}_{\pi_1}(g_B)$ such that $X_i = Y_i$ for every $i \in [m]$ for which $a_i \neq b_i$, where $a = g_A(X)$ and $b = g_B(Y)$ (or the same statement holds when exchanging g_A with g_B).

To see that this property is indeed weaker, observe that the original intersection property can be rephrased similarly, but requires that $X_i = Y_i$ for *every* $i \in [m]$, rather than just those rows i for which $a_i \neq b_i$.

Remark 1.9. After writing the first version of this paper, we realized that we could also prove our main theorem by lower bounding the clique number of \mathcal{G}_{π_1} and applying Lemma 1.7. Nevertheless, the proof that uses Lemma 1.8 is somewhat simpler. Moreover, the graph-theoretic perspective is still crucial even when using Lemma 1.7 instead of Lemma 1.8. See Remark 1.13 for more details.

1.2.4 Prefix-thick sets

Given Lemma 1.8, the main challenge in our proof is to upper bound the size of independent sets in the graph \mathcal{G}_{π_1} , where π_1 is a live transcript of a protocol Π . In particular, this requires us to prove the existence of edges in \mathcal{G}_{π_1} . To this end, we introduce a notion of “prefix-thick sets” and prove the existence of (many) such sets.

As motivation for this notion, fix two functions $g_A, g_B \in \mathcal{V}_{\pi_1}$, and suppose we would like to prove that they are neighbors in the graph \mathcal{G}_{π_1} . In order to do so, we should construct matrices X, Y and strings a, b as in the definition of the weak intersection property. We will see later that we can construct strings a, b that agree on $(1 - \gamma)$ fraction of the rows. Thus, in order to show that g_A and g_B are neighbors, it remains to construct matrices $X \in \mathcal{X}_{\pi_1}(g_A, a)$ and $Y \in \mathcal{Y}_{\pi_1}(g_B, b)$ that agree on the remaining γ fraction of the rows where $a_i \neq b_i$. To this end, recall that the sets $\mathcal{X}_{\pi_1}(g_A, a)$ and $\mathcal{Y}_{\pi_1}(g_B, b)$ are relatively large, i.e., they have density at least $2^{-\gamma \cdot m}$ (since π_1 is alive).

We now take a step back and consider the following simpler combinatorial question: given two large sets of $m \times n$ matrices \mathcal{X} and \mathcal{Y} , can we prove the existence of matrices $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$

that agree on many rows? More generally, let Σ be a finite alphabet, and let $\mathcal{X}, \mathcal{Y} \subseteq \Sigma^m$ be sets of strings of density at least $2^{-\gamma \cdot m}$. Can we prove the existence of strings $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$ that agree on an at least γ fraction of their coordinates?

This natural question leads us to the notion of prefix-thick sets. Recall that the *prefix tree* of a set of strings $\mathcal{X} \subseteq \Sigma^m$ is the rooted tree of depth m that is defined as follows: the vertices of depth i are the prefixes of length i of the strings in \mathcal{X} ; and a string x of depth i is the parent of a string y of depth $(i + 1)$ if and only if x is a prefix of y . We say that the set \mathcal{X} is *prefix thick* if its prefix tree has minimum degree greater than $\frac{1}{2} |\Sigma|$ (or if \mathcal{X} contains a set with this property). We have the following easy observation.

Proposition 1.10. *Let $\mathcal{X}, \mathcal{Y} \subseteq \Sigma^m$. If \mathcal{X} and \mathcal{Y} are both prefix thick, then they intersect.*

In particular, if we can find a large set of coordinates $I \subseteq [m]$ such that both $\mathcal{X}|_I$ and $\mathcal{Y}|_I$ are prefix thick, then $\mathcal{X}|_I$ intersects $\mathcal{Y}|_I$, and hence there exist $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$ that agree on many coordinates. Our second main technical contribution is to deduce the following lemma from a result of Salo and Törmä [ST14]:

Lemma 1.11 (informal version of Lemma 5.7). *If a set of strings $\mathcal{X} \subseteq \Sigma^m$ is sufficiently large, then there exist many sets $I \subseteq [m]$ such that $\mathcal{X}|_I$ is prefix thick.*

Using a simple counting argument, this result can be used to find a set of $\gamma \cdot m$ coordinates $I \subseteq [m]$ such that both $\mathcal{X}|_I$ and $\mathcal{Y}|_I$ are prefix thick, as required.

1.2.5 The proof of the structure theorem

Finally, we return to the structure theorem, and explain how to prove it using the above machinery. Let $f : \{0, 1\}^m \rightarrow \{0, 1\}$ and let $n \in \mathbb{N}$. Fix a protocol Π that solves $KW_f \otimes MUX_n$, and let π_1 be a live transcript. We would like to prove that after transmitting π_1 , the protocol Π must still transmit at least $n - O(\log(mn))$ additional bits. Let \mathcal{G}_{π_1} be the graph defined over \mathcal{V}_{π_1} according to the weak intersection property. That is, two functions $g_A, g_B \in \mathcal{V}_{\pi_1}$ are neighbors in \mathcal{G}_{π_1} if and only if they satisfy the following property:

- there exist matrices $X \in \mathcal{X}_{\pi_1}(g_A)$ and $Y \in \mathcal{Y}_{\pi_1}(g_B)$ such that $X_i = Y_i$ for every $i \in [m]$ for which $a_i \neq b_i$, where $a = g_A(X)$ and $b = g_B(Y)$ (or the same statement holds when exchanging g_A with g_B).

In our proof, we construct a certain sub-graph \mathcal{G}' of \mathcal{G}_{π_1} and prove that every independent set contains at most $2^{-(\Omega(2^n) - O(m \cdot n))}$ fraction of the vertices of that sub-graph. This implies that the chromatic number of \mathcal{G}' is at least $2^{\Omega(2^n) - O(m \cdot n)}$, and hence the protocol must transmit at least

$$\log \log \chi(\mathcal{G}_{\pi_1}) \geq \log \log \chi(\mathcal{G}') \geq n - O(\log(mn))$$

additional bits, as required.

The construction of \mathcal{G}' . Let $\mathcal{V} \subseteq \mathcal{V}_{\pi_1}$ be the set of balanced functions that is associated with π_1 . In order to construct \mathcal{G}' , we first show that, for every $g \in \mathcal{V}$, there exist strings $a_g \in \mathcal{A}_{\pi_1}(g)$ and $b_g \in \mathcal{B}_{\pi_1}(g)$, and a set of rows $I_g \subseteq [m]$, such that:

- The sets $\mathcal{X}_{\pi_1}(g, a_g)|_{I_g}$ and $\mathcal{Y}_{\pi_1}(g, b_g)|_{I_g}$ are prefix thick.
- $a_g|_{[m]-I_g} = b_g|_{[m]-I_g}$.

In order to construct such a triplet (a_g, b_g, I_g) , we apply Lemma 1.11 to show that, for every a and b , there are *many* subsets $I \subseteq [n]$ such that $\mathcal{X}_{\pi_1}(g, a)|_I$ and $\mathcal{Y}_{\pi_1}(g, b)|_I$ are prefix thick. Here, we use the assumption that $\mathcal{X}_{\pi_1}(g, a)$ and $\mathcal{Y}_{\pi_1}(g, b)$ are large. We then deduce the existence of a_g, b_g, I_g that satisfy the above properties using the assumption that it is hard to solve $\text{CC}(KW_f)$ on inputs from the set $\mathcal{A}_{\pi_1}(g) \times \mathcal{B}_{\pi_1}(g)$, combined with the fortification theorem of [DM16].

Having found such a triplet (a_g, b_g, I_g) for each $g \in \mathcal{V}$, we construct \mathcal{G}' as follows: Fix (a, b, I) to be the most popular triplet among all the triplets (a_g, b_g, I_g) . We define \mathcal{V}' to be the set of all functions $g \in \mathcal{V}$ whose triplet (a_g, b_g, I_g) is equal to (a, b, I) , and define \mathcal{G}' to be the sub-graph of \mathcal{G}_{π_1} induced by \mathcal{V}' .

It remains to upper bound the size of independent sets in \mathcal{G}' . In what follows, we use the convenient abbreviations $\mathcal{X}_g^* = \mathcal{X}_{\pi_1}(g, a)|_I$, and $\mathcal{Y}_g^* = \mathcal{Y}_{\pi_1}(g, b)|_I$. Observe that two functions $g_A, g_B \in \mathcal{V}'$ are neighbors in \mathcal{G}' if $\mathcal{X}_{g_A}^* \cap \mathcal{Y}_{g_B}^* \neq \emptyset$ or $\mathcal{X}_{g_B}^* \cap \mathcal{Y}_{g_A}^* \neq \emptyset$: To see it, observe that if $\mathcal{X}_{g_A}^* \cap \mathcal{Y}_{g_B}^* \neq \emptyset$ then there exist matrices $X \in \mathcal{X}_{g_A}^*$ and $Y \in \mathcal{Y}_{g_B}^*$ such that $X_i = Y_i$ for every $i \in I$. Moreover, for every $i \in [m] - I$ it holds that $a_i = b_i$. Hence, the strings a, b and the matrices X, Y satisfy the requirements of the weak intersection property, so g_A and g_B are neighbors in \mathcal{G}' .

In addition, in order to streamline the notation, from now on we make the simplifying assumption that a is the all-ones vector and that b is the all-zeroes vector. In particular, it holds that $g^{-1}(a) = g^{-1}(1)^m$ and $g^{-1}(b) = g^{-1}(0)^m$. In making this assumption, we allow ourselves to ignore the fact that $a|_{[m]-I} = b|_{[m]-I}$, since we will not need this fact again in this overview. We stress that the actual proof does not use this simplifying assumption.

Why is \mathcal{G}' not a clique? At first glance, it may seem that \mathcal{G}' is, in fact, a clique, by the following (flawed) reasoning: For every $g_A, g_B \in \mathcal{V}'$, the sets $\mathcal{X}_{g_A}^*$ and $\mathcal{Y}_{g_B}^*$ are prefix thick by definition, so by Proposition 1.10 it holds that $\mathcal{X}_{g_A}^* \cap \mathcal{Y}_{g_B}^* \neq \emptyset$. Therefore g_A and g_B are neighbors.

The error in the foregoing reasoning is that it deduces that the sets $\mathcal{X}_{g_A}^*$ and $\mathcal{Y}_{g_B}^*$ intersect from the fact that they are prefix thick. To see why this is a mistake, let us consider again our construction of the triplet (a, b, I) . When we applied Lemma 1.11 to $\mathcal{X}_{\pi_1}(g_A, a)$ and $\mathcal{Y}_{\pi_1}(g_B, b)$, we did so based on the assumption that those sets have large density *within the sets* $g_A^{-1}(a) = g_A^{-1}(1)^m$ and $g_B^{-1}(b) = g_B^{-1}(0)^m$ respectively. In particular, we actually viewed $\mathcal{X}_{\pi_1}(g_A, a)$ and $\mathcal{Y}_{\pi_1}(g_B, b)$ as sets of *strings over the “alphabets”* $g_A^{-1}(1)$ and $g_B^{-1}(0)$ respectively. Therefore, while the sets $\mathcal{X}_{g_A}^*$ and $\mathcal{Y}_{g_B}^*$ are indeed prefix thick, *they satisfy this property over different alphabets*. This means that we cannot use Proposition 1.10 to deduce that these sets intersect. Indeed, if $g_A = g_B$, then these sets cannot possibly intersect.

This issue makes it difficult to argue that any two particular vertices g_A and g_B in \mathcal{G}' are neighbors. Our third main technical contribution, described next, is to show how to overcome this difficulty.

Upper bounding the size of independent sets of \mathcal{G}' . We turn to upper bound the size of independent sets in \mathcal{G}' . Consider any set of vertices $\mathcal{S} \subseteq \mathcal{V}'$ that is “too large”. Let g_A and g_B be uniformly distributed functions in \mathcal{S} . We prove that g_A and g_B are neighbors with non-zero probability, and this will imply that \mathcal{S} cannot be an independent set.

We denote the alphabets $\Sigma_A = g_A^{-1}(1)$ and $\Sigma_B = g_B^{-1}(0)$, so $g_A^{-1}(a) = \Sigma_A^m$ and $g_B^{-1}(b) = \Sigma_B^m$. We also let $\Sigma_{AB} = \Sigma_A \cap \Sigma_B$, and let

$$\mathcal{U} = g_A^{-1}(a)|_I \cap g_B^{-1}(b)|_I = (\Sigma_{AB})^I.$$

Now, recall that the sets $\mathcal{X}_{g_A}^*$ and $\mathcal{Y}_{g_B}^*$ are prefix thick over Σ_A and Σ_B respectively. Unfortunately, as explained above, this prefix thickness does not immediately imply that these sets intersect, since they are prefix thick over different alphabets.

The key observation is that, with non-zero probability, *the sets $\mathcal{X}_{g_A}^* \cap \mathcal{U}$ and $\mathcal{Y}_{g_B}^* \cap \mathcal{U}$ are both prefix thick over the alphabet Σ_{AB}* . Since they are both prefix thick over the same alphabet, it follows that the sets $\mathcal{X}_{g_A}^* \cap \mathcal{U}$ and $\mathcal{Y}_{g_B}^* \cap \mathcal{U}$ must intersect. Hence, the sets $\mathcal{X}_{g_A}^*$ and $\mathcal{Y}_{g_B}^*$ intersect with non-zero probability, implying that g_A and g_B are neighbors.

The reason that $\mathcal{X}_{g_A}^* \cap \mathcal{U}$ and $\mathcal{Y}_{g_B}^* \cap \mathcal{U}$ are prefix thick over Σ_{AB} with high probability is that, since \mathcal{S} is relatively large, the sets Σ_A, Σ_B are good samplers. In other words, for every set $W \subseteq \Sigma_A$, the set $W \cap \Sigma_B$ has roughly the same density inside $\Sigma_{AB} = \Sigma_A \cap \Sigma_B$ as the set W inside Σ_A with high probability, and the same holds if we exchange Σ_A with Σ_B . We use this property to show that with sufficiently high probability, for every vertex v in the prefix tree of $\mathcal{X}_{g_A}^*$, the density of v 's children in Σ_{AB} is roughly the same as their density in Σ_A , and the same goes for $\mathcal{Y}_{g_B}^*$. Since the prefix trees of $\mathcal{X}_{g_A}^*$ and $\mathcal{Y}_{g_B}^*$ have minimum degree greater than $\frac{1}{2} \cdot |\Sigma_A|$ and $\frac{1}{2} \cdot |\Sigma_B|$ respectively, it follows that the corresponding trees for $\mathcal{X}_{g_A}^* \cap \mathcal{U}$ and $\mathcal{Y}_{g_B}^* \cap \mathcal{U}$ have minimum degree greater than $\frac{1}{2} |\Sigma_{AB}|$, as required.

Remark 1.12. In the formal proof, we actually require prefix trees of $\mathcal{X}_{g_A}^*$ and $\mathcal{Y}_{g_B}^*$ to have minimum degree greater than $(\frac{1}{2} + \varepsilon) \cdot |\Sigma_A|$ and $(\frac{1}{2} + \varepsilon) \cdot |\Sigma_B|$ respectively, in order to guarantee that the minimum degrees remain larger than $\frac{1}{2} \cdot |\Sigma_{AB}|$ after the intersection with \mathcal{U} .

Remark 1.13. As mentioned in Remark 1.9, in retrospect we realized that it is also possible to lower bound the clique number of \mathcal{G}_{π_1} . To this end, observe that last proof actually shows that the graph \mathcal{G}' is very dense. In other words, the graph \mathcal{G}' has a very large average degree. We now remove the vertices whose degree is slightly smaller than the average degree, and obtain a sub-graph \mathcal{G}'' that has a very large minimum degree. Finally, we construct a clique in \mathcal{G}'' using an iterative greedy algorithm: in each iteration, we add an arbitrary vertex v of \mathcal{G}'' to the clique, and remove from \mathcal{G}'' all the vertices that are not neighbors of v . The clique that is constructed in this way is sufficiently large to imply the structure theorem.

1.2.6 Comparison with previous works

We conclude this section by comparing our techniques with the techniques of the previous works on the subject. As explained above, our proof shares a common proof strategy with several previous works [EIRS91, DM16, KM18, dRMN⁺20]. Recall that in this strategy, we first define a notion of a “live” transcript, in which the protocol is still sufficiently far from solving KW_f on a and b , and in which not too much information has been revealed on the inputs of the players. We then prove that every protocol must have a live transcript of length $\approx \text{CC}(KW_f)$. Finally, we prove a structure theorem, which says that when the protocol reaches a live transcript, it still has to communicate $\approx \text{CC}(KW_g)$ bits. Together, the two last items imply that every protocol must communicate at least $\approx \text{CC}(KW_f) + \text{CC}(KW_g)$.

Our work also shares with the latter works the following general strategy for proving the structure theorem:

1. Given a live transcript π_1 , we measure the amount of information that π_1 reveals on each row of X and Y individually.
2. We partition the rows to “revealed rows”, on which π_1 reveals a lot of information, and to “unrevealed rows”, on which π_1 reveals only a little information.
3. We argue that since the total *amount* of information that π_1 reveals on X and Y is not too large, the number of revealed rows must be relatively small.

4. We restrict the protocol to inputs in which the column vectors a and b agree on the revealed rows. Showing that we can afford to do it is non-trivial, and relies on the assumption that in π_1 , the protocol is still far from solving KW_f , as well as on the fact that there are not too many revealed rows (as we showed in the previous step).
5. We force the protocol to output a solution (i, j) for which $a_i \neq b_i$. This implies in particular that the solution must belong to the unrevealed rows (since we forced the revealed rows to satisfy $a_i = b_i$ in the previous step).
6. We argue that since π_1 reveals only little information on the unrevealed rows, the protocol must spend $\approx \text{CC}(KW_g)$ bits in order to find a solution in such rows.

In particular, Steps 5 and 6 correspond to what we referred to in Section 1.1 as “the first and second obstacles” respectively. The main differences between the various proofs are the ways in which they measure information and implement the above steps. We now describe the analogues of the above steps in our proof:

- The information revealed on individual rows is measured using the notion of “prefix-thickness”. In particular, the unrevealed rows are those in which the sets \mathcal{X} and \mathcal{Y} are prefix-thick. More specifically, the set of unrevealed rows is the set we denoted by I in the construction of \mathcal{G}' in Section 1.2.5.
- Step 3 is implemented via Lemma 1.11, which implies that there is a large prefix-thick set of rows (in other words, there are many unrevealed rows).
- Step 4 is implemented in the construction of \mathcal{G}' in Section 1.2.5. Specifically, when we define \mathcal{G}' to be the induced subgraph of vertices that correspond to a fixed triplet (a, b, I) , with I being the set of unrevealed rows, we in fact restrict the protocol to inputs satisfying $a|_{[m]-I} = b|_{[m]-I}$.
- Step 5 is given to us for free by the definition of strong composition. Indeed, as discussed in Section 1.1, the whole motivation for the notion of strong composition is that we would like to avoid dealing with this obstacle for now.
- Step 6 is implemented by upper bounding the size of independent sets in \mathcal{G}' (as described in Section 1.2.5), and then deducing a lower bound on the communication complexity from the chromatic number of \mathcal{G}' (using Lemma 1.8).

Taking the above perspective, we believe that the this work makes three important contributions to the literature on the KRW conjecture, which we discuss in detail below.

Implementing Step 6 for MUX . The first contribution is developing a technique for implementing the above Step 6 for $KW_f \diamond MUX$. In the aforementioned works, the implementation of this step is tailored to the choice of the inner relation KW_g and its specific properties, whether it was the universal relation [EIRS91, KM18], the KW relation of the parity function [DM16], or a monotone KW relation with a certain “lifted” structure [dRMN⁺20]. In our case, the inner relation is the multiplexor relation MUX , and for years, it has not been clear how we can implement Step 6 for it.

While some earlier works suggested ideas for how such implementation might be carried out [EIRS91, Mei20], the first tangible progress on this question was made by [MS21] in their proof of the KRW conjecture for $U \diamond MUX$ (where U is the universal relation). Unfortunately, their proof as a whole seemed tailored to the case where the outer relation is the universal relation, and it was

not clear how to adapt it to the case where the outer relation is a KW relation KW_f . In particular, their proof did not use the above proof strategy for proving the structure theorem.

Our contribution in this regard consists of both extracting an important technique that was implicit in the proof [MS21] (namely, the “intersection property”, Lemma 1.7), and strengthening it to work with the chromatic number rather than the clique number (Lemma 1.8). This technique is crucial for carrying out Step 6 in the above proof strategy for MUX , and we believe it might be useful for future works.

Bypassing the limitations of existing information measures. The previous works [EIRS91, KM18, dRMN⁺20] measure the information that π_1 reveals on individual rows of X and Y using a notion called *average degree* or *unpredictability* [EIRS91, RM97]. Using this measure of information, [KM18] showed that it is possible to guarantee that π_1 reveals as little as 2 bits of information per unrevealed row. In fact, if one is willing to prove a lower bound that is quantitatively weaker than theirs, this can be pushed down to $(1 + \varepsilon)$ bits of information per unrevealed row. Unfortunately, this technique suffers from two limitations that prevent us from using it in our context:

- Having $(1 + \varepsilon)$ bits of information per unrevealed row is too much: we can only afford less than one bit of information per unrevealed row. To see why, recall that in order to establish the weak intersection property, we want to show that there exist matrices X and Y that agree on the unrevealed rows. In other words, we would like to show that the supports of these rows intersect. Now, in order to make sure, for example, that the supports of X_1 and Y_1 intersect, we need to guarantee that each of these supports contains more than half of all the strings — that is, that less than one bit of information was revealed about X_1 and Y_1 .
- The notion of average degree is inherently an “average-case” notion of information, while our proof seems to require a “worst-case” notion. To see why we seem to need a worst-case notion, recall again that our goal is to force equality in the unrevealed rows of X and Y . If the unrevealed rows are, say, the first two rows, then in order to force them both to be equal in X and Y , we guarantee that π_1 reveals little information about X_2 and Y_2 *even conditioned on $X_1 = Y_1$* . The notion of average degree cannot provide us with such a guarantee.

With regard to the second limitation, we note that worst-case notions of information exist in the literature. For start, [RM97] introduced the notion the “thick” sets, which is a worst-case version of average degree (the name “prefix thick” was chosen to allude to their notion). Moreover, [DM16] measured information using min-entropy, which is a worst-case notion. Nevertheless, the two notions lead to much weaker quantitative statements: they only provide a bound of $O(\log m)$ bits of information per revealed row, which is far from what we can afford.

Our contribution here is identifying a new way to measure information that does provide what we need, namely, prefix-thick sets. These sets guarantee that the support of every row covers more than half of all the strings (i.e. less than one bit revealed), conditioned on every assignment to the previous rows (i.e., a worst case guarantee). This is exactly the guarantee we were looking for, and we believe it might be of use to future works as well.

Equally importantly, in proving Lemma 1.11, we show that this notion of information can indeed be used — that is, we show that we can indeed find a large prefix-thick set of rows. While this lemma follows rather easily from a result of [ST14] (Lemma 5.6), we note that this result was discovered in another community (namely, the community of dynamical systems), and to the best of our knowledge it was not known in the complexity community. Thus, an additional contribution of this work is importing a new tool from a different community.

Conditioning on a and b . The notion of prefix-thick sets is important because it gives us a *worst-case* measure of information, but it is insufficient on its own. While, by definition, π_1 reveals less than one bit of information about each row in the prefix-thick set, we want this property to hold simultaneously with the property that $a_i = b_i$ holds for each revealed row. Nevertheless, if we force the equality $a_i = b_i$ on the revealed rows, this may leak additional information on the unrevealed rows, and violate the prefix-thick guarantee. Indeed, dealing with this issue is exactly where the argument of [KM18] loses the additional one bit of information.

In order to avoid this loss, we change our measure of information once more: instead of measuring the information revealed on X and Y on their own, we measure this information *conditioned on a and b* . This conditioning guarantees that when we force the equality $a_i = b_i$ on the unrevealed rows, this forcing does not leak additional information about the unrevealed rows (since the information on the unrevealed rows was measured conditioned on the whole values of a and b anyway). Formally, this conditioning is implemented in the fact that we work with different sets $\mathcal{X}_{\pi_1}(a)$ and $\mathcal{Y}_{\pi_1}(b)$ for every choice of a and b , whereas previous works worked with single sets \mathcal{X}_{π_1} and \mathcal{Y}_{π_1} (without dependence on a and b) — see also the third item in the definition of a live transcript in Section 1.2.2. Using this conditioning to avoid the loss in information is another new idea of this work.

Unfortunately, this conditioning on a and b creates a new problem: now the unrevealed rows of X and Y are prefix-thick only *when conditioned on a and b* . In other words, the support of each unrevealed row X_i contains more than half of all the strings x such that $g(x) = a_i$ (rather than half of all the strings in $\{0, 1\}^n$). It is therefore unclear how to ensure that X and Y agree on the unrevealed rows when we are only given this weaker guarantee.

Our third main contribution is showing how to bypass this obstacle using the sampling properties of the sets $g^{-1}(1), g^{-1}(0)$, as described in Section 1.2.5 in the part on upper bounding the size of independent sets. While this idea is somewhat technical, it is a key idea that allows us to work with the information *conditioned on a and b* , and therefore allows us to break the barrier of $(1 + \varepsilon)$ bits per row. We therefore believe that this idea, too, may be of use to future works on the KRW conjecture.

Summary. Wrapping up, the key contributions of this work to the literature on the KRW conjecture are the abstraction and strengthening of the intersection technique of [MS21]; identifying the notion of prefix-thick sets as a useful way to define unrevealed rows and showing that they can be found; and showing how to use conditioning on a and b to reduce the amount of information per unrevealed row, while using sampling to mitigate the issues caused by this conditioning. We believe that these new techniques would be of use for future works on the subject.

Remark 1.14. The above description of the general proof strategy for the structure theorem is a bit inaccurate in the case of [EIRS91, KM18]. Specifically, in 5, they do not force the protocol to output a solution (i, j) such that $a_i \neq b_i$, but use a somewhat different argument of similar flavor.

1.3 The organization of the paper

We cover the required preliminaries in Section 2. Then, in Section 3, we state our structure theorem and use it to prove our main theorem (Theorem 1.6). We explain our method for proving multiplexor lower bounds using the chromatic number in Section 4, and introduce the notion of prefix-thick sets in Section 5. We then prove the structure theorem in Section 6. We conclude by showing, in Section 7, an example that demonstrates that improving the constant γ to 0.64 or above would require new ideas.

2 Preliminaries

All the logarithms in this paper are of base 2. For any $n \in \mathbb{N}$, we denote by $[n]$ the set $\{1, \dots, n\}$. If S is a subset of some universe \mathcal{U} , we refer to the fraction $|S| / |\mathcal{U}|$ as *the density of S (within \mathcal{U})*.

Given an alphabet Σ and a set $I \subseteq [n]$, we denote by Σ^I the set of strings of length $|I|$ whose coordinates are indexed by I . Given a string $w \in \Sigma^n$ and a set $I \subseteq [n]$, we denote by $w|_I \in \Sigma^I$ the projection of w to the coordinates in I . Given a set of strings $\mathcal{W} \subseteq \Sigma^n$ and a set $I \subseteq [n]$, we denote by $\mathcal{W}|_I$ the set of projections of strings in \mathcal{W} to I . We denote by \circ the concatenation operator of strings, and denote by $|w|$ the length of the string w .

We denote by $\{0, 1\}^{m \times n}$ the set of Boolean $m \times n$ matrices, and for a set $I \subseteq [m]$, we denote by $\{0, 1\}^{I \times n}$ the set of $|I| \times n$ matrices whose entries are indexed by $I \times [n]$. Given a matrix $X \in \{0, 1\}^{m \times n}$ and a set $I \subseteq [m]$, we denote by $X|_I \in \{0, 1\}^{I \times n}$ the projection of X to the rows in I . Here, too, we extend this notation to sets of matrices $\mathcal{W} \subseteq \{0, 1\}^{m \times n}$. We denote by $X_i \in \{0, 1\}^n$ the i -th row of X .

For a function $g : \{0, 1\}^n \rightarrow \{0, 1\}$ and a matrix $X \in \{0, 1\}^{m \times n}$, we denote by $g(X)$ the column vector that is obtained by applying g to every row of X . Given a column vector $a \in \{0, 1\}^m$, we denote by $g^{-1}(a)$ the set of all matrices $X \in \{0, 1\}^{m \times n}$ such that $g(X) = a$.

The *clique number* of a graph G , denoted $\omega(G)$, is the size of the largest clique in G . The *independence number* of a graph G , denoted $\alpha(G)$, is the size of the largest independent set in G . The *chromatic number* of a graph G , denoted $\chi(G)$, is the the least number of colors required to color the vertices of G . It is easy to prove that if a graph G has n vertices, then $\chi(G)$ is lower bounded by both $\omega(G)$ and $n/\alpha(G)$.

We use the following standard corollaries of the Chernoff-Hoeffding bounds for independent random variables and negatively-associated random variables (see, e.g., Theorems 1.1 and 3.1 in Sections 1.6 and 3.1 of [DP09] respectively for the relevant versions of the Chernoff-Hoeffding bounds):

Fact 2.1. *Let S be a uniformly distributed subset of $[n]$. For every $\beta > 0$ it holds that*

$$\Pr \left[|S| < \left(\frac{1}{2} - \beta \right) \cdot n \right] < 2^{-2 \log e \cdot \beta^2 \cdot n}$$

Fact 2.2. *Let T be a subset of some universe \mathcal{U} , let $k \in \mathbb{N}$, and let S be a uniformly distributed subset of \mathcal{U} of size k . For every $\beta > 0$ it holds that*

$$\Pr \left[\frac{|S \cap T|}{|S|} < \frac{|T|}{|\mathcal{U}|} - \beta \right] < 2^{-2 \log e \cdot \beta^2 \cdot k}$$

and

$$\Pr \left[\frac{|S \cap T|}{|S|} > \frac{|T|}{|\mathcal{U}|} + \beta \right] < 2^{-2 \log e \cdot \beta^2 \cdot k}$$

We also use the binary entropy function, and the following standard approximation that it gives for the binomial coefficients.

Notation 2.3. The *binary entropy function* $H_2 : [0, 1] \rightarrow [0, 1]$ is the function that maps every $0 < p < 1$ to

$$H_2(p) = p \cdot \log \frac{1}{p} + (1 - p) \cdot \log \frac{1}{1 - p},$$

and maps 0 and 1 to 0.

Fact 2.4 (see, e.g., [CT91, Example 12.1.3]). *For every $k, n \in \mathbb{N}$ such that $k \leq n$ it holds that*

$$\frac{1}{n+1} \cdot 2^{H_2(\frac{k}{n}) \cdot n} \leq \binom{n}{k} \leq 2^{H_2(\frac{k}{n}) \cdot n}.$$

2.1 Depth complexity and formula complexity

In this section, we cover the basics of depth complexity. In the introduction, we defined the depth complexity $D(f)$ of a function f as the minimum depth of a circuit computing f with fan-in 2. For our purposes in this paper, however, it is more convenient to use an equivalent definition: namely, the minimum depth of a (de Morgan) formula that computes f .

Definition 2.5. A (*de Morgan*) formula ϕ is a rooted binary tree, whose leaves are identified with literals of the forms x_i and $\neg x_i$, and whose internal vertices are labeled as AND (\wedge) or OR (\vee) gates. Here, the same literal can be associated with more than one leaf. Such a formula ϕ over variables x_1, \dots, x_n computes a function from $\{0, 1\}^n$ to $\{0, 1\}$ in the natural way. The *size* of a formula is the number of its *leaves* (which is the same as the number of its gates up to a factor of 2). The *depth* of a formula is the depth of the tree.

Definition 2.6. The *formula complexity* of a Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, denoted $L(f)$, is the minimal size of a formula that computes f . The *depth complexity* of f , denoted $D(f)$, is the minimal depth of a formula that computes f . If f is a constant function, then we define $L(f) = 0$ and $D(f) = -\infty$.

In what follows, we generalize the latter definition from functions to promise problems, which will be useful when we discuss Karchmer-Wigderson relations later.

Definition 2.7. Let $\mathcal{X}, \mathcal{Y} \subseteq \{0, 1\}^n$ be disjoint sets. We say that a formula ϕ *separates* \mathcal{X} and \mathcal{Y} if $\phi(\mathcal{X}) = 1$ and $\phi(\mathcal{Y}) = 0$. The *formula complexity of the rectangle* $\mathcal{X} \times \mathcal{Y}$, denoted $L(\mathcal{X} \times \mathcal{Y})$, is the size of the smallest formula that separates \mathcal{X} and \mathcal{Y} . The *depth complexity of the rectangle* $\mathcal{X} \times \mathcal{Y}$, denoted $D(\mathcal{X} \times \mathcal{Y})$, is the smallest depth of a formula that separates \mathcal{X} and \mathcal{Y} . If \mathcal{X} or \mathcal{Y} are empty, we define $L(\mathcal{X} \times \mathcal{Y}) = 0$ and $D(\mathcal{X} \times \mathcal{Y}) = -\infty$.

Note that Definition 2.6 is indeed a special case of Definition 2.7 where $\mathcal{X} = f^{-1}(1)$ and $\mathcal{Y} = f^{-1}(0)$. Formula complexity has the following useful sub-additivity property.

Proposition 2.8. *Let \mathcal{X}, \mathcal{Y} be disjoint subsets of $\{0, 1\}^n$. Then, for every two partitions $\mathcal{X} = \mathcal{X}_0 \cup \mathcal{X}_1$ and $\mathcal{Y} = \mathcal{Y}_0 \cup \mathcal{Y}_1$, it holds that*

$$\begin{aligned} L(\mathcal{X} \times \mathcal{Y}) &\leq L(\mathcal{X}_0 \times \mathcal{Y}) + L(\mathcal{X}_1 \times \mathcal{Y}) \\ L(\mathcal{X} \times \mathcal{Y}) &\leq L(\mathcal{X} \times \mathcal{Y}_0) + L(\mathcal{X} \times \mathcal{Y}_1). \end{aligned}$$

We also use the following standard upper bound about the formula complexity of parity, which is obtained by computing the function in the natural way.

Proposition 2.9. *The formula complexity of the parity function over n bits is at most $4n^2$.*

Proof. First, observe that the parity of two bits can be computed by a formula of size 4 at follows:

$$x_1 \oplus x_2 = (x_1 \wedge \neg x_2) \vee (\neg x_1 \wedge x_2).$$

Next, observe that for every $k \in \mathbb{N}$, the parity of 2^k variables can be computed recursively, by first computing the parity of the first and last 2^{k-1} variables separately, and then computing the parity of the resulting two bits. This idea leads to the following recursive formula:

$$\bigoplus_{i=1}^{2^k} x_i = \left(\bigoplus_{i=1}^{2^{k-1}} x_i \wedge \neg \bigoplus_{i=2^{k-1}+1}^{2^k} x_i \right) \vee \left(\neg \bigoplus_{i=1}^{2^{k-1}} x_i \wedge \bigoplus_{i=2^{k-1}+1}^{2^k} x_i \right).$$

It can be easily proved by induction that the size of this formula is at most $(2^k)^2$. Finally, for every natural number n , we can compute the parity of n variables by adding dummy variables (which are fixed to 0) so the total number of variables is the smallest power of two that is at least n . Since this power of two is at most $2n$, the size of the resulting formula is at most $(2n)^2 = 4n^2$. ■

2.2 Communication complexity

Let \mathcal{X} , \mathcal{Y} , and \mathcal{Z} be sets, and let $R \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$ be a relation. The *communication problem* [Yao79] that corresponds to R is the following: two players, Alice and Bob, get inputs $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, respectively. They would like to find $z \in \mathcal{Z}$ such that $(x, y, z) \in R$. To this end, they send bits to each other until they find z , but they would like to send as few bits as possible. The *communication complexity* of R is the minimal number of bits that is transmitted by a protocol that solves R . More formally, we define a protocol as a binary tree, in which every vertex represents a possible state of the protocol, and every edge represents a message that moves the protocol from one state to another:

Definition 2.10. A (*deterministic*) *protocol* from $\mathcal{X} \times \mathcal{Y}$ to \mathcal{Z} is a rooted binary tree with the following structure:

- Every vertex v of the tree is labeled by a rectangle $\mathcal{X}_v \times \mathcal{Y}_v$ where $\mathcal{X}_v \subseteq \mathcal{X}$ and $\mathcal{Y}_v \subseteq \mathcal{Y}$. The root is labeled by the rectangle $\mathcal{X} \times \mathcal{Y}$. The rectangle $\mathcal{X}_v \times \mathcal{Y}_v$ is the set of pairs of inputs that lead the players to the vertex v .
- Each internal vertex v is *owned* by Alice or by Bob. Informally, v is owned by Alice if it is Alice's turn to speak at state v , and same for Bob.
- The two outgoing edges of every internal vertex are labeled by 0 and 1 respectively.
- For every internal vertex v that is owned by Alice, the following holds: Let v_0 and v_1 be the children of v associated with the out-going edges labeled with 0 and 1, respectively. Then,
 - $\mathcal{X}_v = \mathcal{X}_{v_0} \cup \mathcal{X}_{v_1}$, and $\mathcal{X}_{v_0} \cap \mathcal{X}_{v_1} = \emptyset$.
 - $\mathcal{Y}_v = \mathcal{Y}_{v_0} = \mathcal{Y}_{v_1}$.

Informally, when the players are at the vertex v , Alice sends 0 to Bob if her input is in \mathcal{X}_{v_0} and 1 if her input is in \mathcal{X}_{v_1} . An analogous property holds for vertices owned by Bob, while exchanging the roles of \mathcal{X} and \mathcal{Y} .

- Each leaf is labeled with a value $z_\ell \in \mathcal{Z}$. The value z_ℓ is the output of the protocol at ℓ .

We say that the protocol solves a relation $R \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$ if, for every leaf ℓ , it holds that $\mathcal{X}_\ell \times \mathcal{Y}_\ell \times \{z_\ell\} \subseteq R$.

Definition 2.11. Given a protocol Π and a vertex v of Π , the *transcript of v* is the string that is obtained by concatenating the labels of the edges on the path from the root to v . Intuitively, this string consists of the messages that Alice and Bob sent in their conversation until they got to v . Observe that the transcript determines the vertex v uniquely and vice versa, so we often identify the transcript with v . If v is a leaf of the protocol, we say that the transcript is a *full transcript*, and otherwise we say that it is a *partial transcript*.

Definition 2.12. The *communication complexity* of a (deterministic) protocol Π , denoted $\text{CC}(\Pi)$, is the depth of the protocol tree. In other words, it is the maximum number of bits that can be sent in an execution of the protocol on any pair of inputs (x, y) . The *(deterministic) communication complexity* of a relation R , denoted $\text{CC}(R)$, is the minimal communication complexity of a (deterministic) protocol that solves R .

Definition 2.13. We define the *size* of a protocol Π to be its number of leaves. The *protocol size* of a relation R , denoted $\text{L}(R)$, is the minimal size of a protocol that solves it (this is also known as the *protocol partition number* of R).

2.2.1 Non-deterministic protocols

Let $f : \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$, and consider the communication problem in which Alice and Bob compute $f(x, y)$ on inputs $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ respectively. In a non-deterministic protocol for f , there is an untrusted prover who knows both x and y , and whose goal is to convince Alice and Bob that $f(x, y) = 1$. The prover attempts to do so by sending a witness w from some witness set \mathcal{W} to both Alice and Bob. We require that the prover succeeds in convincing both Alice and Bob to accept if and only if it indeed holds that $f(x, y) = 1$, and define the complexity of the protocol as the number of bits that the prover sends.

Definition 2.14. Let $f : \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$, and let \mathcal{W} be a set. A *non-deterministic protocol* Π for f is a pair of functions $a_\Pi : \mathcal{X} \times \mathcal{W} \rightarrow \{0, 1\}$ and $b_\Pi : \mathcal{Y} \times \mathcal{W} \rightarrow \{0, 1\}$ such that for every $(x, y) \in \mathcal{X} \times \mathcal{Y}$ the following holds: $f(x, y) = 1$ if and only if there exists a witness $w \in \mathcal{W}$ such that both $a_\Pi(x, w) = 1$ and $b_\Pi(y, w) = 1$. The *complexity* of the protocol is $\log |\mathcal{W}|$.

Definition 2.15. Let $f : \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$. The *non-deterministic communication complexity* of f , denoted $\text{NCC}(f)$, is the minimal complexity of a non-deterministic protocol for f . The *co-non-deterministic communication complexity* of f , denoted $\text{coNCC}(f)$, is the non-deterministic communication complexity of the negation of f .

The following relationship between the non-deterministic communication complexities is an easy observation.

Proposition 2.16. *Let $f : \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$. Then,*

$$\text{NCC}(f) \leq 2^{\text{coNCC}(f)}.$$

Let \mathcal{X} be a any set. A standard example in communication complexity is the equality function $\text{EQ}_\mathcal{X} : \mathcal{X} \times \mathcal{X} \rightarrow \{0, 1\}$, which outputs 1 if and only if $x = y$. We denote its negation, the inequality function, by $\text{INEQ}_\mathcal{X}$. It is well known that $\text{NCC}(\text{EQ}_\mathcal{X}) = \log |\mathcal{X}|$ and that

$$\log \log |\mathcal{X}| \leq \text{coNCC}(\text{EQ}_\mathcal{X}) \leq \log \log |\mathcal{X}| + 1.$$

Remark 2.17. The above presentation of non-deterministic communication complexity is somewhat non-standard. Usually, non-deterministic communication complexity is defined as the number of monochromatic rectangles required to cover the inputs in $f^{-1}(1)$. Nevertheless, it is not hard to verify that the two definitions are equivalent.

2.3 Karchmer-Wigderson and multiplexor relations

In what follows, we provide some background on Karchmer-Wigderson relations and the related multiplexor relation. We first define KW relations for general rectangles, and then specialize the definition for functions.

Definition 2.18. Let $\mathcal{X}, \mathcal{Y} \subseteq \{0, 1\}^n$ be disjoint sets. The *KW relation* $KW_{\mathcal{X} \times \mathcal{Y}}$ is the communication problem in which Alice’s input is $x \in \mathcal{X}$, Bob’s input is $y \in \mathcal{Y}$, and they would like to find a coordinate $i \in [n]$ such that $x_i \neq y_i$. Note that such a coordinate i always exists, since $x \neq y$.

Definition 2.19. Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ be a non-constant function. The *KW relation of f* , denoted KW_f , is defined by $KW_f \stackrel{\text{def}}{=} KW_{f^{-1}(1) \times f^{-1}(0)}$.

KW relations are related to depth and formula complexity in the following way.

Theorem 2.20 ([KW88], see also [Raz90, KKN92, GMWW14]). *For every two disjoint sets $\mathcal{X}, \mathcal{Y} \subseteq \{0, 1\}^n$ it holds that $D(\mathcal{X} \times \mathcal{Y}) = \text{CC}(KW_{\mathcal{X} \times \mathcal{Y}})$, and $L(\mathcal{X} \times \mathcal{Y}) = L(KW_{\mathcal{X} \times \mathcal{Y}})$. In particular, for every non-constant $f : \{0, 1\}^n \rightarrow \{0, 1\}$, it holds that $D(f) = \text{CC}(KW_f)$, and $L(f) = L(KW_f)$.*

2.3.1 Fortification

Given a protocol solving a KW relation, we sometimes want to relate the amount of information that Alice and Bob transmit about their inputs to the decrease in the complexity of the KW relation. For example, we may want to argue that if Alice transmitted only one bit of information about her input, then the complexity of the KW relation decreased by only one bit. Unfortunately, this is not true in general. Nevertheless, [DM16] showed that every rectangle can be “fortified” such that it satisfies this property. Intuitively, we say that a rectangle $\mathcal{X} \times \mathcal{Y}$ is fortified if when Alice and Bob speak, the complexity decreases in proportion to the amount of information transmitted. Formally, we have the following definition and result.

Definition 2.21. Let $\rho > 0$. We say that a rectangle $\mathcal{X} \times \mathcal{Y}$ is ρ -fortified on Alice’s side if for every $\tilde{\mathcal{X}} \subseteq \mathcal{X}$ it holds that

$$\frac{L(\tilde{\mathcal{X}} \times \mathcal{Y})}{L(\mathcal{X} \times \mathcal{Y})} \geq \rho \cdot \frac{|\tilde{\mathcal{X}}|}{|\mathcal{X}|}.$$

Similarly, we say that $\mathcal{X} \times \mathcal{Y}$ is ρ -fortified on Bob’s side if the same holds for subsets $\tilde{\mathcal{Y}} \subseteq \mathcal{Y}$.

Theorem 2.22 (fortification theorem, [DM16]). *Let $\mathcal{X}, \mathcal{Y} \subseteq \{0, 1\}^n$ be disjoint sets. There exists a subset $\mathcal{X}' \subseteq \mathcal{X}$ such that $\mathcal{X}' \times \mathcal{Y}$ is $\frac{1}{4n}$ -fortified on Alice’s side, and such that $L(\mathcal{X}' \times \mathcal{Y}) \geq \frac{1}{4} \cdot L(\mathcal{X} \times \mathcal{Y})$. An analogous statement holds for Bob’s side.*

2.3.2 Multiplexor relations

As discussed in the introduction, for the most part of the paper we focus on proving lower bounds for a certain “multiplexor composition”. Informally, the multiplexor relation [EIRS91] is a KW relation KW_g in which the function g is given to the players as part of the input. It is often more convenient, however, to allow Alice and Bob to take as inputs (possibly distinct) functions $g_A, g_B : \{0, 1\}^n \rightarrow \{0, 1\}$ respectively, and allow them to reject if they detect that the promise $g_A = g_B$ was violated. It is not hard to see that modifying the relation in this way changes its complexity by at most a constant term.

Definition 2.23 ([EIRS91]). Let $n \in \mathbb{N}$. The (*same-function*) *multiplexor relation*, denoted MUX_n , is the following communication problem: Alice gets a function $g_A : \{0, 1\}^n \rightarrow \{0, 1\}$ and a string $x \in g_A^{-1}(1)$, and Bob gets a function $g_B : \{0, 1\}^n \rightarrow \{0, 1\}$ and a string $y \in g_B^{-1}(0)$. Their goal is to find a coordinate $i \in [n]$ such that $x_i \neq y_i$, and they are allowed to output the special symbol \perp if $g_A \neq g_B$.

It can be shown that $\text{CC}(MUX_n) = n + \Theta(1)$. The upper bound follows from a protocol of Tardos and Zwick [TZ97], and the lower bound follows from the techniques of [MS21] (see Lemma 33 there).

2.4 Half-duplex protocols

One of the important properties of protocols in the standard model of communication complexity is that, at any given point, the players know whose turn it is to speak. This is captured in Definition 2.10, for example, in the assumption that every vertex is owned by one of the players. Hoover et al. [HIMS18] considered a different model, called “half-duplex channels”, in which the players do not necessarily agree on whose turn it is to speak. Intuitively, a half-duplex channel models a walkie-talkie device: such a device has a “push to talk” button, where the user has to push the button in order to speak, and to release it in order to listen. In particular, if both sides try to speak simultaneously, the communication is lost.

A bit more formally, the communication in a half-duplex channel is executed in rounds. At each round, Alice should choose whether she would like to send a bit or to receive, and the same goes for Bob. If Alice chooses to receive and Bob chooses to send a bit $\sigma \in \{0, 1\}$, then Alice receives the bit σ , as in standard protocols (and vice versa). Such rounds are called *classical*. If both Alice and Bob choose to send bits, then the bits are lost, with Alice and Bob being none the wiser. Such rounds are called *wasted*. If both Alice and Bob choose to receive, the round is called *silent*, and there are multiple ways to define the behavior of the channel in such a round (see [HIMS18]). As in [MS21], we use the definition of a “half-duplex channel with adversary”: in silent rounds, Alice and Bob receive bits that are chosen by an adversary.

An interesting property of half-duplex protocols with adversary is that Alice and Bob do not necessarily agree on the “state” of the protocol. Put differently, each player has its own view of what has transpired so far, and these views are not necessarily consistent. Therefore, when defining such protocols formally, we define them as a pair of trees rather than one: each tree defines the behavior and viewpoint of one of the players.

Definition 2.24 (Half-duplex protocols with adversary [HIMS18]). A half-duplex protocol Π from $\mathcal{X} \times \mathcal{Y}$ to \mathcal{Z} is a pair of full 4-ary trees Π_A and Π_B with the following structure.

- Each vertex v of Π_A (respectively, Π_B) is labeled with a set $\mathcal{X}_v \subseteq \mathcal{X}$ (respectively, $\mathcal{Y}_v \subseteq \mathcal{Y}$). Intuitively, \mathcal{X}_v is the set of inputs that can reach the vertex v .
- The four outgoing edges of every internal vertex are labeled by **receive**(0), **receive**(1), **send**(0), and **send**(1) respectively. Intuitively, this label consists of the action that the player chooses to take and its outcome. For every vertex v of Π_A , denote by $v_{\text{receive}(0)}$, $v_{\text{receive}(1)}$, $v_{\text{send}(0)}$, and $v_{\text{send}(1)}$ the children of v that are associated with each of its outgoing edges.
- For every vertex v of Π_A , the following holds:

- $\mathcal{X}_{v_{\text{receive}(0)}} = \mathcal{X}_{v_{\text{receive}(1)}}$.
- $\mathcal{X}_v = \mathcal{X}_{v_{\text{receive}(0)}} \cup \mathcal{X}_{v_{\text{send}(0)}} \cup \mathcal{X}_{v_{\text{send}(1)}}$.
- The sets $\mathcal{X}_{v_{\text{receive}(0)}}$, $\mathcal{X}_{v_{\text{send}(0)}}$, and $\mathcal{X}_{v_{\text{send}(1)}}$ are pairwise disjoint.

Intuitively, when Alice is at the vertex v , Alice sends 0 to Bob if her input is in $\mathcal{X}_{v_{\text{send}(0)}}$, sends 1 if her input is in $\mathcal{X}_{v_{\text{send}(1)}}$, or receives if her input is in $\mathcal{X}_{v_{\text{receive}(0)}} = \mathcal{X}_{v_{\text{receive}(1)}}$. An analogous property holds for the vertices of Π_B , while exchanging the roles of \mathcal{X} and \mathcal{Y} .

- Each leaf ℓ is labeled with a value $z_\ell \in \mathcal{Z}$. Intuitively, z_ℓ is the output of the player at ℓ .

When given an input $(x, y) \in \mathcal{X} \times \mathcal{Y}$, the protocol executes the following algorithm. We initialize a pair of vertices (u, v) to the roots of Π_A and Π_B respectively. Then, at each round, the algorithm performs the following steps:

- For each $\sigma \in \{0, 1\}$, if $x \in \mathcal{X}_{u_{\text{send}(\sigma)}}$ and $y \in \mathcal{Y}_{v_{\text{receive}(0)}} = \mathcal{Y}_{v_{\text{receive}(1)}}$, then the pair (u, v) is updated to $(u_{\text{send}(\sigma)}, v_{\text{receive}(\sigma)})$. In this case, the round is called *classical*.
- The same as the previous step, but exchanging the roles of x, \mathcal{X}, u with y, \mathcal{Y}, v . In this case, too, the round is called *classical*.
- For every $\sigma, \tau \in \{0, 1\}$, if $x \in \mathcal{X}_{u_{\text{send}(\sigma)}}$ and $y \in \mathcal{Y}_{v_{\text{send}(\tau)}}$, then the pair (u, v) is updated to $(u_{\text{send}(\sigma)}, v_{\text{send}(\tau)})$. In this case the round is called *wasted*.
- If $x \in \mathcal{X}_{u_{\text{receive}(0)}} = \mathcal{X}_{u_{\text{receive}(1)}}$ and $y \in \mathcal{Y}_{v_{\text{receive}(0)}} = \mathcal{Y}_{v_{\text{receive}(1)}}$, then the pair (u, v) is updated to $(u_{\text{receive}(\sigma)}, v_{\text{receive}(\tau)})$ for some bits $\sigma, \tau \in \{0, 1\}$. We think of the choice of the bits σ, τ as being taken by an adversary. In this case the round is called *silent*.

We require that for every $(x, y) \in \mathcal{X} \times \mathcal{Y}$ and for every possible choices of the adversary, the vertices u and v reach leaves at the same round, at which point the algorithm halts. We also require that when the algorithm halts, both leaves are labeled with the same output. If, in a particular execution of the protocol, the algorithm reaches a pair of leaves (u, v) that are both labeled by an output $z \in \mathcal{Z}$, then we define z to be the output of that particular execution. We stress that the protocol may output different outputs in different executions on the same input, depending on the choices of the adversary. We say that the protocol *solves* a relation $R \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$ if for every input (x, y) , and in every possible execution of Π on (x, y) , the output z at that execution satisfies $(x, y, z) \in R$.

Remark 2.25. Note that the requirement that the vertices (u, v) always reach leaves at the same round implies that the trees Π_A, Π_B have the same depth.

Definition 2.26. The *communication complexity* of a half-duplex protocol Π , denoted $\text{CC}^{\text{hd}}(\Pi)$, is the depth of the trees Π_A, Π_B . The *half-duplex communication complexity* of a relation $R \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$, denoted $\text{CC}^{\text{hd}}(R)$, is the minimal communication complexity of a half-duplex protocol that solves it.

Observe that every standard deterministic protocol can be viewed as a half-duplex protocol, so $\text{CC}^{\text{hd}}(R) \leq \text{CC}(R)$. On the other hand, as noted in [HIMS18], every half-duplex protocol can be simulated by a standard protocol while incurring a factor of 2 in the complexity, and hence $\text{CC}(R) \leq 2 \cdot \text{CC}^{\text{hd}}(R)$. In short, $\text{CC}(R)$ and $\text{CC}^{\text{hd}}(R)$ are always within a factor of 2 of each other. Nevertheless, as in [HIMS18, MS21], saving this factor of 2 will be important in our application.

We have seen that the players in a half-duplex protocol may have different views. In particular, it follows that there is no well-defined transcript for the execution of the protocol, since Alice and Bob may see different transcripts. Nevertheless, given a transcript $\pi \in \{0, 1\}^*$ and an input x of Alice, we can still say whether π matches Alice's view in the protocol when given the input x . Informally, we say that $x \in \mathcal{X}$ is *consistent* with $\pi \in \{0, 1\}^*$ if there exists some execution of the protocol in which Alice gets the input x such that, for every round i :

- if Alice decides to send a bit on input x , then that bit is π_i ;
- and if Alice decides to receive on input x , then she receives the bit π_i .

More formally, we can define this notion as follows.

Definition 2.27. Fix a half-duplex protocol Π from $\mathcal{X} \times \mathcal{Y}$ to \mathcal{Z} , and let $\pi \in \{0, 1\}^*$. We say that a vertex v of depth $|\pi|$ in the tree Π_A is consistent with π if, for every $i \in [|\pi|]$, the i -th edge in the path from the root to v is labeled with either **receive**(π_i) or **send**(π_i). We say that an input $x \in \mathcal{X}$ is consistent with the π if there exists a vertex v of the tree Π_A that is consistent with π such that $x \in \mathcal{X}_v$. We define the notion that a vertex v in the tree Π_B or an input $y \in \mathcal{Y}$ are consistent with π similarly.

Notation 2.28. Let Π be a half-duplex protocol from $\mathcal{X} \times \mathcal{Y}$ to \mathcal{Z} , and let $\pi \in \{0, 1\}^*$. We denote by \mathcal{X}_π (respectively, \mathcal{Y}_π) the set of inputs $x \in \mathcal{X}$ (respectively $y \in \mathcal{Y}$) that are consistent with π . We say that π is a *transcript* of Π if and only if both \mathcal{X}_π and \mathcal{Y}_π are non-empty.

Remark 2.29. Observe that a transcript π contains the bits that were sent during the execution of the protocol, but does not register who sent them. Hence, there could be inputs x, y of Alice and Bob respectively that are consistent with the same transcript π , but for which Alice and Bob still have “different beliefs” about who sent each bit. For example, it could be the case that on inputs x, y both Alice and Bob send the bit π_1 in the first round.

In particular, recall that in the standard model of communication complexity, the transcript π determines a vertex v of the protocol, and vice versa. Moreover, the rectangle $\mathcal{X}_\pi \times \mathcal{Y}_\pi$ consists of exactly those pairs of inputs (x, y) on which the protocol reaches v . In the half-duplex model, on the other hand, different pairs (x, y) in the rectangle $\mathcal{X}_\pi \times \mathcal{Y}_\pi$ may lead the protocol to different pairs of vertices (u, v) . In fact, depending on the choices of the adversary, even the same pair (x, y) may lead the players to different pairs of vertices (u, v) . This means that some standard arguments in communication complexity require a lot more care when executed in the half-duplex model.

Nevertheless, it is important to note that a transcript π and a consistent input of Alice $x \in \mathcal{X}_\pi$ do determine together a vertex in Π_A . More formally, for every $x \in \mathcal{X}_\pi$, the vertex v corresponding to x and π in Definition 2.27 is unique. To see it, observe that the input x and the transcript π together determine the actions of the above algorithm on the vertex u , and hence determine the vertex that u will reach after $|\pi|$ steps.

2.4.1 Partially half-duplex protocol and multiplexors

Lower bounds for multiplexors in the half-duplex model imply the existence of hard KW relations in the standard model:

Lemma 2.30 ([MS21, Lemma 31]). *For every $n \in \mathbb{N}$ there exists a function $g : \{0, 1\}^n \rightarrow \{0, 1\}$ such that*

$$\text{CC}(KW_g) \geq \text{CC}^{\text{hd}}(\text{MUX}_n) - O(\log n).$$

This result may not seem so impressive on its own, since the existence of a hard KW relation follows immediately from counting arguments. Nevertheless, the importance of this result is that similar lemmas can also be proved for compositions of the form $KW_f \diamond KW_g$ — and this is exactly the kind of lower bounds we need in order to prove the weak KRW conjecture. In fact, this was the original motivation for introducing the half-duplex model (see discussion in [HIMS18]).

Unfortunately, the half-duplex model is rather complicated. In particular, the lower bounds we have for multiplexors in this model are worse than the ones we have in the standard model.

In order to remedy this issue, [MS21] introduced a weaker variant of the model, called “partially half-duplex protocols”, which is easier to analyze. This model is only applicable to a particular kind of communication problems, but the important thing is that it is applicable to multiplexors.

Definition 2.31 ([MS21, Lemma 32]). Let $R \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$. Suppose that there exists a set \mathcal{V} such that the inputs in \mathcal{X} and \mathcal{Y} are pairs of the forms (g_A, x) and (g_B, y) respectively where $g_A, g_B \in \mathcal{V}$. Let Π be a half-duplex protocol that solves R . We say that Π is *partially half-duplex* if whenever it is given an input $((g_A, x), (g_B, y))$ such that $g_A = g_B$, all the rounds in the execution of Π on the input are classical. The *partially half-duplex communication complexity* of R , denoted $\text{CC}^{\text{phd}}(R)$, is the minimal communication complexity of a partially half-duplex protocol that solves R .

[MS21] observed that lower bounds against partially half-duplex protocols are sufficient to imply the existence of hard KW relations.

Lemma 2.32 ([MS21]). *For every $n \in \mathbb{N}$ there exists a function $g : \{0, 1\}^n \rightarrow \{0, 1\}$ such that*

$$\text{CC}(KW_g) \geq \text{CC}^{\text{phd}}(\text{MUX}_n) - O(\log n).$$

In Section 3, we prove a similar lemma for the strong composition $KW_f \otimes \text{MUX}$, using exactly the same ideas as [MS21].

2.5 Linear codes and the Varshamov bound

Let C be a linear subspace of $\{0, 1\}^n$ of dimension k (where we identify $\{0, 1\}$ with the field \mathbb{F}_2), and let $d \in [n]$. We say that C is a (*linear*) *code* with distance d if every non-zero vector $c \in C$ has at least d non-zero coordinates (i.e., its *Hamming weight* is at least d). The following theorem, due to Varshamov, establishes the existence of linear codes with a good trade-off between the distance and the dimension.

Theorem 2.33 (Varshamov’s bound [Var57]). *For every $0 < \delta < \frac{1}{2}$, $\varepsilon > 0$, and for every sufficiently large $n \in \mathbb{N}$, there exists a linear code $C \subseteq \{0, 1\}^n$ with distance $\delta \cdot n$ and dimension at least $(1 - H(\delta) - \varepsilon) \cdot n$.*

3 Main theorem

In this section, we prove our main theorem, stated formally next.

Theorem 3.1. *There exists a constant $\gamma > 0.04$ such that the following holds: for every non-constant function $f : \{0, 1\}^m \rightarrow \{0, 1\}$ and for every $n \in \mathbb{N}$ there exists a function $g : \{0, 1\}^n \rightarrow \{0, 1\}$ such that*

$$\text{CC}(KW_f \otimes KW_g) \geq \log L(KW_f) - (1 - \gamma) \cdot m + n - O(\log(m \cdot n)). \quad (5)$$

The section is organized as follows: We start with reducing our task to proving lower bounds on the multiplexor composition $KW_f \otimes \text{MUX}$ in Section 3.1. Then, in Section 3.2, we define the notion of live transcripts and state our structure theorem. Finally, we derive our main theorem from the structure theorem in Section 3.3.

3.1 Reduction to multiplexor lower bounds

Informally, the multiplexor composition $KW_f \circledast MUX$ is a variant of the composition $KW_f \circledast KW_g$ in which the function g is given to the players as part of the input. In the formal definition, however, it is more convenient to allow Alice and Bob to take as inputs (possibly distinct) functions g_A and g_B respectively, and allow them to reject if they detect that $g_A \neq g_B$. We now define the multiplexor compositions $KW_f \circledast MUX$ and $KW_f \diamond MUX$ formally.

Definition 3.2. Let $f : \{0, 1\}^m \rightarrow \{0, 1\}$ be a non-constant function, and let $n \in \mathbb{N}$. The composition $KW_f \diamond MUX_n$ is the following communication problem: Alice gets as an input a function $g_A : \{0, 1\}^n \rightarrow \{0, 1\}$ and a matrix $X \in \{0, 1\}^{m \times n}$ such that $(f \diamond g_A)(X) = 1$. Bob gets as an input a function $g_B : \{0, 1\}^n \rightarrow \{0, 1\}$ and a matrix $Y \in \{0, 1\}^{m \times n}$ such that $(f \diamond g_B)(Y) = 0$. Their goal is to find an entry $(i, j) \in [m] \times [n]$ such that $X_{i,j} \neq Y_{i,j}$, and they are also allowed to output a special symbol \perp if $g_A \neq g_B$.

The strong composition $KW_f \circledast MUX_n$ is defined similarly, except that the entry (i, j) is also required to satisfy that $a_i \neq b_i$, where $a = g_A(X)$ and $b = g_B(Y)$.

The following lemma allows us to focus, for the rest of the paper, on proving lower bounds for $KW_f \circledast MUX$ in the partially half-duplex model.

Lemma 3.3. *Let $f : \{0, 1\}^m \rightarrow \{0, 1\}$ be a non-constant function, and let $n \in \mathbb{N}$. Then, there exists a function $g : \{0, 1\}^n \rightarrow \{0, 1\}$ such that*

$$\text{CC}(KW_f \diamond KW_g) \geq \text{CC}^{\text{phd}}(KW_f \diamond MUX_n) - \log(m \cdot n) - 3.$$

Similarly, there exists a function $g : \{0, 1\}^n \rightarrow \{0, 1\}$ such that

$$\text{CC}(KW_f \circledast KW_g) \geq \text{CC}^{\text{phd}}(KW_f \circledast MUX_n) - \log(m \cdot n) - 4.$$

Proof. Let $f : \{0, 1\}^m \rightarrow \{0, 1\}$ be a non-constant function, and let $n \in \mathbb{N}$. We prove the statement for $KW_f \circledast KW_g$, and the statement for $KW_f \diamond KW_g$ can be proved similarly. For every function $g : \{0, 1\}^n \rightarrow \{0, 1\}$, fix Π_g to be some optimal (standard) protocol that solves $KW_f \circledast KW_g$, and let c denote the maximal communication complexity of Π_g over all choices of the function g . We show that there exists a partially half-duplex protocol Π that solves $KW_f \circledast MUX_n$ with complexity $c + \lceil \log(m \cdot n) \rceil + 3$, and this will imply the required result.

The protocol Π works as follows. On input (g_A, X) , Alice simulates the protocol Π_{g_A} on the matrix X : specifically, Alice sends a bit whenever it is her turn to send a bit according to Π_A , and she receives whenever it is Bob's turn to send a bit. If Alice finishes in less than c rounds, then she receives until c rounds have passed. Similarly, on input (g_B, Y) , Bob simulates the protocol Π_{g_B} on the matrix Y , and then receives. Finally, after c rounds have passed, Alice and Bob have solutions (i_A, j_A) and (i_B, j_B) to their respective relations. Now, Alice sends (i_A, j_A) , a_{i_A} , and X_{i_A, j_A} to Bob (and Bob receives). Finally, Bob checks that $(i_A, j_A) = (i_B, j_B)$, $a_{i_A} \neq b_{i_A}$, and $X_{i_A, j_A} \neq Y_{i_A, j_A}$. If this is indeed the case, then Bob sends the bit 1 to Alice, and otherwise he sends 0 to Alice (meanwhile, Alice receives). If Bob sent 1, then the output of the protocol is (i_A, j_A) . If Bob sent 0, then the output of the protocol is \perp : the reason is that this case can only happen if Alice and Bob were simulating different protocols, meaning that $g_A \neq g_B$. Clearly, the complexity of Π is $c + \lceil \log(m \cdot n) \rceil + 3$. This concludes the proof. \blacksquare

Remark 3.4. The foregoing proof is a straightforward adaptation of the proofs in [MS21].

Remark 3.5. The composition $KW_f \diamond MUX_n$ was defined by the author in [Mei20]. The definition here is a bit different from the definition there: specifically, in the definition of [Mei20], each player gets m inner functions g_1, \dots, g_m rather than just one. Moreover, Alice and Bob are promised to receive the same inner functions (rather than allowing them to output \perp if they got different functions). This difference is not substantial, and the definition that we use in this paper is more suitable for our purposes.

3.2 The structure theorem

Intuitively, our structure theorem says that every efficient protocol that solves $KW_f \circledast MUX_n$ has roughly the same structure as the obvious protocol. More specifically, this theorem says that if, at a given point of the execution of the protocol, not much information was transmitted about g_A, g_B, X, Y , and the protocol is still sufficiently far from solving KW_f on a and b , then the protocol must transmit at least $n - O(\log(m \cdot n))$ more bits. A point in the execution that satisfies the latter requirements is called a *live transcript*. In what follows, we define live transcripts and state our structure theorem formally. We start by recalling some notation from the introduction.

Notation 3.6. Let $f : \{0, 1\}^m \rightarrow \{0, 1\}$, let $n \in \mathbb{N}$, and let Π be a partially half-duplex protocol that solves $KW_f \circledast MUX_n$. Let π be a (possibly partial) transcript of Π , and let \mathcal{X}_π and \mathcal{Y}_π be the sets of inputs that are consistent with π as in Notation 2.28. For every function $g : \{0, 1\}^n \rightarrow \{0, 1\}$ and strings $a \in f^{-1}(1)$ and $b \in f^{-1}(0)$, we denote

$$\begin{aligned} \mathcal{X}_\pi(g) &= \{X \in \{0, 1\}^{m \times n} : (g, X) \in \mathcal{X}_\pi\} & \mathcal{Y}_\pi(g) &= \{Y \in \{0, 1\}^{m \times n} : (g, Y) \in \mathcal{Y}_\pi\} \\ \mathcal{X}_\pi(g, a) &= \{X \in \mathcal{X}_\pi : g(X) = a\} & \mathcal{Y}_\pi(g, b) &= \{Y \in \mathcal{Y}_\pi : g(Y) = b\} \\ &= \mathcal{X}_\pi(g) \cap g^{-1}(a) & &= \mathcal{Y}_\pi(g) \cap g^{-1}(b). \end{aligned}$$

We also denote by $\mathcal{A}_\pi(g)$ and $\mathcal{B}_\pi(g)$ the sets of all strings $a \in f^{-1}(1)$ and $b \in f^{-1}(0)$ such that $\mathcal{X}_\pi(g, a)$ and $\mathcal{Y}_\pi(g, b)$ are non-empty, respectively. Finally, we denote by \mathcal{V}_π the set of all functions $g : \{0, 1\}^n \rightarrow \{0, 1\}$ such that both $\mathcal{A}_\pi(g)$ and $\mathcal{B}_\pi(g)$ are non-empty.

Notation 3.7. We denote by \mathcal{V}_0 be the set of all *balanced* functions from $\{0, 1\}^n$ to $\{0, 1\}$ (i.e., the functions that take the value 1 on exactly half of the inputs).

Definition 3.8. Let $\kappa \in \mathbb{N}$ be a universal constant to be fixed later ($\kappa \geq 8$ is enough), and let $\gamma > 0$. Let f , n , and Π be as in Notation 3.6. We say that a (possibly partial) transcript π_1 of Π is γ -*alive* if and only if there exists a set $\mathcal{V} \subseteq \mathcal{V}_{\pi_1} \cap \mathcal{V}_0$ of *balanced* functions that satisfies the following conditions:

- $|\mathcal{V}| \geq 2^{-m} \cdot |\mathcal{V}_0|$.
- For every $g \in \mathcal{V}$, it holds that $\log L(\mathcal{A}_{\pi_1}(g) \times \mathcal{B}_{\pi_1}(g)) \geq (1 - \gamma) \cdot m + \kappa \log m + \kappa$.
- For every $g \in \mathcal{V}$, $a \in \mathcal{A}_{\pi_1}(g)$, and $b \in \mathcal{B}_{\pi_1}(g)$, it holds that $|\mathcal{X}_{\pi_1}(g, a)| \geq 2^{-\gamma \cdot m + 1} \cdot |g^{-1}(a)|$ and $|\mathcal{Y}_{\pi_1}(g, b)| \geq 2^{-\gamma \cdot m + 1} \cdot |g^{-1}(b)|$.

When γ is clear from the context, we drop it and just write that π_1 is alive.

We are now ready to state our structure theorem.

Theorem 3.9 (structure theorem). *There exists a constant $\gamma > 0.04$ such that the following holds: Let $f : \{0, 1\}^m \rightarrow \{0, 1\}$, let $n \in \mathbb{N}$, and let Π be a partially half-duplex protocol that solves $KW_f \circledast MUX_n$. For every γ -live transcript π_1 of Π , the complexity of Π is at least*

$$|\pi_1| + n - O(\log(m \cdot n)).$$

3.3 Proof of main theorem from structure theorem

Let $\gamma > 0.04$ be the constant from Theorem 3.9, and let κ be the universal constant from Definition 3.8. Let $f : \{0, 1\}^m \rightarrow \{0, 1\}$ and let $n \in \mathbb{N}$. We may assume without loss of generality that $\log L(KW_f) > (1 - \gamma) \cdot m$, since otherwise the desired result follows easily from a counting argument. By Lemma 3.3, it suffices to prove that

$$\text{CC}^{\text{phd}}(KW_f \otimes \text{MUX}_n) \geq \log L(KW_f) - (1 - \gamma) \cdot m + n - O(\log(m \cdot n)),$$

and this will imply the desired result. Fix a partially half-duplex protocol Π that solves $KW_f \otimes \text{MUX}_n$. We prove that the communication complexity of Π is at least the right-hand side of the last equation. To this end, we construct a sufficiently long γ -live transcript π_1 , and then apply the structure theorem to π_1 .

In order to construct the live transcript π_1 , we first construct, for each $g \in \mathcal{V}_0$, a ‘‘candidate transcript’’ $\pi_{1,g}$ that satisfies the following properties:

1. $|\pi_{1,g}| = \log L(KW_f) - (1 - \gamma) \cdot m - \kappa \log m - \kappa$.
2. $g \in \mathcal{V}_{\pi_{1,g}}$.
3. $\log L(\mathcal{A}_{\pi_1}(g) \times \mathcal{B}_{\pi_1}(g)) \geq (1 - \gamma) \cdot m + \kappa \log m + \kappa$.
4. For every $a \in \mathcal{A}_{\pi_1}(g)$ and $b \in \mathcal{B}_{\pi_1}(g)$, it holds that $|\mathcal{X}_{\pi_1}(g, a)| \geq 2^{-\gamma \cdot m + 1} \cdot |g^{-1}(a)|$ and $|\mathcal{Y}_{\pi_1}(g, b)| \geq 2^{-\gamma \cdot m + 1} \cdot |g^{-1}(b)|$.

We then choose π_1 to be the most popular value of $\pi_{1,g}$ over all functions $g \in \mathcal{V}_0$, and choose \mathcal{V} to be the set of functions $g \in \mathcal{V}_0$ such that $\pi_{1,g} = \pi_1$. A simple counting argument then shows that $|\mathcal{V}| \geq 2^{-m} \cdot |\mathcal{V}_0|$, and hence π_1 is alive.

We turn to describing the construction of a single candidate transcript $\pi_{1,g}$. Fix a function $g \in \mathcal{V}_0$. In what follows, we abbreviate and write $\mathcal{A}_{\pi_{1,g}} = \mathcal{A}_{\pi_{1,g}}(g)$ and $\mathcal{X}_{\pi_{1,g}}(a) = \mathcal{X}_{\pi_{1,g}}(g, a)$, and similarly for $\mathcal{B}_{\pi_{1,g}}$ and $\mathcal{Y}_{\pi_{1,g}}$. Let Π_g be the protocol that is obtained from Π by hard-wiring the inputs of the players such that $g_A = g_B = g$. Observe that since the original protocol Π is *partially* half-duplex, it holds that all the rounds in the hardwired protocol Π_g are classical. In other words, we can think of Π_g as a standard deterministic protocol rather than as a half-duplex protocol. We choose the candidate transcript $\pi_{1,g}$ to be a transcript of Π_g , and construct it iteratively, bit-by-bit. Intuitively, at each iteration we choose the next bit of $\pi_{1,g}$ such that it transmits at most one bit of information about X, Y and decreases the complexity of solving KW_f on a and b by at most one bit. Formally, we initialize $\pi_{1,g}$ to the empty transcript, and then, in each iteration, we perform the following steps:

- Without loss of generality, assume that it is Alice’s turn to speak next in $\pi_{1,g}$.
- For every $a \in \mathcal{A}_{\pi_{1,g}}$ and $X \in \mathcal{X}_{\pi_{1,g}}(a)$, let $\sigma_{a,X}$ be the bit that Alice sends in Π_g at $\pi_{1,g}$ on input X .
- For every $a \in \mathcal{A}_{\pi_{1,g}}$, let σ_a be majority value of $\sigma_{a,X}$ over all $X \in \mathcal{X}_{\pi_{1,g}}(a)$.
- For each $\sigma \in \{0, 1\}$, let $\mathcal{A}_\sigma = \{a \in \mathcal{A}_{\pi_{1,g}} \mid \sigma_a = \sigma\}$. By the sub-additivity property of formula complexity, it holds that

$$L(\mathcal{A}_{\pi_{1,g}} \times \mathcal{B}_{\pi_{1,g}}) \leq L(\mathcal{A}_0 \times \mathcal{B}_{\pi_{1,g}}) + L(\mathcal{A}_1 \times \mathcal{B}_{\pi_{1,g}}).$$

In particular, there exists a bit $\sigma \in \{0, 1\}$ such that $L(\mathcal{A}_\sigma \times \mathcal{B}_{\pi_{1,g}}) \geq L(\mathcal{A}_{\pi_{1,g}} \times \mathcal{B}_{\pi_{1,g}})/2$.

- Append σ to $\pi_{1,g}$.
- If it is Bob's turn to speak next in $\pi_{1,g}$, perform the foregoing steps while exchanging $a \in \mathcal{A}_{\pi_{1,g}}$ with $b \in \mathcal{B}_{\pi_{1,g}}$ and $X \in \mathcal{X}_{\pi_{1,g}}(a)$ with $Y \in \mathcal{Y}_{\pi_{1,g}}(b)$.

We repeat these iterations until $\pi_{1,g}$ is either of length

$$\log L(f) - (1 - \gamma) \cdot m - \kappa \log m - \kappa \quad (6)$$

or a leaf of Π_g . It remains to show that $\pi_{1,g}$ satisfies the above properties of a candidate transcript.

First, observe that the transcript $\pi_{1,g}$ satisfies Property 2 (i.e., $g \in \mathcal{V}_{\pi_{1,g}}$) since it is a transcript of Π_g . Next, we show that $\pi_{1,g}$ satisfies Property 3. At the beginning of the construction, we have $\mathcal{A}_{\pi_{1,g}} = f^{-1}(1)$ and $\mathcal{B}_{\pi_{1,g}} = f^{-1}(0)$. In particular, it holds that $L(\mathcal{A}_{\pi_{1,g}} \times \mathcal{B}_{\pi_{1,g}}) = L(f)$. Observe that at each iteration of the construction, the formula complexity $L(\mathcal{A}_{\pi_{1,g}} \times \mathcal{B}_{\pi_{1,g}})$ is decreased by at most a factor of 2. Hence, at the end of the construction, it holds that

$$L(\mathcal{A}_{\pi_{1,g}} \times \mathcal{B}_{\pi_{1,g}}) \geq L(f) \cdot 2^{-|\pi_{1,g}|} \geq L(f) \cdot 2^{-(\log L(f) - (1-\gamma) \cdot m - \kappa \log m - \kappa)} \geq 2^{(1-\gamma) \cdot m + \kappa \log m + \kappa},$$

so $\pi_{1,g}$ indeed satisfies Property 3.

We now show that the transcript $\pi_{1,g}$ satisfies Property 1. To this end, it suffices to show that $\pi_{1,g}$ cannot be a leaf of Π_g . Suppose for the sake of contradiction that $\pi_{1,g}$ is a leaf of Π_g . By definition, Π_g is a protocol that solves $KW_f \otimes KW_g$, and therefore $\pi_{1,g}$ is labeled with a solution (i, j) for that relation. This means that for every $a \in \mathcal{A}_{\pi_{1,g}}$ and $b \in \mathcal{B}_{\pi_{1,g}}$, and for every $X \in \mathcal{X}_{\pi_{1,g}}(a)$ and $Y \in \mathcal{Y}_{\pi_{1,g}}(b)$, it should hold that $a_i \neq b_i$ and $X_{i,j} \neq Y_{i,j}$. Nevertheless, the requirement that $a_i \neq b_i$ for every $a \in \mathcal{A}_{\pi_{1,g}}$ and $b \in \mathcal{B}_{\pi_{1,g}}$ implies that i is a solution for KW_f on $\mathcal{A}_{\pi_{1,g}} \times \mathcal{B}_{\pi_{1,g}}$. In other words, the relation KW_f is solved on $\mathcal{A}_{\pi_{1,g}} \times \mathcal{B}_{\pi_{1,g}}$, and thus $L(\mathcal{A}_{\pi_{1,g}} \times \mathcal{B}_{\pi_{1,g}}) = 1$. This, however, contradicts the above lower bound on $L(\mathcal{A}_{\pi_{1,g}} \times \mathcal{B}_{\pi_{1,g}})$. It follows that $\pi_{1,g}$ cannot be a leaf.

Finally, we prove that $\pi_{1,g}$ satisfies Property 4. Observe that at the beginning of the construction, for every $a \in \mathcal{A}_{\pi_{1,g}}$ and $b \in \mathcal{B}_{\pi_{1,g}}$ it holds that $\mathcal{X}_{\pi_{1,g}}(a) = g^{-1}(a)$ and $\mathcal{Y}_{\pi_{1,g}}(b) = g^{-1}(b)$. Furthermore, observe that in each iteration of the construction, the sizes of the sets $\mathcal{X}_{\pi_{1,g}}(a)$ and $\mathcal{Y}_{\pi_{1,g}}(b)$ decrease by at most a factor of 2 for every $a \in \mathcal{A}_{\pi_{1,g}}$ and $b \in \mathcal{B}_{\pi_{1,g}}$. Moreover, the construction performs at most $\gamma \cdot m - \kappa \log m - \kappa$ iterations, since $L(f) \leq 2^m$ for every function f . It follows that at the end of the construction we have

$$\begin{aligned} |\mathcal{X}_{\pi_{1,g}}(a)| &\geq 2^{-\gamma \cdot m + 1} \cdot |g^{-1}(a)| \\ |\mathcal{Y}_{\pi_{1,g}}(b)| &\geq 2^{-\gamma \cdot m + 1} \cdot |g^{-1}(b)|. \end{aligned}$$

We conclude the proof by choosing the transcript π_1 of Π to be the most popular transcript $\pi_{1,g}$ over all the functions $g \in \mathcal{V}_0$. We prove that π_1 is γ -alive. First, we choose the set \mathcal{V} that is associated with π_1 to be the set of functions $g \in \mathcal{V}_0$ for which $\pi_{1,g} = \pi_1$. Since the transcripts $\pi_{1,g}$ are of the same length, and since this length is smaller than m , the number of those transcripts is at most 2^m . By an averaging argument, it follows that $|\mathcal{V}| \geq 2^{-m} \cdot |\mathcal{V}_0|$. Furthermore, since π_1 satisfies Properties 3 and 4 for every $g \in \mathcal{V}$, it also satisfies the other requirements of the definition of live transcripts. Hence, π_1 is γ -alive. By the structure theorem, the complexity of the protocol Π is at least

$$|\pi_1| + n - O(\log(m \cdot n)) \geq \log L(f) + n - (1 - \gamma) \cdot m - O(\log(m \cdot n)),$$

as required.

4 Multiplexor lower bounds via graph coloring

In this section, we develop a general method for proving lower bounds on multiplexor-composition problems, by extending a previous method of Mihajlin and Smal [MS21]. We first develop the method for the composition $KW_f \diamond MUX_n$, and then for the strong composition $KW_f \otimes MUX_n$. We note that in this paper we only apply the method for $KW_f \otimes MUX_n$. Nevertheless, we hope that the method will also be useful in the future for proving lower bounds on $KW_f \diamond MUX_n$, and therefore state and prove it for this relation as well.

Definition 4.1. Let Π be a partially half-duplex protocol that solves a relation $KW_f \diamond MUX_n$ for some $f : \{0, 1\}^m \rightarrow \{0, 1\}$ and $n \in \mathbb{N}$, and let π_1 be a (possibly partial) transcript of Π such that $\mathcal{V}_{\pi_1} \neq \emptyset$. The *characteristic graph* of π_1 , denoted \mathcal{G}_{π_1} , is the graph whose vertices are the functions in \mathcal{V}_{π_1} , and whose edges are defined as follows: two functions $g_A, g_B \in \mathcal{V}_{\pi_1}$ are neighbors in \mathcal{G}_{π_1} if and only if either $\mathcal{X}_{\pi_1}(g_A) \cap \mathcal{Y}_{\pi_1}(g_B) \neq \emptyset$ or $\mathcal{X}_{\pi_1}(g_B) \cap \mathcal{Y}_{\pi_1}(g_A) \neq \emptyset$.

Lemma 4.2. *Let Π, f, n be as in Definition 4.1. For every transcript π_1 of Π such that $\mathcal{V}_{\pi_1} \neq \emptyset$, the complexity of Π is at least*

$$|\pi_1| + \log \log \chi(\mathcal{G}_{\pi_1}) - \log \log \log \chi(\mathcal{G}_{\pi_1}) - 4.$$

As discussed in the introduction, [MS21] implicitly proved a similar result, in which the chromatic number $\chi(\mathcal{G}_{\pi_1})$ is replaced by the clique number $\omega(\mathcal{G}_{\pi_1})$. Since it is always the case that $\chi(\mathcal{G}_{\pi_1}) \geq \omega(\mathcal{G}_{\pi_1})$, our lemma generalizes their result. We now state the analogous definition and lemma for strong composition.

Definition 4.3. Let Π be a partially half-duplex protocol that solves a relation $KW_f \otimes MUX_n$ for some $f : \{0, 1\}^m \rightarrow \{0, 1\}$ and $n \in \mathbb{N}$, and let π_1 be a (possibly partial) transcript of Π such that $\mathcal{V}_{\pi_1} \neq \emptyset$. The *characteristic graph* of π_1 , denoted \mathcal{G}_{π_1} , is the graph whose vertices are the functions in \mathcal{V}_{π_1} , and in which two functions $g_A, g_B \in \mathcal{V}_{\pi_1}$ are neighbors if and only if they satisfy the following property:

- **Weak Intersection Property:** there exist matrices $X \in \mathcal{X}_{\pi_1}(g_A)$ and $Y \in \mathcal{Y}_{\pi_1}(g_B)$ such that $X_i = Y_i$ for every $i \in [m]$ for which $a_i \neq b_i$, where $a = g_A(X)$ and $b = g_B(Y)$ (or the same statement holds when exchanging g_A and g_B).

Lemma 4.4. *Let Π, f, n be as in Definition 4.3. For every transcript π_1 of Π such that $\mathcal{V}_{\pi_1} \neq \emptyset$, the complexity of Π is at least*

$$|\pi_1| + \log \log \chi(\mathcal{G}_{\pi_1}) - \log \log \log \chi(\mathcal{G}_{\pi_1}) - 4.$$

Example 4.5. If π_1 is the empty transcript, then the set \mathcal{V}_{π_1} consists of all functions from $\{0, 1\}^n$ to $\{0, 1\}$. In this case, it can be verified that a sufficient condition for two functions $g_A, g_B \in \mathcal{V}_{\pi_1}$ to be neighbors in \mathcal{G}_{π_1} is that $g_A^{-1}(\sigma) \cap g_B^{-1}(\tau) \neq \emptyset$ for every $\sigma, \tau \in \{0, 1\}$ (according to both Definitions 4.1 and 4.3). Therefore, the set of all balanced functions $g \in \mathcal{V}_{\pi_1}$ such that $g(\bar{1}) = 1$ is a clique in \mathcal{G}_{π_1} . This clique has at least $2^{2^n - O(n)}$ vertices, and therefore Lemmas 4.2 and 4.4 yield a lower bound of $n - O(\log n)$ in this case.

Example 4.6. Suppose that Π is a standard protocol that solves the relation $KW_f \diamond MUX_n$ (the case of $KW_f \otimes MUX_n$ is similar). Let π_1 be a full transcript of Π . Then, π_1 is labeled with an output that is either \perp or an entry (i, j) . We consider the two cases separately:

- If π_1 is labeled with \perp , then $\mathcal{V}_{\pi_1} = \emptyset$, and therefore Lemmas 4.2 and 4.4 are inapplicable. The reason is that, in this case, for every function $g : \{0, 1\}^n \rightarrow \{0, 1\}$, either $\mathcal{X}_{\pi_1}(g)$ or $\mathcal{Y}_{\pi_1}(g)$ must be empty: otherwise, there would have existed pairs $(g, X) \in \mathcal{X}_{\pi_1}$ and $(g, Y) \in \mathcal{Y}_{\pi_1}$, and the protocol cannot output \perp when given these inputs.
- Suppose that π_1 is labeled with an entry (i, j) . We claim that in this case the graph \mathcal{G}_{π_1} contains no edges. To see why, observe that there must exist a bit $\sigma \in \{0, 1\}$ such that for every $(g_A, X) \in \mathcal{X}_{\pi_1}$ and $(g_B, Y) \in \mathcal{Y}_{\pi_1}$ it holds that $X_{i,j} = \sigma$ and $Y_{i,j} = 1 - \sigma$: otherwise, there would have existed pairs $(g_A, X) \in \mathcal{X}_{\pi_1}$ and $(g_B, Y) \in \mathcal{Y}_{\pi_1}$ for which $X_{i,j} = Y_{i,j}$, and on those inputs the protocol cannot output (i, j) . Nevertheless, this implies that $\mathcal{X}_{\pi_1}(g_A) \cap \mathcal{Y}_{\pi_1}(g_B) = \emptyset$ for every two functions $g_A, g_B \in \mathcal{V}_{\pi_1}$, and therefore the graph \mathcal{G}_{π_1} does not contain any edges. It follows that $\chi(\mathcal{G}_{\pi_1}) = 1$, and hence Lemma 4.2 does not give a meaningful lower bound

In the rest of this section, we prove Lemmas 4.2 and 4.4. Mihajlin and Smal [MS21] proved their result by reduction from the co-non-deterministic communication complexity of the equality function EQ. We generalize their result by considering a certain “graph equality” problem, and using it as the starting point of the reduction. We define the graph equality problem and lower bound its co-non-deterministic complexity in Section 4.1. Then, we prove the above lemmas by reducing from that problem in Section 4.2.

4.1 The graph equality problem

De Wolf [dW01] studied the following promise version of the equality problem.

Definition 4.7. Fix a graph $G = (V, E)$. The *graph equality problem of G* , denoted GRAPHEQ_G , is the following promise problem: Alice and Bob get as inputs vertices $v_A, v_B \in V$ respectively. Their goal is to output 1 if $v_A = v_B$, and 0 if v_A and v_B are neighbors. If neither is the case, then the output of Alice and Bob may be either 0 or 1. We denote the negation of the problem by GRAPHINEQ_G .

Ilango, Loff, and Oliveira [ILO20] proved the following result on the non-deterministic complexity of GRAPHEQ_G .

Proposition 4.8 ([ILO20]). *For every graph $G = (V, E)$, it holds that $\text{NCC}(\text{GRAPHEQ}_G) = \log \chi(G)$.*

Using the latter result, we deduce the following bounds on the co-non-deterministic complexity of the graph equality problem.

Proposition 4.9. *For every graph $G = (V, E)$, it holds that*

$$\log \log \chi(G) \leq \text{NCC}(\text{GRAPHINEQ}_G) \leq \log \log \chi(G) + 1.$$

Proof. The lower bound follows from Proposition 4.8 by using the fact that the non-deterministic communication complexity of a problem is at most exponential in its co-non-deterministic communication complexity (Proposition 2.16). We prove the upper bound by reduction from GRAPHINEQ_G to $\text{INEQ}_{\lfloor \chi(G) \rfloor}$.

Consider the following non-deterministic protocol for GRAPHINEQ_G : Suppose that Alice and Bob get vertices v_A and v_B that are neighbors in G , and Merlin wants to convince them that $v_A \neq v_B$. We fix in advance some optimal coloring of G that is known to both players. Since v_A and v_B are neighbors, they must be colored with different colors. Merlin now proves to Alice

and Bob that the colors of v_A and v_B are different using an optimal non-deterministic protocol for $\text{INEQ}_{[\chi(G)]}$. Clearly, if v_A and v_B are neighbors then Merlin can always convince Alice and Bob to accept, whereas if $v_A = v_B$ then Merlin would never succeed in convincing them that v_A and v_B have different colors. The complexity of this protocol is $\text{NCC}(\text{INEQ}_{[\chi(G)]}) \leq \log \log \chi(G) + 1$. ■

Related work. The graph equality problem was defined by de Wolf [dW01] in the context of quantum fingerprinting. He showed that the one-round deterministic communication complexity of this problem is $\log \chi(G)$. This result was extended to the case of unbounded rounds by Briët et al. [BBL⁺15], and to non-deterministic communication complexity by Ilango et al. [ILO20].

We note that the graph equality problem is also related to the works of Alon and Orlitsky on dual-source coding [AO95, AO96]. The results on the non-deterministic and co-non-deterministic complexities of this problem are also similar in spirit to the work of Alon on edge coloring [Alo87]. A few additional similar results are the characterization of the communication complexity of the EXACTLY-N problem in the Number-on-Forehead model in terms of a certain chromatic number [CFL83], a related observation of [LY93], and a variant of Proposition 4.8 for 3-regular hypergraphs [AB23, Thm. 9].

4.2 Proof of Lemmas 4.2 and 4.4

We turn to proving Lemmas 4.2 and 4.4. Our proof follows the main ideas of [MS21], with the following main difference: The reduction of [MS21] starts from the problem $\text{INEQ}_{\mathcal{C}}$, where \mathcal{C} is some clique in the characteristic graph \mathcal{G}_{π_1} . Our reduction, on the other hand, starts from the problem $\text{GRAPHINEQ}_{\mathcal{G}_{\pi_1}}$.

As a warm-up, we first prove a simpler result that only holds for standard protocols, rather than partially half-duplex ones.

Proposition 4.10. *Let Π be a (standard) deterministic protocol that solves a relation $KW_f \diamond \text{MUX}_n$ for some $f : \{0, 1\}^m \rightarrow \{0, 1\}$ and $n \in \mathbb{N}$. For every transcript π_1 of Π , the complexity of Π is at least $|\pi_1| + \log \log \chi(\mathcal{G}_{\pi_1}) - 2$.*

Proof. Let Π, f, n be as in the proposition, and let π_1 be a transcript of Π . Let c be the maximal length of a string π_2 such that the concatenation $\pi_1 \circ \pi_2$ is a transcript of Π . We construct a non-deterministic protocol for $\text{GRAPHINEQ}_{\mathcal{G}_{\pi_1}}$ with complexity $c + 1$, and this will imply that

$$c + 1 \geq \text{coNCC}(\text{GRAPHINEQ}_{\mathcal{G}_{\pi_1}}) \geq \log \log \chi(\mathcal{G}_{\pi_1}) - 1.$$

and this will yield the desired result. Consider the following protocol for $\text{GRAPHINEQ}_{\mathcal{G}_{\pi_1}}$: Suppose that Alice and Bob get as inputs $g_A, g_B \in \mathcal{V}_{\pi_1}$ respectively. Merlin would like to convince Alice and Bob that the functions g_A, g_B are neighbors in \mathcal{G}_{π_1} . If this is indeed the case, then by the definition of \mathcal{G}_{π_1} , it holds that either $\mathcal{X}_{\pi_1}(g_A) \cap \mathcal{Y}_{\pi_1}(g_B) \neq \emptyset$ or $\mathcal{X}_{\pi_1}(g_B) \cap \mathcal{Y}_{\pi_1}(g_A) \neq \emptyset$. Merlin begins by telling Alice and Bob which is the case, at the cost of one bit. Without loss of generality, assume that we are in the first case.

Next, Merlin takes a matrix $X \in \mathcal{X}_{\pi_1}(g_A) \cap \mathcal{Y}_{\pi_1}(g_B)$, and computes the transcript of Π when Alice and Bob get (g_A, X) and (g_B, X) respectively. This transcript must be of the form $\pi_1 \circ \pi_2$ (by definition of $\mathcal{X}_{\pi_1}(g_A)$ and $\mathcal{Y}_{\pi_1}(g_B)$). Furthermore, the transcript $\pi_1 \circ \pi_2$ must output \perp , as the protocol cannot find an entry (i, j) where the matrices of Alice and Bob differ (since they are the same matrix). Merlin now sends the string π_2 to Alice and Bob as the rest of the witness. Observe that the complexity of the protocol is $c + 1$.

When Alice receives such a witness from Merlin, she accepts if and only if $\pi_1 \circ \pi_2$ is indeed a transcript of Π that outputs \perp , and there exists a matrix X such that the input (g_A, X) is consistent

with $\pi_1 \circ \pi_2$. Bob does the same on his side. Clearly, if g_A and g_B are indeed neighbors in \mathcal{G}_{π_1} , then Merlin can convince Alice and Bob that this is the case.

We turn to prove the soundness of the protocol. Suppose that Alice and Bob accept some witness w from Merlin. We prove that $g_A \neq g_B$. Let π_2 be the string contained in the witness w . Then, $\pi_1 \circ \pi_2$ is a transcript of Π that outputs \perp , and there exist matrices $X, Y \in \{0, 1\}^{m \times n}$ such that the inputs (g_A, X) and (g_B, Y) are consistent with $\pi_1 \circ \pi_2$. In particular, this implies that when given inputs (g_A, X) and (g_B, Y) , the protocol outputs \perp . It follows that $g_A \neq g_B$, since the protocol Π is only allowed to output \perp on the inputs (g_A, X) and (g_B, Y) if $g_A \neq g_B$. ■

In order to prove Lemma 4.2, we would like to apply a similar argument to partially half-duplex protocols. In this setting, however, we encounter a new issue. Recall that in the above proof, the transcript π_2 was obtained by executing the protocol Π on the inputs (g_A, X) and (g_B, X) . In the half-duplex setting, Alice and Bob do not share the same view of the protocol, so there may not be a single transcript π_2 that they both agree on. In principle, we could replace π_2 in the above argument with two transcripts $\pi_{2,A}$ and $\pi_{2,B}$, corresponding to the views of Alice and Bob respectively. Unfortunately, that would cost us a factor of 2 in the lower bound, and we are not willing to lose such a factor.

[MS21] resolve this issue using the following nice observation: the viewpoints of Alice and Bob on π_2 diverge only if non-classical rounds took place in the execution of the protocol. Nevertheless, since the protocol is *partially* half-duplex, such non-classical rounds can take place only if $g_A \neq g_B$. Hence, the presence of non-classical rounds is, on its own, a convincing proof that $g_A \neq g_B$. This leads to a natural strategy for Merlin: if the execution of the protocol on (g_A, X) and (g_B, X) encountered non-classical rounds, use them as a proof that $g_A \neq g_B$, and otherwise, send the transcript π_2 as in the proof Proposition 4.10. We now provide the formal proof.

Lemma 4.2. *Let Π be a partially half-duplex protocol that solves a relation $KW_f \diamond MUX_n$ for some $f : \{0, 1\}^m \rightarrow \{0, 1\}$ and $n \in \mathbb{N}$. For every transcript π_1 of Π , the complexity of Π is at least*

$$|\pi_1| + \log \log \chi(\mathcal{G}_{\pi_1}) - \log \log \log \chi(\mathcal{G}_{\pi_1}) - 4.$$

Proof. Let Π, f, n be as in the lemma, and let Π_A and Π_B be the protocol trees that are associated with Π according to Definition 2.24. Let π_1 be a transcript of Π . Observe that since π_1 is a transcript of Π , there exists at least one pair of vertices (u, v) of Π_A and Π_B that are consistent with π_1 . Let c be the maximum, over all such pairs, of the number of rounds that the protocol may execute after reaching (u, v) . We prove that

$$c \geq \log \log \chi(\mathcal{G}_{\pi_1}) - \log \log \log \chi(\mathcal{G}_{\pi_1}) - 4,$$

and this will yield the desired result. If $c \geq \log \log \chi(\mathcal{G}_{\pi_1})$, then we are done. Suppose otherwise. We construct a non-deterministic protocol for $\text{GRAPHINEQ}_{\mathcal{G}_{\pi_1}}$ with complexity at most $c + \log c + 3$. This will imply that

$$c + \log c + 3 \geq \text{NCC}(\text{GRAPHINEQ}_{\mathcal{G}_{\pi_1}}) \geq \log \log \chi(\mathcal{G}_{\pi_1}) - 1,$$

and hence

$$c \geq \log \log \chi(\mathcal{G}_{\pi_1}) - \log c - 4 \geq \log \log \chi(\mathcal{G}_{\pi_1}) - \log \log \log \chi(\mathcal{G}_{\pi_1}) - 4,$$

as required.

Consider the following protocol: Suppose that Alice and Bob get as inputs $g_A, g_B \in \mathcal{V}_{\pi_1}$ respectively. Merlin would like to convince Alice and Bob that the functions g_A, g_B are neighbors in \mathcal{G}_{π_1} .

If this is indeed the case, then by the definition of \mathcal{G}_{π_1} , it holds that either $\mathcal{X}_{\pi_1}(g_A) \cap \mathcal{Y}_{\pi_1}(g_B) \neq \emptyset$ or $\mathcal{X}_{\pi_1}(g_B) \cap \mathcal{Y}_{\pi_1}(g_A) \neq \emptyset$. Merlin begins by telling Alice and Bob which is the case, at the cost of one bit. Without loss of generality, assume that we are in the first case.

Let X be a matrix in the intersection $\mathcal{X}_{\pi_1}(g_A) \cap \mathcal{Y}_{\pi_1}(g_B)$. By definition, the input (g_A, X) of Alice is consistent with π_1 , and therefore there exists a (unique) vertex u of Π_A that is consistent with π_1 such that $(g_A, X) \in \mathcal{X}_u$ (see Remark 2.29). Similarly, there exists a (unique) vertex v of Π_B that is consistent with π_1 such that $(g_B, X) \in \mathcal{Y}_v$. Merlin continues to execute the protocol Π on inputs (g_A, X) and (g_B, X) starting from the pair of vertices (u, v) . Note that this execution performs at most c rounds. There are now two cases: either this execution encountered a non-classical round, or not. Merlin sends Alice and Bob an additional bit saying which is the case.

If we are in the first case, Merlin sends to Alice and Bob the following information:

- The index i of the first non-classical round.
- The transcript π_2 of the first $i - 1$ rounds. Note that since these rounds are classical, there is a single transcript that describes both Alice's and Bob's viewpoints.
- A bit saying whether the i -th round is a wasted or a silent round.

Note that this information can be encoded using at most $\lceil \log c \rceil + (c - 1) + 1$ bits. Together with the two bits that were used to tell Alice and Bob the cases in which we are, the complexity of the protocol in this case is at most $c + \log c + 3$.

When Alice receives such a witness from Merlin, she accepts if and only if there exists a matrix X such that the following holds: Let u be the unique vertex of Π_A that is determined by (g_A, X) and π_1 . Then, the input (g_A, X) is required to be consistent with the transcript π_2 in the first $i - 1$ rounds when starting from u . Furthermore, if Merlin claimed that the i -th round is wasted (respectively, silent), then Alice checks that she chooses to send (respectively, receive) at the i -th round when given the input (g_A, X) and having seen the transcript π_2 after starting from u . Similarly, Bob accepts if and only if there exists a matrix Y that satisfies the same conditions while exchanging (g_A, X) and Π_A with (g_B, Y) and Π_B . The completeness of the protocol in this case is obvious. For the soundness, observe that if $g_A = g_B$, then all the rounds of the protocol are classical regardless of the choice of X, Y (since it is *partially* half-duplex), and therefore there is no convincing transcript that Merlin may send.

If we are in the second case (i.e., the execution did not encounter a non-classical round), then Merlin proceeds as in the proof of Proposition 4.10: Merlin sends the (classical) transcript π_2 of the execution to Alice and Bob. The players accept if and only if the transcript $\pi_1 \circ \pi_2$ outputs \perp and there are matrices X, Y such that (g_A, X) and (g_B, Y) are consistent with $\pi_1 \circ \pi_2$. The analysis of the completeness and the soundness is the same as before, and the complexity in this case is at most $c + 2$. ■

We turn to proving Lemma 4.4, which is the analogue of Lemma 4.2 for strong composition. Here, we use exactly the same argument, where the only difference is that the matrix X is replaced with two matrices X, Y that witness the weak intersection property. For completeness, we provide the full proof below.

Lemma 4.4. *Let Π be a partially half-duplex protocol that solves a relation $KW_f \otimes MUX_n$ for some $f : \{0, 1\}^m \rightarrow \{0, 1\}$ and $n \in \mathbb{N}$. For every transcript π_1 of Π , the complexity of Π is at least*

$$|\pi_1| + \log \log \chi(\mathcal{G}_{\pi_1}) - \log \log \log \chi(\mathcal{G}_{\pi_1}) - 4.$$

Proof. Let Π, f, n be as in the lemma, and let Π_A and Π_B be the protocol trees that are associated with Π according to Definition 2.24. Let π_1 be a transcript of Π . Observe that since π_1 is a transcript of Π , there exists at least one pair of vertices (u, v) of Π_A and Π_B that are consistent with π_1 . Let c be the maximum, over all such pairs, of the number of rounds that the protocol may execute after reaching (u, v) . We prove that

$$c \geq \log \log \chi(\mathcal{G}_{\pi_1}) - \log \log \log \chi(\mathcal{G}_{\pi_1}) - 4,$$

and this will yield the required result. If $c \geq \log \log \chi(\mathcal{G}_{\pi_1})$, then we are done. Suppose otherwise. We construct a non-deterministic protocol for $\text{GRAPHINEQ}_{\mathcal{G}_{\pi_1}}$ with complexity at most $c + \log c + 3$. This will imply that

$$c + \log c + 3 \geq \text{NCC}(\text{GRAPHINEQ}_{\mathcal{G}_{\pi_1}}) \geq \log \log \chi(\mathcal{G}_{\pi_1}) - 1,$$

and hence

$$c \geq \log \log \chi(\mathcal{G}_{\pi_1}) - \log c - 4 \geq \log \log \chi(\mathcal{G}_{\pi_1}) - \log \log \log \chi(\mathcal{G}_{\pi_1}) - 4,$$

as required.

Consider the following protocol: Suppose that Alice and Bob get as inputs $g_A, g_B \in \mathcal{V}_{\pi_1}$ respectively. Merlin would like to convince Alice and Bob that the functions g_A, g_B are neighbors in \mathcal{G}_{π_1} . If this is indeed the case, then by the definition of \mathcal{G}_{π_1} , one of the following cases holds:

- There exist matrices $X \in \mathcal{X}_{\pi_1}(g_A)$ and $Y \in \mathcal{Y}_{\pi_1}(g_B)$ such that $X_i = Y_i$ for every $i \in [m]$ for which $a_i \neq b_i$ (where $a = g_A(X)$ and $b = g_B(Y)$).
- The same statement holds when exchanging g_A and g_B .

Merlin begins by telling Alice and Bob which is the case, at the cost of one bit. Without loss of generality, assume that we are in the first case.

Let X, Y, a, b be matrices and strings as in the first case above. By definition, the input (g_A, X) of Alice is consistent with π_1 , and therefore there exists a (unique) vertex u of Π_A that is consistent with π_1 such that $(g_A, X) \in \mathcal{X}_u$ (see Remark 2.29). Similarly, there exists a (unique) vertex v of Π_B that is consistent with π_1 such that $(g_B, Y) \in \mathcal{Y}_v$. Merlin continues to execute the protocol Π on inputs (g_A, X) and (g_B, Y) starting from the pair of vertices (u, v) . Note that this execution performs at most c rounds. There are now two cases: either this execution encountered a non-classical round, or not. Merlin sends Alice and Bob an additional bit saying which is the case.

If we are in the first case, Merlin sends to Alice and Bob the following information:

- The index i of the first non-classical round.
- The transcript π_2 of the first $i - 1$ rounds. Note that since these rounds are classical, there is a single transcript that describes both Alice's and Bob's viewpoints.
- A bit saying whether the i -th round is a wasted or a silent round.

Note that this information can be encoded using at most $\lceil \log c \rceil + (c - 1) + 1$ bits. Together with the two bits that were used to tell Alice and Bob the cases in which we are, the complexity of the protocol in this case is at most $c + \log c + 3$.

When Alice receives such a witness from Merlin, she accepts if and only if there exists a matrix X such that the following holds: Let u be the vertex of Π_A that is determined by (g_A, X) and π_1 . Then, the input (g_A, X) is required to be consistent with the transcript π_2 in the first $i - 1$ rounds when starting from u . Furthermore, if Merlin claimed that the i -th round is wasted (respectively,

silent), then Alice checks that she chooses to send (respectively, receive) at the i -th round, when given the input (g_A, X) and having seen the transcript π_2 after starting from u . Similarly, Bob accepts if and only if there exists a matrix Y that satisfies the same conditions while exchanging (g_A, X) and Π_A with (g_B, Y) and Π_B . The completeness of the protocol in this case is obvious. For the soundness, observe that if $g_A = g_B$, then all the rounds of the protocol are classical regardless of the choice of X, Y (since it is *partially* half-duplex), and therefore there is no convincing transcript that Merlin may send.

Suppose that we are in the second case, i.e., the execution did not encounter a non-classical round. In this case, the transcript $\pi_1 \circ \pi_2$ must output \perp , as the protocol cannot find an entry (i, j) where $a_i \neq b_i$ and $X_{i,j} \neq Y_{i,j}$ (by definition of X, Y, a, b). Merlin now sends π_2 to Alice and Bob as the rest of the witness. Observe that the complexity of the protocol in this case is $c + 1$.

When Alice receives such a witness from Merlin, she accepts if and only if $\pi_1 \circ \pi_2$ is indeed a transcript of Π that outputs \perp , and there exists a matrix X such that the input (g_A, X) is consistent with $\pi_1 \circ \pi_2$. Bob does the same on his side. Again, the completeness of the protocol in this case is obvious.

We turn to prove the soundness of the protocol. Suppose that Alice and Bob accept some witness w from Merlin. We prove that $g_A \neq g_B$. Let π_2 be the string contained in the witness w . Then, $\pi_1 \circ \pi_2$ is a transcript of Π that outputs \perp , and there exist matrices $X, Y \in \{0, 1\}^{m \times n}$ such that the inputs (g_A, X) and (g_B, Y) are consistent with $\pi_1 \circ \pi_2$. In particular, this implies that when given inputs (g_A, X) and (g_B, Y) , the protocol may transmit the transcript $\pi_1 \circ \pi_2$, and hence may output \perp . It follows that $g_A \neq g_B$, as required, since the protocol Π is only allowed to output \perp on the inputs (g_A, X) and (g_B, Y) if $g_A \neq g_B$. ■

5 Prefix-thick sets

In this section, we introduce the notion of prefix-thick sets, and show how to lower bound their number using a result of Salo and Törmä [ST14]. Let Σ be a finite alphabet of size q . We start by recalling the definition of a prefix tree.

Definition 5.1. Let $\mathcal{X} \subseteq \Sigma^m$. The *prefix tree* of \mathcal{X} is a rooted tree $T_{\mathcal{X}}$ of depth m that satisfies the following properties:

- The vertices of $T_{\mathcal{X}}$ at depth i are all the length- i prefixes of strings in \mathcal{X} . In particular, the root of T is the empty string, and its leaves are the strings in \mathcal{X} .
- A vertex $y \in \Sigma^{i+1}$ at depth $i + 1$ is a child of a vertex $x \in \Sigma^i$ at depth i if and only if x is a prefix of y . In this case, we label the edge from x to y with the symbol $\sigma \in \Sigma$ that satisfies $y = x \circ \sigma$ (i.e., $\sigma = y_{i+1}$), and say that y is the σ -child of x .

Definition 5.2. Let $\mathcal{X} \subseteq \Sigma^m$ be a set of strings, and let $t \in \mathbb{R}$. We say that \mathcal{X} is *prefix thick with degree t* if there is a subset $\mathcal{X}' \subseteq \mathcal{X}$ whose prefix tree has minimum degree that is greater than t . When Σ is clear from the context, we abbreviate and say that \mathcal{X} is *prefix thick* if it is prefix thick with degree $q/2$.

As discussed in the introduction, our motivation for using prefix-thick sets is the following easy observation.

Proposition 1.10. *Let $\mathcal{X}, \mathcal{Y} \subseteq \Sigma^m$. If \mathcal{X} and \mathcal{Y} are both prefix thick, then $\mathcal{X} \cap \mathcal{Y} \neq \emptyset$.*

Proof. Let $\mathcal{X}, \mathcal{Y} \subseteq \Sigma^m$ be as in the proposition. Assume that \mathcal{X} and \mathcal{Y} are both prefix thick with degree $q/2$, and let \mathcal{X}' and \mathcal{Y}' be the corresponding subsets of \mathcal{X} and \mathcal{Y} respectively. Let $T_{\mathcal{X}'}$ and $T_{\mathcal{Y}'}$ be the prefix trees of \mathcal{X}' and \mathcal{Y}' respectively, and note that the degree of every vertex in these trees is at most q . Since the degrees of the roots in both $T_{\mathcal{X}'}$ and $T_{\mathcal{Y}'}$ is greater than $q/2$, there must be some vertex $\sigma_1 \in \Sigma$ that is a child of the root in both trees. Similarly, since the degree of σ_1 is greater than $q/2$ in both $T_{\mathcal{X}'}$ and $T_{\mathcal{Y}'}$, there must be some vertex $(\sigma_1, \sigma_2) \in \Sigma^2$ that is a child of σ_1 in both trees. Continuing in this fashion, we obtain a leaf $(\sigma_1, \dots, \sigma_m) \in \Sigma^m$ that belongs to both trees, and this leaf is a string in the intersection $\mathcal{X} \cap \mathcal{Y}$. ■

Remark 5.3. In Section 6, we use the following generalization of the above definitions: we consider m different alphabets $\Sigma_1, \dots, \Sigma_m$ of the same size q , and work with strings in $\Sigma_1 \times \dots \times \Sigma_m$. It is not hard to check that all the definitions and results in this section work in this setting without a change.

In fact, we could even work with m different alphabets of *different* sizes. The results in this section would continue to hold in such a case, but with a somewhat more cumbersome phrasing. For the simplicity of the presentation, we decided to skip this more general version.

Given a set $\mathcal{X} \subseteq \Sigma^m$, we would like to lower bound the number of coordinate sets $I \subseteq [m]$ for which $\mathcal{X}|_I$ is prefix thick. To this end, we use a result of Salo and Törmä [ST14] (see also [Sal21]). We first introduce a few additional terms.

Definition 5.4 ([ST14, Def. 5.6]). Let $\mathcal{X} \subseteq \Sigma^m$ be a set of strings and let $T_{\mathcal{X}}$ be its prefix tree. We say that a vector $w \in [q]^m$ is a *branching structure* of \mathcal{X} if $T_{\mathcal{X}}$ has a sub-tree of depth m in which, for every $i \in [m]$, the degree of all the vertices at depth i is w_i . The *winning set* of \mathcal{X} is the set of all branching structures of \mathcal{X} .

Remark 5.5. A “winning set” is called that way since [ST14] present their result in terms of winning possibilities in a certain game over strings. We describe their result in terms of prefix trees since this view is more useful for our purposes.

Next, observe that $\mathcal{X}|_I$ is prefix thick with degree t if and only if there is a branching structure w of \mathcal{X} such that $w_i > t$ for every $i \in I$. Thus, we can obtain lower bounds on the number of such sets I from the number of branching vectors. The following result of [ST14] tells us exactly the size of the winning set.

Lemma 5.6 ([ST14, Prop. 5.7]). *Let $\mathcal{X} \subseteq \Sigma^m$ be a set of strings and let $\mathcal{W}(\mathcal{X})$ be its winning set. Then, $|\mathcal{W}(\mathcal{X})| = |\mathcal{X}|$.*

Proof. We prove the proposition by induction on m . For base case of $m = 1$, observe that for every $\mathcal{X} \subseteq \Sigma$ it holds that $\mathcal{W}(\mathcal{X}) = \{1, \dots, |\mathcal{X}|\}$. We assume that the lemma holds for m , and prove it for $m + 1$. Let $\mathcal{X} \subseteq \Sigma^{m+1}$ be a set of strings and let $\mathcal{W}(\mathcal{X}) \subseteq [q]^{m+1}$ be its winning set. For every symbol $\sigma \in \Sigma$, let

$$\mathcal{X}_\sigma = \{x \in \Sigma^m : \sigma \circ x \in \mathcal{X}\},$$

and let $\mathcal{W}(\mathcal{X}_\sigma) \subseteq [q]^m$ be the winning set of \mathcal{X}_σ . By the induction assumption, it holds that $|\mathcal{W}(\mathcal{X}_\sigma)| = |\mathcal{X}_\sigma|$ for every $\sigma \in \Sigma$. For every $k \in [q]$, let \mathcal{W}_k be the set of vectors $w \in [q]^m$ such that $k \circ w \in \mathcal{W}(\mathcal{X})$, and for every $w \in [q]^m$, let $c_w \in [q]$ denote the number of symbols $\sigma \in \Sigma$ such that $w \in \mathcal{W}(\mathcal{X}_\sigma)$. The key observation is that for every $k \in [q]$ and $w \in [q]^m$, it holds that $w \in \mathcal{W}_k$ if and only if $k \leq c_w$. We prove this observation next:

- **The “only if” direction:** Assume that $w \in \mathcal{W}_k$, so $k \circ w \in \mathcal{W}(x)$. This implies that the prefix tree $T_{\mathcal{X}}$ of \mathcal{X} has a sub-tree $T_{k \circ w}$ that corresponds to the branching structure $k \circ w$. In particular, the root of $T_{k \circ w}$ has k children. Let $\sigma_1, \dots, \sigma_k$ be the labels of the outgoing edges of the root of $T_{k \circ w}$, and let T_{σ_i} be the sub-tree that is rooted at the σ_i -child of the root. Then, each tree T_{σ_i} is a sub-tree of the prefix tree of \mathcal{X}_{σ_i} , and has branching structure w . Therefore, for each σ_i , it holds that $w \in \mathcal{W}(\mathcal{X}_{\sigma_i})$, so $c_w \geq k$, as required.
- **The “if” direction:** Assume that $k \leq c_w$. Then, there exist k symbols $\sigma_1, \dots, \sigma_k$ such that $w \in \mathcal{W}(\mathcal{X}_{\sigma_i})$ for each $i \in [k]$. Therefore, for each symbol σ_i , the prefix tree of \mathcal{X}_{σ_i} has a sub-tree T_{σ_i} that corresponds to the branching structure w . Now, let $T_{k \circ w}$ be the prefix tree in which the children of the root are the symbols $\sigma_1, \dots, \sigma_k$, and the sub-tree that is rooted at σ_i is T_{σ_i} . It is not hard to see that $T_{k \circ w}$ is a sub-tree of the prefix tree $T_{\mathcal{X}}$ of \mathcal{X} , and has branching structure $k \circ w$. Hence, $k \circ w \in \mathcal{W}(\mathcal{X})$, or alternatively $w \in \mathcal{W}_k$.

Given the above key observation, we prove the lemma as follows:

$$\begin{aligned}
|\mathcal{W}(\mathcal{X})| &= \sum_{k=1}^q |\mathcal{W}_k| \\
&= \sum_{k=1}^q \sum_{w \in \mathcal{W}_k} 1 \\
&= \sum_{w \in [q]^m} \sum_{k \in [q]: w \in \mathcal{W}_k} 1 \\
&= \sum_{w \in [q]^m} \sum_{k=1}^{c_w} 1 && \text{(the key observation)} \\
&= \sum_{w \in [q]^m} c_w \\
&= \sum_{w \in [q]^m} \sum_{\sigma \in \Sigma: w \in \mathcal{W}(\mathcal{X}_{\sigma})} 1 && \text{(definition of } c_w) \\
&= \sum_{\sigma \in \Sigma} \sum_{w \in \mathcal{W}(\mathcal{X}_{\sigma})} 1 \\
&= \sum_{\sigma \in \Sigma} |\mathcal{W}(\mathcal{X}_{\sigma})| \\
&= \sum_{\sigma \in \Sigma} |\mathcal{X}_{\sigma}| && \text{(induction hypothesis)} \\
&= |\mathcal{X}|,
\end{aligned}$$

as required. ■

We now use Lemma 5.6 to lower bound the number of subsets for which $\mathcal{X}|_I$ is prefix thick. Specifically, the following result says that the fraction of such sets I is not much smaller than the density of \mathcal{X} inside Σ^m .

Lemma 5.7. *Let $\mathcal{X} \subseteq \Sigma^m$, let $\varepsilon \geq 0$, and let \mathcal{F} be the family of subsets $I \subseteq [m]$ such that $\mathcal{X}|_I$ is prefix thick with degree $(\frac{1}{2} + \varepsilon) \cdot q$. Then,*

$$\frac{|\mathcal{F}|}{2^m} \geq 2^{-2 \log e \cdot \varepsilon \cdot m} \cdot \frac{|\mathcal{X}|}{|\Sigma^m|}.$$

Proof. Let \mathcal{X} , ε , and \mathcal{F} be as in Lemma 5.7, and let \mathcal{W} be the winning set of \mathcal{X} . Let $\phi : \mathcal{W} \rightarrow 2^{[m]}$ be the mapping that maps every branching structure $w \in \mathcal{W}$ to the set $I \subseteq [m]$ of coordinates i such that $w_i > (\frac{1}{2} + \varepsilon) \cdot q$. Now, observe that for every $I \subseteq [m]$, the preimage $\phi^{-1}(I)$ is contained in the set of all vectors in $[q]^m$ whose entries that are greater than $(\frac{1}{2} + \varepsilon) \cdot q$ are exactly those in I . Hence, for every $I \subseteq [m]$, it holds that

$$\begin{aligned} |\phi^{-1}(I)| &= \left[q - \left\lfloor \left(\frac{1}{2} + \varepsilon \right) \cdot q \right\rfloor \right]^{|I|} \left[\left(\frac{1}{2} + \varepsilon \right) \cdot q \right]^{m-|I|} \\ &\leq \left[\left(\frac{1}{2} + \varepsilon \right) \cdot q \right]^m \\ &\leq \left(\left(\frac{1}{2} + \varepsilon \right) \cdot q \right)^m \\ &= (1 + 2\varepsilon)^m \cdot \frac{q^m}{2^m} \\ &\leq e^{2\varepsilon m} \cdot \frac{q^m}{2^m} \qquad \text{(since } 1 + x \leq e^x \text{).} \end{aligned}$$

Now, by Lemma 5.6, it follows that

$$\begin{aligned} |\mathcal{X}| &= |\mathcal{W}| \\ &= \sum_{I \in \mathcal{F}} |\phi^{-1}(I)| \\ &\leq \sum_{I \in \mathcal{F}} \frac{q^m}{2^m} \cdot e^{2\varepsilon m} \\ &= \frac{q^m}{2^m} \cdot 2^{2 \log e \cdot \varepsilon \cdot m} \cdot |\mathcal{F}|. \end{aligned}$$

The lemma follows from by dividing both sides by $2^{2 \log e \cdot \varepsilon \cdot m} \cdot q^m$. ■

Remark 5.8. We note that the last proof was suggested to us by Ville Salo (over MathOverflow), although we also came up with it independently.

Remark 5.9. Lemma 5.7 can be viewed as a generalization of the Sauer-Shelah lemma [Sau72, She72], and in particular of its strengthening due to Pajor [Paj85]. Specifically, if $\Sigma = \{0, 1\}$, then $\mathcal{X}|_I$ is a prefix-thick set if and only if I is shattered by \mathcal{X} . If we substitute $q = 2$ and $\varepsilon = 0$ in Lemma 5.7, we get that the number of shattered sets $|\mathcal{F}|$ is at least $|\mathcal{X}|$, as shown by [Paj85]. Furthermore, Lemma 5.6 of [ST14] itself is a generalization of a result of [ARS02] on ordered-shattering sets, which is another strengthening of [Paj85].

6 Proof of the structure theorem

In this section, we prove our structure theorem, restated next.

Theorem 3.9. *There exists a constant $\gamma > 0.04$ such that the following holds: Let $f : \{0, 1\}^m \rightarrow \{0, 1\}$, let $n \in \mathbb{N}$, and let Π be a partially half-duplex protocol that solves $KW_f \otimes MUX_n$. For every γ -live transcript π_1 of Π , the complexity of Π is at least*

$$|\pi_1| + n - O(\log(m \cdot n)).$$

Let $\gamma > 0$ be a universal constant that will be fixed later to a value greater than 0.04. Let f, n, Π be as in the structure theorem, and π_1 be a γ -live transcript π_1 of Π , and let \mathcal{G}_{π_1} be the characteristic graph of π_1 as in Definition 4.3. By Lemma 4.4, in order to prove the desired lower bound, it is sufficient to prove that

$$\log \log \chi(\mathcal{G}_{\pi_1}) \geq n - O(\log(m \cdot n)).$$

We prove the latter bound by proving an upper bound on the independence number of a sub-graph \mathcal{G}' of \mathcal{G}_{π_1} . The following two lemmas construct that sub-graph and bound its independence number respectively.

Lemma 6.1. *There exist a universal constant $\varepsilon > 0$, a set of balanced functions $\mathcal{V}' \subseteq \mathcal{V}_0$, a subset $I \subseteq [m]$, and strings $a \in f^{-1}(1)$ and $b \in f^{-1}(0)$, that satisfy the following properties:*

- $|\mathcal{V}'| \geq 2^{-4m} \cdot |\mathcal{V}_0|$.
- $a|_{[m]-I} = b|_{[m]-I}$.
- For every $g \in \mathcal{V}'$, the sets $\mathcal{X}_{\pi_1}(g, a)|_I$ and $\mathcal{Y}_{\pi_1}(g, a)|_I$ are prefix thick with degree $(\frac{1}{2} + \varepsilon) \cdot 2^{n-1}$.

Lemma 6.2. *Let ε and \mathcal{V}' be the constant and set from Lemma 6.1, and let \mathcal{G}' be the sub-graph of \mathcal{G}_{π_1} induced by \mathcal{V}' . Then,*

$$\alpha(\mathcal{G}') \leq 2^{-\left(\frac{\log e}{32} \cdot \varepsilon^2 \cdot 2^n - \frac{1}{2} \cdot m \cdot n - 4 \cdot m - 1\right)} \cdot |\mathcal{V}'|.$$

We prove Lemmas 6.1 and 6.2 in Sections 6.1 and 6.2 respectively. We now derive the structure theorem from those lemmas.

Proof of Theorem 3.9. Let γ, f, n, Π, π_1 be as above. Let ε and \mathcal{V}' be the constant and set whose existence is guaranteed by Lemma 6.1, and let \mathcal{G}' be the sub-graph of \mathcal{G}_{π_1} induced by \mathcal{V}' . By Lemma 6.2, it holds that

$$\chi(\mathcal{G}') \geq |\mathcal{V}'| / \alpha(\mathcal{G}') \geq 2^{\frac{\log e}{32} \cdot \varepsilon^2 \cdot 2^n - \frac{1}{2} \cdot m \cdot n - 4 \cdot m - 1}.$$

Clearly, $\chi(\mathcal{G}_{\pi_1}) \geq \chi(\mathcal{G}')$, and therefore

$$\log \log \chi(\mathcal{G}_{\pi_1}) \geq n + \log \left(\frac{\log e}{32} \cdot \varepsilon^2 \right) - \log \left(\frac{1}{2} \cdot m \cdot n \right) - \log(4 \cdot m) - 1 = n - O(\log(m \cdot n)).$$

By Lemma 4.4, it follows that the communication complexity of the protocol Π is at least

$$|\pi_1| + \log \log \chi(\mathcal{G}_{\pi_1}) - \log \log \log \chi(\mathcal{G}_{\pi_1}) - 4 \geq |\pi_1| + n - O(\log(m \cdot n)),$$

as required. ■

6.1 The construction of \mathcal{G}'

In this section, we prove Lemma 6.1. Recall that π_1 is γ -alive, and hence there exists a set of balanced functions $\mathcal{V} \subseteq \mathcal{V}_{\pi_1}$ that satisfies the following conditions:

- $|\mathcal{V}| \geq 2^{-m} \cdot |\mathcal{V}_0|$.
- For every $g \in \mathcal{V}$, it holds that $\log L(\mathcal{A}_{\pi_1}(g) \times \mathcal{B}_{\pi_1}(g)) \geq (1 - \gamma) \cdot m + \kappa \log m + \kappa$.

- For every $g \in \mathcal{V}$, $a \in \mathcal{A}_{\pi_1}(g)$, and $b \in \mathcal{B}_{\pi_1}(g)$, it holds that $|\mathcal{X}_{\pi_1}(g, a)| \geq 2^{-\gamma \cdot m+1} \cdot |g^{-1}(a)|$ and $|\mathcal{Y}_{\pi_1}(g, b)| \geq 2^{-\gamma \cdot m+1} \cdot |g^{-1}(b)|$.

The crux of our proof is showing that for every $g \in \mathcal{V}$, there exists a subset $I_g \subseteq [m]$ and strings $a_g \in f^{-1}(1)$ and $b_g \in f^{-1}(0)$ that satisfy the following properties:

- $a_g|_{[m]-I_g} = b_g|_{[m]-I_g}$.
- The sets $\mathcal{X}_{\pi_1}(g, a_g)|_{I_g}$ and $\mathcal{Y}_{\pi_1}(g, a_g)|_{I_g}$ are prefix thick with degree $(\frac{1}{2} + \varepsilon) \cdot 2^{n-1}$.

Having shown that, we will choose the triplet (I, a, b) to be the most popular triplet among all triplets (I_g, a_g, b_g) for all $g \in \mathcal{V}$, and choose \mathcal{V}' to be the set of those functions $g \in \mathcal{V}$ whose corresponding triplet is (I, a, b) . By an averaging argument, it will follow that

$$|\mathcal{V}'| \geq 2^{-3m} \cdot |\mathcal{V}| \geq 2^{-4m} \cdot |\mathcal{V}_0|,$$

thus completing the proof.

For the rest of this section, we focus on proving the existence of such I_g, a_g, b_g for every $g \in \mathcal{V}_{\pi_1}$. Fix a function $g \in \mathcal{V}_{\pi_1}$, and let $0 < \beta < \frac{1}{2}$ be some universal constant to be fixed later. For convenience, we abbreviate $\mathcal{A} = \mathcal{A}_{\pi_1}(g)$, $\mathcal{B} = \mathcal{B}_{\pi_1}(g)$, $\mathcal{X}(a) = \mathcal{X}_{\pi_1}(g, a)$, and $\mathcal{Y}(b) = \mathcal{Y}(g, b)$. The high-level idea of our construction of I_g, a_g, b_g is the following:

1. For every $a \in \mathcal{A}$, the set $\mathcal{X}(a)$ is large (since π_1 is alive), and hence there are many subsets $I \subseteq [m]$ such that $\mathcal{X}(a)|_I$ is prefix thick by Lemma 5.7.
2. In particular, for every $a \in \mathcal{A}$ many of those subsets I are large (this follows by a standard concentration bound).
3. Hence, by an averaging argument, there exists a single large set $I_1 \subseteq [m]$ such that $\mathcal{X}(a)|_{I_1}$ is prefix thick for many strings $a \in \mathcal{A}$. We denote the set of those strings by \mathcal{A}_1 .
4. Repeating the same argument for $b \in \mathcal{B}$ and $\mathcal{Y}(b)$, but *this time restricting ourselves to coordinates in I_1* , we get that there exists a large set $I_2 \subseteq I_1$ such that $\mathcal{Y}(b)|_{I_2}$ is prefix thick for many strings $b \in \mathcal{B}$. We denote the set of those strings by \mathcal{B}_1 .
5. Note that it also holds that $\mathcal{X}(a)|_{I_2}$ is prefix thick for every $a \in \mathcal{A}_1$ (since $I_2 \subseteq I_1$). We therefore choose $I_g = I_2$.
6. Recall that the formula complexity $L(\mathcal{A} \times \mathcal{B})$ is large (since π_1 is alive). Using the fact that the sets \mathcal{A}_1 and \mathcal{B}_1 are dense within \mathcal{A} and \mathcal{B} respectively, we deduce that the formula complexity $L(\mathcal{A}_1 \times \mathcal{B}_1)$ is large too (this deduction is non-trivial, as we explain below).
7. Finally, since I_g is large and $L(\mathcal{A}_1 \times \mathcal{B}_1)$ is large, we claim that there must exist strings $a_g \in \mathcal{A}_1$ and $b_g \in \mathcal{B}_1$ such that $a_g|_{[m]-I_g} = b_g|_{[m]-I_g}$: otherwise, Alice and Bob could have solved $KW_{\mathcal{A}_1 \times \mathcal{B}_1}$ too efficiently by sending their bits in the coordinates of $[m] - I_g$.

There is one issue in the argument as described above. In Step 6, we would like to deduce that the complexity $L(\mathcal{A}_1 \times \mathcal{B}_1)$ is large using the fact that the sets \mathcal{A}_1 and \mathcal{B}_1 are large. Nevertheless, such a deduction is false in general: it is easy to construct examples of sets \mathcal{A}_1 and \mathcal{B}_1 that are dense in \mathcal{A} and \mathcal{B} but for which $L(\mathcal{A}_1 \times \mathcal{B}_1)$ is small. In order to remedy this issue, we apply the fortification theorem of [DM16] (see Theorem 2.22) on Alice's and Bob's sides before Steps 1 and 4 respectively. This theorem guarantees that the largeness of \mathcal{A}_1 and \mathcal{B}_1 will imply that the formula complexity $L(\mathcal{A}_1 \times \mathcal{B}_1)$ is large, as required by Step 6. We now execute the above argument formally.

First, we apply Theorem 2.22 (the fortification theorem) to the rectangle $\mathcal{A} \times \mathcal{B}$ on Alice's side. This yields a subset $\mathcal{A}_0 \subseteq \mathcal{A}$ such that the rectangle $\mathcal{A}_0 \times \mathcal{B}$ is $\frac{1}{4m}$ -fortified on Alice's side and satisfies

$$\mathsf{L}(\mathcal{A}_0 \times \mathcal{B}) \geq \frac{1}{4} \cdot \mathsf{L}(\mathcal{A} \times \mathcal{B}).$$

Second, let $a \in \mathcal{A}_0$. Since π_1 is γ -alive, it holds that $|\mathcal{X}(a)| \geq 2 \cdot 2^{-\gamma \cdot m} \cdot |g^{-1}(a)|$. We now apply Lemma 5.7 in order to find sets I for which $\mathcal{X}(a)|_I$ is prefix thick. To this end, we view $\mathcal{X}(a)$ as a subset of strings in $g^{-1}(a) = g^{-1}(a_1) \times \dots \times g^{-1}(a_m)$, where $g^{-1}(a_1), \dots, g^{-1}(a_m)$ are alphabets of size 2^{n-1} since g is balanced. Let \mathcal{F}_a denote the family of subsets $I \subseteq [m]$ such that $\mathcal{X}(a)|_I$ is prefix-thick with degree $(\frac{1}{2} + \varepsilon) \cdot 2^{n-1}$. By Lemma 5.7, it holds that $|\mathcal{F}_a| \geq 2 \cdot 2^{-(\gamma+2 \log e \cdot \varepsilon) \cdot m} \cdot 2^m$.

Next, for every string $a \in \mathcal{A}_0$, let \mathcal{F}'_a denote the family of subsets $I \in \mathcal{F}_a$ of size at least $(\frac{1}{2} - \beta) \cdot m$. By Fact 2.1, the number of subsets $I \subseteq [m]$ whose size is less than $(\frac{1}{2} - \beta) \cdot m$ is at most $2^{-2 \log e \cdot \beta^2 \cdot m} \cdot 2^m$, and therefore

$$|\mathcal{F}'_a| \geq |\mathcal{F}_a| - 2^{-2 \log e \cdot \beta^2 \cdot m} \cdot 2^m \geq \left(2 \cdot 2^{-(\gamma+2 \log e \cdot \varepsilon) \cdot m} - 2^{-2 \log e \cdot \beta^2 \cdot m}\right) \cdot 2^m$$

We will later choose the constants $\beta, \gamma, \varepsilon$ such that they satisfy that $\gamma + 2 \log e \cdot \varepsilon \leq 2 \log e \cdot \beta^2$, which implies that $|\mathcal{F}'_a| \geq 2^{-(\gamma+2 \log e \cdot \varepsilon) \cdot m} \cdot 2^m$.

It now follows by an averaging argument that there exists a set $I_1 \subseteq [m]$ such that $I_1 \in \mathcal{F}'_a$ for at least $2^{-(\gamma+2 \log e \cdot \varepsilon) \cdot m}$ fraction of the strings $a \in \mathcal{A}_0$. Let \mathcal{A}_1 be the set of those strings $a \in \mathcal{A}_0$, so $|\mathcal{A}_1| \geq 2^{-(\gamma+2 \log e \cdot \varepsilon) \cdot m} \cdot |\mathcal{A}_0|$. Since the rectangle $\mathcal{A}_0 \times \mathcal{B}$ is $\frac{1}{4m}$ -fortified on Alice's side, it holds that

$$\mathsf{L}(\mathcal{A}_1 \times \mathcal{B}) \geq \frac{1}{4m} \cdot 2^{-(\gamma+2 \log e \cdot \varepsilon) \cdot m} \cdot \mathsf{L}(\mathcal{A}_0 \times \mathcal{B}) \geq 2^{-(\gamma+2 \log e \cdot \varepsilon) \cdot m - \log m - 4} \cdot \mathsf{L}(\mathcal{A} \times \mathcal{B}).$$

We turn to apply the same argument to the inputs on Bob's side, but this time we restrict ourselves to the coordinates in the set I_1 . We first apply Theorem 2.22 (the fortification theorem) to the rectangle $\mathcal{A}_1 \times \mathcal{B}$ on Bob's side. This yields a subset $\mathcal{B}_0 \subseteq \mathcal{B}$ such that the rectangle $\mathcal{A}_1 \times \mathcal{B}_0$ is $\frac{1}{4m}$ -fortified on Bob's side and satisfies

$$\mathsf{L}(\mathcal{A}_1 \times \mathcal{B}_0) \geq \frac{1}{4} \cdot \mathsf{L}(\mathcal{A}_1 \times \mathcal{B}) \geq 2^{-(\gamma+2 \log e \cdot \varepsilon) \cdot m - \log m - 6} \cdot \mathsf{L}(\mathcal{A} \times \mathcal{B}).$$

Second, observe that for every $b \in \mathcal{B}_0$, it holds that $|\mathcal{Y}(b)| \geq 2 \cdot 2^{-\gamma \cdot m} \cdot |g^{-1}(b)|$, and hence $|\mathcal{Y}(b)|_{I_1} \geq 2 \cdot 2^{-\gamma \cdot m} \cdot |g^{-1}(b)|_{I_1}$. Following similar steps as before, we denote by \mathcal{F}_b , for every $b \in \mathcal{B}$, the family of subsets $I \subseteq I_1$ such that $\mathcal{Y}(b)|_I$ is prefix-thick with degree $(\frac{1}{2} + \varepsilon) \cdot 2^{n-1}$. Again, Lemma 5.7 implies that

$$|\mathcal{F}_b| \geq 2 \cdot 2^{-\gamma \cdot m - 2 \log e \cdot \varepsilon \cdot |I_1|} \cdot 2^{|I_1|} \geq 2 \cdot 2^{-(\gamma+2 \log e \cdot \varepsilon) \cdot m} \cdot 2^{|I_1|}.$$

Next, we let \mathcal{F}'_b be the family of subsets $I \in \mathcal{F}_b$ of size at least

$$\left(\frac{1}{2} - \beta\right) \cdot |I_1| \geq \left(\frac{1}{2} - \beta\right)^2 \cdot m,$$

and observe, as before, that

$$|\mathcal{F}'_b| \geq |\mathcal{F}_b| - 2^{-2 \log e \cdot \beta^2 \cdot |I_1|} \cdot 2^{|I_1|} \geq \left(2 \cdot 2^{-(\gamma+2 \log e \cdot \varepsilon) \cdot |I_1|} - 2^{-2 \log e \cdot \beta^2 \cdot |I_1|}\right) \cdot 2^{|I_1|}.$$

Again, since we will choose the constants $\beta, \gamma, \varepsilon$ such that they satisfy that $\gamma + 2 \log e \cdot \varepsilon \leq 2 \log e \cdot \beta^2$, it follows that

$$|\mathcal{F}'_b| \geq 2^{-(\gamma+2 \log e \cdot \varepsilon) \cdot |I_1|} \cdot 2^{|I_1|} \geq 2^{-(\gamma+2 \log e \cdot \varepsilon) \cdot m} \cdot 2^{|I_1|}.$$

It now follows by an averaging argument that there exists a set $I_g \subseteq I_1$ such that $I_g \in \mathcal{F}'_b$ for at least $2^{-(\gamma+2\log e \cdot \varepsilon) \cdot m}$ fraction of the strings $b \in \mathcal{B}_0$. Let \mathcal{B}_1 be the set of those strings $b \in \mathcal{B}_0$, so $|\mathcal{B}_1| \geq 2^{-(\gamma+2\log e \cdot \varepsilon) \cdot m} \cdot |\mathcal{B}_0|$. Since the rectangle $\mathcal{A}_1 \times \mathcal{B}_0$ is $\frac{1}{4m}$ -fortified on Alice's side, it holds that

$$\mathsf{L}(\mathcal{A}_1 \times \mathcal{B}_1) \geq \frac{1}{4m} \cdot 2^{-(\gamma+2\log e \cdot \varepsilon) \cdot m} \cdot \mathsf{L}(\mathcal{A}_1 \times \mathcal{B}_0) \geq 2^{-(2\gamma+4\log e \cdot \varepsilon) \cdot m - 2\log m - 8} \cdot \mathsf{L}(\mathcal{A} \times \mathcal{B})$$

Observe that $\mathcal{Y}(b)|_{I_g}$ is prefix thick with degree $(\frac{1}{2} + \varepsilon) \cdot 2^{n-1}$ for every $b \in \mathcal{B}_1$ by definition. Moreover, for every $a \in \mathcal{A}_1$, the set $\mathcal{X}(a)|_{I_g}$ is prefix thick with degree $(\frac{1}{2} + \varepsilon) \cdot 2^{n-1}$, because $I_g \subseteq I_1$ and $\mathcal{X}(a)|_{I_1}$ is prefix thick with degree $(\frac{1}{2} + \varepsilon) \cdot 2^{n-1}$.

Finally, we show that there exist strings $a_g \in \mathcal{A}_1$ and $b_g \in \mathcal{B}_1$ such that $a_g|_{[m]-I_g} = b_g|_{[m]-I_g}$. To this end, we will choose the constants $\beta, \gamma, \varepsilon, \kappa$ later such that

$$\mathsf{L}(\mathcal{A}_1 \times \mathcal{B}_1) > 2^{|[m]-I_g| + \log m}.$$

We claim that this implies that there exist strings $a_g \in \mathcal{A}_1$ and $b_g \in \mathcal{B}_1$ such that $a_g|_{[m]-I_g} = b_g|_{[m]-I_g}$. Indeed, suppose otherwise, namely, assume that for every $a \in \mathcal{A}_1$ and $b \in \mathcal{B}_1$ it holds that $a|_{[m]-I_g} \neq b|_{[m]-I_g}$. Consider now the following protocol for $KW_{\mathcal{A}_1 \times \mathcal{B}_1}$: on inputs $a \in \mathcal{A}_1$ and $b \in \mathcal{B}_1$, Alice sends $a|_{[m]-I_g}$, and Bob replies with the coordinate $i \in [m] - I_g$ such that $a_i \neq b_i$ (which exists by the assumption). This protocol transmits at most $|[m] - I_g| + \log m$ bits, and hence its size is at most $2^{|[m]-I_g| + \log m}$. This contradicts, however, our lower bound on $\mathsf{L}(\mathcal{A}_1 \times \mathcal{B}_1)$ above. It follows that there exist strings $a_g \in \mathcal{A}_1$ and $b_g \in \mathcal{B}_1$ such that $a_g|_{[m]-I_g} = b_g|_{[m]-I_g}$, and this concludes the proof.

Choosing the universal constants. We now show how to choose the universal constants $\beta, \gamma, \varepsilon, \kappa$. We choose $\kappa = 8$. The constants β, γ , and ε , are required to satisfy the following constraints:

$$\begin{aligned} \gamma + 2\log e \cdot \varepsilon &\leq 2\log e \cdot \beta^2 \\ \mathsf{L}(\mathcal{A}_1 \times \mathcal{B}_1) &> 2^{|[m]-I_g| + \log m}. \end{aligned}$$

Observe that by our choice of κ it holds that

$$\mathsf{L}(\mathcal{A}_1 \times \mathcal{B}_1) \geq 2^{-(2\gamma+4\log e \cdot \varepsilon) \cdot m - 2\log m - 8} \cdot \mathsf{L}(\mathcal{A} \times \mathcal{B}) > 2^{(1-3\cdot\gamma-4\log e \cdot \varepsilon) \cdot m + \log m},$$

and that $|I_g| \geq (\frac{1}{2} - \beta)^2 \cdot m$. Therefore, a sufficient condition for the second constraint above to be satisfied is

$$1 - 3 \cdot \gamma - 4\log e \cdot \varepsilon \geq 1 - \left(\frac{1}{2} - \beta\right)^2.$$

By rearranging the equations, it follows that it suffices to choose β, γ , and ε such that

$$\gamma \leq \min \left\{ 2\log e \cdot (\beta^2 - \varepsilon), \frac{1}{3} \left(\frac{1}{2} - \beta\right)^2 - 4\log e \cdot \varepsilon \right\}. \quad (7)$$

In fact, it suffices to choose β and γ such that

$$\gamma < \min \left\{ 2\log e \cdot \beta^2, \frac{1}{3} \left(\frac{1}{2} - \beta\right)^2 \right\},$$

and then we can choose ε to be sufficiently small such that Equation (7) holds. It can now be checked that $\beta = 0.12$ and $\gamma = 0.041$ satisfy the last requirement.

6.2 The independence number of \mathcal{G}'

In this section, we prove Lemma 6.2. Let $\gamma, \varepsilon, \mathcal{V}', I, a,$ and $b,$ be as in Lemma 6.1, and let \mathcal{G}' be the sub-graph of \mathcal{G}_{π_1} induced by \mathcal{V}' . We prove that the independence number of \mathcal{V}' is at most

$$2^{-\left(\frac{\log e}{32} \cdot \varepsilon^2 \cdot 2^n - \frac{1}{2} \cdot m \cdot n - 4m - 1\right)} \cdot |\mathcal{V}'|.$$

To this end, let $\mathcal{S} \subseteq \mathcal{V}'$ be a subset of functions that is larger than the above bound. We prove that \mathcal{S} is not an independent set of \mathcal{G}' .

We start by simplifying the notation. Without loss of generality, assume that $I = \{1, \dots, |I|\}$. For every $g \in \mathcal{V}'$, let us denote $\mathcal{X}_g^* = \mathcal{X}_{\pi_1}(g, a)|_I$ and $\mathcal{Y}_g^* = \mathcal{Y}_{\pi_1}(g, b)|_I$. Observe that, for every two functions $g_A, g_B \in \mathcal{S}$, if the sets $\mathcal{X}_{g_A}^*$ and $\mathcal{Y}_{g_B}^*$ intersect then g_A and g_B are neighbors in \mathcal{G}' : indeed, if the two sets intersect then there exist matrices $X \in \mathcal{X}_{\pi_1}(g_A, a)$ and $Y \in \mathcal{Y}_{\pi_1}(g_B, b)$ such that $X_i = Y_i$ for every $i \in I$, and this implies that g_A and g_B are neighbors because $a|_{[m]-I} = b|_{[m]-I}$. Thus, it suffices to prove that there exist functions $g_A, g_B \in \mathcal{S}$ such that $\mathcal{X}_{g_A}^* \cap \mathcal{Y}_{g_B}^* \neq \emptyset$.

In order to prove it, we consider two random functions g_A and g_B that are uniformly distributed over the entire set of balanced functions \mathcal{V}_0 (rather than just \mathcal{S}). For every function $g \in \mathcal{V}_0 \setminus \mathcal{S}$, we define \mathcal{X}_g^* and \mathcal{Y}_g^* to be some arbitrary subsets of $g^{-1}(a)|_I$ and $g^{-1}(b)|_I$ that are prefix thick with degree $\left(\frac{1}{2} + \varepsilon\right) \cdot 2^{n-1}$. We will prove that

$$\Pr[\mathcal{X}_{g_A}^* \cap \mathcal{Y}_{g_B}^* = \emptyset] < \left(\frac{|\mathcal{S}|}{|\mathcal{V}_0|}\right)^2 = \Pr[g_A, g_B \in \mathcal{S}],$$

and this will imply that there exist some functions $g_A, g_B \in \mathcal{S}$ such that $\mathcal{X}_{g_A}^* \cap \mathcal{Y}_{g_B}^* \neq \emptyset$, as required.

To this end, observe that if all the following three events happen simultaneously, then $\mathcal{X}_{g_A}^*$ and $\mathcal{Y}_{g_B}^*$ are *not* disjoint:

1. For every $i \in I$, it holds that $|g_A^{-1}(a_i) \cap g_B^{-1}(b_i)| \leq \left(1 + \frac{\varepsilon}{2}\right) \cdot 2^{n-2}$.
2. The set $\mathcal{X}_{g_A}^* \cap g_B^{-1}(b)|_I$ is prefix thick with degree $\left(\frac{1}{2} + \frac{\varepsilon}{2}\right) \cdot 2^{n-2}$.
3. The set $\mathcal{Y}_{g_B}^* \cap g_A^{-1}(a)|_I$ is prefix thick with degree $\left(\frac{1}{2} + \frac{\varepsilon}{2}\right) \cdot 2^{n-2}$.

The reason is that if all the above three events happen, then $\mathcal{X}_{g_A}^* \cap g_B^{-1}(b)|_I$ and $\mathcal{Y}_{g_B}^* \cap g_A^{-1}(a)|_I$ are both prefix-thick subsets of

$$g_A^{-1}(a)|_I \cap g_B^{-1}(b)|_I = \prod_{i=1}^{|I|} (g_A^{-1}(a_i) \cap g_B^{-1}(b_i))$$

(where we view the sets $g_A^{-1}(a_i) \cap g_B^{-1}(b_i)$ as alphabets), and hence they intersect by Proposition 1.10. Therefore, in order to upper bound the probability that $\mathcal{X}_{g_A}^*$ and $\mathcal{Y}_{g_B}^*$ are disjoint, it is sufficient to upper bound the probability that one of the foregoing events does *not* happen. We will upper bound the probability of each of the three events separately, and then apply the union bound.

We start with upper bounding the probability that Event 2 does not happen. Fix a function g_A . Since $\mathcal{X}_{g_A}^*$ is prefix thick, there exists a subset $\mathcal{X}' \subseteq \mathcal{X}_{g_A}^*$ whose a prefix tree T has minimal degree is greater than $\left(\frac{1}{2} + \varepsilon\right) \cdot 2^{n-1}$. Let T' be the prefix tree of the subset $\mathcal{X}' \cap g_B^{-1}(b)|_I$ of $\mathcal{X}_{g_A}^* \cap g_B^{-1}(b)|_I$. We prove that with high probability over the choice of g_B , the minimal degree of T' is greater than $\left(\frac{1}{2} + \frac{\varepsilon}{2}\right) \cdot 2^{n-2}$, and this will imply the required bound. Fix an internal vertex v of T at

depth i . Observe that if v is also a vertex of T' , then its children in T' are exactly its children in T whose labels belong to $g_B^{-1}(b_{i+1})$. By assumption, v has at least

$$\left(\frac{1}{2} + \varepsilon\right) \cdot 2^{n-1} = \left(\frac{1}{4} + \frac{\varepsilon}{2}\right) \cdot 2^n$$

children in T . Therefore, by Fact 2.2, the probability over the choice of g_B that less than $\left(\frac{1}{4} + \frac{\varepsilon}{4}\right) \cdot 2^{n-1}$ of their labels belong to $g_B^{-1}(b_{i+1})$ is less than

$$2^{-2 \log e \cdot \left(\frac{\varepsilon}{4}\right)^2 \cdot 2^{n-1}} = 2^{-\frac{\log e}{16} \cdot \varepsilon^2 \cdot 2^n}$$

(as $g_B^{-1}(b_{i+1})$ is a uniformly distributed subset of $\{0, 1\}^n$ of size 2^{n-1}). There are at most $2^{m \cdot n}$ vertices in T . By taking a union bound over all of them, we get that the probability that for some internal vertex v in T , it is a vertex of T' and less than

$$\left(\frac{1}{4} + \frac{\varepsilon}{4}\right) \cdot 2^{n-1} = \left(\frac{1}{2} + \frac{\varepsilon}{2}\right) \cdot 2^{n-2}$$

of its children of v belong to $g_B^{-1}(b_{i+1})$ is at most $2^{-\left(\frac{\log e}{16} \cdot \varepsilon^2 \cdot 2^{n-m \cdot n}\right)}$. Now, observe that whenever the latter event does not happen, it holds that the minimum degree of T' is greater than $\left(\frac{1}{2} + \frac{\varepsilon}{2}\right) \cdot 2^{n-2}$. Hence, the probability that Event 2 above does not happen is at most $2^{-\left(\frac{\log e}{16} \cdot \varepsilon^2 \cdot 2^{n-m \cdot n}\right)}$. Similarly, the probability that Event 3 does not happen is at most $2^{-\left(\frac{\log e}{16} \cdot \varepsilon^2 \cdot 2^{n-m \cdot n}\right)}$.

Finally, we upper bound the probability that Event 1 does not happen. Fix any choice of g_A and an index i , and observe that $|g_A^{-1}(a_i)| = 2^{n-1} = \frac{1}{2} \cdot 2^n$. Therefore, by Fact 2.2, the probability over the random choice of g_B that

$$|g_A^{-1}(1) \cap g_B^{-1}(0)| > \left(1 + \frac{\varepsilon}{2}\right) \cdot 2^{n-2} = \left(\frac{1}{2} + \frac{\varepsilon}{4}\right) \cdot 2^{n-1}$$

is at most

$$2^{-2 \log e \cdot \left(\frac{\varepsilon}{4}\right)^2 \cdot 2^{n-1}} = 2^{-\frac{\log e}{16} \cdot \varepsilon^2 \cdot 2^n},$$

and thus the probability that this happens for any $i \in [m]$ is at most $m \cdot 2^{-\frac{\log e}{16} \cdot \varepsilon^2 \cdot 2^n}$. By the union bound, the probability that one of the three events does not happen is at most

$$\begin{aligned} & m \cdot 2^{-\frac{\log e}{16} \cdot \varepsilon^2 \cdot 2^n} + 2^{-\left(\frac{\log e}{16} \cdot \varepsilon^2 \cdot 2^{n-m \cdot n}\right)} + 2^{-\left(\frac{\log e}{16} \cdot \varepsilon^2 \cdot 2^{n-m \cdot n}\right)} \\ & \leq 2^{-\left(\frac{\log e}{16} \cdot \varepsilon^2 \cdot 2^{n-m \cdot n-2}\right)} \\ & = \left(2^{-\left(\frac{\log e}{32} \cdot \varepsilon^2 \cdot 2^{n-\frac{1}{2} \cdot m \cdot n-4 \cdot m-1}\right)} \cdot 2^{-4 \cdot m}\right)^2 \\ & < \left(\frac{|\mathcal{S}|}{|\mathcal{V}'|} \cdot \frac{|\mathcal{V}'|}{|\mathcal{V}_0|}\right)^2 && \text{(assumptions on } \mathcal{S} \text{ and } \mathcal{V}'\text{)} \\ & = \left(\frac{|\mathcal{S}|}{|\mathcal{V}_0|}\right)^2, \end{aligned}$$

as required.

Remark 6.3. There is a small subtlety in the application of Proposition 1.10 above. Specifically, we applied the proposition to strings in which different coordinates belong to different alphabets. The proposition, on the other hand, was only stated for the case of a single alphabet. As noted in Remark 5.3, this is not a problem if all the alphabets are of the same size, but here the alphabets may be of different sizes. Nevertheless, since all the alphabets are of size at most $\left(1 + \frac{\varepsilon}{2}\right) \cdot 2^{n-2}$, we can pretend that they all have exactly that size by adding dummy symbols, and the argument would proceed without a change.

7 A barrier to improving γ

It would have been nice if we could improve the value of the constant γ in the structure theorem to a larger value. In particular, if we could prove our structure theorem with $\gamma = 1$, it would have implied an almost-optimal composition theorem, analogous to the weak KRW conjecture. Indeed, such structure theorems have been proved for other variants of the KRW conjecture [EIRS91, DM16, KM18, dRMN⁺20]. In this section we show that improving the value of γ even to 0.64 would require significant new ideas. Specifically, recall that proof of structure theorem works by showing a lower bound on the chromatic number $\chi(\mathcal{G}_{\pi_1})$. We show that such a lower bound cannot be established for 0.64-live transcripts:

Proposition 7.1. *For every $n \in \mathbb{N}$, for every sufficiently large $m \in \mathbb{N}$ such that $m \geq 2n$, and for every $f : \{0, 1\}^m \rightarrow \{0, 1\}$ such that $L(f) \geq \frac{2^m}{m}$, the following holds: there exists a (standard) protocol Π that solves $KW_f \otimes MUX_n$ and a 0.64-live transcript π_1 of Π such that $\chi(\mathcal{G}_{\pi_1}) = 1$.*

Remark 7.2. We note that the constraint that $m \geq 2n$ in Proposition 7.1 is not very significant, since the interesting regime for deriving formula lower bounds typically assumes that $m \approx 2^n$. Furthermore, we note that the fact that the protocol Π is a standard protocol rather than a partially half-duplex one strengthens our result, since standard protocols can be viewed as partially half-duplex protocols.

The rest of this section is dedicated to proving Proposition 7.1. Let m , n , and f be as in the proposition. Informally, we construct the transcript π_1 such that, for all functions $g_A, g_B : \{0, 1\}^n \rightarrow \{0, 1\}$, the matrices X and Y always disagree on almost all the rows, and the column vectors a and b always disagree on many rows. Such a choice guarantees that there is always a row i such that $X_i \neq Y_i$ and $a_i \neq b_i$, so there are no edges in \mathcal{G}_{π_1} . We construct such a transcript π_1 as follows:

- In order to guarantee that the matrices X and Y always disagree on almost all the rows, we force X to always have $0.08 \cdot m$ ones in its first column, and force Y to always have $0.92 \cdot m$ ones in its first column. A straightforward calculation shows that the fraction of such matrices is at least $2^{-0.64 \cdot m}$.

We still need to guarantee that there are sufficiently many such matrices in $(f \diamond g_A)^{-1}(1)$ and $(f \diamond g_B)^{-1}(0)$. To this end, we choose the associated set $\mathcal{V} \subseteq \mathcal{V}_{\pi_1}$ of balanced functions g such that the number of matrices in $(f \diamond g)^{-1}(1)$ and $(f \diamond g)^{-1}(0)$ with a given first column is always the same. It is not hard to show that the fraction of such functions g is at least 2^{-m} .

- A natural way to guarantee that the column vectors a, b disagree on many rows would be to force a and b to belong to a code with a large distance. Nevertheless, it is not clear how to argue in such a case that the formula complexity $L(\mathcal{A}_{\pi_1}(g) \times \mathcal{B}_{\pi_1}(g))$ is large. Instead, we force both a and b to belong to the same *coset* of a linear code C with a large distance *and a large dimension*.

It is not hard to see that such a construction guarantees that a and b disagree on many rows. Moreover, it can be shown that there exists a choice of a coset of C for which $L(\mathcal{A}_{\pi_1}(g) \times \mathcal{B}_{\pi_1}(g))$ is at least $L(f)$ divided by the number of cosets. This is done by using the fact that the cosets of C form a partition of $\{0, 1\}^m$, and by applying the sub-additivity of formula complexity. We upper bound the number of cosets using the assumption that C has large dimension, thus obtaining the desired lower bound on the formula complexity.

Details follow. We start with setting up some parameters and notation. By applying Varshamov's bound (Theorem 2.33) with $\delta = 0.161$ and $\varepsilon = 0.001$, it follows that there exists a code $C \subseteq \{0, 1\}^m$

with distance $0.161 \cdot m$ and dimension at least $0.362 \cdot m$. Recall that \mathcal{V}_0 denotes the set of all balanced functions from $\{0, 1\}^n$ to $\{0, 1\}$.

We choose the protocol Π to be any protocol that solves $KW_f \otimes MUX_n$, such that in the first messages Alice and Bob communicate the numbers of ones in the first columns of X and Y , and the cosets of C to which a and b belong. Consider a partial transcript π_1 of Π in which Alice and Bob say the following:

1. The first column of X has exactly $\lfloor 0.08 \cdot m \rfloor$ ones.
2. The first column of Y has exactly $\lceil 0.92 \cdot m \rceil$ ones.
3. The string $a = g_A(X)$ belongs to the coset W of C such that the set $\mathcal{A}_W = W \cap f^{-1}(1)$ maximizes the formula complexity $L(\mathcal{A}_W \times f^{-1}(0))$.
4. The string b belongs to the same coset W of C as a , so $b \in \mathcal{B}_W = W \cap f^{-1}(0)$.

Observe that \mathcal{V}_{π_1} is the set of all functions from $\{0, 1\}^n$ to $\{0, 1\}$, and that for every $g \in \mathcal{V}_{\pi_1}$, it holds that $\mathcal{A}_{\pi_1}(g) = \mathcal{A}_W$ and $\mathcal{B}_{\pi_1} = \mathcal{B}_W$. Moreover, for every $g \in \mathcal{V}_{\pi_1}$ and every $a \in \mathcal{A}_W$, it holds that the set $\mathcal{X}_{\pi_1}(g, a)$ is the set of all matrices $X \in g^{-1}(a)$ with exactly $\lfloor 0.08 \cdot m \rfloor$ ones in their first column, and a similar claim holds with $b \in \mathcal{B}_W$, $\mathcal{Y}_{\pi_1}(g, b)$, and $\lceil 0.92 \cdot m \rceil$.

We show that the graph \mathcal{G}_{π_1} does not contain any edge, and therefore $\chi(\mathcal{G}_{\pi_1}) = 1$. Assume for the sake of contradiction that there were two neighbors $g_A, g_B \in \mathcal{V}_{\pi_1}$ in \mathcal{G}_{π_1} . By definition, this implies that g_A and g_B satisfy the weak intersection property: namely, there exist matrices $X \in \mathcal{X}_{\pi_1}(g_A, a)$ and $Y \in \mathcal{Y}_{\pi_1}(g_B, b)$, such that $X_i = Y_i$ for every $i \in [m]$ for which $a_i \neq b_i$, where $a = g_A(X)$ and $b = g_B(Y)$ (or the same statement holds while exchanging g_A and g_B ; without loss of generality, assume that we are in the former case). We know that the first columns of X and Y have exactly $\lfloor 0.08 \cdot m \rfloor$ and $\lceil 0.92 \cdot m \rceil$ ones respectively, and hence X and Y disagree on at least $0.84 \cdot m$ rows. This implies that a and b have to agree on at least $0.84 \cdot m$ coordinates, so the vector $a - b$ (over \mathbb{F}_2) contains at least $0.84 \cdot m$ zeroes. Nevertheless, we assumed that vectors a, b belong to the same coset W of C , and therefore $a - b \in C$. It follows that $a - b$ contains at least $0.161 \cdot m$ ones by the definition of C , which contradicts the conclusion that $a - b$ contains at least $0.84 \cdot m$ zeroes. We reached a contradiction, and therefore there are no edges in \mathcal{G}_{π_1} , as required.

It remains to show that π_1 is γ -alive for $\gamma = 0.64$. To this end, we choose the set \mathcal{V} associated with π_1 to be the set of all functions $g \in \mathcal{V}_0$ that satisfy the following condition: after the first input bit of g is fixed to either 0 or 1, the function g remains balanced. We will show that for every $g \in \mathcal{V}$, $a \in \mathcal{A}_{\pi_1}(g)$, and $b \in \mathcal{B}_{\pi_1}(g)$, it holds that

$$\begin{aligned} |\mathcal{V}| &\geq 2^{-m} \cdot |\mathcal{V}_0| \\ |\mathcal{X}_{\pi_1}(g, a)| &\geq 2^{-\gamma m + 1} \cdot g^{-1}(a) \\ |\mathcal{Y}(g, b)| &\geq 2^{-\gamma m + 1} \cdot g^{-1}(b) \\ \log L(\mathcal{A}_{\pi_1}(g) \times \mathcal{B}_{\pi_1}(g)) &\geq (1 - \gamma) \cdot m + \kappa \log m + \kappa, \end{aligned}$$

where κ is any universal constant. We start with the first equation. In order to choose a function $g \in \mathcal{V}$, it is sufficient to choose the two balanced functions from $\{0, 1\}^{n-1}$ to $\{0, 1\}$ that are obtained by fixing the first input bit of g to 0 and 1 respectively. By Fact 2.4, the number of balanced functions from $\{0, 1\}^{n-1}$ to $\{0, 1\}$ is at least

$$\binom{2^{n-1}}{\frac{1}{2} \cdot 2^{n-1}} \geq \frac{1}{2^{n-1} + 1} \cdot 2^{H_2(\frac{1}{2}) \cdot 2^{n-1}} \geq 2^{2^{n-1} - n}.$$

It follows that

$$\begin{aligned}
|\mathcal{V}| &= \binom{2^{n-1}}{2^{n-1}/2}^2 \\
&\geq \binom{2^{2^{n-1}-n}}{2^{2^{n-1}-n}}^2 \\
&= 2^{2^n-2n}. \\
&\geq 2^{-2n} \cdot |\mathcal{V}_0| \\
&\geq 2^{-m} \cdot |\mathcal{V}_0|,
\end{aligned}$$

where the last inequality follows from the assumption that $m \geq 2n$.

Next, we show that for every $g \in \mathcal{V}$ and $a \in \mathcal{A}_{\pi_1}(g)$ it holds that

$$|\mathcal{X}_{\pi_1}(g, a)| \geq 2^{-\gamma \cdot m+1} \cdot |g^{-1}(a)|,$$

and the analogous inequality for $b \in \mathcal{B}_{\pi_1}(g)$ can be proved similarly. Let $g \in \mathcal{V}$ and $a \in \mathcal{A}_{\pi_1}(g)$, and recall that $\mathcal{X}_{\pi_1}(g, a)$ is the set of matrices $X \in g^{-1}(a)$ such that the first column of X has exactly $\lfloor 0.08 \cdot m \rfloor$ ones. By the assumption that $g \in \mathcal{V}_1$, it follows that for every column vector $v \in \{0, 1\}^m$, there are exactly $2^{m \cdot (n-2)}$ matrices $X \in g^{-1}(a)$ whose first column is v . Hence, it is sufficient to estimate the number of possible first columns of X , namely, the number of binary strings of length m with $\lfloor 0.08 \cdot m \rfloor$ ones. The number of such binary strings is

$$\begin{aligned}
\binom{m}{\lfloor 0.08 \cdot m \rfloor} &\geq \binom{m}{0.079 \cdot m} && \text{(for a sufficiently large } m) \\
&\geq \frac{1}{m+1} \cdot 2^{H_2(0.079) \cdot m} && \text{(Fact 2.4)} \\
&\geq 2^{H_2(0.078) \cdot m} && \text{(for a sufficiently large } m) \\
&= 2^{-(1-H_2(0.078)) \cdot m} \cdot 2^m \\
&\geq 2^{-0.61 \cdot m} \cdot 2^m \\
&\geq 2^{-\gamma \cdot m+1} \cdot 2^m && \text{(for a sufficiently large } m).
\end{aligned}$$

It follows that

$$|\mathcal{X}_{\pi_1}(g, a)| \geq 2^{-\gamma \cdot m+1} \cdot 2^m \cdot 2^{m \cdot (n-2)} = 2^{-\gamma \cdot m+1} \cdot 2^{m \cdot (n-1)} = 2^{-\gamma \cdot m+1} \cdot |g^{-1}(a)|,$$

as required.

Finally, we prove that for every $g \in \mathcal{V}_1$ it holds that

$$\log \mathsf{L}(\mathcal{A}_{\pi_1}(g) \times \mathcal{B}_{\pi_1}(g)) \geq (1 - \gamma) \cdot m + \kappa \log m + \kappa.$$

Let $g \in \mathcal{V}_1$, and recall that $\mathcal{A}_{\pi_1}(g) = \mathcal{A}_W$ and $\mathcal{B}_{\pi_1}(g) = \mathcal{B}_W$. We start by lower bounding the complexity $\mathsf{L}(\mathcal{A}_W \times f^{-1}(0))$. Let \mathcal{W} be the set of cosets of C . Since C has dimension at least $0.362 \cdot m$, it follows that $|\mathcal{W}| \leq 2^{0.638 \cdot m}$. For every $W' \in \mathcal{W}$, let $\mathcal{A}_{W'} = f^{-1}(1) \cap W'$, and recall that W was chosen such that $\mathsf{L}(\mathcal{A}_W \times f^{-1}(0))$ is maximal. Hence, by the sub-additivity of formula complexity, it holds that

$$\mathsf{L}(f) = \mathsf{L}(f^{-1}(1) \times f^{-1}(0)) \leq \sum_{W' \in \mathcal{W}} \mathsf{L}(\mathcal{A}_{W'} \times f^{-1}(0)) \leq |\mathcal{W}| \cdot \mathsf{L}(\mathcal{A}_W \times f^{-1}(0)).$$

It follows that

$$L(\mathcal{A}_W \times f^{-1}(0)) \geq L(f)/|\mathcal{W}|,$$

and therefore

$$\begin{aligned} \log L(\mathcal{A}_W \times f^{-1}(0)) &\geq \log L(f) - \log |\mathcal{W}| \\ &\geq \log \left(\frac{2^m}{m} \right) - \log |\mathcal{W}| && \text{(By choice of } f) \\ &= m - \log(m) - \log |\mathcal{W}| \\ &\geq m - \log(m) - 0.638 \cdot m && (|\mathcal{W}| \leq 2^{0.638 \cdot m}) \\ &= (1 - 0.638) \cdot m - \log(m). \end{aligned}$$

Next, let $\mathcal{B}' = f^{-1}(0) - \mathcal{B}_W$. By the sub-additivity of formula complexity, it holds that

$$L(\mathcal{A}_W \times \mathcal{B}_W) \geq L(\mathcal{A}_W \times f^{-1}(0)) - L(\mathcal{A}_W \times \mathcal{B}').$$

Hence, we can lower bound the formula complexity of the rectangle $\mathcal{A}_W \times \mathcal{B}_W$ by upper bounding the complexity of the rectangle $\mathcal{A}_W \times \mathcal{B}'$. In order to bound the latter complexity, recall that since W is a coset of the linear subspace C , it is the solution space of a (possibly non-homogeneous) system of linear equations. In particular, each string $b \in \mathcal{B}'$ violates at least one of the equations. Let $m' = m - \log |C|$ be the number of the equations in the system, and for each $i \in [m']$, let \mathcal{B}'_i be the set of strings $b \in f^{-1}(0)$ that violate the i -th equation. It holds that $\mathcal{B}' = \bigcup_{i=1}^{m'} \mathcal{B}'_i$, and therefore

$$L(\mathcal{A}_W \times \mathcal{B}') \leq \sum_{i=1}^{m'} L(\mathcal{A}_W \times \mathcal{B}'_i).$$

Now, observe that to determine whether a string $w \in \{0, 1\}^m$ belongs to \mathcal{A}_W or to \mathcal{B}'_i , it suffices to check whether w satisfies the i -th equation or not. This, in turn, amounts to computing the parity function over at most m bits, and therefore has formula complexity at most $4m^2$ by Proposition 2.9. It follows that $L(\mathcal{A}_W \times \mathcal{B}'_i) \leq 4m^2$ for every $i \in [m']$, and hence

$$L(\mathcal{A}_W \times \mathcal{B}') \leq m' \cdot 4m^2 \leq 4m^3.$$

This implies that

$$\begin{aligned} L(\mathcal{A}_W \times \mathcal{B}_W) &\geq L(\mathcal{A}_W \times f^{-1}(0)) - L(\mathcal{A}_W \times \mathcal{B}') \\ &\geq L(\mathcal{A}_W \times f^{-1}(0)) - 4m^3 \\ &\geq 2^{(1-0.638) \cdot m - \log(m)} - 4m^3 \\ &\geq 2^{(1-0.638) \cdot m - \log(m) - 1} && \text{(for a sufficiently large } m). \end{aligned}$$

We conclude that for every universal constant κ and for every sufficiently large value of m it holds that

$$\log L(\mathcal{A}_W \times \mathcal{B}_W) \geq (1 - \gamma) \cdot m + \kappa \cdot \log m + \kappa,$$

where $\gamma = 0.64$, as required.

Acknowledgement. The author is grateful to Ronen Shaltiel for many useful discussions and ideas, and to Sajin Koroth for suggesting the question of whether the KRW conjecture holds for the notion of strong composition. The author would like to thank Yahel Manor, Lê Minh Quý, and

anonymous referees for comments that improved the presentation of this paper, and to Noga Alon, Ishay Haviv, Lianna Hambardzumyan, Alon Orlitzky, Avi Wigderson, and an anonymous referee, for references on the graph equality problem and related problems in the literature. The author would also like to thank Alexander Smal for an in-depth explanation of his joint works with Artur Ignatiev and Ivan Mihajlin [MS21, IMS22], and to Ben Barber for his kind explanations on the version of Harper’s theorem for the general Hamming graph. Finally, the author is grateful to Ville Salo for pointing out his joint work with Ilkka Törmä [ST14] and its connection to prefix-thick sets.

References

- [AB23] Josh Alman and Jaroslaw Blasiok. Matrix multiplication and number on the forehead communication. *CoRR*, abs/2302.11476, 2023.
- [Alo87] Noga Alon. Monochromatic directed walks in arc-colored directed graphs. *Acta Mathematica Hungarica*, 49:163–167, 1987.
- [And87] Alexander E. Andreev. On a method for obtaining more than quadratic effective lower bounds for the complexity of π -schemes. *Moscow University Mathematics Bulletin*, 42(1):24–29, 1987.
- [AO95] Noga Alon and Alon Orlitsky. Repeated communication and ramsey graphs. *IEEE Trans. Inf. Theory*, 41(5):1276–1289, 1995.
- [AO96] Noga Alon and Alon Orlitsky. Source coding and graph entropies. *IEEE Trans. Inf. Theory*, 42(5):1329–1339, 1996.
- [ARS02] Richard P. Anstee, Lajos Rónyai, and Attila Sali. Shattering news. *Graphs Comb.*, 18(1):59–73, 2002.
- [BBL⁺15] Jop Briët, Harry Buhrman, Debbie W. Leung, Teresa Piovesan, and Florian Speelman. Round elimination in exact communication complexity. In Salman Beigi and Robert König, editors, *10th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2015, May 20-22, 2015, Brussels, Belgium*, volume 44 of *LIPICs*, pages 206–225. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015.
- [CFL83] Ashok K. Chandra, Merrick L. Furst, and Richard J. Lipton. Multi-party protocols. In David S. Johnson, Ronald Fagin, Michael L. Fredman, David Harel, Richard M. Karp, Nancy A. Lynch, Christos H. Papadimitriou, Ronald L. Rivest, Walter L. Ruzzo, and Joel I. Seiferas, editors, *Proceedings of the 15th Annual ACM Symposium on Theory of Computing, 25-27 April, 1983, Boston, Massachusetts, USA*, pages 94–99. ACM, 1983.
- [CT91] Thomas M. Cover and Joy A. Thomas. *Elements of information theory*. Wiley-Interscience, 1991.
- [DM16] Irit Dinur and Or Meir. Toward the KRW composition conjecture: Cubic formula lower bounds via communication complexity. In *31st Conference on Computational Complexity, CCC 2016, May 29 to June 1, 2016, Tokyo, Japan*, pages 3:1–3:51, 2016.
- [DP09] Devdatt P. Dubhashi and Alessandro Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009.

- [dRMN⁺20] Susanna F. de Rezende, Or Meir, Jakob Nordström, Toniann Pitassi, and Robert Robere. KRW composition theorems via lifting. In Sandy Irani, editor, *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 43–49. IEEE, 2020. Also available as ECCC TR23-078.
- [dRMN⁺24] Susanna F. de Rezende, Or Meir, Jakob Nordström, Toniann Pitassi, and Robert Robere. KRW composition theorems via lifting. *Comput. Complex.*, 33(1):4, 2024. Also available as ECCC TR23-078.
- [dW01] Ronald de Wolf. *Quantum Computing and Communication Complexity*. PhD thesis, Universiteit van Amsterdam, 2001.
- [EIRS91] Jeff Edmonds, Russell Impagliazzo, Steven Rudich, and Jiří Sgall. Communication complexity towards lower bounds on circuit depth. In *32nd Annual Symposium on Foundations of Computer Science, San Juan, Puerto Rico, 1-4 October 1991*, pages 249–257. IEEE Computer Society, 1991.
- [FMT21] Yuval Filmus, Or Meir, and Avishay Tal. Shrinkage under random projections, and cubic formula lower bounds for AC⁰. In James R. Lee, editor, *12th Innovations in Theoretical Computer Science Conference, ITCS 2021, January 6-8, 2021, Virtual Conference*, volume 185 of *LIPICs*, pages 89:1–89:7. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- [GMWW14] Dmitry Gavinsky, Or Meir, Omri Weinstein, and Avi Wigderson. Toward better formula lower bounds: an information complexity approach to the KRW composition conjecture. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 213–222, 2014.
- [Hås93] Johan Håstad. The shrinkage exponent is 2. In *34th Annual Symposium on Foundations of Computer Science, Palo Alto, California, USA, 3-5 November 1993*, pages 114–123. IEEE Computer Society, 1993.
- [HIMS18] Kenneth Hoover, Russell Impagliazzo, Ivan Mihajlin, and Alexander V. Smal. Half-duplex communication complexity. In *29th International Symposium on Algorithms and Computation, ISAAC 2018, December 16-19, 2018, Jiaoxi, Yilan, Taiwan*, pages 10:1–10:12, 2018.
- [HW93] Johan Håstad and Avi Wigderson. Composition of the universal relation. In *Advances in computational complexity theory, AMS-DIMACS*, 1993.
- [ILO20] Rahul Ilango, Bruno Loff, and Igor C. Oliveira. Np-hardness of circuit minimization for multi-output functions. In Shubhangi Saraf, editor, *35th Computational Complexity Conference, CCC 2020, July 28-31, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 169 of *LIPICs*, pages 22:1–22:36. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [IMS22] Artur Ignatiev, Ivan Mihajlin, and Alexander Smal. Super-cubic lower bound for generalized karchmer-wigderson games. In Sang Won Bae and Heejin Park, editors, *33rd International Symposium on Algorithms and Computation, ISAAC 2022, December 19-21, 2022, Seoul, Korea*, volume 248 of *LIPICs*, pages 66:1–66:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

- [IN93] Russell Impagliazzo and Noam Nisan. The effect of random restrictions on formula size. *Random Struct. Algorithms*, 4(2):121–134, 1993.
- [Khr72] V. M. Khrapchenko. A method of obtaining lower bounds for the complexity of π -schemes. *Mathematical Notes Academy of Sciences USSR*, 10:474–479, 1972.
- [KKN92] Mauricio Karchmer, Eyal Kushilevitz, and Noam Nisan. Fractional covers and communication complexity. In *Proceedings of the Seventh Annual Structure in Complexity Theory Conference, Boston, Massachusetts, USA, June 22-25, 1992*, pages 262–274. IEEE Computer Society, 1992.
- [KM18] Sajin Koroth and Or Meir. Improved composition theorems for functions and relations. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM '18)*, volume 116 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 48:1–48:18, August 2018.
- [KRW91] Mauricio Karchmer, Ran Raz, and Avi Wigderson. Super-logarithmic depth lower bounds via direct sum in communication complexity. In *Proceedings of the Sixth Annual Structure in Complexity Theory Conference, Chicago, Illinois, USA, June 30 - July 3, 1991*, pages 299–304. IEEE Computer Society, 1991.
- [KW88] Mauricio Karchmer and Avi Wigderson. Monotone circuits for connectivity require super-logarithmic depth. In Janos Simon, editor, *Proceedings of the 20th Annual ACM Symposium on Theory of Computing, May 2-4, 1988, Chicago, Illinois, USA*, pages 539–550. ACM, 1988.
- [LY93] Carsten Lund and Mihalis Yannakakis. On the hardness of approximating minimization problems. In S. Rao Kosaraju, David S. Johnson, and Alok Aggarwal, editors, *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, May 16-18, 1993, San Diego, CA, USA*, pages 286–293. ACM, 1993.
- [Mei20] Or Meir. Toward better depth lower bounds: Two results on the multiplexor relation. *Computational Complexity*, 29(1):4, 2020. Available on ECCC as TR19-120.
- [MS21] Ivan Mihajlin and Alexander Smal. Toward better depth lower bounds: The XOR-KRW conjecture. In Valentine Kabanets, editor, *36th Computational Complexity Conference, CCC 2021, July 20-23, 2021, Toronto, Ontario, Canada (Virtual Conference)*, volume 200 of *LIPIcs*, pages 38:1–38:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- [Paj85] Alain Pajor. Sous-espaces l_1^n des espaces de banach. *Éditions Hermann*, 16, 1985.
- [PZ93] Mike Paterson and Uri Zwick. Shrinkage of De Morgan formulae under restriction. *Random Struct. Algorithms*, 4(2):135–150, 1993.
- [Raz90] Alexander A. Razborov. Applications of matrix methods to the theory of lower bounds in computational complexity. *Combinatorica*, 10(1):81–93, 1990.
- [RM97] Ran Raz and Pierre McKenzie. Separation of the monotone NC hierarchy. In *38th Annual Symposium on Foundations of Computer Science, FOCS '97, Miami Beach, Florida, USA, October 19-22, 1997*, pages 234–243, 1997.
- [Sal21] Ville Salo. Trees in positive entropy subshifts. *Axioms*, 10(2):77, 2021.

- [Sau72] Norbert Sauer. On the density of families of sets. *J. Comb. Theory, Ser. A*, 13(1):145–147, 1972.
- [She72] Saharon Shelah. A combinatorial problem; stability and order for models and theories in infinitary languages. *Pacific Journal of Mathematics*, 41:247–261, 1972.
- [ST14] Ville Salo and Ilkka Törmä. Playing with subshifts. *Fundam. Informaticae*, 132(1):131–152, 2014.
- [Sub61] Bella Abramovna Subbotovskaya. Realizations of linear functions by formulas using +, ·, -, . *Soviet Mathematics Doklady*, 2:110–112, 1961.
- [Tal14] Avishay Tal. Shrinkage of De Morgan formulae by spectral techniques. In *Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS '14)*, pages 551–560, 2014.
- [TZ97] Gábor Tardos and Uri Zwick. The communication complexity of the universal relation. In *Proceedings of the Twelfth Annual IEEE Conference on Computational Complexity*, pages 247–259, 1997.
- [Var57] R. R. Varshamov. Estimate of the number of signals in error correcting codes. *Doklady Akadonii Nauk*, pages 739–741, 1957.
- [Yao79] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing (preliminary report). In *STOC*, pages 209–213, 1979.