

Improved Composition Theorems for Functions and Relations

Sajin Korothe¹

Department of Computer Science, University of Haifa, Haifa 3498838, Israel

sajin@csweb.haifa.ac.il

 <https://orcid.org/0000-0002-7989-1963>

Or Meir²

Department of Computer Science, University of Haifa, Haifa 3498838, Israel

ormeir@cs.haifa.ac.il

 <https://orcid.org/0000-0001-5031-0750>

Abstract

One of the central problems in complexity theory is to prove super-logarithmic depth bounds for circuits computing a problem in P , i.e., to prove that P is not contained in NC^1 . As an approach for this question, Karchmer, Raz and Wigderson [5] proposed a conjecture called the KRW conjecture, which if true, would imply that P is not contained in NC^1 .

Since proving this conjecture is currently considered an extremely difficult problem, previous works by Edmonds, Impagliazzo, Rudich and Sgall [1], Håstad and Wigderson [3] and Gavinsky, Meir, Weinstein and Wigderson [2] considered weaker variants of the conjecture. In this work we significantly improve the parameters in these variants, achieving almost tight lower bounds.

2012 ACM Subject Classification Theory of computation → Communication complexity, Theory of computation → Circuit complexity, Theory of computation → Complexity classes

Keywords and phrases circuit complexity, communication complexity, KRW conjecture, composition

Digital Object Identifier 10.4230/LIPIcs.APPROX-RANDOM.2018.48

1 Introduction

The holy grail question in circuit complexity is to prove super-polynomial size lower bounds for NP , i.e., to show that NP is not contained in $P/poly$. This is considered to be an extremely difficult problem and even after years of research we do not even know a super-linear size lower bound. Hence, a natural approach is to prove lower bounds for more restricted classes of circuits. One such restricted class of circuits are NC^1 , which are circuits of polynomial size, logarithmic depth and bounded fan-in. It is widely believed that NC^1 does not contain P . However even this problem is deemed very hard. In particular, we do not even know super-linear lower bounds for NC^1 circuits computing a function even in $NEXP$.

An approach for separating P from NC^1 was suggested by Karchmer, Raz and Wigderson [5]. They conjectured that the depth complexity of Boolean functions adds up under a certain composition of Boolean functions. We refer to this as the KRW conjecture, and if it is true, would give an explicit function in P which does not have NC^1 circuits. In order to state the conjecture we define the following composition of functions.

¹ Supported by the Israel Science Foundation (grant No. 1445/16)

² Partially supported by the Israel Science Foundation (grant No. 1445/16)



© Sajin Korothe and Or Meir;

licensed under Creative Commons License CC-BY

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2018).

Editors: Eric Blais, Klaus Jansen, José D. P. Rolim, and David Steurer; Article No. 48; pp. 48:1–48:18



Leibniz International Proceedings in Informatics

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

► **Definition 1** (Composition). Given two arbitrary Boolean functions $f : \{0, 1\}^m \rightarrow \{0, 1\}$, $g : \{0, 1\}^n \rightarrow \{0, 1\}$ we define their *composition* to be the Boolean function $f \diamond g : (\{0, 1\}^n)^m \rightarrow \{0, 1\}$ obtained as follows:

$$f \diamond g(x_1, x_2, \dots, x_m) = f(g(x_1), g(x_2), \dots, g(x_m))$$

where $x_i \in \{0, 1\}^n, 1 \leq i \leq m$.

We recall a standard definition of *depth complexity*. For a Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, its depth complexity denoted by $D(f)$, is the smallest depth of a circuit of AND, OR and NOT gates of fan-in 2 that computes f . We now state the KRW conjecture.

► **Conjecture 2** (The KRW conjecture, [5]). *Let $f : \{0, 1\}^m \rightarrow \{0, 1\}$ and $g : \{0, 1\}^n \rightarrow \{0, 1\}$ be two arbitrary non-constant Boolean functions. Then the following holds true³,*

$$D(f \diamond g) \approx D(f) + D(g)$$

Since this is a difficult conjecture previous works [1, 2, 3] proved lower bounds for two simplified variants of the KRW conjecture. In this work we improve the parameters of these lower bounds, motivated by the value of these parameters in the target application. Though the two simplified variants of the KRW conjecture are important milestones, we would like to point out that settling these two variants of the conjecture do not imply any circuit lower bound.

1.1 Background

1.1.1 Karchmer-Wigderson relations

Karchmer and Wigderson [6] established an interesting connection between the depth complexity of a Boolean function f and the communication complexity of an associated relation, which is called the Karchmer-Wigderson relation, and is denoted by KW_f . They proved that the depth complexity of f is equal to the deterministic communication complexity of KW_f .

The Karchmer-Wigderson relation associated with a Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is the following communication problem. Alice is given an input $x \in f^{-1}(0)$ and Bob is given an input $y \in f^{-1}(1)$. The objective of the players is to find an index $i \in [n]$ such that $x_i \neq y_i$.

In the rest of this paper we refer to the Karchmer-Wigderson relation as “KW relation”. We also denote the deterministic communication complexity of KW_f by $CC(KW_f)$. In the rest of this paper, when we say communication complexity, we mean deterministic communication complexity.

1.1.2 The KW relation of composition

We study the KRW conjecturing using the Karchmer-Wigderson framework. To this end, we describe how the KW relation related to the composition $f \diamond g$ looks like for arbitrary Boolean functions f, g . In the KW relation $KW_{f \diamond g}$, Alice and Bob’s inputs are conveniently viewed as $m \times n$ matrices X, Y , respectively. Given an $m \times n$ binary matrix X , we define $g(X)$ to be an m bit binary vector obtained by applying g to the rows of X .

The KW relation $KW_{f \diamond g}$ associated with the composition $f \diamond g$ is the following communication problem.

³ The approximate equality in the statement of the conjecture is intentionally left vague, since there are multiple possible definitions which are weaker than strict equality, but which would still imply the $\mathbf{P} \not\subseteq \mathbf{NC}^1$ separation.

- Alice gets a matrix $X \in \{0, 1\}^{m \times n}$ with the promise that the vector $a = g(X)$ is such that $f(a) = 0$.
- Bob gets a matrix $Y \in \{0, 1\}^{m \times n}$ with the promise that the vector $b = g(Y)$ is such that $f(b) = 1$.
- The goal of the players is to find a pair of indices $(i, j) \in [m] \times [n]$ such that $X_{i,j} \neq Y_{i,j}$. It is easy to see that $CC(KW_{f \circ g}) \leq CC(KW_f) + CC(KW_g)$. The KRW conjecture says that the upper bound is essentially optimal for $KW_{f \circ g}$.

1.1.3 Universal relation and its composition

Because of the difficulty of the original conjecture, Karchmer, Raz and Wigderson [5] suggested studying a simpler variant of the conjecture. To this end, they defined a simplified variant of KW relations called the universal relation. The *universal relation* on n bits, denoted by U_n , is a promise problem where Alice is given an input $x \in \{0, 1\}^n$ and Bob is given an input $y \in \{0, 1\}^n$ with the guarantee that $x \neq y$. The goal of the players is to find an index $i \in [n]$ such that $x_i \neq y_i$. This relation is called the universal relation in the sense that KW relation of any Boolean function reduces to it. Tardos and Zwick [9] proved that deterministic communication complexity of U_n is $n + 1$.

Karchmer, Raz and Wigderson [5] suggested to study the composition of universal relations defined next. The composition of the universal relation U_m with the universal relation U_n , denoted by $U_m \diamond U_n$, is defined as the following communication problem.

- Alice gets a matrix $X \in \{0, 1\}^{m \times n}$ and a vector $a \in \{0, 1\}^m$.
- Bob gets a matrix $Y \in \{0, 1\}^{m \times n}$ and a vector $b \in \{0, 1\}^m$.
- They are guaranteed that $a \neq b$.
- They are guaranteed that for any $i \in [m]$, whenever $a_i \neq b_i$, the corresponding rows of the matrices, denoted by X_i and Y_i , are not equal.
- The goal of the players is to find a pair of indices $(i, j) \in [m] \times [n]$ such that $X_{i,j} \neq Y_{i,j}$. This problem generalizes the KW relation corresponding to the composition of Boolean functions. As a first step towards the KRW conjecture, [5] suggested to prove the following:

► **Conjecture 3** (Analogue of KRW conjecture for $U_m \diamond U_n$, [5]). $CC(U_m \diamond U_n) \approx CC(U_m) + CC(U_n) \approx m + n$

The first progress towards this conjecture on the composition of universal relations was made by Edmonds et al. [1] who proved that $CC(U_n \diamond U_n) \geq 2n - O(\sqrt{n})$. Later Håstad and Wigderson [3] improved on the results of [1] using a completely different proof strategy. They [3] proved that for sufficiently large n , $CC(U_n \diamond U_n) \geq 2n - 1$.

1.1.4 Composition of a function with the universal relation

A variant of the above conjecture, which is closer to the original KRW conjecture, is a conjecture dealing with composition of an arbitrary Boolean function with a universal relation. It was suggested by Gavinsky et al. [2]. Given an arbitrary Boolean function $f : \{0, 1\}^m \rightarrow \{0, 1\}$, they define the KW relation $KW_{f \circ U_n}$ as the following promise problem.

- Alice gets a matrix $X \in \{0, 1\}^{m \times n}$ and a vector $a \in \{0, 1\}^m$ such that $f(a) = 0$.
- Bob gets a matrix $Y \in \{0, 1\}^{m \times n}$ and a vector $b \in \{0, 1\}^m$ such that $f(b) = 1$.
- They are guaranteed that for any $i \in [m]$ whenever $a_i \neq b_i$, the corresponding rows of the matrices, denoted by X_i and Y_i , are not equal.
- The goal of the players is to find a pair of indices $(i, j) \in [m] \times [n]$ such that $X_{i,j} \neq Y_{i,j}$.

This definition is a natural candidate for the composition of a function with the universal relation as any instance of $KW_{f \diamond U_n}$ as defined above is also a legal instance of $U_m \diamond U_n$. In addition, any instance of $KW_{f \diamond g}$ where $f : \{0, 1\}^m \rightarrow \{0, 1\}$ and $g : \{0, 1\}^n \rightarrow \{0, 1\}$ is also a legal instance of $KW_{f \diamond U_n}$ if we set $a = g(X)$ and $b = g(Y)$.

We now describe the definition of the promise version of the composition of a function f with a universal relation, due to [2]. As in the above version, goal of Alice and Bob is to find an entry (i, j) on which X and Y differ, but they are allowed to reject if there exists an index $i \in [m]$ such that $a_i \neq b_i$ but $X_i = Y_i$.

► **Remark.** It is easy to see that the complexity of the problem over all inputs is larger by at most two bits than the complexity of the promise problem.

Gavinsky et al. [2] proposed the following analogous conjecture for $KW_{f \diamond U_n}$:

► **Conjecture 4** (Conjecture for $f \diamond U_n$, [2]). $CC(KW_{f \diamond U_n}) \approx CC(KW_f) + CC(U_n) \approx CC(KW_f) + n$

They [2] also proved the following lower bound on the composition⁴.

► **Theorem 5** (Lower bound for $f \diamond U_n$, [2]). *For any $m, n \in \mathbb{N}$, and any non-constant function $f : \{0, 1\}^m \rightarrow \{0, 1\}$, $CC(KW_{f \diamond U_n}) \geq \log(L(f)) + n - O\left(1 + \frac{m}{n}\right) \log m$.*

1.1.5 Known results

In this work, we are interested in improving the dependence on m in the known lower bounds, and in particular in the setting where $m \gg n$. The early results of Edmonds et al. [1] and Håstad and Wigderson [3], established lower bounds for the composition of universal relation with itself when $m = n$. However, these results can be generalized in a straightforward manner for proving lower bounds for communication complexity of $U_m \diamond U_n$.

A straightforward generalization of the [1] bound for $m \neq n$ yields the following result: $CC(U_m \diamond U_n) \geq m + n - O(\sqrt{m})$

A straightforward generalization of the Håstad and Wigderson [3] result only gives a lower bound of $2n - o(1)$ for $CC(U_m \diamond U_n)$. Note that this lower bound is independent of m and hence is far from the conjectured lower bound of $m + n$ when $m \gg n$.

Gavinsky et al. studied $CC(f \diamond U_n)$ and proved a lower bound of $\log(L(f)) + n - O\left(1 + \frac{m}{n}\right) \log m$. The lower bound proved by [2] is for m and n which are not necessarily equal. However, similar to the other known lower bounds it also has a loss term depending on m , $O\left(\frac{m}{n} \log m\right)$, which becomes significant if $m \gg n$.

1.2 Our results

We overcome the additive losses in the above lower bounds and obtain bounds which are optimal except for an $O(\log^* m)$ additive term.

► **Theorem 6.** *For any $m, n \in \mathbb{N}$ with $n \geq 6 \log m$, and any non-constant function $f : \{0, 1\}^m \rightarrow \{0, 1\}$, $CC(KW_{f \diamond U_n}) \geq \log L(f) + n - O(\log^* m)$ where $L(f)$ is the formula complexity of the function f .*

⁴ Note that the lower bound depends on the logarithm of the formula complexity of f , denoted by $\log(L(f))$, instead of the depth complexity of f denoted by $D(f)$. However, the parameters $\log(L(f))$ and $D(f)$ are tightly related.

We also prove a similar lower bound for the deterministic communication complexity of $U_m \diamond U_n$.

► **Theorem 7.** *For any $m, n \in \mathbb{N}$ with $n \geq 6 \log m$, $CC(U_m \diamond U_n) \geq m + n - O(\log^* m)$*

But for sake of simplicity, we prove the following lower bounds with $O(\log m)$ additive losses and defer the proof of $O(\log^* m)$ improvement to the full version of the paper.

► **Theorem 8.** *For any $m, n \in \mathbb{N}$, and any non-constant function $f : \{0, 1\}^m \rightarrow \{0, 1\}$, $CC(KW_{f \diamond U_n}) \geq \log L(f) + n - O(\log m)$ where $L(f)$ is the formula complexity of the function f .*

As a corollary we get the following lower bound for the deterministic communication complexity of $U_m \diamond U_n$.

► **Corollary 9.** *For any $m, n \in \mathbb{N}$, $CC(U_m \diamond U_n) \geq m + n - O(\log m)$*

Proof. As noted earlier, for any non-constant Boolean function $f : \{0, 1\}^m \rightarrow \{0, 1\}$, the KW relation associated with $f \diamond U_n$ reduces to the communication problem $U_m \diamond U_n$. In particular, this implies that for any non-constant $f : \{0, 1\}^m \rightarrow \{0, 1\}$, it holds that $CC(U_m \diamond U_n) \geq CC(KW_{f \diamond U_n})$. Since there exists Boolean functions with formula complexity $\frac{2^m}{\log m}$ among the set of all m -bit Boolean functions, we get by Theorem 8 that, $CC(U_m \diamond U_n) \geq \log\left(\frac{2^m}{\log m}\right) + n - O(\log m) = m + n - O(\log m)$. ◀

Note that in the improved lower bounds, since the losses are only $O(\log^* m)$, we do not get the improved lower bound for $U_m \diamond U_n$ as a corollary from the lower bound for $f \diamond U_n$. But using a very similar proof, an independent lower bound can be derived for $U_m \diamond U_n$.

Motivation

Bridging the gap between the known lower bounds and the conjectured lower bounds is an important pursuit in itself. More importantly, the KRW conjecture implies the two weaker conjectures up to sub-logarithmic additive losses. Thus it is necessary to prove the weaker variants with such losses if one it to settle the KRW conjecture. Our work obtains the first known lower bounds on the two variants of the conjecture, with additive losses smaller than the losses implied if the KRW conjecture is true.

However, the improved parameters in our lower bounds are also significant for the following reason. As we mentioned above, the main motivation for studying the KRW conjecture is that it implies $\mathbf{P} \not\subseteq \mathbf{NC}^1$. In fact, in order to obtain this implication, it suffices to prove a relaxed variant of the conjecture in which g is a random function (but f is an arbitrary function). This relaxation seems to be closer to our reach, since [2] have already proved the conjecture when g is replaced with the universal relation and f is an arbitrary function.

However, in order to derive $\mathbf{P} \not\subseteq \mathbf{NC}^1$ from the relaxation, we need to have lower bounds that are meaningful for values of m that satisfy $m = n^{\omega(1)}$. Unfortunately, the result of [2] does not give a meaningful bound when $m \geq n^2$. If we aspire to prove this relaxation of the KRW conjecture, we need to have lower bounds that work for larger values of m . Our Theorem 8 achieves exactly that: it gives a bound on $KW_{f \diamond U_n}$ that is meaningful as long as $m = o(2^n)$, which is good enough for our purposes.

We should mention that there is another relaxation of the KRW conjecture that implies $\mathbf{P} \neq \mathbf{NC}^1$, in which f is a random function and g is an arbitrary function. However, this relaxation seems farther from our reach, since we do not have any result on $KW_{U_m \diamond g}$ when g is an arbitrary function.

Results for formula lower bounds

We also have a similar (though somewhat weaker) result for formula lower bounds. The analogue of the KRW conjecture for formula lower bounds says that $L(f \diamond g) \approx L(f) \cdot L(g)$ (see [2] for details). This can also be stated in the Karchmer-Wigderson framework: To this end, for any communication problem R , let us denote by $L(R)$ the smallest number of distinct transcripts in a deterministic protocol that solves R . Then, the analogue of the KRW conjecture is to say that $L(KW_{f \diamond g}) \approx L(KW_f) \cdot L(KW_g)$ or equivalently $\log L(KW_{f \diamond g}) \approx \log L(KW_f) + \log L(KW_g)$. We prove that $\log L(KW_{f \diamond U_n}) \geq \log L(KW_f) + n - O\left(\sqrt{\log L(KW_f)} + \log m\right)$. While this result is weaker than our result for depth complexity due to the loss of $O(\sqrt{\log L(KW_f)})$, it is still stronger than the corresponding results of [3, 2] when $\log L(KW_f) \gg n^5$. We defer the proof to the full version of this paper.

1.3 Our techniques

In this section we provide an overview of the proof of our main result (Theorem 8). Our approach is based on a proof strategy due to [1]. We begin by describing this proof strategy, and then describe our new ideas. Recall that the inputs of the parties are pairs (X, a) and (Y, b) where X, Y are $m \times n$ matrices and $a \in f^{-1}(0), b \in f^{-1}(1)$ are m bit column vectors. Fix a deterministic protocol Π that solves $KW_{f \diamond U_n}$. We wish to prove that Π must transmit $\approx \text{CC}(KW_f) + n$ bits.

The basic intuition that underlies the proof consists of the following three observations:

- Morally, in order to solve $KW_{f \diamond U_n}$, the parties have to solve the universal relation on one of the rows of X, Y . To this end, they have to transmit at least n bits about this row.
- However, solving the universal relation on a row X_i, Y_i only makes sense if the parties are guaranteed that $X_i \neq Y_i$, since otherwise they might waste their communication on equal rows.
- Intuitively, in order to find rows X_i, Y_i that are guaranteed to be different, the parties must solve KW_f , and to this end they have to transmit at least $\text{CC}(KW_f)$ bits.

Therefore, the total amount of communication must be at least $\approx \text{CC}(KW_f) + n$.

High-level idea

We now sketch the argument that is based on this intuition. We partition the communication of the protocol Π into two stages – intuitively, the parties should solve KW_f on a, b in the first stage, and then solve the universal relation on a row of X, Y in the second stage. Formally, the first stage is defined as the first $\text{CC}(KW_f) - \alpha$ bits that are transmitted in the protocol, where α is some small “slack term”, and the second stage consists of the remaining bits. We prove that the parties must transmit approximately n bits during the second stage, and this implies that the protocol must transmit approximately $\text{CC}(KW_f) + n$ bits in total.

We start by making the following observation: If, during the first stage, the parties only “talk about” a, b , then it is easy to prove that they have to transmit at least n bits in the second stage: morally, this is because in the second stage they still have to solve the universal relation on one of the rows of X, Y , and must do so “from scratch” (since they did not talk about this row before).

⁵ the work of [1] does not provide a result for formula lower bounds

The challenging case is when the parties “talk about” X, Y during the first stage. We deal with this case by observing that talking about X, Y during the first stage is useless, for the following reason: during the first stage, the parties have not finished solving KW_f yet, since they communicated less than $\text{CC}(KW_f)$ bits. Thus, at this point, they do not know any row i where $a_i \neq b_i$, and therefore they do not know any row i for which they are guaranteed that $X_i \neq Y_i$. This means that any communication about X, Y during the first stage is likely to be wasted on rows where $X_i = Y_i$, and hence will not help the parties to solve the problem.

In short, we observe that any communication about X, Y during the first stage is useless. Therefore, any optimal protocol should behave roughly as if the parties do not talk about X, Y at all during the first stage. However, in such case the parties must transmit n bits during the second stage, and this is what we want to prove.

An adversary argument

The foregoing idea is formalized using an adversary argument. This means that we view the inputs of the parties as if they are chosen by an adversary that can adapt to the messages sent by the parties. We now describe the behavior of our adversary. The adversary starts by letting the parties talk throughout the first stage, and chooses the messages such that no more than $\text{CC}(KW_f) - \alpha$ bits of information are revealed. At the end of the first stage, the adversary looks at the transcript of the communication so far, and partitions the rows of X, Y into two types:

- “Revealed rows”, about which the parties talked much (i.e., more than τ bits for some small parameter τ).
- “Unrevealed rows”, about which the parties talked a little (i.e., at most τ bits).

Intuitively, if the parties end up solving the universal relation on one of the unrevealed rows, then they have to transmit about $n - \tau \approx n$ bits in the second stage, which is what we want to prove.

Hence, the adversary only needs to worry about the revealed rows. In order to deal with the revealed rows, the adversary chooses the inputs of the parties such that $a_i = b_i$ for every revealed row X_i, Y_i . This allows the adversary to prevent the parties from solving the universal relation on a revealed row X_i, Y_i , since the adversary is free to set $X_i = Y_i$ whenever they attempt to do so. It follows that the parties must solve the universal relation on an unrevealed row, and therefore they must transmit roughly n bits during the second stage, as required.

In order for this argument to go through, we must make sure that the adversary is indeed capable of setting $a_i = b_i$ for every revealed row. In principle, the adversary should be able to set $a_i = b_i$ for some rows because the parties have not yet finished solving KW_f . However, the *number* of rows for which the adversary can set $a_i = b_i$ depends on how far the parties are from solving KW_f exactly. Morally, it can be shown that if the parties talked at most $\text{CC}(KW_f) - k$ bits about a and b , then the adversary can set $a_i = b_i$ for about k rows. Therefore, in order for the argument to work, the number of revealed rows should be at most k . This condition depends, in turn, on the choice of the threshold τ (which determines which rows are considered “revealed”) and on the choice of the slack term α (which determines the length of the first stage). This is our point of departure from the work of [1].

The analysis of [1]

The analysis of [1] sets both the threshold τ and the slack term α to⁶ $O(\sqrt{\text{CC}(KW_f)})$ (so the length of the first stage is $\text{CC}(KW_f) - O(\sqrt{\text{CC}(KW_f)})$ bits). The analysis proceeds as follows: In the first stage, the parties talked at most $\text{CC}(KW_f)$ bits about the matrices X, Y (since they talked at most $\text{CC}(KW_f) - O(\sqrt{\text{CC}(KW_f)})$ bits overall). By Markov's inequality, this implies that the number of revealed rows is at most

$$\frac{\text{CC}(KW_f)}{\tau} = O(\sqrt{\text{CC}(KW_f)}).$$

On the other hand, the parties talked at most $\text{CC}(KW_f) - O(\sqrt{\text{CC}(KW_f)})$ bits about a and b (again, since they talked at most $\text{CC}(KW_f) - O(\sqrt{\text{CC}(KW_f)})$ bits overall), and hence the adversary can set $a_i = b_i$ for $O(\sqrt{\text{CC}(KW_f)})$ rows. Thus, for an appropriate choice of the parameters, the adversary can set $a_i = b_i$ for all the revealed rows.

Note that this analysis loses a term of $O(\sqrt{\text{CC}(KW_f)})$ twice: once because it sets $\alpha = O(\sqrt{\text{CC}(KW_f)})$ (thus losing $O(\sqrt{\text{CC}(KW_f)})$ bits in the first stage), and once because it sets $\tau = O(\sqrt{\text{CC}(KW_f)})$ (thus losing $O(\sqrt{\text{CC}(KW_f)})$ in the second stage).

Our analysis

Our goal is to avoid the loss of the $O(\sqrt{\text{CC}(KW_f)})$ term in the lower bound, and therefore we set $\tau = O(1)$ and $\alpha = O(1)$. In order to make the analysis go through with this choice of parameters, we need a new idea.

Our first key idea is the following observation: the analysis of [1] works as if the parties transmit during the first stage $\text{CC}(KW_f)$ bits about X, Y , and another $\text{CC}(KW_f)$ bits about a, b . However, this cannot happen, since the total amount of communication in the first stage is less than $\text{CC}(KW_f)$. Hence, if the parties talked much about X, Y , then they can talk only a little about a, b , and vice versa.

For concreteness, let us denote by ℓ the number of bits that the parties talked about X, Y . Then, by Markov's inequality, the number of revealed rows is at most $\frac{\ell}{\tau}$. On the other hand, the parties talked at most $\text{CC}(KW_f) - \alpha - \ell$ bits about a and b (since they talked at most $\text{CC}(KW_f) - \alpha$ bits overall). Hence, the adversary can set $a_i = b_i$ for ℓ rows, which is more than the number of revealed rows as long as $\tau > 1$.

A further complication

The foregoing simple idea almost works. However, there is still another complication that we have not discussed: when the adversary fixes $a_i = b_i$ for a row, it may leak a bit of information to the parties, and in total, if there are k revealed rows then k bits may be leaked. This leakage causes us to lose k bits in the lower bound of the second stage: in the worst case, after the leakage the parties know $\tau + k$ bits about some unrevealed row, and therefore, in the second stage, they may solve the universal relation on this row using only $n - \tau - k$ bits. In such case, we will lose a term of k in the lower bound.

In the work of [1], this loss was not an issue, since they had $k = O(\sqrt{\text{CC}(KW_f)})$ revealed rows, and they were losing a term of $O(\sqrt{\text{CC}(KW_f)})$ anyhow. However, in our argument above the number of revealed rows is about $\frac{\ell}{\tau}$, which could be almost as large as $\text{CC}(KW_f)$. This loss is therefore unacceptable.

⁶ Note that the original paper of [1] proves the result for $U_n \diamond U_n$ rather than $KW_{f \diamond U_n}$, and therefore all the occurrences of $\text{CC}(KW_f)$ in the following description were originally equal to n .

An iterative adversary

Our second key idea is to modify the adversary in order to deal with the above complication as follows: After the adversary fixed $a_i = b_i$ for each of the revealed rows, it checks if new revealed rows were created in the process. In other words, the adversary checks if the leakage of information caused the parties to know more than τ bits on some of the previously unrevealed rows. If this is the case, the adversary fixes $a_i = b_i$ for each of these new revealed rows as well. The adversary now repeats the process until there are no more revealed rows – i.e., until the parties know at most τ bits of information on each of the unrevealed rows.

Obviously, if such an adversary can be implemented, then it avoids losing the term of k in the second stage. However, we need to show that such an adversary can indeed be implemented, i.e., that the adversary is allowed to fix $a_i = b_i$ for all the revealed rows encountered throughout all the iterations.

To this end, we bound the total number of revealed rows that are encountered in the process using the following potential argument: Whenever the adversary fixes $a_i = b_i$ for a revealed row X_i, Y_i , the parties may gain one bit of information that is leaked about X, Y . However, the parties also lose τ bits of information about X, Y , in the sense that the τ bits of information they knew about X_i, Y_i become useless after the fixing. In total, whenever the adversary fixes $a_i = b_i$ for a revealed row, the parties lose $(\tau - 1)$ bits of information about X, Y . Since at the beginning of the iterative process the parties knew at most ℓ bits of information about X, Y , the total number of revealed rows cannot exceed $\ell/(\tau - 1)$.

Finally, recall the adversary is allowed to fix $a_i = b_i$ for ℓ rows, and observe that this is more than $\ell/(\tau - 1)$ as long as $\tau \geq 2$. Therefore, we can implement the above adversary. This concludes the analysis.

1.4 Organization of the paper

In Section 2 we review the required preliminaries. In Section 3 we discuss the general adversarial strategy and prove a lower bound for the composition of a function with the universal relation.

2 Preliminaries

We use $[n]$ to denote the set $\{1, \dots, n\}$. We denote the set of $m \times n$ binary matrices by $\{0, 1\}^{m \times n}$. For every binary $m \times n$ matrix X , we denote by $X_i \in \{0, 1\}^n$ the i -th row of X .

2.1 Formulas

In this paper we consider (*de-Morgan*) *formulas* of whose internal gates are 2 bit, AND (\wedge) or OR (\vee) gates.

► **Definition 10.** The *formula complexity* of a Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, denoted $L(f)$, is the size of the smallest formula⁷ that computes f . The *depth complexity* of f , denoted $D(f)$, is the smallest depth of a formula that computes f .

⁷ Note that we define here the depth complexity of a function as the depth of a *formula* that computes f , while in the introduction we defined it as the depth of a *circuit with fan-in 2* that computes f . However, for our purposes, this distinction does not matter, since every circuit with fan-in 2 can be converted into a formula with the same depth.

The following definition generalizes the above definitions from functions to promise problems, which will be useful when we discuss Karchmer-Wigderson relations.

► **Definition 11.** 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 $\phi(\mathcal{X}) = 0$ and $\phi(\mathcal{Y}) = 1$. 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} .

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 [10] 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 communicate and find $z \in \mathcal{Z}$ such that $(x, y, z) \in R$. At each round, one of the players sends a bit that depends on her/his input and on the previous messages, until they find z . The *communication complexity of R* is the minimal number of bits that is transmitted by a protocol that solves R .

We now define a notion of protocol size that is analogous to the notion of formula size.

► **Definition 12.** We define the *size* of a protocol Π to be its number of leaves. Note that this is also the number of distinct transcripts of the protocol. We define the *protocol size* of a relation R , denoted $L(R)$, as the size of the smallest protocol that solves it (this is also known as the *protocol partition number* of R).

2.3 Karchmer-Wigderson relations

In this section, we define KW relations formally, and state the correspondence between KW relations and formulas. We start by defining KW relations for general rectangles, and then specialize the definition to functions.

► **Definition 13.** Let $\mathcal{X}, \mathcal{Y} \subseteq \{0, 1\}^n$ be two disjoint sets. The *KW relation $KW_{\mathcal{X} \times \mathcal{Y}} \subseteq \mathcal{X} \times \mathcal{Y} \times [n]$* is defined by $KW_{\mathcal{X} \times \mathcal{Y}} = \{(x, y, i) : x_i \neq y_i\}$. Intuitively, $KW_{\mathcal{X} \times \mathcal{Y}}$ corresponds to the communication problem in which Alice gets $x \in \mathcal{X}$, Bob gets $y \in \mathcal{Y}$, and they would like to find a coordinate $i \in [n]$ such that $x_i \neq y_i$ (note that $x \neq y$ since $\mathcal{X} \cap \mathcal{Y} = \emptyset$).

► **Definition 14.** Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ be a non-constant function. The *KW relation of f* , denoted KW_f , is defined by $KW_f = KW_{f^{-1}(0) \times f^{-1}(1)}$.

We now state the connection between Boolean functions and Karchmer-Wigderson relations.

► **Theorem 15** ([6]⁸). *For every two disjoint sets $\mathcal{X}, \mathcal{Y} \subseteq \{0, 1\}^n$ it holds that $D(\mathcal{X} \times \mathcal{Y}) = CC(KW_{\mathcal{X} \times \mathcal{Y}})$, 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 $D(f) = CC(KW_f)$, and $L(f) = L(KW_f)$.*

We note that for a pair of disjoint sets $\mathcal{X}, \mathcal{Y} \subseteq \{0, 1\}^n$ where either \mathcal{X} or \mathcal{Y} or both is the empty set, the leaf complexity $L(\mathcal{X} \times \mathcal{Y})$ is defined to 0.

Throughout this work, we will rely extensively on the following *subadditivity property* of protocol size and formula complexity: for every $\mathcal{X}, \mathcal{Y} \subseteq \{0, 1\}^n$ such that $\mathcal{X} = \mathcal{X}_0 \cup \mathcal{X}_1$ and $\mathcal{Y} = \mathcal{Y}_0 \cup \mathcal{Y}_1$, it holds that

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

⁸ The connection for formula complexity is implicit in [6], and is discussed explicitly in [8, 4, 2].

2.4 Subadditive measures on protocol trees

We rely upon a property of subadditive measures defined on binary trees, like the protocol tree of a deterministic protocol Π . To this end, we define subadditive measures on binary trees as follows:

► **Definition 16.** Given a rooted binary tree $T = (V, E)$, we say that $\phi : V \rightarrow \mathbb{N}$ is a *subadditive measure on T* if for every vertex u with children v and w in T it holds that $\phi(u) \leq \phi(v) + \phi(w)$.

We state without proof the following easy lemma about such measures.

► **Lemma 17.** *Let $T = (V, E)$ be a rooted binary tree with root r and depth d , and let ϕ be a subadditive measure on T . Then there exists at least one leaf l of the tree for which,*

$$\phi(l) \geq \left\lfloor \frac{\phi(r)}{2^d} \right\rfloor$$

2.5 Predictability and Average degree

In this paper we consider matrices of order $m \times n$ and partial matrices formed by a subset of rows, say σ , of the original set of rows, $[m]$. Let $S \subseteq (\{0, 1\}^n)^\sigma$ denote a set of binary matrices whose rows are indexed by σ . We measure the information revealed on a typical row of S , conditioned on other rows, using the notion of predictability as in [1] but using the presentation in Raz-McKenzie [7].

► **Definition 18 (Projections).** Let $S \subseteq (\{0, 1\}^n)^\rho$. Given a matrix $X \in S$ and a subset $\sigma \subseteq \rho$, we denote by X_σ the projection of X into rows indexed by σ . We extend the definition to a set of matrices S : $S_\sigma = \{X_\sigma \mid X \in S\}$

For a subset σ of $[m]$, we denote by $-\sigma$ its complement set in $[m]$, i.e., $[m] \setminus \sigma$. When σ is a singleton set say $\{i\}$, we denote it with i and $-\sigma$ with $-i$. Similarly, given a bit string $a \in \{0, 1\}^m$ and a subset $\sigma \subseteq [m]$, we denote by a_σ the projection of a to coordinates in σ .

As mentioned earlier, we measure information using the layered graph corresponding to the set of matrices obtained as follows. We interpret a subset of matrices $S \subseteq \{0, 1\}^{m \times n}$ as a $[m]$ bipartite graphs $G_S^i(U, V, E)$ for each $i \in [m]$, defined as follows. The left partition U is the set S_{-i} , the projection of S onto $[m] \setminus \{i\}$. The right partition V is the set S_i , the projection of S onto row i . The edge set E is defined as the set $\{(X, Y) \mid \exists Z \in S, X = Z_{-i}, Y = Z_i\}$. For any $X \in S_{-i}$, we denote by $\deg_i(X, S)$, the degree of the left node X in the graph G_S^i .

We use both average-degree and min-degree as information measures for measuring conditional information. More specifically, we use them to measure information on specific row, conditioned on the information about the other rows.

► **Definition 19 (Average degree of row).** Let S be a subset of $\{0, 1\}^{m \times n}$. We define the average degree of the i th row in the set S to be,

$$\text{avgdeg}_i(S) = \frac{\sum_{X \in S_{-i}} \deg_i(X, S)}{|S_{-i}|} = \frac{|S|}{|S_{-i}|}.$$

Similarly we define min-degree as.

► **Definition 20 (Minimum degree of row).** Let S be a subset of $\{0, 1\}^{m \times n}$. We define the average degree of the i th row in the set S to be,

$$\text{mindeg}_i(S) = \min_{X \in S_{-i}} \deg_i(X, S).$$

48:12 Improved Composition Theorems for Functions and Relations

Collision probability denoted by $CP_i(S)$ is defined to be the probability that two vectors chosen uniformly at random (with replacement) from S has the same i th element, i.e., $CP_i(S) = \Pr_{X, X' \in S} [X_i = X'_i]$. The following is known about $CP_i(S)$,

► **Lemma 21** (Lemma 4.4 from [1]). $CP_i(S) \leq \frac{1}{\text{avgdeg}_i(S)}$.

In fact [1] proves something stronger. It also proves that for an individual $X \in S$, the collision probability with a randomly sampled X' is bounded by inverse of the average degree. We state it as the following lemma.

► **Lemma 22** (follows from proof of Lemma 4.4 of [1]). *For any $X \in S$, for an X' chosen uniformly at random (with replacement) from S has,*

$$\Pr_{X' \in S} [X_i = X'_i] \leq \frac{1}{\text{avgdeg}_i(S)}.$$

The next lemma describes how average degree changes when we remove a set of matrices from S .

► **Lemma 23** (Lemma 4.5 from [1]). *Let $S' \subseteq S$,*

$$\text{avgdeg}_i(S') \geq \frac{|S'|}{|S|} \text{avgdeg}_i(S).$$

3 Depth lower bound

In this section we prove our main result. That is, we prove the following theorem, restated from the introduction.

► **Theorem 8.** *For any $m, n \in \mathbb{N}$, and any non-constant Boolean function $f : \{0, 1\}^m \rightarrow \{0, 1\}$ it holds that,*

$$CC(KW_{f \circ U_n}) \geq \log L(f) + n - 2 \log m - 10$$

where $L(f)$ is the formula complexity of the smallest formula computing f .

Let Π be a deterministic protocol solving the totalized variant of $KW_{f \circ U_n}$ whose communication complexity is less than $\log L(f) + n - 2 \log m + 10$ bits.

Recall that in the communication problem $KW_{f \circ U_n}$, Alice gets (X, a) and Bob gets (Y, b) where X, Y are $m \times n$ matrices and a, b are m bit column vectors. We refer to X, Y as the matrix part, and a, b as the column vector part of players input, respectively. Recall that in the totalized variant of $KW_{f \circ U_n}$, players can get inputs where the promises are not kept, and the protocol solving the totalized variant is supposed to reject such inputs.

We defer the technical details of the argument to the appendix. The appendix section is organized as follows : In Appendix A we show the existence of a partial transcript on which the players have not learned too much information about their inputs. In Section B, we formalize the intuition that the information transmitted by an average partial transcript is divided between the matrix and the column vector part of the players input. In Section B.1, we do a cleanup to ensure that the players are forced to solve the second stage on unrevealed rows, i.e., rows in which the parties have learned very little information in the first stage.

References

- 1 J. Edmonds, R. Impagliazzo, S. Rudich, and J. Sgall. Communication complexity towards lower bounds on circuit depth. *computational complexity*, 10(3):210–246, Dec 2001. doi:10.1007/s00037-001-8195-x.
- 2 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.
- 3 Johan Håstad and Avi Wigderson. Composition of the universal relation. In *Advances In Computational Complexity Theory*, volume 13 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 119–134. DIMACS/AMS, 1990.
- 4 Mauricio Karchmer, Eyal Kushilevitz, and Noam Nisan. Fractional covers and communication complexity. *SIAM J. Discrete Math.*, 8(1):76–92, 1995.
- 5 Mauricio Karchmer, Ran Raz, and Avi Wigderson. Super-logarithmic depth lower bounds via the direct sum in communication complexity. *Computational Complexity*, 5(3/4):191–204, 1995.
- 6 Mauricio Karchmer and Avi Wigderson. Monotone circuits for connectivity require super-logarithmic depth. *SIAM J. Discrete Math.*, 3(2):255–265, 1990.
- 7 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.
- 8 Alexander A. Razborov. Applications of matrix methods to the theory of lower bounds in computational complexity. *Combinatorica*, 10(1):81–93, 1990.
- 9 G. Tardos and U. Zwick. The communication complexity of the universal relation. In *Proceedings of Computational Complexity. Twelfth Annual IEEE Conference*, pages 247–259, Jun 1997. doi:10.1109/CCC.1997.612320.
- 10 Andrew Chi-Chih Yao. Some complexity questions related to distributive computing (preliminary report). In *STOC*, pages 209–213, 1979.

A First Stage Communication

We let the players speak for $t_1 = \log(L(f)) - \alpha$ bits in the first stage, where α is a slack variable to be fixed later. We define a formal measure that measures the progress that the players make throughout the protocol towards solving $KW_{f \circ U_n}$. We then show that this measure does not increase by much during the first stage.

We first define the set of matrices which are “alive” with respect to a partial transcript π of Π . We say that a matrix X is *alive* with respect to partial transcript π , if there exists a pair $(a, b) \in f^{-1}(0) \times f^{-1}(1)$ such that the input $((X, a), (X, b))$ is consistent with the partial transcript π . We define A_π^X as the set of all column vectors $a \in f^{-1}(0)$ for which (X, a) is an input for Alice consistent the partial transcript π . Similarly define B_π^X as the set of all column vectors $b \in f^{-1}(1)$ for which (X, b) is an input for Bob consistent with the partial transcript π . We let the set S_π be the set of alive matrices with respect partial transcript π . To quantify the total progress the players have made given a partial transcript π , we use the following measure.

$$\phi(\pi) = \sum_{X \in S_\pi} L(A_\pi^X \times B_\pi^X)$$

It is easy to note that the measure $\phi(\pi)$ defined above is a subadditive measure on the protocol tree of Π . By applying Lemma 17 on T' we get that, there exists a leaf $l \in T'$

and a corresponding partial transcript π' of Π , with $\phi(l) \geq \frac{1}{2^{\log(L(f))-\alpha}} \phi(\lambda)$. At the root $\phi(\lambda) = 2^{mn} L(f)$. Thus we get the following proposition,

- **Proposition 24.** *There exists a partial transcript π such that,*
- $\phi(\pi') \geq \frac{1}{2^{t_1}} 2^{mn+\log(L(f))} = 2^{mn+\alpha}$.
 - *The length of π' is at most t_1 , and length of π' is strictly less than t_1 if and only if π' is a leaf of the protocol Π .*

B Cleanup after first stage

In this section we do a cleanup so that players are forced to solve U_n on rows where not too much information was revealed.

At this point, we fix the partial transcript π' and differentiate between the information the players have learned about the matrices in $S_{\pi'}$, and the progress players have made on solving KW_f problem associated with an average $X \in S_{\pi'}$. To this end, we describe two measures of information. We first measure the information about the matrix. This is done by measuring what fraction of matrices from $\{0, 1\}^{m \times n}$ are alive in $S_{\pi'}$. Throughout the argument we would like to discard some of the inputs which are consistent with the transcript, and measure the information with respect to remaining inputs. Hence, we also define these measures for arbitrary subsets of $\{0, 1\}^{m \times n}$, and their associated rectangles.

Let $S \subseteq \{0, 1\}^{m \times n}$ be such that every $X \in S$ has an associated rectangle $A_S^X \times B_S^X \subseteq f^{-1}(0) \times f^{-1}(1)$. The subadditive measure ϕ defined earlier extends naturally to such arbitrary as, $\phi(S) = \sum_{X \in S} L(A_S^X \times B_S^X)$. When $S = S_{\pi'}$, for every $X \in S$, $A_S^X \times B_S^X$ is defined to be $A_{\pi'}^X \times B_{\pi'}^X$. The information on the matrix part is measured as:

$$T_X^S = \log \left(\frac{2^{mn}}{|S|} \right) = mn - \log(|S|).$$

To measure the progress players have made on solving KW_f problem conditioned on an average X , we define:

$$T_{a|X}^S = \log \left(\frac{L(f)}{\frac{1}{|S|} \sum_{X \in S} L(A_S^X \times B_S^X)} \right) = \log(L(f)) - \log \left(\frac{1}{|S|} \phi(S) \right).$$

Observe that,

$$T_X^{\pi'} + T_{a|X}^{\pi'} = mn + \log(L(f)) - \log(\phi(\pi')). \quad (1)$$

Since we chose π' such that $\phi(\pi') \geq 2^{mn} 2^\alpha$, we get that,

$$T_X^{\pi'} + T_{a|X}^{\pi'} = mn + \log(L(f)) - mn - \alpha = \log(L(f)) - \alpha.$$

Intuitively, this means that the total amount of information the players have learned in the first stage is bounded by the number of bits communicated in the first stage.

B.1 Regularization

In order to facilitate the analysis, we would like if the matrices $X \in S_{\pi'}$ has roughly the same formula complexity for its associated rectangle $A_{\pi'}^X \times B_{\pi'}^X$ as an average X . In particular, we need that for every $X \in S_{\pi'}$, it holds that $L(A_{\pi'}^X \times B_{\pi'}^X) \approx 2^{\log(L(f)) - T_{a|X}^{\pi'}}$. However, we do not have this property for every matrix X in $S_{\pi'}$. We thus construct a subset $S' \subseteq S_{\pi'}$, such that for every $X \in S'$, $L(A_{S'}^X \times B_{S'}^X) \approx 2^{\log(L(f)) - T_{a|X}^{S'}}$ while maintaining that, $T_X^{S'} + T_{a|X}^{S'} \leq \log(L(f)) - \alpha + O(\log m)$.

We defer the details of the construction of S' to the full version. We denote the set S' by S^{reg} , signifying that this is a regularized set. From now on we work with the set S^{reg} .

B.2 Fixing revealed rows

In this section we describe an iterative process to classify the rows into revealed and unrevealed rows based on a threshold τ , and fix the corresponding bits in the column vector to be same for both players. In this process whenever we find a row whose average degree is smaller than $2^{n-\tau}$, we ensure that $a_i = b_i$ for this row. This specific row ceases to matter as we have maintained the promise irrespective of the value of this row for either players. We call such a row, identified by the iterative process and for which we fix $a_i = b_i$, an *inactive* row. Any row which is not an inactive row is called an *active* row. At any point during the iterative process, we would like to work with a set of matrices that consist only of the active rows. When we set $a_i = b_i$ for a new row i , we remove the row i from the set of active rows to the set of inactive rows. Now, formally, since the players expect to receive $m \times n$ matrices, we maintain a bijection between the set of matrices over active rows and the set of legal $m \times n$ matrices. Finally, when we remove a row from the set of active rows, we update the bijection. This is done by fixing the row as a function of the other rows. When the iterative process terminates, the set of active rows correspond to the unrevealed rows and the set of inactive rows correspond to the revealed rows with respect to threshold τ .

We fix a threshold τ on information learned about a row to classify the rows into “revealed” and “unrevealed rows”. Under our average-degree measure, we say that a row is revealed with respect to a set of matrices $S \subseteq \{0, 1\}^{m \times n}$, if $\text{avgdeg}_i(S) \leq 2^{n-\tau}$. Initially we look for revealed rows with respect to S^{reg} . We fix a revealed row to be a function of the remaining rows, and also fix the corresponding bit in the column vector. Fixing of a revealed row potentially reduces the set of alive matrices by eliminating the matrices whose extension into the revealed row is not according to the fixing chosen. The adversary proceeds by looking for revealed rows with respect to the current set and repeats the process of fixing such a row and corresponding column vector until there are no revealed rows with respect to the current set of alive matrices.

We have to make sure that the above process terminates before reducing the formula complexity of associated rectangles to zero. The following intuitive argument illustrates why the adversary can guarantee such a convergence. Since the players learned $T_X^{S^{\text{reg}}}$ amount of information about the entire matrix, when the adversary fixes a revealed row, the adversary accounts for τ amount of information from the total $T_X^{S^{\text{reg}}}$. However, fixing the corresponding bit in the column vector, could potentially reveal a bit of information about the remaining rows. Thus during one such operation, the adversary accounts for at least $\tau - 1$ bits of information from $T_X^{S^{\text{reg}}}$. Hence number of such operations, and consequently, the number of fixed rows can be at most $\frac{T_X^{S^{\text{reg}}}}{\tau-1}$. Since this is less than total number of rows, not all rows are fixed. We also need to ensure that at the end of all fixings every matrix has non-zero leaf complexity for its associated rectangle. We ensure that when fixing a bit corresponding to a revealed row, the formula complexity of the rectangle associated with a matrix reduces at most by a fraction of $\frac{1}{4}$ th. Since every matrix X in S^{reg} had $2^{\log(L(f)) - T_{a|X}^{S^{\text{reg}}} - 1}$ formula complexity for the associated rectangle, we can ensure that after at most $\frac{T_X^{S^{\text{reg}}}}{\tau-1}$ fixings, the rectangle associated with a matrix has non-zero formula complexity.

The following lemma formalizes the effect of a single iteration of our iterative adversary.

► **Lemma 25.** *Let S' be a subset of $(\{0, 1\}^n)^{m-|\sigma|}$, with an extension function $E_\sigma : S' \rightarrow \{0, 1\}^{m \times n}$, along with rectangles $A_{S'}^X \times B_{S'}^X \subseteq f^{-1}(0) \times f^{-1}(1)$ for every matrix $X \in E_\sigma(S')$. Let $i \in [m] \setminus \sigma$ be a row such that $\text{avgdeg}_i(S') \leq 2^{n-\tau}$.*

Then there exists a set $S'' \subseteq S'_{-i}$, an extension function $E_{\sigma \cup \{i\}} : S'' \rightarrow \{0, 1\}^{m \times n}$, along with rectangles $A_{S''}^X \times B_{S''}^X$, for every matrix $X \in E_{\sigma \cup \{i\}}(S'')$ such that,

- $\frac{|S''|}{2^{(m-(|\sigma|+1))n}} \geq \frac{|S'|}{2^{(m-|\sigma|)n}} 2^{\tau-1}$. That is the amount of information in the matrix part, when restricted to rows $[m] \setminus (\sigma \cup \{i\})$, is decreased by $\tau - 1$ bits (from the amount of information in the matrix part, when restricted to rows $[m] \setminus \sigma$).
- For any $X, Y \in E_{\sigma \cup \{i\}}(S'')$ and for any $a \in A_{S''}^X$, and for any $b \in B_{S''}^Y$, it holds that $a_i = b_i$.
- For any $X \in E_{\sigma \cup \{i\}}(S'')$, $L(A_{S''}^X \times B_{S''}^X) \geq \frac{1}{4} L(A_{S'}^X \times B_{S'}^X)$ as long as $L(A_{S'}^X \times B_{S'}^X) \geq 4$.

We defer the proof of the Lemma to the full version of the paper and prove Theorem 26 which shows how the adversary can iteratively fix revealed rows using Lemma 25 and get a set of matrices with the desired properties for the second stage.

► **Theorem 26.** *Let $S^{\text{reg}} \subseteq \{0, 1\}^{m \times n}$ be the set from Section B.1. Then for any $\tau \in \mathbb{N}$, $n > \tau > 3$, there exists a set of indices $\sigma \subseteq [m]$, a subset $S^{\text{clean}} \subseteq S_{-\sigma}^{\text{reg}}$, an extension function $E_\sigma : S^{\text{clean}} \rightarrow \{0, 1\}^{m \times n}$ and associated rectangles $R_{S^{\text{clean}}}^X = A_{S^{\text{clean}}}^X \times B_{S^{\text{clean}}}^X$ for every $X \in E_\sigma(S^{\text{clean}})$, such that,*

- The set σ of revealed rows is such that $|\sigma| < m$.
- For any row outside σ players have learned at most τ bits of information. More formally for any $i \in [m] \setminus \sigma$, $\text{avgdeg}_i(S^{\text{clean}}) > 2^{n-\tau}$.
- For any $X \in E_\sigma(S^{\text{clean}})$, $A_{S^{\text{clean}}}^X \times B_{S^{\text{clean}}}^X$ is non-empty.
- For any $X, Y \in E_\sigma(S^{\text{clean}})$ not necessarily distinct, for any $a \in A_{S^{\text{clean}}}^X$, and any $b \in B_{S^{\text{clean}}}^Y$, we have that $a_\sigma = b_\sigma$.

Proof. The adversary uses Lemma 25 to fix revealed rows until there is none. Initially the adversary sets $\sigma \leftarrow \emptyset$, and $S' \leftarrow S^{\text{reg}}$. If there is any row $i \in [m] \setminus \sigma$, for which $\text{avgdeg}_i(S') \leq 2^{n-\tau}$, the adversary invokes Lemma 25 to obtain a set S'' of matrices, an extension function $E_{\sigma \cup \{i\}} : S'' \rightarrow \{0, 1\}^{m \times n}$, and an updated set of fixed rows $\sigma \leftarrow \sigma \cup \{i\}$. Adversary sets $S' \leftarrow S''$, and checks again for a row $i \in [m] \setminus \sigma$, for which $\text{avgdeg}_i(S') \leq 2^{n-\tau}$. If no such row exists, the adversary stops. Otherwise the adversary invokes Lemma 25 on S' and repeats the procedure. Note that when the iterative process stops, the resulting set clearly satisfies the requirement.

We claim that the procedure terminates after at most $\frac{T_X^{\text{Sreg}}}{\tau-1}$ steps. Equivalently, $|\sigma| \leq \frac{T_X^{\text{Sreg}}}{\tau-1}$, as each invocation of the Lemma 25 increases $|\sigma|$ by 1. We use the following claim, whose proof is deferred to the full version, to show that $|\sigma| \leq \frac{T_X^{\text{Sreg}}}{\tau-1}$.

► **Claim 27.** *Between every invocation of Lemma 25, the following invariant holds,*

$$\frac{|S'|}{2^{mn-|\sigma|n}} \geq 2^{-(T_X^{\text{Sreg}} - |\sigma|(\tau-1))}$$

Note that $S' \subseteq \{0, 1\}^{(m-|\sigma|) \times n}$ and thus $\frac{|S'|}{2^{mn-|\sigma|n}}$ is always upper bounded by 1. This along with the invariant implies that, $1 \geq 2^{-(T_X^{\text{Sreg}} - |\sigma|(\tau-1))}$. That is $T_X^{\text{Sreg}} \geq |\sigma|(\tau-1)$. Since $\tau > 1$, this proves the claim that $|\sigma| \leq \frac{T_X^{\text{Sreg}}}{\tau-1}$. Since $T_X^{\text{Sreg}} + T_{a|X}^{\text{Sreg}} \leq \log(L(f)) - 4$, and since $\log(L(f)) \leq m$, we get that $|\sigma| < m$.

We claim that S' obtained at the end of such iterative fixing has all the properties claimed by the theorem for S^{clean} . That is we set $S^{\text{clean}} \leftarrow S'$.

Note that Lemma 25 ensures that the rows indexed by σ are fixed as a function of other rows for matrices in S^{clean} . The adversary stops invoking Lemma 25 on S' only when every row $i \in [m] \setminus \sigma$ has $\text{avgdeg}_i(S') > 2^{n-\tau}$. Hence every row $i \in [m] \setminus \sigma$ has $\text{avgdeg}_i(S^{\text{clean}}) > 2^{n-\tau}$.

Suppose at any point during the iterative fixing of the rows, every X that we consider, has formula complexity at least 4. Then, each invocation of Lemma 25 reduces the formula complexity of the associated rectangle by at most a quarter. Since we started with formula complexity $2^{\log(L(f)) - T_{a|X}^{S^{\text{reg}}} - 1}$ for every matrix X , after $|\sigma| \leq \frac{T_X^{S^{\text{reg}}}}{\tau - 1}$ many invocations, any $X \in E_\sigma(S^{\text{clean}})$, has formula complexity $2^{-\left(2 \cdot \frac{T_X^{S^{\text{reg}}}}{\tau - 1}\right)}$ fraction of the formula complexity in S^{reg} . That is,

$$\begin{aligned} L(A_{S^{\text{clean}}}^X \times B_{S^{\text{clean}}}^X) &\geq 2^{-\left(2 \cdot \frac{T_X^{S^{\text{reg}}}}{\tau - 1}\right)} L(A_{S^{\text{reg}}}^X \times B_{S^{\text{reg}}}^X) \\ &\geq 2^{\log(L(f)) - T_{a|X}^{S^{\text{reg}}} - 1 - T_X^{S^{\text{reg}}}}. \end{aligned}$$

Note that this is at least 2^3 as the statement of the theorem assumes that $T_X^{S^{\text{reg}}} + T_{a|X}^{S^{\text{reg}}} \leq \log(L(f)) - 4$ and $\tau > 3$. Thus in all of $|\sigma|$ invocations of the Lemma 25, all the rectangles associated with the all the matrices have formula complexity at least 8. This also ensures that for any $X \in E_\sigma(S^{\text{clean}})$, $A_{S^{\text{clean}}}^X \times B_{S^{\text{clean}}}^X$ is non-empty at the end as $L(A_{S^{\text{clean}}}^X \times B_{S^{\text{clean}}}^X) \geq 8$. This concludes the proof the theorem. \blacktriangleleft

The fact that for any $X \in E_\sigma(S^{\text{clean}})$, $L(A_{S^{\text{clean}}}^X \times B_{S^{\text{clean}}}^X) > 0$ implies that, there is at least one (a, b) pair in $A_{S^{\text{clean}}}^X \times B_{S^{\text{clean}}}^X$. From now on, for any $X \in E_\sigma(S^{\text{clean}})$, we denote by (a^X, b^X) an arbitrary but fixed (a, b) pair in $A_{S^{\text{clean}}}^X \times B_{S^{\text{clean}}}^X$.

C Second Stage Communication

In this section we describe how to handle the second stage of communication and complete the lower bound.

In the second stage, we let the players communicate at most $t_2 = n - \log m - 5$ bits of communication. At the end of the second stage the adversary would like to ensure that on any unrevealed row from the first stage, the players have not learned more than $\tau + 2 + t_2 = n - \log m$ bits of information. Intuitively this should be possible, because after the cleanup at the end of first stage, the players knew at most $\tau + 2$ bits of information about any unrevealed row in S'' . To this end, we use Lemma 23 to ensure that with the additional t_2 bits of communication, the players learn at most t_2 bits of information on any row.

We consider the sub-tree $T_{\pi'}''$ of the protocol tree Π rooted at the node π' (the partial transcript from first stage). We would like to use Lemma 23 on the subtree $T_{\pi'}''$ to find a leaf of the protocol that is supported by at least 2^{-t_2} fraction of the matrices in S^{clean} . To this end, we need to prove that the depth of the tree is at most t_2 . Recall that communication complexity of Π was less than $t_1 + t_2$, where t_1 is an upper bound on the number of bits communicated in the first stage. If π' itself is a leaf of the protocol Π , the tree $T_{\pi'}''$ is just a leaf and the claim is true. If π' is not a leaf of the protocol Π , the way we chose π' ensures that depth of π' in Π is equal to t_1 . Hence there cannot be a leaf of $T_{\pi'}''$, which is at depth more than t_2 , as the corresponding leaf in Π would be at depth more than $t_1 + t_2$, which is a contradiction to the fact $\text{CC}(\Pi) \leq t_1 + t_2$.

We define the following subadditive measure on the nodes of $T_{\pi'}''$. For a node $v \in T_{\pi'}''$, corresponding to a partial transcript π'' of the second stage, define $\phi(\pi'') = |P_{\pi''}|$ where $P_{\pi''} = S_{\pi''} \cap E_\sigma(S^{\text{clean}})$. That is, $P_{\pi''}$ is the set of matrices consistent with π'' and are from the set $E_\sigma(S^{\text{clean}})$. It is easy to verify that this is indeed a subadditive measure on the protocol sub-tree $T_{\pi'}''$. Note that at the root of $T_{\pi'}''$, $\phi(\pi') = |E_\sigma(S^{\text{clean}})|$. Applying Lemma 17, we get that there is a leaf of $T_{\pi'}''$ corresponding to a transcript π'' of the second stage

for which $\phi(\pi'') \geq \frac{1}{2^{t_2}} |S^{\text{clean}}|$. Thus at the end of the second stage, we have a partial transcript π'' such that, $|P_{\pi''}| \geq 2^{-t_2} |S^{\text{clean}}|$. By Lemma 23, we get that for any $i \in [m]$, $\text{avgdeg}_i(P_{\pi'}) \geq 2^{-t_2} \text{avgdeg}_i(E_\sigma(S^{\text{clean}}))$.

Hence at the end of the second stage, we have a full transcript and a set of inputs satisfying the following conditions.

► **Lemma 28.** *Let Π be any deterministic protocol solving $f \diamond U_n$. Let $t = t_1 + t_2$ where $t_1 = \log(L(f)) - \log m - 5$ and $t_2 = n - \log m - 5$. Then there are:*

- *A full transcript π of length at most t of Π .*
 - *A subset $\sigma \subseteq [m]$ of indices where $|\sigma| < m$.*
 - *A set $T \subseteq \{0, 1\}^{(m-|\sigma|) \times n}$.*
 - *An extension function $E_\sigma : T \rightarrow \{0, 1\}^{m \times n}$.*
 - *For any $X \in E_\sigma(T)$, a pair of column vectors a^X, b^X .*
- satisfying the following,*
- *For any $X, Y \in E_\sigma(T)$, $(X, a^X) \in S_\pi^A$ and $(Y, b^Y) \in S_\pi^B$.*
 - *For any $X \in E_\sigma(T)$, $f(a^X) = 0$ and $f(b^X) = 1$.*
 - *For any $X, Y \in E_\sigma(T)$, for any $i \in \sigma$, $a_i^X = b_i^Y$.*
 - *For any row outside σ players have learned at most $t_2 + \tau$ bits of information. More formally for any $i \in [m] \setminus \sigma$, it holds that $\text{avgdeg}_i(T) \geq 2^{n-\tau-2-t_2} = 2^{2+\log m}$.*

C.1 Choosing the inputs

We show that the protocol Π must err on π , by showing that it must reject some inputs as well as accept some other inputs.

Since there is at least one input $((X, a), (X, b))$ consistent with π at the end of second stage, and since π is a leaf of the protocol Π , the transcript π has to reject.

We show that Π cannot reject on π by showing the existence of an input $((X, a^X), (Y, b^Y))$ consistent with π which satisfies all the promises. That is, for any $i \in [m] \setminus \sigma$, if $a_i^X \neq b_i^Y$ then $X_i \neq Y_i$. Recall that for $i \in \sigma$, we have already ensured that $a_i^X = b_i^Y$. We show that there exists inputs where for all $i \in [m] \setminus \sigma$, $X_i \neq Y_i$ using the probabilistic method. This would establish that, for any $i \in [m]$, if $a_i^X \neq b_i^Y$ then $X_i \neq Y_i$. By Lemma 21 we know that for two matrices U, V chosen uniformly at random (with replacement) from T , the probability that $U_i = V_i$ is at most $(\text{avgdeg}_i(T))^{-1}$. We set $X = E_\sigma(U)$ and $Y = E_\sigma(V)$. Thus for any $i \in [m] \setminus \sigma$, we have that $\Pr_{X, Y \sim E_\sigma(T)} [X_i = Y_i] \leq \frac{1}{2^{\log m + 2}} = \frac{1}{4m}$. By a union bound, the probability for two matrices U, V chosen uniformly at random from T , the probability that all the rows indexed by $[m] \setminus \sigma$ are different is at least $1 - \frac{m-|\sigma|}{4m}$. For the rows indexed by σ , we already have that $a_i^X = b_i^Y$ for any $X, Y \in E_\sigma(T)$. Thus with probability $\frac{1}{4}$, the matrices X, Y satisfy all the promises when players are given the input $((X, a^X), (Y, b^Y))$. Hence the protocol cannot reject on π , establishing that Π errs on π .

Thus we get any protocol Π of length at most $t = \log(L(f)) + n - 2 \log m - 10$ cannot correctly solve totalized variant of $KW_{f \diamond U_n}$, by choosing $\tau = 3$. This proves the main theorem from the introduction.