The Direct Sum of Universal Relations

Or Meir*

September 27, 2017

Abstract

The universal relation is the communication problem in which Alice and Bob get as inputs two distinct strings, and they are required to find a coordinate on which the strings differ. The study of this problem is motivated by its connection to Karchmer-Wigderson relations [KW90], which are communication problems that are tightly related to circuit-depth lower bounds.

In this paper, we prove a direct sum theorem for the universal relation, namely, we prove that solving m independent instances of the universal relation is m times harder than solving a single instance. More specifically, it is known that the deterministic communication complexity of the universal relation is at least n. We prove that the deterministic communication complexity of solving m independent instances of the universal relation is at least $m \cdot (n - O(\log m))$.

1 Introduction

The direct-sum question is a classical question that asks whether performing a task on m independent inputs is m times harder than performing it on a single input. This natural question was studied in a variety of computational models (see, e.g., [Uhl74, Pau76, GF81, Sto86, Bsh89, Bsh98]), and the answer turns out to be positive in some models and negative in others. Karchmer, Raz, and Wigderson [KRW95] initiated the study of this question in the setting of communication complexity. One motivation was a connection that they observed between the direct-sum question for the *deterministic* communication complexity of *relations* and the circuit-depth complexity of functions.

Later works have made considerable progress in the study of direct sum for *randomized* communication complexity [BBCR10, GKR14] and for the deterministic communication complexity of *functions* [FKNN95]. However, there is only one[1] known result on the direct-sum question in the original setting of [KRW95] — *deterministic* protocols for *relations*: a direct-sum theorem for a relation that is connected to the set covering problem, which appears in the original paper of [KRW95]. In this work, we provide another example for such a direct-sum theorem, namely, for the universal relation.

---

[1]There are also a few examples of relations for which direct-sum theorems on the deterministic complexity follow trivially from the corresponding results on randomized complexity: This happens when the deterministic complexity of the relation is equal to its randomized complexity, and there is a direct-sum theorem for the randomized complexity. This is the case, for example, for the monotone Karchmer-Wigderson relations of the clique and matching functions [RW92].

However, we are interested in direct-sum theorems on deterministic complexity that are "non-trivial" in the sense that they do not follow directly from results on randomized complexity. This is the case for the universal relation, whose randomized complexity is much smaller than its deterministic complexity ($O(\log n)$ vs. n [RW89]).

The universal relation is the following communication problem: Alice and Bob get two *distinct* strings $x, y \in \{0,1\}^n$, and they are required to find a coordinate $j \in [n]$ such that $x_j \neq y_j$. This problem is a simplified version of Karchmer-Wigderson relations [KW90], which are communication problems that are tightly related to circuit-depth lower bounds. The universal relation was introduced by [KRW95] in the hope that a better understanding of the universal relation would lead to progress in the study of Karchmer-Wigderson relations, and hence to better circuit-depth lower bounds. It is known that the deterministic communication complexity of the universal relation is at least n. We prove the following result.

Theorem 1. *The deterministic communication complexity of solving m independent instances of the universal relation over n bits is at least $m \cdot (n - 2 \log(m) - 8)$.*

Our proof is based on the works of Edmonds et. al. [EIRS01] and Raz and McKenzie [RM97] on composition theorems. In this work, we show how their techniques can be applied to the setting of direct-sum theorems. We hope that our ideas will lead to more direct-sum results in the future.

Remark 2. Following an earlier version of this work [Mei17], Alexander Kozachinsky [Koz17] pointed out a simple rank argument that proves a much stronger lower bound of $m \cdot (n-1) - 1$. We describe this argument in Remark 8 below. We maintain our original (weaker and more complicated) proof in this version since we believe its ideas are valuable and might lead to new direct-sum theorems for other relations in the future.

Remark 3. There is a common misperception that a direct-sum theorem for Karchmer-Wigderson relations would imply better circuit-depth lower bounds. This is inaccurate: in order to prove better circuit-depth lower bounds, one needs to a prove a "composition theorem" rather than a direct-sum theorem. [KRW95] showed that a direct-sum theorem implies a composition theorem in the setting of *monotone* circuits, but this is not necessarily true for *non-monotone* circuits. For example, our direct-sum theorem for the universal relation does not imply a composition theorem for the universal relation (such a composition theorem was proved by [EIRS01, HW93]). See [KRW95] for more details.

The paper is organized as follows: In Section 2 we discuss the universal relation and its direct sum, as well as "totalized" versions of these problems which are important for our proof. We then prove Theorem 1 in Section 3.

Preliminaries: For $n \in \mathbb{N}$, we denote $[n] \stackrel{\text{def}}{=} \{1, \ldots n\}$. We denote by $\{0,1\}^{m \times n}$ the set of $m \times n$ binary matrices. Given a set $I \subseteq [m]$, we denote by $\{0,1\}^{I \times n}$ the set of $|I| \times n$ binary matrices whose rows are labeled by the indices in I. Given a subset of rows $I' \subseteq I$ and a matrix $Z \in \mathcal{Z}$, we denote by $Z|_{I'}$ the projection of Z to the rows in I', and we say that Z is an *extension of the matrix $Z|_{I'}$ (to $\mathcal{Z}$)*. Given a set of matrices $\mathcal{Z} \subseteq \{0,1\}^{m \times n}$ and a set of rows $I \subseteq [m]$, we denote by $\mathcal{Z}|_I$ the set of projections of matrices in $\mathcal{Z}$ to rows in I. We use the standard definitions of communication complexity — see the book of Kushilevitz and Nisan [KN97] for more details.

2 The universal relation, its direct sum, and their totalizations

As explained in the introduction, the *universal relation (on n bits)*, denoted U_n, is the following communication problem: Alice and Bob get two *distinct* strings $x, y \in \{0,1\}^n$, and they are required to find a coordinate $j \in [n]$ such that $x_j \neq y_j$. It is not hard to prove that the deterministic

communication complexity of this problem is at least n. On the other hand, it is interesting to note that its randomized communication complexity is at most $O(\log n)$ [RW89].

The direct sum of the universal relation consists of solving m independent instances of the problem. In order to streamline the presentation, it is convenient to represent the inputs to the direct sum by matrices. This leads to the following definition of the direct sum.

Definition 4. Let $m, n \in \mathbb{N}$. The *m-fold direct sum* of the universal relation on n bits, denoted $U_n^{\otimes m}$ is the communication problem in which Alice and Bob get matrices $X, Y \in \{0, 1\}^{m \times n}$ that differ on every row. They are required to output a tuple $(j_1, \ldots, j_m) \in [n]^m$, such that for every row $i \in [m]$ it holds that $X_{i,j_i} \neq Y_{i,j_i}$.

Håstad and Wigderson [HW93] observed that it is useful to consider a variant of the universal relation, which is a total relation rather than a promise problem: In the *totalized universal relation*, denoted U_n', Alice and Bob may be given identical strings as inputs, and in this case they should output a special "reject" symbol $\perp$. The totalized universal relation is often easier to work with than the non-totalized one. In particular, it is trivial to prove a lower bound of n on its deterministic communication complexity by a reduction from the equality function.

It is not hard to see that this modification does not decrease the complexity of the universal relation by more than two bits. To see it, suppose that there is a protocol Π that solves U_n. Then, there is a protocol Π' that solves U_n' using two more bits: given inputs x and y which may be equal, the players invoke the protocol Π on x, y. Suppose Π outputs a coordinate j. Now, the players check whether $x_j \neq y_j$ by exchanging two more bits. If they find that $x_j = y_j$, they reject, and otherwise they output j.

Similarly, it is useful to consider a "totalization" of the direct sum of the universal relation: Alice and Bob get two arbitrary matrices $X, Y \in \{0, 1\}^{m \times n}$. The parties should reject if X and Y agree on any single row, and otherwise they should output a tuple $(j_1, \ldots, j_m)$ as before.

Definition 5. Let $m, n \in \mathbb{N}$. The *totalized m-fold direct sum* of the universal relation on n bits, denoted $U_n^{\otimes m'}$, is the communication problem in which Alice and Bob get as inputs matrices $X, Y \in \{0, 1\}^{m \times n}$ and behave as follows:

- If X and Y differ on every row, then Alice and Bob behave as in the (non-totalized) direct sum $U_n^{\otimes m}$.

- Otherwise, Alice and Bob output the "reject" symbol $\perp$.

It is not hard to see that this modification does not decrease the complexity of the direct sum by more than $2m$ bits. To see it, suppose there is a protocol for $U_n^{\otimes m}$. Then, there is a protocol Π' that solves $U_n^{\otimes m'}$ using $2m$ more bits: Given X and Y, Alice and Bob would invoke the protocol for $U_n^{\otimes m}$, thus obtaining a tuple $(j_1, \ldots, j_m)$. Then, Alice and Bob would send X_{i,j_i} and Y_{i,j_i} for each row. If they disagree on all these values they would output $(j_1, \ldots, j_m)$, and otherwise they would reject. Hence, to prove Theorem 1, it suffices to prove that the communication complexity of the relation $U_n^{\otimes m'}$ is at least $m \cdot (n - 2 \log m - 6)$.

Remark 6. It is tempting to attempt to prove Theorem 1 in a different way. Consider the direct-sum problem for the *totalized* universal relation U_n': This is the problem where Alice and Bob get two arbitrary matrices $X, Y \in \{0, 1\}^{m \times n}$. The parties should output a tuple $(j_1, \ldots, j_m) \in ([n] \cup \{\perp\})^m$, such that if $j_i \in [n]$ then $X_{i,j_i} \neq Y_{i,j_i}$, and if $j_i = \perp$ then $X_i = Y_i$. The difference between this problem and $U_n^{\otimes m'}$ is that in the problem $U_n^{\otimes m'}$, if the matrices agree on any single row, the players reject the whole input and not just that row.

It is not hard to prove a lower bound of $m \cdot n$ on the direct sum of U_n' by reduction to the direct sum of the equality function. At first glance, it may seem as if we can use it to prove Theorem 1, by reducing the latter direct sum to $U_n^{\otimes m}$. The reduction would work as follows: Given X and Y, Alice and Bob would invoke the protocol for $U_n^{\otimes m}$, thus obtaining a tuple $(j_1, \ldots, j_m)$. Then, Alice and Bob would send X_{i,j_i} and Y_{i,j_i} for each row, and output j_i if $X_{i,j_i} \neq Y_{i,j_i}$ and $\perp$ otherwise.

Unfortunately, this reduction fails. The reason is that if X and Y agree on any row, there is no guarantee on the behavior of the protocol for $U_n^{\otimes m}$. In particular, this protocol may output coordinates j_i for which $X_{i,j_i} = Y_{i,j_i}$ even if $X_i \neq Y_i$. In such case, the above reduction would output $\perp$ for such rows even though it is not allowed to reject them.

Remark 7. We note that our lower bound on the totalized direct sum $U_n^{\otimes m'}$ also implies a lower bound for a related natural communication problem. This problem is defined as follows: Alice and Bob get matrices $X, Y \in \{0,1\}^{m \times n}$ respectively, and they would like to determine whether they agree on at least one row, or differ on all the rows. Our proof implies a lower bound of $m \cdot (n - 2\log m - 6)$ on the communication complexity of this problem as well.

Remark 8. Alexander Kozachinsky [Koz17] pointed out a simple argument that proves a lower bound of $m \cdot n$ for the totalized direct sum $U_n^{\otimes m'}$ (which in turn implies a lower bound of $m \cdot (n-1) - 1$ for the non-totalized direct sum). The argument is as follows. First, observe that it suffices to prove a lower bound of $m \cdot n$ for the problem described in Remark 7 — this can be established by a simple reduction. Next, observe that the matrix of this problem is the m-wise tensor product of the matrix of the inequality function. The rank of the latter matrix is 2^n, and therefore the rank of its m-wise tensor product is $2^{m \cdot n}$. This implies a lower bound of $m \cdot n$ on the communication problem, as required.

3 Proof of Theorem 1

3.1 Proof overview

We prove that the totalized m-fold direct sum of the universal relation on n bits has communication complexity at least $m \cdot (n - O(\log m))$. The proof is by an adversary argument: we describe an adversary that takes a protocol that is "too efficient", and constructs a transcript in which the protocol errs. The adversary constructs the transcript bit by bit, while maintaining the following invariant: at any given point, the adversary can choose an input matrix X for Alice and an input matrix Y for Bob that disagree on all rows and are consistent with the partial transcript that was constructed so far, and the adversary can also choose X and Y to agree on some rows. When the protocol ends, this invariant implies that there are inputs X, Y on which the transcript errs: if the transcript outputs $\perp$ then there is an error since we can choose X and Y that disagree on all rows, and if the output is a tuple $(j_1, \ldots, j_m)$, then there is an error since we can choose X and Y that agree on some rows.

More specifically, throughout the protocol, the adversary maintains a set $I \subseteq [m]$ of *active rows*, which are rows on which the input matrices may agree or disagree. The adversary also maintains a set $\mathcal{Z} \subseteq \{0,1\}^{I \times n}$ of possible assignments to those rows, such that it is possible to find two matrices $Z^1, Z^2 \in \mathcal{Z}$ that disagree on all rows. Finally, the adversary maintains two functions $X, Y : \mathcal{Z} \to \{0,1\}^{m \times n}$ that have the following properties:

- For every $Z \in \mathcal{Z}$, it is consistent with the transcript constructed so far to give Alice the input matrix $X(Z)$ and to Bob the input matrix $Y(Z)$.

- For every $Z \in \mathcal{Z}$, the matrices $X(Z)$ and $Y(Z)$ are extensions of Z

- For every two matrices $Z^1, Z^2 \in \mathcal{Z}$ (not necessarily distinct), the matrices $X(Z^1)$ and $Y(Z^2)$ disagree on all the *inactive* rows in $[m] - I$.

Observe that given all these properties, the invariant that was discussed above is indeed maintained: If the adversary wishes to give the parties two matrices that differ on all the rows, she can choose two such matrices $Z^1, Z^2 \in \mathcal{Z}$, and give the parties the inputs $X(Z^1)$ and $Y(Z^2)$. Otherwise, the adversary can choose a single matrix $Z \in \mathcal{Z}$, and give the parties the inputs $X(Z), Y(Z)$, which agree on the active rows.

When the protocol starts, the set I of active rows is the set $[m]$ of all rows, and the set $\mathcal{Z}$ is the set $\{0,1\}^{m \times n}$ of all matrices. Then, the adversary constructs the transcript bit by bit. When choosing the next bit to be added to the transcript, the adversary chooses the bit in $\{0,1\}$ such that the set $\mathcal{Z}$ decreases by a factor of at most 2 (in other words, at most one bit of information is transmitted). The adversary will continue in this way until almost n bits of information are revealed about a particular row $i \in I$, which runs the risk of fixing the i-th row to a single value (and hence preventing the adversary from choosing matrices that disagree on this row). At this point, the adversary will remove the i-th row from the set of active rows, and will extend the functions $X(Z)$ and $Y(Z)$ such that they always assign the i-th row two different values. The adversary will proceed as until the protocol halts.

It remains to show that if the protocol is "too efficient", then the adversary can maintain the invariant until the protocol ends. To this end, we observe that the protocol has to communicate about n bits in order to turn an active row into an inactive one (more precisely, $n - O(\log m)$ bits). Thus, if the protocol transmitted less than $m \cdot (n - O(\log m))$ bits, then there is still at least one active row when the protocol ends, and therefore the adversary can use the set $\mathcal{Z}$ to show to construct inputs on which the transcript errs.

3.2 Setting up machinery

In order to implement the above argument, we use machinery that was introduced in the work of Edmonds et. al. [EIRS01] on the composition of universal relations, and was further refined by Raz and McKenzie [RM97] in the proof of their simulation theorem. We now introduce the basic notations and propositions of this machinery.

Given a set I of active rows and a set of matrices $\mathcal{Z} \subseteq \{0,1\}^{I \times n}$, we denote by $\delta(\mathcal{Z}) \stackrel{\text{def}}{=} |\mathcal{Z}| / 2^{|I| \cdot n}$ the *density* of $\mathcal{Z}$. We define a sequence of bipartite graphs G_i (for each $i \in I$) that correspond to $\mathcal{Z}$ as follows: The sets of vertices of G_i are $\mathcal{Z}|_i$ and $\mathcal{Z}|_{I-\{i\}}$, and there is an edge between a row $z \in \mathcal{Z}|_i$ and a matrix $Z' \in \mathcal{Z}|_{I-\{i\}}$ if and only if they form together a matrix in $\mathcal{Z}$. We denote by $\mathrm{AvgDeg}_i(\mathcal{Z})$ the average degree of a vertex $\tilde{Z} \in \mathcal{Z}|_{I-\{i\}}$ in G_i, i.e.,

$$\mathrm{AvgDeg}_i(\mathcal{Z}) = \frac{|\mathcal{Z}|}{\left|\mathcal{Z}|_{I-\{i\}}\right|}.$$

We denote by $\mathrm{MinDeg}_i(\mathcal{Z})$ the minimal degree of such a vertex $\tilde{Z} \in \mathcal{Z}|_{I-\{i\}}$. Intuitively, $\mathrm{AvgDeg}_i(\mathcal{Z})$ and $\mathrm{MinDeg}_i(\mathcal{Z})$ are two ways to measure the information that is known about the i-th row conditioned on the other rows — the larger the degree, the less information is known. The advantage of the average degree as a measure of information is that it behaves nicely when additional information is revealed about $\mathcal{Z}$, which is captured by the following easy observation.

Claim 9 ([EIRS01]). *Let $I \subseteq [m]$ and let $\mathcal{Z}' \subseteq \mathcal{Z} \subseteq \{0,1\}^{I \times n}$ be sets of matrices. Then for every $i \in I$ it holds that $\mathrm{AvgDeg}_i(\mathcal{Z}') \geq \frac{|\mathcal{Z}'|}{|\mathcal{Z}|} \cdot \mathrm{AvgDeg}_i(\mathcal{Z})$.*

Proof. Let $i \in I$. It holds that

$$\mathrm{AvgDeg}_i(\mathcal{Z}') = \frac{|\mathcal{Z}'|}{\left|\mathcal{Z}'|_{I-\{i\}}\right|} \geq \frac{|\mathcal{Z}'|}{\left|\mathcal{Z}|_{I-\{i\}}\right|} = \frac{|\mathcal{Z}'|}{|\mathcal{Z}|} \cdot \frac{|\mathcal{Z}|}{\left|\mathcal{Z}|_{I-\{i\}}\right|} = \frac{|\mathcal{Z}'|}{|\mathcal{Z}|} \cdot \mathrm{AvgDeg}_i(\mathcal{Z}),$$

as required. ∎

Another nice property of the average degree is that when we remove a row with a small average degree from I (i.e., when we deactivate the row), the density of $\mathcal{Z}$ *increases*. Intuitively, this means that when we deactivate a row, the information that Alice and Bob sent about this row becomes useless.

Claim 10 ([RM97]). *Let $I \subseteq [m]$, let $\mathcal{Z} \subseteq \{0,1\}^{I \times n}$, and let $i \in I$. Let $\Delta = \mathrm{AvgDeg}_i(\mathcal{Z})$ denote the average degree of the i-th row. Then*

$$\delta(\mathcal{Z}|_{I-\{i\}}) = \frac{2^n}{\Delta} \cdot \delta(\mathcal{Z}).$$

Proof. It holds that

$$\delta(\mathcal{Z}|_{I-\{i\}}) = \frac{\left|\mathcal{Z}|_{I-\{i\}}\right|}{2^{|I-\{i\}| \cdot n}} = \frac{|\mathcal{Z}|}{2^{|I| \cdot n}} \cdot 2^n \cdot \frac{\left|\mathcal{Z}|_{I-\{i\}}\right|}{|\mathcal{Z}|} = \delta(\mathcal{Z}) \cdot 2^n \cdot \frac{1}{\mathrm{AvgDeg}_i(\mathcal{Z})} = \frac{2^n}{\Delta} \cdot \delta(\mathcal{Z}),$$

as required. ∎

The advantage of the minimal degree as a measure of information is that it is preserved when we project the set $\mathcal{Z}$ to a subset of the rows, which is captured by the following observation.

Claim 11 ([RM97]). *Let $I' \subseteq I \subseteq [m]$ and let $\mathcal{Z} \subseteq \{0,1\}^{I \times n}$ be a set of matrices. Then for every $i \in I'$ it holds that $\mathrm{MinDeg}_i(\mathcal{Z}|_{I'}) \geq \mathrm{MinDeg}_i(\mathcal{Z})$.*

Proof. Let $G_1, \ldots, G_m$ be the graphs defined above, and for each $i \in I'$ let G'_i be the corresponding graph for $\mathcal{Z}|_{I'}$. Let $i \in I'$ and let $Z' \in \mathcal{Z}|_{I'-\{i\}}$ be a vertex of G'_i. We prove that the degree of Z' in G'_i is at least $\mathrm{MinDeg}_i(\mathcal{Z})$. Let $\tilde{Z}$ be any extension of Z' in $\mathcal{Z}|_{I-\{i\}}$. Then, it is easy to see that any neighbor $z|$ of $\tilde{Z}$ in G_i is also a neighbor of Z' in G'_i. Therefore, the degree of Z' in G'_i is at least the degree of $\tilde{Z}$ in G_i, which is at least $\mathrm{MinDeg}_i(\mathcal{Z})$. ∎

The following useful lemma, due to [RM97], allows us to switch from the average degree to the minimal degree.

Lemma 12 (The thickness lemma [RM97]). *Let $I \subseteq [m]$ and $\mathcal{Z} \subseteq \{0,1\}^{I \times n}$. Then, there exists a set $\tilde{\mathcal{Z}} \subseteq \mathcal{Z}$ such that $\left|\tilde{\mathcal{Z}}\right| \geq \frac{1}{2} \cdot |\mathcal{Z}|$, and such that $\mathrm{MinDeg}_i(\tilde{\mathcal{Z}}) \geq \frac{1}{2m} \cdot \mathrm{AvgDeg}_i(\mathcal{Z})$ for every $i \in [m]$.*

Let $\Delta = \log m + 2$. We say that a set $\mathcal{Z} \in \{0,1\}^{I \times n}$ is *thick* if $\mathrm{MinDeg}_i(\mathcal{Z}) \geq \Delta$ for every $i \in I$. This property is useful since it allows us to deactivate a row, as will be shown momentarily in Lemma 13. Observe that Lemma 12 implies that any set $\mathcal{Z} \subseteq \{0,1\}^{I \times n}$ can be transformed into a thick set provided that $\mathrm{AvgDeg}_i(\mathcal{Z}) \geq 2 \cdot m \cdot \Delta$ for every $i \in I$ (at the cost of decreasing the size of $\mathcal{Z}$ by a factor of 2). The adversary will strive to maintain the latter property. To this end, the adversary will make sure that all the active rows have average degree at least $4 \cdot m \cdot \Delta$ (in order to have some slackness). If at any point in time, the average degrees of some rows drop below $4 \cdot m \cdot \Delta$, the adversary will remove these rows from the set of active rows. This is done using the following lemma, whose proof combines ideas from [EIRS01] and [RM97].

Lemma 13. *Let $I \subseteq [m]$ and $\mathcal{Z} \subseteq \{0,1\}^{I \times n}$ be a thick set. Then, there exists a set of rows $I' \subseteq I$, a set $\mathcal{Z}' \subseteq \mathcal{Z}|_{I'}$ and functions $\phi_0, \phi_1 : \mathcal{Z}' \to \mathcal{Z}$ that satisfy the following properties:*

- *For every $i' \in I'$, it holds that $\mathrm{AvgDeg}_{i'}(\mathcal{Z}') \geq 4 \cdot m \cdot \Delta$.*

- *For every $Z' \in \mathcal{Z}'$, the matrices $\phi_0(Z'), \phi_1(Z')$ are extensions of Z'.*

- *For every two matrices $Z^{1'}, Z^{2'} \in \mathcal{Z}'$ (not necessarily distinct), it holds that $\phi_0(Z^{1'}), \phi_1(Z^{2'})$ disagree on all rows in $I - I'$.*

- *$\delta(\mathcal{Z}') \geq \delta(\mathcal{Z}) \cdot 2^{(n - 2\log(m) - 4) \cdot (|I| - |I'|)}$.*

Remark 14. The last property is used to capture the idea that in order to deactivate $(|I| - |I'|)$ rows, Alice and Bob had to transmit $(n - O(\log(m))) \cdot (|I| - |I'|)$ bits of information.

Proof. We construct I' using the following iterative process: as long as there exists an index $i \in I$ such that $\mathrm{AvgDeg}_i(I) < 8 \cdot m \cdot \Delta$, we remove it from I. We denote I' by the final set (and denote by I the original set). Let $I - I' = \{i_1, \ldots, i_t\}$ be the indices that were removed. We first observe that

$$\delta(\mathcal{Z}|_{I'}) \geq 2^{(n - \log(m) - \log \Delta - 3) \cdot t} \cdot \delta(\mathcal{Z}) \geq 2^{(n - 2\log(m) - 3) \cdot t} \cdot \delta(\mathcal{Z}),$$

which can be proved by an iterative application of Claim 10. Next, we construct the set $\mathcal{Z}'$ as follows: We first choose uniformly distributed functions $f_1, \ldots, f_t : \{0,1\}^n \to \{0,1\}$. We then put a matrix $Z' \in \mathcal{Z}|_{I'}$ in $\mathcal{Z}'$ if and only if it has two extensions Z^0 and Z^1 in $\mathcal{Z}$ such that for every row $i_j \in I - I'$, the function f_j outputs 0 on the i_j-th row of Z^0, and outputs 1 on the i_j-th row of Z^1. For each such matrix Z', we define $\phi_0(Z') = Z^0$ and $\phi_1(Z') = Z^1$ (if there is more than one possible choice for Z^0, we choose arbitrarily, and the same for Z^1).

It is easy to see that for every choice of $f_1, \ldots, f_t$ this choice of $\mathcal{Z}'$ and ϕ_0, ϕ_1 satisfies the second and third requirements of the lemma. We show that there exists a choice of $f_1, \ldots, f_t$ for which $\mathcal{Z}', \phi_0, \phi_1$ satisfy the first and fourth requirements. To this end we use the following claim, whose proof is deferred to the end of this section.

Claim 15. *Let $Z' \in \mathcal{Z}|_{I'}$. The probability that Z' is put in $\mathcal{Z}'$ (over the choice of $f_1, \ldots, f_t$) is at least $\frac{1}{2}$.*

The latter claim implies that there exists a choice of $f_1, \ldots, f_t$ for which $|\mathcal{Z}'| \geq \frac{1}{2} \cdot |\mathcal{Z}|_{I'}|$. Now, for this choice of $f_1, \ldots, f_t$ it holds that

$$\delta(\mathcal{Z}') \geq \frac{1}{2} \cdot \delta(\mathcal{Z}|_{I'}) \geq 2^{(n - 2\log(m) - 3) \cdot (|I| - |I'|) - 1} \geq 2^{(n - 2\log(m) - 4) \cdot (|I| - |I'|)}.$$

Furthermore, for every $i \in I'$, it follows by Claim 9 that

$$\mathrm{AvgDeg}_i(\mathcal{Z}') \geq \frac{1}{2} \cdot \mathrm{AvgDeg}_i(\mathcal{Z}|_{I'}) \geq 4 \cdot m \cdot \Delta,$$

where the last inequality holds since otherwise, i has average degree in $\mathcal{Z}|_{I'}$ that is less than $8 \cdot m \cdot \Delta$, and therefore it would have been removed in the process of constructing I'. It follows that there exists a subset $\mathcal{Z}'$ that satisfies all four requirements of the lemma, as required. ∎

Proof of Claim 15. We use the same notations as in the proof of Lemma 13. Let $Z' \in \mathcal{Z}|_{I'}$. For every set of indices $\bar{I} \subseteq I - I'$, we say that a matrix $Z \in \mathcal{Z}|_{I' \cup \bar{I}}$ is a *0-extension of Z' to $\bar{I}$* (respectively, *1-extension*) if it is an extension of Z', and for every $i_j \in \bar{I}$ it holds that f_j outputs 0 (respectively 1) on the i_j-th row of Z. Using this notation, our goal is to prove that with probability

7

at least $\frac{1}{2}$, there exists both a 0-extension and a 1-extension of Z' to $I - I'$. We prove a stronger claim: for every $k \in [t]$, let E_j be the event that there exists both a 0-extension and a 1-extension of Z' to $\{i_1, \ldots, i_k\}$. We prove that the probability of E_j is at least $(1 - \frac{1}{2 \cdot m})^k$. Observe that this claim implies that for $k = t$ this probability is at least

$$(1 - \frac{1}{2 \cdot m})^t \geq 1 - \frac{t}{2m} \geq \frac{1}{2},$$

which is what we need to prove.

In order to lower bound the probability of E_k, it suffices to show that $\Pr[E_k] \geq 1 - \frac{1}{2 \cdot m}$ and that $\Pr[E_{k+1}|E_k] \geq 1 - \frac{1}{2 \cdot m}$ for every $k \in [t-1]$. We prove the latter lower bound, and the former lower bound is similar.

Let $k \in [t-1]$. Observe that the event E_k depends only on the choice of $f_1, \ldots, f_k$ — fix such a choice for which E_k occurs. We prove that the probability that there is no 0-extension of Z' to $\{1, \ldots, i_{k+1}\}$ is at most $\frac{1}{4 \cdot m}$. The same argument shows that this holds for a 1-extension, and by the union bound we deduce that the probability that E_{k+1} occurs is at least $1 - \frac{1}{2 \cdot m}$.

Let $Z^{0,j}$ be a 0-extension of Z' to $\{i_1, \ldots, i_k\}$ (such a 0-extension exists since E_k occurs). Let $G'_{i_{k+1}}$ be the graph that corresponds to i_{k+1} and $\mathcal{Z}|_{I' \cup \{i_1, \ldots, i_{k+1}\}}$. Observe that $Z^{0,k}$ is a vertex of $G'_{i_{k+1}}$, and therefore its degree is at least

$$\mathrm{MinDeg}_{i_{k+1}}(\mathcal{Z}|_{I' \cup \{i_1, \ldots, i_{k+1}\}}) \geq \mathrm{MinDeg}_{i_{k+1}}(\mathcal{Z}) \geq \Delta,$$

where the first inequality is due to Claim 11 and the second inequality holds since $\mathcal{Z}$ is thick. This means that there are at least Δ possible values $z \in \mathcal{Z}|_{i_{k+1}}$ that extend $Z^{0,k}$ to $\{i_1, \ldots, i_{k+1}\}$. If f_{j+1} outputs 0 on at least one of these values, then we can add it to $Z^{0,k}$ to obtain a 0-extension of Z' to the i_{k+1}-th row. Now, the probability that f_{k+1} does not output 0 on all these values is at most

$$2^{-\Delta} = 2^{-(\log m + 2)} \geq \frac{1}{4 \cdot m},$$

and therefore the probability that there is no 0-extension of Z' to $\{i_1, \ldots, i_{k+1}\}$ is at most $\frac{1}{4 \cdot m}$, which is what we need to prove. $\blacksquare$

3.3 The adversary argument

We finally turn to prove our Theorem 1, restated next.

Theorem 1. *The deterministic communication complexity of solving the m-fold direct sum of the universal relation over n bits is at least $m \cdot (n - 2\log(m) - 8)$.*

As noted in Section 2, in order to prove Theorem 1 it suffices to prove a lower bound of $m \cdot (n - 2\log(m) - 6)$ on the *totalized m-fold direct sum*. Let $m, n \in \mathbb{N}$ and fix a protocol for the *totalized m-fold direct sum* of the universal relation over n bits. Let c be the maximal number of bits that the protocol transmits, and assume for the sake of contradiction that $c < m \cdot (n - 2\log(m) - 6)$. We design an adversary that finds a transcript π of the protocol and a pair of inputs (X, Y) on which this transcript errs. The adversary constructs the transcript π iteratively, bit by bit. Let π^j be the partial transcript that was constructed after j iterations (so $|\pi^j| = j$). In each iteration, the adversary constructs a set of active rows $I_j \subseteq [m]$, a set of matrices $\mathcal{Z}^j \subseteq \{0,1\}^{I_j \times n}$, and a pair of functions $X^j, Y^j : \mathcal{Z}^j \to \{0,1\}^{m \times n}$ that satisfy the following invariants:

- For every $Z \in \mathcal{Z}^j$, it is consistent with π^j to give Alice the input $X^j(Z)$, and it is also consistent with π^j to give Bob the input $Y^j(Z)$.

- For every $Z \in \mathcal{Z}^j$, the matrices $X^j(Z)$ and $Y^j(Z)$ are extensions of Z.

- For every two matrices $Z^1, Z^2 \in \mathcal{Z}^j$ (not necessarily distinct), the matrices $X^j(Z^1)$ and $Y^j(Z^2)$ disagree on all the rows in $[m] - I_j$.

- For every $i \in I_j$, it holds that $\mathrm{AvgDeg}_i(\mathcal{Z}^j) \geq 4 \cdot \Delta \cdot m$.

- The density of $\mathcal{Z}^j$ is at least $2^{-j+(m-|I_j|)\cdot(n-2\log(m)-6)}$.

The adversary stops when the protocol ends.

We turn to describe how the adversary performs a single iteration. Suppose that the adversary has already performed the first j iterations, and it now performs the $(j+1)$-th iteration. Assume that after the players transmitted the partial transcript π^j, it is Alice's turn to speak (if it is Bob's turn to speak, the adversary's behavior is similar). For every $Z \in \mathcal{Z}^j$, if we give Alice the input $X^j(Z)$, she speaks either 0 or 1. Let $\mathcal{Z}^{j,0}$ and $\mathcal{Z}^{j,1}$ the sets of matrices Z that correspond to the former and latter cases. Without loss of generality, assume that $|\mathcal{Z}^{j,0}| \geq |\mathcal{Z}^{j,1}|$, which also implies that $|\mathcal{Z}^{j,0}| \geq \frac{1}{2} \cdot |\mathcal{Z}^j|$. Then, the adversary appends 0 to the transcript π^j to obtain the new transcript π^{j+1}. Now, if it holds that no average degree in $\mathcal{Z}^{j,0}$ is too low (i.e., $\mathrm{AvgDeg}_i(\mathcal{Z}^{j,0}) \geq 4 \cdot \Delta \cdot m$ for every $i \in I_j$), then the adversary sets $\mathcal{Z}^{j+1} = \mathcal{Z}^{j,0}$, $I_{j+1} = I_j$, $X^{j+1} = X^j$ and $Y^{j+1} = Y^j$, and proceeds to the next iteration. It is not hard to check that the above invariants are maintained.

Suppose that this is not the case, i.e., that there is some row $i \in I_j$ such that $\mathrm{AvgDeg}_i(\mathcal{Z}^{j,0}) < 4 \cdot \Delta \cdot m$. In this case, we are going to use the thickness lemma (Lemma 12) to transform $\mathcal{Z}^{j,0}$ into a thick set, and then use Lemma 13 to deactivate the rows with low average degree. Specifically, observe that since the average degrees were large enough in $\mathcal{Z}^j$ (i.e., $\mathrm{AvgDeg}_i(\mathcal{Z}^j) \geq 4 \cdot \Delta \cdot m$ for every $i \in I_j$) and since $|\mathcal{Z}^{j,0}| \geq \frac{1}{2} \cdot |\mathcal{Z}^j|$, it follows by Claim 9 that all the average degrees of $\mathcal{Z}^{j,0}$ are at least $2 \cdot \Delta \cdot m$. By the thickness lemma, there exists a thick subset $\tilde{\mathcal{Z}} \subseteq \mathcal{Z}^{j,0}$ such that $\delta(\tilde{\mathcal{Z}}) \geq \frac{1}{2} \cdot \delta(\mathcal{Z}^{j,0})$. Next, the adversary applies Lemma 13 to $\tilde{\mathcal{Z}}$ to obtain a new set of rows $I_{j+1} \subseteq I_j$, a new set of matrices $\mathcal{Z}^{j+1} \subseteq \tilde{\mathcal{Z}}|_{I_{j+1}}$, and functions $\phi_0, \phi_1 : \mathcal{Z}^{j+1} \to \tilde{\mathcal{Z}}$ such that

- For every $Z \in \mathcal{Z}^{j+1}$, the matrices $\phi_0(Z), \phi_1(Z)$ are extensions of Z.

- For every two matrices $Z^1, Z^2 \in \mathcal{Z}'$ (not necessarily distinct), it holds that $\phi_0(Z^1), \phi_1(Z^2)$ disagree on all rows in $I_j - I_{j+1}$.

- For every $i \in I_{j+1}$, it holds that $\mathrm{AvgDeg}_i(\mathcal{Z}^{j+1}) \geq 4 \cdot \Delta \cdot m$.

- $\delta(\mathcal{Z}^{j+1}) \geq \delta(\tilde{\mathcal{Z}}) \cdot 2^{(n-2\log(m)-4)\cdot(|I_j|-|I_{j+1}|)}$.

The adversary now chooses $X^{j+1}, Y^{j+1} : \mathcal{Z}^{j+1} \to \{0,1\}^{m \times n}$ by setting $X^{j+1}(Z) = X^j(\phi_0(Z))$ and $Y^{j+1}(Z) = Y^j(\phi_1(Z))$ for every $Z \in \mathcal{Z}^{j+1}$. It is not difficult to see that first four among the above invariants are maintained. To see that the last invariant is maintained, observe that

$$
\begin{aligned}
\delta(\mathcal{Z}^{j+1}) &\geq \delta(\tilde{\mathcal{Z}}) \cdot 2^{(n-2\log(m)-4)\cdot(|I_j|-|I_{j+1}|)} \\
&\geq \frac{1}{2} \cdot \delta(\mathcal{Z}^{j,0}) \cdot 2^{(n-2\log(m)-4)\cdot(|I_j|-|I_{j+1}|)} \\
&\geq \frac{1}{4} \cdot \delta(\mathcal{Z}^j) \cdot 2^{(n-2\log(m)-4)\cdot(|I_j|-|I_{j+1}|)} \\
&\geq \delta(\mathcal{Z}^j) \cdot 2^{(n-2\log(m)-6)\cdot(|I_j|-|I_{j+1}|)} \\
&\geq 2^{-j+(m-|I_j|)\cdot(n-2\log(m)-6)+(n-2\log(m)-7)\cdot(|I_j|-|I_{j+1}|)} \\
&= 2^{-j+(m-|I_j|+|I_j|-|I_{j+1}|)\cdot(n-2\log(m)-6)} \\
&= 2^{-j+(m-|I_{j+1}|)\cdot(n-\log(m)-9)},
\end{aligned}
$$

as required.

Finally, we show that when the protocol halts, the adversary can find a pair of inputs on which the transcript π errs. Let us denote by $\pi, I, \mathcal{Z}, X, Y$ the appropriate objects when the protocol halts. We first observe that $I \neq \emptyset$. To see why, recall that the density of $\mathcal{Z}$ is at least

$$2^{-c+(m-|I|)\cdot(n-2\log(m)-6)}$$

and at most 1. This implies that

$$
\begin{aligned}
2^{-c+(m-|I|)\cdot(n-2\log(m)-6)} &\leq& 1 \\
(m-|I|)\cdot(n-2\log(m)-6) &\leq& c \\
(m-|I|)\cdot(n-2\log(m)-6) &<& m\cdot(n-2\log(m)-6) \\
m-|I| &<& m \\
|I| &>& 0.
\end{aligned}
$$

Hence, I is not empty. Next, we use the following claim, which we prove at the end of this section.

Claim 16. *The set $\mathcal{Z}$ contains two matrices Z^1, Z^2 that disagree on all the rows in I.*

Let $Z^1, Z^2 \in \mathcal{Z}$ be the matrices from the claim. Now, we consider two cases:

- If the transcript π outputs $\perp$, then the adversary gives the input $X(Z^1)$ to Alice and the input $Y(Z^2)$ to Bob. It can be verified that these two matrices disagree on all the rows in $[m]$, and therefore the transcript π errs on this pair of inputs.

- Otherwise, the adversary gives the input $X(Z^1)$ to Alice and the input $Y(Z^1)$ to Bob. These to matrices agree on the rows in I, which is a non-empty set of rows, and therefore the transcript π errs on these inputs.

Thus, the adversary manages to find a transcript π that errs, as required.

Proof of Claim 16. Since all the average degrees of $\mathcal{Z}$ are at least $4 \cdot \Delta \cdot m$, it follows from the thickness lemma 12 that $\mathcal{Z}$ contains a thick subset $\tilde{\mathcal{Z}}$. We prove that $\tilde{\mathcal{Z}}$ contains two matrices Z^1, Z^2 that disagree on all the rows in I. Let us denote $I = \{i_1, \dots, i_t\}$. We construct Z^1, Z^2 iteratively, where in the j-th iteration we choose the values of the i_j-th row of Z^1, Z^2. We now describe the j-th iteration: let us denote $I_{j-1} = \{i_1, \dots, i_{j-1}\}$ and $I_j \stackrel{\text{def}}{=} I_{j-1} \cup \{i_j\}$, and let $Z^{1,j-1}, Z^{2,j-1} \in \tilde{\mathcal{Z}}|_{I_{j-1}}$ be the matrices constructed so far. For every $i \in I_j$, let G'_{i_j} denote the corresponding graph of $\tilde{\mathcal{Z}}|_{I_j}$. Since $\tilde{\mathcal{Z}}$ is thick, it follows from Claim 11 that the minimal degree of G'_{i_j} is

$$\text{MinDeg}_{i_j}(\tilde{\mathcal{Z}}|_{I_j}) \geq \text{MinDeg}_{i_j}(\tilde{\mathcal{Z}}) \geq \Delta.$$

Observe that $Z^{1,j-1}, Z^{2,j-1}$ are vertices of G'_{i_j}, so their degree is at least 2. This means there are two different extensions of $Z^{1,j-1}$ to the i_j-th row, and the same holds for $Z^{2,j-1}$. In particular, it follows that we can extend $Z^{1,j-1}$ and $Z^{2,j-1}$ to the i_j-th using different strings, so they disagree on the i_j-th row. We choose these extensions to be the input $Z^{1,j}, Z^{2,j}$ to the next iteration.

Clearly, when this iterative process ends, the resulting matrices Z^1, Z^2 disagree on all the rows in I. ∎

References

[BBCR10] Boaz Barak, Mark Braverman, Xi Chen, and Anup Rao. How to compress interactive communication. In *STOC*, pages 67–76, 2010.

[Bsh89] Nader H. Bshouty. On the extended direct sum conjecture. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing, May 14-17, 1989, Seattle, Washigton, USA*, pages 177–185, 1989.

[Bsh98] Nader H. Bshouty. On the direct sum conjecture in the straight line model. *J. Complexity*, 14(1):49–62, 1998.

[EIRS01] Jeff Edmonds, Russell Impagliazzo, Steven Rudich, and Jiri Sgall. Communication complexity towards lower bounds on circuit depth. *Computational Complexity*, 10(3):210–246, 2001.

[FKNN95] Tomás Feder, Eyal Kushilevitz, Moni Naor, and Noam Nisan. Amortized communication complexity. *SIAM J. Comput.*, 24(4):736–750, 1995.

[GF81] Giulia Galbiati and Michael J. Fischer. On the complexity of 2-output boolean networks. *Theor. Comput. Sci.*, 16:177–185, 1981.

[GKR14] Anat Ganor, Gillat Kol, and Ran Raz. Exponential separation of information and communication. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 176–185, 2014.

[HW93] Johan Håstad and Avi Wigderson. Composition of the universal relation. In *Advances in computational complexity theory, AMS-DIMACS*, 1993.

[KN97] Eyal Kushilevitz and Noam Nisan. *Communication complexity*. Cambridge University Press, 1997.

[Koz17] Alexander Kozachinsky. Comment on meir's paper "the direct sum of universal relations". *Electronic Colloquium on Computational Complexity (ECCC)*, 24:128, 2017.

[KRW95] 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.

[KW90] Mauricio Karchmer and Avi Wigderson. Monotone circuits for connectivity require super-logarithmic depth. *SIAM J. Discrete Math.*, 3(2):255–265, 1990.

[Mei17] Or Meir. The direct sum of universal relations. *Electronic Colloquium on Computational Complexity (ECCC)*, 24:128, 2017.

[Pau76] Wolfgang J. Paul. Realizing boolean functions on disjoint sets of variables. *Theor. Comput. Sci.*, 2(3):383–396, 1976.

[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.

[RW89] Ran Raz and Avi Wigderson. Probabilistic communication complexity of boolean relations (extended abstract). In *FOCS*, pages 562–567, 1989.

[RW92]   Ran Raz and Avi Wigderson. Monotone circuits for matching require linear depth. *J. ACM*, 39(3):736–744, 1992.

[Sto86]   Quentin F. Stout. Meshes with multiple buses. In *27th Annual Symposium on Foundations of Computer Science, Toronto, Canada, 27-29 October 1986*, pages 264–273, 1986.

[Uhl74]   D Uhlig. On the synthesis of self-correcting schemes from functional elements with a small number of reliable elements. *Matematicheskie Zametki*, 15(6):937–944, 1974.