Electronic Colloquium on Computational Complexity, Report No. 96 (2022)

Communication Complexity of Collision

Mika Göös
EPFL

Siddhartha Jain
EPFL

July 8, 2022

Abstract

The *Collision problem* is to decide whether a given list of numbers $(x_1, \ldots, x_n) \in [n]^n$ is 1-to-1 or 2-to-1 when promised one of them is the case. We show an $n^{\Omega(1)}$ randomised communication lower bound for the natural two-party version of Collision where Alice holds the first half of the bits of each x_i and Bob holds the second half. As an application, we also show a similar lower bound for a weak bit-pigeonhole search problem, which answers a question of Itsykson and Riazanov (CCC 2021).

1 Introduction

Collision problem. The *Collision problem* $\textsc{Col}_N \colon [N]^N \to \{0,1,*\}$ is the following partial (promise) function. The input is a list of numbers $z = (z_1, \ldots, z_N) \in [N]^N$ where N is even. The goal is to distinguish between the following two cases, when promised that z satisfies one of them.

- $\textsc{Col}_N(z) = 0$ iff z is 1-to-1, that is, every number in the list z appears in the list once.
- $\textsc{Col}_N(z) = 1$ iff z is 2-to-1, that is, every number in the list z appears in the list twice.

The Collision problem has been studied exhaustively in quantum query complexity [BHT98, Aar02, AS04, GR04, Kut05, Amb05, Aar12, Aar13, BT16]. It was initially introduced to model the task of breaking collision resistant hash functions, a central problem in cryptanalysis. A robust variant of Collision is complete for NISZK [BSMP91], and consequently it has been featured in black-box oracle separations [LZ17, BCH+19]. The problem has also been used in reductions to show hardness of other problems such as set-equality [Mid04] and various problems in property testing [BHH11]. Upper bounds for Collision has been used to design quantum algorithms for triangle finding [MSS07] and approximate counting [AKKT20].

In this paper, we consider a natural bipartite communication version of this problem, where we split the binary encoding of each number between two parties, Alice and Bob. Specifically, for $N = 2^n$ where n is even, we will define a bipartite function

$$\textsc{BiCol}_N \colon (\{0,1\}^{n/2})^N \times (\{0,1\}^{n/2})^N \to \{0,1,*\}.$$

Here Alice gets as input a list of half-numbers $x = (x_1, \ldots, x_N) \in (\{0,1\}^{n/2})^N$, Bob gets a list of half-numbers $y = (y_1, \ldots, y_N) \in (\{0,1\}^{n/2})^N$, and we view their concatenation $z \coloneqq x \centerdot y$, defined by $z_i \coloneqq x_i y_i$, as an input to $\textsc{Col}_N$. Their goal is to compute $\textsc{BiCol}_N(x, y) \coloneqq \textsc{Col}_N(x \centerdot y)$.

ISSN 1433-8092

Upper bounds. We first observe that BiCol_N admits a deterministic protocol that communicates at most $O(\sqrt{N}\log N)$ bits. Indeed, if $x \centerdot y$ is 1–1, then since Alice's half-numbers are $n/2$ bits long, there are $\sqrt{N}$ distinct half-numbers, each appearing $\sqrt{N}$ many times in x. We may assume this is true also if $x \centerdot y$ is 2–1 (as otherwise it is easy to tell that we are in case 2–1). Consider the set of indices $I \coloneqq \{i \in [N] : x_i = 0^{n/2}\}$, $|I| = \sqrt{N}$. Then $x \centerdot y$ restricted to indices I is 1–1 (resp. 2–1) if the original unrestricted input is 1–1 (resp. 2–1). Hence Alice can send the indices I to Bob, who can determine the value of the function.

If we are allowed randomness, we can do slightly better: there is a randomised protocol of cost $O(N^{1/4}\log N)$. In this protocol, Alice samples a subset $I' \subseteq I$ of size $|I'| = \Theta(N^{1/4})$ uniformly at random and sends it to Bob, who checks for a collision in his part of the input. If the original input was 2–1, then by the birthday paradox, Bob will observe a collision with high probability.

Lower bound. As our main result, we prove a small polynomial lower bound for BiCol_N, which shows that the above randomised protocol cannot be improved too dramatically.

Theorem 1. BiCol_N *has randomised (and even quantum) communication complexity* $\Omega(N^{1/12})$.

We conjecture that the $O(N^{1/4}\log N)$-bit protocol for BiCol_N is essentially optimal (up to logarithmic factors) for randomised protocols. It is an interesting open problem to close this gap.

1.1 Application

Bit-pigeonhole principle. We also show a lower bound for a search problem associated with the *pigeonhole principle*. We define PHP_N^M where $M > N$ as the following search problem: On input $z = (z_1, \dots, z_M) \in [N]^M$ the goal is to output a collision, that is, a pair of distinct indices $i, j \in [M]$ such that $z_i = z_j$. We note that PHP_N^M is a *total* search problem (not a promise problem); it always has a solution since we require $M > N$. As before, we can turn PHP_N^M naturally into a bipartite communication search problem BiPHP_N^M where $N = 2^n$ so that

- Alice holds $x = (x_1, \dots, x_M) \in (\{0,1\}^{n/2})^M$;
- Bob holds $y = (y_1, \dots, y_M) \in (\{0,1\}^{n/2})^M$; and
- the goal is find a collision, that is, distinct $i, j \in [M]$ such that $x_i y_i = x_j y_j$.

Lower bounds. Itsykson and Riazanov [IR21] proved that BiPHP_N^{N+1} requires $\Omega(\sqrt{N})$ bits of randomised communication. Their proof was via a randomised reduction from set-disjoitness. A corollary of their result is that any proof system that can be efficiently simulated by randomised protocols (most notably, tree-like Res($\oplus$) [IS20]) requires exponential size to refute bit-pigeonhole formulas featuring $N + 1$ pigeons and N holes. They asked whether a similar communication lower bound could be proved for the *weak* pigeonhole principle with $M = 2N$ pigeons and N holes. We answer their question in the affirmative in the following theorem.

Theorem 2. BiPHP_N^{2N} *has randomised (and even quantum) communication complexity* $\Omega(N^{1/12})$.

Previously, Hrubeš and Pudlák [HP17] showed a small polynomial lower bound for BiPHP_N^M for every $M > N$ against deterministic (and even dag-like) protocols. By contrast, Theorem 2 is the first randomised lower bound in the $M = 2N$ regime.

1.2 Techniques

Our proof of Theorem 1 proceeds as follows. A popular method to prove communication lower bounds is to start with a partial boolean function $f\colon \{0,1\}^n \to \{0,1,*\}$ that is hard to compute for decision trees and then apply a *lifting theorem* (we use one due to Sherstov [She11]) to conclude that the function $f \circ g$ obtained by composing f with a small gadget $g\colon \Sigma \times \Sigma \to \{0,1\}$ is hard for communication protocols. Here $f \circ g\colon \Sigma^n \times \Sigma^n \to \{0,1,*\}$ is the communication problem where Alice holds $x \in \Sigma^n$, Bob holds $y \in \Sigma^n$, and their goal is to output

$$(f \circ g)(x, y) \;\coloneqq\; f(g(x_1, y_1), \ldots, g(x_n, y_n)).$$

A straightforward application of lifting often produces communication problems that are "artificial" since they are of the composed form. In particular, at first blush, it seems that the $\textsc{BiCol}_N$ problem cannot be written in the form $f \circ g$ for any f and any g for which a lifting theorem holds. To address this issue, our main technical innovation is to show how the composed function $\textsc{Col}_N \circ g$, where g is a sufficiently "regular" gadget, can indeed be *reduced* to the natural problem $\textsc{BiCol}_{N'}$. In this reduction, the input length will blow up polynomially, $N' = N^{\Theta(1)}$, which is the main reason why we only get a small polynomial lower bound. Our new reduction generalises a previous reduction from [IR21, §6], which was tailored for the 2-bit $\textsc{Xor}$ gadget.

To prove Theorem 2 we give a randomised *decision-to-search* reduction from $\textsc{BiCol}_N$ to $\textsc{Bi}\mathrm{PHP}^{2N}_N$. That is, we show that if there is an efficient randomised protocol for solving the *total* search problem $\textsc{Bi}\mathrm{PHP}^{2N}_N$, then there is an efficient randomised protocol for solving the *promise* problem $\textsc{BiCol}_N$. Given this reduction, Theorem 2 then follows from Theorem 1. Similar style of randomised reductions have been considered in prior works [RW92, HN12, GP18, IR21], although they have always reduced from set-disjointness.

2 Reductions and regular functions

We assume some familiarity with communication complexity; see, e.g., the textbooks [KN97, RY20]. In particular, it is often useful to view a bipartite function $f\colon \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ as a 2^n-by-2^n boolean matrix. We now give several definitions for the purposes of the proof of our main result.

Definition 3 (Rectangular reduction)**.** For bipartite functions f, g with domains $\{0,1\}^n \times \{0,1\}^n$ and $\{0,1\}^m \times \{0,1\}^m$, we write $f \leq g$ if there is a *rectangular reduction* from f to g, that is, there exist $a\colon \{0,1\}^n \to \{0,1\}^m$ and $b\colon \{0,1\}^n \to \{0,1\}^m$ such that $f(x,y) = g(a(x), b(y))$ for all x, y.

Next, using basic language from group theory, we define a new class of highly symmetric boolean functions that we call *regular*. (We borrow the term *regular* from group theory where group actions satisfying the property in Definition 4 below are called *regular*.)

Let Π_n denote the symmetric group on $[n]$, that is, the set of all permutations $[n] \to [n]$. Let $S \subseteq \Pi_n \times \Pi_n$ be any group. We let S act on the set $[n] \times [n]$ by permuting the rows and columns, that is, an element $s = (s^A, s^B) \in S$ acts on $(x,y) \in [n] \times [n]$ by $s \cdot (x,y) \coloneqq (s^A(x), s^B(y))$. For $(x,y) \in [n] \times [n]$, we define its *orbit* by $S \cdot (x,y) \coloneqq \{s \cdot (x,y) : s \in S\}$.

Definition 4 (Regular function)**.** A bipartite function $f\colon \{0,1\}^k \times \{0,1\}^k \to \{0,1\}$ is *regular* if there is a group $S \subseteq \Pi_{2^k} \times \Pi_{2^k}$ acting on the domain of f such that the orbit of any $(x,y) \in f^{-1}(b)$, where $b \in \{0,1\}$, equals $f^{-1}(b)$, and, moreover, for every pair of inputs $(x_1,y_1), (x_2,y_2) \in f^{-1}(b)$ there is a unique $s \in S$ such that $s \cdot (x_1, y_1) = s \cdot (x_2, y_2)$.

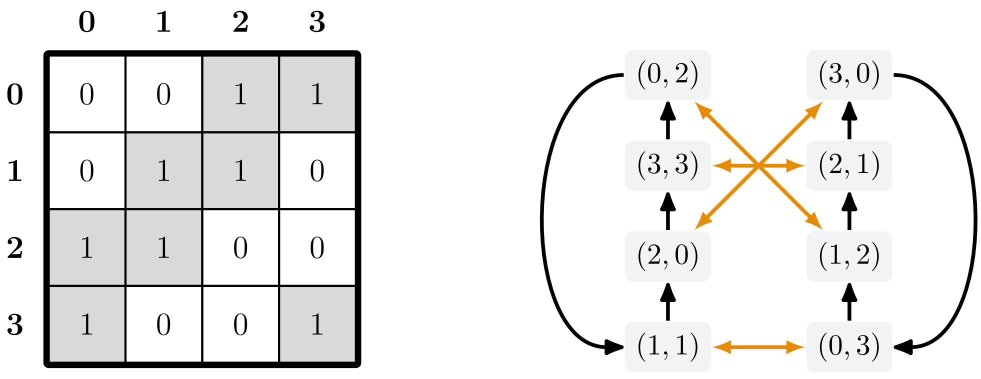

(a) **(b)**

Figure 1: **(a)** The bipartite function $\textsc{Ver}\colon \mathbb{Z}_4 \times \mathbb{Z}_4 \to \{0,1\}$. **(b)** The group relative to which $\textsc{Ver}$ is regular is generated by two elements whose actions on $\textsc{Ver}^{-1}(1)$ are illustrated here. The first generator is $(x, y) \mapsto (x+1, y-1)$ (black arrows) and the second is $(x, y) \mapsto (1-x, -y)$ (orange arrows).

It follows from the definition that $|S| = |f^{-1}(b)| = 2^{2k-1}$ for both $b \in \{0,1\}$. A simple example of a regular function is the 2-bit $\textsc{Xor}$ function together with the 2-element group consisting of the identity map and the map $(x, y) \mapsto (\neg x, \neg y)$. However, the $\textsc{Xor}$ function does not satisfy a fully general lifting theorem. This is why we consider the following more complicated gadget, called a *versatile* gadget, which has been shown to satisfy various lifting theorems [She11, GP18, ABK21].

Definition 5. $\textsc{Ver}\colon \mathbb{Z}_4 \times \mathbb{Z}_4 \to \{0,1\}$ is defined by $\textsc{Ver}(x, y) \coloneqq 1$ iff $x + y \pmod 4 \in \{2,3\}$.

Lemma 6. $\textsc{Ver}$ *is regular.*

Proof. Consider the group $S \subseteq \Pi_4 \times \Pi_4$ generated by the elements $(x, y) \mapsto (x+1, y-1)$ and $(x, y) \mapsto (1-x, -y)$ where we use modulo 4 arithmetic. By explicit computations, we see that (here we list each element as a function of (x, y))

$$S = \left\{ \begin{array}{llll} (x, y), & (x+1, y-1), & (x+2, y-2), & (x+3, y-3), \\ (1-x, -y), & (2-x, 3-y), & (3-x, 2-y), & (-x, 1-y) \end{array} \right\}.$$

It is straightforward to check that S gives rise to orbits $\textsc{Ver}^{-1}(0)$ and $\textsc{Ver}^{-1}(1)$; see Figure 1. Moreover, since $|S| = 8 = |\textsc{Ver}^{-1}(b)|$ for $b \in \{0,1\}$, the uniqueness property holds, too. $\square$

Previously, [GP18] showed that $\textsc{Ver}$ is *random self-reducible*, that is, it admits a randomised reduction that maps any fixed input $(x, y) \in \textsc{Ver}^{-1}(b)$ into a uniform random input in $\textsc{Ver}^{-1}(b)$. It is easy to see that if a function is regular, then it is also random self-reducible (the random self-reduction is to apply a random symmetry from S). The converse, however, is unclear to us: If f is random self-reducible and balanced (meaning $|f^{-1}(0)| = |f^{-1}(1)|$), is it necessarily regular?

3 Lower bound for bipartite collision

In this section we prove Theorem 1. We start with a standard application of a lifting theorem to establish a lower bound for the (somewhat artificial) composed function $\textsc{Col}_N \circ \textsc{Ver}$. Here we think of $\textsc{Col}_N$ as a boolean function $(\{0,1\}^n)^N \to \{0,1\}$ where $N = 2^n$.

Lemma 7. $\textsc{Col}_N \circ \textsc{Ver}$ *has randomised (and even quantum) communication complexity* $\Omega(N^{1/3})$.

Proof. Aaronson and Shi [AS04] (building on [Aar02]) showed that $\deg_{1/3}(\textsc{Col}_N) \geq \Omega(N^{1/3})$ where $\deg_{1/3}(f)$ for a partial boolean function f is the least degree of a multivariate polynomial $p(x)$ such that $p(x) = f(x) \pm 1/3$ for all x such that $f(x) \in \{0,1\}$ and $|p(x)| \leq 4/3$ for all x with $f(x) = *$. Sherstov [She11, §12] proved that for any partial boolean function f, we have that the randomised (and even quantum) communication complexity of $f \circ \textsc{Ver}$ is at least $\Omega(\deg_{1/3}(f))$. Combining these two results proves the lemma. $\square$

The challenging part of the proof is to find a reduction from $\textsc{Col}_N \circ g$ to $\textsc{BiCol}_{N'}$ where g is a regular gadget and N' is polynomially larger than N. Choosing $g \coloneqq \textsc{Ver}$ in the following theorem and combining it with Lemma 7 completes the proof of Theorem 1. Note that the input length becomes $N' \coloneqq N^4$ so that we obtain the lower bound $\Omega(N^{1/3}) = \Omega(N'^{1/12})$, as claimed.

Theorem 8. *Let $g\colon \{0,1\}^k \times \{0,1\}^k \to \{0,1\}$ be a regular gadget. For every $N = 2^n$ we have*

$$\textsc{Col}_N \circ g \;\leq\; \textsc{BiCol}_{N^{2k}}.$$

Proof. Consider the bipartite function $\textsc{Col}_N \circ g$. Alice's input here is an N-tuple $(a^{(1)}, \ldots, a^{(N)})$ where $a^{(j)} \in (\{0,1\}^k)^n$ for each $j \in [N]$. Bob's input $(b^{(1)}, \ldots, b^{(N)})$ has a similar form. These bipartite inputs encode, via the gadgets, the input $(z^{(1)}, \ldots, z^{(N)})$ to $\textsc{Col}_N$ such that

$$z^{(j)} \;\coloneqq\; g^n(a^{(j)}, b^{(j)}) \;\coloneqq\; (g(a_1^{(j)}, b_1^{(j)}), \ldots, g(a_n^{(j)}, b_n^{(j)})) \;\in\; \{0,1\}^n \quad \text{where} \quad a_i^{(j)}, b_i^{(j)} \in \{0,1\}^k.$$

Let $S \subseteq \Pi_{2^k} \times \Pi_{2^k}$ be the symmetry group relative to which g is regular. Recall that $|S| = 2^{2k-1}$ and each $s \in S$ has the form $s = (s^A, s^B)$ with $s^A, s^B \in \Pi_{2^k}$. We fix an arbitrary ordering of the elements of S and write $S(i)$ for the i-th element in this ordering. Thus $S = \{S(1), \ldots, S(2^{2k-1})\}$.

We first describe how the reduction expands each individual input $(a, b) \coloneqq (a^{(j)}, b^{(j)})$ to g^n into an ordered list of inputs to g^n. In more detail, the reduction

- takes an input $(a, b) = (a_1, \ldots, a_n, b_1, \ldots, b_n) \in (\{0,1\}^k)^{2n}$ to g^n, and
- returns $\textsc{Unfold}(a, b) \in (\{0,1\}^{2kn})^{N^{2k-1}}$, an *ordered list* of N^{2k-1} many inputs to g^n.

For any n-tuple of indices $I = (i_1, \ldots i_n) \in [|S|]^n$, we define the I-th pair in $\textsc{Unfold}(a, b)$ by

$$\textsc{Unfold}(a, b)_I \;\coloneqq\; (\underbrace{s_1^A(a_1) s_2^A(a_2) \ldots s_n^A(a_n)}_{\text{Alice's half}},\; \underbrace{s_1^B(b_1) s_2^B(b_2) \ldots s_n^B(b_n)}_{\text{Bob's half}}) \quad \text{where} \quad s_j \coloneqq S(i_j).$$

Besides each pair in the list $\textsc{Unfold}(a, b)$ being an input to g^n, we will also soon interpret them as pairs of half-numbers that are part of the input to $\textsc{BiCol}_{N^{2k}}$.

Below, we write $\textsc{SetUnfold}(a, b) \subseteq \{0,1\}^{2kn}$ for the *set* of elements in the list $\textsc{Unfold}(a, b)$, that is, ignoring the ordering and multiplicity of elements.

Claim 9. *We have the following properties.*

(i) $\textsc{SetUnfold}(a, b) = (g^n)^{-1}(z) = g^{-1}(z_1) \times \cdots \times g^{-1}(z_n)$ *where* $z_i \coloneqq g(a_i, b_i)$.

(ii) All pairs in $\textsc{Unfold}(a, b)$ *are distinct.*

(iii) Suppose $g^n(a, b) \neq g^n(a', b')$. *Then* $\textsc{SetUnfold}(a, b) \cap \textsc{SetUnfold}(a', b') = \emptyset$.

(iv) Suppose $g^n(a, b) = g^n(a', b')$. *Then* $\textsc{SetUnfold}(a, b) = \textsc{SetUnfold}(a', b')$.

Proof. Item (i): Up to reordering of bits, the set equals $(S\cdot(a_1,b_1))\times(S\cdot(a_2,b_2))\times\cdots\times(S\cdot(a_n,b_n))$. By regularity, the orbit $S\cdot(a_i,b_i)$ is equal to $g^{-1}(z_i)$ for any i. Item (ii): The uniqueness property of the regular group action ensures that we do not get any repeated elements. Item (iii): If $z \coloneqq g^n(a,b) \neq g^n(a',b') =: z'$ then there is some i such that $z_i \neq z'_i$. The i-th component of every pair in $\textsc{Unfold}(a,b)$ lies in $g^{-1}(z_i)$ while the i-th component of every pair in $\textsc{Unfold}(a,b)$ lies in $g^{-1}(z'_i)$. The claim follows since these preimage sets are disjoint. Item (iv): If $g^n(a,b)=g^n(a',b')$, then (i) shows $\textsc{Unfold}$ produces the same set for both (a,b) and (a',b'). ☐

Our final reduction from $\textsc{Col}_N \circ g$ maps Alice's $(a^{(1)},\dots,a^{(N)})$ and Bob's $(b^{(1)},\dots,b^{(N)})$ (which together encode the input $z=(z^{(1)},\dots,z^{(N)})$ to $\textsc{Col}_N$) to an input to $\textsc{BiCol}_{N^{2k}}$ given by

$$\textsc{Unfold}(a^{(1)},b^{(1)}),\dots,\textsc{Unfold}(a^{(N)},b^{(N)}).$$

Note that the reduction is rectangular: Alice can compute her part of the input, and Bob his.

It remains to check that the reduction treats 1–1 and 2–1 inputs correctly. If the input z to $\textsc{Col}_N$ is 1–1, then the reduction produces a 1–1 input by (ii) and (iii). If the input z to $\textsc{Col}_N$ is 2–1 then for every index i there is exactly one more index j such that $z^{(i)} \coloneqq g^n(a^{(i)},b^{(i)}) = g^n(a^{(j)},b^{(j)}) =: z^{(j)}$. Hence, by (iv) the lists $\textsc{Unfold}(a^{(i)},b^{(i)})$ and $\textsc{Unfold}(a^{(j)},b^{(j)})$ have every element colliding with each other. This produces a 2–1 input. ☐

4 Lower bound for bipartite pigeonhole

In this section we prove Theorem 2. We do it by describing a reduction from the decision problem $\textsc{BiCol}_N$ to the search problem $\textsc{BiPHP}_N^{2N}$.

Theorem 10. *If there is a randomised protocol for $\textsc{BiPHP}_N^{2N}$ of communication cost d, then there is a randomised protocol for $\textsc{BiCol}_N$ of cost $O(d)$.*

Proof. The proof idea is to start with an input to $\textsc{BiCol}_N$ and then append it with more numbers to construct an input to $\textsc{BiPHP}_N^{2N}$. Adding more numbers will create some new collisions in the input list, but our reduction will remember which collisions where "planted" during the reduction. We then randomly shuffle the input list so as to make the planted collisions indistinguishable from collisions (if any) coming from the original input to $\textsc{BiCol}_N$. We now explain this in more detail.

Let (x,y) be an input to $\textsc{BiCol}_N$. That is, Alice holds $x=(x_1,\dots,x_N)\in(\{0,1\}^{n/2})^N$ and Bob holds $y=(y_1,\dots,y_N)\in(\{0,1\}^{n/2})^N$. In the reduction, we first append Alice's input by the *planted* half-numbers $(a_1,\dots,a_N)\in(\{0,1\}^{n/2})^N$ and Bob's input by the *planted* half-numbers $(b_1,\dots,b_N)\in(\{0,1\}^{n/2})^N$ where the concatenated strings a_ib_i, $i\in[N]$, range lexicographically over all binary numbers in $\{0,1\}^n$.

Next, Alice and Bob use public randomness to sample a permutation $\pi\colon[2N]\to[2N]$ uniformly at random, which they then use to permute their lists of length $2N$. While doing so, they remember which positions in the permuted list occupy planted numbers (namely, those in positions $\pi(\{N+1,\dots,2N\})$). Call the resulting list (x',y'). We now let Alice and Bob run the hypothesised protocol $\mathcal{P}$ for $\textsc{BiPHP}_N^{2N}$ on input (x',y') to find some collision $x'_iy'_i = x'_jy'_j$ where $i\neq j$. (We assume for simplicity that $\mathcal{P}$ finds a collision with probability 1. The following analysis can be adapted even when $\mathcal{P}$ errs with bounded probability.)

We have two cases depending on whether (x,y) was 1–1 or 2–1 (see Figure 2):

- If (x,y) was 1–1 then (x',y') is 2–1. Moreover, each collision in (x',y') involves a planted number. In particular, the collision $\{i,j\}$ found by the protocol always features at least one planted number.

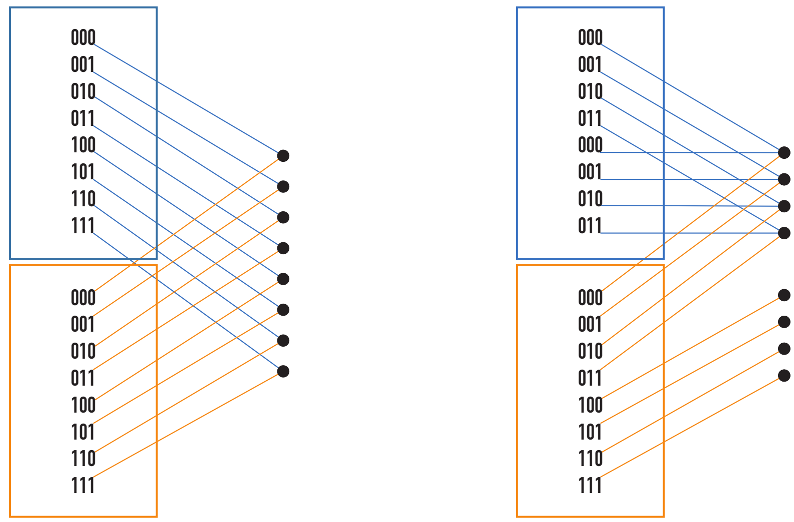

Figure 2: Illustration of collisions in 1–1 and 2–1 inputs. The original input (x, y) is drawn at the top, and the planted numbers (a, b) are drawn at the bottom.

- If (x, y) was 2–1 then (x', y') is an input where $N/2$ many numbers appear thrice, and $N/2$ numbers appear once. We claim that the collision $\{i, j\}$ found by $\mathcal{P}$ will not feature a planted number with probability at least $1/3$ (over the random choice of π). Indeed, let $k \notin \{i, j\}$ be the third position such that $x'_i y'_i = x'_j y'_j = x'_k y'_k$. Then conditioned on π having produced the input (x', y'), each position in $\{i, j, k\}$ is equally likely to occupy a planted number. Thus, with probability $1/3$, the planted number lies in position k and not in $\{i, j\}$.

Our protocol for BiCol_N guesses that (x, y) is 2–1 if the collision $\{i, j\}$ returned by $\mathcal{P}$ does not involve a planted number. We can further reduce the error probability down to $(2/3)^t$ by repeating the randomised reduction and $\mathcal{P}$ some $t = O(1)$ times and seeing if any one of these runs finds a collision without a planted number. $\square$

Acknowledgements

We thank anonymous RANDOM reviewers for their helpful comments.

References

[Aar02] Scott Aaronson. Quantum lower bound for the collision problem. In *Proceedings of the 34th Symposium on Theory of Computing (STOC)*, pages 635–642. ACM, 2002. doi:10.1145/509907.509999.

[Aar12] Scott Aaronson. Impossibility of succinct quantum proofs for collision-freeness. *Quantum Information and Computation*, 12(1-2):21–28, 2012. doi:10.26421/QIC12.1-2-3.

[Aar13] Scott Aaronson. The collision lower bound after 12 years, 2013. QStart talk. URL: https://scottaaronson.blog/?p=1458.

[ABK21] Anurag Anshu, Shalev Ben-David, and Srijita Kundu. On query-to-communication lifting for adversary bounds. In *Proceedings of the 36th Computational Complexity Conference (CCC)*, volume 200, pages 30:1–30:39. Schloss Dagstuhl, 2021. doi:10.4230/LIPICS.CCC.2021.30.

[AKKT20] Scott Aaronson, Robin Kothari, William Kretschmer, and Justin Thaler. Quantum lower bounds for approximate counting via laurent polynomials. In Shubhangi Saraf, editor, *35th Computational Complexity Conference, CCC 2020, July 28-31, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 169 of *LIPIcs*, pages 7:1–7:47. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.CCC.2020.7.

[Amb05] Andris Ambainis. Polynomial degree and lower bounds in quantum complexity: collision and element distinctness with small range. *Theory Comput.*, 1:37–46, 2005. doi:10.4086/toc.2005.v001a003.

[AS04] Scott Aaronson and Yaoyun Shi. Quantum lower bounds for the collision and the element distinctness problems. *Journal of the ACM*, 51(4):595–605, jul 2004. doi:10.1145/1008731.1008735.

[BCH+19] Adam Bouland, Lijie Chen, Dhiraj Holden, Justin Thaler, and Prashant Nalini Vasudevan. On the power of statistical zero knowledge. *SIAM Journal on Computing*, 49(4):FOCS17–1–FOCS17–58, 2019. doi:10.1137/17m1161749.

[BHH11] Sergey Bravyi, Aram Harrow, and Avinatan Hassidim. Quantum algorithms for testing properties of distributions. *IEEE Transactions on Information Theory*, 57(6):3971–3981, 2011. doi:10.1109/TIT.2011.2134250.

[BHT98] Gilles Brassard, Peter Høyer, and Alain Tapp. Quantum cryptanalysis of hash and claw-free functions. In *Proceedings of the 3rd Latin American Symposium on Theoretical Informatics (LATIN)*, pages 163–169. Springer, 1998.

[BSMP91] Manuel Blum, Alfredo De Santis, Silvio Micali, and Giuseppe Persiano. Noninteractive zero-knowledge. *SIAM Journal on Computing*, 20(6):1084–1118, 1991. doi:10.1137/0220068.

[BT16] Mark Bun and Justin Thaler. Dual polynomials for collision and element distinctness. *Theory Comput.*, 12(1):1–34, 2016. doi:10.4086/toc.2016.v012a016.

[GP18] Mika Göös and Toniann Pitassi. Communication lower bounds via critical block sensitivity. *SIAM Journal on Computing*, 47(5):1778–1806, 2018. doi:10.1137/16M1082007.

[GR04] Lov K. Grover and Terry Rudolph. How significant are the known collision and element distinctness quantum algorithms? *Quantum Inf. Comput.*, 4(3):201–206, 2004. doi:10.26421/QIC4.3-5.

[HN12] Trinh Huynh and Jakob Nordström. On the virtue of succinct proofs: Amplifying communication complexity hardness to time–space trade-offs in proof complexity. In *Proceedings of the 44th Symposium on Theory of Computing (STOC)*, pages 233–248. ACM, 2012. doi:10.1145/2213977.2214000.

[HP17] Pavel Hrubeš and Pavel Pudlák. Random formulas, monotone circuits, and interpolation. In *Proceedings of the 58th Symposium on Foundations of Computer Science (FOCS)*, pages 121–131, 2017. doi:10.1109/FOCS.2017.20.

[IR21] Dmitry Itsykson and Artur Riazanov. Proof complexity of natural formulas via communication arguments. In *Proceedings of 36th Computational Complexity Conference (CCC)*, volume 200, pages 3:1–3:34. Schloss Dagstuhl, 2021. doi:10.4230/LIPIcs.CCC.2021.3.

[IS20] Dmitry Itsykson and Dmitry Sokolov. Resolution over linear equations modulo two. *Annals of Pure and Applied Logic*, 171(1):1–31, 2020. doi:10.1016/j.apal.2019.102722.

[KN97] Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, 1997. doi:10.1017/CBO9780511574948.

[Kut05] Samuel Kutin. Quantum lower bound for the collision problem with small range. *Theory of Computing*, 1(2):29–36, 2005. doi:10.4086/toc.2005.v001a002.

[LZ17] Shachar Lovett and Jiapeng Zhang. On the impossibility of entropy reversal, and its application to zero-knowledge proofs. In *Proceedings of the 15th Theory of Cryptography Conference (TCC)*, pages 31–55. Springer, 2017. doi:10.1007/978-3-319-70500-2_2.

[Mid04] Gatis Midrijānis. A polynomial quantum query lower bound for the set equality problem. In *Proceedings of the 31st International Conference on Automata, Languages and Programming (ICALP)*, volume 3142, pages 996–1005. Springer, 2004. doi:10.1007/978-3-540-27836-8_83.

[MSS07] Frédéric Magniez, Miklos Santha, and Mario Szegedy. Quantum algorithms for the triangle problem. *SIAM Journal on Computing*, 37(2):413–424, January 2007. doi:10.1137/050643684.

[RW92] Ran Raz and Avi Wigderson. Monotone circuits for matching require linear depth. *Journal of the ACM*, 39(3):736–744, jul 1992. doi:10.1145/146637.146684.

[RY20] Anup Rao and Amir Yehudayoff. *Communication Complexity: And Applications*. Cambridge University Press, 2020. doi:10.1017/9781108671644.

[She11] Alexander Sherstov. The pattern matrix method. *SIAM Journal on Computing*, 40(6):1969–2000, 2011. doi:10.1137/080733644.

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il