

Quantum Communication Advantage in TFNP

Mika Göös
EPFL

Tom Gur
Cambridge

Siddhartha Jain
UT Austin

Jiawei Li
UT Austin

November 5, 2024

Abstract

We exhibit a total search problem whose communication complexity in the quantum SMP (simultaneous message passing) model is exponentially smaller than in the classical two-way randomized model. Moreover, the quantum protocol is computationally efficient and its solutions are classically verifiable, that is, the problem lies in the communication analogue of the class TFNP. Our problem is a bipartite version of a query complexity problem recently introduced by Yamakawa and Zhandry (JACM 2024). We prove the classical lower bound using the structure-vs-randomness paradigm for analyzing communication protocols.

1 Introduction

A foremost goal in quantum communication complexity is to demonstrate an exponential advantage of quantum protocols over classical protocols. Such quantum speed-up results come in two flavors:

- **Partial problems** have a promise about the structure of the problem input. A long line of work [BCW99, Raz99, BJK04, GKK⁺08, KR11, Gav16, Gav19, Gav21, GRT22, LLY24] has culminated in strong separations results: Gavinsky [Gav21] (for a search problem) and Girish, Raz, and Tal [GRT22] (for a function problem) exhibited partial problems with an exponential separation between the quantum SMP (simultaneous message passing) model and the classical two-way randomized communication model; that is, separations between the weakest standard quantum model and the strongest standard classical model.
- **Total problems** come with no structural promises—the problem is defined for all possible input strings. Obtaining exponential quantum speed-ups is notoriously difficult for total problems. In some settings, this is even known to be impossible: In decision tree complexity, classical and quantum query complexities for any total function are polynomially related [ABK⁺21]. In communication complexity, the seminal work of Bar-Yossef, Jayram, and Kerenidis [BJK04] showed an exponential quantum–classical separation in the restricted setting of one-way protocols for a total search problem (the *hidden matching* problem). Since then, essentially no progress has been made for total problems.

In a recent breakthrough, Yamakawa and Zhandry [YZ24a] introduced a new total search problem that exhibits an exponential quantum–classical advantage in the query complexity model. In this work, we define a communication variant of their problem, called *Bipartite NullCodeword* and obtain the strongest quantum–classical communication separation so far for any *total problem*.

Theorem 1 (Main result). *There exists a total two-party search problem $S \subseteq \{0, 1\}^N \times \{0, 1\}^N \times Q$ for $N = 2^{\Theta(n)}$ that admits a poly(n)-bit protocol in the quantum SMP model, but requires $2^{n^{\Omega(1)}}$ bits of communication in the randomized two-way communication model.*

The separation in [Theorem 1](#) is qualitatively as strong as we can hope to get for any total relation. Recall that in a quantum SMP protocol, Alice and Bob send a quantum state to a third-party referee, Charlie, who outputs the answer without further interaction. Even the one-way model can simulate this since Bob can emulate Charlie. Furthermore, our separation has two additional desirable properties.

Verifiability: The relation we construct is *verifiable* using polylog(N) bits of classical communication. It thus lies in the communication analogue of the complexity class TFNP (Total Functional NP) [[GKRS19](#), [dRGR22](#)].

Efficiency: Our quantum SMP protocol is *computationally efficient*, i.e, all parties can be implemented using polylog(N) size quantum circuits with oracle access to the input. This is the same notion of efficiency as in the result of [[GRT22](#)].

Besides being fundamental considerations in theoretical computer science, verifiability and efficiency are also desirable properties for a quantum–classical communication complexity separation to be used as an experiment to demonstrate *unconditional* quantum advantage (a.k.a. quantum information supremacy [[ABK24](#)]). We see our result as a theoretical feasibility result for designing such an experiment based on communication complexity. We compare our result to several other notable quantum–classical communication separations in [Table 1](#).

1.1 Other related work

The Yamakawa–Zhandry problem. The Yamakawa–Zhandry problem [[YZ24a](#)] was a breakthrough average-case query separation. Since then, it has proven to be a useful in other settings with modifications to the problem. Before our work, it was tweaked to make progress on the question of getting an oracle separation of QMA and QCMA [[LLPY24](#), [BK24a](#)]. More relevant to our work, Li et al. [[LLPY24](#)] also use the Yamakawa–Zhandry problem to get a new separation between one-way quantum and one-way randomized communication complexity. However, their problem is easy in the classical two-way model since Bob’s input is short.

Query-to-communication lifting. Most techniques in communication complexity literature which can distinguish quantum and classical communication complexity are tailored to promise problems. One technique that works for total problems (functions and relations) is query-to-communication lifting [[GPW20](#), [CFK⁺21](#)].¹ The current strongest separation for a *total boolean function* between the quantum and classical two-way communication is cubic: This is achieved by lifting an analogous cubic separation in query complexity obtained by Sherstov, Storozhenko, and Wu [[SSW23](#), Theorem 1.5] and independently by Bansal and Sinha [[BS21](#), Corollary 1.5].

The exponential quantum query separation from the Yamakawa–Zhandry problem [[YZ24a](#)] can also be converted to one for a total problem via lifting (Lifted NULLCODEWORD in [Table 1](#)). Given this, we can get an exponential separation between quantum and classical two-way communication for a *total relation*. However, we do not know how to simulate queries to the input of the original relation for the lifted relation by using less than two rounds of communication. This limits us to a two-way vs. two-way separation, as opposed to a quantum SMP vs. classical two-way separation.

¹Lifting theorems show that $CC(f \circ g) = \Omega(Q(f))$ where g is a boolean function gadget, CC is a communication complexity measure and Q is a query complexity measure.

Candidate problem	Reference	Quantum u.b.	Classical l.b.	f / R	Totality
Vector in Subspace	[Raz99, KR11]	one-way	two-way	function	partial
Gap Hamming Relation	[Gav21]	SMP	two-way	relation	partial
FORRELATION \circ XOR	[GRT22]	SMP	two-way	function	partial
Hidden Matching	[BJK04]	one-way	one-way	relation	total
Lifted NULLCODEWORD	[YZ24a, GPW20]	two-way	two-way	relation	total
Bipartite NULLCODEWORD	This work	SMP	two-way	relation	total

Table 1: Several notable exponential quantum–classical separations. **Green text** indicates a strong result and **red text** indicates a weak result.

Relations vs. partial functions. Our separation for a total relation and the separation of [GRT22] for a partial function are formally incomparable. However, it is interesting to note that, in a different context, a separation for a total relation can be automatically translated to a separation for a partial function. Namely, Ben-David and Kundu [BK24b] show how to do this in the query complexity model by using a pointer-based construction. If we tried to implement a similar trick naively for our problem, we would lose the quantum easiness in the SMP model.

2 Technical overview

We construct a relation to prove [Theorem 1](#) in two stages. First, we define a natural bipartite analogue of the relation studied by Yamakawa–Zhandry which we call *Bipartite NullCodeword* (BiNC). We show in [Sections 4](#) and [5](#) that for BiNC there is an average case separation between QSMP and two-way randomized communication for a certain code and distribution. We then convert this into a worst-case separation for a total relation in [Section 6](#).

2.1 The Yamakawa–Zhandry problem and its bipartite variant

We build upon the Yamakawa–Zhandry query complexity problem [YZ24a], which we call the NULLCODEWORD (NC) problem. In this problem, an error correcting code $C \subseteq \Sigma^n$ is chosen, where Σ is an exponential-size alphabet, namely $|\Sigma| = 2^{\Theta(n)}$. The inputs are n random oracles $H_1, \dots, H_n: \Sigma \rightarrow \{0, 1\}$, and the goal is to find a codeword $x \in C$ such that $H(x) = 0^n$, where $H: \Sigma^n \rightarrow \{0, 1\}^n$ is the concatenation of all H_i s.

Lemma 2 (Informal, [YZ24a]). *When C is instantiated by a folded Reed–Solomon code with certain parameters, there exists a $\text{poly}(n)$ -query quantum algorithm (called the Yamakawa–Zhandry algorithm) solving NULLCODEWORD, but any classical algorithm requires $2^{n^{\Omega(1)}}$ queries.*

Our communication problem is a bipartite analogue of NULLCODEWORD which we call BiNC. In BiNC, we give the first half of the oracles $H_1, \dots, H_{n/2}$ as input to Alice, and the second half $H_{n/2+1}, \dots, H_n$ as input to Bob. Their goal is still to find a codeword $x \in C$ such that $H(x) = 0^n$. The input length for Alice and Bob is $N := n|\Sigma|/2 = 2^{\Theta(n)}$.

Quantum easiness of BiNC. One immediate advantage of our definition of BiNC is that it is fairly easy in the quantum communication setting. The key observation is that the Yamakawa–Zhandry

algorithm for `NULLCODEWORD` can be divided into two stages—it first makes *non-adaptive queries* to each H_1, \dots, H_n individually, and then *processes* the results without making any more queries. Therefore, we can get a QSMP protocol of `BiNC` by simulating the Yamakawa–Zhandry algorithm: Alice and Bob first simulate the non-adaptive queries to their halves of inputs and send their results to Charlie; Charlie then simulates the processing stage of the Yamakawa–Zhandry algorithm and outputs the solution. Furthermore, this QSMP protocol is computationally efficient and uses no shared entanglement/randomness.

2.2 Classical lower bound for `BiNC`

Proving a classical communication lower bound for `BiNC` turns out to be a more challenging task. First, most of the communication lower bound techniques we currently have only work for boolean functions (or relations with few or unique solutions), but `BiNC` has exponentially many solutions for a typical random input. Furthermore, the choice of the error correcting code C —which is a folded Reed–Solomon code—complicates the analysis for proving a classical lower bound. For example, information complexity seems like a viable approach if C is a random code. But the analysis would become intimidating if a folded Reed–Solomon code is considered.

Below, we give an overview of our proof. For $X \subseteq \{0, 1\}^N$ and $I \subseteq [N]$, we use $X_I := \{x_I \in \{0, 1\}^I : x \in X\}$ to denote the set of strings in X restricted to the coordinates in I .

Simple special case: Subcube protocols. We start by proving a lower bound for highly structured protocols. A rectangle $X \times Y \subseteq \{0, 1\}^N \times \{0, 1\}^N$ is said to be a subcube rectangle if there exist two index sets $I, J \subseteq [N]$ such that X_I and Y_J are *fixed* strings and the sets $X_{\bar{I}} = \{0, 1\}^{\bar{I}}$ and $Y_{\bar{J}} = \{0, 1\}^{\bar{J}}$ contain all possible strings. A subcube protocol is such that every node of the protocol tree corresponds to a subcube rectangle. This model is closely related to decision trees (but strictly more general). The below property of the code C will be central to the lower bound.

Definition 3 (List-recoverability: Simplified definition). A code $C \subseteq \Sigma^n$ is *list-recoverable* if for any collection of subsets $S_1, \dots, S_n \subseteq \Sigma$ such that $\sum_i |S_i| \leq \ell$, it holds that

$$|\{(x_1, \dots, x_n) \in C : |i \in [n] : x_i \in S_i| \geq 0.4n\}| \leq 2^{o(n)}. \quad (1)$$

Suppose Π is a subcube protocol that communicates at most $\ell \leq 2^{o(n)}$ bits. Assume also for simplicity that all rectangles in the subcube protocol fix at most ℓ coordinates. It follows from (1) that for any rectangle R in the protocol tree of Π , the number of codewords $x \in C$ such that more than $0.4n$ bits of $H(x)$ are fixed in R is at most $2^{o(n)}$.

Let us now sketch how to use this property to bound the success probability of Π . Say that a codeword $x \in C$ is *dangerous* for a rectangle R if at least $0.4n$ of the bits of x are fixed by R ; intuitively, this means that the protocol has communicated a lot of information about x . Consider any leaf rectangle R and suppose it is labeled with solution $x \in C$. Assuming for contradiction that Π never errs, then x is a correct solution (all-0 codeword) for all inputs $H \in R$; in particular, x is dangerous for R . We have two key facts:

- *Few dangerous codewords:* By list-recoverability (1), the number of codewords x that ever become dangerous in a single execution of Π is at most $2^{o(n)}$.
- *Dangerous codewords rarely become solutions:* Consider a single execution of the protocol and suppose Alice sends a message that causes a given codeword x to become dangerous for the first time at rectangle R . Then there are at least $0.1n$ unfixed input bits of x on Bob’s side. Thus, at most $2^{-0.1n}$ fraction of inputs in R would map x to 0^n .

We can conclude our lower bound for subcube protocols by a union-bound argument: None of the dangerous codewords will likely be a valid solution at the leaf and hence Π errs often.

Structure-vs-randomness dichotomy. To prove a lower bound for general protocols, we use the structure-vs-randomness framework developed by Wang, Yang, and Zhang [YZ24b, WYZ23] who build upon query-to-communication lifting techniques [GPW20]. We use this framework to convert any general communication protocol into a *subcube-like protocol*. The nodes of a subcube-like protocol correspond to *subcube-like rectangles* $X \times Y \subseteq \{0, 1\}^N \times \{0, 1\}^N$ defined such that

1. X_I, Y_J are *fixed* strings for some index sets $I, J \subseteq [N]$. (structure)
2. $X_{\bar{I}}, Y_{\bar{J}}$ are pseudorandom (they “look” like the uniform distribution). (randomness)

We show that subcube-like protocols behave similarly enough to subcube protocols that we may simply re-do the proof sketched above for subcube protocols.

Biased input distribution. So far, we have been ignoring a non-trivial issue—the list-recoverability property in Definition 3 required by our argument is not known to hold for the folded Reed–Solomon code C used for the original Yamakawa–Zhandry problem. To improve the parameters in the list-recoverability of the folded Reed–Solomon code, one could decrease its degree parameter k . However, a smaller degree k would also weaken the list-decodability of its dual code, which is essential for the quantum upper bound of the NULLCODEWORD problem. This inherent tension between the list-recoverability and the dual-decodability poses the following dilemma: if we decrease the degree k to an extent that list-recoverability is strong enough for our lower bound argument, the dual-decodability would be too weak to retain our quantum upper bound.

To resolve this, we tweak the input distribution of NULLCODEWORD and BINC. Recall that the NULLCODEWORD problem in [YZ24a] is defined over uniform random oracles. We now consider the scenario where the input bits can be *biased*, i.e., each symbol is (independently) mapped to 1 with probability p for a constant $p \in (0, 1)$. When p gets smaller, it becomes easier to find an all-0 codeword. By generalising the original analysis in [YZ24a], we show that a weaker dual-decodability suffices for the quantum upper bound when p is a small constant. Therefore, we can decrease the degree k of the code to achieve the list-recoverability parameters we need for our lower bound analysis while maintaining the quantum upper bound under the biased input distribution.

It might be of independent interest that our analysis broadens the range of parameters for which there is an exponential quantum advantage in the Yamakawa–Zhandry [YZ24a] problem.

Conversion to a total problem. Our discussion so far has been about an average case separation of quantum and classical communication complexity. Our input distribution has a non-zero probability of having *no solution*. Thus, we do not have totality yet. To fix this, we borrow a trick from the original [YZ24a] paper. The idea begins with the observation that an $O(n)$ -wise independent distribution suffices for the quantum upper bound. Thus, instead of a uniform random input we could give arbitrary inputs and allow Alice and Bob to XOR their inputs with a string that they can choose from an $O(n)$ -wise independent hash family. The canonical family is that of low-degree polynomials over finite fields [Vad12], which can be efficiently implemented.

However, modifying the problem in this way also allows for a *classical* upper bound! The protocol can simply pick an arbitrary codeword $x \in C$ and run Gaussian elimination to find a function h from the family which maps the codeword x to 0^n . To get around this, we further modify the problem and require the protocol to solve t instances of BINC that are XORed with the *same* hash function.

To prove this lower bound, we first apply a direct product theorem for product distributions [BRWY13] and show that for each fixed choice of hash function, the success probability of any sub-exponential cost classical protocol is less than $\approx 2^{-t}$. Thus, by taking t as a large polynomial in n , we can use a *union bound* over all hash function from the family and show that the success probability is still $o(1)$ even when the protocol is given the freedom of picking a hash function.

3 Preliminaries

In this work, we study *search* problems, formalized by relations. We introduce terminology below.

A search problem R is a relation $R \subseteq [k]^N \times O$ where for any given $x \in [k]^N$ we want to find $y \in O$ such that $(x, y) \in R$. A two-party search problem R is a bipartite relation $R \subseteq [k]^N \times [k]^N \times O$ where given $x, y \in [k]^N$ the two parties want to find a z such that $(x, y, z) \in R$. When the two-party search problem is total and efficiently verifiable, it is said to be in communication-TFNP (see [GKRS19, dRGR22] for a discussion of this class).

3.1 Communication complexity

We assume some familiarity with quantum computing and communication complexity. For more background refer to textbooks [NC00, KN97, RY20]. We define QSMP and two-way randomized communication protocols below.

Definition 4 (QSMP). A QSMP (without shared randomness/entanglement) protocol Π for relation $R \subseteq X \times Y \times O$ is pair of quantum states $|\phi_x\rangle, |\phi_y\rangle$ for every input $(x, y) \in X \times Y$. Alice and Bob send $|\phi_x\rangle$ and $|\phi_y\rangle$ respectively to Charlie, and he outputs a solution $\Pi(x, y) = \mathbf{o}_{x,y}$. The protocol is said to succeed with probability $1 - \epsilon$ if the probability that $(x, y, \mathbf{o}_{x,y}) \notin R$ is at most ϵ . The communication cost (denoted $|\Pi|$) is the total number of qubits sent by Alice and Bob on the worst-case input.

Definition 5 (Two-way randomized). A (public-coin) two-way randomized protocol Π is a sequence of messages exchanged between Alice and Bob which depend on their input, previous messages, and publicly sampled randomness after which Bob outputs $\Pi(x, y) = \mathbf{o}_{x,y}$. The protocol is said to succeed with probability $1 - \epsilon$ if the probability that $(x, y, \mathbf{o}_{x,y}) \notin R$ is at most ϵ . The communication cost (denoted $|\Pi|$) is the total number of bits sent by Alice and Bob on the worst-case input.

3.2 Error-correcting codes

For a prime power q , we denote by \mathbb{F}_q the finite field of order q and denote by $\mathbb{F}_q^{\leq k}[x]$ the set of polynomials over \mathbb{F}_q with degree at most k . For any vector x , define $\text{hw}(x)$ as the Hamming weight of x , i.e., the number of non-zero elements in x .

An error correcting code C of length N and distance d over the alphabet Σ is a subset of Σ^N such that the Hamming distance of every $c_1, c_2 \in C$ is lower bounded by d . We will use the notation N or n for the code length in order to distinguish with the input length N of the communication problem.

We will focus on linear codes, which are codes that are linear subspaces.

Definition 6 (Linear code and its dual code). A code $C \subseteq \mathbb{F}_q^N$ is a linear code if C is a linear subspace of \mathbb{F}_q^N . The dual code C^\perp of a linear code C is defined as its orthogonal complement:

$$C^\perp = \{x \in \mathbb{F}_q^N : x \cdot y = 0, \forall y \in C\}.$$

Definition 7 (Generalized Reed–Solomon code). The Reed–Solomon code $\text{RS}_{\mathbb{F}_q, \gamma, k}$ over \mathbb{F}_q with generator $\gamma \in \mathbb{F}_q^*$ and degree $0 \leq k \leq N$ has length $N := q - 1$, and is defined as:

$$\text{RS}_{\mathbb{F}_q, \gamma, k} := \{(f(0), f(\gamma), \dots, f(\gamma^{N-1})) : f \in \mathbb{F}_q^{\leq k}[x]\}.$$

For any vector $v = (v_0, \dots, v_{N-1}) \in \mathbb{F}_q^{*N}$, the generalized Reed–Solomon code $\text{GRS}_{\mathbb{F}_q, \gamma, k, v}$ is defined as:

$$\text{GRS}_{\mathbb{F}_q, \gamma, k, v} := \{(v_0 f(0), v_1 f(\gamma), \dots, v_{N-1} f(\gamma^{N-1})) : f \in \mathbb{F}_q^{\leq k}[x]\}.$$

The Reed–Solomon code is a special case of the generalized Reed–Solomon code with $v = (1, \dots, 1)$. The generalized Reed–Solomon code is a linear code, and the dual code of $\text{RS}_{\mathbb{F}_q, \gamma, k}$ is $\text{GRS}_{\mathbb{F}_q, \gamma, N-k-2, v}$ for some $v \in \mathbb{F}_q^N$.

Since multiplying a scalar on each coordinate does not affect the decodability, the classical list-decoding algorithm for the Reed–Solomon code also applies to the generalized Reed–Solomon code.

Lemma 8 (List-decodability of the GRS code [GS99]). *There is a deterministic list-decoding algorithm ListDecode for $\text{GRS}_{\mathbb{F}_q, \gamma, k, v}$ such that for any $z \in \mathbb{F}_q^N$, $\text{ListDecode}(z)$ returns a list of all codewords with hamming distance at most $N - \sqrt{kN}$ to z in $\text{poly}(N)$ time.*

Definition 9 (Folded linear codes). For a linear code $C \subseteq \mathbb{F}_q^N$ and an integer m that divides N , the m -folded code $C^{(m)}$ is a linear code over alphabet $\Sigma = \mathbb{F}_q^m$ of length $n := N/m$, defined as follow.

$$C^{(m)} = \{((x_1, \dots, x_m), (x_{m+1}, \dots, x_{2m}), \dots, (x_{N-m+1}, \dots, x_N)) : (x_1, \dots, x_N) \in C\}.$$

Note that $(C^\perp)^{(m)} = (C^{(m)})^\perp$. In particular, the dual code of folded Reed–Solomon code $\text{RS}_{\mathbb{F}_q, \gamma, k}^{(m)}$ is the folded generalized Reed–Solomon code $\text{GRS}_{\mathbb{F}_q, \gamma, d, N-k-2, v}^{(m)}$ for some $v \in \mathbb{F}_q^N$.

Definition 10 (List-recoverability). A code $C \subseteq \Sigma^n$ is (ζ, ℓ, L) list-recoverable if for any collection of subsets $S_1, \dots, S_n \subseteq \Sigma$ such that $|S_i| \leq \ell$ for any $i \in [n]$, it holds that

$$|\{(x_1, \dots, x_n) \in C : |i \in [n] : x_i \in S_i| \geq \zeta n\}| \leq L.$$

The folded Reed–Solomon code is known to be list-recoverable in certain parameter regimes.

Lemma 11 ([GR08, Rud07]). *For positive integers ℓ, r, s, m , where $s \leq m$, and a real number $\zeta \in (0, 1)$, the code $\text{RS}_{\mathbb{F}_q, \gamma, k}^{(m)}$ is (ζ, ℓ, q^s) list-recoverable if the following inequalities hold:*

$$\frac{\zeta N}{m} \geq \left(1 + \frac{s}{r}\right) \cdot \frac{(N\ell k^s)^{\frac{1}{s+1}}}{m - s + 1}. \quad (2)$$

$$(r + s) \cdot (N\ell/k)^{\frac{1}{s+1}} < q. \quad (3)$$

Quantum fourier transform over the code space. The quantum Fourier transform over the finite field \mathbb{F}_q where $q = p^r$ for a prime p , is a unitary defined as follows:

$$\text{QFT}_{\mathbb{F}_q} |x\rangle = \frac{1}{\sqrt{q}} \sum_{z \in \mathbb{F}_q} \omega_p^{\text{Tr}(x \cdot z)} |z\rangle,$$

where the trace function is defined as $\text{Tr}(x) := \sum_{i=0}^{r-1} x^{p^i}$ for any $x \in \mathbb{F}_q$.

The quantum Fourier transform over the alphabet $\Sigma = \mathbb{F}_q^m$ is defined as the m -tensor product of $\text{QFT}_{\mathbb{F}_q}$:

$$\text{QFT}_{\Sigma} |x\rangle := \text{QFT}_{\mathbb{F}_q}^{\otimes m} |x_1\rangle |x_2\rangle \dots |x_m\rangle = \frac{1}{\sqrt{|\Sigma|}} \sum_{z \in \Sigma} \omega_p^{\text{Tr}(x \cdot z)} |z\rangle.$$

Note that the second equality uses the linearity of $\text{Tr}(\cdot)$. Similarly, the quantum Fourier transform over the code space Σ^n is

$$\text{QFT} |x\rangle := \text{QFT}_{\Sigma}^{\otimes n} |x\rangle = \frac{1}{|\Sigma|^{n/2}} \sum_{z \in \Sigma^n} \omega_p^{\text{Tr}(x \cdot z)} |z\rangle.$$

For a quantum state $|\phi\rangle = \sum_x f(x) |x\rangle$, we use \hat{f} to denote its fourier coefficient, i.e.,

$$\text{QFT} |\phi\rangle = \sum_z \hat{f}(z) |z\rangle.$$

3.3 The Yamakawa–Zhandry problem

The Yamakawa–Zhandry problem is a search problem in which, given a fixed code and query access to random oracles, the goal is to find a codeword that is null for the oracles.

Definition 12 (NULLCODEWORD [YZ24a]). Fix a family of codes $C = \{C_n\}$, where $C_n \subseteq \Sigma^n$. For each $i \in [n]$, let $H_i: \Sigma \rightarrow \{0, 1\}$, and let $H: \Sigma^n \rightarrow \{0, 1\}^n$ be the concatenation $H(x) = (H_1(x_1), \dots, H_n(x_n))$. The relation $\text{NC}_n^C \subseteq \{0, 1\}^{n|\Sigma|} \times \Sigma^n$ is defined as:

$$\text{NC}_n^C = \{(H, x) : x \in C \wedge H(x) = 0^n\}.$$

The NULLCODEWORD problem NC^C is defined as the following search problem: Given n and the oracle access to $H \in \{0, 1\}^{n|\Sigma|}$, find a string x such that $(H, x) \in \text{NC}_n^C$.

A folded Reed–Solomon code C with certain parameters is chosen in [YZ24a], which allows NC^C to have a quantum algorithm with $\text{poly}(n)$ queries and at the same time be exponentially hard for any classical (randomized) algorithm, when the input H is given by a uniformly random oracle.

4 The Bipartite NullCodeword problem

We present the formal definition of the Bipartite NullCodeword communication problem (BiNC) in [Section 4.1](#). We then, in [Section 4.2](#), instantiate the problem with the folded Reed–Solomon code, with a careful parameterization that provides the properties needed for our techniques. Finally, we provide an efficient quantum SMP protocol for BiNC in [Section 4.3](#). Throughout this section, we will work with a biased-input distribution. In [Section 6](#), we will convert this into a total problem.

4.1 Definition of BiNC

To transform the NULLCODEWORD problem ([Definition 12](#)) into a communication problem, the mapping $H: \Sigma^n \rightarrow \{0, 1\}^n$ is divided into two halves $H_A, H_B: \Sigma^{n/2} \rightarrow \{0, 1\}^{n/2}$ as the input of Alice and Bob respectively.

Definition 13 (Bipartite NullCodeword, BiNC). Fix a family of codes $C = \{C_n\}$, where $C_n \subseteq \Sigma^n$. Let H_1, \dots, H_n be n mappings from Σ to $\{0, 1\}$, define $H_A, H_B: \Sigma^{n/2} \rightarrow \{0, 1\}^{n/2}$, where $H_A = (H_1, \dots, H_{n/2})$ and $H_B = (H_{n/2+1}, \dots, H_n)$. The bipartite relation $\text{BiNC}_n^C \subseteq \{0, 1\}^{n|\Sigma|/2} \times \Sigma^n$ is defined as:

$$\text{BiNC}_n^C = \{(H_A, H_B, x) : x \in C_n \wedge H_A(x_A) = H_B(x_B) = 0^{n/2}\},$$

where $x_A = (x_1, \dots, x_{n/2})$ and $x_B = (x_{n/2+1}, \dots, x_n)$.

The superscript C may be omitted in BiNC^C if the choice of code C is clear from the context. We will consider a distributional version of BiNC^C where the input is given by a *biased random oracle*.

Definition 14 (p -biased input distribution). A distribution over $H = (H_1, H_2, \dots, H_n)$ where each H_i maps every element of Σ to 1 with probability p independently.

We will pick the constant $p := 2^{-6}$ for our setting, though we note that for any constant value below $1/2$, we can get a non-trivial tradeoff and preserve the classical lower bound.

4.2 A suitable code: Trading off dual-decodability for list-recoverability

We modify the code used in the NULLCODEWORD problem [YZ24a] to achieve a stronger list-recoverability, which we will need to prove the classical two-way communication lower bound. We do this by reducing the *degree* of the polynomials in the construction of the Reed–Solomon code. This hampers the decodability of the dual code, but we show that over the biased input distribution (Definition 14), we can still decode the dual code with high probability, which is necessary for our upper bound. We instantiate the code we will use for the remainder of this paper below.

Definition 15. Define code $C = \{C_n\}$ as the folded Reed–Solomon codes $\text{RS}_{\mathbb{F}_q, \gamma, k}^{(m)}$ with the following choice of parameters:

- the length of the code $n = 2^t - 1$ for some $t \in \mathbb{N}$;
- $q = 2^{2t}$, and thus, $\mathbb{N} = q - 1$ and $m = \mathbb{N}/n = 2^t + 1$;
- γ is an arbitrary generator of \mathbb{F}_q^* ;
- degree $k = 0.1\mathbb{N}$.

By definition, $|C_n| = q^{k+1} = 2^{\Theta(n^2)}$ and $|\Sigma| = |\mathbb{F}_q^m| = q^m = 2^{\Theta(n)}$.

Only the degree parameter k is different compared with the original NULLCODEWORD problem [YZ24a], in which k is chosen to be $\alpha\mathbb{N}$ for some constant $\alpha \in (5/6, 1)$. Below, we state a lemma with the parameters for which our code is list-recoverable and dual-decodable over the p -biased input distribution.

Lemma 16 (Modified from Lemma 4.2 in [YZ24a]). *The code C satisfies:*

1. C_n is (ζ, ℓ, L) -list-recoverable, where $\zeta = 0.4, \ell = 2^{n^{0.8}}, L = 2^{n^{0.9}}$. (List-Recoverability)
2. Let \mathcal{D}_n be the product distribution on Σ^n where each symbol is 0 with probability $1 - p$ and otherwise a uniformly random element of $\Sigma \setminus \{0\}$. If $p \leq 0.04$, there exists a poly(n) time deterministic decoding algorithm Decode for the dual code C_n^\perp such that

$$\Pr_{e \sim \mathcal{D}_n} [\forall x \in C_n^\perp, \text{Decode}(x + e) = x] = 1 - 2^{-\Omega(n)}. \quad (4)$$

(Unique-Decodability of the Dual)

Proof. The proof of (1), the list-recoverability property, follows immediately from [Lemma 11](#). Namely, we apply [Lemma 11](#) with $s := n^{0.9}, r := n^{0.95}$. It is straightforward to verify that both inequalities (2) and (3) are satisfied with the parameters specified in [Definition 15](#).

It remains to show the unique decodability of the dual code. Note that the dual code C_n^\perp is the folded generalized Reed–Solomon code $\text{GRS}_{\mathbb{F}_q, \gamma, d, v}^{(m)}$ for some $v \in \mathbb{F}_q^{\mathbb{N}}$, where $d := \mathbb{N} - k - 2 = 0.9\mathbb{N} - 2$. Let $\varepsilon := 0.01$. For any vector $x \in \mathbb{F}_q^{\mathbb{N}}$, let $x^f \in \Sigma^n$ be the m -folded version of x ; similarly, for any vector $z \in \Sigma^n$, define $z^u \in \mathbb{F}_q^{\mathbb{N}}$ by unfolding each letter in z . So, if $z = x_1, \dots, x_m \in \mathbb{F}_q^n$ then $z^u = (x_1, \dots, x_m) \in \mathbb{F}_q^{\mathbb{N}}$.

The algorithm $\text{Decode}(z)$ works as follows: On input $z \in \Sigma^n$, it first unfolds z into $z^u \in \mathbb{F}_q^{\mathbb{N}}$. It then runs the list-decoding algorithm $\text{ListDecode}(z^u)$ ([Lemma 8](#)) for the generalized Reed–Solomon code $\text{GRS}_{\mathbb{F}_q, \gamma, d, v}$ to get a list of codewords. If there is a unique $x \in \mathbb{F}_q^{\mathbb{N}}$ in the list such that $\text{hw}(z^u - x) \leq (p + \varepsilon)\mathbb{N}$, it outputs x^f (interpreted as a codeword of $\text{GRS}_{\mathbb{F}_q, \gamma, \mathbb{N} - k - 2, v}^{(m)}$), and otherwise outputs \perp .

We now analyze the success probability of Decode . Define a subset $\mathcal{G} \subseteq \Sigma^n$ of *good* error terms, leading to unique decoding, as follows.

$$\mathcal{G} := \{e \in \Sigma^n : \text{hw}(e) \leq (p + \varepsilon)n\}.$$

Applying a Chernoff bound to the Hamming weight of e , we know e is good with high probability.

$$\Pr_{e \sim \mathcal{D}_n}[e \in \mathcal{G}] \geq 1 - 2^{-\Omega(n)}. \quad (5)$$

Therefore, it suffices to prove that the algorithm Decode always succeeds on a good error. Towards this end, we first show a useful property of set \mathcal{G} .

Claim 17. *For any $e \in \mathcal{G}$ and $y \in C_n^\perp \setminus \{0\}$, we have $\text{hw}(e^u - y^u) > (p + \varepsilon)\mathbb{N}$.*

Proof. For any $e \in \mathcal{G}$, $\text{hw}(e^u) \leq (p + \varepsilon)\mathbb{N} < 0.05\mathbb{N}$. Note that for any non-zero codeword $y \in C_n^\perp$, it holds that $\text{hw}(y^u) \geq \mathbb{N} - d \geq 0.1\mathbb{N}$, because a degree d non-zero polynomial could have at most d roots. By the triangle inequality, we have $\text{hw}(e^u - y^u) > 0.05\mathbb{N} \geq (p + \varepsilon)\mathbb{N}$. \square

Using [Claim 17](#), we proceed to show the correctness of the algorithm Decode .

Claim 18. *For any $x \in C_n^\perp$ and $e \in \mathcal{G}$, $\text{Decode}(x + e) = x$.*

Proof. Fix any $x \in C_n^\perp$ and $e \in \mathcal{G}$. Recall that $d = 0.9\mathbb{N} - 2$, it holds that $\mathbb{N} - \sqrt{d\mathbb{N}} > 0.05\mathbb{N} \geq (p + \varepsilon)\mathbb{N}$. By [Lemma 8](#), x^u is contained in the list of codewords outputted by the list-decoding algorithm for $\text{GRS}_{\mathbb{F}_q, \gamma, d, v}$. For the uniqueness, consider any $y \in C_n^\perp, y \neq x$, since C_n^\perp is a linear code, we have

$$(x + e) - y = e - (y - x) = e - y'$$

for some $y' \in C_n^\perp \setminus \{0\}$. By [Claim 17](#), we have that $\text{hw}((x + e)^u - y^u) = \text{hw}(e^u - (y')^u) > (p + \varepsilon)\mathbb{N}$. Thus, x^u is the only codeword in the list whose Hamming distance from $(x + e)^u$ is smaller than or equal to $(p + \varepsilon)\mathbb{N}$. We conclude that $\text{Decode}(x + e) = x$. \square

The proof of the second part of [Lemma 16](#) then follows from [Claim 18](#) and inequality (5). \square

Remark. In the following subsection and [Section 6](#), we will denote BiNC to be the problem with C fixed to be the code in [Definition 15](#), dropping the superscript.

4.3 Quantum SMP easiness of BiNC

We prove the quantum easiness of BiNC by adapting the query algorithm for the NULLCODEWORD problem to the biased bipartite setting, noting that the protocol can be implemented under the constraints of the quantum SMP model.

Theorem 19. *BiNC has a quantum SMP protocol using $\text{poly}(n)$ qubits over the p -biased distribution.*

Algorithm 1 Quantum SMP Protocol for BiNC

Notation:

For each $i \in [n]$, define the state

$$|\phi_i\rangle \propto \sum_{e \in \Sigma: H_i(e)=0} |e\rangle.$$

Let $U_{\text{add}}, U_{\text{decode}}$ be unitaries defined as follows:

$$|x\rangle |e\rangle \xrightarrow{U_{\text{add}}} |x\rangle |x+e\rangle \xrightarrow{U_{\text{decode}}} |x - \text{Decode}(x+e)\rangle |x+e\rangle.$$

Alice:

Prepare $|\phi_1\rangle, \dots, |\phi_{n/2}\rangle$ using $H_1, \dots, H_{n/2}$.

Send $|\phi_1\rangle, \dots, |\phi_{n/2}\rangle$ to Charlie.

Bob:

Prepare $|\phi_{n/2+1}\rangle, \dots, |\phi_n\rangle$ using $H_{n/2+1}, \dots, H_n$.

Send $|\phi_{n/2+1}\rangle, \dots, |\phi_n\rangle$ to Charlie.

Charlie:

After receiving $|\phi\rangle := |\phi_1\rangle \otimes \dots \otimes |\phi_n\rangle$, perform the following steps:

1. Prepare a state:

$$|\psi\rangle \propto \sum_{x \in \mathcal{C}} |x\rangle.$$

2. Apply QFT to both $|\psi\rangle$ and $|\phi\rangle$. Now, Charlie holds the state:

$$|\eta\rangle := \text{QFT} |\psi\rangle \otimes \text{QFT} |\phi\rangle.$$

3. Apply $(I \otimes \text{QFT}^{-1})U_{\text{decode}}U_{\text{add}}$ to $|\eta\rangle$.
 4. Measure the second register and output the measurement result.
-

We provide a proof sketch here and refer the reader to [Appendix A](#) for the full proof of the correctness of our quantum SMP protocol, as captured in [Theorem 19](#).

Proof Sketch. As pointed out by Liu [[Liu23](#)], the original NULLCODEWORD problem can be solved by a two-stage quantum algorithm \mathcal{A}^{YZ} :

Query Make $\text{poly}(n)$ queries to each of H_1, \dots, H_n , and store the raw results from those queries into states $|\phi_1\rangle, \dots, |\phi_n\rangle$, where $|\phi_i\rangle$ stores the query results from H_i .

Process Process $|\phi_1\rangle, \dots, |\phi_n\rangle$ and output the solution without making any more queries to H .

Our quantum SMP protocol for BiNC^C emulates the query protocol as follows. Alice and Bob first prepare states $|\phi_1\rangle, \dots, |\phi_{n/2}\rangle$ and $|\phi_{n/2+1}\rangle, \dots, |\phi_n\rangle$ by simulating the **Query** stage of \mathcal{A}^{YZ} on their half of inputs, and then send them to Charlie. Then, Charlie runs the same **Process** stage of \mathcal{A}^{YZ} with states $|\phi_1\rangle, \dots, |\phi_n\rangle$ and outputs the solution. See [Algorithm 1](#) for the quantum SMP protocol in detail.

It remains to argue that the algorithm \mathcal{A}^{YZ} still works for NC^C after the following two changes:

1. The code C specified in [Definition 15](#) has a lower degree.
2. Each H_i is drawn from a p -biased distribution outputting 1 with probability $p = 2^{-6}$ rather than being uniformly random.

As shown in [Lemma 16](#), the first change weakens the unique-decodability of the dual code C^\perp . In particular, we can only guarantee the unique-decodability against a milder distribution \mathcal{D}_n of errors, where each location is scrambled with probability $p = 2^{-6} \leq 0.04$. However, in the analysis of the original Yamakawa–Zhandry problem, the dual code needs to deal with the distribution of errors where each location is scrambled with probability $1/2$.

Fortunately, the second change fixes this problem once the biased input distribution is plugged into the analysis. Here, we briefly sketch the idea: Recall that the state $|\hat{\phi}_i\rangle = \text{QFT} |\phi_i\rangle$ is generated in the second step of Charlie for each $i \in [n]$. In expectation over H_i drawn from the biased distribution, the state $|\hat{\phi}_i\rangle$ has weight $1 - p$ on 0, and the remaining weight of p is uniformly distributed over $\Sigma \setminus \{0\}$. Therefore, at a high level, before Charlie applies U_{decode} in the third step, one can think of the error e is now drawn from the distribution \mathcal{D}_n defined in [Lemma 16](#), which can be uniquely decoded with high probability by [Lemma 16](#).

The remaining analysis in [\[YZ24a\]](#) is not affected by these two changes. This concludes the sketch of proof of correctness of our quantum SMP protocol. \square

5 Classical lower bound for BiNC

In this section, we prove the distributional lower bound for the BiNC problem.

Theorem 20. *Let $C \subseteq \Sigma^n$ be a (ζ, ℓ, L) -list-recoverable code with $\zeta \leq 0.4$ and $L = 2^{o(n)}$. Then BiNC_n^C has randomized communication complexity $\Omega(\ell)$ over the p -biased input distribution ([Definition 14](#)) for any constant $p \in (0, 1)$.*

In particular, we can later instantiate [Theorem 20](#) with the code from [Lemma 16](#) (for which we proved the quantum upper bound). We present the proof of [Theorem 20](#) in stages. We first ([Section 5.1](#)) prove the theorem for highly structured communication protocols that we call *subcube protocols*, which are closely related to decision trees. We then ([Section 5.2](#)) generalize that proof to a more powerful class of *subcube-like protocols*. Finally ([Section 5.3](#)), we use structure-vs-randomness paradigm [[GPW20](#), [YZ24b](#)] to show that any communication protocol can be converted (with modest loss) into a subcube-like protocol.

Input distribution: Simplification. Recall that the p -biased input distribution assigns all bits *i.i.d.* such that the probability of a 1-bit is p . We first prove [Theorem 20](#) for the uniform distribution, $p = 1/2$, and then indicate ([Section 5.4](#)) how to generalize it to any $p \in (0, 1)$.

5.1 Lower bound for subcube protocols

We start by defining the notion of subcube protocols. We study exclusively *deterministic* protocols—run on a random distribution over inputs—as this is enough to prove a lower bound against randomised protocols (by Yao’s principle).

Definition 21 (Subcubes). A set $X \subseteq \{0, 1\}^N$ is a *subcube* if there exists $I \subseteq [N]$ such that $X = \{x \in \{0, 1\}^N : x_I = a\}$ for some partial assignment $a \in \{0, 1\}^I$. The *codimension* of X is defined as $|I|$. Moreover, if $i \in I$ then we say the corresponding variable x_i is *fixed* in X .

Definition 22 (Subcube protocols). A rectangle $R = X \times Y \subseteq \{0, 1\}^N \times \{0, 1\}^N$ is said to be a *subcube rectangle* if X and Y are subcubes. The codimension of R is then the sum of codimensions of X and Y . A communication protocol Π is said to be a *subcube protocol* if for every node in the protocol tree of Π , the associated rectangle is a subcube rectangle. We denote by $|\Pi|$ the worst-case communication cost of Π (maximum length of the transcript).

The simplest example of a subcube protocol is when the players send individual input bits to each other. For example, Alice can send $x_1 \in \{0, 1\}$ to Bob, which corresponds to splitting Alice’s input domain $\{0, 1\}^n$ to two subcubes where $x_1 = 1$ or $x_1 = 0$. Subcube protocols that send one bit at a time correspond precisely to decision trees.

More generally, in each communication round of a subcube protocol, the length of the message can be larger than 1 bit if v has large arity (number of its children). The correspondence between subcube protocols and decision trees breaks when the protocol tree has nodes with arity larger than 2. For example, suppose Alice on input $x \in \{0, 1\}^N$ sends the index $i \in [N]$ of her first 1-bit (or indicates her input is all-0). This is a $\log(N + 1)$ -bit message that induces a partition of $\{0, 1\}^N$ into $N + 1$ subcubes. Such a protocol is not efficiently simulated by a shallow subcube protocol sending one bit at a time, that is, a decision tree. (See [AKK16] for more separations between subcube partitions and decision trees.)

Cleanup. We first show that any subcube protocol for BiNC_n^C (or for any communication-TFNP problem) can be converted into a *zero-error* protocol with bounded codimension at the leaves.

Lemma 23 (Subcube protocol cleanup). *Suppose Π is a subcube protocol solving BiNC_n^C with error $\epsilon > 0$ over the uniform distribution. Then there exists another subcube protocol Π' such that*

- (C1) $|\Pi'| \leq O(|\Pi|)$.
- (C2) Π' is zero-error: on a uniform random input, it outputs the symbol \perp (“I don’t know”) with probability at most 2ϵ , and otherwise it outputs a valid solution.
- (C3) Every leaf rectangle of Π' has codimension $O(|\Pi|/\epsilon)$.

Proof. Since Π communicates $|\Pi|$ bits, the protocol tree has at most $2^{|\Pi|}$ leaves. The leaves partition the input domain into at most $2^{|\Pi|}$ rectangles. Denote by \mathbf{R} the rectangle of the leaf reached when Π is run on a uniform random input. A fixed leaf rectangle $R \subseteq \{0, 1\}^N \times \{0, 1\}^N$ is reached with probability $\Pr[\mathbf{R} = R] = |R|/2^{2N} \leq 2^{-\text{codim}(R)}$. Letting $\mathbf{H}(\cdot)$ denote Shannon entropy, we have

$$\mathbb{E}[\text{codim}(\mathbf{R})] \leq \mathbb{E}[\log(1/\Pr[\mathbf{R} = R])] = \mathbf{H}(\mathbf{R}) \leq |\Pi|.$$

By Markov’s inequality, $\Pr[\text{codim}(\mathbf{R}) \leq |\Pi|/\epsilon] \geq 1 - \epsilon$. If we reach a node fixing more than $|\Pi|/\epsilon$ indices, we modify Π to output \perp . This produces a protocol Π' satisfying (C1) and (C3).

Finally, we assume that the protocol Π ends with Bob outputting the solution. Alice can then check if his solution is valid for her input, and send a 1 if yes. If she sends a 0, Bob outputs \perp . This

means our protocol can be converted into one which is zero-error via this round of interaction which verifies the solution outputted by Bob. In the modified protocol, the success probability is the same and we can assume Π always outputs a correct solution or \perp . The total probability of outputting \perp is the error probability of the original protocol plus the probability of the leaf codimension being greater than $|\Pi|/\epsilon$. Thus, the probability is 2ϵ , which satisfies (C2). \square

Proof of Theorem 20 for subcube protocols. By Lemma 23 we may assume w.l.o.g. that Π is zero-error and has bounded codimension at the leaves. For the sake of contradiction, assume that $|\Pi| \leq o(\ell)$. Sample a uniform random input $\mathbf{H} \sim \{0, 1\}^N \times \{0, 1\}^N$ and run the protocol to get a random transcript $\Pi(\mathbf{H})$. In this run, we encounter a sequence of nodes $\mathbf{v}_1, \dots, \mathbf{v}_d$ of the protocol tree of Π where \mathbf{v}_1 is the root node, \mathbf{v}_d is a leaf, and d is the number of communication rounds.

We call a codeword $c \in C$ *dangerous* in a partial assignment $\rho \in \{0, 1, *\}^N \times \{0, 1, *\}^N$ if at least $0.4n$ bits of c are fixed by ρ . In communication round i , we denote by $\mathbf{Q}_i \subseteq C$ the set of codewords that are dangerous for the partial assignment corresponding to the subcube rectangle at node \mathbf{v}_i . We observe three properties of $\mathbf{Q} := (\mathbf{Q}_1, \dots, \mathbf{Q}_d)$:

- (i) \mathbf{Q} is monotone: $\mathbf{Q}_i \subseteq \mathbf{Q}_{i+1}$ for all i .
- (ii) If $\Pi(\mathbf{H})$ outputs $c \neq \perp$, then $c \in \mathbf{Q}_d$.
- (iii) $|\mathbf{Q}_d| \leq L$.

Here (i) is immediate. Item (ii) follows because Π is zero-error. To see (iii), note that since $|\Pi| \leq o(\ell)$ we have from Lemma 23 that at most $O(|\Pi|) \leq o(\ell)$ bits of \mathbf{H} are fixed. By the $(0.4, \ell, L)$ -list-recoverability of the code C , we obtain (iii). Using (iii) we can write $\mathbf{Q}_d = \{\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_L\}$ where every $\mathbf{c}_i \in C$ is a random variable.

Claim 24. *For every $j \in [L]$ we have $\Pr[\mathbf{H}(\mathbf{c}_j) = 0^n] \leq 2^{-\Omega(n)}$.*

Proof. Fix $j \in [L]$ and consider the first node $\mathbf{v} \in \{\mathbf{v}_1, \dots, \mathbf{v}_d\}$ where \mathbf{c}_j becomes dangerous. That is, $\mathbf{v} = \mathbf{v}_i$ such that $\mathbf{c}_j \in \mathbf{Q}_i$ but $\mathbf{c}_j \notin \mathbf{Q}_{i-1}$. Let R_v denote the rectangle corresponding to node v . Note that for $v \in \text{supp}(\mathbf{v})$, the events “ $\mathbf{v} = v$ ” and “ $\mathbf{H} \in R_v$ ” are the same: monotonicity (i) states that once a codeword $c \in C$ becomes dangerous at node v , it will stay dangerous for all leaves in v ’s subtree. Suppose Alice sent a message at \mathbf{v} ’s parent \mathbf{v}_{i-1} . Since $\mathbf{c}_j \notin \mathbf{Q}_{i-1}$, Bob has at most $0.4n$ bits of \mathbf{c}_j fixed at $R_{\mathbf{v}_{i-1}}$, and hence also at R_v . Because of these $0.1n$ unfixed bits of \mathbf{c}_j on Bob’s side, we have

$$\Pr[\mathbf{H}(\mathbf{c}_j) = 0^n \mid \mathbf{v} = v] = \Pr[\mathbf{H}(\mathbf{c}_j) = 0^n \mid \mathbf{H} \in R_v] \leq 2^{-0.1n} \leq 2^{-\Omega(n)}. \quad (6)$$

We now calculate

$$\begin{aligned} \Pr[\mathbf{H}(\mathbf{c}_j) = 0^n] &= \sum_v \Pr[\mathbf{H}(\mathbf{c}_j) = 0^n \mid \mathbf{v} = v] \cdot \Pr[\mathbf{v} = v] \\ &\leq 2^{-\Omega(n)} \sum_v \Pr[\mathbf{v} = v] && \text{(by (6))} \\ &= 2^{-\Omega(n)}. && \square \end{aligned}$$

Using [Claim 24](#), we can now show that Π errs with high probability, which finishes the proof.

$$\begin{aligned}
\Pr[\Pi(\mathbf{H}) \text{ is correct}] &\leq \Pr[\exists c \in \mathbf{Q}_d: \mathbf{H}(c) = 0^n] && \text{(by (ii))} \\
&\leq \sum_{j \in [L]} \Pr[\mathbf{H}(c_j) = 0^n] && \text{(union bound)} \\
&\leq L \cdot 2^{-\Omega(n)} && \text{(Claim 24)} \\
&\leq 2^{o(n)} \cdot 2^{-\Omega(n)} \leq o(1). && \square
\end{aligned}$$

5.2 Subcube-like protocols

We will now prove [Theorem 20](#) for a more general class of *subcube-like protocols*. To set this up, we introduce the notion of *dense* random variables, which is a notion of pseudorandomness first considered in [\[GLM⁺16\]](#). We first recall the definition of min-entropy.

Definition 25 (Min-entropy). Given a random variable $\mathbf{X} \in \mathcal{X}$, its min-entropy is defined to be

$$\mathbf{H}_\infty(\mathbf{X}) := \log \frac{1}{\max_{x \in \mathcal{X}} \Pr[\mathbf{X} = x]}$$

For example, note that $\mathbf{H}_\infty(\mathbf{X}) \geq t$ iff we have $\Pr[\mathbf{X} = x] \leq 2^{-t}$ for all x .

Definition 26. A random variable $\mathbf{X} \in \{0, 1\}^N$ is said to be γ -dense if for every $I \subseteq [N]$, we have

$$\mathbf{H}_\infty(\mathbf{X}_I) \geq \gamma|I|$$

Definition 27 (Subcube-like rectangle). A rectangle $R = X \times Y$ is said to be γ -subcube-like with respect to (I, J) if (here we denote $\bar{I} := [N] \setminus I$)

- X_I and Y_J are fixed strings.
- $\mathbf{X}_{\bar{I}}$ and $\mathbf{Y}_{\bar{J}}$ are γ -dense where $\mathbf{X} \sim X$ and $\mathbf{Y} \sim Y$ are sampled uniformly at random.

Similar to subcubes, the codimension of a subcube-like rectangle w.r.t. (I, J) is said to be $|I| + |J|$. A communication protocol Π is said to be a *subcube-like protocol* if for every node in the protocol tree, the associated rectangle is subcube-like. We can now re-prove the cleanup lemma ([Lemma 23](#)) for subcube-like protocols; the original proof works verbatim.

Lemma 28 (Subcube-like protocol cleanup). *Suppose Π is a γ -subcube-like protocol solving $\text{BIN}C_n^C$ with error $\epsilon > 0$ over the uniform distribution. Then there exist another γ -subcube-like protocol Π' such that [Items \(C1\) to \(C3\)](#) are satisfied.* \square

We are now ready to prove that [Theorem 20](#) holds for 0.8-subcube-like protocols.

Proof of [Theorem 20](#) for subcube-like protocols. The argument is nearly identical to that for subcube protocols. Define $\mathbf{Q} := (\mathbf{Q}_1, \dots, \mathbf{Q}_d)$ as before for the candidate subcube-like protocol Π . The only crucial difference is in [Claim 24](#): We show that it continues to hold (albeit with a smaller constant hidden by the Ω -notation). Indeed, if \mathbf{c}_j becomes dangerous for the first time at $\mathbf{v} = v$, then Bob (or Alice) has at least $0.1n$ unfixed bits of \mathbf{c}_j in his input set Y where $R_v = X \times Y$. These unfixed bits come from a 0.8-dense distribution $\mathbf{Y} \sim Y$ where the min-entropy of the marginal distribution of these unfixed bits is at least $0.8 \cdot 0.1n = 0.08n$. This implies $\Pr[\mathbf{H}(\mathbf{c}_j) = 0^n \mid \mathbf{H} \in R_v] \leq 2^{-0.08n} \leq 2^{-\Omega(n)}$ showing that the bound [\(6\)](#) continues to hold, which re-proves [Claim 24](#). We can finally conclude that $\Pr[\Pi(\mathbf{H}) \text{ is correct}] = o(1)$ using the same calculation as before. \square

5.3 From general protocols to subcube-like protocols

We will now prove [Theorem 20](#) for general protocols. The proof is by a reduction to the case of subcube-like protocols using the following theorem.

Theorem 29 (Protocol preprocessing). *Suppose Π is a protocol solving some communication problem with error $\epsilon > 0$ over the uniform distribution. Then for any constant $\gamma < 1$ there exist a γ -subcube-like protocol Π' such that*

1. $|\Pi'| \leq O(|\Pi|/\epsilon)$
2. Π' has error 2ϵ over the uniform distribution.

Simply put, [Theorem 29](#) states that any protocol can be assumed to be subcube-like without loss of generality. A similar theorem was proved in [[YZ24b](#), [WYZ23](#)], but instead of a bound on the communication cost ([Item 1](#)), their version produced a subcube-like protocol with a bound on the codimension of the leafs. This is a slightly weaker guarantee than [Item 1](#): Recall from [Lemma 23\(C3\)](#) that bounded communication cost implies a bound on the leaf codimension. We prefer our version in [Theorem 29](#) as it is more elegant and potentially easier to use in subsequent work.

Proof sketch. [Theorem 29](#) is only a slight strengthening of the techniques in [[GPW20](#), [YZ24b](#), [WYZ23](#)] and so we are content with only sketching the crucial differences. To begin, we assume the protocol Π is in the normal form where the players send one bit in each communication round. Protocol Π' is now defined by [Algorithm 2](#). The only difference to [[YZ24b](#), [WYZ23](#)] is that we send the index i encoded using a Huffman code. That is, we use the following fact.

Fact 30 ([[Huf52](#)]). *Let \mathbf{k} be any random variable. There exist an encoding function $C(\cdot)$ (called the Huffman code) such that $\mathbb{E}[|C(\mathbf{k})|] \leq \mathbf{H}(\mathbf{k}) + 1$.*

Otherwise the algorithm is the one stated in [[YZ24b](#), [WYZ23](#)]. It crucially relies on the *density-restoring partition* from [[GPW20](#)], which states:

Lemma 31 (Density-restoring partition [[GPW20](#), [CFK⁺21](#)]). *Let $\mathbf{X} \in \{0, 1\}^N$ be a random variable and let $0 \leq \gamma < 1$ be a constant. Then there exists a partition*

$$\mathcal{X} := \mathcal{X}^1 \cup \dots \cup \mathcal{X}^\ell$$

where every \mathcal{X}^j is associated with a set $I_j \subseteq [n]$ and a value $x_j \in \{0, 1\}^{I_j}$ such that:

- $(\mathbf{X}_{I_j} | \mathbf{X} \in \mathcal{X}^j)$ is fixed to x_j .
- $(\mathbf{X}_{\bar{I}_j} | \mathbf{X} \in \mathcal{X}^j)$ is γ -dense.
- If we pick i at random according to $\Pr[i = i] = |\mathcal{X}^i|/|\mathcal{X}|$, then the expected codimension of \mathcal{X}^i is $n - \mathbf{H}_\infty(\mathbf{X}) + O(1)$.

The proof of this lemma is by analysing a greedy algorithm that constructs the \mathcal{X}^j s one by one. We refer the reader to [[GPW20](#), [CFK⁺21](#)] for a proof. We now state the key fact about [Algorithm 2](#): over the uniform input distribution, the the probabilities of typical *transcripts* of Π' (concatenation of Π' 's messages) are similar to those of Π . We write $\Pi' := \Pi'(\mathbf{X}, \mathbf{Y})$ for the transcript generated on a uniform random input $(\mathbf{X}, \mathbf{Y}) \sim \{0, 1\}^N \times \{0, 1\}^N$. The following is a version of [[YZ24b](#), [Lemma 3.3](#)]. Their original statement bounds the expected codimension of the leaf ℓ reached, but the same proof implicitly bounds $\log(1/\Pr[\Pi' = \ell])$, whose average value is $\mathbf{H}(\Pi')$ by definition. (Note that our use of Huffman encoding only relabels messages, which does not change the entropy of the transcript.)

Algorithm 2 Protocol Π' on input (x, y)

initialize: $v \leftarrow$ root of Π , $X \times Y = \{0, 1\}^N \times \{0, 1\}^N$, $I, J \leftarrow \emptyset$
while v is not at leaf level **do** [invariant: $X \times Y$ is γ -subcube-like w.r.t. (I, J)]
 let v_0, v_1 be the children of v
 Suppose Alice sends a bit at v (*otherwise swap roles of X and Y*)
 let $X = X^0 \cup X^1$ be the partition according to Alice's message function at v
 update $X \leftarrow X^b$ where b is such that $x \in X^b$
 let $X = \bigcup_i X^i$ be a density-restoring partition for X (with associated sets I_i)
 Let $C(\cdot)$ be the Huffman encoding of \mathbf{k} where $\Pr[\mathbf{k} = i] = |X_i|/|X|$
 update $X \leftarrow X^i$ and $I \leftarrow I \cup I_i$ where i is such that $x \in X^i$
 Alice sends $(b, C(i))$ to Bob
 update $v \leftarrow v_b$
Output the same as leaf v

Lemma 32 ([YZ24b]). $\mathbf{H}(\Pi') \leq O(|\Pi|)$.

We can now verify [Items 1](#) and [2](#). Denote by $\mathbf{v}_1, \dots, \mathbf{v}_d$ the nodes of the protocol tree of Π' encountered when running it on a uniform random input. That is, \mathbf{v}_1 is the root node, \mathbf{v}_d is a leaf, and d is the number of communication rounds. We can thus calculate,

$$\begin{aligned} \mathbb{E}[|\Pi'|] &= \sum_{k \in [d]} \mathbb{E}[|\mathbf{b}, C(\mathbf{i}) \text{ at } \mathbf{v}_k|] \\ &\leq d + \sum_k \mathbb{E}[|C(\mathbf{i}) \text{ at } \mathbf{v}_k|] \\ &\leq d + \sum_k (\mathbf{H}(\mathbf{i} \text{ at } \mathbf{v}_k) + 1) && \text{(Fact 30)} \\ &\leq 2d + \mathbf{H}(\Pi') \\ &\leq O(|\Pi|). && \text{(Lemma 32)} \end{aligned}$$

If the communication cost exceeds $O(|\Pi|)$ by a factor larger than $1/\epsilon$, we make Alice and Bob abort and output \perp . By Markov's inequality, this increases the probability of \perp by at most ϵ . For the new protocol, we have [Item 1](#) by definition and [Item 2](#) since the new error is at most $\epsilon + \epsilon = 2\epsilon$. This concludes the proof sketch of [Theorem 29](#).

Remark. Even though we assumed Π sends one bit per communication round, Π' can send long *variable-length* messages. This is important in order to ensure every node in Π' is subcube-like.

5.4 Proof for any $p \in (0, 1)$

To extend the proof of [Theorem 20](#) to p -biased distributions, we will do the analysis over a modified definition of the input. For any $p = 2^{-k}$ for some integer k , define $H := (\text{AND}_k^N(\overline{H}_1), \text{AND}_k^N(\overline{H}_2))$ where $\overline{H} := \overline{H}_1, \overline{H}_2 \in \{0, 1\}^{kN}$. That is, H is defined by dividing the input \overline{H} into blocks of size k and performing a logical AND on those bits. Sampling \overline{H} uniformly from $\{0, 1\}^{kN} \times \{0, 1\}^{kN}$ then results in a p -biased distribution for H . Thus, doing the analysis for this modified problem over the uniform distribution gives us a lower bound for BINCC over the p -biased distribution.

Since our input is now a uniform distribution, we can assume we have a subcube-like protocol by [Theorem 29](#). Note that for any $p_2 > p_1$, BINCC is easier with the p_1 -biased distributions than p_2 -biased distributions. If Alice and Bob are given a sample H from a p_1 -biased distribution, then Alice and Bob can flip each 0-bit independently to 1 with probability $(p_2 - p_1)/(1 - p_1)$. The result is

another random variable H' which is p_2 -biased. If they run the algorithm on H' and get a solution c , then that c is also valid for H . Since for any $p \in (0, 1)$ there is an integer k such that $p \geq 2^{-k}$, it suffices to prove the lower bound for such distributions. We claim that [Theorem 20](#) also continues to hold for subcube-like protocols when we have a p -biased distribution for any constant $p = 2^{-k}$.

Proof of [Theorem 20](#) for p -biased distributions. We sample $\overline{H} \sim \{0, 1\}^{kN} \times \{0, 1\}^{kN}$, which induces a p -biased distribution on H , where $p = 2^{-k}$. We preprocess the protocol over \overline{H} using [Theorem 29](#). Define $\mathcal{Q} = \mathcal{Q}_1, \dots, \mathcal{Q}_d$ the same way for the candidate subcube-like protocol Π over this distribution. Note that [Claim 24](#) continues to hold. For any c_j , if it becomes dangerous for the first time at v , then Bob has at least $0.1n$ unfixed bits in his input. Since Bob's rectangle is subcube-like, the underlying bits of \overline{H} come from a 0.8-dense distribution. The min-entropy of the marginal distribution for these bits is at least $0.08kn$. This implies $\Pr[H(c_j) = 0^n \mid \overline{H} \in R_v] \leq 2^{-0.08kn} = 2^{-\Omega(n)}$. We can then conclude that $\Pr[\Pi(\overline{H}) \text{ is correct}] = o(1)$ using the same calculation as before. \square

6 Converting to a total relation

So far, we have shown that BiNC achieves an exponential separation between quantum SMP and classical two-way communication *over a biased input distribution*. In this section, we convert BiNC into a total problem and prove our main theorem.

Theorem 1 (Main result). *There exists a total two-party search problem $S \subseteq \{0, 1\}^N \times \{0, 1\}^N \times \mathcal{Q}$ for $N = 2^{\Theta(n)}$ that admits a $\text{poly}(n)$ -bit protocol in the quantum SMP model, but requires $2^{n^{\Omega(1)}}$ bits of communication in the randomized two-way communication model.*

6.1 Construction

Following the approach of Yamakawa and Zhandry [[YZ24a](#), Section 6] towards obtaining a separation for *total* problems, we consider a total problem variant of BiNC, to which we refer to as Total Bipartite NullCodeword (T-BiNC). At a high level, we allow Alice and Bob to pick a common hash function from a $\lambda(n)$ -wise independent hash function family, and then XOR it with their own input. The quantum easiness still holds for any $\lambda(n)$ -wise independent random inputs. However, the freedom of choosing a common hash function also makes the problem too easy, so we further require Alice and Bob to solve t instances of BiNC while using the same common hash function.

To provide a formal definition, we first introduce some notation. Recall that we can simulate the p -biased distribution ([Definition 14](#)), for $p = 2^{-6}$, by splitting each bit into the conjunction of six uniformly random bits. For any function $f: \mathcal{X} \rightarrow \{0, 1\}^m$, where \mathcal{X} is any input domain, we use notation $\overline{f}: \mathcal{X} \rightarrow \{0, 1\}^{6m}$ for the expanded version. Though, we often give $\overline{f}: \mathcal{X} \rightarrow \{0, 1\}^{6m}$ first, and then define $f: \mathcal{X} \rightarrow \{0, 1\}^m$ by taking the AND of every six consecutive bits of the output of \overline{f} .

Fix C to be the code defined in [Definition 15](#). We abuse notation and write $\overline{h}: \Sigma^n \rightarrow \{0, 1\}^{6n}$ as the concatenation of $\overline{h}(\cdot, 1), \dots, \overline{h}(\cdot, n)$. We take $\lambda(n) = n^2$ and assume $\{\overline{h}_k\}$ is implemented by the standard low-degree polynomial construction. See, e.g., [[Vad12](#), Section 3.5.5]. With this construction, $|\mathcal{K}| = 2^r$ with $r = \text{poly}(n)$.

Definition 33 (Total Bipartite NullCodeword, T-BiNC). Take $t := nr$. Let $\overline{H}^{(1)}, \dots, \overline{H}^{(t)}$ be t mappings $\Sigma^n \rightarrow \{0, 1\}^{6n}$, where each $\overline{H}^{(i)}$ is the concatenation of n mappings $\overline{H}_1^{(i)}, \dots, \overline{H}_n^{(i)}$ of type $\Sigma \rightarrow \{0, 1\}^6$. Let $\overline{H}_A^{(i)}, \overline{H}_B^{(i)}: \Sigma^{n/2} \rightarrow \{0, 1\}^{3n}$ be the first half and the second half of $\overline{H}^{(i)}$.

T-BiNC is defined as the following communication problem: Alice and Bob are given t mappings $(\overline{H}_A^{(1)}, \dots, \overline{H}_A^{(t)})$ and $(\overline{H}_B^{(1)}, \dots, \overline{H}_B^{(t)})$, respectively, and they need to find a key $k \in \mathcal{K}$ and t strings

$x^{(1)}, \dots, x^{(t)}$ such that

$$\forall i \in [t] : \quad x^{(i)} \in C_n \quad \text{and} \quad H^{(i)}(x^{(i)}) \oplus h_k(x^{(i)}) = 0^n.$$

Remark (totality). The totality of T-BiNC is implied by the fact that for any input, the quantum SMP protocol can always find a solution with high probability, which will be proved later. A more direct proof can be shown as follows: Fixing any input $\overline{H}^{(1)}, \dots, \overline{H}^{(t)}$. For each $i \in [t]$, by the $\lambda(n)$ -wise independence of \overline{h}_k , it holds that

$$\Pr_{k \sim \mathcal{K}}[\{x \in C_n : H^{(i)}(x^{(i)}) \oplus h_k(x^{(i)}) = 0^n\} = \emptyset] = 2^{-\Omega(n)}.$$

Since $t = \omega(n)$, there exists a key k such that the set of codewords mapped to 0^n is non-empty.

6.2 Worst-case quantum easiness for T-BiNC

We here present a proof sketch for the quantum easiness of T-BiNC.

Lemma 34. T-BiNC has a quantum SMP protocol using $\text{poly}(n)$ qubits over the worst-case input.

Proof Sketch. As in [YZ24a], the query algorithm for the total version of the NULLCODEWORD (T-NC^C) will first randomly pick a key $k \sim \mathcal{K}$, and then solve t copies of NC^C, where the i -th NC^C instance is given by oracle $H^{(i)} \oplus h_k$. Let us call this algorithm \mathcal{A}^T . Similar to the proof of Theorem 19, the algorithm \mathcal{A}^T can also be captured by the same two-stage framework, **Query** and **Process**. Therefore, our SMP protocol for T-BiNC can simulate it in the same manner: Alice and Bob simulate the **Query** stage on their own inputs and send their results to Charlie; Charlie then simulates the **Process** stage and output the solution. See Algorithm 3 for the details of the quantum SMP protocol for T-BiNC.

Now we have to show that for the code C specified in Definition 15, the algorithm \mathcal{A}^T correctly solves T-NC^C. It was shown in [YZ24a, Zha12, Lemma 6.10 using Lemma 2.5] that since \mathcal{A}^T uses the algorithm \mathcal{A}^{YZ} for solving NC^C as a sub-procedure, the correctness of \mathcal{A}^{YZ} (in the average case) would imply the correctness of \mathcal{A}^T in the worst case. Recall that the correctness of \mathcal{A}^{YZ} for NC^C is already proven in Theorem 19, we conclude Lemma 34. \square

6.3 Classical lower bound for T-BiNC

The classical lower bound follows via a union bound after applying a direct product theorem. We shall use the following direct product theorem for randomized communication complexity by Braverman, Rao, Weinstein, and Yehudayoff [BRWY13].

Theorem 35 ([BRWY13]). *Let f be any communication problem (function or relation), and μ be a product distribution of f 's input. Denote the maximum success probability of a classical two-way communication protocol for solving f by $\text{suc}(f, \mu, C)$. Let f^n denote the problem of solving n copies of f and μ^n denote the product distribution on n inputs.*

Then, if $T \log^2 T = o(nC)$ and $\text{suc}(f, \mu, C) \leq 2/3$, it holds that $\text{suc}(f^n, \mu^n, T) \leq 2^{-\Omega(n)}$.

Remark (functions vs relations). Theorem 35 was originally proven for boolean functions f in [BRWY13]. However, their technique does not depend on whether there is a unique solution or many valid solutions. Specifically for product distributions, the technique is to

- (i) extract a protocol with low *external information* cost, and then

Algorithm 3 Quantum SMP Protocol for T-BiNC

Notation:

For each $i \in [t], j \in [n]$, define the states

$$|\phi_j^{(i)}\rangle \propto \sum_{e \in \Sigma} |e\rangle |H_j^{(i)}(e)\rangle, \quad |\bar{\phi}_j^{(i)}\rangle \propto \sum_{e \in \Sigma} |e\rangle |H_j^{(i)}(e) \oplus h_k(e, j)\rangle$$

and define the set

$$T_j^{(i)} = \{e \in \Sigma : H_j^{(i)}(e) \oplus h_k(e, j) = 0\}.$$

Let $U_{\text{add}}, U_{\text{decode}}$ be unitaries defined as follows:

$$|x\rangle |e\rangle \xrightarrow{U_{\text{add}}} |x\rangle |x+e\rangle \xrightarrow{U_{\text{decode}}} |x - \text{Decode}(x+e)\rangle |x+e\rangle.$$

Alice:

For each $i \in [t]$:

1. Prepare n copies of $|\phi_1^{(i)}\rangle, \dots, |\phi_{n/2}^{(i)}\rangle$ using $H_1^{(i)}, \dots, H_{n/2}^{(i)}$.
2. Send those copies of $|\phi_1^{(i)}\rangle, \dots, |\phi_{n/2}^{(i)}\rangle$ to Charlie.

Bob:

For each $i \in [t]$:

1. Prepare n copies of $|\phi_{n/2+1}^{(i)}\rangle, \dots, |\phi_n^{(i)}\rangle$ using $H_{n/2+1}^{(i)}, \dots, H_n^{(i)}$.
2. Send those copies of $|\phi_{n/2+1}^{(i)}\rangle, \dots, |\phi_n^{(i)}\rangle$ to Charlie.

Charlie:

Randomly draw a key $k \sim \mathcal{K}$ and prepare t copies of the state $|\psi\rangle \propto \sum_{x \in C} |x\rangle$.

for $i \in [t]$ **do**

1. For each $j \in [n]$, generate the state $|\bar{\phi}_j^{(i)}\rangle$ using $|\phi_j^{(i)}\rangle$ and then measure the second register of it. If the measurement returns 0, then Charlie successfully generates the state

$$|\phi_j^k\rangle \propto \sum_{e \in T_j^{(i)}} |e\rangle;$$

otherwise, retry with another copy of $|\phi_j^{(i)}\rangle$.

2. Apply QFT to $|\phi^k\rangle := |\phi_1^k\rangle \otimes \dots \otimes |\phi_n^k\rangle$ and $|\psi\rangle$. Now, Charlie holds the state:

$$|\eta\rangle := \text{QFT} |\psi\rangle \otimes \text{QFT} |\phi^k\rangle.$$

3. Apply $(I \otimes \text{QFT}^{-1})U_{\text{decode}}U_{\text{add}}$ to $|\eta\rangle$.
4. Measure the second register and store the measurement result as $x^{(i)}$.

Output the solution $(k, x^{(1)}, \dots, x^{(t)})$.

(ii) compress the protocol using ideas from [BBCR13].

Step (i) is easily seen to be true regardless of the number of solutions. Step (ii) works for relations because the key idea is to simulate the transcript of the protocol, which is oblivious to whether we are solving a function or relation.

Now we are ready to establish the classical lower bound of T-BINC.

Lemma 36. T-BINC requires $2^{n^{\Omega(1)}}$ communication for any classical two-way randomized communication protocol.

Proof. By Yao’s principle, it suffices to prove a lower bound when the inputs $(\overline{H}_A^{(1)}, \dots, \overline{H}_A^{(t)})$ and $(\overline{H}_B^{(1)}, \dots, \overline{H}_B^{(t)})$ are drawn from the uniform distribution. By definition, for each $i \in [t]$, $H_A^{(i)}$ and $H_B^{(i)}$ are drawn from the p -biased distribution.

Now consider a fixed hash function \overline{h}_k where $k \in \mathcal{K}$. Note that for each $i \in [t]$, $H_A^{(i)} \oplus h_k$ and $H_B^{(i)} \oplus h_k$ are also drawn from the p -biased distribution, and they can be viewed as t copies of input to BINC. By Theorem 20, we know that there exist a number $C = 2^{n^{\Omega(1)}}$ such that $\text{suc}(\text{BINC}, \mu, C) \leq 2/3$, where μ is the p -biased distribution. Then, we apply Theorem 35, it holds that $\text{suc}(\text{BINC}^t, \mu^t, C) \leq 2^{-\Omega(t)}$.

Finally, by a union bound over all $k \in \mathcal{K}$, the success probability of a communication protocol for T-BINC using C bits is at most

$$|\mathcal{K}| \cdot 2^{-\Omega(t)} = 2^r \cdot 2^{-\Omega(nr)} = 2^{-\Omega(n)},$$

since $|\mathcal{K}| = 2^r$ and $t = nr$. □

7 Future directions

We highlight three directions for future investigation.

Total boolean functions. The most immediate problems left open by our work are showing any exponential separations between classical and quantum communication for *total boolean functions*, or proving that they are always polynomially related. We do not expect any modification of the Yamakawa–Zhandry problem to work since it is a *search* problem with inherently exponential number of solutions.

Multiparty communication. Beyond the 2-party model, it would be interesting to show that k -party quantum NOF (Number on Forehead) communication can be exponentially separated from k -party randomized NOF communication. We expect that a natural 3-party version of the problem we study achieves this separation, where each player gets $2/3$ of the input oracles. However, in the NOF model we lack techniques to lower bound randomized communication which don’t also lower bound quantum communication (see [LSS09]). Note that lower bounds in the 2-party model already imply lower bounds in the k -party NIH (Number in Hand) model, and it is easy to verify that our quantum upper bound does indeed work in the NIH model when the inputs to NULLCODEWORD are split into k equal parts among the k parties. Hence, our result gives a separation in the k -party NIH model.

NISQ separation. The era of fault-tolerant quantum computers has not arrived yet, but we do have so-called NISQ (noisy intermediate-scale quantum) devices. Chen et al. [CCHL23] defined NISQ as a computational class and studied its complexity. Can we find a quantum–classical communication separation where the quantum upper bound is in an analogue of their class?

A Correctness of the quantum SMP protocol for BiNC

In this section, we prove the correctness of our quantum SMP protocol described in Algorithm 1.

Lemma 37. *When each H_i is drawn from a p -biased distribution, Charlie outputs a valid solution of BiNC^C with high probability.*

Recall that for each $i \in [n]$,

$$|\phi_i\rangle \propto \sum_{e_i \in \Sigma: H_i(e_i)=0} |e_i\rangle, \quad |\phi\rangle := |\phi_1\rangle \otimes \dots \otimes |\phi_n\rangle,$$

and

$$|\psi\rangle \propto \sum_{x \in C} |x\rangle.$$

After receiving the state $|\phi\rangle$, Charlie is essentially solving the NC^C problem without any additional queries to H_1, \dots, H_n . Recall that in the original algorithm \mathcal{A}^{YZ} for NC^C, the first part is preparing the state $|\phi\rangle$ using queries to H_1, \dots, H_n , and the second part is processing $|\phi\rangle$ without queries. Charlie’s algorithm is the same as the second part of \mathcal{A}^{YZ} . Therefore, it suffices to check that the original analysis of the \mathcal{A}^{YZ} in [YZ24a] still holds for the new code C and biased input distribution.

Analysis of the Yamakawa–Zhandry Algorithm. We follow the original analysis in [YZ24a] and highlight where it has to be changed. We will use $\text{negl}(n)$ to denote the class of functions which are asymptotically smaller than an inverse-polynomial in n , which is the standard notation in cryptographic literature. The main technical lemma in [YZ24a] is the following.

Lemma 38 (Lemma 5.1 in [YZ24a]). *Let $|\psi\rangle$ and $|\phi\rangle$ be quantum states on a quantum system over an alphabet $\Sigma = \mathbb{F}_q^m$ written as*

$$|\psi\rangle = \sum_{x \in \Sigma^n} V(x) |x\rangle, \quad |\phi\rangle = \sum_{e \in \Sigma^n} W(e) |e\rangle.$$

Let $F: \Sigma^n \rightarrow \Sigma^n$ be a function. Let $\text{GOOD} \subseteq \Sigma^n \times \Sigma^n$ be a subset such that for any $(x, e) \in \text{GOOD}$, we have $F(x + e) = x$.

Let BAD be the complement of GOOD , i.e., $\text{BAD} := (\Sigma^n \times \Sigma^n) \setminus \text{GOOD}$. Suppose that we have

$$\sum_{(x,e) \in \text{BAD}} |\hat{V}(x)\hat{W}(e)|^2 \leq \epsilon \tag{7}$$

$$\sum_{z \in \Sigma^n} \left| \sum_{(x,e) \in \text{BAD}: x+e=z} \hat{V}(x)\hat{W}(e) \right|^2 \leq \delta. \tag{8}$$

Let U_{add} and U_F be unitaries defined as follows:

$$|x\rangle |e\rangle \xrightarrow{U_{\text{add}}} |x\rangle |x+e\rangle \xrightarrow{U_F} |x - F(x+e)\rangle |x+e\rangle.$$

Then we have

$$(I \otimes \text{QFT}^{-1})U_F U_{\text{add}}(\text{QFT} \otimes \text{QFT})|\psi\rangle|\phi\rangle \approx_{\sqrt{\epsilon}+\sqrt{\delta}} |\Sigma|^{n/2} \sum_{z \in \Sigma^n} (V \cdot W)(z) |0\rangle|z\rangle.$$

Define set $T_i := \{e \in \Sigma : H_i(e) = 0\}$, $T := T_1 \times T_2 \times \dots \times T_n \in \Sigma^n$, one can represent

$$|\psi\rangle = \sum_{x \in \Sigma^n} V(x) |x\rangle, \quad |\phi\rangle = \sum_{e \in \Sigma^n} W(e) |e\rangle$$

with

$$V(x) = \begin{cases} \frac{1}{\sqrt{|C|}} & x \in C \\ 0 & \text{otherwise} \end{cases}$$

$$W(e) = \begin{cases} \frac{1}{\sqrt{|T|}} & e \in T \\ 0 & \text{otherwise} \end{cases}$$

By definition, a string z is a valid solution if and only if $V(z) \cdot W(z) > 0$. Therefore, by taking $F(\cdot) := \text{Decode}(\cdot)$ and assuming inequalities (7) and (8) are satisfied, [Lemma 38](#) implies that Charlie will get a valid solution with high probability when he measures the second register of the resulted quantum state of the step (3).

It remains to prove inequalities (7) and (8) hold with high probability.

Claim 39 (Modified from Claim 6.4 in [[YZ24a](#)]). *With high probability over $H = (H_1, \dots, H_n)$ drawn from a p -biased distribution, there is a subset $\text{GOOD} \subseteq \Sigma^n \times \Sigma^n$ such that $\text{Decode}_{C^\perp}(x + e) = x$ for any $(x, e) \in \text{GOOD}$ and we have*

$$\sum_{(x,e) \in \text{BAD}} |\hat{V}(x)\hat{W}(e)|^2 \leq \text{negl}(n),$$

$$\sum_{z \in \Sigma^n} \left| \sum_{(x,e) \in \text{BAD}: x+e=z} \hat{V}(x)\hat{W}(e) \right|^2 \leq \text{negl}(n).$$

where $\text{BAD} = (\Sigma^n \times \Sigma^n) \setminus \text{GOOD}$.

To simplify [Claim 39](#), we reuse the definition of set \mathcal{G} in [Lemma 16](#),

$$\mathcal{G} := \{e \in \Sigma^n : \text{hw}(e) \leq (p + \epsilon)n\}.$$

Take $\mathcal{B} := \Sigma^n \setminus \mathcal{G}$. By [Claim 18](#), we can set $\text{GOOD} := C^\perp \times \mathcal{G}$, and $\text{BAD} := (\Sigma^n \times \Sigma^n) \setminus \text{GOOD}$. Plug them into [Claim 39](#), and note that $\hat{V}(x) = 0$ for all $x \notin C^\perp$, we have the following:

$$\sum_{(x,e) \in \text{BAD}} |\hat{V}(x)\hat{W}(e)|^2 = \sum_{e \in \mathcal{B}} |\hat{W}(e)|^2,$$

$$\sum_{z \in \Sigma^n} \left| \sum_{(x,e) \in \text{BAD}: x+e=z} \hat{V}(x)\hat{W}(e) \right|^2 = \sum_{z \in \Sigma^n} \left| \sum_{\substack{x \in C^\perp, e \in \mathcal{B} \\ :x+e=z}} \hat{V}(x)\hat{W}(e) \right|^2.$$

Then, by an averaging argument, it suffices bound their expected values to be negligible, i.e.,

$$\mathbf{E}_H \left[\sum_{e \in \mathcal{B}} |\hat{W}(e)|^2 \right] \leq \text{negl}(n), \quad (9)$$

$$\mathbf{E}_H \left[\sum_{z \in \Sigma^n} \left| \sum_{\substack{x \in C^\perp, e \in \mathcal{B} \\ :x+e=z}} \hat{V}(x) \hat{W}(e) \right|^2 \right] \leq \text{negl}(n). \quad (10)$$

We only prove inequality (9) here, while the proof for inequality (10) is the same as in [YZ24a]. For any $i \in [n]$, let

$$W_i(e_i) = \begin{cases} \frac{1}{\sqrt{|T_i|}} & e_i \in T_i \\ 0 & \text{otherwise} \end{cases}$$

We now prove the following claim.

Claim 40 (Modified from the first half of Claim 6.6 in [YZ24a]). *For all $i \in [n]$, it holds that*

$$\mathbf{E}_{H_i} \left[|\hat{W}_i(0)|^2 \right] = 1 - p,$$

Proof.

$$\mathbf{E}_{H_i} \left[|\hat{W}_i(0)|^2 \right] = \mathbf{E}_{H_i} \left[\left| \frac{1}{\sqrt{|\Sigma|}} \sum_{z \in \Sigma} W_i(z) \right|^2 \right] = \frac{\mathbf{E}_{H_i} [|T_i|]}{|\Sigma|} = 1 - p.$$

□

The next claim follows from the same proof as in [YZ24a] by exploiting the symmetry.

Claim 41 (Second half of Claim 6.6 in [YZ24a]). *For all $i \in [n]$ and $e, e' \in \Sigma \setminus \{0\}$, it holds that*

$$\mathbf{E}_{H_i} \left[|\hat{W}_i(e)|^2 \right] = \mathbf{E}_{H_i} \left[|\hat{W}_i(e')|^2 \right].$$

Note that

$$\hat{W}^H(e) = \prod_{i=1}^n \hat{W}_i^{H_i}(e_i).$$

Therefore, by combining Claim 40 and Claim 41, we get that for all $e \in \Sigma^n$,

$$\mathbf{E}_H \left[|\hat{W}(e)|^2 \right] = \mathcal{D}_n(e),$$

where \mathcal{D}_n is the same distribution defined in Lemma 16—each symbol is 0 with probability $1 - p$ and otherwise a uniformly random element of $\Sigma \setminus \{0\}$.

Finally, by the linearity of expectation and Lemma 16, we conclude that

$$\mathbf{E}_H \left[\sum_{e \in \mathcal{B}} |\hat{W}(e)|^2 \right] = \Pr_{e \sim \mathcal{D}_n} [e \in \mathcal{B}] \leq \text{negl}(n).$$

Acknowledgments

MG is supported by the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract number MB22.00026. TG is supported by ERC Starting Grant 101163189 and UKRI Future Leaders Fellowship MR/X023583/1. SJ and JL are supported by Scott Aaronson’s Quantum Systems Accelerator grant.

References

- [ABK⁺21] Scott Aaronson, Shalev Ben-David, Robin Kothari, Shravas Rao, and Avishay Tal. Degree vs. approximate degree and quantum implications of Huang’s sensitivity theorem. In *Proceedings of the 53rd Symposium on Theory of Computing (STOC)*, pages 1330–1342. ACM, 2021. doi:10.1145/3406325.3451047.
- [ABK24] Scott Aaronson, Harry Buhrman, and William Kretschmer. A qubit, a coin, and an advice string walk into a relational problem. In *Proceedings of the 15th Innovations in Theoretical Computer Science Conference (ITCS)*, volume 287 of *LIPICs*. Schloss Dagstuhl, 2024. doi:10.4230/lipics.itcs.2024.1.
- [AKK16] Andris Ambainis, Martins Kokainis, and Robin Kothari. Nearly optimal separations between communication (or query) complexity and partitions. In *Proceedings of the 31st Conference on Computational Complexity (CCC)*, volume 50 of *LIPICs*, pages 4:1–4:14. Schloss Dagstuhl, 2016. doi:10.4230/LIPICs.CCC.2016.4.
- [BBCR13] Boaz Barak, Mark Braverman, Xi Chen, and Anup Rao. How to compress interactive communication. *SIAM Journal on Computing*, 42(3):1327–1363, 2013. doi:10.1137/100811969.
- [BCW99] Harry Buhrman, Richard Cleve, and Avi Wigderson. Quantum vs. classical communication and computation. In *Proceedings of the 30th Symposium on Theory of Computing (STOC)*, pages 63–68. ACM, 1999. doi:10.1145/276698.276713.
- [BJK04] Ziv Bar-Yossef, T. S. Jayram, and Iordanis Kerenidis. Exponential separation of quantum and classical one-way communication complexity. In *Proceedings of the 36th Symposium on Theory of Computing (STOC)*, pages 128–137. ACM, 2004. doi:10.1145/1007352.1007379.
- [BK24a] Shalev Ben-David and Srijita Kundu. Oracle separation of QMA and QCMA with bounded adaptivity. In *Proceedings of the 51st International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 297 of *LIPICs*, pages Art. No. 21, 18. Schloss Dagstuhl, 2024. doi:10.4230/lipics.icalp.2024.21.
- [BK24b] Shalev Ben-David and Srijita Kundu. Separations in query complexity for total search problems. Technical report, arXiv, 2024. doi:10.48550/ARXIV.2410.16245.
- [BRWY13] Mark Braverman, Anup Rao, Omri Weinstein, and Amir Yehudayoff. Direct products in communication complexity. In *Proceedings of the 54th Symposium on Foundations of Computer Science (FOCS)*, pages 746–755. IEEE, 2013. doi:10.1109/FOCS.2013.85.
- [BS21] Nikhil Bansal and Makrand Sinha. k -forrelation optimally separates quantum and classical query complexity. In *Proceedings of the 53rd Symposium on Theory of Computing (STOC)*, pages 1303–1316. ACM, 2021. doi:10.1145/3406325.3451040.

- [CCHL23] Sitan Chen, Jordan Cotler, Hsin-Yuan Huang, and Jerry Li. The complexity of NISQ. *Nature Communications*, 14(1), 2023. doi:10.1038/s41467-023-41217-6.
- [CFK⁺21] Arkadev Chattopadhyay, Yuval Filmus, Sajin Koroth, Or Meir, and Toniann Pitassi. Query-to-communication lifting using low-discrepancy gadgets. *SIAM Journal on Computing*, 50(1):171–210, 2021. doi:10.1137/19M1310153.
- [dRGR22] Susanna de Rezende, Mika Göös, and Robert Robere. Proofs, circuits, and communication. *SIGACT News*, 53(1), 2022. doi:10.1145/3532737.3532745.
- [Gav16] Dmitry Gavinsky. Entangled simultaneity versus classical interactivity in communication complexity. In *Proceedings of the 48th Symposium on Theory of Computing (STOC)*, pages 877–884. ACM, New York, 2016. doi:10.1145/2897518.2897545.
- [Gav19] Dmitry Gavinsky. Quantum versus classical simultaneity in communication complexity. *IEEE Transactions on Information Theory*, 65(10):6466–6483, 2019. doi:10.1109/TIT.2019.2918453.
- [Gav21] Dmitry Gavinsky. Bare quantum simultaneity versus classical interactivity in communication complexity. *IEEE Transactions on Information Theory*, 67(10):6583–6605, 2021. doi:10.1109/TIT.2021.3050528.
- [GKK⁺08] Dmitry Gavinsky, Julia Kempe, Iordanis Kerenidis, Ran Raz, and Ronald de Wolf. Exponential separation for one-way quantum communication complexity, with applications to cryptography. *SIAM Journal on Computing*, 38(5):1695–1708, 2008. doi:10.1137/070706550.
- [GKRS19] Mika Göös, Prithish Kamath, Robert Robere, and Dmitry Sokolov. Adventures in monotone complexity and TFNP. In *Proceedings of the 10th Innovations in Theoretical Computer Science Conference (ITCS)*, pages 38:1–38:19, 2019. doi:10.4230/LIPIcs.ITCS.2019.38.
- [GLM⁺16] Mika Göös, Shachar Lovett, Raghu Meka, Thomas Watson, and David Zuckerman. Rectangles are nonnegative juntas. *SIAM Journal on Computing*, 45(5):1835–1869, 2016. doi:10.1137/15M103145X.
- [GPW20] Mika Göös, Toniann Pitassi, and Thomas Watson. Query-to-communication lifting for BPP. *SIAM Journal on Computing*, 49(4):441–461, 2020. doi:10.1137/17M115339X.
- [GR08] Venkatesan Guruswami and Atri Rudra. Explicit codes achieving list decoding capacity: Error-correction with optimal redundancy. *IEEE Transactions on Information Theory*, 54(1):135–150, 2008. doi:10.1109/TIT.2007.911222.
- [GRT22] Uma Girish, Ran Raz, and Avishay Tal. Quantum versus randomized communication complexity, with efficient players. *Computational Complexity*, 31(2):17, 2022. doi:10.1007/S00037-022-00232-7.
- [GS99] Venkatesan Guruswami and Madhu Sudan. Improved decoding of reed-solomon and algebraic-geometry codes. *IEEE Transactions on Information Theory*, 45(6):1757–1767, 1999. doi:10.1109/18.782097.
- [Huf52] David Huffman. A method for the construction of minimum-redundancy codes. *Proceedings of the IRE*, 40(9):1098–1101, 1952. doi:10.1109/JRPROC.1952.273898.
- [KN97] Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, 1997. doi:10.1017/CBO9780511574948.

- [KR11] Bo'az Klartag and Oded Regev. Quantum one-way communication can be exponentially stronger than classical communication. In *Proceedings of the 43rd Symposium on Theory of Computing (STOC)*, pages 31–40. ACM, New York, 2011. doi:10.1145/1993636.1993642.
- [Liu23] Qipeng Liu. Non-uniformity and quantum advice in the quantum random oracle model. In *Proceedings of the 42nd Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, volume 14004 of *Lecture Notes in Computer Science*, pages 117–143. Springer, 2023. doi:10.1007/978-3-031-30545-0\5.
- [LLPY24] Xingjian Li, Qipeng Liu, Angelos Pelecanos, and Takashi Yamakawa. Classical vs quantum advice and proofs under classically-accessible oracle. In *Proceedings of the 15th Innovations in Theoretical Computer Science Conference (ITCS)*, volume 287:72 of *LIPICs*, pages 1–19. Schloss Dagstuhl, 2024. doi:10.4230/lipics.itcs.2024.72.
- [LSS09] Troy Lee, Gideon Schechtman, and Adi Shraibman. Lower bounds on quantum multiparty communication complexity. In *Proceedings of the 24th Conference on Computational Complexity (CCC)*, pages 254–262. IEEE, 2009. doi:10.1109/CCC.2009.24.
- [NC00] Michael Nielsen and Isaac Chuang. *Quantum computation and quantum information*. Cambridge University Press, 2000. doi:10.1017/CBO9780511976667.
- [Raz99] Ran Raz. Exponential separation of quantum and classical communication complexity. In *Proceedings of the 31st Symposium on Theory of Computing (STOC)*, page 358–367. ACM, 1999. doi:10.1145/301250.301343.
- [Rud07] Atri Rudra. *List decoding and property testing of error-correcting codes*. PhD thesis, University of Washington, 2007.
- [RY20] Anup Rao and Amir Yehudayoff. *Communication Complexity: And Applications*. Cambridge University Press, 2020. doi:10.1017/9781108671644.
- [SSW23] Alexander Sherstov, Andrey Storozhenko, and Pei Wu. An optimal separation of randomized and quantum query complexity. *SIAM Journal on Computing*, 52(2):525–567, 2023. doi:10.1137/22M1468943.
- [Vad12] Salil Vadhan. Pseudorandomness. *Foundations and Trends in Theoretical Computer Science*, 7(1-3):1–336, 2012. doi:10.1561/0400000010.
- [WYZ23] Shuo Wang, Guangxu Yang, and Jiapeng Zhang. Communication complexity of set-intersection problems and its applications. Technical Report TR23-164, Electronic Colloquium on Computational Complexity (ECCC), 2023. URL: <https://eccc.weizmann.ac.il/report/2023/164/>.
- [YZ24a] Takashi Yamakawa and Mark Zhandry. Verifiable quantum advantage without structure. *Journal of the ACM*, 71(3):1–50, 2024. doi:10.1145/3658665.
- [YZ24b] Guangxu Yang and Jiapeng Zhang. Communication lower bounds for collision problems via density increment arguments. In *Proceedings of the 56th Symposium on Theory of Computing (STOC)*, 2024. doi:10.1145/3618260.3649607.
- [Zha12] Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. In *Proceedings of the 32nd Cryptology Conference (CRYPTO)*, volume 7417 of *Lecture Notes in Computer Science*, pages 758–775. Springer, 2012. doi:10.1007/978-3-642-32009-5\44.