

Communication Lower Bounds of Key-Agreement Protocols via Density Increment Arguments

Mi-Ying (Miryam) Huang * Xinyu Mao * Guangxu Yang * Jiapeng Zhang *

September 10, 2023

Abstract

Constructing key-agreement protocols in the random oracle model (ROM) is a viable method to assess the feasibility of developing public-key cryptography within Minicrypt. Unfortunately, as shown by Impagliazzo and Rudich (STOC 1989) and Barak and Mahmoody (Crypto 2009), such protocols can only guarantee limited security: any ℓ -query protocol can be attacked by an $O(\ell^2)$ -query adversary. This quadratic gap matches the key-agreement protocol proposed by Merkle (CACM 78), known as Merkle’s Puzzles.

Besides query complexity, the communication complexity of key-agreement protocols in the ROM is also an interesting question in the realm of fine-grained cryptography, even though only limited security is achievable. Haitner et al. (ITCS 2019) first observed that in Merkle’s Puzzles, to obtain secrecy against an eavesdropper with $O(\ell^2)$ queries, the honest parties must exchange $\Omega(\ell)$ bits. Therefore, they conjectured that high communication complexity is unavoidable, any ℓ -query protocols with c bits of communication could be attacked by an $O(c \cdot \ell)$ -query adversary. This, if true, will suggest that Merkle’s Puzzle is also optimal regarding communication complexity. Building upon techniques from communication complexity, Haitner et al. (ITCS 2019) confirmed this conjecture for two types of key agreement protocols with certain natural properties.

This work affirms the above conjecture for all non-adaptive protocols with perfect completeness. Our proof uses a novel idea called *density increment argument*. This method could be of independent interest as it differs from previous communication lower bounds techniques (and bypasses some technical barriers).

1 Introduction

Key-agreement protocols [DH76] allow two parties, Alice and Bob, to agree on a shared private key by communicating over an insecure public channel. Its security requires that any (efficient) eavesdropper cannot learn the key from the transcript. In an early work, Merkle [Mer78] first proposed an ingenious key-agreement protocol, known as *Merkle’s Puzzles*, as follows.

*Research supported by NSF CAREER award 2141536.

Thomas Lord Department of Computer Science, University of Southern California.
Email: {miying.huang, xinyumao, guangxuy, jiapengz}@usc.edu

Protocol 1.1 (Merkle’s Puzzles). Let $f : [N] \rightarrow [M]$ be a cryptographic hash function and let ℓ be a parameter measuring the query complexity of this protocol. Alice and Bob first agree on a set $W \subseteq [N]$ of size ℓ^2 . Then, at the beginning of the protocol, Alice makes ℓ random queries in W , i.e., $f(w_1), \dots, f(w_\ell)$. Similarly, Bob makes another ℓ random queries $f(w'_1), \dots, f(w'_\ell)$. By the birthday paradox, there is a good chance that $\{w_1, \dots, w_\ell\} \cap \{w'_1, \dots, w'_\ell\} \neq \emptyset$. Alice then sends $z_1 = f(w_1), \dots, z_\ell = f(w_\ell)$ to Bob, and Bob checks if there is a w'_j in his query such that $f(w'_j) = z_i$ for some $i \in [\ell]$. If such a pair (w'_j, z_i) exists, then Bob sends z_i back to Alice and sets w'_j as his key; otherwise, Bob aborts. Finally, according to z_i , Alice chooses w_i as her key.

As long as the function f is collision-free on W , Alice and Bob will agree on the same key with high probability. In terms of security, if f is modeled as a random function, we can show that any eavesdropper that breaks this protocol with constant probability has to query a constant fraction of inputs in W ; consequently, the query complexity of any eavesdropper must be $\Omega(\ell^2)$.

On the other hand, Impagliazzo and Rudich [IR89], followed by Barak and Mahmoody [BMG09], showed that key-agreement protocol is essentially a public-key primitive and is unlikely to be based only on hardness assumptions for symmetric cryptography—any key-agreement protocol only guarantees limited security as long as the symmetric hardness is used in a black-box way. Specifically, they studied key-agreement protocols in the *random oracle model* (ROM). In the ROM, all parties, including the eavesdropper, have oracle access to a random function $f : [N] \rightarrow [M]$, which is an idealization of symmetric primitives like collision-resistant hash function. The efficiency of parties is measured by the number of queries they make to the oracle (in the worst case). [IR89] proved that any key-agreement protocols in the ROM with ℓ queries can be attacked by an eavesdropper with $O(\ell^6)$ queries. [BMG09] further improved the efficiency of the eavesdropper to $O(\ell^2)$ queries. This result indicates that Merkle’s puzzle is optimal in terms of the number of oracle queries since it reaches quadratic security. Despite its limited security, the complexity of key-agreement protocols in the ROM is still an interesting question of fine-grained cryptography. A long line of research has been conducted on the limitation and possibility of key-agreement protocols in the ROM, in both classical setting [DH76, Mer78, IR89, BMG09, HMO⁺19, ACMS23], distributed setting [DH21] and quantum setting [ACC⁺22].

Besides oracle queries, another important cost in key-agreement protocols is the communication cost between Alice and Bob. The communication complexity of (multi-party) protocols, such as key-agreement, optimally-fair coin tossing, statistically hiding commitment schemes, and multi-party computation, has garnered considerable attention recently [DSLMM11, HHR15, HMO⁺19, Cou19, AHMS20, CN22].

In this paper, we focus on the communication complexity of key-agreement protocols: a problem initiated by Haitner et al. [HMO⁺19]. Concretely, they observed that the communication complexity of Merkle’s Puzzle is also $\tilde{\Omega}(\ell)$ ¹, and they conjectured that high communication cost is unavoidable.

Conjecture 1.2 ([HMO⁺19], informal). *Let $\Pi = (A, B)$ be a key-agreement protocol such that:*

1. *A and B agree on the same key with high probability;*
2. *A and B each make at most ℓ queries to the random function (oracle);*
3. *Π is secure against any adversary with q queries to the random oracle.*

¹We drop low order terms such as $\log N$ and $\log M$ here.

Then **A** and **B** must communicate $\Omega(q/\ell)$ bits.

As we discussed, Merkle’s puzzle matches the lower bound in this conjecture for $q = \Theta(\ell^2)$. For $q = o(\ell^2)$, an asymmetric version of Merkle’s puzzle also matches this lower bound.

Protocol 1.3 (Asymmetric version of Merkle’s Puzzles). Alice and Bob first fix a domain W of size q . Then Alice makes $c := q/\ell$ random queries in W and sends them to Bob. Bob also makes ℓ random queries (in W) and checks if there is a common query in accordance with the original Merkle’s Puzzles.

[HMO⁺19] partly tackled this conjecture for two types of key-agreement protocols. We say a protocol is *non-adaptive* if both parties choose all their queries at the beginning of the protocol (before querying the oracle and communicating); that is, their queries are determined by their internal randomness. Haitner et al. [HMO⁺19] proved that for any protocol $\Pi = (\mathbf{A}, \mathbf{B})$ that satisfies the conditions in conjecture 1.2:

- If Π is non-adaptive and has only two rounds, **A** and **B** must exchange $\Omega(q/\ell)$ bits.
- If the queries are uniformly sampled, then **A** and **B** must communicate $\Omega(q^2/\ell^3)$ bits.

Note that protocols with uniform queries are also special non-adaptive protocols.

In this paper, we affirm conjecture 1.2 for non-adaptive protocols with *perfect completeness*, i.e., Alice and Bob agree on the same key with probability 1. Specifically, we prove the following theorem.

Theorem 1.4 (Informal). *Let $\Pi = (\mathbf{A}, \mathbf{B})$ be a non-adaptive key-agreement protocol such that:*

1. ***A** and **B** agree on the same key with probability 1;*
2. ***A** and **B** each make at most ℓ queries to the random oracle;*
3. *Π is secure against any adversary with q queries to the random oracle.*

*Then **A** and **B** must communicate $\Omega(q/\ell)$ bits.*

Our proof is built on the density increment argument introduced by Yang and Zhang [YZ22, YZ23], which they used to prove communication lower bounds for the *unique disjointness problem*. Looking at our main theorem carefully, we acknowledge two non-trivial requirements in our statement: non-adaptivity and perfect completeness. However, these limitations are not inherent in this method. Therefore, we are optimistic that our method has a good chance to overcome these two limitations; more details will be discussed in section 1.2.

It is worth noting that Mazor [Maz23] recently devised a non-adaptive protocol with perfect completeness and quadratic security guarantee. We observed that this protocol, with minor adjustments, allows a trade-off between communication and security in a similar fashion to protocol 1.3. Our result shows that Mazor’s construction is optimal among non-adaptive protocols with perfect completeness.

1.1 Proof overview

Now we give a high-level overview of our proof. Since the execution of key-agreement protocols and the attacking process involve many random variables, we first explain our notations.

- We use bold and uppercase letters for random variables and corresponding regular letters for samples and values, such as $f, r_A, r_B, Q_A, Q_B, \tau, Q_E$ and f_E (uppercase for sets and lowercase for elements and functions).
- Let F the RO that the parties have access to, which is a random function from $[N]$ to $[M]$. Moreover, let R_A, R_B be Alice's and Bob's internal randomness. (R_A, R_B, F) determines the entire execution of key-agreement protocols.
- Let $Q_A, Q_B \subseteq [N]$ be the queries made by Alice and Bob in the execution, respectively. Notice that Q_A, Q_B is fully determined by R_A, R_B for non-adaptive protocols. Q_A and Q_B are usually ordered sets since Alice and Bob make oracle queries one at a time. For the sake of notation convenience, we sometimes regard Q_A and Q_B as unordered sets.
- Let T be the communication transcript between Alice and Bob. Notice that T is observed by the attacker Eve.
- Let $Q_E \subseteq [N]$ be Eve's queries. Let $F_E = F(Q_E)$ be Eve's observations of the random oracle F . We interpret F_E as a partial function: for every $x \in Q_E$, $F_E(x) = F(x)$; for all other x , $F_E(x) = \perp$.

To study the security of key-agreement protocols, Impagliazzo and Rudich [IR89] observed that the advantage of Alice and Bob over Eve mainly comes from their intersection queries which have not been queried by Eve, i.e., the knowledge from $(Q_A \cap Q_B) \setminus Q_E$ and $F((Q_A \cap Q_B) \setminus Q_E)$. Based on this insight, they devised an attacker that aims to guess (and query) the set $(Q_A \cap Q_B)$. In order to learn intersection queries more efficiently, [BMG09] introduced the notion of *heavy query*. Given Eve's current observation, which consists of a transcript τ and a partial function f_E , an input $w \in [N] \setminus Q_E$ is said to be ε -heavy with respect to (τ, f_E) if

$$\Pr[w \in (Q_A \cap Q_B) \mid \tau, f_E] \geq \varepsilon.$$

Now we give an informal description of Eve's strategy ²:

- **Stage I.** Eve checks if there exists a heavy query conditioned on transcript τ and her observations of the random oracle f_E . If yes, then query them, update f_E , and repeat until there are no heavy queries.
- **Stage II.** Eve simulates Alice and Bob based on observed information and outputs Alice's key in her simulation. In other words, Eve simply outputs a sample from the distribution of Alice's key conditioned on observed information.

Suppose that Alice and Bob each make at most ℓ queries and set $\varepsilon = \Theta(1/\ell)$. A standard technique can prove that Stage I stops within $O(\ell/\varepsilon) = O(\ell^2)$ queries. We can also show that in order to clean up all heavy queries (Stage I), $\Omega(\ell^2)$ queries are inevitable. This querying process does not explore strong connections to communication complexity.

²This is not exactly the same as [BMG09] due to some technical challenges in [BMG09].

Our approach. Our main observation is that if Alice and Bob communicate too little, they cannot utilize their common queries and thus have no advantage over Eve! Hence, we focus on queries correlated with the transcript τ instead of all intersection queries. With this in mind, we introduce *correlated query* as a refined notion of heavy query.

Definition 1.5 (ε -correlated set, informal; see definition 3.2). Eve’s view consists of a transcript τ and a partial function f_E . We say a set $S = \{w_1, \dots, w_r\} \subseteq [N]$ is ε -*correlated* with respect to (τ, f_E) if

$$\mathbf{H}(F(w_1), \dots, F(w_r) \mid R_A, R_B, f_E) - \mathbf{H}(F(w_1), \dots, F(w_r) \mid R_A, R_B, f_E, \tau) \geq \varepsilon,$$

where $\mathbf{H}(\cdot)$ denotes the Shannon entropy.

We use $F(S)$ to denote $(F(w_1), \dots, F(w_r))$ in the future, and $F(S)$ can also be viewed as a partial function with domain S . A main difference between our attacker and [BMG09] is that: instead of making ε -heavy queries, we clean up all ε -correlated sets of size at most 2ℓ . Another difference is that we choose $\varepsilon = \Theta(1)$ and [BMG09] set $\varepsilon = \Theta(1/\ell)$. Intuitively, this is because a correlated set of size ℓ is as effective as ℓ single heavy queries. Along these lines, we then have to prove two things:

- **Success.** Eve can guess the key of Alice/Bob if there is no ε -correlated set of size at most 2ℓ .
- **Efficiency.** Eve can remove all ε -correlated sets (of size at most 2ℓ) after querying $O(c)$ correlated sets, where c is the number of communication bits between Alice and Bob. Thus, the query complexity of Eve is $O(c \cdot \ell)$.

Eve can guess the key if there are no small ε -correlated sets. Assume that the protocol Π is non-adaptive, i.e., Q_A (or Q_B) is determined by r_A (resp., r_B). To study the success probability of Eve, we consider a rectangle $\mathcal{X} \times \mathcal{Y}$ as follows. Every $x \in \mathcal{X}$ has the form $x = (r_A, f_A)$ (Alice’s view) and every $y \in \mathcal{Y}$ has the form $y = (r_B, f_B)$ (Bob’s view), where f_A, f_B have domain Q_A, Q_B respectively. Note that we enumerate x and y independently in the rectangle. Consequently, some pairs (x, y) in this rectangle may be *inconsistent*. Concretely, we say that a pair $x = (r_A, f_A)$ and $y = (r_B, f_B)$ is inconsistent if there exists an input $w \in Q_A \cap Q_B$ such that $f_A(w) \neq f_B(w)$. Define an output table as follows:

$$\mathcal{M}(x, y) \stackrel{\text{def}}{=} \begin{cases} \text{Alice's key output by } \Pi(r_A, r_B, f_A \cup f_B), & \text{if } f_A \text{ and } f_B \text{ are consistent;} \\ *, & \text{otherwise.} \end{cases}$$

This table indeed captures all possible executions of the protocol Π . This table is a partial function because many entries are undefined (the $*$ entries).

During the attack, Eve observes the transcript τ and makes queries to f . Whenever Eve has observed (τ, f_E) , we update the table \mathcal{M} by removing the entries that are inconsistent with Eve’s observation, i.e., we update the table to

$$\mathcal{M}_{\tau, f_E}(x, y) \stackrel{\text{def}}{=} \begin{cases} \mathcal{M}(x, y), & \text{if } (x, y) \text{ are consistent with } (\tau, f_E); \\ *, & \text{otherwise.} \end{cases}$$

Given this observation (τ, f_E) , the defined entries of \mathcal{M}_{τ, f_E} capture all possible views of Alice and Bob. Now we say \mathcal{M}_{τ, f_E} is almost monochromatic if almost all defined entries of \mathcal{M}_{τ, f_E} are equal

to the same output $b \in \{0, 1\}$.³ A key step in our proof is to show \mathcal{M}_{τ, f_E} is almost monochromatic provided that the following conditions are met:

1. Π has perfect completeness;
2. there is no small ε -correlated set respect to (τ, f_E) .

Once Eve realizes \mathcal{M}_{τ, f_E} is almost monochromatic, she knows that Alice’s key is b with high probability.

Upper bound the number of Eve’s queries via density increment argument This part of our proof is based on the density increment argument in [YZ22, YZ23]. We first define a density function to capture the amount of hidden information in the transcript τ about the random function F , which is not known by Eve. For every τ and f_E , its density function $\Phi(\tau, f_E)$ is defined as

$$\Phi(\tau, f_E) \stackrel{\text{def}}{=} \mathbf{H}(F \mid R_A, R_B, f_E) - \mathbf{H}(F \mid R_A, R_B, f_E, \tau).$$

If we replace τ and f_E with corresponding random variables, T and F_E , then $\Phi(T, F_E)$ equals to $\mathbf{I}(F; T \mid R_A, R_B, F_E)$, the mutual information of F and T conditioned on R_A, R_B, F_E . This quantity is strongly related to the *information complexity* (IC), a powerful tool for proving lower bounds in communication complexity [CSWY01, BCCR10]. IC usually refers to the mutual information of Alice’s input and Bob’s input conditioned on the transcript, so the IC for key-agreement should look like $\mathbf{I}(R_A, R_B; T)$. However, in the ROM, the random function F is another random resource involved in the computation. Therefore, we cannot use IC *as a black box* to study such key-agreement protocols. Instead, we use the density increment argument proposed by [YZ23], which reinterprets IC in a white-box manner.

Let us turn back to our proof. The key idea is that whenever Eve *queries an ε -correlated set*, the density function *decreases by at least ε* in expectation. To make things clearer, we first explain our sampling procedure. There are several random variables involved in the analysis, including $(R_A, R_B, F, T, S_1, S_2, \dots)$. Here S_i is the query set made by Eve in the i -th round. In our analysis, we *do not* sample (R_A, R_B, F) all at once. Instead, we consider these random variables to be sampled in the following order.

1. We first sample the transcript $\tau \leftarrow T$ and send it to the attacker.
2. In the i -th round of the attack,
 - Eve samples her next query set S_i conditioned on $(\tau, S_1, f(S_1), \dots, S_{i-1}, f(S_{i-1}))$.
 - We sample $f(S_i)$ conditional on $(\tau, S_1, f(S_1), \dots, S_{i-1}, f(S_{i-1}), S_i)$, and Eve receives $f(S_i)$.

Suppose that at some point, Eve has already observed f_E , e.g., $f_E = f(S_1 \cup \dots \cup S_{i-1})$ and decided to query S_i next. By definition, Eve only queries correlated sets, i.e., S_i is ε -correlated w.r.t. (τ, f_E) .

³More precisely, ‘almost all’ means if we sample an entry (x, y) according to the probability that it appears in real execution (conditioned on τ, f_E), we have $\mathcal{M}(x, y) = b$ with high probability.

And we prove that for any ε -correlated set S_i ,

$$\mathbf{E}_{f(S_i) \leftarrow F(S_i) |_{\tau, f_E}} [\Phi(\mathbf{T}, f_E \cup f(S_i))] \leq \Phi(\mathbf{T}, f_E) - \varepsilon, \quad (1)$$

where $f_E \cup f(S_i)$ is Eve's updated observation after making oracle queries on S_i . We then finish our argument by observing the following two properties of Φ :

- In the beginning, $\Phi(\mathbf{T}, f_\emptyset) \leq c$. Here f_\emptyset denotes the all-empty function since Eve has no information about the oracle before making any queries.
- Φ is non-negative: $\Phi(\tau, f_E) \geq 0$ for all τ, f_E .

Equation (1) says that each time Eve queries an ε -correlated set, Φ decreases by ε (in expectation), so Eve can query at most $O(c/\varepsilon) = O(c)$ sets (in expectation), as we set $\varepsilon = \Theta(1)$. Since each set queried by Eve is of size at most 2ℓ , we conclude that the total number of Eve's queries is $O(c\ell)$.

Comparison with [HMO⁺19]. The paper by Haitner et al. uses mostly direct calculations to derive an upper bound of the mutual information characterizing the advantage of Alice and Bob over Eve. An important ingredient in their proof is that conditioning on Eve's view does not introduce significant dependency between Alice and Bob; this is true for two-round protocols but fails for multi-round protocols. Even with perfect completeness, their approach encounters similar barriers. In contrast, our proof mainly depends on the investigation of the structure of the table \mathcal{M}_{τ, f_E} , and hence the number of rounds is no longer a restriction.

1.2 Discussions and open problems

In this section, we discuss some open problems and future directions. An immediate question is how to remove the restrictions in our main theorem. We briefly discuss some potential ways to solve them below.

Protocols with imperfect completeness. In our proof, the property of perfect completeness is used in lemma 3.6. The perfect completeness restriction is an analog of proving *deterministic* communication complexity, while key-agreement protocols with imperfect completeness can be likened to *randomized* communication protocols. The density increment argument used in this paper was originally inspired by the proofs of *query-to-communication lifting theorems* in communication complexity [RM97, GPW15, GPW17, YZ22]. In communication complexity, past experience suggests that the density increment argument is robust in the sense that it usually extends to proving randomized communication lower bounds. For example, the deterministic query-to-communication lifting theorem was formalized by [GPW15], then [GPW17] proved the extension to the randomized query-to-communication lifting theorem, even though it took several years.

Protocols with adaptive queries. The density increment argument has a good chance of proving communication lower bounds for adaptive protocols. Particularly, our efficiency proof directly applies to adaptive protocols. Our proof only utilized the non-adaptivity in lemma 3.6. The round-by-round analysis introduced by Barak and Mahmoody [BMG09] might be helpful to circumvent this obstacle. Admittedly, the analysis might be slightly more complicated, but we do not see a fundamental barrier here.

Further potential applications. The heavy query technique used in the proof of [BMG09] has found applications in the context of black-box separations and black-box security in the random oracle model (see, e.g., [DSLMM11, KSY11, BKS11, MP12, HOZ16]). Likewise, it will be interesting to check if our approach offers fresh perspectives and potential solutions to some open problems. The following is a list of potential questions.

1. Devise an $O(\ell)$ -round and $O(\ell^2)$ -query attack for key-agreement protocols in the ROM [BMG09, MMV11].
2. Consider an M -party protocol where all pairs among M players agree on secret keys. Given an attack that recovers a constant fraction of the $\binom{M}{2}$ keys with $O(M \cdot \ell^2)$ oracle queries [DH21].
3. In the quantum setting, Alice and Bob are capable of conducting quantum computation and classical communication, and the random oracle allows quantum queries. [ACC+22] introduced the Polynomial Compatibility Conjecture and gave an attack (only for protocols with perfect completeness) assuming this conjecture holds. Devise an attack that has better efficiency or fewer restrictions.

2 Preliminary

2.1 Notations

For a random variable X , denote x is sampled from (the distribution of) X as $x \leftarrow X$; the support of X is defined as $\text{supp}(X) \stackrel{\text{def}}{=} \{x : \Pr[X = x] > 0\}$.

Partial functions There are many ways to view a partial function $f : [N] \rightarrow [M] \cup \{\perp\}$ with domain $Q \stackrel{\text{def}}{=} \{w \in [N] : f(w) \neq \perp\}$: It can be viewed as a function $f_Q : Q \rightarrow [M]$, or a list $((w_i, f(w_i))_{i \in [Q]}$. We say two partial functions are *consistent* if they agree on the intersection of their domains. For consistent partial functions f_1 and f_2 , we use $f_1 \cup f_2$ to denote the partial function with domain $Q_1 \cup Q_2$ and is consistent with f_1 and f_2 .

2.2 Key-agreement protocols

Let $\Pi = (A, B)$ be a two-party protocol consisting of a pair of probabilistic interactive Turing machines, where the two parties A and B are often referred to as Alice and Bob. A protocol is called ℓ -*oracle-aided* if Alice and Bob have access to an oracle $f : [N] \rightarrow [M]$ and each party makes at most ℓ queries to f . An oracle-aided protocol is called *non-adaptive* when both parties choose their queries before querying the oracle and communicating. Π produces a *transcript* τ which is the communication bits between players. The *communication complexity* of Π , denoted by $\text{CC}(\Pi)$, is the length of the transcript of Π in the worst case.

We focus on oracle-aided key-agreement protocols in the random oracle model, where the oracle f is uniformly sampled from the collection of all functions from $[N]$ to $[M]$. Note that the execution of the key-agreement protocol is completely determined by r_A, r_B and f , where r_A (resp., r_B) is Alice's (resp., Bob's) internal randomness. We call the tuple (r_A, r_B, f) an *extended view*. Let $\text{EV} = (\mathbf{R}_A, \mathbf{R}_B, F)$ denote the distribution of the extended view in a random execution. For every extended view $v = (r_A, r_B, f)$, let $\text{tran}(v), \text{out}_A(v), \text{out}_B(v)$ be the communication transcript, A 's output, and B 's output respectively, given the extended view v .

Definition 2.1 (Key-agreement protocols). Let $\alpha, \gamma \in [0, 1], q \in \mathbb{N}$. A protocol $\Pi = (\mathbf{A}, \mathbf{B})$ is a (α, q, γ) -*key-agreement* if the following conditions hold:

1. $(1 - \alpha)$ -*completeness*. $\Pr_{v \leftarrow EV} [\text{out}_A(v) = \text{out}_B(v)] \geq 1 - \alpha$.
2. (q, γ) -*security*. For any q -oracle-aided adversary \mathbf{E} ,

$$\Pr_{v=(r_A, r_B, f) \leftarrow EV} \left[\mathbf{E}^f(\text{tran}(v)) = \text{out}_A(v) \right] \leq \gamma.$$

Since we aim to prove lower bounds, we assume each party outputs one bit, as per [HMO⁺19]. Moreover, [HMO⁺19] proved that studying the following normalized key-agreement protocols suffices.

Normalized key-agreement protocols. Following [HMO⁺19], to simplify the proof of the lower bound, we can transform the key-agreement protocol Π into a normalized protocol called Π' , such that the secret key output by Bob in Π' is the first bit of his last query. Formally,

Proposition 2.2. *Let Π be a non-adaptive, ℓ -oracle-aided (α, q, γ) -key-agreement protocol with communication complexity c . Then there is a non-adaptive $(\ell + 1)$ -oracle-aided (α, q, γ) -key-agreement protocol Π' with communication complexity $c + 1$, in which Bob's output is the first bit of his last query.*

2.3 Basic information theory

The Shannon entropy of a random variable X is defined as

$$\mathbf{H}(X) \stackrel{\text{def}}{=} \sum_{x \in \text{supp}(X)} \Pr[X = x] \log \left(\frac{1}{\Pr[X = x]} \right).$$

The conditional entropy of a random variable X given Y is defined as

$$\mathbf{H}(X | Y) \stackrel{\text{def}}{=} \mathbf{E}_{y \leftarrow Y} [\mathbf{H}(X | Y = y)].$$

We often use (conditional) entropy conditioned on some event E , which is defined by the same formula where the probability measure $\Pr[\cdot]$ is replaced by $\Pr'[\cdot] \stackrel{\text{def}}{=} \Pr[\cdot|E]$. Entropy conditioned on event E is denoted as $\mathbf{H}(X|E), \mathbf{H}(X|Y, E)$.

Let X and Y be two (possibly correlated) random variables. The mutual information of X and Y is defined by

$$\mathbf{I}(X; Y) \stackrel{\text{def}}{=} \mathbf{H}(X) - \mathbf{H}(X | Y) = \mathbf{H}(Y) - \mathbf{H}(Y | X).$$

The conditional mutual information is:

$$\mathbf{I}(X_i; Y | X_1, \dots, X_{i-1}) \stackrel{\text{def}}{=} \mathbf{H}(X_i | X_1, \dots, X_{i-1}) - \mathbf{H}(X_i | Y, X_1, \dots, X_{i-1})$$

Proposition 2.3 (Entropy chain rule). *For random variables X_1, X_2, \dots, X_n , it holds that*

$$\mathbf{H}(X_1, X_2, \dots, X_n) = \sum_{i=1}^n \mathbf{H}(X_i | X_1, X_2, \dots, X_{i-1}).$$

Proposition 2.4. (Chain rule for mutual information) For X_1, X_2, \dots, X_n are n random variables and Y is another random variable,

$$I(X_1, X_2, \dots, X_n; Y) = \sum_{i=1}^n I(X_i; Y \mid X_1, X_2, \dots, X_{i-1}).$$

Proposition 2.5. (Data processing inequality) For two random variables X, Y and a function f ,

$$H(f(X)) \leq H(X) \text{ and } I(f(X); Y) \leq I(X; Y)$$

3 Communication complexity of key-agreement protocols

This section proves the main theorem:

Theorem 3.1 (Formal version of theorem 1.4). Let $\Pi = (A, B)$ be an ℓ -query-aided, non-adaptive $(0, q, \gamma)$ -key-agreement (i.e., Π enjoys perfect completeness), then

$$\text{CC}(\Pi) \geq \frac{q}{2(\ell+1)} \cdot \frac{(1-\gamma)^3}{27} - 1 = \Omega\left(\frac{q}{\ell}\right).$$

By proposition 2.2, it suffices to show that

$$\text{CC}(\Pi) \geq \frac{q}{2\ell} \cdot \frac{(1-\gamma)^3}{27}, \tag{2}$$

for all normalized protocol Π that satisfies the conditions in theorem 3.1.

Correlated sets play a central role in our proof; here we give the formal definition.

Definition 3.2 (ε -correlated). Let τ be a transcript and f_E be a partial function with domain Q_E . We say a set $S \subseteq [N]$ is ε -*correlated* with respect to (τ, f_E) if

$$H(F(S) \mid R_A, R_B, F(Q_E) = f_E) - H(F(S) \mid R_A, R_B, F(Q_E) = f_E \wedge T = \tau) \geq \varepsilon,$$

where (R_A, R_B, F) is a random extended view and $T \stackrel{\text{def}}{=} \text{tran}(R_A, R_B, F)$.

3.1 Description of the attacker

The attacker is described in alg. 1. In the algorithm, $f_E^{(i)}$ stands for the observations of Eve till the end of the i -th iteration. Moreover, we use $EV(\tau, f_E^{(i)})$ to denote the distribution of the extended view EV conditioned on the following two events: (1) the random oracle is consistent with $f_E^{(i)}$; (2) the transcript is τ .

Algorithm 1: The attacker E

Input: transcript τ
Oracle : $f : [N] \rightarrow [M]$
Output: $b \in \{0, 1, \perp\}$
Set $\varepsilon := (1 - \gamma)^2/9$
Initialize $i := 0$ and $f_E^{(0)}$ as the empty function
while $\exists S \subseteq [N]$ **s.t.** $|S| \leq 2\ell$ **and is ε -correlated w.r.t.** $(\tau, f_E^{(i)})$ **do**
 Let S_{i+1} be any ε -correlated set of size at most 2ℓ
 Query f on S_{i+1} and receive $f(S_{i+1})$
 Set $f_E^{(i+1)} := f_E^{(i)} \cup f(S_{i+1})$.
 $i := i + 1$
if $\exists b \in \{0, 1\}$ **s.t.** $\Pr_{v \leftarrow EV(\tau, f_E^{(i)})} [\text{out}_A(v) = b] \geq 1 - \sqrt{2\varepsilon}$ **then**
 Output b
else
 Output \perp

3.2 Success probability of the attacker

This subsection analyzes the attacker's success probability for perfect completeness. We will first introduce the language of the combinatorial rectangle and then use it to analyze the attacker's success probability.

3.2.1 Through the lens of rectangles

Combinatorial rectangle is a standard tool in communication complexity. We thus develop this language for key-agreement protocols in the following.

Let Π be a non-adaptive key-agreement protocol, meaning that queries of Alice is a function $Q_A(r_A)$ of her internal randomness r_A . If f_A is a partial function with domain $Q_A(r_A)$, we call the pair (r_A, f_A) a *profile* of Alice. The profile space of Alice, denoted by \mathcal{X} , consists of all possible profiles of Alice, namely,

$$\mathcal{X} \stackrel{\text{def}}{=} \{(r_A, f_A) : f_A \text{ is a partial function with domain } Q_A(r_A)\}.$$

For Bob, we analogously define Q_B and $\mathcal{Y} \stackrel{\text{def}}{=} \{(r_B, f_B) : f_B \text{ is a partial function with domain } Q_B(r_B)\}$. Given a profile pair $(x = (r_A, f_A), y = (r_B, f_B)) \in \mathcal{X} \times \mathcal{Y}$, Alice and Bob can run the protocol by using f_A and f_B respectively as oracle answers: when Alice needs to issue an oracle query w , she takes $f_A(w)$ as oracle answer; similarly, Bob takes $f_B(w)$ as oracle answer when querying w . Hence, we can still define the transcript $\text{tran}(x, y)$ and output $\text{out}_A(x, y), \text{out}_B(x, y)$.

Note that some profile pairs are imaginary in the sense that the oracle answers of Alice and Bob are inconsistent. We say $x = (r_A, f_A) \in \mathcal{X}$ and $y = (r_B, f_B) \in \mathcal{Y}$ are *consistent* if f_A and f_B are consistent. Define the output table $\mathcal{M}_\Pi \in \{0, 1, *\}^{\mathcal{X} \times \mathcal{Y}}$ via

$$\mathcal{M}_\Pi(x, y) \stackrel{\text{def}}{=} \begin{cases} \text{out}_A(x, y), & \text{if } x, y \text{ are consistent;} \\ *, & \text{otherwise.} \end{cases}$$

Let $D \stackrel{\text{def}}{=} \{(x, y) \in \mathcal{X} \times \mathcal{Y} : \mathcal{M}_\Pi(x, y) \neq *\}$ be the set of all consistent profile pairs; such profile pairs can be witnessed in real execution.

A set $R \subseteq \mathcal{X} \times \mathcal{Y}$ is called a *rectangle* if $R = \mathcal{X}_R \times \mathcal{Y}_R$ for some $\mathcal{X}_R \subseteq \mathcal{X}$ and $\mathcal{Y}_R \subseteq \mathcal{Y}$. Let τ be a transcript and f_E be a partial function with domain Q_E . We care about the profiles that are consistent with f_E and produce transcript τ ; formally, we consider the rectangle $\mathcal{X}_{\tau, f_E} \times \mathcal{Y}_{\tau, f_E}$ where

$$\mathcal{X}_{\tau, f_E} \stackrel{\text{def}}{=} \{x = (r_A, f_A) \in \mathcal{X} : \exists y = (r_B, f_B) \in \mathcal{Y} \text{ s.t. } f_A, f_B, f_E \text{ are consistent and } \text{tran}(x, y) = \tau\},$$

and

$$\mathcal{Y}_{\tau, f_E} \stackrel{\text{def}}{=} \{y = (r_B, f_B) \in \mathcal{Y} : \exists x = (r_A, f_A) \in \mathcal{X} \text{ s.t. } f_A, f_B, f_E \text{ are consistent and } \text{tran}(x, y) = \tau\}.$$

If Π has perfect completeness, the rectangle $\mathcal{X}_{\tau, f_E} \times \mathcal{Y}_{\tau, f_E}$ has the following simple but useful property.

Claim 3.3. *Assume that Π has perfect completeness. Let $(x, y), (x', y') \in \mathcal{X}_{\tau, f_E} \times \mathcal{Y}_{\tau, f_E}$ for some τ and f_E . If $\mathcal{M}_\Pi(x, y) = 0$ and $\mathcal{M}_\Pi(x', y') = 1$, then $\mathcal{M}_\Pi(x, y') = \mathcal{M}_\Pi(x', y) = *$.*

Proof. Assume $\mathcal{M}_\Pi(x, y') \neq *$. Since (x, y') appears in some execution of Π , by perfect completeness, we have $\text{out}_A(x, y') = \text{out}_B(x, y')$. However, $\text{out}_A(x, y') = \text{out}_A(x, y) = 0$ while $\text{out}_B(x, y') = \text{out}_B(x', y) = 1$, a contradiction. The argument for (x', y) is similar. \square

Let $\mathcal{QV}(\tau, f_E)$ denote the query set of Alice and Bob conditioned on (τ, f_E) , namely, $\mathcal{QV}(\tau, f_E) \stackrel{\text{def}}{=} (\mathcal{Q}_A(\mathcal{R}_A), \mathcal{Q}_B(\mathcal{R}_B))$, where $(\mathcal{R}_A, \mathcal{R}_B, \cdot) = \text{EV}(\tau, f_E)$. Given $(Q_A, Q_B) \in \text{supp}\mathcal{QV}(\tau, f_E)$, we obtain a subrectangle of $\mathcal{X}_{\tau, f_E} \times \mathcal{Y}_{\tau, f_E}$ by adding the restriction that Alice's (resp., Bob's) queries is Q_A (resp., Q_B). That is, we consider $\mathcal{X}_{\tau, f_E}(Q_A) \times \mathcal{Y}_{\tau, f_E}(Q_B)$ where

$$\mathcal{X}_{\tau, f_E}(Q_A) \stackrel{\text{def}}{=} \{x = (r_A, f_A) \in \mathcal{X}_{\tau, f_E} : \mathcal{Q}_A(r_A) = Q_A\},$$

$$\mathcal{Y}_{\tau, f_E}(Q_B) \stackrel{\text{def}}{=} \{y = (r_B, f_B) \in \mathcal{Y}_{\tau, f_E} : \mathcal{Q}_B(r_B) = Q_B\}.$$

Definition 3.4 (Monochromatic Rectangle). A rectangle $R \subseteq \mathcal{X} \times \mathcal{Y}$ is *b-monochromatic* if $R \cap D \neq \emptyset$ and for every $(x, y) \in R \cap D$, $\mathcal{M}_\Pi(x, y) = b$; R is said to be *monochromatic* if it is *b-monochromatic* for some $b \in \{0, 1\}$.

The following lemma shows that if the protocol is normalized and has perfect completeness, the rectangle $\mathcal{X}_{\tau, f_E} \times \mathcal{Y}_{\tau, f_E}$ has a special structure: It can be partitioned into monochromatic rectangles according to the queries.

Lemma 3.5. *Suppose Π is normalized and has perfect completeness. Let τ be a transcript and f_E be a partial function. For all $(Q_A, Q_B) \in \text{supp}\mathcal{QV}(\tau, f_E)$, the rectangle $\mathcal{X}_{\tau, f_E}(Q_A) \times \mathcal{Y}_{\tau, f_E}(Q_B)$ is monochromatic.*

Proof. Since Π is normalized, for any $(x, y) \in \mathcal{X}_{\tau, f_E}(Q_A) \times \mathcal{Y}_{\tau, f_E}(Q_B)$, $\text{out}_B(x, y)$ is determined by Q_B . Moreover, because of perfect completeness, $\text{out}_A(x, y) = \text{out}_B(x, y)$ for all $(x, y) \in \mathcal{X}_{\tau, f_E}(Q_A) \times \mathcal{Y}_{\tau, f_E}(Q_B)$. Thus, $\mathcal{X}_{\tau, f_E}(Q_A) \times \mathcal{Y}_{\tau, f_E}(Q_B)$ is monochromatic. \square

3.2.2 Analyzing the attacker's success probability

Next, we show that alg. 1 breaks the security of normalized protocols. The following lemma characterizes what happens after all small ε -correlated sets are cleaned up; it roughly says that if there exists no small ε -correlated set, the key is almost determined conditioned on Eve's information.

Lemma 3.6. *Let τ be a transcript and f_E be a partial function with domain Q_E . If there exists no ε -correlated set of size at most 2ℓ w.r.t. (τ, f_E) , then $\exists b \in \{0, 1\}$ s.t. $\Pr_{v \leftarrow EV(\tau, f_E)} [\text{out}_A(v) = b] \geq 1 - \sqrt{2\varepsilon}$.*

Proof. Write $\delta \stackrel{\text{def}}{=} \sqrt{2\varepsilon}$. Assume towards contradiction that

$$\Pr_{v \leftarrow EV(\tau, f_E)} [\text{out}_A(v) = b] > \delta, \forall b \in \{0, 1\}.$$

For $b \in \{0, 1\}$, define

$$\mathcal{G}_b \stackrel{\text{def}}{=} \{(Q_A, Q_B) \in \text{supp} QV(\tau, f_E) : \mathcal{X}_{\tau, f_E}(Q_A) \times \mathcal{Y}_{\tau, f_E}(Q_B) \text{ is } b\text{-monochromatic}\}.$$

By lemma 3.5, $\forall b \in \{0, 1\}$,

$$\Pr_{v \leftarrow EV(\tau, f_E)} [(Q_A(v), Q_B(v)) \in \mathcal{G}_b] = \Pr_{v \leftarrow EV(\tau, f_E)} [\text{out}_A(v) = b] > \delta. \quad (3)$$

For $Q = (Q_A, Q_B)$, $Q' = (Q'_A, Q'_B)$, define

$$h(Q, Q') \stackrel{\text{def}}{=} \mathbf{H}(F(Q_A \cup Q_B) \mid F(Q_E) = f_E) - \mathbf{H}(F(Q_A \cup Q_B) \mid Q_A(R_A) = Q'_A \wedge Q_B(R_B) = Q'_B \wedge F(Q_E) = f_E \wedge T = \tau),$$

where (R_A, R_B, F) is a random extended view and $T = \text{tran}(R_A, R_B, F)$ as usual. Then, we have

Claim 3.7. *For all $Q_0 = (Q_A^0, Q_B^0) \in \mathcal{G}_0$ and $Q_1 = (Q_A^1, Q_B^1) \in \mathcal{G}_1$, we have $h(Q_b, Q_{1-b}) \geq 1$ for some $b \in \{0, 1\}$.*

The above claim suggests some kind of correlation with the transcript exists; next, we prove such correlation gives rise to an ε -correlated set.

Consider the following complete bipartite graph, denoted by G :

1. The left vertex set is V_0 and each vertex $v \in V_0$ is associated with some $Q(v) \in \mathcal{G}_0$.
2. The right vertex set is V_1 and each vertex $v \in V_1$ is associated with some $Q(v) \in \mathcal{G}_1$.
3. For each $Q \in \mathcal{G}_0 \cup \mathcal{G}_1$, the number of vertices associated with Q is proportional to $\Pr_{QV(\tau, f_E)} [Q]$.

We assign an orientation to G as follows: for all $v_0 \in V_0, v_1 \in V_1$, if $h(Q(v_0), Q(v_1)) \geq 1$, then the edge $\{v_0, v_1\}$ is directed towards v_1 ; otherwise, $\{v_0, v_1\}$ is directed towards v_0 . Let $E(G)$ denote the set of all directed edges. By claim 3.7, each directed edge $v \rightarrow v'$ satisfies $h(Q(v), Q(v')) \geq 1$. Let $\Gamma(v) \stackrel{\text{def}}{=} \{v' : (v \rightarrow v') \in E(G)\}$ denote the set of out-neighbors of v . WLOG, assume that $|V_0| \leq |V_1|$. By average argument, there exists some $v^* \in V_0 \cup V_1$ such that $|\Gamma(v^*)| \geq \frac{|V_0| \cdot |V_1|}{|V_0| + |V_1|} \geq |V_0|/2$. Say $v^* \in V_{b^*}$, then we have

$$\Pr_{v \leftarrow V_{1-b^*}} [(v^* \rightarrow v) \in E(G)] = \frac{|\Gamma(v^*)|}{|V_{1-b^*}|} \geq \frac{|V_0|}{2|V_{1-b^*}|} = \frac{1}{2} \cdot \frac{\Pr_{v \leftarrow EV(\tau, f_E)} [(Q_A(v), Q_B(v)) \in \mathcal{G}_0]}{\Pr_{v \leftarrow EV(\tau, f_E)} [(Q_A(v), Q_B(v)) \in \mathcal{G}_{1-b^*}]} > \frac{\delta}{2}.$$

Let $Q^* \stackrel{\text{def}}{=} Q(v^*)$. Then we have

$$\begin{aligned}
\mathbf{E}_{Q \leftarrow \mathcal{QV}(\tau, f_E)} [h(Q^*, Q)] &\geq \mathbf{E}_{Q \leftarrow \mathcal{QV}(\tau, f_E)} [h(Q^*, Q) \mid Q \in \mathcal{G}_{1-b^*}] \Pr_{Q \leftarrow \mathcal{QV}(\tau, f_E)} [Q \in \mathcal{G}_{1-b^*}] \\
&\geq \mathbf{E}_{v \leftarrow V_{1-b^*}} [h(Q(v^*), Q(v))] \cdot \delta \\
&\geq \Pr_{v \leftarrow V_{1-b^*}} [(v^* \rightarrow v) \in E(G)] \cdot \delta \\
&= \frac{\delta^2}{2} = \varepsilon,
\end{aligned} \tag{4}$$

where the second inequality follows from eq. (3) and the construction of G , and the third inequality holds because $h(Q(v^*), Q(v)) \geq \mathbb{1}[(v^* \rightarrow v) \in E(G)]$.

Note that $\mathbf{E}_{Q \leftarrow \mathcal{QV}(\tau, f_E)} [h(Q^*, Q)] \geq \varepsilon$ means that

$$\mathbf{H}(F(Q_A^* \cup Q_B^*) \mid F(Q_E) = f_E) - \mathbf{H}(F(Q_A^* \cup Q_B^*) \mid \mathcal{Q}_A(R_A), \mathcal{Q}_B(R_B), F(Q_E) = f_E \wedge T = \tau) \geq \varepsilon,$$

where $Q^* = (Q_A^*, Q_B^*)$. Thus, letting $\widehat{Q} = Q_A^* \cup Q_B^*$, we have

$$\begin{aligned}
&\mathbf{H}(F(\widehat{Q}) \mid R_A, R_B, F(Q_E) = f_E) - \mathbf{H}(F(\widehat{Q}) \mid R_A, R_B, F(Q_E) = f_E \wedge T = \tau) \\
&\geq \mathbf{H}(F(\widehat{Q}) \mid R_A, R_B, F(Q_E) = f_E) - \mathbf{H}(F(\widehat{Q}) \mid \mathcal{Q}_A(R_A), \mathcal{Q}_B(R_B), F(Q_E) = f_E \wedge T = \tau) \\
&= \mathbf{H}(F(\widehat{Q}) \mid F(Q_E) = f_E) - \mathbf{H}(F(\widehat{Q}) \mid \mathcal{Q}_A(R_A), \mathcal{Q}_B(R_B), F(Q_E) = f_E \wedge T = \tau) \\
&\geq \varepsilon,
\end{aligned}$$

where the first inequality is by data processing inequality and the second step holds as $F(\widehat{Q})$, R_A , R_B are independent. That is, \widehat{Q} is ε -correlated w.r.t. (τ, f_E) , a contradiction. \square

Proof of claim 3.7. Define

$$R_b \stackrel{\text{def}}{=} \mathcal{X}_{\tau, f_E}(Q_A^b) \times \mathcal{Y}_{\tau, f_E}(Q_B^b) \text{ where } b \in \{0, 1\}.$$

For all $(x, y) \in R_0$, $(x', y') \in R_1$, we have $\mathcal{M}_\Pi(x, y) = 0$ and $\mathcal{M}_\Pi(x', y') = 1$, and hence $\mathcal{M}_\Pi(x, y') = *$ according to claim 3.3. This means that oracle answers in profile x and profile y' are inconsistent.

Note that all inconsistent queries are in $S \stackrel{\text{def}}{=} Q_A^0 \cap Q_B^1$. Therefore,

$$\text{supp}(F(S) \mid \mathcal{Q}_A(R_A) = Q_A^0 \wedge \mathcal{Q}_B(R_B) = Q_B^0 \wedge T = \tau \wedge F(Q_E) = f_E) \cap \text{supp}(F(S) \mid \mathcal{Q}_A(R_A) = Q_A^1 \wedge \mathcal{Q}_B(R_B) = Q_B^1 \wedge T = \tau \wedge F(Q_E) = f_E) = \emptyset.$$

A simple average argument shows that for some $b^* \in \{0, 1\}$,

$$\left| \text{supp}(F(S) \mid \mathcal{Q}_A(R_A) = Q_A^{b^*} \wedge \mathcal{Q}_B(R_B) = Q_B^{b^*} \wedge T = \tau \wedge F(Q_E) = f_E) \right| \leq \frac{|\text{supp}(F(S) \mid F(Q_E) = f_E)|}{2}. \tag{5}$$

Consequently,

$$\begin{aligned}
\Delta &\stackrel{\text{def}}{=} \mathbf{H}(F(S) \mid F(Q_E) = f_E) - \mathbf{H}(F(S) \mid \mathcal{Q}_A(R_A) = Q_A^{b^*} \wedge \mathcal{Q}_B(R_B) = Q_B^{b^*} \wedge F(Q_E) = f_E \wedge T = \tau) \\
&\geq \mathbf{H}(F(S) \mid F(Q_E) = f_E) - \log \left| \text{supp}(F(S) \mid \mathcal{Q}_A(R_A) = Q_A^{b^*} \wedge \mathcal{Q}_B(R_B) = Q_B^{b^*} \wedge T = \tau \wedge F(Q_E) = f_E) \right| \\
&\geq \log \left| \text{supp}(F(S) \mid F(Q_E) = f_E) \right| - \log \frac{|\text{supp}(F(S) \mid F(Q_E) = f_E)|}{2} \\
&= 1,
\end{aligned}$$

where the second inequality follows from eq. (5) and the fact that $F(S)|_{F(Q_E)=f_E}$ is uniform distribution.

Now that $\Delta \geq 1$, it suffice to show $h(Q_{1-b^*}, Q_{b^*}) \geq \Delta$. Since $S \subseteq Q_A^{1-b^*} \cup Q_B^{1-b^*}$, this follows from chain rule:

$$\begin{aligned} & h(Q_{1-b^*}, Q_{b^*}) - \Delta \\ &= \mathbf{H}\left(F(\bar{S}) \mid F(Q_E) = f_E\right) - \mathbf{H}\left(F(\bar{S}) \mid F(S), Q_A(R_A) = Q_A^{b^*} \wedge Q_B(R_B) = Q_B^{b^*} \wedge F(Q_E) = f_E \wedge T = \tau\right) \\ &\geq 0, \end{aligned}$$

where $\bar{S} \stackrel{\text{def}}{=} (Q_A^{1-b^*} \cup Q_B^{1-b^*}) \setminus S$ and the inequality holds since $F(\bar{S})|_{F(Q_E)=f_E}$ is uniform distribution (and uniform distribution has maximum entropy). \square

Corollary 3.8 (Accuracy of E). *Let Π be an ℓ -oracle-aided, non-adaptive $(1, q, \gamma)$ -key-agreement. Assume the Π is normalized, then alg. 1 guesses the key correctly with probability at least $1 - \sqrt{2\varepsilon}$, i.e.,*

$$\Pr_{v=(r_A, r_B, f) \leftarrow \text{EV}} \left[\mathbf{E}^f(\text{tran}(v)) = \text{out}_A(v) \right] > 1 - \sqrt{2\varepsilon}.$$

Proof. By lemma 3.6, E outputs $\text{out}_A(v)$ except with probability less than $\sqrt{2\varepsilon}$. \square

3.3 Efficiency of the attacker

In this subsection, we analyze the efficiency of the attacker Eve (alg. 1) via the density increment argument [YZ22, YZ23]. We first introduce the density function. Intuitively, the density function $\Phi(\tau, f_E)$ captures the amount of hidden information contained in the transcript τ about the random function F given Eve's observation of oracle f_E . As Eve makes effective queries, she learns (a constant amount of) information in each iteration, so the density function decreases by a constant.

Definition 3.9 (Density function). Let τ be a transcript and f_E be a partial function with domain Q_E . Define density function Φ via

$$\Phi(\tau, f_E) \stackrel{\text{def}}{=} \mathbf{H}(F \mid R_A, R_B, F(Q_E) = f_E) - \mathbf{H}(F \mid R_A, R_B, F(Q_E) = f_E \wedge T = \tau),$$

where (R_A, R_B, F) is a random extended view and $T \stackrel{\text{def}}{=} \text{tran}(R_A, R_B, F)$.

Lemma 3.10. *The density function Φ satisfies the following properties:*

1. Φ is non-negative.
2. $\mathbf{E}_{\tau \leftarrow T} [\Phi(\tau, f_\emptyset)] \leq \text{CC}(\Pi)$, where f_\emptyset denotes the empty function.
3. If S is ε -correlated w.r.t. (τ, f_E) , then $\mathbf{E}_{f_S \leftarrow F(S) \mid T=\tau, F(Q_E)=f_E} [\Phi(\tau, f_E \cup f_S)] \leq \Phi(\tau, f_E) - \varepsilon$.

Proof. We prove these statements as follows.

1. F is uniform distribution given R_A, R_B and conditioned on $f(Q_E) = f_E$. Hence Φ is non-negative.

2. By definition, we have that

$$\begin{aligned}
\mathbf{E}_{\tau \leftarrow T} [\Phi(\tau, f_\emptyset)] &= \mathbf{E}_{\tau \leftarrow T} [\mathbf{H}(F \mid R_A, R_B) - \mathbf{H}(F \mid R_A, R_B, T = \tau)] \\
&= \mathbf{H}(F \mid R_A, R_B) - \mathbf{H}(F \mid R_A, R_B, T) \\
&= \mathbf{I}(F; T \mid R_A, R_B) \\
&\leq \mathbf{H}(T) \\
&\leq \text{CC}(\Pi).
\end{aligned}$$

3. Write $Q'_E \stackrel{\text{def}}{=} Q_E \cup S$. We decompose $\Phi(\tau, f_E) - \mathbf{E}_{f_S \leftarrow F(S) \mid T=\tau, F(Q_E)=f_E} [\Phi(\tau, f_E \cup f_S)] = \phi_1 - \phi_2$, where

$$\phi_1 \stackrel{\text{def}}{=} \mathbf{H}(F \mid R_A, R_B, F(Q_E) = f_E) - \mathbf{E}_{f_S \leftarrow F(S) \mid T=\tau, F(Q_E)=f_E} [\mathbf{H}(F \mid R_A, R_B, F(Q'_E) = (f_E \cup f_S))],$$

and

$$\phi_2 \stackrel{\text{def}}{=} \mathbf{H}(F \mid R_A, R_B, F(Q_E) = f_E \wedge T = \tau) - \mathbf{E}_{f_S \leftarrow F(S) \mid \tau, f_E} [\mathbf{H}(F \mid R_A, R_B, F(Q'_E) = (f_E \cup f_S) \wedge T = \tau)].$$

Since R_A, R_B, F are independent, we have (by chain rule)

$$\phi_1 = \mathbf{H}(F(S) \mid R_A, R_B, F(Q_E) = f_E).$$

Observe that by the definition of conditional entropy,

$$\mathbf{E}_{f_S \leftarrow F(S) \mid \tau, f_E} [\mathbf{H}(F \mid R_A, R_B, F(Q'_E) = (f_E \cup f_S) \wedge T = \tau)] = \mathbf{H}(F \mid R_A, R_B, F(S), F(Q_E) = f_E \wedge T = \tau).$$

By the chain rule,

$$\begin{aligned}
\phi_2 &= \mathbf{H}(F \mid R_A, R_B, F(Q_E) = f_E \wedge T = \tau) - \mathbf{H}(F \mid R_A, R_B, F(S), F(Q_E) = f_E \wedge T = \tau) \\
&= \mathbf{H}(F(S) \mid R_A, R_B, F(Q_E) = f_E \wedge T = \tau).
\end{aligned} \tag{6}$$

Since S is ε -correlated, $\mathbf{H}(F(S) \mid R_A, R_B, F(Q_E) = f_E) - \mathbf{H}(F(S) \mid R_A, R_B, F(Q_E) = f_E \wedge T = \tau) \geq \varepsilon$, and hence

$$\Phi(\tau, f_E) - \mathbf{E}_{f_S \leftarrow F(S) \mid T=\tau, F(Q_E)=f_E} [\Phi(\tau, f_E \cup f_S)] = \phi_1 - \phi_2 \geq \varepsilon.$$

□

Following lemma 3.10, we can deduce that our attacker E (alg. 1) makes at most $\text{CC}(\Pi)/\varepsilon$ iterations in expectation.

Lemma 3.11 (Efficiency of E). $\mathbf{E}[\# \text{ of iterations in the running of } E] \leq \frac{\text{CC}(\Pi)}{\varepsilon}$.

Proof. Recall the sampling procedure in section 1.1. Then, we define some random variables in a random execution for analysis. Let $F_E^{(i)} = F_E^{(i-1)} \cup F(S_i)$ be the observations of Eve until the end of the i -th iteration, where $F_E^{(0)}$ is the empty function. If E does not enter the i -th iteration, we define $F_E^{(i)} = F_E^{(i-1)}$. Define a counter variable to record the number of iterations as follows: $C_0 \stackrel{\text{def}}{=} 0$ and for $i \geq 0$

$$C_{i+1} \stackrel{\text{def}}{=} \begin{cases} C_i + 1, & \text{if } E \text{ enters the } i\text{-th iteration;} \\ C_i, & \text{otherwise.} \end{cases}$$

We claim that for every τ and f_E ,

$$\mathbf{E} \left[\underbrace{\Phi(T, F_E^{(i)}) - \Phi(T, F_E^{(i+1)}) - \varepsilon(C_{i+1} - C_i)}_{\geq 0} \mid T = \tau \wedge F_E^{(i)} = f_E \right] \geq 0. \quad (7)$$

To see this, consider the event $\text{Enter}_i \stackrel{\text{def}}{=} \text{'E enters the } i\text{-th iteration'}$. Conditioned on Enter_i , $C_{i+1} - C_i = 1$ and by the third item of lemma 3.10, the underlined part is non-negative; conditioned on $\neg \text{Enter}_i$, the underlined part equals zero by definition.

Since eq. (7) holds for all (τ, f_E) , we get $\mathbf{E} \left[\Phi(T, F_E^{(i)}) - \Phi(T, F_E^{(i+1)}) - \varepsilon(C_{i+1} - C_i) \right] \geq 0$. Summing over $i = 0, \dots, N - 1$, we obtain

$$\mathbf{E}[\Phi(T, F_E^{(0)})] - \mathbf{E}[\Phi(T, F_E^{(N)})] - \varepsilon \mathbf{E}[C_N - C_0] \geq 0.$$

By the first and second items of lemma 3.10, we have $\mathbf{E}[\Phi(T, F_E^{(N)})] \geq 0$ and $\mathbf{E}[\Phi(T, F_E^{(0)})] \leq \text{CC}(\Pi)$. Note that $C_0 = 0$ and C_N equals the total number of iterations because there can never be more than N iterations. Therefore, we get

$$\mathbf{E}[\# \text{ of iterations in the running of E}] = \mathbf{E}[C_N] \leq \frac{\text{CC}(\Pi)}{\varepsilon}.$$

□

So far, we have bounded the expected number of iterations of alg. 1 from above; however, alg. 1 could make too many queries in the worst case. To prove our main theorem, we need an attacker who makes a bounded number of queries in the worst case. We construct such an attacker by running E for a limited number of iterations.

Theorem 3.12. *Let E' be an attacker who runs E but aborts when the number of iterations exceeds $\frac{\text{CC}(\Pi)}{\varepsilon^{3/2}}$. Then the following statements hold:*

1. *Efficiency: E' makes at most $q_{E'} = 2\ell \cdot \text{CC}(\Pi)/\varepsilon^{3/2}$ oracle queries.*
2. *Accuracy: The success probability of E' is at least γ .*

Proof. Efficiency holds because E' queries at most $\text{CC}(\Pi)/\varepsilon^{3/2}$ sets and each set has size at most 2ℓ . As for accuracy, let β, β' be the success probability of E, E' respectively. By the definition of E' , we have

$$\begin{aligned} |\beta' - \beta| &\leq \Pr[E' \text{ aborts}] \\ &= \Pr \left[\# \text{ of iterations in the running of E is more than } \text{CC}(\Pi)/\varepsilon^{3/2} \right]. \end{aligned}$$

Lemma 3.11 together with Markov's inequality shows that this quantity is at most $\sqrt{\varepsilon}$. Therefore, we have $\beta' \geq \beta - \sqrt{\varepsilon}$. By the accuracy of E (corollary 3.8) and our choice of ε (i.e., $\varepsilon = (1 - \gamma)^2/9$), we obtain $\beta' \geq 1 - \sqrt{2\varepsilon} - \sqrt{\varepsilon} > 1 - 3\sqrt{\varepsilon} = \gamma$. □

Proving the main theorem Theorem 3.1 immediately follows from the above lemma.

Proof of theorem 3.1. Let Π be a protocol that satisfies the conditions of theorem 3.1. It suffices to prove $\text{CC}(\Pi) \geq \frac{q}{2\ell} \cdot \frac{(1-\gamma)^3}{27}$ (eq. (2)), provided that Π is normalized. Since E' in theorem 3.12 succeeds with probability γ and Π is a (q, γ) -secure by assumption, we must have $q_{E'} > q$, which implies

$$\text{CC}(\Pi) > \frac{q}{2\ell} \cdot \varepsilon^{3/2} = \frac{q}{2\ell} \cdot \frac{(1-\gamma)^3}{27}.$$

□

Acknowledgements

We are grateful to the anonymous reviewers for their insightful comments and suggestions, which have significantly enhanced the presentation of this work.

References

- [ACC⁺22] Per Austrin, Hao Chung, Kai-Min Chung, Shiuan Fu, Yao-Ting Lin, and Mohammad Mahmoody. On the impossibility of key agreements from quantum random oracles. In *Advances in Cryptology—CRYPTO 2022: 42nd Annual International Cryptology Conference, CRYPTO 2022, Santa Barbara, CA, USA, August 15–18, 2022, Proceedings, Part II*, pages 165–194. Springer, 2022. 2, 8
- [ACMS23] Abtin Afshar, Geoffroy Couteau, Mohammad Mahmoody, and Elahe Sadeghi. Fine-grained non-interactive key-exchange: Constructions and lower bounds. In *Advances in Cryptology—EUROCRYPT 2023: 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Lyon, France, April 23–27, 2023, Proceedings, Part I*, pages 55–85. Springer, 2023. 2
- [AHMS20] Benny Applebaum, Thomas Holenstein, Manoj Mishra, and Ofer Shayevitz. The communication complexity of private simultaneous messages, revisited. *Journal of Cryptology*, 33(3):917–953, 2020. 2
- [BBCR10] Boaz Barak, Mark Braverman, Xi Chen, and Anup Rao. How to compress interactive communication. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 67–76, 2010. 6
- [BKSY11] Zvika Brakerski, Jonathan Katz, Gil Segev, and Arkady Yerukhimovich. Limits on the power of zero-knowledge proofs in cryptographic constructions. In *Theory of Cryptography: 8th Theory of Cryptography Conference, TCC 2011, Providence, RI, USA, March 28–30, 2011. Proceedings 8*, pages 559–578. Springer, 2011. 8
- [BMG09] Boaz Barak and Mohammad Mahmoody-Ghidary. Merkle puzzles are optimal—an $O(n^2)$ -query attack on any key exchange from a random oracle. In *Advances in cryptology—CRYPTO 2009*, volume 5677 of *Lecture Notes in Comput. Sci.*, pages 374–390. Springer, Berlin, 2009. 2, 4, 5, 7, 8

- [CN22] Shahar P Cohen and Moni Naor. Low communication complexity protocols, collision resistant hash functions and secret key-agreement protocols. *Cryptology ePrint Archive*, 2022. [2](#)
- [Cou19] Geoffroy Couteau. A note on the communication complexity of multiparty computation in the correlated randomness model. In *Advances in Cryptology—EUROCRYPT 2019: 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19–23, 2019, Proceedings, Part II* 38, pages 473–503. Springer, 2019. [2](#)
- [CSWY01] Amit Chakrabarti, Yaoyun Shi, Anthony Wirth, and Andrew Yao. Informational complexity and the direct sum problem for simultaneous message complexity. In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, pages 270–278. IEEE, 2001. [6](#)
- [DH76] W. Diffie and M. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, 1976. [1](#), [2](#)
- [DH21] Itai Dinur and Ben Hasson. Distributed merkle puzzles. In *Theory of Cryptography: 19th International Conference, TCC 2021, Raleigh, NC, USA, November 8–11, 2021, Proceedings, Part II* 19, pages 310–332. Springer, 2021. [2](#), [8](#)
- [DSLMM11] Dana Dachman-Soled, Yehuda Lindell, Mohammad Mahmoody, and Tal Malkin. On the black-box complexity of optimally-fair coin tossing. In *Theory of Cryptography: 8th Theory of Cryptography Conference, TCC 2011, Providence, RI, USA, March 28–30, 2011. Proceedings* 8, pages 450–467. Springer, 2011. [2](#), [8](#)
- [GPW15] Mika Göös, Toniann Pitassi, and Thomas Watson. Deterministic communication vs. partition number. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 1077–1088. IEEE, 2015. [7](#)
- [GPW17] Mika Göös, Toniann Pitassi, and Thomas Watson. Query-to-communication lifting for bpp. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 132–143. IEEE, 2017. [7](#)
- [HHRS15] Iftach Haitner, Jonathan J Hoch, Omer Reingold, and Gil Segev. Finding collisions in interactive protocols—tight lower bounds on the round and communication complexities of statistically hiding commitments. *SIAM Journal on Computing*, 44(1):193–242, 2015. [2](#)
- [HMO⁺19] Iftach Haitner, Noam Mazon, Rotem Oshman, Omer Reingold, and Amir Yehudayoff. On the communication complexity of key-agreement protocols. In *10th Innovations in Theoretical Computer Science*, volume 124 of *LIPICs. Leibniz Int. Proc. Inform.*, pages Art. No. 40, 16. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2019. [2](#), [3](#), [7](#), [9](#)
- [HOZ16] Iftach Haitner, Eran Omri, and Hila Zarosim. Limits on the usefulness of random oracles. *Journal of Cryptology*, 29(2):283–335, 2016. [8](#)

- [IR89] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *Proceedings of the twenty-first annual ACM symposium on Theory of computing*, pages 44–61, 1989. 2, 4
- [KSY11] Jonathan Katz, Dominique Schröder, and Arkady Yerukhimovich. Impossibility of blind signatures from one-way permutations. In *Theory of Cryptography: 8th Theory of Cryptography Conference, TCC 2011, Providence, RI, USA, March 28-30, 2011. Proceedings 8*, pages 615–629. Springer, 2011. 8
- [Maz23] Noam Mazor. Key-agreement with perfect completeness from random oracles. *Cryptology ePrint Archive*, 2023. 3
- [Mer78] Ralph C Merkle. Secure communications over insecure channels. *Communications of the ACM*, 21(4):294–299, 1978. 1, 2
- [MMV11] Mohammad Mahmoody, Tal Moran, and Salil Vadhan. Time-lock puzzles in the random oracle model. In *Advances in Cryptology–CRYPTO 2011: 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings 31*, pages 39–50. Springer, 2011. 8
- [MP12] Mohammad Mahmoody and Rafael Pass. The curious case of non-interactive commitments—on the power of black-box vs. non-black-box use of primitives. In *Advances in Cryptology–CRYPTO 2012: 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*, pages 701–718. Springer, 2012. 8
- [RM97] Ran Raz and Pierre McKenzie. Separation of the monotone nc hierarchy. In *Proceedings 38th Annual Symposium on Foundations of Computer Science*, pages 234–243. IEEE, 1997. 7
- [YZ22] Guangxu Yang and Jiapeng Zhang. Simulation methods in communication complexity, revisited. In *Electron. Colloquium Comput. Complex., TR22-019*, 2022. 3, 6, 7, 15
- [YZ23] Guangxu Yang and Jiapeng Zhang. Lifting theorems meet information complexity: Known and new lower bounds of set-disjointness. In *Manuscript*, 2023. 3, 6, 15