Unconditional Security using (Random) Anonymous Bulletin Board

Abstract. In a seminal work, Ishai et al. (FOCS–2006) study the viability of designing unconditionally secure protocols for key agreement and multi-party computation (MPC) using an anonymous bulletin board (ABB) as a building block. While their results establish the feasibility of key agreement and honest-majority MPC in the ABB model, the optimality of protocols in terms of communication is not studied. This paper enhances the study of unconditional security in the ABB model in multiple ways.

- We present a key agreement protocol with a novel combinatorial insight to offer a 100% increase in throughput over the (FOCS–2006) study; i.e., using the same amount of communication, we can almost double the bit-length of the agreed key. We also demonstrate near optimality of our approach.
- We also offer unconditionally secure protocols for the binary erasure channel and oblivious transfer problems. Our protocols employ one-shot access to an ABB and a single message from a sender and a receiver. We demonstrate the round optimality of our constructions.

1 Introduction

Securely realizing unconditionally secure cryptographic primitives is a topic of immense value and has a rich history. This work revisits a particularly surprising work by Ishai et al. [19] that analyzes the possibility of realizing cryptography with unconditional security by assuming access to an *anonymous bulletin board* (ABB). Ishai et al. establish unconditional security for prominent cryptographic tasks such as key agreement and honest-majority multiparty computation (MPC) based solely on access to an ABB that allows a sender to publish her message without revealing her identity. In particular, they demonstrate that ABB is sufficient to implement unconditionally secure point-to-point channels between two parties without making any other assumption. Ishai et al. then extend it to achieve MPC with unconditional security in the presence of an honest majority. Interestingly, on the negative side, they also show that anonymous broadcast cannot offer unconditional security when there is no honest majority in an MPC computation.

Since the publication of the paper by Ishai et al. in 2006, the field of anonymous communication has witnessed tremendous growth: the anonymous communication network Tor [12] serves more than two million unique users daily using an overlay network of several thousand nodes all over the Internet. As the use of blockchains brings users financial dealing to the (public) Internet, there have

been significant efforts towards introducing and improving anonymity over the Internet. Startups such Nym [11] and xx.network [34,35] are developing generic anonymous communication networks to break the link between users' identity and their transactions, and several blockchain projects have started incorporating anonymous communications, such as Tor and I2P, in their designs. [18] Academic literature on anonymous communication as well as protocol implementations have significantly expanded in the last two decades [1, 5, 10, 13, 23]. It is safe to say that ABBs are prevalent on the Internet today. And our goal is to understand the efficacy and efficiency of developing cryptography assuming access to such as an ABB.

The utility of the ABB towards unconditional security is easy to explain using the simple key agreement protocol between Alice and Bob against an honest-but-curious adversary. In this protocol proposed by Ishai et al. [19], Alice and Bob both pick a random integer each (say r_A and r_B respectively) and publish those to the anonymous broadcast channel. The agreed single secret bit can be 0 if $r_A > r_B$ and 1 if $r_A < r_B$. If $r_A = r_B$, rerun the protocol. Notice that Alice and Bob know their respective input and thus can compute the secret bit; however, others cannot distinguish r_A from r_B and have no information about the agreed bit.

This "indistinguishability property" can be abstracted as a multi-set. Conceptually, we observe that the use of ABB converts a vector (or key-value store) of user inputs to a multi-set. This allows us to ask the question, what if Alice and Bob send multiple (say k) messages each? Can we agree on more than k bit using this $2k$-sized multi-set? We answer this question affirmatively to demonstrate that Alice and Bob can indeed agree on close to $2k$ secret bits, which improve the throughput of the key agreement by 100%, as compared to Ishai et al. [19]. The aim of this work is to make cryptography based on anonymity *optimal* w.r.t. various communication measures, such as the round and communication complexity.

Our work is a quantitative investigation – emphasizing optimal round and communication complexity. To this end, we establish connections of implementing functionalities using ABB in our context with various well-known communication-limited MPC models (like *non-interactive correlation distillation* [2, 21, 29, 30] and *one-way secure computation* [15]). Our problems translate into analytically tractable instances of these MPC models, which have been challenging to analyze in general. These connections lead us to several (near-optimal) protocol constructions. Our practically motivated problems lead to new research questions in these MPC models, potentially of interest to the broader community.

1.1 Our Contributions

From the modeling perspective, this work assumes the existence of an anonymous broadcast, which we model as an Anonymous Bulletin Board (ABB). There are four parties $\mathcal{A}, \mathcal{B}, \mathcal{C}$, and $\mathcal{D}$. The bulletin board ideal functionality, represented as $\mathsf{ABB}_{m_A,m_B,m_C}$, takes as input three multi-sets: (1) $A := \{a_1, \ldots, a_{m_A}\}$ from party $\mathcal{A}$, (2) $B := \{b_1, \ldots, b_{m_B}\}$ from $\mathcal{B}$, and (3) $C := \{c_1, \ldots, c_{m_C}\}$ from $\mathcal{C}$.

Note that party $\mathcal{D}$ does not provide any input. The functionality outputs the multi-set $\Gamma = A \cup B \cup C := \{\gamma_1, \ldots, \gamma_{m_A+m_B+m_C}\}$ to all four parties. In particular, it does not output a set containing the three multi-sets; but outputs their set union. We call party $\mathcal{C}$ the helper and party $\mathcal{D}$ an eavesdropper. In the randomized version of bulletin board (rABB), the three multi-sets A, B, C are sampled according to some distributions P, Q, R respectively. See Section 3 for a more formal definition of ABB as well as its randomized version (rABB).

We summarize our contribution in the following three corollaries.

Corollary 1 (Key-agreement Protocol: Informal). *We present a secure two-party key-agreement protocol using* $\mathsf{rABB}_{m,m}$ *with message length* n *that establishes (near-optimal)* $2m$*-bit keys with* mn*-bit communication complexity.*

Theorem 1 provides the formal statement for this corollary. We achieve security against the eavesdropper $\mathcal{D}$ with unbounded computational power. We prove that our protocol achieves near-optimal key length. The achieved key length in our protocol is close to the maximal key length in any protocol using rABB *even when allowing arbitrary interactions* (over private authenticated channel). We prove this fact using mutual information, entropy-based arguments, and the results in [24, 25]. Our result implies that the message length n has negligible impact on the key's length, all that matters are the number of messages m. Note that message length n helps in achieving a smaller failure probability. We also present a duplicate-recovery variant of the protocol in Corollary 1 that is suitable for different parameters.

Compared to previous state-of-the-art [19], our protocol approaches a 100% increase in throughput for large values of k as shown in Figure 1.

Corollary 2 (Binary Erasure Channel: Informal). *We present a round-optimal protocol securely realizing a binary erasure channel from* $\mathcal{B}$ *to* $\mathcal{A}$ *utilizing the* rABB *(with a helper) and a private authenticated channel from* $\mathcal{A}$ *and* $\mathcal{B}$*.*

Theorem 6 provides the formal statement of this corollary. Out protocol achieves unconditional security against semi-host adversary. Our protocol uses only one round of communication from party $\mathcal{B}$ to $\mathcal{A}$. For a rational erasure probability e/d, where $\gcd(e, d) = 1$, the communication cost of our protocol is proportional to the denominator d. Determining the minimum communication cost remains open. We prove the round optimality of our protocol by showing that the realizing BEC in the rABB-hybrid without communication is impossible. We employ the techniques developed recently in the SNIS/SNIR leterature [2,21] to prove this impossibility.

Ishai et al. showed that it is impossible to realize oblivious transfer (OT) with statistical security in the dishonest majority settings. This implies that it is impossible to realize OT (as well as randomized input OT) in the ABB without the helper $\mathcal{C}$. We prove that randomized OT can be securely realized in the rABB with the helper.

Corollary 3. *We present a one-round protocol for establishing binary oblivious transfer between* $\mathcal{A}$ *and* $\mathcal{B}$ *utilizing the* rABB *with the helper* $\mathcal{C}$ *and private authenticated channel between parties* $\mathcal{A}$ *and* $\mathcal{B}$*.*

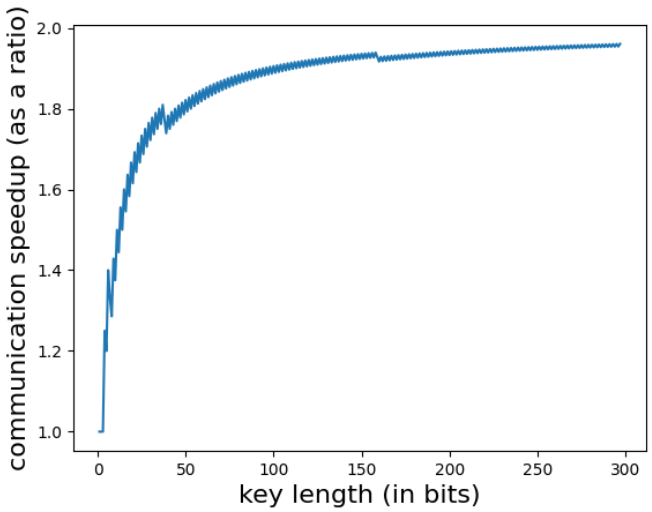

Fig. 1. We plot the ratio of communication required by [19] over our protocol's communication for various values of key-length k. For example, for 50-bit keys, [19] requires roughly 1.8× our communication cost.

Theorem 8 provides the formal statement of this corollary. This protocol achieves unconditional security against semi-host adversary. The protocol is similar to the BEC protocol and can be seen as an implementation of two correlated BEC.

1.2 Related Works

Key-Agreement There are many works focused on key-agreement or developing secure point-to-point links based on anonymous communication. [4] performs key agreements by having each party send a set of position-labeled bits, and discard any identical bits to use the remaining bits as the key. [31] considers key-agreement when the receivers (instead of the senders) are anonymous. [14] considers key-agreement in a similar setting, where a "deck of cards" is dealt such that each party has several cards from the deck. The remaining cards are dealt to the adversary. Using this setup, the parties would like to agree on a secret key. Finally, Gilad and Herzberg [16] demonstrate the practical utility of [19] for the IP-level security protocol IPSec.

There has been extensive study of establishing fixed length secret key in the source model in which parties observe i.i.d samples from a joint distribution and the eavesdropper possibly observes some side information from these samples [3, 9, 17, 26–28]. The main objective is to study the achievable key rate when the number of samples tend to infinity.

[20] study the question of boostrapping anonymous communication. The objective is to communicate a large amount of data using non-anonymous communication and only a small amount of anonymous broadcasts.

Non-interactive Correlation Distillation In information theory and theoretical computer science, non-interactive correlation distillation (NICD) is a well-studied analytically-tractable problem [6,7,29,30,36]. NICD also aims for establishing secure key agreements. In NICD, each party holds a noise version of some source bits, a particular form of correlated private randomness. It is common in NICD that the failure probability for the key-agreement instances is high. On the other hand, in our model, parties have access to ABB that generates a different form of conditional distribution. We are the first to choose this distribution and achieve near-optimal key-length.

1.3 Outline of Our Paper

The rest of the paper is organized as follows. We first start by introducing notations and backgrounds in Section 2. Then Section 3 formally defines the anonymous bulletin board and its variants, which we use throughout the rest of the paper. The first protocol we present is our key-agreement protocol in Section 4. Then in Section 5 we consider constructions of binary erasure channel BEC. Section 6 presents our one-round protocols realizing random oblivious transfer ROT.

2 Preliminaries

This section introduces some notations and basic background that will be useful in the later sections.

2.1 Sets

Throughout the paper, we may use the word "set" when we mean "multi-set". We will often use capital letters to denote multi-sets. Alternatively, we will define a multi-set by listing its elements in curly braces. We denote elements of the multi-set using lowercase letters. Whenever we talk about multi-sets, especially the union of sets, we assume that all elements are randomized. For example, for multi-set A and B, by only looking at $A \cup B$, it should be impossible to determine which elements came from A. Additionally, we assume that there exists some canonical ordering of elements in a multi-set. For simplicity, it may help to assume that sets are automatically sorted by their canonical ordering. We note that properties such as not being able to determine which elements came from A by only looking at $A \cup B$ still hold.

2.2 Binary Erasure Channel (BEC)

In a binary erasure channel (BEC) with erasure probability p, a sender S sends a binary message $m \in \{0,1\}$ to a receiver R. With probability $1-p$, R receives the message m. With probability p, the message is "erased" and R receives nothing. In this case, we say that R receives $\perp$. The sender is unaware of whether the erasure happened or not.

Security of BEC requires that the sender does not learn whether the bit was erased or not, and the receiver not learning anything about the bit if it is erased (and it received $\perp$).

2.3 Oblivious Transfer

Oblivious Transfer. Oblivious transfer, denoted as OT is, is a two-party functionality that takes as input $(x_0, x_1) \in \{0,1\}^2$ from Alice, a bit $b \in \{0,1\}$ from Bob, and outputs x_b to Bob. Security of OT requires that Alice learns nothing about the bit b, and Bob learns nothing about x_{1-b}.

Randomized Oblivious Transfer. Random oblivious transfer, denoted as ROT, is a correlation that samples x_0, x_1, b uniformly and independently at random. It provides Alice the secret share $r_A = (x_0, x_1)$ and provides Bob the secret share $r_B = (b, x_b)$.

2.4 Entropy and Mutual Information

We shall use mutual information and entropy-based arguments to prove the optimality of our key-agreement protocols.

Definition 1 (Mutual Information). *Let X and Y be a pair of discrete random variables over the space $\mathcal{X} \times \mathcal{Y}$. If their joint probability distribution is $P_{XY}(x,y)$, the mutual information between them, denoted as $I(X,Y)$, is*

$$I(X,Y) := \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} P_{XY}(x,y) \log \frac{P_{XY}(x,y)}{P_X(x) P_Y(y)}.$$

Moreover, the conditional mutual information of $(X, Y|Z)$ is defined as follows.

$$I(X,Y|Z) := \sum_{z \in \mathcal{Z}} P_Z(z) \sum_{y \in \mathcal{Y}} \sum_{x \in \mathcal{X}} P_{X,Y|Z}(x,y|z) \log \frac{P_{X,Y|Z}(x,y|z)}{P_{X|Z}(x|z) P_{Y|Z}(y|z)}.$$

Definition 2 (Entropy). *Let X be a discrete random variable distributed according to $P \colon \mathcal{X} \to (0,1)$. The entropy of X, denoted as $H(X)$, is defined as*

$$H(X) := -\sum_{x \in \mathcal{X}} P(x) \log P(x).$$

Definition 3 (Conditional Entropy). *The conditional entropy of X given Y is defined as*

$$H(X|Y) := -\sum_{x\in\mathcal{X},y\in\mathcal{Y}} P_{XY}(x,y)\log\frac{P_{XY}(x,y)}{P_Y(y)}.$$

Fact 1 *It holds that*

$$I(X,Y) = H(X) - H(X|Y) = H(Y) - H(Y|X).$$

Furthermore,

$$I(X,Y|Z) = H(X|Z) - H(X|Y,Z) = H(Y|Z) - H(Y|X,Z).$$

3 Anonymous Bulletin Board Formalism

In this section, we first formally define Anonymous Bulletin Board (ABB) and Random Anonymous Bulletin Board (rABB). We then discuss the similarities and differences between ABB and rABB, and crucially, show that our protocol is equivalent in the ABB-hybrid and the rABB-hybrid. Additionally, we present a connection of rABB with the offline phase of the offline-online model.

3.1 Anonymous Bulletin Board

We assume all messages are from the domain $\mathbb{Z}_{2^n}$, where $n \in \mathbb{N}$ is called the *message length*. In an anonymous bulletin board (ABB), parties can privately send multiple messages to the ABB. Then, the ABB will broadcast the "set" of messages, with order and sender information removed. Additionally, we note that the ABB waits until it receives all messages before publishing. Therefore, a rushing adversary is impossible.

We define the ideal functionality in Figure 2. It is defined over a four-party setting, $\mathcal{A}$ and $\mathcal{B}$ represent the main participants that are trying to achieve something through interaction with the ABB. $\mathcal{C}$ represents the facilitators that are trying to assist $\mathcal{A}$ and $\mathcal{B}$ through interaction with the ABB. $\mathcal{D}$ represents eavesdroppers that do not send messages to the ABB, but receive the output of the ABB.

We define this ideal functionality as $\mathsf{ABB}_{m_A,m_B,m_C}$. Formally, $\mathsf{ABB}_{m_A,m_B,m_C}$ takes multi-set of inputs $A := \{a_1,\ldots,a_{m_A}\}$ from $\mathcal{A}$, multi-set of inputs $B := \{b_1,\ldots,b_{m_B}\}$ from $\mathcal{B}$, multi-set of inputs $C := \{c_1,\ldots,c_{m_C}\}$ from $\mathcal{C}$, and outputs the multi-set $\Gamma = A \cup B \cup C := \{\gamma_1,\ldots,\gamma_{m_A+m_B+m_C}\}$. In particular, note that there are no "links" or associations between the different messages that A sends.

We note that this model is more powerful than the random public anonymous bulletin board functionality presented next since ABB allows messages to be chosen adaptively, that is, dependent on previous messages.

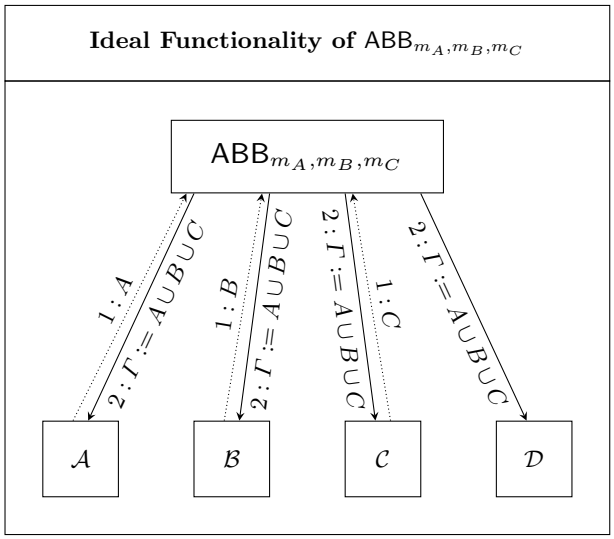

Fig. 2. Ideal Functionality of Anonymous Bulletin Board (ABB). Dotted lines represent private authenticated channels, while solid lines represent public channels. The number in front of messages shows the order/round in which the messages are sent. A is the multi-set $\{a_1, \ldots, a_{m_A}\}$, B and C are similarly defined.

3.2 Random Anonymous Bulletin Board

We also define a random anonymous bulletin board (rABB) in Figure 3, which takes additional parameters P, Q, R, which are independent distributions, and samples set of messages A according to distribution P, set of messages B according to distribution Q, set of messages C according to distribution R, privately outputs A to $\mathcal{A}$, B to $\mathcal{B}$, C to $\mathcal{C}$ (using private authenticated channels denoted with dashed lines), and outputs the multi-set Γ to all parties (using public channels denoted with solid lines). We define the ideal functionality as $\mathsf{rABB}^{P,Q,R}_{m_A,m_B,m_C}$.

3.3 Comparison between ABB and rABB

We note that the rABB functionality is as powerful as the ABB functionality when messages are not chosen adaptively. Essentially, since messages are not chosen adaptively, the parties should be able to determine the distribution of the messages they want to send before interacting with ABB. Therefore, parties can simply program a rABB with the appropriate P, Q, R in order to mimic the messages they are going to send. We note that P, Q, R can be distributions with sample space of size 1. Effectively making them deterministic.

In all of our protocols, parties sample their inputs uniformly at random (without replacement) and send them to the ABB. We note that this is precisely equivalent to rABB with P, Q, R set to output values uniformly at random without

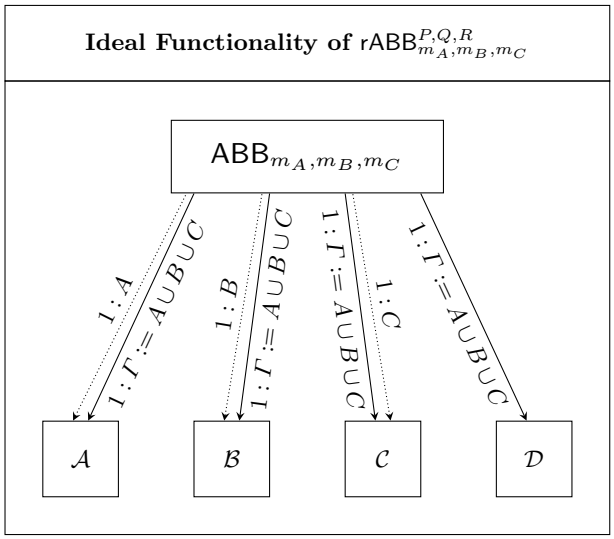

Fig. 3. Ideal Functionality of Random Anonymous Bulletin Board (rABB). Dotted lines represent private authenticated channels, while solid lines represent public channels. The number in front of messages shows the order/round in which the messages are sent. A is the multi-set $\{a_1, \ldots, a_{m_A}\}$ sampled according to P, B is sampled according to Q, and C is sampled according to R.

replacement. Intuitively, since the messages are chosen at random by semi-honest parties, it does not matter if the parties chose them and sent them to ABB, or if rABB chose them and sent them to the parties. The end result is both the ABB/rABB and parties will have the same random values.

3.4 An Equivalent Reformulation of rABB

In the semi-honest setting, the random anonymous bulletin board can be reformulated as the preprocessing step (offline phase) of the offline-online paradigm. In this step, parties will receive private correlated randomness from a conditional distribution.

Reformulation of rABB$^{P,Q,R}_{m_A,m_B,m_C}$. The bulletin board samples (A, B, C) from the conditional distribution $(P, Q, R)|\Gamma)$, where $A = \{a_1, a_2, \ldots, a_{m_A}\}, B = \{b_1, b_2, \ldots, b_{m_B}\}$, and $C = \{c_1, c_2, \ldots, c_{m_C}\}$. Alice receives the multi-set A, Bob receives B, and $\mathcal{C}$ receives C. The bulletin board sends $\Gamma = A \cup B \cup C$ to all parties $\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{D}$. So, party $\mathcal{D}$ gets some side information about the correlated randomness of $\mathcal{A}, \mathcal{B}, \mathcal{C}$. We note that (A, B) is sampled according to the conditional distributions $(P, Q|\Gamma)$.

Discussion. Unlike the standard offline-online model in which correlated private randomness is sampled from a joint distribution, our model samples them from a conditional distribution, which is a family of joint distributions.

4 Key Agreement Protocols

This section presents our optimal length key agreement protocols in the rABB-hybrid in which party $\mathcal{C}$ does not send any messages to the bulletin board. We start by defining the problem setting.

4.1 Problem Setting

Suppose parties are in $\mathsf{rABB}^{P,Q}_{m_A,m_B}$-hybrid (without the helper $\mathcal{C}$). That is, party $\mathcal{A}$ has $A = \{a_1, a_2, \ldots, a_{m_A}\}$ sampled according to P, party $\mathcal{B}$ has $B = \{b_1, b_2, \ldots, b_{m_B}\}$ sampled according to Q, and party $\mathcal{D}$ has $A \cup B$. Every message is n-bit. Parties $\mathcal{A}$ and $\mathcal{B}$ are allowed to communicate with each other through a public noiseless channel. Party $\mathcal{D}$ (the eavesdropper) can see the messages sent between $\mathcal{A}$ and $\mathcal{B}$. At the end of the protocol, two parties $\mathcal{A}$ and $\mathcal{B}$ agree on a sample space Ω_L for the key, and $\mathcal{A}$ outputs a (variable-length) key $K_A \in \Omega_L$, and $\mathcal{B}$ outputs a key $K_B \in \Omega_L$. So, the key length is $\log \Omega_L$, where log denotes the logarithmic with base two. We define the *expected key length* of the protocol as the expectation of $\log \Omega_L$, where the expectation is taken over the randomness of samples A and B.

The protocol is correct if $K_A = K_B = K$ with high probability and K is close to a uniform distribution over Ω_L. It is secure if the eavesdropper $\mathcal{D}$ learns almost nothing about the key K. More formally, the statistical distance between two distributions $(K_A, K_B, T, A \cup B)$ and $(U_\Omega, U_\Omega, T, A \cup B)$ is small, where T is the transcript of the protocol (messages sent between $\mathcal{A}$ and $\mathcal{B}$).

Given a desired key length k, we define the *failure probability* as

$$\min\left(\Pr[\mathsf{length}(K_A) < k], \Pr[\mathsf{length}(K_B < k)]\right),$$

where the probabilities are taken over the randomness of K_A and K_B, respectively. Typically, the failure probability is negligible. We define the *communication cost* as $m \cdot n + \mathsf{length}(T)$, where $\mathsf{length}(T)$ denotes the length of the protocol's transcript.

Objectives. In this work, we focus on the following objectives. Given a desired length $k \in \mathbb{N}$ and a failure probability δ, we are interested in constructing key agreement protocols that output a (variable-length) key of length at least k with failure probability at most δ and with least communication costs.

Remark. Our protocols always achieve perfect correctness and perfect secrecy even when the key length is less than the desired threshold k. Our problem setting is similar to the key agreement model considered in [25]. The main difference is that parties get components A and B, respectively, of a conditional distribution of the form $(P, Q|Z)$ in our setting; while parties get A and B sampled according to a joint distribution of the form (P, Q) in their setting. Similar to the setting considered in another line of work [3, 26, 27], the eavesdropper has some side information about the samples of $\mathcal{A}$ and $\mathcal{B}$. The difference is that parties have access to multiple i.i.d samples in their setting, while parties have access to only one sample of a (large) joint distribution.

Remark. In our problem setting, parties agree on a variable-length key. One can rely on the asymptotic equipartition property and apply standard extraction procedures to obtain a fixed-length key.

4.2 Our Protocol

Next, we present our protocol in Figure 4.2, as well as provide an overview of the protocol below.

The protocol is similar to the example presented in [19]. At a high level, parties $\mathcal{A}$ and $\mathcal{B}$ will each receive m random values from the rABB, which are sampled from P and Q respectively. For security, $P = Q$, and for optimal performance, P is a uniform distribution over $(\mathbb{Z}_{2^n})^m$ under the constraint that elements do not repeat. That is, P produces a set $A := \{a_1, \ldots, a_m\}$ such that all elements are equally likely to be in the set, and that for all $i \neq j, a_i \neq a_j$. Q produces a similar set B. Once they see the set of values, $\mathcal{A}$ and $\mathcal{B}$ can easily distinguish between $\mathcal{A}$'s values and $\mathcal{B}$'s values, while the eavesdropper cannot. Using this information, $\mathcal{A}$ and $\mathcal{B}$ can agree on a key k determined by the positions of $\mathcal{A}$'s values. Intuitively, there are $\binom{2m}{m}$ possible cases of which ones are $\mathcal{A}$'s values, and only $\mathcal{A}$ and $\mathcal{B}$ can identify one of the $\binom{2m}{m}$ possible cases.

We also specify an algorithm to determine the value of k. We do so by assuming a canonical ordering of the elements in the set, and the more $\mathcal{A}$'s values are towards the "front" of the set, the higher the value of k is.

Algorithm 1:

Parameters:
Input : A', Γ'
Output : k

We assume there exists some canonical ordering of elements in a set.
Identifies which of $\gamma'_1, \ldots, \gamma'_{2m'}$ are in the set A', and create a bitstring x such that if the i^{th} element is in A', then the i^{th} bit is set to 1.
Compute $m' = size(A')$
$k = 0$
$c = m'$
for i *in* range$(2 \cdot m')$*:* **do**
    if $x\,[i] == 1$ **then**
        $k = k +$ binomial$(2 \cdot m' - i - 1\ ,\ c)$
        $c = c - 1$
return k

Explanation of Algorithm 1 The idea behind algorithm 1 is to first convert the set of elements into a binary string using the canonical ordering and assigning 1s at $\mathcal{A}$'s inputs and 0s to $\mathcal{B}$'s input. The algorithm can determine which bit string

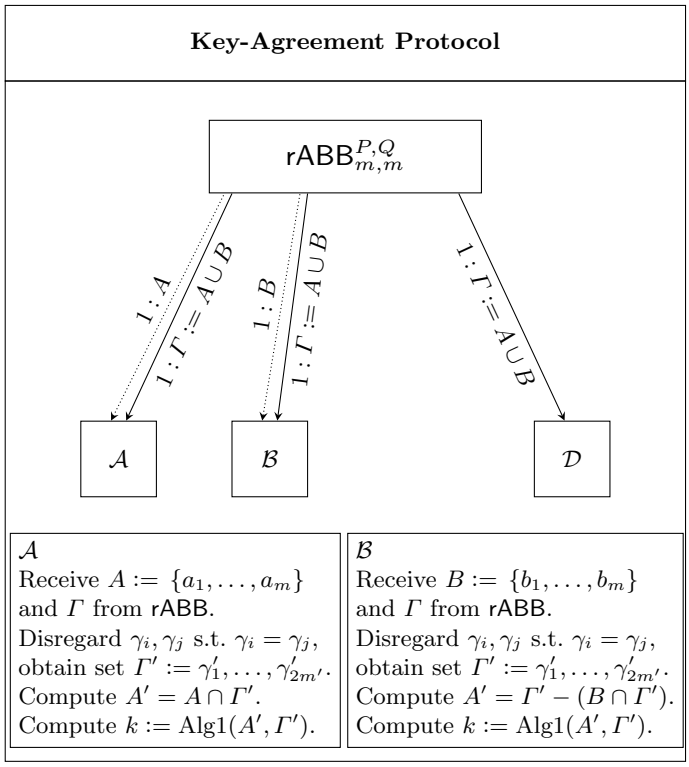

Fig. 4. Key-Agreement between parties $\mathcal{A}$ and $\mathcal{B}$ in presence of an eavesdropper $\mathcal{D}$

this is by examing all the 1 bits that appear and counting how many bit strings are skipped. For example, let us look at a simple case of $m = 3$, and the bit-string being 101010. Upon seeing the first 1, the algorithm knows all bit-strings of the form $0 * * * * *$ have been skipped over, where the $* * * * *$ consists of 3 1s, and 2 0s. There are $\binom{5}{3}$ of them. Then, upon seeing the next 1, the algorithm knows all bit-strings of the form $100 * * *$ have been skipped over, where the $* * *$ represents 2 1s and 1 0s. There are $\binom{3}{2}$ of them. By continuing through the entire bit-string, the algorithm can determine the value of this bit-string, thus determining the value of the key k.

4.3 Performance Analysis and Parameter Choices

Now, we use concentration bounds to show that the expected key length in our protocol is highly concentrated around the mean that is large. Therefore, our protocol outputs a long key with high probability.

Theorem 1. *Fix $m, n \in \mathbb{N}$, and $0 \leqslant \varepsilon \leqslant 1$. Our protocol in Figure 4 uses m messages (each being n-bit strings) to help parties agree on a key of length*

at least $k = \log\binom{2m'}{m'} = (2m') \cdot (1 - o(1))$ *with failure probability at most* $\exp\left(-\varepsilon^2 \cdot m \cdot \left(1 - \frac{m}{2^n}\right)/2\right)$, *where* $m' = m \cdot \left(1 - \frac{m}{2^n}\right) \cdot (1 - \varepsilon)$.

Proof Sketch. Let us fix the elements of $\mathcal{B}$. For each element of $\mathcal{A}$, there is at least $\left(1 - \frac{m}{2^n}\right)$ probability that such an element of $\mathcal{A}$ will not be a duplicate of an element of $\mathcal{B}$. Therefore, we can apply a simplified Chernoff Bound. For $2m'$ unique elements, $\mathcal{A}$ and $\mathcal{B}$ can agree on a $\log\binom{2m'}{m'} \approx \log\left(\frac{4^m}{\sqrt{m\pi}}\right)$-bit key.

Theorem 2. *Fix* $m, n \in \mathbb{N}$, *satisfying* $m = o\left(2^{n/3}\right)$. *Using* m *messages of size* n*-bit each, our protocol in Figure 4 allows parties to agree on a* $k = \log\binom{2m}{m} = 2m \cdot (1 - o(1))$ *bit key with probability at least* $1 - \exp\left(-\frac{m^2}{2 \cdot 2^n}\right)$.

Proof Sketch. By the birthday bound, with probability at least $1 - \exp\left(-\frac{m^2}{2 \cdot 2^n}\right)$, all $2m$ values will be unique. For $2m$ unique elements, $\mathcal{A}$ and $\mathcal{B}$ can agree on a $\log\binom{2m}{m}$-bit key.

Parameter Choices. Our protocol has two main parameters, m and n. The expected key length increases with m and n. At the same time, the communication cost increases with m and n as well. However, we note that this is a simple optimization problem and that automatic searches for optimal parameters can be done. Furthermore, for common key-length such as 128 or 256 bits, the search only has to be performed once.

We perform this automated search and present our results in Figure 5. Concretely, using 702-bits of communication, $\mathcal{A}$ and $\mathcal{B}$ can agree on a 128-bit key in expectation, and using 1550-bits of communication, $\mathcal{A}$ and $\mathcal{B}$ can agree on a 256-bit key in expectation.

4.4 On the Optimality of Key Length

This section shows that our protocol in *Figure* 4 is near optimal in terms of the expected key length. In fact, we will prove a much stronger result. That is, the expected key length of any interactive protocol for key agreement in the $\mathsf{rABB}^{P,Q}_{m,m}$ is at most $2m - \mathsf{poly}(\log \mathsf{m})$, where P, Q are arbitrary independent distributions over $(\mathbb{Z}_{2^n})^m$, and n is the message length. By Theorem 1, for any $\varepsilon > 0$, our protocol achieves $2(1 - \varepsilon)m$-bit key length with an exponentially small failure probability (depending on ε). This means that our *non-interactive protocol* asymptotically achieves the optimal key length of the best interactive protocol. We provide detailed proof below.

First, we upper bound the expected key length by the mutual information. Then, we show that the mutual information of any $\mathsf{rABB}^{P,Q}_{m,m}$ is at most $\log\binom{2m}{m}$.

Theorem 3. *Let* $m, n \in \mathbb{N}$ *and* P, Q *be independent distributions over* $(\mathbb{Z}_{2^n})^m$ *Suppose parties are in the random public anonymous bulletin board hybrid* $\mathsf{rABB}^{P,Q}_{m,m}$. *Then, the expected key length in any key agreement protocol (allowing interaction) is at most* $I(\mathsf{rABB}^{P,Q}_{m,m}) + 1 + \log 3$.

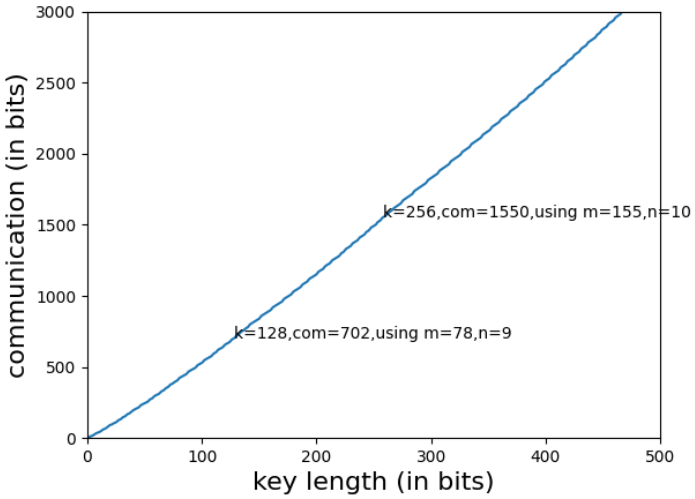

Fig. 5. For various key lengths k, we plot the communication required (in bits) such that the expected key length is at least k. We also label two points, one for 128-bit keys and one for 256-bit keys. We label the key length and communication required, as well as the parameter choices of m and n used.

We shall employ the techniques developed recently in [24, 25] to prove the theorem above. We say that Alice and Bob are in (X, Y)-correlation hybrid if Alice has x and Bob has y, where (x, y) is sampled according to the joint distribution (X, Y). The following result shall be useful for the proof.

Theorem 4. *[24, 25] Let (X, Y) be a joint distribution. Then, the maximal expected key length in the (X, Y)-correlation hybrid (allowing an arbitrary amount of communication) is at most $I(X, Y) + 1 + \log 3$.*

Proof Sketch of Theorem 3. Recall the reformulation of rABB as correlated private randomness in Section 3.4. The correlation $\mathsf{rABB}^{P,Q}_{m,m}$ is a conditional distribution of the form $(X, Y|Z)$, where Z the random variable denoting the eavesdropper's view (the set $A \cup B \cup C$). Conditioned on fixing the eavesdropper's view $(Z = z)$, applying Theorem 4 to the joint distribution $(X, Y|Z = z)$ yields that the key length is at most $I(X, Y|Z = z) + 1 + \log 3$. Thus, the expected key length is at most

$$\mathbb{E}_z[I(X, Y|Z = z) + 1 + \log 3] = I(X, Y|Z) + 1 + \log 3.$$

Next, we bound the mutual information of the rABB.

Lemma 1. *Let $(X, Y|Z)$ be the correlation corresponding to the random public bulletin board $\mathsf{rABB}^{P,Q}_{m,m}$. For each z in the sample space of the random variable*

Z, let ℓ_z be the length of z after removing all duplicate elements. Then

$$I(X,Y|Z) = \sum_z p_Z(z) \cdot \log \binom{2\ell_z}{\ell_z} = \mathbb{E}_z \left[\log \binom{2\ell_z}{\ell_z} \right].$$

Proof. First, note that $Z = X \cup Y$. Thus, $H(X|Y,Z) = 0$ since X is completely determined conditioned on knowing Y and Z. We have

$$\begin{aligned} I(X,Y|Z) &= \sum_z p_Z(z) \cdot I(X,Y|Z=z) \\ &= \sum_z p_Z(z) \cdot (H(X|Z=z) - H(X|Y,Z=z)) \qquad (Fact\ 1) \\ &= \sum_z p_Z(z) \cdot H(X|Z=z) \end{aligned}$$

For each $x = \{a_1, a_2, \ldots, a_m\}$ in the sample space of X, there is no duplicates in x; that is $a_i \neq a_j$ for every $i \neq j$. Conditioned on $Z = z = \{a_1, \ldots, a_m, b_1, \ldots, b_m\}$, which might contains duplicates, the number of x that are consistent with z is $\binom{2\ell_z}{\ell_z}$. Thus, the support's size of the random variable $(X|Z=z)$ is $\binom{2\ell_z}{\ell_z}$. Observe that the random variable $(X|Z=z)$ is uniform over its support. This implies that $H(X|Z=z) = \log \binom{2\ell_z}{\ell_z}$, for every z such that $p_Z(z) > 0$. Therefore, we have

$$H(X,Y|Z) = \sum_z p_Z(z) \cdot \log \binom{2\ell_z}{\ell_z},$$

which completes the proof.

By our construction in Figure 4, it is clear that the expected key length of our protocol is the quantity $E_z \log \binom{2\ell_z}{\ell_z}$ defined above. The following results are consequences of Lemma 1.

Corollary 4. *The expected key length of the protocol in Figure 4 is exactly $I(\mathsf{rABB}^{P,Q}_{m,m})$, where P and Q are the distribution that samples m messages randomly without replacement.*

Corollary 5. *Let $m, n \in \mathbb{N}$ and let P, Q be arbitrary distributions over $(\mathbb{Z}_{2^n})^m$. Then, the expected key length of any protocol in the $\mathsf{rABB}^{P,Q}_{m,m}$ is at most $\log \binom{2m}{m}$.*

Remark. The message length does not influence the key length. It helps only in lowering the failure probability.

4.5 Security Analysis

Theorem 5. *The key-agreement protocol in Figure 4 securely establishes a shared key k between $\mathcal{A}$ and $\mathcal{B}$, with $\mathcal{D}$ learning no information regarding the key.*

Proof. We give an outline of the security proof. The correctness comes from the fact that $\mathcal{A}$ can determine which elements were received by $\mathcal{A}$ and thus belong to A, while $\mathcal{B}$ can determine which elements were received by $\mathcal{B}$ and thus belong to B. Since $\Gamma := A \cup B$, $\mathcal{B}$ is also able to determine which elements were received by $\mathcal{A}$ and thus belong to A. Therefore, the information that $\mathcal{A}$ and $\mathcal{B}$ have are the same, and will allow them to agree on the same key k. Regarding privacy, note that due to the property of rABB, only $\mathcal{A}$ and $\mathcal{B}$ can determine which elements were received by $\mathcal{A}$ and thus belong to A. To $\mathcal{D}$, elements belonging to A and B look indistinguishable. Therefore, only $\mathcal{A}$ and $\mathcal{B}$ will know the value of the key k. We also note that since A and B are chosen according to the same distribution (uniform distribution in this case), all key values are equally likely to occur and the adversary gains no information.

4.6 Duplicate Recovery

We discuss a variant of our protocol named duplicate-recovery variant that is especially useful in settings where m is relatively close to 2^n, which means that many duplicates are likely to occur. This protocol allows parties to consider duplicates as opposed to disregarding them and is optimal in these settings. (We note that under the same m, having duplicates will decrease the key length k. This variant is designed to "recover" slightly from cases when duplicates exist. However, increasing n such that no duplicates occur will result in a larger key length.)

We also highlight a combinatorial problem that naturally arises in this variant that may be of independent interest.

Protocol Overview. This variant of the protocol is very similar to the original protocol, with the key difference being P and Q allows sample with replacement, i.e. duplicates are possible and parties do not discard the duplicates. Then, when parties want to determine the key k, they have to consider all possible cases. (Note that for security, P cannot sample values independently and uniformly at random, instead, P samples in a way such that every "set" occurs with equal probability. For example, $\{0, 0\}$ and $\{0, 1\}$ have the same probability of occurring. With this, the security of this variant closely follows the security of the original protocol.)

In general, given a multi-set of values Γ, one can list all possible values of A and B that can produce such a multi-set. However, a direct listing requires exponential computation as there are exponentially many possible cases. Given a generic algorithm for counting the number of possible cases, one can apply the same idea as Algorithm 1 to recursively determine the k value of a given set of inputs. We note that this generic algorithm for counting the number of possible cases given Γ may be of independent interest in the field of combinatorics.

Counting Solutions. We elaborate more on the problem of counting the number of possible sets that exist for a given Γ. One can easily do so by enumerating all possible solutions. For example, let us look at the simple case of when $m = 2$ and $n = 1$. Figure 6 shows all possible Γ and the corresponding possible A and

B. We can clearly see that when $\Gamma = \{0, 0, 1, 1\}$, there are 3 possible solutions for A, B, and when $\Gamma = \{0, 1, 1, 1\}$, there are 2 possible solutions for A, B.

Alternatively, we can model it as an Integer Programming (IP) problem. Let us define $z_1, \ldots, z_{N-1}$ as z_i being the amount of i showing up in Γ. Similarly, define $x_1, \ldots, x_{N-1}$ as x_i being the amount of i showing up in A, and $y_1, \ldots, y_{N-1}$ being the amount of i in B. Note that $\forall i, x_i \in \mathbb{Z}, y_i \in \mathbb{Z}, z_i \in \mathbb{Z}$.

We need to find the number of possible solutions to the x_is and y_js satisfying the equations and constraints

$$\sum_{i=0}^{N-1} x_i = m$$

$$\sum_{i=0}^{N-1} y_i = m$$

$$\forall i, x_i + y_i = z_i$$

$$\forall i, x_i \geqslant 0$$

$$\forall i, y_i \geqslant 0$$

To the best of our knowledge, there are no works specifically answering this question. We believe this question may be of independent interest. We also note that when there are no duplicates, that is $\forall i, z_i \leqslant 1$, this reduces to a simple binomial problem with the solution being $\binom{2m}{m}$.

5 Binary Erasure Channel

Rabin and Crépeau [8,32,33] showed that binary erasure channels suffice for general secure computation using interaction. These elegant noise sources provide uncluttered access to abstract the primary hurdles in achieving security. This section focuses on constructing binary erasure channels in the rABB-hybrid. We present our round-optimal secure protocols.

Problem Setting. Suppose parties are in the $\mathsf{rABB}^{P,Q,R}_{m_A,m_B,m_C}$ hybrid (with the helper $\mathcal{C}$). That is, party $\mathcal{A}$ has $A = \{a_1, a_2, \ldots, a_{m_A}\}$ sampled according to P, party $\mathcal{B}$ has $B = \{b_1, b_2, \ldots, b_{m_B}\}$ sampled according to Q, and $\mathcal{C}$ has $C = \{c_1, c_2, \ldots, c_{m_C}\}$ sampled according to R. Furthermore, all parties have the set $A \cup B \cup C$. Parties $\mathcal{A}$ and $\mathcal{B}$ want to establish a binary erasure channel $\mathsf{BEC}(p)$ between them, with $\mathcal{A}$ being the receiver and $\mathcal{B}$ being the sender. Parties $\mathcal{A}$ and $\mathcal{B}$ can communicate via an authenticated channel. We are in the semi-honest adversary model; that is, parties follow the protocol description but are curious to learn more from the protocol's transcript. Unlike in the key agreement protocols, the adversary here can corrupt a party.

For the binary erasure channel, without loss of generality, we assume that $\mathcal{B}$ is the sender and $\mathcal{A}$ is the receiver. So $\mathcal{A}$ will send a bit β to $\mathcal{A}$. The protocol is

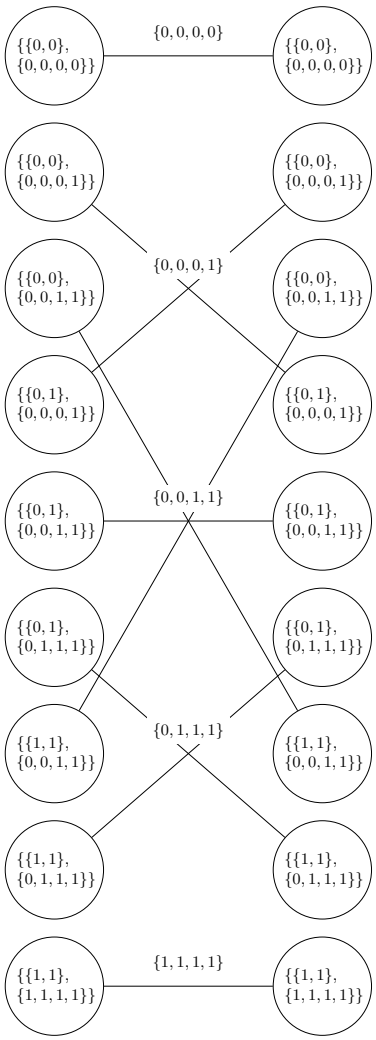

Fig. 6. Bipartite-graph for parties view and transcript for $m = 2$, $n = 1$. The nodes at the left represent the view of $\mathcal{A}$, which includes its private input (A) and the transcript it sees (Γ). Similarly, the nodes at the right represent the view of $\mathcal{B}$. Edges represent consistent views (transcript matches the private inputs produce correct transcript), with the transcript (Γ) labeled above the edges.

correct if $\mathcal{A}$ receives β with probability $(1-p)$, and $\mathcal{A}$ receives $\perp$ (nothing) with erasure probability p. We define security following the standard simulation-based definition. Intuitively, it is secure against the corrupted sender $\mathcal{B}$ if $\mathcal{B}$ does not know whether the sender bit gets erased or not; it is secure against the corrupted receiver if the sender bit is uniformly random in the receiver's view whenever she outputs $\perp$. We say that the protocol is ε-statistical secure if the simulation error is at most ε, and perfect secure if $\varepsilon = 0$.

We assume that no duplicates exist throughout the rest of the section. In the event that duplicates occur, parties can abort and rerun the protocol. Practically, by setting n to be large enough, we can ensure that with a high probability, no duplicates exist.

Remark. We note that with the key-agreement protocols in the previous section, $\mathcal{A}$ and $\mathcal{B}$ can establish a private authenticated channel from an authenticated channel.

Theorem 6. *Let $p \in (0,1)$ be a rational number. There is a perfectly secure one-round protocol for $\mathsf{BEC}(p)$ in the rABB-hybrid.*

Remark. The communication cost in our protocol is proportional to the denominator of the erasure probability. We left determining the minimum communication cost as an open problem.

Note that any irrational number can be approximated by a rational number with arbitrary precision. For example, this can be done by using Drichilet's approximation algorithm. Therefore, we have the following result as a corollary.

Corollary 6. *For any erasure probability $p \in (0,1)$ and $\varepsilon \in (0,1)$, there is a ε-statistical secure one-round protocol for $\mathsf{BEC}(p)$ in the rABB-hybrid.*

5.1 Construction

We present our protocol in Figure 7, as well as provide an overview of the protocol below.

The main idea behind the protocol is that given a set of two values, one from A and one from B, both $\mathcal{B}$ and $\mathcal{A}$ can distinguish and identify the owner of the value, and agree on a one-bit key used to send a one-bit message. On the other hand, if the set of two inputs is from B and C, then $\mathcal{A}$ learns nothing about the key and thus nothing about the message.

Therefore, in the protocol, $\mathcal{A}$ and $\mathcal{C}$ will each receive several values from the rABB, while $\mathcal{B}$ receives one value. $\mathcal{B}$ will then select two values from the multi-set published by the rABB, ensuring that one of them is his value, and encrypt the bit message using the key derived from those two values. If he selected his value and one of $\mathcal{C}$'s values, then the message is erased. If he selected his value and one of $\mathcal{A}$'s values, then $\mathcal{A}$ receives the message. Note that $\mathcal{B}$ cannot distinguish between $\mathcal{A}$'s and $\mathcal{C}$'s value, so $\mathcal{B}$ will not know if the message was erased, and $\mathcal{A}$ cannot distinguish between $\mathcal{B}$ and $\mathcal{C}$'s value, so the message can indeed be erased.

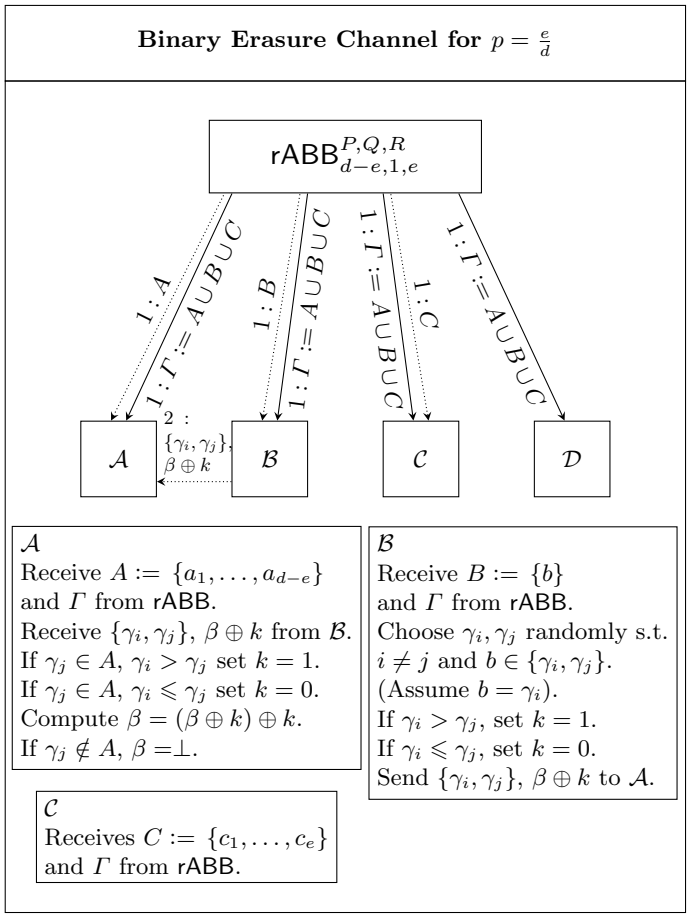

Fig. 7. Binary Erasure Channel from $\mathcal{B}$ to $\mathcal{A}$ using a helper $\mathcal{C}$ and in presence of an eavesdropper $\mathcal{D}$

5.2 Correctness and Security Proofs

Correctness. Observe that $\gamma_i \neq \gamma_j$ since there are no collisions at all. Thus, $\gamma_i > \gamma_j$ with probability $1/2$ and $\gamma_i < \gamma_j$ with probability $1/2$. Thus, the bit k is a uniformly random bit. Observe that $\gamma_i \in A$ with probability $(d-e)/d$. So, $\mathcal{A}$'s output is β with probability $1 - e/d$ and $\perp$ with probability e/d. Therefore, the protocol is perfectly correct.

Security. For security against a corrupted $\mathcal{B}$, whether the bit gets erased or not depends entirely on the event $\gamma_j \in A$ that $\mathcal{B}$ knows nothing about. Therefore, $\mathcal{B}$ does not know whether $\mathcal{A}$'s output is $\perp$ (erased) or β. For security against a corrupted $\mathcal{A}$, we need to show that when $\mathcal{A}$ outputs $\perp$, the bit β is uniformly random in the view of $\mathcal{A}$. $\mathcal{A}$ outputs $\perp$ when $\gamma_j \notin A$. In $\mathcal{A}$'s view, the event

$\gamma_i > \gamma_j$ is uniformly random. It means that $\mathcal{A}$ has β masking with a uniformly random bit. Hence, it follows from the property of the one-time pad that the bit β is uniformly random in $\mathcal{A}$'s view.

5.3 On the Round Optimality of Our Protocols

Our protocols use only one round of communication from $\mathcal{B}$ to $\mathcal{A}$. This section will show that it is impossible to securely implement BEC in the rABB-hybrid without communication. In fact, we will show that it is impossible to implement the BEC with randomized inputs, a weaker functionality. We shall employ the techniques from secure non-interactive simulation (SNIS/SNIR), recently introduced in [2, 21, 22], to prove the following theorem.

Theorem 7. *Let $\varepsilon \in (0,1)$ be the erasure probability. It is impossible to securely implement $\mathsf{BEC}(p)$ in the $\mathsf{rABB}^{P,Q,R}_{m_A,m_B,m_C}$ hybrid without communication.*

Proof Sketch. We prove this by contradiction. Suppose that it is possible to get $\mathsf{BEC}(p)$ from the $\mathsf{rABB}^{P,Q,R}_{m_A,m_B,m_C}$. It follows from [2, 21] that if it is possible to implement the randomized inputs $\mathsf{BEC}(p)$ from some other distribution (X, Y), then the eigenvalues of $\mathsf{BEC}(p)$ must be a subset of eigenvalues of the distribution (X, Y). Note that the eigenvalues of $\mathsf{BEC}(p)$ are 1 and $\sqrt{1-p}$. The correlation $\mathsf{rABB}_{m_A,m_B,m_C}$ is a family of joint distributions of the form $(X, Y|Z)$. Therefore, it must be the case that $\sqrt{1-p}$ is an eigenvalue of the correlation $(X, Y|Z = z)$, for every z in support of the random variable Z. This implies that $\sqrt{1-p}$ is an eigenvalue of all the conditional distributions $(X, Y|Z = z)$, which is impossible.

6 Random Oblivious Transfer

In this section, we consider the construction of randomized oblivious transfer (ROT) (see Section 2 for definitions) in the rABB-bybrid. ROT is the randomized version of the prominent OT functionality. In the rABB-hybrid with the helper $\mathcal{C}$, we construct an efficient one-round protocol. We start by describing the problem setting as follows.

Problem Setting. $\mathcal{A}$ and $\mathcal{B}$ will like to establish a random oblivious transfer (ROT) between them with $\mathcal{A}$ being the receiver and $\mathcal{B}$ being the sender. They have access to a rABB and a helper party $\mathcal{C}$, as well as a private authenticated channel from $\mathcal{B}$ to $\mathcal{A}$. We note that with the key-agreement protocol, $\mathcal{A}$ and $\mathcal{B}$ can establish a private authenticated channel from an authenticated channel. In the random oblivious transfer, $\mathcal{B}$ receives two random binary messages $\beta_0 \in \{0,1\}$ and $\beta_1 \in \{0,1\}$. $\mathcal{A}$ will receive a random choice bit $\psi \in \{0,1\}$, as well as the message β_ψ. Security is defined as $\mathcal{B}$ learning nothing about ψ, and $\mathcal{A}$ learning anything about $\beta_{1-\psi}$.

We also note that in our setting, we assume that no duplicates exist. Practically, by setting n to be large enough, we can ensure that with a high probability, no duplicates exist. In the event that duplicates occur, parties can simply abort

and rerun the protocol. Therefore, we assume that no duplicates exist throughout the rest of the section.

We prove the following theorem.

Theorem 8. *There is a perfectly secure one-round protocol for* ROT *in the* rABB*-hybrid (with the helper).*

We present our protocol in Figure 8, as well as provide an overview of the protocol in the next section.

6.1 Construction

The main idea behind the protocol is that given a set of two values, one from A and one from B, both $\mathcal{B}$ and $\mathcal{A}$ can distinguish and identify the owner of the values, and agree on a random one-bit message. On the other hand, if the set of two values is from B and C, then $\mathcal{A}$ learns nothing about the message.

Therefore, in the protocol, $\mathcal{A}$ and $\mathcal{C}$ will each receive one value from the rABB, while $\mathcal{B}$ receives two values. $\mathcal{B}$ will then randomly select two pairs of two values each from the multi-set published by the rABB, ensuring that one of the values in each pair is one of his values. Note that this allows $\mathcal{B}$ to obtain two random messages, with the random message being decided by if $\mathcal{B}$'s value is the larger one in the pair. The pair containing his input and one of $\mathcal{C}$'s values will cause the associated message to be erased. While the pair containing his values and one of $\mathcal{A}$'s values will allow the associated message to be delivered. Thus, $\mathcal{A}$ will receive exactly one of the two messages, and also know which message she received. Note that $\mathcal{B}$ cannot distinguish between $\mathcal{A}$'s and $\mathcal{C}$'s values, so $\mathcal{B}$ will not know which message was delivered, and $\mathcal{A}$ cannot distinguish between $\mathcal{B}$ and $\mathcal{C}$'s values, so $\mathcal{A}$ will not learn anything about the message that was erased.

We note that this is very similar to the BEC protocol, and that essentially we are running two "correlated BEC" with erasure probability $\frac{1}{2}$, and correlated in a way such that when one BEC is erased, the other BEC will not be.

Role of $\mathcal{C}$. We briefly discuss the role of $\mathcal{C}$, and why $\mathcal{C}$ is necessary for BEC and ROT, but not for key agreement. Essentially, $\mathcal{C}$ serves to create confusion and cause some information to be lost/erased. In BEC, with a certain probability, the message needs to be erased, and that happens precisely when $\mathcal{C}$ influences the protocol ($\mathcal{C}$'s value was selected by $\mathcal{B}$). In ROT, one of the two messages needs to be erased, and it's the message that is affected by $\mathcal{C}$ that ends up being lost. On the other hand, for key agreement, we want to preserve as much information as possible in order to obtain a larger key. Thus removing the participation of $\mathcal{C}$ from the key agreement allows the best protocol performance.

6.2 Correctness and Security Proofs

For correctness, since $\mathcal{B}$ cannot distinguish between a and c, $\Pr[a \in \{\gamma_{i_0}, \gamma_{j_0}\}] = \Pr[a \in \{\gamma_{i_1}, \gamma_{j_1}\}] = \frac{1}{2}$. Therefore, $\Pr[\psi = 0] = \Pr[\psi = 1] = \frac{1}{2}$. Additionally, since both $\mathcal{A}$ and $\mathcal{B}$ can distinguish between a and b_i, the message received by $\mathcal{A}$ and

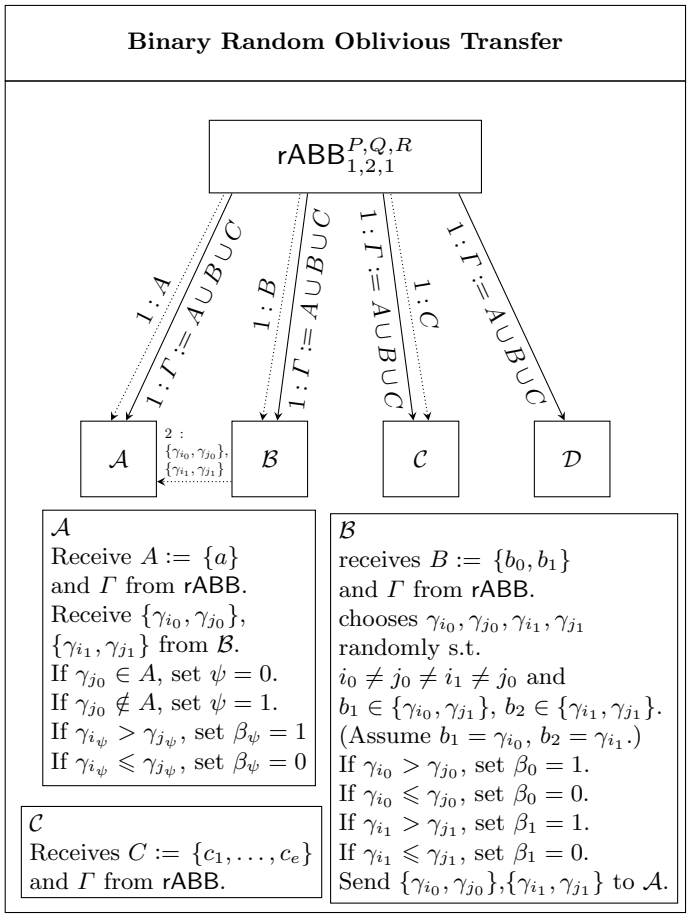

Fig. 8. Binary Random Oblivious Transfer from $\mathcal{B}$ to $\mathcal{A}$ using a helper $\mathcal{C}$ and in presence of an eavesdropper $\mathcal{D}$

$\mathcal{B}$ will be consistent. For security against a corrupted $\mathcal{B}$, $\mathcal{B}$ cannot distinguish a from c and therefore will learn nothing about ψ. For security against a corrupted $\mathcal{A}$, since $\mathcal{A}$ cannot distinguish between c and b_i, $\mathcal{A}$ learns nothing about $\beta_{1-\psi}$. In $\mathcal{A}$'s view, the event $b_i > c$ is uniformly random.

References

1. Ittai Abraham, Benny Pinkas, and Avishay Yanai. Blinder - scalable, robust anonymous committed broadcast. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *CCS '20: 2020 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, USA, November 9-13, 2020*, pages 1233–1252. ACM, 2020. doi:10.1145/3372297.3417261.
2. Pratyush Agarwal, Varun Narayanan, Shreya Pathak, Manoj Prabhakaran, Vinod M. Prabhakaran, and Mohammad Ali Rehan. Secure non-interactive reduction and spectral analysis of correlations. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part III*, volume 13277 of *LNCS*, pages 797–827. Springer, Heidelberg, May / June 2022. doi:10.1007/978-3-031-07082-2_28.
3. Rudolf Ahlswede and Imre Csiszár. Common randomness in information theory and cryptography - I: secret sharing. *IEEE Trans. Inf. Theory*, 39(4):1121–1132, 1993. doi:10.1109/18.243431.
4. Bowen Alpern and Fred B. Schneider. Key exchange using 'keyless cryptography'. *Information Processing Letters*, 16(2):79–81, 1983. URL: https://www.sciencedirect.com/science/article/pii/0020019083900297, doi:https://doi.org/10.1016/0020-0190(83)90029-7.
5. Megumi Ando, Anna Lysyanskaya, and Eli Upfal. On the complexity of anonymous communication through public networks. In Stefano Tessaro, editor, *2nd Conference on Information-Theoretic Cryptography, ITC 2021, July 23-26, 2021, Virtual Conference*, volume 199 of *LIPIcs*, pages 9:1–9:25. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.ITC.2021.9.
6. Andrej Bogdanov and Elchanan Mossel. On extracting common random bits from correlated sources. *IEEE Trans. Inf. Theory*, 57(10):6351–6355, 2011. doi:10.1109/TIT.2011.2134067.
7. Siu On Chan, Elchanan Mossel, and Joe Neeman. On extracting common random bits from correlated sources on large alphabets. *IEEE Trans. Inf. Theory*, 60(3):1630–1637, 2014. doi:10.1109/TIT.2014.2301155.
8. Claude Crépeau. Equivalence between two flavours of oblivious transfers. In Carl Pomerance, editor, *CRYPTO'87*, volume 293 of *LNCS*, pages 350–354. Springer, Heidelberg, August 1988. doi:10.1007/3-540-48184-2_30.
9. Imre Csiszár and Prakash Narayan. Secrecy capacities for multiple terminals. *IEEE Trans. Inf. Theory*, 50(12):3047–3061, 2004. doi:10.1109/TIT.2004.838380.
10. Debajyoti Das, Sebastian Meiser, Esfandiar Mohammadi, and Aniket Kate. Anonymity trilemma: Strong anonymity, low bandwidth overhead, low latency - choose two. In *2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings, 21-23 May 2018, San Francisco, California, USA*, pages 108–126. IEEE Computer Society, 2018. doi:10.1109/SP.2018.00011.
11. Claudia Diaz, Harry Halpin, and Aggelos Kiayias. The nym network the next generation of privacy infrastructure. Technical report, Nym Technologies SA, 2021 [Online]. URL: https://nymtech.net/nym-whitepaper.pdf.
12. Roger Dingledine, Nick Mathewson, and Paul F. Syverson. Tor: The second-generation onion router. In Matt Blaze, editor, *USENIX Security 2004*, pages 303–320. USENIX Association, August 2004.
13. Saba Eskandarian, Henry Corrigan-Gibbs, Matei Zaharia, and Dan Boneh. Express: Lowering the cost of metadata-hiding communication with cryptographic

privacy. In Michael Bailey and Rachel Greenstadt, editors, *30th USENIX Security Symposium, USENIX Security 2021, August 11-13, 2021*, pages 1775–1792. USENIX Association, 2021. URL: https://www.usenix.org/conference/usenixsecurity21/presentation/eskandarian.

14. Michael J. Fischer and Rebecca N. Wright. Multiparty secret key exchange using a random deal of cards. In Joan Feigenbaum, editor, *CRYPTO'91*, volume 576 of *LNCS*, pages 141–155. Springer, Heidelberg, August 1992. doi:10.1007/3-540-46766-1_10.
15. Sanjam Garg, Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Cryptography with one-way communication. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 191–208. Springer, Heidelberg, August 2015. doi:10.1007/978-3-662-48000-7_10.
16. Yossi Gilad and Amir Herzberg. Plug-and-play IP security - anonymity infrastructure instead of PKI. In Jason Crampton, Sushil Jajodia, and Keith Mayes, editors, *Computer Security - ESORICS 2013 - 18th European Symposium on Research in Computer Security, Egham, UK, September 9-13, 2013. Proceedings*, volume 8134 of *Lecture Notes in Computer Science*, pages 255–272. Springer, 2013. doi:10.1007/978-3-642-40203-6_15.
17. Amin Aminzadeh Gohari and Venkat Anantharam. Information-theoretic key agreement of multiple terminals: part I. *IEEE Trans. Inf. Theory*, 56(8):3973–3996, 2010. doi:10.1109/TIT.2010.2050832.
18. Ryan Henry, Amir Herzberg, and Aniket Kate. Blockchain access privacy: Challenges and directions. *IEEE Secur. Priv.*, 16(4):38–45, 2018. doi:10.1109/MSP.2018.3111245.
19. Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Cryptography from anonymity. In *47th FOCS*, pages 239–248. IEEE Computer Society Press, October 2006. doi:10.1109/FOCS.2006.25.
20. Sune K. Jakobsen and Claudio Orlandi. How to bootstrap anonymous communication. In Madhu Sudan, editor, *ITCS 2016*, pages 333–344. ACM, January 2016. doi:10.1145/2840728.2840743.
21. Hamidreza Amini Khorasgani, Hemanta K. Maji, and Hai H. Nguyen. Secure non-interactive simulation: Feasibility and rate. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part III*, volume 13277 of *LNCS*, pages 767–796. Springer, Heidelberg, May / June 2022. doi:10.1007/978-3-031-07082-2_27.
22. Hamidreza Amini Khorasgani, Hemanta K. Maji, and Hai H. Nguyen. Secure non-interactive simulation from arbitrary joint distributions. In *TCC 2022, Part II*, LNCS, pages 378–407. Springer, Heidelberg, November 2022. doi:10.1007/978-3-031-22365-5_14.
23. Albert Kwon, Henry Corrigan-Gibbs, Srinivas Devadas, and Bryan Ford. Atom: Horizontally scaling strong anonymity. In *Proceedings of the 26th Symposium on Operating Systems Principles, Shanghai, China, October 28-31, 2017*, pages 406–422. ACM, 2017. doi:10.1145/3132747.3132755.
24. Cheuk Ting Li and Venkat Anantharam. One-shot variable-length secret key agreement approaching mutual information. In *56th Annual Allerton Conference on Communication, Control, and Computing, Allerton 2018, Monticello, IL, USA, October 2-5, 2018*, pages 259–266. IEEE, 2018. doi:10.1109/ALLERTON.2018.8635830.
25. Cheuk Ting Li and Venkat Anantharam. One-shot variable-length secret key agreement approaching mutual information. *IEEE Trans. Inf. Theory*, 67(8):5509–5525, 2021. doi:10.1109/TIT.2021.3087963.

26. Ueli M. Maurer. Protocols for secret key agreement by public discussion based on common information. In Ernest F. Brickell, editor, *CRYPTO'92*, volume 740 of *LNCS*, pages 461–470. Springer, Heidelberg, August 1993. doi:10.1007/3-540-48071-4_32.
27. Ueli M. Maurer and Stefan Wolf. Unconditionally secure key agreement and the intrinsic conditional information. *IEEE Trans. Inf. Theory*, 45(2):499–514, 1999. doi:10.1109/18.748999.
28. Ueli M. Maurer and Stefan Wolf. Information-theoretic key agreement: From weak to strong secrecy for free. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 351–368. Springer, Heidelberg, May 2000. doi:10.1007/3-540-45539-6_24.
29. Elchanan Mossel and Ryan O'Donnell. Coin flipping from a cosmic source: On error correction of truly random bits. *Random Structures & Algorithms*, 26(4):418–436, 2005. doi:10.1002/rsa.20062.
30. Elchanan Mossel, Ryan O'Donnell, Oded Regev, Jeffrey E Steif, and Benny Sudakov. Non-interactive correlation distillation, inhomogeneous markov chains, and the reverse bonami-beckner inequality. *Israel Journal of Mathematics*, 154(1):299–336, 2006.
31. Andreas Pfitzmann and Michael Waidner. Networks without user observability — design options. In Franz Pichler, editor, *Advances in Cryptology — EUROCRYPT' 85*, pages 245–253, Berlin, Heidelberg, 1986. Springer Berlin Heidelberg.
32. Michael O. Rabin. How to exchange secrets by oblivious transfer. *Technical Memo TR-81*, 1981.
33. Michael O. Rabin. How to exchange secrets with oblivious transfer. Cryptology ePrint Archive, Report 2005/187, 2005. https://eprint.iacr.org/2005/187.
34. xx Foundation. xx network white paper. Technical report, xx Foundation, 2021 [Online]. URL: https://xx.network/wp-content/uploads/2021/10/xx-whitepaper-v2.0.pdf.
35. xx Foundation. xx network white paper xx cmix. Technical report, xx Foundation, 2021 [Online]. URL: https://xx.network/wp-content/uploads/2021/10/xx-whitepaper-v2.0.pdf.
36. Ke Yang. On the (im)possibility of non-interactive correlation distillation. In Martin Farach-Colton, editor, *LATIN 2004*, volume 2976 of *LNCS*, pages 222–231. Springer, Heidelberg, April 2004.