

Unconditional Security using (Random) Anonymous Bulletin Board

Albert Yu*, Hai H. Nguyen[†], Aniket Kate*[‡], Hemanta K. Maji*

*Department of Computer Science, Purdue University, USA

[†]ETH Zurich

[‡]Supra Research

Abstract—In a seminal work, Ishai et al. (FOCS–2006) studied the viability of designing unconditionally secure protocols for key agreement and secure 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 with respect to their round and communication complexity is not studied. This paper enriches this 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 200% throughput over the (FOCS–2006) study; i.e., using the same number of messages, we can (almost) double the bit-length of the agreed key. We also prove the near optimality of our approach.
- We offer unconditionally secure protocols for the (random) string oblivious transfer functionalities. We present a 1-round chosen message random string oblivious transfer and show how to extend it to a non-interactive (random) string oblivious transfer protocol and a 2-round chosen message string oblivious transfer.
- We prove a 1-round lower bound for BEC under certain conditions.

Central to our technical contributions is the abstraction of a distributional variant of the random ABB functionality. Investigating the concrete efficiency of founding MPC from this primitive leads to fascinating new mathematical challenges in well-established MPC models, which will be of broader interest to the community.

I. 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. [24] that analyzes the possibility of performing cryptography with unconditional security using an *anonymous bulletin board* (ABB). Ishai et al. establish unconditional security for prominent cryptographic tasks such as key agreement and honest-majority secure 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, diversifying the primitives that facilitate secure computation. Interestingly, they complement these constructions by showing the impossibility of unconditional secure

computation using anonymous broadcast in the absence of an honest majority.

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 [16] 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 as Nym [15] and xx.network [39], [40] 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 [23]. Academic literature on anonymous communication, as well as protocol implementations, have significantly expanded in the last two decades [1], [6], [13], [17], [28]. It is safe to say that ABBs are prevalent on the Internet today. Motivated by these real-world applications, our goal is to understand the efficacy and concrete efficiency of developing cryptography assuming access to such an ABB.

The utility of the ABB towards unconditional security is easy to illustrate using Alpern and Schneider’s [5] elegant key agreement protocol between Alice and Bob against an honest-but-curious adversary, which is referenced in Ishai et al. [24]. Alice and Bob independently pick random integers (say r_A and r_B , respectively) and publish those to the anonymous broadcast channel. The agreed single secret bit is 1 if $r_A > r_B$ and 0 if $r_A < r_B$. If $r_A = r_B$, then Alice and Bob fail to establish the secret bit and rerun the protocol. Notice that Alice and Bob know their respective input and thus can compute the secret bit; however, eavesdroppers cannot distinguish r_A from r_B and have no information about the agreed bit. Moreover, the failure probability (using the birthday bound) depends on the size of the sample space of the integers.

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 brings us to the question: what if Alice and Bob send multiple (say m) messages each? Can we agree on more than m bits using this $2m$ -sized multi-set? We answer this question affirmatively to demonstrate that Alice and Bob can indeed agree on close to $2m$ secret bits, which

improves the throughput of the key agreement to 200%, as compared to Alpern and Schneider [5]. This work aims to determine the concrete communication and round complexity of key cryptographic functionalities based on anonymity. This investigation leads to both qualitative and quantitative research questions in this context.

To this end, we establish connections of implementing functionalities using ABB in our context with various well-studied communication-limited MPC models (like *non-interactive correlation distillation* [34], [35], *secure non-interactive simulation/reduction* [2], [26], *one-way secure computation* [20], and *private simultaneous messages* [18]). Our problems translate into analytically tractable instances of these MPC models, which have generally been challenging to analyze. These connections lead us to several (near-optimal) protocol constructions. Our practically-motivated research objectives lead to fascinating research questions in these MPC models, potentially of interest to the broader cryptographic community.

A. Our Contributions

From the modeling perspective, this work assumes the existence of an anonymous broadcast, which we model as an Anonymous Bulletin Board (ABB) hybrid. There are four parties $\mathcal{A}, \mathcal{B}, \mathcal{C}$, and \mathcal{D} . The bulletin board ideal functionality, represented as $\text{ABB}_{m_A, m_B, m_C}$, takes as input three multi-sets: (1) $A := \{a_1, \dots, a_{m_A}\}$ from party \mathcal{A} , (2) $B := \{b_1, \dots, b_{m_B}\}$ from \mathcal{B} , and (3) $C := \{c_1, \dots, c_{m_C}\}$ from \mathcal{C} . Note that party \mathcal{D} does not provide any input. The functionality outputs the multi-set $\Gamma = A \cup B \cup C := \{\gamma_1, \dots, \gamma_{m_A+m_B+m_C}\}$ to all four parties.

Example 1 (Clarification on our ABB model). *For illustrative purposes, consider $m_A = m_B = m_C = 2$. Party \mathcal{A} sends the vector (x_1, x_2) to the ABB, party \mathcal{B} sends the vector (y_1, y_2) to the ABB, and party \mathcal{C} sends the vector (z_1, z_2) to the ABB. Our ABB publishes the multi-set of all the received elements $\{x_1, x_2, y_1, y_2, z_1, z_2\}$. More concretely, interpret this multi-set represented as the sorted vector containing its elements (with multiplicities). In particular, this multi-set is different from the following alternative interpretations.*

- 1) *The set $\{\{x_1, y_1, z_1\}, \{x_2, y_2, z_2\}\}$. In this alternative version, messages with identical indices are linked to each other. If parties wish to link messages using our ABB, they need to explicitly encode the indices into their message, which causes a logarithmic increase in their length.*
- 2) *The set $\{\{x_1, x_2\}, \{y_1, y_2\}, \{z_1, z_2\}\}$. This alternative version links messages sent by the same party while hiding the identity of the party. Such linking is achieved using our ABB by encoding anonymized identities of the parties into their respective messages.*

We refer to party \mathcal{C} as the *helper* and party \mathcal{D} as the *eavesdropper*. In the randomized version of bulletin board (rABB), the three multi-sets A, B, C are sampled according to some independent distributions P, Q, R , respectively. See

Section IV for a formal definition of ABB and its randomized version (rABB).

In addition to the bulletin board, parties also have public authenticated channels between them. In the ABB setting, we define each party's *communication complexity* as the number of bits that the party sends to the ABB plus the number of bits it sends to other parties through the public authenticated channels. For example, the communication complexity of party \mathcal{A} is the sum of the following quantities.

- The bit length of A (the message that party \mathcal{A} sends to the ABB)
- The bit length of the messages that party \mathcal{A} sends to other parties (\mathcal{B}, \mathcal{C} , and \mathcal{D}).

We define the communication complexity in the rABB setting in a similar manner. For example, the communication complexity of \mathcal{A} is the sum of the following quantities.

- Bit length of A (the message that that party \mathcal{A} receives from rABB)
- Bit length of the messages that party \mathcal{A} sends to other parties (\mathcal{B}, \mathcal{C} , and \mathcal{D}).

The sequel summarizes our contributions.

Result 1 (Key-agreement Protocol: Informal). *We present a non-interactive two-party key-agreement protocol using $\text{rABB}_{m, m, 0}$ with individual message length n that establishes (near-optimal) $2m$ -bit keys with $(m \cdot n)$ -bit communication complexity.*

Theorem 1 provides the formal statement for this result. Our construction is secure against a computationally unbounded eavesdropper \mathcal{D} . Our construction is straightforward to implement, and the key length (i.e., throughput) is near-optimal. Here, throughput is the ratio of the key length to the number of messages. The length n of the individual messages affects our algorithm's failure probability, the event where parties fail to agree on a key. Small messages would result in close-to-1 failure probability. Surprisingly, when n is larger than a particular threshold, it has essentially no impact on the key length. We also present a duplicate-recovery variant of the protocol in Result 1, which is suitable for other parameter regimes. Details on the duplicate-recovery variant can be found in Section V-G.

Remark 1 (Upper bound on our key length: additional comments). *Our proof of the optimality of our key length considers a wide family of protocols. In these protocols, parties can interact over multiple rounds using the public authenticated channels after the rABB invocation. The parties \mathcal{A}, \mathcal{B} , and \mathcal{C} receive messages from arbitrary independent message distributions P, Q , and R , respectively (not necessarily the uniform distribution). In our protocol, rABB delivers random independent messages to the parties. We prove this result using mutual information, entropy-based arguments, and the recent results of [29], [30].*

In our protocol, parties have access to a *single* ABB or rABB that they call once. Many other protocols (such as [9],

[14], and some protocols in [24]) either require additional assumptions, such as on the synchrony of the system model or require multiple independent instances of ABB to be implemented. We emphasize that this is qualitatively different from our protocol setting, and a direct comparison of the communication costs of these protocols against ours results in an inaccurate representation of both their protocols and ours. Therefore, we focus our concrete communication cost analysis on comparison with the state-of-the-art protocol in this setting, which is [5]. The result of this concrete communication cost comparison is presented in Figure 1. For typical values of k such as 128, [5] requires roughly $2.9\times$ our communication cost.




Fig. 1. Plot of the ratio of communication required by [5] over our protocol's communication needed to achieve various expected values of key-length k . For example, to generate a 128-bit key (on average), [5] requires roughly $2.9\times$ our communication cost.

In the context of implementing oblivious transfers, Ishai et al. [24] proved the impossibility of realizing oblivious transfer (OT) using ABB when honest parties are not in the majority. This implies that it is impossible to realize oblivious transfers (as well as their randomized versions) in the ABB-hybrid without the helper party \mathcal{C} . We construct oblivious transfer protocols that achieve a few different functionality variants – a step towards diversifying setups for oblivious transfers.

Result 2. *We present a 1-round (round-optimal) protocol for establishing (chosen message) random string oblivious transfer (cmROT $^\ell$) from sender \mathcal{B} to receiver \mathcal{A} with the helper \mathcal{C} .*

The cmROT $^\ell$ functionality takes as input two ℓ -bit messages x_0 and x_1 from the sender and delivers the tuple (b, x_b) to the receiver, where the bit b is chosen uniformly at random. Theorem 6 provides the formal statement of this result.

The round optimality of this construction is a consequence of Result 3 and the fact that one can use (chosen message)

random string OT to implement an erasure channel.¹

Corollary I.1. *We extend Result 2 to a non-interactive protocol for establishing random string oblivious transfer (ROT $^\ell$) from sender \mathcal{B} to receiver \mathcal{A} with the helper \mathcal{C} .*

The ROT $^\ell$ functionality samples two uniformly random messages x_0, x_1 , a uniformly random bit b , delivers the tuple (x_0, x_1) to the sender, and delivers the tuple (b, x_b) to the receiver. Discussions on the non-interactive random string oblivious transfer can be found in Section VI-D.

Corollary I.2. *We extend Result 2 to a two-round protocol for establishing (chosen message) string oblivious transfer (cmOT $^\ell$) from sender \mathcal{B} to receiver \mathcal{A} with the helper \mathcal{C} .*

The cmOT $^\ell$ functionality takes as input two ℓ -bit messages x_0 and x_1 from the sender, one bit b from the receiver, and delivers x_b to the receiver. This protocol achieves unconditional security against semi-honest adversaries. In our protocol, the sender \mathcal{B} sends a message to \mathcal{A} using a private authenticated channel. Discussions on the string oblivious transfer can be found in Section VI-D.

Corollary I.3. *We extend these protocols to 1-out-of- N OT (where the sender chooses N inputs) in Section VI-D.*

Corollary I.4. *Using results from [20], we can realize any one-way secure computation for arbitrary (possibly randomized) sender-receiver functions from our ROT protocol.*

Corollary I.5. *We can implement a Binary Erasure Channel (BEC) with erasure probability $\frac{e}{d}$ using 1-out-of- d ROT from Corollary I.3.*

This can be done by having e of the messages be a special “erased” message while the remaining $d - e$ messages are the actual bit being sent.

Remark 2 (New research problems in interaction-limited MPC models). *Our use of rABB P,Q,R can be interpreted as sampling from the joint distribution $(P, Q, R|\Gamma)$ in a preprocessing step, where Γ is the union of the three. Under this interpretation, our research problems translate into research questions in the NICD [34], [35], SNIS [26], SNIR [2], OWSC [20], and PSM [18] models.*

- 1) For key agreement, we prove that the uniform distribution achieves the optimal result even against arbitrary independent distributions.
- 2) For random string oblivious transfer, we show that by using specialized distributions P, Q, R , we are able to obtain non-interactive random string oblivious transfer.

Result 3. *We prove a 1-round lower bound for implementing a binary erasure channel from \mathcal{B} to \mathcal{A} utilizing the rABB (with a helper) and a public authenticated channel from \mathcal{B} (the sender) to \mathcal{A} (the receiver) when the messages are*

¹The sender can choose to send $x_0 = 11$ and $x_1 = 0m$ for a bit $m \in \{0, 1\}$. The receiver receives the bit m with a probability of $1/2$; otherwise, it is erased with a probability of $1/2$. Therefore, the impossibility of implementing an erasure channel extends to this case.

sampled from uniform distributions P, Q, R (see Theorem 7 for details). Proving the optimality for arbitrary independent distributions P, Q, R remains open. Recall P, Q, R represent the distribution of the messages sent by rABB to the parties \mathcal{A}, \mathcal{B} , and \mathcal{C} , respectively. Analyzing this distributional variant of rABB motivates new research directions in interaction-limited MPC models, like SNIS and SNIR. This problem is challenging even when P, Q, R are flat distributions over a sparse subset of the message space. Typically, these models (like NICD [34], [35], SNIS [26], SNIR [2], OWSC [20], and PSM [18]) have strong hardness-of-computation results. However, for our application scenarios, there are non-trivial and practically useful construction as well.

B. Related Works

1) *Key-Agreement*: There are many works focused on key-agreement or developing secure point-to-point links based on anonymous communication. [5] 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. There are also several works that expand upon or improve [5]. For example, [42] expands [5] to work over semi-honest channels. [9] proposes a protocol that only requires k total messages for a k -bit key by utilizing the fact that parties can set the source of the message to be honest or false, and also send messages in random order. [14] similarly proposes a protocol that requires the parties to send messages in a random order by implementing random wait times. [36] considers key-agreement when the receivers (instead of the senders) are anonymous. [19] 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 [21] demonstrate the practical utility of [24] 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 [4], [12], [22], [31]–[33]. The main objective is to study the achievable key rate when the number of samples tend to infinity.

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

2) *Communication-limited MPC Models: Non-interactive Correlation Distillation*. In information theory and theoretical computer science, non-interactive correlation distillation (NICD) is a well-studied analytically-tractable problem [8], [10], [34], [35], [41]. NICD also aims to establish 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, parties have access to ABB that generates a different form of conditional

distribution in the rABB-hybrid model. We are the first to choose this distribution and achieve near-optimal key length.

Secure Non-interactive Simulation/Reduction. Secure non-interactive simulation/reduction (SNIS/SNIR) is a cryptographic primitive introduced recently [2], [26]. In this model, parties have i.i.d samples of a source correlated private randomness; the objective is to non-interactively and securely transform these samples into i.i.d samples of another target correlated private randomness. This line of work investigates both the feasibility and efficiency of SNIS/SNIR constructions. We shall employ the techniques to prove the impossibility results in their settings to show the round-complexity of realizing BEC or OT using rABB-hybrid.

One-way Secure Computation. One-way secure computation [3], [20] uses one round of communication to securely transform the samples of the source distribution to the samples of the target distribution.

II. TECHNICAL OVERVIEW

This section provides a technical overview of our results. For a complete list of notations and backgrounds, refer to Section III. The formal definition of the anonymous bulletin board (ABB) and its variant are in Section IV.

A. Key Agreement

We first present an overview of our (near-optimal) key-agreement protocol in Figure 4. We construct a key agreement protocol in which parties \mathcal{A} and \mathcal{B} receive a set of m messages of n bits each (A and B respectively). Additionally, all parties ($\mathcal{A}, \mathcal{B}, \mathcal{D}$) receive the set of $2m$ messages ($\Gamma = A \cup B$) from the rABB. The parties first discard any duplicate messages in Γ , resulting in $2m'$ total messages where m' messages belong to each set A and B . Since no duplicate messages exist, only parties \mathcal{A} and \mathcal{B} can identify which of the $2m'$ messages belong to each set A and B . By using a canonical ordering of the $2m'$ messages and assigning messages belonging to A as 1 and messages belonging to B as 0, the two parties can agree on a $2m'$ bit string that is known only to them. Then, by using standard techniques in combinatorics, the two parties can index the agreed upon bit string out of the $\binom{2m'}{m'}$ possible bit strings and agree on a key of length $\log \binom{2m'}{m'}$.

We discuss how the parameter choices m and n affect the expected key length and the failure probability using standard techniques in probability. Additionally, we perform brute force searches to identify the optimal parameters for various key lengths and compare those results with previous state-of-the-art results.

Finally, using techniques on mutual information, we prove that under the setting of arbitrary/unlimited message length, our protocol achieves the optimal expected key length given parameter m .

Additionally, we present a variant of our protocol called duplicate recovery, which is suitable for small values of n . In duplicate recovery, instead of removing all duplicates, the protocol considers the duplicates as part of the possible

distributions. We note that in this case, indexing the possible distributions becomes non-trivial. We present such a problem as a new problem in combinatorics, as well as reformulate it as an Integer Programming (IP) problem. We believe this problem may be of independent interest.

The complete description and analysis of the key agreement protocol can be found in Section V.

B. Chosen Message Random String Oblivious Transfer

We present an overview of our construction of cmROT in Figure 7. A single bit of random oblivious transfer can be seen as two BEC instances that are correlated in a way such that whenever one of the messages is erased, the other message is delivered.

We use a set of four elements, one belonging to \mathcal{A} , one belonging to \mathcal{C} , and two belonging to \mathcal{B} , that is divided into two subsets that each contain an element belonging to \mathcal{B} . \mathcal{B} is able to identify both messages that belong to \mathcal{B} in the two subsets, and can therefore obtain two bits. On the other hand, \mathcal{A} can only identify the element belonging to \mathcal{B} in the subset that contains \mathcal{A} 's element. This creates a setting where \mathcal{B} is able to identify two messages while \mathcal{A} is only able to identify one of them.

When we directly perform the above step multiple times, a natural issue arises in which \mathcal{B} is unable to identify what messages \mathcal{A} can obtain, but will instead get a cartesian product of all the possible bits.

The key observation is that security still holds if we set all elements belonging to \mathcal{A} to be even (or odd), all elements belonging to \mathcal{C} to be odd (or even, respectively), and half the elements belonging to \mathcal{B} to be even and half to be odd. This will allow \mathcal{B} to “link” the bits that form the same message, thus identifying the two possible messages that \mathcal{A} can obtain without learning which message \mathcal{A} obtains.

We can also compress the multiple calls to rABB into a single call using sequence identifiers and parallel identifiers (full detail can be found in Section IV-D).

Finally, to ensure that \mathcal{C} learns nothing about either message, \mathcal{B} sends two “correction messages” that get xored with the original message to create the final message to \mathcal{A} through a private authenticated channel (such a private channel can be established in parallel with no additional round using our key-agreement protocol).

The complete description and analysis of the random string oblivious transfer protocol can be found in Section VI.

III. PRELIMINARIES

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

A. 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 the desired properties such as not being able to determine which elements came from A by only looking at $A \cup B$ hold for when sets are randomized or when they are sorted according to the canonical ordering.

B. 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.

Roughly, 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).

C. String Oblivious Transfer

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

Note that when $\ell = 1$, the functionality OT^1 is the (standard) bit oblivious transfer.

Random String Oblivious Transfer. Random oblivious transfer, denoted as ROT^ℓ , is a correlation that samples $x_0 \in \{0, 1\}^\ell, x_1 \in \{0, 1\}^\ell, b \in \{0, 1\}$ uniformly and independently at random. It provides Bob with the secret share $r_B = (x_0, x_1)$ and provides Alice the secret share $r_A = (b, x_b)$.

Chosen Message String Random Oblivious Transfer. Chosen Message Random oblivious transfer, denoted as cmROT^ℓ , is a functionality that takes as input $(x_0, x_1) \in (\{0, 1\}^\ell)^2$ from Bob, samples a bit b uniformly at random, and outputs (b, x_b) to Alice.

D. Entropy and Mutual Information

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

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

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

Definition III.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,

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

IV. 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.

A. 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 $\text{ABB}_{m_A, m_B, m_C}$. Formally, $\text{ABB}_{m_A, m_B, m_C}$ takes multi-set of inputs $A := \{a_1, \dots, a_{m_A}\}$ from \mathcal{A} , multi-set of inputs $B := \{b_1, \dots, b_{m_B}\}$ from \mathcal{B} , multi-set of inputs $C := \{c_1, \dots, c_{m_C}\}$ from \mathcal{C} , and outputs the multi-set $\Gamma = A \cup B \cup C := \{\gamma_1, \dots, \gamma_{m_A+m_B+m_C}\}$. In particular,

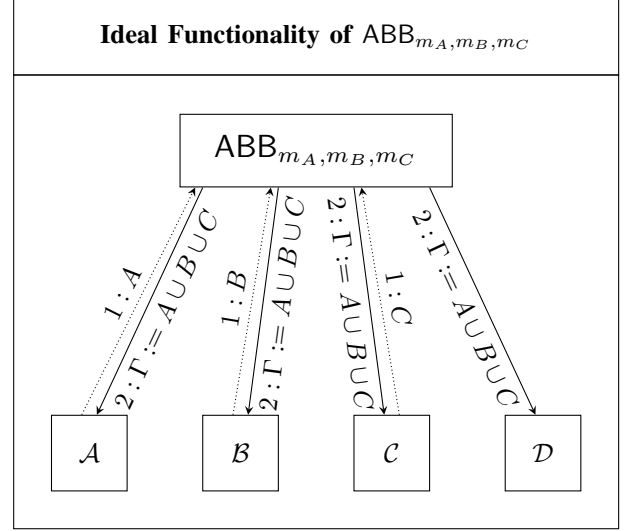


Fig. 2. Ideal Functionality of Anonymous Bulletin Board (ABB). Dotted lines represent private authenticated channels, while solid lines represent public channels. The number in front of messages shows the order/round in which the messages are sent. A is the multi-set $\{a_1, \dots, a_{m_A}\}$, B and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

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.

Additionally, we emphasize that one call to $\text{ABB}_{m, m, m}$ is different from m calls to $\text{ABB}_{1, 1, 1}$, even when we do not consider adaptive message choosing. Specifically, for m calls to $\text{ABB}_{1, 1, 1}$, all parties will know that the first 3 messages did not come from the same parties. That is, parties gain additional information on subsets of messages that definitely did not come from the same party. Whereas in a single call to $\text{ABB}_{m, m, m}$, parties do not gain such information.

B. 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 $\text{rABB}_{m_A, m_B, m_C}^{P, Q, R}$.

C. 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.

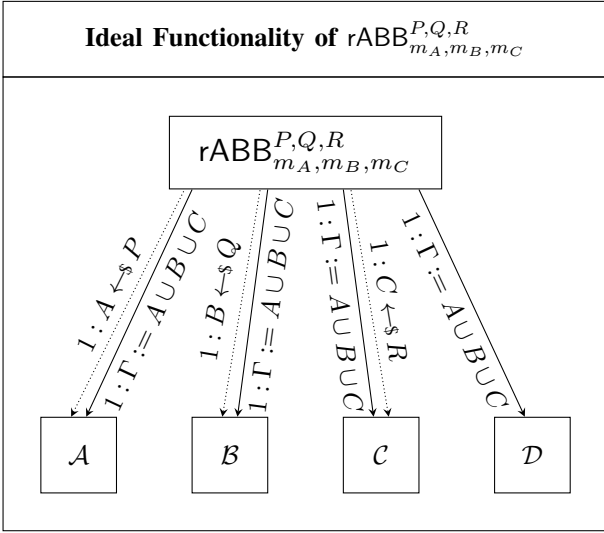


Fig. 3. Ideal Functionality of Random Anonymous Bulletin Board (rABB). Dotted lines represent private authenticated channels, while solid lines represent public channels. The number in front of messages shows the order/round in which the messages are sent. A is the multi-set $\{a_1, \dots, a_{m_A}\}$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

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 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.

D. Compression of Non-Adaptive Sequential and Parallel Calls

We also note that when messages are not chosen adaptively, we can compress multiple sequential calls to the ABB or rABB into a single call to the ABB/rABB by including a “sequence identifier” in each message. Furthermore, we can also compress multiple parallel calls to the ABB or rABB into a single call to the ABB or rABB by including a “parallel identifier” in each message.

Using the rABB as an example. Let $\text{rABB}_{i,j}$ represent $\text{rABB}_{m_{A_{i,j}}, m_{B_{i,j}}, m_{C_{i,j}}}^{P_{i,j}, Q_{i,j}, R_{i,j}}$.

Let us assume we want to call $\text{rABB}_{1,1}, \text{rABB}_{1,2}, \dots, \text{rABB}_{1,\pi}$ in the first round, $\text{rABB}_{2,1}, \text{rABB}_{2,2}, \dots, \text{rABB}_{2,\pi}$ in the second round, up to $\text{rABB}_{\sigma,1}, \text{rABB}_{\sigma,2}, \dots, \text{rABB}_{\sigma,\pi}$ in the σ^{th} round.

We can instead call $\text{rABB}'_{i,j}$, where $\text{rABB}'_{i,j}$ is $\text{rABB}_{m_{A_{i,j}}, m_{B_{i,j}}, m_{C_{i,j}}}^{P'_{i,j}, Q'_{i,j}, R'_{i,j}}$, where $P'_{i,j}$ is the same distribution as

$P_{i,j}$, except that every message is prefixed by i, j . $Q'_{i,j}$ and $R'_{i,j}$ are similarly defined.

Given this modification, we can now call $\text{rABB}'' := \text{rABB}_{\sum_{i,j} m'_A, \sum_{i,j} m'_B, \sum_{i,j} m'_C}^{P'', Q'', R''}$ where P'' is the union of all $P'_{i,j}$. Q'' and R'' are similarly defined. m'_A is defined as the sum of all $m'_{A_{i,j}}$, and m'_B and m'_C are similarly defined.

Given that all messages in rABB'' are prefixed by their sequence identifier and parallel identifier, the parties can locally divide them into the appropriate $\text{rABB}'_{i,j}$ s, and the result remains the same as the original.

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

V. KEY AGREEMENT PROTOCOLS

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

A. Problem Setting

Suppose parties are in $\text{rABB}_{m_A, m_B}^{P, Q}$ -hybrid (without the helper C). That is, the parties have access to a single instance of $\text{rABB}_{m_A, m_B}^{P, Q}$, which they can call one time. We note that this is different from having access to $\max(m_A, m_B)$ different $\text{rABB}_{1,1}^{P', Q'}$ -hybrid. That is, party A has $A = \{a_1, a_2, \dots, a_{m_A}\}$ sampled according to P , party B has $B = \{b_1, b_2, \dots, b_{m_B}\}$ sampled according to Q , and party D has $A \cup B$. Every message is n -bit. Parties A and B are allowed to communicate with each other through a public noiseless channel. Party D (the eavesdropper) can see the messages sent between A and B . At the end of the protocol, two parties A and B agree on a sample space Ω_L for the key, and A outputs a (variable-length) key $K_A \in \Omega_L$, and 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

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

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 $m \cdot n + \text{length}(T)$, where $\text{length}(T)$ 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 [30]. 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 [4], [31], [32], the eavesdropper has some side information about the samples of A and 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.

B. Our Protocol

Next, we present our protocol in Figure 4, as well as provide an overview of the protocol below.

The protocol is similar to the example presented in [24]. 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 and Q are independent identical copies, 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 $A := \{a_1, \dots, a_m\}$ 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 can also efficiently assign key values to the identified cases. 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.

To efficiently compute the value of K , we employ standard techniques for analyzing fixed-weight bitstrings using combinatorial². These techniques and algorithms have long been known and described in various places such as [7].

For completeness, we present such an algorithm in Algorithm 1 and briefly explain the logic behind the algorithm.

Algorithm 1: Key Determination

Parameters:

Input : A', Γ'

Output : K

- 1 We assume there exists some canonical ordering of elements in a set.
 - 2 Identifies which of $\gamma'_1, \dots, \gamma'_{2m'}$ 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.
 - 3 Compute $m' = \text{size}(A')$
 - 4 $K = 0$
 - 5 $c = m'$
 - 6 **for** i in $\text{range}(2 \cdot m')$: **do**
 - 7 **if** $x[i] == 1$ **then**
 - 8 $K = K + \text{binomial}(2 \cdot m' - i - 1, c)$
 - 9 $c = c - 1$
 - 10 **return** K
-

a) 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 this is by examining 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 .

C. 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

²We thank our anonymous reviewers for pointing us in the direction of these techniques.

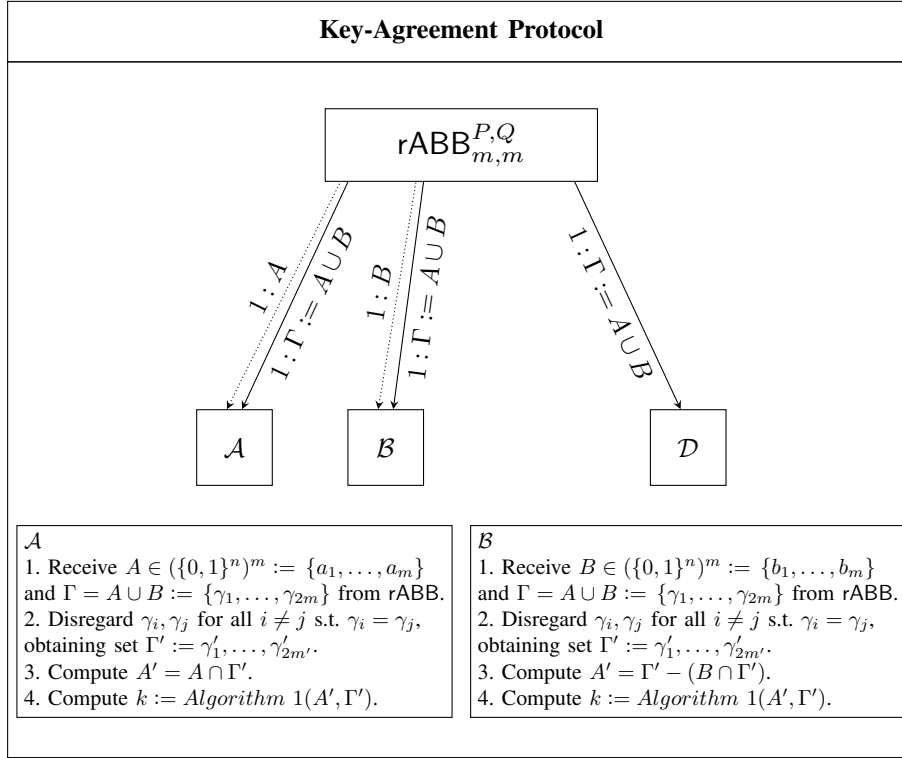


Fig. 4. Key-agreement protocol between parties \mathcal{A} and \mathcal{B} in presence of an eavesdropper \mathcal{D} in the $rABB_{m,m}^{P,Q}$ -hybrid, where P, Q are independent uniform distributions.

mean that is large. Therefore, our protocol outputs a long key with high probability.

Theorem 1. Fix $m, n \in \mathbb{N}$, and $0 \leq \varepsilon \leq 1$. Our non-interactive 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 $\exp(-\varepsilon^2 \cdot m \cdot (1 - \frac{m}{2^n})/2)$, where $m' = m \cdot (1 - \frac{m}{2^n}) \cdot (1 - \varepsilon)$.

Proof. Let us fix the elements of \mathcal{B} . For each element of \mathcal{A} , there is at least $(1 - \frac{m}{2^n})$ 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 $\log \binom{2m'}{m'} \approx \log \left(\frac{4^m}{\sqrt{m^n}} \right)$ -bit key.

Theorem 2. Fix $m, n \in \mathbb{N}$, satisfying $m = o(2^{n/3})$. 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. 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.

D. Comparison With previous State-of-the-Art

Recall that we are in the setting where parties have access to a single ABB or $rABB$ that they can call once. To the best of our knowledge, the best previous state-of-the-art key agreement protocol that works in this setting is [5]. Other protocols such as [9] and [14] require that parties send their message in a random order (which requires additional assumptions on the synchrony of the system model), and are often presented as using sequential calls to the ABB in order to perform key agreement on more than one bit. The key agreement protocols presented in [24] either require more communication rounds and communication costs than the one in [5], or are fundamentally similar to the one in [5].

Therefore, we focus our concrete communication cost analysis on comparison with [5]. Additionally, we present our analysis using the expected value of the key length. Similar analyses can be done on achieving the desired key length with high probability instead of in expectation by using standard techniques on concentration bounds, etc.




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.

To achieve a k -bit key in expectation, [5] requires each party to send $2k$ messages, each of length $\log(2k)+1$, where $\log(2k)$ bits are used to represent the sequence identifier and 1-bit is used to represent the random bit chosen. This results in a total communication cost of $2k \cdot (\log(2k) + 1)$. Even when we assume that parties do not send the leading 0s in the sequence identifiers, the total communication cost is still $2k + \sum_{i=1}^{\log(2k)} i \cdot 2^{i-1} = 2k \log(2k) + 1$.

We then compare this communication cost for various key lengths against the result from our automated search results presented in Figure 5 by taking a ratio of their communication cost divided by our communication cost, which produced Figure 1 shown in our contribution at the beginning of the paper.

For example, for 128-keys, their protocol required $2 \cdot 128 \log(2 \cdot 128) + 1 = 2049$ bits of communication while our protocol required 702 bits of communication, resulting in a ratio of $\frac{2049}{702} = 2.9188$.

For completeness, we also analyze the sequential version of the key-agreement protocol in [5]. In the sequential protocol, parties use number of rounds linear in the length of the key to achieve a lower communication cost. To obtain a k bit key, each party sends 1 bit in each of the $2k$ rounds, resulting in $2k$ total bits of communication per party.

An anonymous reviewer also proposed an alternate protocol for key agreement. We will briefly present their proposed protocol here and then compare our protocol to it.

Recall that in [24], Ishai et al. presented a non-interactive SUM protocol, in which two parties each send the additive shares of their values to the ABB. The two parties can sum up all messages (shares), and then subtract their value to learn the other's value, while both values are statistically hidden from

the adversary. This protocol can be used to allow one party to send a secret/private message to another. To achieve key agreement, Alice will pick a secret key, create additive shares of the key, and send them to the ABB, while Bob can send multiple random values to the ABB. Bob can then identify Alice's messages, and use their sum to compute the shared key.

Firstly, we want to point out that the SUM protocol is only statistically secure, while our protocol is unconditionally secure. While having the same round complexity (both being non-interactive), our protocol achieves a lower total communication cost (in bits). As stated in Lemma 4.1 from [24], they require each party share the messages into at least κ shares such that $\log\binom{2\kappa}{\kappa} > \ell$, where ℓ is the length of the message in bits. This translates to a bound of k shares of length at least k bits each in order to obtain a k bit key. This means the proposed protocol will require k^2 bits of communication, while our protocol only requires $2k \cdot (\log(2k) + 1)$ bits of total communication.

E. On the Optimality of Key Length

In this section, we analyze protocols under the setting of having arbitrary/unlimited message length and show that our protocol in Figure 4 is near optimal in terms of the expected key length based on the number of messages. In fact, we will prove a much stronger result. That is, the expected key length of any interactive protocol for key agreement in the $\text{rABB}_{m,m}^{P,Q}$ is at most $2m - \text{poly}(\log m)$, where P, Q are arbitrary independent distributions over $(\{0, 1\}^n)^m$, 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 $\text{rABB}_{m,m}^{P,Q}$ is at most $\log\binom{2m}{m}$.

Theorem 3. *Let $m, n \in \mathbb{N}$ and P, Q be independent distributions over $(\{0, 1\}^n)^m$. Suppose parties are in the random public anonymous bulletin board hybrid $\text{rABB}_{m,m}^{P,Q}$. Then, the expected key length in any key agreement protocol (allowing interaction) is at most $I(\text{rABB}_{m,m}^{P,Q}) + 1 + \log 3$.*

We shall employ the techniques developed recently in [29], [30] 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. [29], [30] *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 of Theorem 3. Recall the reformulation of rABB as correlated private randomness in Section IV-E. The correlation

$\text{rABB}_{m,m}^{P,Q}$ 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 $\text{rABB}_{m,m}^{P,Q}$. 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

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

For each $x = \{a_1, a_2, \dots, a_m\}$ 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 $Z = z = \{a_1, \dots, a_m, b_1, \dots, b_m\}$, which might contain 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. \square

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

Corollary V.1. *The expected key length of the protocol in Figure 4 is exactly $I(\text{rABB}_{m,m}^{P,Q})$, where P and Q are the distribution that samples m messages randomly without replacement.*

Corollary V.2. *Let $m, n \in \mathbb{N}$ and let P, Q be arbitrary distributions over $(\{0, 1\}^n)^m$. Then, the expected key length of any protocol in the $\text{rABB}_{m,m}^{P,Q}$ is at most $\log \binom{2m}{m}$.*

F. 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. \square

G. 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 key-agreement protocol, with the key difference being P and Q allows sample with replacement, i.e. duplicate messages from the same party 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.




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.

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, \dots, z_{N-1} as z_i being the amount of the element i showing up in Γ . Similarly, define x_1, \dots, x_{N-1} as x_i being the amount of i showing up in A , and y_1, \dots, 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_i s

and y_j s satisfying the equations and constraints

$$\begin{aligned} \sum_{i=0}^{N-1} x_i &= m \\ \sum_{i=0}^{N-1} y_i &= m \\ \forall i, x_i + y_i &= z_i \\ \forall i, x_i &\geq 0 \\ \forall i, y_i &\geq 0 \end{aligned}$$

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 $\forall i, z_i \leq 1$, this reduces to a simple binomial problem with the solution being $\binom{2m}{m}$.

VI. RANDOM STRING OBLIVIOUS TRANSFER

In this section, we consider the construction of chosen message random oblivious transfer (cmROT) (see Section III for definitions) in the rABB-hybrid.

In the rABB-hybrid with the helper \mathcal{C} , we construct an efficient 1-round protocol. We start by describing the problem setting as follows.

Problem Setting. \mathcal{A} and \mathcal{B} will like to establish a chosen message random string oblivious transfer (cmROT^ℓ) for ℓ -bit strings 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 public authenticated channel from \mathcal{B} to \mathcal{A} . We note that with the key-agreement protocol, \mathcal{A} and \mathcal{B} can turn the public authenticated channel into a private authenticated channel. Furthermore, by using the technique in Section IV-D, this key-agreement can be done in parallel with the first step of the cmROT^ℓ protocol and does not require any additional round.

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 shall prove the following theorem.

Theorem 6. *For any $\ell \in \{1, 2, \dots\}$, there is a perfectly secure 1-round protocol for cmROT^ℓ in the rABB-hybrid (with the helper).*

A. Construction

Intuition. We begin with some intuition. 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. Additionally, observe that this still holds if we randomly set

all messages from A to be even (or odd) and C to be odd (or even respectively), while messages from B contain both even and odd values.

Additionally, by using the technique discussed in Section IV-D, we can effectively perform several parallel calls to the rABB in the same round.

Formally, the parties utilize rABB ^{P,Q,R} , where $P = P_1 \cup P_2 \cup \dots \cup P_\sigma$, $Q = Q_1 \cup Q_2 \cup \dots \cup Q_\sigma$, and $R = R_1 \cup R_2 \cup \dots \cup R_\sigma$ for some parameter $\sigma \in \{1, 2, \dots\}$ (chosen appropriately later).

Independently, each P_i is in “even mode” with probability $\frac{1}{2}$, which means $P_i = P_{i,1}^{(EVEN)} \cup P_{i,2}^{(EVEN)} \cup \dots \cup P_{i,\ell}^{(EVEN)}$, where each $P_{i,j}^{(EVEN)}$ samples a n bit even value (with least significant bit being 0) uniformly at random, denoted as $\alpha_{i,j}^{(EVEN)}$, and outputs the tuple $(i, j, \alpha_{i,j}^{(EVEN)})$. Similarly, each P_i is in “odd mode” with probability $\frac{1}{2}$, which means $P_i = P_{i,1}^{(ODD)} \cup P_{i,2}^{(ODD)} \cup \dots \cup P_{i,\ell}^{(ODD)}$, where each $P_{i,j}^{(ODD)}$ samples a n bit odd value (with least significant bit being 1) uniformly at random, denoted as $\alpha_{i,j}^{(ODD)}$, and outputs the tuple $(i, j, \alpha_{i,j}^{(ODD)})$.

R_i is defined as an independent copy of P_i . Similarly, we denote the output of $Q_{i,j}^{(EVEN)}$ as $(i, j, \omega_{i,j}^{(EVEN)})$.

$Q_i = Q_{i,1} \cup Q_{i,2} \cup \dots \cup Q_{i,\ell}$, where each $Q_{i,j}$ independently samples one even n -bit value uniformly at random and samples one odd n -bit value uniformly at random, and outputs the set of the two tuples $\left\{ (i, j, \beta_{i,j}^{(EVEN)}), (i, j, \beta_{i,j}^{(ODD)}) \right\}$.

Parties first invoke rABB ^{P,Q,R} , which samples A according to distribution P , samples B according to distribution Q , and samples C according to distribution R . rABB ^{P,Q,R} then sends A to \mathcal{A} , B to \mathcal{B} , and C to \mathcal{C} , as well as send $\Gamma = A \cup B \cup C$ to $\mathcal{A}, \mathcal{B}, \mathcal{C}$.

A now has $\sigma\ell$ values in the form of $(i, j, \alpha_{i,j}^0)$. \mathcal{C} now has $\sigma\ell$ values in the form of $(i, j, \omega_{i,j}^0)$. \mathcal{B} now has $2\sigma\ell$ values in the form of $(i, j, \beta_{i,j}^0)$. All parties also see the set containing all values.

Using the index information, parties can locally separate the values according to i, j . Each i, j should now contain $(i, j, \alpha_{i,j}^0)$, $(i, j, \beta_{i,j}^{(EVEN)})$, $(i, j, \beta_{i,j}^{(ODD)})$, and $(i, j, \omega_{i,j}^0)$.

The parties will first look at $i, j = 1$, and find the lexicographically smallest i such that the four values they see contain exactly two even and two odd values. We denote this as i^* . The parties will now disregard all $i \neq i^*$, and focus only on i^*, j . Note that this means that P_{i^*} and R_{i^*} were in different modes, that is, if P_{i^*} was in even mode, then R_{i^*} was in odd mode, or vice versa. Furthermore, \mathcal{A} , having access to $\alpha_{i^*,j}^0$, can identify whether P_{i^*} was in even mode or odd mode. Without loss of generality, let us assume P_{i^*} was in even mode, and R_{i^*} is in odd mode. Each i^*, j now contains $(i^*, j, \alpha_{i^*,j}^{(EVEN)})$, $(i^*, j, \beta_{i^*,j}^{(EVEN)})$, $(i^*, j, \beta_{i^*,j}^{(ODD)})$, and $(i^*, j, \omega_{i^*,j}^{(ODD)})$.

\mathcal{B} now computes the two “intermediate message” y_{EVEN} and y_{ODD} (we can equivalently think of them as y_0 and y_1). To compute the j^{th} bit of y_{EVEN} , denoted as $y_{EVEN,j}$, \mathcal{B} looks at $(i^*, j, \beta_{i^*,j,EVEN})$ and $(i^*, j, \alpha_{i^*,j,EVEN})$. If $\beta_{i^*,j,EVEN} \geq \alpha_{i^*,j,EVEN}$, then $y_{EVEN,j} = 0$, else $y_{EVEN,j} = 1$. Note that \mathcal{B} simply compares the other even value against $\beta_{i^*,j,EVEN}$. In particular, \mathcal{B} does not know whether he is comparing against $\alpha_{i^*,j,EVEN}$ or $\omega_{i^*,j,EVEN}$. Similarly, if $\beta_{i^*,j,ODD}$ is the greater of the two odd values ($\beta_{i^*,j,ODD} \geq \omega_{i^*,j,ODD}$ in this case), then $y_{ODD,j} = 0$, else $y_{ODD,j} = 1$.

\mathcal{B} can do this for all $i^*, 1$ to i^*, ℓ , and obtain two ℓ -bit messages y_{EVEN} and y_{ODD} . \mathcal{B} then computes ℓ -bit “key” r_{EVEN} such that $y_{EVEN} \oplus r_{EVEN} = x_{EVEN} = x_0$ and r_{ODD} such that $y_{ODD} \oplus r_{ODD} = x_{ODD} = x_1$.

Since \mathcal{A} has $\alpha_{i^*,j,EVEN}$, \mathcal{A} similarly computes y_{EVEN} by comparing the two even values against each other and setting $y_{EVEN,j} = 0$ if $\alpha_{i^*,j}^{(EVEN)}$ is smaller than $\beta_{i^*,j}^{(EVEN)}$, and 1 otherwise. Note that as \mathcal{A} cannot distinguish between $\beta_{i^*,j}^{(ODD)}$ and $\omega_{i^*,j}^{(ODD)}$, and thus cannot compute y_{ODD} .

Then, using a private authenticated channel, \mathcal{B} sends (r_{EVEN}, r_{ODD}) to \mathcal{A} , in that order. \mathcal{A} can now compute $x_0 = x_{EVEN} = y_{EVEN} \oplus r_{EVEN}$. Additionally, \mathcal{A} computes $b = 0$ if $\alpha_{i^*,j}$ is even, and $b = 1$ otherwise.

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.

Our Construction. Figure 7 presents our protocol constructing cmROT ^{ℓ} in the rrABB-hybrid. For every message of the form (i, j, u) , we call that (i, j) the identifier of the message and u the message's payload. The distributions P, Q, R are defined as follows. Let $\sigma, \ell \in \{1, 2, \dots\}$. Define $S_n^{(EVEN)} := \{x : x \in \{0, 1\}^n, x_n = 0\}$ – the set containing all n -bit even value and $S_n^{(ODD)} = \{x : x \in \{0, 1\}^n, x_n = 1\}$ – the set containing all n -bit odd value.

For every $1 \leq i \leq \sigma$ and $1 \leq j \leq \ell$, define $P_{i,j}^{(EVEN)}$ is the uniform distribution over the sample space

$$S_{i,j,n}^{(EVEN)} = \{(i, j, \alpha) : \alpha \in S_n^{(EVEN)}\}.$$

Similarly, define $P_{i,j}^{(ODD)}$ as the uniform distribution over the sample space

$$S_{i,j,n}^{(ODD)} = \{(i, j, \alpha) : \alpha \in S_n^{(ODD)}\},$$

and define $Q_{i,j}$ as the uniform distribution over the sample space

$$T_{i,j,n} := S_{i,j,n}^{(EVEN)} \times S_{i,j,n}^{(ODD)}.$$

Now, for every $1 \leq i \leq \sigma$, define P_i is the uniform distribution over the sample space

$$S_{i,n} := S_{i,1,n}^{(EVEN)} \times S_{i,2,n}^{(EVEN)} \times \dots \times S_{i,\ell,n}^{(EVEN)} \cup S_{i,1,n}^{(ODD)} \times S_{i,2,n}^{(ODD)} \times \dots \times S_{i,\ell,n}^{(ODD)}.$$

Define R_i as an i.i.d of P_i . For every $1 \leq i \leq \sigma$, define Q_i as the joint distribution of independent random variables X_1, X_2, \dots, X_ℓ distributed according to $Q_{i,1}, Q_{i,2}, \dots, Q_{i,\ell}$, respectively.

Next, define P as the joint distribution $(P_1, P_2, \dots, P_\ell)$. Similarly, $Q := (Q_1, Q_2, \dots, Q_\ell)$ and $R := (R_1, R_2, \dots, R_\ell)$.

B. Correctness and Security Proofs

We provide a high-level proof overview of the security and correctness of our protocol.

\mathcal{A} and \mathcal{B} can both distinguish between $\beta_{i^*,j,EVEN}$ and $\alpha_{i^*,j,EVEN}$ (or $\beta_{i^*,j,ODD}$ and $\alpha_{i^*,j,ODD}$) and can thus agree on the same message x_b .

A corrupt \mathcal{A} cannot distinguish between $\beta_{i^*,j,EVEN}$ and $\omega_{i^*,j,EVEN}$ (or $\beta_{i^*,j,ODD}$ and $\omega_{i^*,j,ODD}$), therefore \mathcal{A} does not learn anything about the x_{1-b} .

A corrupt \mathcal{B} cannot distinguish between the case where the two values are $\alpha_{i^*,j,EVEN}$ and $\omega_{i^*,j,ODD}$, or the case where the two values are $\alpha_{i^*,j,ODD}$ and $\omega_{i^*,j,EVEN}$. Therefore, \mathcal{B} learns nothing about which message \mathcal{A} received (the bit b).

A corrupt \mathcal{C} does not see r_{EVEN} and r_{ODD} . Therefore, although \mathcal{C} can learn one of y_{EVEN}, y_{ODD} , \mathcal{C} learns nothing about x_{EVEN} or x_{ODD} .

While our protocol does have a small failure probability, all such failures are publicly detectable. Conditioned on the fact that failure does not occur, our protocol is unconditionally secure. Furthermore, given the protocol output, a reverse sampling of the view of the parties is efficient. Therefore, we can trivially construct a simulator that simulates the view of the corrupt party.

C. Performance Analysis

We provide a brief discussion on how our protocol compares to the protocol presented in [24]. As stated in [24], once we can obtain key agreement using rABB, we can then implement general honest majority MPC to obtain primitives such as random OT.

Concretely, to get random OT assuming shared keys between all pairs of \mathcal{A} , \mathcal{B} , and \mathcal{C} , at least two rounds of communication are required. In contrast, our protocol achieves chosen message random OT in a single round of communication.

Regarding our communication cost, each party receives at most $2 \cdot \sigma \cdot \ell \cdot n \cdot \log(\sigma) \cdot \log(\ell)$ bits from the rABB (excluding Γ) and sends at most $2 \cdot \ell$ bits through the private authenticated channel. Therefore, our communication cost is $O(\sigma \ell \log(\sigma) \log(\ell) n)$.

The failure probability of our protocol is upper bounded by $2^{-\sigma}$ (where all P_i and Q_i are in the same mode) plus

Input. \mathcal{B} has input $(x_0, x_1) \in (\{0, 1\}^\ell)^2$ and \mathcal{A} has no inputs.

Hybrid. Parties are in $rABB^{P,Q,R}$ -hybrid with appropriate P, Q, R such that

- 1) \mathcal{A} has a set A containing $\sigma \ell$ messages of the form (i, j, u) with distinct identifiers (i, j) and payloads (u 's) are all even or all odd, where $1 \leq i \leq \sigma$ and $1 \leq j \leq \ell$.
- 2) \mathcal{B} has a set B containing $2\sigma \ell$ messages such that, for any $1 \leq i \leq \sigma$ and $1 \leq j \leq \ell$, there are two messages with identifier (i, j) such that their payloads have different parity.
- 3) \mathcal{C} has a set C containing $\sigma \ell$ messages of the form (i, j, u) with distinct identifiers (i, j) and payloads (u 's) are all even or all odd.

Every party also receives $\Gamma = A \cup B \cup C$.

1-Round Protocol.

- Both \mathcal{A} and \mathcal{B} identify the smallest $i^* \in \{1, 2, \dots, \sigma\}$ such that the payloads of any message in A and any message in C have different parity.
- For $1 \leq j \leq \ell$, from the two sets B and Γ , party \mathcal{B} identifies the four messages with identifier (i^*, j) . There are exactly two of them in B and exactly two of them whose payloads are even. Then, he sets $y_{0,j} = 0$ if his even payload is bigger than the other even one and $y_{0,j} = 1$ otherwise. Similarly, he sets $y_{1,j} = 0$ if his odd payload is bigger than the other odd and $y_{1,j} = 1$ otherwise.
- \mathcal{B} sends (r_0, r_1) to \mathcal{A} using the private authenticated channel, where $r_0 = y_0 \oplus x_0, r_1 = y_1 \oplus x_1$.

Output Computation. \mathcal{A} receives (r_0, r_1) from \mathcal{B} . For $1 \leq j \leq \ell$, party \mathcal{A} identifies the two messages with identifier (i^*, j) such that (1) one of them is in A and (2) their payloads have the same parity.

- 1) Case 1: If the two payloads are even, \mathcal{A} sets $\tilde{y}_{0,j} = 0$ if his payload is smaller than the other and sets $\tilde{y}_{0,j} = 1$ otherwise.
- 2) Case 2: If the two payloads are odd, \mathcal{A} sets $\tilde{y}_{1,j} = 0$ if his payload is smaller than the other and sets $\tilde{y}_{1,j} = 1$ otherwise.

Then \mathcal{A} outputs $(b = 0, \tilde{y}_0 \oplus r_0)$ in case 1, and \mathcal{A} outputs $(b = 1, \tilde{y}_1 \oplus r_1)$ in case 2. In any case, \mathcal{B} outputs nothing.

Fig. 7. Realizing cmROT^ℓ in $rABB^{P,Q,R}$ -hybrid with appropriately chosen independent distributions P, Q, R . The parameter σ is chosen large enough so that the probability for the existence of such i^* is negligible.

$1 - (1 - 2^{n-1})^{2^\ell}$ (probability of having at least 1 collision in the messages).

An anonymous reviewer also proposed an alternate protocol for ROT. We will briefly present their proposed protocol here and then compare our protocol to it.

Similar to the reviewer proposed key agreement protocol, recall that in [24], Ishai et al. presented a non-interactive SUM protocol, in which two parties each send the additive shares of their values to the ABB. The two parties can sum up all messages (shares), and then subtract their value to learn the other's value, while both values are statistically hidden from the adversary. This protocol can be used to allow one party to send a secret/private message to another. To achieve ROT, the helper party \mathcal{C} can pick random m_0, m_1, b , and send (m_0, m_1) to \mathcal{B} and (b, m_b) to \mathcal{A} . We note that although in this case, the helper party \mathcal{C} will know m_0, m_1, b , if we follow Corollary I.2, then both the chosen message as well as the choice bit can be hidden from \mathcal{C} .

As stated in the key agreement, we want to point out that the SUM protocol is only statistically secure, while our protocol is unconditionally secure. While having the same round complexity (both being non-interactive), our protocol achieves a lower total communication cost (in bits). As stated in Lemma 4.1 from [24], they require each party share the messages into at least κ shares such that $\log \binom{2^\kappa}{\kappa} > \ell$, where ℓ is the length of the message in bits. This translates to a bound of k shares of length at least ℓ bits each in order to send a ℓ bit message. This means the proposed protocol will require $\Omega(\ell^2)$ bits of communication, while our protocol only requires $O(\ell \log(\ell))$ bits of total communication.

Remark 3 (Optimality of Round Complexity). *We note that as we can trivially get $\text{BEC}(0.5)$ from a chosen message random oblivious transfer; the existence of a non-interactive chosen message random oblivious transfer will imply the existence of a non-interactive BEC protocol, which contradicts Theorem 7. This proves that our chosen message random oblivious transfer is round optimal. We note that as this is based on Theorem 7, this round optimality only applies to uniform distributions. We note that the reviewer proposed protocol above is able to bypass this bound by using non-uniform distribution, and do not contradict this bound.*

D. Generalizations and Extensions

Construction for Corollary I.1 (Non-interactive Random String Oblivious Transfer) We extend our protocol to a non-interactive random string oblivious transfer where both the message and the choice are random. In our protocol, \mathcal{B} and \mathcal{A} used our non-interactive key-agreement protocol to agree on a random shared key that is used to establish a private authenticated channel in which \mathcal{B} is able to send r_{EVEN} and r_{ODD} to \mathcal{A} in order to determine the message. Instead of sending r_{EVEN} and r_{ODD} using the shared key, \mathcal{B} and \mathcal{A} can directly derive r_{EVEN} and r_{ODD} from the key in a non-interactive way. Since the shared key is random, the resulting r_{EVEN} and r_{ODD} will also be random. This results in a non-interactive random oblivious transfer protocol.

Construction for Corollary I.2 (2-round Chosen Message String Oblivious Transfer) Following standard techniques of obtaining cmOT from cmROT, the receiver will first send a bit to the sender, which tells the sender if he needs to swap the two correction messages or not. This allow us to obtain a 2-round chosen message oblivious transfer from a 1-round chosen message random oblivious transfer.

1-out-of- N Oblivious Transfer In our protocol, we partitioned values into even and odd, which results in a 1-out-of-2 cmROT $^\ell$. Our protocol can be generalized to any partitioning scheme Π_N where values are partitioned into N subsets. Our protocol will be modified to give \mathcal{A} one value belonging to one of the subsets, \mathcal{B} N value where each value belonging to a different subset (one value each from each subset), and \mathcal{C} $N - 1$ values that belong to all subsets except the one \mathcal{A} received. The rest of the protocol is modified accordingly. This allows us to achieve one round 1-out-of- N OT. However, note that the failure probability will now be $N^{-\sigma}$. In order to keep the failure probability low, σ needs to be increased accordingly, which increases the communication complexity of our protocol.

Alternatively, we can also use a standard technique of using $\log(N)$ 1-out-of-2 OTs to achieve 1-out-of- N OT.

Note that this result applies to all the different types of OT and ROT.

VII. ON THE LOWER BOUND OF BEC

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], [26], [27], to prove the following theorem.

Theorem 7. *Let $p \in (0, 1)$ be the erasure probability. Any zero round protocol implementing $\text{BEC}(p)$ in $\text{rABB}_{m_A, m_B, m_C}^{U_A, U_B, U_C}$ -hybrid has constant insecurity, where U_A, U_B, U_C are uniform distribution over $(\{0, 1\}^n)^{m_A}, (\{0, 1\}^n)^{m_B}, (\{0, 1\}^n)^{m_C}$ respectively, and n is the message length.*

Proof Sketch. We prove this by contradiction. Suppose that it is possible to get $\text{BEC}(p)$ from the $\text{rABB}_{m_A, m_B, m_C}^{U_A, U_B, U_C}$. It follows from [2], [26] that if it is possible to implement the randomized inputs $\text{BEC}(p)$ from some other distribution (X, Y) , then the eigenvalues of $\text{BEC}(p)$ must be a subset of eigenvalues of the distribution (X, Y) . Note that the eigenvalues of $\text{BEC}(p)$ are 1 and $\sqrt{1-p}$. The correlation $\text{rABB}_{m_A, m_B, m_C}$ 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.

We provide elaborated arguments on Appendix A. Additionally, we propose an alternative round-optimal protocol for achieving BEC in Appendix A

REFERENCES

- [1] Ittai Abraham, Benny Pinkas, and Avishay Yanai. Blinder - scalable, robust anonymous committed broadcast. In Jay Ligatti, Xinning 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] Shweta Agrawal, Yuval Ishai, Eyal Kushilevitz, Varun Narayanan, Manoj Prabhakaran, Vinod M. Prabhakaran, and Alon Rosen. Cryptography from one-way communication: On completeness of finite channels. In Shihō Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part III*, volume 12493 of *LNCS*, pages 653–685. Springer, Heidelberg, December 2020. doi:10.1007/978-3-030-64840-4_22.
- [4] 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.
- [5] 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.
- [6] 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.
- [7] Daniel Beer, Nov 2015. URL: <https://dlbeer.co.nz/articles/kwbs.html>.
- [8] 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.
- [9] Claude Castelluccia and Pars Mutaf. Shake them up! a movement-based pairing protocol for CPU-constrained devices. In *Third International Conference on Mobile Systems, Applications, and Services (MobiSys2005)*, Seattle, WA, June 2005. USENIX Association. URL: <https://www.usenix.org/conference/mobisys2005/shake-them-movement-based-pairing-protocol-cpu-constrained-devices>.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] 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.
- [14] Roberto Di Pietro and Gabriele Oliveri. Coke crypto-less over-the-air key establishment. *IEEE Transactions on Information Forensics and Security*, 8(1):163–173, 2013. doi:10.1109/TIFS.2012.2226718.
- [15] 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>.
- [16] 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.
- [17] 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>.
- [18] Uriel Feige, Joe Kilian, and Moni Naor. A minimal model for secure computation (extended abstract). In *26th ACM STOC*, pages 554–563. ACM Press, May 1994. doi:10.1145/195058.195408.
- [19] 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.
- [20] 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.
- [21] 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.
- [22] 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.
- [23] 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.
- [24] 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.
- [25] 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.
- [26] 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.
- [27] Hamidreza Amini Khorasgani, Hemanta K. Maji, and Hai H. Nguyen. Secure non-interactive simulation from arbitrary joint distributions. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part II*, volume 13748 of *LNCS*, pages 378–407. Springer, Heidelberg, November 2022. doi:10.1007/978-3-031-22365-5_14.
- [28] 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.
- [29] 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.
- [30] 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.
- [31] 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.
- [32] 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.
- [33] 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.
- [34] 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.
- [35] 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.
- [36] 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.
- [37] Michael O. Rabin. How to exchange secrets by oblivious transfer. *Technical Memo TR-81*, 1981.
- [38] Michael O. Rabin. How to exchange secrets with oblivious transfer. Cryptology ePrint Archive, Report 2005/187, 2005. <https://eprint.iacr.org/2005/187>.
- [39] 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>.
- [40] 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>.
- [41] 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.
- [42] Mordechai M. Yung. A secure and useful “keyless cryptosystem”. *Information Processing Letters*, 21(1):35–38, 1985. URL: <https://www.sciencedirect.com/science/article/pii/0020019085901061>, doi: [https://doi.org/10.1016/0020-0190\(85\)90106-1](https://doi.org/10.1016/0020-0190(85)90106-1).

On the Lowerbound of BEC

Theorem 8. *Suppose $p = e/d$ for $e, d \in \{1, 2, \dots\}$ and $e < d$. Any zero round protocol implementing $\text{BEC}(p)$ in $\text{rABB}_{d-e, 1, e}^{U_A, U_B, U_C}$ -hybrid has constant insecurity, where U_A, U_B, U_C are uniform distribution over $(\{0, 1\}^n)^{d-e}, \{0, 1\}^n, (\{0, 1\}^n)^e$ respectively, and n is the message length.*

Proof. We prove this by contradiction. Suppose there is a non-interactive secure protocol for $\text{BEC}(p)$ in rABB -hybrid. We introduce some terminologies and notations. Let $m = m_A + m_B + m_C$. Let A, B, C be random variables sampled according to the distributions U_A, U_B, U_C , respectively. Recall that $A \cup B \cup C$ contains no duplicate with high probability over the random choices of A, B, C . Let Γ be a set containing m distinct elements in $\{0, 1\}^n$. Consider the conditional distribution $(A, B | A \cup B \cup C = \Gamma)$. After removing all zero rows and columns, the probability mass function of $(A, B | A \cup B \cup C = \Gamma)$ is a matrix of size $\binom{m}{m_A} \times \binom{m}{m_B}$. From now, we refer to $(A, B | A \cup B \cup C = \Gamma)$ as the distribution after removing all these zero rows and columns. Observe that any protocol realizing $\text{BEC}(p)$ from $(A, B | A \cup B \cup C = \Gamma)$ has constant insecurity if there is no perfectly secure protocol realizing $\text{BEC}(p)$ from that distribution. Now, observe that for any two Γ and Γ' each containing m distinct elements, the probability mass functions $(A, B | A \cup B \cup C = \Gamma)$ and $(A, B | A \cup B \cup C = \Gamma')$ are the same (up to permutations of rows and columns). Therefore, there must exist a perfect secure protocol for $\text{BEC}(p)$ from $(A, B | A \cup B \cup C = \Gamma)$.

Now, we employ the technique developed recently in secure non-interactive simulation/reduction [26]. Let T and \bar{T} be the Markov and the adjoint Markov operator associated with the conditional distribution $(A, B | A \cup B \cup C = \Gamma)$ (refer to [26] for definitions).

There is a perfectly secure protocol if and only if there are functions f, g such that

$$Tg = f, \text{ and } \bar{T}f = (1 - p)g.$$

Combining two equations together yields $T\bar{T}f = (1 - p)f$. This implies that $(1 - p)$ is an eigenvalue of the operator $T\bar{T}$ with associated eigenvector f . We shall show that these two conditions yield a contradiction. Observe that any column of $T\bar{T}$ is a permutation of any other column of $T\bar{T}$. Following the approach in [26], [27], the function f must have only two output values 1 or -1 . These two facts together give the contradiction. \square

Binary Erasure Channel Rabin and Crépeau [11], [37], [38] 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.

A. Problem Setting.

Suppose parties are in the $rABB_{m_A, m_B, m_C}^{P, Q, R}$ hybrid (with the helper C). That is, party \mathcal{A} has $A = \{a_1, a_2, \dots, a_{m_A}\}$ sampled according to P , party \mathcal{B} has $B = \{b_1, b_2, \dots, b_{m_B}\}$ sampled according to Q , and \mathcal{C} has $C = \{c_1, c_2, \dots, c_{m_C}\}$ 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 $\text{BEC}(p)$ 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{B} will send a bit β to \mathcal{A} . The protocol is 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 perfectly 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. Furthermore, using the technique in Section IV-D, we can establish the private authenticated channel in parallel with the BEC protocol without using any additional rounds of communication.

Theorem 9. *Let $p \in (0, 1)$ be a rational number. There is a perfectly secure one-round protocol for $\text{BEC}(p)$ 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 Dirichlet's approximation algorithm. Therefore, we have the following result as a corollary.

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

B. Construction

We present our protocol in Figure 8, 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.

C. 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.

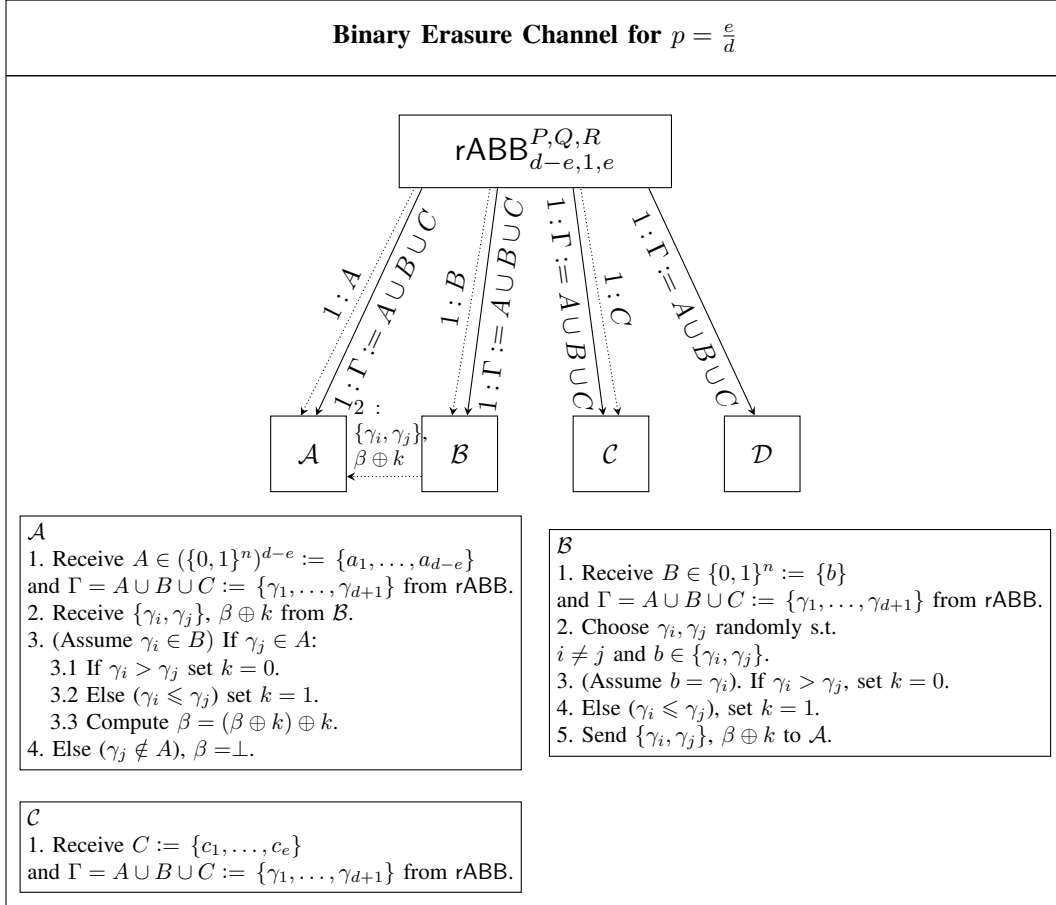


Fig. 8. Binary Erasure Channel from \mathcal{B} to \mathcal{A} using a helper \mathcal{C} and in presence of an eavesdropper \mathcal{D}