


Ad Hoc Broadcast, Trace, and Revoke

Plus Time-Space Trade-Offs for Attribute-Based Encryption

罗辑 (Ji Luo) 

Paul G. Allen School of Computer Science & Engineering,
University of Washington, Seattle, USA
luoji@cs.washington.edu

8 July 2024

Abstract

Traitor tracing schemes [Chor–Fiat–Naor, Crypto ’94] help content distributors fight against piracy and are defined with the content distributor as a trusted authority having access to the secret keys of all users. While the traditional model caters well to its original motivation, its centralized nature makes it unsuitable for many scenarios. For usage among mutually untrusted parties, a notion of *ad hoc* traitor tracing (naturally with the capability of broadcast and revocation) is proposed and studied in this work. Such a scheme allows users in the system to generate their own public/secret key pairs, without trusting any other entity. To encrypt, a list of public keys is used to identify the set of recipients, and decryption is possible with a secret key for any of the public keys in the list. In addition, there is a tracing algorithm that given a list of recipients’ public keys and a pirate decoder capable of decrypting ciphertexts encrypted to them, identifies at least one recipient whose secret key must have been used to construct the said decoder.

Two constructions are presented. The first is based on functional encryption for circuits (conceptually, obfuscation) and has constant-size ciphertext, yet its decryption time is linear in the number of recipients. The second is a generic transformation that reduces decryption time at the cost of increased ciphertext size. A matching lower bound on the trade-off between ciphertext size and decryption time is shown, indicating that the two constructions achieve all possible optimal trade-offs, i.e., they fully demonstrate the Pareto front of efficiency. The lower bound also applies to broadcast encryption (hence all mildly expressive attribute-based encryption schemes) and is of independent interest.

Keywords. *ad hoc*, decentralized, distributed, flexible, traitor tracing, broadcast encryption, attribute-based encryption, functional encryption, obfuscation.

Contents

1	Introduction	1
1.1	Overview	4
2	Preliminaries	7
3	<i>Ad Hoc</i> Broadcast, Trace, and Revoke	12
3.1	Simplified Security Notions	13
4	<i>Ad Hoc</i> Private Linear Broadcast Encryption	15
4.1	Construction	16
4.2	Security	18
5	AH-BTR from AH-PLBE	21
6	Trading Ciphertext Size for Decryption Time	24
7	Lower Bound on Ciphertext Size and Decryption Time	26
	References	30

1 Introduction

Traitor tracing schemes [CFN94] enable content distributors to fight against piracy. A content distributor such as a media streaming service can generate a public key and many different secret keys for individual subscribers, all of which can decrypt the ciphertexts created using the public key. Given a pirate decoder capable of decrypting, which could have been created from the secret keys of multiple subscribers, the tracing algorithm can find at least one subscriber (a traitor) whose key was used to create the said decoder. A long line of subsequent works [BSW06, BW06, BN08, BZ14, NWZ16, GKW18, GKRW18, CVW⁺18, GQWW19, GKW19, Zha20a, Zha21, GLW23] proposed the different security definitions, extended the functionality, and presented new constructions.

While the traditional model caters well to the needs of content distributors, its centralized nature makes it unsuitable for many scenarios, e.g., when a group of individuals want to communicate among themselves and trace traitors who provide decoders to outsiders. For example [Zha21], in an encrypted group chat among protesters, the users are worried about potential infiltration by government agents. To mitigate this concern, they want to trace traitors and remove them from the group. If they used a traditional traitor tracing scheme, whoever set it up would be able to impersonate anyone since they would know all the secret keys. Moreover, as words are spread and the protest gets wider support, more people need to join the group. The joining process should as simple as possible, preferably without interaction. This motivation naturally calls for a decentralized notion of traitor tracing, termed *ad hoc* traitor tracing in this work.

The first question is thus naturally the following:

What is the right notion of a secure ad hoc traitor tracing scheme?

Having formalized its syntax and security, we study its constructions:

*How can such a scheme be constructed,
from what assumptions, and with what efficiency?*

Efficiency improvement (in both size and time) never ends until we reach the optimum, for which it is necessary to understand where the limit stands:

What bounds are there on the efficiency of such schemes?

Our Contributions. We provide answers to all three questions.

- *Conceptually*, we pose the question of *ad hoc* traitor tracing, develop from the ideas thereof, and eventually arrive at the definitions for *ad hoc* broadcast, trace, and revoke (AH-BTR). We prove the relation among the security notions considered in this work.
- *Construction-wise*, we present secure AH-BTR schemes based on functional encryption for general circuits [BSW11]. With $\text{poly}(\lambda)$ factors ignored, they achieve

$$\begin{array}{ll} \text{encryption time} & T_{\text{Enc}} = O(N), \\ \text{ciphertext size} & |\text{ct}| = O(N^{1-\gamma}), \\ \text{decryption time} & T_{\text{Dec}} = O(N^\gamma), \end{array}$$

for any constant $0 \leq \gamma \leq 1$, where N is the number of recipients.

- *Questing for the ultimate efficiency*, we prove that for all secure AH-BTR,

$$|\text{ct}| \cdot T_{\text{Dec}} = \Omega(N),$$

so our schemes offer *all possible optimal* trade-offs between $|\text{ct}|$ and T_{Dec} , fully demonstrating the Pareto front of AH-BTR efficiency. Better yet, the bound holds for a restricted kind of weakly secure broadcast encryption [FN94], which is a specific case of attribute-based encryption [SW05, GPSW06]. Our result is the *first* space-time lower bound applicable to any computationally secure BE scheme, shedding new insights into the efficiency of ABE and BE.

A final addition is that our scheme is *compatible* with the existing public-key encryption schemes, i.e., the keys of such a scheme can be those of any secure public-key encryption, and there is no need to regenerate keys to take advantage of our scheme.

More on the Lower Bound. Most works on broadcast encryption have been focused on minimizing component sizes, motivating shorter ciphertexts with savings on broadcaster storage and bandwidth. Decryption time has been largely neglected. However, by our lower bound, a BE scheme with constant-size ciphertext could force each recipient to spend $\Omega(N)$ time on decryption and drive the total computational cost to $\Omega(N^2)$. In contrast, the naïve scheme encrypting to recipients individually has total cost only $O(N)$ in storage, communication, and computation.

This urges us to *rethink about the goals* of broadcast encryption — are short component sizes still the ultimate desideratum, given the high total cost? The lower bound, as an integral part of our results, shows that optimizing for one efficiency parameter might bring inefficiency in another, and calls the question of the *trade-offs* among multiple efficiency parameters of advanced forms of encryption into attention.

Open Questions. The tracing model in this work is black-box and classical, and recent works [Zha21, Zha20c] have studied white-box or quantum traitor tracing. It would be interesting to understand the *ad hoc* versions of those tracing models.

Another question for future investigation is whether AH-BTR can be constructed from more lightweight assumptions, such as group- or lattice-based ones, without going through obfuscation. This appears to require significant deviation from existing paradigms, as typical group-based or lattice-based constructions share public parameters among all parties so that their keys can be correlated and ciphertexts compressed, yet the motivation of AH-BTR repels any use of public parameters. See related works for more discussion on the technical challenges.

Related Works. We discuss them by aspects.

*Ad Hoc, Decentralized, Distributed, or Flexible Broadcast Encryption.*¹ Decentralized BE with interactive management of recipient sets was studied in [PPS12, DPP07]. *Ad hoc* (also known as distributed or flexible) BE was studied in several prior/later works. Schemes based on pairing [DHMR08, WQZD10, KMW23] or witness encryption [FWW23] require global set-up, and the obfuscation-based one [BZ14] do not.

¹While the details of the definition in each work differ, their common theme is that each recipient generates their own key pair.

Broadcast, Trace, and Revoke. BTR [NP01,NNL01] is also known as *broadcast and trace* or *trace and revoke*. Non-AH version of BTR supporting *public tracing* with optimal size currently is only known from functional encryption for general circuits (polynomial hardness, same as in this work) [AKYY23,JLL23] or witness encryption or obfuscation (non-falsifiable assumptions) [NWZ16,GVW19]. BTR is also known from pairing (standard assumption) [BW06,GKSW10] with $\Theta(\sqrt{N})$ -size components supporting *public tracing*, or from pairing (generic group model) [Zha20a] with various size trade-offs supporting *secret tracing* (still $\Omega(\sqrt{N})$ when size is balanced across components), or from both pairing (standard assumption) and LWE [GQWW19] with $O(N^\epsilon)$ -size ciphertexts for any constant $\epsilon > 0$ but having $\Omega(N)$ -size keys and only supporting *secret tracing*. Regardless of public or secret traceability, these schemes generate recipients’ decryption keys correlated by the master secret key, the major downside that AH-BTR intends to address.

Continuing the discussion of technical challenges in open questions, AH-BTR implies BTR supporting *public tracing* with the same ciphertext size. Therefore, it makes more sense to survey the techniques for public-tracing BTR schemes, than secret-tracing or non-BTR traitor tracing ones, in search of non-obfustopia instantiations. Filtered as such, the only schemes [BW06,GKSW10] with non-trivial (i.e., sublinear) ciphertext size are pairing-based and heavily rely on shared public parameter for key correlation (enabling ciphertext compression) — antithetical to the fully decentralized nature of AH-BTR. Even if we cease the insistence of having no centrally generated public parameters, the only known *ad hoc* type of (pairing-based) schemes are broadcast encryption [WQZD10,KMW23] without tracing, which are clever modifications of non-AH BE schemes. However, it is unclear how such adaptations can be done for [BW06, GKSW10] or, more generally, how AH-BTR can be constructed following any known pairing-based paradigms.

Registration-Based (Registered) Encryption. RBE [GHMR18,GHM⁺19,GV20,CES21,GKMR23,HLWW23,FWW23,FKdP23,FFM⁺23,ZZGQ23,ZLZ⁺24,GLWW24] is an emerging paradigm to decentralize functional encryption, where users generate their own key pairs and their public keys are aggregated for use during encryption. AH-BTR and RBE share similarities in motivation and techniques — e.g., typical constructions of both rely on laconic cryptography [CDG⁺17] to compress public keys.

We can conceive casting *ad hoc* private linear broadcast encryption, our building block of AH-BTR, as RBE for compare-index-and-reveal, yet there is no study of this functionality in the literature. Even under this view, RBE is not “ergonomic” to the usage pattern of AH-BTR and such reduction may suffer efficiency issues. The reason is that RBE requires distributing decryption updates (public information that, when used with user-generated secret keys, helps with decrypting ciphertexts encrypted using the aggregated public key) as the public keys are aggregated. RBE aggregation corresponds to the choice of recipients in AH-BTR, which happens at encryption time. Consequently, decryption updates from RBE would have to be included in every ciphertext, or every recipient must redo aggregation. Without further investigation, it is unclear whether the issue can be resolved. In this work, we study AH-BTR directly and do not try casting it under RBE.

Efficiency Parameters. Existing works on BE [FN94,BGW05,GW09,BWZ14,AY20,AWY20,BV22,Zha20a,Wee22] and its extensions [DPP07,Del07,SF07,BZ14] have been focused on improving the sizes of various components, and the time complexity has been largely

overlooked. In a rare exception, the work of [AL10] reduces the number of pairing operations during decryption down to constant, yet the overall decryption time is not among its concerns. This work brings the total decryption time into the picture.

Lower Bounds. Previous works [BC95,LS98,KYDB98,AK08,KY09,GKW15,DLY21] show a few efficiency lower bounds related to ABE and BE, yet they only apply to information-theoretically secure primitives and even specific construction techniques. Moreover, all of them prove space (ciphertext or secret key size, or their trade-off) lower bounds, whereas this work is about space-time trade-offs. Based on obfuscation [BWZ14] or both LWE and pairing [AY20], broadcast encryption with $|ct|, |sk| = O(1)$ can be achieved, circumventing all previously known bounds. A concurrent work [JLL23] proves lower bounds on (partially hiding) functional encryption, which is more expressive than BE and ABE and hence subject to stricter lower bounds than that in this work. The two works complement each other in understanding the efficiency trade-offs of advanced forms of encryption.

1.1 Overview

Developing Definitions. We start with the first principles of *ad hoc* traitor tracing. Syntactically, there should be a key generation algorithm that is run by each user of the system.² To encrypt, a list of public keys is used to identify the set of recipients. Decryption should only require one secret key from the list of public keys. In addition, the decryptor gets random access to all the recipients' public keys as well as the ciphertext. The choice to give random access to these inputs is based on performance concerns, as the decryptor might not have to read all of the public keys or the ciphertext.

It should be clear that such a scheme would automatically have the functionality of broadcast encryption [FN94]. There is no event prior to encryption that “binds” the system to a specific, fixed set of possible recipients, and the encryptor is free to use whatever public keys it sees fit. Similarly, the encryptor can remove any public key when it encrypts a second ciphertext, i.e., the scheme supports revocation. Therefore, the object is named *ad hoc* broadcast, trace, and revoke (AH-BTR).

As usual with broadcast encryption, we do not hide the list of recipients and provide the recipient list for free during decryption. Hiding the recipients makes ciphertext grow at least linearly with the number of recipients, diminishing the potential of efficiency. As we shall see, it is possible to construct AH-BTR with short ciphertexts.

Due to the decentralized nature of such systems, an adversary might indistinguishably generate malformed keys, which could potentially evade tracers that only take well-formed keys into account. To make it worse, a malformed key could be used to mount a denial-of-service attack against (other) honest users if it appears in the list of recipients' public keys during encryption — the encryption algorithm might have been carelessly designed and the presence of certain malformed keys could make it impossible to decrypt for anyone, including the recipients with honestly generated public keys.

In order to protect against such attacks by definition, we require correctness be *robust* against malformed keys — however, for performance reasons, namely to be able to index into any particular public key in constant time, we reject *blatantly* malformed keys, e.g., those of incorrect lengths, in the definition of correctness. This restriction does not hamper the usefulness of such a scheme.

²We aim for a scheme without any trusted party, so there should be no global set-up.

As for security, we naturally consider *public traceability*, i.e., no secret key is required to run the tracing algorithm. When attacking the scheme, the adversary is allowed to supply an arbitrary list of recipients' public keys, generated honestly by the challenger or (adversarially) by the adversary, so that the definition covers the scenario when (blatantly or not) malformed keys are present in the list of recipients' public keys. The tracing algorithm is given oracle access to a stateless³ decoder. It must *not accuse* an honest user, defined as one whose public key is generated by the challenger without its secret key revealed to the adversary. It *must find* a traitor as long as the decoder breaks semantic security (i.e., succeeds in decrypting with non-negligible probability), where *traitors* are associated with public keys in the recipient list that are either generated by the challenger yet having their secret keys revealed to the adversary or crafted by the adversary in any manner (e.g., skewed distribution, or even without a well-defined secret key).

The issues above must be identified and conceptually resolved (as done here) to arrive at suitable definitions accurately capturing the decentralized nature of AH-BTR.

Simplifying Security Notions. Traditionally [BSW06], traceability has been defined using one comprehensive *interactive* experiment,⁴ which is complicated to work with. Intuitively, the notion requires that *i*) a traitor should be found from a decoder with sufficient advantage and *ii*) no honest user should be identified as a traitor, regardless of the decoder advantage.

We therefore define two security notions for AH-BTR capturing the requirements separately. The former is called *completeness* and the latter is called *soundness*. Their conjunction is equivalent to *traceability*. Since only one requirement is considered in each notion, both of them can be vastly simplified and the security experiments become *non-interactive*. They are much more convenient for reductionist proofs.

Construction. Our first construction of AH-BTR adapts the blueprint of obtaining traitor tracing schemes from private linear broadcast encryption (PLBE) schemes introduced in [BSW06]. We consider *ad hoc* PLBE:⁵

- Everyone generates their own public and secret key pair (pk, sk) .
- Encryption uses a list $\{pk_j\}_{j \in [N]}$ of N public keys of the recipients as well as a cut-off index $0 \leq i_{\perp} \leq N$.
- Decryption is possible with sk_j if (and only if) $j > i_{\perp}$.

There are two security requirements. Message-hiding requires that the plaintext is hidden if $i_{\perp} = N$. Index-hiding requires that an adversary without sk_j for an *honest* pk_j cannot distinguish between cut-off index being $(j - 1)$ versus j .

Colloquially, the cut-off index i_{\perp} disables $sk_1, \dots, sk_{i_{\perp}}$, and the only way to detect whether an index is disabled is to have control over the corresponding key pair (by

³The general transformation [KY01,BSW06] to deal with stateful decoder applies to our definition of AH-BTR, *mutatis mutandis*.

⁴While some previous works [BF99,GKW18,Zha20a] separate traceability into multiple notions, each notion still requires interaction in its security experiment, due to the centralized nature of the set-up process of traditional traitor tracing.

⁵AH-PLBE can be cast as multi-authority attribute-based encryption [Cha07] for 1-local monotone functions without global set-up.

knowing sk or generating a malformed pk). When $i_{\perp} = N$, the plaintext should be hidden since all keys are disabled.

An AH-PLBE scheme gives rise to an AH-BTR scheme. The AH-BTR inherits the key generation and decryption algorithms from AH-PLBE. To perform AH-BTR encryption, simply encrypt using AH-PLBE with $i_{\perp} = 0$, disabling no key so that every recipient can decrypt. Given a pirate decoder with advantage at least ε , the tracing algorithm estimates its advantages with the cut-off index i_{\perp} being $0, 1, 2, \dots, N$, and identifies the recipient associated with pk_{i^*} as a traitor if the advantage changes by $\Omega(\varepsilon/N)$ when i_{\perp} increases from $(i^* - 1)$ to i^* .

For security, the message-hiding property translates to completeness, and index-hiding to soundness. It now remains to construct an AH-PLBE.

Constructing AH-PLBE. It is folklore that any public-key encryption (PKE) scheme can be used to construct a naïve PLBE by encrypting individually to each recipient. The individual ciphertext that corresponds to a disabled key encrypts garbage instead of the actual plaintext. This scheme is also *ad hoc*. The downside of it is that the ciphertext is of size $\Omega(N)$.

Our scheme uses obfuscation to help compressing the naïve PLBE ciphertext. The ciphertext will contain an obfuscated program, which, when evaluated at $j \in [N]$, allows us to recover the PKE ciphertext under pk_j . However, the obfuscated program itself cannot simply compute each PKE ciphertext if we want AH-PLBE ciphertexts of size $o(N)$, as there is not enough space in the program to encode all the public keys that have been independently generated. Instead, the program encodes a short hash bound to the long list of public keys while supporting computation on them.

Laconic oblivious transfer (OT) [CDG⁺17] serves the purpose. It allows compressing an arbitrarily long string D down to a fixed-length hash h with which one can efficiently perform oblivious transfer. The sender can encrypt messages L_0, L_1 to a hash h and an index m into D . The time to encrypt is independent of the length of D . The receiver will be able to obtain $L_{D[m]}$ by decrypting the laconic OT ciphertext.

During AH-PLBE encryption, we use laconic OT to compress the list of public keys. The obfuscated program in our AH-PLBE ciphertext, when evaluated at $j \in [N]$, will output *i*) a garbled circuit whose input (resp. output) is a PKE public key (resp. ciphertext) and *ii*) a bunch of laconic OT ciphertexts that decrypts to the labels so that the garbled circuit is evaluated at pk_j . Decryption proceeds in the obvious manner.

The obfuscated program size, thus the ciphertext size, can be made constant,⁶ because both the time to garble a PKE encryption circuit and the time of laconic OT encryptions are constant.

YOU CAN (NOT) OPTIMIZE. While our first construction enjoys constant-size ciphertext, its decryption algorithm runs in time $\Omega(N)$. Concretely, the laconic OT hash is a Merkle tree, and before performing laconic OT decryption, it is necessary to reconstruct the tree as it is not stored in the ciphertext. In contrast, the decryption time of the scheme implied by the naïve PLBE is constant in the RAM model, as it only looks at the relevant piece of the underlying PKE ciphertext.

⁶We ignore *fixed* polynomial factors in the security parameter. The point is that the size does not grow with N , the number of recipients. Furthermore, exact dependency on λ is only meaningful for concrete security, whereas this work focuses on polynomial security, in which scenario one can arbitrarily tune down such dependency by setting $\lambda' = \lambda^\varepsilon$ for any constant $\varepsilon > 0$, where λ' is the actual value of the security parameter to use for the algorithms without affecting polynomial security.

We can trade ciphertext size for decryption time by using the naïve PLBE on top of our construction. By grouping the recipients into $\Theta(N^{1-\gamma})$ sets of size $\Theta(N^\gamma)$ and using our basic construction over each set, we obtain a scheme with ciphertext size $\Theta(N^{1-\gamma})$ and decryption time $\Theta(N^\gamma)$. The core idea of this transformation was formalized as the user expansion compiler [Zha20a] in the context of traditional traitor tracing.

All the constructions we now know suffer from $|\text{ct}| \cdot T_{\text{dec}} = \Omega(N)$, where $|\text{ct}|$ is the ciphertext length and T_{dec} is the decryption time. It turns out that this bound necessarily holds for all secure AH-BTR, and the blame is on the functionality of broadcast encryption (not traitor tracing). Indeed, it is possible to make both $|\text{ct}|$ and T_{dec} constant in a traditional traitor tracing scheme [BZ14]. In existing broadcast encryption (or revocation) schemes [BGW05, Del07, GW09, BZ14, AY20, AWY20, BV22] for N users, encrypting to arbitrary subsets of size S or $(N - S)$ makes $|\text{ct}| \cdot T_{\text{dec}} = \Omega(S)$. It is precisely the capability to encrypt to many $(N/2)$ -subsets among N users that is the deal breaker, as we shall see in the formal proof. Interestingly, the adversary used in the proof simply runs the decryption algorithm with a *non-decrypting* key (while *lying* about the recipient set), so the bound holds as long as the scheme is not *blatantly* insecure.

We explain the ideas of our proof based on a corollary⁷ of a result [Unr07] dealing with random oracles in the presence of non-uniform advice. Let $S, T \geq 0$ be such that $ST \ll N$. The corollary says that for any adversary learning any S -bit function (advice) of a random string $R \xleftarrow{\$} \{0, 1\}^N$ and additionally (adaptively) querying at most T bits in R , it is “indistinguishable” to flip a bit in R at a random location after the advice is computed (using the non-flipped R) and before queries are answered, even if the index of the potentially flipped bit is revealed to the adversary after the advice is computed.

Back to AH-BTR. Imagine that there are $2N$ users in the system, associated with key pairs $(\text{pk}_{j,s}, \text{sk}_{j,s})$ for $j \in [N]$ and $s \in \{0, 1\}$. Consider a ciphertext ct encrypting a random plaintext to $\{\text{pk}_{j,R[j]}\}_{j \in [N]}$ for a random string R and regard ct as the advice. Suppose Y is either R itself or R flipped at index $i^* \xleftarrow{\$} [N]$. Let’s try decrypting ct using $\text{sk}_{i^*, Y[i^]}$ while *pretending* that ct is generated for Y . Each time the AH-BTR decryption algorithm wants to read pk_j , we probe $Y[j]$ and respond with $\text{pk}_{j, Y[j]}$. By way of contradiction, suppose $|\text{ct}| \cdot T_{\text{dec}} \ll N$, which would translate to $ST \ll N$ in the corollary.

By the correctness of AH-BTR, when Y is R itself, the attempted decryption should successfully recover the plaintext. From the corollary it follows that the other case (Y is R flipped at i^*) should also lead to successful recovery. But if $Y[i^*] = \neg R[i^*]$, by the security of AH-BTR, the attempted decryption must fail to recover the plaintext except for negligible probability, yielding a contradiction.

2 Preliminaries

We denote by $\lambda \in \mathbb{N}$ the security parameter, by $\text{poly}(\cdot)$ a polynomial function, and by $\text{negl}(\lambda)$ a negligible function of λ . Efficient algorithms are probabilistic random-access machines $M^w(x)$ of running time $\text{poly}(|x|, |w|)$. Efficient adversaries (in interactive experiments) are probabilistic Turing machines of (total) running time $\text{poly}(\lambda)$, with or without $\text{poly}(\lambda)$ -long advices. (All of the proofs in this work are uniform.) The advantage of \mathcal{A} in distinguishing Exp_0 and Exp_1 is $(\Pr[\text{Exp}_0^{\mathcal{A}}(1^\lambda) = 1] - \Pr[\text{Exp}_1^{\mathcal{A}}(1^\lambda) = 1])$.

⁷This corollary is also a lower bound of a probabilistic variant of Yao’s box problem [Yao90] (generalized and studied in [CHK22]), on which our proof can be alternatively based.

We write $\approx, \approx_s, \equiv$ for computational indistinguishability, statistical indistinguishability, and identity.

Under the standard assumption that a pseudorandom generator (with polynomial security) exists, we can assume, whenever convenient, that a randomized algorithm uses a uniformly random λ -bit string as its randomness (without losing polynomial security considered in this work or degrading its efficiency).

For $n, n' \in \mathbb{N}$, we write $[n..n']$ for the set $\{n, \dots, n'\}$, and $[n]$ for $[1..n]$. For a bit-string D , we denote by $|D|$ its bit-length, and given an index $m \in [|D|]$, we denote by $D[m]$ the m^{th} bit of D . For two bit-strings D, D' , their concatenation is $D \| D'$. Given a circuit $C : \{0, 1\}^{n+M_0} \rightarrow \{0, 1\}^{n'}$ and $w \in \{0, 1\}^n$, we define $C[w]$ to be C with w hardwired as its first portion of input, so $C[w](x) = C(w \| x)$. For an event X , its indicator random variable is $\mathbb{1}_X$. For events X, Y in the same probability space, “ X implies Y ” means $X \subseteq Y$.

Garbled Circuits. The following version of partially hiding garbling [IW14] suffices for the purpose of this work.

Definition 1 (garbled circuit [Yao86,LP09,BHR12,IW14]). A *circuit garbling scheme* consists of 2 efficient algorithms.

- $\text{Garble}(1^\lambda, C, w)$ takes as input a circuit $C : \{0, 1\}^{n+M_0} \rightarrow \{0, 1\}^{n'}$ and some hardwired input $w \in \{0, 1\}^n$. It outputs a garbled circuit \widehat{C} and M_0 pairs of labels $L_{m_0,b} \in \{0, 1\}^\lambda$ for $m_0 \in [M_0], b \in \{0, 1\}$.
- $\text{Eval}(1^\lambda, \widehat{C}, x, \{L_{m_0}\}_{m_0 \in [M_0]})$ takes as input a garbled circuit, a non-hardwired input, and M_0 labels. It outputs an n' -bit string.

The scheme must be *correct*, i.e., for all $\lambda \in \mathbb{N}$, $n, M_0, n' \in \mathbb{N}$, $C : \{0, 1\}^{n+M_0} \rightarrow \{0, 1\}^{n'}$, $w \in \{0, 1\}^n$, $x \in \{0, 1\}^{M_0}$,

$$\Pr \left[\begin{array}{l} (\widehat{C}, \{L_{m_0,b}\}_{m_0 \in [M_0], b \in \{0,1\}}) \stackrel{\$}{\leftarrow} \text{Garble}(1^\lambda, C, w) \\ : \text{Eval}(1^\lambda, \widehat{C}, x, \{L_{m_0,x[m_0]}\}_{m_0 \in [M_0]}) = C[w](x) \end{array} \right] = 1.$$

Definition 2 (garbled circuit security [Yao86,LP09,BHR12,IW14]). Let $(\text{Garble}, \text{Eval})$ be a circuit garbling scheme (Definition 1). A *simulator* is an efficient algorithm

$$\text{SimGarble}(1^\lambda, C : \{0, 1\}^{n+M_0} \rightarrow \{0, 1\}^{n'}, x \in \{0, 1\}^{M_0}, y \in \{0, 1\}^{n'}) \rightarrow (\widehat{C}, \{L_{m_0}\}_{m_0 \in [M_0]})$$

taking as input a circuit, a non-hardwired input, and a circuit output, and producing a simulated garbled circuit and M_0 simulated labels. The scheme is *w-hiding* (or *secure* for the purpose of this work) if there exists a simulator SimGarble such that $\text{Exp}_{\text{GC}}^0 \approx \text{Exp}_{\text{GC}}^1$, where $\text{Exp}_{\text{GC}}^b(1^\lambda)$ with adversary \mathcal{A} proceeds as follows.

- **Challenge.** Launch $\mathcal{A}(1^\lambda)$ and receive a circuit $C : \{0, 1\}^{n+M_0} \rightarrow \{0, 1\}^{n'}$, a hardwired input $w \in \{0, 1\}^n$, and a non-hardwired input $x \in \{0, 1\}^{M_0}$ from it. Run

$$\begin{array}{ll} \text{if } b = 0, & (\widehat{C}, \{L_{m_0,b}\}_{m_0 \in [M_0], b \in \{0,1\}}) \stackrel{\$}{\leftarrow} \text{Garble}(1^\lambda, C, w); \\ \text{if } b = 1, & (\widehat{C}, \{L_{m_0,x[m_0]}\}_{m_0 \in [M_0]}) \stackrel{\$}{\leftarrow} \text{SimGarble}(1^\lambda, C, x, C[w](x)); \end{array}$$

and send $(\widehat{C}, \{L_{m_0,x[m_0]}\}_{m_0 \in [M_0]})$ to \mathcal{A} .

- **Guess.** \mathcal{A} outputs a bit b' , which is the output of the experiment.

Puncturable Pseudorandom Function. We rely on PPRF [BW13,KPTZ13,BGI14,SW14].

Definition 3 (PPRF [BW13,KPTZ13,BGI14,SW14]). A *puncturable pseudorandom function (PPRF) family* (with key space, domain, and codomain $\{0,1\}^\lambda$) consists of 2 efficient algorithms.

- $\text{Puncture}(1^\lambda, k \in \{0,1\}^\lambda, x)$ takes as input a non-punctured key and a point. It outputs a punctured key $\overset{\circ}{k}_x$.
- $\text{Eval}(1^\lambda, k, x \in \{0,1\}^\lambda)$ takes as input a (punctured or non-punctured) key and a point. It is deterministic and outputs a λ -bit string.

The scheme must be *correct*, i.e., for all $\lambda \in \mathbb{N}$, $x, x' \in \{0,1\}^\lambda$ such that $x \neq x'$,

$$\Pr \left[\begin{array}{l} k \xleftarrow{\$} \{0,1\}^\lambda \\ \overset{\circ}{k}_x \xleftarrow{\$} \text{Puncture}(1^\lambda, k, x) \end{array} : \text{Eval}(1^\lambda, k, x') = \text{Eval}(1^\lambda, \overset{\circ}{k}_x, x') \right] = 1.$$

Definition 4 (PPRF security [BW13,KPTZ13,BGI14,SW14]). A PPRF (Puncture, Eval) per Definition 3 is *pseudorandom at the punctured point* (or *secure* for the purpose of this work) if $\text{Exp}_{\text{PPRF}}^0 \approx \text{Exp}_{\text{PPRF}}^1$, where $\text{Exp}_{\text{PPRF}}^b(1^\lambda)$ with adversary \mathcal{A} proceeds as follows.

- **Challenge.** Launch $\mathcal{A}(1^\lambda)$ and receive from it a point $x \in \{0,1\}^\lambda$. Run

$$k \xleftarrow{\$} \{0,1\}^\lambda, \quad \overset{\circ}{k}_x \xleftarrow{\$} \text{Puncture}(1^\lambda, k, x), \quad r_0 \xleftarrow{\$} \{0,1\}^\lambda, \quad r_1 \xleftarrow{\$} \{0,1\}^\lambda,$$

and send $(\overset{\circ}{k}_x, r_b)$ to \mathcal{A} .

- **Guess.** \mathcal{A} outputs a bit b' , which is the output of the experiment.

Public-Key Encryption. Our *ad hoc* broadcast, trace, and revoke scheme can be based on any public-key encryption scheme.

Definition 5 (PKE). A *public-key encryption (PKE) scheme* (with message space $\{0,1\}^\lambda$ and public key length $M_0(\lambda)$) consists of 3 efficient algorithms.

- $\text{Gen}(1^\lambda)$ outputs a pair (pk, sk) of public and secret keys with $|\text{pk}| = M_0(\lambda)$.
- $\text{Enc}(1^\lambda, \text{pk}, \mu \in \{0,1\}^\lambda)$ takes as input the public key and a message. It outputs a ciphertext ct.
- $\text{Dec}(1^\lambda, \text{sk}, \text{ct})$ takes as input the secret key and a ciphertext. It outputs a message.

The scheme must be *correct*, i.e., for all $\lambda \in \mathbb{N}$, $\mu \in \{0,1\}^\lambda$,

$$\Pr \left[\begin{array}{l} (\text{pk}, \text{sk}) \xleftarrow{\$} \text{Gen}(1^\lambda) \\ \text{ct} \xleftarrow{\$} \text{Enc}(1^\lambda, \text{pk}, \mu) \end{array} : \text{Dec}(1^\lambda, \text{sk}, \text{ct}) = \mu \right] = 1.$$

Definition 6 (PKE security). A PKE scheme (Gen, Enc, Dec) per Definition 5 is *semantically secure for random messages* (or *secure* for the purpose of this work) if

$$\{(1^\lambda, \mu_0, \mu_1, \text{pk}, \text{ct}_0)\}_{\lambda \in \mathbb{N}} \approx \{(1^\lambda, \mu_0, \mu_1, \text{pk}, \text{ct}_1)\}_{\lambda \in \mathbb{N}},$$

where $(\text{pk}, \text{sk}) \xleftarrow{\$} \text{Gen}(1^\lambda)$ and $\mu_b \xleftarrow{\$} \{0,1\}^\lambda$, $\text{ct}_b \xleftarrow{\$} \text{Enc}(1^\lambda, \text{pk}, \mu_b)$ for $b \in \{0,1\}$.

Laconic Oblivious Transfer. We rely on laconic oblivious transfer [CDG⁺17].

Definition 7 (laconic OT [CDG⁺17]). A *laconic oblivious transfer (OT) scheme* (with message space $\{0, 1\}^\lambda$) consists of 4 efficient algorithms.

- $\text{Gen}(1^\lambda, M \in \mathbb{N})$ takes the database length as input and outputs a hash key hk .
- $\text{Hash}(1^\lambda, \text{hk}, D \in \{0, 1\}^M)$ takes as input a hash key and a database. The algorithm is deterministic, runs in time $(M + 1) \text{poly}(\lambda, \log(M + 1))$, and outputs a hash h of length $\text{poly}(\lambda, \log(M + 1))$ and a processed database \widehat{D} .
- $\text{Send}(1^\lambda, \text{hk}, h, m \in [M], L_0 \in \{0, 1\}^\lambda, L_1 \in \{0, 1\}^\lambda)$ takes as input a hash key, a hash, an index, and two labels (messages). It outputs a ciphertext ct .
- $\text{Recv}^{\widehat{D}}(1^\lambda, \text{hk}, h, m \in [M], \text{ct})$ is given random access to a processed database, and takes as input a hash key, a hash, an index, and a ciphertext. The algorithm runs in time $\text{poly}(\lambda, \log(M + 1))$ and outputs a label (message).

The scheme must be *correct*, i.e., for all $\lambda, M \in \mathbb{N}, D \in \{0, 1\}^M, m \in [M], L_0, L_1 \in \{0, 1\}^\lambda$,

$$\Pr \left[\begin{array}{l} \text{hk} \xleftarrow{\$} \text{Gen}(1^\lambda, M) \\ (h, \widehat{D}) \leftarrow \text{Hash}(1^\lambda, \text{hk}, D) \\ \text{ct} \xleftarrow{\$} \text{Send}(1^\lambda, \text{hk}, h, m, L_0, L_1) \end{array} : \text{Recv}^{\widehat{D}}(1^\lambda, \text{hk}, h, m, \text{ct}) = L_{D[m]} \right] = 1.$$

We only need database-selective security [AL18]. The following indistinguishability-based definition is equivalent to the usual simulation-based formulation.

Definition 8 (laconic OT security [CDG⁺17, AL18, KNTY19]). A laconic OT scheme $(\text{Gen}, \text{Hash}, \text{Send}, \text{Recv})$ per Definition 7 is *database-selectively sender-private* (or *secure* for the purpose of this work) if $\text{Exp}_{\text{LOT}}^0 \approx \text{Exp}_{\text{LOT}}^1$, where $\text{Exp}_{\text{LOT}}^b(1^\lambda)$ with adversary \mathcal{A} proceeds as follows.

- **Setup.** Launch $\mathcal{A}(1^\lambda)$ and receive from it some $M \in \mathbb{N}$ and a database $D \in \{0, 1\}^M$.
Run

$$\text{hk} \xleftarrow{\$} \text{Gen}(1^\lambda, M), \quad (h, \widehat{D}) \leftarrow \text{Hash}(1^\lambda, \text{hk}, D),$$

and send hk to \mathcal{A} .

- **Challenge.** \mathcal{A} submits an index $m \in [M]$ and two labels (messages) $L_0, L_1 \in \{0, 1\}^\lambda$.
Run

$$\text{ct} \xleftarrow{\$} \begin{cases} \text{Send}(1^\lambda, \text{hk}, h, m, L_0, L_1), & \text{if } b = 0; \\ \text{Send}(1^\lambda, \text{hk}, h, m, L_{D[m]}, L_{D[m]}), & \text{if } b = 1; \end{cases}$$

and send ct to \mathcal{A} .

- **Guess.** \mathcal{A} outputs a bit b' , which is the output of the experiment.

Obfuscation. We rely on indistinguishability obfuscator for polynomial-sized domain.

Definition 9 ((circuit) obfuscator [BGI⁺01]). A (circuit) obfuscator is an efficient algorithm $\text{Obf}(1^\lambda, C)$ taking a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^{n'}$ as input and producing an obfuscated circuit $\tilde{C} : \{0, 1\}^n \rightarrow \{0, 1\}^{n'}$ as output. The scheme must be *correct*, i.e., for all $\lambda \in \mathbb{N}$, $n, n' \in \mathbb{N}$, $C : \{0, 1\}^n \rightarrow \{0, 1\}^{n'}$, $x \in \{0, 1\}^n$,

$$\Pr[\text{Obf}(1^\lambda, C)(x) = C(x)] = 1.$$

Definition 10 (*iO* [BGI⁺01] for $\text{poly}(\lambda)$ -sized domain). An obfuscator Obf (Definition 9) is an *indistinguishability obfuscator for polynomial-sized domain* (*iO* for $\text{poly}(\lambda)$ -sized domain) if $\text{Exp}_{iO}^0 \approx \text{Exp}_{iO}^1$, where $\text{Exp}_{iO}^b(1^\lambda)$ with adversary \mathcal{A} proceeds as follows.

- **Challenge.** Launch $\mathcal{A}(1^\lambda)$ and receive from it the domain size 1^{2^n} and two circuits $C_0, C_1 : \{0, 1\}^n \rightarrow \{0, 1\}^{n'}$. Send $\text{Obf}(1^\lambda, C_b)$ to \mathcal{A} .
- **Guess.** \mathcal{A} outputs a bit b' . The output of the experiment is b' if C_0, C_1 have the same (description) size and $C_0(x) = C_1(x)$ for all $x \in \{0, 1\}^n$. Otherwise, the output is set to 0.

Assumption. All of the primitives defined in this section are implied by the existence of weakly selectively secure, single key, and sublinearly succinct public-key functional encryption for general circuits (so-called *obfuscation-minimum PKFE*), of which we refer the reader to [KNTY19] for the precise definition.

Lemma 1. *Suppose there exists an obfuscation-minimum PKFE with polynomial security, then there exist*

- [Yao86,LP09,BHR12] a secure circuit garbling scheme (Definitions 1 and 2),
- [GGM84,BW13,KPTZ13,BGI14] a secure PPRF (Definitions 3 and 4),
- [folklore] a secure PKE scheme (Definitions 5 and 6),
- [CDG⁺17,LZ17,AL18,KNTY19] a secure laconic OT scheme (Definitions 7 and 8), and
- [LT17,LZ17] an *iO* for $\text{poly}(\lambda)$ -sized domain (Definitions 9 and 10),

with polynomial security.

Alternatively, those primitives can be based on the existence of *iO* and one-way function. However, *iO* security (for circuits whose domains are not necessarily $\text{poly}(\lambda)$ -sized) is not known to be *falsifiable* [GW11] and it is hard to conceive [GGSW13] a reduction of *iO* security to *complexity assumptions* [GK16]. Since all of the security notions defined in this section are falsifiable, it is unsatisfactory to base them on *iO* from a theoretical point of view.

In contrast, obfuscation-minimum PKFE security is falsifiable and there are constructions [JLS21,JLS22] from well-studied complexity assumptions. The point of Lemma 1 is to base our constructions solely on one falsifiable assumption, or even complexity assumptions.

3 Ad Hoc Broadcast, Trace, and Revoke

This section concerns the definitions for *ad hoc* broadcast, trace, and revoke. After formally defining the syntax and correctness of AH-BTR, we present an intuitive definition of its security. While that definition is comprehensive, it is not the easiest to work with, so we turn to define two simpler security notions, whose conjunction is equivalent to the comprehensive definition. The proof of their equivalence follows the definitions. Later in this paper, we will only work with the simpler notions.

Definition 11 (AH-BTR). An *ad hoc broadcast, trace, and revoke (AH-BTR) scheme* (with message space $\{0, 1\}^\lambda$ and public key length $M_0(\lambda)$) consists of 4 efficient algorithms.

- $\text{Gen}(1^\lambda)$ outputs a pair (pk, sk) of public and secret keys with $|\text{pk}| = M_0(\lambda)$.
- $\text{Enc}(1^\lambda, \{\text{pk}_j\}_{j \in [N]}, \mu \in \{0, 1\}^\lambda)$ takes as input a list of public keys and a message. It outputs a ciphertext ct .
- $\text{Dec}^{\{\text{pk}_j\}_{j \in [N]}, \text{ct}}(1^\lambda, N, i \in [N], \text{sk}_i)$ is given random access to a list of public keys and a ciphertext, and takes as input the length of the list, an index, and a secret key. It outputs a message.
- $\text{Trace}^{\mathcal{D}}(1^\lambda, \{\text{pk}_j^*\}_{j \in [N]}, 1^{1/\epsilon^*})$ is given oracle access to a (stateless randomized) distinguisher \mathcal{D} and takes as input a list of public keys and an error bound (in unary). It outputs an index $i^* \in \{\perp\} \cup [N]$.⁸

The scheme must be *robustly correct*, i.e., for all $\lambda \in \mathbb{N}$, $N \in \mathbb{N}$, $i \in [N]$, $\{\text{pk}_j\}_{j \in [N] \setminus \{i\}}$ ⁹ such that $|\text{pk}_j| = M_0(\lambda)$ for all $j \in [N] \setminus \{i\}$, and $\mu \in \{0, 1\}^\lambda$,

$$\Pr \left[\begin{array}{l} (\text{pk}_i, \text{sk}_i) \stackrel{\$}{\leftarrow} \text{Gen}(1^\lambda) \\ \text{ct} \stackrel{\$}{\leftarrow} \text{Enc}(1^\lambda, \{\text{pk}_j\}_{j \in [N]}, \mu) \end{array} : \text{Dec}^{\{\text{pk}_j\}_{j \in [N]}, \text{ct}}(1^\lambda, N, i, \text{sk}_i) = \mu \right] = 1.$$

Definition 12 (traceability). An AH-BTR scheme $(\text{Gen}, \text{Enc}, \text{Dec}, \text{Trace})$ per Definition 11 is *traceable* if all efficient adversaries win $\text{Exp}_{\text{trace}}$ only with negligible probability, where $\text{Exp}_{\text{trace}}(1^\lambda)$ with adversary \mathcal{B} proceeds as follows.

- **Setup.** Launch $\mathcal{B}(1^\lambda)$. Initialize the set S to \emptyset and let $Q \leftarrow 0$.
- **Query.** Repeat the following for arbitrarily many rounds determined by \mathcal{B} . In each round, \mathcal{B} has two options.
 - \mathcal{B} can request that a new user be initialized and obtain the newly generated public key. Upon this request, let $Q \leftarrow Q + 1$, run

$$(\text{pk}_Q, \text{sk}_Q) \stackrel{\$}{\leftarrow} \text{Gen}(1^\lambda),$$

insert Q into S , and send pk_Q to \mathcal{B} .

⁸Considering an index instead of a set of indices does not lose currently provable properties. The issue with existing formalism is that there is no known way to define the “true” set of traitors (e.g., a user whose secret key is revealed to then immediately discarded by the adversary is not a “true” traitor, which should not and cannot be identified by Trace), hence the security definition cannot require Trace to catch all “true” traitors. Consequently, we can only require it to and prove that it does identify at least one traitor. Our constructions can be modified to potentially find multiple traitors in the usual way [BSW06].

⁹These public keys could be out of the support of Gen, i.e., malformed.

- \mathcal{B} can query for sk_t by submitting $t \in [Q]$. Upon this query, remove t from S and send sk_t to \mathcal{B} .
- **Challenge.** \mathcal{B} outputs a (probabilistic) circuit \mathcal{D} , a list $\{pk_j^*\}_{j \in [N]}$ of public keys, and an error bound $1^{1/\epsilon^*}$ in unary. Run

$$i^* \stackrel{\$}{\leftarrow} \text{Trace}^{\mathcal{D}}(1^\lambda, \{pk_j^*\}_{j \in [N]}, 1^{1/\epsilon^*}).$$

Let

- FalsePos be the event that $i^* \in [N]$ and $pk_{i^*}^* = pk_s$ for some $s \in S$,
- GoodDist the event that

$$\left| \Pr \left[\begin{array}{l} \mu_0 \stackrel{\$}{\leftarrow} \{0,1\}^\lambda, \quad \mu_1 \stackrel{\$}{\leftarrow} \{0,1\}^\lambda \\ \beta \stackrel{\$}{\leftarrow} \{0,1\} \\ \text{ct} \stackrel{\$}{\leftarrow} \text{Enc}(1^\lambda, \{pk_j^*\}_{j \in [N]}, \mu_\beta) \end{array} : \mathcal{D}(\mu_0, \mu_1, \text{ct}) = \beta \right] - \frac{1}{2} \right| \geq \epsilon^*,$$

- and NotFound the event that $i^* \notin [N]$ (i.e., $i^* = \perp$).

\mathcal{B} wins if and only if $\text{FalsePos} \vee (\text{GoodDist} \wedge \text{NotFound})$.

AH-BTR as defined above is a key-encapsulation mechanism, following [Zha20a]. Using hybrid encryption, such a scheme can be easily adapted for arbitrarily long messages with traceability under adversarially chosen messages. As noted in Remark 3 of [Zha20b], traceability implies KEM security (or IND-CPA when combined with hybrid encryption).

3.1 Simplified Security Notions

The traceability of AH-BTR guarantees that a traitor must be found (if the decoder has high advantage) and innocent users must not be accused (regardless of the advantage of the decoder). Decomposing the two requirements (plus some apparent weakening) makes each of them simpler (in particular, *non-interactive*) in the decentralized setting.¹⁰ The first requirement is called *completeness*, and the second *soundness*.

Definition 13 (completeness). An AH-BTR scheme (Gen, Enc, Dec, Trace) per Definition 11 is *complete* if all efficient adversaries win $\text{Exp}_{\text{complete}}$ only with negligible probability, where $\text{Exp}_{\text{complete}}(1^\lambda)$ with adversary \mathcal{C} proceeds as follows.

- **Challenge.** Launch $\mathcal{C}(1^\lambda)$, which outputs a (probabilistic) circuit \mathcal{D} , a list $\{pk_j^*\}_{j \in [N]}$ of public keys, and an error bound $1^{1/\epsilon^*}$ in unary. Run

$$i^* \stackrel{\$}{\leftarrow} \text{Trace}^{\mathcal{D}}(1^\lambda, \{pk_j^*\}_{j \in [N]}, 1^{1/\epsilon^*}).$$

Let

- GoodDist be the event that

$$\left| \Pr \left[\begin{array}{l} \mu_0 \stackrel{\$}{\leftarrow} \{0,1\}^\lambda, \quad \mu_1 \stackrel{\$}{\leftarrow} \{0,1\}^\lambda \\ \beta \stackrel{\$}{\leftarrow} \{0,1\} \\ \text{ct} \stackrel{\$}{\leftarrow} \text{Enc}(1^\lambda, \{pk_j^*\}_{j \in [N]}, \mu_\beta) \end{array} : \mathcal{D}(\mu_0, \mu_1, \text{ct}) = \beta \right] - \frac{1}{2} \right| \geq \epsilon^*,$$

¹⁰Similar simplification to non-interactive security experiments also works, *mutatis mutandis*, for the usual definitions considering a set of traitors identified by Trace.

- and NotFound the event that $i^* \notin [N]$ (i.e., $i^* = \perp$).

\mathcal{C} wins if and only if $\text{GoodDist} \wedge \text{NotFound}$.

Definition 14 (soundness). An AH-BTR scheme $(\text{Gen}, \text{Enc}, \text{Dec}, \text{Trace})$ per Definition 11 is *sound* if all efficient adversaries win $\text{Exp}_{\text{sound}}$ only with negligible probability, where $\text{Exp}_{\text{sound}}(1^\lambda)$ with adversary \mathcal{C} proceeds as follows.

- **Challenge.** Run $(\text{pk}, \text{sk}) \xleftarrow{\$} \text{Gen}(1^\lambda)$, then run $\mathcal{C}(1^\lambda, \text{pk})$, which outputs a (probabilistic) circuit \mathcal{D} , some $N \in \mathbb{N}$, a challenge index $i_\perp^* \in [N]$, a list $\{\text{pk}_j^*\}_{j \in [N] \setminus \{i_\perp^*\}}$ of public keys, and an error bound $1^{1/\epsilon^*}$ in unary. Let $\text{pk}_{i_\perp^*}^* \leftarrow \text{pk}$ and run

$$i^* \xleftarrow{\$} \text{Trace}^{\mathcal{D}}(1^\lambda, \{\text{pk}_j^*\}_{j \in [N]}, 1^{1/\epsilon^*}).$$

\mathcal{C} wins if and only if $i^* = i_\perp^*$ (the event FalsePos).

Theorem 2 (♣). An AH-BTR scheme is traceable if and only if it is both complete and sound.

Proof (Theorem 2). The reductionist proof of necessity is straight-forward – the query phase is unused by the reduction algorithm for completeness, and used only for creating the public key given to the adversary as input for soundness.

To show sufficiency, suppose the AH-BTR scheme $(\text{Gen}, \text{Enc}, \text{Dec}, \text{Trace})$ is both complete and sound and let \mathcal{B} be an efficient adversary against its traceability. We consider two efficient adversaries. \mathcal{C}_1 is against the completeness of the scheme. It works by internally simulating the traceability game for \mathcal{B} and outputting whatever \mathcal{B} outputs. Consider the coupling between $\text{Exp}_{\text{complete}}$ for \mathcal{C}_1 and the simulated $\text{Exp}_{\text{trace}}$ for \mathcal{B} inside, writing the events for adversary \mathcal{X} in its security experiment with subscript \mathcal{X} ,

$$\text{GoodDist}_{\mathcal{C}_1} \iff \text{GoodDist}_{\mathcal{B}} \quad \text{and} \quad \text{NotFound}_{\mathcal{C}_1} \iff \text{NotFound}_{\mathcal{B}}.$$

Therefore,

$$\Pr[\text{GoodDist}_{\mathcal{B}} \wedge \text{NotFound}_{\mathcal{B}}] = \Pr[\text{GoodDist}_{\mathcal{C}_1} \wedge \text{NotFound}_{\mathcal{C}_1}].$$

\mathcal{C}_2 is against the soundness of the scheme. Let $B = \text{poly}(\lambda) > 1$ be an upper bound of the running time of \mathcal{B} . The adversary \mathcal{C}_2 does the following.

- $\mathcal{C}_2(\text{pk})$ launches \mathcal{B} , initializes S to \emptyset , lets $Q \leftarrow 0$, and samples and sets

$$s^* \xleftarrow{\$} [B], \quad \text{pk}_{s^*} \leftarrow \text{pk}, \quad (\text{pk}_t, \text{sk}_t) \xleftarrow{\$} \text{Gen}() \quad \text{for } t \in [B] \setminus \{s^*\}.$$

- \mathcal{C}_2 answers queries from \mathcal{B} and updates Q, S as stipulated by the query phase of the traceability experiment, except that it aborts if \mathcal{B} queries for sk_{s^*} .
- After the query phase, \mathcal{B} outputs

$$\mathcal{D}, \quad \{\text{pk}_j^*\}_{j \in [N]}, \quad 1^{1/\epsilon^*},$$

and \mathcal{C}_2 samples or sets

$$i_\perp^* \begin{cases} \xleftarrow{\$} I_\perp^*, & \text{if } I_\perp^* \leftarrow \{i \in [N] : \text{pk}_i^* = \text{pk}\} \neq \emptyset; \\ \leftarrow \perp & \text{otherwise.} \end{cases}$$

It aborts if $i_\perp^* = \perp$. Otherwise, \mathcal{C}_2 outputs

$$\mathcal{D}, \quad N, \quad i_\perp^*, \quad \{\text{pk}_j^*\}_{j \in [N] \setminus \{i_\perp^*\}}, \quad 1^{1/\epsilon^*}.$$

Consider the coupling between $\text{Exp}_{\text{sound}}$ for \mathcal{C}_2 and the simulated $\text{Exp}_{\text{trace}}$ for \mathcal{B} inside. Routine calculation yields

$$\Pr[\text{FalsePos}_{\mathcal{C}_2}] \geq \frac{1}{B^2} \Pr[\text{FalsePos}_{\mathcal{B}}].$$

By the union bound,

$$\begin{aligned} & \Pr[\text{FalsePos}_{\mathcal{B}} \vee (\text{GoodDist}_{\mathcal{B}} \wedge \text{NotFound}_{\mathcal{B}})] \\ & \leq \Pr[\text{FalsePos}_{\mathcal{B}}] + \Pr[\text{GoodDist}_{\mathcal{B}} \wedge \text{NotFound}_{\mathcal{B}}] \\ & \leq B^2 \Pr[\text{FalsePos}_{\mathcal{C}_2}] + \Pr[\text{GoodDist}_{\mathcal{C}_1} \wedge \text{NotFound}_{\mathcal{C}_1}] \\ & = (\text{poly}(\lambda))^2 \text{negl}(\lambda) + \text{negl}(\lambda) = \text{negl}(\lambda). \quad \square \end{aligned}$$

4 Ad Hoc Private Linear Broadcast Encryption

Our construction of AH-BTR follows that of traitor tracing schemes in [BSW06]. We define *ad hoc* private broadcast linear encryption (AH-PLBE) by adapting the notion of PLBE [BSW06] to the *ad hoc* setting.

Definition 15 (AH-PLBE). An *ad hoc private linear broadcast encryption (AH-PLBE) scheme* (with message space $\{0, 1\}^\lambda$ and public key length $M_0(\lambda)$) consists of 3 efficient algorithms.

- $\text{Gen}(1^\lambda)$ outputs a pair (pk, sk) of public and secret keys with $|\text{pk}| = M_0(\lambda)$.
- $\text{Enc}(1^\lambda, \{\text{pk}_j\}_{j \in [N]}, i_\perp \in [0..N], \mu \in \{0, 1\}^\lambda)$ takes as input a list of public keys, a cut-off index, and a message. It outputs a ciphertext ct .
- $\text{Dec}^{\{\text{pk}_j\}_{j \in [N]}, \text{ct}}(1^\lambda, N, i \in [N], \text{sk}_i)$ is given random access to a list of public keys and a ciphertext, and takes as input the length of the list, an index, and a secret key. It outputs a message.

The scheme must be *robustly correct*, i.e., for all $\lambda \in \mathbb{N}$, $N \in \mathbb{N}$, $i \in [N]$, $\{\text{pk}_j\}_{j \in [N] \setminus \{i\}}$ ¹¹ such that $|\text{pk}_j| = M_0(\lambda)$ for all $j \in [N] \setminus \{i\}$, and $\mu \in \{0, 1\}^\lambda$,

$$\Pr \left[\begin{array}{l} (\text{pk}_i, \text{sk}_i) \stackrel{\$}{\leftarrow} \text{Gen}(1^\lambda) \\ \text{ct} \stackrel{\$}{\leftarrow} \text{Enc}(1^\lambda, \{\text{pk}_j\}_{j \in [N]}, 0, \mu) \end{array} : \text{Dec}^{\{\text{pk}_j\}_{j \in [N]}, \text{ct}}(1^\lambda, N, i, \text{sk}_i) = \mu \right] = 1.$$

Security. We define security notions of AH-PLBE analogously to those in [BSW06], except “mode indistinguishability” (Game 1 in [BSW06]), which is for private tracing thus not needed here (public tracing). The two security definitions have a one-to-one correspondence to the simplified security notions of AH-BTR in Section 3.1. Namely, *message-hiding* translates to *completeness*, and *index-hiding* translates to *soundness*.

Definition 16 (message-hiding). An AH-PLBE scheme $(\text{Gen}, \text{Enc}, \text{Dec})$ per Definition 15 is *message-hiding* if $\text{Exp}_{\text{MH}}^0 \approx \text{Exp}_{\text{MH}}^1$, where $\text{Exp}_{\text{MH}}^b(1^\lambda)$ with adversary \mathcal{A} proceeds as follows.

¹¹These public keys could be out of the support of Gen , i.e., malformed.

- **Challenge.** Launch $\mathcal{A}(1^\lambda)$ and receive from it a list $\{\text{pk}_j^*\}_{j \in [N]}$ of public keys. Run

$$\mu_0 \xleftarrow{\$} \{0, 1\}^\lambda, \quad \mu_1 \xleftarrow{\$} \{0, 1\}^\lambda, \quad \text{ct} \xleftarrow{\$} \text{Enc}(1^\lambda, \{\text{pk}_j^*\}_{j \in [N]}, N, \mu_b),$$

and send $(\mu_0, \mu_1, \text{ct})$ to \mathcal{A} .

- **Guess.** \mathcal{A} outputs a bit b' , which is the output of the experiment.

Definition 17 (index-hiding). An AH-PLBE scheme $(\text{Gen}, \text{Enc}, \text{Dec})$ per Definition 15 is *index-hiding* if $\text{Exp}_{\text{IH}}^0 \approx \text{Exp}_{\text{IH}}^1$, where $\text{Exp}_{\text{IH}}^b(1^\lambda)$ with adversary \mathcal{A} proceeds as follows.

- **Challenge.** Run $(\text{pk}, \text{sk}) \xleftarrow{\$} \text{Gen}(1^\lambda)$, launch $\mathcal{A}(1^\lambda, \text{pk})$, which chooses some $N \in \mathbb{N}$, a cut-off index $i_\perp^* \in [N]$, and a list $\{\text{pk}_j^*\}_{j \in [N] \setminus \{i_\perp^*\}}$ of public keys. Let $\text{pk}_{i_\perp^*}^* \leftarrow \text{pk}$, run

$$\mu \xleftarrow{\$} \{0, 1\}^\lambda, \quad \text{ct} \xleftarrow{\$} \text{Enc}(1^\lambda, \{\text{pk}_j^*\}_{j \in [N]}, i_\perp^* - 1 + b, \mu),$$

and send (μ, ct) to \mathcal{A} .

- **Guess.** \mathcal{A} outputs a bit b' , which is the output of the experiment.

4.1 Construction

Ingredients of Construction 1. Let

- GC = (GC.Garble, GC.Eval, GC.SimGarble) be a circuit garbling scheme such that the algorithm GC.Garble uses λ -bit randomness,
- PPRF = (PPRF.Puncture, PPRF.Eval) a PPRF,
- PKE = (PKE.Gen, PKE.Enc, PKE.Dec) a PKE scheme whose PKE.Enc uses λ -bit randomness and whose public keys are (exactly) of polynomial length M_0 ,
- LOT = (LOT.Gen, LOT.Hash, LOT.Send, LOT.Recv) a laconic OT scheme,
- Obf an obfuscator.

Construction 1 (AH-PLBE). Our AH-PLBE works as follows.

- Gen is the same as PKE.Gen.
- $\text{Enc}(\{\text{pk}_j\}_{j \in [N]}, i_\perp, \mu)$ first checks whether $|\text{pk}_j| = M_0$ for all $j \in [N]$. If not, it outputs $\text{ct} = \perp$ and terminates. Otherwise, the algorithm hashes down the public keys by running

$$\begin{aligned} M &\leftarrow NM_0, & D &\leftarrow \text{pk}_1 \parallel \dots \parallel \text{pk}_N, \\ \text{hk} &\xleftarrow{\$} \text{LOT.Gen}(M), & (h, \widehat{D}) &\leftarrow \text{LOT.Hash}(\text{hk}, D). \end{aligned}$$

It samples a placeholder message $\mu_\perp \xleftarrow{\$} \{0, 1\}^\lambda$ and PPRF keys

$$k^{\text{GC}} \xleftarrow{\$} \{0, 1\}^\lambda, \quad k^{\text{PKE}} \xleftarrow{\$} \{0, 1\}^\lambda, \quad k_{m_0}^{\text{LOT}} \xleftarrow{\$} \{0, 1\}^\lambda \quad \text{for } m_0 \in [M_0],$$

and obfuscates C_{GC} (Figure 1) by running

$$\widetilde{C}_{\text{GC}} \xleftarrow{\$} \text{Obf}(C_{\text{GC}}[N, \text{hk}, h, i_\perp, \mu_\perp, \mu, k^{\text{GC}}, k^{\text{PKE}}, \{k_{m_0}^{\text{LOT}}\}_{m_0 \in [M_0]}]).$$

The algorithm outputs $\text{ct} = (\text{hk}, \widetilde{C}_{\text{GC}})$ as the ciphertext.

$C_{GC}[N, hk, h, i_{\perp}, \mu_{\perp}, \mu, k^{GC}, k^{PKE}, \{k_{m_0}^{LOT}\}_{m_0 \in [M_0]}](i)$	
<p>Hardwired.</p>	<p>N, number of users; hk, laconic OT hash key; h, laconic OT hash of $D = pk_1 \parallel \dots \parallel pk_N$; i_{\perp}, cut-off index; μ_{\perp}, placeholder message; μ, message; k^{GC}, PPRF key for circuit garbling; k^{PKE}, PPRF key for public-key encryption; $k_{m_0}^{LOT}$, PPRF key for sending the m_0^{th} label using laconic OT.</p>
<p>Input.</p>	<p>$i \in [N]$, index of recipient.</p>
<p>Output.</p>	<p>Computed as follows.</p> $r_i^{GC} \leftarrow \text{PPRF.Eval}(k^{GC}, i)$ $r_i^{PKE} \leftarrow \text{PPRF.Eval}(k^{PKE}, i)$ $r_{i,m_0}^{LOT} \leftarrow \text{PPRF.Eval}(k_{m_0}^{LOT}, i) \quad \text{for } m_0 \in [M_0]$ $(\widehat{C}_{ct,i}, \{L_{i,m_0,b}\}_{m_0 \in [M_0], b \in \{0,1\}})$ $\leftarrow \begin{cases} \text{GC.Garble}(\widehat{C}_{ct}, (\mu_{\perp}, r_i^{PKE}); r_i^{GC}), & \text{if } i \leq i_{\perp}; \\ \text{GC.Garble}(\widehat{C}_{ct}, (\mu, r_i^{PKE}); r_i^{GC}), & \text{if } i > i_{\perp}; \end{cases}$ $\text{LOT.ct}_{i,m_0} \leftarrow \text{LOT.Send}(hk, h, (i-1)M_0 + m_0, L_{i,m_0,0}, L_{i,m_0,1}; r_{i,m_0}^{LOT}) \quad \text{for } m_0 \in [M_0]$ <p>output $(\widehat{C}_{ct,i}, \{\text{LOT.ct}_{i,m_0}\}_{m_0 \in [M_0]})$</p>
$C_{ct}[\mu'_i, r_i^{PKE}](pk_i)$	
<p>Hardwired.</p>	<p>μ'_i, message or placeholder message; r_i^{PKE}, public-key encryption randomness.</p>
<p>Input.</p>	<p>pk_i, public key of recipient.</p>
<p>Output.</p>	<p>$\text{PKE.ct}_i \leftarrow \text{PKE.Enc}(pk_i, \mu'_i; r_i^{PKE})$.</p>

Figure 1. The circuits C_{GC} and C_{ct} in Construction 1.

- $\text{Dec}^{\{\text{pk}_j\}_{j \in [N]}, \text{ct}}(N, i, \text{sk}_i)$ first parses $\text{ct} = (\text{hk}, \widetilde{C}_{\text{GC}})$ and recomputes

$$M \leftarrow NM_0, \quad D \leftarrow \text{pk}_1 \parallel \dots \parallel \text{pk}_N, \quad (h, \widehat{D}) \leftarrow \text{LOT.Hash}(\text{hk}, D).$$

The algorithm next runs the obfuscated circuit,

$$(\widehat{C}_{\text{ct}, i}, \{\text{LOT.ct}_{i, m_0}\}_{m_0 \in [M_0]}) \leftarrow \widetilde{C}_{\text{GC}}(i),$$

to obtain the garbled C_{ct} (Figure 1) for the decryptor and the laconic OT ciphertexts sending its labels. It then receives the labels,

$$L_{i, m_0, \text{pk}_i[m_0]} \leftarrow \text{LOT.Recv}^{\widehat{D}}(\text{hk}, h, (i-1)M_0 + m_0, \text{LOT.ct}_{i, m_0}) \quad \text{for } m_0 \in [M_0],$$

and evaluates the garbled circuit,

$$\text{PKE.ct}_i \leftarrow \text{GC.Eval}(\widehat{C}_{\text{ct}, i}, \text{pk}_i, \{L_{i, m_0, \text{pk}_i[m_0]}\}_{m_0 \in [M_0]}),$$

to obtain the PKE ciphertext under the decryptor's public key. Lastly, the algorithm runs and outputs (as the decrypted message)

$$\mu \leftarrow \text{PKE.Dec}(\text{sk}_i, \text{PKE.ct}_i).$$

Robust Correctness. It follows from the correctness of the ingredients.

Efficiency. By laconic OT efficiency, the call to LOT.Gen takes time $\text{poly}(\lambda, \log(N+1))$, that to LOT.Hash takes time $(N+1) \text{poly}(\lambda, \log(N+1))$, and $|\text{hk}|, |h| = \text{poly}(\lambda, \log(N+1))$. As we shall see later, it suffices to pad C_{GC} to size $\text{poly}(\lambda, \log(N+1))$ for the security proofs to go through. Putting these together,

$$T_{\text{Enc}}, T_{\text{Dec}} = (N+1) \text{poly}(\lambda, \log(N+1)), \quad |\text{ct}| = \text{poly}(\lambda, \log(N+1)).$$

In practice and for security reasons, we always assume $N \leq 2^\lambda$ and $\log(N+1)$ is absorbed by λ .¹² Therefore, with $\text{poly}(\lambda)$ factors ignored, both encryption and decryption take linear time, and the ciphertext is constant-size.

Compatibility. Since the key generation algorithm of Construction 1 is just the key generation algorithm of the underlying PKE scheme (which only has to be semantically secure for random messages), it is compatible with the existing public-key encryption schemes, i.e., existing users possessing PKE key pairs can utilize our AH-PLBE without regenerating their keys.

4.2 Security

Theorem 3 (¶). *Suppose in Construction 1, the obfuscator Obf is an $i\mathcal{O}$ for $\text{poly}(\lambda)$ -sized domain, then the resultant AH-PLBE is message-hiding.*

Theorem 4 (¶). *Suppose in Construction 1, all of the ingredients are secure, then the resultant AH-PLBE is index-hiding.*

¹²A scheme can always set the ciphertext to the message itself whenever $N > 2^\lambda$ and remain correct and asymptotically secure. See also Footnote 6.

Proof (Theorem 3). For Construction 1, the only difference between Exp_{MH}^0 and Exp_{MH}^1 is whether C_{GC} used to create $\text{ct} = (\text{hk}, \widetilde{C}_{\text{GC}})$ has μ_0 or μ_1 hardwired as μ . In C_{GC} (Figure 1), μ is used only in the branch $i > i_{\perp}$, which is never taken in Exp_{MH}^0 or Exp_{MH}^1 because i_{\perp} is hardwired to be N and the domain of i is $[N]$. Therefore, the two C_{GC} 's in Exp_{MH}^0 and Exp_{MH}^1 being obfuscated are functionally equivalent and have the same size. Moreover, their domain size is N (polynomially large). Therefore, $\text{Exp}_{\text{MH}}^0 \approx \text{Exp}_{\text{MH}}^1$ reduces to the $i_{\mathcal{O}}$ security for $\text{poly}(\lambda)$ -sized domain of Obf . \square

Proof (Theorem 4). The only difference between Exp_{IH}^0 and Exp_{IH}^1 is whether the C_{GC} being obfuscated hardwires μ (in Exp_{IH}^0) or μ_{\perp} (in Exp_{IH}^1) into $C_{\text{ct}, i_{\perp}^*}$, which only affects the output of C_{GC} at $i = i_{\perp}^*$. We consider the following hybrids, each (except the first) described by the changes from the previous one.

- H_0^b (for $b \in \{0, 1\}$) is Exp_{IH}^b , where

$$\begin{aligned} \text{hk} &\stackrel{\$}{\leftarrow} \text{LOT.Gen}(NM_0), & (h, \widehat{D}) &\stackrel{\$}{\leftarrow} \text{LOT.Hash}(\text{hk}, \text{pk}_1^* \parallel \dots \parallel \text{pk}_N^*), \\ k^{\text{GC}} &\stackrel{\$}{\leftarrow} \{0, 1\}^{\lambda}, & k^{\text{PKE}} &\stackrel{\$}{\leftarrow} \{0, 1\}^{\lambda}, & k_{m_0}^{\text{LOT}} &\stackrel{\$}{\leftarrow} \{0, 1\}^{\lambda} \text{ for } m_0 \in [M_0], \\ \widetilde{C}_{\text{GC}} &\stackrel{\$}{\leftarrow} \text{Obf}(C_{\text{GC}}[N, \text{hk}, h, i_{\perp}^* - 1 + b, \mu_{\perp}, \mu, k^{\text{GC}}, k^{\text{PKE}}, \{k_{m_0}^{\text{LOT}}\}_{m_0 \in [M_0]}]), \\ \text{ct} &= (\text{hk}, \widetilde{C}_{\text{GC}}). \end{aligned}$$

- H_1^b alters the obfuscation into

$$\begin{aligned} \widetilde{C}_{\text{GC}} &\stackrel{\$}{\leftarrow} \text{Obf}(C'_{\text{GC}}[N, \text{hk}, h, \mu_{\perp}, \mu, \\ & i_{\perp}^*, \overset{\circ}{k}_{i_{\perp}^*}^{\text{GC}}, \overset{\circ}{k}_{i_{\perp}^*}^{\text{PKE}}, \{\overset{\circ}{k}_{m_0, i_{\perp}^*}^{\text{LOT}}\}_{m_0 \in [M_0]}, \widehat{C}_{\text{ct}, i_{\perp}^*}, \{\text{LOT.ct}_{i_{\perp}^*, m_0}\}_{m_0 \in [M_0]}]), \end{aligned}$$

where

- C'_{GC} is defined in Figure 2,
- the PPRF keys are punctured at i_{\perp}^* by running

$$\begin{aligned} \overset{\circ}{k}_{i_{\perp}^*}^{\text{GC}} &\stackrel{\$}{\leftarrow} \text{PPRF.Puncture}(k^{\text{GC}}, i_{\perp}^*), \\ \overset{\circ}{k}_{i_{\perp}^*}^{\text{PKE}} &\stackrel{\$}{\leftarrow} \text{PPRF.Puncture}(k^{\text{PKE}}, i_{\perp}^*), \\ \overset{\circ}{k}_{m_0, i_{\perp}^*}^{\text{LOT}} &\stackrel{\$}{\leftarrow} \text{PPRF.Puncture}(k_{m_0}^{\text{LOT}}, i_{\perp}^*) \quad \text{for } m_0 \in [M_0], \end{aligned}$$

- and the output $(\widehat{C}_{\text{ct}, i_{\perp}^*}, \{\text{LOT.ct}_{i_{\perp}^*, m_0}\}_{m_0 \in [M_0]})$ of C'_{GC} at $i = i_{\perp}^*$ is computed as

$$\begin{aligned} r^{\text{GC}} &\leftarrow \text{PPRF.Eval}(k^{\text{GC}}, i_{\perp}^*), & r^{\text{PKE}} &\leftarrow \text{PPRF.Eval}(k^{\text{PKE}}, i_{\perp}^*), \\ r_{i_{\perp}^*, m_0}^{\text{LOT}} &\leftarrow \text{PPRF.Eval}(k_{m_0}^{\text{LOT}}, i_{\perp}^*) \quad \text{for } m_0 \in [M_0], \\ (\widehat{C}_{\text{ct}, i_{\perp}^*}, \{L_{i_{\perp}^*, m_0, b}\}_{m_0 \in [M_0], b \in \{0, 1\}}) & \\ &\leftarrow \begin{cases} \text{GC.Garble}(C_{\text{ct}}, (\mu, r_{i_{\perp}^*}^{\text{PKE}}); r_{i_{\perp}^*}^{\text{GC}}), & \text{if } b = 0; \\ \text{GC.Garble}(C_{\text{ct}}, (\mu_{\perp}, r_{i_{\perp}^*}^{\text{PKE}}); r_{i_{\perp}^*}^{\text{GC}}), & \text{if } b = 1; \end{cases} \\ \text{LOT.ct}_{i_{\perp}^*, m_0} &\leftarrow \text{LOT.Send}(\text{hk}, h, (i_{\perp}^* - 1)M_0 + m_0, \\ & L_{i_{\perp}^*, m_0, 0}, L_{i_{\perp}^*, m_0, 1}; r_{i_{\perp}^*, m_0}^{\text{LOT}}) \quad \text{for } m_0 \in [M_0]. \end{aligned}$$

- H_2^b changes $r_{i_\perp^*}^{\text{GC}}$, $r_{i_\perp^*}^{\text{PKE}}$, and $r_{i_\perp^*, m_0}^{\text{LOT}}$'s into true randomness, i.e.,

$$r^{\text{GC}} \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda, \quad r^{\text{PKE}} \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda, \quad r_{i_\perp^*, m_0}^{\text{LOT}} \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda \quad \text{for } m_0 \in [M_0].$$

- H_3^b removes the unused labels from $\text{LOT.ct}_{i_\perp^*, m_0}$'s by setting

$$\begin{aligned} \text{LOT.ct}_{i_\perp^*, m_0} &\leftarrow \text{LOT.Send}(\text{hk}, h, (i_\perp^* - 1)M_0 + m_0, \\ &\quad L_{i_\perp^*, m_0, \text{pk}_{i_\perp^*}^* [m_0]}, L_{i_\perp^*, m_0, \text{pk}_{i_\perp^*}^* [m_0]}; r_{i_\perp^*, m_0}^{\text{LOT}}) \quad \text{for } m_0 \in [M_0]. \end{aligned}$$

- H_4^b changes $\widehat{C}_{\text{ct}, i_\perp^*}$ into simulation, i.e.,

$$\begin{aligned} \text{PKE.ct}_{i_\perp^*} &\leftarrow \begin{cases} \text{PKE.Enc}(\text{pk}_{i_\perp^*}^*, \mu; r^{\text{PKE}}), & \text{if } b = 0; \\ \text{PKE.Enc}(\text{pk}_{i_\perp^*}^*, \mu_\perp; r^{\text{PKE}}), & \text{if } b = 1; \end{cases} \\ (\widehat{C}_{\text{ct}, i_\perp^*}, \{L_{i_\perp^*, m_0, \text{pk}_{i_\perp^*}^* [m_0]}\}_{m_0 \in [M_0]}) &\stackrel{\$}{\leftarrow} \text{GC.SimGarble}(C_{\text{ct}}, \text{pk}_{i_\perp^*}^*, \text{PKE.ct}_{i_\perp^*}), \end{aligned}$$

where $\text{pk}_{i_\perp^*}^* = \text{pk}$ is sampled by the experiment (not adversarially controlled).

$C'_{\text{GC}}[N, \text{hk}, h, \mu_\perp, \mu, i_\perp^*, \hat{k}_{i_\perp^*}^{\text{GC}}, \hat{k}_{i_\perp^*}^{\text{PKE}}, \{\hat{k}_{m_0, i_\perp^*}^{\text{LOT}}\}_{m_0 \in [M_0]}, \widehat{C}_{\text{ct}, i_\perp^*}, \{\text{LOT.ct}_{i_\perp^*, m_0}\}_{m_0 \in [M_0]}](i)$

Hardwired. $N, \text{hk}, h, \mu_\perp, \mu,$ see Figure 1;
 i_\perp^* , challenge cut-off index;
 $\hat{k}_{\dots, i_\perp^*}^{\dots}$, PPRF keys punctured at i_\perp^* ;
 $\widehat{C}_{\text{ct}, i_\perp^*}, \text{LOT.ct}_{i_\perp^*, \dots},$ hardwired output of C'_{GC} at $i = i_\perp^*$.

Input. $i \in [N],$ index of recipient.

Output. Computed as follows.

if $i = i_\perp^*$:

output $(\widehat{C}_{\text{ct}, i_\perp^*}, \{\text{LOT.ct}_{i_\perp^*, m_0}\}_{m_0 \in [M_0]})$ as hardwired

else:

$r_i^{\text{GC}} \leftarrow \text{PPRF.Eval}(\hat{k}_{i_\perp^*}^{\text{GC}}, i)$
 $r_i^{\text{PKE}} \leftarrow \text{PPRF.Eval}(\hat{k}_{i_\perp^*}^{\text{PKE}}, i)$
 $r_{i, m_0}^{\text{LOT}} \leftarrow \text{PPRF.Eval}(\hat{k}_{m_0, i_\perp^*}^{\text{LOT}}, i) \quad \text{for } m_0 \in [M_0]$

$(\widehat{C}_{\text{ct}, i}, \{L_{i, m_0, b}\}_{m_0 \in [M_0], b \in \{0, 1\}})$
 $\leftarrow \begin{cases} \text{GC.Garble}(\widehat{C}_{\text{ct}}, (\mu_\perp, r_i^{\text{PKE}}); r_i^{\text{GC}}), & \text{if } i < i_\perp^*; \\ \text{GC.Garble}(\widehat{C}_{\text{ct}}, (\mu, r_i^{\text{PKE}}); r_i^{\text{GC}}), & \text{if } i > i_\perp^*; \end{cases}$

$\text{LOT.ct}_{i, m_0} \leftarrow \text{LOT.Send}(\text{hk}, h, (i - 1)M_0 + m_0,$
 $\quad L_{i, m_0, 0}, L_{i, m_0, 1}; r_{i, m_0}^{\text{LOT}}) \quad \text{for } m_0 \in [M_0]$

output $(\widehat{C}_{\text{ct}, i}, \{\text{LOT.ct}_{i, m_0}\}_{m_0 \in [M_0]})$

Figure 2. The circuit C'_{GC} in the proof of Theorem 4.

The following claims hold, all of which are immediate by inspection.

Claim 5. $H_0^b \approx H_1^b$ for $b \in \{0, 1\}$ if Obf is an iO for $\text{poly}(\lambda)$ -sized domain.

Claim 6. $H_1^b \approx H_2^b$ for $b \in \{0, 1\}$ if PPRF is pseudorandom at the punctured point.

Claim 7. $H_2^b \approx H_3^b$ for $b \in \{0, 1\}$ if LOT is database-selectively sender-private.

Claim 8. $H_3^b \approx H_4^b$ for $b \in \{0, 1\}$ if GC is w -hiding.

Claim 9. $H_4^0 \approx H_4^1$ if PKE is semantically secure for random messages.

$\text{Exp}_{\text{IH}}^0 \approx \text{Exp}_{\text{IH}}^1$ follows from a hybrid argument. □

5 AH-BTR from AH-PLBE

Ingredient of Construction 2. Let $\text{ahPLBE} = (\text{ahPLBE.Gen}, \text{ahPLBE.Enc}, \text{ahPLBE.Dec})$ be an AH-PLBE scheme.

Construction 2 (adapted from [BSW06; Section 2.2]). Our AH-BTR works as follows.

- Gen is the same as ahPLBE.Gen .
- $\text{Enc}(\{\text{pk}_j\}_{j \in [N]}, \mu)$ runs and outputs $\text{ct} \stackrel{\$}{\leftarrow} \text{ahPLBE.Enc}(\{\text{pk}_j\}_{j \in [N]}, 0, \mu)$.
- Dec is the same as ahPLBE.Dec .
- $\text{Trace}^{\mathcal{D}}(\{\text{pk}_j^*\}_{j \in [N]}, 1^{1/\epsilon^*})$ defines for $i \in [0..N]$,

$$\epsilon_i = \Pr \left[\underbrace{\begin{array}{l} \mu_0 \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda, \quad \mu_1 \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda, \quad \beta \stackrel{\$}{\leftarrow} \{0, 1\} \\ \text{ct} \stackrel{\$}{\leftarrow} \text{ahPLBE.Enc}(1^\lambda, \{\text{pk}_j^*\}_{j \in [N]}, i, \mu_\beta) \end{array}} : \mathcal{D}(\mu_0, \mu_1, \text{ct}) = \beta \right] - \frac{1}{2}.$$

experiment \mathcal{E}_i (sampling and testing) and event E_i (correct guessing)

Setting $\delta \leftarrow \frac{\epsilon^*}{10N}$ and $\eta \leftarrow \left\lceil \frac{\lambda + \log(2N+2)}{2\delta^2} \right\rceil$, for each $i \in [0..N]$, the algorithm runs \mathcal{E}_i for η times independently, counts the absolute frequency $\xi_i \in [0..\eta]$ of E_i , and computes $\widehat{\epsilon}_i = \frac{\xi_i}{\eta} - \frac{1}{2}$. It outputs

$$i^* = \begin{cases} \min T, & \text{if } T \leftarrow \{i \in [N] : |\widehat{\epsilon}_i - \widehat{\epsilon}_{i-1}| \geq 3\delta\} \neq \emptyset; \\ \perp, & \text{if } T = \emptyset. \end{cases}$$

Robust Correctness, Efficiency, Compatibility. These are inherited from the underlying AH-PLBE. When based on Construction 1, the resultant AH-BTR has

$$T_{\text{Enc}} = (N+1) \text{poly}(\lambda), \quad |\text{ct}| = \text{poly}(\lambda), \quad T_{\text{Dec}} = (N+1) \text{poly}(\lambda),$$

and is compatible with the existing public-key encryption schemes.

Theorem 10 (¶). *Suppose in Construction 2, the AH-PLBE scheme ahPLBE is message-hiding, then the resultant AH-BTR is complete.*

Theorem 11 (¶). *Suppose in Construction 2, the AH-PLBE scheme ahPLBE is index-hiding, then the resultant AH-BTR is sound.*

Proof (Theorem 10). Consider any efficient adversary \mathcal{C} against the completeness of Construction 2. Let GoodEst be the event that $|\widehat{\varepsilon}_i - \varepsilon_i| \leq \delta$ for all $i \in [0..N]$. By the Chernoff bound, the union bound, and the law of total probability,

$$\Pr[\neg\text{GoodEst}] = \mathbb{E}[\Pr[\neg\text{GoodEst} \mid \varepsilon^*, N]] \leq \mathbb{E}[2(N+1) \exp(-2\delta^2\eta)] \leq 2^{-\lambda}.$$

Let BadEnd be the event that $|\varepsilon_N| > \frac{\varepsilon^*}{2}$, then $\text{GoodDist} \wedge \neg\text{BadEnd}$ implies

$$\begin{aligned} \max_{i \in [N]} |\varepsilon_{i-1} - \varepsilon_i| &\geq \frac{1}{N} \sum_{i=1}^N |\varepsilon_{i-1} - \varepsilon_i| \geq \frac{1}{N} \left| \sum_{i=1}^N (\varepsilon_{i-1} - \varepsilon_i) \right| = \frac{1}{N} |\varepsilon_0 - \varepsilon_N| \\ &\geq \frac{1}{N} (|\varepsilon_0| - |\varepsilon_N|) \geq \frac{1}{N} \left(\varepsilon^* - \frac{\varepsilon^*}{2} \right) = \frac{\varepsilon^*}{2N} = 5\delta. \end{aligned}$$

\uparrow
GoodDist
 \uparrow
 \neg BadEnd

Therefore, $\text{GoodDist} \wedge \neg\text{BadEnd} \wedge \text{GoodEst}$ implies

$$\max_{i \in [N]} |\widehat{\varepsilon}_{i-1} - \widehat{\varepsilon}_i| \geq \max_{i \in [N]} (|\varepsilon_{i-1} - \varepsilon_i| - 2\delta) \geq 5\delta - 2\delta = 3\delta,$$

\uparrow
GoodEst
 \uparrow
GoodDist \wedge \neg BadEnd

which in turn implies $T \neq \emptyset$ hence $i^* \in [N]$, i.e., $\neg\text{NotFound}$. By contraposition,

$$\text{GoodDist} \wedge \text{NotFound} \wedge \text{GoodEst} \implies \text{BadEnd}.$$

By the union bound,

$$\begin{aligned} \Pr[\mathcal{C} \text{ wins}] &\leq \Pr[\neg\text{GoodEst}] + \Pr[(\mathcal{C} \text{ wins}) \wedge \text{GoodEst}] \\ &= \Pr[\neg\text{GoodEst}] + \Pr[\text{GoodDist} \wedge \text{NotFound} \wedge \text{GoodEst}] \\ &\leq 2^{-\lambda} + \Pr[\text{BadEnd}], \end{aligned}$$

so it remains to show $\Pr[\text{BadEnd}] = \text{negl}(\lambda)$.

Consider the following efficient adversary \mathcal{A} against the message-hiding property of ahPLBE.

- \mathcal{A} runs \mathcal{C} to obtain

$$\mathcal{D}, \quad \{\text{pk}_j^*\}_{j \in [N]}, \quad 1^{1/\varepsilon^*}.$$

- \mathcal{A} runs \mathcal{E}_N once and notes down $\alpha \in \{0, 1\}$ indicating whether E_N happened, i.e., $\alpha = 1$ if and only if \mathcal{D} guessed correctly in the trial.
- \mathcal{A} submits $\{\text{pk}_j^*\}_{j \in [N]}$ to the message-hiding experiment, receives $(\mu_0, \mu_1, \text{ct})$ back, and runs and outputs $b' \stackrel{\$}{\leftarrow} \mathcal{D}(\mu_0, \mu_1, \text{ct}) \oplus \alpha$.

Routine calculation shows that the advantage of \mathcal{A} is $\mathbb{E}[4\varepsilon_N^2]$, which must be negligible by the message-hiding property of ahPLBE. Let $B = \text{poly}(\lambda)$ be an upper bound of $1/\varepsilon^*$ (B exists since \mathcal{C} outputs $1^{1/\varepsilon^*}$ in polynomial time). By Markov's inequality,

$$\begin{aligned} \Pr[\text{BadEnd}] &= \Pr[4\varepsilon_N^2 > (\varepsilon^*)^2] \leq \Pr[4\varepsilon_N^2 > B^{-2}] \\ &\leq B^2 \mathbb{E}[4\varepsilon_N^2] = (\text{poly}(\lambda))^2 \text{negl}(\lambda) = \text{negl}(\lambda). \quad \square \end{aligned}$$

Proof (Theorem 11). Consider any efficient adversary \mathcal{C} against the soundness of Construction 2. Similarly to the **proof** of Theorem 10, define the event GoodEst and recall that $\Pr[\neg\text{GoodEst}] \leq 2^{-\lambda}$. We have

$$\begin{aligned} \Pr[\mathcal{C} \text{ wins}] &\leq \Pr[\neg\text{GoodEst}] + \Pr[(\mathcal{C} \text{ wins}) \wedge \text{GoodEst}] \\ &= \Pr[\neg\text{GoodEst}] + \Pr[\text{FalsePos} \wedge \text{GoodEst}] \\ &\leq 2^{-\lambda} + \Pr[\text{FalsePos} \wedge \text{GoodEst}], \end{aligned}$$

and it suffices to prove $\Pr[\text{FalsePos} \wedge \text{GoodEst}] = \text{negl}(\lambda)$.

Let α be a random element in an execution of Trace with

$$\alpha = \begin{cases} 0, & \text{if } i^* \in [N] \text{ and } \widehat{\varepsilon}_{i^*-1} - \widehat{\varepsilon}_{i^*} \geq 3\delta; \\ 1, & \text{if } i^* \in [N] \text{ and } \widehat{\varepsilon}_{i^*-1} - \widehat{\varepsilon}_{i^*} \leq -3\delta; \\ \perp, & \text{if } i^* = \perp. \end{cases}$$

Consider the following efficient adversary \mathcal{A} against the index-hiding property of ahPLBE .

- $\mathcal{A}(\text{pk})$ runs $\mathcal{C}(\text{pk})$ to obtain

$$\mathcal{D}, \quad N, \quad i_{\perp}^*, \quad \{\text{pk}_j^*\}_{j \in [N] \setminus \{i_{\perp}^*\}}, \quad 1^{1/\varepsilon^*},$$

and sets $\text{pk}_{i_{\perp}^*}^* \leftarrow \text{pk}$.

- \mathcal{A} runs

$$i^* \xleftarrow{\$} \text{Trace}^{\mathcal{D}}(\{\text{pk}_j^*\}_{j \in [N]}, 1^{1/\varepsilon^*}),$$

and aborts if $i^* \neq i_{\perp}^*$.

- \mathcal{A} notes down $\alpha \in \{0, 1\}$ from the above execution of Trace , submits

$$N, \quad i_{\perp}^*, \quad \{\text{pk}_j^*\}_{j \in [N] \setminus \{i_{\perp}^*\}}$$

to the index-hiding experiment, gets (μ, ct) back, samples and sets

$$\beta \xleftarrow{\$} \{0, 1\}, \quad \mu_{\beta} \leftarrow \mu, \quad \mu_{\neg\beta} \xleftarrow{\$} \{0, 1\}^{\lambda},$$

and runs and outputs $b' \xleftarrow{\$} \mathcal{D}(\mu_0, \mu_1, \text{ct}) \oplus \neg\beta \oplus \alpha$.

Routine calculation shows that the advantage of \mathcal{A} is

$$\mathbb{E}[\mathbb{1}_{\text{FalsePos}} \cdot (-1)^{\alpha} (\varepsilon_{i^*-1} - \varepsilon_{i^*})],$$

which must be negligible by the index-hiding property of ahPLBE .

Let $B = \text{poly}(\lambda)$ be an upper bound of $10N/\varepsilon^*$ (B exists since \mathcal{C} outputs 1^N and $1^{1/\varepsilon^*}$ in polynomial time). The event $\text{FalsePos} \wedge \text{GoodEst}$ implies

$$\begin{aligned} |(\varepsilon_{i^*-1} - \varepsilon_{i^*}) - (\widehat{\varepsilon}_{i^*-1} - \widehat{\varepsilon}_{i^*})| &\leq 2\delta < 3\delta \leq |\widehat{\varepsilon}_{i^*-1} - \widehat{\varepsilon}_{i^*}| \\ \implies (-1)^{\alpha} (\varepsilon_{i^*-1} - \varepsilon_{i^*}) &= |\varepsilon_{i^*-1} - \varepsilon_{i^*}| \geq 3\delta - 2\delta = \frac{\varepsilon^*}{10N} \geq B^{-1}. \end{aligned}$$

Moreover, $(-1)^{\alpha} (\varepsilon_{i^*-1} - \varepsilon_{i^*}) \geq -1$ always holds. These together show that

$$\begin{aligned} &\Pr[\text{FalsePos} \wedge \text{GoodEst}] \\ &= B \mathbb{E}[\mathbb{1}_{\text{FalsePos}} \cdot \mathbb{1}_{\text{GoodEst}} \cdot B^{-1}] \end{aligned}$$

$$\begin{aligned}
&\leq B \mathbb{E}[\mathbb{1}_{\text{FalsePos}} \cdot \mathbb{1}_{\text{GoodEst}} \cdot (-1)^\alpha (\varepsilon_{i^*-1} - \varepsilon_{i^*})] \\
&\leq B \left(\mathbb{E}[\mathbb{1}_{\text{FalsePos}} \cdot \mathbb{1}_{\text{GoodEst}} \cdot (-1)^\alpha (\varepsilon_{i^*-1} - \varepsilon_{i^*})] \right. \\
&\quad \left. + \mathbb{E}[\mathbb{1}_{\text{FalsePos}} \cdot \mathbb{1}_{\neg\text{GoodEst}} \cdot (-1)^\alpha (\varepsilon_{i^*-1} - \varepsilon_{i^*})] + \mathbb{E}[\mathbb{1}_{\text{FalsePos}} \cdot \mathbb{1}_{\neg\text{GoodEst}}] \right) \\
&= B \left(\mathbb{E}[\mathbb{1}_{\text{FalsePos}} \cdot (-1)^\alpha (\varepsilon_{i^*-1} - \varepsilon_{i^*})] + \Pr[\text{FalsePos} \wedge \neg\text{GoodEst}] \right) \\
&\leq B \left(\mathbb{E}[\mathbb{1}_{\text{FalsePos}} \cdot (-1)^\alpha (\varepsilon_{i^*-1} - \varepsilon_{i^*})] + 2^{-\lambda} \right) \\
&= \text{poly}(\lambda) (\text{negl}(\lambda) + 2^{-\lambda}) = \text{negl}(\lambda). \quad \square
\end{aligned}$$

6 Trading Ciphertext Size for Decryption Time

While Construction 2 achieves constant ciphertext size, it takes time $\Omega(N)$ to decrypt. In contrast, the naïve scheme that encrypts to each user separately has $\Omega(N)$ -size ciphertext, yet decryption only takes constant time. By grouping the recipients and encrypting to each group separately, we can trade ciphertext size for decryption time.¹³ Previous work [Zha20a] already systemizes the idea of grouping in the context of traditional traitor tracing.

Ingredients of Construction 3. Let $\text{old} = (\text{old.Gen}, \text{old.Enc}, \text{old.Dec}, \text{old.Trace})$ be an AH-BTR scheme and γ some¹⁴ constant ($0 < \gamma < 1$).

Construction 3 (adapted from [Zha20a; Theorem 1]). Our new AH-BTR works as follows.

- Gen is the same as old.Gen.
- $\text{Enc}(\{\text{pk}_j\}_{j \in [N]}, \mu)$ sets $N_1 = \lceil N^\gamma \rceil$ and $N_2 = \lceil N/N_1 \rceil$. It runs

$$\text{old.ct}_{j_1} \stackrel{\$}{\leftarrow} \text{old.Enc}(\{\text{pk}_j\}_{(j_1-1)N_2 < j \leq j_1 N_2}, \mu) \quad \text{for } j_1 \in [N_1].$$

The algorithm outputs $\text{ct} = \{\text{old.ct}_{j_1}\}_{j_1 \in [N_1]}$.

- $\text{Dec}^{\{\text{pk}_j\}_{j \in [N]}, \text{ct}}(N, i, \text{sk}_i)$ sets $N_1 = \lceil N^\gamma \rceil$, $N_2 = \lceil N/N_1 \rceil$. It parses ct as $\{\text{old.ct}_{j_1}\}_{j_1 \in [N_1]}$, finds $i_1 \in [N_1]$ such that $(i_1-1)N_2 < i \leq i_1 N_2$, and sets $N'_2 = \min\{N_2, N - (i_1-1)N_2\}$. The algorithm runs and outputs

$$\text{old.Dec}^{\{\text{pk}_j\}_{(i_1-1)N_2 < j \leq i_1 N_2}, \text{old.ct}_{i_1}}(N'_2, i - (i_1-1)N_2, \text{sk}_i).$$

- $\text{Trace}^{\mathcal{D}}(\{\text{pk}_j^*\}_{j \in [N]}, 1^{1/\varepsilon^*})$ sets $N_1 = \lceil N^\gamma \rceil$ and $N_2 = \lceil N/N_1 \rceil$. It runs

$$i_{j_1}^* \stackrel{\$}{\leftarrow} \text{old.Trace}^{\mathcal{D}_{j_1}}(\{\text{pk}_j^*\}_{(j_1-1)N_2 < j \leq j_1 N_2}, 1^{N_1/\varepsilon^*}) \quad \text{for } j_1 \in [N_1],$$

¹³Alternatively, one can reformulate Construction 2 as a compiler that trades decryption time for ciphertext size, by grouping the recipients and compressing the groups. We refrained from such a formulation because the “transformation” uses a quite strong additional assumption, namely functional encryption for general circuits.

¹⁴We require that $N \mapsto \lceil N^\gamma \rceil$ can be computed in (deterministic) time $\text{poly}(\log(N+1))$.

where $\mathcal{D}_{j_1}(\mu_0, \mu_1, \text{old.ct}^*)$ runs and outputs $\mathcal{D}(\mu_0, \mu_1, \{\text{old.ct}_{j'_1}\}_{j'_1 \in [N_1]})$ with

$$\text{old.ct}_{j'_1} \begin{cases} \stackrel{\$}{\leftarrow} \text{old.Enc}(\{\text{pk}_j^*\}_{(j'_1-1)N_2 < j \leq j'_1 N_2}, \mu_0), & \text{if } j'_1 < j_1; \\ \leftarrow \text{old.ct}^*, & \text{if } j'_1 = j_1; \\ \stackrel{\$}{\leftarrow} \text{old.Enc}(\{\text{pk}_j^*\}_{(j'_1-1)N_2 < j \leq j'_1 N_2}, \mu_1), & \text{if } j'_1 > j_1. \end{cases}$$

The algorithm outputs

$$\begin{cases} (j_1 - 1)N_2 + i_{j_1}^*, & \text{if } i_{j'_1}^* = \perp \text{ for all } j'_1 < j_1 \text{ and } i_{j_1}^* \neq \perp; \\ \perp, & \text{if } i_{j'_1}^* = \perp \text{ for all } j'_1 \in [N_1]. \end{cases}$$

Robust Correctness and Compatibility. These are inherited from the underlying AH-BTR. When based on Construction 2, the resultant AH-BTR is compatible with the existing public-key encryption schemes.

Efficiency. Let $\gamma_1, \gamma_2, \gamma_3$ be constants such that the AH-BTR efficiency is

$$T_{\text{Enc}} = (N + 1)^{\gamma_1} \text{poly}(\lambda), \quad |\text{ct}| = (N + 1)^{\gamma_2} \text{poly}(\lambda), \quad T_{\text{Dec}} = (N + 1)^{\gamma_3} \text{poly}(\lambda),$$

then the underlying efficiency is mapped to the resultant efficiency¹⁵ by

$$(\gamma_1, \gamma_2, \gamma_3) \mapsto (1 - \gamma + \gamma\gamma_1, 1 - \gamma + \gamma\gamma_2, \gamma\gamma_3).$$

When based on Construction 2, the resultant AH-BTR enjoys

$$T_{\text{Enc}} = (N + 1) \text{poly}(\lambda), \quad |\text{ct}| = (N + 1)^{1-\gamma} \text{poly}(\lambda), \quad T_{\text{Dec}} = (N + 1)^\gamma \text{poly}(\lambda).$$

Theorem 12 (¶). *Suppose in Construction 3, the underlying AH-BTR scheme old is complete, then so is the resultant AH-BTR.*

Theorem 13 (¶). *Suppose in Construction 3, the underlying AH-BTR scheme old is sound, then so is the resultant AH-BTR.*

Proof (Theorem 12). Let \mathcal{C} be an efficient adversary against the completeness of the resultant scheme. Consider the following efficient adversary \mathcal{C}_{old} against the completeness of old.

- \mathcal{C}_{old} launches \mathcal{C} to obtain

$$\mathcal{D}, \quad \{\text{pk}_j^*\}_{j \in [N]}, \quad 1^{1/\varepsilon^*}.$$

It computes N_1, N_2 as specified by the resultant scheme.

- \mathcal{C}_{old} samples $j_1^* \stackrel{\$}{\leftarrow} [N_1]$, prepares $\mathcal{D}_{j_1^*}$ (using \mathcal{D} , as specified by the resultant scheme), and outputs

$$\mathcal{D}_{j_1^*}, \quad \{\text{pk}_j^*\}_{(j_1^*-1)N_2 < j \leq j_1^* N_2}, \quad 1^{N_1/\varepsilon^*}.$$

¹⁵We assume that old.ct's are of deterministic length so Dec knows the location of each particular old.ct. Alternatively, Enc can store a look-up table of their locations in ct.

Let $B = \text{poly}(\lambda)$ be an upper bound of N_1 . Routine calculation shows

$$\Pr[\mathcal{C}_{\text{old}} \text{ wins}] \geq \frac{1}{B} \Pr[\mathcal{C} \text{ wins}],$$

hence by the completeness of old,

$$\Pr[\mathcal{C} \text{ wins}] \leq B \Pr[\mathcal{C}_{\text{old}} \text{ wins}] = \text{poly}(\lambda) \text{negl}(\lambda) = \text{negl}(\lambda). \quad \square$$

Proof (Theorem 13). Let \mathcal{C} be an efficient adversary against the soundness of the resultant scheme. Consider the following efficient adversary \mathcal{C}_{old} against the soundness of old.

- $\mathcal{C}_{\text{old}}(\text{pk})$ launches $\mathcal{C}(\text{pk})$ to obtain

$$\mathcal{D}, \quad N, \quad i_{\perp}^*, \quad \{\text{pk}_j^*\}_{j \in [N] \setminus \{i_{\perp}^*\}}, \quad 1^{1/\epsilon^*}.$$

It computes N_1, N_2 as specified by the resultant scheme.

- \mathcal{C}_{old} computes $j_1^* = \lceil i_{\perp}^*/N_2 \rceil$ and outputs

$$\mathcal{D}_{j_1^*}, \quad \min\{N_2, N - (j_1^* - 1)N_2\}, \quad i_{\perp}^* - (j_1^* - 1)N_2, \quad \{\text{pk}_j^*\}_{(j_1^* - 1)N_2 < j \leq j_1^* N_2, j \neq i_{\perp}^*}, \quad 1^{N_1/\epsilon^*}.$$

Routine calculation and the soundness of old yield

$$\Pr[\mathcal{C} \text{ wins}] \leq \Pr[\mathcal{C}_{\text{old}} \text{ wins}] = \text{negl}(\lambda). \quad \square$$

7 Lower Bound on Ciphertext Size and Decryption Time

Ideally, we would like a scheme satisfying $|\text{ct}|, T_{\text{Dec}} = \Theta(1)$, yet curiously, even with the heavy hammer of obfuscation, we fail to achieve $|\text{ct}| \cdot T_{\text{Dec}} = o(N)$. It turns out that this limitation is inherent. In this section, we prove that for all secure AH-BTR,

$$|\text{ct}| \cdot T_{\text{Dec}} = \Omega(N),$$

and therefore, we have constructed all the optimal (ignoring $\text{poly}(\lambda)$ factors) AH-BTR schemes in this work, completely pinning down the Pareto front of its efficiency. In fact, we will show a related bound against a restricted kind of broadcast encryption,¹⁶ which can be implemented using AH-BTR in a straight-forward manner.

The scheme is *restricted* in the sense that the users are paired and encryption only broadcasts to those sets for which there is precisely one recipient from each pair. The required security notion is also weaker — it does not consider collusion among multiple non-recipients nor adaptive attacks.

Definition 18 (restricted broadcast encryption and its security). A *restricted broadcast encryption (BE) scheme* (for the purpose of this work) consists of 3 efficient algorithms.

- $\text{Gen}(1^\lambda, 1^N)$ takes a length parameter as input. It outputs a master public key mpk and a list $\{\text{sk}_{j,s}\}_{j \in [N], s \in \{0,1\}}$ of secret keys.
- $\text{Enc}(1^\lambda, \text{mpk}, R, \mu)$ takes as input the master public key mpk , an N -bit string $R \in \{0, 1\}^N$, and a message $\mu \in \{0, 1\}^\lambda$. It outputs a ciphertext ct_R .

¹⁶The lower bound thus also applies to all mildly expressive attribute-based encryption schemes.

- $\text{Dec}^{\text{mpk},i,r,\text{sk}_{i,r},R,\text{ct}_R}(1^\lambda)$ is given random access to the master public key mpk , a secret key with its description $(i, r, \text{sk}_{i,r})$, a ciphertext with its attribute (R, ct_R) . It is supposed to recover μ if and only if $R[i] = r$.

The scheme must be *correct*, i.e., for all $\lambda, N \in \mathbb{N}, R \in \{0, 1\}^N, i \in [N], \mu \in \{0, 1\}^\lambda$,

$$\Pr \left[\begin{array}{l} (\text{mpk}, \{\text{sk}_{j,s}\}_{j \in [N], s \in \{0,1\}}) \xleftarrow{\$} \text{Gen}(1^\lambda, 1^N) \\ \text{ct}_R \xleftarrow{\$} \text{Enc}(1^\lambda, \text{mpk}, R, \mu) \\ : \quad \text{Dec}^{\text{mpk},i,R[i],\text{sk}_{i,R[i]},R,\text{ct}_R}(1^\lambda) = \mu \end{array} \right] = 1.$$

The scheme is *1-key secure for random challenge against uniform adversaries* (or *secure for the purpose of this work*) if

$$\{(1^\lambda, 1^N, \text{mpk}, R, i^*, \mu_0, \text{sk}_{i^*, -R[i^*]}, \boxed{\text{ct}_0})\}_{\lambda \in \mathbb{N}} \approx \{(1^\lambda, 1^N, \text{mpk}, R, i^*, \mu_0, \text{sk}_{i^*, -R[i^*]}, \boxed{\text{ct}_1})\}_{\lambda \in \mathbb{N}}$$

with the components being

$$\begin{array}{lll} (\text{mpk}, \{\text{sk}_{j,s}\}_{j \in [N], s \in \{0,1\}}) \xleftarrow{\$} \text{Gen}(1^\lambda, 1^N), & R \xleftarrow{\$} \{0, 1\}^N, & i^* \xleftarrow{\$} [N], \\ \text{for } b \in \{0, 1\}, & \mu_b \xleftarrow{\$} \{0, 1\}^\lambda, & \text{ct}_b \xleftarrow{\$} \text{Enc}(1^\lambda, \text{mpk}, R, \mu_b). \end{array}$$

for all polynomially bounded $N = N(\lambda)$,¹⁷ where the computational indistinguishability only has to hold against *uniform* adversaries.

Theorem 14 (♣). *For all secure restricted BE,*

$$\max |\text{ct}| \cdot \max T_{\text{Dec}} \geq \frac{N}{1000}$$

for all polynomially bounded $N = N(\lambda)$ and sufficiently large λ , where ct runs through all possible ciphertexts and T_{Dec} the time to probe R and produce output by Dec , both for R of length N .

We remark that “for sufficiently large λ ” is necessary because asymptotic security, by definition, is a tail property unaffected by finitely many λ ’s. The bound starts to hold once the scheme starts to be secure against the adversary used in the proof. While the statement and the proof here apply to perfectly correct schemes with polynomial security, it is straight-forward to adapt them for schemes with sufficient (say, constant) gap between correctness and security.

To prove Theorem 14, we need the following lemma:

Lemma 15 (adapted from [Unr07; Theorem 2]). *For all $N, P \in \mathbb{N}$ subject to $1 \leq P \leq N$, distribution D supported over a finite set Z , function $F : Z \times \{0, 1\}^N \rightarrow \{0, 1\}^S$, there exists a function $G : Z \times \{0, 1\}^N \rightarrow \{0, 1, \perp\}^N$ such that*

$$|\{j \in [N] : G(z, R)[j] \neq \perp\}| \leq P \quad \text{for all } z \in Z \text{ and } R \in \{0, 1\}^N$$

¹⁷ N need not be a computable function of λ . This does not make the security definition “non-uniform”, as a standard guessing argument (with advantage sign correction) applies to an interactive formulation in which the uniform and efficient \mathcal{A} chooses N .

and for all¹⁸ oracle (randomized) algorithm \mathcal{B}^Y making at most T queries to Y ,

$$|\Pr[\mathcal{B}^R(z, F(z, R)) \rightarrow 1] - \Pr[\mathcal{B}^H(z, F(z, R)) \rightarrow 1]| \leq \sqrt{\frac{ST}{2P}},$$

where

$$R \stackrel{\$}{\leftarrow} \{0, 1\}^N, \quad z \stackrel{\$}{\leftarrow} D, \quad H[j] \begin{cases} = G(z, R)[j], & \text{if } G(z, R)[j] \neq \perp; \\ \stackrel{\$}{\leftarrow} \{0, 1\}, & \text{if } G(z, R)[j] = \perp. \end{cases}$$

Proof (Theorem 14). Define

$$S = 1 + \max |\text{ct}|, \quad T = 1 + \max \{\text{number of bits in } R \text{ probed by Dec}\}.$$

For $\lambda, N \geq 1$, it is necessary that $|\text{ct}| \geq 1$ because ct can encode any string μ of length λ , and that $\max T_{\text{Dec}} \geq T$ because Dec performs all the probes and, in addition, produces at least 1 bit of output. Therefore,

$$\max |\text{ct}| \cdot \max T_{\text{Dec}} \geq \frac{\max |\text{ct}| + 1}{2} \cdot \max T_{\text{Dec}} \geq \frac{ST}{2}.$$

It remains to prove $ST \geq \frac{2N}{1000}$ for sufficiently large λ . It suffices to consider the case when $N \geq 2$ and $ST \leq 2N$.

We prepare for Lemma 15. Let P be determined later, and

$$z = \left(\begin{array}{c} \mu, z_{\text{Enc}}, \text{mpk}, \\ \{\text{sk}_{j,s}\}_{j \in [N], s \in \{0,1\}} \end{array} \right) \sim D = \left\{ \begin{array}{c} \mu \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda \\ z_{\text{Enc}} : \text{randomness for Enc} \\ (\text{mpk}, \{\text{sk}_{j,s}\}_{j \in [N], s \in \{0,1\}}) \stackrel{\$}{\leftarrow} \text{Gen}(1^N) \end{array} \right\},$$

$$F(z, R) = 0^{S-|\text{ct}|-1} \|1\| \text{ct}, \quad \text{where } \text{ct} \leftarrow \text{Enc}(\text{mpk}, R, \mu; z_{\text{Enc}}).$$

Let G be the function guaranteed by Lemma 15 and make $\mathcal{B}^Y(z, f)$ do the following.

- Sample $i^* \stackrel{\$}{\leftarrow} [N]$ and query $r^* \leftarrow Y[i^*]$.
- Read $\mu, \text{mpk}, \text{sk}_{i^*, r^*}$ from z . Read ct from f .
- Run $\mu' \stackrel{\$}{\leftarrow} \text{Dec}^{\text{mpk}, i^*, r^*, \text{sk}_{i^*, r^*}, Y, \text{ct}}()$.
- Output 1 if and only if $\mu = \mu'$.

Note that \mathcal{B} indeed makes at most T queries to Y , the first to obtain r^* and the rest to run Dec .

For $w \in \{1, 2, 3, 4, 5\}$, write p_w for $\Pr[\mathcal{B}^{Y_w}(z, f; i^*) \rightarrow 1]$, where

$$i^* \stackrel{\$}{\leftarrow} [N], \quad Y_1 = R,$$

$$Y_2[j] \begin{cases} = G(z, F(z, R))[j], & \text{if } G(z, F(z, R))[j] \neq \perp; \\ \stackrel{\$}{\leftarrow} \{0, 1\}, & \text{if } G(z, F(z, R))[j] = \perp; \end{cases}$$

$$Y_3[j] \begin{cases} = G(z, F(z, R))[j], & \text{if } j \neq i^* \text{ and } G(z, F(z, R))[j] \neq \perp; \\ \stackrel{\$}{\leftarrow} \{0, 1\}, & \text{if } j \neq i^* \text{ and } G(z, F(z, R))[j] = \perp; \\ \stackrel{\$}{\leftarrow} \{0, 1\}, & \text{if } j = i^*; \end{cases}$$

$$Y_4[j] \begin{cases} = R[j], & \text{if } j \neq i^*; \\ \stackrel{\$}{\leftarrow} \{0, 1\}, & \text{if } j = i^*; \end{cases} \quad Y_5[j] \begin{cases} = R[j], & \text{if } j \neq i^*; \\ = \neg R[i^*], & \text{if } j = i^*. \end{cases}$$

¹⁸Here, \mathcal{B}^Y need not be efficient for the lemma to hold. The particular \mathcal{B}^Y used in this work is efficient.

By the correctness of the restricted BE scheme, $p_1 = 1$.

From Lemma 15,

$$|p_1 - p_2| \leq \sqrt{\frac{ST}{2P}}, \quad |p_4 - p_3| \leq \sqrt{\frac{ST}{2P}}.$$

Here, the second inequality is obtained by applying the lemma to

$$\mathcal{C}^Y(z, f) = \mathcal{B}^{Y'}(z, f; i^*), \quad \text{where } i^* \xleftarrow{\$} [N], \quad Y'[j] \begin{cases} = Y[j], & \text{if } j \neq i^*; \\ \xleftarrow{\$} \{0, 1\}, & \text{if } j = i^*. \end{cases}$$

Clearly, $|p_2 - p_3| \leq \frac{P}{N}$. Setting $P = \left\lceil \sqrt[3]{\frac{STN^2}{2}} \right\rceil$, we have

$$\begin{aligned} |p_1 - p_4| &\leq |p_1 - p_2| + |p_2 - p_3| + |p_3 - p_4| \\ &\leq \sqrt{\frac{ST}{2P}} + \frac{P}{N} + \sqrt{\frac{ST}{2P}} \leq 3\sqrt{\frac{ST}{2N}} + \frac{1}{N} < 4\sqrt{\frac{ST}{2N}}, \end{aligned}$$

where the last inequality follows from $N \geq 2$. By how $Y[i^*]$ is set,

$$p_4 = \frac{p_1 + p_5}{2} \implies p_5 = p_1 - 2(p_1 - p_4) \geq p_1 - 2|p_1 - p_4| > 1 - 8\sqrt{\frac{ST}{2N}}.$$

Consider the following adversary $\mathcal{A}(\text{mpk}, R, i^*, \mu_0, \text{sk}_{i^*, -R[i^*]}, \text{ct})$ against the security of the restricted BE scheme.

- Construct Y_5 from R and let $r^* \leftarrow Y_5[i^*] = -R[i^*]$.
- Run $\mu' \xleftarrow{\$} \text{Dec}^{\text{mpk}, i^*, r^*, \text{sk}_{i^*, r^*}, Y_5, \text{ct}}()$, i.e., pretend $R[i^*]$ were $-R[i^*]$ and try decrypting using the (supposedly non-decrypting) key given to \mathcal{A} .
- Output 1 if and only if $\mu' = \mu_0$.

If $\text{ct} = \text{ct}_1$ is an encryption of μ_1 , then μ_0 is uniformly random and independent of everything else, hence

$$\Pr[\mathcal{A}(\dots) \rightarrow 1 \text{ with } \text{ct} = \text{ct}_1] \leq 2^{-\lambda}.$$

Note that \mathcal{A} is a uniform adversary. By the security of the restricted BE scheme,

$$p_5 = \Pr[\mathcal{B}^{Y_5}(z, f; i^*) \rightarrow 1] = \Pr[\mathcal{A}(\dots) \rightarrow 1 \text{ with } \text{ct} = \text{ct}_0] \leq 2^{-\lambda} + \text{negl}(\lambda) < \frac{1}{5}$$

for sufficiently large λ , which gives

$$1 - 8\sqrt{\frac{ST}{2N}} < p_5 < \frac{1}{5} \implies ST > \frac{2N}{1000}. \quad \square$$

Corollary 16 (¶). *For all secure AH-BTR,*

$$\max |\text{ct}| \cdot \max T_{\text{Dec}} \geq \frac{N}{1000}$$

for all polynomially bounded $N = N(\lambda)$ and sufficiently large λ ,¹⁹ where T_{Dec} only counts the time to probe pk_j 's and produce output. Ignoring $\text{poly}(\lambda)$ factors, Construction 3 achieves all possible optimal trade-offs in terms of the exponents over N in the dependency of ciphertext size and (actual) decryption time, fully demonstrating the Pareto front of AH-BTR efficiency.

¹⁹The remarks following Theorem 14 are also applicable here.

Proof (Corollary 16). Suppose $\text{ahBTR} = (\text{ahBTR.Gen}, \text{ahBTR.Enc}, \text{ahBTR.Dec}, \text{ahBTR.Trace})$ is a secure AH-BTR and construct the following restricted BE scheme.

- $\text{Gen}(1^N)$ runs

$$(\text{pk}_{j,s}, \text{sk}_{j,s}) \stackrel{\$}{\leftarrow} \text{ahBTR.Gen}() \quad \text{for } j \in [N], s \in \{0, 1\}$$

and outputs $\text{mpk} = \{\text{pk}_{j,s}\}_{j \in [N], s \in \{0,1\}}$ with $\{\text{sk}_{j,s}\}_{j \in [N], s \in \{0,1\}}$.

- $\text{Enc}(\text{mpk}, R, \mu)$ runs and outputs

$$\text{ct} \stackrel{\$}{\leftarrow} \text{ahBTR.Enc}(\{\text{pk}_{j,R[j]}\}_{j \in [N]}, \mu).$$

- $\text{Dec}^{\text{mpk}, i, r, \text{sk}_{i,r}, R, \text{ct}}()$ runs $\text{ahBTR.Dec}^{K, \text{ct}}(N, i, \text{sk}_{i,r})$, where K is an oracle implemented by Dec for ahBTR.Dec to probe pk_j 's. Whenever ahBTR.Dec probes $\text{pk}_j[m_0]$, we make Dec probe $R[j]$ and answer $\text{pk}_{j,R[j]}[m_0]$.

It is straight-forward to verify that the constructed scheme is correct and secure. Since a restricted BE ciphertext is precisely an AH-BTR ciphertext, each probe to pk_j 's by ahBTR.Dec translates to exactly one probe to $R[j]$ by Dec with no more additional probes by Dec on its own, and Dec outputs whatever ahBTR.Dec outputs, the corollary follows from Theorem 14. \square

Acknowledgement. The author was supported by NSF grants CNS-1936825 (CAREER), CNS-2026774, a J.P. Morgan AI Research Award, and a Simons Collaboration on the Theory of Algorithmic Fairness. The views expressed in this work are those of the author and do not reflect the official policy or position of any of the supporters. The author thanks Huijia Lin for her encouragement and support for him to pursue an independent research project (culminating this paper), Hoeteck Wee for helpful suggestions on writing from another context, Daniel Wichs for helpful discussions, and anonymous reviewers for their constructive feedback. He started researching this topic after rereading a post [sil21] on V2EX.

References

- [AK08] Per Austrin and Gunnar Kreitz. Lower bounds for subset cover based broadcast encryption. In Serge Vaudenay, editor, *AFRICACRYPT 08*, volume 5023 of *LNCS*, pages 343–356. Springer, Heidelberg, June 2008.
- [AKYY23] Shweta Agrawal, Simran Kumari, Anshu Yadav, and Shota Yamada. Broadcast, trace and revoke with optimal parameters from polynomial hardness. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part III*, volume 14006 of *LNCS*, pages 605–636. Springer, Heidelberg, April 2023.
- [AL10] Nuttapon Attrapadung and Benoît Libert. Functional encryption for inner product: Achieving constant-size ciphertexts with adaptive security or support for negation. In Phong Q. Nguyen and David Pointcheval, editors, *PKC 2010*, volume 6056 of *LNCS*, pages 384–402. Springer, Heidelberg, May 2010.

- [AL18] Prabhanjan Ananth and Alex Lombardi. Succinct garbling schemes from functional encryption through a local simulation paradigm. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 455–472. Springer, Heidelberg, November 2018.
- [AWY20] Shweta Agrawal, Daniel Wichs, and Shota Yamada. Optimal broadcast encryption from LWE and pairings in the standard model. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part I*, volume 12550 of *LNCS*, pages 149–178. Springer, Heidelberg, November 2020.
- [AY20] Shweta Agrawal and Shota Yamada. Optimal broadcast encryption from pairings and LWE. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 13–43. Springer, Heidelberg, May 2020.
- [BC95] Carlo Blundo and Antonella Cresti. Space requirements for broadcast encryption. In Alfredo De Santis, editor, *EUROCRYPT'94*, volume 950 of *LNCS*, pages 287–298. Springer, Heidelberg, May 1995.
- [BF99] Dan Boneh and Matthew K. Franklin. An efficient public key traitor tracing scheme. In Michael J. Wiener, editor, *CRYPTO'99*, volume 1666 of *LNCS*, pages 338–353. Springer, Heidelberg, August 1999.
- [BGI⁺01] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 1–18. Springer, Heidelberg, August 2001.
- [BGI14] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In Hugo Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 501–519. Springer, Heidelberg, March 2014.
- [BGW05] Dan Boneh, Craig Gentry, and Brent Waters. Collusion resistant broadcast encryption with short ciphertexts and private keys. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 258–275. Springer, Heidelberg, August 2005.
- [BHR12] Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Foundations of garbled circuits. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *ACM CCS 2012*, pages 784–796. ACM Press, October 2012.
- [BN08] Dan Boneh and Moni Naor. Traitor tracing with constant size ciphertext. In Peng Ning, Paul F. Syverson, and Somesh Jha, editors, *ACM CCS 2008*, pages 501–510. ACM Press, October 2008.
- [BSW06] Dan Boneh, Amit Sahai, and Brent Waters. Fully collusion resistant traitor tracing with short ciphertexts and private keys. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 573–592. Springer, Heidelberg, May / June 2006.
- [BSW11] Dan Boneh, Amit Sahai, and Brent Waters. Functional encryption: Definitions and challenges. In Yuval Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 253–273. Springer, Heidelberg, March 2011.

- [BV22] Zvika Brakerski and Vinod Vaikuntanathan. Lattice-inspired broadcast encryption and succinct ciphertext-policy ABE. In Mark Braverman, editor, *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, volume 215 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 28:1–28:20, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [BW06] Dan Boneh and Brent Waters. A fully collusion resistant broadcast, trace, and revoke system. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 2006*, pages 211–220. ACM Press, October / November 2006.
- [BW13] Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 280–300. Springer, Heidelberg, December 2013.
- [BWZ14] Dan Boneh, Brent Waters, and Mark Zhandry. Low overhead broadcast encryption from multilinear maps. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 206–223. Springer, Heidelberg, August 2014.
- [BZ14] Dan Boneh and Mark Zhandry. Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 480–499. Springer, Heidelberg, August 2014.
- [CDG⁺17] Chongwon Cho, Nico Döttling, Sanjam Garg, Divya Gupta, Peihan Miao, and Antigoni Polychroniadou. Laconic oblivious transfer and its applications. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 33–65. Springer, Heidelberg, August 2017.
- [CES21] Kelong Cong, Karim Eldefrawy, and Nigel P. Smart. Optimizing registration based encryption. In Maura B. Paterson, editor, *18th IMA International Conference on Cryptography and Coding*, volume 13129 of *LNCS*, pages 129–157. Springer, Heidelberg, December 2021.
- [CFN94] Benny Chor, Amos Fiat, and Moni Naor. Tracing traitors. In Yvo Desmedt, editor, *CRYPTO'94*, volume 839 of *LNCS*, pages 257–270. Springer, Heidelberg, August 1994.
- [Cha07] Melissa Chase. Multi-authority attribute based encryption. In Salil P. Vadhan, editor, *TCC 2007*, volume 4392 of *LNCS*, pages 515–534. Springer, Heidelberg, February 2007.
- [CHK22] Henry Corrigan-Gibbs, Alexandra Henzinger, and Dmitry Kogan. Single-server private information retrieval with sublinear amortized time. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 3–33. Springer, Heidelberg, May / June 2022.

- [CVW⁺18] Yilei Chen, Vinod Vaikuntanathan, Brent Waters, Hoeteck Wee, and Daniel Wichs. Traitor-tracing from LWE made simple and attribute-based. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 341–369. Springer, Heidelberg, November 2018.
- [Del07] Cécile Delerablée. Identity-based broadcast encryption with constant size ciphertexts and private keys. In Kaoru Kurosawa, editor, *ASIACRYPT 2007*, volume 4833 of *LNCS*, pages 200–215. Springer, Heidelberg, December 2007.
- [DHMR08] Vanesa Daza, Javier Herranz, Paz Morillo, and Carla Ràfols. *Ad-hoc* threshold broadcast encryption with shorter ciphertexts. *Electronic Notes in Theoretical Computer Science*, 192(2):3–15, 2008. Proceedings of the Third Workshop on Cryptography for Ad-hoc Networks (WCAN 2007).
- [DLY21] Ivan Bjerre Damgård, Kasper Green Larsen, and Sophia Yakubov. Broadcast secret-sharing, bounds and applications. In Stefano Tessaro, editor, *2nd Conference on Information-Theoretic Cryptography (ITC 2021)*, volume 199 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 10:1–10:20, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [DPP07] Cécile Delerablée, Pascal Paillier, and David Pointcheval. Fully collusion secure dynamic broadcast encryption with constant-size ciphertexts or decryption keys. In Tsuyoshi Takagi, Tatsuaki Okamoto, Eiji Okamoto, and Takeshi Okamoto, editors, *PAIRING 2007*, volume 4575 of *LNCS*, pages 39–59. Springer, Heidelberg, July 2007.
- [FFM⁺23] Danilo Francati, Daniele Friolo, Monosij Maitra, Giulio Malavolta, Ahmadreza Rahimi, and Daniele Venturi. Registered (inner-product) functional encryption. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part V*, volume 14442 of *LNCS*, pages 98–133. Springer, Heidelberg, December 2023.
- [FKdP23] Dario Fiore, Dimitris Kolonelos, and Paola de Perthuis. Cuckoo commitments: Registration-based encryption and key-value map commitments for large spaces. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part V*, volume 14442 of *LNCS*, pages 166–200. Springer, Heidelberg, December 2023.
- [FN94] Amos Fiat and Moni Naor. Broadcast encryption. In Douglas R. Stinson, editor, *CRYPTO'93*, volume 773 of *LNCS*, pages 480–491. Springer, Heidelberg, August 1994.
- [FWW23] Cody Freitag, Brent Waters, and David J. Wu. How to use (plain) witness encryption: Registered ABE, flexible broadcast, and more. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part IV*, volume 14084 of *LNCS*, pages 498–531. Springer, Heidelberg, August 2023.
- [GGM84] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions (extended abstract). In *25th FOCS*, pages 464–479. IEEE Computer Society Press, October 1984.

- [GGSW13] Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 467–476. ACM Press, June 2013.
- [GHM⁺19] Sanjam Garg, Mohammad Hajiabadi, Mohammad Mahmoody, Ahmadreza Rahimi, and Sruthi Sekar. Registration-based encryption from standard assumptions. In Dongdai Lin and Kazue Sako, editors, *PKC 2019, Part II*, volume 11443 of *LNCS*, pages 63–93. Springer, Heidelberg, April 2019.
- [GHMR18] Sanjam Garg, Mohammad Hajiabadi, Mohammad Mahmoody, and Ahmadreza Rahimi. Registration-based encryption: Removing private-key generator from IBE. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part I*, volume 11239 of *LNCS*, pages 689–718. Springer, Heidelberg, November 2018.
- [GK16] Shafi Goldwasser and Yael Tauman Kalai. Cryptographic assumptions: A position paper. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A, Part I*, volume 9562 of *LNCS*, pages 505–522. Springer, Heidelberg, January 2016.
- [GKMR23] Noemi Glaeser, Dimitris Kolonelos, Giulio Malavolta, and Ahmadreza Rahimi. Efficient registration-based encryption. In Weizhi Meng, Christian Damsgaard Jensen, Cas Cremers, and Engin Kirda, editors, *ACM CCS 2023*, pages 1065–1079. ACM Press, November 2023.
- [GKRW18] Rishab Goyal, Venkata Koppula, Andrew Russell, and Brent Waters. Risky traitor tracing and new differential privacy negative results. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part I*, volume 10991 of *LNCS*, pages 467–497. Springer, Heidelberg, August 2018.
- [GKSW10] Sanjam Garg, Abishek Kumarasubramanian, Amit Sahai, and Brent Waters. Building efficient fully collusion-resilient traitor tracing and revocation schemes. In Ehab Al-Shaer, Angelos D. Keromytis, and Vitaly Shmatikov, editors, *ACM CCS 2010*, pages 121–130. ACM Press, October 2010.
- [GKW15] Romain Gay, Iordanis Kerenidis, and Hoeteck Wee. Communication complexity of conditional disclosure of secrets and attribute-based encryption. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 485–502. Springer, Heidelberg, August 2015.
- [GKW18] Rishab Goyal, Venkata Koppula, and Brent Waters. Collusion resistant traitor tracing from learning with errors. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *50th ACM STOC*, pages 660–670. ACM Press, June 2018.
- [GKW19] Rishab Goyal, Venkata Koppula, and Brent Waters. New approaches to traitor tracing with embedded identities. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019, Part II*, volume 11892 of *LNCS*, pages 149–179. Springer, Heidelberg, December 2019.
- [GLW23] Junqing Gong, Ji Luo, and Hoeteck Wee. Traitor tracing with $N^{1/3}$ -size ciphertexts and $O(1)$ -size keys from k -Lin. In Carmit Hazay and Martijn

- Stam, editors, *EUROCRYPT 2023, Part III*, volume 14006 of *LNCS*, pages 637–668. Springer, Heidelberg, April 2023.
- [GLWW24] Rachit Garg, George Lu, Brent Waters, and David J. Wu. Reducing the CRS size in registered ABE systems. Cryptology ePrint Archive, Report 2024/749, 2024. To appear at Crypto 2024, available at <https://eprint.iacr.org/2024/749>.
- [GPSW06] Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 2006*, pages 89–98. ACM Press, October / November 2006. Available as Cryptology ePrint Archive Report 2006/309.
- [GQWW19] Rishab Goyal, Willy Quach, Brent Waters, and Daniel Wichs. Broadcast and trace with N^ϵ ciphertext size from standard assumptions. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 826–855. Springer, Heidelberg, August 2019.
- [GV20] Rishab Goyal and Satyanarayana Vusirikala. Verifiable registration-based encryption. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part I*, volume 12170 of *LNCS*, pages 621–651. Springer, Heidelberg, August 2020.
- [GVW19] Rishab Goyal, Satyanarayana Vusirikala, and Brent Waters. Collusion resistant broadcast and trace from positional witness encryption. In Dongdai Lin and Kazue Sako, editors, *PKC 2019, Part II*, volume 11443 of *LNCS*, pages 3–33. Springer, Heidelberg, April 2019.
- [GW09] Craig Gentry and Brent Waters. Adaptive security in broadcast encryption systems (with short ciphertexts). In Antoine Joux, editor, *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 171–188. Springer, Heidelberg, April 2009.
- [GW11] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In Lance Fortnow and Salil P. Vadhan, editors, *43rd ACM STOC*, pages 99–108. ACM Press, June 2011.
- [HLWW23] Susan Hohenberger, George Lu, Brent Waters, and David J. Wu. Registered attribute-based encryption. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part III*, volume 14006 of *LNCS*, pages 511–542. Springer, Heidelberg, April 2023.
- [IW14] Yuval Ishai and Hoeteck Wee. Partial garbling schemes and their applications. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *ICALP 2014, Part I*, volume 8572 of *LNCS*, pages 650–662. Springer, Heidelberg, July 2014.
- [JLL23] Aayush Jain, Huijia Lin, and Ji Luo. On the optimal succinctness and efficiency of functional encryption and attribute-based encryption. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part III*, volume 14006 of *LNCS*, pages 479–510. Springer, Heidelberg, April 2023.

- [JLS21] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from well-founded assumptions. In Samir Khuller and Virginia Vassilevska Williams, editors, *53rd ACM STOC*, pages 60–73. ACM Press, June 2021.
- [JLS22] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from LPN over \mathbb{F}_p , DLIN, and PRGs in NC^0 . In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part I*, volume 13275 of *LNCS*, pages 670–699. Springer, Heidelberg, May / June 2022.
- [KMW23] Dimitris Kolonelos, Giulio Malavolta, and Hoeteck Wee. Distributed broadcast encryption from bilinear groups. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part V*, volume 14442 of *LNCS*, pages 407–441. Springer, Heidelberg, December 2023.
- [KNTY19] Fuyuki Kitagawa, Ryo Nishimaki, Keisuke Tanaka, and Takashi Yamakawa. Adaptively secure and succinct functional encryption: Improving security and efficiency, simultaneously. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 521–551. Springer, Heidelberg, August 2019.
- [KPTZ13] Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 2013*, pages 669–684. ACM Press, November 2013.
- [KY01] Aggelos Kiayias and Moti Yung. On crafty pirates and foxy tracers. In *ACM Workshop on Security and Privacy in Digital Rights Management, DRM '01*, pages 22–39, Berlin, Heidelberg, 2001. Springer-Verlag.
- [KY09] Jonathan Katz and Arkady Yerukhimovich. On black-box constructions of predicate encryption from trapdoor permutations. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 197–213. Springer, Heidelberg, December 2009.
- [KYDB98] Kaoru Kurosawa, Takuya Yoshida, Yvo Desmedt, and Mike Burmester. Some bounds and a construction for secure broadcast encryption. In Kazuo Ohta and Dingyi Pei, editors, *ASIACRYPT'98*, volume 1514 of *LNCS*, pages 420–433. Springer, Heidelberg, October 1998.
- [LP09] Yehuda Lindell and Benny Pinkas. A proof of security of Yao's protocol for two-party computation. *Journal of Cryptology*, 22(2):161–188, April 2009.
- [LS98] Michael Luby and Jessica Staddon. Combinatorial bounds for broadcast encryption. In Kaisa Nyberg, editor, *EUROCRYPT'98*, volume 1403 of *LNCS*, pages 512–526. Springer, Heidelberg, May / June 1998.
- [LT17] Huijia Lin and Stefano Tessaro. Indistinguishability obfuscation from trilinear maps and block-wise local PRGs. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 630–660. Springer, Heidelberg, August 2017.

- [LZ17] Qipeng Liu and Mark Zhandry. Decomposable obfuscation: A framework for building applications of obfuscation from polynomial hardness. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 138–169. Springer, Heidelberg, November 2017.
- [NNL01] Dalit Naor, Moni Naor, and Jeffery Lotspiech. Revocation and tracing schemes for stateless receivers. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 41–62. Springer, Heidelberg, August 2001.
- [NP01] Moni Naor and Benny Pinkas. Efficient trace and revoke schemes. In Yair Frankel, editor, *FC 2000*, volume 1962 of *LNCS*, pages 1–20. Springer, Heidelberg, February 2001.
- [NWZ16] Ryo Nishimaki, Daniel Wichs, and Mark Zhandry. Anonymous traitor tracing: How to embed arbitrary information in a key. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 388–419. Springer, Heidelberg, May 2016.
- [PPS12] Duong Hieu Phan, David Pointcheval, and Mario Strefler. Decentralized dynamic broadcast encryption. In Ivan Visconti and Roberto De Prisco, editors, *SCN 12*, volume 7485 of *LNCS*, pages 166–183. Springer, Heidelberg, September 2012.
- [SF07] Ryuichi Sakai and Jun Furukawa. Identity-based broadcast encryption. Cryptology ePrint Archive, Report 2007/217, 2007. <https://eprint.iacr.org/2007/217>.
- [sil21] sillydaddy. gpg 加密文件：一份加密文件，可以被不同的密码解密 (GPG file encryption: One encrypted file can be decrypted by many keys). <https://v2ex.com/t/759538>, March 2021. Retrieved on 20 May 2022, archived at <https://web.archive.org/web/20220520040245/https://v2ex.com/t/759538>.
- [SW05] Amit Sahai and Brent R. Waters. Fuzzy identity-based encryption. In Ronald Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 457–473. Springer, Heidelberg, May 2005.
- [SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In David B. Shmoys, editor, *46th ACM STOC*, pages 475–484. ACM Press, May / June 2014.
- [Unr07] Dominique Unruh. Random oracles and auxiliary input. In Alfred Menezes, editor, *CRYPTO 2007*, volume 4622 of *LNCS*, pages 205–223. Springer, Heidelberg, August 2007.
- [Wee22] Hoeteck Wee. Optimal broadcast encryption and CP-ABE from evasive lattice assumptions. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 217–241. Springer, Heidelberg, May / June 2022.
- [WQZD10] Qianhong Wu, Bo Qin, Lei Zhang, and Josep Domingo-Ferrer. Ad hoc broadcast encryption (poster presentation). In Ehab Al-Shaer, Angelos D.

- Keromytis, and Vitaly Shmatikov, editors, *ACM CCS 2010*, pages 741–743. ACM Press, October 2010.
- [Yao86] Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *27th FOCS*, pages 162–167. IEEE Computer Society Press, October 1986.
- [Yao90] Andrew Chi-Chih Yao. Coherent functions and program checkers (extended abstract). In *22nd ACM STOC*, pages 84–94. ACM Press, May 1990.
- [Zha20a] Mark Zhandry. New techniques for traitor tracing: Size $N^{1/3}$ and more from pairings. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part I*, volume 12170 of *LNCS*, pages 652–682. Springer, Heidelberg, August 2020.
- [Zha20b] Mark Zhandry. New techniques for traitor tracing: Size $N^{1/3}$ and more from pairings. Cryptology ePrint Archive, Report 2020/954, 2020. <https://eprint.iacr.org/2020/954>.
- [Zha20c] Mark Zhandry. Schrödinger’s pirate: How to trace a quantum decoder. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part III*, volume 12552 of *LNCS*, pages 61–91. Springer, Heidelberg, November 2020.
- [Zha21] Mark Zhandry. White box traitor tracing. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part IV*, volume 12828 of *LNCS*, pages 303–333, Virtual Event, August 2021. Springer, Heidelberg.
- [ZLZ⁺24] Ziqi Zhu, Jiangtao Li, Kai Zhang, Junqing Gong, and Haifeng Qian. Registered functional encryptions from pairings. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part II*, volume 14652 of *LNCS*, pages 373–402. Springer, Heidelberg, May 2024.
- [ZZGQ23] Ziqi Zhu, Kai Zhang, Junqing Gong, and Haifeng Qian. Registered ABE via predicate encodings. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part V*, volume 14442 of *LNCS*, pages 66–97. Springer, Heidelberg, December 2023.