

Kent Academic Repository

Full text document (pdf)

Citation for published version

Agrawal, Shweta and Bhattacharjee, Sanjay and Phan, Duong Hieu and Stehlé, Damien and Yamada, Shota (2017) Efficient Public Trace and Revoke from Standard Assumptions. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security. . pp. 2277-2293. Association for Computing Machinery, New York- United States ISBN 978-1-4503-4946-8.

DOI

<https://doi.org/10.1145/3133956.3134041>

Link to record in KAR

<https://kar.kent.ac.uk/83284/>

Document Version

Author's Accepted Manuscript

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

Efficient Public Trace and Revoke from Standard Assumptions

Shweta Agrawal
IIT Madras, India,
shweta@iitm.ac.in

and

Sanjay Bhattacharjee
Turing Lab (ASU), ISI Kolkata, India,
sanjay.bhattacharjee@gmail.com

and

Duong Hieu Phan
XLIM (U. Limoges, CNRS), France,
duong-hieu.phan@unilim.fr

and

Damien Stehlé
ENS de Lyon, LIP (U. Lyon, CNRS, ENSL, INRIA, UCBL), France,
damien.stehle@ens-lyon.fr

and

Shota Yamada
National Institute of Advanced Industrial Science and Technology (AIST), Japan,
yamada-shota@aist.go.jp

Abstract

We provide efficient constructions for trace-and-revoke systems with public traceability in the black-box confirmation model. Our constructions achieve adaptive security, are based on standard assumptions and achieve significant efficiency gains compared to previous constructions.

Our constructions rely on a generic transformation from inner product functional encryption (IPFE) schemes to trace-and-revoke systems. Our transformation requires the underlying IPFE scheme to only satisfy a very weak notion of security – the attacker may only request a bounded number of *random* keys – in contrast to the standard notion of security where she may request an unbounded number of arbitrarily chosen keys. We exploit the much weaker security model to provide a new construction for bounded collusion and random key IPFE from the learning with errors assumption (LWE), which enjoys improved efficiency compared to the scheme of Agrawal *et al.* [CRYPTO'16].

Together with IPFE schemes from Agrawal *et al.*, we obtain trace and revoke from LWE, Decision Diffie Hellman and Decision Composite Residuosity.

Keywords: Inner-product functional encryption; Trace-and-revoke; Public traceability.

1 Introduction

A traitor tracing system Chor et al. (1994) is a multi-receiver encryption system, which aids content distributors in identifying malicious receivers that construct pirate decryption boxes. In more detail, data is encrypted under some public key pk and each legitimate user of the system is provided a secret key sk_i that allows her to decrypt

the content. Since nothing prevents a user from making copies of her key and selling them for profit, traitor tracing systems provide the following security guarantee to deter such behavior: if a coalition of users pool together their keys and construct a pirate decoder box capable of decrypting the ciphertext, then there is an efficient “trace” algorithm which, given access to any such decoder, outputs the identity of at least one guilty user.

An orthogonal functionality is that of broadcast encryption Fiat and Naor (1993), where the content provider encrypts data to some subset S of users. Functionality requires that any user in S be able decrypt the content and security posits that no collusion of users outside S can do so. *Trace-and-revoke* systems combine these two functionalities – when the system is deployed, the content is encrypted to all users on the channel. However, if copyright infringement occurs, then tracing is used to detect the malicious users, or “traitors”, and future content is encrypted using broadcast encryption to all users except the traitors.

Trace-and-revoke systems have been studied extensively Naor and Pinkas (2000); Naor et al. (2001); Dodis and Fazio (2003); Kim et al. (2003); Phan and Trinh (2011) and are notoriously hard to construct (please see Boneh and Waters (2006) for a detailed discussion). A desirable attribute for trace-and-revoke systems is *public traceability*, meaning that the tracing algorithm does not require any additional secrets. Due to this property, the overall system remains secure even if the tracing party is compromised. Moreover, the tracing capability can be outsourced to an untrusted party in this setting.

To the best of our knowledge, trace-and-revoke systems with public traceability have only been achieved by Boneh and Waters Boneh and Waters (2006), and quite recently by Nishimaki, Wichs and Zhandry (NWZ) Nishimaki et al. (2016). The Boneh-Waters construction is quite powerful in that it supports malicious collusions of unbounded size but its ciphertexts are very large (their size grows proportionally to \sqrt{N} , where N is the total number of users) and the scheme relies on pairing groups of composite order. To achieve a ciphertext size that does not depend on the total number of users in the system, we consider the bounded collusion model, where the number of possible traitors is a priori bounded by some t that is polynomial in the security parameter λ . The bounded collusion model is quite standard in traitor tracing schemes and has received significant attention; however, until the work of Nishimaki *et al.* (NWZ) Nishimaki et al. (2016), all known schemes in this model Boneh and Franklin (1999); Hofheinz and Striecks (2014); Ling et al. (2014) support either revocation or public traceability but not both.

Recently, Nishimaki *et al.* (NWZ) Nishimaki et al. (2016) provided a generic construction for traitor tracing systems from functional encryption schemes. Functional encryption Sahai and Waters (2005); Boneh et al. (2011) is a generalization of public key encryption allowing fine grained access to encrypted data. We note that the strongest constructions in Nishimaki et al. (2016) are based on the existence of indistinguishability obfuscation Barak et al. (2012), for which we do not at present have any candidate construction based on well established hardness assumptions. Since our focus is on efficient constructions based on well established hardness assumptions, we do not consider these in this work. One may also instantiate the NWZ compiler with a bounded collusion functional encryption scheme which can be based on standard assumptions such as the existence of public key encryption Gorbunov et al. (2012) or subexponential time hardness of learning with errors (LWE) Goldwasser et al. (2013); Agrawal and Rosen (2016). For trace and revoke, this results in a construction that supports public black box traceability and adaptive security in addition to anonymity of honest users and an exponential size universe of identities.

However, the generic nature of their construction results in loss of concrete efficiency. For instance, when based on the bounded collusion FE of Gorbunov et al. (2012), the resulting scheme has a ciphertext size growing at least as $O((r+t)^5 \text{Poly}(\lambda))$ where r is the maximum size of the list of revoked users and t the maximum coalition size (please see Appendix 6 for an explanation of the bound). By relying on learning with errors, this blowup can be improved to $O((r+t)^4 \text{Poly}(\lambda))$ but at the cost of relying on heavy machinery such as attribute based encryption Gorbunov et al. (2013) and fully homomorphic encryption Goldwasser et al. (2013). Additionally, this construction must also rely on complexity leveraging for adaptive security and learning with errors with subexponential error rates. The bounded collusion FE of Agrawal and Rosen (2016) leads to better asymptotic

bounds $O(r + t)^3 \text{Poly}(\lambda)$) but suffers from large polynomial factors which hurt concrete efficiency.

Our Approach. In this work, we revisit the connection between functional encryption and trace-and-revoke systems and observe that the notion of FE required for bounded collusion trace-and-revoke schemes is significantly weaker than that considered by Nishimaki et al. (2016). To begin, we show that the functionality required from the underlying functional encryption scheme may be significantly weakened; rather than FE for polynomial sized circuits,¹ we show that *inner product* functional encryption (IPFE) Abdalla et al. (2015); Agrawal et al. (2016) suffices. Efficient constructions for IPFE satisfying adaptive security are available Agrawal et al. (2016), leading to trace-and-revoke systems which are significantly simpler and more efficient than those implied by Nishimaki et al. (2016). We further improve our constructions by observing that for the application of trace and revoke, the underlying IPFE schemes must be secure in a much weaker security model than full fledged IPFE: the adversary may be restricted to only make a bounded number of key queries, and only key queries for *randomly* chosen vectors. We exploit the much weaker security model to provide new constructions for bounded collusion and random key IPFE from LWE and Decision Composite Residuosity (DCR), which enjoy substantial benefits over using those of Agrawal et al. (2016) in terms of parameter sizes. The improvement is greatest for the LWE construction, as the LWE modulus can be slightly super-polynomial rather than subexponential, itself allowing to choose a smaller LWE dimension.

Our Results. We construct efficient trace-and-revoke systems with bounded collusion resistance, from standard assumptions. Our schemes support public, black-box traceability and achieve the strongest notion of adaptive security as defined by Boneh and Waters (2006). Our construction is generic and leverages recent constructions of modular inner product functional encryption (IPFE) Abdalla et al. (2015); Agrawal et al. (2016). Moreover, by targeting the weak security game required by our application, we obtain more efficient versions of IPFE schemes that suffice for our purposes. While Nishimaki et al. (2016) achieves trace-and-revoke in the strong security model under the existence of public-key encryption, our approach leads to significantly more efficient schemes under the DCR, LWE and DDH assumptions. In particular, we achieve ciphertext and key sizes that are *linear* in the sum of revoked list size r and maximum coalition size t . Our DDH-based construction achieves ciphertext and key sizes $O((r + t)\lambda)$, our DCR-based construction achieves ciphertext and key sizes $\tilde{O}((r + t)\lambda^3)$, while our LWE-based construction has ciphertext size $\tilde{O}(r + t + \lambda)$ and key size $\tilde{O}((r + t + \lambda)\lambda)$. We note that our security definition considers the strongest notion of “usefulness” Boneh and Waters (2006) of the pirate decoder, which is not satisfied by most other constructions. Indeed some schemes Naor and Pinkas (2000); Dodis and Fazio (2003) are actually insecure in this strong game (see Appendix 6 for a detailed discussion). Finally, we give a DDH-based traitor tracing construction (without revocation) that supports encryption of k messages with ciphertext and key sizes $O((k + t)\lambda)$. This improves ciphertext expansion over the trace-and-revoke construction, as the plaintext messages are binary.

Our Techniques. Let $\mathcal{FE} = (\mathcal{FE}.\text{Setup}, \mathcal{FE}.\text{KeyGen}, \mathcal{FE}.\text{Enc}, \mathcal{FE}.\text{Dec})$ be a functional encryption scheme for the inner-product functionality over \mathbb{Z}_p^ℓ . Recall the inner product functionality: the ciphertext encodes a vector $\mathbf{v} \in \mathbb{Z}_p^\ell$, the secret key encodes a vector $\mathbf{x} \in \mathbb{Z}_p^\ell$ and decryption recovers the inner product $\langle \mathbf{x}, \mathbf{v} \rangle \bmod p$.

To construct a trace-and-revoke scheme, we proceed as follows. At the time of key generation, a user id is first assigned a uniformly sampled vector $\mathbf{x}_{\text{id}} \in \mathbb{Z}_p^\ell$ and the entry $p_{\text{id}} = (\text{id}, \mathbf{x}_{\text{id}})$ is stored in the public directory pd for full public traceability. We may consider revocation and tracing as two distinct functionalities that need to be combined so that neither interferes with the security properties of the other. We employ two different techniques to implement these functionalities.

¹More accurately, the circuits required by the NWZ compiler are relatively simple, but ones for which we do not know any better FE constructions than the general case.

To revoke a set \mathcal{R} of users with $|\mathcal{R}| \leq r$, we first deterministically compute a vector $\mathbf{v}_{\mathcal{R}} \in \mathbb{Z}_p^\ell$ such that for all $\text{id} \in \mathcal{R}$, we have $\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle = 0$ (modulo p). Note that this can be implemented only if $r < \ell$. At the same time, for a user $\text{id} \notin \mathcal{R}$, the probability that $\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle = 0$ must be negligible, as otherwise it would de facto be handled as a revoked user. To guarantee this, we require that p is $\lambda^{\omega(1)}$. Since we choose \mathbf{x}_{id} uniformly random, we have $\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle \neq 0$ for $\text{id} \notin \mathcal{R}$ with overwhelming probability.

Using the underlying \mathcal{FE} scheme, we would like to encrypt the message $m \in \mathbb{Z}_p^*$ such that the users in the set \mathcal{R} are not able to decrypt the message, but users not in \mathcal{R} are able to decrypt. We achieve it as follows:

$$C = (\mathcal{FE}.\text{Enc}(\text{pk}, m \cdot \mathbf{v}_{\mathcal{R}}), \mathcal{R}) = (C_1, C_2).$$

Here the operation \cdot denotes the scalar multiplication of each component of $\mathbf{v}_{\mathcal{R}}$ with m . To decrypt, the user id with the vector \mathbf{x}_{id} and the \mathcal{FE} secret key $sk_{\mathbf{x}_{\text{id}}}$ proceeds as follows:

(a) Compute $\mathbf{v}_{\mathcal{R}}$ from \mathcal{R} and abort if $\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle = 0$.

(b) If $\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle \neq 0$, compute

$$\frac{\mathcal{FE}.\text{Dec}(sk_{\mathbf{x}_{\text{id}}}, C_1)}{\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle} = \frac{\langle \mathbf{x}_{\text{id}}, m \cdot \mathbf{v}_{\mathcal{R}} \rangle}{\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle} = m.$$

A non-revoked user will be able to correctly decrypt this ciphertext with overwhelming probability. On the other hand, a revoked user cannot implement Step (b).

We now consider the (public) tracing procedure. We will show that given an oracle access to a pirate decoder \mathcal{D} and a set $\mathcal{S} = \{\text{id}_1, \text{id}_2, \dots\}$ of suspected traitors with $|\mathcal{S}| \leq t$, it is possible to find an identity id in the set \mathcal{T} of traitors, as long as $\mathcal{T} \subseteq \mathcal{S}$. Here, we assume $\mathcal{R} \cap \mathcal{S} = \emptyset$ for simplicity.

Given a pirate decoder \mathcal{D} , our tracing algorithm first finds a pair of messages m and m' such that \mathcal{D} can distinguish the encryption of m and m' with noticeable probability. As we will show in the main body, such a pair can be found efficiently. Then, the tracing algorithm proceeds as follows. Let us consider a subset of suspect traitors $\mathcal{S}_i = \{\text{id}_i, \text{id}_{i+1}, \dots\}$ for $i = 1, \dots, |\mathcal{S}| + 1$. We then generate a probe ciphertext $C^{\mathcal{S}_i}$ associated to \mathcal{S}_i with the following properties:

- The distribution of $C^{\mathcal{S}}$ corresponds to the normal encryption of m .
- The distribution of C^{\emptyset} corresponds to the normal encryption of m' .
- The probes $C^{\mathcal{S}_{i-1}}$ and $C^{\mathcal{S}_i}$ are indistinguishable without a secret key for id_{i-1} .

The tracing algorithm then estimates the distinguishing advantage of the decoder \mathcal{D} for $C^{\mathcal{S}_{i-1}}$ and $C^{\mathcal{S}_i}$ for all $i \in \{2, \dots, |\mathcal{S}| + 1\}$. It outputs the identity id_{i-1} of the user that is excluded from \mathcal{S}_{i-1} to get \mathcal{S}_i such that the distinguishing advantage between them is non-negligible.

We prove that the tracing algorithm always outputs some user in \mathcal{T} . To see this, we first observe that by the first and second properties above, the decoder \mathcal{D} distinguishes $C^{\mathcal{S}_1} = C^{\mathcal{S}}$ and $C^{\mathcal{S}_{|\mathcal{S}|+1}} = C^{\emptyset}$ with non-negligible advantage. Therefore, by the triangle inequality, there exists at least one index i such that \mathcal{D} distinguishes $C^{\mathcal{S}_{i-1}}$ and $C^{\mathcal{S}_i}$ with non-negligible advantage. By the third property above, the identity id_{i-1} indeed corresponds to a traitor.

The above idea is implemented using inner product functional encryption. To create the probe ciphertext, we first set $\mathbf{v}_{\mathcal{S}} \in \mathbb{Z}_p^\ell$ as follows: If $i = 1$, we set $\mathbf{v}_{\mathcal{S}} = \mathbf{0}$; If $i = |\mathcal{S}| + 1$, we set $\mathbf{v}_{\mathcal{S}_i} = (m' - m) \cdot \mathbf{v}_{\mathcal{R}}$ where $\mathbf{v}_{\mathcal{R}}$ is chosen as in the ordinary encryption algorithm; Otherwise, we set $\mathbf{v}_{\mathcal{S}_i}$ so that

- $\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{S}_i} \rangle = 0$ for every $\text{id} \in \mathcal{S}_i \cup \mathcal{R}$,
- $\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{S}_i} \rangle = (m' - m) \cdot \langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle$ for every $\text{id} \in \mathcal{S}_1 \setminus \mathcal{S}_i$.

Note that this can be implemented only if $r + t < \ell$. We then set the probe ciphertext as follows:

$$C^{\mathcal{S}_i} = (C_1, C_2) = (\mathcal{FE}.\text{Enc}(\text{pk}, \mathbf{v}_{\mathcal{S}_i} + m \cdot \mathbf{v}_{\mathcal{R}}), \mathcal{R}).$$

We will show that by setting the probe ciphertext for tracing as above, we can satisfy the three requirements. By construction, the first and the second requirements are satisfied. To see the third property, we consider the decryption result of the ciphertext using a secret key $\text{sk}_{\mathbf{x}_{\text{id}}}$ for id . We have

$$\frac{\mathcal{FE}.\text{Dec}(\text{sk}_{\mathbf{x}_{\text{id}}}, C_1)}{\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle} = \frac{\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{S}_i} + m \cdot \mathbf{v}_{\mathcal{R}} \rangle}{\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle} = \frac{\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{S}_i} \rangle}{\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle} + m.$$

Therefore, the decryption result of the probe ciphertext $C^{\mathcal{S}_i}$ is m if $\text{id} \in \mathcal{S}_i$ and m' if $\text{id} \in \mathcal{S} \setminus \mathcal{S}_i$. Then we observe that the decryption results of $C^{\mathcal{S}_i}$ and $C^{\mathcal{S}_{i-1}}$ are the same, as long as we use a secret key for $\text{id} \in \mathcal{S} \cup \mathcal{R}$ with $\text{id} \neq \text{id}_{i-1}$. By the security property of inner product functional encryption, this implies that any coalition of users $\subseteq \mathcal{S}$ cannot distinguish two ciphertexts without having $\text{sk}_{\mathbf{x}_{\text{id}_{i-1}}}$. Namely, the third requirement regarding the probe ciphertext also holds.

Our LWE-based IPFE Here, we give the overview of our direct construction of LWE-based IPFE scheme that enjoys improved efficiency compared to Agrawal et al. (2016). Let ℓ and p be the dimension and modulus of the space on which inner-products are taken. Furthermore, let $q = p^k$ be the LWE modulus, where k is some integer. In our scheme, the master secret key is $\mathbf{Z} \in \mathbb{Z}^{\ell \times n}$, chosen from a Gaussian distribution with standard deviation σ . The public key is of the form $\text{pk} = (\mathbf{A} \in \mathbb{Z}_q^{m \times n}, \mathbf{U} = \mathbf{Z}\mathbf{A} \in \mathbb{Z}_q^{\ell \times n})$. To generate a secret key for the vector $\mathbf{x} \in \mathbb{Z}_p^\ell$, we first pick a vector $\bar{\mathbf{x}} \in \mathbb{Z}^\ell$ from a short Gaussian distribution over \mathbb{Z}^ℓ conditioned on $\bar{\mathbf{x}} \equiv \mathbf{x} \pmod{p}$. Then, the secret key is set as $\text{sk}_{\mathbf{x}} = (\bar{\mathbf{x}}^t, \bar{\mathbf{x}}^t \cdot \mathbf{Z})$. One may wonder why do we set $\bar{\mathbf{x}}$ like this instead of just setting $\bar{\mathbf{x}} = \mathbf{x}$. This is because we will use some nice properties of the Gaussian distribution in our security proof, as will be explained later. The ciphertext for a vector $\mathbf{y} \in \mathbb{Z}_p^\ell$ is of the form $(\mathbf{c}_0 \approx \mathbf{A}\mathbf{s}, \mathbf{c}_1 \approx \mathbf{U}\mathbf{s} + p^{k-1} \cdot \mathbf{y})$ where $\mathbf{x} \approx \mathbf{y}$ means that $\|\mathbf{x} - \mathbf{y}\|$ is small.

Here, we skip the explanation of the decryption algorithm and directly go to the intuition for the security proof. We first observe that since all entries of \mathbf{Z} are small, $\mathbf{c}_1 \approx \mathbf{Z}\mathbf{A}\mathbf{s} \approx \mathbf{Z}\mathbf{c}_0$. Given this observation, we can change the distribution of the ciphertext as \mathbf{c}_0 being a random vector $\mathbf{u} \leftarrow \mathbb{Z}_q^\ell$ and $\mathbf{c}_1 \approx \mathbf{Z}\mathbf{u} + p^{k-1} \cdot \mathbf{y}$ without being detected by the adversary, assuming the LWE assumption.

The main difficulty in the proof is in showing that $\mathbf{c}_1 \approx \mathbf{Z}\mathbf{u} + p^{k-1}\mathbf{y}$ does not leak any information *more than necessary*. Note that \mathbf{c}_1 *does* leak some information. Namely, given a secret key $\text{sk}_{\mathbf{x}}$ for \mathbf{x} , we can still decrypt the modified ciphertext to obtain $\langle \mathbf{x}, \mathbf{y} \rangle \pmod{p}$. What we have to prove is that the ciphertext does not leak any information of $\langle \mathbf{x}', \mathbf{y} \rangle \pmod{p}$ for all $\mathbf{x}' \notin \text{Span}_{\mathbb{Z}_p}(\{\mathbf{x}_i\}_{i \in [L]})$, where L is the number of key queries and $\{\mathbf{x}_1, \dots, \mathbf{x}_L\}$ is the set of vectors for which the adversary is given corresponding secret keys.

This will be shown by an information theoretic argument using the fact that certain amount of information on \mathbf{Z} is hidden from the adversary. In particular, we explain that an attempt to obtain any information of $\langle \mathbf{x}', \mathbf{y} \rangle \pmod{p}$ by computing $\langle \mathbf{x}', \mathbf{c}_1 \rangle \approx \mathbf{x}'^t \mathbf{Z}\mathbf{u} + p^{k-1} \cdot \langle \mathbf{x}', \mathbf{y} \rangle \pmod{q}$ fails because $\mathbf{x}'^t \mathbf{Z}$ retains sufficiently high min-entropy and thus $\mathbf{x}'^t \mathbf{Z}\mathbf{u}$ is uniformly random modulo q by the leftover hash lemma.

To see this, let $\mathbf{X}_{\text{top}} \in \mathbb{Z}^{L \times \ell}$ be the matrix obtained by vertically concatenating $\{\bar{\mathbf{x}}_i \in \mathbb{Z}^\ell\}_{i \in [L]}$. Via secret keys, the adversary learns the value of $\mathbf{X}_{\text{top}}\mathbf{Z}$. Let us ignore the additional leakage on \mathbf{Z} from the public key in this overview. Note that in $\mathbf{X}_{\text{top}}\mathbf{Z}$, the matrix \mathbf{X}_{top} acts in parallel on the columns of \mathbf{Z} . We can hence restrict ourselves to the distribution of \mathbf{z}_i conditioned on $\mathbf{b}_i := \mathbf{X}_{\text{top}}\mathbf{z}_i$. It can be seen that \mathbf{z}_i is distributed on the shifted kernel lattice Λ , defined as

$$\Lambda = \{\mathbf{v} \in \mathbb{Z}^m : \mathbf{X}_{\text{top}} \cdot \mathbf{v} = 0\}.$$

If the standard deviation σ is sufficiently large (i.e., larger than the smoothing parameter of Λ), the vector \mathbf{z}_i behaves like the continuous Gaussian even though it is sampled from the discrete Gaussian. In particular, it

spreads all directions under the only constraint that $\mathbf{X}_{top}\mathbf{z}_i = \mathbf{b}_i$, and thus $\langle \mathbf{x}', \mathbf{z}_i \rangle$ has sufficiently high entropy, allowing us to conclude. In Agrawal et al. (2016), the equivalent of \mathbf{X}_{top} in their proof is arbitrarily chosen by the adversary and $\ell = L + 1$. This results in exponentially large smoothing parameter for corresponding Λ . Therefore, they have to take σ exponentially large, which is exactly the source of the inefficiency in their scheme. In our case, the matrix \mathbf{X}_{top} is chosen uniformly at random from a small-width Gaussian distribution. (Recall that in our weakened security definition, the adversary does not have control over \mathbf{x}_i .) Furthermore, we set ℓ large compared to L . We can then invoke the result of Agrawal et al. (2013), which says that the smoothing parameter of Λ corresponding to such \mathbf{X}_{top} is small. This allows us to choose σ much smaller and significantly improve the efficiency.

Organization of the paper. The remainder of the paper is organized as follows. In Section 2, we provide definitions and preliminaries required for our work. In Section 3, we provide our generic construction of trace-and-revoke systems from inner product functional encryption. In Section 4, we provide our new construction of bounded collusion IPFE from LWE and in Section 5 we provide concrete instantiations of trace-and-revoke systems from the DDH and DCR assumptions. We provide a generic transformation from an inner product functional encryption scheme to a traitor tracing scheme that supports multi-message encryption in the full version of this work that is on ePrint Agrawal et al. (2017).

2 Definitions and Preliminaries

Notation. The set $\{1, \dots, n\}$ of natural numbers is denoted by $[n]$. A set is denoted by an uppercase letter. The cardinality of a set X is denoted as $|X|$. If X is finite, we let $U(X)$ denote the uniform distribution over X , and we may write $x \leftarrow X$ to refer to x being sampled from $U(X)$. Vectors will be denoted by bold letters. By default, we treat a vector as a column vector. For two vectors \mathbf{x} and \mathbf{y} , we let $\langle \mathbf{x}, \mathbf{y} \rangle$ denote the canonical inner product between them and $(\mathbf{x} \parallel \mathbf{y})$ denote the vertical concatenation of them. For a positive integer N , we let \mathbb{Z}_N denote the ring of integers with addition and multiplication modulo N . The set of all functions that run in polynomial time is denoted by $\text{Poly}(\cdot)$.

In our scheme descriptions, a user's identifying information is denoted by id . A set of users is thus represented by a set of their respective identifying information. A set of users is denoted by an uppercase calligraphic letter. The set of revoked users is denoted by \mathcal{R} . The set of traitors is denoted by \mathcal{T} and the set of users that are suspected to be traitors is denoted by \mathcal{S} .

In this section, we recall the notions of trace-and-revoke systems and inner product functional encryption.

2.1 Trace-and-Revoke Systems

In a public key traitor tracing encryption scheme, there is a single public key for encryption and many users with decryption capabilities, each having its own unique secret key. Additionally, the encryption scheme provides a feature to identify at least one user from a coalition of malicious users (traitors) that built an unauthorized decryption device \mathcal{D} . Let \mathcal{T} be the set of traitors and we assume that the size $|\mathcal{T}|$ of the traitor coalition is at most t . The tracing algorithm aims at disclosing the identity of at least one user from the set \mathcal{T} of traitors.

In Boneh and Franklin (1999), the minimal black-box access model was considered where the tracing procedure has access to the pirate decryption device \mathcal{D} only through an oracle $\mathcal{O}^{\mathcal{D}}$. The oracle $\mathcal{O}^{\mathcal{D}}$ takes as input any message-ciphertext pair (M, C) and returns 1 if $\mathcal{D}(C) = M$ and 0 otherwise. Hence, it only tells whether the decoder decrypts C to M or not. If the decoder fails to decrypt correctly, the tracing algorithm knows nothing about the decrypted value returned by the decoder. A practical example supporting this assumption is that a pirated media player will only indicate if it is able to play some encrypted media and nothing more about the results of his attempts of decryption.

The decryption device \mathcal{D} is assumed to decrypt correctly with significant probability all messages that have been properly encrypted, as otherwise the decryption device is not very useful. Let \mathcal{R} be any set of revoked users, of cardinality $\leq r$. Let the message m be sampled uniformly at random from the message space \mathcal{M} and let $C^{(\mathcal{R})}$ be the output of the encryption algorithm Enc using the public encryption key pk and \mathcal{R} as the set of revoked users. With $C^{(\mathcal{R})}$ as input, the device \mathcal{D} outputs m with probability significantly more than $1/|\mathcal{M}|$:

$$\Pr_{\substack{m \leftarrow U(\mathcal{M}) \\ C^{(\mathcal{R})} \leftarrow \text{Enc}(\text{pk}, \mathcal{R}, m)}} \left[\mathcal{O}^{\mathcal{D}}(C^{(\mathcal{R})}, m) = 1 \right] \geq \frac{1}{|\mathcal{M}|} + \frac{1}{\lambda^c}, \quad (1)$$

for some constant $c > 0$.² The probability of decryption for a decoder \mathcal{D} can be estimated by repeatedly querying the oracle $\mathcal{O}^{\mathcal{D}}$ with plaintext-ciphertext pairs, using Hoeffding's inequality. Alternatively, we may force the correct decryption probability to be non-negligibly close to 1, by using an all-or-nothing transform (see Kiayias and Yung (2002)). We also assume that the decoder \mathcal{D} is stateless/resettable, i.e., it cannot see and adapt to it being tested, and replies independently to successive queries. Handling stateful pirate boxes has been investigated in Kiayias and Yung (2001b,a).

We let the identity space ID and the message space \mathcal{M} be implicit arguments to the setup algorithm below. We let the secret key space \mathcal{K} and the ciphertext space \mathcal{C} (along with ID and \mathcal{M}) be implicit public parameters output by the setup algorithm.

Definition 1 *A dynamic identity-based trace-and-revoke scheme (t, r) - \mathcal{TR} in black-box confirmation model is a tuple $\mathcal{TR} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec}, \text{Trace})$ of five probabilistic polynomial-time algorithms with the following specifications.*

- $\text{Setup}(1^\lambda, 1^t, 1^r)$ takes as input the security parameter λ , the bound t on the size of traitor coalitions and the bound r on the number of revoked users. It outputs $(\text{msk}, \text{pk}, \text{pd})$ containing the master secret key, the encryption key and the (initially empty) public directory pd . We will implicitly assume that pd is available to all algorithms.
- $\text{KeyGen}(\text{msk}, \text{id})$ takes as input the master secret msk and an identity $\text{id} \in \text{ID}$ of a user, and outputs a secret key sk_{id} and some public information p_{id} for id . It also updates the public directory pd to include p_{id} .³
- $\text{Enc}(\text{pk}, \mathcal{R}, m)$ takes as input the public key pk , a set \mathcal{R} of cardinality $\leq r$ which contains the p_{id} of each revoked user in pd , and a plaintext message $m \in \mathcal{M}$. It outputs a ciphertext $C \in \mathcal{C}$.
- $\text{Dec}(\text{sk}_{\text{id}}, C)$ takes as input a secret key sk_{id} of a user with identity id and a ciphertext $C \in \mathcal{C}$. It outputs a plaintext $m \in \mathcal{M}$.
- $\text{Trace}(\text{pd}, \mathcal{R}, \mathcal{S}, \mathcal{O}^{\mathcal{D}})$ is a black-box confirmation tracing algorithm that takes as input the public directory pd , a set \mathcal{R} of $\leq r$ revoked users, a set \mathcal{S} of $\leq t$ suspect users, and has black-box access to the pirate decoder \mathcal{D} through the oracle $\mathcal{O}^{\mathcal{D}}$. It outputs an identity id or \perp .

The correctness requirement is that, with overwhelming probability over the randomness used by the algorithms, we have:

$$\forall m \in \mathcal{M}, \forall \text{id} \in \text{ID} : \text{Dec}(\text{sk}_{\text{id}}, \text{Enc}(\text{pk}, \mathcal{R}, m)) = m,$$

for any set \mathcal{R} of $\leq r$ revoked users and for any id such that $\text{id} \notin \mathcal{R}$.

²In Nishimaki et al. (2016), a weaker notion of usefulness is considered (leading to a better security guarantee): the box is considered useful if it distinguishes between encryptions of two adversarially chosen plaintexts. We note that our security proof actually handles this weaker usefulness. In fact, we show in Lemma 8 that the notion of usefulness given here implies that it is possible to efficiently find two plaintexts whose ciphertext distributions can be distinguished by the decryption box. The rest of the security proof carries over in an identical way for both usefulness notions.

³We emphasize that p_{id} does not need to contain id .

Public Traceability. It is required that, when \mathcal{S} contains the set \mathcal{T} of traitors who produced the pirate decoder \mathcal{D} , then the id output by the tracing algorithm belongs to \mathcal{T} . This requirement is formalized using the following game, denoted by AD-TT, between an adversary \mathcal{A} and a challenger:

- The challenger runs $\text{Setup}(1^\lambda, 1^t, 1^r)$ and gives pk to \mathcal{A} .
- Adversary \mathcal{A} may ask the challenger to add polynomially many users in the system. Adversary \mathcal{A} may choose the id 's of the users, but does not obtain the corresponding sk_{id} . Nevertheless, the public directory pd is updated accordingly.
- Adversary \mathcal{A} is allowed to make up to t arbitrary traitor key queries. It may observe the database pd to choose its queries in an adaptive way. If it queries $\text{id} \in \text{ID}$ to the challenger, then:
 - If the key for id was previously generated, i.e., if p_{id} is found in the database pd , then the challenger responds with sk_{id} . The challenger records the identity query id in a list \mathcal{T} .
 - Otherwise (i.e., user id is a new user in the system), the challenger runs $\text{KeyGen}(\text{msk}, \text{id})$, responds with sk_{id} and updates the directory pd with the public information p_{id} for id . The challenger also records the identity query id in the list \mathcal{T} .
- Adversary \mathcal{A} is allowed to (adaptively) choose a set \mathcal{R} of up to r revoked users in pd . The challenger gives \mathcal{A} all the corresponding sk_{id} . These queries can be interleaved with extensions of the number of users and user corruption queries, in an adaptive manner.
- Adversary \mathcal{A} finally produces a pirate decoder \mathcal{D} . It chooses a suspect set \mathcal{S} of cardinality $\leq t$ that contains \mathcal{T} , and sends \mathcal{S} to the challenger.
- The challenger then runs $\text{Trace}(\text{pd}, \mathcal{R}, \mathcal{S}, \mathcal{O}^{\mathcal{D}})$. The adversary wins if both of the following hold:
 - Equation (1) is satisfied for the set of revoked users \mathcal{R} chosen by the adversary (i.e., decoder \mathcal{D} is useful),
 - the execution of Trace outputs \perp or outputs an id that does not belong to \mathcal{T} with probability $\geq 1/\lambda^c$.

No probabilistic polynomial-time adversary \mathcal{A} should be able to win game AD-TT with non-negligible probability.

Almost Public Traceability. This is the same as public traceability, except that Trace only outputs the associated information about the traitors instead of their identities, namely p_{id} instead of id . Consequently, the second winning condition of the adversary should be adapted so that it only requires the execution of Trace to output a p_{id} that does not belong to $\text{pd}_{\mathcal{T}}$, which is the set of all $p_{\text{id}'}$ for $\text{id}' \in \mathcal{T}$.

This restriction does not change much the functionality of the tracing because, from p_{id} , the authority can immediately map back to id and the authority can still delegate the tracing procedure to untrusted parties. On the other side, this variant may be useful in practice as we do not leak the information of users in the public directory.

We note that our proposed schemes satisfy the public traceability instead of the almost public traceability. However, it is easy to modify them so that they satisfy the latter. Hereafter, we will not discuss about almost public traceability.

Traitor Tracing Scheme. A traitor tracing scheme is simply a trace-and-revoke scheme without the capacity of revoking users. It corresponds to the above definition where the revoked set is always set to be empty, in the encryption as well as in the security game.

Semantic Security. The IND-CPA security of a trace-and-revoke scheme \mathcal{TR} is defined based on the following game.

- The challenger runs $\text{Setup}(1^\lambda, 1^\ell, 1^r)$ and gives the produced public key pk to the adversary \mathcal{A} . The adversary may ask the challenger to add polynomially many users in the system.
- The adversary (adaptively) chooses a set \mathcal{R} of $\leq r$ revoked users in pd . The challenger gives \mathcal{A} all the sk_{id} such that $p_{\text{id}} \in \mathcal{R}$.
- The adversary then chooses two messages $m_0, m_1 \in \mathcal{M}$ of equal length and gives them to the challenger.
- The challenger samples $b \leftarrow \{0, 1\}$ and provides $C_{m_b} \leftarrow \text{Enc}(\text{pk}, \mathcal{R}, m_b)$ to \mathcal{A} .
- Finally, the adversary returns its guess $b' \in \{0, 1\}$ for the b chosen by the challenger. The adversary wins this game if $b = b'$.

The advantage of the adversary is defined as $\text{Adv}_{\mathcal{TR}, \mathcal{A}}^{\text{IND-CPA}} = |\Pr[b = b'] - 1/2|$. The scheme \mathcal{TR} is said semantically secure if there is no probabilistic polynomial-time adversary \mathcal{A} that wins this game with non-negligible advantage.

2.2 Inner Product Functional Encryption

In this section, we define functional encryption for the functionality of inner products over \mathbb{Z}_p .

Definition 2 A functional encryption scheme \mathcal{FE} for the inner product functionality over \mathbb{Z}_p is a tuple $\mathcal{FE} = (\mathcal{FE}.\text{Setup}, \mathcal{FE}.\text{KeyGen}, \mathcal{FE}.\text{Enc}, \mathcal{FE}.\text{Dec})$ of four probabilistic polynomial-time algorithms with the following specifications:

- $\mathcal{FE}.\text{Setup}(1^\lambda, 1^\ell)$ takes as input the security parameter λ and outputs the public key and the master secret key pair (pk, msk) ;
- $\mathcal{FE}.\text{KeyGen}(\text{msk}, \mathbf{x})$ takes as input the master secret key msk and a vector $\mathbf{x} \in \mathbb{Z}_p^\ell$ and outputs the secret key $\text{sk}_{\mathbf{x}}$;
- $\mathcal{FE}.\text{Enc}(\text{pk}, \mathbf{y})$ takes as input the public key pk and a message $\mathbf{y} \in \mathbb{Z}_p^\ell$ and outputs the ciphertext $\text{ct}_{\mathbf{y}}$;
- $\mathcal{FE}.\text{Dec}(\text{sk}_{\mathbf{x}}, \text{ct}_{\mathbf{y}})$ takes as input the secret key of a user $\text{sk}_{\mathbf{x}}$ and the ciphertext $\text{ct}_{\mathbf{y}}$, and outputs an element from $\mathbb{Z}_p \cup \{\perp\}$.

The correctness requirement is that, with overwhelming probability over the randomness used by the algorithms, for $(\text{pk}, \text{msk}) \leftarrow \mathcal{FE}.\text{Setup}(1^\lambda, 1^\ell)$ and $\forall \mathbf{x}, \mathbf{y} \in \mathbb{Z}_p^\ell$:

$$\mathcal{FE}.\text{Dec}(\mathcal{FE}.\text{KeyGen}(\text{msk}, \mathbf{x}), \mathcal{FE}.\text{Enc}(\text{pk}, \mathbf{y})) = \langle \mathbf{x}, \mathbf{y} \rangle \bmod p.$$

Security of \mathcal{FE} . We consider security of functional encryption in the standard indistinguishability setting Boneh et al. (2011).

Definition 3 A functional encryption scheme $\mathcal{FE} = (\mathcal{FE}.\text{Setup}, \mathcal{FE}.\text{KeyGen}, \mathcal{FE}.\text{Enc}, \mathcal{FE}.\text{Dec})$ provides semantic security under chosen-plaintext attacks (or IND-CPA security) if no probabilistic polynomial-time adversary \mathcal{A} has non-negligible advantage in the following game:

- The challenger runs $\mathcal{FE}.\text{Setup}(1^\lambda, 1^\ell)$ and the master public key mpk is given to \mathcal{A} .
- The adversary adaptively makes secret key queries to the challenger. At each query, adversary \mathcal{A} chooses a vector $\mathbf{x} \in \mathbb{Z}_p^\ell$ and obtains the corresponding secret key $\text{sk}_{\mathbf{x}} \leftarrow \mathcal{FE}.\text{KeyGen}(\text{msk}, \mathbf{x})$.

- Adversary \mathcal{A} chooses distinct messages $\mathbf{y}_0, \mathbf{y}_1 \in \mathbb{Z}_p^\ell$ subject to the restriction that, for every vector \mathbf{x} queried in the previous step, it holds that $\langle \mathbf{x}, \mathbf{y}_0 \rangle = \langle \mathbf{x}, \mathbf{y}_1 \rangle \pmod p$ and sends them to the challenger. In response, the challenger samples $b \leftarrow \{0, 1\}$ and sends $\text{ct}^* \leftarrow \mathcal{FE}.\text{Enc}(\text{pk}, \mathbf{y}_b)$ to \mathcal{A} .
- Adversary \mathcal{A} makes further secret key queries for arbitrary vectors $\mathbf{x} \in \mathbb{Z}_p^\ell$ of its choice. As before, it is required that $\langle \mathbf{x}, \mathbf{y}_0 \rangle = \langle \mathbf{x}, \mathbf{y}_1 \rangle \pmod p$ for each query \mathbf{x} made by \mathcal{A} .
- Adversary \mathcal{A} eventually outputs a bit $b' \in \{0, 1\}$ and wins if $b' = b$.

The adversary's advantage is defined as $\text{Adv}_{\mathcal{A}}(\lambda) := |\Pr[b' = b] - 1/2|$.

The Random-Key Bounded-Collusion Model. In bounded collusion functional encryption Gorbunov et al. (2012), the adversary \mathcal{A} is restricted to ask at most Q secret key queries for some fixed polynomial Q , which is input to the setup algorithm. Additionally, our application permits an additional weakening of the security model for inner product functional encryption: we are only required to show security against an adversary who first sees arbitrarily many *random* vectors $\mathbf{x} \leftarrow \mathbb{Z}_p^\ell$, requests secret keys for an adaptively chose subset of them, and does not make secret key queries after it gets the challenge ciphertext. The above definition of security against such a restricted adversary will be called Q -IND-CPA.

2.3 Lattice background

A lattice Λ is a (non-zero) discrete subgroup of \mathbb{R}^m . A *basis* of Λ is a linearly independent set of vectors whose \mathbb{Z} -span is Λ . We recall that the *smoothing parameter* of Λ is defined as

$$\eta_\varepsilon(\Lambda) = \min \left(\sigma > 0 : \sum_{\widehat{\mathbf{b}} \in \widehat{\Lambda}} \exp(-\pi \|\widehat{\mathbf{b}}\|^2 / \sigma^2) \leq 1 + \varepsilon \right),$$

where $\widehat{\Lambda} = \{\widehat{\mathbf{b}} \in \text{Span}_{\mathbb{R}}(\Lambda) : \widehat{\mathbf{b}}^T \cdot \Lambda \subseteq \mathbb{Z}\}$ refers to the dual of Λ . Note that if $\sigma = \Omega(\sqrt{\lambda})$, we have that there exists $\varepsilon = 2^{-\Omega(\lambda)}$ such that $\sigma \geq \eta_\varepsilon(\mathbb{Z})$.

For a lattice $\Lambda \subseteq \mathbb{R}^m$, a vector $\mathbf{c} \in \mathbb{R}^m$, and an invertible $\Sigma \in \mathbb{R}^{m \times m}$, we define the Gaussian distribution of parameter Λ , \mathbf{c} , and Σ by $D_{\Lambda, \Sigma, \mathbf{c}}(\mathbf{b}) \sim \rho_{\Sigma, \mathbf{c}}(\mathbf{b}) = \exp(-\pi \|\Sigma^{-1}(\mathbf{b} - \mathbf{c})\|^2)$ for all $\mathbf{b} \in \Lambda$. When $\Sigma = \sigma \mathbf{I}_m$, we simply write $D_{\Lambda, \sigma, \mathbf{c}}$. Sometimes, for convenience, we use the notation $D_{\Lambda + \mathbf{c}, \Sigma}$ as a shorthand for $\mathbf{c} + D_{\Lambda, \Sigma, -\mathbf{c}}$.

For $m \geq n$ and a rank- n matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$, denote $U_{\mathbf{X}} = \{\|\mathbf{X}\mathbf{u}\| : \mathbf{u} \in \mathbb{R}^n, \|\mathbf{u}\| = 1\}$. The least singular value of \mathbf{X} is then defined as $s_n(\mathbf{X}) := \inf(U_{\mathbf{X}})$ and similarly the largest singular value of \mathbf{X} is $s_1(\mathbf{X}) := \sup(U_{\mathbf{X}})$. For a matrix $\mathbf{Y} \in \mathbb{R}^{n' \times m'}$ with $n' > m'$, the least singular value and the largest singular value are defined as $s_1(\mathbf{Y}) := s_1(\mathbf{Y}^t)$ and $s_{m'}(\mathbf{Y}) := s_{m'}(\mathbf{Y}^t)$ respectively.

For the rest of this section, we assume that lattices are full-rank, i.e., the dimensions of the span and the ambient space match.

Lemma 1 (Corollary 2.8 in Gentry et al. (2008)) *Let $\Lambda' \subseteq \Lambda \subseteq \mathbb{R}^m$ be two lattices with the same dimension. Let $\varepsilon \in (0, 1/2)$. Then for any $\mathbf{c} \in \mathbb{R}^m$ and any Σ such that $s_m(\Sigma) \geq \eta_\varepsilon(\Lambda')$, the distribution $D_{\Lambda, \Sigma, \mathbf{c}} \pmod{\Lambda'}$ is within statistical distance 2ε from the uniform distribution over Λ/Λ' .*

Lemma 2 (Lemma 1 in Katsumata and Yamada (2016)) *Let $r \geq \Omega(\sqrt{\lambda})$ and $q, \ell, m > 0$ integers. Let $\mathbf{b} \in \mathbb{Z}_q^m$ be arbitrary and \mathbf{x} chosen from $D_{\mathbb{Z}^m, r}$. Then for any $\mathbf{V} \in \mathbb{Z}^{\ell \times m}$ and positive real $r' > s_1(\mathbf{V})$, there exists a probabilistic polynomial-time algorithm $\text{ReRand}(\mathbf{V}, \mathbf{b} + \mathbf{x}, r, r')$ that outputs $\mathbf{b}' = \mathbf{V}\mathbf{b} + \mathbf{x}' \in \mathbb{Z}_q^\ell$ where \mathbf{x}' is within statistical distance $2^{-\Omega(\lambda)}$ from $D_{\mathbb{Z}^\ell, 2rr'}$.*

We use the following variant of the leftover hash lemma, adapted from Micciancio and Mol (2011) (see also Lemma 11 in Agrawal et al. (2016)).

Lemma 3 (Micciancio and Mol (2011)) Let $m \geq n \geq 1$ and $q = p^k$ for p prime and $k \geq 1$. Take \mathcal{X} a distribution over \mathbb{Z}^m . Let D_0 be a uniform distribution over $\mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^n$ and D_1 be the distribution of $(\mathbf{A}, \mathbf{A} \cdot \mathbf{x})$, where sampling $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$ and $\mathbf{x} \leftarrow \mathcal{X}$. Then,

$$\Delta(D_0, D_1) \leq \frac{1}{2} \sqrt{\sum_{i=1}^k p^{i \cdot n} \cdot \text{Pr}_i}.$$

where Pr_i is the collision probability of two independent samples from $(\mathcal{X} \bmod p^i)$.

The above lemma implies that if the distribution $(\mathcal{X} \bmod p)$ is within statistical distance ε from the uniform distribution over \mathbb{Z}_p^m , then we have

$$\Delta(D_0, D_1) \leq \varepsilon + \sqrt{q^n/p^m}.$$

This can be seen by considering a distribution \mathcal{X}' such that $(\mathcal{X} \bmod p)$ is uniform distribution over \mathbb{Z}_p^m and $\Delta(\mathcal{X}, \mathcal{X}') \leq \varepsilon$.

Lemma 4 (Special case of Lemma 8 in Agrawal et al. (2016)) There exists a universal constant $K > 1$ such that for all $m \geq 2n$, $\varepsilon > 0$ and $\sigma \geq K\eta_\varepsilon(\mathbb{Z})$, the following holds for $\mathbf{X} \leftarrow D_{\mathbb{Z}, \sigma}^{n \times m}$:

$$\begin{aligned} \Pr [\sigma\sqrt{2\pi m}/K < s_n(\mathbf{X}) \leq s_1(\mathbf{X}) < \sigma K\sqrt{2\pi m}] \\ > 1 - 4m\varepsilon + O(\exp(-m/K)). \end{aligned}$$

We will also require the following theorem, adapted from Theorem 17 in Ling et al. (2014).

Theorem 1 (Ling et al. (2014)) Let n , m_1 , m_2 , and λ be integers satisfying $m_2 \geq m_1 > 100$ and σ_1, σ_2 be positive real numbers. Let $n' = \max\{\lambda, n\}$ and assume that $n' > 100$. We also assume that they satisfy $\sigma_1 \geq \Omega(\sqrt{m_1 n' \log m_1})$, $m_1 \geq \Omega(n' \log(\sigma_1 n'))$, and $\sigma_2 \geq \Omega(n'^{5/2} \sqrt{m_1} \sigma_1^2 \log^{3/2}(m_1 \sigma_1))$. Then, there exists a probabilistic polynomial-time algorithm that given n , m_1 , m_2 , λ (in unary), σ_1 , and σ_2 , returns $\mathbf{X}_1 \in \mathbb{Z}^{n \times m_1}$, $\mathbf{X}_2 \in \mathbb{Z}^{n \times m_2}$, and $\mathbf{U} \in \mathbb{Z}^{m \times m}$ with $m = m_1 + m_2$ such that:

- the distribution of $(\mathbf{X}_1, \mathbf{X}_2)$ is within statistical distance $2^{-\Omega(n')}$ of the distribution $D_{\mathbb{Z}, \sigma_1}^{n \times m_1} \times (D_{\mathbb{Z}^{m_2}, \sigma_2, \delta_1} \times \dots \times D_{\mathbb{Z}^{m_2}, \sigma_2, \delta_n})^t$, where δ_i denotes the i th canonical unit vector in \mathbb{Z}^{m_2} whose i th coordinate is 1 and whose remaining coordinates are 0,
- we have $|\det \mathbf{U}| = 1$ and $(\mathbf{X}_1 | \mathbf{X}_2) \cdot \mathbf{U} = (\mathbf{I}_n | \mathbf{0})$,
- every column of \mathbf{U} has norm $\leq O(\sqrt{n' m_1} \sigma_2)$ with probability $\geq 1 - 2^{-\Omega(n')}$.

Three remarks are in order regarding the theorem. First, we take the transpose of the theorem in Ling et al. (2014). This is just for a notational convenience. Secondly, the distribution of $\mathbf{X} = (\mathbf{X}_1 | \mathbf{X}_2)$ in Theorem 17 in Ling et al. (2014) is slightly different from the above in that all entries of the first column of \mathbf{X} equal to 1. As noted right after Lemma 7 in Ling et al. (2014), the theorem still holds even with the change. Finally, in the above theorem, we introduce the statistical security parameter λ and differentiate it from the lattice dimension n , while the theorem in Ling et al. (2014) assigns the same variable n for both. This change is introduced because we will invoke the theorem for possibly small n for which 2^{-n} is no longer negligible.

In our security analysis, we need a variant of the above theorem where \mathbf{X} is chosen from a slightly different distribution and \mathbf{U} need not be efficiently samplable.

Lemma 5 Let n , m_1 , m_2 , m , λ , n' , σ_1 , σ_2 be as in Theorem 1. Then, for all but $2^{-\Omega(n')}$ probability over $(\mathbf{X}_1, \mathbf{X}_2) \in \mathbb{Z}^{n \times m_1} \times \mathbb{Z}^{n \times m_2}$ chosen from $D_{\mathbb{Z}, \sigma_1}^{n \times m_1} \times D_{\mathbb{Z}, \sigma_2}^{n \times m_2}$, there exists $\mathbf{U} \in \mathbb{Z}^{m \times m}$ such that $|\det \mathbf{U}| = 1$, $(\mathbf{X}_1 | \mathbf{X}_2) \cdot \mathbf{U} = (\mathbf{I}_n | \mathbf{0})$, and every column of \mathbf{U} has norm $\leq O(\sqrt{n' m_1} \sigma_2)$.

To prepare for the proof of Lemma 5, we define Rényi Divergence (RD) and review its properties following Bai et al. (2015). For any two probability distributions P and Q such that the support of P is a subset of the support of Q over a countable domain X , we define the RD (of order 2) by $R(P\|Q) = \sum_{x \in X} P(x)^2/Q(x)$, with the convention that the fraction is zero when both the numerator and denominator are zero. We will use the following property: if P (resp. Q) is a direct product of independent distributions P_1 and P_2 (resp. Q_1 and Q_2), then we have $RD(P\|Q) = RD(P_1 \times P_2\|Q_1 \times Q_2) = RD(P_1\|Q_1) \cdot RD(P_2\|Q_2)$.

Lemma 6 (Lemma 2.9 in Bai et al. (2015)) *Let P and Q denote distributions with $\text{Supp}(P) \subseteq \text{Supp}(Q)$ and $A \subseteq \text{Supp}(Q)$ be arbitrary set. Then, we have $Q(A) \geq P(A)^2/R(P\|Q)$ where $P(A)$ and $Q(A)$ are measure of A under the distribution P and Q respectively.*

We also recall that the RD between two offset discrete Gaussians is bounded as follows.

Lemma 7 (Lemma 4.2 in Langlois et al. (2014)) *For any n -dimensional lattice $L \subseteq \mathbb{R}^n$ and invertible matrix Σ , set $P = D_{\Lambda, \Sigma, \mathbf{w}}$ and $Q = D_{\Lambda, \Sigma, \mathbf{z}}$ for some fixed $\mathbf{w}, \mathbf{z} \in \Lambda$. Then, $R(P\|Q) \leq \exp(2\pi\|\mathbf{w} - \mathbf{z}\|^2/s_n(\Sigma)^2)$.*

Then, we proceed to the proof of Lemma 5.

Proof: [Proof of Lemma 5] Let $A \subseteq \mathbb{Z}^{n \times m}$ be the set of $\mathbf{X} = (\mathbf{X}_1|\mathbf{X}_2)$ such that \mathbf{U} satisfying the properties listed in the statement does not exist. Theorem 1 implies that when \mathbf{X} is sampled from the distribution $Q := D_{\mathbb{Z}, \sigma_1}^{m_1 \times m_2} \times (D_{\mathbb{Z}^{m_2}, \sigma_2, \delta_1} \times \cdots \times D_{\mathbb{Z}^{m_2}, \sigma_2, \delta_n})^t$, we have $Q(A) \leq 2^{-\Omega(n')}$. We want to prove that $P(A) = 2^{-\Omega(n')}$ for the distribution $P := D_{\mathbb{Z}, \sigma_1}^{n \times m_1} \times D_{\mathbb{Z}, \sigma_2}^{n \times m_2}$. By Lemma 6, we have $P(A) \leq \sqrt{Q(A) \cdot R(P\|Q)} \leq \sqrt{R(P\|Q)} \cdot 2^{-\Omega(n')}$. To complete the proof, it suffices to show $R(P\|Q) = O(1)$. We have

$$\begin{aligned} R(P\|Q) &= R(D_{\mathbb{Z}, \sigma_1}^{n \times m_1} \times D_{\mathbb{Z}, \sigma_2}^{n \times m_2} \| D_{\mathbb{Z}, \sigma_1}^{m_1 \times m_2} \\ &\quad \times (D_{\mathbb{Z}^{m_2}, \sigma_2, \delta_1} \times \cdots \times D_{\mathbb{Z}^{m_2}, \sigma_2, \delta_n})) \\ &= R((D_{\mathbb{Z}, \sigma_2})^n \| (D_{\mathbb{Z}, \sigma_2, 1})^n) \\ &\leq \exp(2\pi n/\sigma_2^2), \end{aligned}$$

where we use Lemma 7 in the last inequality. Since $\sigma_2 \geq \Omega(n^{1/2})$, we conclude that $R(P\|Q) = O(1)$. This completes the proof of Lemma 5. \square

Next, we define the learning with errors (LWE) assumption. It was shown that the assumption holds as long as certain lattice problems are hard in the worst case Regev (2005); Peikert (2009); Brakerski et al. (2013).

Definition 4 *For an integers $n = n(\lambda)$, $m = m(\lambda)$, $q = q(\lambda)$, a real number $\alpha(\lambda) \in (0, 1)$, and an algorithm \mathcal{A} , the advantage for the learning with errors problem $\text{LWE}_{n, m, q, \alpha}$ of \mathcal{A} is defined as follows:*

$$|\Pr[\mathcal{A}(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{x}) \rightarrow 1] - \Pr[\mathcal{A}(\mathbf{A}, \mathbf{w} + \mathbf{x}) \rightarrow 1]|$$

where $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{w} \leftarrow \mathbb{Z}_q^m$, and $\mathbf{x} \leftarrow D_{\mathbb{Z}, \alpha q}^m$. We say that $\text{LWE}_{n, m, q, \alpha}$ assumption holds if the advantage is negligible for every probabilistic polynomial-time \mathcal{A} .

3 Trace and Revoke from Inner-Product Functional Encryption

In this section, we provide a generic transformation from a bounded collusion, random keys inner-product functional encryption scheme \mathcal{FE} to a trace-and-revoke scheme \mathcal{TR} . Since intuition was provided in Section 1, we proceed directly to the formal construction.

3.1 The Scheme

We construct a trace-and-revoke scheme \mathcal{TR} following the specifications of Definition 1. Our scheme assumes the existence of a public directory pd which contains the identities of the users that have been assigned keys in the system. The public directory is initially empty. We assume that pd can only be modified by a central authority (the key generator).

1. $\text{Setup}(1^\lambda, 1^t, 1^r)$. Upon input the security parameter λ , the bound t on the number of traitors and the bound r on the number of revoked users, proceed as follows:
 - (a) Let $(\text{pk}, \text{msk}) \leftarrow \mathcal{FE}.\text{Setup}(1^\lambda, 1^\ell)$, where $\ell = t + r + 1$.
 - (b) Output the public key pk and master secret key msk .
2. $\text{KeyGen}(\text{msk}, \text{id})$. Upon input the master secret key msk and a user identity $\text{id} \in \text{ID}$, proceed as follows:
 - (a) Sample $\mathbf{x}_{\text{id}} \leftarrow \mathbb{Z}_p^\ell$. The pair $p_{\text{id}} = (\text{id}, \mathbf{x}_{\text{id}})$ is appended to the public directory pd .
 - (b) Let $\text{sk}_{\text{id}} \leftarrow \mathcal{FE}.\text{KeyGen}(\text{msk}, \mathbf{x}_{\text{id}})$.
 - (c) Output sk_{id} .
3. $\text{Enc}(\text{pd}, \text{pk}, \mathcal{R}, m)$. Upon input the public key pk , a set of revoked users \mathcal{R} of cardinality $\leq r$ and a plaintext messages $m \in \mathcal{M} = \mathbb{Z}_p$, proceed as follows:
 - (a) Compute $\mathbf{v}_{\mathcal{R}} \in \mathbb{Z}_p^\ell \setminus \{\vec{0}\}$ such that $\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle = 0$ for every $\text{id} \in \mathcal{R}$.
 - (b) Compute $\mathbf{y}_{\mathcal{R}} = m \cdot \mathbf{v}_{\mathcal{R}}$.
 - (c) Output $C = (C_1, C_2) = (\mathcal{FE}.\text{Enc}(\text{pk}, \mathbf{y}_{\mathcal{R}}), \mathcal{R})$.
4. $\text{Dec}(\text{pd}, \text{sk}_{\text{id}}, C)$. Upon input the secret key sk_{id} for user id and a ciphertext $C = (C_1, C_2)$, proceed as follows:
 - (a) Parse C_2 as $C_2 = \mathcal{R}$. If $\text{id} \in \mathcal{R}$, then abort.
 - (b) Compute $\mathbf{v}_{\mathcal{R}} \in \mathbb{Z}_p^\ell \setminus \{\vec{0}\}$ such that $\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle = 0$ for every $\text{id} \in \mathcal{R}$.
 - (c) Compute and output $m = \mathcal{FE}.\text{Dec}(\text{sk}_{\text{id}}, C_1) / \langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle$.
5. $\text{Trace}(\text{pd}, \mathcal{R}, \mathcal{S}, \mathcal{O}^{\mathcal{D}})$. Upon input the public directory pd , a revoked set of users \mathcal{R} , a suspect set \mathcal{S} of users and given access to the oracle $\mathcal{O}^{\mathcal{D}}$, first proceed as follows:
 - (a) Find $m, m' \in \mathcal{M}$ such that the following quantity is non-negligible:
$$\left| \Pr_{C \leftarrow \text{Enc}(\text{pd}, \text{pk}, \mathcal{R}, m)} [\mathcal{O}^{\mathcal{D}}(C, m) = 1] - \Pr_{C' \leftarrow \text{Enc}(\text{pd}, \text{pk}, \mathcal{R}, m')} [\mathcal{O}^{\mathcal{D}}(C', m) = 1] \right|.$$

(b) Set $\mathcal{S}_1 = \{\text{id}_1, \text{id}_2, \dots\} = \mathcal{S} \setminus \mathcal{R}$.

(c) Compute $\mathbf{v}_{\mathcal{R}} \in \mathbb{Z}_p^\ell \setminus \{\vec{0}\}$ such that $\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle = 0$ for every $\text{id} \in \mathcal{R}$.

Then execute the following steps with $i = 1, 2, \dots$:

(d) If $i = 1$, set $\mathbf{v}_{\mathcal{S}_i} = \vec{0}$. If $\mathcal{S}_i = \emptyset$, set $\mathbf{v}_{\mathcal{S}_i} = (m' - m) \cdot \mathbf{v}_{\mathcal{R}}$. Else compute $\mathbf{v}_{\mathcal{S}_i} \in \mathbb{Z}_p^\ell$ such that:

- i. $\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{S}_i} \rangle = 0$ for every $\text{id} \in \mathcal{S}_i \cup \mathcal{R}$.

- ii. $\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{S}_i} \rangle = (m' - m) \cdot \langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle$ for every $\text{id} \in \mathcal{S}_1 \setminus \mathcal{S}_i$.
- (e) Repeat the following steps sufficiently many times (as dictated by Hoeffding's inequality) to compute an approximation of the probability p_i that the response from $\mathcal{O}^{\mathcal{D}}$ is $b_i = 1$.
 - i. Construct $\mathbf{y} = \mathbf{v}_{\mathcal{S}_i} + m \cdot \mathbf{v}_{\mathcal{R}} \in \mathbb{Z}_p^\ell$;
 - ii. The probe ciphertext is $C^{\mathcal{S}_i} = (\mathcal{FE}.\text{Enc}(\text{pk}, \mathbf{y}), \mathcal{R})$;
 - iii. Provide the oracle $\mathcal{O}^{\mathcal{D}}$ with $(C^{\mathcal{S}_i}, m)$ as input and get a binary value b_i as output.
- (f) If $i > 1$ and $|p_i - p_{i-1}|$ is non-negligible, then output id_{i-1} and abort;
- (g) If $\mathcal{S}_i = \emptyset$, then output \perp and abort; else, set $\mathcal{S}_{i+1} = \mathcal{S}_i \setminus \{\text{id}_i\}$.

For the correctness and the tracing security proof, we require that in Step (a) of Algorithm Enc, in Step (b) of Algorithm Dec and in Step (c) of Algorithm Trace, the vector $\mathbf{v}_{\mathcal{R}}$ be uniquely determined by \mathcal{R} , in the same unique way across all algorithms. One way of achieving this property is to order the \mathbf{x}_{id} 's for $\text{id} \in \mathcal{R}$ lexicographically, and run a deterministic linear system solver. We proceed in the same way (using always the same deterministic algorithm) for vector $\mathbf{v}_{\mathcal{S}_i}$ at Step (d) of Algorithm Trace.

We remark that one can send \mathcal{R} instead of $\mathbf{v}_{\mathcal{R}}$ in the encryption algorithm. This will make the ciphertext longer, but make the encryption and decryption algorithms slightly more efficient.

We first check the correctness of the scheme.

Theorem 2 *Assume that $p = \lambda^{\omega(1)}$. Let \mathcal{R} be a set of revoked users of cardinality $\leq r$. Then, for every $\text{id} \notin \mathcal{R}$ and every $m \in \mathcal{M} = \mathbb{Z}_p$, we have*

$$\text{Dec}(\text{pd}, \text{sk}_{\text{id}}, \text{Enc}(\text{pd}, \text{pk}, \mathcal{R}, m)) = m,$$

with probability $\geq 1 - \lambda^{-\omega(1)}$.

Proof: As \mathbf{x}_{id} is uniform in \mathbb{Z}_p^ℓ , and thanks to the parameter choices of $p = \lambda^{\omega(1)}$ and $\ell > r$, we have that $\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle \neq 0$, with overwhelming probability. The execution of $\text{Dec}(\text{pd}, \text{sk}_{\text{id}}, C)$, with $C = (C_1, C_2) = \text{Enc}(\text{pd}, \text{pk}, \mathcal{R}, m)$, proceeds to Step (b) and computes (with overwhelming probability):

$$\text{Dec}(\text{pd}, \text{sk}_{\text{id}}, C) = \frac{\mathcal{FE}.\text{Dec}(\text{sk}_{\text{id}}, C_1)}{\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle} = \frac{\langle \mathbf{x}_{\text{id}}, m \cdot \mathbf{v}_{\mathcal{R}} \rangle}{\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle} = m,$$

by correctness of \mathcal{FE} . □

Now, we consider the implementation of Step (a) of Algorithm Trace. The aim is to find $m, m' \in \mathbb{Z}_p$ such that an encryption of m has a non-negligible probability difference of decrypting to m and m' via $\mathcal{O}^{\mathcal{D}}$. These plaintexts are used for tracing as follows: the first probe ciphertext distribution will be a genuine encryption of m , while the last probe ciphertext distribution will be a genuine encryption of m' . (To see this, observe that for the last probe ciphertext, we have $\mathcal{S}_i = \emptyset$ and $\mathbf{v}_{\mathcal{S}_i} = (m' - m) \cdot \mathbf{v}_{\mathcal{R}}$. Consequently, we have $C^{\mathcal{S}_i} = (\mathcal{FE}.\text{Enc}(\text{pk}, \mathbf{y}), \mathcal{R})$ where $\mathbf{y} = \mathbf{v}_{\mathcal{S}_i} + m \cdot \mathbf{v}_{\mathcal{R}} = m' \cdot \mathbf{v}_{\mathcal{R}}$.) The fact that $\mathcal{O}^{\mathcal{D}}$ behaves differently for these two distributions ensures that there will be an i such that $|p_i - p_{i-1}|$ is non-negligible. Now, if the oracle $\mathcal{O}^{\mathcal{D}}$ was perfect, i.e., a genuine encryption of m always decrypts to m for all m , then the existence of a pair (m, m') as in Step (a) would be immediate. The difficulty is that the oracle $\mathcal{O}^{\mathcal{D}}$ only achieves correct decryption with non-negligible advantage.

Lemma 8 *Let \mathcal{R} be arbitrary and assume that Equation (1) holds for \mathcal{R} . Then, with probability $\geq 1/(4\lambda^c)$ over the choice of $m, m' \leftarrow \mathcal{M}$, we have:*

$$\left| \Pr_{C \leftarrow \text{Enc}(\text{pk}, \mathcal{R}, m)} [\mathcal{O}^{\mathcal{D}}(C, m) = 1] - \Pr_{C' \leftarrow \text{Enc}(\text{pk}, \mathcal{R}, m')} [\mathcal{O}^{\mathcal{D}}(C', m) = 1] \right| \geq \frac{1}{2\lambda^c}.$$

Based on Lemma 8, Step (a) of Algorithm `Trace` can be implemented by repeatedly sampling $m, m' \leftarrow \mathcal{M}$ and estimating the probabilities that $\mathcal{O}^{\mathcal{D}}(C, m) = 1$ and $\mathcal{O}^{\mathcal{D}}(C', m) = 1$ using Hoeffding's bound, until the probability difference is sufficiently large.

Proof: For $m, m' \in \mathcal{M}$, let $P(m', m)$ denote the probability that $\mathcal{O}^{\mathcal{D}}(C', m) = 1$, where $C' \leftarrow \text{Enc}(\text{pd}, \text{pk}, \mathcal{R}, m')$. Equation (1) states that

$$\Pr_{m \leftarrow \mathcal{M}} [P(m, m)] \geq \frac{1}{|\mathcal{M}|} + \frac{1}{\lambda^c}.$$

Let us assume by contradiction (of the statement to be proved), that

$$\Pr_{m, m' \leftarrow \mathcal{M}} [|P(m, m) - P(m', m)| < \frac{1}{2\lambda^c}] > 1 - \frac{1}{4\lambda^c}. \quad (2)$$

We show that if (2) holds, then the following inequality holds as well.

$$\Pr_{m' \leftarrow \mathcal{M}} [\Pr_{m \leftarrow \mathcal{M}} [|P(m, m) - P(m', m)| < \frac{1}{2\lambda^c}] > 1 - \frac{1}{2\lambda^c}] > \frac{1}{2}. \quad (3)$$

By contradiction of (3) above, let us assume that

$$\Pr_{m' \leftarrow \mathcal{M}} [\Pr_{m \leftarrow \mathcal{M}} [|P(m, m) - P(m', m)| < \frac{1}{2\lambda^c}] > 1 - \frac{1}{2\lambda^c}] \leq \frac{1}{2}.$$

We consider two types of m' , depending whether $\Pr_m [|P(m, m) - P(m', m)| < \frac{1}{2\lambda^c}]$ is greater than $1 - \frac{1}{2\lambda^c}$ (Type 1) or not (Type 2). Let $x \leq 1/2$ be the proportion of m' 's of the first type. Then we would have

$$\begin{aligned} & \Pr_{m, m'} [|P(m, m) - P(m', m)| < \frac{1}{2\lambda^c}] \\ &= \Pr_{m'} [\Pr_m [|P(m, m) - P(m', m)| < \frac{1}{2\lambda^c}]] \\ &= \frac{1}{|\mathcal{M}|} \sum_{\substack{m' \\ \text{of Type 1}}} \Pr_m [|P(m, m) - P(m', m)| < \frac{1}{2\lambda^c}] \\ &+ \frac{1}{|\mathcal{M}|} \sum_{\substack{m' \\ \text{of Type 2}}} \Pr_m [|P(m, m) - P(m', m)| < \frac{1}{2\lambda^c}] \\ &\leq \frac{1}{|\mathcal{M}|} \sum_{\substack{m' \\ \text{of Type 1}}} 1 + \frac{1}{|\mathcal{M}|} \sum_{\substack{m' \\ \text{of Type 2}}} (1 - \frac{1}{2\lambda^c}) \\ &= x + (1 - x)(1 - \frac{1}{2\lambda^c}) \leq 1 - \frac{1}{4\lambda^c}, \end{aligned}$$

which would contradict (2) above.

We consider an m' of Type 1. Using the fact that $\sum_m P(m', m) \leq 1$, we obtain:

$$\sum_m P(m, m) < \frac{|\mathcal{M}|}{2\lambda^c} + \sum_m \left(P(m', m) + \frac{1}{2\lambda^c} \right) \leq 1 + \frac{|\mathcal{M}|}{\lambda^c}.$$

This contradicts Equation (1). □

3.2 Semantic Security

We start by proving IND-CPA security of our scheme.

Theorem 3 *If \mathcal{FE} is r -IND-CPA secure, then \mathcal{TR} is IND-CPA secure.*

Proof: Let $\mathcal{A}_{\mathcal{TR}}$ be a probabilistic polynomial-time adversary that breaks semantic security of \mathcal{TR} . We construct a probabilistic polynomial-time adversary $\mathcal{A}_{\mathcal{FE}}$ that breaks semantic security of \mathcal{FE} . Adversary $\mathcal{A}_{\mathcal{FE}}$ proceeds as follows.

- It first obtains the public key \mathbf{pk} output by the \mathcal{FE} challenger (who runs the $\mathcal{FE}.\text{Setup}(1^\lambda, 1^\ell)$ algorithm) and relays it to $\mathcal{A}_{\mathcal{TR}}$.
- The adversary $\mathcal{A}_{\mathcal{TR}}$ adaptively chooses at most r identities id (that forms the revoked set \mathcal{R}) and are included in pd . The adversary $\mathcal{A}_{\mathcal{FE}}$ then queries the \mathcal{FE} challenger for each \mathbf{x}_{id} for all $\text{id} \in \mathcal{R}$ and receives the corresponding sk_{id} . Adversary $\mathcal{A}_{\mathcal{FE}}$ relays all sk_{id} for each $\text{id} \in \mathcal{R}$ to $\mathcal{A}_{\mathcal{TR}}$.
- When $\mathcal{A}_{\mathcal{TR}}$ chooses two messages $m_0, m_1 \in \mathcal{M}$ and provides them to $\mathcal{A}_{\mathcal{FE}}$, adversary $\mathcal{A}_{\mathcal{FE}}$ proceeds as follows:
 - It computes $\mathbf{v}_{\mathcal{R}} \in \mathbb{Z}_p^\ell \setminus \{\vec{0}\}$ such that $\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle = 0$ for every $\text{id} \in \mathcal{R}$.
 - It sends $\mathbf{y}_{\mathcal{R},0} = m_0 \cdot \mathbf{v}_{\mathcal{R}}$ and $\mathbf{y}_{\mathcal{R},1} = m_1 \cdot \mathbf{v}_{\mathcal{R}}$ to the \mathcal{FE} challenger who samples $b \leftarrow \{0, 1\}$ and encrypts $\mathbf{y}_{\mathcal{R},b}$ as $C_{\mathbf{y}_{\mathcal{R},b}} \leftarrow \mathcal{FE}.\text{Enc}(\mathbf{pk}, \mathbf{y}_{\mathcal{R},b})$.
 - Adversary $\mathcal{A}_{\mathcal{FE}}$ receives $C_{\mathbf{y}_{\mathcal{R},b}}$ from the \mathcal{FE} challenger and sends $C = (C_{\mathbf{y}_{\mathcal{R},b}}, \mathcal{R})$ to $\mathcal{A}_{\mathcal{TR}}$.
- Finally, adversary $\mathcal{A}_{\mathcal{TR}}$ outputs its guess $b' \in \{0, 1\}$ and $\mathcal{A}_{\mathcal{FE}}$ also outputs b' as its own guess of b .

Note that adversary $\mathcal{A}_{\mathcal{FE}}$ behaves as an IND-CPA challenger in the view of $\mathcal{A}_{\mathcal{TR}}$. Further, it is a valid adversary against \mathcal{FE} , as $\langle \mathbf{y}_{\mathcal{R},0}, \mathbf{x}_{\text{id}} \rangle = \langle \mathbf{y}_{\mathcal{R},1}, \mathbf{x}_{\text{id}} \rangle$ for every vector \mathbf{x}_{id} queried to the \mathcal{FE} challenger (i.e., each $\text{id} \in \mathcal{R}$). The advantage of $\mathcal{A}_{\mathcal{FE}}$ is exactly the same as the advantage of $\mathcal{A}_{\mathcal{TR}}$. \square

We may observe that for \mathcal{TR} to be IND-CPA secure, an r -IND-CPA secure \mathcal{FE} scheme is sufficient. However, as we see below, for traceability with up to t colluding traitors along with r already revoked users, we need an \mathcal{FE} scheme that is $(t+r)$ -IND-CPA secure.

3.3 Traceability

Here, we prove the traceability of the scheme. To start with, we first prove the following lemma.

Lemma 9 *Assume that a pirate decoder \mathcal{D} satisfies Equation (1) for some \mathcal{R} and \mathcal{S} . Then, the execution of Trace does not return \perp but returns some $\text{id} \in \mathcal{S}$ with overwhelming probability.*

Proof: We consider a variant of Trace that continues its execution until it exhausts $\mathcal{S} \setminus \mathcal{R}$, even if it has already output an id . We consider the probabilities p_i at the start and end of that modified execution.

1. At the beginning, algorithm Trace considers $\mathcal{S}_1 = \mathcal{S} \setminus \mathcal{R}$ and $\mathbf{v}_{\mathcal{S}_1} = \mathbf{0}$. Hence, the genuine ciphertext output by the Enc algorithm and the probe ciphertext created by the Trace algorithm for the suspect subset \mathcal{S}_1 are exactly the same.
2. When $i = |\mathcal{S} \setminus \mathcal{R}| + 1$, we have $\mathcal{S}_i = \emptyset$ and $\mathbf{v}_{\mathcal{S}_i} = (m' - m) \cdot \mathbf{v}_{\mathcal{R}}$. In Step (a) of the Trace algorithm, the messages m and m' were chosen such that the difference in the probabilities p_1 and $p_{|\mathcal{S} \setminus \mathcal{R}|+1}$ is $\geq 1/(2\lambda^c)$.

Note that the two latter observations imply, via the triangle inequality, that there exists an i such that $|p_i - p_{i-1}|$ is non-negligible. By the Hoeffding bound, Trace algorithm outputs id_{i-1} with overwhelming probability. \square

Then, we prove the following theorem.

Theorem 4 *If \mathcal{FE} is $(t + r)$ -IND-CPA secure, then \mathcal{TR} satisfies public traceability.*

Proof: Let us assume by contradiction that an adversary \mathcal{A} can break the public traceability of \mathcal{TR} with non-negligible probability. We then construct a probabilistic polynomial-time adversary $\mathcal{A}_{\mathcal{FE}}$ that breaks the semantic security of \mathcal{FE} . Adversary $\mathcal{A}_{\mathcal{FE}}$ proceeds as follows.

- It first obtains the public key pk output by the \mathcal{FE} challenger (who runs the $\mathcal{FE}.\text{Setup}(1^\lambda, 1^\ell)$ algorithm) and relays it to the adversary \mathcal{A} .
- When \mathcal{A} asks $\mathcal{A}_{\mathcal{FE}}$ to create a p_{id} for some id , adversary $\mathcal{A}_{\mathcal{FE}}$ in turn asks the \mathcal{FE} challenger to do the same. The \mathcal{FE} challenger randomly chooses a vector $\mathbf{x}_{\text{id}} \leftarrow \mathbb{Z}_p^\ell$ and sends it to $\mathcal{A}_{\mathcal{FE}}$ who further relays it to \mathcal{A} .
- When \mathcal{A} makes a key query for an identity id , adversary $\mathcal{A}_{\mathcal{FE}}$ queries the \mathcal{FE} challenger for a secret key. Adversary $\mathcal{A}_{\mathcal{FE}}$ receives the corresponding sk_{id} from the \mathcal{FE} challenger and relays it to \mathcal{A} .
- When \mathcal{A} chooses a set \mathcal{R} of up to r revoked users, adversary $\mathcal{A}_{\mathcal{FE}}$ makes $|\mathcal{R}|$ key queries to the \mathcal{FE} challenger. Adversary $\mathcal{A}_{\mathcal{FE}}$ is given the set sk_{id} 's of corresponding secret keys that is relayed to \mathcal{A} . Recall that by the definition of the public traceability game, these queries can be interleaved with extensions of the number of users and user corruption queries, in an adaptive manner.

Note that since \mathcal{A} makes at most t key queries and $|\mathcal{R}| \leq r$, adversary $\mathcal{A}_{\mathcal{FE}}$ makes at most $t + r$ key queries for the \mathcal{FE} challenger.

- Adversary \mathcal{A} finally produces a pirate decoder \mathcal{D}^4 and chooses a suspect set \mathcal{S} of cardinality $\leq t$ that contains \mathcal{T} . Then, the adversary $\mathcal{A}_{\mathcal{FE}}$ executes the Trace algorithm on $\mathcal{O}^{\mathcal{D}}$ to find i such that $|p_i - p_{i-1}|$ is non-negligible. If Trace outputs \perp or index i such that $\text{id}_{i-1} \in \mathcal{T}$, then $\mathcal{A}_{\mathcal{FE}}$ outputs a random bit. We say that the event Abort occurs in such a case. Otherwise, it sets $\mathbf{y}_0 = \mathbf{v}_{\mathcal{S}_{i-1}} + m \cdot \mathbf{v}_{\mathcal{R}}$ and $\mathbf{y}_1 = \mathbf{v}_{\mathcal{S}_i} + m \cdot \mathbf{v}_{\mathcal{R}}$, and sends them as challenge messages to the \mathcal{FE} challenger.⁵
- The \mathcal{FE} challenger samples $b \leftarrow \{0, 1\}$ and then sends $\mathcal{FE}.\text{Enc}(\text{pk}, \mathbf{y}_b)$ to $\mathcal{A}_{\mathcal{FE}}$. The adversary $\mathcal{A}_{\mathcal{FE}}$ runs $\mathcal{O}^{\mathcal{D}}$ on input (C_b, m) , where $C_b = (\mathcal{FE}.\text{Enc}(\text{pk}, \mathbf{y}_b), \mathcal{R})$. Then $\mathcal{O}^{\mathcal{D}}$ outputs the bit $b' \in \{0, 1\}$.
- Finally, adversary $\mathcal{A}_{\mathcal{FE}}$ outputs the same bit $b' \in \{0, 1\}$ if $p_i - p_{i-1} > 0$ and $1 - b'$ otherwise.

We first argue that $\mathcal{A}_{\mathcal{FE}}$ is a valid adversary against the \mathcal{FE} challenger. Recall that when Abort does not occur, we have $\text{id}_{i-1} \notin \mathcal{T}$ but $\text{id}_{i-1} \in \mathcal{S}$. The keys queried by $\mathcal{A}_{\mathcal{FE}}$ are for $\text{id} \in \mathcal{R} \cup \mathcal{T}$. This set $\mathcal{R} \cup \mathcal{T}$ can be partitioned into $\mathcal{R} \cup (\mathcal{T} \cap \mathcal{S}_{i-1})$ and $\mathcal{T} \cap (\mathcal{S}_1 \setminus \mathcal{S}_{i-1})$. Note that since $\text{id}_{i-1} \notin \mathcal{T}$, we have $\mathcal{T} \cap \mathcal{S}_{i-1} = \mathcal{T} \cap \mathcal{S}_i$ and thus $\mathcal{R} \cup (\mathcal{T} \cap \mathcal{S}_{i-1}) = \mathcal{R} \cup (\mathcal{T} \cap \mathcal{S}_i)$.

1. For $\text{id} \in \mathcal{R}$, we have $\langle \mathbf{x}_{\text{id}}, \mathbf{y}_0 \rangle = \langle \mathbf{x}_{\text{id}}, \mathbf{y}_1 \rangle = 0$. For $\text{id} \in \mathcal{S}_{i-1} \cap \mathcal{T} = \mathcal{S}_i \cap \mathcal{T}$, we have $\langle \mathbf{x}_{\text{id}}, \mathbf{y}_0 \rangle = \langle \mathbf{x}_{\text{id}}, \mathbf{y}_1 \rangle = m \cdot \langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle$. Hence for all $\text{id} \in \mathcal{R} \cup (\mathcal{T} \cap \mathcal{S}_{i-1})$ for which the sk_{id} was queried by $\mathcal{A}_{\mathcal{FE}}$, the inner products $\langle \mathbf{x}_{\text{id}}, \mathbf{y}_0 \rangle$ and $\langle \mathbf{x}_{\text{id}}, \mathbf{y}_1 \rangle$ have the same value.

⁴Recall that we assume that \mathcal{D} is stateless/resettable and replies independently to successive queries.

⁵Here, m and m' are chosen as in Step (a), $\mathbf{v}_{\mathcal{R}} \in \mathbb{Z}_p^\ell$ is chosen as in Step (c), and $\mathbf{v}_{\mathcal{S}_{i-1}}, \mathbf{v}_{\mathcal{S}_i} \in \mathbb{Z}_p^\ell$ are chosen as in Step (d) of algorithm Trace.

2. Similarly, for $\text{id} \in \mathcal{T} \cap (\mathcal{S}_1 \setminus \mathcal{S}_{i-1})$, we have $\langle \mathbf{x}_{\text{id}}, \mathbf{y}_0 \rangle = \langle \mathbf{x}_{\text{id}}, \mathbf{y}_1 \rangle = m' \cdot \langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle$.

Hence, $\mathcal{A}_{\mathcal{FE}}$ is a valid adversary against the \mathcal{FE} challenger.

We recollect that in the AD-TT game, we say that \mathcal{A} wins if the decryption box \mathcal{D} output by it is such that when **Trace** is executed on input $\mathcal{O}^{\mathcal{D}}$, it fails to identify a traitor. In such a case, **Trace** either outputs \perp or it outputs an $\text{id}_{i-1} \notin \mathcal{T}$ with probability at least $1/\lambda^c$. We next argue that if \mathcal{A} outputs \mathcal{D} that satisfies this winning condition of the AD-TT game, then $\mathcal{A}_{\mathcal{FE}}$ has non-negligible advantage in the above game. To see this, we first observe that when **Abort** occurs, $\mathcal{A}_{\mathcal{FE}}$ returns a random bit and it correctly guesses b with probability $1/2$. Then, it suffices to show the following:

- In the above game, **Abort** does not occur with non-negligible probability.
- Conditioned on **Abort** not occurring, **Trace** outputs id_{i-1} such that $|p_i - p_{i-1}|$ is non-negligible.

Indeed, the combination of them implies that the advantage of $\mathcal{A}_{\mathcal{FE}}$ is non-negligible, since $|p_i - p_{i-1}|$ is the advantage of $\mathcal{A}_{\mathcal{FE}}$ conditioned on **Abort** not occurring.

The second item follows because if $|p_i - p_{i-1}|$ is not sufficiently large, **Trace** does not output id_{i-1} at Step (f) of **Trace** except for a negligible probability (because of the Hoeffding bound). Next, we prove the first item. Since we are assuming \mathcal{D} satisfies the winning condition, when **Trace** is executed on input $\mathcal{O}^{\mathcal{D}}$, it outputs \perp or it outputs an $\text{id}_{i-1} \notin \mathcal{T}$ with probability at least $1/\lambda^c$. The claim now follows since the former event occurs only with negligible probability by Lemma 9. \square

4 Trace and Revoke from Learning with Errors

Recall that Agrawal et al. (2016) provided a construction for inner product functional encryption from LWE. Instantiating our generic transformation of Section 3 with this scheme is possible, but leads to reliance on LWE with subexponential error rates. In Subsection 4.2, we provide a new construction of an inner product functional encryption scheme from LWE in a much weaker model than that considered in Agrawal et al. (2016). We restrict to the setting of bounded collusions and also crucially exploit the fact that the adversary's key requests are random vectors for our application as described in Section 3. The performances of both resulting trace-and-revoke systems are discussed in Subsection 4.1.

4.1 Two Trace-and-Revoke Constructions

Our IPFE to trace-and-revoke generic transformation cannot be directly instantiated with the LWE-based IPFE over \mathbb{Z}_p from Agrawal et al. (2016), because the key generation algorithm of the latter is stateful: it keeps track of all the secret keys it has generated. The statefulness necessity may be explained as follows. The master secret key is an *integer* matrix with small entries. When the attacker makes a key query for a vector *modulo* p , it learns the *integer* product between a conversion to the integers of that vector and the master secret key. If the key generation algorithm does not maintain a state, then it does not seem possible to prevent an adversary from making key queries for vectors that are linearly dependent *modulo* p but linearly independent over the *integers*: the attacker could then make valid key queries but still learn the master secret key.

The Key Generation State is Unnecessary. In Agrawal et al. (2016), it was noted that if the vectors queried by the adversary are guaranteed to be linearly independent modulo p , then there is no need for a stateful key generation algorithm. In our case, there are as many vectors as users, each vector is uniformly sampled from \mathbb{Z}_p^ℓ and the adversary has access to $\leq r + t < \ell$ vectors. By setting $p = 2^{\Omega(\lambda)}$, the probability that there

exists a subset of t key vectors that are linearly independent is $2^{-\Omega(\lambda)}$. We can then remove the state in the LWE-based IPFE over \mathbb{Z}_p^ℓ , and apply the transformation from the previous section.

The resulting trace-and-revoke scheme inherits the unsatisfactory performance of its underlying IPFE (see (Agrawal et al., 2016, Section 4.2) for further details), stemming from the subexponential error rate in the LWE hardness assumption.

Large LWE Errors are Unnecessary. In Subsection 4.2 below, we exploit the randomness of the key queries further, as well as the bounded number of queries (as allowed by our trace-and-revoke application). We obtain a random-key bounded-collusion FE for inner products from LWE with significantly better parameters. In particular, we rely on slightly super-polynomial error rates for LWE, which allows to take smaller parameters.⁶

Both the public key and master secret key of the resulting trace-and-revoke scheme consist of $\tilde{O}((t+r+\lambda)\lambda)$ bits. To every user id corresponds a secret key sk_{id} of bit-length $\tilde{O}(t+r+\lambda)$ and a vector p_{id} of bit-length $\tilde{O}(t+r+\lambda)$. Algorithm Enc maps a plaintext in $\{0, 1\}$ to a ciphertext of bit-length $\tilde{O}(t+r+\lambda)$.

4.2 Bounded Collusion FE for Inner Products from LWE

The construction we provide here relies on LWE with a small error rate and hence small modulus and dimension. Our construction is quite close to Agrawal et al. (2016) except the key generation algorithm. In Agrawal et al. (2016), the key generation algorithm is deterministic, whereas in our scheme it is randomized and involves certain Gaussian distribution. This change allows us to prove the security of our scheme in the improved parameter setting compared with Agrawal et al. (2016).

Construction. Let p be the modulus of the scheme, 2ℓ be the dimension of the scheme, and L be the upper bound on the size of the collusion.

- $\mathcal{FE}.\text{Setup}(1^\lambda, 1^L, 1^{2\ell})$. Set integers $n, m, q = p^k$ for some integer $k \geq 2$, and reals $\alpha \in (0, 1)$ and $\sigma_0, \sigma_1, \sigma_2 > 0$, as explained below. Sample $\mathbf{A} \leftarrow \mathbb{Z}_q^{m \times n}$ and $\mathbf{Z} \leftarrow D_{\mathbb{Z}, \sigma_0}^{2\ell \times m}$. Compute $\mathbf{U} = \mathbf{Z} \cdot \mathbf{A} \in \mathbb{Z}_q^{2\ell \times n}$. Define

$$\text{msk} := \mathbf{Z} \quad \text{and} \quad \text{pk} := (\mathbf{A}, \mathbf{U}).$$

- $\mathcal{FE}.\text{KeyGen}(\text{msk}, \mathbf{x})$. Given $\mathbf{x} = (x_1, \dots, x_{2\ell})^t \in \mathbb{Z}_p^{2\ell}$, sample $\bar{x}_i \leftarrow D_{p\mathbb{Z}+x_i, \sigma_1}$ for $i \in [\ell]$ and $\bar{x}_i \leftarrow D_{p\mathbb{Z}+x_i, \sigma_2}$ for $i \in [\ell+1, 2\ell]$. Set $\bar{\mathbf{x}} := (\bar{x}_1, \dots, \bar{x}_{2\ell})^t \in \mathbb{Z}^{2\ell}$ and $\mathbf{z}_{\mathbf{x}} = \bar{\mathbf{x}}^t \cdot \mathbf{Z} \in \mathbb{Z}^m$. Note that we have $\bar{\mathbf{x}} \equiv \mathbf{x} \pmod{p}$ by construction. Finally, return $\text{sk}_{\mathbf{x}} = (\bar{\mathbf{x}}, \mathbf{z}_{\mathbf{x}})$.

- $\mathcal{FE}.\text{Enc}(\text{pk}, \mathbf{y})$. To encrypt a vector $\mathbf{y} \in \mathbb{Z}_p^{2\ell}$, sample $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{e}_0, \mathbf{e}_1 \leftarrow D_{\mathbb{Z}, \alpha q}^m$ and compute

$$\mathbf{c}_0 = \mathbf{A}\mathbf{s} + \mathbf{e}_0 \in \mathbb{Z}_q^m, \quad \mathbf{c}_1 = \mathbf{U}\mathbf{s} + \mathbf{e}_1 + p^{k-1} \cdot \mathbf{y} \in \mathbb{Z}_q^{2\ell}.$$

Then, return the ciphertext $C = (\mathbf{c}_0, \mathbf{c}_1)$.

- $\mathcal{FE}.\text{Dec}(\text{sk}_{\text{id}}, C)$. Given $C = (\mathbf{c}_0, \mathbf{c}_1)$ and a secret key $(\bar{\mathbf{x}}, \mathbf{z}_{\mathbf{x}})$ for $\mathbf{x} \in \mathbb{Z}_p^{2\ell}$, compute $\mu' = \langle \bar{\mathbf{x}}, \mathbf{c}_1 \rangle - \langle \mathbf{z}_{\mathbf{x}}, \mathbf{c}_0 \rangle \pmod{q}$ and output the value $\mu \in \mathbb{Z}_p$ that minimizes $|p^{k-1} \cdot \mu - \mu'|$.

⁶We observe that the scheme from Subsection 4.2 allows for polynomial error rates, but the correctness of our trace-and-revoke construction requires $p \geq \lambda^{\omega(1)}$, which leads to a $\lambda^{\omega(1)}$ LWE error rate in our IPFE.

Setting the Parameters. We have to set the parameters so that the correctness requirement is satisfied and the security reduction from $\text{LWE}_{n,m,q,\alpha'}$ works, for some non-trivial error rate α' . We require that

- $p^{k-1}/4 > \sigma_0(\sigma_1 + \sigma_2)\alpha q\sqrt{\ell m} \cdot \omega(\log^{3/2} \lambda)$, to ensure that the error term in decryption has magnitude less than $p^{k-1}/4$ with probability $1 - \lambda^{-\omega(1)}$,
- $\sigma_1, \sigma_2 \geq p \cdot \Omega(\sqrt{\lambda})$, to be able to apply Lemma 1 in the security proof,
- $\alpha/\alpha' \geq \Omega(\sigma_0\sqrt{m})$ and $\sigma_0, \alpha'q \geq \Omega(\sqrt{\lambda})$, to be able to apply Lemma 2 in the security proof,
- $\kappa \geq \Omega(\lambda + L \log \lambda)$, to ensure the (overwhelmingly likely) existence of a \mathbf{U} as in Lemma 5 in the security proof,
- $\sigma_1 \geq \Omega(\sqrt{\ell \kappa \log \ell})$, $\ell \geq \Omega(\kappa \log(\sigma_1 \kappa))$, and $\sigma_2 \geq \Omega(\kappa^{5/2} \sqrt{\ell} \sigma_1^2 \log^{3/2}(\ell \sigma_1))$, to be able to apply Lemma 5 in the security proof with $\kappa \geq \Omega(\lambda + L \log \lambda)$ as above,
- $\sigma_0 \geq \Omega(p\kappa\ell\sigma_2)$ and $q^{n+1}/p^m \leq 2^{-\Omega(\kappa)}$, to be able to apply Lemma 10 in the security proof.

To satisfy the above requirements and rely on LWE parameters for which all known attacks cost $2^{o(\lambda)}$, we may set the parameters as follows. We choose $\kappa = \Theta(\lambda + L \log \lambda)$, $p = \lambda^{\omega(1)}$, and:⁷

$$\begin{aligned} \ell &= \tilde{\Theta}((\lambda + L) \log p) \\ \sigma_0 &= \tilde{\Theta}((\lambda + L)^5 (p \log p)^3 \lambda) & \sigma_1 &= \Theta(p\sqrt{\lambda}) \\ \sigma_2 &= \tilde{\Theta}((\lambda + L)^3 (p \log p)^2 \lambda) \\ 1/\alpha &= \tilde{\Theta}((\lambda + L)^9 (p \log p)^6 \lambda^2) \\ 1/\alpha' &= \tilde{\Theta}((\lambda + L)^{14.5} (p \log p)^9 \lambda^3) & m &= \tilde{\Theta}(\lambda + L) \\ q &= \tilde{\Theta}((\lambda + L)^{15} (p \log p)^9 \lambda^3) & k &= \Theta(1) \\ n &= \tilde{\Theta}(\lambda) \end{aligned}$$

Decryption Correctness. To show the correctness of the scheme, we first observe that, modulo q :

$$\mu' = \langle \bar{\mathbf{x}}, \mathbf{c}_1 \rangle - \langle \mathbf{z}_x, \mathbf{c}_0 \rangle = p^{k-1} \cdot \langle \mathbf{x}, \mathbf{y} \rangle + \langle \bar{\mathbf{x}}, \mathbf{e}_1 \rangle - \langle \mathbf{z}_x, \mathbf{e}_0 \rangle.$$

Below, we show that the magnitude of the term $\langle \bar{\mathbf{x}}, \mathbf{e}_1 \rangle - \langle \mathbf{z}_x, \mathbf{e}_0 \rangle$ is $\leq \sigma_0(\sigma_1 + \sigma_2)\alpha q\sqrt{\ell m} \cdot \omega(\log^{3/2} \lambda)$ with probability $1 - \lambda^{-\omega(1)}$. Thanks to the parameter choices, the latter upper bound is smaller than $p^{k-1}/4$, which suffices to guarantee decryption correctness.

Note that $\bar{x}_i \in \mathbb{Z}^{2\ell}$ is chosen from $D_{p\mathbb{Z}+x_i, \sigma_1}$ if $i \in [\ell]$ and $D_{p\mathbb{Z}+x_i, \sigma_2}$ otherwise. We thus have $\|\bar{\mathbf{x}}\| \leq (\sigma_1 + \sigma_2)\sqrt{\ell} \cdot \omega(\sqrt{\log \lambda})$ with probability $1 - \lambda^{-\omega(1)}$. This, together with $\mathbf{e}_1 \sim D_{\mathbb{Z}, \alpha q}^{2\ell}$, implies that $|\langle \bar{\mathbf{x}}, \mathbf{e}_1 \rangle| \leq \alpha q(\sigma_1 + \sigma_2)\sqrt{\ell} \cdot \omega(\log \lambda)$ with probability $1 - \lambda^{-\omega(1)}$. Furthermore, since each column of \mathbf{Z} is chosen from $D_{\mathbb{Z}, \sigma_0}^{2\ell}$, we have $\|\mathbf{z}_x\| \leq \sigma_0(\sigma_1 + \sigma_2)\sqrt{\ell m} \cdot \omega(\log \lambda)$ with probability $1 - \lambda^{-\omega(1)}$. As a result, we have $|\langle \mathbf{z}_x, \mathbf{e}_0 \rangle| \leq \sigma_0(\sigma_1 + \sigma_2)\alpha q\sqrt{\ell m} \cdot \omega(\log^{3/2} \lambda)$ with probability $1 - \lambda^{-\omega(1)}$.

Security. We now show that the scheme above is secure, for our relaxed notion of L -IND-CPA security. The proof is similar to Agrawal et al. (2016), but we exploits the weaker security model of bounded number of random key queries. In particular, we perform a much more careful analysis on the conditional distribution of \mathbf{Z} from the view of the adversary.

Theorem 5 *If the parameters are set as above, the above scheme is L -IND-CPA secure under the $\text{LWE}_{n,m,q,\alpha'}$ assumption.*

⁷We note that it is possible to choose parameters that allow to take p as low as $p = 2$, but in our trace-and-revoke application we use $p = \lambda^{\omega(1)}$ to guarantee correctness.

Proof: The proof proceeds with a sequence of games that starts with the real game and ends with a game in which the adversary's advantage is negligible. For each i , we call S_i the event that the adversary wins in Game i .

Game 0: This is the ordinary security game. Namely, at the outset of the game, the adversary \mathcal{A} is given the master public key pk . Then, it sees Q random vectors $\{\mathbf{x}_i\}_{i \in [Q]}$, where $\mathbf{x}_i \leftarrow \mathbb{Z}_p^{2\ell}$ and Q is an arbitrary polynomial specified by \mathcal{A} . Then, it makes secret key queries for these vectors. The number of the key queries is bounded by L . Note that the adversary can only make key queries for random vectors chosen as $\mathbf{x} \leftarrow \mathbb{Z}_p^{2\ell}$. In the challenge phase, the adversary \mathcal{A} comes up with two distinct vectors $\mathbf{y}_0, \mathbf{y}_1$ and receives an encryption C of \mathbf{y}_β for $\beta \leftarrow \{0, 1\}$ sampled by the challenger. The adversary is not allowed to make secret key queries after the challenge phase. When \mathcal{A} halts, it outputs $\beta' \in \{0, 1\}$ and S_0 is the event that $\beta' = \beta$. Note that for any vector \mathbf{x} for which \mathcal{A} makes a secret key query, we must have $\langle \mathbf{x}, \mathbf{y}_0 \rangle \equiv \langle \mathbf{x}, \mathbf{y}_1 \rangle \pmod p$ if \mathcal{A} is a legitimate adversary.

Game 1: We modify the generation of \mathbf{x} and $\bar{\mathbf{x}}$ for all secret key queries. Namely, instead of choosing $\mathbf{x} \leftarrow \mathbb{Z}_p^{2\ell}$ and then sampling $\bar{\mathbf{x}}$, the challenger first chooses $\bar{\mathbf{x}} = (\bar{x}_1, \dots, \bar{x}_{2\ell})^t$ as $\bar{x}_i \leftarrow D_{\mathbb{Z}, \sigma_1}$ for $i \in [\ell]$ and $\bar{x}_i \leftarrow D_{\mathbb{Z}, \sigma_2}$ for $i \in [\ell + 1, 2\ell]$ and then sets $\mathbf{x} := \bar{\mathbf{x}} \pmod p$. We claim that this changes the joint distribution of $(\mathbf{x}, \bar{\mathbf{x}})$ only negligibly. To see this, we observe that the distribution of \bar{x}_i conditioned on $\bar{x}_i \equiv x_i \pmod p$ is $D_{p\mathbb{Z} + x_i, \sigma_1}$ for $i \in [\ell]$ and $D_{p\mathbb{Z} + x_i, \sigma_2}$ for $i \in [\ell + 1, 2\ell]$. Therefore, it suffices to show that $\bar{x}_i \pmod p$ is statistically close to the uniform distribution over \mathbb{Z}_p when \bar{x}_i is chosen from $D_{\mathbb{Z}, \sigma_1}$ or $D_{\mathbb{Z}, \sigma_2}$. This follows from $\sigma_1, \sigma_2 \geq p \cdot \Omega(\sqrt{\lambda})$ and Lemma 1. Therefore, we have that $|\Pr[S_1] - \Pr[S_0]| \leq 2^{-\Omega(\lambda)}$.

Game 2: We modify the generation of $C = (\mathbf{c}_0, \mathbf{c}_1)$ in the challenge phase. Namely at the outset of the game, the challenger picks $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{e} \leftarrow D_{\mathbb{Z}, \alpha'q}^m$ (which may be chosen ahead of time) as well as $\mathbf{Z} \leftarrow D_{\mathbb{Z}, \sigma_0}^{2\ell \times m}$. Let $\mathbf{V} \in \mathbb{Z}^{(m+2\ell) \times m}$ be the matrix that is obtained by putting \mathbf{I}_m on top of \mathbf{Z} , where \mathbf{I}_m is the unit matrix of size m . We then set the ciphertext $C = (\mathbf{c}_0, \mathbf{c}_1) \in \mathbb{Z}_q^m \times \mathbb{Z}_q^{2\ell}$ as

$$\begin{aligned} \mathbf{b} &= \mathbf{A}\mathbf{s} + \mathbf{e} \\ (\mathbf{c}_0 \parallel \mathbf{c}_1) &= \text{ReRand}(\mathbf{V}, \mathbf{b}, \alpha'q, \alpha/(2\alpha')) + p^{k-1} \cdot \mathbf{y}_\beta \end{aligned} \quad (4)$$

where ReRand is from Lemma 2. We claim that this change alters the view of the adversary only negligibly. To show this, we first observe that $s_1(\mathbf{V}) \leq \sqrt{1 + s_1(\mathbf{Z})^2} \leq O(\sigma_0 \sqrt{m})$ holds with all but $2^{-\Omega(m)} \leq 2^{-\Omega(\lambda)}$ probability by Lemma 4. By Lemma 2 and our parameter choices, we have

$$\begin{aligned} \mathbf{c}_0 &= \mathbf{I}_m \cdot \mathbf{A}\mathbf{s} + \mathbf{e}_0 = \mathbf{A}\mathbf{s} + \mathbf{e}_0, \\ \mathbf{c}_1 &= \mathbf{Z} \cdot \mathbf{A}\mathbf{s} + \mathbf{e}_1 + p^{k-1} \cdot \mathbf{y}_\beta = \mathbf{U}\mathbf{s} + \mathbf{e}_1 + p^{k-1} \cdot \mathbf{y}_\beta, \end{aligned}$$

where \mathbf{e}_0 and \mathbf{e}_1 are within statistical distance $2^{-\Omega(\lambda)}$ from $D_{\mathbb{Z}, \alpha'q}^m$. Therefore, we have that $|\Pr[S_2] - \Pr[S_1]| \leq 2^{-\Omega(\lambda)}$.

Game 3: We further modify the generation of $C = (\mathbf{c}_0, \mathbf{c}_1)$ in the challenge phase. Instead of setting $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$, we choose $\mathbf{b} = \mathbf{u} + \mathbf{e}$, where $\mathbf{u} \leftarrow \mathbb{Z}_q^m$. Then, the ciphertext is set as in Equation (4). Under the LWE assumption, we have that $|\Pr[S_3] - \Pr[S_2]|$ is negligible.

Game 4: We modify the generation of $C = (\mathbf{c}_0, \mathbf{c}_1)$ once more. Namely, the ciphertext is now set as

$$\begin{aligned} \mathbf{c}_0 &= \mathbf{u} + \mathbf{e}_0, \\ \mathbf{c}_1 &= \mathbf{Z} \cdot \mathbf{u} + \mathbf{e}_1 + p^{k-1} \cdot \mathbf{y}_\beta, \end{aligned}$$

where $\mathbf{u} \leftarrow \mathbb{Z}_q^m$ and $\mathbf{e}_0, \mathbf{e}_1 \leftarrow D_{\mathbb{Z}, \alpha'q}^m$. Similarly to Game 2, this change does not alter the view of the adversary much. By Lemma 2 and our parameter choices, we have that $|\Pr[S_4] - \Pr[S_3]| \leq 2^{-\Omega(\lambda)}$. Below, we prove that $\Pr[S_4]$ is exponentially close to $1/2$, which will complete the proof.

Define $\mathbf{y} = \mathbf{y}_1 - \mathbf{y}_0 \in \mathbb{Z}_p^{2\ell}$. Let $\{\mathbf{x}_{i_j} \in \mathbb{Z}_p^{2\ell}\}_{j \in [L']}$ be the vectors corresponding to the secret key queries made by \mathcal{A} , where $L' \leq L$. As \mathcal{A} is a legitimate adversary, we have $\langle \mathbf{x}_{i_j}, \mathbf{y} \rangle = 0 \pmod p$ for each secret key query \mathbf{x}_{i_j} . The view of the adversary contains L' tuples $\{\mathbf{x}_{i_j}, \bar{\mathbf{x}}_{i_j}, \mathbf{z}_{\mathbf{x}_{i_j}}\}_{j \in [L']}$, where the vectors $\{\mathbf{x}_{i_j}\}_{j \in [L']}$ form a \mathbb{Z}_p -basis of a subspace of the $(2\ell - 1)$ -dimensional vector space $\mathbf{y}^\perp := \{\mathbf{x} \in \mathbb{Z}_p^{2\ell} : \langle \mathbf{x}, \mathbf{y} \rangle = 0 \pmod p\}$. We define $\mathbf{X}_{top} \in \mathbb{Z}^{L' \times 2\ell}$ as the matrix whose j -th row is $\bar{\mathbf{x}}_{i_j}^t$ for $j \in [L']$.

We say that $\mathbf{X}_{top} \in \mathbb{Z}^{L' \times 2\ell}$ is *good* when we can choose $\mathbf{U} \in \mathbb{Z}^{2\ell \times 2\ell}$ such that $|\det \mathbf{U}| = 1$, $\mathbf{X}_{top} \cdot \mathbf{U} = (\mathbf{I}_{L'} | \mathbf{0})$, and every row of \mathbf{U} has norm $\leq O(\sqrt{\kappa\ell}\sigma_2)$ (see Lemma 5). For a good \mathbf{X}_{top} , we can define $\mathbf{X} \in \mathbb{Z}^{2\ell \times 2\ell}$ as $\mathbf{X} := \mathbf{U}^{-1}$. It can be seen that the upper L' rows of \mathbf{X} corresponds to \mathbf{X}_{top} . We denote the lower $2\ell - L'$ rows of the matrix as \mathbf{X}_{bot} . We note that since \mathbf{X} is invertible over \mathbb{Z} , so is it modulo q . Without loss of generality, we assume that \mathbf{U} and \mathbf{X}_{bot} are deterministically determined from \mathbf{X}_{top} . (If there are more than one matrix satisfying the required properties, we sort them in the lexicographical order and pick the first one.) Note that it might be infeasible to efficiently compute \mathbf{U} and \mathbf{X}_{bot} from \mathbf{X}_{top} . This does not cause any problem in our proof because all the following arguments are information theoretic.

We state the following lemmas:

Lemma 10 *Assume that $\sigma_0 \geq \Omega(p\kappa\ell\sigma_2)$ and $q^{n+1}/p^m \leq 2^{-\Omega(\kappa)}$. Then the following distributions are within $2^{-\Omega(\kappa)}$ statistical distance:*

$$\begin{aligned} & (\mathbf{A}, \mathbf{u}, \mathbf{Z}\mathbf{A}, \mathbf{X}_{top}, \mathbf{X}_{top}\mathbf{Z}, \mathbf{X}_{bot}\mathbf{Z}\mathbf{u}) \approx \\ & (\mathbf{A}, \mathbf{u}, \mathbf{Z}\mathbf{A}, \mathbf{X}_{top}, \mathbf{X}_{top}\mathbf{Z}, \mathbf{v}) \end{aligned}$$

where $\mathbf{A} \leftarrow \mathbb{Z}_q^{m \times n}$, $\mathbf{u} \leftarrow \mathbb{Z}_q^m$, $\mathbf{Z} \leftarrow D_{\mathbb{Z}, \sigma_0}^{2\ell \times m}$, each row of $\mathbf{X}_{top} \in \mathbb{Z}_q^{L' \times 2\ell}$ is chosen from $D_{\mathbb{Z}, \sigma_1}^\ell \times D_{\mathbb{Z}, \sigma_2}^\ell$, and $\mathbf{v} \leftarrow \mathbb{Z}_q^{2\ell - L'}$. Note that if \mathbf{X}_{bot} is not good, then \mathbf{X}_{bot} is undefined. In such a case, the term $\mathbf{X}_{bot}\mathbf{Z}\mathbf{u}$ is replaced with \perp .

Lemma 11 *If there exists an adversary \mathcal{A} whose advantage in **Game 4** is ϵ , then there exists another (unbounded) adversary \mathcal{B} whose distinguishing advantage between the two distributions in Lemma 10 is $\epsilon/Q^{L'}$.*

Given these two lemmas, we can conclude the proof of Theorem 5 since these imply $\epsilon/Q^{L'} < 2^{-\Omega(\kappa)}$ and thus $\epsilon < Q^L \cdot 2^{-\Omega(\kappa)} = 2^{O(L \log \lambda) - \Omega(\kappa)} \leq 2^{-\Omega(\lambda)}$. \square

It remains to prove Lemmas 10 and 11.

Proof: [Proof of Lemma 10] By Lemma 5, matrix \mathbf{X}_{top} is good with all but $2^{-\Omega(\kappa)}$ probability. In the following, let us fix good \mathbf{X}_{top} and prove that the above two distributions are $2^{-\Omega(\kappa)}$ -close. We first consider the distribution $\mathbf{X}_{bot}\mathbf{Z}$ conditioned on $\mathbf{X}_{top}\mathbf{Z}$. Note that in $\mathbf{X}_{top}\mathbf{Z}$ and $\mathbf{X}_{bot}\mathbf{Z}$, matrices \mathbf{X}_{top} and \mathbf{X}_{bot} act in parallel on the columns of \mathbf{Z} . We can hence restrict ourselves to the distribution of $\mathbf{X}_{bot}\mathbf{z}_i$ conditioned on $\mathbf{X}_{top}\mathbf{z}_i$, with \mathbf{z}_i sampled from $D_{\mathbb{Z}^{2\ell}, \sigma_0}$. Let $\mathbf{b}_i = \mathbf{X}_{top}\mathbf{z}_i \in \mathbb{Z}^{L'}$ and fix $\mathbf{z}_i^* \in \mathbb{Z}^{2\ell}$ arbitrary such that $\mathbf{b}_i = \mathbf{X}_{top}\mathbf{z}_i^*$. The distribution of \mathbf{z}_i given $(\mathbf{X}_{top}, \mathbf{X}_{top}\mathbf{z}_i)$ is $D_{\Lambda + \mathbf{z}_i^*, \sigma_0}$, with $\Lambda = \{\mathbf{x} \in \mathbb{Z}^{2\ell} : \mathbf{X}_{top}\mathbf{x} = \mathbf{0}\}$. Therefore, we have that given $\mathbf{X}_{top}\mathbf{z}_i$, the vector $\mathbf{X}_{bot}\mathbf{z}_i$ is distributed as $\mathbf{X}_{bot} \cdot D_{\Lambda + \mathbf{z}_i^*, \sigma_0}$.

Let \mathbf{U}_{lef} (resp. \mathbf{U}_{rig}) denote the left L' (resp. right $2\ell - L'$) columns of \mathbf{U} . We are to show that the distribution $\mathbf{X}_{bot} \cdot D_{\Lambda + \mathbf{z}_i^*, \sigma_0}$ is $D_{\mathbb{Z}^{2\ell - L'}, \sigma_0 \sqrt{\Sigma}^{-1}, \mathbf{w}}$, where $\Sigma = \mathbf{U}_{rig}^t \mathbf{U}_{rig}$ and $\mathbf{w} = -\sqrt{\Sigma}^{-t} \mathbf{U}_{rig}^t \mathbf{U}_{lef} \mathbf{b}_i$. To see this, we first show that the supports of both distributions are the same. More specifically, we prove $\mathbf{X}_{bot} \cdot \Lambda = \mathbb{Z}^{2\ell - L'}$. To do so, it suffices to show that for any $\mathbf{a} \in \mathbb{Z}^{2\ell - L'}$, we have $\mathbf{a} \in \mathbf{X}_{bot} \cdot \Lambda$. By the construction of \mathbf{U} , we have $\mathbf{X}_{top}\mathbf{U}_{rig} = \mathbf{0}$ and $\mathbf{X}_{bot}\mathbf{U}_{rig} = \mathbf{I}_{2\ell - L'}$. Now, $\mathbf{a} \in \mathbf{X}_{bot} \cdot \Lambda$ follows because we have $\mathbf{a} = \mathbf{X}_{bot} \cdot (\mathbf{U}_{rig} \mathbf{a})$ and $\mathbf{X}_{top} \cdot \mathbf{U}_{rig} \mathbf{a} = \mathbf{0}$. We next evaluate the probability of $\mathbf{a} \in \mathbb{Z}^{2\ell - L'}$ being output by $\mathbf{X}_{bot} \cdot D_{\Lambda + \mathbf{z}_i^*, \sigma_0}$. This probability equals to the probability of $\mathbf{a}' \in \mathbb{Z}^{2\ell}$ being output by $D_{\Lambda + \mathbf{z}_i^*, \sigma_0}$ for \mathbf{a}' that is the unique vector in $\Lambda + \mathbf{z}_i^*$ satisfying $\mathbf{a} = \mathbf{X}_{bot} \cdot \mathbf{a}'$.

Since $\mathbf{a}' \in \Lambda + \mathbf{z}_i^*$, we have $\mathbf{X}_{top} \cdot \mathbf{a}' = \mathbf{X}_{top}(\mathbf{a}' - \mathbf{z}_i^*) + \mathbf{X}_{top}\mathbf{z}_i^* = \mathbf{0} + \mathbf{b}_i = \mathbf{b}_i$. Therefore, the vector \mathbf{a}' can be written as $\mathbf{a}' = \mathbf{X}^{-1}(\mathbf{b}_i \parallel \mathbf{a}) = \mathbf{U}(\mathbf{b}_i \parallel \mathbf{a}) = \mathbf{U}_{lef}\mathbf{b}_i + \mathbf{U}_{rig}\mathbf{a}$. The probability we consider is proportional to

$$\begin{aligned} & \exp(-\pi\|\mathbf{a}'\|^2/\sigma_0^2) \\ &= \exp(-\pi\|\mathbf{U}_{lef}\mathbf{b}_i + \mathbf{U}_{rig}\mathbf{a}\|^2/\sigma_0^2) \\ &= \exp\left(-\pi\|\sqrt{\Sigma}\mathbf{a} + \sqrt{\Sigma}^{-t}\mathbf{U}_{rig}^t\mathbf{U}_{lef}\mathbf{b}_i\|^2/\sigma_0^2\right) \\ &\cdot \underbrace{\exp\left(-\pi(\|\mathbf{U}_{lef}\mathbf{b}_i\|^2 - \|\sqrt{\Sigma}^{-t}\mathbf{U}_{rig}^t\mathbf{U}_{lef}\mathbf{b}_i\|^2)/\sigma_0^2\right)}_{\text{does not depend on } \mathbf{a}} \\ &\propto \exp\left(-\pi\|\sqrt{\Sigma}(\mathbf{a} - \mathbf{w})\|^2/\sigma_0^2\right). \end{aligned}$$

This implies this equals to the probability of \mathbf{a} being output by $D_{\mathbb{Z}^{2\ell-L'}, \sigma_0\sqrt{\Sigma}^{-1}, \mathbf{w}}$. To sum up, conditioned on $\mathbf{X}_{top}\mathbf{Z}$, the matrix $\mathbf{X}_{bot}\mathbf{Z}$ is distributed as $(D_{\mathbb{Z}^{2\ell-L'}, \sigma_0\sqrt{\Sigma}^{-1}, \mathbf{w}})^m$.

We then consider the joint distribution of $(\mathbf{A}, \mathbf{u}, \mathbf{Z}\mathbf{A}, \mathbf{X}_{bot}\mathbf{Z}\mathbf{u})$ (conditioned on $(\mathbf{X}_{top}, \mathbf{X}_{top}\mathbf{Z})$). In the following, let us consider the distribution of $\mathbf{XZ}\mathbf{A}$ instead of $\mathbf{Z}\mathbf{A}$. We do not lose any information by doing this since \mathbf{X} is invertible modulo q and the latter distribution can be recovered from the former by just multiplying \mathbf{X}^{-1} from the left. Furthermore, we observe that $\mathbf{XZ}\mathbf{A}$ is the vertical concatenation of $\mathbf{X}_{top}\mathbf{Z}\mathbf{A}$ and $\mathbf{X}_{bot}\mathbf{Z}\mathbf{A}$. Since the former can be recovered from $\mathbf{X}_{top}\mathbf{Z}$ and \mathbf{A} , which are already included in the tuples, we ignore the former.

Let us denote $\mathbf{Y} := \mathbf{X}_{bot}\mathbf{Z} \sim (D_{\mathbb{Z}^{2\ell-L'}, \sigma_0\sqrt{\Sigma}^{-1}, \mathbf{w}})^m$. To complete the proof, we will show that the following distributions are statistically close:

$$(\mathbf{A}, \mathbf{u}, \mathbf{Y}\mathbf{A}, \mathbf{Y}\mathbf{u}) \approx (\mathbf{A}, \mathbf{u}, \mathbf{B}, \mathbf{v})$$

where $\mathbf{B} \leftarrow \mathbb{Z}_q^{(2\ell-L') \times n}$, and $\mathbf{v} \leftarrow \mathbb{Z}_q^{2\ell-L'}$. We first show that $(\mathbf{Y} \bmod p)$ is within $2^{\Omega(-\lambda)}$ statistical distance from the uniform distribution over $\mathbb{Z}_p^{(2\ell-L') \times m}$. This follows by setting $\Lambda = \mathbb{Z}^{2\ell-L'}$ and $\Lambda' = p \cdot \mathbb{Z}^{2\ell-L'}$ and applying Lemma 1 to \mathbf{Y} in a column-wise manner. We check that the parameters satisfy the required condition of Lemma 1. We have

$$\begin{aligned} s_{2\ell-L'}(\sqrt{\Sigma}^{-1}) &= s_1(\Sigma)^{-1/2} \geq ((2\ell - L')^2 \cdot \|\Sigma\|_\infty)^{-1/2} \\ &\geq \Omega((\kappa^{1/2}\ell\sigma_2)^{-1}), \end{aligned}$$

where the last inequality follows from the upper bound on the norms of the rows of \mathbf{U} . We therefore have $\sigma_0 \cdot s_{2\ell-L'}(\sqrt{\Sigma}^{-1}) \geq p \cdot \Omega(\sqrt{\kappa})$ by our choice of σ_0 . We then finally apply Lemma 3 in a row-wise manner to obtain that $\mathbf{Y}(\mathbf{A} \parallel \mathbf{u})$ is almost uniformly random. We note that the lemma can be applicable because $q^{n+1}/p^m \leq 2^{-\Omega(\kappa)}$. This completes the proof of Lemma 10. \square

Proof: [Proof of Lemma 11] The reduction works as follows. Given $(\mathbf{A}, \mathbf{u}, \mathbf{Z}\mathbf{A}, \mathbf{X}_{top}, \mathbf{X}_{top}\mathbf{Z}, \mathbf{v})$, algorithm \mathcal{B} randomly guesses indices $\{i_j\}_{j \in [L']} \in [Q]^{L'}$ for which the adversary makes key queries. The public key $\text{pk} = (\mathbf{A}, \mathbf{U} = \mathbf{Z} \cdot \mathbf{A})$ and the master key $\text{msk} = \mathbf{Z}$ are determined by the given problem instance. (Note that \mathbf{Z} is not given to \mathcal{B} , so it is implicitly chosen.) Then \mathcal{B} chooses $\{\bar{\mathbf{x}}_i\}_{i \in [Q]}$ as follows. When $i \in \{i_j\}_{j \in [L']}$, there exists j such that $i = i_j$. Then algorithm \mathcal{B} sets $\bar{\mathbf{x}}_i^t$ to be the j -th row of the given matrix \mathbf{X}_{top} . Otherwise, it chooses $\bar{\mathbf{x}}_i$ as in **Game 4**. The key queries are handled as follows. Whenever \mathcal{A} queries key for \mathbf{x}_i such that $i \notin \{i_j\}_{j \in [L']}$, algorithm \mathcal{B} aborts and outputs a random bit. Other queries can be handled using $\mathbf{X}_{top}\mathbf{Z}$ in the problem instance. To create the challenge ciphertext \mathcal{B} sets

$$\mathbf{c}_0 = \mathbf{u}, \quad \mathbf{c}_1 = \mathbf{X}^{-1} \cdot (\mathbf{X}_{top}\mathbf{Z}\mathbf{u} \parallel \mathbf{v}) + \mathbf{e}_1 + \mathbf{y}_\beta.$$

We can observe that when $\mathbf{v} = \mathbf{X}_{bot}\mathbf{Z}\mathbf{u}$, we have $\mathbf{c}_1 = \mathbf{Z}\mathbf{u} + \mathbf{e}_1 + \mathbf{y}_\beta$ and the distribution of the challenge ciphertext corresponds to that of **Game 4**.

We then consider the case of \mathbf{v} is random. We will show that the distribution of $\mathbf{X} \cdot \vec{c}_1 \bmod q$ is independent of β . As the matrix \mathbf{X} is independent of $\beta \in \{0, 1\}$ and invertible over \mathbb{Z}_q , this implies that the distribution of \mathbf{c}_1 is independent of β as well (recall that \mathbf{X} is information theoretically known to \mathcal{A} , which means that, if \mathbf{c}_1 carries any information on β , so does $\mathbf{X} \cdot \mathbf{c}_1 \bmod q$). The first L' entries of $\mathbf{X} \cdot \mathbf{c}_1$ (namely, $\mathbf{X}_{top} \cdot \mathbf{c}_1$) do not depend on β because we have the equality $p^{k-1} \cdot \mathbf{X}_{top} \cdot \mathbf{y}_0 = p^{k-1} \cdot \mathbf{X}_{top} \cdot \mathbf{y}_1 \bmod q$, by construction of \mathbf{X}_{top} . The last $2\ell - L'$ entries are uniformly random, since they are masked by the random vector \mathbf{v} .

At the end of the game, algorithm \mathcal{B} outputs the same bit as \mathcal{A} .

It can be seen that \mathcal{B} perfectly simulates **Game 4** when $\mathbf{v} = \mathbf{X}_{bot}\mathbf{Z}\mathbf{u}$ and a game that is independent of β when \mathbf{v} is random. Therefore, conditioned on \mathcal{B} not aborting, the distinguishing advantage of \mathcal{B} is the same as \mathcal{A} . Since \mathcal{B} aborts and outputs a random bit with probability $1/Q^{L'}$, the advantage of \mathcal{B} is $\epsilon/Q^{L'}$. This completes the proof of Lemma 11. \square

5 Trace and Revoke from DDH and Paillier

In this section, we describe two (near) instantiations of the generic construction presented in the last section. We are not aware of existing IPFE schemes that meet the requirements for our generic Trace-and-Revoke construction, but some existing ones can be made to fit the framework.

5.1 Trace and Revoke from DDH

Following the work of Abdalla *et al.* Abdalla et al. (2015), two DDH-based adaptively secure IPFEs modulo the group size q have been proposed Agrawal et al. (2016); Benhamouda et al. (2017). However, these schemes enjoy limited correctness: as decryption involves the computation of a discrete logarithm, one restricts the set of exponents to be small. For instance, one may design the schemes so that inner products that are small compared to q can be decrypted. This restriction seems incompatible with the requirements of our trace-and-revoke scheme, as the inner product $m \cdot \langle \mathbf{x}_{id}, \mathbf{v}_{\mathcal{R}} \rangle$ occurring in the decryption algorithm has no reason to be small compared to p , even if the plaintext m is small. In the DDH-based trace-and-revoke scheme below, we circumvent the issue for the scheme from Agrawal et al. (2016) by removing the $\langle \mathbf{x}_{id}, \mathbf{v}_{\mathcal{R}} \rangle$ component before taking the discrete logarithm.

- **Setup**($1^\lambda, 1^t, 1^r, L$). Choose a cyclic group \mathbb{G} of prime order q along with two generators $g, h \leftarrow \mathbb{G}$. DDH in \mathbb{G} should be 2^λ -hard, but taking base- g logarithms of elements g^x with $x \in \{1, \dots, L\}$ should be tractable. Set $\ell = t + r + 1$. For each $i \leq \ell$, sample $s_i, t_i \leftarrow \mathbb{Z}_q$ and compute $h_i = g^{s_i} \cdot h^{t_i}$. Define

$$\text{msk} := (\mathbf{s}, \mathbf{t}) \text{ and } \text{pk} := (\mathbb{G}, g, h, \{h_i\}_{i \in [\ell]}).$$

- **KeyGen**(msk, id). Sample $\mathbf{x}_{id} \leftarrow \mathbb{Z}_q^\ell$. Set $\text{sk}_{id} = (\langle \mathbf{x}_{id}, \mathbf{s} \rangle, \langle \mathbf{x}_{id}, \mathbf{t} \rangle) \in \mathbb{Z}_q^2$ and $p_{id} = \mathbf{x}_{id}$.
- **Enc**(pk, \mathcal{R} , m) proceeds as follows to encrypt $m \in \{1, \dots, L\}$.
 1. Compute $\mathbf{v}_{\mathcal{R}} \in \mathbb{Z}_q^\ell \setminus \{\vec{0}\}$ such that $\langle \mathbf{x}_{id}, \mathbf{v}_{\mathcal{R}} \rangle = 0$ for every $\text{id} \in \mathcal{R}$.
 2. Set $\mathbf{y} = m \cdot \mathbf{v}_{\mathcal{R}}$ and sample $r \leftarrow \mathbb{Z}_q$.
 3. Compute $D_0 = g^r$, $D_1 = h^r$ and $E_i = g^{y_i} \cdot h_i^r$ for all $i \leq \ell$.

The ciphertext C is $(D_0, D_1, E_1, \dots, E_\ell, \mathcal{R})$.

- $\text{Dec}(\text{sk}_{\text{id}}, C)$. Write $C = (D_0, D_1, E_1, \dots, E_\ell, \mathcal{R})$ and let \mathbf{x}_{id} denote the vector corresponding to $\text{sk}_{\text{id}} = (s_x, t_x)$. Compute:

$$C_{\mathbf{x}_{\text{id}}} = \left(\prod_{i=1}^{\ell} E_i^{x_{\text{id},i}} \right) / (D_0^{s_x} \cdot D_1^{t_x}).$$

Then, recover $\mathbf{v}_{\mathcal{R}}$ from \mathcal{R} and output the base- g logarithm of $C_{\mathbf{x}_{\text{id}}}^{1/\langle \mathbf{v}_{\mathcal{R}}, \mathbf{x}_{\text{id}} \rangle}$.

- $\text{Trace}(\text{pd}, \mathcal{S}, \mathcal{R}, \mathcal{O}^{\mathcal{D}})$ proceeds as described in Section 3.

Correctness follows from elementary computations. The only difference with the direct instantiation of our trace-and-revoke construction is that the division by $\langle \mathbf{v}_{\mathcal{R}}, \mathbf{x}_{\text{id}} \rangle$ occurs before the computation of the discrete logarithm, hence enabling efficient decryption.

Key and Ciphertext Sizes. Both the public key and master secret key consist of $O((t+r)\log q)$ bits. To every user id corresponds a secret key sk_{id} of bit-length $O(\log q)$ and a vector p_{id} of bit-length $O((t+r)\log q)$. Algorithm Enc maps a plaintext in $\mathbb{Z}_q \setminus \{0\}$ to a ciphertext of bit-length $O((t+r)\log q)$. If we choose the DDH group as an elliptic curve group (without pairings), we may set $\log q = O(\lambda)$.

5.2 Trace and Revoke from Paillier

In Agrawal et al. (2016), Agrawal *et al.* described two IPFEs relying on the algebraic framework of Paillier’s encryption scheme Paillier (1999). One scheme handles inner products of short integers vectors, while the other handles inner products modulo a product $N = p \cdot q$ of two safe primes. Both are proved secure under the Decision Composite Residuosity (DCR) hardness assumption. We explain here how to instantiate our trace-and-revoke construction using this IPFE over \mathbb{Z}_N^ℓ .

A first difficulty is the fact that the Key Generation algorithm is stateful. However, this issue can be handled by noticing that for random queries, the key generation algorithm can be made stateless (see Agrawal et al. (2016) and Subsection 4.1 for more details). A further difficulty is the non-primality of N : our transformation requires the modulus to be prime. We may actually apply the transformation and “pretend” that N is prime, both in the scheme and in its security proof. The non-primality of N can be noticed only when finding vectors orthogonal modulo N to some specified vectors. When such a task is performed, either the linear algebra operations proceed and find such a vector, or they fail. In the latter case, a non-trivial factor of N has been found, which leads to an algorithm against DCR. Hence, under the DCR hardness assumption, such an event is unlikely. We now describe the resulting DCR-based trace-and-revoke scheme.

- $\text{Setup}(1^\lambda, 1^t, 1^r)$. Choose safe prime numbers $p = 2p' + 1, q = 2q' + 1$ with sufficiently large primes $p', q' > 2^{\text{Poly}(\lambda)}$, and compute $N = pq$. Then, sample $g' \leftarrow \mathbb{Z}_{N^2}^*$ and compute $g = g'^{2N} \bmod N^2$, which generates the subgroup of $(2N)$ -th residues in $\mathbb{Z}_{N^2}^*$ with overwhelming probability. Set $\ell = t + r + 1$ and sample \mathbf{s} from the integer Gaussian distribution $D_{\mathbb{Z}^\ell, \sigma}$ with standard deviation parameter σ satisfying $\sigma \geq \sqrt{\ell N \text{Poly}(\lambda)}$. Compute $h_i = g^{s_i} \bmod N^2$ for all $i \leq \ell$. Define

$$\text{msk} := \mathbf{s} \quad \text{and} \quad \text{pk} := (N, g, \mathcal{G}, \{h_i\}_{i \in [\ell]}).$$

- $\text{KeyGen}(\text{msk}, \text{id})$. Sample $\mathbf{x}_{\text{id}} \in \mathbb{Z}^\ell$ with coefficients i.i.d. uniform in $\{0, \dots, N-1\}$. Set $\text{sk}_{\text{id}} = \langle \mathbf{x}_{\text{id}}, \mathbf{s} \rangle \in \mathbb{Z}$ and $p_{\text{id}} = \mathbf{x}_{\text{id}}$.
- $\text{Enc}(\text{pk}, \mathcal{R}, m)$ proceeds as follows to encrypt $m \in \mathbb{Z}_N \setminus \{0\}$.
 1. Compute $\mathbf{v}_{\mathcal{R}} \in \mathbb{Z}_p^\ell \setminus \{\vec{0}\}$ such that $\langle \mathbf{x}_{\text{id}}, \mathbf{v}_{\mathcal{R}} \rangle = 0$ for every $\text{id} \in \mathcal{R}$.

2. Set $\mathbf{y} = m \cdot \mathbf{v}_{\mathcal{R}}$ and sample $r \leftarrow \{0, \dots, \lfloor N/4 \rfloor\}$.
3. Compute $C_0 = g^r \bmod N^2$ and $C_i = (1 + y_i N) \cdot h_i^r \bmod N^2$ for all $i \leq \ell$.

The ciphertext C is $(C_0, C_1, \dots, C_\ell, \mathcal{R})$.

- $\text{Dec}(\text{sk}_{\text{id}}, C)$. Write $C = (C_0, C_1, \dots, C_\ell, \mathcal{R})$ and let \mathbf{x}_{id} denote the vector corresponding to sk_{id} . Compute:

$$C_{\mathbf{x}_{\text{id}}} = C_0^{-\text{sk}_{\text{id}}} \cdot \left(\prod_{i=1}^{\ell} C_i^{x_{\text{id},i}} \right) \bmod N^2.$$

Then, recover $\mathbf{v}_{\mathcal{R}}$ from \mathcal{R} and output

$$\left(\frac{C_{\mathbf{x}_{\text{id}}} - 1 \bmod N^2}{N} \right) / \langle \mathbf{v}_{\mathcal{R}}, \mathbf{x}_{\text{id}} \rangle.$$

- $\text{Trace}(\text{pd}, \mathcal{S}, \mathcal{R}, \mathcal{O}^{\mathcal{D}})$ proceeds as described in Section 3.

We note that by exploiting the fact that the attacker makes random queries, we may improve the parameter sizes provided by Agrawal et al. (2016) exactly as in Subsection 4.2. In more detail, the proof of Theorem 5 (Appendix F) in Agrawal et al. (2016) can be modified to show that the advantage of the adversary in Game 3 is negligible even when σ is chosen as above, exactly as described in Subsection 4.2.

Key and Ciphertext Sizes. The public key and master secret key respectively consist of $O((t+r) \log N)$ and $O((t+r) \log N)$ bits. Note that the master secret key bit-length can be shrunk to $O(\lambda)$ by only storing the seed of the pseudo-random generator used to create it. In that case, the master secret key may be recomputed every time the KeyGen algorithm is called. Further, to every user id corresponds a secret key sk_{id} of bit-length $O((t+r) \log N)$ and a vector p_{id} of bit-length $O((t+r) \log N)$. Algorithm Enc maps a plaintext in $\mathbb{Z}_N \setminus \{0\}$ to a ciphertext of bit-length $O((t+r) \log N)$. To compensate for the number field sieve, we must choose $\log N = \tilde{\Omega}(\lambda^3)$.

Acknowledgements: We thank Benoît Libert for helpful discussions. The authors would also like to thank the anonymous referees for their valuable comments and helpful suggestions. Sanjay Bhattacharjee was funded by the “Programme Avenir Lyon Saint-Etienne de l’Université de Lyon” in the framework of the programme “Investissements d’Avenir” (ANR-11-IDEX-0007). Damien Stehlé was supported by the ERC Starting Grant ERC-2013-StG-335086-LATTAC. Duong Hieu Phan was supported by the ANR ALAMBIC (ANR-16-CE39-0006). Shota Yamada was supported by JSPS KAKENHI Grant Number 16K16068 and JST CREST Grant Number JPMJCR1688, Japan.

References

- Michel Abdalla, Florian Bourse, Angelo De Caro, and David Pointcheval. 2015. Simple Functional Encryption Schemes for Inner Products. In *PKC 2015 (LNCS)*. Springer, 733–751. https://doi.org/10.1007/978-3-662-46447-2_33
- Shweta Agrawal, Sanjay Bhattacharjee, Duong Hieu Phan, Damien Stehle, and Shota Yamada. 2017. Efficient Public Trace and Revoke from Standard Assumptions. Cryptology ePrint Archive, Report 2017/650. (2017). <http://eprint.iacr.org/2017/650>.

- Shweta Agrawal, Craig Gentry, Shai Halevi, and Amit Sahai. 2013. Discrete Gaussian Leftover Hash Lemma over Infinite Domains. In *Advances in Cryptology – ASIACRYPT 2013, Part I (LNCS)*, Vol. 8269. Springer, 97–116. https://doi.org/10.1007/978-3-642-42033-7_6
- Shweta Agrawal, Benoît Libert, and Damien Stehlé. 2016. Fully Secure Functional Encryption for Inner Products, from Standard Assumptions. In *Advances in Cryptology – CRYPTO 2016, Part III (LNCS)*, Vol. 9816. Springer, 333–362. https://doi.org/10.1007/978-3-662-53015-3_12
- Shweta Agrawal and Alon Rosen. 2016. Functional Encryption for Bounded Collusions, Revisited. Eprint. (2016).
- Shi Bai, Adeline Langlois, Tancrede Lepoint, Damien Stehlé, and Ron Steinfeld. 2015. Improved Security Proofs in Lattice-Based Cryptography: Using the Rényi Divergence Rather Than the Statistical Distance. In *Advances in Cryptology – ASIACRYPT 2015, Part I (LNCS)*, Vol. 9452. Springer, 3–24. https://doi.org/10.1007/978-3-662-48797-6_1
- Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil Vadhan, and Ke Yang. 2012. On the (Im)Possibility of Obfuscating Programs. *J. ACM* 59, 2 (May 2012).
- F. Benhamouda, F. Bourse, and H. Lipmaa. 2017. CCA-Secure Inner-Product Functional Encryption from Projective Hash Functions. In *Proc. of PKC (LNCS)*, Vol. 10175. Springer, 36–66.
- D. Boneh and M. K. Franklin. 1999. An Efficient Public Key Traitor Tracing Scheme. In *Proc. of CRYPTO (LNCS)*, Vol. 1666. Springer, 338–353.
- Dan Boneh, Amit Sahai, and Brent Waters. 2011. Functional Encryption: Definitions and Challenges. In *TCC 2011 (LNCS)*, Yuval Ishai (Ed.), Vol. 6597. Springer, 253–273.
- D. Boneh and B. Waters. 2006. A fully collusion resistant broadcast, trace, and revoke system. In *Proc. of ACM CCS*. ACM, 211–220.
- Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. 2013. Classical hardness of learning with errors. In *45th ACM STOC*. ACM Press, 575–584.
- B. Chor, A. Fiat, and M. Naor. 1994. Tracing Traitors. In *Proc. of CRYPTO (LNCS)*, Vol. 839. Springer, 257–270.
- Yevgeniy Dodis and Nelly Fazio. 2003. Public Key Trace and Revoke Scheme Secure against Adaptive Chosen Ciphertext Attack. In *PKC 2003 (LNCS)*, Yvo Desmedt (Ed.), Vol. 2567. Springer, 100–115.
- Amos Fiat and Moni Naor. 1993. Broadcast Encryption. In *Proceedings of the 13th Annual International Cryptology Conference on Advances in Cryptology (CRYPTO '93)*.
- Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. 2008. Trapdoors for hard lattices and new cryptographic constructions. In *40th ACM STOC*, Richard E. Ladner and Cynthia Dwork (Eds.). ACM Press, 197–206.
- S. Goldwasser, Y. Tauman Kalai, R. Popa, V. Vaikuntanathan, and N. Zeldovich. 2013. Reusable garbled circuits and succinct functional encryption. In *Proc. of STOC*. ACM Press, 555–564.
- Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. 2012. Functional Encryption with Bounded Collusions from Multiparty Computation. In *CRYPTO*.
- Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. 2013. Attribute-based encryption for circuits. In *45th ACM STOC*. ACM Press, 545–554.

- Dennis Hofheinz and Christoph Striecks. 2014. A Generic View on Trace-and-Revoke Broadcast Encryption Schemes. In *CT-RSA 2014 (LNCS)*. Springer, 48–63. https://doi.org/10.1007/978-3-319-04852-9_3
- Shuichi Katsumata and Shota Yamada. 2016. Partitioning via Non-linear Polynomial Functions: More Compact IBEs from Ideal Lattices and Bilinear Maps. In *Advances in Cryptology – ASIACRYPT 2016, Part II (LNCS)*, Vol. 10032. Springer, 682–712. https://doi.org/10.1007/978-3-662-53890-6_23
- A. Kiayias and M. Yung. 2001a. On Crafty Pirates and Foxy Tracers. In *Proc. of DRM Workshop (LNCS)*, Vol. 2320. Springer, 22–39.
- A. Kiayias and M. Yung. 2001b. Self Protecting Pirates and Black-Box Traitor Tracing. In *Proc. of CRYPTO (LNCS)*, Vol. 2139. Springer, 63–79.
- Aggelos Kiayias and Moti Yung. 2002. Traitor Tracing with Constant Transmission Rate. In *EUROCRYPT 2002 (LNCS)*, Lars R. Knudsen (Ed.), Vol. 2332. Springer, 450–465.
- Chong Hee Kim, Yong Ho Hwang, and Pil Joong Lee. 2003. An Efficient Public Key Trace and Revoke Scheme Secure against Adaptive Chosen Ciphertext Attack. In *ASIACRYPT 2003 (LNCS)*, Chi-Sung Laih (Ed.), Vol. 2894. Springer, 359–373. https://doi.org/10.1007/978-3-540-40061-5_23
- A. Langlois, D. Stehlé, and R. Steinfeld. 2014. GGHLite: More Efficient Multilinear Maps from Ideal Lattices. In *Proc. of EUROCRYPT (LNCS)*. Springer, 239–256.
- S. Ling, D. H. Phan, D. Stehlé, and R. Steinfeld. 2014. Hardness of k-LWE and Applications in Traitor Tracing. In *Proc. of CRYPTO (LNCS)*, Vol. 8616. Springer, 315–334.
- Daniele Micciancio and Petros Mol. 2011. Pseudorandom Knapsacks and the Sample Complexity of LWE Search-to-Decision Reductions. In *CRYPTO 2011 (LNCS)*, Phillip Rogaway (Ed.), Vol. 6841. Springer, 465–484.
- Dalit Naor, Moni Naor, and Jeffery Lotspiech. 2001. Revocation and Tracing Schemes for Stateless Receivers. In *CRYPTO 2001 (LNCS)*, Joe Kilian (Ed.), Vol. 2139. Springer, 41–62.
- M. Naor and B. Pinkas. 2000. Efficient Trace and Revoke Schemes. In *Proc. of Financial Cryptography (LNCS)*, Vol. 1962. Springer, 1–20.
- Hung Q. Ngo, Duong Hieu Phan, and David Pointcheval. 2013. Black-Box Trace&Revoke Codes. *Algorithmica* 67, 3 (2013), 418–448.
- Ryo Nishimaki, Daniel Wichs, and Mark Zhandry. 2016. Anonymous traitor tracing: how to embed arbitrary information in a key. In *Eurocrypt*. Springer, 388–419.
- Pascal Paillier. 1999. Public-Key Cryptosystems Based on Composite Degree Residuosity Classes. In *EUROCRYPT’99 (LNCS)*, Jacques Stern (Ed.), Vol. 1592. Springer, 223–238.
- Chris Peikert. 2009. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In *41st ACM STOC*, Michael Mitzenmacher (Ed.). ACM Press, 333–342.
- Duong Hieu Phan and Viet Cuong Trinh. 2011. Identity-Based Trace and Revoke Schemes. In *ProvSec 2011 (LNCS)*, Xavier Boyen and Xiaofeng Chen (Eds.), Vol. 6980. Springer, 204–221.
- Oded Regev. 2005. On lattices, learning with errors, random linear codes, and cryptography. In *37th ACM STOC*, Harold N. Gabow and Ronald Fagin (Eds.). ACM Press, 84–93.

- Amit Sahai and Brent R. Waters. 2005. Fuzzy Identity-Based Encryption. In *EUROCRYPT 2005 (LNCS)*, Ronald Cramer (Ed.), Vol. 3494. Springer, 457–473.
- D. R. Stinson and R. Wei. 1998a. Combinatorial Properties and Constructions of Traceability Schemes and Frameproof Codes. *SIAM J. Discrete Math.* 11, 1 (1998), 41–53.
- D. R. Stinson and R. Wei. 1998b. Key Preassigned Traceability Schemes for Broadcast Encryption. In *Proc. of SAC (LNCS)*, Vol. 1556. Springer, 144–156.

6 Additional Relevant Work

There are multiple parameters in trace-and-revoke systems that one desires to optimize such as security definition, hardness assumption, public traceability, collusion size, efficiency. The most general adaptive security definition for trace and revoke was provided by Boneh and Waters Boneh and Waters (2006). Here, the adversary is permitted to adaptively make key requests, and must finally submit a pirate decoder. For the adversary to win the game, pirate decoder must be *useful*, i.e., the challenger must be allowed to test it with various “probe” ciphertexts and these must be decrypted with non-negligible probability, and the tracing algorithm must be able to output at least one user whose key was not requested by the adversary.

Strong Security for Trace and Revoke. The definition of usefulness of the pirate decoder involves a subtlety – in the strongest definition, the pirate decoder may be queried with ciphertexts that may encode a set of maliciously chosen revoked users Boneh and Waters (2006). Most constructions do not satisfy this strong notion of security, indeed some schemes are actually insecure in this strong game.

For instance in the schemes Naor and Pinkas (2000); Dodis and Fazio (2003), a probe ciphertext may be distinguished from a normal ciphertext using a revoked key. In the polynomial interpolation based method in Naor and Pinkas (2000), in order to run tracing on a suspect set, the authority chooses a probe polynomial which agrees with the original polynomial on all the points in the suspected set. Therefore, if the suspected set contains all the traitor keys, then the pirate cannot detect this change from the original polynomial to the probe polynomial and the tracing works well. However, if the pirate knows one key (an evaluation of the original polynomial) in the revoked set, then it can detect this change. This means that a revoked key is useless in decrypting ciphertexts but useful in detecting the presence of a tracing procedure. Therefore, the tracing algorithm from Naor and Pinkas (2000) does not allow the adversary to choose and corrupt keys of the revoked set in the tracing game.

Combinatorial Schemes We remark that another line of work constructs *combinatorial schemes* Chor et al. (1994); Naor and Pinkas (2000); Stinson and Wei (1998a,b); Naor et al. (2001); Ngo et al. (2013), in contrast to the algebraic ones we have discussed so far; however these are usually less efficient than the algebraic candidates and the combination of trace and revoke is often studied in weaker security models.

Parameters Obtained with the NWZ Compiler. The NWZ compiler Nishimaki et al. (2016) may be instantiated with the bounded collusion functional encryption scheme from Gorbunov et al. (2012). This results in a scheme that has a ciphertext size that depends polynomially on the size of the circuit used by NWZ, as well as quartically on the collusion bound $r + t$. Since the circuit used by NWZ has an input size of $O(r + t)$, the ciphertext size grows at least as $O((r + t)^5 \text{Poly}(\lambda))$.

If the compiler is instantiated with the bounded collusion scheme of Goldwasser et al. (2013) (compiled with Gorbunov et al. (2012)), then the ciphertext size still grows as $O((r + t)^4 \text{Poly}(\lambda))$, and moreover relies on the subexponential hardness of learning with errors in addition to heavy hammers such as fully homomorphic

encryption and attribute based encryption. We note that the $\mathcal{Poly}(\lambda)$ factors above are unspecified, and possibly large: for instance, the circuit in Gorbunov et al. (2012) is represented using randomizing polynomials which adds a polynomial factor blow-up. Similarly, using the bounded collusion FE of Agrawal and Rosen (2016) leads to better asymptotic bounds $O(r + t)^3 \mathcal{Poly}(\lambda)$ but also suffers from large polynomial factors, since again the circuit is represented using randomizing polynomials. Here, a quadratic factor $(r + t)^2$ is incurred by the query dependence of Agrawal and Rosen (2016) and an additional factor $(r + t)$ is incurred due to circuit size dependence.