



Research Article

Fine-Grained Secure Attribute-Based Encryption*

Yuyu Wang

University of Electronic Science and Technology of China, Chengdu, China
wangyuyu@uestc.edu.cn

Jiixin Pan

Department of Mathematical Sciences, NTNU - Norwegian University of Science and Technology,
Trondheim, Norway
jiixin.pan@ntnu.no

Yu Chen

School of Cyber Science and Technology, Shandong University, Qingdao 266237, China
State Key Laboratory of Cryptology, P.O. Box 5159, Beijing 100878, China
Key Laboratory of Cryptologic Technology and Information Security, Ministry of Education, Shandong
University, Qingdao 266237, China
yuchen@sdu.edu.cn

Communicated by Alon Rosen

Received 29 August 2021 / Revised 18 July 2023 / Accepted 19 July 2023

Online publication 3 August 2023

Abstract. Fine-grained cryptography is constructing cryptosystems in a setting where an adversary's resource is a-prior bounded and an honest party has less resource than an adversary. Currently, only simple form of encryption schemes, such as secret-key and public-key encryption, are constructed in this setting. In this paper, we enrich the available tools in fine-grained cryptography by proposing the *first* fine-grained secure attribute-based encryption (ABE) scheme. Our construction is adaptively secure under the widely accepted worst-case assumption, $\text{NC}^1 \subseteq \oplus\text{L}/\text{poly}$, and it is presented in a generic manner using the notion of predicate encodings (Wee, TCC'14). By properly instantiating the underlying encoding, we can obtain different types of ABE schemes, including identity-based encryption. Previously, all of these schemes were unknown in fine-grained cryptography. Our main technical contribution is constructing ABE schemes without using pairing or the Diffie-Hellman assumption. Hence, our results show that, even if one-way functions do not exist, we still have ABE schemes with meaningful security. For more application of our techniques, we construct an efficient (quasi-adaptive) non-interactive zero-knowledge proof system.

Keywords. Fine-grained cryptography, Identity-based encryption, Attribute-based encryption, Quasi-adaptive non-interactive zero-knowledge proof.

*A preliminary version of this paper appeared at Crypto 2021 [34], this is the full version

1. Introduction

1.1. Motivation

Modern cryptography bases the security of schemes on assumptions, including the basic ones (such as the existence of one-way functions (OWFs)), the more advanced ones (such as the hardness of factoring, discrete logarithms, and some lattice problems), and the much more exotic ones (such as the existence of generic groups [25, 30] or algebraic groups [16]). Although there is some analysis on these assumptions, it is less desirable. We are interested in how to construct cryptography based on much mild assumptions or which form of security cryptography can be achieved if all classical assumptions (such as the existence of OWFs) do not hold.

Fine-grained cryptography is a direction in approaching the aforementioned problems. It aims at cryptography with weaker security in a setting where adversaries have only bounded resources and honest users have less resources than the adversaries. Under this setting it is possible to make the underlying assumption extremely mild, for instance, assuming $\text{NC}^1 \subsetneq \oplus\text{L}/\text{poly}$. This is a widely accepted worst-case assumption. As $\oplus\text{L}/\text{poly}$ is the class of languages with polynomial-sized branching programs and all languages in NC^1 have polynomial-sized branching programs of constant width [3], this assumption holds if there exists one language having only polynomial-sized branching programs of non-constant width. This is different from assuming the existence of OWFs which is an average-case assumption. It requires that the OWF be hard to invert on a random input. Hence, $\text{NC}^1 \subsetneq \oplus\text{L}/\text{poly}$ is more likely to be true.

The study on fine-grained cryptography was initialized by Merkle [26]. In the recent years, we are interested in which kind of cryptosystems can be constructed in this setting. We highlight the recent constructions of OWFs [8], symmetric-key and (leveled fully homomorphic) public-key encryption [9, 13], verifiable computation [9], hash proof systems (HPS) [14], and non-interactive zero-knowledge (NIZK) proof systems [2]. However, due to the restriction on running resources, many important primitives remain unknown. Surprisingly, digital signature schemes are among them, although they are implied by OWFs in the classical setting.

Our Goal: Fine-Grained Secure ABEs. We focus on constructing attribute-based encryption (ABE) schemes [18] with fine-grained security, since it has many applications and implies important primitives, including digital signatures. In an ABE scheme, messages are encrypted under descriptive values \mathbf{x} , secret keys are associated with values \mathbf{y} , and a secret key decrypts the ciphertext if and only if $\mathbf{p}(\mathbf{x}, \mathbf{y}) = 1$ for some boolean predicate \mathbf{p} . Here the predicate \mathbf{p} may express arbitrary access policy. This is in contrast to traditional public-key encryption (PKE) schemes without access control on data. Identity-based encryption [6, 12, 29] is a simplified version of ABE, where \mathbf{p} is the equality predicate, and it implies signatures in a natural manner (even in the fine-grained setting).

In general, it is challenging to construct ABEs. For instance, in the classical setting, it is shown that IBEs cannot be constructed using trapdoor permutations (TDP) or CCA-secure PKE schemes in a black-box manner [7]. Moreover, many pairing-based constructions of ABE and IBE (for instance, [5, 10]) heavily rely on the algebraic structures of pairing groups. These necessary structures are not available in fine-grained cryptography. Thus, in this paper, we develop new techniques to improve on the state of

the art of fine-grained cryptography, which only provides primitives related to TDP and CCA-secure PKE.

1.2. Our Contributions

We construct the *first* fine-grained secure ABE scheme. In particular, our scheme is computable in $\text{AC}^0[2]$ and secure against adversaries in NC^1 . Note that $\text{AC}^0[2] \subsetneq \text{NC}^1$ [28,31]. Similar to several existing NC^1 fine-grained primitives [9,13,14], the security of our scheme is based on the same worst-case assumption $\text{NC}^1 \subsetneq \oplus\text{L}/\text{poly}$. This is a widely accepted, weak assumption. For simplicity, we consider fine-grained cryptography as schemes with NC^1 honest users and adversaries and security based on $\text{NC}^1 \subsetneq \oplus\text{L}/\text{poly}$ in the rest of this paper.

Previously, fine-grained cryptography can only achieve symmetric-key and public-key encryption and HPS. Our work enriches its available tools and brings fine-grained cryptography closer to classical cryptography in terms of functionality.

In particular, our construction is presented in a generic manner using predicate encodings [10,36]. Hence, by suitably instantiating the underlying encoding, we directly obtain a fine-grained IBE scheme (which in turn implies a fine-grained signature scheme), fine-grained ABEs for inner-product encryption, non-zero inner-product encryption, spatial encryption, doubly spatial encryption, boolean span programs, and arithmetic span programs, and also fine-grained broadcast encryption and fuzzy IBE schemes. Prior to this work, it was unknown whether these primitives can be constructed in NC^1 based on a worst-case complexity assumption.

Finally, we use our technique to construct an efficient quasi-adaptive NIZK [22] with fine-grained security. Here “quasi-adaptive” means that common reference strings (CRSs) may depend on the language of the NIZK system.

Applications of Security Against NC^1 . Other than only relying on weak assumptions and running with low complexity, our results have the following applications.

Since security against NC^1 captures adversaries with limited parallel running-time, our constructions are well-suited for systems where attacks make sense only if they succeed in a short period of time. For example, our ABEs (and other fine-grained encryption primitives) can be used to protect messages that are only valuable in a short period of time, and that can be published or deleted later. As another example, the fine-grained signature (implied by our fine-grained IBE) and fine-grained QA-NIZK prevent adversaries from forging signatures and proofs with the running-time of an honest user, thereby ensuring security by letting the system reject users who have timed out when attempting to generate signatures or proofs. Moreover, as noted in [13], combining fine-grained primitives with standard ones immediately yield hybrids that are secure against NC^1 adversaries under $\text{NC}^1 \subsetneq \oplus\text{L}/\text{poly}$ and secure against polynomial time adversaries under stronger assumptions.

1.3. Technique Overview

We borrow the frameworks of the pairing-based constructions of IBEs in [5] and ABEs in [10] to upgrade the available fine-grained techniques [1,14,21] in achieving our goal. In a

nutshell, we transform a suitable symmetric-key primitive to an ABE in the fine-grained setting.

Previous frameworks in [5, 10] use pairings and the Diffie-Hellman assumptions. In contrast to them, our work develops new techniques to build ABEs without pairings or the Diffie-Hellman assumptions, but only under the mild assumption that $\text{NC}^1 \subsetneq \oplus\text{L}/\text{poly}$. For simplicity, we mostly focus on our techniques in the context of IBE here, and give some ideas about how they can be extended to construct ABEs. In this paper, we consider adaptive security where adversaries can adaptively request user secret keys and a challenge ciphertext.

The Approach of Blazy, Kiltz, and Pan, and Its Limitations in NC^1 . The “MAC \rightarrow IBE” transformation of BKP [5] is an abstraction of the Chen-Wee (CW) IBE scheme [11], and it also implements the “PRF \rightarrow Signature” framework by Bellare and Goldwasser (BG) [4] in the IBE context. The BKP transformation requires an “affine MAC”, namely, a MAC whose verification is done by checking a particular system of affine equations. Variables in these affine equations are included in the MAC secret key, and the (public) coefficients are derived from the message (which will be the identity of the resulting IBE scheme) to be signed. Such a MAC scheme can be constructed based on the Diffie-Hellman assumption which is generalized as the MDDH assumption.

We give some ideas about how an affine MAC can be turned into an IBE scheme. The master public key of an IBE scheme, $\text{pk} = \text{Com}(\text{sk}_{\text{MAC}})$, is a commitment of the MAC secret key, sk_{MAC} . A user secret key $\text{usk}[\text{id}]$ of an identity id consists of a BG signature, namely, a MAC tag τ_{id} on the message id and a NIZK proof of the validity of τ_{id} w.r.t. the secret key committed in pk .

Since the MAC verification consists of only affine equations, after implementing the aforementioned commitments and NIZK proofs with a (tuned) Groth-Sahai (GS) proof system [19],¹ the BKP IBE ciphertext ct_{id} can be viewed as a randomized linear combination of pk w.r.t. id . This is the key observation of BKP. The BKP framework can be further improved and extended to construct ABEs using predicate encodings [36] as in the CGW framework [10] by Chen, Gay, and Wee.

The MDDH assumption and the pairing-based GS proofs are two key ingredients for the BKP framework which are not available in fine-grained cryptography. One direction to resolve this is to develop a fine-grained GS proof system, but it is not clear what the counterpart of “pairing-product equations” will be. Instead, we achieve our goal with a simpler and more direct approach.

A Hard Subset Membership Problem for NC^1 circuits. We first need to find a counterpart of the MDDH assumption in NC^1 , since the separation assumption $\text{NC}^1 \subsetneq \oplus\text{L}/\text{poly}$ does not directly give us tools in constructing cryptographic schemes. In the work of [1, 21], it is shown that, if $\text{NC}^1 \subsetneq \oplus\text{L}/\text{poly}$ holds, then the following two distributions are indistinguishable for NC^1 circuits:

$$\underbrace{\{\mathbf{M}_0 \in \{0, 1\}^{n \times n} : \mathbf{M}_0 \xleftarrow{\$} \text{ZeroSamp}(n)\}}_{=D_0} \text{ and } \underbrace{\{\mathbf{M}_1 \in \{0, 1\}^{n \times n} : \mathbf{M}_1 \xleftarrow{\$} \text{OneSamp}(n)\}}_{=D_1}$$

¹Essentially, the BKP framework used the GS proof for linear equations and replaced the GS commitment with the Pedersen commitment.

where $n = n(\lambda)$ is some polynomial in security parameter λ , and the randomized sampling algorithms **ZeroSamp** and **OneSamp** output matrices with rank $n - 1$ and full rank, respectively. Concrete definitions of these algorithms are given in Sect. 2.2, and they are not relevant in this section.

This indistinguishability implies a hard subset membership problem in NC^1 implicitly given by Egashira, Wang, and Tanaka [15] for their HPS: Given a matrix \mathbf{M}^\top from D_0 and a random vector \mathbf{t} in two specific distributions represented by \mathbf{M} , the task of the problem is to tell whether \mathbf{t} is in the span of \mathbf{M} .

Our IBE in NC^1 . Our main technical contribution is a new approach of using the subset membership problem to transform an affine MAC to IBEs in the fine-grained setting. Our starting point is constructing a secure affine MAC in NC^1 . We prove that, if the subset membership problem is hard in NC^1 , then our MAC is secure for NC^1 adversaries.

Next, we propose a generic construction of IBE based on affine MACs, following the BKP framework. In stark contrast to the BKP, our construction does not require pairings. Essentially, we develop a Groth-Sahai-like proof system in NC^1 to prove the validity of our affine MAC. This proof system allows us to show that if our affine MAC is secure then our resulting IBE is secure in NC^1 . At the core of our proof system is a new commitment scheme in NC^1 , for which we achieve the hiding property by exploiting the concrete structure of matrices in D_0 .

We give more details about the security proof. Firstly, the zero-knowledge property allows us to generate user secret keys for adversaries without knowing the MAC secret key. Secondly, we show that if an adversary can break the adaptive security of our IBE, then we can construct a reduction to break the security of our affine MAC. This is a crucial step, and we require some extractability of the proof system to extract the MAC forgery from the IBE adversary. In the BKP framework, this extractability can be achieved by computing the inversion of some matrix $\mathbf{A} \in \mathbb{Z}_q^{k \times k}$ for some positive integer k . However, in our setting, inverting a matrix in $\{0, 1\}^{n \times n}$ is impossible, otherwise, this will lead to a distinguisher for the subset membership problem in NC^1 . Also, there is no known way to sample a matrix with its inverse efficiently [14]. To solve it, our proof system develop a new method in achieving this extractability without inverting any matrix. Our core idea is to prove that with a fresh random string $\mathbf{r} \xleftarrow{\$} \{0\} \times \{0, 1\}^{n-1}$, it is possible to extract the forgery from our NC^1 -commitments by switching the distribution of the public parameter $\mathbf{A} \in D_0$ twice (from D_0 to D_1 and then back to D_0) and changing the distribution of \mathbf{r} during the switching procedure.

Dual System Methodology in NC^1 and ABE. Our techniques for IBE can also be viewed as the dual system encryption methodology [35] in NC^1 , which is an alternative interpretation of our approach. In our proof, there are two important technical steps, switching ciphertexts to invalid and randomizing MAC tags in the user secret keys. These correspond to switching ciphertexts and user secret keys from functional to semi-functional in the dual system encryption methodology [5, 10, 24, 35]. Dual system methodology is very useful in constructing predicate encryption and it was only known with pairings. Our work is for the first time implementing the dual system methodology without pairings.

Similar to the extension from BKP-IBE [5] to CGW-ABE [10], we further extend our techniques in constructing ABEs. We first use (part of) a predicate encoding scheme [10, 36] to generalize the notion of affine MAC and make it useful for constructing ABEs.

After that, we upgrade our IBE techniques, and transform the generalized affine MAC to an adaptively secure ABE in NC^1 via the rest part of the predicate encoding scheme. Here, the predicate encoding scheme is to construct ABEs in a modular and generic way, in particular, it can generalize the encoding of different ABEs (such as identity-based encryption and inner-product encryption).

More Extension and Open Problem. We are optimistic that our approach can yield many more new public-key schemes in fine-grained cryptography. In particular, we show that our techniques can also be used to construct an efficient QA-NIZK in NC^1 with adaptive soundness in ‘‘Appendix B’’. Roughly, we use the technique for proving the hiding property of the underlying commitment scheme in our IBE scheme to achieve adaptive soundness.

Also, we are optimistic that our approach can be used to construct hierarchical IBE [17,20]. We leave a detailed treatment of it as an open problem.

1.4. Comparison with the Proceedings Version

This is the full version of the paper appeared at Crypto 2021 [34]. In this full version, we give the full proof for the security of the fine-grained IBE scheme in Sect.4.2. Additionally, we give the proof of Theorem 2.13, the definition, constructions, and security proofs of the fine-grained QA-NIZKs (with the comparison to previous and subsequent fine-grained NIZKs [2,32,33]), and the instantiations of predicate encodings, in Appendices A to C.

2. Preliminaries

Notations. We note that all arithmetic computations are over $GF(2)$ in this work. Namely, all arithmetic computations are performed with a modulus of 2. We write $a \stackrel{\$}{\leftarrow} \mathcal{A}(b)$ (respectively, $a = \mathcal{A}(b)$) to denote the random variable outputted by a probabilistic (respectively, deterministic) algorithm \mathcal{A} on input b . By $x \stackrel{\$}{\leftarrow} \mathcal{S}$ we denote the process of sampling an element x from a set or distribution \mathcal{S} uniformly at random. By $\mathbf{x} \in \{0, 1\}^n$ we denote a column vector with size n and by, say, $\mathbf{x} \in \{1\} \times \{0, 1\}^{n-1}$ we mean that the first element of \mathbf{x} is 1. By $[n]$ we denote the set $\{1, \dots, n\}$. By x_i (respectively, \mathbf{x}_i) we denote the i th element of a vector \mathbf{x} (respectively, \mathbf{x}). By negl we denote an unspecified negligible function.

For a matrix $\mathbf{A} \in \{0, 1\}^{n \times t}$ with $\text{rank } t' < n$, we denote the sets $\{\mathbf{y} \mid \exists \mathbf{x} \text{ s.t. } \mathbf{y} = \mathbf{A}\mathbf{x}\}$ and $\{\mathbf{x} \mid \mathbf{A}\mathbf{x} = \mathbf{0}\}$ by $\text{Im}(\mathbf{A})$ (i.e., the span of \mathbf{A}) and $\text{Ker}(\mathbf{A})$ respectively. By $\mathbf{A}^\perp \in \{0, 1\}^{n \times (n-t')}$ we denote a matrix consisting of $n - t'$ linear independent column vectors in the kernel of \mathbf{A}^\top . Note that for any $\mathbf{y} \notin \text{Im}(\mathbf{A})$, we have $\mathbf{y}^\top \mathbf{A}^\perp \neq \mathbf{0}$. By $(a_{ij})_{i \in [l], j \in [m]}$

we denote the matrix $\begin{pmatrix} a_{11} & \cdots & a_{1m} \\ \vdots & \ddots & \vdots \\ a_{l1} & \cdots & a_{lm} \end{pmatrix}$. Let $\mathbf{A} = (a_{ij})_{i \in [l], j \in [m]}$ be an $l \times m$ matrix and

$\mathbf{B} = (\mathbf{B}_{ij})_{i \in [m], j \in [n]}$ be a large matrix consisting of $m \times n$ matrices \mathbf{B}_{ij} for all $i \in [m]$

and $j \in [n]$. By $h \odot \mathbf{A}$ we denote $(h \cdot a_{ij})_{i \in [l], j \in [m]}$ and by $\mathbf{A} \odot \mathbf{B}$ we denote

$$\left(\sum_{k=1}^m a_{ik} \odot \mathbf{B}_{kj} \right)_{i \in [l], j \in [n]} .$$

By \mathbf{M}_0^n , \mathbf{M}_1^n , and \mathbf{N}^n , we denote the following $n \times n$ matrices:

$$\mathbf{M}_0^n = \begin{pmatrix} 0 & \cdots & 0 & 0 \\ 1 & 0 & & 0 \\ 0 & 1 & \ddots & \vdots \\ \vdots & \vdots & \ddots & 0 \\ 0 & \cdots & 0 & 1 \end{pmatrix}, \mathbf{M}_1^n = \begin{pmatrix} 0 & \cdots & 0 & 1 \\ 1 & 0 & & 0 \\ 0 & 1 & \ddots & \vdots \\ \vdots & \vdots & \ddots & 0 \\ 0 & \cdots & 0 & 1 \end{pmatrix},$$

$$\mathbf{N}^n = \begin{pmatrix} 0 & \cdots & 0 \\ \vdots & 0 & \cdots & 0 \\ 0 & & \ddots & \vdots \\ 1 & 0 & \cdots & 0 \end{pmatrix},$$

and by $\mathbf{0}$ we denote a zero vector $(0, \dots, 0)^\top$.

Games. We follow [5] to use code-based games for defining and proving security. A game \mathbf{G} contains procedures INIT and FINALIZE, and some additional procedures P_1, \dots, P_n , which are defined in pseudo-code. All variables in a game are initialized as 0, and all sets are empty (denote by \emptyset). An adversary $\mathcal{A} = \{a_\lambda\}_{\lambda \in \mathbb{N}}$ is executed in game \mathbf{G} w.r.t. the security parameter λ (denote by \mathbf{G}^{a_λ}) if a_λ first calls INIT, obtaining its output. Next, it may make arbitrary queries to P_i (according to their specification) and obtain their output. Finally, it makes one single call to FINALIZE(\cdot) and stops. We use $\mathbf{G}^{a_\lambda} \Rightarrow d$ to denote that \mathbf{G} outputs d after interacting with a_λ , and d is the output of FINALIZE.

2.1. Function Families

In this section, we recall the definitions of function families, NC^1 circuits, $\text{AC}^0[2]$ circuits, and $\oplus\text{L}/\text{poly}$. Note that $\text{AC}^0[2] \subsetneq \text{NC}^1$ [28, 31].

Definition 2.1. (*Function Family*) A function family is a family of (possibly randomized) functions $\mathcal{F} = \{f_\lambda\}_{\lambda \in \mathbb{N}}$, where for each λ , f_λ has a domain D_λ^f and a range R_λ^f .

Definition 2.2. (NC^1) The class of (non-uniform) NC^1 function families is the set of all function families $\mathcal{F} = \{f_\lambda\}_{\lambda \in \mathbb{N}}$ for which there is a polynomial $p(\cdot)$ and constant c such that for each λ , f_λ can be computed by a (randomized) circuit of size $p(\lambda)$, depth $c \log(\lambda)$, and fan-in 2 using AND, OR, and NOT gates.

LSamp(n): For all $i, j \in [n]$ and $i < j$: $r_{i,j} \stackrel{\$}{\leftarrow} \{0, 1\}$ Return $\begin{pmatrix} 1 & r_{1,2} & \cdots & r_{1,n-1} & r_{1,n} \\ 0 & 1 & r_{2,3} & \cdots & r_{2,n} \\ 0 & 0 & \ddots & & \vdots \\ \vdots & \vdots & \ddots & 1 & r_{n-1,n} \\ 0 & \cdots & 0 & 0 & 1 \end{pmatrix}$	RSamp(n): For $i = 1, \dots, n-1$ $r_i \stackrel{\$}{\leftarrow} \{0, 1\}$ Return $\begin{pmatrix} 1 & \cdots & 0 & r_1 \\ 0 & 1 & & r_2 \\ 0 & 0 & \ddots & \vdots \\ \vdots & \vdots & \ddots & 1 & r_{n-1} \\ 0 & \cdots & 0 & 0 & 1 \end{pmatrix}$	ZeroSamp(n): $\mathbf{R}_0 \stackrel{\$}{\leftarrow} \text{LSamp}(n) \in \{0, 1\}^{n \times n}$ $\mathbf{R}_1 \stackrel{\$}{\leftarrow} \text{RSamp}(n) \in \{0, 1\}^{n \times n}$ Return $\mathbf{R}_0 \mathbf{M}_0^* \mathbf{R}_1 \in \{0, 1\}^{n \times n}$ OneSamp(n): $\mathbf{R}_0 \stackrel{\$}{\leftarrow} \text{LSamp}(n)$ $\mathbf{R}_1 \stackrel{\$}{\leftarrow} \text{RSamp}(n)$ Return $\mathbf{R}_0 \mathbf{M}_1^* \mathbf{R}_1 \in \{0, 1\}^{n \times n}$
--	--	---

Fig. 1. Definitions of LSamp, RSamp, ZeroSamp, and OneSamp. $n = n(\lambda)$ is a polynomial in the security parameter λ .

Definition 2.3. ($\text{AC}^0[2]$) The class of (non-uniform) $\text{AC}^0[2]$ function families is the set of all function families $\mathcal{F} = \{f_\lambda\}_{\lambda \in \mathbb{N}}$ for which there is a polynomial $p(\cdot)$ and constant c such that for each λ , f_λ can be computed by a (randomized) circuit of size $p(\lambda)$, depth c , and unbounded fan-in using AND, OR, NOT, and PARITY gates.

One can see that multiplication of a constant number of matrices can be performed in $\text{AC}^0[2]$, since it can be done in constant depth with PARITY gates.

Definition 2.4. ($\oplus\text{L}/\text{poly}$) $\oplus\text{L}/\text{poly}$ is the set of all boolean function families $\mathcal{F} = \{f_\lambda\}_{\lambda \in \mathbb{N}}$ for which there is a constant c such that for each λ , there is a non-deterministic Turing machine \mathcal{M}_λ such that for each input x with length λ , $\mathcal{M}_\lambda(x)$ uses at most $c \log(\lambda)$ space, and $f_\lambda(x)$ is equal to the parity of the number of accepting paths of $\mathcal{M}_\lambda(x)$.

2.2. Sampling Procedure

We now recall the definitions of four sampling procedures LSamp, RSamp, ZeroSamp, and OneSamp in Fig. 1. Note that the output of ZeroSamp(n) is always a matrix of rank $n - 1$ and the output of OneSamp(n) is always a matrix of full rank [13].

We now recall several assumptions and lemmata on ZeroSamp and OneSamp given in [13].

Definition 2.5. (*Fine-grained matrix linear assumption* [13]) There exists a polynomial $n = n(\lambda)$ in the security parameter λ such that for any family $\mathcal{A} = \{a_\lambda\}_{\lambda \in \mathbb{N}}$ in NC^1 , we have

$$\begin{aligned} & \left| \Pr [a_\lambda(\mathbf{M}) = 1 \mid \mathbf{M} \stackrel{\$}{\leftarrow} \text{ZeroSamp}(n)] \right. \\ & \quad \left. - \Pr [a_\lambda(\mathbf{M}') = 1 \mid \mathbf{M}' \stackrel{\$}{\leftarrow} \text{OneSamp}(n)] \right| \leq \text{negl}(\lambda). \end{aligned}$$

Lemma 2.6. (Lemma 4.3 in [13]) *If $\text{NC}^1 \subsetneq \oplus\text{L}/\text{poly}$, then the fine-grained matrix linear assumption holds.*

Remark. Notice that for any polynomial $n = n(\lambda)$, we have $\{f_n\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$ iff $\{f_\lambda\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$ since $O(\log(n(\lambda))) = O(\log(\lambda))$. Hence, in the above lemma, we

can also set $n(\cdot)$ as an identity function, i.e., $n = \lambda$. For simplicity, in the rest of the paper, we always let $\text{ZeroSamp}(\cdot)$ and $\text{OneSamp}(\cdot)$ take as input λ .

The following lemma implies that for a matrix \mathbf{M}^\top sampled by $\text{ZeroSamp}(\lambda)$, there is a unique non-zero vector with the first (respectively, last) element being 1 in the kernel of \mathbf{M} (respectively, \mathbf{M}^\top).

Lemma 2.7. (Lemma 3 in [15]) *For all $\lambda \in \mathbb{N}$ and all $\mathbf{M}^\top \in \text{ZeroSamp}(\lambda)$, it holds that $\text{Ker}(\mathbf{M}^\top) = \{\mathbf{0}, \mathbf{k}\}$ where \mathbf{k} is a vector such that $\mathbf{k} \in \{0, 1\}^{\lambda-1} \times \{1\}$.*

Lemma 2.8. (Lemma 4 in [15]) *For all $\lambda \in \mathbb{N}$ and all $\mathbf{M}^\top \in \text{ZeroSamp}(\lambda)$, it holds that $\text{Ker}(\mathbf{M}) = \{\mathbf{0}, \mathbf{k}\}$ where \mathbf{k} is a vector such that $\mathbf{k} \in \{1\} \times \{0, 1\}^{\lambda-1}$.*

The following lemma indicates a simple relation between the distributions of the outputs of $\text{ZeroSamp}(\lambda)$ and $\text{OneSamp}(\lambda)$.

Lemma 2.9. (Lemma 7 in [15]) *For all $\lambda \in \mathbb{N}$, the distributions of $\mathbf{M} + \mathbf{N}^\lambda$ and \mathbf{M}' are identical, where $\mathbf{M}^\top \stackrel{\$}{\leftarrow} \text{ZeroSamp}(\lambda)$ and $\mathbf{M}'^\top \stackrel{\$}{\leftarrow} \text{OneSamp}(\lambda)$.*

We now give two lemmata showing that when sampling a random vector \mathbf{w} from $\{0, 1\}^\lambda$, the first element of \mathbf{w} does not affect the distribution of $\mathbf{M}\mathbf{w}$ for $\mathbf{M}^\top \in \text{ZeroSamp}(\lambda)$.

Lemma 2.10. (Lemma 5 in [15]) *For all $\lambda \in \mathbb{N}$ and all $\mathbf{M}^\top \in \text{ZeroSamp}(\lambda)$, it holds that*

$$\text{Im}(\mathbf{M}) = \left\{ \mathbf{x} \mid \mathbf{w} \in \{0\} \times \{0, 1\}^{\lambda-1}, \mathbf{x} = \mathbf{M}\mathbf{w} \right\} = \left\{ \mathbf{x} \mid \mathbf{w} \in \{1\} \times \{0, 1\}^{\lambda-1}, \mathbf{x} = \mathbf{M}\mathbf{w} \right\}.$$

Lemma 2.11. *For all $\lambda \in \mathbb{N}$ and all $\mathbf{M}^\top \in \text{ZeroSamp}(\lambda)$, the distributions of \mathbf{x} and \mathbf{x}' are identical, where $\mathbf{w} \stackrel{\$}{\leftarrow} \{0\} \times \{0, 1\}^{\lambda-1}$, $\mathbf{w}' \stackrel{\$}{\leftarrow} \{1\} \times \{0, 1\}^{\lambda-1}$, $\mathbf{x} = \mathbf{M}\mathbf{w}$, and $\mathbf{x}' = \mathbf{M}\mathbf{w}'$.*

Proof. According to Lemma 2.8, for any $\mathbf{M}^\top \in \text{ZeroSamp}(\lambda)$, there exists $\mathbf{k} \in \text{Ker}(\mathbf{M})$ such that $\mathbf{k} \in \{1\} \times \{0, 1\}^{\lambda-1}$. Therefore, the distributions of $(\mathbf{w} + \mathbf{k})$, where $\mathbf{w} \stackrel{\$}{\leftarrow} \{0\} \times \{0, 1\}^{\lambda-1}$, and $\mathbf{w}' \stackrel{\$}{\leftarrow} \{1\} \times \{0, 1\}^{\lambda-1}$ are identical. Moreover, we have $\mathbf{M}\mathbf{w} = \mathbf{M}(\mathbf{w} + \mathbf{k})$. Hence, the distributions of $\mathbf{M}\mathbf{w}$ and $\mathbf{M}\mathbf{w}'$ are identical, completing the proof of Lemma 2.11. \square

Below we recall the a theorem implicitly given in [15] as the subset membership problem for an HPS. Roughly, it shows that for $\mathbf{M}^\top \stackrel{\$}{\leftarrow} \text{ZeroSamp}(\lambda)$, a vector sampled from the span of \mathbf{M} is indistinguishable from one sampled outside the span of \mathbf{M} for any adversary in NC^1 . The proof of this theorem is given in ‘‘Appendix A’’ for completeness.

Definition 2.12. (Fine-grained subset membership problem [15]) Let $\text{SY} = \{\text{SampYes}_\lambda\}_{\lambda \in \mathbb{N}}$ and $\text{SN} = \{\text{SampNo}_\lambda\}_{\lambda \in \mathbb{N}}$ be function families described in Fig. 2. For all $\lambda \in \mathbb{N}$, all $\mathbf{M}^\top \in \text{ZeroSamp}(\lambda)$, and all $\mathbf{x} \in \text{SampNo}_\lambda(\mathbf{M})$, we have $\mathbf{x} \in \{0, 1\}^\lambda \setminus \text{Im}(\mathbf{M})$, then

$\text{SampYes}_\lambda(\mathbf{M} \in \{0, 1\}^{\lambda \times \lambda}):$ $\mathbf{w} \stackrel{\$}{\leftarrow} \{1\} \times \{0, 1\}^{\lambda-1}$ Return $\mathbf{x} = \mathbf{M}\mathbf{w}$	$\text{SampNo}_\lambda(\mathbf{M} \in \{0, 1\}^{\lambda \times \lambda}):$ $\mathbf{w} \stackrel{\$}{\leftarrow} \{1\} \times \{0, 1\}^{\lambda-1}$ Return $\mathbf{x} = (\mathbf{M} + \mathbf{N}^\lambda)\mathbf{w}$.
---	---

Fig. 2. Definitions of SY and SN. Note that SY, SN $\in \text{AC}^0[2]$, since they only involve operations including sampling random bits and multiplication of a matrix and a vector.

for $\mathbf{M}^\top \stackrel{\$}{\leftarrow} \text{ZeroSamp}(\lambda)$ and any adversary $\mathcal{A} = \{a_\lambda\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$, we have

$$\left| \Pr [a_\lambda(\mathbf{x}, \mathbf{M}) = 1 \mid \mathbf{x} \stackrel{\$}{\leftarrow} \text{SampYes}_\lambda(\mathbf{M})] - \Pr [a_\lambda(\mathbf{x}, \mathbf{M}) = 1 \mid \mathbf{x} \stackrel{\$}{\leftarrow} \text{SampNo}_\lambda(\mathbf{M})] \right| \leq \text{negl}(\lambda).$$

Theorem 2.13. ([15]) *If $\text{NC}^1 \subsetneq \oplus\text{L}/\text{poly}$, then the fine-grained subset membership problem (see Definition 2.12) holds.*

Remark. Note that the subset membership problem in [15] gives a stronger result additionally showing that the output distributions of $\text{SampYes}_\lambda(\mathbf{M})$ and $\text{SampNo}_\lambda(\mathbf{M})$ are identical to the uniform distributions over $\text{Im}(\mathbf{M})$ and $\{0, 1\}^\lambda \setminus \text{Im}(\mathbf{M})$ respectively. We only need a weak form of it in this work.

2.3. Predicate Encodings

We now recall the definition of predicate encodings. As in [10], our resulting construction of ABE is generally based on a predicate encoding. By exploiting various types of encodings, we can achieve a broad class of ABEs.

Our definitions are slightly different from the original definition in [10], in that our definition is over $GF(2)$ rather than $GF(p)$, and we require that the encodings are performed in a circuit class \mathcal{C}_1 .

Definition 2.14. (*Predicate Encoding* [10]) Let $\mathbf{P} = \{\mathbf{p}_\lambda\}_{\lambda \in \mathbb{N}}$ with $\mathbf{p}_\lambda : \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$ be a predicate, where \mathcal{X} and \mathcal{Y} are polynomial-sized spaces associated with λ . An \mathcal{C}_1 -predicate encoding for \mathbf{P} is a function family $\text{PE} = \{\mathbf{rE}_\lambda, \mathbf{kE}_\lambda, \mathbf{sE}_\lambda, \mathbf{sD}_\lambda, \mathbf{rD}_\lambda\}_{\lambda \in \mathbb{N}} \in \mathcal{C}_1$ with

- $\mathbf{rE}_\lambda : \mathcal{Y} \times \{0, 1\}^\ell \rightarrow \{0, 1\}^\eta$,
- $\mathbf{kE}_\lambda : \mathcal{Y} \times \{0, 1\} \rightarrow \{0, 1\}^\eta$,
- $\mathbf{sE}_\lambda : \mathcal{X} \times \{0, 1\}^\ell \rightarrow \{0, 1\}^\zeta$,
- $\mathbf{sD}_\lambda : \mathcal{X} \times \mathcal{Y} \times \{0, 1\}^\zeta \rightarrow \{0, 1\}$,
- $\mathbf{rD}_\lambda : \mathcal{X} \times \mathcal{Y} \times \{0, 1\}^\eta \rightarrow \{0, 1\}$,

where $\ell = \ell(\lambda)$, $\eta = \eta(\lambda)$, and $\zeta = \zeta(\lambda)$ are polynomials in λ .

Linearity is satisfied is for all $\lambda \in \mathbb{N}$ and all $(\mathbf{x}, \mathbf{y}) \in \mathcal{X} \times \mathcal{Y}$, $\mathbf{rE}_\lambda(\mathbf{y}, \cdot)$, $\mathbf{kE}_\lambda(\mathbf{y}, \cdot)$, $\mathbf{sE}_\lambda(\mathbf{x}, \cdot)$, $\mathbf{sD}_\lambda(\mathbf{x}, \mathbf{y}, \cdot)$, and $\mathbf{rD}_\lambda(\mathbf{x}, \mathbf{y}, \cdot)$ are $\{0, 1\}$ -linear. Namely, for any $\mathbf{y} \in \mathcal{Y}$, any $\mathbf{w}_0, \mathbf{w}_1 \in \{0, 1\}^\ell$, and any $c \in \{0, 1\}$, we have $\mathbf{rE}_\lambda(\mathbf{y}, \mathbf{w}_0 + \mathbf{w}_1 \cdot c) = \mathbf{rE}_\lambda(\mathbf{y}, \mathbf{w}_0) + \mathbf{rE}_\lambda(\mathbf{w}_1) \cdot c$, and the same argument can be made for \mathbf{kE}_λ , \mathbf{sE}_λ , \mathbf{sD}_λ , and \mathbf{rD}_λ .

$\mathcal{X} = \{0, 1\}^n, \mathcal{Y} = \{0, 1\}^n$ $\ell = (1 + n), \eta = 1, \zeta = 1$ $p_\lambda(\mathbf{x}, \mathbf{y})$: Return 1 iff $\mathbf{x} = \mathbf{y}$	$sE_\lambda(\mathbf{x}, (u, \mathbf{w}^\top)^\top) = u + \mathbf{x}^\top \mathbf{w}$ $rE_\lambda(\mathbf{y}, (u, \mathbf{w}^\top)^\top) = u + \mathbf{y}^\top \mathbf{w}$ $kE_\lambda(\mathbf{y}, \alpha) = \alpha$ $sD_\lambda(\mathbf{x}, \mathbf{y}, c) = c$ $rD_\lambda(\mathbf{x}, \mathbf{y}, d) = d$
--	---

Fig. 3. Definitions of $P_{\text{eq}} = \{p_\lambda\}_{\lambda \in \mathbb{N}}$ and $PE_{\text{eq}} = \{rE_\lambda, kE_\lambda, sE_\lambda, sD_\lambda, rD_\lambda\}$.

Restricted α -reconstruction is satisfied if for all $\lambda \in \mathbb{N}$, all $(\mathbf{x}, \mathbf{y}) \in \mathcal{X} \times \mathcal{Y}$ such that $p_\lambda(\mathbf{x}, \mathbf{y}) = 1$, all $\mathbf{w} \in \{0, 1\}^\ell$, and all $\alpha \in \{0, 1\}$, we have

$$rD_\lambda(\mathbf{x}, \mathbf{y}, rE_\lambda(\mathbf{y}, \mathbf{w})) = sD_\lambda(\mathbf{x}, \mathbf{y}, sE_\lambda(\mathbf{x}, \mathbf{w})) \text{ and } rD_\lambda(\mathbf{x}, \mathbf{y}, kE_\lambda(\mathbf{y}, \alpha)) = \alpha.$$

α -privacy is satisfied if for all $\lambda \in \mathbb{N}$, all $(\mathbf{x}, \mathbf{y}) \in \mathcal{X} \times \mathcal{Y}$ such that $p_\lambda(\mathbf{x}, \mathbf{y}) = 0$, and all $\alpha \in \{0, 1\}$, the following distributions are identical:

$$(\mathbf{x}, \mathbf{y}, \alpha, sE_\lambda(\mathbf{x}, \mathbf{w}), rE_\lambda(\mathbf{y}, \mathbf{w}) + kE_\lambda(\mathbf{y}, \alpha)) \text{ and } (\mathbf{x}, \mathbf{y}, \alpha, sE_\lambda(\mathbf{x}, \mathbf{w}), rE_\lambda(\mathbf{y}, \mathbf{w})),$$

where $\mathbf{w} \xleftarrow{\$} \{0, 1\}^\ell$.

Intuitively, in a modularly designed attribute-based encryption (ABE) scheme, the attribute value in the user's key is encoded by rE_λ and kE_λ , while that in the ciphertext is encoded by sE_λ . The decryption algorithm uses the associated decoding algorithms rD_λ and sD_λ to decode (rE_λ, kE_λ) and sE_λ respectively. The difference between the decoding results of rD_λ and sD_λ for (rE_λ, kE_λ) and sE_λ is the decoding result of rD_λ for kE_λ only, which is used to yield the session key. These encoding algorithms can be instantiated according to the predicates of different types of ABEs, thus allowing for a modular and generic approach to ABE construction. Namely, different ABE schemes can be constructed by plugging in different encoding algorithms, based on the desired access structure for the scheme.

Remark on Notions for Predicate Encodings. Similar to [10], we abuse the notion

$$rE_\lambda(\mathbf{x}, \mathbf{W}) \text{ where } \mathbf{W} = (\mathbf{w}_{ij})_{i \in [l], j \in [m]} \text{ and } \mathbf{w}_{ij} \in \{0, 1\}^\ell$$

for all i, j to denote the matrix

$$(rE_\lambda(\mathbf{x}, \mathbf{w}_{ij}))_{i \in [l], j \in [m]}.$$

The same argument is made for $(kE_\lambda, sE_\lambda, sD_\lambda, rD_\lambda)$.

Encoding for Equality. We now give an example of predicate encoding PE_{eq} for equality P_{eq} in Fig. 3. By instantiating our ABKEM given later in Sect. 5 with this encoding, we immediately achieve an IBKEM. Linearity is straightforward. Restricted α -reconstruction follows from the fact that $u + \mathbf{x}^\top \mathbf{w} = u + \mathbf{y}^\top \mathbf{w}$ when $\mathbf{x} = \mathbf{y}$, and α -privacy follows from the fact that $u + \mathbf{x}^\top \mathbf{w}$ and $u + \mathbf{y}^\top \mathbf{w}$ are pairwise independent if $\mathbf{x} \neq \mathbf{y}$.

<p><u>INIT:</u> $(pk, sk) \stackrel{\\$}{\leftarrow} \text{Gen}_\lambda$ Return pk</p> <p><u>USKGEN(y):</u> $\overline{//k(\lambda) \times l(\lambda)}$ queries $Q_y \stackrel{\\$}{\leftarrow} Q_y \cup \{y\}$ Return $usk[id] \stackrel{\\$}{\leftarrow} \text{USKGen}_\lambda(sk, y)$</p>	<p><u>ENC(x):</u> $\overline{//one query}$ $(K^*, ct^*) \stackrel{\\$}{\leftarrow} \text{Enc}_\lambda(pk, x)$ $K^* \stackrel{\\$}{\leftarrow} \mathcal{K}$ Return (K^*, ct^*)</p> <p><u>FINALIZE(β):</u> If $(p_\lambda(x, y) \neq 1$ for all $y \in Q_y$, return β Else return 0</p>
---	---

Fig. 4. Security Games $\text{PR-AT-CPA}_{\text{real}}$ and $\text{PR-AT-CPA}_{\text{rand}}$ for defining PR-AT-CPA security for ABKEM. The boxed statement redefining K^* is only executed in game $\text{PR-AT-CPA}_{\text{rand}}$.

2.4. Attribute-Based Key Encapsulation

We now give the definition of fine-grained ABKEM, the instantiation of which can be easily converted into ABEs by using a one-time symmetric cipher.

Definition 2.15. (*Attribute-based key encapsulation*) A \mathcal{C}_1 -attribute-based key encapsulation (ABKEM) scheme for a predicate $\mathbf{P} = \{p_\lambda\}_\lambda$ is a function family $\text{ABKEM} = \{\text{Gen}_\lambda, \text{USKGen}_\lambda, \text{Enc}_\lambda, \text{Dec}_\lambda\}_{\lambda \in \mathbb{N}} \in \mathcal{C}_1$ with the following properties.

- Gen_λ returns the (master) public/secret key (pk, sk) . We assume that pk implicitly defines value spaces \mathcal{X} and \mathcal{Y} , a key space \mathcal{K} , and a ciphertext space \mathcal{C} .
- $\text{USKGen}_\lambda(sk, y)$ returns a user secret-key $usk[y]$ for a value $y \in \mathcal{Y}$.
- $\text{Enc}_\lambda(pk, x)$ returns a symmetric key $K \in \mathcal{K}$ together with a ciphertext $ct \in \mathcal{C}$ w.r.t. $x \in \mathcal{X}$.
- $\text{Dec}_\lambda(usk[y], y, x, ct)$ deterministically returns a decapsulated key $K \in \mathcal{K}$ or the reject symbol \perp .

Perfect correctness is satisfied if for all $\lambda \in \mathbb{N}$, all $(pk, sk) \in \text{Gen}_\lambda$, all $y \in \mathcal{Y}$, all $x \in \mathcal{X}$, all $usk[y] \in \text{USKGen}_\lambda(sk, y)$, and all $(K, ct) \in \text{Enc}_\lambda(pk, x)$, if $p_\lambda(x, y) = 1$, we have

$$\Pr[\text{Dec}_\lambda(usk[y], y, x, ct) = K] = 1.$$

The security requirement we consider is indistinguishability against chosen plaintext and attribute attacks (PR-AT-CPA) defined as follows.

Definition 2.16. (*PR-AT-CPA Security for ABKEM*) Let $k(\cdot)$ and $l(\cdot)$ be functions in λ . ABKEM is \mathcal{C}_2 - (k, l) -PR-AT-CPA secure if for any $\mathcal{A} = \{a_\lambda\}_{\lambda \in \mathbb{N}} \in \mathcal{C}_2$, where a_λ is allowed to make k rounds of adaptive queries to $\text{USKGEN}(\cdot)$ and each round it queries l inputs, we have

$$\left| \Pr[\text{PR-AT-CPA}_{\text{real}}^{a_\lambda} \Rightarrow 1] - \Pr[\text{PR-AT-CPA}_{\text{rand}}^{a_\lambda} \Rightarrow 1] \right| \leq \text{negl}(\lambda),$$

where the experiments are defined in Fig. 4.

3. Generalized Affine MAC

In this section, we give the definition of generalized affine MAC, which generalizes the notion of standard affine MAC [5] by using predicate encodings, and show how to construct it in the fine-grained setting under the assumption $\text{NC}^1 \not\subseteq \oplus\text{L}/\text{poly}$.

3.1. Definitions

The definition of generalized affine MAC is as follows.

Definition 3.1. (*Generalized Affine MAC*) Let $\text{PE} = \{\text{sE}_\lambda, \text{rE}_\lambda, \text{kE}_\lambda, \text{sD}_\lambda, \text{rD}_\lambda\}_{\lambda \in \mathbb{N}} \in \mathcal{C}_1$ be a predicate encoding for $\text{P} = \{\text{p}_\lambda\}_{\lambda \in \mathbb{N}}$, where $\text{rE}_\lambda : \mathcal{Y} \times \{0, 1\}^\ell \rightarrow \{0, 1\}^\eta$, $\text{kE}_\lambda : \mathcal{Y} \times \{0, 1\} \rightarrow \{0, 1\}^\eta$, and $\text{sE}_\lambda : \mathcal{X} \times \{0, 1\}^\ell \rightarrow \{0, 1\}^\xi$.

A \mathcal{C}_1 -generalized affine message authentication code for PE is a function family $\text{MAC}_{\text{GA}} = \{\text{Gen}_{\text{MAC}_\lambda}, \text{Tag}_\lambda, \text{Ver}_{\text{MAC}_\lambda}\}_{\lambda \in \mathbb{N}} \in \mathcal{C}_1$.

1. $\text{Gen}_{\text{MAC}_\lambda}$ returns sk_{MAC} containing $(\mathbf{B}, \mathbf{X}, x')$, where $\mathbf{B}^\top \in \text{ZeroSamp}(\lambda)$, $\mathbf{X} \in \{0, 1\}^{\lambda \times \ell}$, and $x' \in \{0, 1\}$.
2. $\text{Tag}_\lambda(\text{sk}_{\text{MAC}}, \mathbf{m} \in \mathcal{Y})$ returns a tag $\tau = (\mathbf{t}, \mathbf{u}) \in \{0, 1\}^\lambda \times \{0, 1\}^\eta$, computed as

$$\mathbf{t} \stackrel{\$}{\leftarrow} \text{SampYes}_\lambda(\mathbf{B}) \quad (1)$$

$$\mathbf{u} = \text{rE}_\lambda(\mathbf{m}, \mathbf{X}^\top \mathbf{t}) + \text{kE}_\lambda(\mathbf{m}, x') \in \{0, 1\}^\eta. \quad (2)$$

3. $\text{Ver}_{\text{MAC}_\lambda}(\text{sk}_{\text{MAC}}, \mathbf{m}, \tau = (\mathbf{t}, \mathbf{u}))$ verifies if Eq. (2) holds.

Correctness is satisfied if for any $\text{sk}_{\text{MAC}} \in \text{Gen}_{\text{MAC}_\lambda}$, $\mathbf{m} \in \mathcal{Y}$, and $\tau \in \text{Tag}_\lambda(\text{sk}_{\text{MAC}}, \mathbf{m})$, we have $1 = \text{Ver}_{\text{MAC}_\lambda}(\text{sk}_{\text{MAC}}, \mathbf{m}, \tau)$.

The security requirement we consider is pseudorandomness against chosen message attacks (PR-CMA) defined as follows.

Definition 3.2. (*PR-CMA Security*) Let $k = k(\lambda)$ and $l = l(\lambda)$ be polynomials in λ . MAC_{GA} is \mathcal{C}_2 - (k, l) -PR-CMA secure if for any $\mathcal{A} = \{a_\lambda\}_{\lambda \in \mathbb{N}} \in \mathcal{C}_2$, where a_λ is allowed to make k rounds of adaptive queries to $\text{EVAL}(\cdot)$ and each round it queries l inputs, we have

$$\Pr [\text{PR-CMA}_{\text{real}}^{a_\lambda} \Rightarrow 1] - \Pr [\text{PR-CMA}_{\text{rand}}^{a_\lambda} \Rightarrow 1] \leq \text{negl}(\lambda),$$

where the experiments are defined in Fig. 5.

Roughly, the PR-CMA security says that in the presence of many tags and a challenge token (\mathbf{h}_0, h_1) , an adversary cannot tell whether the h_1 is honestly generated or randomness.

<p><u>INIT</u>: $sk_{MAC} = (\mathbf{B}, \mathbf{X}, x') \xleftarrow{\\$} \text{Gen}_{MAC_\lambda}(\text{par})$ Return ε</p> <p><u>EVAL</u>(m): $\mathcal{Q}_m = \mathcal{Q}_m \cup \{m\}$ Return $(t, \mathbf{u}) \xleftarrow{\\$} \text{Tag}_\lambda(sk_{MAC}, m)$</p>	<p><u>CHAL</u>(m^*): $\mathbf{h}_0 = sE_\lambda(m^*, \mathbf{X}^\top) \in \{0, 1\}^{c \times \lambda}$ $h_1 = x' \in \{0, 1\}$ $h_1 \xleftarrow{\\$} \{0, 1\}$ Return (\mathbf{h}_0, h_1)</p> <p><u>FINALIZE</u>($\beta \in \{0, 1\}$): If $p_\lambda(m^*, m) \neq 1$ for all $m \in \mathcal{Q}_m$, return β Else return 0</p>
---	--

Fig. 5. Games $\text{PR-CMA}_{\text{real}}$ and $\text{PR-CMA}_{\text{rand}}$ for defining PR-CMA security. The boxed statement re-defining h_1 is only executed in game $\text{PR-CMA}_{\text{rand}}$.

<p><u>Gen</u>$_{MAC_\lambda}(\text{par})$: $\mathbf{B}^\top \xleftarrow{\\$} \text{ZeroSamp}(\lambda)$ $\mathbf{X} \xleftarrow{\\$} \{0, 1\}^{\lambda \times \ell}$ $x' \xleftarrow{\\$} \{0, 1\}$ Return $sk_{MAC} = (\mathbf{B}, \mathbf{X}, x')$</p>	<p><u>Tag</u>$_\lambda(sk_{MAC}, m \in \mathcal{Y})$: $\mathbf{t} \xleftarrow{\\$} \text{SampYes}_\lambda(\mathbf{B})$ $\mathbf{u} = rE_\lambda(m, \mathbf{X}^\top \mathbf{t}) + kE_\lambda(m, x') \in \{0, 1\}^n$ Return $\tau = (t, \mathbf{u})$</p> <p><u>Ver</u>$_{MAC_\lambda}(sk_{MAC}, m \in \mathcal{Y}, \tau)$: If $\mathbf{u} = rE_\lambda(m, \mathbf{X}^\top \mathbf{t}) + kE_\lambda(m, x')$ return 1 Else return 0</p>
--	--

Fig. 6. Definition of $\text{MAC}_{GA} = \{\text{Gen}_{MAC_\lambda}, \text{Tag}_\lambda, \text{Ver}_{MAC_\lambda}\}_{\lambda \in \mathbb{N}}$.

Standard Affine MAC. Let $\mathbf{X} = (\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n) \xleftarrow{\$} \{0, 1\}^{\lambda \times (n+1)}$. When $p_\lambda(\cdot)$ is an identity function, \mathbf{u} is computed as

$$\mathbf{u} = \mathbf{x}_0^\top \mathbf{t} + \sum_{i=1}^n m_i \mathbf{x}_i^\top \mathbf{t} + x' \in \{0, 1\}^n \quad (3)$$

in Eq. (2), and \mathbf{h}_0 is computed as

$$\mathbf{h}_0 = h \cdot \left(\mathbf{x}_0^\top + \sum_{i=1}^n m_i^* \mathbf{x}_i^\top \right) \in \{0, 1\}^{1 \times \lambda} \quad (4)$$

in Fig. 5, i.e., the predicate encoding is the one for equality (see Fig. 3), the above definition becomes exactly the same as that of affine MAC given in [5] for the HPS based IBKEM, except that we only consider computations over $GF(2)$ and \mathbf{t} is sampled by SampYes_λ . We give the definition as below.

Definition 3.3. (*Affine MAC* [5]) A *Generalized affine MAC* for the predicate P_{eq} and encoding PE_{eq} defined as in Fig. 3 is said to be an affine MAC.

3.2. Construction

In this section, we give our construction of $\text{AC}^0[2]$ -generalized affine MAC based on $\text{NC}^1 \subseteq \oplus\text{L}/\text{poly}$. It is a natural extension of the standard affine MAC from an HPS in [5].

<pre> INIT: // Games G₀-G₂ B^T $\stackrel{\\$}{\leftarrow}$ ZeroSamp(λ), $x' \stackrel{\\$}{\leftarrow} \{0, 1\}$ For X $\stackrel{\\$}{\leftarrow} \{0, 1\}^{\lambda \times \ell}$ Return ε CHAL($m^* \in \mathcal{X}$): // Games G₀-G₁, G₁, G₂ h₀ = sE_{λ}(m^*, X^T) $\in \{0, 1\}^{c \times \lambda}$ h₁ = $x' \in \{0, 1\}$ [h₁ $\stackrel{\\$}{\leftarrow} \{0, 1\}$] Return (h₀, h₁) FINALIZE($\beta \in \{0, 1\}$): // Games G₀-G₂ If p_{λ}(m^*, m) $\neq 1$ for all m $\in \mathcal{Q}_m$ return β Else return 0 EVAL(m): // Game G₂ Q_m = Q_m $\cup \{m\}$ t $\stackrel{\\$}{\leftarrow}$ SampNo_{λ}(B) u = rE_{λ}(m, X^Tt) $\in \{0, 1\}^\eta$ Return (t, u) </pre>	<pre> EVAL(m): // Game G₀ Q_m = Q_m $\cup \{m\}$ t $\stackrel{\\$}{\leftarrow}$ SampYes_{λ}(B) u = rE_{λ}(m, X^Tt) + kE_{λ}(m, x') $\in \{0, 1\}^\eta$ Return (t, u) EVAL(m): // Games G₁, i, [G₁, i] Q_m = Q_m $\cup \{m\}$ // Let m be the c-th query (1 $\leq c \leq k \cdot l$) If c < i then t $\stackrel{\\$}{\leftarrow}$ SampNo_{λ}(B) u = rE_{λ}(m, X^Tt) $\in \{0, 1\}^\eta$ If c > i then t $\stackrel{\\$}{\leftarrow}$ SampYes_{λ}(B) u = rE_{λ}(m, X^Tt) + kE_{λ}(m, x') $\in \{0, 1\}^\eta$ If c = i then t $\stackrel{\\$}{\leftarrow}$ SampYes_{λ}(B) [t $\stackrel{\\$}{\leftarrow}$ SampNo_{λ}(B)] u = rE_{λ}(m, X^Tt) + kE_{λ}(m, x') $\in \{0, 1\}^\eta$ Return (t, u) </pre>
--	---

Fig. 7. Games G_0 , $(G_{1,i}, G'_{1,i})_{1 \leq i \leq k \cdot l}$, $G_{1,k \cdot l + 1}$, G_2 for the proof of Theorem 3.4.

Theorem 3.4. *If $\text{NC}^1 \subsetneq \oplus\text{L}/\text{poly}$ and $\text{PE} = \{\text{sE}_\lambda, \text{rE}_\lambda, \text{kE}_\lambda, \text{sD}_\lambda, \text{rD}_\lambda\}_{\lambda \in \mathbb{N}} \in \text{AC}^0[2]$ is a predicate encoding, where $\text{rE}_\lambda : \mathcal{Y} \times \{0, 1\}^\ell \rightarrow \{0, 1\}^\eta$, $\text{kE}_\lambda : \mathcal{Y} \times \{0, 1\} \rightarrow \{0, 1\}^\eta$, and $\text{sE}_\lambda : \mathcal{X} \times \{0, 1\}^\ell \rightarrow \{0, 1\}^c$, then MAC_{GA} is an $\text{AC}^0[2]$ -generalized affine MAC that is NC^1 - (k, l) -PR-CMA secure, where k is any constant and $l = l(\lambda)$ is any polynomial in λ .*

Proof. First, we note that $(\{\text{Gen}_{\text{MAC}_\lambda}\}_{\lambda \in \mathbb{N}}, \{\text{Tag}_\lambda\}_{\lambda \in \mathbb{N}}, \{\text{Ver}_{\text{MAC}_\lambda}\}_{\lambda \in \mathbb{N}})$ are computable in $\text{AC}^0[2]$, since they only involve operations including sampling random bits and multiplication of a constant number of matrices, which can be done in constant depth with PARITY gates. Also, it is straightforward that MAC_{GA} satisfies correctness.

We now prove that MAC_{GA} is NC^1 - (k, l) -PR-CMA secure by defining a sequence of intermediate games as in Fig. 7.

Let $\mathcal{A} = \{a_\lambda\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$ be any adversary against the PR-CMA-security of MAC_{GA} . Game G_0 is the real attack game. In games $G_{1,i}$, the first $i - 1$ queries to the EVAL oracle are answered with (\mathbf{t}, \mathbf{u}) , where $\mathbf{t} \stackrel{\$}{\leftarrow} \text{SampNo}_\lambda(\mathbf{B})$ and \mathbf{u} contains no information on $\text{kE}_\lambda(m, x')$, and the remaining are answered as in the real scheme. To interpolate between $G_{1,i}$ and $G_{1,i+1}$, we also define $G'_{1,i}$, which answers the i -th query to EVAL by picking $\mathbf{t} \stackrel{\$}{\leftarrow} \text{SampNo}_\lambda(\mathbf{B})$. By definition, we have $G_0 = G_{1,1}$. \square

Lemma 3.5. $\Pr[\text{PR-CMA}_{\text{real}}^{a_\lambda} \Rightarrow 1] = \Pr[G_0^{a_\lambda} \Rightarrow 1] = \Pr[G'_{1,1}^{a_\lambda} \Rightarrow 1]$.

Lemma 3.6. *There exists an adversary $\mathcal{B}_{1,i} = \{b_\lambda^{1,i}\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$ such that $b_\lambda^{1,i}$ breaks the fine-grained subset membership problem (see Definition 2.12), which holds under $\text{NC}^1 \subsetneq \oplus\text{L}/\text{poly}$ according to Theorem 2.13, with probability*

$$\left| \Pr \left[\mathcal{G}'_{1,i}^{a_\lambda} \Rightarrow 1 \right] - \Pr \left[\mathcal{G}_{1,i}^{a_\lambda} \Rightarrow 1 \right] \right|.$$

Proof. Games $\mathbf{G}_{1,i}$ and $\mathbf{G}'_{1,i}$ only differ in the distribution of \mathbf{t} returned by the EVAL oracle for its i -th query. We build $b_\lambda^{1,i}$ as follows.

The distinguisher $b_\lambda^{1,i}$ runs in exactly the same way as the challenger in $\mathbf{G}_{1,i}$ except that for its i -th query, it obtains \mathbf{t} which is sampled as $\mathbf{t} \stackrel{\$}{\leftarrow} \text{SampYes}_\lambda(\mathbf{B})$ or $\mathbf{t} \stackrel{\$}{\leftarrow} \text{SampNo}_\lambda(\mathbf{B})$. When a_λ outputs $\beta \in \{0, 1\}$, b_λ outputs β if no \mathbf{m} such that $\mathbf{p}_\lambda(\mathbf{m}^*, \mathbf{m}) = 1$ was queried to EVAL. Otherwise, b_λ outputs 0.

Since a_λ only makes constant rounds of queries, all the operations in b_λ are performed in NC^1 . Hence, we have $\mathcal{B}_{1,i} \in \text{NC}^1$.

When \mathbf{t} is sampled as $\mathbf{t} \stackrel{\$}{\leftarrow} \text{SampYes}_\lambda(\mathbf{B})$ (respectively, $\mathbf{t} \stackrel{\$}{\leftarrow} \text{SampNo}_\lambda(\mathbf{B})$), the view of a_λ is exactly the same as its view in $\mathbf{G}_{1,i}$ (respectively, $\mathbf{G}'_{1,i}$). Thus the advantage of $b_\lambda^{1,i}$ in breaking the subset membership problem is $|\Pr[\mathbf{G}_{1,i}^{a_\lambda} \Rightarrow 1] - \Pr[\mathbf{G}'_{1,i} \Rightarrow 1]|$, completing this part of proof. \square

Lemma 3.7. $\Pr[\mathbf{G}_{1,i+1}^{a_\lambda} \Rightarrow 1] = \Pr[\mathbf{G}_{1,i}^{a_\lambda} \Rightarrow 1]$.

Proof. Let \mathbf{m} be the i -th query to EVAL such that $\mathbf{p}_\lambda(\mathbf{m}^*, \mathbf{m}) \neq 1$ and let (\mathbf{t}, \mathbf{u}) be its tag. We have $\mathbf{t} \notin \text{Im}(\mathbf{B})$ due to Theorem 2.13. We use an information-theoretic argument to show that in $\mathbf{G}'_{1,i}$, \mathbf{u} does not reveal any information on x' . Information-theoretically, a_λ may learn $\mathbf{B}^\top \mathbf{X}$ from each c -th query with $c > i$. Thus, for $\mathbf{X} \stackrel{\$}{\leftarrow} \{0, 1\}^{\lambda \times \ell}$ and $\mathbf{w} \stackrel{\$}{\leftarrow} \{0, 1\}^{\ell \times 1}$, a_λ information-theoretically obtains the distribution of

$$\begin{aligned} & \left(\begin{array}{c} \mathbf{X}^\top \mathbf{B} \\ \mathbf{h}_0 = h \odot \mathbf{sE}_\lambda(\mathbf{m}^*, \mathbf{X}^\top) \\ \mathbf{u} = \mathbf{rE}_\lambda(\mathbf{m}, \mathbf{X}^\top \mathbf{t}) + \mathbf{kE}_\lambda(\mathbf{m}, x') \end{array} \right) \\ &= \left(\begin{array}{c} (\mathbf{X}^\top + \mathbf{wB}^{\perp\top})\mathbf{B} \\ \mathbf{h}_0 = \mathbf{sE}_\lambda(\mathbf{m}^*, \mathbf{X}^\top + \mathbf{wB}^{\perp\top}) \\ \mathbf{u} = \mathbf{rE}_\lambda(\mathbf{m}, (\mathbf{X}^\top + \mathbf{wB}^{\perp\top})\mathbf{t}) + \mathbf{kE}_\lambda(\mathbf{m}, x') \end{array} \right) \\ &= \left(\begin{array}{c} \mathbf{X}^\top \mathbf{B} \\ \mathbf{h}_0 = \mathbf{sE}_\lambda(\mathbf{m}^*, \mathbf{X}^\top) + \mathbf{sE}_\lambda(\mathbf{m}^*, \mathbf{wB}^{\perp\top}) \\ \mathbf{u} = \mathbf{rE}_\lambda(\mathbf{m}, \mathbf{X}^\top \mathbf{t}) + \mathbf{rE}_\lambda(\mathbf{m}, \mathbf{w}) + \mathbf{kE}_\lambda(\mathbf{m}, x') \end{array} \right) (\because \mathbf{t} \notin \text{Im}(\mathbf{B})). \end{aligned}$$

This distribution is identical to the distribution of

$$\left(\begin{array}{c} \mathbf{X}^\top \mathbf{B} \\ \mathbf{h}_0 = \mathbf{sE}_\lambda(\mathbf{m}^*, \mathbf{X}^\top) + \mathbf{sE}_\lambda(\mathbf{m}^*, \mathbf{wB}^{\perp\top}) \\ \mathbf{u} = \mathbf{rE}_\lambda(\mathbf{m}, \mathbf{X}^\top \mathbf{t}) + \mathbf{rE}_\lambda(\mathbf{m}, \mathbf{w}) \end{array} \right),$$

since the distribution of

$$(\mathbf{m}^*, \mathbf{m}, x', \mathbf{sE}_\lambda(\mathbf{m}^*, \mathbf{w}), \mathbf{rE}_\lambda(\mathbf{m}, \mathbf{w}) + \mathbf{kE}_\lambda(\mathbf{m}, x'))$$

and

$$(\mathbf{m}^*, \mathbf{m}, x', \mathbf{sE}_\lambda(\mathbf{m}^*, \mathbf{w}), \mathbf{rE}_\lambda(\mathbf{m}, \mathbf{w})),$$

<p><u>INIT:</u> // Games H_0-H_2</p> <p>$\mathbf{B}^\top \stackrel{\\$}{\leftarrow} \text{ZeroSamp}(\lambda); x' \stackrel{\\$}{\leftarrow} \{0, 1\}$</p> <p>$\mathbf{X} \stackrel{\\$}{\leftarrow} \{0, 1\}^{\lambda \times \ell}$</p> <p>Return ε</p> <p><u>CHAL</u>(m^*): // Games H_0-H_2</p> <p>$\mathbf{h}_0 = \text{sE}_\lambda(m^*, \mathbf{X}^\top) \in \{0, 1\}^{c \times \lambda}$</p> <p>$h_1 \stackrel{\\$}{\leftarrow} \{0, 1\}$</p> <p>Return (\mathbf{h}_0, h_1)</p> <p><u>FINALIZE</u>($\beta \in \{0, 1\}$): // Games H_0-H_2</p> <p>If $\mathbf{p}_\lambda(m^*, m) \neq 1$ for all $y \in \mathcal{Q}_m$</p> <p style="padding-left: 20px;">return β</p> <p>Else return 0</p> <p><u>EVAL</u>(m): // Game H_0</p> <p>$\mathcal{Q}_m = \mathcal{Q}_m \cup \{m\}$</p> <p>$\mathbf{t} \stackrel{\\$}{\leftarrow} \text{SampNo}_\lambda(\mathbf{B})$</p> <p>$\mathbf{u} = \text{rE}_\lambda(m, \mathbf{X}^\top \mathbf{t}) \in \{0, 1\}^\eta$</p> <p>Return (\mathbf{t}, \mathbf{u})</p>	<p><u>EVAL</u>(m): // Games $H_{1,i}, \boxed{H'_{1,i}}$</p> <p>$\mathcal{Q}_m = \mathcal{Q}_m \cup \{m\}$ // Let m be the c-th query</p> <p>$(1 \leq c \leq k \cdot l)$</p> <p>If $c > i$ then</p> <p style="padding-left: 20px;">$\mathbf{t} \stackrel{\\$}{\leftarrow} \text{SampNo}_\lambda(\mathbf{B})$</p> <p style="padding-left: 20px;">$\mathbf{u} = \text{rE}_\lambda(m, \mathbf{X}^\top \mathbf{t}) \in \{0, 1\}^\eta$</p> <p>If $c < i$ then</p> <p style="padding-left: 20px;">$\mathbf{t} \stackrel{\\$}{\leftarrow} \text{SampYes}_\lambda(\mathbf{B})$</p> <p style="padding-left: 20px;">$\mathbf{u} = \text{rE}_\lambda(m, \mathbf{X}^\top \mathbf{t}) + \text{kE}_\lambda(m, x') \in \{0, 1\}^\eta$</p> <p>If $c = i$ then</p> <p style="padding-left: 20px;">$\mathbf{t} \stackrel{\\$}{\leftarrow} \text{SampNo}_\lambda(\mathbf{B})$</p> <p style="padding-left: 20px;">$\boxed{\mathbf{t} \stackrel{\\$}{\leftarrow} \text{SampYes}_\lambda(\mathbf{B})}$</p> <p style="padding-left: 20px;">$\mathbf{u} = \text{rE}_\lambda(m) \mathbf{X}^\top \mathbf{t} + \text{kE}_\lambda(m, x') \in \{0, 1\}^\eta$</p> <p>Return (\mathbf{t}, \mathbf{u})</p> <p><u>EVAL</u>(m): // Game H_2</p> <p>$\mathcal{Q}_m = \mathcal{Q}_m \cup \{m\}$</p> <p>$\mathbf{t} \stackrel{\\$}{\leftarrow} \text{SampYes}_\lambda(\mathbf{B})$</p> <p>$\mathbf{u} = \text{rE}_\lambda(m, \mathbf{X}^\top \mathbf{t}) + \text{kE}_\lambda(m, x') \in \{0, 1\}^\eta$</p> <p>Return (\mathbf{t}, \mathbf{u})</p>
--	---

Fig. 8. Games $H_0, (H_{1,i}, H'_{1,i})_{1 \leq i \leq k \cdot l}, H_{1,k \cdot l+1}, H_2$ for the proof of Lemma 3.9.

are identical due to the α -privacy of PE, completing this part of proof. \square

Lemma 3.8. $\Pr[\mathbf{G}_2^{a_\lambda} \Rightarrow 1] = \Pr[\mathbf{G}_{1,k \cdot l+1}^{a_\lambda} \Rightarrow 1]$.

Proof. Note that a_λ can ask at most $k \cdot l$ -many EVAL queries. In both $\mathbf{G}_{1,k \cdot l+1}$ and \mathbf{G}_2 , all the answers of EVAL are independent of x' . Hence, h_1 from $\mathbf{G}_{1,k \cdot l+1}$ is uniform in the view of a_λ . \square

We now do all the previous steps in the reverse order as in Fig. 8. Then, by using the above arguments in a reverse order, we have the following lemma.

Lemma 3.9. *There exists an adversary $\mathcal{B}_2 = \{b_\lambda^2\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$ such that b_λ^2 breaks the fine-grained subset membership problem with probability at least*

$$\left(\Pr[\text{PR-CMA}_{\text{rand}}^{a_\lambda} \Rightarrow 1] - \Pr[\mathbf{G}_2^{a_\lambda} \Rightarrow 1] \right) / (k \cdot l).$$

Putting all above together, Theorem 3.4 immediately follows. \square

An Affine MAC. By instantiating the underlying predicate encoding in Fig. 6 with the encoding for equality (see Fig. 3), we immediately obtain an affine MAC $\text{MAC} = \{\text{Gen}_{\text{MAC}\lambda}, \text{Tag}_\lambda, \text{Ver}_{\text{MAC}\lambda}\}_{\lambda \in \mathbb{N}}$ as in Fig. 9 for message space $\{0, 1\}^\ell$, which will be used to construct an IBE scheme in NC^1 later. Formally, we have the following corollary derived from Theorem 3.4.

Corollary 3.10. *If $\text{NC}^1 \not\subseteq \oplus\text{L}/\text{poly}$, then MAC is an $\text{AC}^0[2]$ -affine MAC that is NC^1 - (k, l) -PR-CMA secure, where k is any constant and $l = l(\lambda)$ is any polynomial in λ .*

$\text{Gen}_{\text{MAC}_\lambda}(\text{par}):$ $\mathbf{B}^\top \xleftarrow{\$} \text{ZeroSamp}(\lambda)$ $\mathbf{x}_0, \dots, \mathbf{x}_\ell \xleftarrow{\$} \{0, 1\}^\lambda$ $x' \xleftarrow{\$} \{0, 1\}$ Return $\text{sk}_{\text{MAC}} = (\mathbf{B}, \mathbf{x}_0, \dots, \mathbf{x}_\ell, x')$	$\text{Tag}_\lambda(\text{sk}_{\text{MAC}}, \mathbf{m} \in \{0, 1\}^\ell):$ $\mathbf{t} \xleftarrow{\$} \text{SampYes}_\lambda(\mathbf{B})$ $u = (\mathbf{x}'_0 + \sum_{i=1}^{\ell} \mathbf{m}_i \cdot \mathbf{x}_i^\top) \mathbf{t} + x' \in \{0, 1\}$ Return $\tau = (\mathbf{t}, u)$ $\text{Ver}_{\text{MAC}_\lambda}(\text{sk}_{\text{MAC}}, \tau, \mathbf{m}):$ If $u = (\mathbf{x}'_0 + \sum_{i=1}^{\ell} \mathbf{m}_i \cdot \mathbf{x}_i^\top) \mathbf{t} + x'$ return 1 Else return 0
--	---

Fig. 9. Definition of $\text{MAC} = \{\text{Gen}_{\text{MAC}_\lambda}, \text{Tag}_\lambda, \text{Ver}_{\text{MAC}_\lambda}\}_{\lambda \in \mathbb{N}}$.

4. Fine-Grained Secure Identity-Based Encryption

In this section, we present our fine-grained IBE scheme, which captures the core techniques of our ABE scheme given later in Sect. 5.

4.1. Definition

We now give the definition of fine-grained IBKEM, which is a special case of fine-grained ABKEM (see Definition 2.15) where the boolean predicate is restricted to be the equality predicate.

Definition 4.1. (*Identity-based key encapsulation*) A \mathcal{C}_1 -identity key encapsulation (IBKEM) scheme is a function family $\text{IBKEM} = \{\text{Gen}_\lambda, \text{USKGen}_\lambda, \text{Enc}_\lambda, \text{Dec}_\lambda\}_{\lambda \in \mathbb{N}} \in \mathcal{C}_1$ with the following properties.

- Gen_λ returns the (master) public/secret key (pk, sk) . We assume that pk implicitly defines an identity space \mathcal{ID} , a key space \mathcal{K} , and a ciphertext space \mathcal{C} .
- $\text{USKGen}_\lambda(\text{sk}, \text{id})$ returns a user secret-key $\text{usk}[\text{id}]$ for an identity $\text{id} \in \mathcal{ID}$.
- $\text{Enc}_\lambda(\text{pk}, \text{id})$ returns a symmetric key $\mathbf{K} \in \mathcal{K}$ together with a ciphertext $\text{ct} \in \mathcal{C}$ w.r.t. $\text{id} \in \mathcal{ID}$.
- $\text{Dec}_\lambda(\text{usk}[\text{id}], \text{id}, \text{ct})$ deterministically returns a decapsulated key $\mathbf{K} \in \mathcal{K}$ or the reject symbol \perp .

Perfect correctness is satisfied if for all $\lambda \in \mathbb{N}$, all $(\text{pk}, \text{sk}) \in \text{Gen}_\lambda$, all $\text{id} \in \mathcal{ID}$, all $\text{usk}[\text{id}] \in \text{USKGen}_\lambda(\text{sk}, \text{id})$, and all $(\mathbf{K}, \text{ct}) \in \text{Enc}_\lambda(\text{pk}, \text{id})$, we have

$$\Pr[\text{Dec}_\lambda(\text{usk}[\text{id}], \text{id}, \text{ct}) = \mathbf{K}] = 1.$$

The security requirement we consider is indistinguishability against chosen plaintext and identity attacks (PR-ID-CPA) defined as follows.

Definition 4.2. (*PR-ID-CPA Security for IBKEM*) Let $k(\cdot)$ and $l(\cdot)$ be functions in λ . IBKEM is \mathcal{C}_2 - (k, l) -PR-ID-CPA secure if for any $\mathcal{A} = \{a_\lambda\}_{\lambda \in \mathbb{N}} \in \mathcal{C}_2$, where a_λ is allowed to make k rounds of adaptive queries to $\text{USKGen}(\cdot)$ and each round it queries l inputs, we have

$$\left| \Pr[\text{PR-ID-CPA}_{\text{real}}^{a_\lambda} \Rightarrow 1] - \Pr[\text{PR-ID-CPA}_{\text{rand}}^{a_\lambda} \Rightarrow 1] \right| \leq \text{negl}(\lambda),$$

<p>Procedure INIT:</p> $(pk, sk) \stackrel{\$}{\leftarrow} \text{Gen}_\lambda$ Return pk <p>Procedure USKGEN(id):</p> $//k(\lambda) \times l(\lambda)$ queries $Q_{id} \stackrel{\$}{\leftarrow} Q_{id} \cup \{id\}$ Return $usk[id] \stackrel{\$}{\leftarrow} \text{USKGen}_\lambda(sk, id)$	<p>Procedure ENC(id*):</p> $//$ one query $(K^*, ct^*) \stackrel{\$}{\leftarrow} \text{Enc}_\lambda(pk, id^*)$ $K^* \stackrel{\$}{\leftarrow} \mathcal{K}$ Return (K^*, ct^*) <p>Procedure FINALIZE(β):</p> Return $(id^* \notin Q_{id}) \wedge \beta$
--	--

Fig. 10. Security Games PR-ID-CPA_{real} and PR-ID-CPA_{rand} for defining PR-ID-CPA-security for IBKEM. The boxed statement redefining K^* is only executed in game PR-ID-CPA_{rand}.

<p>Gen$_\lambda$:</p> $A^T \stackrel{\$}{\leftarrow} \text{ZeroSamp}(\lambda)$ $sk_{MAC} = (B, x_0, \dots, x_\ell, x')$ $\stackrel{\$}{\leftarrow} \text{Gen}_{MAC_\lambda}(\text{par})$ For $i = 0, \dots, \ell$: $Y_i \stackrel{\$}{\leftarrow} \{0, 1\}^{(\lambda-1) \times \lambda}$ $Z_i = (Y_i^T x_i)A \in \{0, 1\}^{\lambda \times \lambda}$ $y' \stackrel{\$}{\leftarrow} \{0, 1\}^{\lambda-1}$ $z' = (y'^T x')A \in \{0, 1\}^{1 \times \lambda}$ $pk = (A, (Z_i)_{0 \leq i \leq \ell}, z')$ $sk = (sk_{MAC}, (Y_i)_{0 \leq i \leq \ell}, y')$ Return (pk, sk) <p>USKGen$_\lambda$(sk, id $\in \{0, 1\}^\ell$):</p> $(t, u) \stackrel{\$}{\leftarrow} \text{Tag}_\lambda(sk_{MAC}, id)$ $v = t^T (Y_0^T + \sum_{i=1}^\ell id_i \odot Y_i^T) + y'^T \in \{0, 1\}^{1 \times (\lambda-1)}$ Return $usk[id] = (t, u, v)$	<p>Enc$_\lambda$(pk, id):</p> $r \stackrel{\$}{\leftarrow} \{0\} \times \{0, 1\}^{\lambda-1}$ $c_0 = Ar \in \{0, 1\}^\lambda$ $c_1 = (Z_0 + \sum_{i=1}^\ell id_i \odot Z_i)r \in \{0, 1\}^\lambda$ $K = z' \cdot r \in \{0, 1\}$ Return K and $ct = (c_0, c_1)$ <p>Dec$_\lambda$(usk[id], id, ct):</p> Parse $usk[id] = (t, u, v)$ Parse $ct = (c_0, c_1) \in \{0, 1\}^\lambda \times \{0, 1\}^\lambda$ $K = (v u)c_0 - t^T c_1$ Return K
---	--

Fig. 11. Definition of our IBKEM = {Gen $_\lambda$, USKGen $_\lambda$, Enc $_\lambda$, Dec $_\lambda$ } $_{\lambda \in \mathbb{N}}$ with identity space $\{0, 1\}^\ell$ and key space $\{0, 1\}$. id_i denotes the i th bit of id for all $i \in [\ell]$.

where the experiments are defined in Fig. 10.

4.2. Construction

Let $MAC = \{\text{Gen}_{MAC_\lambda}, \text{Tag}_\lambda, \text{Ver}_{MAC_\lambda}\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$ be an affine MAC over $\{0, 1\}^\lambda$ with message space \mathcal{ID} in Fig. 9. Our IBKEM $\text{IBKEM} = \{\text{Gen}_\lambda, \text{USKGen}_\lambda, \text{Enc}_\lambda, \text{Dec}_\lambda\}_{\lambda \in \mathbb{N}}$ for key-space $\mathcal{K} = \{0, 1\}$ and identity space $\{0, 1\}^\ell$ is defined as in Fig. 11.²

Theorem 4.3. *Under the assumption $\text{NC}^1 \subsetneq \oplus \text{L}/\text{poly}$ and the NC^1 - (k, l) -PR-CMA security of MAC, where k is any constant and $l = l(\lambda)$ is any polynomial in λ , IBKEM is an $\text{AC}^0[2]$ -IBKEM that is NC^1 - (k, l) -PR-ID-CPA secure against NC^1 .*

Proof. First, we note that $\{\text{Gen}_\lambda\}_{\lambda \in \mathbb{N}}$, $\{\text{USKGen}_\lambda\}_{\lambda \in \mathbb{N}}$, $\{\text{Enc}_\lambda\}_{\lambda \in \mathbb{N}}$, and $\{\text{Dec}_\lambda\}_{\lambda \in \mathbb{N}}$ are computable in $\text{AC}^0[2]$, since they only involve operations including multiplication of a constant number of matrices, sampling random bits, and running $MAC \in \text{AC}^0[2]$.

²The IBKEM can be straightforwardly extended to one with large key space as we will discuss later in this section.

Correctness follows from the fact that by Eq. (3) in Sect. 3.1, we have

$$\begin{aligned} (\mathbf{v}||u)\mathbf{c}_0 &= \left(\mathbf{t}^\top \left(\mathbf{Y}_0^\top + \sum_{i=1}^{\ell} \text{id}_i \odot \mathbf{Y}_i^\top \right) + \mathbf{y}'^\top || \mathbf{t}^\top \left(\mathbf{x}_0 + \sum_{i=1}^{\ell} \text{id}_i \odot \mathbf{x}_i \right) + x' \right) \mathbf{A} \mathbf{r} \\ \mathbf{t}^\top \mathbf{c}_1 &= \mathbf{t}^\top \left(\mathbf{Y}_0^\top || \mathbf{x}_0 + \sum_{i=1}^{\ell} \text{id}_i \odot (\mathbf{Y}_i^\top || \mathbf{x}_i) \right) \mathbf{A} \mathbf{r} \end{aligned}$$

and the difference of the two elements yields $\mathbf{K} = (\mathbf{y}'^\top || x') \mathbf{A} \mathbf{r} = \mathbf{z}' \cdot \mathbf{r}$.

Let $\mathcal{A} = \{a_\lambda\}_{\lambda \in \mathbb{N}}$ be any adversary against the NC^1 - (k, l) -PR-ID-CPA security of IBKEM. We now prove the NC^1 - (k, l) -PR-ID-CPA security by defining a sequence of games \mathbf{G}_0 - \mathbf{G}_6 as in Fig. 12. Roughly, in the first four games, we show how to extract a challenge token for MAC from the challenge session key and ciphertext by switching the distribution of \mathbf{A} twice and changing the distribution of the randomness \mathbf{r} during the switching procedure. In the last two games, we show that the commitments \mathbf{Z}_i and \mathbf{z}' perfectly hide the secrets, and the answers of queries made by a_λ reveal no useful information other than the tags and token for MAC. \square

Lemma 4.4. $\Pr[\text{PR-ID-CPA}_{\text{real}}^{a_\lambda} \Rightarrow 1] = \Pr[\mathbf{G}_1^{a_\lambda} \Rightarrow 1] = \Pr[\mathbf{G}_0^{a_\lambda} \Rightarrow 1]$.

Proof. \mathbf{G}_0 is the real attack game. In game \mathbf{G}_1 , we change the simulation of \mathbf{c}_0^* , \mathbf{c}_1^* and \mathbf{K}^* in $\text{ENC}(\text{id}^*)$ by substituting \mathbf{Z}_i and \mathbf{z}' with their respective definitions and substituting \mathbf{A} with $\mathbf{A} + \mathbf{N}^\lambda$. Since we have

$$\mathbf{N}^\lambda \mathbf{r} = \begin{pmatrix} 0 & \cdots & 0 \\ \vdots & 0 & \cdots & 0 \\ 0 & & \ddots & \vdots \\ 1 & 0 & \cdots & 0 \end{pmatrix} \begin{pmatrix} 0 \\ r_2 \\ \vdots \\ r_\lambda \end{pmatrix} = \mathbf{0},$$

the view of a_λ in \mathbf{G}_1 is identical to its view in \mathbf{G}_0 , completing this part of proof. \square

Lemma 4.5. *There exists an adversary $\mathcal{B}_1 = \{b_\lambda^1\}_{\lambda \in \mathbb{N}}$ such that b_λ^1 breaks the fine-grained matrix linear assumption (see Definition 2.5), which holds under $\text{NC}^1 \subsetneq \oplus \text{L/poly}$ according to Lemma 2.6, with advantage*

$$|\Pr[\mathbf{G}_2^{a_\lambda} \Rightarrow 1] - \Pr[\mathbf{G}_1^{a_\lambda} \Rightarrow 1]|.$$

Proof. \mathbf{G}_1 and \mathbf{G}_2 only differ in the distribution of \mathbf{A} , namely, $\mathbf{A}^\top \stackrel{\$}{\leftarrow} \text{ZeroSamp}(\lambda)$ or $\mathbf{A}^\top \stackrel{\$}{\leftarrow} \text{OneSamp}(\lambda)$, and we build the distinguisher b_λ^1 as follows.

b_λ^1 runs in exactly the same way as the challenger of \mathbf{G}_1 except that in INIT, instead of generating \mathbf{A} by itself, it takes as input \mathbf{A}^\top generated as $\mathbf{A}^\top \stackrel{\$}{\leftarrow} \text{ZeroSamp}(\lambda)$ or $\mathbf{A}^\top \stackrel{\$}{\leftarrow} \text{OneSamp}(\lambda)$ from its own challenger. When a_λ outputs β , b_λ^1 outputs β as well if id^* was not queried to USKGEN . Otherwise, b_λ^1 outputs 0.

```

INIT: //Games G0-G1, G2-G3, G4, G5-G6
AT  $\xleftarrow{\$}$  ZeroSamp( $\lambda$ ), AT  $\xleftarrow{\$}$  OneSamp( $\lambda$ ), AT  $\xleftarrow{\$}$  ZeroSamp( $\lambda$ )
R1 =  $\begin{pmatrix} \mathbf{I}_{\lambda-1} & \mathbf{0} \\ \tilde{\mathbf{r}}^T & 1 \end{pmatrix}^T \xleftarrow{\$}$  RSamp( $\lambda$ ), R0  $\xleftarrow{\$}$  LSamp( $\lambda$ ), AT = R0M0TR1
skMAC = (B, x0, ..., xℓ, x')  $\xleftarrow{\$}$  GenMAC,λ(G)
For i = 0, ..., ℓ :
  Yi  $\xleftarrow{\$}$  {0, 1}(λ-1)×λ, Zi = (YiT || xi)A ∈ {0, 1}λ×λ
  Di = YiT + xi ·  $\tilde{\mathbf{r}}^T$  ∈ {0, 1}λ×(λ-1), Zi = (0 || Di)R0T ∈ {0, 1}λ×λ
y'  $\xleftarrow{\$}$  {0, 1}λ-1, z' = (y'T || x')A ∈ {0, 1}1×λ
d' = y'T + x' ·  $\tilde{\mathbf{r}}^T$  ∈ {0, 1}1×(λ-1), z' = (0 || d')R0T ∈ {0, 1}1×λ
pk = (A, (Zi)0≤i≤ℓ, z')
sk = (skMAC, (Yi)0≤i≤ℓ, y')
Return pk

FINALIZE( $\beta$ ): //Games G0-G6
Return (id*  $\notin$  Qid) ∧  $\beta$ 

USKGEN(id): //Games G0-G4, G5-G6
Qid = Qid ∪ {id}, (t, u)  $\xleftarrow{\$}$  Tagλ(skMAC, id)
v = tT(Y0T +  $\sum_{i=1}^{\ell}$  idi ◊ YiT) + y'T ∈ {0, 1}1×(λ-1)
v = tT(D0 +  $\sum_{i=1}^{\ell}$  idi ◊ Di) + d' - u ·  $\tilde{\mathbf{r}}^T$  ∈ {0, 1}1×(λ-1)
usk[id] = (t, u, v)
Return usk[id]

ENC(id*): //Games G0, G1-G4, G3-G4, G5, G6
r  $\xleftarrow{\$}$  {0} × {0, 1}λ-1, r  $\xleftarrow{\$}$  {1} × {0, 1}λ-1
c0* = Ar ∈ {0, 1}λ, c0* = (A + Nλ)r
c1* = Z0r +  $\sum_{i=1}^{\ell}$  idi* ◊ Zir ∈ {0, 1}λ
c1* = (Y0T || x0)(A + Nλ)r +  $\sum_{i=1}^{\ell}$  idi* ◊ (YiT || xi)(A + Nλ)r
c1* = Z0r + x0 +  $\sum_{i=1}^{\ell}$  idi* ◊ (Zir + xi)
K* = z' · r ∈ {0, 1}, K* = (y'T || x')(A + Nλ)r, K* = z' · r + x'
K*  $\xleftarrow{\$}$  {0, 1}
Return K* and ct* = (c0*, c1*)

```

Fig. 12. Games G₀-G₆ for the proof of Theorem 4.3.

If \mathbf{A} is generated as $\mathbf{A}^T \xleftarrow{\$} \text{ZeroSamp}(\lambda)$ (respectively, $\mathbf{A}^T \xleftarrow{\$} \text{OneSamp}(\lambda)$), the view of a_λ is the same as its view in \mathbf{G}_1 (respectively, \mathbf{G}_2). Hence, the probability that b_λ^1 breaks the fine-grained matrix linear assumption is

$$|\Pr[\mathbf{G}_2^{a_\lambda} \Rightarrow 1] - \Pr[\mathbf{G}_1^{a_\lambda} \Rightarrow 1]|.$$

Moreover, since a_λ only makes constant rounds of queries, all operations in b_λ^1 are performed in NC^1 . Hence, we have $\mathcal{B}_1 = \{b_\lambda^1\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$, completing this part of proof. \square

Lemma 4.6. $\Pr[\mathbf{G}_3^{a_\lambda} \Rightarrow 1] = \Pr[\mathbf{G}_2^{a_\lambda} \Rightarrow 1]$.

Proof. In this game, we sample \mathbf{r} in $\text{ENC}(\text{id}^*)$ as $\mathbf{r} \stackrel{\$}{\leftarrow} \{1\} \times \{0, 1\}^{\lambda-1}$ instead of $\mathbf{r} \stackrel{\$}{\leftarrow} \{0\} \times \{0, 1\}^{\lambda-1}$. According to Lemma 2.9, the distribution of $\mathbf{A} + \mathbf{N}^\lambda$ in \mathbf{G}_2 and \mathbf{G}_3 is identical to that of a matrix sampled from ZeroSamp . Then this lemma follows from Lemma 2.11 immediately. \square

Lemma 4.7. *There exists an adversary $\mathcal{B}_2 = \{b_\lambda^2\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$ such that b_λ^2 breaks the fine-grained matrix linear assumption with advantage*

$$|\Pr[\mathbf{G}_4^{a_\lambda} \Rightarrow 1] - \Pr[\mathbf{G}_3^{a_\lambda} \Rightarrow 1]|.$$

Proof. \mathbf{G}_3 and \mathbf{G}_4 only differ in the distribution of \mathbf{A} , namely, $\mathbf{A}^\top \stackrel{\$}{\leftarrow} \text{OneSamp}(\lambda)$ or $\mathbf{A}^\top \stackrel{\$}{\leftarrow} \text{ZeroSamp}(\lambda)$, and we build the distinguisher b_λ^2 as follows.

b_λ^2 runs in exactly the same way as the challenger of \mathbf{G}_3 except that in INIT , instead of generating \mathbf{A} by itself, it takes as input \mathbf{A}^\top generated as $\mathbf{A}^\top \stackrel{\$}{\leftarrow} \text{ZeroSamp}(\lambda)$ or $\mathbf{A}^\top \stackrel{\$}{\leftarrow} \text{OneSamp}(\lambda)$ from its own challenger. When a_λ outputs β , b_λ^2 outputs β as well if id^* was not queried to USKGEN . Otherwise, b_λ^2 outputs 0.

If \mathbf{A} is generated as $\mathbf{A}^\top \stackrel{\$}{\leftarrow} \text{OneSamp}(\lambda)$ (respectively, $\mathbf{A}^\top \stackrel{\$}{\leftarrow} \text{ZeroSamp}(\lambda)$), the view of a_λ is the same as its view in \mathbf{G}_3 (respectively, \mathbf{G}_4). Hence, the probability that b_λ^2 breaks the fine-grained matrix linear assumption is

$$|\Pr[\mathbf{G}_4^{a_\lambda} \Rightarrow 1] - \Pr[\mathbf{G}_3^{a_\lambda} \Rightarrow 1]|.$$

Moreover, since a_λ only makes constant rounds of queries, all operations in b_λ^2 are performed in NC^1 . Hence we have $\mathcal{B}_2 = \{b_\lambda^2\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$, completing this part of proof. \square

Lemma 4.8. $\Pr[\mathbf{G}_5^{a_\lambda} \Rightarrow 1] = \Pr[\mathbf{G}_4^{a_\lambda} \Rightarrow 1]$.

Proof. In \mathbf{G}_5 , we do not use $(\mathbf{Y}_i)_{i=0}^\ell$ and \mathbf{y}' in $\text{USKGEN}(\text{id})$ or $\text{ENC}(\text{id}^*)$ any more. We give the sampling procedure for \mathbf{A} in an explicit way and change the simulation of \mathbf{Z}_i , \mathbf{z}' , \mathbf{v} , \mathbf{c}_1^* , and \mathbf{K}^* as in Fig. 12. We now show that all the changes are purely conceptual.

In \mathbf{G}_5 , we generate \mathbf{A} by sampling $\mathbf{R}_1 = \begin{pmatrix} \mathbf{I}_{\lambda-1} & 0 \\ \tilde{\mathbf{r}}^\top & 1 \end{pmatrix}^\top \stackrel{\$}{\leftarrow} \text{RSamp}(\lambda)$ and $\mathbf{R}_0 \stackrel{\$}{\leftarrow} \text{LSamp}(\lambda)$, and setting $\mathbf{A}^\top = \mathbf{R}_0 \mathbf{M}_0^\lambda \mathbf{R}_1$. This is exactly the “zero-sampling” procedure, in which case, we have

$$\begin{aligned}
\mathbf{Z}_i &= (\mathbf{Y}_i^\top \parallel \mathbf{x}_i) \mathbf{A} = (\mathbf{Y}_i^\top \parallel \mathbf{x}_i) \mathbf{R}_1^\top \mathbf{M}_0^{\lambda \top} \mathbf{R}_0^\top \\
&= (\mathbf{Y}_i^\top \parallel \mathbf{x}_i) \begin{pmatrix} \mathbf{I}_{\lambda-1} & \mathbf{0} \\ \tilde{\mathbf{r}}^\top & 1 \end{pmatrix} \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \ddots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & & 1 \\ 0 & & \cdots & & 0 \end{pmatrix} \mathbf{R}_0^\top \\
&= (\mathbf{Y}_i^\top + \mathbf{x}_i \cdot \tilde{\mathbf{r}}^\top \parallel \mathbf{x}_i) \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \ddots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & & 1 \\ 0 & & \cdots & & 0 \end{pmatrix} \mathbf{R}_0^\top \\
&= (\mathbf{0} \parallel \mathbf{Y}_i^\top + \mathbf{x}_i \cdot \tilde{\mathbf{r}}^\top) \mathbf{R}_0^\top = (\mathbf{0} \parallel \mathbf{D}_i) \mathbf{R}_0^\top
\end{aligned}$$

and

$$\mathbf{c}_1^* = \sum_{i=1}^{\ell} \text{id}_i^* \odot (\mathbf{Y}_i^\top \parallel \mathbf{x}_i) (\mathbf{A} + \mathbf{N}^\lambda) \mathbf{r} = \sum_{i=1}^{\ell} \text{id}_i^* \odot (\mathbf{Z}_i \mathbf{r} + \mathbf{x}_i).$$

Hence, the distributions of \mathbf{Z}_i and \mathbf{c}_1^* in \mathbf{G}_5 remain the same, and the distributions of \mathbf{z}' and \mathbf{K}^* can be analyzed in the same way. The distribution of \mathbf{v} does not change as well since

$$\begin{aligned}
\mathbf{v} &= \mathbf{t}^\top \left(\mathbf{Y}_0^\top + \sum_{i=1}^{\ell} \text{id}_i \odot \mathbf{Y}_i^\top \right) + \mathbf{y}'^\top \\
&= \mathbf{t}^\top \left(\mathbf{Y}_0^\top + \mathbf{x}_0 \cdot \tilde{\mathbf{r}}^\top + \sum_{i=1}^{\ell} \text{id}_i \odot (\mathbf{Y}_i^\top + \mathbf{x}_i \cdot \tilde{\mathbf{r}}^\top) \right) + (\mathbf{y}'^\top + \mathbf{x}' \cdot \tilde{\mathbf{r}}^\top) \\
&\quad - \left(\mathbf{t}^\top \left(\mathbf{x}_0 + \sum_{i=1}^{\ell} \text{id}_i \odot \mathbf{x}_i \right) + \mathbf{x}' \right) \cdot \tilde{\mathbf{r}}^\top \\
&= \mathbf{t}^\top \left(\mathbf{D}_0 + \sum_{i=1}^{\ell} \text{id}_i \odot \mathbf{D}_i \right) + \mathbf{d}' - \mathbf{u} \cdot \tilde{\mathbf{r}}^\top.
\end{aligned}$$

Putting all above together, Lemma 4.8 immediately follows. \square

Lemma 4.9. *There exists an adversary $\mathcal{B}_3 = \{b_\lambda^3\}_{\lambda \in \mathbb{N}} \in \mathbf{NC}^1$ such that b_λ^3 breaks the \mathbf{NC}^1 -(k, l)-PR-CMA-security of MAC with advantage*

$$|\Pr[\mathbf{G}_6^{a_\lambda} \Rightarrow 1] - \Pr[\mathbf{G}_5^{a_\lambda} \Rightarrow 1]|.$$

<p><u>INIT:</u></p> $\mathbf{R}_1 = \begin{pmatrix} \mathbf{I}_{\lambda-1} & \mathbf{0} \\ \tilde{\mathbf{r}}^\top & 1 \end{pmatrix}^\top \stackrel{\$}{\leftarrow} \text{RSamp}(\lambda),$ $\mathbf{R}_0 \stackrel{\$}{\leftarrow} \text{LSamp}(\lambda), \mathbf{A}^\top = \mathbf{R}_0 \mathbf{M}_0^\top \mathbf{R}_1$ <p>For $i = 0, \dots, \ell$:</p> $\mathbf{D}_i \stackrel{\$}{\leftarrow} \{0, 1\}^{\lambda \times (\lambda-1)}$ $\mathbf{Z}_i = (\mathbf{0} \parallel \mathbf{D}_i) \mathbf{R}_1^\top \in \{0, 1\}^{\lambda \times \lambda}$ $\mathbf{d}' \stackrel{\$}{\leftarrow} \{0, 1\}^{1 \times (\lambda-1)}$ $\mathbf{z}' = (\mathbf{0} \parallel \mathbf{d}') \mathbf{R}_0^\top \in \{0, 1\}^{1 \times \lambda}$ <p>Return $\text{pk} = (\mathbf{A}, \{\mathbf{Z}_i\}_{0 \leq i \leq \ell}, \mathbf{z}')$</p> <p><u>USKGEN(id):</u></p> $\mathcal{Q}_{\text{id}} = \mathcal{Q}_{\text{id}} \cup \{\text{id}\}$ $(\mathbf{t}, u) \stackrel{\$}{\leftarrow} \text{EVAL}(\text{id})$ $\mathbf{v} = (\mathbf{D}_0 + \sum_{i=1}^{\ell} \text{id}_i \odot \mathbf{D}_i)^\top \mathbf{t} + \mathbf{d}'^\top - \tilde{\mathbf{r}} \cdot u \in \{0, 1\}^{\lambda-1}$ <p>Return $\text{usk}[\text{id}] = (\mathbf{t}, u, \mathbf{v})$</p>	<p><u>ENC(id*):</u> //one query</p> $(\mathbf{h}_0, h_1) \stackrel{\$}{\leftarrow} \text{CHAL}(\text{id}^*)$ $r_2, \dots, r_n \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda$ $\mathbf{r} = (1, r_2, \dots, r_n)^\top$ $\mathbf{c}_0^* = (\mathbf{A} + \mathbf{N}^\lambda) \mathbf{r} \in \{0, 1\}^\lambda$ $\mathbf{c}_1^* = \mathbf{Z}_0 \mathbf{r} + \sum_{i=1}^{\ell} \text{id}_i^* \odot \mathbf{Z}_i \mathbf{r} + \mathbf{h}_0^\top \in \{0, 1\}^\lambda$ $\mathbf{K}^* = \mathbf{z} \cdot \mathbf{r} + h_1 \in \{0, 1\}$ <p>Return \mathbf{K}^* and $\text{ct}^* = (\mathbf{c}_0^*, \mathbf{c}_1^*)$</p> <p><u>FINALIZE($\beta$):</u></p> <p>Return $(\text{id}^* \notin \mathcal{Q}_{\text{id}}) \wedge \beta$</p>
---	--

Fig. 13. Description of $\mathcal{B}_3 = \{b_\lambda^3\}_{\lambda \in \mathbb{N}}$ (having access to the oracles $\text{INIT}_{\text{MAC}}, \text{EVAL}, \text{CHAL}, \text{FINALIZE}_{\text{MAC}}$ of the $\text{PR-CMA}_{\text{real}}/\text{PR-CMA}_{\text{rand}}$ games of Fig. 5 (instantiated with the encoding for equality predicate)) for the proof of Lemma 4.9.

Proof. The challenger of \mathbf{G}_6 answers the $\text{ENC}(\text{id}^*)$ query by choosing random \mathbf{K}^* . We build b_λ^3 as in Fig. 13 to show that the differences between \mathbf{G}_6 and \mathbf{G}_5 can be bounded by the advantage of breaking the PR-CMA security of MAC .

b_λ^3 runs in the same way as the challenger of \mathbf{G}_5 except that it samples \mathbf{D}_i and \mathbf{d}' uniformly at random from $\{0, 1\}^{\lambda \times (\lambda-1)}$ and $\{0, 1\}^{1 \times (\lambda-1)}$ respectively. This does not change the view of a_λ since \mathbf{Y}_i and \mathbf{y}' were uniformly sampled in \mathbf{G}_5 . Moreover, every time on receiving a query id to USKGEN , b_λ^3 forwards id to its evaluation oracle EVAL to obtain the answer (\mathbf{t}, u) , and on receiving the query id^* to ENC , b_λ^3 forwards id^* to its challenge oracle CHAL and uses the answer (\mathbf{h}_0, h_1) to simulate \mathbf{r} , \mathbf{c}_1^* , and \mathbf{K}^* as in Fig. 13. When a_λ outputs β , b_λ^3 outputs β as well if id^* was not queried to USKGEN . Otherwise, b_λ^3 outputs 0.

If h_1 is uniform (i.e., b_λ^3 is in Game $\text{PR-CMA}_{\text{rand}}$) then the view of a_λ is identical to its view in \mathbf{G}_6 . If h_1 is real (i.e., b_λ^3 is in Game $\text{PR-CMA}_{\text{real}}$), then the view of a_λ is identical to its view in \mathbf{G}_5 . Thus the advantage of b_λ^3 is exactly

$$|\Pr[\mathbf{G}_6^{a_\lambda} \Rightarrow 1] - \Pr[\mathbf{G}_5^{a_\lambda} \Rightarrow 1]|.$$

Moreover, since all operations in b_λ^3 are performed in NC^1 , we have $\mathcal{B}_3 = \{b_\lambda^3\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$, completing this part of proof. \square

We now do all the previous steps in the reverse order as in Fig. 14. Note that the view of the adversary in \mathbf{H}_0 (respectively, \mathbf{H}_4) is identical to its view in \mathbf{G}_6 (respectively, $\text{PR-ID-CPA}_{\text{rand}}$). By using the above arguments in a reverse order, we have the following lemma.

Lemma 4.10. *There exists an adversary $\mathcal{B}_4 = \{b_\lambda^4\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$ such that b_λ^4 breaks the fine-grained matrix linear assumption with advantage*

$$(|\Pr[\mathbf{H}_4^{a_\lambda} \Rightarrow 1] - \Pr[\mathbf{H}_0^{a_\lambda} \Rightarrow 1]|) / 2.$$

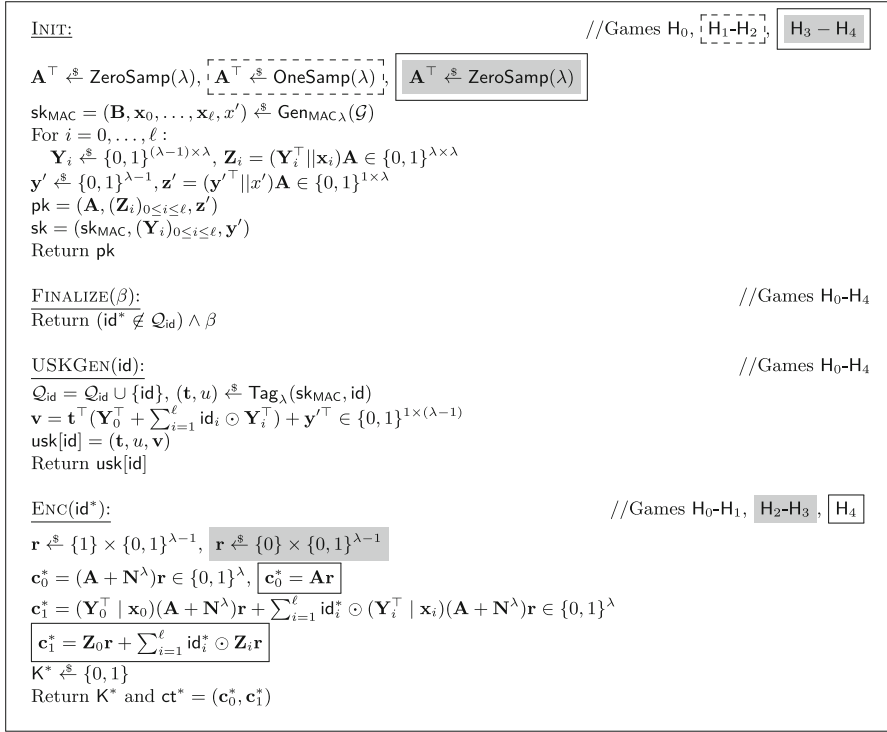


Fig. 14. Games H_0-H_4 for the proof of Theorem 4.3.

Putting all above together, Theorem 4.3 immediately follows. \square

Extension to IBKEM with Large Key Space. The key space of the above IBKEM is $\{0, 1\}$, while by running it in parallel, we can easily extend it to an IBKEM with large key space. The resulting scheme can still be performed in $\text{AC}^0[2]$ since running in parallel does not increase the circuit depth. The same extension can be also made for our fine-grained secure ABKEM given later in Sect. 5.

Extension to QA-NIZK. Our techniques for proving the hiding property of the underlying commitment scheme in our IBKEM can also be used to construct an efficient fine-grained QA-NIZK in NC^1 with adaptive soundness. We refer the reader to ‘‘Appendix C’’ for details.

5. Fine-Grained Secure Attribute-Based Encryption

In this section, we generalize our IBE scheme as a fine-grained ABE scheme by using predicate encodings [10, 36]. By instantiating the underlying encodings in different ways, we can achieve ABEs for inner product, non-zero inner product, spatial encryption, doubly spatial encryption, boolean span programs, and arithmetic span programs, and also broadcast encryption and fuzzy IBE schemes, which are computable in $\text{AC}^0[2]$ and secure against NC^1 under $\text{NC}^1 \subsetneq \oplus \text{L/poly}$. We refer the reader to ‘‘Appendix C’’ for

<p><u>Gen$_{\lambda}$</u>:</p> <p>$\mathbf{A}^T \xleftarrow{\\$} \text{ZeroSamp}(\lambda)$</p> <p>$\text{sk}_{\text{MAC}} = (\mathbf{B}, \mathbf{X}, x') \xleftarrow{\\$} \text{Gen}_{\text{MAC}_{\lambda}}(\text{par})$</p> <p>For $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_{\ell})$ and $i = 1, \dots, \ell$:</p> <p style="padding-left: 20px;">$\mathbf{Y}_i \xleftarrow{\\$} \{0, 1\}^{(\lambda-1) \times \lambda}$</p> <p style="padding-left: 20px;">$\mathbf{Z}_i = (\mathbf{Y}_i^T \mathbf{x}_i) \mathbf{A} \in \{0, 1\}^{\lambda \times \lambda}$</p> <p>$\mathbf{y}' \xleftarrow{\\$} \{0, 1\}^{(\lambda-1)}$</p> <p>$\mathbf{z}' = (\mathbf{y}'^T x') \mathbf{A} \in \{0, 1\}^{1 \times \lambda}$</p> <p>$\text{pk} = (\mathbf{A}, (\mathbf{Z}_i)_{1 \leq i \leq \ell}, \mathbf{z}')$</p> <p>$\text{sk} = (\text{sk}_{\text{MAC}}, (\mathbf{Y}_i)_{1 \leq i \leq \ell}, \mathbf{y}')$</p> <p>Return (pk, sk)</p> <p><u>USKGen$_{\lambda}$</u>($\text{sk}, \mathbf{y} \in \mathcal{Y}$):</p> <p>$(\mathbf{t}, \mathbf{u}) \xleftarrow{\\$} \text{Tag}_{\lambda}(\text{sk}_{\text{MAC}}, \mathbf{y})$</p> <p>$\mathbf{v} = \text{rE}_{\lambda}(\mathbf{y}, \begin{pmatrix} \mathbf{t}^T \mathbf{Y}_1^T \\ \vdots \\ \mathbf{t}^T \mathbf{Y}_{\ell}^T \end{pmatrix})$</p> <p style="padding-left: 20px;">$+ \text{kE}_{\lambda}(\mathbf{y}, \mathbf{y}'^T) \in \{0, 1\}^{\eta \times (\lambda-1)}$</p> <p>Return $\text{usk}[\mathbf{y}] = (\mathbf{t}, \mathbf{u}, \mathbf{v})$</p>	<p><u>Enc$_{\lambda}$</u>($\text{pk}, \mathbf{x} \in \mathcal{X}$):</p> <p>$\mathbf{r} \xleftarrow{\\$} \{0\} \times \{0, 1\}^{\lambda-1}$</p> <p>$\mathbf{c}_0 = \mathbf{A} \mathbf{r} \in \{0, 1\}^{\lambda}$</p> <p>$\mathbf{C}_1 = \text{sE}_{\lambda}(\mathbf{x}, \begin{pmatrix} \mathbf{r}^T \mathbf{Z}_1^T \\ \vdots \\ \mathbf{r}^T \mathbf{Z}_{\ell}^T \end{pmatrix}) \in \{0, 1\}^{\zeta \times \lambda}$</p> <p>$\mathbf{K} = \mathbf{z}' \cdot \mathbf{r} \in \{0, 1\}$.</p> <p>Return \mathbf{K} and $\text{ct} = (\mathbf{c}_0, \mathbf{C}_1)$</p> <p><u>Dec$_{\lambda}$</u>($\text{usk}[\mathbf{y}], \mathbf{y}, \mathbf{x}, \text{ct}$):</p> <p>Parse $\text{usk}[\mathbf{y}] = (\mathbf{t}, \mathbf{u}, \mathbf{v})$</p> <p>Parse $\text{ct} = (\mathbf{c}_0, \mathbf{C}_1)$</p> <p>$\mathbf{K} = \text{rD}_{\lambda}(\mathbf{x}, \mathbf{y}, \mathbf{v} \mathbf{u}) \mathbf{c}_0$</p> <p style="padding-left: 20px;">$- \text{sD}_{\lambda}(\mathbf{x}, \mathbf{y}, \mathbf{C}_1 \mathbf{t}) \in \{0, 1\}$</p> <p>Return \mathbf{K}</p>
--	---

Fig. 15. Construction of $\text{ABKEM} = \{\text{Gen}_{\lambda}, \text{USKGen}_{\lambda}, \text{Enc}_{\lambda}, \text{Dec}_{\lambda}\}_{\lambda \in \mathbb{N}}$.

several instances of the encodings and also to [10] for more instances. We note that the encodings in [10] are defined over $GF(p)$, while the ours are over $GF(2)$. However, the proofs for encodings in [10] can be adopted in our case, since the linearity and α -reconstruction properties hold in $GF(p)$ also hold in $GF(2)$ and by the standard linear-independence arguments in $GF(2)$, the α -privacy also holds in our case.

Let $\text{PE} = \{\text{rE}_{\lambda}, \text{kE}_{\lambda}, \text{sE}_{\lambda}, \text{sD}_{\lambda}, \text{rD}_{\lambda}\}_{\lambda \in \mathbb{N}} \in \text{AC}^0[2]$ be a predicate encoding for $\text{P} = \{\text{p}_{\lambda}\}_{\lambda \in \mathbb{N}}$ with $\text{rE}_{\lambda} : \mathcal{Y} \times \{0, 1\}^{\ell} \rightarrow \{0, 1\}^{\eta}$, $\text{kE}_{\lambda} : \mathcal{Y} \times \{0, 1\} \rightarrow \{0, 1\}^{\eta}$, $\text{sE}_{\lambda} : \mathcal{X} \times \{0, 1\}^{\ell} \rightarrow \{0, 1\}^{\zeta}$, $\text{sD}_{\lambda} : \mathcal{X} \times \mathcal{Y} \times \{0, 1\}^{\zeta} \rightarrow \{0, 1\}$, and $\text{rD}_{\lambda} : \mathcal{X} \times \mathcal{Y} \times \{0, 1\}^{\eta} \rightarrow \{0, 1\}$. Let $\text{MAC}_{\text{GA}} = \{\text{Gen}_{\text{MAC}_{\lambda}}, \text{Tag}_{\lambda}, \text{Ver}_{\text{MAC}_{\lambda}}\}_{\lambda \in \mathbb{N}} \in \text{AC}^0[2]$ be a PE-generalized affine MAC over $\{0, 1\}^{\lambda}$ with message space \mathcal{Y} . Our ABKEM $\text{ABKEM} = \{\text{Gen}_{\lambda}, \text{USKGen}_{\lambda}, \text{Enc}_{\lambda}, \text{Dec}_{\lambda}\}_{\lambda \in \mathbb{N}}$ is defined as in Fig. 15.

Theorem 5.1. *Under the assumption $\text{NC}^1 \not\subseteq \oplus \text{L}/\text{poly}$ and the NC^1 - (k, l) - sE_{λ} -PR-CMA-security of MAC_{GA} , where k is any constant and $l = l(\lambda)$ is any polynomial in λ , ABKEM is an $\text{AC}^0[2]$ -ABKEM that is NC^1 - (k, l) -PR-AT-CPA secure against NC^1 .*

Proof. First, we note that $\{\text{Gen}_{\lambda}\}_{\lambda \in \mathbb{N}}$, $\{\text{USKGen}_{\lambda}\}_{\lambda \in \mathbb{N}}$, $\{\text{Enc}_{\lambda}\}_{\lambda \in \mathbb{N}}$, and $\{\text{Dec}_{\lambda}\}_{\lambda \in \mathbb{N}}$ are computable in $\text{AC}^0[2]$, since they only involve operations including multiplication of a constant number of matrices, sampling random bits, and running $\text{MAC}_{\text{GA}} \in \text{AC}^0[2]$.

By Eq. (2) in Sect. 3.1, we have

$$\begin{aligned} & \text{rD}_{\lambda}(\mathbf{x}, \mathbf{y}, \mathbf{v} || \mathbf{u}) \mathbf{c}_0 \\ &= \text{rD}_{\lambda}(\mathbf{x}, \mathbf{y}, \text{rE}_{\lambda} \left(\mathbf{y}, \begin{pmatrix} \mathbf{t}^T \mathbf{Y}_1^T \\ \vdots \\ \mathbf{t}^T \mathbf{Y}_{\ell}^T \end{pmatrix} + \text{kE}_{\lambda}(\mathbf{y}, \mathbf{y}'^T) \right) || \begin{pmatrix} \mathbf{t}^T \mathbf{x}_1 \\ \vdots \\ \mathbf{t}^T \mathbf{x}_{\ell} \end{pmatrix} + \text{kE}_{\lambda}(\mathbf{y}, x')) \mathbf{A} \mathbf{r} \end{aligned}$$

and

$$\text{sD}_\lambda(\mathbf{x}, \mathbf{y}, \mathbf{C}_1 \mathbf{t}) = \text{sD}_\lambda(\mathbf{x}, \mathbf{y}, \text{sE}_\lambda \left(\mathbf{x}, \begin{pmatrix} \mathbf{t}^\top (\mathbf{Y}_1^\top || \mathbf{x}_1) \\ \vdots \\ \mathbf{t}^\top (\mathbf{Y}_\ell^\top || \mathbf{x}_\ell) \end{pmatrix} \right) \mathbf{A} \mathbf{r}).$$

Then, due to restricted α -reconstruction (see Definition 2.14), the difference of the above equations yields $\mathbf{K} = (\mathbf{y}'^\top || \mathbf{x}') \mathbf{A} \mathbf{r} = \mathbf{z}' \cdot \mathbf{r}$, i.e., correctness is satisfied.

Let $\mathcal{A} = \{a_\lambda\}_{\lambda \in \mathbb{N}}$ be any adversary against the NC^1 - (k, l) -PR-AT-CPA security of ABKEM. We now prove the NC^1 - (k, l) -PR-AT-CPA security by defining a sequence of games \mathbf{G}_0 - \mathbf{G}_6 as in Fig. 16. Roughly, in the first four games, we show how to extract a challenge token for MAC_{GA} from the challenge session key and ciphertext by switching the distribution of \mathbf{A} twice and changing the distribution of the randomness \mathbf{r} during the switching procedure. In the last two games, we show that the commitments \mathbf{Z}_i and \mathbf{z}' perfectly hide the secrets, and the answers of queries made by a_λ reveal no useful information other than the tags and token for MAC. \square

Lemma 5.2. $\Pr[\text{PR-AT-CPA}_{\text{real}}^{a_\lambda} \Rightarrow 1] = \Pr[\mathbf{G}_1^{a_\lambda} \Rightarrow 1] = \Pr[\mathbf{G}_0^{a_\lambda} \Rightarrow 1].$

Proof. \mathbf{G}_0 is the real attack game. In game \mathbf{G}_1 , we change the simulation of \mathbf{c}_0^* , \mathbf{C}_1^* and \mathbf{K}^* in $\text{ENC}(\mathbf{x})$ by substituting \mathbf{Z}_i and \mathbf{z}' with their respective definitions and substituting \mathbf{A} with $\mathbf{A} + \mathbf{N}^\lambda$. Since we have

$$\mathbf{N}^\lambda \mathbf{r} = \begin{pmatrix} 0 & \cdots & 0 \\ \vdots & 0 & \cdots & 0 \\ 0 & & \ddots & \vdots \\ 1 & 0 & \cdots & 0 \end{pmatrix} \begin{pmatrix} 0 \\ r_2 \\ \vdots \\ r_\lambda \end{pmatrix} = 0,$$

the view of a_λ in \mathbf{G}_1 is identical to its view in \mathbf{G}_0 , completing this part of proof. \square

Lemma 5.3. *There exists an adversary $\mathcal{B}_1 = \{b_\lambda^1\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$ such that b_λ^1 breaks the fine-grained matrix linear assumption (see Definition 2.5), which holds under $\text{NC}^1 \subsetneq \oplus \mathbb{L}/\text{poly}$ according to Theorem 2.13, with advantage*

$$|\Pr[\mathbf{G}_2^{a_\lambda} \Rightarrow 1] - \Pr[\mathbf{G}_1^{a_\lambda} \Rightarrow 1]|.$$

Proof. \mathbf{G}_1 and \mathbf{G}_2 only differ in the distribution of \mathbf{A} , namely, $\mathbf{A}^\top \stackrel{\$}{\leftarrow} \text{ZeroSamp}(\lambda)$ or $\mathbf{A}^\top \stackrel{\$}{\leftarrow} \text{OneSamp}(\lambda)$, and we build the distinguisher b_λ^1 as follows.

b_λ^1 runs in exactly the same way as the challenger of \mathbf{G}_1 except that in INIT, instead of generating \mathbf{A} by itself, it takes as input \mathbf{A}^\top generated as $\mathbf{A}^\top \stackrel{\$}{\leftarrow} \text{ZeroSamp}(\lambda)$ or $\mathbf{A}^\top \stackrel{\$}{\leftarrow} \text{OneSamp}(\lambda)$ from its own challenger. When a_λ outputs β , b_λ^1 outputs β as well if no \mathbf{y} such that $\mathbf{p}_\lambda(\mathbf{x}, \mathbf{y}) = 1$ was queried to USKGEN. Otherwise, b_λ^1 outputs 0.

INIT: //Games G_0 - G_1 , G_2 - G_3 , G_4 , G_5 - G_6

$A^\top \xleftarrow{\$} \text{ZeroSamp}(\lambda)$, $A^\top \xleftarrow{\$} \text{OneSamp}(\lambda)$, $A^\top \xleftarrow{\$} \text{ZeroSamp}(\lambda)$

$R_1 = \begin{pmatrix} I_{\lambda-1} & 0 \\ \tilde{r}^\top & 1 \end{pmatrix}^\top \xleftarrow{\$} \text{RSamp}(\lambda)$, $R_0 \xleftarrow{\$} \text{LSamp}(\lambda)$, $A^\top = R_0 M_0^\dagger R_1$

$\text{sk}_{\text{MAC}} = (\mathbf{B}, \mathbf{X}, x') \xleftarrow{\$} \text{Gen}_{\text{MAC}\lambda}(\mathcal{G})$

For $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_\ell)$ and $i = 1, \dots, \ell$:

$\mathbf{Y}_i \xleftarrow{\$} \{0, 1\}^{(\lambda-1) \times \lambda}$, $\mathbf{Z}_i = (\mathbf{Y}_i^\top \parallel \mathbf{x}_i) \mathbf{A} \in \{0, 1\}^{\lambda \times \lambda}$

$\mathbf{D}_i = \mathbf{Y}_i^\top + \mathbf{x}_i \cdot \tilde{r}^\top \in \{0, 1\}^{\lambda \times (\lambda-1)}$, $\mathbf{Z}_i = (0 \parallel \mathbf{D}_i) \mathbf{R}_0^\top \in \{0, 1\}^{\lambda \times \lambda}$

$\mathbf{y}' \xleftarrow{\$} \{0, 1\}^{\lambda-1}$, $\mathbf{z}' = (\mathbf{y}'^\top \parallel x') \mathbf{A} \in \{0, 1\}^{1 \times \lambda}$

$\mathbf{d}' = \mathbf{y}'^\top + x' \cdot \tilde{r}^\top \in \{0, 1\}^{1 \times (\lambda-1)}$, $\mathbf{z}' = (0 \parallel \mathbf{d}') \mathbf{R}_0^\top \in \{0, 1\}^{1 \times \lambda}$

$\text{pk} = (\mathbf{A}, (\mathbf{Z}_i)_{1 \leq i \leq \ell}, \mathbf{z}')$, $\text{sk} = (\text{sk}_{\text{MAC}}, (\mathbf{Y}_i)_{1 \leq i \leq \ell}, \mathbf{y}')$

Return pk

FINALIZE(β): //Games G_0 - G_6

If $(\text{pk}(x, y) \neq 1)$ for all $y \in Q_y$, return β

Else return 0

USKGEN(y): //Games G_0 - G_4 , G_5 - G_6

$Q_y = Q_y \cup \{y\}$, $(\mathbf{t}, \mathbf{u}) \xleftarrow{\$} \text{Tag}_\lambda(\text{sk}_{\text{MAC}}, y)$

$\mathbf{v} = r \mathbf{E}_\lambda(y, \begin{pmatrix} \mathbf{t}^\top \mathbf{Y}_1^\top \\ \vdots \\ \mathbf{t}^\top \mathbf{Y}_\ell^\top \end{pmatrix}) + k \mathbf{E}_\lambda(y, \mathbf{y}'^\top) \in \{0, 1\}^{\eta \times (\lambda-1)}$

$\mathbf{v} = r \mathbf{E}_\lambda(y, (\mathbf{D}_1^\top \mathbf{t}, \dots, \mathbf{D}_\ell^\top \mathbf{t})^\top) + k \mathbf{E}_\lambda(y, \mathbf{d}') - \mathbf{u} \cdot \tilde{r}^\top \in \{0, 1\}^{\eta \times (\lambda-1)}$

$\text{usk}[y] = (\mathbf{t}, \mathbf{u}, \mathbf{v})$

Return $\text{usk}[y]$

ENC(x): //Games G_0 , G_1 - G_4 , G_3 - G_4 , G_5 , G_6

$\mathbf{r} \xleftarrow{\$} \{0\} \times \{0, 1\}^{\lambda-1}$, $\mathbf{r} \xleftarrow{\$} \{1\} \times \{0, 1\}^{\lambda-1}$

$\mathbf{c}_0^* = \mathbf{A} \mathbf{r} \in \{0, 1\}^\lambda$, $\mathbf{c}_0^* = (\mathbf{A} + \mathbf{N}^\lambda) \mathbf{r}$

$\mathbf{C}_1^* = \text{sE}_\lambda(x, \begin{pmatrix} \mathbf{r}^\top \mathbf{Z}_1^\top \\ \vdots \\ \mathbf{r}^\top \mathbf{Z}_\ell^\top \end{pmatrix}) \in \{0, 1\}^{c_\lambda}$

$\mathbf{C}_1^* = \text{sE}_\lambda(x, ((\mathbf{Y}_1^\top \parallel \mathbf{x}_1)(\mathbf{A} + \mathbf{N}^\lambda) \mathbf{r}, \dots, (\mathbf{Y}_\ell^\top \parallel \mathbf{x}_\ell)(\mathbf{A} + \mathbf{N}^\lambda) \mathbf{r})^\top)$

$\mathbf{C}_1^* = \text{sE}_\lambda(x, (\mathbf{Z}_1 \mathbf{r}, \dots, \mathbf{Z}_\ell \mathbf{r})^\top) + \text{sE}_\lambda(x, (\mathbf{x}_1, \dots, \mathbf{x}_\ell)^\top)$

$K^* = \mathbf{z}' \cdot \mathbf{r} \in \{0, 1\}$, $K^* = (\mathbf{y}'^\top \parallel x')(\mathbf{A} + \mathbf{N}^\lambda) \mathbf{r}$, $K^* = \mathbf{z}' \cdot \mathbf{r} + x'$

$K^* \xleftarrow{\$} \{0, 1\}$

Return K^* and $\text{ct}^* = (\mathbf{c}_0^*, \mathbf{C}_1^*)$

Fig. 16. Games G_0 - G_6 for the proof of Theorem 5.1.

If \mathbf{A} is generated as $\mathbf{A}^\top \xleftarrow{\$} \text{ZeroSamp}(\lambda)$ (respectively, $\mathbf{A}^\top \xleftarrow{\$} \text{OneSamp}(\lambda)$), the view of a_λ is the same as its view in G_1 (respectively, G_2). Hence, the probability that b_λ^1 breaks the fine-grained matrix linear assumption is

$$\left| \Pr [G_2^{a_\lambda} \Rightarrow 1] - \Pr [G_1^{a_\lambda} \Rightarrow 1] \right|.$$

Moreover, since a_λ only makes constant rounds of queries, all operations in b_λ^1 are performed in NC^1 . Hence, we have $\mathcal{B}_1 = \{b_\lambda^1\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$, completing this part of proof. \square

Lemma 5.4. $\Pr[\mathbf{G}_3^{a_\lambda} \Rightarrow 1] = \Pr[\mathbf{G}_2^{a_\lambda} \Rightarrow 1]$.

Proof. In this game, we sample \mathbf{r} in $\text{ENC}(\mathbf{x})$ as $\mathbf{r} \xleftarrow{\$} \{0, 1\}^\lambda$ instead of $\mathbf{r} \xleftarrow{\$} \{0\} \times \{0, 1\}^{\lambda-1}$. According to Lemma 2.9, the distributions of $\mathbf{A} + \mathbf{N}^\lambda$ in both \mathbf{G}_2 and \mathbf{G}_3 are identical to that of a matrix sampled from ZeroSamp . Then this lemma follows from Lemma 2.11 immediately. \square

Lemma 5.5. *There exists an adversary $\mathcal{B}_2 = \{b_\lambda^2\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$ such that b_λ^2 breaks the fine-grained matrix linear assumption with advantage*

$$|\Pr[\mathbf{G}_4^{a_\lambda} \Rightarrow 1] - \Pr[\mathbf{G}_3^{a_\lambda} \Rightarrow 1]|.$$

Proof. \mathbf{G}_1 and \mathbf{G}_2 only differ in the distribution of \mathbf{A} , namely, $\mathbf{A}^\top \xleftarrow{\$} \text{OneSamp}(\lambda)$ or $\mathbf{A}^\top \xleftarrow{\$} \text{ZeroSamp}(\lambda)$, and we build the distinguisher b_λ^2 against Lemma 2.6 as follows.

b_λ^2 runs in exactly the same way as the challenger of \mathbf{G}_3 except that in INIT , instead of generating \mathbf{A} by itself, it takes as input \mathbf{A}^\top generated as $\mathbf{A}^\top \xleftarrow{\$} \text{ZeroSamp}(\lambda)$ or $\mathbf{A}^\top \xleftarrow{\$} \text{OneSamp}(\lambda)$ from its own challenger. When a_λ outputs β , b_λ^2 outputs β as well if no \mathbf{y} such that $\mathbf{p}_\lambda(\mathbf{x}, \mathbf{y}) = 1$ was queried to USKGEN . Otherwise, b_λ^2 outputs 0.

If \mathbf{A} is generated as $\mathbf{A}^\top \xleftarrow{\$} \text{OneSamp}(\lambda)$ (respectively, $\mathbf{A}^\top \xleftarrow{\$} \text{ZeroSamp}(\lambda)$), the view of a_λ is the same as its view in \mathbf{G}_3 (respectively, \mathbf{G}_4). Hence, the probability that b_λ^2 breaks the fine-grained matrix linear assumption is

$$|\Pr[\mathbf{G}_4^{a_\lambda} \Rightarrow 1] - \Pr[\mathbf{G}_3^{a_\lambda} \Rightarrow 1]|.$$

Moreover, since a_λ only makes constant rounds of queries, all operations in b_λ^2 are performed in NC^1 . Hence, we have $\mathcal{B}_2 = \{b_\lambda^2\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$, completing this part of proof. \square

Lemma 5.6. $\Pr[\mathbf{G}_5^{a_\lambda} \Rightarrow 1] = \Pr[\mathbf{G}_4^{a_\lambda} \Rightarrow 1]$.

Proof. In \mathbf{G}_5 , we do not use $(\mathbf{Y}_i)_{i=1}^\ell$ and \mathbf{y}' in $\text{USKGEN}(\mathbf{y})$ or $\text{ENC}(\mathbf{x})$ any more. We give the sampling procedure for \mathbf{A} in an explicit way and change the simulation of \mathbf{Z}_i , \mathbf{z}' , \mathbf{v} , \mathbf{C}_1^* , and \mathbf{K}^* as in Fig. 16. We now show that all the changes are purely conceptual.

In \mathbf{G}_5 , we generate \mathbf{A} by sampling $\mathbf{R}_1 = \begin{pmatrix} \mathbf{I}_{\lambda-1} & \mathbf{0} \\ \tilde{\mathbf{r}}^\top & 1 \end{pmatrix}^\top \stackrel{\mathcal{L}}{\leftarrow} \text{RSamp}(\lambda)$ and $\mathbf{R}_0 \stackrel{\mathcal{L}}{\leftarrow} \text{LSamp}(\lambda)$, and setting $\mathbf{A}^\top = \mathbf{R}_0 \mathbf{M}_0^\lambda \mathbf{R}_1$. This is exactly the “zero-sampling” procedure, in which case, we have

$$\begin{aligned} \mathbf{Z}_i &= (\mathbf{Y}_i^\top \parallel \mathbf{x}_i) \mathbf{A} = (\mathbf{Y}_i^\top \parallel \mathbf{x}_i) \mathbf{R}_1^\top \mathbf{M}_0^{\lambda \top} \mathbf{R}_0^\top \\ &= (\mathbf{Y}_i^\top \parallel \mathbf{x}_i) \begin{pmatrix} \mathbf{I}_{\lambda-1} & \mathbf{0} \\ \tilde{\mathbf{r}}^\top & 1 \end{pmatrix} \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \ddots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & & 1 \\ 0 & & \cdots & & 0 \end{pmatrix} \mathbf{R}_0^\top \\ &= (\mathbf{Y}_i^\top + \mathbf{x}_i \cdot \tilde{\mathbf{r}}^\top \parallel \mathbf{x}_i) \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \ddots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & & 1 \\ 0 & & \cdots & & 0 \end{pmatrix} \mathbf{R}_0^\top \\ &= (\mathbf{0} \parallel \mathbf{Y}_i^\top + \mathbf{x}_i \cdot \tilde{\mathbf{r}}^\top) \mathbf{R}_0^\top = (\mathbf{0} \parallel \mathbf{D}_i) \mathbf{R}_0^\top \end{aligned}$$

and

$$\begin{aligned} \mathbf{C}_1^* &= \mathbf{sE}_\lambda(\mathbf{x}, ((\mathbf{Y}_1^\top \parallel \mathbf{x}_1)(\mathbf{A} + \mathbf{N}^\lambda) \mathbf{r}, \dots, (\mathbf{Y}_\ell^\top \parallel \mathbf{x}_\ell)(\mathbf{A} + \mathbf{N}^\lambda) \mathbf{r})^\top) \\ &= \mathbf{sE}_\lambda(\mathbf{x}, (\mathbf{Z}_1 \mathbf{r} + \mathbf{x}_1, \dots, \mathbf{Z}_\ell \mathbf{r} + \mathbf{x}_\ell)^\top) \\ &= \mathbf{sE}_\lambda(\mathbf{x}, (\mathbf{Z}_1 \mathbf{r}, \dots, \mathbf{Z}_\ell \mathbf{r})^\top) + \mathbf{sE}_\lambda(\mathbf{x}, (\mathbf{x}_1, \dots, \mathbf{x}_\ell)^\top). \end{aligned}$$

Hence, the distributions of \mathbf{Z}_i in \mathbf{G}_5 remain the same, and the distributions of \mathbf{z}' and \mathbf{K}^* can be analyzed in the same way. The distribution of \mathbf{v} does not change as well since

$$\begin{aligned} \mathbf{v} &= \mathbf{rE}_\lambda(\mathbf{y}, (\mathbf{Y}_1 \mathbf{t}, \dots, \mathbf{Y}_\ell \mathbf{t})^\top) + \mathbf{kE}_\lambda(\mathbf{y}, \mathbf{y}'^\top) \\ &= \mathbf{rE}_\lambda(\mathbf{y}, ((\mathbf{Y}_1 + \tilde{\mathbf{r}} \cdot \mathbf{x}_1^\top) \mathbf{t}, \dots, (\mathbf{Y}_\ell + \tilde{\mathbf{r}} \cdot \mathbf{x}_\ell^\top) \mathbf{t})^\top) + \mathbf{kE}_\lambda(\mathbf{y}, \mathbf{y}'^\top + \mathbf{x}' \cdot \tilde{\mathbf{r}}^\top) \\ &\quad - \left(\mathbf{rE}_\lambda(\mathbf{y}, (\tilde{\mathbf{r}} \cdot \mathbf{x}_1^\top \cdot \mathbf{t}, \dots, \tilde{\mathbf{r}} \cdot \mathbf{x}_\ell^\top \cdot \mathbf{t})^\top) + \mathbf{kE}_\lambda(\mathbf{y}, \mathbf{x}' \cdot \tilde{\mathbf{r}}^\top) \right) \\ &= \mathbf{rE}_\lambda(\mathbf{y}, (\mathbf{D}_1^\top \mathbf{t}, \dots, \mathbf{D}_\ell^\top \mathbf{t})^\top) + \mathbf{kE}_\lambda(\mathbf{y}, \mathbf{d}') - \mathbf{u} \cdot \tilde{\mathbf{r}}^\top. \end{aligned}$$

Putting all above together, Lemma 5.6 immediately follows. \square

<p>INIT:</p> $\mathbf{R}_1 = \begin{pmatrix} \mathbf{I}_{\lambda-1} & \mathbf{0} \\ \tilde{\mathbf{r}}^\top & 1 \end{pmatrix} \stackrel{\$}{\leftarrow} \text{RSamp}(\lambda),$ $\mathbf{R}_0 \stackrel{\$}{\leftarrow} \text{LSamp}(\lambda), \mathbf{A}^\top = \mathbf{R}_0 \mathbf{M}_0^\dagger \mathbf{R}_1$ <p>For $i = 1, \dots, \ell$:</p> $\mathbf{D}_i \stackrel{\$}{\leftarrow} \{0, 1\}^{\lambda \times (\lambda-1)}$ $\mathbf{Z}_i = (0 \parallel \mathbf{D}_i) \mathbf{R}_0^\top \in \{0, 1\}^{\lambda \times \lambda}$ $\mathbf{d}' \stackrel{\$}{\leftarrow} \{0, 1\}^{1 \times (\lambda-1)}, \mathbf{z}' = (0 \parallel \mathbf{d}') \mathbf{R}_0^\top \in \{0, 1\}^{1 \times \lambda}$ $\text{pk} = (\mathbf{A}, (\mathbf{Z}_i)_{1 \leq i \leq \ell}, \mathbf{z}')$ <p>Return pk</p> <p>USKGEN(y):</p> $\mathcal{Q}_y = \mathcal{Q}_y \cup \{y\}$ $(\mathbf{t}, \mathbf{u}) \stackrel{\$}{\leftarrow} \text{EVAL}(y)$ $\mathbf{v} = \text{rE}_\lambda(y, (\mathbf{D}_1^\top \mathbf{t}, \dots, \mathbf{D}_\ell^\top \mathbf{t})^\top) + \text{kE}_\lambda(y, \mathbf{d}') - \mathbf{u} \cdot \tilde{\mathbf{r}}^\top \in \{0, 1\}^{\eta \times (\lambda-1)}$ $\text{usk}[y] = (\mathbf{t}, \mathbf{u}, \mathbf{v})$ <p>Return usk[y]</p>	<p>ENC(x): //one query</p> $(\mathbf{h}_0, h_1) \stackrel{\$}{\leftarrow} \text{CHAL}(x)$ $r_2, \dots, r_n \stackrel{\$}{\leftarrow} \{0, 1\}$ $\mathbf{r} = (1, r_2, \dots, r_n)^\top$ $\mathbf{c}_0^* = (\mathbf{A} + \mathbf{N}^\lambda) \mathbf{r} \in \{0, 1\}^\lambda$ $\mathbf{C}_1^* = \text{sE}_\lambda(x, \begin{pmatrix} \mathbf{r}^\top \mathbf{Z}_1^\top \\ \vdots \\ \mathbf{r}^\top \mathbf{Z}_\ell^\top \end{pmatrix}) + \mathbf{h}_0 \in \{0, 1\}^{c \times \lambda}$ $\mathbf{K}^* = \mathbf{z}' \cdot \mathbf{r} + h_1 \in \{0, 1\}$ <p>Return \mathbf{K}^* and $\text{ct}^* = (\mathbf{c}_0^*, \mathbf{C}_1^*)$</p> <p>FINALIZE($\beta$):</p> <p>If $(\text{p}_\lambda(x, y) \neq 1$ for all $y \in \mathcal{Q}_y$)</p> <p> return β</p> <p>Else return 0</p>
---	---

Fig. 17. Description of $\mathcal{B}_3 = \{b_\lambda^3\}_{\lambda \in \mathbb{N}}$ (having access to the oracles INIT_{MAC} , EVAL , CHAL , $\text{FINALIZE}_{\text{MAC}}$ of the $\text{PR-CMA}_{\text{real}}/\text{PR-CMA}_{\text{rand}}$ games of Fig. 5) for the proof of Lemma 5.7.

Lemma 5.7. *There exists an adversary $\mathcal{B}_3 = \{b_\lambda^3\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$ such that b_λ^3 breaks the NC^1 - (k, l) - PR-CMA security of MAC_{GA} with advantage*

$$\left| \Pr [\mathbf{G}_6^{a_\lambda} \Rightarrow 1] - \Pr [\mathbf{G}_5^{a_\lambda} \Rightarrow 1] \right|.$$

Proof. The challenger of \mathbf{G}_6 answers the $\text{ENC}(\mathbf{x})$ query by choosing random \mathbf{K}^* . We build b_λ^3 as in Fig. 17 to show that the differences between \mathbf{G}_6 and \mathbf{G}_5 can be bounded by its advantage of breaking the PR-CMA security of MAC_{GA} .

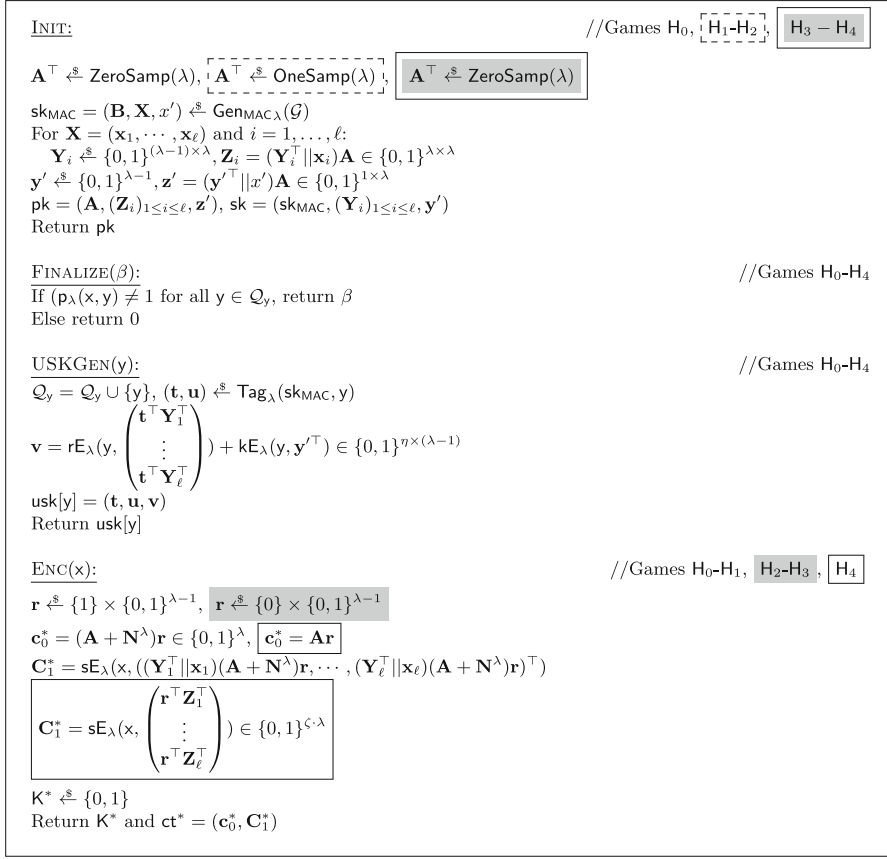

b_λ^3 runs in the same way as the challenger of \mathbf{G}_5 except that it samples \mathbf{D}_i and \mathbf{d}' uniformly at random from $\{0, 1\}^{\lambda \times (\lambda-1)}$ and $\{0, 1\}^{1 \times (\lambda-1)}$ respectively. This does not change the view of a_λ since \mathbf{Y}_i and \mathbf{y}' were uniformly sampled in \mathbf{G}_5 . Moreover, every time on receiving a query \mathbf{y} to USKGEN , b_λ^3 forwards \mathbf{y} to its evaluation oracle EVAL to obtain the answer (\mathbf{t}, \mathbf{u}) , and on receiving the query \mathbf{x} to ENC , b_λ^3 forwards \mathbf{x} to its challenge oracle CHAL and uses the answer (h, \mathbf{h}_0, h_1) to simulate \mathbf{r} , \mathbf{C}_1^* , and \mathbf{K}^* as in Fig. 17. When a_λ outputs β , b_λ^3 outputs β as well if no \mathbf{y} such that $\text{p}_\lambda(\mathbf{x}, \mathbf{y}) = 1$ was queried to USKGEN . Otherwise, b_λ^3 outputs 0.

If h_1 is uniform (i.e., b_λ^3 is in Game $\text{PR-CMA}_{\text{rand}}$) then the view of a_λ is identical to its view in \mathbf{G}_6 . If h_1 is real (i.e., b_λ^3 is in Game $\text{PR-CMA}_{\text{real}}$) then the view of \mathcal{A} is identical to its view in \mathbf{G}_5 . Hence, the advantage of b_λ^3 in breaking the PR-CMA security is

$$\left| \Pr [\mathbf{G}_6^{a_\lambda} \Rightarrow 1] - \Pr [\mathbf{G}_5^{a_\lambda} \Rightarrow 1] \right|.$$

Moreover, since a_λ only makes constant rounds of queries, all operations in b_λ^3 are performed in NC^1 . Hence, we have $\mathcal{B}_3 = \{b_\lambda^3\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$, completing this part of proof.

We now do all the previous steps in the reverse order as in Fig. 18. Note that the view of the adversary in \mathbf{H}_0 (respectively, \mathbf{H}_4) is identical to its view in \mathbf{G}_6 (respectively, $\text{PR-AT-CPA}_{\text{rand}}$). By using the above arguments in a reverse order, we have the following lemma. \square

Fig. 18. Games H_0-H_4 for the proof of Theorem 5.1.Fig. 19. Games G_0-G_5 for the proof of Proposition A.2.

Lemma 5.8. *There exists an adversary $\mathcal{B}_4 = \{b_\lambda^4\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$ such that b_λ^4 breaks the fine-grained matrix linear assumption with advantage*

$$(|\Pr [H_4^{a_\lambda} \Rightarrow 1] - \Pr [H_0^{a_\lambda} \Rightarrow 1]|) / 2.$$

Putting all above together, Theorem 5.1 immediately follows. \square

Acknowledgements

We would like to thank the anonymous reviewers for their valuable comments on a previous version of this paper. Parts of Yuyu Wang's work was supported by the National Natural Science Foundation for Young Scientists of China under Grant Number 62002049, the Natural Science Foundation of Sichuan under Grant Number 2023NSFSC0472, the Sichuan Science and Technology Program under Grant Number 2022YFG0037, and the National Key Research and Development Program of China under Grant Number 2022YFB3104600. Parts of Jiaxin Pan's work was supported by the Research Council of Norway under Project No. 324235. Parts of Yu Chen's work was supported by the National Key Research and Development Program of China under Grant Number 2021YFA1000600, the National Natural Science Foundation of China under Grant Number 62272269, Taishan Scholar Program of Shandong Province.

Appendices

A. The Proof of Theorem 2.13

Proof. We prove Theorem 2.13 by the following two propositions.

Proposition A.1. *For all $\mathbf{M}^\top \in \text{ZeroSamp}(\lambda)$ and $\mathbf{x} \in \text{SampNo}_\lambda(\mathbf{M})$, we have $\mathbf{x} \in \{0, 1\}^\lambda \setminus \text{Im}(\mathbf{M})$.*

Proof of Proposition A.1. According to Lemma 2.10, we have $\text{Im}(\mathbf{M}) = \{\mathbf{x} \mid \mathbf{w} \in \{0\} \times \{0, 1\}^{\lambda-1}, \mathbf{x} = \mathbf{M}\mathbf{w}\}$. Since $\mathbf{N}^\lambda \mathbf{w} = \mathbf{0}$ for any $\mathbf{w} \in \{0\} \times \{0, 1\}^{\lambda-1}$, we have $\text{Im}(\mathbf{M}) = \{\mathbf{x} \mid \mathbf{w} \in \{0\} \times \{0, 1\}^{\lambda-1}, \mathbf{x} = (\mathbf{M} + \mathbf{N}^\lambda)\mathbf{w}\}$.³ Moreover, $(\mathbf{M} + \mathbf{N}^\lambda)$ is of full rank according to Lemma 2.9. Hence, for any $\mathbf{w} \in \{1\} \times \{0, 1\}^{\lambda-1}$ and any $\mathbf{x} \in \text{Im}(\mathbf{M})$, we have $(\mathbf{M} + \mathbf{N}^\lambda)\mathbf{w} \neq \mathbf{x}$. Namely, for any $\mathbf{w} \in \{1\} \times \{0, 1\}^{\lambda-1}$, we have $(\mathbf{M} + \mathbf{N}^\lambda)\mathbf{w} \in \{0, 1\}^\lambda \setminus \text{Im}(\mathbf{M}^\top)$, completing the proof of Proposition A.1. \square

Proposition A.2. *For any $\mathcal{A} = \{a_\lambda\} \in \text{NC}^1$,*

$$|\Pr [a_\lambda(\mathbf{M}, \mathbf{x}_0) = 1] - \Pr [a_\lambda(\mathbf{M}, \mathbf{x}_1) = 1]| \leq \text{negl}(\lambda)$$

where $\mathbf{M}^\top \xleftarrow{\$} \text{ZeroSamp}(\lambda)$, $\mathbf{x}_0 \xleftarrow{\$} \text{SampYes}_\lambda(\mathbf{M})$, and $\mathbf{x}_1 \xleftarrow{\$} \text{SampNo}_\lambda(\mathbf{M})$.

³See Sect. 2 for the notion of \mathbf{N}^λ .

Proof of Proposition A.2. Let $\mathcal{A} = \{a_\lambda\}$ be any adversary in NC^1 . We give intermediate games in Fig. 19 to show that the advantage of \mathcal{A} in breaking Proposition A.2 is negligible.

Lemma A.3. $\Pr[\mathbf{G}_1^{a_\lambda} \Rightarrow 1] = \Pr[\mathbf{G}_0^{a_\lambda} \Rightarrow 1]$.

Proof. In \mathbf{G}_1 we sample $\mathbf{w} \xleftarrow{\$} \{0\} \times \{0, 1\}^{\lambda-1}$ instead of $\mathbf{w} \xleftarrow{\$} \{1\} \times \{0, 1\}^{\lambda-1}$. Then Lemma A.3 follows from the fact that the distributions of $\mathbf{x} = \mathbf{M}\mathbf{w}$ and $\mathbf{x}' = \mathbf{M}\mathbf{w}'$ are identical where $\mathbf{M}^\top \xleftarrow{\$} \text{ZeroSamp}(\lambda)$, $\mathbf{w} \xleftarrow{\$} \{1\} \times \{0, 1\}^{\lambda-1}$, and $\mathbf{w}' \xleftarrow{\$} \{0\} \times \{0, 1\}^{\lambda-1}$, according to Lemma 2.11. \square

Lemma A.4. $\Pr[\mathbf{G}_2^{a_\lambda} \Rightarrow 1] = \Pr[\mathbf{G}_1^{a_\lambda} \Rightarrow 1]$.

Proof. In \mathbf{G}_2 , we compute \mathbf{x} as $\mathbf{x} = (\mathbf{M} + \mathbf{N}^\lambda)\mathbf{w}$ instead. Then Lemma A.4 follows from the fact that for any $\mathbf{w} \in \{0\} \times \{0, 1\}^{\lambda-1}$, we have $\mathbf{N}^\lambda \mathbf{w} = \mathbf{0}$. \square

Lemma A.5. *There exists an adversary $\mathcal{B}_1 = \{b_\lambda^1\} \in \text{NC}^1$ such that b_λ^1 breaks Definition 2.5, which holds under $\text{NC}^1 \not\subseteq \oplus\text{L}/\text{poly}$ according to Lemma 2.6, with advantage*

$$|\Pr[\mathbf{G}_3^{a_\lambda} \Rightarrow 1] - \Pr[\mathbf{G}_2^{a_\lambda} \Rightarrow 1]|.$$

Proof. \mathbf{G}_2 and \mathbf{G}_3 only differ in the distribution of \mathbf{M} , namely, $\mathbf{M}^\top \xleftarrow{\$} \text{ZeroSamp}(\lambda)$ or $\mathbf{M}^\top \xleftarrow{\$} \text{OneSamp}(\lambda)$, and we build the distinguisher b_λ^1 as follows.

b_λ^1 runs in exactly the same way as the challenger of \mathbf{G}_2 except that in INIT, instead of generating \mathbf{M} by itself, it takes as input \mathbf{M}^\top generated as $\mathbf{M}^\top \xleftarrow{\$} \text{ZeroSamp}(\lambda)$ or $\mathbf{M}^\top \xleftarrow{\$} \text{OneSamp}(\lambda)$ from its own challenger. When a_λ outputs β , b_λ^1 outputs β as well. If \mathbf{M} is generated as $\mathbf{M}^\top \xleftarrow{\$} \text{ZeroSamp}(\lambda)$ (respectively, $\mathbf{M}^\top \xleftarrow{\$} \text{OneSamp}(\lambda)$), the view of a_λ is the same as its view in \mathbf{G}_2 (respectively, \mathbf{G}_3). Hence, the probability that b_λ^1 breaks the fine-grained matrix linear assumption is

$$|\Pr[\mathbf{G}_3^{a_\lambda} \Rightarrow 1] - \Pr[\mathbf{G}_2^{a_\lambda} \Rightarrow 1]|.$$

Moreover, since all operations in b_λ^1 are performed in NC^1 , we have $\mathcal{B}_1 = \{b_\lambda^1\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$, completing this part of proof. \square

Lemma A.6. $\Pr[\mathbf{G}_4^{a_\lambda} \Rightarrow 1] = \Pr[\mathbf{G}_3^{a_\lambda} \Rightarrow 1]$.

Proof. In \mathbf{G}_4 we sample $\mathbf{w} \xleftarrow{\$} \{1\} \times \{0, 1\}^{\lambda-1}$ instead of $\mathbf{w} \xleftarrow{\$} \{0\} \times \{0, 1\}^{\lambda-1}$. Let $\mathbf{M}^\top \xleftarrow{\$} \text{OneSamp}(\lambda)$. The distribution of $\mathbf{M} + \mathbf{N}^\lambda$ is identical to the output distribution of $\text{ZeroSamp}(\lambda)$ according to Lemma 2.9. Therefore, according to Lemma 2.11, the distributions of $\mathbf{x} = (\mathbf{M} + \mathbf{N}^\lambda)\mathbf{w}$ and $\mathbf{x}' = (\mathbf{M} + \mathbf{N}^\lambda)\mathbf{w}'$ are identical for $\mathbf{w} \xleftarrow{\$} \{1\} \times \{0, 1\}^{\lambda-1}$ and $\mathbf{w}' \xleftarrow{\$} \{0\} \times \{0, 1\}^{\lambda-1}$, completing this part of proof. \square

Lemma A.7. *There exists an adversary $\mathcal{B}_2 = \{b_\lambda^2\}_{\lambda \in \mathbb{N}}$ such that b_λ^2 breaks the fine-grained matrix linear assumption with advantage*

$$|\Pr[\mathbf{G}_5^{a_\lambda} \Rightarrow 1] - \Pr[\mathbf{G}_4^{a_\lambda} \Rightarrow 1]|.$$

Proof. \mathbf{G}_5 and \mathbf{G}_4 only differ in the distribution of \mathbf{M} , namely, $\mathbf{M}^\top \stackrel{\$}{\leftarrow} \text{OneSamp}(\lambda)$ or $\mathbf{M}^\top \stackrel{\$}{\leftarrow} \text{ZeroSamp}(\lambda)$, and we build the distinguisher b_λ^2 as follows. b_λ^2 runs in exactly the same way as the challenger of \mathbf{G}_2 except that in INIT, instead of generating \mathbf{M} by itself, it takes as input \mathbf{M}^\top generated as $\mathbf{M}^\top \stackrel{\$}{\leftarrow} \text{OneSamp}(\lambda)$ or $\mathbf{M}^\top \stackrel{\$}{\leftarrow} \text{ZeroSamp}(\lambda)$ from its own challenger. When a_λ outputs β , b_λ^2 outputs β as well. If \mathbf{M} is generated as $\mathbf{M}^\top \stackrel{\$}{\leftarrow} \text{OneSamp}(\lambda)$ (respectively, $\mathbf{M}^\top \stackrel{\$}{\leftarrow} \text{ZeroSamp}(\lambda)$), the view of a_λ is the same as its view in \mathbf{G}_4 (respectively, \mathbf{G}_5). Hence, the probability that b_λ^2 breaks the fine-grained matrix linear assumption is

$$|\Pr[\mathbf{G}_5^{a_\lambda} \Rightarrow 1] - \Pr[\mathbf{G}_4^{a_\lambda} \Rightarrow 1]|.$$

Moreover, since all operations in b_λ^2 are performed in NC^1 , we have $\mathcal{B}_2 = \{b_\lambda^2\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$, completing this part of proof. \square

Then Proposition A.2 follows from the fact that \mathbf{G}_0 and \mathbf{G}_5 are the real games of Proposition A.2, where the values \mathbf{x} are sampled from SampYes_λ and SampNo_λ respectively. \square

Putting all above together, Theorem 2.13 immediately follows. \square

B. Fine-Grained Secure Quasi-Adaptive NIZK

In this section, we construct fine-grained QA-NIZK with adaptive soundness. We first give the definition of NC^1 -QA-NIZK with adaptive soundness. Then we prove an NC^1 version of the Kernel Matrix Diffie-Hellman assumption [27], based on which we give a warm-up QA-NIZK in NC^1 with relatively low efficiency. Finally, we show how to achieve a more efficient construction.

B.1. Definitions

We now recall the definition of fine-grained QA-NIZK. Let \mathcal{D}_λ be a probability distribution over a collection of relations $\mathbf{R} = \{\mathbf{R}_\mathbf{M}\}_{\mathbf{M} \in \mathcal{D}_\lambda}$ parametrized by a matrix $\mathbf{M} \in \{0, 1\}^{n \times t}$ of rank $t' < n$ generated as $(\mathbf{M}^\top, \mathbf{M}^\perp) \stackrel{\$}{\leftarrow} \mathcal{D}_\lambda$ with the associated language

$$\mathcal{L}_\mathbf{M} = \{\mathbf{t} : \exists \mathbf{w} \in \{0, 1\}^t, \text{ s.t. } \mathbf{t} = \mathbf{M}\mathbf{w}\}.$$

Witness Sampleability. Notice that similar to witness sampleable distribution in the classical world [22], we require that \mathcal{D}_λ additionally outputs a non-zero matrix $\mathbf{M}^\perp \in \{0, 1\}^{n \times (n-t')}$ in the kernel of \mathbf{M}^\top . An example of sampleable distribution is $\text{ZeroSamp}(n)$, which can additionally sample a non-zero vector in the kernel of its output.⁴

⁴In fact, the rightmost vector $(r_1, \dots, r_{n-1}, 1)^\top$ of the intermediate matrix generated by $\text{RSamp}(n)$ (see Fig. 1) forms a vector in the kernel of \mathbf{M}^\top . See the proof of Lemma 3 in [15] for more details.

$\text{INIT}(\mathbf{M}):$ $(\text{crs}, \text{td}) \stackrel{\$}{\leftarrow} \text{Gen}_\lambda(\mathbf{M})$ Return crs.	$\text{FINALIZE}(\mathbf{y}^*, \pi^*):$ $\text{If } \mathbf{y}^* \notin \mathcal{L}_{\mathbf{M}} \text{ then}$ $\quad \text{return } \text{Ver}_\lambda(\text{crs}, \mathbf{y}^*, \pi^*)$ Else return 0
--	--

Fig. 20. The AS security game for QANIZK.

Definition B.1. (Quasi-adaptive non-interactive zero-knowledge proof) A \mathcal{C}_1 -quasi-adaptive non-interactive zero-knowledge proof (QA-NIZK) for a set of language distributions $\{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$ is a function family $\text{QANIZK} = \{\text{Gen}_\lambda, \text{Prove}_\lambda, \text{Ver}_\lambda, \text{Sim}_\lambda\}_{\lambda \in \mathbb{N}} \in \mathcal{C}_1$ with the following properties.

- $\text{Gen}_\lambda(\mathbf{M})$ returns a CRS crs and a simulation trapdoor td .
- $\text{Prove}_\lambda(\text{crs}, \mathbf{y}, \mathbf{w})$ returns a proof π .
- $\text{Ver}_\lambda(\text{crs}, \mathbf{y}, \pi)$ deterministically returns 1 (accept) or 0 (reject).
- $\text{Sim}_\lambda(\text{crs}, \text{td}, \mathbf{y})$ returns a simulated proof π .

Perfect completeness is satisfied if for all $(\mathbf{M}^\top, \mathbf{M}^\perp) \in \mathcal{D}_\lambda$, all vectors (\mathbf{y}, \mathbf{w}) such that $\mathbf{y} = \mathbf{M}\mathbf{w}$, all $(\text{crs}, \text{td}) \in \text{Gen}_\lambda(\mathbf{M})$, and all $\pi \in \text{Prove}_\lambda(\text{crs}, \mathbf{y}, \mathbf{w})$, we have

$$\text{Ver}_\lambda(\text{crs}, \mathbf{y}, \pi) = 1.$$

Perfect zero knowledge is satisfied if for all λ , all $(\mathbf{M}^\top, \mathbf{M}^\perp) \in \mathcal{D}_\lambda$, all (\mathbf{y}, \mathbf{w}) with $\mathbf{y} = \mathbf{M}\mathbf{w}$, and all $(\text{crs}, \text{td}) \in \text{Gen}_\lambda(\mathbf{M})$, the following two distributions are identical:

$$\text{Prove}_\lambda(\text{crs}, \mathbf{y}, \mathbf{w}) \text{ and } \text{Sim}_\lambda(\text{crs}, \text{td}, \mathbf{y}).$$

Definition B.2. (Adaptive soundness for QANIZK) QANIZK is said to satisfy \mathcal{C}_2 -adaptive soundness if for any adversary $\mathcal{A} = \{a_\lambda\}_{\lambda \in \mathbb{N}} \in \mathcal{C}_2$,

$$\Pr[\text{AS}^{a_\lambda} \Rightarrow 1] \leq \text{negl}(\lambda),$$

where Game AS is defined in Fig. 20.

We note that in the above definition, the term “quasi-adaptive” means that the construction of the CRS depends on the statement \mathbf{M} . On the other hand, “adaptive” in the context of adaptive soundness means that in the soundness experiment, the adversary can choose the statement adaptively after seeing the CRS.

B.2. A Warm-Up Construction

A New Lemma. We now prove the following lemma under the assumption $\text{NC}^1 \not\subseteq \oplus\text{L}/\text{poly}$, based on which we can achieve adaptively sound QA-NIZKs in NC^1 . It can be thought of as the counterpart of the Kernel Matrix Diffie-Hellman assumption [27] in NC^1 .

Definition B.3. (*Fine-grained kernel matrix assumption*) If $\text{NC}^1 \subsetneq \oplus\text{L}/\text{poly}$, then for all $\lambda \in \mathbb{N}$ and any adversary $\mathcal{A} = \{a_\lambda\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$, we have

$$\Pr[\mathbf{c}^\top \mathbf{M} = \mathbf{0} \wedge \mathbf{c} \neq \mathbf{0} | \mathbf{c} \stackrel{\$}{\leftarrow} a_\lambda(\mathbf{M})] \leq \text{negl}(\lambda),$$

where $\mathbf{M}^\top \stackrel{\$}{\leftarrow} \text{ZeroSamp}(\lambda)$.

Lemma B.4. *If $\text{NC}^1 \subsetneq \oplus\text{L}/\text{poly}$, then the fine-grained kernel matrix assumption (see Definition B.3) holds.*

Proof. Let $\mathcal{A} = \{a_\lambda\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$ be an adversary such that a_λ breaks the fine-grained kernel matrix assumption with probability ϵ , we construct another adversary $\mathcal{B} = \{b_\lambda\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$ such that b_λ breaks the fine-grained subset membership problem (see Definition 2.12), which holds under $\text{NC}^1 \subsetneq \oplus\text{L}/\text{poly}$ according to Theorem 2.13, with the same probability as follows.

On input (\mathbf{M}, \mathbf{u}) where $\mathbf{M}^\top \stackrel{\$}{\leftarrow} \text{ZeroSamp}(\lambda)$ and $\mathbf{u} \stackrel{\$}{\leftarrow} \text{SampYes}_\lambda(\mathbf{M})$ or $\mathbf{u} \stackrel{\$}{\leftarrow} \text{SampNo}_\lambda(\mathbf{M})$, b_λ forwards \mathbf{M} to a_λ . When a_λ outputs \mathbf{c} , b_λ outputs 1 iff the last element in \mathbf{c} is 1, $\mathbf{c}^\top \mathbf{M} = \mathbf{0}$, and $\mathbf{c}^\top \mathbf{u} = 0$.

When $\mathbf{u} \stackrel{\$}{\leftarrow} \text{SampYes}_\lambda(\mathbf{M})$, the probability that b_λ outputs 1 is ϵ . The reason is that when a_λ succeeds, we must have $\mathbf{c}^\top \mathbf{u} = 0$ when $\mathbf{c}^\top \mathbf{M} = \mathbf{0}$, and the last element of \mathbf{c} must be 1 according to Lemma 2.7. Moreover, when $\mathbf{u} \stackrel{\$}{\leftarrow} \text{SampNo}_\lambda(\mathbf{M})$, we have $\mathbf{u} = (\mathbf{M} + \mathbf{N}^\lambda) \mathbf{w}$ for some $\mathbf{w} \in \{1\} \times \{0, 1\}^{\lambda-1}$. If $\mathbf{c}^\top \mathbf{M} = \mathbf{0}$, we have $\mathbf{c}^\top \mathbf{u} = \mathbf{c}^\top \mathbf{N}^\lambda \mathbf{w} = \mathbf{c}^\top (0, \dots, 0, 1)^\top$, i.e., either $\mathbf{c}^\top \mathbf{u} = 1$ or the last element of \mathbf{c} is 0. Hence, b_λ outputs 0 anyway when $\mathbf{u} \stackrel{\$}{\leftarrow} \text{SampNo}_\lambda(\mathbf{M})$. Therefore, we have

$$\left| \Pr [b_\lambda(\mathbf{x}) = 1 \mid \mathbf{x} \stackrel{\$}{\leftarrow} \text{SampYes}_\lambda(\lambda)] - \Pr [b_\lambda(\mathbf{x}) = 1 \mid \mathbf{x} \stackrel{\$}{\leftarrow} \text{SampNo}_\lambda(\lambda)] \right| = \epsilon.$$

Moreover, since all operations in b_λ are performed in NC^1 , we have $\mathcal{B} = \{b_\lambda\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$, completing the proof of Lemma B.4. \square

Constructing QA-NIZK Based on Lemma B.4. Based on the above lemma, we can easily achieve NC^1 -QA-NIZKs with adaptive soundness, one-time simulation soundness, and unbounded simulation soundness against NC^1 by adopting the techniques in [23].⁵ Specifically, we only have to move the algorithms in [23] from $GF(p)$ for a large prime p to $GF(2)$, change the matrix Diffie-Hellman distributions to $\text{SampYes}_\lambda(\lambda)$, and generate a large number of proofs in parallel to bound the advantage of the adversary. We now give an adaptively sound QA-NIZK QANIZK_0 w.r.t. a set of (sampleable) distributions $\{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$ in Fig. 21 as an instance.⁶

Theorem B.5. *If $\text{NC}^1 \subsetneq \oplus\text{L}/\text{poly}$, then QANIZK_0 is an $\text{AC}^0[2]$ -QA-NIZK that is NC^1 -adaptively sound for all \mathbf{M} in the distributions $\{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$ (see Appendix B.1 for the definition of $\{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$).*

⁵One-time (respectively, unbounded) simulation soundness prevents the adversary from proving a false statement after seeing a single simulated proof for a statement (respectively, multiple simulated proofs for statements) of its choice. We refer the reader to [23] for the formal definitions.

⁶We do not exploit the sampleability of the distribution for this construction.

$\text{Gen}_\lambda(\mathbf{M} \in \{0, 1\}^{n \times t}):$ $\mathbf{A}^\top \xleftarrow{\$} \text{ZeroSamp}(\lambda)$ $\text{For } i = 1, \dots, \ell$ $\quad \mathbf{K}_i \xleftarrow{\$} \{0, 1\}^{n \times \lambda}$ $\quad \mathbf{P}_i = \mathbf{M}^\top \mathbf{K}_i \in \{0, 1\}^{t \times \lambda}$ $\quad \mathbf{C}_i = \mathbf{K}_i \mathbf{A} \in \{0, 1\}^{n \times \lambda}$ $\text{Return } \text{crs} = (\mathbf{A}, (\mathbf{P}_i, \mathbf{C}_i)_{i=1}^\ell) \text{ and } \text{td} = (\mathbf{K}_i)_{i=1}^\ell$	$\text{Prove}_\lambda(\text{crs}, \mathbf{y} \in \{0, 1\}^n, \mathbf{x} \in \{0, 1\}^t):$ $\text{For } i = 1, \dots, \ell$ $\quad \pi_i = \mathbf{x}^\top \mathbf{P}_i \in \{0, 1\}^{1 \times \lambda}$ $\text{Return } \pi = (\pi_i)_{i=1}^\ell$ $\text{Sim}_\lambda(\text{crs}, \text{td}, \mathbf{y}):$ $\text{For } i = 1, \dots, \ell$ $\quad \pi_i = \mathbf{y}^\top \mathbf{K}_i \in \{0, 1\}^{1 \times \lambda}$ $\text{Return } \pi = (\pi_i)_{i=1}^\ell$ $\text{Ver}_\lambda(\text{crs}, \mathbf{y}, \pi):$ $\text{If } \pi_i \mathbf{A} = \mathbf{y}^\top \mathbf{C}_i \text{ for } i = 1, \dots, \ell$ $\quad \text{return } 1$ $\text{Else return } 0$
--	---

Fig. 21. Definition of $\text{QANIZK}_0 = \{\text{Gen}_\lambda, \text{Prove}_\lambda, \text{Ver}_\lambda, \text{Sim}_\lambda\}_{\lambda \in \mathbb{N}}$. We require that $2^{\ell(\cdot)}$ is some super-polynomial in λ .

Proof. First, we note that $\{\text{Gen}_\lambda\}_{\lambda \in \mathbb{N}}$, $\{\text{Prove}_\lambda\}_{\lambda \in \mathbb{N}}$, $\{\text{Sim}_\lambda\}_{\lambda \in \mathbb{N}}$, and $\{\text{Ver}_\lambda\}_{\lambda \in \mathbb{N}}$ are computable in $\text{AC}^0[2]$, since they only involve operations including multiplication of a constant number of matrices and sampling random bits. Perfect correctness and perfect zero-knowledge follow from the fact that for all $\mathbf{y} = \mathbf{M}\mathbf{x}$ and $\mathbf{P}_i = \mathbf{M}^\top \mathbf{K}_i$, we have

$$\mathbf{x}^\top \mathbf{P}_i = \mathbf{x}^\top \mathbf{M}^\top \mathbf{K}_i = \mathbf{y}^\top \mathbf{K}_i.$$

Let $\mathcal{A} = \{a_\lambda\}_{\lambda \in \mathbb{N}}$ be an adversary breaking the adaptive soundness of QANIZK_0 with advantage ϵ , we have the following lemma. \square

Lemma B.6. *There exists an adversary $\mathcal{B} = \{b_\lambda\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$ such that b_λ breaks the fine-grained kernel matrix assumption (see Definition B.3), which holds under $\text{NC}^1 \subsetneq \oplus \mathbb{L}/\text{poly}$ according to Lemma B.4, with probability $\epsilon - 1/2^\ell$.*

Proof. We construct b_λ as follows.

b_λ on input \mathbf{A} samples $(\mathbf{M}^\top, \mathbf{M}^\perp) \xleftarrow{\$} \mathcal{D}_\lambda$ and $\mathbf{K}_i \xleftarrow{\$} \{0, 1\}^{n \times \lambda}$, and sets $\mathbf{P}_i = \mathbf{M}^\top \mathbf{K}_i$ and $\mathbf{C}_i = \mathbf{K}_i \mathbf{A}$ for all $i \in [\ell]$. Then it sends $\text{crs} = (\mathbf{A}, (\mathbf{P}_i, \mathbf{C}_i)_{i=1}^\ell)$ to a_λ . When a_λ outputs $(\pi = (\pi_i)_{i=1}^\ell, \mathbf{y})$, b_λ searches j such that

$$\pi_j \mathbf{A} = \mathbf{y}^\top \mathbf{C}_j = \mathbf{y}^\top \mathbf{K}_j \mathbf{A}$$

and

$$\pi_j - \mathbf{y}^\top \mathbf{K}_j \neq \mathbf{0}.$$

If the searching procedure fails, b_λ aborts. b_λ then outputs $\pi_j - \mathbf{y}^\top \mathbf{K}_j$.

When a_λ succeeds, we have $\pi_j \mathbf{A} = \mathbf{y}^\top \mathbf{C}_j$ for all j and $\mathbf{y} \notin \text{Im}(\mathbf{M})$. Let $\hat{\mathbf{a}}$ be a fixed non-zero vector such that $\hat{\mathbf{a}} \notin \text{Im}(\mathbf{A})$. For each i , since a_λ learns no information on \mathbf{K}_i other than $\mathbf{M}^\top \mathbf{K}_i$ and $\mathbf{K}_i \mathbf{A}$, $\mathbf{y}^\top \mathbf{K}_i \hat{\mathbf{a}}$ is information-theoretically hidden in the view of a_λ , i.e., the probability that there exists j such that $\pi_j \hat{\mathbf{a}} - \mathbf{y}^\top \mathbf{K}_j \hat{\mathbf{a}} \neq 0$ is at least $1/2^\ell$.

$\text{Gen}_\lambda(\mathbf{M} \in \{0, 1\}^{n \times t}):$ $\mathbf{A}^\top \stackrel{\$}{\leftarrow} \text{ZeroSamp}(\lambda)$ For $i = 1, \dots, n - t'$ $\mathbf{K}_i \stackrel{\$}{\leftarrow} \{0, 1\}^{n \times (\lambda-1)}$ $\mathbf{P}_i = \mathbf{M}^\top \mathbf{K}_i \in \{0, 1\}^{t \times (\lambda-1)}$ $\mathbf{C}_i = (\mathbf{K}_i \ \mathbf{0}) \mathbf{A} \in \{0, 1\}^{n \times (\lambda-1)}$ Return $\text{crs} = (\mathbf{A}, (\mathbf{P}_i, \mathbf{C}_i)_{i=1}^{n-t'})$ and $\text{td} = (\mathbf{K}_i)_{i=1}^{n-t'}$	$\text{Prove}_\lambda(\text{crs}, \mathbf{y} \in \{0, 1\}^n, \mathbf{x} \in \{0, 1\}^{t'}):$ For $i = 1, \dots, n - t'$ $\pi_i = \mathbf{x}^\top \mathbf{P}_i \in \{0, 1\}^{1 \times (\lambda-1)}$ Return $\pi = (\pi_i)_{i=1}^{n-t'}$ $\text{Sim}_\lambda(\text{crs}, \text{td}, \mathbf{y}):$ For $i = 1, \dots, n - t'$ $\pi_i = \mathbf{y}^\top \mathbf{K}_i \in \{0, 1\}^{1 \times (\lambda-1)}$ Return $\pi = (\pi_i)_{i=1}^{n-t'}$ $\text{Ver}_\lambda(\text{crs}, \mathbf{y}, \pi):$ If $(\pi_i \ \mathbf{0}) \mathbf{A} = \mathbf{y}^\top \mathbf{C}_i$ for $i = 1, \dots, n - t'$, return 1 Else return 0
---	---

Fig. 22. Definition of $\text{QANIZK}_1 = \{\text{Gen}_\lambda, \text{Prove}_\lambda, \text{Ver}_\lambda, \text{Sim}_\lambda\}_{\lambda \in \mathbb{N}}$.

Since $\pi_j \hat{\mathbf{a}} - \mathbf{y}^\top \mathbf{K}_j \hat{\mathbf{a}} \neq 0$ implies $\pi_j - \mathbf{y}^\top \mathbf{K}_j \neq \mathbf{0}$, the probability that b_λ breaks the fine-grained kernel matrix assumption is at least $\epsilon - 1/2^\ell$, completing this part of proof. \square

Since the fine-grained kernel matrix assumption holds if $\text{NC}^1 \not\subseteq \oplus \text{L}/\text{poly}$ according to Lemma B.4, putting all above together, Theorem B.5 immediately follows. \square

B.3. A More Efficient Construction

A disadvantage of the scheme in Appendix B.2 is that we have to generate a large number of proofs in parallel. In this section, we give a more efficient NC^1 -adaptively sound QA-NIZK $\text{QANIZK}_1 = \{\text{Gen}_\lambda, \text{Prove}_\lambda, \text{Ver}_\lambda, \text{Sim}_\lambda\}_{\lambda \in \mathbb{N}}$ w.r.t. a set of distributions $\{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$ in Fig. 22. As in Definition B.1, we require that \mathcal{D}_λ be witness sampleable, i.e., it outputs a matrix $\mathbf{M} \in \{0, 1\}^{n \times t}$ of rank $t' < n$ additionally with a matrix (or vector) $\mathbf{M}^\perp \in \{0, 1\}^{n \times (n-t')}$ with rank $n - t'$ in its kernel.

The proof size of this construction is $(\lambda - 1) \cdot (n - t')$. Since \mathbf{M} (or \mathbf{M}^\top) is usually a combination of matrices sampled from $\text{ZeroSamp}(\lambda)$ in NC^1 , $n - t'$ is typically a constant number. For instance, when proving that two ciphertexts of the PKE scheme in [13] correspond to the same message or proving the validity of a public key of the PKE scheme in [15], the proof size is only $\lambda - 1$ in contrast to $\lambda \cdot \ell$ for a large number ℓ in the warm-up construction.

Theorem B.7. *If $\text{NC}^1 \not\subseteq \oplus \text{L}/\text{poly}$, then QANIZK_1 is an $\text{AC}^0[2]$ -QA-NIZK that is NC^1 -adaptively sound for all \mathbf{M} in the distributions $\{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$ (see Appendix B.1 for the definition of $\{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$).*

Proof. First, we note that $\{\text{Gen}_\lambda\}_{\lambda \in \mathbb{N}}$, $\{\text{Prove}_\lambda\}_{\lambda \in \mathbb{N}}$, $\{\text{Sim}_\lambda\}_{\lambda \in \mathbb{N}}$, and $\{\text{Ver}_\lambda\}_{\lambda \in \mathbb{N}}$ are computable in $\text{AC}^0[2]$, since they only involve operations including multiplications of a constant number of matrices and sampling random bits.

Perfect correctness and perfect zero-knowledge follow from the fact that for all $\mathbf{y} = \mathbf{M}\mathbf{x}$ and $\mathbf{P}_i = \mathbf{M}^\top \mathbf{K}_i$, we have

$$\mathbf{x}^\top \mathbf{P}_i = \mathbf{x}^\top \mathbf{M}^\top \mathbf{K}_i = \mathbf{y}^\top \mathbf{K}_i.$$

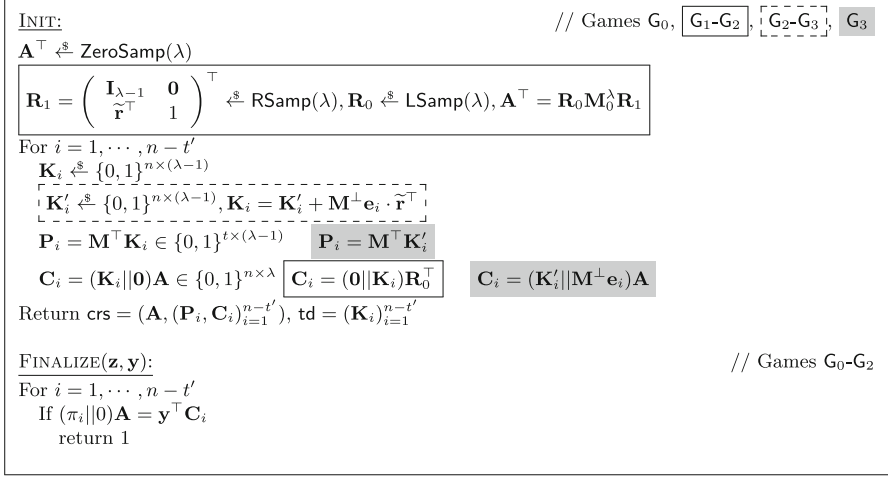


Fig. 23. Games $\mathbf{G}_0, \mathbf{G}_1, \mathbf{G}_2$ for the proof of Theorem B.7. $\mathbf{e}_i \in \{0, 1\}^{n-t'}$ denotes the vector with the i th element being 1 and the others being 0.

Let $\mathcal{A} = \{a_\lambda\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$ be any adversary against the NC^1 -adaptive soundness of QANIZK_1 . We now show that QANIZK_1 is adaptively sound against NC^1 via a sequence of hybrid games as in Fig. 23. The crucial step is to use the technique exploited by our IBKEM to switch $(\mathbf{K}_i \| \mathbf{0}) \mathbf{A}$ to $(\mathbf{0} \| \mathbf{K}_i) \mathbf{R}_0^\top$, and then switch it back to $(\mathbf{K}'_i \| \mathbf{M}^\perp \mathbf{e}_i) \mathbf{A}$ for $\mathbf{K}_i = \mathbf{K}'_i + \mathbf{M}^\perp \mathbf{e}_i \cdot \tilde{\mathbf{r}}^\top$, where \mathbf{R}_0 and $\tilde{\mathbf{r}}$ are intermediate values generated during the sampling procedure for \mathbf{A} . \square

Lemma B.8. $\Pr[\text{AS}^{a_\lambda} \Rightarrow 1] = \Pr[\mathbf{G}_1^{a_\lambda} \Rightarrow 1] = \Pr[\mathbf{G}_0^{a_\lambda} \Rightarrow 1]$.

Proof. In \mathbf{G}_1 , we generate \mathbf{A} by sampling $\mathbf{R}_1 = \begin{pmatrix} \mathbf{I}_{\lambda-1} & \mathbf{0} \\ \tilde{\mathbf{r}}^\top & 1 \end{pmatrix}^\top \xleftarrow{\$} \text{RSamp}(\lambda)$ and $\mathbf{R}_0 \xleftarrow{\$} \text{LSamp}(\lambda)$, and setting $\mathbf{A}^\top = \mathbf{R}_0 \mathbf{M}_0^\lambda \mathbf{R}_1$. Moreover, for all i , we replace $\mathbf{C}_i = (\mathbf{K}_i \| \mathbf{0}) \mathbf{A}$ by $\mathbf{C}_i = (\mathbf{0} \| \mathbf{K}_i) \mathbf{R}_0^\top$. The view of \mathcal{A} in this game is identical to its view in \mathbf{G}_0 since the way we generate \mathbf{A} is exactly the “zero-sampling” procedure, and we have

$$\begin{aligned} \mathbf{C}_i &= (\mathbf{K}_i \| \mathbf{0}) \mathbf{A} = (\mathbf{K}_i \| \mathbf{0}) \mathbf{R}_1^\top \mathbf{M}_0^\lambda \mathbf{R}_0^\top \\ &= (\mathbf{K}_i \| \mathbf{0}) \begin{pmatrix} \mathbf{I}_{\lambda-1} & \mathbf{0} \\ \tilde{\mathbf{r}}^\top & 1 \end{pmatrix} \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \ddots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & & 1 \\ 0 & & \cdots & & 0 \end{pmatrix} \mathbf{R}_0^\top \end{aligned}$$

$$\begin{aligned}
 &= (\mathbf{K}_i \| \mathbf{0}) \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \ddots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & & 1 \\ 0 & & \cdots & & 0 \end{pmatrix} \mathbf{R}_0^\top \\
 &= (\mathbf{0} \| \mathbf{K}_i) \mathbf{R}_0^\top.
 \end{aligned}$$

□

Lemma B.9. $\Pr[\mathbf{G}_2^{a_\lambda} \Rightarrow 1] = \Pr[\mathbf{G}_1^{a_\lambda} \Rightarrow 1]$.

Proof. In \mathbf{G}_2 , for all i , instead of generating \mathbf{K}_i as a uniformly random matrix, we generate \mathbf{K}_i by randomly sampling $\mathbf{K}'_i \xleftarrow{\$} \{0, 1\}^{n \times (\lambda-1)}$ and setting $\mathbf{K}_i = \mathbf{K}'_i + \mathbf{M}^\perp \mathbf{e}_i \cdot \tilde{\mathbf{r}}^\top$, where $\mathbf{e}_i \in \{0, 1\}^{n-i'}$ denotes the vector with the i th element being 1 and the other bits being 0. Since the distribution of \mathbf{K}_i is still uniform, the view of \mathcal{A} remains the same. □

Lemma B.10. $\Pr[\mathbf{G}_3^{a_\lambda} \Rightarrow 1] = \Pr[\mathbf{G}_2^{a_\lambda} \Rightarrow 1]$.

Proof. This lemma follows from the fact that for all i , we have

$$\mathbf{M}\mathbf{K}_i = \mathbf{M}(\mathbf{K}'_i + \mathbf{M}^\perp \mathbf{e}_i \cdot \tilde{\mathbf{r}}^\top) = \mathbf{M}\mathbf{K}'_i$$

and

$$\begin{aligned}
 \mathbf{C}_i &= (\mathbf{0} \| \mathbf{K}'_i + \mathbf{M}^\perp \mathbf{e}_i \cdot \tilde{\mathbf{r}}^\top) \mathbf{R}_0^\top \\
 &= (\mathbf{K}'_i + \mathbf{M}^\perp \mathbf{e}_i \cdot \tilde{\mathbf{r}}^\top \| \mathbf{M}^\perp \mathbf{e}_i) \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \ddots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & & 1 \\ 0 & & \cdots & & 0 \end{pmatrix} \mathbf{R}_0^\top \\
 &= (\mathbf{K}'_i \| \mathbf{M}^\perp \mathbf{e}_i) \begin{pmatrix} \mathbf{I}_{\lambda-1} & \mathbf{0} \\ \tilde{\mathbf{r}}^\top & 1 \end{pmatrix} \mathbf{M}_0^{\lambda \top} \mathbf{R}_0^\top \\
 &= (\mathbf{K}'_i \| \mathbf{M}^\perp \mathbf{e}_i) \mathbf{A}
 \end{aligned}$$

□

Lemma B.11. *There exists an adversary $\mathcal{B} = \{b_\lambda\}_{\lambda \in \mathbb{N}} \in \text{NC}^1$ such that b_λ breaks the fine-grained kernel matrix assumption (see Definition B.3), which holds under $\text{NC}^1 \subsetneq \oplus \text{L/poly}$ according to Lemma B.4, with probability $\Pr[\mathbf{G}_3^{a_\lambda} \Rightarrow 1]$.*

Proof. We construct b_λ as follows.

b_λ on input \mathbf{A} samples $(\mathbf{M}^\top, \mathbf{M}^\perp) \xleftarrow{\$} \mathcal{D}_\lambda$ and $\mathbf{K}'_i \xleftarrow{\$} \{0, 1\}^{n \times (\lambda-1)}$, and sets $\mathbf{P}_i = \mathbf{M}^\top(\mathbf{K}'_i || \mathbf{0})$ and $\mathbf{C}_i = (\mathbf{K}'_i || \mathbf{M}^\perp \mathbf{e}_i) \mathbf{A}$ for all i . Then it sends $\text{crs} = (\mathbf{A}, (\mathbf{P}_i, \mathbf{C}_i)_{i=1}^{n-t'})$ to a_λ . When a_λ outputs $(\pi = (\pi_i)_{i=1}^{n-t'}, \mathbf{y})$, b_λ searches j such that

$$(\pi_j || 0) \mathbf{A} = \mathbf{y}^\top \mathbf{C}_j = \mathbf{y}^\top (\mathbf{K}'_j || \mathbf{M}^\perp \mathbf{e}_j) \mathbf{A}$$

and

$$\pi_j || 0 - \mathbf{y}^\top (\mathbf{K}'_j || \mathbf{M}^\perp \mathbf{e}_j) \neq \mathbf{0}.$$

If the searching procedure fails, b_λ aborts. b then outputs $\pi_j || 0 - \mathbf{y}^\top (\mathbf{K}'_j || \mathbf{M}^\perp \mathbf{e}_j)$.

Since all the operations performed by b_λ are in NC^1 , we have $\mathcal{B} \in \text{NC}^1$.

When a_λ succeeds, we have $(\pi_j || 0) \mathbf{A} = \mathbf{y}^\top \mathbf{C}_j$ for all j and $\mathbf{y} \notin \text{Span}(\mathbf{M})$. In this case, $\mathbf{y}^\top \mathbf{M}^\perp \neq \mathbf{0}$, i.e.,

there must exist j such that $\mathbf{y}^\top \mathbf{M}^\perp \mathbf{e}_i = 1$. Hence the probability that b_λ breaks the fine-grained kernel matrix assumption is exactly $\Pr[\mathcal{G}_3^{a_\lambda} \Rightarrow 1]$, completing this part of proof. \square

Putting all above together, Theorem B.7 immediately follows. \square

Concurrent Fine-Grained NIZKs. Assuming $\text{NC}^1 \subsetneq \oplus\text{L}/\text{poly}$, our work presents an efficient QA-NIZK that achieves perfect zero-knowledge and can handle languages expressible as linear subspaces. Below we compare our QA-NIZK to other existing fine-grained NIZKs [2, 32, 33].

Ball, Dachman-Soled, and Kulkarni [2] previously constructed a NIZK for circuit satisfiability against NC^1 adversaries in the uniform random string (URS) model, where the setup only samples public coins. Their scheme achieves offline zero-knowledge, meaning that the distribution of honest URSs and proofs is computationally indistinguishable from that of the output of a simulator drawn from a specific distribution. However, their construction is not in the fully fine-grained setting, since their prover requires more computational resources than NC^1 (even for statements represented as NC^1 circuits). This requirement is inherent in their construction, since the underlying NIZK for $\oplus\text{L}/\text{poly}$ in their construction requires computing the determinant of a matrix, which cannot be done in NC^1 .

$\mathcal{X} = \{0, 1\}^n, \mathcal{Y} = \{0, 1\}^n$ $\ell = (1 + n), \eta = 1, \zeta = n$ $\text{p}_\lambda(\mathbf{x}, \mathbf{y})$: Return 1 iff $\mathbf{x}^\top \mathbf{y} = 0$	$\text{sE}_\lambda(\mathbf{x}, \begin{pmatrix} u \\ \mathbf{w} \end{pmatrix}) = \mathbf{x} \cdot u + \mathbf{w}$ $\text{rE}_\lambda(\mathbf{y}, \begin{pmatrix} u \\ \mathbf{w} \end{pmatrix}) = \mathbf{y}^\top \mathbf{w}$ $\text{kE}_\lambda(\mathbf{y}, \alpha) = \alpha$ $\text{sD}_\lambda(\mathbf{x}, \mathbf{y}, \mathbf{c}) = \mathbf{c}^\top \mathbf{y}$ $\text{rD}_\lambda(\mathbf{x}, \mathbf{y}, d) = d$
--	---

Fig. 24. Definitions of the predicate and encoding of an ABE scheme for inner product (with short secret keys).

$\mathcal{X} = \{0, 1\}^n, \mathcal{Y} = \{0, 1\}^n$ $\ell = n, \eta = n, \zeta = 1$ $\text{p}\lambda(\mathbf{x}, \mathbf{y})$: Return 1 iff $\mathbf{x}^\top \mathbf{y} = 1$	$\text{sE}_\lambda(\mathbf{x}, \mathbf{w}) = \mathbf{x}^\top \mathbf{w}$ $\text{rE}_\lambda(\mathbf{y}, \mathbf{w}) = \mathbf{w}$ $\text{kE}_\lambda(\mathbf{y}, \alpha) = \mathbf{y}\alpha$ $\text{sD}_\lambda(\mathbf{x}, \mathbf{y}, \mathbf{c}) = \mathbf{c} \cdot (\mathbf{x}^\top \mathbf{y})$ $\text{rD}_\lambda(\mathbf{x}, \mathbf{y}, \mathbf{d}) = (\mathbf{x}^\top \mathbf{y}) \mathbf{x}^\top \mathbf{d}$
---	--

Fig. 25. Definitions of the predicate and encoding of an ABE scheme for non-zero inner product (with short secret ciphertexts).

$\mathcal{X} = \{0, 1\}^n$ $\mathcal{Y} = \{0, 1\}^{n \times n'}$ $\ell = n + (n' - 1)$ $\eta = n$ $\zeta = n$ $\text{p}\lambda((\mathbf{x}, \omega), \mathbf{M})$: Return 1 iff \mathbf{x} satisfies \mathbf{M}	$\text{sE}_\lambda((\mathbf{x}, \omega), \begin{pmatrix} \mathbf{w} \\ \mathbf{u} \end{pmatrix}) = (x_1 w_1, \dots, x_n w_n)$ $\text{rE}_\lambda(\mathbf{M}, \begin{pmatrix} \mathbf{w} \\ \mathbf{u} \end{pmatrix}) = (\mathbf{M}_1 \begin{pmatrix} 0 \\ \mathbf{u} \end{pmatrix} + w_1, \dots, \mathbf{M}_n \begin{pmatrix} 0 \\ \mathbf{u} \end{pmatrix} + w_n)$ $\text{kE}_\lambda(\mathbf{y}, \alpha) = (\mathbf{M}_1 \begin{pmatrix} \alpha \\ \mathbf{0} \end{pmatrix}, \dots, \mathbf{M}_n \begin{pmatrix} \alpha \\ \mathbf{0} \end{pmatrix})$ $\text{sD}_\lambda((\mathbf{x}, \omega), \mathbf{M}, \mathbf{c}) = \sum_{i=1}^n \omega_i c_i$ $\text{rD}_\lambda((\mathbf{x}, \omega), \mathbf{M}, \mathbf{d}) = \sum_{i=1}^n x_i \omega_i d_i$
---	---

Fig. 26. Definitions of the predicate and encoding of an ABE scheme for boolean span programs. \mathbf{x} satisfies \mathbf{M} w.r.t. some ω iff $\sum_{i: x_i=1} \omega_i \mathbf{M}_i = (1, 0, \dots, 0)^\top$, where \mathbf{M}_i denotes the i th row of \mathbf{M} . Notice that in the original definition in [10], ω is not part of the attribute and is computed from \mathbf{M} and \mathbf{x} in decryption. We put ω into the attribute since computing it from \mathbf{M} and \mathbf{x} involves Gaussian elimination which cannot be executed in NC^1 . This does not affect the security of the resulting ABE since ω can be efficiently computed publicly in the original encoding scheme anyway.

More recently, Wang and Pan [32] proposed a fully fine-grained NIZK protocol for circuit satisfiability in NC^1 , where all algorithms (including the CRS generator, prover, verifier, and simulator) are in NC^1 . Their scheme can achieve either perfect soundness or perfect zero-knowledge and can be converted into a NIZK in the URS model and a non-interactive zap. Notably, their underlying NIZK for linear languages supports the same class of languages as our QA-NIZK. However, their construction has larger proving/verification cost and proof size than ours. Especially, their proof size is dependent of the statement size, while ours is not.

Another fine-grained NIZK is recently proposed by Wang and Pan [33] in a different fine-grained setting under no assumption. Specifically, it treats adversaries in AC^0 and requires that all algorithms run in AC^0 .

C. Instantiations of Encodings

In this section, other than the one in Fig. 3, we give several examples of predicate encodings in Figs. 24, 25, and 26. By instantiating our resulting ABE in Sect. 5 with these encodings, we immediately achieve ABEs for inner product, non-zero inner product, and boolean span programs. All the encodings can be performed in $\text{AC}^0[2]$ since they only involve multiplication of a constant number of matrices.

References

- [1] B. Applebaum, Y. Ishai, E. Kushilevitz, Cryptography in NC^0 , in *45th FOCS*. (IEEE Computer Society Press, 2004), pp. 166–175
- [2] M. Ball, D. Dachman-Soled, M. Kulkarni, New techniques for zero-knowledge: Leveraging inefficient provers to reduce assumptions, interaction, and trust, in Micciancio, D., Ristenpart, T. (eds.) *CRYPTO 2020, Part III*. LNCS, vol. 12172 (Springer, Heidelberg, 2020), pp. 674–703
- [3] D.A.M. Barrington, Bounded-width polynomial-size branching programs recognize exactly those languages in NC^1 , in *18th ACM STOC* (ACM Press, 1986), pp. 1–5
- [4] M. Bellare, S. Goldwasser, New paradigms for digital signatures and message authentication based on non-interactive zero knowledge proofs, in Brassard, G. (ed.) *CRYPTO'89*. LNCS, vol. 435 (Springer, Heidelberg, 1990), pp. 194–211
- [5] O. Blazy, E. Kiltz, J. Pan, (Hierarchical) identity-based encryption from affine message authentication, in Garay, J.A., Gennaro, R. (eds.) *CRYPTO 2014, Part I*. LNCS, vol. 8616 (Springer, Heidelberg, 2014), pp. 408–425
- [6] D. Boneh, M.K. Franklin, Identity-based encryption from the Weil pairing, in Kilian, J. (ed.) *CRYPTO 2001*. LNCS, vol. 2139 (Springer, Heidelberg, 2001), pp. 213–229
- [7] D. Boneh, P.A. Papakonstantinou, C. Rackoff, Y. Vahlis, B. Waters, On the impossibility of basing identity based encryption on trapdoor permutations, in *49th FOCS* (IEEE Computer Society Press, 2008), pp. 283–292
- [8] C. Brzuska, G. Couteau, Towards fine-grained one-way functions from strong average-case hardness. *IACR Cryptol. ePrint Arch.* **2020**, 1326 (2020)
- [9] M. Campanelli, R. Gennaro, Fine-grained secure computation, in Beimel, A., Dziembowski, S. (eds.) *TCC 2018, Part II*. LNCS, vol. 11240 (Springer, Heidelberg, 2018), pp. 66–97
- [10] J. Chen, R. Gay, H. Wee, Improved dual system ABE in prime-order groups via predicate encodings, in Oswald, E., Fischlin, M. (eds.) *EUROCRYPT 2015, Part II*. LNCS, vol. 9057 (Springer, Heidelberg, 2015), pp. 595–624
- [11] J. Chen, H. Wee, Fully, (almost) tightly secure IBE and dual system groups, in Canetti, R., Garay, J.A. (eds.) *CRYPTO 2013, Part II*. LNCS, vol. 8043 (Springer, Heidelberg, 2013), pp. 435–460
- [12] C. Cocks, An identity based encryption scheme based on quadratic residues, in Honary, B. (ed.) *8th IMA International Conference on Cryptography and Coding*. LNCS, vol. 2260 (Springer, Heidelberg, 2001), pp. 360–363
- [13] A. Degwekar, V. Vaikuntanathan, P.N. Vasudevan, Fine-grained cryptography, in Robshaw, M., Katz, J. (eds.) *CRYPTO 2016, Part III*. LNCS, vol. 9816 (Springer, Heidelberg, 2016), pp. 533–562
- [14] S. Egashira, Y. Wang, K. Tanaka, Fine-grained cryptography revisited, in Galbraith, S.D., Moriai, S. (eds.) *ASIACRYPT 2019, Part III*. LNCS, vol. 11923 (Springer, Heidelberg, 2019), pp. 637–666
- [15] S. Egashira, Y. Wang, K. Tanaka, Fine-grained cryptography revisited. *J. Cryptol.* **34**(3), 23 (2021)
- [16] G. Fuchsbauer, E. Kiltz, J. Loss, The algebraic group model and its applications, in Shacham, H., Boldyreva, A. (eds.) *CRYPTO 2018, Part II*. LNCS, vol. 10992 (Springer, Heidelberg, 2018), pp. 33–62
- [17] C. Gentry, A. Silverberg, Hierarchical ID-based cryptography, in Zheng, Y. (ed.) *ASIACRYPT 2002*. LNCS, vol. 2501 (Springer, Heidelberg, 2002), pp. 548–566
- [18] V. Goyal, O. Pandey, A. Sahai, B. Waters, Attribute-based encryption for fine-grained access control of encrypted data, in: Juels, A., Wright, R.N., De Capitani di Vimercati, S. (eds.) *ACM CCS 2006* (ACM Press, 2006), Available as Cryptology ePrint Archive Report 2006/309, pp. 89–98
- [19] J. Groth, A. Sahai, Efficient non-interactive proof systems for bilinear groups, in Smart, N.P. (ed.) *EUROCRYPT 2008*. LNCS, vol. 4965 (Springer, Heidelberg, 2008), pp. 415–432
- [20] J. Horwitz, B. Lynn, Toward hierarchical identity-based encryption, in Knudsen, L.R. (ed.) *EUROCRYPT 2002*. LNCS, vol. 2332 (Springer, Heidelberg, 2002), pp. 466–481
- [21] Y. Ishai, E. Kushilevitz, Randomizing polynomials: A new representation with applications to round-efficient secure computation, in *41st FOCS* (IEEE Computer Society Press, 2000), pp. 294–304
- [22] C.S. Jutla, A. Roy, Shorter quasi-adaptive NIZK proofs for linear subspaces, in Sako, K., Sarkar, P. (eds.) *ASIACRYPT 2013, Part I*. LNCS, vol. 8269 (Springer, Heidelberg, 2013), pp. 1–20
- [23] E. Kiltz, H. Wee, Quasi-adaptive NIZK for linear subspaces revisited, in Oswald, E., Fischlin, M. (eds.) *EUROCRYPT 2015, Part II*. LNCS, vol. 9057 (Springer, Heidelberg, 2015), pp. 101–128

- [24] A.B. Lewko, B. Waters, New techniques for dual system encryption and fully secure HIBE with short ciphertexts, in: Micciancio, D. (ed.) *TCC 2010*. LNCS, vol. 5978 (Springer, Heidelberg, 2010), pp. 455–479
- [25] U.M. Maurer, Abstract models of computation in cryptography (invited paper), in: Smart, N.P. (ed.) *10th IMA International Conference on Cryptography and Coding*. LNCS, vol. 3796 (Springer, Heidelberg, 2005), pp. 1–12
- [26] R.C. Merkle, Secure communications over insecure channels. *Commun. ACM* **21**(4), 294–299 (1978)
- [27] P. Morillo, C. Ràfols, J.L. Villar, The kernel matrix Diffie–Hellman assumption, in Cheon, J.H., Takagi, T. (eds.) *ASIACRYPT 2016, Part I*. LNCS, vol. 10031 (Springer, Heidelberg, 2016), pp. 729–758
- [28] A.A. Razborov, Lower bounds on the size of bounded depth circuits over a complete basis with logical addition, in *Mathematical notes of the Academy of Sciences of the USSR* 41(4) (1987)
- [29] A. Shamir, Identity-based cryptosystems and signature schemes, in Blakley, G.R., Chaum, D. (eds.) *CRYPTO '84*. LNCS, vol. 196 (Springer, Heidelberg, 1984), pp. 47–53
- [30] V. Shoup, Lower bounds for discrete logarithms and related problems, in Fumy, W. (ed.) *EUROCRYPT '97*. LNCS, vol. 1233 (Springer, Heidelberg, 1997), pp. 256–266
- [31] R. Smolensky, Algebraic methods in the theory of lower bounds for Boolean circuit complexity, in: Aho, A. (ed.) *19th ACM STOC* (ACM Press, 1987), pp. 77–82
- [32] Y. Wang, J. Pan, Non-interactive zero-knowledge proofs with fine-grained security, in Dunkelman, O., Dziembowski, S. (eds.) *EUROCRYPT 2022, Part II*. LNCS, vol. 13276 (Springer, Heidelberg, 2022), pp. 305–335
- [33] Y. Wang, J. Pan, Unconditionally secure NIZK in the fine-grained setting, in Agrawal, S., Lin, D. (eds.) *ASIACRYPT 2022, Part II*. LNCS, vol. 13792 (Springer, Heidelberg, 2022), pp. 437–465
- [34] Y. Wang, J. Pan, Y. Chen, Fine-grained secure attribute-based encryption, in Malkin, T., Peikert, C. (eds.) *CRYPTO 2021, Part IV*. LNCS, vol. 12828 (Springer, Heidelberg, 2021), pp. 179–207
- [35] B. Waters, Dual system encryption: Realizing fully secure IBE and HIBE under simple assumptions, in Halevi, S. (ed.) *CRYPTO 2009*. LNCS, vol. 5677 (Springer, Heidelberg, 2009), pp. 619–636
- [36] H. Wee, Dual system encryption via predicate encodings, in Lindell, Y. (ed.) *TCC 2014*. LNCS, vol. 8349 (Springer, Heidelberg, 2014), pp. 616–637

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.