

# Adaptive Security via Deletion in Attribute-Based Encryption: Solutions from Search Assumptions in Bilinear Groups

Rishab Goyal<sup>1\*</sup>, Jiahui Liu<sup>2†</sup>, and Brent Waters<sup>2,3‡</sup>

<sup>1</sup> MIT, Cambridge, MA

<sup>2</sup> University of Texas at Austin, Austin, TX

<sup>3</sup> NTT Research, Sunnyvale, CA

**Abstract.** One of the primary research challenges in Attribute-Based Encryption (ABE) is constructing and proving cryptosystems that are adaptively secure. To date the main paradigm for achieving adaptive security in ABE is dual system encryption. However, almost all such solutions in bilinear groups rely on (variants of) either the subgroup decision problem over composite order groups or the decision linear assumption. Both of these assumptions are decisional rather than search assumptions and the target of the assumption is a source or bilinear group element. This is in contrast to earlier selectively secure ABE systems which can be proven secure from either the decisional or search Bilinear Diffie-Hellman assumption. In this work we make progress on closing this gap by giving a new ABE construction for the subset functionality and prove security under the Search Bilinear Diffie-Hellman assumption.

We first provide a framework for proving adaptive security in Attribute-Based Encryption systems. We introduce a concept of ABE with deletable attributes where any party can take a ciphertext encrypted under the attribute string  $x \in \{0, 1\}^n$  and modify it into a ciphertext encrypted under any string  $x' \in \{0, 1, \perp\}^n$  where  $x'$  is derived by replacing any bits of  $x$  with  $\perp$  symbols (i.e. “deleting” attributes of  $x$ ). The semantics of the system are that any private key for a circuit  $C$  can be used to decrypt a ciphertext associated with  $x'$  if none of the input bits read by circuit  $C$  are  $\perp$  symbols and  $C(x') = 1$ .

We show a pathway for combining ABE with deletable attributes with constrained pseudorandom functions to obtain adaptively secure ABE building upon the recent work of Tsabary [30]. Our new ABE system will be adaptively secure and be a ciphertext-policy ABE that supports the same functionality as the underlying constrained PRF as long as the

---

\*Email: [goyal@utexas.edu](mailto:goyal@utexas.edu). Work done in part while at UT Austin supported by IBM PhD Fellowship, and at the Simons Institute for the Theory of Computing supported by Simons-Berkeley research fellowship. Research supported in part by NSF CNS Award #1718161, an IBM-MIT grant, and by the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR00112020023. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Government or DARPA.

†Email: [jiahui@cs.utexas.edu](mailto:jiahui@cs.utexas.edu).

‡Email: [bwaters@cs.utexas.edu](mailto:bwaters@cs.utexas.edu). Supported by NSF CNS-1908611, CNS-1414082, DARPA SafeWare, Packard Foundation Fellowship, and Simons Investigator Award.

PRF is “deletion conforming”. Here we also provide a simple constrained PRF construction that gives subset functionality.

Our approach enables us to access a broader array of Attribute-Based Encryption schemes support deletion of attributes. For example, we show that both the Goyal et al. (GPSW) [19] and Boyen [6] ABE schemes can trivially handle a deletion operation. And, by using a hardcore bit variant of GPSW scheme we obtain an adaptively secure ABE scheme under the *Search* Bilinear Diffie-Hellman assumption in addition to pseudo random functions in NC1. This gives the first adaptively secure ABE from a *search* assumption as all prior work relied on decision assumptions over source group elements.

## 1 Introduction

Attribute-Based Encryption (ABE), since its introduction by Sahai and Waters [29], has significantly propelled the concept of secure communication. The traditional notion of Public Key Encryption (PKE) [11, 28, 14] was meant to enable a one-to-one private communication channel with a specific targeted user over an insecure network. ABE, on the other hand, provides a more fine-grained access control over plaintext delivery where it allows the encryptor to specify a policy  $f$  which is attached to the ciphertext. In such systems, each user decryption key is associated with an attribute string  $x$  such that it can recover the encrypted message only when  $f(x) = 1$ , that is when the policy  $f$  accepts the attribute  $x$ .<sup>4</sup>

Since its inception in 2005, the notion of Attribute-Based Encryption has received tremendous amount of attention. Initial developments in the context of provably secure ABE constructions as well as new proof techniques were driven by bilinear map-based realizations. The earliest such constructions (e.g. [29, 19]) were proven secure under only a relaxed notion of security called *selective* security where an attacker is required to declare the descriptor  $f^*$  that will be associated with the challenge ciphertext at the beginning of the game, i.e. even *before* seeing the public parameters. This relaxation enabled the use of a so-called “partitioning” strategy for proving security. Intuitively, availability of the challenge descriptor  $f^*$  to the reduction algorithm, before it needs to sample the system public-secret parameters, enables the reduction algorithm to shape its view of the system parameters into a *partition*. Such a partitioned view of the parameters allows the reduction algorithm to generate a secret key  $sk_x$  for every attribute  $x$  as long as  $f^*(x) = 0$  (that is, whenever  $f^*$  rejects the attribute  $x$ ), while simultaneously being able to translate a distinguishing attack on a challenge ciphertext associated with  $f^*$  into breaking a number theoretic assumption. Unfortunately, such a partitioning strategy does not naturally translate [25] to the case of full or adaptive security where an attacker gets to choose the challenge function  $f^*$  after it sees the public parameters as well as makes a polynomial

---

<sup>4</sup>For readers familiar with the notions of “ciphertext-policy” ABE and “key-policy” ABE, we will be using the ciphertext-policy vernacular in the sequel.

number of secret key queries. In this scenario the best known partitioning-style reductions will simply have to guess the function  $f^*$  to be chosen by the attacker and abort the reduction if the guess does not align with the actual choice of the attacker. This guessing approach incurs a security loss in the reduction proportional to the number of functions to choose from, and thus necessitates the use of a subexponentially secure variant of the underlying number theoretic assumption.

The shortcomings of the partitioning paradigm suggested the need for a new set of proof techniques for attaining adaptive security. The most well-known proof technique in that direction is Waters’ dual system methodology [31] which led to the first adaptively secure ABE scheme whose security was proven under a static assumption by Lewko et al. [22]. Their approach allowed for adaptive security by moving beyond partitioning proofs.<sup>5</sup> Subsequently, several other works achieved adaptive security in ABE systems with various desiderata [26, 23, 24, 32, 2]. One prominent trait of all these dual system solutions is that they almost exclusively rely on (variants of) the decision subgroup decision assumption or the decision linear assumption. Briefly, the decision linear assumption over a prime order bilinear group  $\mathbb{G}$  states that given  $g, v, w, v^a, w^b \in \mathbb{G}$  it is hard to distinguish between  $g^{a+b}$  and a random group element in  $\mathbb{G}$ . This is a potentially stronger assumption due to the facts that (1) it is decisional and (2) the target of the assumption  $g^{a+b}$  is in the bilinear group.<sup>6</sup> In contrast earlier selectively secure schemes (such as [29, 19]) can be proven secure under the Search Bilinear Diffie-Hellman assumption which states that given  $g, g^a, g^b, g^c$  it is difficult to compute  $e(g, g)^{abc}$ . In our work we work toward closing this gap by constructing new ABE systems provably secure from search assumptions.

We start by building upon a recent breakthrough due to Tsabary [30] for proving adaptively secure ABE systems from the Learning with Errors (LWE) assumption [27]. Until this work all prior ABE systems (that go beyond Identity-Based Encryption) from the LWE assumption (e.g. [15, 6, 4, 16]) relied on a partitioning argument and were thus selectively secure. Tsabary’s ABE construction is for the family of subset predicates where both private keys and ciphertexts are associated with subsets over  $[N]$  and a secret key for subset  $S$  can decrypt a ciphertext for subset  $T$  iff  $S \subseteq T$ .<sup>7</sup> While the subset predicate class is rather limited in comparison to the functionalities mentioned earlier, the work is remarkable given the lack of progress towards realizing adaptive security from LWE for so many years. (It was known [8, 17] how to prove security in a slightly weaker model where the attacker sees the public parameters, but is not allowed

<sup>5</sup>Notably, earlier works of Gentry [12] and Gentry-Halevi [13] moved beyond partitioning for IBE and Hierarchical IBE.

<sup>6</sup>If  $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$  is a bilinear map, then we refer to elements in  $\mathbb{G}$  as being in the source group or bilinear group.

<sup>7</sup>Tsabary actually presents their construction as realizing  $t$ -CNF for any constant  $t$ . However, this can be viewed as a special application of ABE for subsets. For this reason we will interpret their construction in terms of subset semantics for the purposes of this introduction.

any private key queries before committing to  $f^*$ ; however, these works do not appear to give any further insight into achieving full/adaptive security.)

Tsabary’s idea is to start with a *selectively* secure Attribute-Based Encryption scheme with certain special partial evaluation properties, and combine it with an *adaptively* secure Constrained Pseudorandom Function (CPRF) [7, 5, 21] that satisfies complementary “conforming” properties. Intuitively, the central idea in the work can be interpreted as a mechanism to leverage adaptive security of the CPRF for proving adaptive security of the resulting ABE system, while relying on the underlying selectively secure ABE scheme mostly for the encryption-decryption capability. Tsabary cleverly executed the above idea, and showed that combining these primitives in the right manner the resulting ABE system is adaptively secure, and the policy class it supports matches the constraint class of the underlying CPRF. For instantiating the entire framework, Tsabary derived a simple construction for constrained PRFs for subset constraints with requisite conforming properties from CPRF construction by [9], thereby giving an adaptively secure ABE scheme for subset predicates.

The framework requires the starting selectively secure ABE system to support partial ciphertext evaluation. Such a partial computation feature is not supported in many existing ABE systems, with the Boneh et al. [4] construction being the only known construction providing requisite capability. In particular, none of the bilinear map schemes such as [19], or the simpler (albeit less powerful) LWE-based ABE scheme of Boyen [6] support partial evaluation.

*This work.* In this work, we provide a framework to both broaden and simplify the adaptively secure ABE transformation. At the core of our work is the observation that while [30] relies on the partial ciphertext evaluation framework of Boneh et al. [4], there is hardly any computation performed on the ciphertext. Concretely, the transformation the partial evaluation performed on the ciphertext exactly corresponds to the CPRF constrain operation. Now in a CPRF scheme for subset constraints over a universe of elements  $[N]$ , the CPRF master key  $\text{msk}$  consists of  $N + 1$  regular PRF keys  $k_0, k_1, \dots, k_N$ . And, to evaluate the CPRF on a set  $S \subseteq [N]$ , the evaluator computes the following:

$$\text{CPRF}(\text{msk}, S) = F(k_0, S) \bigoplus_{i \in S} F(k_i, S).$$

For constraining the master key to a constraint set  $T \subseteq [N]$  such that evaluation works on all input sets  $S \subseteq T$ , all we need to do is “delete” all the regular PRF keys  $k_j$  for which  $j \notin T$  — thus no elaborate computation is required in constraining the key.

Our work builds around this key insight wherein we introduce the complementary notions of ABE with *deletable attributes* and *deletion conforming* CPRFs. At a high level, a *key-policy* ABE with deletable attributes allows encryption to a non-binary attribute string  $x \in \{0, 1, \perp\}^n$ , where  $\perp$  represents a “deleted” attribute. The semantics of such an ABE scheme are that a user decryption key for a policy circuit  $C$  can decrypt the ciphertext associated with

attribute  $x$  as long as the circuit  $C$  does not *touch* any of the deleted input wires and  $C(x) = 1$ .<sup>8</sup> Moreover, any user given just the public parameters can take a ciphertext  $\text{ct}$  for attribute string  $x$  and produce another ciphertext  $\text{ct}'$  encrypting the same message but for an attribute string  $x'$ , where  $x'$  is the same as  $x$  except it can have some further attribute bits deleted (i.e., changed to  $\perp$ s). Armed with these abstractions we are able to compile these into an adaptively secure ciphertext-policy ABE scheme using a transformation that follows [30] in spirit.

The potential benefits of our approach are twofold. First, we show that the framework of ABE with deletable attributes encompasses a much broader range of ABE systems. Notably, this includes the early bilinear map based GPSW construction [19] as well as the LWE-based scheme of Boyen [6].<sup>9</sup> As it turns out, showing that these schemes support attribute deletion is extremely simple — e.g., in GPSW one just has to literally “delete” ciphertext components associated with the corresponding attributes. Furthermore, following this paradigm leads to the first fully secure ABE scheme from a search problem in bilinear map setting. This is done by applying a very minor tweak to original GPSW which is to hide the message under a hardcore bit. With this tweak, we can show that the resulting scheme is adaptively secure under the *Search* Bilinear Diffie-Hellman (BDH) assumption [3] in addition to assuming pseudorandom functions in NC1 which is a minicrypt assumption. We also note that pseudorandom functions in NC1 are implied by the Bilinear decisional Diffie-Hellman assumption; thus we can alternatively base our security entirely on that assumption. We emphasize that all prior work on adaptively secure ABE from bilinear maps relied on *decision* assumptions over the source group.

A second (and perhaps more nuanced) benefit of trading off partial evaluation for deletion is in simplicity. Given that deletion is a more restricted operation arguing security is inherently simpler when we only perform deletion on input wires, compared to arbitrary partial circuit evaluation. We remark that there can be a tradeoff in the direction of functionality. While our construction using deletion matches the subset functionality given in [30], it is entirely possible that in the future we may find a larger class of functionalities that are supported by a partial computation framework and not by deletion. Doing so is an intriguing open question.

## 1.1 Technical Overview

Following the framework developed in [30], our work provides a mechanism to leverage adaptive security of a constrained PRF for upgrading the security of an

<sup>8</sup>Here by not touching an input wire, we mean that the circuit must not read/depend upon that particular input wire.

<sup>9</sup>We recently learned of the existence of an attack [1] on Boyen’s ABE scheme. We still include the proof that it is deletable to demonstrate wider applicability of our framework, but do not claim extension of Boyen’s scheme as an instantiation from LWE. To instantiate our framework under LWE, we believe that one could show the [4] scheme to be deletable.

ABE system from selective to adaptive. Concretely, we show that starting with a selectively secure *key-policy* ABE (KP-ABE) system that permits attribute deletion, we could pair it with an adaptively secure CPRF scheme to build an adaptively secure *ciphertext-policy* ABE (CP-ABE) system. Such a pairing mandates the CPRF scheme to satisfy certain special properties that we refer to as deletion conforming. The transformation flips the semantics of the underlying ABE system from key-policy to ciphertext-policy, and the constraint class associated with the CPRF maps directly to the predicate class for the resulting ciphertext-policy ABE system.

We now provide an overview of our framework and techniques. The overview is broken into four parts — first, we introduce the concept of attribute deletion for key-policy ABE systems; second, we define the complementary notion of deletion conforming CPRFs, and describe a simple construction for the family of subset constraints; third, we show how these aforementioned KP-ABE and CPRFs systems (for the right functionalities) be combined to construct an adaptively secure CP-ABE scheme; and lastly, we provide two concrete instantiations for KP-ABE with deletable attributes from standard assumptions.

*A Key-Policy ABE with Deletable Attributes.* We begin by informally introducing the concept of attribute deletion with formal definitions provided in Section 3. Recall that in the key-policy setting, the semantics of an ABE scheme are that every ciphertext  $\text{ct}_x$  is associated with an attribute string  $x \in \{0, 1\}^n$ , while every secret decryption key  $\text{sk}_C$  is associated with a policy circuit  $C : \{0, 1\}^n \rightarrow \{0, 1\}$ . Here the functionality provided by the scheme is that decryption recovers the encrypted messages whenever the policy circuit accepts the attribute (i.e.,  $C(x) = 1$ ). An ABE system with deletable attributes provides two additional capabilities — (1) the encryption algorithm can now compute ciphertexts for non-binary attribute strings  $x \in \{0, 1, \perp\}^n$  as well, where the ‘ $\perp$ ’ symbol is interpreted as an ‘unset’ attribute bit, (2) given any ciphertext  $\text{ct}_x$ , one can publicly reduce it to another ciphertext  $\text{ct}_{x'}$  encrypting the same message with the associated attribute string  $x'$  so long as  $x'$  can be obtained by having some attribute bits of  $x$  deleted (i.e., changed from *set* to *unset*).

Formally, such schemes have a special *Delete* algorithm that take as input the public parameters  $\text{pp}$ , a ciphertext  $\text{ct}_x$ , and an index set  $\mathcal{I} \subseteq [n]$  and it outputs a modified ciphertext  $\text{ct}'$ . Here the set  $\mathcal{I}$  denotes the indices of attribute bits that the user wants to delete, and let  $\text{Restrict}(x, \mathcal{I})$  denote the string  $x'$  that is obtained by deleting attribute bits of  $x$  that lie in set  $\mathcal{I}$ . The correctness requirement in presence of attribute deletion is expanded as follows: a secret key  $\text{sk}_C$  can decrypt a ciphertext  $\text{ct}_x$  if the circuit  $C$  does not *read* any of the *unset* input wires in attribute  $x$ , and evaluating  $C$  on  $x$  outputs 1. (For example, consider the following circuit:  $C(x) = x_2 \oplus x_3$ , where  $x_i$  denotes the  $i$ -th bit of  $x$ . For such a circuit  $C$ , we have that a corresponding secret key  $\text{sk}_C$  can not be used to decrypt a ciphertext  $\text{ct}_x$  whenever either  $x_2/x_3 = \perp$ , or  $x_2 \oplus x_3 \neq 1$ . That is, if  $x_2 = x_3 \neq \perp$ , then decryption succeeds irrespective of how other attribute bits are set.)

For security, such schemes must satisfy a special deletion indistinguishability property (in addition to the regular IND-CPA security). Briefly, deletion indistinguishability states that the distributions of ciphertexts generated by either running the encryption algorithm directly, or the encryption algorithm followed by the deletion algorithm should be computationally indistinguishable as long as they encrypt the same message and w.r.t. the same attribute string. That is, we have the following:

$$\{\text{Delete}(\text{pp}, \text{Enc}(\text{pp}, m, x), \mathcal{I})\} \approx_c \{\text{Enc}(\text{pp}, m, x')\}, \quad \text{where } x' = \text{Restrict}(x, \mathcal{I}).$$

Here the distributions must remain indistinguishable even if the distinguisher gets the ABE master key.

Intuitively, the goal of such a deletable key-policy ABE system is to enable arbitrary attribute deletion on ciphertexts while extending the usual policy circuit evaluation functionality over to partial/incomplete input strings. Typically, evaluating circuits on incomplete inputs is regarded as an invalid operation, but here our abstraction relies on the fact that as long as all the input wires actually used by the circuit are set (i.e., are 0/1), then we could still legally evaluate the circuit and define its output for partial inputs. As we describe later on, such a attribute deletion framework is already powerful enough for realizing adaptive security in ABE systems for subset predicates.

*Deletion Conforming CPRFs.* A regular constrained PRF (CPRF) [7, 5, 21] consists of a pseudorandom function (PRF)  $\text{CPRF}(\cdot, \cdot)$  with a key  $\text{msk}$ . The constrained property states that given master key  $\text{msk}$ , there is a way to generate a constrained key  $\text{ck}_f$  for any constraint function  $f$  such that  $\text{CPRF}(\text{msk}, x) = \text{CPRF}(\text{ck}_f, x)$  whenever  $f(x) = 1$ . Also, the standard constrained pseudorandomness property states that an attacker cannot distinguish PRF evaluations  $\text{CPRF}(\text{msk}, x_i)$  from uniformly random values on all inputs  $x_i$  for which  $f(x_i) = 0$ , even after it gets to see the constrained key  $\text{ck}_f$ . The CPRF scheme is said to be adaptively secure if the adversary can choose the challenge constraint function  $f$  after making polynomially many PRF evaluation queries. In this work, similar to [30], we instead require the CPRF to achieve adaptive key simulation security. Key simulation property states that there exists an efficient key simulation algorithm  $\text{KeySim}$  such that an attacker cannot distinguish a simulated key  $\text{ck}_f \leftarrow \text{KeySim}(f)$  from a honestly constrained key  $\text{ck}_f$  for any adaptively chosen challenge constraint  $f$  as long as all its PRF evaluation queries  $x_i$  are not satisfied by the constraint  $f$ , i.e.  $f(x_i) = 0$  for all evaluation queries  $x_i$ . Tsabary provided a CPRF construction for subset constraints which satisfies both adaptive pseudorandomness and key simulation security properties.<sup>10</sup> As a side contribution, in the main body we show that the standard constrained pseudorandomness already implies key simulation security.

Inspired by our deletable attribute framework for ABE systems, we define the notion of deletion conforming CPRFs, or DCCPRF in short. Intuitively, it

<sup>10</sup>As we pointed out before, Tsabary gives a construction for  $t$ -CNF (for constant  $t$ ) constraint functions, but this can be viewed as a special case of subset constraints.

states a CPRF system is deletion conforming if any constrained key  $\text{ck}_f$  in such a scheme can be *deterministically* computed by simply “deleting” specific bits of the master key  $\text{msk}$  (i.e., replacing some bits of the master key with the special  $\perp$  symbol). Additionally, it must be the case that the PRF evaluation algorithm for any given input  $x$  be simplified into a circuit  $C_x$  such that evaluating  $C_x$  on a master key  $\text{msk}$  and a constrained key  $\text{ck}_f$  matches on all valid inputs (i.e., all  $x$  such that  $f(x) = 1$ ). Here evaluating the circuit on a constrained key is defined similar to that for partial inputs as in the deletable KP-ABE setting, since a constrained key could have partially *unset* key bits (i.e., contain  $\perp$  symbols). All these notions are formally defined later in Section 4.

As mentioned previously, here we construct a deletion conforming CPRF for subset constraints. A subset constraint family is defined over a universe of elements  $[N] := \{1, \dots, N\}$ , where input to the PRF is a set  $S \subseteq [N]$  (which could be represented as an  $N$ -bit binary string), and each constraint function is associated with another set  $T \subseteq [N]$  such that an input set  $S$  satisfies the constraint iff  $S \subseteq T$ . A CPRF scheme for such a constraint family can be built using a combinatorial strategy as introduced in [9], where the CPRF master key  $\text{msk}$  consists of  $N + 1$  regular PRF keys  $k_0, k_1, \dots, k_N$ , and the CPRF output on a set  $S$  is computed by first selecting all PRF keys  $k_i$  such that the associated index  $i \in S$ , which is then followed by independent PRF evaluation under all selected keys and finally XORing all the evaluations together.<sup>11</sup> Concretely, the evaluator proceeds as follows:

$$\text{CPRF}(\text{msk}, S) = F(k_0, S) \bigoplus_{i \in S} F(k_i, S).$$

Note that a constrained key for a subset  $T$  can be simply set as the corresponding subset of underlying PRF keys, that is  $\text{ck}_T = \{k_0\} \cup \{k_i\}_{i \in T}$ . Observe that for every input set  $S$  satisfying the constraint set  $T$  (i.e.,  $S \subseteq T$ ), the constrained key  $\text{ck}_T$  already contains the necessary PRF keys for performing the PRF evaluation, thus correctness of evaluation for constrained keys follows immediately. Next, the proof of adaptive constrained pseudorandomness security follows from a simple observation that a reduction algorithm can simply guess an index  $i \in \{0, 1, \dots, N\}$  which is meant to denote the index of the regular PRF key that is not required for answering the constrained key query, but is needed for evaluating the CPRF on the challenge input. Since  $N$  is a polynomial, thus such a reduction strategy gives a proof of adaptive security with just polynomial security loss.

Finally, to complete our overview of CPRFs, we just need to argue that our CPRF construction satisfies the desired deletion conforming properties. This mostly follows by inspection of our aforementioned construction thereby aligning with our goal of simplicity and precision. Concretely, note that a constrained key  $\text{ck}_T$  can simply be deterministically obtained by “deleting” all the regular PRF keys  $k_i$  for which  $i \notin T$ . Also, for any input set  $S$ , the corresponding CPRF evaluation circuit can be described as: first, it reads the input wires (encoding

<sup>11</sup>In the construction the master key consists of  $N + 1$  PRF keys instead of  $N$  keys just so that pseudorandomness holds for empty set as well.



the appropriate PRF key) corresponding to set  $S$ , and then evaluates the circuit  $F(\cdot, S)$  on each block of input wires, which is finally followed up by XORing them together. Observe that since this circuit does not even read/touch the input wires corresponding to PRF keys  $k_i$  for which  $i \notin S$ , thus evaluating the circuit on a master key  $\text{msk}$  and constrained key  $\text{ck}_T$  is well-defined and gives the same output whenever  $S \subseteq T$ . Thus, this completes the proof sketch for the above CPRF to be deletion conforming. More details on our construction are provided in the full version.

*Building adaptively secure Ciphertext-Policy ABE.* Moving on to our main transformation, our approach is to decouple the adaptivity and functionality (delivering the message to users) requirements of a CP-ABE scheme, and deal with them separately. Following Tsabary's paradigm, we rely on our deletion conforming CPRFs for enabling the reduction algorithm to be able to answer the adaptive key queries, while still using the selectively secure deletable KP-ABE system for guaranteeing that the message is hidden. At a very high level, the idea is to handle the adaptivity problem outside of the underlying KP-ABE system, while using its attribute deletion capabilities to compute the CP-ABE challenge ciphertext from a KP-ABE challenge ciphertext that was selectively obtained. Below we sketch our transformation.

The public parameters of the CP-ABE system contains the deletable (KP-ABE) parameters  $\text{del.pp}$ , while the master secret key consists of a DCCPRF master key  $\text{prf.msk}$  as well as the deletable ABE master key  $\text{del.msk}$ . Recall that in a CP-ABE system, each secret key is associated with an attribute string  $x \in \{0, 1\}^N$ . To sample a secret key for attribute  $x$ , the key generator first computes a tag value  $t$  as the CPRF evaluation with input  $x$ , i.e.  $t = \text{CPRF}(\text{prf.msk}, x)$ . Let  $C_x$  denote the simplified explicit circuit that performs the CPRF evaluation on input  $x$ , i.e.  $C_x(\text{key}) = \text{CPRF}(\text{key}, x)$ . The key generator then creates a policy circuit  $f_{x,t}$ , given the tag value  $t$  and circuit description  $C_x$ , as:

$$f_{x,t}(z) = \begin{cases} 1 & \text{if } C_x(z) \neq t, \\ 0 & \text{otherwise.} \end{cases}$$

The (CP-ABE) secret key  $\text{sk}_x$  for attribute  $x$  now corresponds to a (KP-ABE) secret key for the above policy circuit, i.e.  $\text{sk}_x = \text{del.sk}_{f_{x,t}}$ . To encrypt a message  $m$  under a policy circuit  $g$ , the encryptor first samples a *simulated* constrained key  $\widetilde{\text{prf.sk}}_g$  with  $g$  being used as the constraint function, and then it computes the ciphertext as an KP-ABE encryption of message  $m$  with attribute string set as  $\text{sk}_g$ . The resulting decryption algorithm is exactly the decryption algorithm of the underlying KP-ABE scheme.

First, note that, by the deletion conforming properties, evaluating  $C_x$  is well-defined and accurately matches the corresponding CPRF output on every accepting constrained key. Thus with this observation we get that correctness of the above construction follows from the fact that whenever  $g(x) = 1$ , then  $f_{x,t}(\widetilde{\text{prf.sk}}_g) = 1$  with all but negligible probability, since  $C_x(\text{prf.sk}_g) = t =$

$C_x(\text{prf.msk})$  happens only with negligible probability by pseudorandomness of the underlying CPRF.

Next we describe the intuition behind the proof of adaptive security. Note that initially the challenge ciphertext for policy  $g^*$  with message  $m$  is computed as KP-ABE encryption of message  $m$  with a simulated CPRF constrained key  $\widetilde{\text{prf.sk}}_{g^*}$  as the attribute string. As a first step, we instead switch this to be a honestly constrained key  $\text{prf.sk}_{g^*} = \text{Constrain}(\text{prf.msk}, g^*)$ . Since the CPRF satisfies the adaptive key simulation property, thus this change will be indistinguishable. Note that it is important that the CPRF is adaptively secure for this reduction to work since to answer the pre-challenge key queries, the reduction algorithm needs to query for the respective CPRF evaluations. Next, by the deletion conforming property of the constrained PRF scheme, we have that the constrained key  $\text{prf.sk}_{g^*}$  can be computed by simply deleting certain specific key bits of the master key  $\text{prf.msk}$ . Let  $\mathcal{I}_{g^*}$  denote such a set of indices, i.e.  $\text{prf.sk}_{g^*} = \text{Restrict}(\text{prf.msk}, \mathcal{I}_{g^*})$ . By relying on the deletion indistinguishability property of the KP-ABE scheme, we get that the challenge ciphertext can instead be computed as first encrypting the message  $m$  under attribute string  $\text{prf.msk}$ , and then deleting the attribute bits as specified by set  $\mathcal{I}_{g^*}$  by running the KP-ABE deletion algorithm. Finally, since the attribute string  $\text{prf.msk}$  is sampled at the beginning of the security game, thus  $\text{prf.msk}$  can be selectively specified to the KP-ABE challenger thereby allowing us to argue that the message is also hidden. Our construction and its proof is formally provided in Section 6.

*Perfect correctness?* Although at first glance it may seem that imperfect correctness is an inherent and unavoidable feature of the above framework, we show in the full version that this is not the case where we provide an alternate construction which is perfectly correct. Very briefly, our idea is to have two deletable ABE sub-systems working in parallel, instead of just one, where both the ciphertexts and secret keys contain two copies (one under each ABE sub-system). The only difference is that while sampling a secret key under both the systems independently, the key generator uses two distinct tag values, where one of the tag values is computed as is now, whereas the other tag value will be its complement. Such a trick gets around the imperfect correctness problem since it can never happen that  $C_x(\widetilde{\text{prf.sk}}_g) = t_0$  as well as  $C_x(\widetilde{\text{prf.sk}}_g) = t_1$  where  $t_0, t_1$  are the complementary tag pairs. It turns out that the proof of adaptive security now is more involved, as we need to first use the existing proof structure to erase the information about the challenge message from the first deletable ABE sub-ciphertext, then we would have to undo correlations created between parts of the challenge ciphertext and secret keys, and finally use a similar proof structure to erase the information about the challenge message from the second deletable ABE sub-ciphertext as well.

*Another interpretation.* Abstractly, the deletion paradigm described above can be interpreted as a mechanism to selectively activate the trapdoors embedded inside the secret keys such that whenever trapdoor is activated then the challenger can simulate the secret keys for all possible attributes. The property such

simulated keys satisfy is that they are indistinguishable from honestly sampled secret keys as long as the challenge policy does not accept the corresponding key attribute. On a more intuitive level, one could also observe some similarities between the above framework and the Dual System methodology [31], where switching from a simulated CPRF key to an honestly constrained CPRF key could be comparable to moving from a *normal* to a *semi-functional* ciphertext, and the secret keys are already sampled in the semi-functional mode.

*Deleting attributes in [19, 6].* Finally, we show that existing ABE schemes by Goyal et al. (GPSW) [19] from bilinear maps, and by Boyen [6] from LWE<sup>12</sup> already lie in the class of ABE schemes with deletable attributes, thereby displaying the generality of our framework. Below we give an overview of our deletion algorithms. More details are provided later in Section 7, where we also show that a KP-ABE scheme with deletable attributes for monotonic access structures can be generically upgraded to non-monotonic log-depth circuits (i.e.,  $\mathbf{NC}^1$ ).

*Deletions in [19].* First, we look at the bilinear map based ABE construction by GPSW. They proposed a KP-ABE scheme for monotone access structures and proved its security under the Decisional Bilinear Diffie-Hellman (DBDH) assumption that can also be readily adapted to a scheme provably secure under the Search Bilinear Diffie-Hellman assumption. The public parameters in the GPSW scheme contain  $n$  group elements in the base group  $\{T_i\}_{i \in [n]}$  and one group element in the target group  $K$ , where  $n$  denotes the length of the attributes. A ciphertext encrypting a message  $m$  under an attribute  $x \in \{0, 1\}^n$  is of the following form:

$$\text{ct} = (m \cdot K^s, \{T_i^s\}_{i \in [n]: x_i=1}),$$

where  $s$  is a random exponent. Basically the term  $T_i^s$  encodes the  $i$ -th bit of the attribute, and during decryption the algorithm pairs the ciphertext component  $T_i^s$  with a corresponding key component (iff the policy circuit reads the  $i$ -th input wire) and performs a polynomial interpolation in the exponent to reconstruct the masking term  $K^s$ .

Our observation is that to delete an attribute bit, say  $j$ , one could simply drop the term  $T_j^s$  from ciphertext (if it exists). Multiple attribute bits could be deleted analogously. As long as the policy circuit does not read the deleted input wire, the correctness for deleted ciphertexts follows immediately from the correctness of GPSW scheme itself. Similarly, to encrypt a message  $m$  under a non-binary attribute string  $x \in \{0, 1, \perp\}^n$ , we simply treat each  $\perp$  symbol as a 0 bit, and therefore do not encode it in the ciphertext. Clearly, the distributions of freshly computed ciphertexts and deleted ciphertexts (encrypting the same message  $m$  and attribute  $x$ ) are identical, thus deletion indistinguishability for GPSW is merely a statistical property. Combining this with the fact that GPSW

<sup>12</sup>We want to remind the reader the existence of an attack [1] on Boyen’s ABE scheme. Deletions in Boyen’s scheme are merely provided for illustrative purposes in the full version.

provides selective IND-CPA security, we obtain that GPSW augmented with the deletion procedure is KP-ABE scheme with deletable attributes. Later in the full version [18] we also describe a hardcore bit variant of the above scheme whose security relies on the Computational Bilinear Diffie-Hellman (CBDH) assumption.

*Deletions in [6].* Next, we look at the LWE-based ABE construction by Boyen. Boyen’s scheme is also for monotone access structures and its security relies on the LWE assumption. The public parameters in Boyen’s scheme consist  $\ell + 1$  matrices of appropriate dimensions  $(\mathbf{A}_0, \{\mathbf{A}_i\}_{i \in [\ell]})$  and a vector  $\mathbf{u}$ , where  $\ell$  denotes the length of the attributes. Now a ciphertext  $\text{ct}$  encrypting a message bit  $\text{msg}$  under an attribute  $x \in \{0, 1\}^\ell$  is of the following form  $\text{ct} = (c_0, \mathbf{c}_{1,0}, \mathbf{c}_{1,1}, \dots, \mathbf{c}_{1,\ell})$ , where

$$c_0 = \mathbf{s}^\top \cdot \mathbf{u} + \nu_0 + \lfloor \frac{q}{2} \rfloor \cdot \text{msg},$$

$$\forall i \in [0, \ell], \quad \mathbf{c}_{1,i} = \begin{cases} \mathbf{s}^\top \cdot \mathbf{A}_i + \nu_{1,i} & \text{if } i = 0 \text{ or } x_i = 1, \\ \nu_{1,i} & \text{otherwise.} \end{cases}$$

and  $\mathbf{s}$  is a random secret vector, and  $\nu_0, \{\nu_{1,i}\}_{i \in [\ell]}$  are sampled i.i.d. according to the LWE noise distribution. Here the vector  $\mathbf{c}_{1,i}$  encodes the  $i$ -th bit of the attribute, and during decryption the algorithm combines the ciphertext component  $\mathbf{c}_{1,i}$  with a corresponding key component (iff the policy circuit reads the  $i$ -th input wire).

For deleting attributes in Boyen’s scheme, instead of dropping the respective ciphertext component, we replace with freshly sampled noise. Concretely, to delete an attribute bit, say  $j$ , we replace the vector  $\mathbf{c}_{1,j}$  in the ciphertext with a freshly sampled noise vector  $\nu'_{1,i}$ . Multiple attribute bits could be deleted analogously.<sup>13</sup> And as for our augmented GPSW scheme, during encryption we treat each  $\perp$  symbol as a 0 bit, and the correctness and deletion indistinguishability of the resultant follows either immediately from Boyen’s scheme or by inspection.

*Recent Independent Work* Recently, Katsumata, Nishimaki, Yamada, and Yamakawa (KNYY) [20] gave an exciting construction showing how to expand the framework of [30] to encompass an inner product encryption and Fuzzy IBE functionality within the LWE setting. An important insight was showing that a specific cryptosystem could relax the earlier conforming property to just functional equivalence and thus leverage a particular constrained PRF of [10] to achieve greater functionality.

In contrast, our work shows how to relax the conforming property to deletion so that it is realizable in a broader setting that includes bilinear maps. But we show that is still sufficient to maintain the  $t$ -CNF functionality. KNYY show that in the LWE setting one can strengthen the framework to handle a broader class of LWE specific constrained PRFs. The works were performed independently.

<sup>13</sup>We could also drop the deleted ciphertext components instead of replacing them with LWE noise, however for ensuring consistency with Boyen’s scheme we keep it this way.

*Comparing techniques with [30].* We conclude by giving some further technical comparisons between our framework and the earlier work of Tsabary [30] that we build upon. Our work follows a similar pathway which is to leverage adaptive security of constrained PRFs (with special properties) inside a key-policy ABE scheme (with special properties) to achieve an adaptively secure ciphertext-policy ABE scheme, but differences lie in the flavour of these special properties required from the underlying constrained PRF and key-policy ABE systems. Tsabary started with the LWE-based ABE construction of Boneh et al. [4], and using the homomorphic properties of the underlying ABE scheme, Tsabary developed a framework for partial ciphertext evaluation and a circuit splitting/composition abstraction, wherein the ABE scheme allows a user to encrypt messages under *partially evaluated* attributes such that they are indistinguishable from *partially evaluated* ciphertexts encrypting same message under the original (unevaluated) attribute. Concretely, [30] relies on the fact that for any attribute  $x$  and circuit  $C$ , one could compute ciphertexts of the form:  $\text{ct}_0 = \text{Enc}(\text{pp}, m, x)$ ,  $\text{ct}_1 = \text{Enc}(\text{pp}, m, C(x))$  such that given a secret key  $\text{sk}_{\tilde{C}}$  for some circuit  $\tilde{C}$  s.t.  $\tilde{C}(x) = 1$ , a user can not only decrypt ciphertexts of the form  $\text{ct}_0$ , but it can also decrypt ciphertexts of the form  $\text{ct}_1$  as long as there exists another circuit  $C'$  with the semantics that  $\tilde{C}(\cdot) = C'(C(\cdot))$  that the decryptor knows. Here the equality between the circuit  $\tilde{C}$  and the composition of  $C, C'$  mandates the resultant ‘gate-by-gate’ circuit descriptions must be *identical*. With such an ABE scheme with these special properties as the centerpiece, [30] built a constrained PRF that conforms with the necessary circuit splitting/composition semantics. Very briefly, [30] required that the PRF evaluation circuit with the input hardwired can be split into two sub-circuits such that one of those sub-circuits can be used during generating the CP-ABE ciphertext. Combining all these things in an extremely careful manner gives the desired result of an adaptively secure CP-ABE scheme for subset policies.

Our approach, on the other hand, is to skip the entire partial evaluation and circuit splitting/composition framework, and instead go with a simpler abstraction of input deletion while also demanding (as part of our definitional framework) an explicit descriptions for all the circuits used throughout the analysis.

## 2 Preliminaries

*Notation.* Let PPT denote probabilistic polynomial-time. We denote the set of all positive integers upto  $n$  as  $[n] := \{1, \dots, n\}$ . Also, we use  $[0, n]$  to denote the set of all non-negative integers upto  $n$ , i.e.  $[0, n] := \{0\} \cup [n]$ . Throughout this paper, unless specified, all polynomials we consider are positive polynomials. For any finite set  $S$ ,  $x \leftarrow S$  denotes a uniformly random element  $x$  from the set  $S$ . Similarly, for any distribution  $\mathcal{D}$ ,  $x \leftarrow \mathcal{D}$  denotes an element  $x$  drawn from distribution  $\mathcal{D}$ . The distribution  $\mathcal{D}^n$  is used to represent a distribution over vectors of  $n$  components, where each component is drawn independently from the distribution  $\mathcal{D}$ .

For any  $n \in \mathbb{N}$ , string  $x \in \{0, 1, \perp\}^n$  and index set  $\mathcal{I} \subseteq [n]$ , let  $\text{Restrict}(x, \mathcal{I})$  denote the string  $\tilde{x} \in \{0, 1, \perp\}^n$  such that

$$\forall i \in [n], \quad \tilde{x}_i = \begin{cases} x_i & \text{if } i \notin \mathcal{I}, \\ \perp & \text{otherwise.} \end{cases}$$

where  $x_i$  and  $\tilde{x}_i$  denote the  $i$ th elements of strings  $x$  and  $\tilde{x}$ , respectively. For any string  $x \in \{0, 1, \perp\}^n$ , let  $\text{BotSet}(x)$  denote the subset of indices in  $[n]$  such that for every  $i \in \text{BotSet}(x)$ ,  $x_i = \perp$  and for every  $i \notin \text{BotSet}(x)$ ,  $x_i \in \{0, 1\}$ . Formally,  $\text{BotSet}(x) := \{i \in [n] : x_i = \perp\}$ .

*Circuit notation.* Also, throughout the paper we use the circuit model of computation. Consider any circuit  $C : \{0, 1\}^n \rightarrow \{0, 1\}$  that takes  $n$ -bits of input and outputs a single bit. For any circuit  $C$ , we define  $\text{Unsupported}(C) \subseteq [n]$  to be set of indices  $i \in [n]$  such that the circuit  $C$  does not use on the  $i$ th input wire (i.e.,  $C$  does not read the  $i$ th input bit).<sup>14</sup>

Lastly, we use  $\text{CEval}$  to denote an “expanded” notion of circuit evaluation. The algorithm  $\text{CEval}$  takes as input a circuit  $C : \{0, 1\}^n \rightarrow \{0, 1\}^m$ , and a string  $x \in \{0, 1, \perp\}^n$ , and it first checks that  $\text{BotSet}(x) \subseteq \text{Unsupported}(C)$ . If the check fails, it outputs the all-zeros string  $0^m$ ; otherwise it evaluates the circuit  $C$  on string  $x$ , and outputs the same result as the circuit which is  $C(x)$ . Note that evaluating the circuit  $C$  on string  $x$  (that could possibly contain non-binary input bits) is well-defined in the last step, since the evaluator  $\text{CEval}$  only runs the circuit  $C$  after its checks that  $\text{BotSet}(x) \subseteq \text{Unsupported}(C)$ , and thus we know that if the check succeeds then all the input wires/bits read by circuit  $C$  are defined and not set as  $\perp$ . Formally,  $\text{CEval}$  can be defined as:

$$\text{CEval}(C, x) = \begin{cases} C(x) & \text{if } \text{BotSet}(x) \subseteq \text{Unsupported}(C), \\ 0^m & \text{otherwise.} \end{cases}$$

Due to space constraints, we move the definition of regular pseudorandom functions to the full version.

### 3 Key Policy Attribute-Based Encryption with Deletable Attributes

In this section, we introduce the notion of Key Policy Attribute-Based Encryption (KP-ABE) with deletable attributes. First, we provide the syntax, and later describe our definitions for KP-ABE with deletable attributes.

<sup>14</sup>Note that our definition of the *unsupported indices* for a circuit  $C$  is very restrictive. Concretely, we say that an index  $i \in \text{Unsupported}(C)$  iff as per the circuit description of  $C$  the  $i$ th input wire is unused/untouched. For instance, consider two circuits  $C, \tilde{C}$  which takes length 2-bit strings as inputs:  $C(x) = (x_1 \vee \neg x_1) \wedge x_2$  and  $\tilde{C}(x) = x_2$ . Here  $\text{Unsupported}(C) = \emptyset$  and  $\text{Unsupported}(\tilde{C}) = \{1\}$ , i.e. circuits  $C, \tilde{C}$  have different unsupported indices even though they are functionally identical. This is because as per the circuit description of  $C$ , it does use both input wires/bits; whereas  $\tilde{C}$  ignores the first input wire/bit.

*Syntax.* A key-policy attribute based encryption (KP-ABE) scheme with deletable attributes for a class of circuits  $\mathcal{C} = \{\mathcal{C}_n\}_{n \in \mathbb{N}}$  and message space  $\mathcal{M}$  consists of the following PPT algorithms:

- Setup** $(1^\lambda, 1^n) \rightarrow (\text{pp}, \text{msk})$ . On input the security parameter  $\lambda$  and attribute length  $n$ , the setup algorithm outputs a set of public parameters  $\text{pp}$ , and master secret key  $\text{msk}$ .
- KeyGen** $(\text{msk}, f) \rightarrow \text{sk}_f$ . On input the master secret key  $\text{msk}$  and a circuit  $f \in \mathcal{C}_n$ , the key generation algorithm outputs a predicate key  $\text{sk}_f$ .
- Enc** $(\text{pp}, x, m) \rightarrow \text{ct}$ . On input the public parameters  $\text{pp}$ , an attribute string  $x \in \{0, 1, \perp\}^n$ , and a message  $m \in \mathcal{M}$ , the encryption algorithm outputs a ciphertext  $\text{ct}$ . Note that here the attribute string  $x$  is possibly a non-binary string as it could contain  $\perp$  symbols.
- Dec** $(\text{sk}_f, \text{ct}) \rightarrow m/\text{fail}$ . On input a secret key  $\text{sk}_f$  and a ciphertext  $\text{ct}$ , the decryption algorithm either outputs a message  $m$  or a special string **fail** (to denote decryption failure).
- Delete** $(\text{pp}, \text{ct}, \mathcal{I}) \rightarrow \text{ct}'$ . On input of the public parameters  $\text{pp}$ , a ciphertext  $\text{ct}$  and a set of indices  $\mathcal{I} \subseteq [n]$ , the deletion algorithm outputs a modified ciphertext  $\text{ct}'$ .

We require such an ABE scheme to satisfy the following properties.

*Correctness.* Intutively, it says that the above scheme is correct if decrypting a ciphertext, which was either directly computed using the encryption algorithm or generated by the ciphertext deletion algorithm, outputs the correct message as long as the policy circuit accepts the attribute associated with the ciphertext.

Formally, an KP-ABE scheme with deletable attributes is said to be correct if for all  $\lambda, n \in \mathbb{N}$ ,  $f \in \mathcal{C}_n$ ,  $m \in \mathcal{M}$ ,  $x_0 \in \{0, 1, \perp\}^n$  and a sequence of indices sets  $\mathcal{I}_1, \mathcal{I}_2 \dots, \mathcal{I}_k \subseteq [n]$ , for any  $k \geq 0$ , the following holds:

$$\text{CEval}(f, x_k) = 1 \implies \Pr \left[ \begin{array}{l} (\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^n) \\ \text{Dec}(\text{sk}_f, \text{ct}_k) = m : \text{sk}_f \leftarrow \text{KeyGen}(\text{msk}, f), \text{ct}_0 \leftarrow \text{Enc}(\text{pp}, x_0, m) \\ (\forall i \in [k]) \text{ct}_i \leftarrow \text{Delete}(\text{pp}, \text{ct}_{i-1}, \mathcal{I}_i) \end{array} \right] = 1,$$

where  $x_k$  is defined by the following sequence of operations:  $x_i \leftarrow \text{Restrict}(x_{i-1}, \mathcal{I}_i)$  for all  $i \in [k]$ .

*Security.* For security, we have two requirements. First, we require the scheme to provide standard semantic security as for standard ABE schemes. Here we consider both selective and adaptive IND-CPA security definitions. Second, we introduce a notion of indistinguishability for ciphertexts with deleted attributes, in which the adversary cannot distinguish between a ciphertext modified by the Delete algorithm and a ciphertext directly encrypted from the same message with respect to the same attribute string after deletion. Formally, they are defined as below.

**Definition 3.1 (Adaptive IND-CPA Security).** A KP-ABE scheme is adaptively secure if for every stateful admissible PPT adversary  $\mathcal{A}$ , there exists a negligible function  $\text{negl}(\cdot)$  such that for all  $\lambda, n \in \mathbb{N}$ , the following holds

$$\Pr \left[ \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{ct}) = b : \begin{array}{l} (\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^n) \\ ((m_0, m_1), x^* \in \{0, 1\}^n) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(1^\lambda, 1^n, \text{pp}) \\ b \leftarrow \{0, 1\}; \text{ct} \leftarrow \text{Enc}(\text{pp}, x^*, m_b) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda),$$

where the adversary  $\mathcal{A}$  is admissible as long as every secret key query  $f$  made by  $\mathcal{A}$  to the oracle  $\text{KeyGen}(\text{msk}, \cdot)$  satisfies the condition that  $f(x^*) = 0$ . Here  $x^*$  is the challenge attribute chosen by  $\mathcal{A}$ . Note that the adversary must choose  $x^*$  as a binary string, that is it must not contain any  $\perp$  symbols.<sup>15</sup>

**Definition 3.2 (Selective IND-CPA Security).** A KP-ABE scheme is said to be selectively secure if in the above security game (see Definition 3.1), the adversary must instead declare the challenge attribute  $x^* \in \{0, 1\}^n$  at the beginning of the game, that is even before it receives the public parameters  $\text{pp}$  from the challenger.

**Definition 3.3 (Deletion Indistinguishability).** A KP-ABE scheme with deletable attributes satisfies deletion indistinguishability property if for every stateful PPT adversary  $\mathcal{A}$ , there exists a negligible function  $\text{negl}(\cdot)$ , such that for all  $\lambda, n \in \mathbb{N}$ , the following holds

$$\Pr \left[ \mathcal{A}(\text{ct}_b) = b : \begin{array}{l} (\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^n), b \leftarrow \{0, 1\} \\ (m, x \in \{0, 1, \perp\}^n, \mathcal{I} \subseteq [n]) \leftarrow \mathcal{A}(1^\lambda, 1^n, \text{pp}, \text{msk}) \\ \tilde{\text{ct}} \leftarrow \text{Enc}(\text{pp}, x, m), \text{ct}_0 \leftarrow \text{Delete}(\text{pp}, \tilde{\text{ct}}, \mathcal{I}) \\ \tilde{x} \leftarrow \text{Restrict}(x, \mathcal{I}), \text{ct}_1 \leftarrow \text{Enc}(\text{pp}, \tilde{x}, m) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda).$$

Note that the attribute vector  $x$  chosen by the adversary  $\mathcal{A}$  can contain  $\perp$  symbols.

## 4 Constrained PRFs: Defining Deletion Conformity

In this section, we first recall the notion of constrained PRFs (CPRFs) [7, 5, 21], and later introduce our notion of *deletion conforming* CPRFs.

*Syntax.* A constrained PRF (CPRF) for constraint class  $\mathcal{F} = \{\mathcal{F}_N\}_{N \in \mathbb{N}}$  consists of three PPT algorithms (Setup, Constrain, Eval) with the following syntax:

**Setup**( $1^\lambda, 1^N$ )  $\rightarrow$   $\text{msk}$ . On input the security parameter  $\lambda$  and input length  $N$ , the setup algorithm outputs a master secret key  $\text{msk} \in \{0, 1\}^k$ . Let  $k = k(\lambda, N)$  denote the length of secret key, where  $k(\cdot, \cdot)$  is an a-priori fixed polynomial.

**Constrain**( $\text{msk}, f$ )  $\rightarrow$   $\text{sk}_f$ . On input a constraint function  $f \in \mathcal{F}_N$  and master secret key  $\text{msk}$ , the constrain algorithm outputs a constrained key  $\text{sk}_f$ .

<sup>15</sup>Note that since  $x^*$  does not contain  $\perp$  symbols, thus  $f(x^*)$  is always well-defined and we do not need define the admissibility constraint as  $\text{CEval}(f, x^*) = 0$  instead.



$\text{Eval}(\text{sk}, x) \rightarrow y$ . The evaluation algorithm takes as input a (possibly constrained) secret key  $\text{sk}$  and a string  $x \in \{0, 1\}^N$ , and outputs a string  $y$ . Let  $m = m(\lambda, N)$  denote the length of the output string  $y$  for some polynomial  $m(\cdot, \cdot)$ .<sup>16</sup>

*Correctness of CPRF evaluation.* A CPRF scheme is said to be correct if for all  $\lambda, N \in \mathbb{N}$ ,  $f \in \mathcal{F}_N$ , and  $x \in \{0, 1\}^N$ , the following holds:

$$f(x) = 1 \implies \Pr \left[ \text{Eval}(\text{msk}, x) = \text{Eval}(\text{sk}_f, x) : \begin{array}{l} \text{msk} \leftarrow \text{Setup}(1^\lambda, 1^N) \\ \text{sk}_f \leftarrow \text{Constrain}(\text{msk}, f) \end{array} \right] = 1$$

*Security.* Next, we recall the notion of single-key adaptive pseudorandomness security for constrained PRFs. Later on we also define the notion of key simulation security as defined in [30].

**Definition 4.1 (Adaptive single-key constrained pseudorandomness).** We say that a CPRF = (Setup, Constrain, Eval) satisfies adaptive single-key constrained pseudorandomness security if for any stateful admissible PPT adversary  $\mathcal{A}$  there exists a negligible function  $\text{negl}(\cdot)$ , such that for all  $\lambda, N \in \mathbb{N}$ , the following holds:

$$\Pr \left[ \mathcal{A}^{\text{Eval}(\text{msk}, \cdot), \text{Constrain}(\text{msk}, \cdot)}(r_b) = b : \begin{array}{l} \text{msk} \leftarrow \text{Setup}(1^\lambda, 1^N), b \leftarrow \{0, 1\} \\ x^* \leftarrow \mathcal{A}^{\text{Eval}(\text{msk}, \cdot), \text{Constrain}(\text{msk}, \cdot)}(1^\lambda, 1^N) \\ r_0 \leftarrow \{0, 1\}^m, r_1 = \text{Eval}(\text{msk}, x^*) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda).$$

Here the adversary  $\mathcal{A}$  is said to be admissible as long as it satisfies the following conditions — (1) it makes at most one query to the constrain oracle  $\text{Constrain}(\text{msk}, \cdot)$ , and its queried function  $f$  must be such that  $f(x^*) = 0$ , (2) it must not send  $x^*$  as one of its evaluation queries to  $\text{Eval}(\text{msk}, \cdot)$ .

The above pseudorandomness security could be extended to collusion-resistant notions where the adversary could make polynomially many constrain queries, however in this work we only require single-key security. Next, we define key simulation security for CPRFs.

**Definition 4.2 (Adaptive key simulation).** We say that a CPRF = (Setup, Constrain, Eval) satisfies adaptive key simulation security if there exists a PPT algorithm  $\text{KeySim}$  such that for any stateful admissible PPT adversary  $\mathcal{A}$ , there exists a negligible function  $\text{negl}(\cdot)$ , such that for all  $\lambda, N \in \mathbb{N}$ , the following holds:

$$\Pr \left[ \mathcal{A}^{\text{Eval}(\text{msk}, \cdot)}(sk_b) = b : \begin{array}{l} \text{msk} \leftarrow \text{Setup}(1^\lambda, 1^N), b \leftarrow \{0, 1\} \\ f^* \leftarrow \mathcal{A}^{\text{Eval}(\text{msk}, \cdot)}(1^\lambda, 1^N) \\ \text{sk}_0 \leftarrow \text{KeySim}(1^\lambda, 1^N, f^*), \text{sk}_1 \leftarrow \text{Constrain}(\text{msk}, f^*) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda).$$

Here the adversary  $\mathcal{A}$  is said to be admissible if all its evaluation queries  $x \in \{0, 1\}^N$  satisfy the condition that  $f^*(x) = 0$ . That is, none of the queried inputs are satisfied by the constraint  $f^*$ .

<sup>16</sup>Here we consider a single PRF evaluation algorithm that could take as input a master key as well as a constrained key. Thus, both the master and constrained keys are of same length  $k$ . Note that one could instead split it into two separate evaluation algorithms, however for ease of exposition we avoid it.

*Non-colliding property.* A constrained PRF CPRF that satisfies key simulation security (Definition 4.2) is said to be *non-colliding* if there exists a negligible function  $\text{negl}(\cdot)$  such that for all  $\lambda, N \in \mathbb{N}$ , every input  $x \in \{0, 1\}^N$ , constraint function  $f \in \mathcal{F}_N$ , the following holds:

$$\Pr \left[ \text{Eval}(\text{msk}, x) = \text{Eval}(\text{sk}'_f, x) : \begin{array}{l} \text{msk} \leftarrow \text{Setup}(1^\lambda, 1^N) \\ \text{sk}'_f \leftarrow \text{KeySim}(1^\lambda, 1^N, f) \end{array} \right] \leq \text{negl}(\lambda).$$

Later on in the full version, we show that if the CPRF satisfies (0-key) pseudorandomness security, then it also satisfies the non-colliding property as long as the output length of the PRF is large enough. Additionally, we also show that the adaptive single-key constrained pseudorandomness security in fact implies adaptive key simulation security.

#### 4.1 Deletion Conforming CPRFs

Now we define the deletion conforming property for CPRFs. Intuitively, it states that a constrained key in such a CPRF scheme must be deterministically computable by simply deleting specific bits of the master key (i.e., replacing some bits of the master key with a special  $\perp$  symbol). Formally we define it below.

**Definition 4.3 (Deletion Conforming CPRF).** *We say that a constrained PRF scheme  $\text{CPRF} = (\text{Setup}, \text{Constrain}, \text{Eval})$  for a function class  $\mathcal{F} = \{\mathcal{F}_N\}_{N \in \mathbb{N}}$  is a deletion conforming CPRF if the constrain algorithm  $\text{Constrain}$  is deterministic, and there exists two polynomial time algorithms  $(\text{CircuitGen}, \text{DeleteFunc})$  with the following syntax and properties:*

$\text{CircuitGen}(1^\lambda, 1^N, x) \rightarrow C_x$ . *The circuit generation algorithm is a deterministic algorithm that takes as input the security parameter  $\lambda$ , length parameter  $N$ , and input string  $x \in \{0, 1\}^N$ . It outputs the description of a circuit  $C_x$ .*

$\text{DeleteFunc}(1^\lambda, 1^N, f) \rightarrow \mathcal{I}_f$ . *The key deletion algorithm is a deterministic algorithm that takes as input the security parameter  $\lambda$ , length parameter  $N$ , and a constraint function  $f \in \mathcal{F}_N$ . It outputs a set of indices  $\mathcal{I}_f \subseteq [k]$ , where  $k$  denotes the length of the master secret key.*

*We say that  $\text{DCCPRF} = (\text{Setup}, \text{Constrain}, \text{Eval}, \text{CircuitGen}, \text{DeleteFunc})$  is a deletion conforming CPRF if for all  $\lambda, N \in \mathbb{N}$ , every function  $f \in \mathcal{F}_N$ , input  $x \in \{0, 1\}^N$ , and master key  $\text{msk} \leftarrow \text{Setup}(1^\lambda, 1^N)$ , the following properties are satisfied.*

1. **Function deletion property:**  $\text{Constrain}(\text{msk}, f) = \text{Restrict}(\text{msk}, \mathcal{I}_f)$ , where index set  $\mathcal{I}_f$  is computed as  $\mathcal{I}_f = \text{DeleteFunc}(1^\lambda, 1^N, f)$ .
2. **Circuit evaluation property:** Let  $C_x = \text{CircuitGen}(1^\lambda, 1^N, x)$ . It states that  $\text{Eval}(\text{msk}, x) = C_x(\text{msk})$  irrespective of whether  $f(x) = 0/1$ , and  $\text{Eval}(\text{sk}_f, x) = \text{CEval}(C_x, \text{sk}_f)$  whenever  $f(x) = 1$  where  $\text{sk}_f = \text{Constrain}(\text{msk}, f)$  or  $\text{sk}_f \leftarrow \text{KeySim}(1^\lambda, 1^N, f)$ .

Here recall that the `Restrict` and `CEval` operations are as defined in Section 2 — `Restrict( $s, \mathcal{I}$ )` denotes a string after replacing the bits in  $s$  with indices corresponding to indices in set  $\mathcal{I}$  with  $\perp$ ; and `CEval( $C, x$ )` denotes evaluating the circuit  $C$  on input  $x$ , but setting the circuit output to be the all zeros string  $0^m$  if the circuit  $C$  depends on the input wires whose indices have  $\perp$  symbol in  $x$ .

## 5 Ciphertext Policy Attribute-Based Encryption

In this section, we recall the notion of Ciphertext Policy Attribute-Based Encryption (CP-ABE). First, we provide the syntax and definitions, and later define the predicate class we study in this work.

*Syntax.* A ciphertext-policy attribute based encryption (CP-ABE) scheme for a class of predicates  $\mathcal{F} = \{\mathcal{F}_N\}_{N \in \mathbb{N}}$  and message space  $\mathcal{M}$  consists of the following PPT algorithms:

`Setup( $1^\lambda, 1^N$ )`  $\rightarrow$  `(pp, msk)`. On input the security parameter  $\lambda$  and attribute length  $N$ , the setup algorithm outputs a set of public parameters `pp`, and master secret key `msk`.

`KeyGen(msk,  $x$ )`  $\rightarrow$  `sk $_x$` . On input the master secret key `msk` and a key attribute  $x \in \{0, 1\}^N$ , the key generation algorithm outputs a predicate key `sk $_x$` .

`Enc(pp,  $f, m$ )`  $\rightarrow$  `ct`. On input the public parameters `pp`, a predicate  $f \in \mathcal{F}_N$ , and a message  $m \in \mathcal{M}$ , the encryption algorithm outputs a ciphertext `ct`.

`Dec(sk $_x$ , ct)`  $\rightarrow$   `$m$ /fail`. On input a secret key `sk $_x$`  and a ciphertext `ct`, the decryption algorithm either outputs a message  $m$  or a special string `fail` (to denote decryption failure).

*Correctness.* A CP-ABE scheme is said to be correct if for all  $\lambda, N \in \mathbb{N}$ ,  $f \in \mathcal{F}_N$ ,  $m \in \mathcal{M}$ ,  $x \in \{0, 1\}^N$ , the following holds:

$$f(x) = 1 \implies \Pr \left[ \text{Dec}(\text{sk}_x, \text{ct}) = m : \begin{array}{l} (\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^N) \\ \text{sk}_x \leftarrow \text{KeyGen}(\text{msk}, x), \text{ct} \leftarrow \text{Enc}(\text{pp}, x, m) \end{array} \right] = 1.$$

*Security.* For security, we require the scheme to achieve adaptive security (see Definition 3.1). Note that the admissibility condition for the adversary  $\mathcal{A}$  in the security game is modified as follows. The adversary  $\mathcal{A}$  is admissible as long as every secret key query  $x \in \{0, 1\}^N$  made by  $\mathcal{A}$  to the oracle `KeyGen(msk,  $\cdot$ )` satisfies the condition that  $f^*(x) = 0$ , where  $f^*$  is the challenge predicate chosen by  $\mathcal{A}$ .

## 6 Building Adaptively Secure CP-ABE

In this section, we build an adaptively secure CP-ABE scheme from a selectively secure KP-ABE scheme with deletable attributes `DelABE` and a single-key adaptively secure deletion conforming CPRF scheme `DCCPRF`.

## 6.1 Construction

Let  $\text{DelABE} = (\text{DelABE.Setup}, \text{DelABE.KeyGen}, \text{DelABE.Enc}, \text{DelABE.Dec}, \text{DelABE.Delete})$  be a KP-ABE scheme with deletable attributes for predicate class  $\mathcal{C} = \{\mathcal{C}_n\}_{n \in \mathbb{N}}$ , and  $\text{DCCPRF} = (\text{PRF.Setup}, \text{PRF.Constrain}, \text{PRF.Eval}, \text{PRF.CircuitGen}, \text{PRF.DeleteFunc}, \text{PRF.KeySim})$  be a deletion conforming CPRF for constraint class  $\mathcal{F} = \{\mathcal{F}_N\}_{N \in \mathbb{N}}$ . We require the predicate class  $\mathcal{C}$  to be sufficiently expressive such that it contains circuits which perform comparison on top of a circuit generated by the  $\text{PRF.CircuitGen}$  algorithm. The requirement will become evident after the construction.

Below we describe our CP-ABE scheme  $\text{ABE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$  for predicate class  $\mathcal{F} = \{\mathcal{F}_N\}_{N \in \mathbb{N}}$ .

$\text{Setup}(1^\lambda, 1^N) \rightarrow (\text{pp}, \text{msk})$ . The setup algorithm first runs DCCPRF setup to generate the corresponding master secret key:  $\text{prf.msk} \leftarrow \text{PRF.Setup}(1^\lambda, 1^N)$ . Let  $k = k(\lambda, N)$  denote the length of the master secret key  $\text{prf.msk}$ . Next, it runs the deletable ABE setup algorithm  $\text{DelABE.Setup}$  to get deletable ABE public parameters and master secret key as:  $(\text{del.msk}, \text{del.pp}) \leftarrow \text{DelABE.Setup}(1^\lambda, 1^k)$ . It sets public parameters and master key as  $\text{pp} = \text{del.pp}$ ,  $\text{msk} = (\text{prf.msk}, \text{del.msk})$ .

$\text{KeyGen}(\text{msk}, x) \rightarrow \text{sk}$ . Let  $\text{msk} = (\text{prf.msk}, \text{del.msk})$ . The key generation algorithm first computes  $t = \text{PRF.Eval}(\text{prf.msk}, x)$  and generates a circuit  $C_x : \{0, 1\}^k \rightarrow \{0, 1\}^m$  as  $C_x = \text{PRF.CircuitGen}(x)$ . Next, it creates the following circuit  $(f_{x,t} : \{0, 1\}^k \rightarrow \{0, 1\})$

$$f_{x,t}(z) = \begin{cases} 1 & \text{if } C_x(z) \neq t, \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

Finally, the algorithm runs the deletable ABE key generation to sample the secret key  $\text{sk}$  as  $\text{sk} \leftarrow \text{DelABE.KeyGen}(\text{del.msk}, f_{x,t})$ .

$\text{Enc}(\text{pp}, f, m) \rightarrow \text{ct}$ . The encryption algorithm runs the CPRF key simulation to generate a simulated key as  $\text{sk}'_f \leftarrow \text{PRF.KeySim}(1^\lambda, 1^N, f)$ . Next, it runs the deletable ABE encryption algorithm with attribute  $\text{sk}'_f$  as  $\text{ct} \leftarrow \text{DelABE.Enc}(\text{pp}, m, \text{sk}'_f)$ , and outputs ciphertext  $\text{ct}$ .

$\text{Dec}(\text{sk}, \text{ct}) \rightarrow m/\text{fail}$ . The decryption algorithm runs the deletable ABE decryption as  $z = \text{DelABE.Dec}(\text{sk}, \text{ct})$ , and outputs  $z$  as decryption output.

## 6.2 Correctness and Efficiency

We start by proving that our construction satisfies the CP-ABE correctness condition, and also discuss the efficiency of the resulting scheme. First, we prove correctness.

**Lemma 6.1 (Correctness).** *If the deletable KP-ABE scheme  $\text{DelABE}$  satisfies correctness, and the deletion conforming CPRF scheme  $\text{DCCPRF}$  satisfies non-colliding and circuit evaluation properties, then the CP-ABE scheme  $\text{ABE}$  described above is correct.*

*Proof.* We show that the scheme decrypts correctly with all but negligible probability. In the full version, we will discuss how to boost the imperfect correctness to perfect correctness.

Fix any security parameter  $\lambda$  and attribute length  $N$ . For every predicate  $f \in \mathcal{F}_N$ , message  $m \in \mathcal{M}$ , and attribute  $x \in \{0,1\}^N$ , we have that the decryption algorithm  $\text{Dec}$ , on inputs ciphertext  $\text{ct}$  and secret key  $\text{sk}$ , simply outputs  $z = \text{DelABE.Dec}(\text{sk}, \text{ct})$ . Consider  $(\text{del.msk}, \text{del.pp})$  and  $\text{prf.msk}$  to be the deletable KP-ABE and CPRF parameters sampled during setup. Note that the ciphertext  $\text{ct}$  is computed as  $\text{ct} \leftarrow \text{DelABE.Enc}(\text{del.pp}, m, \text{sk}'_f)$ , where  $\text{sk}'_f \leftarrow \text{PRF.KeySim}(1^\lambda, 1^N, f)$ . Also, the secret key  $\text{sk}$  is sampled as  $\text{sk} \leftarrow \text{DelABE.KeyGen}(\text{del.msk}, f_{x,t})$ , where  $t = \text{PRF.Eval}(\text{prf.msk}, x)$  and  $f_{x,t}$  is as defined in the construction. First, observe that by correctness of the deletable KP-ABE scheme, if  $\text{CEval}(f_{x,t}, \text{sk}'_f) = 1$ , then the decryption algorithm outputs message  $m$  correctly, i.e.  $z = m$ . Thus, to complete the completeness argument, we just need to show that whenever  $f(x) = 1$ , then  $\text{CEval}(f_{x,t}, \text{sk}'_f) = 1$  as well with all but negligible probability (over the choice of random coins used during setup and encryption).

Recall that circuit  $f_{x,t}(\text{sk}'_f) = 1$  if and only if  $C_x(\text{sk}'_f) \neq t$ , where  $C_x = \text{PRF.CircuitGen}(x)$ . Now if  $f(x) = 1$ , by the circuit evaluation property of deletion conforming CPRF, we get that  $C_x(\text{sk}'_f) = \text{PRF.Eval}(\text{sk}'_f, x)$ . Since  $t = \text{PRF.Eval}(\text{prf.msk}, x)$ , thus by the non-colliding property, we know that the event  $C_x(\text{sk}'_f) = t$  happens with only negligible probability. Therefore, whenever  $f(x) = 1$ , the decryption algorithm outputs message  $m$  with all but negligible probability. This completes the correctness argument.  $\square$

Next, we state the depth of the circuit  $f_{x,t}$  for which we run the KP-ABE key generation algorithm.

**Lemma 6.2 (Circuit depth).** *For every  $\lambda, N \in \mathbb{N}$ , predicate  $f \in \mathcal{F}_N$  and attribute  $x \in \{0,1\}^N$ , we have that  $\text{depth}(f_{x,t}) = \text{depth}(C_x) + O(\log \lambda)$ .*

*Proof.* This follows immediately from our construction. Note that the circuit depth of  $f_{x,t}$  is depth of  $C_x$  plus the depth of a circuit to check equality on two strings in  $\{0,1\}^m$ . Since  $m$  is a polynomial in the security parameter  $\lambda$ , and equality check on two strings in  $\{0,1\}^m$  can be efficiently performed in depth  $O(\log m) = O(\log \lambda)$  using XOR gates and OR gates, thus the lemma follows.  $\square$

### 6.3 Security

Next, we prove that the CP-ABE scheme constructed above is adaptively secure. Formally, we prove the following.

**Theorem 6.3.** *If the deletion KP-ABE scheme  $\text{DelABE}$  satisfies selective IND-CPA security and deletion indistinguishability (Definitions 3.2 and 3.3), and the deletion conforming CPRF scheme  $\text{DCCPRF}$  satisfies adaptive key simulation security, and circuit evaluation and function deletion properties (Definitions 4.2 and 4.3), then the CP-ABE scheme  $\text{ABE}$  satisfies adaptive IND-CPA security as per Definition 3.1.*

*Proof.* We prove the security via a sequence of hybrid games. We will first define the sequence of hybrid games, and then show that they are indistinguishable for any PPT adversary.

*Game 0.* This corresponds to the original adaptive IND-CPA security game.

- **Setup Phase.** The challenger runs  $\text{prf.msk} \leftarrow \text{PRF.Setup}(1^\lambda, 1^N)$  and  $(\text{del.msk}, \text{del.pp}) \leftarrow \text{DelABE.Setup}(1^\lambda, 1^k)$ . Next, it sets  $\text{pp} = \text{del.pp}$  and  $\text{msk} = (\text{prf.msk}, \text{del.msk})$  and sends  $\text{pp}$  to the adversary  $\mathcal{A}$ .
- **Pre-Challenge Query Phase.** The adversary  $\mathcal{A}$  makes polynomially many key queries on attributes it chooses. For each key query on attribute  $x \in \{0, 1\}^N$ , the challenger proceeds as follows:
  1. It computes  $t = \text{PRF.Eval}(\text{prf.msk}, x)$ , and generates a circuit  $C_x : \{0, 1\}^k \rightarrow \{0, 1\}^m$  as  $C_x = \text{PRF.CircuitGen}(1^\lambda, 1^N, x)$ . Next, it creates a circuit  $f_{x,t}$  as described in Eq. (1).
  2. Then it computes a secret key as  $\text{sk} \leftarrow \text{DelABE.KeyGen}(\text{del.msk}, f_{x,t})$ , and sends  $\text{sk}$  to  $\mathcal{A}$ .
- **Challenge Phase.**  $\mathcal{A}$  sends two messages  $(m_0, m_1)$  and a predicate function  $f^* \in \mathcal{F}_N$  as its challenge to the challenger. The challenger responds with ciphertext  $\text{ct}^*$  to  $\mathcal{A}$ , where  $\text{ct}^*$  is computed as follows:
  1. The challenger generates a simulated key as  $\text{sk}_{f^*} \leftarrow \text{PRF.KeySim}(1^\lambda, 1^N, f^*)$ .
  2. Next, it chooses a random bit  $b \leftarrow \{0, 1\}$ , and computes the challenge ciphertext as  $\text{ct}^* \leftarrow \text{DelABE.Enc}(\text{del.pp}, \text{sk}'_{f^*}, m_b)$ .
- **Post-Challenge Query Phase.** This is identical to the pre-challenge query phase.
- **Guess.** The adversary  $\mathcal{A}$  finally sends the guess  $b'$ , and wins if  $b = b'$ .

*Game 1.* This game is identical to **Game 0** except that in the **Challenge Phase** step 1, the challenger encrypts the challenge ciphertext to a *real constrained PRF key* with respect to challenge function  $f^*$  instead of the simulated key.

- **Challenge Phase.**  $\mathcal{A}$  sends two messages  $(m_0, m_1)$  and a predicate function  $f^* \in \mathcal{F}_N$  as its challenge to the challenger. The challenger responds with ciphertext  $\text{ct}^*$  to  $\mathcal{A}$ , where  $\text{ct}^*$  is computed as follows:
  1. **The challenger generates a constrained key as  $\text{sk}_{f^*} \leftarrow \text{PRF.Constrain}(\text{prf.msk}, f^*)$ .**

*Game 2.* This game is identical to **Game 1** except that in the **Challenge Phase** step 1, the challenger generates the real constrained PRF key  $\text{sk}_{f^*}$  with respect to  $f^*$  directly using the  $\text{PRF.DeleteFunc}$  and  $\text{Restrict}$  algorithms on the PRF master secret key  $\text{prf.msk}$ .

- **Challenge Phase.**  $\mathcal{A}$  sends two messages  $(m_0, m_1)$  and a predicate function  $f^* \in \mathcal{F}_N$  as its challenge to the challenger. The challenger responds with ciphertext  $\text{ct}^*$  to  $\mathcal{A}$ , where  $\text{ct}^*$  is computed as follows:
  1. **The challenger first computes a set of indices  $\mathcal{I}_{f^*} := \text{PRF.DeleteFunc}(1^\lambda, 1^N, f^*)$ , and then it computes the constrained key as  $\text{sk}_{f^*} = \text{Restrict}(\text{prf.msk}, \mathcal{I}_{f^*})$ .**

*Game 3.* This game is identical to **Game 2** except that in the **Challenge Phase** step **2**, the challenger encrypts the message to the *PRF master secret key*  $\text{prf.msk}$  and then uses  $\text{DelABE.Delete}$  to modify the ciphertext according to the indices set  $\mathcal{I}_{f^*}$ .

- **Challenge Phase.**  $\mathcal{A}$  sends two messages  $(m_0, m_1)$  and a predicate function  $f^* \in \mathcal{F}_N$  as its challenge to the challenger. The challenger responds with ciphertext  $\text{ct}^*$  to  $\mathcal{A}$ , where  $\text{ct}^*$  is computed as follows:
  1. The challenger first computes a set of indices  $\mathcal{I}_{f^*} := \text{PRF.DeleteFunc}(1^\lambda, 1^N, f^*)$ .
  2. Next, it chooses a random bit  $b \leftarrow \{0, 1\}$ , and **computes a KP-ABE ciphertext as  $\text{ct}' \leftarrow \text{DelABE.Enc}(\text{del.pp}, \text{prf.msk}, m_b)$ . Then it computes challenge ciphertext as  $\text{ct}^* \leftarrow \text{DelABE.Delete}(\text{del.pp}, \text{ct}', \mathcal{I}_{f^*})$ .**

*Analysis.* Next, we show by a sequence of lemmas that no PPT adversary can distinguish between any two adjacent games with non-negligible advantage. In the last game, we show that the advantage of any PPT adversary is negligible. This completes the proof of adaptive security of our CP-ABE scheme ABE.

Let  $\mathcal{A}$  denote the PPT attacker playing the adaptive IND-CPA security game with the ABE challenger. In the sequel, we denote the advantage of adversary  $\mathcal{A}$  in **Game  $i$**  as  $\text{Adv}_{\mathcal{A}}^i(\lambda) = \Pr[\mathcal{A} \text{ wins in Game } i] - \frac{1}{2}$ , where recall that  $\mathcal{A}$  wins in **Game  $i$**  if it guesses the challenger's bit  $b$  correctly.

**Lemma 6.4.** *Assuming the key simulation security of the deletion conforming CPRF DCCPRF holds, then for any PPT adversary  $\mathcal{A}$ , there exists a negligible function  $\text{negl}_1(\cdot)$ , such that for all  $\lambda, N \in \mathbb{N}$ , we have that  $\text{Adv}_{\mathcal{A}}^0(\lambda) - \text{Adv}_{\mathcal{A}}^1(\lambda) \leq \text{negl}_1(\lambda)$ .*

*Proof.* Suppose there exists an adversary  $\mathcal{A}$  and a non-negligible function  $\epsilon(\cdot)$  such that  $\text{Adv}_{\mathcal{A}}^0(\lambda) - \text{Adv}_{\mathcal{A}}^1(\lambda) \geq \epsilon(\lambda)$ , then we construct a reduction algorithm  $\mathcal{B}$  such that  $\mathcal{B}$  has non-negligible advantage in the *key simulation game* of the deletion conforming CPRF. Below we describe our reduction algorithm  $\mathcal{B}$ .

- In the setup phase, the key simulation challenger  $\mathcal{K}$  runs  $\text{PRF.Setup}$ , and  $\mathcal{B}$  runs  $\text{DelABE.Setup}$  to sample a key pair as  $(\text{del.pp}, \text{del.msk}) \leftarrow \text{DelABE.Setup}(1^\lambda, 1^k)$ .  $\mathcal{B}$  then sends  $\text{del.pp}$  to  $\mathcal{A}$  as the public parameters.
- In the pre-challenge query phase, when  $\mathcal{A}$  sends a key query on attribute  $x$  to  $\mathcal{B}$ ,  $\mathcal{B}$  sends  $x$  to the key simulation challenger  $\mathcal{K}$  as its PRF evaluation query.  $\mathcal{K}$  answers  $\mathcal{B}$  with  $t$ , where  $t = \text{PRF.Eval}(\text{prf.msk}, x)$ .  $\mathcal{B}$  uses  $t$  and  $x$  to generate circuit  $C_x$  and circuit  $f_{x,t}$  as in Game 0; then it computes the secret key  $\text{sk} \leftarrow \text{DelABE.KeyGen}(\text{del.msk}, f_{x,t})$ , and sends  $\text{sk}$  to  $\mathcal{A}$  as the secret key for attribute  $x$ .
- In the challenge phase,  $\mathcal{A}$  sends the predicate function  $f^*$  and messages  $m_0, m_1$  to the reduction algorithm  $\mathcal{B}$ .  $\mathcal{B}$  then forwards  $f^*$  to  $\mathcal{K}$  as its challenge constraint function. Let  $\text{sk}_{f^*}$  denote  $\mathcal{K}$ 's response.  $\mathcal{B}$  chooses a random bit  $b \leftarrow \{0, 1\}$ , and computes the challenge ciphertext as  $\text{ct}^* \leftarrow \text{DelABE.Enc}(\text{del.pp}, \text{sk}_{f^*}, m_b)$ , and sends  $\text{ct}^*$  to  $\mathcal{A}$ .

- The post-challenge phase is identical to the pre-challenge query phase. Finally,  $\mathcal{A}$  outputs its guess  $b'$ , and if  $b = b'$  then  $\mathcal{B}$  outputs 0 as its guess (to denote that  $\text{sk}_{f^*}$  was a simulated key). Otherwise,  $\mathcal{B}$  outputs 1 as its guess.

First, note that  $\mathcal{A}$  must be an admissible adversary in the CP-ABE security game, thus it must hold that  $f^*(x) = 0$  for all attributes  $x$  queried by  $\mathcal{A}$ . Therefore,  $\mathcal{B}$  is also an admissible adversary in the key simulation game since it also satisfies condition that  $f^*(x) = 0$  for all inputs  $x$  queried by  $\mathcal{B}$ . Next, observe that if the challenger  $\mathcal{K}$  samples  $\text{sk}_{f^*}$  as a simulated key, then  $\mathcal{B}$  perfectly simulates Game 0 for  $\mathcal{A}$ , otherwise it simulates Game 1. Thus,  $\mathcal{B}$ 's advantage in the key simulation game is at least  $\epsilon(\lambda)$ , which is non-negligible and contradicts the key simulation security.  $\square$

**Lemma 6.5.** *Assuming the function deletion property of the deletion conforming CPRF DCCPRF holds, then for any adversary  $\mathcal{A}$ , parameters  $\lambda, N \in \mathbb{N}$ , we have that  $\text{Adv}_{\mathcal{A}}^1(\lambda) = \text{Adv}_{\mathcal{A}}^2(\lambda)$ .*

*Proof.* This follows immediately from the function deletion property. Recall that function deletion property states that for all  $\lambda, N \in \mathbb{N}$ , every constraint function  $f^* \in \mathcal{F}_N$ , and master key  $\text{prf.msk} \leftarrow \text{PRF.Setup}(1^\lambda, 1^N)$ , we have that:

$$\Pr \left[ \begin{array}{l} \mathcal{I}_{f^*} = \text{PRF.DeleteFunc}(1^\lambda, 1^N, f^*) \\ \text{sk}_{f^*}^{(1)} = \text{sk}_{f^*}^{(2)} : \text{sk}_{f^*}^{(1)} = \text{PRF.Constrain}(\text{prf.msk}, f^*) \\ \text{sk}_{f^*}^{(2)} = \text{Restrict}(\text{prf.msk}, \mathcal{I}_{f^*}) \end{array} \right] = 1.$$

Note that  $\text{sk}_{f^*}^{(1)}$  and  $\text{sk}_{f^*}^{(2)}$  exactly correspond to the CPRF keys as generated in **Game 1** and **Game 2**, respectively. Since they are identical, thus the adversary's advantage is also identical in these two games.  $\square$

**Lemma 6.6.** *Assuming the deletion indistinguishability security of the deletable KP-ABE DelABE holds, then for any PPT adversary  $\mathcal{A}$ , there exists a negligible function  $\text{negl}_2(\cdot)$ , such that for all  $\lambda, N \in \mathbb{N}$ , we have that  $\text{Adv}_{\mathcal{A}}^2(\lambda) - \text{Adv}_{\mathcal{A}}^3(\lambda) \leq \text{negl}_2(\lambda)$ .*

*Proof.* Suppose there exists an adversary  $\mathcal{A}$  and a non-negligible function  $\epsilon(\cdot)$  such that  $\text{Adv}_{\mathcal{A}}^2(\lambda) - \text{Adv}_{\mathcal{A}}^3(\lambda) \geq \epsilon(\lambda)$ , then we construct a reduction algorithm  $\mathcal{B}$  such that  $\mathcal{B}$  has non-negligible advantage in the *deletion indistinguishability game* of the deletable KP-ABE. Below we describe our reduction algorithm  $\mathcal{B}$ .

- In the setup phase, the deletion indistinguishability challenger  $\mathcal{D}$  runs  $\text{DelABE.Setup}$  and sends the deletable ABE parameters ( $\text{del.pp}, \text{del.msk}$ ) to  $\mathcal{B}$ .  $\mathcal{B}$  then samples a CPRF master key as  $\text{prf.msk} \leftarrow \text{PRF.Setup}(1^\lambda, 1^N)$ , and sends  $\text{del.pp}$  to  $\mathcal{A}$  as the CP-ABE public parameters.
- In the pre-challenge query phase,  $\mathcal{A}$  sends a key query on attribute  $x$  to  $\mathcal{B}$ .  $\mathcal{B}$  first evaluates the CPRF as  $t = \text{PRF.Eval}(\text{prf.msk}, x)$ , and uses  $t$  and  $x$  to generate circuits  $C_x$  and  $f_{x,t}$  as in Game 2. It then computes the secret key  $\text{sk} \leftarrow \text{DelABE.KeyGen}(\text{del.msk}, f_{x,t})$ , and sends  $\text{sk}$  to  $\mathcal{A}$  as the secret key for attribute  $x$ .



- In the challenge phase,  $\mathcal{A}$  sends the predicate function  $f^*$  and messages  $m_0, m_1$  to  $\mathcal{B}$ . The reduction algorithm  $\mathcal{B}$  samples a random bit  $b \leftarrow \{0, 1\}$ , and computes a set of indices  $\mathcal{I}_{f^*} = \text{PRF.DeleteFunc}(1^\lambda, 1^N, f^*)$ , and sends  $(m_b, \text{prf.msk}, \mathcal{I}_{f^*})$  to the deletion challenger  $\mathcal{D}$ . Let  $\text{ct}^*$  denote the challenger's response.  $\mathcal{B}$  forwards  $\text{ct}^*$  to  $\mathcal{A}$  as its challenge ciphertext.
- The post-challenge phase is identical to the pre-challenge query phase. Finally,  $\mathcal{A}$  outputs its guess  $b'$ , and if  $b = b'$  then  $\mathcal{B}$  outputs 0 as its guess (to denote that  $\text{ct}^*$  was a freshly encrypted ciphertext). Otherwise,  $\mathcal{B}$  outputs 1 as its guess.

Note that if the challenger  $\mathcal{D}$  computes  $\text{ct}^*$  by first restricting the attribute to the constrained key and then encrypting it directly using the KP-ABE encryption algorithm, then  $\mathcal{B}$  perfectly simulates Game 2 for  $\mathcal{A}$ , otherwise it simulates Game 3. Thus,  $\mathcal{B}$ 's advantage in the deletion indistinguishability game is at least  $\epsilon(\lambda)$ , which is non-negligible and contradicts the deletion indistinguishability security.  $\square$

**Lemma 6.7.** *Assuming the selective IND-CPA security of the deletable KP-ABE DelABE holds and the deletion conforming CPRF DCCPRF satisfies circuit evaluation property, then for any PPT adversary  $\mathcal{A}$ , there exists a negligible function  $\text{negl}_3(\cdot)$ , such that for all  $\lambda, N \in \mathbb{N}$ , we have that  $\text{Adv}_{\mathcal{A}}^3(\lambda) \leq \text{negl}_3(\lambda)$ .*

*Proof.* Suppose there exists an adversary  $\mathcal{A}$  and a non-negligible function  $\epsilon(\cdot)$  such that  $\text{Adv}_{\mathcal{A}}^3(\lambda) \geq \epsilon(\lambda)$ , then we construct a reduction algorithm  $\mathcal{B}$  such that  $\mathcal{B}$  has non-negligible advantage in the *selective IND-CPA game* of the deletable KP-ABE. Below we describe our reduction algorithm  $\mathcal{B}$ .

- In the setup phase,  $\mathcal{B}$  first samples a CPRF master key as  $\text{prf.msk} \leftarrow \text{PRF.Setup}(1^\lambda, 1^N)$ , and sends  $\text{prf.msk}$  as its challenge attribute to the selective IND-CPA challenger  $\mathcal{D}$ . The challenger runs  $\text{DelABE.Setup}$  and sends the deletable public parameters  $\text{del.pp}$  to  $\mathcal{B}$ .  $\mathcal{B}$  simply forwards  $\text{del.pp}$  to  $\mathcal{A}$  as the CP-ABE public parameters.
- In the pre-challenge query phase, when  $\mathcal{A}$  sends a key query on attribute  $x$  to  $\mathcal{B}$ ,  $\mathcal{B}$  computes  $t_x = \text{PRF.Eval}(\text{prf.msk}, x)$  and generates the circuit  $f_{x, t_x}$  using  $x$  and  $t_x$ . Next,  $\mathcal{B}$  sends secret key query on predicate  $f_{x, t_x}$  to the challenger  $\mathcal{D}$ .  $\mathcal{D}$  replies  $\mathcal{B}$ 's query with  $\text{sk}$  and  $\mathcal{B}$  forwards  $\text{sk}$  to  $\mathcal{A}$  as the secret key for attribute  $x$ .
- In the challenge phase,  $\mathcal{A}$  sends the predicate function  $f^*$  and messages  $m_0, m_1$  to  $\mathcal{B}$ .  $\mathcal{B}$  sends  $(m_0, m_1)$  to  $\mathcal{D}$ . Let  $\text{ct}'$  denote the KP-ABE challenge ciphertext sent by  $\mathcal{D}$ .  $\mathcal{B}$  first computes the index set  $\mathcal{I}_{f^*} = \text{PRF.DeleteFunc}(1^\lambda, 1^N, f^*)$ , and then computes challenge ciphertext as  $\text{ct}^* \leftarrow \text{DelABE.Delete}(\text{del.pp}, \text{ct}', \mathcal{I}_{f^*})$ .  $\mathcal{B}$  sends  $\text{ct}^*$  to  $\mathcal{A}$  as its challenge ciphertext.
- The post-challenge phase is identical to the pre-challenge query phase. Finally,  $\mathcal{A}$  outputs its guess  $b'$ , and  $\mathcal{B}$  outputs the same bit  $b'$  as its guess.

First, note that for each key query on attribute  $x$  made by  $\mathcal{A}$ , we have that  $C_x(\text{prf.msk}) = \text{PRF.Eval}(\text{prf.msk}, x)$ . This follows from the circuit evaluation

property of the deletion conforming CPRF. Since  $t_x = \text{PRF.Eval}(\text{prf.msk}, x)$ , thus by definition of the circuit  $f_{x,t_x}$  (see Eq. (1)), we have that  $f_{x,t_x}(\text{prf.msk}) = 0$  for every attribute  $x$ . Thus, the reduction algorithm  $\mathcal{B}$  is an admissible adversary in the selective IND-CPA game. Next, observe that  $\mathcal{B}$  perfectly simulates Game 3 for  $\mathcal{A}$ , therefore  $\mathcal{B}$ 's advantage in the selective IND-CPA game is at least  $\epsilon(\lambda)$ , which is non-negligible and contradicts the selective IND-CPA security of the deletable KP-ABE system.  $\square$

Combining Lemmas 6.4 to 6.7, the Theorem 6.3 follows.  $\square$

## 7 Deletable ABE from standard assumptions

In this section we show that [19] is already a KP-ABE scheme with deletable attributes. First, we show that the KP-ABE schemes for monotone access structures in [19] have efficient deletion algorithms such that the resulting scheme satisfies both the semantic security as well deletion indistinguishability properties. Later on, we briefly elaborate the well-known approach for building a KP-ABE scheme for  $\text{NC}^1$  (i.e., log-depth circuits) from any KP-ABE scheme for monotone access structures, and describe that it preserves the deletion property of the underlying system.

### 7.1 Deletable ABE from Bilinear Maps via [19]

Goyal et al. (GPSW) [19] proposed a KP-ABE scheme for monotone access structures and proved its security under the Decisional Bilinear Diffie-Hellman (DBDH) assumption [3]. Here we show that the GPSW scheme, described in [19, Section 4], is also a deletable KP-ABE scheme for the same predicate class. Let  $\text{GPSW} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$  represent the KP-ABE construction provided in [19, Section 4]. Formally, they proved the following.

**Theorem 7.1 ([19, Theorem 1, Paraphrased]).** *If the Decisional Bilinear Diffie-Hellman (DBDH) assumption holds, then the scheme GPSW is a selective IND-CPA secure scheme as per Definition 3.2.*

Now we describe a simple deletion algorithm for the GPSW scheme, and argue that the augmented GPSW scheme satisfies all the required properties described in Section 3. We start by briefly discussing some notational changes that we make to the GPSW syntax.

*Notation.* For consistency with our ABE definitions, we interpret the attribute string as a bit string  $x \in \{0, 1\}^n$ , where as is the GPSW construction [19, Section 4] the attribute was parsed as a set of subset of the attribute universe  $\mathcal{U} = \{1, 2, \dots, n\}$ . Here  $n$  denotes the length of the attributes selected during system setup. Note that this is mostly a syntactic change, and does not affect the GPSW scheme in any significant way.

Below we recall the **Setup** and **Enc** algorithms as provided in [19, Section 4], and also describe our **Delete** algorithm.

**Setup**( $1^\lambda, 1^n$ )  $\rightarrow$  (**pp**, **msk**). The setup algorithm chooses a bilinear group  $\mathbb{G}_1$  of prime order  $p$ . Let  $g$  denote the generator of the group  $\mathbb{G}_1$ , and  $e : \mathbb{G}_1 \times \mathbb{G}_1 \rightarrow \mathbb{G}_2$  be associated the bilinear map. It chooses a random key exponent  $\alpha \in \mathbb{Z}_p$ , and also chooses a random exponent per bit position of the attribute, that is  $t_i \leftarrow \mathbb{Z}_p$  for  $i \in [n]$ .

It outputs the public parameters and master secret key as **pp** =  $(g, e(g, g)^\alpha, \{g^{t_i}\}_{i \in [n]})$  and **msk** =  $(\alpha, \{t_i\}_{i \in [n]})$ .<sup>17</sup>

**Enc**(**pp**,  $x, m$ )  $\rightarrow$  **ct**. The encryption algorithm parses the public parameters as **pp** =  $(g, K, \{T_i\}_{i \in [n]})$ , and an attribute  $x \in \{0, 1\}^n$ . It chooses a random exponent  $s \in \mathbb{Z}_p$ , and publishes the ciphertext as

$$\mathbf{ct} = (x, m \cdot K^s, \{T_i^s\}_{i \in [n]: x_i=1}).$$

*Encrypting to attributes with  $\perp$  symbols.* First, note that in the GPSW encryption algorithm the input attribute string  $x$  is a binary string, that is  $x \in \{0, 1\}^n$ . However, in our deletable ABE framework, we allow the encryptor to choose attributes with  $\perp$  symbols, thus the attribute string  $x$  now lies in  $\{0, 1, \perp\}^n$  instead of  $\{0, 1\}^n$ . Now our augmented encryption algorithm is identical to the above encryption algorithm, that is the ciphertext is computed as

$$\mathbf{ct} = (x, m \cdot K^s, \{T_i^s\}_{i \in [n]: x_i=1}).$$

Note that previously the algorithm does not compute  $T_i^s$  for all  $i$  wherever  $x_i = 0$ . Now the augmented encryption algorithm also does not compute  $T_i^s$  for all  $i$  wherever  $x_i = \perp$ . That is, it treats  $\perp$  symbols as a 0 bit during encryption. Therefore, the deletion algorithm can simply *delete* the  $T_i^s$  terms from the ciphertext wherever  $i \in \mathcal{I}$  to compute a corresponding deleted ciphertext. Formally, we describe it below.

**Delete**(**pp**, **ct**,  $\mathcal{I}$ )  $\rightarrow$  **ct'**. The algorithm parses the ciphertext as **ct** =  $(x, E', \{E_i\}_{i \in [n]: x_i=1})$ . It sets the output ciphertext **ct'** as

$$\mathbf{ct}' = (\text{Restrict}(x, \mathcal{I}), E', \{E_i\}_{i \in [n] \setminus \mathcal{I}: x_i=1}).$$

*Deletion Indistinguishability.* First, we show that the augmented GPSW scheme **AugGPSW** = (**Setup**, **KeyGen**, **Enc**, **Dec**, **Delete**) satisfies the deletion indistinguishability property. Below we prove a much stronger statement which in turn implies deletion indistinguishability. Intuitively, we argue that, for every choice of system parameters, the distribution of a freshly encrypted ciphertext and a (corresponding) deleted ciphertext are *identical*.

<sup>17</sup>The parameters also contain the bilinear map parameters, but here we don't explicitly write it for simplicity.

**Lemma 7.2.** For every  $\lambda, n \in \mathbb{N}$ , parameters  $(\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^n)$ , attribute  $x \in \{0, 1\}^n$ , message  $m \in \mathcal{M}$ , and index set  $\mathcal{I} \in [n]$ , the following two distributions are identical:

$$\mathcal{D}_1 = \left\{ \text{ct} : \begin{array}{l} x' = \text{Restrict}(x, \mathcal{I}) \\ \text{ct} \leftarrow \text{Enc}(\text{pp}, x', m) \end{array} \right\}, \quad \mathcal{D}_2 = \left\{ \text{ct}' : \begin{array}{l} \text{ct} \leftarrow \text{Enc}(\text{pp}, x, m) \\ \text{ct}' \leftarrow \text{Delete}(\text{pp}, \text{ct}, \mathcal{I}) \end{array} \right\}.$$

That is,  $\mathcal{D}_1 \equiv \mathcal{D}_2$ .

*Proof.* The proof of this lemma immediately follows by inspection of the encryption and deletion algorithms described above. Consider any  $\lambda, n$ , key pair  $(\text{pp}, \text{msk})$ , attribute  $x$ , message  $m$  and index set  $\mathcal{I}$ . First, note that the distributions  $\mathcal{D}_1$  and  $\mathcal{D}_2$  can be expanded as follows:

$$\begin{aligned} \mathcal{D}_1 &= \{(x', m \cdot K^s, \{T_i^s\}_{i \in S_1}) : x' = \text{Restrict}(x, \mathcal{I}), s \leftarrow \mathbb{Z}_p, S_1 = \{i \in [n] : x'_i = 1\}\}, \\ \mathcal{D}_2 &= \{(x', m \cdot K^s, \{T_i^s\}_{i \in S_2}) : x' = \text{Restrict}(x, \mathcal{I}), s \leftarrow \mathbb{Z}_p, S_2 = \{i \in [n] \setminus \mathcal{I} : x_i = 1\}\}. \end{aligned}$$

Recall by definition of `Restrict`, we have that  $x'_i = 1$  if and only if  $x_i = 1$  and  $i \notin \mathcal{I}$ . Therefore, it follows that  $\mathcal{D}_1 \equiv \mathcal{D}_2$ .  $\square$

*Correctness.* Note that since a deleted ciphertext is *identically* distributed to a freshly encrypted ciphertext, and also `GPSW` is a perfectly correct ABE scheme, thus correctness of our `AugGPSW` scheme follows.

*Selective IND-CPA Security.* Note that even though in our scheme, attribute vectors could contain  $\perp$  symbols (i.e., lie in  $\{0, 1, \perp\}^n$ ), the IND-CPA attacker is only allowed to specify a binary string as a challenge attribute (i.e., it must lie in  $\{0, 1\}^n$ ) in the selective security game (Definition 3.2). Therefore, the selective IND-CPA security proof of `AugGPSW` follows from selective IND-CPA security proof of `GPSW`.

Hence, combining above facts, Lemma 7.2 and Theorem 7.1, we obtain the following:

**Theorem 7.3.** *If the Decisional Bilinear Diffie-Hellman (DBDH) assumption holds, then the scheme `AugGPSW` is a KP-ABE scheme with deletable attributes that satisfies selective IND-CPA security as well as deletion indistinguishability (Definitions 3.2 and 3.3).*

Also, later on in the full version we describe how to get deletable ABE from Computational Bilinear Diffie-Hellman (CBDH) assumption. It follows from a straightforward use of hardcore predicate on top of the `GPSW` scheme.

## 7.2 Deletable ABE: Monotonic Access Structures to $\text{NC}^1$

Suppose we start with a KP-ABE scheme for arbitrary polynomial-sized monotone boolean formulas, then there is a well-known folklore transformation that gives us a KP-ABE scheme for log-depth circuits ( $\text{NC}^1$ ) generically from the

underlying scheme. The idea can be described as follows. First, the key generation algorithm, on input a log-depth (non-monotone) circuit  $C$ , generates a polynomial-sized (non-monotone) boolean formula  $f_C$  that evaluates the same circuit. (Note that size of the formula  $f_C$  grows exponentially with the depth of circuit  $C$ , thus the same transformation does not work for larger depth circuits.) Now the formula  $f_C$  is a possibly *non-monotone* boolean formula, thus it could apply negation ( $\neg$ ) gates on non-atomic formulae. Next, one using De Morgan’s identities can translate the non-monotone boolean formula  $f_C$  into another formula  $\tilde{f}_C$  such that in the description of formula  $\tilde{f}_C$ , negation gates are only applied on input wires. In other words, formula  $\tilde{f}_C$  can alternatively be interpreted as a monotone boolean formula being applied on the literals. (Recall that a literal is an atomic formula or its negation, i.e. either an input wire or its negation). With this observation, one could use KP-ABE scheme for monotone boolean formulas to obtain a KP-ABE scheme for  $\mathbf{NC}^1$ .

A more concrete description of above transformation is provided later in the full version.

## References

1. Agrawal, S., Biswas, R., Nishimaki, R., Xagawa, K., Xie, X., Yamada, S.: Attacks on boyen’s attribute-based encryption scheme in tcc 2013. Personal communication (2020)
2. Attrapadung, N.: Dual system encryption via doubly selective security: Framework, fully secure functional encryption for regular languages, and more. In: EUROCRYPT (2014)
3. Boneh, D., Franklin, M.K.: Identity-based encryption from the Weil Pairing. In: CRYPTO ’01 (2001)
4. Boneh, D., Gentry, C., Gorbunov, S., Halevi, S., Nikolaenko, V., Segev, G., Vaikuntanathan, V., Vinayagamurthy, D.: Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In: EUROCRYPT (2014)
5. Boneh, D., Waters, B.: Constrained pseudorandom functions and their applications. In: ASIACRYPT 2013 (2013)
6. Boyen, X.: Attribute-based functional encryption on lattices. In: TCC (2013)
7. Boyle, E., Goldwasser, S., Ivan, I.: Functional signatures and pseudorandom functions. In: PKC 2014 (2014)
8. Brakerski, Z., Vaikuntanathan, V.: Circuit-abe from lwe: unbounded attributes and semi-adaptive security. In: Crypto (2016)
9. Davidson, A., Katsumata, S., Nishimaki, R., Yamada, S.: Constrained prfs for bit-fixing (and more) from owfs with adaptive security and constant collusion resistance. Cryptology ePrint Archive, Report 2018/982 (2018)
10. Davidson, A., Katsumata, S., Nishimaki, R., Yamada, S., Yamakawa, T.: Adaptively secure constrained pseudorandom functions in the standard model. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020 (2020)
11. Diffie, W., Hellman, M.E.: New directions in cryptography (1976)
12. Gentry, C.: Practical identity-based encryption without random oracles. In: EUROCRYPT ’06. vol. 4004 of LNCS, pp. 445–464 (2006)
13. Gentry, C., Halevi, S.: Hierarchical identity based encryption with polynomially many levels. In: TCC (2009)

14. Goldwasser, S., Micali, S.: Probabilistic encryption. *J. Comput. Syst. Sci.* 28(2), 270–299 (1984)
15. Gorbunov, S., Vaikuntanathan, V., Wee, H.: Attribute-based encryption for circuits. In: *STOC* (2013)
16. Gorbunov, S., Vaikuntanathan, V., Wee, H.: Predicate encryption for circuits from lwe. In: *Annual Cryptology Conference* (2015)
17. Goyal, R., Koppula, V., Waters, B.: Semi-adaptive security and bundling functionalities made generic and easy. In: *TCC 2016-B* (2016)
18. Goyal, R., Liu, J., Waters, B.: Adaptive security via deletion in attribute-based encryption: Solutions from search assumptions in bilinear groups. *Cryptology ePrint Archive, Report 2021/343* (2021), <https://ia.cr/2021/343>
19. Goyal, V., Pandey, O., Sahai, A., Waters, B.: Attribute-based encryption for fine-grained access control of encrypted data. In: *CCS '06* (2006)
20. Katsumata, S., Nishimaki, R., Yamada, S., Yamakawa, T.: Adaptively secure inner product encryption from LWE. In: Moriai, S., Wang, H. (eds.) *ASIACRYPT 2020* (2020)
21. Kiayias, A., Papadopoulos, S., Triandopoulos, N., Zacharias, T.: Delegatable pseudorandom functions and applications. In: *CCS* (2013)
22. Lewko, A.B., Okamoto, T., Sahai, A., Takashima, K., Waters, B.: Fully secure functional encryption: Attribute-based encryption and (hierarchical) inner product encryption. In: *EUROCRYPT*. pp. 62–91 (2010)
23. Lewko, A.B., Waters, B.: Decentralizing attribute-based encryption. In: *EUROCRYPT*. pp. 568–588 (2011)
24. Lewko, A.B., Waters, B.: New proof methods for attribute-based encryption: Achieving full security through selective techniques. In: *CRYPTO* (2012)
25. Lewko, A.B., Waters, B.: Why proving HIBE systems secure is difficult. In: *EUROCRYPT* (2014)
26. Okamoto, T., Takashima, K.: Fully secure functional encryption with general relations from the decisional linear assumption. In: *CRYPTO*. pp. 191–208 (2010)
27. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: *STOC* (2005)
28. Rivest, R.L., Shamir, A., Adleman, L.M.: A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM* 21(2), 120–126 (1978)
29. Sahai, A., Waters, B.: Fuzzy identity-based encryption. In: *EUROCRYPT*. pp. 457–473 (2005)
30. Tsabary, R.: Fully secure attribute-based encryption for t-cnf from LWE. In: *CRYPTO* (2019)
31. Waters, B.: Dual system encryption: Realizing fully secure ibe and hibe under simple assumptions. In: *CRYPTO*. pp. 619–636 (2009)
32. Wee, H.: Dual system encryption via predicate encodings. In: *TCC* (2014)