

Compact Adaptively Secure ABE for NC^1 from k -Lin

Lucas Kowalczyk^{1,*} and Hoeteck Wee^{2,**}

¹ Columbia University
luke@cs.columbia.edu

² CNRS, ENS, PSL
wee@di.ens.fr

Abstract. We present compact attribute-based encryption (ABE) schemes for NC^1 that are adaptively secure under the k -Lin assumption with polynomial security loss. Our KP-ABE scheme achieves ciphertext size that is linear in the attribute length and independent of the policy size even in the many-use setting, and we achieve an analogous efficiency guarantee for CP-ABE. This resolves the central open problem posed by Lewko and Waters (CRYPTO 2011). Previous adaptively secure constructions either impose an attribute “one-use restriction” (or the ciphertext size grows with the policy size), or require q -type assumptions.

1 Introduction

Attribute-based encryption (ABE) [SW05,GPSW06] is a generalization of public-key encryption to support fine-grained access control for encrypted data. Here, ciphertexts and keys are associated with descriptive values which determine whether decryption is possible. In a key-policy ABE (KP-ABE) scheme for instance, ciphertexts are associated with attributes like ‘(author:Waters), (inst:UT), (topic:PK)’ and keys with access policies like ((topic:MPC) OR (topic:SK)) AND (NOT(inst:UCL)), and decryption is possible only when the attributes satisfy the access policy. A ciphertext-policy (CP-ABE) scheme is the dual of KP-ABE with ciphertexts associated with policies and keys with attributes.

Over past decade, substantial progress has been made in the design and analysis of ABE schemes, leading to a large families of schemes that achieve various trade-offs between efficiency, security and underlying assumptions. Meanwhile, ABE has found use as a tool for providing and enhancing privacy in a variety of settings from electronic medical records to messaging systems and online social networks. Moreover, we expect further deployment of ABE, thanks to the recent standardization efforts of the European Telecommunications Standards Institute (ETSI).³

In this work, we consider KP-ABE schemes for access policies in NC^1 that simultaneously:

- (1) enjoy compact ciphertexts whose size grows only with the length of the attribute and is independent of the policy size, even for complex policies that refer to each attribute many times;
- (2) achieve adaptive security (with polynomial security loss);
- (3) rely on simple hardness assumptions in the standard model;
- (4) can be built with asymmetric prime-order bilinear groups.

* Supported in part by an NSF Graduate Research Fellowship DGE-16-44869; The Leona M. & Harry B. Helmsley Charitable Trust; ERC Project aSCEND (H2020 639554); the Defense Advanced Research Project Agency (DARPA) and Army Research Office (ARO) under Contract W911NF-15-C-0236; and NSF grants CNS-1445424, CNS-1552932 and CCF-1423306. Any opinions, findings and conclusions or recommendations expressed are those of the authors and do not necessarily reflect the views of the the Defense Advanced Research Projects Agency, Army Research Office, the National Science Foundation, or the U.S. Government.

** Supported in part by ERC Project aSCEND (H2020 639554).

³ <https://www.etsi.org/news-events/news/1328-2018-08-press-etsi-releases-cryptographic-standards-for-secure-access-control>

reference	adaptive compact assumption unbounded			
GPSW [GPSW06]		✓	static ✓	
LOSTW [LOS ⁺ 10,OT10]	✓		static ✓	
LW [LW12]	✓	✓	q -type	
OT [OT12]	✓		2-Lin ✓	✓
Att [Att16]	✓	✓	q -type	✓
CGKW [CGKW18]	✓		k -Lin ✓	✓
ours, Section 6	✓	✓	static ✓	
ours, Section B	✓	✓	static ✓	✓

Fig. 1. Summary of KP-ABE schemes for NC1

We also consider the analogous question for CP-ABE schemes with compact keys. In both KP and CP-ABE, all four properties are highly desirable from both a practical and theoretical standpoint and moreover, properties (1), (2) and (4) are crucial for many real-world applications of ABE. In addition, properties (2), (3) and (4) are by now standard cryptographic requirements pertaining to speed and efficiency, strong security guarantees under realistic and natural attack models, and minimal hardness assumptions. There is now a vast body of works on ABE (e.g. [GPSW06,LOS⁺10,OT10,LW12], see Fig 1) showing how different combinations of (1) – (4), culminating in several unifying frameworks that provide a solid understanding of the design and analysis of these schemes [Att14,Wee14,CGW15,Att16,AC17]. Nonetheless, prior to this work, it was not known how to even simultaneously realize (1) – (3) for NC¹ access policies⁴; indeed, this is widely regarded one of the main open problems in pairing-based ABE.

Our results. We present the first KP-ABE and CP-ABE schemes for NC¹ that simultaneously realize properties (1) – (4). Our KP-ABE scheme achieves ciphertext size that is linear in the attribute length and independent of the policy size even in the many-use setting; the same holds for the key size in our CP-ABE. Both schemes achieve adaptive security under the k -Lin assumption in asymmetric prime-order bilinear groups with polynomial security loss. We also present an “unbounded” variant of our compact KP-ABE scheme with constant-size public parameters.

As an immediate corollary, we obtain delegation schemes for NC¹ with public verifiability and adaptive soundness under the k -Lin assumption [PRV12,LW12,CW14].

Our construction leverages a refinement of the recent “partial selectivization” framework for adaptive security [JKK⁺17] (which in turn builds upon [FKPR14,FJP15,HJO⁺16,JW16]) along with the classic dual system encryption methodology [Wat09,LW12].

1.1 Technical overview

Our starting point is the Lewko-Waters framework for constructing compact adaptively secure ABE [LW12] based on the dual system encryption methodology⁵ [Wat09,LW10,LOS⁺10]. Throughout, we focus on monotone NC¹ circuit access policies, and note that the constructions extend readily

⁴ Note that there exist constructions of ABE for more general access policies like monotone span programs / Boolean formulas with threshold gates [GPSW06], and even polynomial-sized Boolean circuits [GVW13,GGH⁺13], but all such constructions sacrifice at least one of the properties (1)-(3).

⁵ Essentially, the dual system proof method provides guidance for transforming suitably-designed functional encryption schemes which are secure for one adversarial secret key request to the multi-key setting where multiple keys may be requested by the adversary. Our main technical contribution involves the analysis of the initial single-key-secure component, which we refer to later as our “Core 1-ABE” component.

to the non-monotone setting⁶. Let (G_1, G_2, G_T) be an asymmetric bilinear group of prime order p , where g, h are generators of G_1, G_2 respectively.

Warm-up. We begin with the prior compact KP-ABE for monotone NC¹ [LW12, LOS⁺10, GPSW06]; this is an adaptively secure scheme that comes with the downside of relying on q -type assumptions (q -type assumptions are assumptions of size that grows with some parameter q . It is known that many q -type assumptions become stronger as q grows [Che06], and in general such complex and dynamic assumptions are not well-understood). The construction uses composite-order groups, but here we'll suppress the distinction between composite-order and prime-order groups for simplicity. We associate ciphertexts $\text{ct}_{\mathbf{x}}$ with attribute vectors⁷ $\mathbf{x} \in \{0, 1\}^n$ and keys sk_f with Boolean formulas f :

$$\begin{aligned} \text{msk} &:= (\mu, w_1, \dots, w_n) \\ \text{mpk} &:= (g, g^{w_1}, \dots, g^{w_n}, e(g, h)^\mu), \\ \text{ct}_{\mathbf{x}} &:= (g^s, \{g^{s w_i}\}_{x_i=1}, e(g, h)^{\mu s} \cdot M) \\ \text{sk}_f &:= (\{h^{\mu_j + r_j w_{\rho(j)}}\}_{j \in [m]}, \{h^{r_j}\}_{j \in [m]}), \rho : [m] \rightarrow [n] \end{aligned} \tag{1}$$

where μ_1, \dots, μ_m are shares of $\mu \in \mathbb{Z}_p$ w.r.t. the formula f ; the shares satisfy the requirement that for any $\mathbf{x} \in \{0, 1\}^n$, the shares $\{\mu_j\}_{x_{\rho(j)}=1}$ determine μ if \mathbf{x} satisfies f (i.e., $f(\mathbf{x}) = 1$), and reveal nothing about μ otherwise; and ρ is a mapping from the indices of the shares (in $[m]$) to the indices of the attributes (in $[n]$) to which they are associated. For decryption, observe that we can compute $\{e(g, h)^{\mu_j s}\}_{x_i=1}$, from which we can compute the blinding factor $e(g, h)^{\mu s}$ via linear reconstruction “in the exponent”.

Here, m is polynomial in the formula size, and we should think of $m = \text{poly}(n) \gg n$. Note that the ciphertext consists only of $O(n)$ group elements and therefore satisfies our compactness requirement.

Proving adaptive security. The crux of the proof of adaptive security lies in proving that μ remains computationally hidden given just a single ciphertext and a single key and no mpk (the more general setting with mpk and multiple keys follows via what is by now a textbook application of the dual system encryption methodology). In fact, it suffices to show that μ is hidden given just

$$\begin{aligned} \text{ct}'_{\mathbf{x}} &:= (\{w_i\}_{x_i=1}) \quad // \text{“stripped down” } \text{ct}_{\mathbf{x}} \\ \text{sk}_f &:= (\{h^{\mu_j + r_j w_{\rho(j)}}\}_{j \in [m]}, \{h^{r_j}\}_{j \in [m]}) \end{aligned}$$

where \mathbf{x}, f are adaptively chosen subject to the constraint $f(\mathbf{x}) = 0$. Henceforth, we refer to $(\text{ct}'_{\mathbf{x}}, \text{sk}_f)$ as our “core 1-ABE component”. Looking ahead to our formalization of adaptive security for this core 1-ABE, we actually require that μ is hidden even if the adversary sees h^{w_1}, \dots, h^{w_n} ; this turns out to be useful for the proof of our KP-ABE (for improved concrete efficiency).

Core technical contribution. The technical novelty of this work lies in proving adaptive security of the core 1-ABE component under the DDH assumption. Previous analysis either relies on a

⁶ Most directly by pushing all NOT gates to the input nodes of each circuit and using new attributes to represent the negation of each original attribute. It is likely that the efficiency hit introduced by this transformation can be removed through more advanced techniques à la [OSW07, LSW10], but we leave this for future work.

⁷ Some works associate ciphertexts with a set $S \subseteq [n]$ where $[n]$ is referred to as the attribute universe, in which case $\mathbf{x} \in \{0, 1\}^n$ corresponds to the characteristic vector of S .

q -type assumption [LW12,BSW07,Att14,AC17], or imposes the one-use restriction (that is, ρ is injective and $m = n$, in which case security can be achieved unconditionally) [LOS⁺10,Wee14]. Our analysis relies on a piecewise guessing framework which refines and simplifies a recent framework of Jafargholi et al. for proving adaptive security via pebbling games [JKK⁺17] (which in turn builds upon [FKPR14,FJP15,HJO⁺16,JW16]).

Let \mathbf{G}_0 denote the view of the adversary $(\text{ct}'_{\mathbf{x}}, \text{sk}_f)$ in the real game, and \mathbf{G}_1 denote the same thing except we replace $\{\mu_j\}$ in sk_f with shares of a random value independent of μ . Our goal is to show that $\mathbf{G}_0 \approx_c \mathbf{G}_1$. First, let us define an additional family of games $\{\mathbf{H}^U\}$ parameterized by $U \subseteq [m]$: \mathbf{H}^U is the same as \mathbf{G}_0 except we replace $\{\mu_j : j \in U\}$ in sk_f with uniformly random values. In particular, $\mathbf{H}^\emptyset = \mathbf{G}_0$.

We begin with the “selective” setting, where the adversary specifies \mathbf{x} at the start of the game. Suppose we can show that $\mathbf{G}_0 \approx_c \mathbf{G}_1$ in this simpler setting via a series of $L + 1$ hybrids of the form:

$$\mathbf{G}_0 = \mathbf{H}^{h_0(\mathbf{x})} \approx_c \mathbf{H}^{h_1(\mathbf{x})} \approx_c \dots \approx_c \mathbf{H}^{h_L(\mathbf{x})} = \mathbf{G}_1$$

where $h_0, \dots, h_L : \{0, 1\}^n \rightarrow \{U \subseteq [m] : |U| \leq R'\}$ are functions of the adversary’s choices \mathbf{x} . Then, the piecewise guessing framework basically tells us that $\mathbf{G}_0 \approx_c \mathbf{G}_1$ in the adaptive setting with a security loss roughly $m^{R'} \cdot L$, where the factor L comes from the hybrid argument and the factor $m^{R'}$ comes from guessing $h_i(\mathbf{x})$ (a subset of $[m]$ of size at most R'). Ideally, we would want $m^{R'} \ll 2^n$, where 2^n is what we achieve from guessing \mathbf{x} itself

First, we describe a straight-forward approach which achieves $L = 2$ and $R' = m$ implicit in [LW12] (but incurs a huge security loss $2^m \gg 2^n$) where

$$h_1(\mathbf{x}) = \{j : x_{\rho(j)} = 0\}.$$

That is, $\mathbf{H}^{h_1(\mathbf{x})}$ is \mathbf{G}_0 with μ_j in sk_f replaced by fresh $\mu'_j \leftarrow \mathbb{Z}_p$ for all j satisfying $x_{\rho(j)} = 0$. Here, we have

- $\mathbf{G}_0 \approx_c \mathbf{H}^{h_1(\mathbf{x})}$ via DDH, since $h^{\mu_j + w_{\rho(j)} r_j}, h^{r_j}$ computationally hides μ_j whenever $x_{\rho(j)} = 0$ and $w_{\rho(j)}$ is not leaked in $\text{ct}_{\mathbf{x}}$;
- $\mathbf{H}^{h_1(\mathbf{x})} \approx_s \mathbf{G}_1$ via security of the secret-sharing scheme since the shares $\{\mu_j : x_{\rho(j)} = 1\}$ leak no information about μ whenever $f(\mathbf{x}) = 0$.

This approach is completely generic and works for any secret-sharing scheme.

In our construction, we use a variant of the secret-sharing scheme for NC^1 in [JKK⁺17] (which is in turn a variant of Yao’s secret-sharing scheme [VNS⁺03,IK02]), for which the authors also gave a hybrid argument achieving $L = 8^d$ and $R' = O(d \log m)$ where d is the depth of the formula; this achieves a security loss $2^{O(d \log m)}$. Recall that the circuit complexity class NC^1 is captured by Boolean formulas of logarithmic depth and fan-in two, so the security loss here is quasi-polynomial in n . We provide a more detailed analysis of the functions h_0, h_1, \dots, h_L used in their scheme, and show that the subsets of size $O(d)$ output by these functions can be described only $O(d)$ bits instead of $O(d \log m)$ bits. Roughly speaking, we show that the subsets are essentially determined by a path of length d from the output gate to an input gate, which can be described using $O(d)$ bits since the gates have fan-in two. Putting everything together, this allows us to achieve adaptive security for the core 1-ABE component with a security loss $2^{O(d)} = \text{poly}(n)$.

Our ABE scheme. To complete the overview, we sketch our final ABE scheme which is secure under the k -Linear Assumption in prime-order bilinear groups.

To obtain prime-order analogues of the composite-order examples, we rely on the previous framework of Chen et al. [CGW15,GDCC16,BKP14] for simulating composite-order groups in

prime-order ones. Let (G_1, G_2, G_T) be a bilinear group of prime order p . We start with the KP-ABE scheme in (1) and carry out the following substitutions:

$$g^s \mapsto [\mathbf{s}^\top \mathbf{A}]_1, h^{r_j} \mapsto [\mathbf{r}_j]_2, w_i \mapsto \mathbf{W}_i \leftarrow \mathbb{Z}_p^{(k+1) \times k}, \mu \mapsto \mathbf{v} \leftarrow \mathbb{Z}_p^{k+1} \quad (2)$$

where

$$\mathbf{A} \leftarrow \mathbb{Z}_p^{k \times (k+1)}, \mathbf{s}, \mathbf{r}_j \leftarrow \mathbb{Z}_p^k,$$

k corresponds to the k -Lin Assumption desired for security⁸, and $[\cdot]_1, [\cdot]_2$ correspond respectively to exponentiations in the prime-order groups G_1, G_2 . We note that the naive transformation following [CGW15] would have required \mathbf{W}_i of dimensions at least $(k+1) \times (k+1)$; here, we incorporated optimizations from [GDCC16, BKP14]. This yields the following prime-order KP-ABE scheme for NC¹:

$$\begin{aligned} \text{msk} &:= (\mathbf{v}, \mathbf{W}_1, \dots, \mathbf{W}_n) \\ \text{mpk} &:= ([\mathbf{A}]_1, [\mathbf{A}\mathbf{W}_1]_1, \dots, [\mathbf{A}\mathbf{W}_n]_1, e([\mathbf{A}]_1, [\mathbf{v}]_2)), \\ \text{ct}_x &:= \left([\mathbf{s}^\top \mathbf{A}]_1, \{[\mathbf{s}^\top \mathbf{A}\mathbf{W}_i]_1\}_{x_i=1}, e([\mathbf{s}^\top \mathbf{A}]_1, [\mathbf{v}]_2) \cdot M \right) \\ \text{sk}_f &:= (\{[\mathbf{v}_j + \mathbf{W}_{\rho(j)} \mathbf{r}_j]_2, [\mathbf{r}_j]_2\}_{j \in [m]}) \end{aligned}$$

where \mathbf{v}_j is the j 'th share of \mathbf{v} . Decryption proceeds as before by first computing

$$\{e([\mathbf{s}^\top \mathbf{A}]_1, [\mathbf{v}_j]_2)\}_{\rho(j)=0 \vee x_{\rho(j)}=1}$$

and relies on the associativity relations $\mathbf{A}\mathbf{W}_i \cdot \mathbf{r}_j = \mathbf{A} \cdot \mathbf{W}_i \mathbf{r}_j$ for all i, j [CW13].

In the proof, in place of the DDH assumption which allows us to argue that $(h^{w_i r_j}, h^{r_j})$ is pseudorandom, we will rely on the fact that by the k -Lin assumption, we have

$$(\mathbf{A}, \mathbf{A}\mathbf{W}_i, [\mathbf{W}_i \mathbf{r}_j]_2, [\mathbf{r}_j]_2) \approx_c (\mathbf{A}, \mathbf{A}\mathbf{W}_i, [\mathbf{W}_i \mathbf{r}_j + \delta_{ij} \mathbf{a}^\perp]_2, [\mathbf{r}_j]_2)$$

where $\mathbf{A} \leftarrow \mathbb{Z}_p^{k \times (k+1)}$, $\mathbf{W}_i \leftarrow \mathbb{Z}_p^{(k+1) \times 2k}$, $\mathbf{r}_j \leftarrow \mathbb{Z}_p^{2k}$ and $\mathbf{a}^\perp \in \mathbb{Z}_p^{k+1}$ satisfies $\mathbf{A} \cdot \mathbf{a}^\perp = \mathbf{0}$.

Organization. We describe the piecewise guessing framework for adaptive security in Section 3 and a pebbling strategy (used to define h_0, \dots, h_L) in Section 4. We describe a secret-sharing scheme and prove adaptive security of the core 1-ABE component in Section 5. We present our full KP-ABE and CP-ABE schemes in Section 6 and Appendix A. We also present our unbounded KP-ABE scheme in Section B.

2 Preliminaries

Notation. We denote by $s \leftarrow S$ the fact that s is picked uniformly at random from a finite set S . By PPT, we denote a probabilistic polynomial-time algorithm. Throughout this paper, we use 1^λ as the security parameter. We use lower case boldface to denote (column) vectors and upper case boldface to denote matrices. We use \equiv to denote two distributions being identically distributed, and \approx_c to denote two distributions being computationally indistinguishable. For any two finite sets (also including spaces and groups) S_1 and S_2 , the notation “ $S_1 \approx_c S_2$ ” means the uniform distributions over them are computationally indistinguishable.

⁸ e.g. $k = 1$ corresponds to security under the Symmetric External Diffie-Hellman Assumption (SXDH), and $k = 2$ corresponds to security under the Decisional Linear Assumption (DLIN).

2.1 Monotone Boolean formulas and NC^1

Monotone Boolean formula. A monotone Boolean formula $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is specified by a directed acyclic graph (DAG) with three kinds of nodes: input gate nodes, gate nodes, and a single output node. Input nodes have in-degree 0 and out-degree 1, AND/OR nodes have in-degree (fan-in) 2 and out-degree (fan-out) 1, and the output node has in-degree 1 and out-degree 0. We number the edges (wires) $1, 2, \dots, m$, and each gate node is defined by a tuple (g, a_g, b_g, c_g) where $g : \{0, 1\}^2 \rightarrow \{0, 1\}$ is either AND or OR, a_g, b_g are the incoming wires, c_g is the outgoing wire and $a_g, b_g < c_g$. The size of a formula m is the number of edges in the underlying DAG and the depth of a formula d is the length of the longest path from the output node.

NC^1 and log-depth formula. A standard fact from complexity theory tells us that the circuit complexity class monotone NC^1 is captured by monotone Boolean formulas of log-depth and fan-in two. This follows from the fact that we can transform any depth d circuit with fan-in two and unbounded fan-out into an equivalent circuit with fan-in two and fan-out one (for all gate nodes) of the same depth, and a 2^d blow-up in the size. To see this, note that one can start with the root gate of an NC^1 circuit and work downward by each level of depth. For each gate g considered at depth i , if either of its two input wires are coming from the output wire of a gate (at depth $i - 1$) with more than one output wire, then create a new copy of the gate at depth $i - 1$ with a single output wire going to g (note that this copy may increase the output wire multiplicity of gates at depth strictly lower than $i - 1$). This procedure preserves the functionality of the original circuit, and has the result that at its end, each gate in the circuit has input wires which come from gates with output multiplicity 1. The procedure does not increase the depth of the circuit (any duplicated gates are added at a level that already exists), so the new circuit is a formula (all gates have fan-out 1) of depth d with fan-in 2, so its size is at most 2^d . d is logarithmic in the size of the input for NC^1 circuits, so the blowup from this procedure is polynomial in n . Hence we will consider the class NC^1 as a set of Boolean formulas (where gates have fan-in 2 and fan-out 1) of depth $O(\log n)$ and refer to $f \in \text{NC}^1$ formulas.

2.2 Secret sharing

A secret sharing scheme is a pair of algorithms (`share`, `reconstruct`) where `share` on input $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and $\mu \in \mathbb{Z}_p$ outputs $\mu_1, \dots, \mu_m \in \mathbb{Z}_p$ together with $\rho : [m] \rightarrow \{0, 1, \dots, n\}$.

- Correctness stipulates that for every $x \in \{0, 1\}^n$ such that $f(x) = 1$, we have

$$\text{reconstruct}(f, x, \{\mu_j\}_{\rho(j)=0 \vee x_{\rho(j)}=1}) = \mu.$$

- Security stipulates that for every $x \in \{0, 1\}^n$ such that $f(x) = 0$, the shares

$$\{\mu_j\}_{\rho(j)=0 \vee x_{\rho(j)}=1}$$

perfectly hide μ .

Note the inclusion of $\rho(j) = 0$ in both correctness and security. All the secret sharing schemes in this work will in fact be linear (in the standard sense): `share` computes a linear function of the secret μ and randomness over \mathbb{Z}_p , and `reconstruct` computes a linear function of the shares over \mathbb{Z}_p , that is, $\mu = \sum_{\rho(j)=0 \vee x_{\rho(j)}=1} \omega_j \mu_j$.

2.3 Attribute-based encryption

An attribute-based encryption (ABE) scheme for a predicate $\text{pred}(\cdot, \cdot)$ consists of four algorithms (Setup, Enc, KeyGen, Dec):

$\text{Setup}(1^\lambda, \mathcal{X}, \mathcal{Y}, \mathcal{M}) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm gets as input the security parameter λ , the attribute universe \mathcal{X} , the predicate universe \mathcal{Y} , the message space \mathcal{M} and outputs the public parameter mpk , and the master key msk .

$\text{Enc}(\text{mpk}, x, m) \rightarrow \text{ct}_x$. The encryption algorithm gets as input mpk , an attribute $x \in \mathcal{X}$ and a message $m \in \mathcal{M}$. It outputs a ciphertext ct_x . Note that x is public given ct_x .

$\text{KeyGen}(\text{mpk}, \text{msk}, y) \rightarrow \text{sk}_y$. The key generation algorithm gets as input msk and a value $y \in \mathcal{Y}$. It outputs a secret key sk_y . Note that y is public given sk_y .

$\text{Dec}(\text{mpk}, \text{sk}_y, \text{ct}_x) \rightarrow m$. The decryption algorithm gets as input sk_y and ct_x such that $\text{pred}(x, y) = 1$. It outputs a message m .

Correctness. We require that for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$ such that $\text{pred}(x, y) = 1$ and all $m \in \mathcal{M}$,

$$\Pr[\text{Dec}(\text{mpk}, \text{sk}_y, \text{Enc}(\text{mpk}, x, m)) = m] = 1,$$

where the probability is taken over $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \mathcal{X}, \mathcal{Y}, \mathcal{M})$, $\text{sk}_y \leftarrow \text{KeyGen}(\text{mpk}, \text{msk}, y)$, and the coins of Enc.

Security definition. For a stateful adversary \mathcal{A} , we define the advantage function

$$\text{Adv}_{\mathcal{A}}^{\text{ABE}}(\lambda) := \Pr \left[\begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \mathcal{X}, \mathcal{Y}, \mathcal{M}); \\ (x^*, m_0, m_1) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{mpk}); \\ b \leftarrow_{\text{R}} \{0, 1\}; \text{ct}_{x^*} \leftarrow \text{Enc}(\text{mpk}, x^*, m_b); \\ b' \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{ct}_{x^*}) \end{array} \right] - \frac{1}{2}$$

with the restriction that all queries y that \mathcal{A} makes to $\text{KeyGen}(\text{msk}, \cdot)$ satisfy $\text{pred}(x^*, y) = 0$ (that is, sk_y does not decrypt ct_{x^*}). An ABE scheme is *adaptively secure* if for all PPT adversaries \mathcal{A} , the advantage $\text{Adv}_{\mathcal{A}}^{\text{ABE}}(\lambda)$ is a negligible function in λ .

2.4 Prime-Order Bilinear Groups and the Matrix Diffie-Hellman Assumption

A generator \mathcal{G} takes as input a security parameter λ and outputs a group description $\mathbb{G} := (p, G_1, G_2, G_T, e)$, where p is a prime of $\Theta(\lambda)$ bits, G_1 , G_2 and G_T are cyclic groups of order p , and $e : G_1 \times G_2 \rightarrow G_T$ is a non-degenerate bilinear map. We require that the group operations in G_1 , G_2 and G_T as well the bilinear map e are computable in deterministic polynomial time with respect to λ . Let $g_1 \in G_1$, $g_2 \in G_2$ and $g_T = e(g_1, g_2) \in G_T$ be the respective generators. We employ the *implicit representation* of group elements: for a matrix \mathbf{M} over \mathbb{Z}_p , we define $[\mathbf{M}]_1 := g_1^{\mathbf{M}}$, $[\mathbf{M}]_2 := g_2^{\mathbf{M}}$, $[\mathbf{M}]_T := g_T^{\mathbf{M}}$, where exponentiation is carried out component-wise. Also, given $[\mathbf{A}]_1, [\mathbf{B}]_2$, we let $e([\mathbf{A}]_1, [\mathbf{B}]_2) = [\mathbf{AB}]_T$.

We define the matrix Diffie-Hellman (MDDH) assumption on G_1 [EHK+13]:

Definition 1 (MDDH $_{k,\ell}^m$ Assumption). Let $\ell > k \geq 1$ and $m \geq 1$. We say that the MDDH $_{k,\ell}^m$ assumption holds if for all PPT adversaries \mathcal{A} , the following advantage function is negligible in λ .

$$\text{Adv}_{\mathcal{A}}^{\text{MDDH}_{k,\ell}^m}(\lambda) := \left| \Pr[\mathcal{A}(\mathbb{G}, [\mathbf{M}]_1, [\mathbf{MS}]_1) = 1] - \Pr[\mathcal{A}(\mathbb{G}, [\mathbf{M}]_1, [\mathbf{U}]_1) = 1] \right|$$

where $\mathbf{M} \leftarrow_{\text{R}} \mathbb{Z}_p^{\ell \times k}$, $\mathbf{S} \leftarrow_{\text{R}} \mathbb{Z}_p^{k \times m}$ and $\mathbf{U} \leftarrow_{\text{R}} \mathbb{Z}_p^{\ell \times m}$.

The MDDH assumption on G_2 can be defined in an analogous way. Escala *et al.* [EHK⁺13] showed that

$$k\text{-Lin} \Rightarrow \text{MDDH}_{k,k+1}^1 \Rightarrow \text{MDDH}_{k,\ell}^m \quad \forall \ell > k, m \geq 1$$

with a tight security reduction (that is, $\text{Adv}_{\mathcal{A}}^{\text{MDDH}_{k,\ell}^m}(\lambda) = \text{Adv}_{\mathcal{A}'}^{k\text{-Lin}}(\lambda)$). In fact, the MDDH assumption is a generalization of the k -Lin Assumption, such that the k -Lin Assumption is equivalent to the $\text{MDDH}_{k,k+1}^1$ Assumption as defined above.

Definition 2 (k -Lin Assumption). *Let $k \geq 1$. We say that the k -Lin Assumption holds if for all PPT adversaries \mathcal{A} , the following advantage function is negligible in λ .*

$$\text{Adv}_{\mathcal{A}}^{k\text{-Lin}}(\lambda) := \text{Adv}_{\mathcal{A}}^{\text{MDDH}_{k,k+1}^1}(\lambda)$$

Henceforth, we will use MDDH_k to denote $\text{MDDH}_{k,k+1}^1$. Lastly, we note that the k -Lin Assumption itself is a generalization, where setting $k = 1$ yields the Symmetric External Diffie-Hellman Assumption (SXDH), and setting $k = 2$ yields the standard Decisional Linear Assumption (DLIN).

3 Piecewise Guessing Framework for Adaptive Security

We now refine the adaptive security framework of [JKK⁺17], making some simplifications along the way to yield the piecewise guessing framework that will support our security proof. We use $\langle \mathbf{A}, \mathbf{G} \rangle$ to denote the output of an adversary \mathbf{A} in an interactive game \mathbf{G} , and an adversary wins if the output is 1, so that the winning probability is denoted by $\Pr[\langle \mathbf{A}, \mathbf{G} \rangle = 1]$.

Suppose we have two adaptive games \mathbf{G}_0 and \mathbf{G}_1 which we would like to show to be indistinguishable. In both games, an adversary \mathbf{A} makes some adaptive choices that define $z \in \{0, 1\}^R$. Informally, the piecewise guessing framework tells us that if we can show that $\mathbf{G}_0, \mathbf{G}_1$ are ϵ -indistinguishable in the selective setting where all choices defining z are committed to in advance via a series of $L + 1$ hybrids, where each hybrid depends only on at most $R' \ll R$ bits of information about z , then $\mathbf{G}_0, \mathbf{G}_1$ are $2^{2R'} \cdot L \cdot \epsilon$ -indistinguishable in the adaptive setting.

Overview. We begin with the selective setting where the adversary commits to $z = z^*$ in advance. Suppose we can show that $\mathbf{G}_0 \approx_c \mathbf{G}_1$ in this simpler setting via a series of $L + 1$ hybrids of the form:

$$\mathbf{G}_0 = \mathbf{H}^{h_0(z^*)} \approx_c \mathbf{H}^{h_1(z^*)} \approx_c \dots \approx_c \mathbf{H}^{h_L(z^*)} = \mathbf{G}_1$$

where $h_0, \dots, h_L : \{0, 1\}^R \rightarrow \{0, 1\}^{R'}$ and $\{\mathbf{H}^u\}_{u \in \{0, 1\}^{R'}}$ is a family of games where the messages sent to the adversary in \mathbf{H}^u depend on u .⁹ In particular, the ℓ 'th hybrid only depends on $h_\ell(z^*)$ where $|h_\ell(z^*)| \ll |z^*|$.

Next, we describe how to slightly strengthen this hybrid sequence so that we can deduce that $\mathbf{G}_0 \approx_c \mathbf{G}_1$ even for an adaptive choice of z . Note that $\{\mathbf{H}^u\}_{u \in \{0, 1\}^{R'}}$ is now a family of adaptive games where z is adaptively defined as the game progresses. We have two requirements:

The first, *end-point equivalence*, just says the two equivalences

$$\mathbf{G}_0 = \mathbf{H}^{h_0(z^*)}, \quad \mathbf{G}_1 = \mathbf{H}^{h_L(z^*)}$$

⁹ Informally, $\{\mathbf{H}^u\}$ describes the simulated games used in the security reduction, where the reduction guesses R' bits of information described by u about some choices z made by the adversary; these R' bits of information are described by $h_\ell(z)$ in the ℓ 'th hybrid. In the ℓ 'th hybrid, the reduction guesses a $u \in \{0, 1\}^{R'}$ and simulates the game according to \mathbf{H}^u and hopes that the adversary will pick an z such that $h_\ell(z) = u$; note that the adversary is not required to pick such an z . One way to think of \mathbf{H}^u is that the reduction is committed to u , but the adversary can do whatever it wants.

hold even in the adaptive setting, that is, even if the adversary’s behavior defines an z different from z^* . In our instantiation, h_0 and h_L are constant functions, so this equivalence will be immediate.

The second, *neighbor indistinguishability*, basically says that for any $\ell \in [L]$, we have

$$\mathbf{H}^{u_0} \approx_c \mathbf{H}^{u_1}, \forall u_0, u_1 \in \{0, 1\}^{R'}$$

as long as the adversary chooses z such that

$$h_{\ell-1}(z) = u_0 \wedge h_\ell(z) = u_1$$

It is easy to see that this is a generalization of $\mathbf{H}^{h_{\ell-1}(z^*)} \approx_c \mathbf{H}^{h_\ell(z^*)}$ if we require $z = z^*$. To formalize this statement, we need to formalize the restriction on the adversary’s choice of z by having the game output 0 whenever the restriction is violated. That is, we define a pair of “selective” games $\widehat{\mathbf{H}}_{\ell,0}(u_0, u_1), \widehat{\mathbf{H}}_{\ell,1}(u_0, u_1)$ for any $u_0, u_1 \in \{0, 1\}^{R'}$, where

$\widehat{\mathbf{H}}_{\ell,b}(u_0, u_1)$ is the same as \mathbf{H}^{u_b} , except we replace the output with 0 whenever $(h_{\ell-1}(z), h_\ell(z)) \neq (u_0, u_1)$.

That is, in both games, the adversary “commits” in advance to u_0, u_1 . Proving indistinguishability here is easier because the reduction knows u_0, u_1 and only needs to handle adaptive choices of z such that $(h_{\ell-1}(z), h_\ell(z)) = (u_0, u_1)$.

Adaptive security lemma. The next lemma tells us that the two requirements above implies that $\mathbf{G}_0 \approx_c \mathbf{G}_1$ with a security loss $2^{2R'} \cdot L$ (stated in the contra-positive). In our applications, $2^{R'}$ and L will be polynomial in the security parameter.

Lemma 1 (adaptive security lemma). Fix $\mathbf{G}_0, \mathbf{G}_1$ along with $h_0, h_1, \dots, h_L : \{0, 1\}^R \rightarrow \{0, 1\}^{R'}$ and $\{\mathbf{H}^u\}_{u \in \{0,1\}^{R'}}$ such that

$$\forall z^* \in \{0, 1\}^R : \mathbf{H}^{h_0(z^*)} = \mathbf{G}_0, \mathbf{H}^{h_L(z^*)} = \mathbf{G}_1$$

Suppose there exists an adversary \mathbf{A} such that

$$\Pr[\langle \mathbf{A}, \mathbf{G}_0 \rangle = 1] - \Pr[\langle \mathbf{A}, \mathbf{G}_1 \rangle = 1] \geq \epsilon$$

then there exists $\ell \in [L]$ and $u_0, u_1 \in \{0, 1\}^{R'}$ such that

$$\Pr[\langle \mathbf{A}, \widehat{\mathbf{H}}_{\ell,0}(u_0, u_1) \rangle = 1] - \Pr[\langle \mathbf{A}, \widehat{\mathbf{H}}_{\ell,1}(u_0, u_1) \rangle = 1] \geq \frac{\epsilon}{2^{2R'}L}$$

This lemma is essentially a restatement of the main theorem of [JKK⁺17, Theorem 2]; we defer a comparison to the end of this section.

Proof. For the proof, we need to define the game $\mathbf{H}_\ell(z^*)$ for all $\ell = 0, 1, \dots, L$ and all $z^* \in \{0, 1\}^R$

$\mathbf{H}_\ell(z^*)$ is the same as $\mathbf{H}^{h_\ell(z^*)}$, except we replace the output with 0 whenever $z \neq z^*$.

Roughly speaking, in $\mathbf{H}_\ell(z^*)$, the adversary “commits” to making choices $z = z^*$ in advance.

– Step 1. We begin the proof by using “random guessing” to deduce that

$$\Pr_{z^* \leftarrow \{0,1\}^R} [\langle \mathbf{A}, \mathbf{H}_0(z^*) \rangle = 1] - \Pr_{z^* \leftarrow \{0,1\}^R} [\langle \mathbf{A}, \mathbf{H}_L(z^*) \rangle = 1] \geq \frac{\epsilon}{2^R}$$

This follows from the fact that $H^{h_0(z)} = G_0$, $H^{h_L(z)} = G_1$ which implies

$$\Pr_{z^* \leftarrow \{0,1\}^R} [\langle A, H_0(z^*) \rangle = 1] = \frac{1}{2^R} \Pr[\langle A, G_0 \rangle = 1]$$

$$\Pr_{z^* \leftarrow \{0,1\}^R} [\langle A, H_L(z^*) \rangle = 1] = \frac{1}{2^R} \Pr[\langle A, G_1 \rangle = 1].$$

– Step 2. Via a standard hybrid argument, we have that there exists ℓ such that

$$\Pr_{z^* \leftarrow \{0,1\}^R} [\langle A, H_{\ell-1}(z^*) \rangle = 1] - \Pr_{z^* \leftarrow \{0,1\}^R} [\langle A, H_\ell(z^*) \rangle = 1] \geq \frac{\epsilon}{2^R L}$$

which implies that:

$$\sum_{z' \in \{0,1\}^R} [\langle A, H_{\ell-1}(z') \rangle = 1] - \sum_{z' \in \{0,1\}^R} [\langle A, H_\ell(z') \rangle = 1] \geq \frac{\epsilon}{L}$$

– Step 3. Next, we relate $\widehat{H}_{\ell,0}$, $\widehat{H}_{\ell,1}$ and $H_{\ell-1}$, H_ℓ . First, we define the set

$$\mathcal{U}_\ell := \{(h_{\ell-1}(z'), h_\ell(z')) : z' \in \{0,1\}^R\} \subseteq \{0,1\}^{R'} \times \{0,1\}^{R'}, \ell \in [L]$$

Observe that for all $(u_0, u_1) \in \mathcal{U}_\ell$, we have

$$\Pr[\langle A, \widehat{H}_{\ell,1}(u_0, u_1) \rangle = 1] = \sum_{z': (h_{\ell-1}(z'), h_\ell(z')) = (u_0, u_1)} \Pr[\langle A, H_\ell(z') \rangle = 1]$$

Then, we have

$$\begin{aligned} & \sum_{z' \in \{0,1\}^R} \Pr[\langle A, H_\ell(z') \rangle = 1] \\ &= \sum_{(u_0, u_1) \in \mathcal{U}_\ell} \left(\sum_{z': (h_{\ell-1}(z'), h_\ell(z')) = (u_0, u_1)} \Pr[\langle A, H_\ell(z') \rangle = 1] \right) \\ &= \sum_{(u_0, u_1) \in \mathcal{U}_\ell} \Pr[\langle A, \widehat{H}_{\ell,1}(u_0, u_1) \rangle = 1] \end{aligned}$$

By the same reasoning, we also have

$$\sum_{z' \in \{0,1\}^R} \Pr[\langle A, H_{\ell-1}(z') \rangle = 1] = \sum_{(u_0, u_1) \in \mathcal{U}_\ell} \Pr[\langle A, \widehat{H}_{\ell,0}(u_0, u_1) \rangle = 1]$$

This means that

$$\begin{aligned} & \sum_{(u_0, u_1) \in \mathcal{U}_\ell} \left(\Pr[\langle A, \widehat{H}_{\ell,0}(u_0, u_1) \rangle = 1] - \Pr[\langle A, \widehat{H}_{\ell,1}(u_0, u_1) \rangle = 1] \right) \\ &= \sum_{z' \in \{0,1\}^R} \Pr[\langle A, H_{\ell-1}(z') \rangle = 1] - \sum_{z' \in \{0,1\}^R} \Pr[\langle A, H_\ell(z') \rangle = 1] \geq \frac{\epsilon}{L} \end{aligned}$$

where the last inequality follows from Step 2.

– Step 4. By an averaging argument, and using the fact that $|\mathcal{U}_\ell| \leq 2^{2R'}$, there exists $(u_0, u_1) \in \mathcal{U}_\ell$ such that

$$\Pr[\langle A, \widehat{H}_{\ell,0}(u_0, u_1) \rangle = 1] - \Pr[\langle A, \widehat{H}_{\ell,1}(u_0, u_1) \rangle = 1] \geq \frac{\epsilon}{2^{2R'} L}$$

This completes the proof. Note that $2^{2R'}$ can be replaced by $\max_\ell |\mathcal{U}_\ell|$. □

Comparison with [JKK⁺17]. Our piecewise guessing framework makes explicit the game H^u which are described implicitly in the applications of the framework in [JKK⁺17]. Starting from H^u and h_0, \dots, h_L , we can generically specify the intermediate games $\widehat{H}_{\ell,0}, \widehat{H}_{\ell,1}$ as well as the games H_0, \dots, H_L used in the proof of security. The framework of [JKK⁺17] does the opposite: it starts with the games H_0, \dots, H_L , and the theorem statement assumes the existence of h_0, \dots, h_L and $\widehat{H}_{\ell,0}, \widehat{H}_{\ell,1}$ that are “consistent” with H_0, \dots, H_L (as defined via a “selectivization” operation). We believe that starting from H^u and h_0, \dots, h_L yields a simpler and clearer framework which enjoys the advantage of not having to additionally construct and analyze $\widehat{H}_{\ell,0}, \widehat{H}_{\ell,1}$ and H_ℓ in the applications.

Finally, we point out that the sets \mathcal{U} and \mathcal{W} in [JKK⁺17, Theorem 2] corresponds to \mathcal{U}_ℓ and $\{0, 1\}^R$ over here (that is, we do obtain the same bounds), and the i 'th function h_i corresponds to the ℓ 'th function $h_{\ell-1} \circ h_\ell$ over here.

4 Pebbling Strategy for NC¹

We now define a pebbling strategy for NC¹ which will be used to define the functions h_0, \dots, h_L we'll use in the piecewise guessing framework. Fix a formula $f : \{0, 1\}^n \rightarrow \{0, 1\}$ of size m and an input $x \in \{0, 1\}^n$ for which $f(x) = 0$. A pebbling strategy specifies a sequence of L subsets of $[m]$, corresponding to subsets of input nodes and gates in f that are pebbled. We refer to each subset in the sequence as a pebbling configuration and the i 'th term in this sequence is the output of $h_i(f, x)$ (where the combination of f, x correspond to the adaptive choices z made in our security game that will be later analyzed in the piecewise guessing framework).

Our pebbling strategy is essentially the same as that in [JKK⁺17, Section 4]; the main difference is that we provide a better bound on the size of the description of each pebbling configuration in Theorem 1.

4.1 Pebbling Rules

Fix a formula $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and an input $x \in \{0, 1\}^n$ for which $f(x) = 0$. We are allowed to place or remove pebbles on input nodes and gates in f , subject to some rules. The goal of a pebbling strategy is to find a sequence of pebbling instructions that follow the rules and starting with the initial configuration (in which there are no pebbles at all), will end up in a configuration where only the root gate has a pebble. Intuitively, the rules say that we can place a pebble a node or a gate if we know that the out-going wire will be 0. More formally,

Definition 3 (Pebbling Rules).

1. Can place or remove a pebble on any AND gate for which (at least) one input wire comes out of a node with a pebble on it.
2. Can place or remove a pebble on any OR gate for which all of the incoming wires come out of nodes which have pebbles on them.
3. Can place or remove a pebble on any input node for which $x_i = 0$.

Given (f, x) , a pebbling strategy returns a sequence of pebbling instructions of the form PEBBLE g or unPEBBLE g for some gate g , with the property that each successively applied instruction follows the pebbling rules in Definition 3.

4.2 Pebbling Strategy

Given an NC¹ formula f (recall Section 2.1) and an input x on which the formula evaluates to 0, consider the pebbling instruction sequence returned by the following recursive procedure, which

maintains the invariant that the output wire evaluates to 0 for each gate that the procedure is called upon. The strategy is described in Figure 2 and begins by calling $\text{Pebble}(f, x, g^*)$ on the root gate g^* . We give an example in Figure 3.

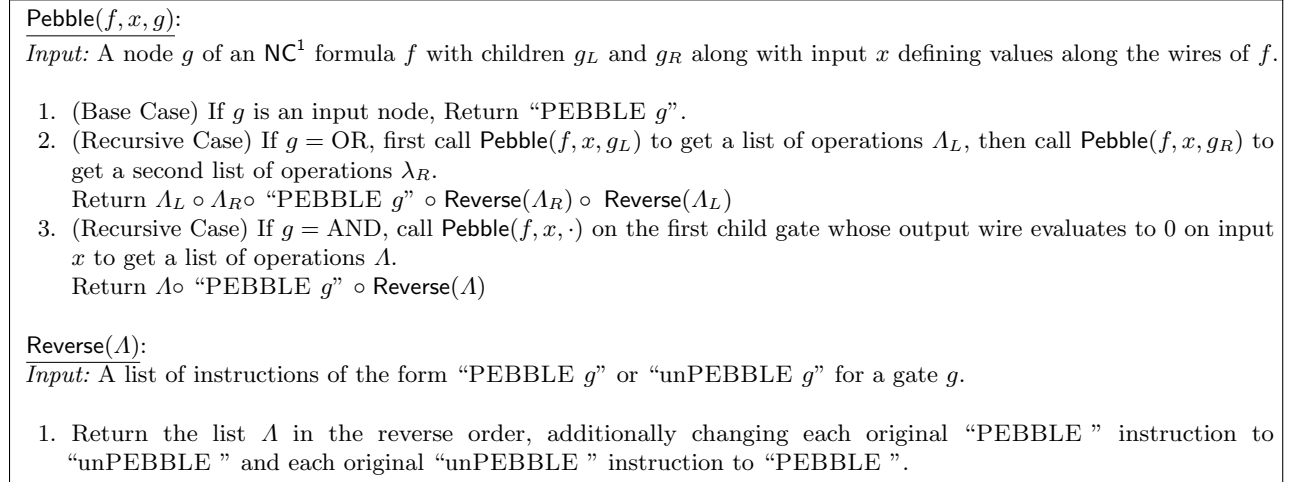


Fig. 2. NC^1 formula pebbling strategy.

Note that if this procedure is called on the root gate of a formula f with an input x such that $f(x) = 0$, then every AND gate on which the $\text{Pebble}()$ procedure is called will have *at least one* child node with an output wire which evaluates to 0, and every OR gate on which the $\text{Pebble}()$ procedure is called will have child nodes with output wires which *both* evaluate to 0. Furthermore, by inspection, $\text{Pebble}(f, x, g^*)$ returns a sequence of pebbling instructions for the circuit that follows the rules in Definition 3.

4.3 Analysis.

To be useful in the piecewise guessing framework, we would like for the sequence of pebbling instructions to have the property that each configuration formed by successive applications of the instructions in the sequence is as short to describe as possible (i.e., minimize the maximum representation size R'). One way to achieve this is to have, at any configuration along the way, as few pebbles as possible. An even more succinct representation can be obtained if we allow many pebbles but have a way to succinctly represent their location. Additionally, we would like to minimize the worst-case length, L , of any sequence produced. We achieve these two goals in the following theorem.

Theorem 1 (pebbling NC^1). *For every input $x \in \{0, 1\}^n$ and any monotone formula f of depth d and fan-in two for which $f(x) = 0$, there exists a sequence of $L(d) = 8^d$ pebbling instructions such that every intermediate pebbling configuration can be described using $R'(d) = 3d$ bits.*

Proof. Follows from the joint statements of Lemma 2 and Lemma 4 applied to the pebbling strategy in Figure 2.

Comparison with [JKK⁺17]. Note that the strategy reproduced in Figure 2 is essentially the same as one analyzed by [JKK⁺17], which argued that every configuration induced by the pebbling




Fig. 3. Intermediate pebbling configurations on input $x = 001$. The thick black outline around a node corresponds to having a pebble on the node. Note that steps 10-17 correspond to “undoing” steps 1-8 so that at the end of step 17, there is exactly one pebble on the \vee node leading to the output node.

instruction sequence it produces can be described using $d(\log m + 2)$ bits, where m is the number of wires in the formula. This follows from the fact that each such pebbling configuration has at most d gates with pebbled children, and we can specify each such gate using $\log m$ bits and the pebble-status of its two children using an additional two bits. Our Lemma 4 analyzes the same pebbling strategy but achieves a more succinct representation by leveraging the fact that not all configurations of d pebbled gates are possible due to the pebbling strategy used, so we don’t need the full generality allowed by $d \cdot \log m$ bits. Instead, Lemmas 3 and 4 show that every configuration produced follows a pattern that can be described using only $3d$ bits.

Lemma 2 ([JKK⁺17]). *The pebbling strategy in Figure 2 called on the root gate g^* for a formula f of depth d with assignment x such that $f(x) = 0$, $\text{Pebble}(f, x, g^*)$, returns a sequence of instructions of length at most $L(d) \leq 8^d$.*

This bound is a special case of that shown in [JKK⁺17, Lemma 2] for fan-in two circuits.

Proof. This statement follows inductively on the depth of the formula on which $\text{Pebble}()$ is called.

For the base case, when $d = 0$ (and Pebble has therefore been called on an input node) there is just one instruction returned, and:

$$1 \leq 8^0$$

When $\text{Pebble}()$ is called on a node at depth $d > 0$, the node is either an OR gate or an AND gate.

When $\text{Pebble}()$ is called on an OR gate, using our inductive hypothesis for the instructions returned for the subformula of depth $d - 1$, notice that the number of instructions returned is:

$$L(d - 1) + L(d - 1) + 1 + L(d - 1) + L(d - 1) = 8^{d-1} + 8^{d-1} + 1 + 8^{d-1} + 8^{d-1} = 4 \cdot 8^{d-1} + 1 \leq 8^d$$

When $\text{Pebble}()$ is called on an AND gate, using our inductive hypothesis for the instructions returned for the subformula of depth $d - 1$, notice that the number of instructions returned is:

$$L(d - 1) + 1 + L(d - 1) = 8^{d-1} + 1 + 8^{d-1} = 2 \cdot 8^{d-1} + 1 \leq 8^d$$

□

We note that the following lemma is new to this work and will be used to bound the representation size $R(d)$ of any configuration produced by application of the instructions output by the pebbling strategy.

Lemma 3 (structure of pebbling configuration). *Every configuration induced by application of the instructions produced by the pebbling strategy in Figure 2 called on the root gate g^* of a formula f of depth d with assignment x such that $f(x) = 0$, $\text{Pebble}(f, x, g^*)$, has the following property for all gates g in f with children g_L, g_R :*

If any node in the sub-tree rooted at g_R is pebbled, then there exists at most one pebble on the sub-tree rooted at g_L , namely a pebble on g_L itself

Proof. Call a node “good” if it satisfies the property above. First, we make the following observation about the behavior of $\text{Reverse}()$: Applying $\text{Reverse}()$ to a list of instructions inducing a list of configurations for which all nodes are “good” produces a new list for which this is true. This holds since $\text{Reverse}()$ does not change the configurations induced by a list of instructions, just the ordering (which is reversed). This follows from a simple proof by induction on the length of the input instruction list and the fact that for an input list of instructions parsed as $L_1 \circ L_2$ for two smaller-length lists, we can implement $\text{Reverse}(L_1 \circ L_2)$ as $\text{Reverse}(L_2) \circ \text{Reverse}(L_1)$.

We proceed with our original proof via a proof by induction on the depth of the formula upon which $\text{Pebble}()$ is called.

Inductive Hypothesis: For formulas f of depth $d - 1$ with root gate g^* and assignment x such that $f(x) = 0$, $\text{Pebble}(f, x, g^*)$ returns a sequence of instructions that induces a sequence of configurations that (1) end with a configuration where g^* is the only pebbled node, and satisfies: (2) in every configuration all nodes are “good.”

Base Case: when $\text{Pebble}(f, x, g^*)$ is called on a formula of depth 0, the formula consists of just an input node g^* . The (single) returned instruction $\text{PEBBLE } g^*$ then satisfies that in both the initial and final configuration, the single node g^* is good. Also, the sequence ends in the configuration where g^* is the only pebbled node.

Inductive Step: when $\text{Pebble}(f, x, g^*)$ is called on formula of depth $d > 0$. Let g_L^*, g_R^* denote the children of the root gate g^* (either an AND or OR gate). Note that the sub-formulas $f_{g_L^*}$ and $f_{g_R^*}$ rooted at g_L^* and g_R^* have depth $d - 1$. We proceed via a case analysis:

If g^* is an AND gate, then suppose the sequence of instructions returned is

$$\text{Pebble}(f_{g_R^*}, x, g_R^*) \circ \text{PEBBLE } g^* \circ \text{Reverse}(\text{Pebble}(f_{g_R^*}, x, g_R^*))$$

(The case with g_L^* instead of g_R^* is handled analogously, even simpler). Suppose $\text{Pebble}(f_{g_R^*}, x, g_R^*)$ (and thus $\text{Reverse}(\text{Pebble}(f_{g_R^*}, x, g_R^*))$) produces L_0 instructions. We proceed via a case analysis:

- Take any of the first L_0 configurations (starting from 0'th). Here, all pebbles are in the subformula rooted at g_R^* . We can then apply part (2) of the inductive hypothesis to the subformula $f_{g_R^*}$ rooted at g_R^* (of depth $d - 1$) to deduce that property “good” holds for all nodes in $f_{g_R^*}$. All nodes in $f_{g_L^*}$ are unpebbled in all configurations, so they are automatically good. Lastly, the root gate g^* has no pebbled nodes in the subformula rooted at g_L , so it is also good.
- For the $(L_0 + 1)$ 'th configuration reached after $\text{PEBBLE } g^*$, there are only two pebbles, one on g^* (from the $\text{PEBBLE } g^*$ instruction) and another on g_R^* (from part (1) of our inductive hypothesis applied to the (depth $d - 1$) subformula $f_{g_R^*}$). It is clear that all nodes in this configuration are good.
- For the last L_0 configurations, there is one pebble on g^* and all remaining pebbles are in the subformula rooted at g_R^* . Clearly, g^* is good. All nodes in $f_{g_L^*}$ are unpebbled in all configurations,

so they are also good. Moreover, we can apply the inductive hypothesis to $f_{g_R^*}$ combined with our observation that `Reverse` preserves property (2) of this hypothesis to deduce that all nodes in the subformula are also good for all configurations.

Lastly, notice that since the last L_0 instructions undo the first L_0 instructions, the final configuration features a single pebble on g^* .

If g^* is an OR gate, then the sequence of instructions returned is

$$\text{Pebble}(f_{g_L^*}, x, g_L^*) \circ \text{Pebble}(f_{g_R^*}, x, g_R^*) \circ \text{PEBBLE } g^* \circ \text{Reverse}(\text{Pebble}(f_{g_R^*}, x, g_R^*)) \circ \text{Reverse}(\text{Pebble}(f_{g_L^*}, x, g_L^*))$$

Suppose $\text{Pebble}(f_{g_R^*}, x, g_R^*)$, $\text{Pebble}(f_{g_L^*}, x, g_L^*)$, and thus $\text{Reverse}(\text{Pebble}(f_{g_R^*}, x, g_R^*))$, $\text{Reverse}(\text{Pebble}(f_{g_L^*}, x, g_L^*))$, produces L_0, L_1 instructions. We proceed via a case analysis:

- Take any of the first L_0 configurations (starting from 0'th). Here, all pebbles are in the subformula $f_{g_L^*}$ rooted at g_L^* . We can then apply part (2) of the inductive hypothesis to (depth $d - 1$) $f_{g_L^*}$ to deduce that property “good” holds for all nodes in $f_{g_L^*}$. All nodes in the subformula rooted at g_R^* , $f_{g_R^*}$, are unpebbled in all configurations, so they are automatically good. Lastly, the root gate g^* has no pebbled nodes in the subformula rooted at g_R^* , so it is also good. Finally, by part (1) of this application of the inductive hypothesis, we know that L_0 th configuration features a single pebble on g_L^* .
- Take any of the next L_1 configurations (starting from the L_0 'th). Here, all pebbles are in the subformula rooted at g_R^* except for the single pebble on g_L^* . We can then apply part (2) of the inductive hypothesis to (depth $d - 1$) $f_{g_R^*}$ (of depth $d - 1$) to deduce that property “good” holds for all nodes in $f_{g_R^*}$. All nodes in the subformula rooted at g_L^* have no pebbles in their own subformulas, so they are automatically good. Lastly, the root gate g^* may have pebbled nodes in the subformula rooted at g_R^* but the only pebbled node in the subformula rooted at g_L^* is g_L^* itself, so it is also good. Finally, we know that the $L_0 + L_1$ th configuration features two pebbles: a pebble on g_L^* (from the first L_0 instructions), and a pebble on g_R^* (by part (1) of this application of the inductive hypothesis).
- For the $(L_0 + L_1 + 1)$ 'th configuration reached after `PEBBLE` g^* , there are only three pebbles, one on g^* (from the `PEBBLE` g^* instruction), one on g_L^* (from the first L_0 instructions), and another on g_R^* (from the next L_1 instructions). It is clear that all nodes in this configuration are good.
- For the next L_1 configurations (reversing the instructions of the set of size L_1), there is one pebble on g^* , one pebble on g_L^* , and all remaining pebbles are in the subformula rooted at g_R^* , $f_{g_R^*}$. g^* is good, since it only has one pebble in the subformula rooted at g_L^* , on g_L^* itself. All nodes in the subformula rooted at g_L^* have no pebbles in their own subformulas, so they are also good. Moreover, we can apply the inductive hypothesis to (depth $d - 1$) $f_{g_R^*}$ combined with our observation that `Reverse` preserves property (2) of this hypothesis to deduce that all nodes in $f_{g_R^*}$ are also good for all configurations. Note the final configuration in this sequence then contains two pebbles, one of g^* and one on g_L^* .
- For the final L_0 configurations (reversing the instructions of the set of size L_0), there is one pebble on g^* , and all remaining pebbles are in the subformula rooted at g_L^* . g^* is good, since it has no pebbles in the subformula rooted at g_R^* . Similarly, all nodes in the subformula rooted at g_R^* are also good. Moreover, we can apply the inductive hypothesis to (depth $d - 1$) $f_{g_L^*}$ combined with our observation that `Reverse` preserves property (2) of this hypothesis to deduce that all nodes in $f_{g_L^*}$ are also good for all configurations.

Lastly, notice that since the last $L_0 + L_1$ instructions undo the first $L_0 + L_1$ instructions, the final configuration features a single pebble on g^* .

□

Lemma 4 ($R'(d) = 3d$). *Every configuration induced by application of the instructions produced by the pebbling strategy in Figure 2 for a formula f of depth d with assignment x such that $f(x) = 0$ can be described using $R'(d) = 3d$ bits.*

Proof. We can interpret $3d$ bits in the following way to specify a pebbling: the first d bits specify a path down the formula starting at the root gate (moving left or right based on the setting of each bit), the next $2(d - 1)$ bits specify, for each of the $(d - 1)$ non-input nodes along the path, which of its children are pebbled. Finally one of the last 2 bits is used to denote if the root node is pebbled.

From Lemma 3, we know that for all gates g with children g_L, g_R , if any node in the sub-tree rooted at g_R is pebbled, then there exists at most one pebble on the sub-tree rooted at g_L , namely a pebble on g_L itself. So, given a pebbling configuration, we can start at the root node and describe the path defined by taking the child with more pebbles on its subtree using d bits. All pebbles in the configuration are either on the root node or on children of nodes on this path and therefore describable in the remaining $2d$ bits. \square

5 Core Adaptive Security Component

In this section, we will describe the secret-sharing scheme (`share`, `reconstruct`) used in our ABE construction. In addition, we describe a core component of our final ABE, and prove adaptive security using the pebbling strategy defined and analyzed in Section 4 to define hybrids in the piecewise guessing framework of Section 3.

Overview. As described in the overview in Section 1.1, we will consider the following “core 1-ABE component”:

$$\begin{aligned} \text{ct}'_{\mathbf{x}} &:= (\{w_i\}_{x_i=1}) \quad // \text{“stripped down” } \text{ct}_{\mathbf{x}} \\ \text{sk}'_f &:= (\{h^{\mu_j}\}_{\rho(j)=0} \cup \{h^{\mu_j+r_j w_{\rho(j)}}, h^{r_j}\}_{\rho(j)\neq 0}) \end{aligned}$$

where $(\{\mu_j\}, \rho) \leftarrow \text{share}(f, \mu)$. We want to show that under the DDH assumption, μ is hidden given just $(\text{ct}'_{\mathbf{x}}, \text{sk}'_f)$ where \mathbf{x}, f are adaptively chosen subject to the constraint $f(\mathbf{x}) = 0$. We formalize this via a pair of games $\mathbf{G}_0^{1\text{-ABE}}, \mathbf{G}_1^{1\text{-ABE}}$ and the requirement $\mathbf{G}_0^{1\text{-ABE}} \approx_c \mathbf{G}_1^{1\text{-ABE}}$. In fact, we will study a more abstract construction based on any CPA-secure encryption with:

$$\begin{aligned} \text{ct}'_{\mathbf{x}} &:= (\{w_i\}_{x_i=1}) \quad // \text{“stripped down” } \text{ct}_{\mathbf{x}} \\ \text{sk}'_f &:= \{\mu_j\}_{\rho(j)=0} \cup \{\text{CPA.Enc}(w_{\rho(j)}, \mu_j)\}_{\rho(j)\neq 0} \text{ where } (\{\mu_j\}, \rho) \leftarrow \text{share}(f, \mu) \end{aligned}$$

5.1 Linear secret sharing for NC^1

We first describe a linear secret-sharing scheme for NC^1 ; this is essentially the information-theoretic version of Yao’s secret-sharing for NC^1 in [JKK⁺17, VNS⁺03, IK02]. It suffices to work with Boolean formulas where gates have fan-in 2 and fan-out 1, thanks to the transformation in Section 2.1. We describe the scheme in Figure 4, and give an example in Figure 5. Note that our non-standard definition of secret-sharing in Section 2.2 allows the setting of $\rho(j) = 0$ for shares that are available for reconstruction for all x . We remark that the output of share satisfies $|\{\mu_j\}| \leq 2m$ since each of the m nodes adds a single μ_j to the output set, except for OR gates which add two: μ_{j_a} and μ_{j_b} .

The reconstruction procedure `reconstruct` of the scheme is essentially applying the appropriate linear operations to get the output wire value $\hat{\mu}_c$ at each node starting from the leaves of the formula to get to the root $\hat{\mu}_m = \mu$.

share(f, μ):

Input: A formula $f : \{0, 1\}^n \rightarrow \{0, 1\}$ of size m and a secret $\mu \in \mathbb{Z}_p$.

1. For each non-output wire $j = 1, \dots, m-1$, pick a uniformly random $\hat{\mu}_j \leftarrow \mathbb{Z}_p$. For the output wire, set $\hat{\mu}_m = \mu$
2. For each outgoing wire j from input node i , add $\mu_j = \hat{\mu}_j$ to the output set of shares and set $\rho(j) = i$.
3. For each AND gate g with input wires a, b and output wire c , add $\mu_c = \hat{\mu}_c + \hat{\mu}_a + \hat{\mu}_b \in \mathbb{Z}_p$ to the output set of shares and set $\rho(c) = 0$.
4. For each OR gate g with input wires a, b and output wire c , add $\mu_{c_a} = \hat{\mu}_c + \hat{\mu}_a \in \mathbb{Z}_p$ and $\mu_{c_b} = \hat{\mu}_c + \hat{\mu}_b \in \mathbb{Z}_p$ to the output set of shares and set $\rho(c_a) = 0$ and $\rho(c_b) = 0$.
5. Output $\{\mu_j\}, \rho$.

Fig. 4. Information-theoretic linear secret sharing scheme **share** for NC^1




Fig. 5. Left: Formula $(x_1 \vee x_2) \vee (x_1 \wedge x_3)$, where the wires are numbered $1, 2, \dots, 7$. Right: Shares $(\mu_1, \dots, \mu_{7_b})$ and mapping ρ for the formula corresponding to secret $\mu \in \mathbb{Z}_p$

- Given $\hat{\mu}_a, \hat{\mu}_b$ associated with the input wires of an AND gate, we recover the gate's output wire value $\hat{\mu}_c$ by subtracting their values from μ_c (which is available since $\rho(c) = 0$).
- Given one of $\hat{\mu}_a, \hat{\mu}_b$ associated with the input wires of an OR gate, we recover the gate's output wire value $\hat{\mu}_c$ by subtracting it from the appropriate choice of μ_{c_a} or μ_{c_b} (which are both available since $\rho(c_a) = \rho(c_b) = 0$).

Note that $\text{reconstruct}(f, x, \{\mu_j\}_{\rho(j)=0 \vee x_{\rho(j)}=1})$ computes a linear operation with respect to the shares μ_j . This follows from the fact that the operation at each gate in reconstruction is a linear operation, and the composition of linear operations is itself a linear operation. Therefore, $\text{reconstruct}(f, x, \{\mu_j\}_{\rho(j)=0 \vee x_{\rho(j)}=1})$ is equivalent to identifying the coefficients ω_j of this linear function, where $\mu = \sum_{\rho(j)=0 \vee x_{\rho(j)}=1} \omega_j \mu_j$.

As with any linear secret-sharing scheme, **share** and **reconstruct** can be extended in the natural way to accommodate vectors of secrets. Specifically, for a vector of secrets $\mathbf{v} \in \mathbb{Z}_p^k$, define:

$$\text{share}(f, \mathbf{v}) := (\{\mathbf{v}_j := (v_{1,j}, \dots, v_{k,j})\}, \rho) \text{ where } (\{v_{i,j}\}, \rho) \leftarrow \text{share}(f, v_i)$$

(note that ρ is identical for all i). **reconstruct** can also be defined component-wise:

$$\text{reconstruct}(f, x, \{\mathbf{v}_j\}_{\rho(j)=0 \vee x_{\rho(j)}=1}) := \sum_{\rho(j)=0 \vee x_{\rho(j)}=1} \omega_j \mathbf{v}_j \text{ where } \omega_j \text{ are computed as above}$$

Our final ABE construction will use this extension.

5.2 Core 1-ABE Security Game

Definition 4 (core 1-ABE security $G_0^{1\text{-abe}}, G_1^{1\text{-abe}}$). For a stateful adversary \mathcal{A} , we define the following games $G_\beta^{1\text{-ABE}}$ for $\beta \in \{0, 1\}$.

$$\langle \mathcal{A}, G_\beta^{1\text{-ABE}} \rangle := \mathbb{I} \left\{ \begin{array}{l} \mu^{(0)}, \mu^{(1)} \leftarrow \mathbb{Z}_p; w_i \leftarrow \text{CPA.Setup}(\lambda) \\ b' \leftarrow \mathcal{A}^{\mathcal{O}_F(\cdot), \mathcal{O}_X(\cdot), \mathcal{O}_E(\cdot, \cdot)}(\mu^{(0)}) \end{array} \right\}$$

where the adversary \mathcal{A} adaptively interacts with three oracles:

$$\begin{aligned} \mathcal{O}_F(f) &:= \{\text{sk}'_f = \{\mu_j\}_{\rho(j)=0} \cup \{\text{CPA.Enc}(w_{\rho(j)}, \mu_j)\}_{\rho(j) \neq 0} \text{ where } (\{\mu_j\}, \rho) \leftarrow \text{share}(f, \mu^{(\beta)}) \\ \mathcal{O}_X(x) &:= (\text{ct}'_x = \{w_i\}_{x_i=1}) \\ \mathcal{O}_E(i, m) &:= \text{CPA.Enc}_{w_i}(m) \end{aligned}$$

with the restrictions that (i) only one query is made to each of $\mathcal{O}_F(\cdot)$ and $\mathcal{O}_X(\cdot)$, and (ii) the queries f and x to $\mathcal{O}_F(\cdot), \mathcal{O}_X(\cdot)$ respectively, satisfy $f(x) = 0$.

To be clear, the β in $G_\beta^{1\text{-ABE}}$ affects only the implementation of the oracle \mathcal{O}_F (where $\mu^{(\beta)}$ is shared). We will show that $G_0^{1\text{-ABE}} \approx_c G_1^{1\text{-ABE}}$ where we instantiate share using the scheme in Section 5.1. That is, Theorem 2 will bound the quantity:

$$\Pr[\langle \mathcal{A}, G_0^{1\text{-ABE}} \rangle = 1] - \Pr[\langle \mathcal{A}, G_1^{1\text{-ABE}} \rangle = 1]$$

Comparison with [JKK⁺17]. Proving adaptive security for the core 1-ABE with share is very similar to the proof for adaptive secret-sharing for circuits in [JKK⁺17]. One main difference is that in our case, the adaptive choices z correspond to both (f, x) , while in the adaptive secret-sharing proof of [JKK⁺17], f is fixed, and the adaptive choices correspond to x , but revealed one bit at a time (that is, $\mathcal{O}_X(i, x_i)$ returns w_i if $x_i = 1$). Another difference is the \mathcal{O}_E oracle included in our core 1-ABE game, which enables the component to be embedded in a standard dual-system hybrid proof for our full ABE systems. Lastly, we leverage our improved analysis in Lemmas 3 and 4 to achieve polynomial security loss, rather than the quasi-polynomial loss we would get from following their proof more directly.

5.3 Adaptive Security for Core 1-ABE Component

We will show that $G_0^{1\text{-ABE}} \approx_c G_1^{1\text{-ABE}}$ as defined in Definition 4 using the piecewise guessing framework. To do this, we need to first define a family of games $\{H^u\}$ along with functions h_0, \dots, h_L , using the pebbling strategy in Section 4. First, we will describe share^u , which will be used to define H^u .

Defining share^u Recall that Lemma 4 describes how to parse a $u \in \{0, 1\}^{3d}$ as a pebbling configuration: a subset of the nodes of f . Further, note that each node contains one output wire, so we can equivalently view u as a subset of $[m]$ denoting the output wires of pebbled gates. Given a pebbling configuration u of an NC^1 formula, the shares are generated as in the secret-sharing scheme in Figure 4, except for each pebbled node with output wire c , we replace μ_c with an independent random $\mu_c \leftarrow \mathbb{Z}_p$ (in the case of a pebbled OR gate, we replace both associated μ_{c_a} and μ_{c_b} with independent random $\mu_{c_a}, \mu_{c_b} \leftarrow \mathbb{Z}_p$, i.e: both μ_{c_a}, μ_{c_b} are associated with wire c). In particular, we get the procedure $\text{share}^u(f, \mu)$ defined in Figure 6.

$\text{share}^u(f, \mu)$:

Input: A formula $f : \{0, 1\}^n \rightarrow \{0, 1\}$, a secret $\mu \in \mathbb{Z}_p$, and a pebbling configuration u of the nodes of f .

1. Compute $(\{\mu'_j\}, \rho) \leftarrow \text{share}(f, \mu)$ as defined in Figure 4
2. For each μ'_j , if $j \in u$ (i.e. if j is the output wire of a pebbled node), then sample $\mu_j \leftarrow \mathbb{Z}_p$. Otherwise, set $\mu_j := \mu'_j$.
3. Output $\{\mu_j\}, \rho$.

Fig. 6. Pebbling-modified secret sharing scheme share^u

Hybrid Distribution \mathbf{H}^u We now define our hybrid games, and remark that Section 3 used $z \in \{0, 1\}^R$ to denote the adaptive choices made by an adversary, and the functions h_ℓ that define our hybrid games will depend on the adaptive choices of both the $f \in \text{NC}^1$ and $x \in \{0, 1\}^n$ chosen during the game, so in our application of the piecewise guessing framework of Section 3, z will be (f, x) . Note that the conclusion of the framework is independent of the size of the adaptive input ($R = |f| + n$), and the framework allows its x to be defined in parts over time, though in our application, x will be defined in one shot.

Definition 5 (\mathbf{H}^u and h_ℓ). Let \mathbf{H}^u be $\mathbf{G}_0^{1\text{-ABE}}$ with $\boxed{\text{share}^u}(f, \mu^{(0)})$ used in the implementation of oracle $\mathcal{O}_F(f)$ (replacing $\boxed{\text{share}}(f, \mu^{(0)})$). Let $h_\ell : \text{NC}^1 \times \{0, 1\}^n \rightarrow \{0, 1\}^{R'}$ denote the function that on formula f with root gate g^* and input $x \in \{0, 1\}^n$ where $f(x) = 0$, outputs the pebbling configuration created from following the first ℓ instructions from $\text{Pebble}(f, x, g^*)$ of Figure 2.

Note that the first 0 instructions specify a configuration with no pebbles, so h_0 is a constant function for all f, x . Also, from the inductive proof in Lemma 3, we know that all sequences of instructions from $\text{Pebble}(f, x, g^*)$ when $f(x) = 0$ result in a configuration with a single pebble on the root gate g^* , so h_L is a constant function for all f, x where $f(x) = 0$. Furthermore, note that for all such f, x :

- $\mathbf{H}^{h_0(f, x)}$ is equivalent to $\mathbf{G}_0^{1\text{-ABE}}$ (since $\text{share}^{h_0(f, x)}(f, \mu^{(0)}) = \text{share}(f, \mu^{(0)})$);
- $\mathbf{H}^{h_L(f, x)}$ is equivalent to $\mathbf{G}_1^{1\text{-ABE}}$ (since $\text{share}^{h_L(f, x)}(f, \mu^{(0)}) = \text{share}(f, \mu^{(1)})$ for an independently random $\mu^{(1)}$ which is implicitly defined by the independently random value associated with the output wire of the pebbled root gate: μ_m).

We now have a series of hybrids $\mathbf{G}_0^{1\text{-ABE}} \equiv \mathbf{H}^{h_0(f, x)}, \mathbf{H}^{h_1(f, x)}, \dots, \mathbf{H}^{h_L(f, x)} \equiv \mathbf{G}_1^{1\text{-ABE}}$ which satisfy end-point equivalence and, according to the piecewise guessing framework described in Section 3, define games $\widehat{\mathbf{H}}_{\ell, 0}(u_0, u_1), \widehat{\mathbf{H}}_{\ell, 1}(u_0, u_1)$ for $\ell \in [0, L]$.

Lemma 5 (neighboring indistinguishability). For all $\ell \in [L]$ and $u_0, u_1 \in \{0, 1\}^{R'}$,

$$\Pr[\langle \mathbf{A}, \widehat{\mathbf{H}}_{\ell, 0}(u_0, u_1) \rangle = 1] - \Pr[\langle \mathbf{A}, \widehat{\mathbf{H}}_{\ell, 1}(u_0, u_1) \rangle = 1] \leq n \cdot \text{Adv}_{\mathcal{B}}^{\text{CPA}}(\lambda)$$

Proof. First, observe that the difference between $\widehat{\mathbf{H}}_{\ell, 0}(u_0, u_1)$ and $\widehat{\mathbf{H}}_{\ell, 1}(u_0, u_1)$ lies in $\mathcal{O}_F(\cdot)$: the former uses $\boxed{\text{share}^{u_0}}$ and the latter, uses $\boxed{\text{share}^{u_1}}$. Now, fix the adaptive query f to \mathcal{O}_F . We consider two cases.

First, suppose there does not exist $x' \in \{0, 1\}^n$ such that $h_{\ell-1}(f, x') = u_0$ and $h_\ell(f, x') = u_1$. Then, both $\langle \mathbf{A}, \widehat{\mathbf{H}}_{\ell, 0}(u_0, u_1) \rangle$ and $\langle \mathbf{A}, \widehat{\mathbf{H}}_{\ell, 1}(u_0, u_1) \rangle$ output 0 (i.e., abort) with probability 1 and then we are done.

In the rest of the proof, we deal with the second case, namely there exists $x' \in \{0, 1\}^n$ such that $h_{\ell-1}(f, x') = u_0$ and $h_\ell(f, x') = u_1$. This means that u_0 and u_1 are neighboring pebbling configurations in $\text{Pebble}(f, x', g^*)$, so they differ by a pebbling instruction that follows one of the

rules in Definition 3. We proceed via a case analysis depending on what the instruction taking configuration u_0 to u_1 is (the instruction is uniquely determined given u_0, u_1, f):

- pebble/unpebble input node with out-going wire \boxed{j} : Here, the only difference from $\text{share}^{u_0}(f, \mu^{(0)})$ to $\text{share}^{u_1}(f, \mu^{(0)})$ is that we change $\boxed{\mu_j}$ to a random element of \mathbb{Z}_p (or vice-versa). The pebbling rule for an input node requires that the input x to $\mathcal{O}_X(\cdot)$ in both $\widehat{H}_{\ell,0}(u_0, u_1)$ and $\widehat{H}_{\ell,1}(u_0, u_1)$ satisfies $x_{\rho(j)} = 0$. Indistinguishability then follows from the CPA security of (CPA.Setup, CPA.Enc, CPA.Dec) under key $w_{\rho(j)}$; this is because $x_{\rho(j)} = 0$ and therefore $w_{\rho(j)}$ will not need to be supplied in the answer to the query to $\mathcal{O}_X(x)$. In fact, the two hybrids are computationally indistinguishable even if the adversary sees all $\{w_i : i \neq \rho(j)\}$ (as may be provided by $\mathcal{O}_X(x)$).
- pebble/unpebble AND gate with out-going wire \boxed{c} and input wires a, b corresponding to nodes g_a, g_b . Here, the only difference from $\text{share}^{u_0}(f, \mu^{(0)})$ to $\text{share}^{u_1}(f, \mu^{(0)})$ is that we change $\boxed{\mu_c}$ from an actual share $\hat{\mu}_a + \hat{\mu}_b + \hat{\mu}_c$ to a random element of \mathbb{Z}_p (or vice-versa). The pebbling rules for an AND gate require that there is a pebble on either g_a or g_b , say g_a . Therefore, μ_a is independent and uniformly random in both distributions $\text{share}^{u_0}(f, \mu^{(0)})$ and $\text{share}^{u_1}(f, \mu^{(0)})$, and thus $\hat{\mu}_a$ is fresh and independently random in both distributions (this uses the fact that g_a has fan-out 1) and makes the distribution of $\mu_c = \hat{\mu}_a + \hat{\mu}_b + \hat{\mu}_c$ in hybrid $\ell - 1$ independently random. We may then deduce that $\text{share}^{u_0}(f, \mu^{(0)})$ and $\text{share}^{u_1}(f, \mu^{(0)})$ are identically distributed, and therefore so is the output $\mathcal{O}_F(f)$. (This holds even if the adversary receives all of $\{w_i : i \in [n]\}$ from its query to $\mathcal{O}_X(x)$).
- pebble/unpebble OR gate with out-going wire \boxed{c} and input wires a, b corresponding to nodes g_a, g_b . Here, the only difference from $\text{share}^{u_0}(f, \mu^{(0)})$ to $\text{share}^{u_1}(f, \mu^{(0)})$ is that we change $\boxed{\mu_{c_a}, \mu_{c_b}}$ from actual shares $(\hat{\mu}_a + \hat{\mu}_c, \hat{\mu}_b + \hat{\mu}_c)$ to random elements of \mathbb{Z}_p (or vice-versa). The pebbling rules for an OR gate require that there are pebbles on both g_a and g_b . Therefore, μ_a and μ_b are independent and uniformly random in both distributions $\text{share}^{u_0}(f, \mu^{(0)})$ and $\text{share}^{u_1}(f, \mu^{(0)})$, and thus $\hat{\mu}_a, \hat{\mu}_b$ are fresh and independently random in both distributions (using the fact that g_a, g_b have fan-out 1), and make the distributions of $\mu_{c_a} = \hat{\mu}_a + \hat{\mu}_c$, $\mu_{c_b} = \hat{\mu}_b + \hat{\mu}_c$ in hybrid $\ell - 1$ both independently random. We may then deduce that $\text{share}^{u_0}(f, \mu^{(0)})$ and $\text{share}^{u_1}(f, \mu^{(0)})$ are identically distributed, and therefore so is the output $\mathcal{O}_F(f)$. (This holds even if the adversary receives all of $\{w_i : i \in [n]\}$ in its query to $\mathcal{O}_X(x)$).

In all cases, the simulator can return an appropriately distributed answer to $\mathcal{O}_X(x) = \{w_i\}_{x_i=1}$ since it has all w_i except in the first case, where it is missing only a w_i such that $x_i = 0$. Additionally, we note that in all cases, a simulator can return appropriately distributed answers to queries to the encryption oracle $\mathcal{O}_E(i, m) = \text{Enc}_{w_i}(m)$, since only in the first case (an input node being pebbled or unpebbled) is there a w_i not directly available to be used to simulate the oracle, and in that case, the simulator has oracle access to an $\text{Enc}_{w_i}(\cdot)$ function in the CPA symmetric-key security game, and it can uniformly guess which of the n variables is associated with the input node being pebbled and answer \mathcal{O}_E requests to that variable with the CPA $\text{Enc}_{w_i}(\cdot)$ oracle (the factor of n due to guessing is introduced here since the simulator may not know which variable is associated with the input node at the time of the oracle request, e.g: for requests to \mathcal{O}_E made before \mathcal{O}_X , so the simulator must guess uniformly and take a security loss of n).

In all but the input node case, the two distributions $\langle \mathbf{A}, \widehat{H}_{\ell,0}(u_0, u_1) \rangle$ and $\langle \mathbf{A}, \widehat{H}_{\ell,1}(u_0, u_1) \rangle$ are identical, and in the input node case, we've bounded the difference by the distinguishing probability of the symmetric key encryption scheme, the advantage function $\text{Adv}_{\mathcal{B}}^{\text{CPA}}(\lambda)$, conditioned on a correct guess of which of the n input variables corresponds to the pebbled/unpebbled input node.

Therefore,

$$\Pr[\langle \mathbf{A}, \widehat{\mathbf{H}}_{\ell,0}(u_0, u_1) \rangle = 1] - \Pr[\langle \mathbf{A}, \widehat{\mathbf{H}}_{\ell,1}(u_0, u_1) \rangle = 1] \leq n \cdot \text{Adv}_{\mathcal{B}}^{\text{CPA}}(\lambda)$$

□

5.4 CPA-secure symmetric encryption

We will instantiate (CPA.Setup, CPA.Enc, CPA.Dec) in our Core 1-ABE of Definition 4 with a variant of the standard CPA-secure symmetric encryption scheme based on k -Lin from [EHK⁺13] that supports messages $[M]_2 \in G_2$ of an asymmetric prime-order bilinear group \mathbb{G} :

$$\begin{aligned} \text{CPA.Setup}(1^\lambda) : & \text{Run } \mathbb{G} \leftarrow \mathcal{G}(1^\lambda). \text{ Sample } \mathbf{M}_0 \leftarrow \mathbb{Z}_p^{k \times k}, \mathbf{m}_1 \leftarrow \mathbb{Z}_p^k, \\ & \text{output } \text{sk} = (\text{sk}_0, \text{sk}_1) := (\mathbf{M}_0, \mathbf{m}_1^\top) \\ \text{CPA.Enc}(\text{sk}, [M]_2) : & \text{Sample } \mathbf{r} \leftarrow \mathbb{Z}_p^k, \text{ output } (\text{ct}_0, \text{ct}_1) := ([M + \mathbf{m}_1^\top \mathbf{r}]_2, [\mathbf{M}_0 \mathbf{r}]_2) \\ \text{CPA.Dec}((\text{sk}_0, \text{sk}_1), (\text{ct}_0, \text{ct}_1)) : & \text{Output } \text{ct}_0 \cdot \text{sk}_0 \cdot \text{sk}_1^{-1} \cdot \text{ct}_1. \end{aligned}$$

Correctness Note that: $\text{ct}_0 \cdot \text{sk}_1 \cdot \text{sk}_0^{-1} \cdot \text{ct}_1 = [M + \mathbf{m}_1^\top \mathbf{r} - \mathbf{m}_1^\top \mathbf{r}]_2 = [M]_2$.

Lemma 6. $\text{Adv}_{\mathcal{B}}^{\text{CPA}}(\lambda) \leq \text{Adv}_{\mathcal{B}'}^{k\text{-LIN}}(\lambda)$

Proof. Consider the following adversary \mathcal{B}^* , which when given MDDH $_k$ challenge: $(\mathbb{G}, [\mathbf{M}]_2, [\mathbf{z}]_2)$ first parses \mathbf{M} into $\mathbf{M}_0 :=$ the first k rows of \mathbf{M} and $\mathbf{m}_1^\top :=$ the $k+1$ th row of \mathbf{M} (here $[\mathbf{z}]_2 = [(\begin{smallmatrix} \mathbf{M}_0 \mathbf{s} \\ \mathbf{m}_1^\top \cdot \mathbf{s} + r \end{smallmatrix})]_2$ where either $r \leftarrow \mathbb{Z}_p$ or $r = 0$) and prepares the following challenge ciphertext using $[\mathbf{z}]_2$ from the MDDH $_k$ challenge: $(\text{ct}_0, \text{ct}_1) := ([M_b + \mathbf{m}_1^\top \mathbf{s} + r]_2, [\mathbf{M}_0 \mathbf{s}]_2)$.

To answer CPA encryption oracle queries for messages $[M]_2$, draw $\mathbf{r} \leftarrow \mathbb{Z}_p^k$ and return $(\text{ct}_0, \text{ct}_1) = ([M + \mathbf{m}_1^\top \mathbf{r}]_2, [\mathbf{M}_0 \mathbf{r}]_2)$. Note that this is an appropriately distributed encryption and can be computed using the MDDH $_k$ challenge element $[\mathbf{M}]_2$.

If $r = 0$, then the challenge ciphertext is distributed as a normal encryption of $[M]_2$ with $\mathbf{r} = \mathbf{s}$. This is the normal CPA security game.

If $r \leftarrow \mathbb{Z}_p$, then the challenge ciphertext is distributed as an encryption of a random message $[M + r]_2$ with $\mathbf{r} = \mathbf{s}$. In this game the adversary's advantage is 0.

Therefore, we have: $\text{Adv}_{\mathcal{B}}^{\text{CPA}}(\lambda) \leq \text{Adv}_{\mathcal{B}^*}^{\text{MDDH}_k}(\lambda)$. From Section 2.4 we have that: $\text{Adv}_{\mathcal{B}^*}^{\text{MDDH}_k}(\lambda) = \text{Adv}_{\mathcal{B}'}^{k\text{-LIN}}(\lambda)$. So, we have: $\text{Adv}_{\mathcal{B}}^{\text{CPA}}(\lambda) \leq \text{Adv}_{\mathcal{B}'}^{k\text{-LIN}}(\lambda)$.

Theorem 2. *The Core 1-ABE component of Definition 4 implemented with (share, reconstruct) from Section 5.1 and the CPA-secure symmetric encryption scheme (CPA.Setup, CPA.Enc, CPA.Dec) from Section 5.4 satisfies:*

$$\Pr[\langle \mathbf{A}, \mathbf{G}_0^{1\text{-ABE}} \rangle = 1] - \Pr[\langle \mathbf{A}, \mathbf{G}_1^{1\text{-ABE}} \rangle = 1] \leq 2^{6d} \cdot 8^d \cdot n \cdot \text{Adv}_{\mathcal{B}^*}^{k\text{-LIN}}(\lambda)$$

Proof. Recall the hybrids $\mathbf{G}_0^{1\text{-ABE}} \equiv \mathbf{H}^{h_0(f,x)}, \mathbf{H}^{h_1(f,x)}, \dots, \mathbf{H}^{h_L(f,x)} \equiv \mathbf{G}_1^{1\text{-ABE}}$ defined in Section 5.3. Lemma 5 tells us that: for all $\ell \in [L]$ and $u_0, u_1 \in \{0, 1\}^{R'}$,

$$\Pr[\langle \mathbf{A}, \widehat{\mathbf{H}}_{\ell,0}(u_0, u_1) \rangle = 1] - \Pr[\langle \mathbf{A}, \widehat{\mathbf{H}}_{\ell,1}(u_0, u_1) \rangle = 1] \leq n \cdot \text{Adv}_{\mathcal{B}}^{\text{CPA}}(\lambda)$$

These hybrids satisfy the end-point equivalence requirement, so Lemma 1 then tells us that:

$$\Pr[\langle \mathbf{A}, \mathbf{G}_0^{1\text{-ABE}} \rangle = 1] - \Pr[\langle \mathbf{A}, \mathbf{G}_1^{1\text{-ABE}} \rangle = 1] \leq 2^{2R'} \cdot L \cdot n \cdot \text{Adv}_{\mathcal{B}}^{\text{CPA}}(\lambda)$$

Lemma 4 tells us that $R' \leq 3d$, and Lemma 2 tells us that $L \leq 8^d$, where d is the depth of the formula. Finally, Lemma 6 tells us that $\text{Adv}_{\mathcal{B}}^{\text{CPA}}(\lambda) \leq \text{Adv}_{\mathcal{B}^*}^{k\text{-LIN}}(\lambda)$. So:

$$\Pr[\langle \mathbf{A}, \mathbf{G}_0^{1\text{-ABE}} \rangle = 1] - \Pr[\langle \mathbf{A}, \mathbf{G}_1^{1\text{-ABE}} \rangle = 1] \leq 2^{6d} \cdot 8^d \cdot n \cdot \text{Adv}_{\mathcal{B}^*}^{k\text{-LIN}}(\lambda)$$

□

6 Our KP-ABE Scheme

In this section, we present our compact KP-ABE for NC^1 that is adaptively secure under the MDDH_k assumption in asymmetric prime-order bilinear groups. For attributes of length n , our ciphertext comprises $O(n)$ group elements, independent of the formula size, while simultaneously allowing attribute reuse in the formula. As mentioned in the overview in Section 1.1, we incorporated optimizations from [GDCC16,BKP14] to shrink \mathbf{W}_i and thus the secret key, and hence the need for the \mathcal{O}_E oracle in the core 1-ABE security game.

6.1 The scheme

Our KP-ABE scheme is as follows:

$\text{Setup}(1^\lambda, 1^n)$: Run $\mathbb{G} = (p, G_1, G_2, G_T, e) \leftarrow \mathcal{G}(1^\lambda)$. Sample

$$\mathbf{A} \leftarrow \mathbb{Z}_p^{k \times (k+1)}, \mathbf{W}_i \leftarrow \mathbb{Z}_p^{(k+1) \times k} \forall i \in [n], \mathbf{v} \leftarrow \mathbb{Z}_p^{k+1}$$

and output:

$$\begin{aligned} \text{msk} &:= (\mathbf{v}, \mathbf{W}_1, \dots, \mathbf{W}_n) \\ \text{mpk} &:= ([\mathbf{A}]_1, [\mathbf{A}\mathbf{W}_1]_1, \dots, [\mathbf{A}\mathbf{W}_n]_1, e([\mathbf{A}]_1, [\mathbf{v}]_2)) \end{aligned}$$

$\text{Enc}(\text{mpk}, x, M)$: Sample $\mathbf{s} \leftarrow \mathbb{Z}_p^k$. Output:

$$\begin{aligned} \text{ct}_x &= (\text{ct}_1, \{\text{ct}_{2,i}\}_{x_i=1}, \text{ct}_3) \\ &:= \left([\mathbf{s}^\top \mathbf{A}]_1, \{[\mathbf{s}^\top \mathbf{A}\mathbf{W}_i]_1\}_{x_i=1}, e([\mathbf{s}^\top \mathbf{A}]_1, [\mathbf{v}]_2) \cdot M \right) \end{aligned}$$

$\text{KeyGen}(\text{mpk}, \text{msk}, f)$: Sample $(\{\mathbf{v}_j\}, \rho) \leftarrow \text{share}(f, \mathbf{v})$, $\mathbf{r}_j \leftarrow \mathbb{Z}_p^k$. Output:

$$\begin{aligned} \text{sk}_f &= (\{\text{sk}_{1,j}, \text{sk}_{2,j}\}) \\ &:= (\{[\mathbf{v}_j + \mathbf{W}_{\rho(j)} \mathbf{r}_j]_2, [\mathbf{r}_j]_2\}) \end{aligned}$$

where $\mathbf{W}_0 = \mathbf{0}$.

$\text{Dec}(\text{mpk}, \text{sk}_f, \text{ct}_x)$: Compute ω_j such that $\mathbf{v} = \sum_{\rho(j)=0 \vee x_{\rho(j)}=1} \omega_j \mathbf{v}_j$ as described in Section 5.1.

Output:

$$\text{ct}_3 \cdot \prod_{\rho(j)=0 \vee x_{\rho(j)}=1} \left(\frac{e(\text{ct}_{2,\rho(j)}, \text{sk}_{2,j})}{e(\text{ct}_1, \text{sk}_{1,j})} \right)^{\omega_j}$$

6.2 Correctness

Correctness relies on the fact that for all j , we have

$$\frac{e(\text{ct}_1, \text{sk}_{1,j})}{e(\text{ct}_{2,\rho(j)}, \text{sk}_{2,j})} = [\mathbf{s}^\top \mathbf{A} \mathbf{v}_j]_T$$

which follows from the fact that

$$\mathbf{s}^\top \mathbf{A} \mathbf{v}_j = \underbrace{\mathbf{s}^\top \mathbf{A}}_{\text{ct}_1} \cdot \underbrace{(\mathbf{v}_j + \mathbf{W}_{\rho(j)} \mathbf{r}_j)}_{\text{sk}_{1,j}} - \underbrace{\mathbf{s}^\top \mathbf{A} \mathbf{W}_{\rho(j)}}_{\text{ct}_{2,\rho(j)}} \cdot \underbrace{\mathbf{r}_j}_{\text{sk}_{2,j}}$$

Therefore, for all f, x such that $f(x) = 1$, we have:

$$\begin{aligned}
\text{ct}_3 \cdot \prod_{\rho(j)=0 \vee x_{\rho(j)}=1} \left(\frac{e(\text{ct}_{2,\rho(j)}, \text{sk}_{2,j})}{e(\text{ct}_{1,\rho(j)}, \text{sk}_{1,j})} \right)^{\omega_j} &= M \cdot [\mathbf{s}^\top \mathbf{A} \mathbf{v}]_T \cdot \prod_{\rho(j)=0 \vee x_{\rho(j)}=1} [\mathbf{s}^\top \mathbf{A} \mathbf{v}_j]_T^{-\omega_j} \\
&= M \cdot [\mathbf{s}^\top \mathbf{A} \mathbf{v}]_T \cdot [-\mathbf{s}^\top \mathbf{A} \sum_{\rho(j)=0 \vee x_{\rho(j)}=1} \omega_j \mathbf{v}_j]_T \\
&= M \cdot [\mathbf{s}^\top \mathbf{A} \mathbf{v}]_T \cdot [-\mathbf{s}^\top \mathbf{A} \mathbf{v}]_T \\
&= M
\end{aligned}$$

6.3 Adaptive Security

Description of hybrids To describe the hybrid distributions, it would be helpful to first give names to the various forms of ciphertext and keys that will be used. A ciphertext can be in one of the following forms:

- Normal: generated as in the scheme.
- SF: same as a Normal ciphertext, except $\mathbf{s}^\top \mathbf{A}$ replaced with $\mathbf{c}^\top \leftarrow \mathbb{Z}_p^{k+1}$. That is,

$$\text{ct}_x := \left(\boxed{\mathbf{c}^\top}_1, \{ \boxed{\mathbf{c}^\top} \mathbf{W}_i \}_1, e(\boxed{\mathbf{c}^\top}_1, [\mathbf{v}]_2) \cdot M \right)$$

A secret key can be in one of the following forms:

- Normal: generated as in the scheme.
- SF: same as a Normal key, except \mathbf{v} replaced with $\mathbf{v} + \delta \mathbf{a}^\perp$, where a fresh $\delta \leftarrow \mathbb{Z}_p$ is chosen per SF key and \mathbf{a}^\perp is any fixed $\mathbf{a}^\perp \in \mathbb{Z}_p^{k+1} \setminus \{\mathbf{0}\}$ such that $\mathbf{A} \mathbf{a}^\perp = \mathbf{0}$. That is,

$$\text{sk}_f := (\{ [\mathbf{v}_j + \mathbf{W}_{\rho(j)} \mathbf{r}_j]_2, [\mathbf{r}_j]_2 \})$$

where $(\{\mathbf{v}_j\}, \rho) \leftarrow \text{share}(f, \boxed{\mathbf{v} + \delta \mathbf{a}^\perp})$, $\mathbf{r}_j \leftarrow \mathbb{Z}_p^k$.

SF stands for semi-functional following the terminology in previous works [LW10, Wat09].

Hybrid sequence. Suppose the adversary \mathbf{A} makes at most Q secret key queries. The hybrid sequence is as follows:

- \mathbf{H}_0 : real game
- \mathbf{H}_1 : same as \mathbf{H}_0 , except we use a SF ciphertext.
- $\mathbf{H}_{2,\ell}$, $\ell = 0, \dots, Q$: same as \mathbf{H}_1 , except the first ℓ keys are SF and the remaining $Q - \ell$ keys are Normal.
- \mathbf{H}_3 : replace M with random \widetilde{M} .

Proof overview.

- We have $\mathbf{H}_0 \approx_c \mathbf{H}_1 \equiv \mathbf{H}_{2,0}$ via k -Lin, which tells us $([\mathbf{A}]_1, [\mathbf{s}^\top \mathbf{A}]_1) \approx_c ([\mathbf{A}]_1, [\mathbf{c}^\top]_1)$. Here, the security reduction will pick $\mathbf{W}_1, \dots, \mathbf{W}_n$ and \mathbf{v} so that it can simulate the mpk, the ciphertext and the secret keys.

- We have $H_{2,\ell-1} \approx_c H_{2,\ell}$, for all $\ell \in [Q]$. The difference between the two is that we switch the ℓ 'th sk_f from Normal to SF using the adaptive security of our core 1-ABE component in $\mathbf{G}^{1\text{-ABE}}$ from Section 5. The idea is to sample

$$\mathbf{v} = \tilde{\mathbf{v}} + \mu \mathbf{a}^\perp, \mathbf{W}_i = \widetilde{\mathbf{W}}_i + \mathbf{a}^\perp \mathbf{w}_i^\top$$

so that mpk can be computed using $\tilde{\mathbf{v}}, \widetilde{\mathbf{W}}_i$ and perfectly hide $\mu, \mathbf{w}_1, \dots, \mathbf{w}_n$. Roughly speaking: the reduction

- uses $\mathcal{O}_X(x)$ in $\mathbf{G}^{1\text{-ABE}}$ to simulate the challenge ciphertext
 - uses $\mathcal{O}_F(f)$ in $\mathbf{G}^{1\text{-ABE}}$ to simulate ℓ 'th secret key
 - uses $\mu^{(0)}$ from $\mathbf{G}^{1\text{-ABE}}$ together with $\mathcal{O}_E(i, \cdot) = \text{Enc}(w_i, \cdot)$ to simulate the remaining $Q - \ell$ secret keys
- We have $H_{2,Q} \equiv H_3$. In $H_{2,Q}$, the secret keys only leak $\mathbf{v} + \delta_1 \mathbf{a}^\perp, \dots, \mathbf{v} + \delta_Q \mathbf{a}^\perp$. This means that $\mathbf{c}^\top \mathbf{v}$ is statistically random (as long as $\mathbf{c}^\top \mathbf{a}^\perp \neq 0$).

Lemma 7 ($H_0 \approx_c H_1 \equiv H_{2,0}$).

$$|\Pr[\langle \mathcal{A}, H_0 \rangle = 1] - \Pr[\langle \mathcal{A}, H_1 \rangle = 1]| \leq \text{Adv}_{\mathcal{A}}^{k\text{-LIN}}(\lambda)$$

Proof. Given MDDH $_k$ challenge ($[\mathbf{A}]_1, [\mathbf{z}^\top]_1$), where either $\mathbf{z}^\top = \mathbf{s}^\top \mathbf{A}$ or $\mathbf{z}^\top = \mathbf{c}^\top$, an adversary \mathcal{A}' could simply choose $\mathbf{W}_i \leftarrow \mathbb{Z}_p^{(k+1) \times k}, \mathbf{v} \leftarrow \mathbb{Z}_p^{k+1}$, form the public parameters with $\mathbf{A}, \mathbf{W}_i, \mathbf{v}$, and choose its own $\mathbf{r}_j \leftarrow \mathbb{Z}_p^k$ when responding to key requests. For the challenge ciphertext, \mathcal{A}' creates:

$$\text{ct}_{\mathbf{x}} := \left([\mathbf{z}^\top]_1, \{[\mathbf{z}^\top \mathbf{W}_i]_1\}_{x_i=1}, e([\mathbf{z}^\top]_1, [\mathbf{v}]_2) \cdot M \right)$$

If $\mathbf{z}^\top = \mathbf{s}^\top \mathbf{A}$, \mathcal{A}' has simulated H_0 ; If $\mathbf{z}^\top = \mathbf{c}^\top$, \mathcal{A}' has simulated $H_1 \equiv H_{2,0}$. \square

Lemma 8 ($H_{2,\ell-1} \approx_c H_{2,\ell}$).

$$|\Pr[\langle \mathcal{A}, H_{2,\ell-1} \rangle = 1] - \Pr[\langle \mathcal{A}, H_{2,\ell} \rangle = 1]| \leq 2^{6d} \cdot 8^d \cdot n \cdot \text{Adv}_{\beta^*}^{k\text{-LIN}}(\lambda)$$

Proof. For each $\beta \in \{0, 1\}$, consider the following adversary \mathcal{A}' in $\mathbf{G}_\beta^{1\text{-ABE}}$ which internally simulates \mathcal{A} and the challenger in the ABE security game:

- First, \mathcal{A}' samples $\mathbf{A} \leftarrow \mathbb{Z}_p^{k \times (k+1)}, \widetilde{\mathbf{W}}_i \leftarrow \mathbb{Z}_p^{(k+1) \times k}, \tilde{\mathbf{v}} \leftarrow \mathbb{Z}_p^{k+1}$, computes $\mathbf{a}^\perp \in \mathbb{Z}_p^{k+1} \setminus \{\mathbf{0}\}$ such that $\mathbf{A} \mathbf{a}^\perp = \mathbf{0}$ and implicitly defines

$$\mathbf{v} := \tilde{\mathbf{v}} + \mu^{(0)} \mathbf{a}^\perp, \mathbf{W}_i := \widetilde{\mathbf{W}}_i + \mathbf{a}^\perp \mathbf{w}_i^\top$$

where $\mathbf{w}_i \in \mathbb{Z}_p^k, \mu^{(0)} \in \mathbb{Z}_p$ are chosen in $\mathbf{G}_\beta^{1\text{-ABE}}$. Then, \mathcal{A}' outputs:

$$\text{mpk} := ([\mathbf{A}]_1, [\mathbf{A} \widetilde{\mathbf{W}}_1]_1, \dots, [\mathbf{A} \widetilde{\mathbf{W}}_n]_1, e([\mathbf{A}]_1, [\tilde{\mathbf{v}}]_2))$$

- When \mathcal{A} requests a challenge ciphertext for attribute \mathbf{x} along with M_0, M_1 , \mathcal{A}' queries $\mathcal{O}_X(\mathbf{x}) \rightarrow (\{\mathbf{w}_i\}_{x_i=1})$ in $\mathbf{G}_\beta^{1\text{-ABE}}$. \mathcal{A}' then samples $\mathbf{c} \leftarrow \mathbb{Z}_p^{k+1}$ and $b \leftarrow \{0, 1\}$ (the challenge bit in the standard ABE security game) and returns the following (SF) challenge ciphertext for \mathcal{A} :

$$\text{ct}_{\mathbf{x}} = \left([\mathbf{c}^\top]_1, \{[\mathbf{c}^\top \underbrace{(\widetilde{\mathbf{W}}_i + \mathbf{a}^\perp \mathbf{w}_i^\top)}_{=\mathbf{W}_i}]_1\}_{x_i=1}, e([\mathbf{c}^\top]_1, \underbrace{[\tilde{\mathbf{v}} + \mu^{(0)} \mathbf{a}^\perp]_2}_{=\mathbf{v}}) \cdot M_b \right)$$

- For the first $\ell - 1$ secret keys requested, say for formula f , \mathcal{A}' computes

$$(\{\mathbf{v}_j\}, \rho) \leftarrow \text{share}(f, \underbrace{\tilde{\mathbf{v}} + \tilde{\delta} \mathbf{a}^\perp}_{=\mathbf{v} + \delta \mathbf{a}^\perp})$$

where $\tilde{\delta} \leftarrow \mathbb{Z}_p$ is drawn independently for each key (here, the per-key $\delta = \tilde{\delta} - \mu^{(0)}$ implicitly). Next, for each j , it queries $\mathcal{O}_E(\rho(j), [0]_2) \rightarrow ([\mathbf{w}_{\rho(j)}^\top \mathbf{r}_j]_2, [\mathbf{r}_j]_2)$ in $\mathbf{G}_\beta^{1\text{-ABE}}$ (since $\mathcal{O}_E(\rho(j), [0]_2) = \text{CPA.Enc}_{\mathbf{w}_{\rho(j)}}([0]_2)$), and forms the following (SF) key:

$$\text{sk}_f = (\{ \underbrace{[\mathbf{v}_j + \widetilde{\mathbf{W}}_{\rho(j)} \mathbf{r}_j + \mathbf{a}^\perp \mathbf{w}_{\rho(j)}^\top \mathbf{r}_j]_2, [\mathbf{r}_j]_2}_{=\mathbf{v}_j + \mathbf{W}_{\rho(j)} \mathbf{r}_j} \})$$

- For the last $Q - \ell$ secret keys requested, say for formula f , \mathcal{A}' proceeds as before for the first $\ell - 1$ keys except

$$(\{\mathbf{v}_j\}, \rho) \leftarrow \text{share}(f, \underbrace{\tilde{\mathbf{v}} + \mu^{(0)} \mathbf{a}^\perp}_{=\mathbf{v}})$$

It is easy to see that it forms a Normal key.

- For the ℓ th secret key requested, say for formula f , \mathcal{A}' computes $(\{\mathbf{v}_j\}, \rho) \leftarrow \text{share}(f, \tilde{\mathbf{v}})$, queries $\mathcal{O}_F(f) \rightarrow (\{ [\mu_j + \mathbf{w}_{\rho(j)}^\top \mathbf{r}_j]_2, [\mathbf{r}_j]_2 \})$ in $\mathbf{G}_\beta^{1\text{-ABE}}$, then uses these components to return:

$$\text{sk}_f = (\{ \underbrace{[\mathbf{v}_j + \widetilde{\mathbf{W}}_{\rho(j)} \mathbf{r}_j + \mathbf{a}^\perp (\mu_j + \mathbf{w}_{\rho(j)}^\top \mathbf{r}_j)]_2, [\mathbf{r}_j]_2}_{=(\mathbf{v}_j + \mu_j \mathbf{a}^\perp) + \mathbf{W}_j \mathbf{r}_j} \})$$

We claim that if $\beta = 0$, then sk_f is a Normal key, and if $\beta = 1$, then sk_f is a SF key. This follows the fact that thanks to linearity, the shares

$$(\{\mathbf{v}_j + \mu_j \mathbf{a}^\perp\}, \rho), \text{ where } (\{\mathbf{v}_j\}, \rho) \leftarrow \text{share}(f, \tilde{\mathbf{v}}), (\{\mu_j\}, \rho) \leftarrow \text{share}(f, \mu^{(\beta)})$$

are identically distributed to $\text{share}(f, \tilde{\mathbf{v}} + \mu^{(\beta)} \mathbf{a}^\perp)$. The claim then follows from the fact that $\tilde{\mathbf{v}} + \mu^{(0)} \mathbf{a}^\perp = \mathbf{v}$ and that $\tilde{\mathbf{v}} + \mu^{(1)} \mathbf{a}^\perp$ is identically distributed to $\mathbf{v} + \delta \mathbf{a}^\perp$ (where $\delta = \mu^{(1)} - \mu^{(0)}$ is a fresh random value for this key).

Putting everything together, for $\beta \in \{0, 1\}$, when \mathcal{A}' interacts with $\mathbf{G}_\beta^{1\text{-ABE}}$, then \mathcal{A}' simulates $\mathbf{H}_{2, \ell-1+\beta}$. It follows then that:

$$|\Pr[\langle \mathcal{A}, \mathbf{H}_{2, \ell-1} \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_{2, \ell} \rangle = 1]| \leq |\Pr[\langle \mathcal{A}', \mathbf{G}_0^{1\text{-ABE}} \rangle = 1] - \Pr[\langle \mathcal{A}', \mathbf{G}_1^{1\text{-ABE}} \rangle = 1]|$$

From Theorem 2, we then have:

$$|\Pr[\langle \mathcal{A}, \mathbf{H}_{2, \ell-1} \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_{2, \ell} \rangle = 1]| \leq 2^{6d} \cdot 8^d \cdot n \cdot \text{Adv}_{\mathcal{B}^*}^{k\text{-LIN}}(\lambda)$$

□

Lemma 9 ($\mathbf{H}_{2, Q} \approx_s \mathbf{H}_3$).

$$|\Pr[\langle \mathcal{A}, \mathbf{H}_{2, Q} \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_3 \rangle = 1]| \leq \frac{1}{p}$$

Proof. These two hybrids are identically distributed conditioned on $\mathbf{c}^\top \mathbf{a}^\perp \neq 0$. To see this, consider two ways of sampling \mathbf{v} : as $\tilde{\mathbf{v}} \leftarrow \mathbb{Z}_p^{k+1}$ and as $\tilde{\mathbf{v}} + \tilde{m} \mathbf{a}^\perp$ for an independent $\tilde{m} \leftarrow \mathbb{Z}_p$. Note that both result in \mathbf{v} having a uniform distribution.

Using $\tilde{\mathbf{v}}$ to simulate hybrid $\mathbf{H}_{2,Q}$ obviously results in $\mathbf{H}_{2,Q}$ (where $\mathbf{v} = \tilde{\mathbf{v}}$). However, using the identically distributed $\mathbf{v} = \tilde{\mathbf{v}} + \tilde{m} \mathbf{a}^\perp$ to simulate $\mathbf{H}_{2,Q}$ results in \mathbf{H}_3 (where $\tilde{M} = M \cdot e([\mathbf{c}^\top]_1, [\tilde{m} \mathbf{a}^\perp]_2)$ is randomly distributed as long as $\mathbf{c}^\top \mathbf{a}^\perp \neq 0$, and for redefined independently random $\tilde{\delta}_i := \delta_i + \tilde{m}$ in the secret keys).

\mathbf{c} is chosen at random and independent from $\mathbf{a}^\perp \neq \mathbf{0}$, so $\mathbf{c}^\top \mathbf{a}^\perp = 0$ with probability $\frac{1}{p}$, and since we know that $\mathbf{H}_{2,Q} \equiv \mathbf{H}_3$ conditioned on $\mathbf{c}^\top \mathbf{a}^\perp \neq 0$, then we have:

$$|\Pr[\langle \mathcal{A}, \mathbf{H}_{2,Q} \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_3 \rangle = 1]| \leq \frac{1}{p}$$

□

Theorem 3 (adaptive KP-ABE). *The KP-ABE construction in Section 6.1 is adaptively secure under the $MDDH_k$ assumption.*

Proof.

$$\begin{aligned} |\Pr[\langle \mathcal{A}, \mathbf{H}_0 \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_3 \rangle = 1]| &\leq |\Pr[\langle \mathcal{A}, \mathbf{H}_0 \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_1 \rangle = 1]| \\ &\quad + \sum_{\ell=1}^Q |\Pr[\langle \mathcal{A}, \mathbf{H}_{2,\ell-1} \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_{2,\ell} \rangle = 1]| \\ &\quad + |\Pr[\langle \mathcal{A}, \mathbf{H}_{2,Q} \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_3 \rangle = 1]| \end{aligned}$$

(Since $\mathbf{H}_1 \equiv \mathbf{H}_{2,0}$). Summing the results of Lemmas 7, 8, and 9, we then have:

$$|\Pr[\langle \mathcal{A}, \mathbf{H}_0 \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_3 \rangle = 1]| \leq \text{Adv}_{\mathcal{B}^*}^{k\text{-LIN}}(\lambda) + Q \cdot 2^{6d} \cdot 8^d \cdot n \cdot \text{Adv}_{\mathcal{B}^*}^{k\text{-LIN}}(\lambda) + \frac{1}{p}$$

If $d = O(\log n)$, then under the k -Lin Assumption this is a negligible function of λ (the number of queries made Q and the attribute vector length n are both polynomial in λ , and $\frac{1}{p}$ is a negligible function of λ). It's easy to see that $\text{Adv}_{\mathcal{A}}^{\text{ABE}}(\lambda) = 0$ in the \mathbf{H}_3 hybrid game (since a random message is encrypted in the challenge ciphertext). So, any adversary in the real game (\mathbf{H}_0) will have advantage negligibly close to 0, and our construction satisfies adaptive security. □

Acknowledgments. We thank Allison Bishop, Sanjam Garg, Rocco Servedio, and Daniel Wichs for helpful discussions.

References

- AC17. Shashank Agrawal and Melissa Chase. Simplifying design and analysis of complex predicate encryption schemes. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part I*, volume 10210 of *LNCS*, pages 627–656. Springer, Heidelberg, April / May 2017.
- Att14. Nuttapon Attrapadung. Dual system encryption via doubly selective security: Framework, fully secure functional encryption for regular languages, and more. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 557–577. Springer, Heidelberg, May 2014.
- Att16. Nuttapon Attrapadung. Dual system encryption framework in prime-order groups via computational pair encodings. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 591–623. Springer, Heidelberg, December 2016.
- BKP14. Olivier Blazy, Eike Kiltz, and Jiaxin Pan. (Hierarchical) identity-based encryption from affine message authentication. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 408–425. Springer, Heidelberg, August 2014.
- BSW07. John Bethencourt, Amit Sahai, and Brent Waters. Ciphertext-policy attribute-based encryption. In *2007 IEEE Symposium on Security and Privacy*, pages 321–334. IEEE Computer Society Press, May 2007.
- CGKW18. Jie Chen, Junqing Gong, Lucas Kowalczyk, and Hoeteck Wee. Unbounded ABE via bilinear entropy expansion, revisited. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 503–534. Springer, Heidelberg, April / May 2018.
- CGW15. Jie Chen, Romain Gay, and Hoeteck Wee. Improved dual system ABE in prime-order groups via predicate encodings. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 595–624. Springer, Heidelberg, April 2015.
- Che06. Jung Hee Cheon. Security analysis of the strong Diffie-Hellman problem. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 1–11. Springer, Heidelberg, May / June 2006.
- CW13. Jie Chen and Hoeteck Wee. Fully, (almost) tightly secure IBE and dual system groups. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 435–460. Springer, Heidelberg, August 2013.
- CW14. Jie Chen and Hoeteck Wee. Semi-adaptive attribute-based encryption and improved delegation for Boolean formula. In Michel Abdalla and Roberto De Prisco, editors, *SCN 14*, volume 8642 of *LNCS*, pages 277–297. Springer, Heidelberg, September 2014.
- EHK⁺13. Alex Escala, Gottfried Herold, Eike Kiltz, Carla Ràfols, and Jorge Villar. An algebraic framework for Diffie-Hellman assumptions. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 129–147. Springer, Heidelberg, August 2013.
- FJP15. Georg Fuchsbauer, Zahra Jafargholi, and Krzysztof Pietrzak. A quasipolynomial reduction for generalized selective decryption on trees. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 601–620. Springer, Heidelberg, August 2015.
- FKPR14. Georg Fuchsbauer, Momchil Konstantinov, Krzysztof Pietrzak, and Vanishree Rao. Adaptive security of constrained PRFs. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 82–101. Springer, Heidelberg, December 2014.
- GDCC16. Junqing Gong, Xiaolei Dong, Jie Chen, and Zhenfu Cao. Efficient IBE with tight reduction to standard assumption in the multi-challenge setting. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 624–654. Springer, Heidelberg, December 2016.
- GGH⁺13. Sanjam Garg, Craig Gentry, Shai Halevi, Amit Sahai, and Brent Waters. Attribute-based encryption for circuits from multilinear maps. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 479–499. Springer, Heidelberg, August 2013.
- GPSW06. Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 06*, pages 89–98. ACM Press, October / November 2006. Available as Cryptology ePrint Archive Report 2006/309.
- GVW13. Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 545–554. ACM Press, June 2013.
- HJO⁺16. Brett Hemenway, Zahra Jafargholi, Rafail Ostrovsky, Alessandra Scafuro, and Daniel Wichs. Adaptively secure garbled circuits from one-way functions. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 149–178. Springer, Heidelberg, August 2016.
- IK02. Yuval Ishai and Eyal Kushilevitz. Perfect constant-round secure computation via perfect randomizing polynomials. In Peter Widmayer, Francisco Triguero Ruiz, Rafael Morales Bueno, Matthew Hennessy, Stephan Eidenbenz, and Ricardo Conejo, editors, *ICALP 2002*, volume 2380 of *LNCS*, pages 244–256. Springer, Heidelberg, July 2002.

- JKK⁺17. Zahra Jafargholi, Chethan Kamath, Karen Klein, Ilan Komargodski, Krzysztof Pietrzak, and Daniel Wichs. Be adaptive, avoid overcommitting. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 133–163. Springer, Heidelberg, August 2017.
- JW16. Zahra Jafargholi and Daniel Wichs. Adaptive security of Yao’s garbled circuits. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part I*, volume 9985 of *LNCS*, pages 433–458. Springer, Heidelberg, October / November 2016.
- LOS⁺10. Allison B. Lewko, Tatsuaki Okamoto, Amit Sahai, Katsuyuki Takashima, and Brent Waters. Fully secure functional encryption: Attribute-based encryption and (hierarchical) inner product encryption. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 62–91. Springer, Heidelberg, May / June 2010.
- LSW10. Allison B. Lewko, Amit Sahai, and Brent Waters. Revocation systems with very small private keys. In *2010 IEEE Symposium on Security and Privacy*, pages 273–285. IEEE Computer Society Press, May 2010.
- LW10. Allison B. Lewko and Brent Waters. New techniques for dual system encryption and fully secure HIBE with short ciphertexts. In Daniele Micciancio, editor, *TCC 2010*, volume 5978 of *LNCS*, pages 455–479. Springer, Heidelberg, February 2010.
- LW11. Allison B. Lewko and Brent Waters. Unbounded HIBE and attribute-based encryption. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 547–567. Springer, Heidelberg, May 2011.
- LW12. Allison B. Lewko and Brent Waters. New proof methods for attribute-based encryption: Achieving full security through selective techniques. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 180–198. Springer, Heidelberg, August 2012.
- OSW07. Rafail Ostrovsky, Amit Sahai, and Brent Waters. Attribute-based encryption with non-monotonic access structures. In Peng Ning, Sabrina De Capitani di Vimercati, and Paul F. Syverson, editors, *ACM CCS 07*, pages 195–203. ACM Press, October 2007.
- OT10. Tatsuaki Okamoto and Katsuyuki Takashima. Fully secure functional encryption with general relations from the decisional linear assumption. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 191–208. Springer, Heidelberg, August 2010.
- OT12. Tatsuaki Okamoto and Katsuyuki Takashima. Fully secure unbounded inner-product and attribute-based encryption. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 349–366. Springer, Heidelberg, December 2012.
- PRV12. Bryan Parno, Mariana Raykova, and Vinod Vaikuntanathan. How to delegate and verify in public: Verifiable computation from attribute-based encryption. In Ronald Cramer, editor, *TCC 2012*, volume 7194 of *LNCS*, pages 422–439. Springer, Heidelberg, March 2012.
- SW05. Amit Sahai and Brent R. Waters. Fuzzy identity-based encryption. In Ronald Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 457–473. Springer, Heidelberg, May 2005.
- VNS⁺03. V. Vinod, Arvind Narayanan, K. Srinathan, C. Pandu Rangan, and Kwangjo Kim. On the power of computational secret sharing. In Thomas Johansson and Subhamoy Maitra, editors, *INDOCRYPT 2003*, volume 2904 of *LNCS*, pages 162–176. Springer, Heidelberg, December 2003.
- Wat09. Brent Waters. Dual system encryption: Realizing fully secure IBE and HIBE under simple assumptions. In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 619–636. Springer, Heidelberg, August 2009.
- Wee14. Hoeteck Wee. Dual system encryption via predicate encodings. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 616–637. Springer, Heidelberg, February 2014.

A Our CP-ABE Scheme

In this section, we present our compact CP-ABE for NC^1 that is adaptively secure under the MDDH_k assumption in asymmetric prime-order bilinear groups. The construction is analogous to our KP-ABE scheme in Section 6. One notable difference is that we introduce \mathbf{Br} in the secret keys, and we need to introduce additional intermediate distributions in the proof of security (this also removes the use of the \mathcal{O}_E oracle in the core 1-ABE security game); the reduction is also a bit more complex as we need to embed the output of \mathcal{O}_F into the CP-ABE ciphertext.

A.1 CP-ABE Construction

Our CP-ABE scheme is as follows:

$\text{Setup}(1^\lambda, 1^n)$: Run $\mathbb{G} = (p, G_1, G_2, G_T, e) \leftarrow \mathcal{G}(1^\lambda)$. Sample

$$\mathbf{A} \leftarrow \mathbb{Z}_p^{k \times 2k}, \mathbf{B} \leftarrow \mathbb{Z}_p^{(k+1) \times k}, \mathbf{U}_0, \mathbf{W}_i \leftarrow \mathbb{Z}_p^{2k \times (k+1)}, \mathbf{v} \leftarrow \mathbb{Z}_p^{2k}$$

and output:

$$\begin{aligned} \text{msk} &:= (\mathbf{v}, \mathbf{B}, \mathbf{U}_0, \mathbf{W}_1, \dots, \mathbf{W}_n) \\ \text{mpk} &:= ([\mathbf{A}]_1, [\mathbf{AU}_0]_1, [\mathbf{AW}_1]_1, \dots, [\mathbf{AW}_n]_1, e([\mathbf{A}]_1, [\mathbf{v}]_2)) \end{aligned}$$

$\text{Enc}(\text{mpk}, f, M)$: Sample $(\{\mathbf{u}_j^\top\}, \rho) \leftarrow \text{share}(f, \mathbf{s}^\top \mathbf{AU}_0)$, $\mathbf{s}, \mathbf{s}_j \leftarrow \mathbb{Z}_p^k$. Output:

$$\begin{aligned} \text{ct}_f &= (\text{ct}_1, \{\text{ct}_{2,j}, \text{ct}_{3,j}\}, \text{ct}_4) \\ &:= \left([\mathbf{s}^\top \mathbf{A}]_1, \{[\mathbf{u}_j^\top + \mathbf{s}_j^\top \mathbf{AW}_{\rho(j)}]_1, [\mathbf{s}_j^\top \mathbf{A}]_1\}, e([\mathbf{s}^\top \mathbf{A}]_1, [\mathbf{v}]_2) \cdot M \right) \end{aligned}$$

where $\mathbf{W}_0 = \mathbf{0}$.

$\text{KeyGen}(\text{mpk}, \text{msk}, f)$: Sample $\mathbf{r} \leftarrow \mathbb{Z}_p^k$. Output:

$$\begin{aligned} \text{sk}_x &= (\text{sk}_1, \text{sk}_2, \{\text{sk}_{3,i}\}_{x_i=1}) \\ &:= ([\mathbf{v} + \mathbf{U}_0 \mathbf{Br}]_2, [\mathbf{Br}]_2, \{[\mathbf{W}_i \mathbf{Br}]_2\}_{x_i=1}) \end{aligned}$$

$\text{Dec}(\text{mpk}, \text{sk}_x, \text{ct}_f)$: Compute ω_j such that $\mathbf{s}^\top \mathbf{AU}_0 = \sum_{\rho(j)=0 \vee x_{\rho(j)}=1} \omega_j \mathbf{u}_j$ as described in Section 5.1.

Output:

$$\frac{\text{ct}_4}{e(\text{ct}_1, \text{sk}_1)} \cdot \prod_{\rho(j)=0 \vee x_{\rho(j)}=1} \left(\frac{e(\text{ct}_{2,j}, \text{sk}_2)}{e(\text{ct}_{3,j}, \text{sk}_{3,\rho(j)})} \right)^{\omega_j}$$

A.2 Correctness

Correctness relies on the fact that for all j , we have

$$\frac{e(\text{ct}_{2,j}, \text{sk}_2)}{e(\text{ct}_{3,j}, \text{sk}_{3,\rho(j)})} = [\mathbf{u}_j^\top \mathbf{Br}]_T$$

which follows from the fact that

$$\mathbf{u}_j^\top \mathbf{Br} = \underbrace{(\mathbf{u}_j^\top + \mathbf{s}_j^\top \mathbf{A} \mathbf{W}_{\rho(j)})}_{\text{ct}_{2,j}} \cdot \underbrace{\mathbf{Br}}_{\text{sk}_2} - \underbrace{\mathbf{s}_j^\top \mathbf{A}}_{\text{ct}_{3,j}} \cdot \underbrace{\mathbf{W}_{\rho(j)} \mathbf{Br}}_{\text{sk}_{3,\rho(j)}}$$

and also from the fact that

$$e(\text{ct}_1, \text{sk}_1) = [\mathbf{s}^\top \mathbf{A} \mathbf{v} + \mathbf{s}^\top \mathbf{A} \mathbf{U}_0 \mathbf{Br}]_T$$

Therefore, for all f, x such that $f(x) = 1$, we have:

$$\begin{aligned} \frac{\text{ct}_4}{e(\text{ct}_1, \text{sk}_1)} \cdot \prod_{\rho(j)=0 \vee x_{\rho(j)}=1} \left(\frac{e(\text{ct}_{2,j}, \text{sk}_2)}{e(\text{ct}_{3,j}, \text{sk}_{3,\rho(j)})} \right)^{\omega_j} &= \frac{M \cdot [\mathbf{s}^\top \mathbf{A} \mathbf{v}]_T}{[\mathbf{s}^\top \mathbf{A} \mathbf{v} + \mathbf{s}^\top \mathbf{A} \mathbf{U}_0 \mathbf{Br}]_T} \cdot \prod_{\rho(j)=0 \vee x_{\rho(j)}=1} [\mathbf{u}_j^\top \mathbf{Br}]_T^{\omega_j} \\ &= \frac{M \cdot [\mathbf{s}^\top \mathbf{A} \mathbf{v}]_T}{[\mathbf{s}^\top \mathbf{A} \mathbf{v} + \mathbf{s}^\top \mathbf{A} \mathbf{U}_0 \mathbf{Br}]_T} \cdot \left[\sum_{\rho(j)=0 \vee x_{\rho(j)}=1} \omega_j \mathbf{u}_j^\top \mathbf{Br} \right]_T \\ &= \frac{M \cdot [\mathbf{s}^\top \mathbf{A} \mathbf{v}]_T}{[\mathbf{s}^\top \mathbf{A} \mathbf{v} + \mathbf{s}^\top \mathbf{A} \mathbf{U}_0 \mathbf{Br}]_T} \cdot [\mathbf{s}^\top \mathbf{A} \mathbf{U}_0 \mathbf{Br}]_T \\ &= \frac{M \cdot [\mathbf{s}^\top \mathbf{A} \mathbf{v} + \mathbf{s}^\top \mathbf{A} \mathbf{U}_0 \mathbf{Br}]_T}{[\mathbf{s}^\top \mathbf{A} \mathbf{v} + \mathbf{s}^\top \mathbf{A} \mathbf{U}_0 \mathbf{Br}]_T} \\ &= M \end{aligned}$$

A.3 Adaptive Security

Description of hybrids A ciphertext can be in one of the following forms:

- Normal: generated as in the scheme.
- SF: same as a Normal ciphertext, except $\mathbf{s}^\top \mathbf{A}, \mathbf{s}_j^\top \mathbf{A}$ replaced with $\mathbf{c}^\top, \mathbf{c}_j^\top$, where $\mathbf{c}, \mathbf{c}_j \leftarrow \mathbb{Z}_p^{2k}$. That is,

$$\text{ct}_f := \left(\boxed{\mathbf{c}^\top}_1, \{[\mathbf{u}_j^\top + \boxed{\mathbf{c}_j^\top} \mathbf{W}_{\rho(j)}]_1, \boxed{\mathbf{c}_j^\top}_1\}, e(\boxed{\mathbf{c}^\top}_1, [\mathbf{v}]_2) \cdot M \right)$$

A secret key can be in one of the following forms:

- Normal: generated as in the scheme.
- SF: same as a Normal key, except \mathbf{v} replaced with $\mathbf{v} + \mathbf{A}^\perp \boldsymbol{\delta}^{(q)}$, where a fresh $\boldsymbol{\delta}^{(q)} \leftarrow \mathbb{Z}_p^k$ is chosen per SF and \mathbf{A}^\perp is any fixed $\mathbf{A}^\perp \in \mathbb{Z}_p^{2k \times k} \setminus \{\mathbf{0}\}$ such that $\mathbf{A} \mathbf{A}^\perp = \mathbf{0}$. That is,

$$\text{sk}_x := (\boxed{\mathbf{v} + \mathbf{A}^\perp \boldsymbol{\delta}^{(q)}} + \mathbf{U}_0 \mathbf{Br}_2, [\mathbf{Br}]_2, \{[\mathbf{W}_i \mathbf{Br}]_2\}_{x_i=1})$$

- P-Normal: same as a Normal key, except \mathbf{Br} replaced with $\mathbf{d} \leftarrow \mathbb{Z}_p^{k+1}$. That is,

$$\text{sk}_x := ([\mathbf{v} + \mathbf{U}_0 \mathbf{d}]_2, \boxed{\mathbf{d}}_2, \{[\mathbf{W}_i \mathbf{d}]_2\}_{x_i=1})$$

- P-SF: same as a SF key, except \mathbf{Br} replaced with $\mathbf{d} \leftarrow \mathbb{Z}_p^{k+1}$. That is,

$$\text{sk}_x := (\boxed{\mathbf{v} + \mathbf{A}^\perp \boldsymbol{\delta}^{(q)}} + \mathbf{U}_0 \boxed{\mathbf{d}}_2, \boxed{\mathbf{d}}_2, \{[\mathbf{W}_i \mathbf{d}]_2\}_{x_i=1})$$

Here, P stands for pseudo following [Wee14,CGW15].

Hybrid sequence. Suppose the adversary A makes at most Q secret key queries. The hybrid sequence is as follows:

- H_0 : real game
- H_1 : same as H_0 , except we use a SF ciphertext.
- $H_{2,\ell,1}, \ell = 0, \dots, Q$: same as H_1 , except the ℓ 'th key is P-Normal, the first $\ell - 1$ keys are SF and the last $Q - \ell$ keys are Normal.
- $H_{2,\ell,2}$: same as $H_{2,\ell,1}$ except the ℓ 'th key is P-SF.
- $H_{2,\ell,3}$: same as $H_{2,\ell,1}$ except the ℓ 'th key is SF.
- H_3 : replace M with random.

Proof overview.

- We have $H_0 \approx_c H_1 \equiv H_{2,0,3}$ via k -Lin (and its self-reducibility), which tells us

$$([\mathbf{A}]_1, [\mathbf{s}^\top \mathbf{A}]_1, \{[\mathbf{s}_j^\top \mathbf{A}]_1\}) \approx_c ([\mathbf{A}]_1, [\mathbf{c}^\top]_1, \{[\mathbf{c}_j^\top]_1\})$$

Here, the security reduction will pick $\mathbf{U}_0, \mathbf{W}_1, \dots, \mathbf{W}_n$ and \mathbf{v} so that it can simulate the mpk, the ciphertext and the secret keys.

- We have $H_{2,\ell-1,3} \approx_c H_{2,\ell,1}$ for all $\ell \in [Q]$. The difference between the two is that we switch the ℓ 'th sk_f from Normal to P-Normal.

This follows again via k -Lin, which tells us $([\mathbf{B}]_2, [\mathbf{Br}]_2) \approx_c ([\mathbf{B}]_2, [\mathbf{d}]_2)$. Again, the security reduction will pick $\mathbf{U}_0, \mathbf{W}_1, \dots, \mathbf{W}_n$ and \mathbf{v} so that it can simulate the mpk, the ciphertext and the secret keys.

- We have $H_{2,\ell,1} \approx_c H_{2,\ell,2}$ for all $\ell \in [Q]$. The difference between the two is that we switch the ℓ 'th sk_f from P-Normal to P-SF. The idea is to program:

$$\mathbf{W}_i = \widetilde{\mathbf{W}}_i + \mathbf{A}^\perp \mathbf{w}_i (\mathbf{b}^\perp)^\top, \mathbf{U}_0 = \widetilde{\mathbf{U}}_0 + \mathbf{A}^\perp \mathbf{u} (\mathbf{b}^\perp)^\top$$

where $\mathbf{w}_i, \mathbf{b}^\perp \in \mathbb{Z}_p^k, \mathbf{A}^\perp \in \mathbb{Z}_p^{2k \times k}, \mathbf{u} \in \mathbb{Z}_p^{k+1}$ and

$$\mathbf{A} \mathbf{A}^\perp = \mathbf{0}, (\mathbf{b}^\perp)^\top \mathbf{B} = \mathbf{0}$$

Note that the public parameters and the normal and SF keys information-theoretically hide a random \mathbf{u} and the \mathbf{w}_i 's from $G^{1\text{-ABE}}$. Next, we focus on the SF ciphertext ct_f , and the ℓ 'th secret key sk_x , which is either P-Normal or P-SF.

First, we argue that if we ignore $\mathbf{v} + \mathbf{U}_0 \mathbf{d}$ in sk_x , then \mathbf{u} remains computationally hidden given ct_f, sk_x using the $G^{1\text{-ABE}}$ security game. Theorem 2 tells us that \mathbf{u} is computationally hidden given

$$\mathbf{c}, \{[\mathbf{c}_j^\top \mathbf{A}^\perp]_1, [\mu_j + \mathbf{c}_j^\top \mathbf{A}^\perp \mathbf{w}_{\rho(j)}]_1\}_j, \{\mathbf{w}_i\}_{i=1}$$

where $(\{\mu_j\}, \rho) \leftarrow \text{share}(f, \mathbf{c}^\top \mathbf{A}^\perp \mathbf{u})$ and we treat $\mathbf{c}_j^\top \mathbf{A}^\perp \in \mathbb{Z}_p^{1 \times k}, \mathbf{c}_j \leftarrow \mathbb{Z}_p^{2k}$ as the randomness used for CPA.Enc, even for adaptive choices of f, x . We can then use the entropy in \mathbf{u} to hide the \mathbf{A}^\perp -component of \mathbf{v} in $\mathbf{v} + \underbrace{\mathbf{A}^\perp \mathbf{u} (\mathbf{b}^\perp)^\top}_{\neq \mathbf{0}} \mathbf{d}$.

- We have $H_{2,\ell,2} \approx_c H_{2,\ell,3}$ for all $\ell \in [Q]$. The difference between the two is that we switch the ℓ 'th sk_f from P-SF to SF.

This follows again via k -Lin, which tells us $([\mathbf{B}]_2, [\mathbf{Br}]_2) \approx_c ([\mathbf{B}]_2, [\mathbf{d}]_2)$, symmetrically to the proof for $H_{2,\ell-1,3} \approx_c H_{2,\ell,1}$, except that $\mathbf{v} + \mathbf{A}^\perp \delta^{(\ell)}$ is used instead of \mathbf{v} in the ℓ th secret key.

- We have $H_{2,Q,3} \equiv H_3$. In $H_{2,Q,3}$, the secret keys only leak $\mathbf{v} + \mathbf{A}^\perp \boldsymbol{\delta}^{(1)}, \dots, \mathbf{v} + \mathbf{A}^\perp \boldsymbol{\delta}^{(Q)}$. This means that $\mathbf{c}^\top \mathbf{v}$ is statistically random (as long as $\mathbf{c}^\top \mathbf{A}^\perp \neq \mathbf{0}$).

Lemma 10 ($H_0 \approx_c H_1 \equiv H_{2,0,3}$).

$$|\Pr[\langle \mathcal{A}, H_0 \rangle = 1] - \Pr[\langle \mathcal{A}, H_1 \rangle = 1]| \leq \text{Adv}_{\mathcal{A}^*}^{k\text{-LIN}}(\lambda)$$

Proof. Given $\text{MDDH}_{k,2k}^{2m+1}$ challenge $([\mathbf{A}]_1, [\mathbf{Z}^\top]_1)$, where either $\mathbf{Z}^\top = \mathbf{S}^\top \mathbf{A}$ for a $\mathbf{S}^\top \leftarrow \mathbb{Z}_p^{(2m+1) \times k}$ or $\mathbf{Z}^\top = \mathbf{C}^\top$ for a $\mathbf{C}^\top \leftarrow \mathbb{Z}_p^{(2m+1) \times 2k}$, an adversary \mathcal{A}' could simply choose $\mathbf{U}_0, \mathbf{W}_i \leftarrow \mathbb{Z}_p^{2k \times (k+1)}, \mathbf{v} \leftarrow \mathbb{Z}_p^{2k}$, form the public parameters with $\mathbf{A}, \mathbf{U}_0, \mathbf{W}_i, \mathbf{v}$, and choose its own $\mathbf{B} \leftarrow \mathbb{Z}_p^{(k+1) \times k}, \mathbf{r} \leftarrow \mathbb{Z}_p^k$ when responding to key requests.

For the challenge ciphertext, \mathcal{A}' computes: $(\{\mathbf{u}_j^\top\}, \rho) \leftarrow \text{share}(f, \mathbf{z}_{2m+1}^\top \mathbf{U}_0)$, parses the rows of \mathbf{Z}^\top as \mathbf{z}_j^\top for $j \in [m+1]$, and returns:

$$\text{ct}_x := \left([\mathbf{z}_{2m+1}^\top]_1, \{[\mathbf{u}_j^\top + \mathbf{z}_j^\top \mathbf{W}_{\rho(j)}]_1, [\mathbf{z}_j^\top]_1\}, e([\mathbf{z}_{2m+1}^\top]_1, [\mathbf{v}]_2) \cdot M_b \right)$$

(note that $|\{\mathbf{u}_j\}| \leq 2m$).

If $\mathbf{Z}^\top = \mathbf{S}^\top \mathbf{A}$, then the challenge ciphertext is **Normal** and \mathcal{A}' has simulated H_0 ;

If $\mathbf{Z}^\top = \mathbf{C}^\top$, then the challenge ciphertext is **SF** and \mathcal{A}' has simulated $H_1 \equiv H_{2,0,3}$.

Finally, recall from Section 2.4 that $\text{Adv}_{\mathcal{A}'}^{\text{MDDH}_{k,2k}^{2m+1}}(\lambda) = \text{Adv}_{\mathcal{A}^*}^{k\text{-LIN}}(\lambda)$ \square

Lemma 11 ($H_{2,\ell-1,3} \approx_c H_{2,\ell,1}$).

$$|\Pr[\langle \mathcal{A}, H_{2,\ell-1,3} \rangle = 1] - \Pr[\langle \mathcal{A}, H_{2,\ell,1} \rangle = 1]| \leq \text{Adv}_{\mathcal{A}^*}^{k\text{-LIN}}(\lambda)$$

Proof. Given MDDH_k challenge $([\mathbf{B}]_2, [\mathbf{z}]_2)$, where either $\mathbf{z} = \mathbf{B}\mathbf{r}$ for $\mathbf{r} \leftarrow \mathbb{Z}_p^k$ or $\mathbf{z} = \mathbf{d}$, for $\mathbf{d} \leftarrow \mathbb{Z}_p^{k+1}$, an adversary \mathcal{A}' could simply choose $\mathbf{A} \leftarrow \mathbb{Z}_p^{k \times 2k}, \mathbf{U}_0, \mathbf{W}_i \leftarrow \mathbb{Z}_p^{2k \times (k+1)}, \mathbf{v} \leftarrow \mathbb{Z}_p^{2k}$, and form the public parameters with $\mathbf{A}, \mathbf{U}_0, \mathbf{W}_i, \mathbf{v}$. \mathcal{A}' could then compute $\mathbf{A}^\perp \in \mathbb{Z}_p^{2k \times k}$ such that $\mathbf{A}\mathbf{A}^\perp = \mathbf{0}$ (to be used in answering secret key queries).

- For the (SF) challenge ciphertext, \mathcal{A}' computes: $(\{\mathbf{u}_j^\top\}, \rho) \leftarrow \text{share}(f, \mathbf{c}^\top \mathbf{U}_0)$, draws $\mathbf{c}, \mathbf{c}_j \leftarrow \mathbb{Z}_p^{2k}$ for each j , and creates:

$$\text{ct}_f := \left([\mathbf{c}^\top]_1, \{[\mathbf{u}_j^\top + \mathbf{c}_j^\top \mathbf{W}_{\rho(j)}]_1, [\mathbf{c}_j^\top]_1\}, e([\mathbf{c}^\top]_1, [\mathbf{v}]_2) \cdot M_b \right)$$

- For the first $\ell - 1$ secret keys requested, say the q th request is for \mathbf{x} , \mathcal{A}' draws $\boldsymbol{\delta}^{(q)}, \mathbf{r}^{(q)} \leftarrow \mathbb{Z}_p^k$, and forms the following (SF) key:

$$\text{sk}_x := ([\mathbf{v} + \mathbf{A}^\perp \boldsymbol{\delta}^{(q)} + \mathbf{U}_0 \mathbf{B}\mathbf{r}^{(q)}]_2, [\mathbf{B}\mathbf{r}^{(q)}]_2, \{[\mathbf{W}_i \mathbf{B}\mathbf{r}^{(q)}]_2\}_{x_i=1})$$

- For the last $Q - \ell$ secret keys requested, \mathcal{A}' proceeds as before for the first $\ell - 1$ keys except using just \mathbf{v} instead of $\mathbf{v} + \mathbf{A}^\perp \boldsymbol{\delta}^{(q)}$. It is easy to see that it forms a **Normal** key.
- For the ℓ th secret key request, \mathcal{A}' forms the following key:

$$\text{sk}_x := ([\mathbf{v} + \mathbf{U}_0 \mathbf{z}]_2, [\mathbf{z}]_2, \{[\mathbf{W}_i \mathbf{z}]_2\}_{x_i=1})$$

If $\mathbf{z} = \mathbf{Br}$, then the ℓ th key is Normal and \mathcal{A}' has simulated $\mathbf{H}_{2,\ell-1,3}$;

If $\mathbf{z} = \mathbf{d}$, then the ℓ th key is P-Normal and \mathcal{A}' has simulated $\mathbf{H}_{2,\ell,1}$. \square

Lemma 12 ($\mathbf{H}_{2,\ell,1} \approx_c \mathbf{H}_{2,\ell,2}$).

$$|\Pr[\langle \mathcal{A}, \mathbf{H}_{2,\ell,1} \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_{2,\ell,2} \rangle = 1]| \leq 2^{6d} \cdot 8^d \cdot n \cdot \text{Adv}_{\mathcal{G}^*}^{k\text{-LIN}}(\lambda)$$

Proof. Consider the following adversary \mathcal{A}' in $\mathbf{G}_\beta^{1\text{-ABE}}$ which internally simulates \mathcal{A} and the challenger in the ABE security game:

- First, \mathcal{A}' samples $\mathbf{A} \leftarrow \mathbb{Z}_p^{k \times 2k}$, $\mathbf{B} \leftarrow \mathbb{Z}_p^{(k+1) \times k}$, $\tilde{\mathbf{U}}_0, \tilde{\mathbf{W}}_i \leftarrow \mathbb{Z}_p^{2k \times (k+1)}$, $\tilde{\mathbf{v}} \leftarrow \mathbb{Z}_p^{2k}$, computes $\mathbf{A}^\perp \in \mathbb{Z}_p^{2k \times k} \setminus \{\mathbf{0}\}$, $\mathbf{b}^\perp \in \mathbb{Z}_p^{(k+1)}$ such that $\mathbf{A}\mathbf{A}^\perp = \mathbf{0}$ and $(\mathbf{b}^\perp)^\top \mathbf{B} = \mathbf{0}$ and implicitly defines

$$\mathbf{v} := \tilde{\mathbf{v}} - \frac{\mu^{(0)}((\mathbf{b}^\perp)^\top \mathbf{d})}{(\mathbf{c}^\top \mathbf{A}^\perp \mathbf{u})} \mathbf{A}^\perp \mathbf{u}, \mathbf{U}_0 := \tilde{\mathbf{U}}_0 + \frac{\mu^{(b)}}{(\mathbf{c}^\top \mathbf{A}^\perp \mathbf{u})} \mathbf{A}^\perp \mathbf{u} (\mathbf{b}^\perp)^\top, \mathbf{W}_i := \tilde{\mathbf{W}}_i + \mathbf{A}^\perp \mathbf{w}_i (\mathbf{b}^\perp)^\top$$

where $\mathbf{w}_i \in \mathbb{Z}_p^k$, $\mu^{(b)} \in \mathbb{Z}_p$ are chosen in $\mathbf{G}_\beta^{1\text{-ABE}}$, $\mathbf{c} \leftarrow \mathbb{Z}_p^{2k}$ is chosen for use in the challenge ciphertext, $\mathbf{d} \leftarrow \mathbb{Z}_p^{k+1}$ is chosen for use in the ℓ th secret key, and $\mathbf{u} \leftarrow \mathbb{Z}_p^k$. Note that \mathcal{A}' can compute \mathbf{v} since it has $\mu^{(0)}$ from $\mathbf{G}_\beta^{1\text{-ABE}}$ and knows all other vectors.

Then, \mathcal{A}' generates the public parameters as:

$$\text{mpk} := ([\mathbf{A}]_1, [\mathbf{A}\tilde{\mathbf{U}}_0]_1, [\mathbf{A}\tilde{\mathbf{W}}_1]_1, \dots, [\mathbf{A}\tilde{\mathbf{W}}_n]_1, e([\mathbf{A}]_1, [\tilde{\mathbf{v}}]_2))$$

- When \mathcal{A} requests a challenge ciphertext for formula f along with M_0, M_1 , \mathcal{A}' queries $\mathcal{O}_F(f) \rightarrow (\{ [\mu_j + \mathbf{r}_j^\top \mathbf{w}_{\rho(j)}]_1, [\mathbf{r}_j]_1 \})$ in $\mathbf{G}_\beta^{1\text{-ABE}}$. \mathcal{A}' then samples $\tilde{\mathbf{c}}_j \leftarrow \mathbb{Z}_p^k$ for each j and $b \leftarrow \{0, 1\}$ (the challenge bit in the standard ABE security game), defines $\mathbf{A}_C^\perp := \begin{bmatrix} (\mathbf{A}^\perp)^\top \\ \mathbf{M} \end{bmatrix} \in \mathbb{Z}_p^{2k \times 2k}$ for a choice of \mathbf{M} that makes \mathbf{A}_C^\perp invertible, computes $[\mathbf{c}_j]_1 := \left[(\mathbf{A}_C^\perp)^{-1} \begin{pmatrix} \mathbf{r}_j \\ \tilde{\mathbf{c}}_j \end{pmatrix} \right]_1$, computes: $(\{\tilde{\mathbf{u}}_j^\top\}, \rho) \leftarrow \text{share}(f, \mathbf{c}^\top \tilde{\mathbf{U}}_0)$, and returns the following appropriately distributed (SF) challenge ciphertext:

$$\begin{aligned} \text{ct}_f &:= \left([\mathbf{c}^\top]_1, \{ [\tilde{\mathbf{u}}_j^\top + (\mu_j + \mathbf{r}_j^\top \mathbf{w}_{\rho(j)}) (\mathbf{b}^\perp)^\top + \mathbf{c}_j^\top \tilde{\mathbf{W}}_{\rho(j)}]_1, [\mathbf{c}_j^\top]_1 \}, e([\mathbf{c}^\top]_1, [\mathbf{v}]_2) \cdot M_b \right) \\ &= \left([\mathbf{c}^\top]_1, \underbrace{\{ [\tilde{\mathbf{u}}_j^\top + \mu_j (\mathbf{b}^\perp)^\top] \}}_{\equiv \text{share}(f, \mathbf{c}^\top \mathbf{U}_0)} + \underbrace{\{ \mathbf{c}_j^\top \tilde{\mathbf{W}}_{\rho(j)} + \mathbf{r}_j^\top \mathbf{w}_{\rho(j)} (\mathbf{b}^\perp)^\top \}}_{=\mathbf{c}_j^\top \mathbf{W}_{\rho(j)}}, [\mathbf{c}_j^\top]_1, e([\mathbf{c}^\top]_1, [\mathbf{v}]_2) \cdot M_b \right) \end{aligned}$$

Note that $\{\mu_j (\mathbf{b}^\perp)^\top\}$ is distributed like the output of $\text{share}(f, \mu^{(b)} (\mathbf{b}^\perp)^\top)$, and therefore due to linearity and the fact that $\mathbf{c}^\top \mathbf{U}_0 = \mathbf{c}^\top \tilde{\mathbf{U}}_0 + \mu^{(b)} (\mathbf{b}^\perp)^\top$, then $\{\tilde{\mathbf{u}}_j^\top + \mu_j (\mathbf{b}^\perp)^\top\}$ is distributed like $\text{share}(f, \mathbf{c}^\top \tilde{\mathbf{U}}_0 + \mu^{(b)} (\mathbf{b}^\perp)^\top) \equiv \text{share}(f, \mathbf{c}^\top \mathbf{U}_0)$. Also, note that $\mathbf{c}_j^\top \mathbf{W}_i = \mathbf{c}_j^\top \tilde{\mathbf{W}}_i + \mathbf{r}_j^\top \mathbf{w}_{\rho(j)} (\mathbf{b}^\perp)^\top$ since $\mathbf{c}_j^\top \mathbf{A}^\perp = \mathbf{r}_j$.

- For the first $\ell - 1$ secret keys requested, say the q th request is for \mathbf{x} , \mathcal{A}' draws $\delta^{(q)}, \mathbf{r}^{(q)} \leftarrow \mathbb{Z}_p^k$, and forms the following (SF) key:

$$\text{sk}_{\mathbf{x}} = ([\mathbf{v} + \mathbf{A}^\perp \delta^{(q)} + \underbrace{\tilde{\mathbf{U}}_0 \mathbf{Br}^{(q)}}_{=\mathbf{U}_0 \mathbf{Br}^{(q)}}]_2, [\mathbf{Br}^{(q)}]_2, \{ \underbrace{[\tilde{\mathbf{W}}_i \mathbf{Br}^{(q)}]_2}_{=\mathbf{W}_i \mathbf{Br}^{(q)}} \}_{x_i=1})$$

- For the last $Q - \ell$ secret keys requested, \mathcal{A}' proceeds as before for the first $\ell - 1$ keys except using just \mathbf{v} instead of $\mathbf{v} + \mathbf{A}^\perp \boldsymbol{\delta}^{(q)}$. It is easy to see that it forms a Normal key.
- For the ℓ th secret key requested, say for \mathbf{x} , queries $\mathcal{O}_X(\mathbf{x}) \rightarrow (\{\mathbf{w}_i\}_{x_i=1})$ in $\mathbf{G}_\beta^{1\text{-ABE}}$, then uses these components to return:

$$\text{sk}_{\mathbf{x}} = ([\underbrace{\tilde{\mathbf{v}} + \tilde{\mathbf{U}}_0 \mathbf{d}}_{=\mathbf{v} + \frac{(\mu^{(0)} - \mu^{(b)})(\mathbf{b}^\perp)^\top \mathbf{d}}{(\mathbf{c}^\top \mathbf{A}^\perp \mathbf{u})} \mathbf{A}^\perp \mathbf{u} + \mathbf{U}_0 \mathbf{d}}]_2, [\mathbf{d}]_2, \{ \underbrace{[(\tilde{\mathbf{W}}_i + \mathbf{A}^\perp \mathbf{w}_i (\mathbf{b}^\perp)^\top) \mathbf{d}]_2}_{=\mathbf{W}_i \mathbf{d}} \}_{x_i=1})$$

If $\beta = 0$, then the ℓ th key is a P-Normal key since $\mathbf{v} + \frac{(\mu^{(0)} - \mu^{(0)})(\mathbf{b}^\perp)^\top \mathbf{d}}{(\mathbf{c}^\top \mathbf{A}^\perp \mathbf{u})} \mathbf{A}^\perp \mathbf{u} = \mathbf{v}$.

If $\beta = 1$, then the ℓ th key is a P-SF key, where $\boldsymbol{\delta}^{(\ell)} = \frac{(\mu^{(0)} - \mu^{(1)})(\mathbf{b}^\perp)^\top \mathbf{d}}{(\mathbf{c}^\top \mathbf{A}^\perp \mathbf{u})} \mathbf{u}$.

Putting everything together, for $\beta \in \{0, 1\}$, when \mathcal{A}' interacts with $\mathbf{G}_\beta^{1\text{-ABE}}$, then \mathcal{A}' simulates $\mathbf{H}_{2,\ell,1+\beta}$. So:

$$|\Pr[\langle \mathcal{A}, \mathbf{H}_{2,\ell,1} \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_{2,\ell,2} \rangle = 1]| \leq |\Pr[\langle \mathcal{A}', \mathbf{G}_0^{1\text{-ABE}} \rangle = 1] - \Pr[\langle \mathcal{A}', \mathbf{G}_1^{1\text{-ABE}} \rangle = 1]|$$

and from Theorem 2, we then have:

$$|\Pr[\langle \mathcal{A}, \mathbf{H}_{2,\ell-1} \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_{2,\ell} \rangle = 1]| \leq 2^{6d} \cdot 8^d \cdot n \cdot \text{Adv}_{\beta^*}^{k\text{-LIN}}(\lambda)$$

□

Lemma 13 ($\mathbf{H}_{2,\ell,2} \approx_c \mathbf{H}_{2,\ell,3}$).

$$|\Pr[\langle \mathcal{A}, \mathbf{H}_{2,\ell,2} \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_{2,\ell,3} \rangle = 1]| \leq \text{Adv}_{\mathcal{A}'}^{k\text{-LIN}}(\lambda)$$

Proof. Omitted, since the proof is completely analogous to that of Lemma 11, using $\mathbf{v} + \mathbf{A}^\perp \boldsymbol{\delta}^{(\ell)}$ for a new random $\boldsymbol{\delta}^{(\ell)}$ instead of \mathbf{v} when creating the ℓ th key.

Lemma 14 ($\mathbf{H}_{2,Q,3} \approx_c \mathbf{H}_3$).

$$|\Pr[\langle \mathcal{A}, \mathbf{H}_{2,Q,3} \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_3 \rangle = 1]| \leq \frac{1}{p}$$

Proof. These two hybrids are identically distributed conditioned on $\mathbf{c}^\top \mathbf{A}^\perp \neq \mathbf{0}$. To see this, consider two ways of choosing \mathbf{v} : $\mathbf{v} = \tilde{\mathbf{v}} \leftarrow \mathbb{Z}_p^{2k}$ and $\mathbf{v} = \tilde{\mathbf{v}} + \mathbf{A}^\perp \tilde{\mathbf{m}}$ for an independently random $\tilde{\mathbf{m}} \leftarrow \mathbb{Z}_p^k$. Note that both result in \mathbf{v} having a uniform distribution.

Using $\tilde{\mathbf{v}}$ to simulate hybrid $\mathbf{H}_{2,Q,3}$ obviously results in $\mathbf{H}_{2,Q,3}$ (where $\mathbf{v} = \tilde{\mathbf{v}}$). However, using the identically distributed $\mathbf{v} = \tilde{\mathbf{v}} + \mathbf{A}^\perp \tilde{\mathbf{m}}$ to simulate $\mathbf{H}_{2,Q,3}$ results in \mathbf{H}_3 (where $M \cdot [\mathbf{c}^\top \mathbf{A}^\perp \tilde{\mathbf{m}}]_T$ is a randomly distributed message as long as $\mathbf{c}^\top \mathbf{A}^\perp \neq \mathbf{0}$, and for redefined independently random $\tilde{\boldsymbol{\delta}}^{(i)} = \boldsymbol{\delta}^{(i)} + \tilde{\mathbf{m}}$ in the secret keys).

\mathbf{c} is chosen at random and independent from $\mathbf{A}^\perp \neq \mathbf{0}$, so $\mathbf{c}^\top \mathbf{A}^\perp = \mathbf{0}$ with probability $\frac{1}{p}$, and since we know that $\mathbf{H}_{2,Q,3} \equiv \mathbf{H}_3$ conditioned on $\mathbf{c}^\top \mathbf{A}^\perp \neq \mathbf{0}$, then we have:

$$|\Pr[\langle \mathcal{A}, \mathbf{H}_{2,Q,3} \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_3 \rangle = 1]| \leq \frac{1}{p}$$

□

Theorem 4 (adaptive CP-ABE). *The CP-ABE construction in Appendix A.1 is adaptively secure under the MDDH_k assumption.*

Proof. Note that since $H_1 \equiv H_{2,0,3}$:

$$\begin{aligned}
|\Pr[\langle \mathcal{A}, H_0 \rangle = 1] - \Pr[\langle \mathcal{A}, H_3 \rangle = 1]| &\leq |\Pr[\langle \mathcal{A}, H_0 \rangle = 1] - \Pr[\langle \mathcal{A}, H_1 \rangle = 1]| \\
&+ \sum_{\ell=1}^Q |\Pr[\langle \mathcal{A}, H_{2,\ell-1,3} \rangle = 1] - \Pr[\langle \mathcal{A}, H_{2,\ell,1} \rangle = 1]| \\
&+ \sum_{\ell=1}^Q |\Pr[\langle \mathcal{A}, H_{2,\ell,1} \rangle = 1] - \Pr[\langle \mathcal{A}, H_{2,\ell,2} \rangle = 1]| \\
&+ \sum_{\ell=1}^Q |\Pr[\langle \mathcal{A}, H_{2,\ell,2} \rangle = 1] - \Pr[\langle \mathcal{A}, H_{2,\ell,3} \rangle = 1]| \\
&+ |\Pr[\langle \mathcal{A}, H_{2,Q,3} \rangle = 1] - \Pr[\langle \mathcal{A}, H_3 \rangle = 1]|
\end{aligned}$$

Summing the results of Lemmas 10, 11, 12, 13, and 14, we then have that:

$$|\Pr[\langle \mathcal{A}, H_0 \rangle = 1] - \Pr[\langle \mathcal{A}, H_3 \rangle = 1]| \leq \text{Adv}_{\mathcal{B}^*}^{k\text{-LIN}}(\lambda) + 2 \cdot Q \cdot \text{Adv}_{\mathcal{B}^*}^{k\text{-LIN}}(\lambda) + Q \cdot 2^{6d} \cdot 8^d \cdot n \cdot \text{Adv}_{\mathcal{B}^*}^{k\text{-LIN}}(\lambda) + \frac{1}{p}$$

If $d = O(\log n)$, then under the k -Lin assumption this quantity is a negligible function of λ (the number of queries made Q and the attribute vector length n are both polynomial in λ , and $\frac{1}{p}$ is a negligible function of λ). It's easy to see that $\text{Adv}_{\mathcal{A}}^{\text{ABE}}(\lambda) = 0$ in the H_3 hybrid game (since a random message is encrypted in the challenge ciphertext). So, any adversary in the real game (H_0) will have advantage negligibly close to 0, and our construction satisfies adaptive security. \square

B Our Unbounded KP-ABE Scheme

In this section, we use the modular technique presented in [CGKW18] to transform our KP-ABE construction from Section 6 (for NC^1 that is compact and adaptively secure under the MDDH_k assumption in asymmetric prime-order bilinear groups) into a construction with the same properties plus an added benefit that the scheme is unbounded (that is, the public parameters are of constant size [LW11]).

B.1 Unbounded KP-ABE Construction

$\text{Setup}(1^\lambda, 1^n)$: Run $\mathbb{G} = (p, G_1, G_2, G_T, e) \leftarrow \mathcal{G}(1^\lambda)$. Sample

$$\mathbf{A}_1 \leftarrow \mathbb{Z}_p^{k \times (2k+1)}, \mathbf{W}, \mathbf{W}_0, \mathbf{W}_1 \leftarrow \mathbb{Z}_p^{(2k+1) \times k}, \mathbf{v} \leftarrow \mathbb{Z}_p^{2k+1}$$

and output:

$$\begin{aligned} \text{msk} &:= (\mathbf{v}, \mathbf{W}, \mathbf{W}_0, \mathbf{W}_1) \\ \text{mpk} &:= ([\mathbf{A}_1]_1, [\mathbf{A}_1 \mathbf{W}]_1, [\mathbf{A}_1 \mathbf{W}_0]_1, [\mathbf{A}_1 \mathbf{W}_1]_1, e([\mathbf{A}_1]_1, [\mathbf{v}]_2)) \end{aligned}$$

$\text{Enc}(\text{mpk}, x, M)$: Sample $\mathbf{s}, \mathbf{s}_i \leftarrow \mathbb{Z}_p^k$. Output:

$$\begin{aligned} \text{ct}_x &= (\text{ct}_1, \{\text{ct}_{2,i}, \text{ct}_{3,i}\}_{x_i=1}, \text{ct}_4) \\ &:= ([\mathbf{s}^\top \mathbf{A}_1]_1, \{[\mathbf{s}^\top \mathbf{A}_1 \mathbf{W} + \mathbf{s}_i^\top \mathbf{A}_1 (\mathbf{W}_0 + i \cdot \mathbf{W}_1)]_1, [\mathbf{s}_i^\top \mathbf{A}_1]_1\}_{x_i=1}, e([\mathbf{s}^\top \mathbf{A}_1]_1, [\mathbf{v}]_2) \cdot M) \end{aligned}$$

$\text{KeyGen}(\text{mpk}, \text{msk}, f)$: Sample $(\{\mathbf{v}_j\}, \rho) \leftarrow \text{share}(f, \mathbf{v})$, $\mathbf{r}_j \leftarrow \mathbb{Z}_p^k$. Output:

$$\begin{aligned} \text{sk}_f &= (\{\text{sk}_{1,j}, \text{sk}_{2,j}, \text{sk}_{3,j}\}, \{\text{sk}_{4,j}\}) \\ &:= (\{[\mathbf{v}_j + \mathbf{W} \mathbf{r}_j]_2, [\mathbf{r}_j]_2, [(\mathbf{W}_0 + \rho(j) \cdot \mathbf{W}_1) \mathbf{r}_j]_2\}_{\rho(j) \neq 0}, \{[\mathbf{v}_j]_2\}_{\rho(j) = 0}) \end{aligned}$$

$\text{Dec}(\text{mpk}, \text{sk}_f, \text{ct}_x)$: Compute ω_j such that $\mathbf{v} = \sum_{\rho(j)=0 \vee x_{\rho(j)}=1} \omega_j \mathbf{v}_j$ as described in Section 5.1.

Output:

$$\text{ct}_4 \cdot \prod_{x_{\rho(j)}=1} \left(\frac{e(\text{ct}_{2,\rho(j)}, \text{sk}_{2,j})}{e(\text{ct}_1, \text{sk}_{1,j}) \cdot e(\text{ct}_{3,\rho(j)}, \text{sk}_{3,j})} \right)^{\omega_j} \cdot \prod_{\rho(j)=0} e(\text{ct}_1, \text{sk}_{4,j})^{-\omega_j}$$

B.2 Correctness

Correctness relies on the fact that for all j , we have

$$\frac{e(\text{ct}_1, \text{sk}_{1,j}) \cdot e(\text{ct}_{3,\rho(j)}, \text{sk}_{3,j})}{e(\text{ct}_{2,\rho(j)}, \text{sk}_{2,j})} = [\mathbf{s}^\top \mathbf{A}_1 \mathbf{v}_j]_T$$

which follows from the fact that

$$\mathbf{s}^\top \mathbf{A}_1 \mathbf{v}_j = \underbrace{\mathbf{s}^\top \mathbf{A}_1}_{\text{ct}_1} \cdot \underbrace{(\mathbf{v}_j + \mathbf{W} \mathbf{r}_j)}_{\text{sk}_{1,j}} - \underbrace{(\mathbf{s}^\top \mathbf{A}_1 \mathbf{W} + \mathbf{s}_{\rho(j)}^\top \mathbf{A}_1 (\mathbf{W}_0 + \rho(j) \cdot \mathbf{W}_1))}_{\text{ct}_{2,\rho(j)}} \cdot \underbrace{\mathbf{r}_j}_{\text{sk}_{2,j}} + \underbrace{\mathbf{s}_{\rho(j)}^\top \mathbf{A}_1}_{\text{ct}_{3,\rho(j)}} \cdot \underbrace{(\mathbf{W}_0 + \rho(j) \cdot \mathbf{W}_1) \mathbf{r}_j}_{\text{sk}_{3,j}}]_2$$

and also the fact that for all j ,

$$e(\text{ct}_1, \text{sk}_{4,j}) = [\mathbf{s}^\top \mathbf{A}_1 \mathbf{v}_j]_T$$

Therefore, for all f, x such that $f(x) = 1$, we have:

$$\begin{aligned}
\text{ct}_4 \cdot \prod_{x_{\rho(j)=1}} \left(\frac{e(\text{ct}_{2,\rho(j)}, \text{sk}_{2,j})}{e(\text{ct}_1, \text{sk}_{1,j}) \cdot e(\text{ct}_{3,\rho(j)}, \text{sk}_{3,j})} \right)^{\omega_j} &\cdot \prod_{\rho(j)=0} e(\text{ct}_1, \text{sk}_{4,j})^{-\omega_j} \\
&= M \cdot [\mathbf{s}^\top \mathbf{A}_1 \mathbf{v}]_T \cdot \prod_{\rho(j)=0 \vee x_{\rho(j)=1}} [\mathbf{s}^\top \mathbf{A}_1 \mathbf{v}_j]_T^{-\omega_j} \\
&= M \cdot [\mathbf{s}^\top \mathbf{A}_1 \mathbf{v}]_T \cdot [-\mathbf{s}^\top \mathbf{A}_1 \sum_{\rho(j)=0 \vee x_{\rho(j)=1}} \omega_j \mathbf{v}_j]_T \\
&= M \cdot [\mathbf{s}^\top \mathbf{A}_1 \mathbf{v}]_T \cdot [-\mathbf{s}^\top \mathbf{A}_1 \mathbf{v}]_T \\
&= M
\end{aligned}$$

B.3 Adaptive Security

Entropy expansion lemma. Our security proof relies on the “entropy expansion lemma” in [CGKW18]. First, we introduce some additional notation. Let \mathbf{A} be a matrix over \mathbb{Z}_p . We use $\text{span}(\mathbf{A})$ to denote the column span of \mathbf{A} , and we use $\text{span}^\ell(\mathbf{A})$ to denote matrices of width ℓ where each column lies in $\text{span}(\mathbf{A})$; this means $\mathbf{M} \leftarrow_{\mathbb{R}} \text{span}^\ell(\mathbf{A})$ is a random matrix of width ℓ where each column is chosen uniformly from $\text{span}(\mathbf{A})$. We use $\text{basis}(\mathbf{A})$ to denote a basis of $\text{span}(\mathbf{A})$, and we use $(\mathbf{A}_1 \mid \mathbf{A}_2)$ to denote the concatenation of matrices $\mathbf{A}_1, \mathbf{A}_2$.

Pick random

$$\mathbf{A}_1 \leftarrow_{\mathbb{R}} \mathbb{Z}_p^{\ell_1 \times \ell}, \mathbf{A}_2 \leftarrow_{\mathbb{R}} \mathbb{Z}_p^{\ell_2 \times \ell}, \mathbf{A}_3 \leftarrow_{\mathbb{R}} \mathbb{Z}_p^{\ell_3 \times \ell}$$

where $\ell := \ell_1 + \ell_2 + \ell_3$. Let $(\mathbf{A}_1^\parallel \mid \mathbf{A}_2^\parallel \mid \mathbf{A}_3^\parallel)^\top$ denote the inverse of $(\mathbf{A}_1^\top \mid \mathbf{A}_2^\top \mid \mathbf{A}_3^\top)$, so that $\mathbf{A}_i \mathbf{A}_i^\parallel = \mathbf{I}$ (known as *non-degeneracy*) and $\mathbf{A}_i \mathbf{A}_j^\parallel = \mathbf{0}$ if $i \neq j$ (known as *orthogonality*). Here, we focus on the case $(\ell_1, \ell_2, \ell_3) = (k, 1, k)$ and so $\ell = 2k + 1$.

Lemma 15 (entropy expansion lemma [CGKW18]). *Under the MDDH_k assumption, we have*

$$\begin{aligned}
\mathbb{D}_0 &:= \left\{ \begin{array}{l} \text{aux} : [\mathbf{A}_1]_1, [\mathbf{A}_1 \mathbf{W}]_1, [\mathbf{A}_1 \mathbf{W}_0]_1, [\mathbf{A}_1 \mathbf{W}_1]_1 \\ \text{ct} : [\mathbf{c}^\top]_1, \{[\mathbf{c}^\top \mathbf{W} + \mathbf{c}_i^\top (\mathbf{W}_0 + i \cdot \mathbf{W}_1)]_1, [\mathbf{c}_i^\top]_1\}_{i \in [n]} \\ \text{sk} : \{[\mathbf{W} \mathbf{D}_i]_2, [\mathbf{D}_i]_2, [(\mathbf{W}_0 + i \cdot \mathbf{W}_1) \mathbf{D}_i]_2\}_{i \in [n]} \end{array} \right\} \\
\approx_c & \\
\mathbb{D}_1 &:= \left\{ \begin{array}{l} \text{aux} : [\mathbf{A}_1]_1, [\mathbf{A}_1 \mathbf{W}]_1, [\mathbf{A}_1 \mathbf{W}_0]_1, [\mathbf{A}_1 \mathbf{W}_1]_1 \\ \text{ct} : [\mathbf{c}^\top]_1, \{[\mathbf{c}^\top (\mathbf{W} + \mathbf{V}_i^{(2)}) + \mathbf{c}_i^\top (\mathbf{W}_0 + i \cdot \mathbf{W}_1 + \mathbf{U}_i^{(2)})]_1, [\mathbf{c}_i^\top]_1\}_{i \in [n]} \\ \text{sk} : \{[(\mathbf{W} + \mathbf{V}_i^{(2)}) \mathbf{D}_i]_2, [\mathbf{D}_i]_2, [(\mathbf{W}_0 + i \cdot \mathbf{W}_1 + \mathbf{U}_i^{(2)}) \mathbf{D}_i]_2\}_{i \in [n]} \end{array} \right\}
\end{aligned}$$

where $\mathbf{W}, \mathbf{W}_0, \mathbf{W}_1 \leftarrow_{\mathbb{R}} \mathbb{Z}_p^{(2k+1) \times k}$, $\mathbf{V}_i^{(2)}, \mathbf{U}_i^{(2)} \leftarrow_{\mathbb{R}} \text{span}^k(\mathbf{A}_2^\parallel)$, $\mathbf{D}_i \leftarrow_{\mathbb{R}} \mathbb{Z}_p^{k \times k}$, and $\mathbf{c}, \mathbf{c}_i \leftarrow_{\mathbb{R}} \text{span}(\mathbf{A}_1^\top)$ in the left distribution while $\mathbf{c}, \mathbf{c}_i \leftarrow_{\mathbb{R}} \text{span}(\mathbf{A}_1^\top, \mathbf{A}_2^\top)$ in the right distribution, where the concrete security loss $|\Pr[\mathcal{A}'(\mathbb{D}_0) = 1] - \Pr[\mathcal{A}'(\mathbb{D}_1) = 1]| \leq (5n + 1) \cdot \text{Adv}_{\mathcal{G}'}^{k\text{-LIN}}(\lambda)$.

This lemma allows us to use a hybrid proof to first transition to a game in which the challenge ciphertext and secret keys have components in the \mathbf{A}_2 space which mirror those of our (bounded) construction of Section 6. We then follow the same proof structure as in Section 6.

Description of hybrids A ciphertext can be in one of the following forms:

- **Normal**: generated as in the scheme.
- **SF**: same as a Normal ciphertext, except $\mathbf{s}^\top \mathbf{A}_1, \mathbf{s}_i^\top \mathbf{A}_1$ replaced with $\mathbf{c}^\top, \mathbf{c}_i^\top \leftarrow \mathbb{Z}_p^{2k+1}$ and we use the substitution:

$$\mathbf{W} \rightarrow \widehat{\mathbf{V}}_i := \mathbf{W} + \mathbf{V}_i^{(2)} \text{ in } i\text{th component, and } \mathbf{W}_0 + i \cdot \mathbf{W}_1 \rightarrow \widehat{\mathbf{U}}_i := \mathbf{W}_0 + i \cdot \mathbf{W}_1 + \mathbf{U}_i^{(2)} \quad (3)$$

where $\mathbf{U}_i^{(2)}, \mathbf{V}_i^{(2)} \leftarrow \text{span}^k(\mathbf{A}_2^\parallel)$. Concretely, a SF ciphertext is of the form:

$$\text{ct}_{\mathbf{x}} := ([\mathbf{c}^\top]_1, \{[\mathbf{c}^\top \widehat{\mathbf{V}}_i + \mathbf{c}_i^\top \widehat{\mathbf{U}}_i]_1, [\mathbf{c}_i^\top]_1\}_{x_i=1}, e([\mathbf{c}^\top]_1, [\mathbf{v}]_2) \cdot M)$$

A secret key can be in one of the following forms:

- **Normal**: generated as in the scheme.
- **P-SF**: same as a Normal key, except we use the same substitution as in (3), concretely making a P-SF key of the form:

$$\text{sk}_f := (\{[\mathbf{v}_j + \widehat{\mathbf{V}}_{\rho(j)} \mathbf{r}_j]_2, [\mathbf{r}_j]_2, [\widehat{\mathbf{U}}_{\rho(j)} \mathbf{r}_j]_2\}_{\rho(j) \neq 0}, \{[\mathbf{v}_j]_2\}_{\rho(j)=0})$$

- **SF**: same as a P-SF key, except \mathbf{v} replaced with $\mathbf{v} + \delta \mathbf{a}^\perp$, where a fresh $\delta \leftarrow \mathbb{Z}_p$ is chosen per SF key and $\mathbf{a}^\perp \leftarrow \text{span}(\mathbf{A}_2^\parallel) \setminus \{\mathbf{0}\}$.

Hybrid sequence. Suppose the adversary \mathbf{A} makes at most Q secret key queries. The hybrid sequence is as follows:

- H_0 : real game
- H_1 : same as H_0 , except all keys are P-SF, and we use a SF ciphertext.
- $\text{H}_{2,\ell}, \ell = 0, \dots, Q$: same as H_1 , except the first ℓ keys are SF and the remaining $Q - \ell$ keys are P-SF.
- H_3 : replace M with random \widetilde{M} .

Proof overview.

- We have $\text{H}_0 \approx_c \text{H}_1 \equiv \text{H}_{2,0}$ via Lemma 15. In the reduction, on input

$$\left\{ \begin{array}{l} \text{aux} : [\mathbf{A}_1^\top]_1, [\mathbf{A}_1^\top \mathbf{W}]_1, [\mathbf{A}_1^\top \mathbf{W}_0]_1, [\mathbf{A}_1^\top \mathbf{W}_1]_1 \\ \text{ct} : [\mathbf{C}_0]_1, \{[\mathbf{C}_{1,i}]_1, [\mathbf{C}_{2,i}]_1\}_{i \in [n]} \\ \text{sk} : \{[\mathbf{K}_{0,i}]_2, [\mathbf{K}_{1,i}]_2, [\mathbf{K}_{2,i}]_2\}_{i \in [n]} \end{array} \right\},$$

we sample $\mathbf{v} \leftarrow_{\mathbb{R}} \mathbb{Z}_p^{2k+1}$, compute $(\{\mathbf{v}_j\}, \rho) \leftarrow \text{share}(f, \mathbf{v})$, draw $\tilde{\mathbf{r}}_{j,\ell} \leftarrow_{\mathbb{R}} \mathbb{Z}_p^k$ for shares j and keys $\ell \in [Q]$, and simulate the game with

$$\left\{ \begin{array}{l} \text{mpk} : \text{aux}, e([\mathbf{A}_1^\top]_1, [\mathbf{v}]_2) \\ \text{ct}_{\mathbf{x}} : [\mathbf{C}_0]_1, \{[\mathbf{C}_{1,i}]_1, [\mathbf{C}_{2,i}]_1\}_{i:x_i=1}, e([\mathbf{C}_0]_1, [\mathbf{v}]_2) \cdot M_b \\ \text{sk}_f^\ell : \{[\mathbf{v}_k + \mathbf{K}_{0,\rho(j)} \tilde{\mathbf{r}}_{j,\ell}]_2, [\mathbf{K}_{1,\rho(j)} \tilde{\mathbf{r}}_{j,\ell}]_2, [\mathbf{K}_{2,\rho(j)} \tilde{\mathbf{r}}_{j,\ell}]_2\} \end{array} \right\}.$$

In both cases, we set $\mathbf{r}_{j,\ell} := \mathbf{D}_i \tilde{\mathbf{r}}_{j,\ell}$ where $\mathbf{D}_i \leftarrow_{\mathbb{R}} \mathbb{Z}_p^{k \times k}$ as defined in the entropy expansion lemma (Lemma 15). Therefore all $\mathbf{r}_{j,\ell}$ are uniformly distributed over \mathbb{Z}_p^k with high probability.

- We have $H_{2,\ell-1} \approx_c H_{2,\ell}$, for all $\ell \in [Q]$. The difference between the two is that we switch the ℓ 'th sk_f from P-SF to SF using the adaptive security of our core 1-ABE component in $\mathbf{G}^{1\text{-ABE}}$ from Section 5.

The idea is to sample

$$\mathbf{v} = \tilde{\mathbf{v}} + \mu \mathbf{a}^\perp$$

where $\mathbf{a}^\perp \leftarrow \text{span}(\mathbf{A}_2^\parallel) \setminus \{\mathbf{0}\}$ so that mpk can be computed using $\tilde{\mathbf{v}}$ and perfectly hides $\mu, \mathbf{w}_1, \dots, \mathbf{w}_n$. Roughly speaking: the reduction

- uses $\mathcal{O}_X(x)$ in $\mathbf{G}^{1\text{-ABE}}$ to simulate the challenge ciphertext
 - uses $\mathcal{O}_F(f)$ in $\mathbf{G}^{1\text{-ABE}}$ to simulate ℓ 'th secret key
 - uses $\mu^{(0)}$ from $\mathbf{G}^{1\text{-ABE}}$ together with $\mathcal{O}_E(i, \cdot) = \text{Enc}(w_i, \cdot)$ to simulate the remaining $Q - \ell$ secret keys
- We have $H_{2,Q} \equiv H_3$. In $H_{2,Q}$, the secret keys only leak $\mathbf{v} + \delta_1 \mathbf{a}^\perp, \dots, \mathbf{v} + \delta_Q \mathbf{a}^\perp$. This means that $\mathbf{c}^\top \mathbf{v}$ is statistically random (as long as $\mathbf{c}^\top \mathbf{a}^\perp \neq 0$).

Lemma 16 ($H_0 \approx_c H_1 \equiv H_{2,0}$).

$$|\Pr[\langle \mathcal{A}, H_0 \rangle = 1] - \Pr[\langle \mathcal{A}, H_1 \rangle = 1]| \leq (5n + 1) \cdot \text{Adv}_{\mathcal{B}'}^{k\text{-LIN}}(\lambda)$$

Proof. Consider the following adversary \mathcal{A}' attempting to distinguish the distributions in the Entropy Expansion Lemma 15, which internally simulates \mathcal{A} and the challenger in the ABE security game:

- \mathcal{A}' receives input:

$$\mathbb{D}_\beta = \left\{ \begin{array}{l} \text{aux} : [\mathbf{A}_1^\top]_1, [\mathbf{A}_1^\top \mathbf{W}]_1, [\mathbf{A}_1^\top \mathbf{W}_0]_1, [\mathbf{A}_1^\top \mathbf{W}_1]_1 \\ \text{ct} : [\mathbf{C}_0]_1, \{[\mathbf{C}_{1,i}]_1, [\mathbf{C}_{2,i}]_1\}_{i \in [n]} \\ \text{sk} : \{[\mathbf{K}_{0,i}]_2, [\mathbf{K}_{1,i}]_2, [\mathbf{K}_{2,i}]_2\}_{i \in [n]} \end{array} \right\}$$

- First, \mathcal{A}' samples $\mathbf{v} \leftarrow_{\mathbb{R}} \mathbb{Z}_p^{2k+1}$ and outputs:

$$\text{mpk} := ([\mathbf{A}_1]_1, [\mathbf{A}_1 \mathbf{W}]_1, [\mathbf{A}_1 \mathbf{W}_0]_1, [\mathbf{A}_1 \mathbf{W}_1]_1, e([\mathbf{A}_1]_1, [\mathbf{v}]_2)),$$

- When \mathcal{A} requests a challenge ciphertext for attribute \mathbf{x} along with M_0, M_1 , \mathcal{A}' samples $b \leftarrow \{0, 1\}$ (the challenge bit in the standard ABE security game) and returns the following challenge ciphertext for \mathcal{A} :

$$\text{ct}_{\mathbf{x}} := [\mathbf{C}_0]_1, \{[\mathbf{C}_{1,i}]_1, [\mathbf{C}_{2,i}]_1\}_{i: x_i=1}, e([\mathbf{C}_0]_1, [\mathbf{v}]_2) \cdot M_b$$

- For any secret keys requested, say for formula f , \mathcal{A}' computes $(\{\mathbf{v}_j\}, \rho) \leftarrow \text{share}(f, \mathbf{v})$, draws $\tilde{\mathbf{r}}_j \leftarrow_{\mathbb{R}} \mathbb{Z}_p^k$ and forms the following key:

$$\text{sk}_f := \{[\mathbf{v}_j + \mathbf{K}_{0,\rho(j)} \tilde{\mathbf{r}}_j]_2, [\mathbf{K}_{1,\rho(j)} \tilde{\mathbf{r}}_j]_2, [\mathbf{K}_{2,\rho(j)} \tilde{\mathbf{r}}_j]_2\}$$

Notice that if $\beta = 0$ (the input to \mathcal{A}' was drawn from distribution \mathbb{D}_0 defined in Lemma 15), then the challenge $\text{ct}_{\mathbf{x}}$ and all sk_f are Normal, and if $\beta = 1$ (the input to \mathcal{A}' was drawn from distribution \mathbb{D}_1), then $\text{ct}_{\mathbf{x}}$ is distributed as a SF ciphertext and all sk_f are distributed as P-SF keys.

Putting everything together, for $\beta \in \{0, 1\}$, when \mathcal{A}' interacts with \mathbb{D}_β , then \mathcal{A}' simulates H_β . It follows then that:

$$|\Pr[\langle \mathcal{A}, H_0 \rangle = 1] - \Pr[\langle \mathcal{A}, H_1 \rangle = 1]| \leq |\Pr[\mathcal{A}'(\mathbb{D}_0) = 1] - \Pr[\mathcal{A}'(\mathbb{D}_1) = 1]|$$

From Lemma 15, we then have:

$$|\Pr[\langle \mathcal{A}, H_0 \rangle = 1] - \Pr[\langle \mathcal{A}, H_1 \rangle = 1]| \leq (5n + 1) \cdot \text{Adv}_{\mathcal{B}'}^{k\text{-LIN}}(\lambda)$$

□

Lemma 17 ($H_{2,\ell-1} \approx_c H_{2,\ell}$).

$$|\Pr[\langle \mathcal{A}, H_{2,\ell-1} \rangle = 1] - \Pr[\langle \mathcal{A}, H_{2,\ell} \rangle = 1]| \leq 2^{6d} \cdot 8^d \cdot n \cdot \text{Adv}_{\mathcal{G}^*}^{k\text{-LIN}}(\lambda)$$

Proof. For each $\beta \in \{0, 1\}$, consider the following adversary \mathcal{A}' in $\mathcal{G}_\beta^{1\text{-ABE}}$ which internally simulates \mathcal{A} and the challenger in the ABE security game:

- First, \mathcal{A}' samples $\mathbf{A}_1, \mathbf{A}_3 \leftarrow \mathbb{Z}_p^{k \times (2k+1)}, \mathbf{A}_2 \leftarrow \mathbb{Z}_p^{1 \times (2k+1)}, \mathbf{W}, \mathbf{W}_0, \mathbf{W}_1 \leftarrow \mathbb{Z}_p^{(2k+1) \times k}, \tilde{\mathbf{v}} \leftarrow \mathbb{Z}_p^{2k+1}$, samples $(\mathbf{U}_i^{(2)} \in \mathbb{Z}_p^{(2k+1) \times k}), (\tilde{\mathbf{V}}_i^{(2)} \in \mathbb{Z}_p^{(2k+1) \times k}) \leftarrow \text{span}^k(\mathbf{A}_2^\parallel)$ and $(\mathbf{a}^\perp \in \mathbb{Z}_p^{2k+1}) \leftarrow \text{span}(\mathbf{A}_2^\parallel) \setminus \{\mathbf{0}\}$, and implicitly defines

$$\mathbf{v} := \tilde{\mathbf{v}} + \mu^{(0)} \mathbf{a}^\perp, \quad \mathbf{V}_i^{(2)} := \tilde{\mathbf{V}}_i^{(2)} + \mathbf{a}^\perp \mathbf{w}_i$$

where $\mu^{(0)} \in \mathbb{Z}_p, \mathbf{w}_i \in \mathbb{Z}_p^k$ is chosen in $\mathcal{G}_\beta^{1\text{-ABE}}$. (Note that \mathbf{v} is distributed randomly in \mathbb{Z}_p^{2k+1} and $\mathbf{V}_i^{(2)}$ is distributed like the output of $\text{span}^k(\mathbf{A}_2^\parallel)$). Then, \mathcal{A}' outputs:

$$\text{mpk} := ([\mathbf{A}_1]_1, [\mathbf{A}_1 \mathbf{W}]_1, [\mathbf{A}_1 \mathbf{W}_0]_1, [\mathbf{A}_1 \mathbf{W}_1]_1, e([\mathbf{A}_1]_1, [\tilde{\mathbf{v}}]_2)),$$

- When \mathcal{A} requests a challenge ciphertext for attribute \mathbf{x} along with M_0, M_1 , \mathcal{A}' queries $\mathcal{O}_X(\mathbf{x}) \rightarrow (\{\mathbf{w}_i\}_{x_i=1})$ in $\mathcal{G}_\beta^{1\text{-ABE}}$. \mathcal{A}' then samples $\mathbf{c}, \mathbf{c}_i \leftarrow \mathbb{Z}_p^{2k+1}$ and $b \leftarrow \{0, 1\}$ (the challenge bit in the standard ABE security game) and returns the following (SF) challenge ciphertext for \mathcal{A} :

$$\text{ct}_{\mathbf{x}} := \left([\mathbf{c}^\top]_1, \left\{ [\mathbf{c}^\top \underbrace{(\mathbf{W} + \tilde{\mathbf{V}}_i^{(2)} + \mathbf{a}^\perp \mathbf{w}_i)}_{\tilde{\mathbf{V}}_i} + \mathbf{c}_i^\top \underbrace{\mathbf{W}_0 + i \cdot \mathbf{W}_1 + \mathbf{U}_i^{(2)}}_{\tilde{\mathbf{U}}_i}]_1, [\mathbf{c}_i^\top]_1 \right\}_{x_i=1}, e([\mathbf{c}^\top]_1, \underbrace{[\tilde{\mathbf{v}} + \mu^{(0)} \mathbf{a}^\perp]_2}_{=\mathbf{v}}) \cdot M_b \right)$$

- For the first $\ell - 1$ secret keys requested, say for formula f , \mathcal{A}' computes

$$(\{\mathbf{v}_j\}, \rho) \leftarrow \text{share}(f, \underbrace{\tilde{\mathbf{v}} + \delta \mathbf{a}^\perp}_{=\mathbf{v} + \delta \mathbf{a}^\perp})$$

where $\tilde{\delta} \leftarrow \mathbb{Z}_p$ is drawn independently for each key (here, the per-key $\delta = \tilde{\delta} - \mu^{(0)}$ implicitly). Next, for each j , it queries $\mathcal{O}_E(\rho(j), [0]_2) \rightarrow ([\mathbf{w}_{\rho(j)}^\top \mathbf{r}_j]_2, [\mathbf{r}_j]_2)$ in $\mathcal{G}_\beta^{1\text{-ABE}}$ (since $\mathcal{O}_E(\rho(j), [0]_2) = \text{CPA.Enc}_{\mathbf{w}_{\rho(j)}}([0]_2)$), and forms the following (SF) key:

$$\text{sk}_f := \left(\left\{ \underbrace{[\mathbf{v}_j + (\mathbf{W} + \tilde{\mathbf{V}}_{\rho(i)}^{(2)}) \mathbf{r}_j + \mathbf{a}^\perp \mathbf{w}_{\rho(j)}^\top \mathbf{r}_j]_2}_{\mathbf{v}_j + \tilde{\mathbf{V}}_{\rho(j)} \mathbf{r}_j}, [\mathbf{r}_j]_2, \left\{ \underbrace{[(\mathbf{W}_0 + \rho(j) \cdot \mathbf{W}_1 + \mathbf{U}_{\rho(j)}^{(2)}) \mathbf{r}_j]_2}_{\tilde{\mathbf{U}}_{\rho(j)} \mathbf{r}_j} \right\}_{\rho(j) \neq 0}, \{[\mathbf{v}_j]_2\}_{\rho(j)=0} \right)$$

- For the last $Q - \ell$ secret keys requested, say for formula f , \mathcal{A}' proceeds as before for the first $\ell - 1$ keys except

$$(\{\mathbf{v}_j\}, \rho) \leftarrow \text{share}(f, \underbrace{\tilde{\mathbf{v}} + \mu^{(0)} \mathbf{a}^\perp}_{=\mathbf{v}})$$

It is easy to see that it forms a P-SF key.

- For the ℓ th secret key requested, say for formula f , \mathcal{A}' computes $(\{\mathbf{v}_j\}, \rho) \leftarrow \text{share}(f, \tilde{\mathbf{v}})$, queries $\mathcal{O}_F(f) \rightarrow (\{[\mu_j + \mathbf{w}_{\rho(j)}^\top \mathbf{r}_j]_2, [\mathbf{r}_j]_2\})$ in $\mathcal{G}_\beta^{1\text{-ABE}}$, then uses these components to return:

$$\text{sk}_f := \left(\left\{ \underbrace{[\mathbf{v}_j + (\mathbf{W} + \tilde{\mathbf{V}}_{\rho(i)}^{(2)}) \mathbf{r}_j + \mathbf{a}^\perp (\mu_j + \mathbf{w}_{\rho(j)}^\top \mathbf{r}_j)]_2}_{(\mathbf{v}_j + \mu_j \mathbf{a}^\perp) + \tilde{\mathbf{V}}_{\rho(j)} \mathbf{r}_j}, [\mathbf{r}_j]_2, \left\{ \underbrace{[(\mathbf{W}_0 + \rho(j) \cdot \mathbf{W}_1 + \mathbf{U}_{\rho(j)}^{(2)}) \mathbf{r}_j]_2}_{\tilde{\mathbf{U}}_{\rho(j)} \mathbf{r}_j} \right\}_{\rho(j) \neq 0}, \{[\mathbf{v}_j]_2\}_{\rho(j)=0} \right)$$

We claim that if $\beta = 0$, then sk_f is a P-SF key, and if $\beta = 1$, then sk_f is a SF key. This follows the fact that thanks to linearity, the shares

$$(\{\mathbf{v}_j + \mu_j \mathbf{a}^\perp\}, \rho), \text{ where } (\{\mathbf{v}_j\}, \rho) \leftarrow \text{share}(f, \tilde{\mathbf{v}}), (\{\mu_j\}, \rho) \leftarrow \text{share}(f, \mu^{(\beta)})$$

are identically distributed to $\text{share}(f, \tilde{\mathbf{v}} + \mu^{(\beta)} \mathbf{a}^\perp)$. The claim then follows from the fact that $\tilde{\mathbf{v}} + \mu^{(0)} \mathbf{a}^\perp = \mathbf{v}$ and that $\tilde{\mathbf{v}} + \mu^{(1)} \mathbf{a}^\perp$ is identically distributed to $\mathbf{v} + \delta \mathbf{a}^\perp$ (where $\delta = \mu^{(1)} - \mu^{(0)}$ is a fresh random value for this key).

Putting everything together, for $\beta \in \{0, 1\}$, when \mathcal{A}' interacts with $\mathbf{G}_\beta^{1\text{-ABE}}$, then \mathcal{A}' simulates $\mathbf{H}_{2, \ell-1+\beta}$. It follows then that:

$$|\Pr[\langle \mathcal{A}, \mathbf{H}_{2, \ell-1} \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_{2, \ell} \rangle = 1]| \leq |\Pr[\langle \mathcal{A}', \mathbf{G}_0^{1\text{-ABE}} \rangle = 1] - \Pr[\langle \mathcal{A}', \mathbf{G}_1^{1\text{-ABE}} \rangle = 1]|$$

From Theorem 2, we then have:

$$|\Pr[\langle \mathcal{A}, \mathbf{H}_{2, \ell-1} \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_{2, \ell} \rangle = 1]| \leq 2^{6d} \cdot 8^d \cdot n \cdot \text{Adv}_{\mathcal{B}^*}^{k\text{-LIN}}(\lambda)$$

□

Lemma 18 ($\mathbf{H}_{2, Q} \approx_s \mathbf{H}_3$).

$$|\Pr[\langle \mathcal{A}, \mathbf{H}_{2, Q} \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_3 \rangle = 1]| \leq \frac{1}{p}$$

Proof. These two hybrids are identically distributed conditioned on $\mathbf{c}^\top \mathbf{a}^\perp \neq 0$. To see this, consider two ways of sampling \mathbf{v} : as $\tilde{\mathbf{v}} \leftarrow \mathbb{Z}_p^{2k+1}$ and as $\tilde{\mathbf{v}} + \tilde{m} \mathbf{a}^\perp$ for an independent $\tilde{m} \leftarrow \mathbb{Z}_p$. Note that both result in \mathbf{v} having a uniform distribution.

Using $\tilde{\mathbf{v}}$ to simulate hybrid $\mathbf{H}_{2, Q}$ obviously results in $\mathbf{H}_{2, Q}$ (where $\mathbf{v} = \tilde{\mathbf{v}}$). However, using the identically distributed $\mathbf{v} = \tilde{\mathbf{v}} + \tilde{m} \mathbf{a}^\perp$ to simulate $\mathbf{H}_{2, Q}$ results in \mathbf{H}_3 (where $\tilde{M} = M \cdot e([\mathbf{c}^\top]_1, [\tilde{m} \mathbf{a}^\perp]_2)$ is randomly distributed as long as $\mathbf{c}^\top \mathbf{a}^\perp \neq 0$, and for redefined independently random $\tilde{\delta}_i := \delta_i + \tilde{m}$ in the secret keys).

\mathbf{c} is chosen at random and independent from $\mathbf{a}^\perp \neq \mathbf{0}$, so $\mathbf{c}^\top \mathbf{a}^\perp = 0$ with probability $\frac{1}{p}$, and since we know that $\mathbf{H}_{2, Q} \equiv \mathbf{H}_3$ conditioned on $\mathbf{c}^\top \mathbf{a}^\perp \neq 0$, then we have:

$$|\Pr[\langle \mathcal{A}, \mathbf{H}_{2, Q} \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_3 \rangle = 1]| \leq \frac{1}{p}$$

□

Theorem 5 (adaptive unbounded KP-ABE). *The unbounded KP-ABE construction in Appendix B.1 is adaptively secure under the MDDH_k assumption.*

Proof.

$$\begin{aligned} |\Pr[\langle \mathcal{A}, \mathbf{H}_0 \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_3 \rangle = 1]| &\leq |\Pr[\langle \mathcal{A}, \mathbf{H}_0 \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_1 \rangle = 1]| \\ &+ \sum_{\ell=1}^Q |\Pr[\langle \mathcal{A}, \mathbf{H}_{2, \ell-1} \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_{2, \ell} \rangle = 1]| \\ &+ |\Pr[\langle \mathcal{A}, \mathbf{H}_{2, Q} \rangle = 1] - \Pr[\langle \mathcal{A}, \mathbf{H}_3 \rangle = 1]| \end{aligned}$$

(Since $H_1 \equiv H_{2,0}$). Summing the results of Lemmas 16, 17, and 18, we then have:

$$|\Pr[\langle \mathcal{A}, H_0 \rangle = 1] - \Pr[\langle \mathcal{A}, H_3 \rangle = 1]| \leq (5n + 1) \cdot \text{Adv}_{\mathcal{B}'}^{k\text{-LIN}}(\lambda) + Q \cdot 2^{6d} \cdot 8^d \cdot n \cdot \text{Adv}_{\mathcal{B}^*}^{k\text{-LIN}}(\lambda) + \frac{1}{p}$$

If $d = O(\log n)$, then under the k -Lin Assumption this is a negligible function of λ (the number of queries made Q and the attribute vector length n are both polynomial in λ , and $\frac{1}{p}$ is a negligible function of λ). It's easy to see that $\text{Adv}_{\mathcal{A}}^{\text{ABE}}(\lambda) = 0$ in the H_3 hybrid game (since a random message is encrypted in the challenge ciphertext). So, any adversary in the real game (H_0) will have advantage negligibly close to 0, and our construction satisfies adaptive security. \square