

A Leakage-Resilient Spatial Encryption Scheme

Jill-Jènn Vie, Michel Abdalla, Crypto Team

August 23, 2011

This report is written in English because it is part of an upcoming research article.

The general context

A scheme is said leakage resilient if it remains secure even when an adversary is able to learn partial information about some secret values used throughout the lifetime of the system.

This recent area appeared because of the need to develop schemes that resist to side-channel attacks, such as power-consumption, fault or time analyses.

Today, there exist a few encryption schemes that implement this security model, notably a leakage-resilient identity-based encryption scheme (IBE): a public-key encryption scheme where the public keys are the identities, and a master key allows the generation of a secret key for each identity.

The research problem

There exists a hierarchical variant of the IBE scheme (HIBE), in which the identities form the nodes of a tree called hierarchy. Any identity can delegate secret keys to its successors and the root identity has a master key. Each node is associated with an identity vector, which represents the sequence of identities from the root to this node.

A spatial encryption scheme is a generalization of HIBE, in which keys are associated to affine spaces. In spatial encryption, any user possessing a secret key for an affine space can delegate secret keys for any of its affine subspaces, for the natural inclusion. To date, none of the existing constructions of spatial encryption are leakage resilient and the goal of this project is to propose one such construction.

Contributions

In this work, we propose the first leakage-resilient spatial encryption scheme. Our scheme works over bilinear groups of composite order, has constant-size ciphertexts, and is reasonably efficient. To achieve this goal, we adapted an existing leakage-resilient HIBE

construction due to Lewko, Rouselakis, and Waters in TCC 2011 to the spatial encryption setting. More precisely, we showed how to combine the Lewko-Rouselakis-Waters technique for achieving leakage resilience with a previous construction of spatial encryption due to Boneh and Hamburg in ASIACRYPT 2008.

Arguments supporting its validity

In order to show evidence that our new construction of leakage-resilient spatial encryption is sound, we also prove it secure in a formal security model. The proof of security is based on existing standard complexity assumptions over bilinear groups of composite order. As we show in the report, the resulting scheme remains secure even when the adversary learns partial information about the secret keys associated with each affine space, including the entire space.

Summary and future work

In this work, we have designed the first leakage-resilient spatial encryption scheme. Since spatial encryption includes identity-based and hierarchical identity-based encryption as special cases, previous solutions such as the Lewko-Rouselakis-Waters leakage-resilient HIBE construction can be seen as a particular instantiation of our construction. In addition to having constant-size ciphertexts and being reasonably efficient, our scheme is proven secure based on existing standard complexity assumptions over bilinear groups of composite order.

Unfortunately, as mentioned in the report, our scheme does not tolerate leakage during key updates. To address this problem and obtain a scheme with stronger security guarantees, a promising future research direction would be to adapt a recent result by Lewko, Lewko and Waters in STOC 2011 to the spatial encryption scenario.

A Leakage-Resilient Spatial Encryption Scheme

Jill-Jênn Vie
Supervisor: Michel Abdalla
Crypto Team

August 23, 2011

Contents

1	Introduction	2
1.1	Leakage Models	2
1.2	Spatial Encryption	2
1.3	Outline	3
2	Preliminaries	3
2.1	Composite Order Bilinear Groups	3
2.2	Complexity Assumptions	4
3	Generalized Identity-Based Encryption	5
3.1	Identity-Based Encryption	5
3.2	Hierarchical Identity-Based Encryption	5
3.3	Spatial Encryption	6
4	Our New Leakage-Resilient Spatial Encryption Scheme	7
4.1	A Few Useful Notations	7
4.2	Construction	7
4.3	Security Definition	9
4.4	Proof Overview	12
5	Concluding Remarks	13
	References	16
A	Security Proof	18
A.1	Semi-functionality	18
A.2	Lemma for Leakage Analysis	18
A.3	Proof	19

1 Introduction

When a theoretical scientist designs a cryptosystem, he is often unpleasantly surprised by the weaknesses its physical implementation can induce¹. Indeed, history has proven that these attempts to break the system, called *side-channel attacks*, can be devastating.

In order to design schemes that can provably resist these attacks, a stronger notion of security has been established. A *leakage-resilient* scheme considers that an attacker is able to learn partial information about some secret values used throughout the lifetime of the system.

1.1 Leakage Models

Our achievement of leakage resilience will be based on the model of *memory attacks*, introduced by Akavia, Goldwasser, and Vaikuntanathan [AGV09]. According to it, the adversary can learn *arbitrary* information about the secret state of a system, by selecting polynomial-time computable functions $f_i : \{0, 1\}^* \rightarrow \{0, 1\}^{\lambda_i}$ and learning the value of f_i applied to the internal state of the system.

As described in [BSW11], there are two main variants of the memory attacks model, known as the bounded-leakage and the continual-leakage models.

In a *bounded-leakage* model, we define a *leakage bound* λ such that the overall amount of information learned by the adversary throughout the entire lifetime of the system is $\sum_i \lambda_i \leq \lambda$.

If a system is used continually for a sufficiently long time, then any static piece of information that stays unmodified on the system can eventually be recovered by the adversary. Thus, the secret keys of such systems must be periodically *refreshed*. Recently, Dodis et al. [DHLAW10] suggested the *continual-leakage* model, in which a scheme periodically self-refreshes its internal secret key, while the corresponding public key remains fixed. Therefore, it is the amount of leakage observed by the adversary in between any two successive refreshes which is bounded by λ . We can think of λ as a *leakage rate*.

1.2 Spatial Encryption

The goal of this work is to consider leakage resilience in the context of spatial encryption, which is a generic public-key cryptosystem first defined in [BH08] where vectors play the role of public keys and secret keys are associated to affine spaces. In a spatial encryption scheme, any user possessing a secret key for a certain space W_1 can *delegate* a secret key for a space W_2 included in W_1 . A message encrypted for a certain vector \mathbf{u} can be decrypted by any user possessing a secret key associated to a space that contains \mathbf{u} .

¹Such as those due to a constant random number (cf. Sony's ECDSA code in PS3), but this is another problem.

In addition to proposing the concept of spatial encryption, Boneh and Hamburg [BH08] also provided a construction based on bilinear maps which is provably secure in the selective security model based on the BDDHE assumption ([BGW05]). In the selective secure model, the adversary decides on the vector upon which he wants to be challenged before seeing the public key of the scheme. Their work was later improved by Moriyama and Doi in [MD11] who proposed a fully secure (as opposed to selective) spatial encryption scheme in the standard model based on bilinear groups of composite order. In fact, the scheme that we propose in this work can be seen as a generalization of Moriyama-Doi spatial encryption scheme to leakage-resilient setting.

1.3 Outline

In Section 2, we will state our complexity assumptions. In Section 3, we will recall the formal definition of spatial encryption along with those of identity-based and hierarchical identity-based encryption, which are special cases of the former. In Section 4, we will present our leakage-resilient spatial encryption scheme. It is worth noting that, to date, none of the existing constructions of spatial encryption are leakage resilient and that our scheme is the first one to achieve this level of security. We will end this report with a few concluding remarks.

2 Preliminaries

2.1 Composite Order Bilinear Groups

Composite order bilinear groups were first used in cryptography by [BGN05] (see also [Bon07]). We suppose the existence of an efficient group generator algorithm \mathbb{G} which takes as input the security parameter λ and outputs a description $\mathbb{I} = (N, \mathbb{G}, \mathbb{G}_T, \hat{e})$ of a bilinear setting, where \mathbb{G} and \mathbb{G}_T are cyclic groups of order N , and $\hat{e} : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ is a map with the following properties:

1. (bilinearity) $\forall g, h \in \mathbb{G}$ and $a, b \in \mathbb{Z}_N$ it holds that $\hat{e}(g^a, h^b) = \hat{e}(g, h)^{ab}$;
2. (non-degeneracy) $\exists g \in \mathbb{G}$ such that $\hat{e}(g, g)$ has order N in \mathbb{G}_T .

We assume that the group descriptions of \mathbb{G} and \mathbb{G}_T include generators of the respective cyclic groups. We require that the group operations in \mathbb{G} and \mathbb{G}_T as well as the bilinear map \hat{e} are computable in deterministic polynomial time in λ . In our construction, we will make hardness assumptions for bilinear settings whose order N is product of four distinct primes each of length $\Theta(\lambda)$. For an integer m dividing N , we let \mathbb{G}_m denote the subgroup of \mathbb{G} of order m . From the fact that the group is cyclic, it is easy to verify that if g and h are group elements of co-prime orders then $\hat{e}(g, h) = 1$. This is called the *orthogonality property* and is a crucial tool in our constructions.

2.2 Complexity Assumptions

To prove the security of our system, we will use the following assumptions in composite order groups, also used in [LOS⁺10, LW10, LRW11]. The first two of them belong to the class of Generalized Subgroup Decision Assumptions described in [BWY11]: in a bilinear group of order $N = p_1 \dots p_n$, for each subset $S \subseteq \{1, \dots, n\}$, there exists $G_S = G_{\prod_{i \in S} p_i}$ subgroup of order $\prod_{i \in S} p_i$. Let S_0, S_1 be two subsets. It is *hard to distinguish* a random element from G_{S_0} from a random element of G_{S_1} , even if one is given random elements from a family of (G_{Z_i}) such that each Z_i satisfy either $S_0 \cap Z_i = \emptyset = S_1 \cap Z_i$ or $S_0 \cap Z_i \neq \emptyset \neq S_1 \cap Z_i$.

Assumption 1. Given $\mathcal{D}_1 = (N, \mathbb{G}, \mathbb{G}_T, \hat{e}, g_1, g_3)$, no PPT (probabilistic polynomial-time) adversary has a non-negligible advantage in distinguishing

$$T_1 = g_1^a \text{ from } T_2 = g_1^a g_3^b$$

where $a, b \xleftarrow{R} \mathbb{Z}_N$.

The advantage of an algorithm \mathcal{A} in breaking Assumption 1 is defined as:

$$\mathbf{Adv}_1^{\mathcal{A}}(\lambda) = |\Pr[\mathcal{A}(\mathcal{D}_1, T_1) = 1] - \Pr[\mathcal{A}(\mathcal{D}_1, T_2) = 1]|$$

We say that Assumption 1 holds if for all PPT \mathcal{A} , $\mathbf{Adv}_1^{\mathcal{A}}(\lambda)$ is a negligible function of λ .

Assumption 2. Given $\mathcal{D}_2 = (N, \mathbb{G}, \mathbb{G}_T, \hat{e}, g_1, g_3, g_1^x g_2^y, g_2^{y'} g_3^{z'})$ where $x, y, y', z' \xleftarrow{R} \mathbb{Z}_N$, no PPT adversary has a non-negligible advantage in distinguishing

$$T_1 = g_1^a g_3^c \text{ from } T_2 = g_1^a g_2^b g_3^c$$

where $a, b, c \xleftarrow{R} \mathbb{Z}_N$.

The advantage of an algorithm \mathcal{A} in breaking Assumption 2 is defined as:

$$\mathbf{Adv}_2^{\mathcal{A}}(\lambda) = |\Pr[\mathcal{A}(\mathcal{D}_2, T_1) = 1] - \Pr[\mathcal{A}(\mathcal{D}_2, T_2) = 1]|$$

We say that Assumption 2 holds if for all PPT \mathcal{A} , $\mathbf{Adv}_2^{\mathcal{A}}(\lambda)$ is a negligible function of λ .

Assumption 3. Given $\mathcal{D}_3 = (N, \mathbb{G}, \mathbb{G}_T, \hat{e}, g_1, g_2, g_3, g_1^x g_2^y, g_1^{x'} g_2^{y'})$ where $x, y, x', y' \xleftarrow{R} \mathbb{Z}_N$, no PPT adversary has a non-negligible advantage in distinguishing

$$T_1 = \hat{e}(g_1, g_1)^{xx'} \in \mathbb{G}_T \text{ from } T_2 \xleftarrow{R} \mathbb{G}_T.$$

The advantage of an algorithm \mathcal{A} in breaking Assumption 3 is defined as:

$$\mathbf{Adv}_3^{\mathcal{A}}(\lambda) = |\Pr[\mathcal{A}(\mathcal{D}_3, T_1) = 1] - \Pr[\mathcal{A}(\mathcal{D}_3, T_2) = 1]|$$

We say that Assumption 3 holds if for all PPT \mathcal{A} , $\mathbf{Adv}_3^{\mathcal{A}}(\lambda)$ is a negligible function of λ .

3 Generalized Identity-Based Encryption

3.1 Identity-Based Encryption

An identity-based encryption system is a public-key cryptosystem which allows users to encrypt messages knowing only the recipient's identity and some public parameters. It consists of four PPT algorithms.

Setup $(1^\lambda) \rightarrow (\text{PP}, \text{MK})$ The setup algorithm takes an integer security parameter λ as input and outputs the public parameters PP and the original master key MK . In the rest of this report, all algorithms will take implicitly the security parameter and the public parameters as inputs.

KeyGen $(\text{MK}', X) \rightarrow K$ The key generation algorithm takes in a master key MK' and either $X = I$, an identity, or $X = \varepsilon$, the empty string. In the former case, it outputs a secret key $K = \text{SK}$, for the identity I . In the latter case, it outputs another master key, $K = \text{MK}''$, such that $|\text{MK}'| = |\text{MK}''|$. This new master key can now be used instead of the original key in calls of **KeyGen**.

Encrypt $(M, I) \rightarrow \text{CT}$ The encryption algorithm takes in a message M and an identity I . It outputs a ciphertext CT .

Decrypt $(\text{CT}, \text{SK}) \rightarrow M$ The decryption algorithm takes in a ciphertext CT for message M and a secret key SK . It outputs the message M .

The correctness requirement is that if the identity I used during the encryption of M is the same as the identity of the secret key used during decryption, then the output of **Decrypt** is the message M . That is, for all PP generated by a call to **Setup**, for all master keys MK' generated by applying the **KeyGen** algorithm and for all M, I ,

$$\text{Decrypt}(\text{Encrypt}(M, I), \text{KeyGen}(\text{MK}', I)) = M.$$

3.2 Hierarchical Identity-Based Encryption

In an HIBE system, identities form a structured hierarchy: a user can delegate keys to its subordinate identities, and thus decrypt any message encrypted to them.

We handle identity vectors $\mathbf{I} = (I_1, \dots, I_j)$, and define the children of \mathbf{I} as all vectors $\mathbf{I}||I_x$, where $||$ denotes concatenation and I_x is an identity. A user that has a secret key for vector \mathbf{I} can create secret keys for all users that have identity vectors prefixed by \mathbf{I} . An example of hierarchy can be seen in Figure 1.

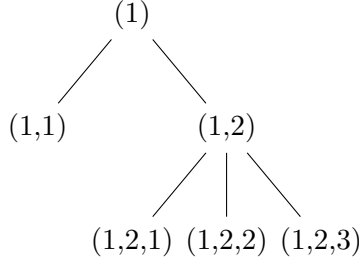


Figure 1: Vector identities in a hierarchy.

The depth D of the tree is the maximum length of the identity vectors and is called the depth of the hierarchy. The master key is formulated as a secret key for level 0, thus the corresponding identity vector is the empty vector.

$\left. \begin{array}{l} \mathbf{Setup}(1^\lambda) \rightarrow (\text{PP}, \text{MK}) \\ \mathbf{KeyGen}(\text{MK}', \mathbf{I}) \rightarrow K \\ \mathbf{Encrypt}(M, \mathbf{I}) \rightarrow \text{CT} \end{array} \right\}$ have the same description than in the IBE scheme.

$\mathbf{Delegate}(\mathbf{I}, \text{SK}_{\mathbf{I}}, I_x) \rightarrow \text{SK}_{\mathbf{I}||I_x}$ The delegation algorithm allows a user that has a secret key for identity vector \mathbf{I} to construct secret keys down the hierarchy. The algorithm takes in an identity vector \mathbf{I} , a secret key for that identity vector and an identity I_x . It outputs a secret key for the identity vector $\mathbf{I}||I_x$.

$\mathbf{Decrypt}(\text{CT}, \text{SK}) \rightarrow M$ The decryption algorithm takes in a ciphertext CT for message M and for identity vector \mathbf{I}_c , and a secret key SK for identity vector \mathbf{I}_k . If \mathbf{I}_k is a prefix of \mathbf{I}_c , it outputs the message M .

3.3 Spatial Encryption

In a spatial encryption scheme, the secret keys are associated to affine subspaces of \mathbb{Z}_N^n , and the delegation relation is defined by subspace inclusion.

Spatial encryption is a generalization of HIBE. Indeed, if we fix $n = D$ and consider an identity vector $\mathbf{I} = (I_1, \dots, I_j)$, we can associate it in a bijective manner to the affine space $\{(I_1, \dots, I_j, x_{j+1}, \dots, x_n) | x_{j+1}, \dots, x_n \in \mathbb{Z}_N\}$. The delegation relation is respected.

$\left. \begin{array}{l} \mathbf{Setup}(1^\lambda) \rightarrow (\text{PP}, \text{MK}) \\ \mathbf{KeyGen}(\text{SK}_E, W) \rightarrow K \\ \mathbf{Encrypt}(m, \mathbf{u}) \rightarrow \text{CT} \end{array} \right\}$ have the same description than in the IBE scheme.

$\mathbf{Delegate}(W_1, \text{SK}_{W_1}, W_2) \rightarrow \text{SK}_{W_2}$ The algorithm takes in an affine space W_1 , a secret key for that space, and a subspace W_2 included in W_1 . It outputs a secret key for W_2 .

Decrypt(CT, \mathbf{u} , W , SK_W) $\rightarrow m$ The decryption algorithm takes in a ciphertext CT for message m and for vector \mathbf{u} , and a secret key SK for space W . If W contains \mathbf{u} , it first delegates SK_W to obtain $SK_{\mathbf{u}}$. Then it outputs the message m .

4 Our New Leakage-Resilient Spatial Encryption Scheme

We now present our construction of a leakage-resilient spatial encryption scheme with constant-size ciphertext, which can be seen as an extension of the leakage-resilient HIBE scheme in [LRW11]. The security proof of this system is available in Appendix A.

4.1 A Few Useful Notations

For a vector $v = (v_1, \dots, v_n)^\top \in \mathbb{Z}_N^n$ of field elements, we use g^v to denote the vector of group elements

$$g^v = (g^{v_1}, \dots, g^{v_n})^\top \in \mathbb{G}^n$$

To simplify the notation, we introduce an operator

$$\begin{aligned} \psi : \mathbb{G}^n \times \mathbb{Z}_N^n &\rightarrow \mathbb{G} \\ (g^v, w) &\mapsto g^{\langle v, w \rangle} \end{aligned}$$

and its extension:

$$\begin{aligned} \psi^* : \mathbb{G}^n \times \mathbb{Z}_N^{n \times d} &\rightarrow \mathbb{G}^d \\ \left(g^v, \left(\mathbf{C}_1, \dots, \mathbf{C}_d \right) \right) &\mapsto \left(\psi(g^v, \mathbf{C}_1), \dots, \psi(g^v, \mathbf{C}_d) \right)^\top \end{aligned}$$

We can notice that ψ is easily computable.

4.2 Construction

The system parameters for our spatial encryption system will be primes p_1, p_2, p_3 (where each $\log p_i$ is approximately the security parameter λ) and two groups \mathbb{G} and \mathbb{G}_T of order $N = p_1 p_2 p_3$, with a bilinear pairing $\hat{e} : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$. Additionally, the public parameters will include group elements $g_1 \in \mathbb{G}_{p_1}, g_3 \in \mathbb{G}_{p_3}, g_1^\varphi \in \mathbb{G}, t \in \mathbb{G}_T$ and a vector $g_1^\alpha \in \mathbb{G}^n$.

A secret key for an affine space $W = \text{Aff}(M, \mathbf{u})$ of dimension d will have the form

$$(\mathbf{k}_\rho, k_r, k_u, \mathbf{k}_{\text{del}}) = \left(g_1^\rho, g_1^r, g_1^{\tau+r(\varphi+\langle \mathbf{u}, \alpha \rangle) - \langle \rho, \sigma \rangle}, g_1^{rM^\top \alpha} \right) \times g_3^\mu \in \mathbb{G}^{d+n+2}.$$

Setup(λ) generates the system parameters $N = p_1 p_2 p_3, \mathbb{G}, \mathbb{G}_T$. It then chooses parameters

$$g_1 \xleftarrow{R} \mathbb{G}_{p_1} \quad g_3 \xleftarrow{R} \mathbb{G}_{p_3} \quad \varphi \xleftarrow{R} \mathbb{Z}_N \quad \boldsymbol{\alpha}, \boldsymbol{\rho}, \boldsymbol{\sigma} \xleftarrow{R} \mathbb{Z}_N^n \quad \boldsymbol{\mu} \xleftarrow{R} \mathbb{Z}_N^{2n+2}$$

and secret parameter $\tau \xleftarrow{R} \mathbb{G}$, then computes $t = \hat{e}(g_1, g_1)^\tau$. It outputs public parameters

$$\text{PP} = (N, g_1, g_3, g_1^\varphi, g_1^\alpha, t, g_1^\sigma)$$

and secret key

$$\text{SK}_E = \left(g_1^\rho, g_1^r, g_1^{\tau+r\varphi-\langle \boldsymbol{\rho}, \boldsymbol{\sigma} \rangle}, g_1^{r\alpha} \right) \times g_3^\mu.$$

Delegate($W_1, \text{SK}_{W_1}, W_2$) takes two subspaces $W_1 = \text{Aff}(M_1, \mathbf{u}_1)$ and $W_2 = \text{Aff}(M_2, \mathbf{u}_2)$ and a key SK_{W_1} under the form $(\mathbf{k}_\rho, k_r, \mathbf{k}_{\mathbf{u}_1}, \mathbf{k}_{\text{del}})$. Let $\boldsymbol{\mu}$ be the \mathbb{G}_{p_3} part of SK_{W_1} . We similarly divide $\boldsymbol{\mu}$ into $(\boldsymbol{\mu}_\rho, \mu_r, \boldsymbol{\mu}_{\mathbf{u}_1}, \boldsymbol{\mu}_{\text{del}})$. Since W_2 is a subspace of W_1 , we must have $M_2 = M_1 T$ and $\mathbf{u}_2 = \mathbf{u}_1 + M_1 \mathbf{v}$ for some matrix T and vector \mathbf{v} . We can then compute a key for W_2 :

$$\begin{aligned} \widehat{\text{SK}}_{W_2} &= (\mathbf{k}_\rho, k_r, \underbrace{\mathbf{k}_{\mathbf{u}_1} \cdot \psi(\mathbf{k}_{\text{del}}, \mathbf{v})}_{k_{\mathbf{u}_2}}, \underbrace{\psi^*(\mathbf{k}_{\text{del}}, T)}_{\mathbf{k}'_{\text{del}}}) \\ &= \left(g_1^\rho, g_1^r, g_1^{\tau+r(\varphi+\langle \mathbf{u}_1, \boldsymbol{\alpha} \rangle) - \langle \boldsymbol{\rho}, \boldsymbol{\sigma} \rangle} \cdot g_1^{r\mathbf{v}^\top M_1^\top \boldsymbol{\alpha}}, g_1^{rT^\top M_1^\top \boldsymbol{\alpha}} \right) \times \left(g_3^{\boldsymbol{\mu}_\rho}, g_3^{\mu_r}, g_3^{\boldsymbol{\mu}_{\mathbf{u}_1} + \langle \mathbf{v}, \boldsymbol{\mu}_{\text{del}} \rangle}, g_3^{T^\top \boldsymbol{\mu}_{\text{del}}} \right) \\ &= \left(g_1^\rho, g_1^r, g_1^{\tau+r(\varphi+\langle \mathbf{u}_1, \boldsymbol{\alpha} \rangle) - \langle \boldsymbol{\rho}, \boldsymbol{\sigma} \rangle + r\langle M_1 \mathbf{v}, \boldsymbol{\alpha} \rangle}, g_1^{rM_2^\top \boldsymbol{\alpha}} \right) \times g_3^{\boldsymbol{\mu}' } \\ &= \left(g_1^\rho, g_1^r, g_1^{\tau+r(\varphi+\langle \mathbf{u}_2, \boldsymbol{\alpha} \rangle) - \langle \boldsymbol{\rho}, \boldsymbol{\sigma} \rangle}, g_1^{rM_2^\top \boldsymbol{\alpha}} \right) \times g_3^{\boldsymbol{\mu}' } \end{aligned}$$

where $\boldsymbol{\mu}' = (\boldsymbol{\mu}_\rho, \mu_r, \boldsymbol{\mu}_{\mathbf{u}_1} + \langle \mathbf{v}, \boldsymbol{\mu}_{\text{del}} \rangle, T^\top \boldsymbol{\mu}_{\text{del}})$.

Then we pick $\Delta r \xleftarrow{R} \mathbb{Z}_p$, $\Delta \boldsymbol{\rho} \xleftarrow{R} \mathbb{Z}_p^n$ and $\Delta \boldsymbol{\mu} \xleftarrow{R} \mathbb{Z}_p^{d+n+2}$ to re-randomize it:

$$\begin{aligned} \text{SK}_{W_2} &= \widehat{\text{SK}}_{W_2} \times \left(g_1^{\Delta \boldsymbol{\rho}}, g_1^{\Delta r}, g_1^{\Delta r(\varphi+\langle \mathbf{u}_2, \boldsymbol{\alpha} \rangle) - \langle \Delta \boldsymbol{\rho}, \boldsymbol{\sigma} \rangle}, g_1^{\Delta r M_2^\top \boldsymbol{\alpha}} \right) \times g_3^{\Delta \boldsymbol{\mu}} \\ &= \left(g_1^{\rho+\Delta \boldsymbol{\rho}}, g_1^{r+\Delta r}, g_1^{\tau+(r+\Delta r)(\varphi+\langle \mathbf{u}_2, \boldsymbol{\alpha} \rangle) - \langle \boldsymbol{\rho}+\Delta \boldsymbol{\rho}, \boldsymbol{\sigma} \rangle}, g_1^{(r+\Delta r)M_2^\top \boldsymbol{\alpha}} \right) \times g_3^{\boldsymbol{\mu}' + \Delta \boldsymbol{\mu}}. \end{aligned}$$

Notice that **Delegate**($W_1, \text{SK}_{W_1}, W_1$) is a re-randomization of SK_{W_1} .

Encrypt(m, \mathbf{u}), where m is encoded as an element of the target group \mathbb{G}_T , picks a random $s \xleftarrow{R} \mathbb{Z}_N^n$ and computes the ciphertext

$$\text{CT} = (\text{CT}^*, c_m) = (\mathbf{c}_\sigma, \mathbf{c}_u, c_s, c_m) = \left(g_1^{s\boldsymbol{\sigma}}, g_1^{-s(\varphi+\langle \mathbf{u}, \boldsymbol{\alpha} \rangle)}, g_1^s, m \cdot t^s \right) \in \mathbb{G}^{n+2} \times \mathbb{G}_T.$$

Decrypt(CT, \mathbf{u} , W , SK_W) where $\text{CT} = (c_\sigma, c_u, c_s, c_m)$ first delegates SK_W to obtain the key $\text{SK}_u = (\mathbf{k}_\rho, k_r, k_u) = \left(g_1^\rho, g_1^r, g_1^{\tau+r(\varphi+\langle \mathbf{u}, \boldsymbol{\alpha} \rangle) - \langle \rho, \boldsymbol{\sigma} \rangle}\right) \times g_3^\mu$. It then recovers

$$\begin{aligned} \frac{c_m}{\hat{e}_3(\text{CT}^*, \text{SK}_u)} &= \frac{c_m}{\hat{e}(c_\sigma, \mathbf{k}_\rho) \cdot \hat{e}(c_u, k_r) \cdot \hat{e}(c_s, k_u)} \\ &= \frac{m \cdot t^s}{\hat{e}(g_1, g_1)^{s\langle \rho, \boldsymbol{\sigma} \rangle - rs(\varphi + \langle \mathbf{u}, \boldsymbol{\alpha} \rangle) + s(\tau + r(\varphi + \langle \mathbf{u}, \boldsymbol{\alpha} \rangle) - \langle \rho, \boldsymbol{\sigma} \rangle)}} \\ &= \frac{m \cdot \hat{e}(g_1, g_1)^{s\tau}}{\hat{e}(g_1, g_1)^{s\tau}} = m. \end{aligned}$$

4.3 Security Definition

The security of our system is based on a game, called **MasterLeakSpatial**. In this game, the challenger first makes a call to **Setup**, to get a secret key for the whole space, and public parameters that it can give to the adversary. In a first phase, the adversary can make a polynomial number of queries, in any order:

- **Create** queries to store keys in an array;
- **Leak** queries to apply a leak function of his choice on any key (as long as the total number of leaked bits for a certain key does not exceed ℓ_{SK});
- or **Reveal** queries to recover a whole key (as long as it is not a key associated to the entire space, obviously).

At the end of this first phase, the adversary chooses a challenge vector \mathbf{u}^* such that none of the revealed spaces (associated to its revealed keys) contain it, as well as two messages m_0, m_1 . It gives them to the challenger, which chooses a random bit β and returns $\text{CT}^* = \mathbf{Encrypt}(m_\beta, \mathbf{u}^*)$. From then on, the second phase starts, and the adversary can make **Create** and **Reveal** queries on keys whose space does not contain \mathbf{u}^* , naturally. It finally has to suggest a value for β . If it guesses right, it wins.

The scheme is considered leakage resilient if no PPT adversary can win the **MasterLeakSpatial** game with probability significantly better than 1/2. Indeed, roles are symmetric with respect to 1/2: if the adversary always loses the game, it is as concerning as if it had always won.

In order to formally define the **MasterLeakSpatial** security game, let us first introduce the following parameters:

- the set $\mathcal{R} \in \mathcal{P}(E)$ of spaces for which a key has been revealed;
- the array $\mathcal{T} \subset E \times \text{SK} \times \mathbb{N}$ that holds tuples of the form (**space**, **key**, **leaked**). An example is shown in Figure 2.

	E	SK_E	0	W	SK_W	1	W'	$SK_{W'}$	0	E	SK'_E	5
	0			1			2			3		

Figure 2: Example of \mathcal{T} array being filled in the MasterLeakSpatial game. Each cell contains the triplet (space, key, leaked).

The formal definition of the game MasterLeakSpatial, which is described in Figure 3, consists of the following phases:

Setup. The challenger makes a call to **Setup**(1^λ) and gets a key SK for the whole space E and the public parameters PP. It gives PP to the adversary.

Phase 1. In this phase, the adversary can make any of the following queries to the challenger, in any possible way: **Create**, **Leak**, **Reveal**, **Delegate**. These algorithms are listed in Figure 3.

Challenge. The adversary chooses a challenge vector \mathbf{u}^* such that no element of \mathcal{R} contains it, as well as two messages m_0, m_1 of equal size. It makes a call to **LR**(\mathbf{u}^*, m_0, m_1) and obtains a ciphertext CT*.

Phase 2. This is the same as **Phase 1**, except the only allowed queries are **Create** and **Reveal** for secret keys whose space does not contain \mathbf{u}^* .

Guess. The adversary chooses a bit β' and calls **Finalize**(β'). If the output is **True**, it succeeds.

Game MasterLeakSpatial	
<p>procedure Initialize</p> $(PP, SK) \xleftarrow{R} \text{Setup}$ $\mathcal{R} \leftarrow \emptyset$ $\mathcal{T}[0] \leftarrow (E, SK, 0)$ $\text{nbKeys} \leftarrow 1$ $\beta \xleftarrow{R} \{0, 1\}$ Return PP <p>procedure Create(h, W)</p> If $\mathcal{T}[h].\text{space} = E$ then $SK_W \xleftarrow{R} \text{KeyGen}(\mathcal{T}[h].\text{key}, W)$ $\mathcal{T}[\text{nbKeys}] \leftarrow (W, SK_W, 0)$ $\text{nbKeys} \leftarrow \text{nbKeys} + 1$ Return SK_W Else Return \perp <p>procedure Delegate(h, W')</p> $SK_{W'} \leftarrow \text{Delegate}(\mathcal{T}[h].\text{space}, \mathcal{T}[h].\text{key}, W')$ $\mathcal{T}[\text{nbKeys}] \leftarrow (W', SK_{W'}, 0)$ $\text{nbKeys} \leftarrow \text{nbKeys} + 1$ Return $SK_{W'}$	<p>procedure Leak(h, f)</p> $SK \leftarrow \mathcal{T}[h].\text{key}$ If $\mathcal{T}[h].\text{leaked} + f(SK) \leq \ell_{SK}$ then $\mathcal{T}[h].\text{leaked} \leftarrow \mathcal{T}[h].\text{leaked} + f(SK) $ Return $f(SK)$ Else Return \perp <p>procedure Reveal(h)</p> If $\mathcal{T}[h].\text{space} \neq E$ then $\mathcal{R} \leftarrow \mathcal{R} \cup \mathcal{T}[h].\text{space}$ Return $\mathcal{T}[h].\text{key}$ Else Return \perp <p>procedure LR(u^*, m_0, m_1)</p> $CT^* \xleftarrow{R} \text{Encrypt}(u^*, m_\beta)$ Return CT^* <p>procedure Finalize(β')</p> Return $(\beta' = \beta)$

Figure 3: Algorithms of the game MasterLeakSpatial.

The security definition we will use is the following:

Definition 1. A spatial encryption system Π is ℓ_{SK} -master-leakage secure if for all PPT adversaries \mathcal{A} it is true that

$$\text{Adv}_{\mathcal{A}, \Pi}^{\text{MasterLeakSpatial}}(\lambda, \ell_{SK}) \leq \text{negl}(\lambda)$$

where $\text{Adv}_{\mathcal{A}, \Pi}^{\text{MasterLeakSpatial}}(\lambda, \ell_{SK})$ is the advantage of \mathcal{A} in game *MasterLeakSpatial* with security parameter λ and leakage parameter $\ell_{SK} = \ell_{SK}(\lambda)$ and is formally defined as:

$$\text{Adv}_{\mathcal{A}, \Pi}^{\text{MasterLeakSpatial}}(\lambda, \ell_{SK}) = \left| \Pr[\mathcal{A} \text{ succeeds}] - \frac{1}{2} \right|,$$

where the probability is over all random bits used by the challenger and the attacker.

4.4 Proof Overview

4.4.1 Dual System Encryption

The Dual System Encryption, first developed by Waters in [Wat09], is a methodology for proving security of encryption systems.

Both ciphertexts and secret keys will take on one of two indistinguishable forms: *normal*, or *semi-functional*. A semi-functional key cannot decrypt a semi-functional ciphertext. All other combinations work, as shown in Table 1.

This technique consists in gradually modifying the secret keys and ciphertexts of the system to make them *semi-functional*. When both secret keys and ciphertexts are semi-functional, proving security is straightforward.

Table 1: Different combinations of secret keys and ciphertexts.

secret key	ciphertext	decryption
normal	normal	correct (obviously)
normal	semi-functional	correct
semi-functional	normal	correct
semi-functional	semi-functional	incorrect

4.4.2 Hybrid Games

In order to prove that the advantage of the adversary in the `MasterLeakSpatial` game is negligible, we will define a sequence of games that are closely related to each other². We will prove that the difference in probability of winning for the adversary is negligible between two consecutive games of this sequence, until we get to a game where both keys and ciphertexts will be semi-functional, so the adversary will not have any chance of winning the game.

First, we define `MasterLeakSpatial*`, which is the same game as `MasterLeakSpatial` except that all `Delegate(h, W')` calls have been replaced by `Create(0, W')` calls. As the keys are identically distributed, the advantage of the adversary in `MasterLeakSpatial` is negligibly close to its advantage in `MasterLeakSpatial*`.

Then we define `MasterLeakC`, similar to `MasterLeakSpatial*` except that the challenge ciphertext is now semi-functional. Using Assumption 1, we will show that the advantages of the adversary in those games are negligibly close.

In `MasterLeakCK`, the semi-functional versions of the secret keys are stored in \mathcal{T} as well: it contains quadruplets of the form `(space, key, keySF, leaked)`. The `Create` calls use the normal versions of the keys while the `Leak` and `Reveal` queries apply to the

²The complete description of the algorithms of these games is available in Figures 5 and 6.

semi-functional versions. We will need Assumption 3 to prove that the advantage of the adversary is negligible in this game.

Finally, in the MasterLeakCK_j game, for the $j - 1$ first keys the attacker creates, it will have access to the semi-functional versions of them via **Leak** or **Reveal** queries. For the other keys, it will have access to the normal versions (see also Figure 4). Notice that MasterLeakCK_0 is exactly MasterLeakC while $\text{MasterLeakCK}_{\text{nbKeys}}$ is exactly MasterLeakCK . We will show that if Assumption 2 holds, the difference of probability of winning for the adversary in MasterLeakCK_j versus $\text{MasterLeakCK}_{j+1}$ is negligible, which will conclude the proof.

	W_0	$\widetilde{\text{SK}}_{W_0}$	4	W_1	$\widetilde{\text{SK}}_{W_1}$	6	W_2	SK_{W_2}	2	W_3	SK_{W_3}	3	
			$j - 2$			$j - 1$			j			$j + 1$	

Figure 4: Example of \mathcal{T} array being filled in the MasterLeakCK_j game. Each cell contains the quadruplet (space, key, keySF, leaked). The adversary has access to the semi-functional versions of the $j - 1$ first keys, and to the normal versions of the other keys.

5 Concluding Remarks

We have designed the first leakage-resilient spatial encryption scheme, which is an important instance of generalized identity-based encryption. Though reasonably efficient, our scheme does not tolerate leakage during key updates. To address this problem and obtain a scheme with stronger security guarantees, a promising direction would be to adapt the result of [LLW11] to the spatial encryption scenario.

Game MasterLeakC

procedure Initialize

$(PP, SK) \xleftarrow{R} \text{Setup}$
 $\mathcal{R} \leftarrow \emptyset$
 $\mathcal{T}[0] \leftarrow (E, SK, 0)$
 $\text{nbKeys} \leftarrow 1$
 $\beta \xleftarrow{R} \{0, 1\}$
 Return PP

procedure Create(h, W)

If $\mathcal{T}[h].\text{space} = E$ then
 $SK_W \xleftarrow{R} \text{KeyGen}(\mathcal{T}[h].\text{key}, W)$
 $\mathcal{T}[\text{nbKeys}] \leftarrow (W, SK_W, 0)$
 $\text{nbKeys} \leftarrow \text{nbKeys} + 1$
 Return SK_W

Else
 Return \perp

procedure Delegate(h, W')

Return **Create**(0, W')

procedure Leak(h, f)

$SK \leftarrow \mathcal{T}[h].\text{key}$
 If $\mathcal{T}[h].\text{leaked} + |f(SK)| \leq \ell_{SK}$ then
 $\mathcal{T}[h].\text{leaked} \leftarrow \mathcal{T}[h].\text{leaked} + |f(SK)|$
 Return $f(SK)$
 Else
 Return \perp

procedure Reveal(h)

If $\mathcal{T}[h].\text{space} \neq E$ then
 $\mathcal{R} \leftarrow \mathcal{R} \cup \mathcal{T}[h].\text{space}$
 Return $\mathcal{T}[h].\text{key}$
 Else
 Return \perp

procedure LR(u^*, m_0, m_1)

$CT^* \xleftarrow{R} \boxed{\text{EncryptSF}}(u^*, m_\beta)$
 Return CT^*

procedure Finalize(β')

Return ($\beta' = \beta$)

Game MasterLeakCK

procedure Initialize

$(PP, SK) \xleftarrow{R} \text{Setup}$
 $\boxed{\widetilde{SK} \leftarrow \text{KeyGenSF}(SK, E)}$
 $\mathcal{R} \leftarrow \emptyset$
 $\mathcal{T}[0] \leftarrow (E, SK, \boxed{\widetilde{SK}}, 0)$
 $\text{nbKeys} \leftarrow 1$
 $\beta \xleftarrow{R} \{0, 1\}$
 Return PP

procedure Create(h, W)

If $\mathcal{T}[h].\text{space} = E$ then
 $SK_W \xleftarrow{R} \text{KeyGen}(\mathcal{T}[h].\text{key}, W)$
 $\boxed{\widetilde{SK}_W \xleftarrow{R} \text{KeyGenSF}(\mathcal{T}[h].\text{key}, W)}$
 $\mathcal{T}[\text{nbKeys}] \leftarrow (W, SK_W, \boxed{\widetilde{SK}_W}, 0)$
 $\text{nbKeys} \leftarrow \text{nbKeys} + 1$
 Return \widetilde{SK}_W

Else
 Return \perp

procedure Delegate(h, W')

Return **Create**(0, W')

procedure Leak(h, f)

$\widetilde{SK} \leftarrow \boxed{\mathcal{T}[h].\text{keySF}}$
 If $\mathcal{T}[h].\text{leaked} + |f(\widetilde{SK})| \leq \ell_{SK}$ then
 $\mathcal{T}[h].\text{leaked} \leftarrow \mathcal{T}[h].\text{leaked} + |f(\widetilde{SK})|$
 Return $f(\widetilde{SK})$
 Else
 Return \perp

procedure Reveal(h)

If $\mathcal{T}[h].\text{space} \neq E$ then
 $\mathcal{R} \leftarrow \mathcal{R} \cup \mathcal{T}[h].\text{space}$
 Return $\boxed{\mathcal{T}[h].\text{keySF}}$
 Else
 Return \perp

procedure LR(u^*, m_0, m_1)

$CT^* \xleftarrow{R} \text{EncryptSF}(u^*, m_\beta)$
 Return CT^*

procedure Finalize(β')

Return ($\beta' = \beta$)

Figure 5: Algorithms of the games MasterLeakC and MasterLeakCK.

Game MasterLeakCK_j

procedure Initialize

$(PP, SK) \xleftarrow{R} \text{Setup}$
 $\widetilde{SK} \leftarrow \text{KeyGenSF}(SK, E)$
 $\mathcal{R} \leftarrow \emptyset$
 $\mathcal{T}[0] \leftarrow (E, SK, \widetilde{SK}, 0)$
 $\text{nbKeys} \leftarrow 1$
 $\beta \xleftarrow{R} \{0, 1\}$
 Return PP

procedure Create(h, W)

If $\mathcal{T}[h].\text{space} = E$ then
 $SK_W \xleftarrow{R} \text{KeyGen}(\mathcal{T}[h].\text{key}, W)$
 $\widetilde{SK}_W \xleftarrow{R} \text{KeyGenSF}(\mathcal{T}[h].\text{key}, W)$
 $\mathcal{T}[\text{nbKeys}] \leftarrow (W, SK_W, \widetilde{SK}_W, 0)$
 $\text{nbKeys} \leftarrow \text{nbKeys} + 1$
 Return SK_W
 Else
 Return \perp

procedure Delegate(h, W')

Return **Create**($0, W'$)

procedure Leak(h, f)

$SK \leftarrow \boxed{\text{nbKeys} > j} ? \mathcal{T}[h].\text{key} : \mathcal{T}[h].\text{keySF}$
 If $\mathcal{T}[h].\text{leaked} + |f(SK)| \leq \ell_{SK}$ then
 $\mathcal{T}[h].\text{leaked} \leftarrow \mathcal{T}[h].\text{leaked} + |f(SK)|$
 Return $f(SK)$
 Else
 Return \perp

procedure Reveal(h)

If $\mathcal{T}[h].\text{space} \neq E$ then
 $\mathcal{R} \leftarrow \mathcal{R} \cup \mathcal{T}[h].\text{space}$
 Return $\boxed{\text{nbKeys} > j} ? \mathcal{T}[h].\text{key} : \mathcal{T}[h].\text{keySF}$
 Else
 Return \perp

procedure LR(u^*, m_0, m_1)

$CT^* \xleftarrow{R} \text{EncryptSF}(PP, u^*, m_\beta)$
 Return CT^*

procedure Finalize(β')

Return ($\beta' = \beta$)

Figure 6: Algorithms of the game MasterLeakCK_j.

References

- [AGV09] Adi Akavia, Shafi Goldwasser, and Vinod Vaikuntanathan. Simultaneous hardcore bits and cryptography against memory attacks. In Omer Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 474–495. Springer, Berlin, Germany, March 15–17, 2009.
- [BFO08] Alexandra Boldyreva, Serge Fehr, and Adam O’Neill. On notions of security for deterministic encryption, and efficient constructions without random oracles. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 335–359, Santa Barbara, CA, USA, August 17–21, 2008. Springer, Berlin, Germany.
- [BGN05] Dan Boneh, Eu-Jin Goh, and Kobbi Nissim. Evaluating 2-DNF formulas on ciphertexts. In Joe Kilian, editor, *TCC 2005*, volume 3378 of *LNCS*, pages 325–341, Cambridge, MA, USA, February 10–12, 2005. Springer, Berlin, Germany.
- [BGW05] Dan Boneh, Craig Gentry, and Brent Waters. Collusion resistant broadcast encryption with short ciphertexts and private keys. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 258–275, Santa Barbara, CA, USA, August 14–18, 2005. Springer, Berlin, Germany.
- [BH08] Dan Boneh and Michael Hamburg. Generalized identity based and broadcast encryption schemes. In Josef Pieprzyk, editor, *ASIACRYPT 2008*, volume 5350 of *LNCS*, pages 455–470, Melbourne, Australia, December 7–11, 2008. Springer, Berlin, Germany.
- [BKKV10] Zvika Brakerski, Yael Tauman Kalai, Jonathan Katz, and Vinod Vaikuntanathan. Overcoming the hole in the bucket: Public-key cryptography resilient to continual memory leakage. In *51st FOCS*, pages 501–510. IEEE Computer Society Press, 2010.
- [Bon07] Dan Boneh. Bilinear groups of composite order (invited talk). In Tsuyoshi Takagi, Tatsuaki Okamoto, Eiji Okamoto, and Takeshi Okamoto, editors, *PAIRING 2007*, volume 4575 of *LNCS*, page 1, Tokyo, Japan, July 2–4, 2007. Springer, Berlin, Germany.
- [BSW11] Elette Boyle, Gil Segev, and Daniel Wichs. Fully leakage-resilient signatures. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 89–108, Tallinn, Estonia, May 15–19, 2011. Springer, Berlin, Germany.
- [BWY11] Mihir Bellare, Brent Waters, and Scott Yilek. Identity-based encryption secure against selective opening attack. In Yuval Ishai, editor, *TCC 2011*,

volume 6597 of *LNCS*, pages 235–252, Providence, RI, USA, March 28–30, 2011. Springer, Berlin, Germany.

- [DHLAW10] Yevgeniy Dodis, Kristiyan Haralambiev, Adriana López-Alt, and Daniel Wichs. Cryptography against continuous memory attacks. In *51st FOCS*, pages 511–520. IEEE Computer Society Press, 2010.
- [LLW11] Allison B. Lewko, Mark Lewko, and Brent Waters. How to leak on key updates. In *43rd ACM STOC*, pages 725–734. ACM Press, 2011.
- [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, French Riviera, May 30 – June 3, 2010. Springer, Berlin, Germany.
- [LRW11] Allison B. Lewko, Yannis Rouselakis, and Brent Waters. Achieving leakage resilience through dual system encryption. In Yuval Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 70–88, Providence, RI, USA, March 28–30, 2011. Springer, Berlin, Germany.
- [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, Zurich, Switzerland, February 9–11, 2010. Springer, Berlin, Germany.
- [MD11] Daisuke Moriyama and Hiroshi Doi. A fully secure spatial encryption scheme. *IEICE Transactions*, 94-A(1):28–35, 2011.
- [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, Santa Barbara, CA, USA, August 16–20, 2009. Springer, Berlin, Germany.

A Security Proof

A.1 Semi-functionality

A *semi-functional* key has \mathbb{G}_{p_2} parts, while a normal key has not. We pick at random g_2 , a generator of \mathbb{G}_{p_2} .

KeyGenSF(SK, W) first calls **Delegate**(E, SK, W) so as to get a normal key $\text{SK}_W = (\mathbf{k}_\rho, k_r, k_u, \mathbf{k}_{\text{del}})$. It then picks $\gamma \xleftarrow{R} \mathbb{Z}_N^{n+2}$, $\theta \xleftarrow{R} \mathbb{Z}_N^d$ and computes

$$\widetilde{\text{SK}}_W = \left((\mathbf{k}_\rho, k_r, k_u) \times g_2^\gamma, \mathbf{k}_{\text{del}} \times g_2^\theta \right)$$

EncryptSF(\mathbf{u}, m) first calls **Encrypt**(\mathbf{u}, m) to get $\text{CT} = (\mathbf{c}_\sigma, \mathbf{c}_u, c_s, c_m)$. It then picks $\delta \xleftarrow{R} \mathbb{Z}_N^{n+2}$ and computes the ciphertext

$$\widetilde{\text{CT}} = \left((\mathbf{c}_\sigma, \mathbf{c}_u, c_s) \times g_2^\delta, c_m \right).$$

γ, θ, δ are called the *semi-functional parameters* of the secret key, the delegation key and the ciphertext, respectively. If someone uses a semi-functional key $\widetilde{\text{SK}}$ for space $W = \text{Aff}(M, \mathbf{u}_1)$ with parameters (γ, θ) to construct a secret key for vector $\mathbf{u}_2 = \mathbf{u}_1 + M\mathbf{v}$ with **KeyGen**, then this will be semi-functional with parameters $\gamma' = \gamma + (0, \dots, 0, \langle \mathbf{v}, \theta \rangle)$.

A semi-functional secret key for vector \mathbf{u} is called *nominal* with respect to a ciphertext for vector \mathbf{u}' if it can correctly decrypt it, thus if and only if:

$$\langle \gamma, \delta \rangle = 0 \pmod{p_2} \text{ and } \mathbf{u} = \mathbf{u}'$$

because we get an extra term $\hat{e}(g_2, g_2)^{\langle \gamma, \delta \rangle}$ by the pairing. If the secret key cannot decrypt it, it is called *truly* semi-functional.

A.2 Lemma for Leakage Analysis

Our analysis of the leakage resilience of our system will rely on the following lemma from [BKKV10, LRW11], which is proven in [BFO08]. Below, we let $\text{dist}(X_1, X_2)$ denote the statistical distance of two random variables X_1 and X_2 .

Lemma 1. *Let $m, l, d \in \mathbb{N}$, $m \geq l \geq 2d$ and let p be a prime. Let $X \xleftarrow{R} \mathbb{Z}_p^{m \times \ell}$, let $Y \xleftarrow{R} \mathbb{Z}_p^{m \times d}$, and let $T \xleftarrow{R} \text{Rk}_d(\mathbb{Z}_p^{\ell \times d})$, where $\text{Rk}_d(\mathbb{Z}_p^{\ell \times d})$ denotes the set of $\ell \times d$ matrices of rank d with entries in \mathbb{Z}_p . Let $f : \mathbb{Z}_p^{m \times d} \rightarrow W$ be some function. Then,*

$$\text{dist}((X, f(X \cdot T)), (X, f(Y))) \leq \varepsilon$$

as long as

$$|W| \leq 4 \cdot \left(1 - \frac{1}{p}\right) \cdot p^{\ell - (2d-1)} \cdot \varepsilon^2.$$

More precisely, we will use the following corollary.

Corollary 1. *Let $m \in \mathbb{N}, m \geq 3$, and let p be a prime. Let $\Delta \xleftarrow{R} \mathbb{Z}_p^m$, $\nu \xleftarrow{R} \mathbb{Z}_p^m$ and let ν' be chosen uniformly randomly from the set of vectors in \mathbb{Z}_p^m which are orthogonal to Δ under the dot product modulo p . Let $f : \mathbb{Z}_p^m \rightarrow W$ be some function. Then,*

$$\text{dist}((\Delta, f(\nu)), (\Delta, f(\nu'))) \leq \varepsilon$$

as long as

$$|W| \leq 4 \cdot \left(1 - \frac{1}{p}\right) \cdot p^{m-2} \cdot \varepsilon^2.$$

Proof. We apply Lemma 1 with $d = 1$ and $\ell = m - 1$. Y then corresponds to ν , while X corresponds to a basis of the orthogonal space of Δ . We note that ν' is then distributed as $X \cdot T$, where $T \xleftarrow{R} \text{Rk}_1(\mathbb{Z}_p^{m-1 \times 1})$. We note that X is determined by Δ , and is distributed as $X \xleftarrow{R} \mathbb{Z}_p^{m \times m-1}$, since Δ is chosen uniformly randomly from \mathbb{Z}_p^m . It follows that

$$\text{dist}((\Delta, f(\nu)), (\Delta, f(\nu'))) = \text{dist}((X, f(X \cdot T)), (X, f(Y))) \leq \varepsilon.$$

□

This corollary allows us to set $\ell_{\text{SK}} = (n - 1 - 2c) \log p_2$ for our construction (we will have $n + 1 = m$), where c is any fixed positive constant (so that $\varepsilon = p_2^{-c}$ is negligible).

A.3 Proof

We will prove the following theorem.

Theorem 1. *Under assumptions 1, 2, 3 and for $\ell_{\text{SK}} = (n - 1 - 2c) \log p_2$, where $c > 0$ is any fixed positive constant, our spatial encryption scheme is ℓ_{SK} -master-leakage secure.*

Table 2: Assumptions that will be used for the proofs.

game	game	assumption
MasterLeakSpatial	MasterLeakSpatial*	trivial
MasterLeakSpatial*	MasterLeakC	2.1
MasterLeakCK _j	MasterLeakCK _{j+1}	2.2
MasterLeakCK		2.3

Theorem 2. *Any polynomial-time attacker \mathcal{A} has only a negligibly different probability of winning in MasterLeakSpatial versus MasterLeakSpatial*.*

Proof. It is easy to verify that the output of **Delegate**($W_1, \text{SK}_{W_1}, W_2$) is identically distributed to the output of **KeyGen**(SK_E, W_2). □

Theorem 3. *If assumption 2.1 holds, any polynomial-time attacker \mathcal{A} has only a negligibly different probability of winning in $\text{MasterLeakSpatial}^*$ versus MasterLeakC .*

Proof. We suppose there exists a PPT attacker \mathcal{A} which attains a non-negligible difference in probability of winning between those two games. We will build a PPT algorithm \mathcal{B} that breaks assumption 2.1 with non-negligible advantage.

\mathcal{B} receives $\mathcal{D}_1 = (N, \mathbb{G}, \mathbb{G}_T, \hat{e}, g_1, g_3)$ and a challenge term $T \in \mathbb{G}_{p_1 p_2^2}$. Then it plays the $\text{MasterLeakSpatial}^*$ or the MasterLeakC game with \mathcal{A} in the following way.

Setup. \mathcal{B} picks $(\tau, \sigma, \alpha, \varphi) \xleftarrow{R} \mathbb{Z}_N \times \mathbb{Z}_N^n \times \mathbb{Z}_N^n \times \mathbb{Z}_N$ then gives

$$\text{PP} = (N, g_1, g_3, g_1^\varphi, g_1^\alpha, \hat{e}(g_1, g_1)^\tau, g_1^\sigma)$$

to \mathcal{A} , where N , g_1 and g_3 are given by the challenger.

Phase 1. Knowing α , the simulator can generate a secret key for the whole space E and use it to execute all secret queries (**Create**, **Leak**, **KeyGen**).

Challenge. The adversary \mathcal{A} gives \mathcal{B} two messages m_0 and m_1 and a challenge vector \mathbf{u}^* . The simulator \mathcal{B} chooses $\beta \xleftarrow{R} \{0, 1\}$ and outputs the ciphertext

$$\text{CT} = (\mathbf{c}_\sigma, \mathbf{c}_{\mathbf{u}^*}, c_s, c_m) = (T^\sigma, T^{-(\varphi + \langle \mathbf{u}^*, \alpha \rangle)}, T, m_\beta \cdot \hat{e}(T, g_1^\tau)).$$

Phase 2. \mathcal{B} works the same way as in **Phase 1**.

If $T = g_1^a g_2^b$, then the ciphertext is semi-functional. This implicitly sets

$$s = a \quad \text{and} \quad \delta = b(\sigma, -\varphi - \langle \mathbf{u}^*, \alpha \rangle, 1).$$

s is properly distributed since $a \xleftarrow{R} \mathbb{Z}_N$ according to the assumption. δ is properly distributed in the attacker's view because the factors $\sigma, -\varphi - \langle \mathbf{u}^*, \alpha \rangle$ are only seen modulo p_1 in the public parameters and not modulo p_2 . Thus they are random modulo p_2 in \mathcal{A} 's view. Therefore, \mathcal{B} has properly simulated the MasterLeakC game.

If $T = g_1^a$, the ciphertext is normal since it has no \mathbb{G}_{p_2} parts. In this case, \mathcal{B} has properly simulated the $\text{MasterLeakSpatial}^*$ game. \square

Theorem 4. *If assumption 2.2 holds, any polynomial-time attacker \mathcal{A} has only a negligibly different probability of winning in MasterLeakCK_j versus $\text{MasterLeakCK}_{j+1}$.*

Proof. \mathcal{B} receives $\mathcal{D}_2 = (N, \mathbb{G}, \mathbb{G}_T, \hat{e}, g_1, g_3, g_1^x g_2^y, g_2^{y'}, g_3^{z'})$ and a challenge term $T \in \mathbb{G}_{p_1 p_2^2 p_3}$. Then it plays the MasterLeakCK_j or the $\text{MasterLeakCK}_{j+1}$ game with \mathcal{A} as follows.

Setup. \mathcal{B} picks $(\tau, \boldsymbol{\sigma}, \boldsymbol{\alpha}, \varphi) \xleftarrow{R} \mathbb{Z}_N \times \mathbb{Z}_N^n \times \mathbb{Z}_N^n \times \mathbb{Z}_N$ then gives

$$\text{PP} = (N, g_1, g_3, g_1^\varphi, g_1^\alpha, \hat{e}(g_1, g_1)^\tau, g_1^\sigma)$$

to \mathcal{A} , where N , g_1 and g_3 are given by the challenger.

Phase 1.

- For the first $j - 1$ keys, \mathcal{B} picks $(\boldsymbol{\rho}', r, \boldsymbol{\mu}) \xleftarrow{R} \mathbb{Z}_N^n \times \mathbb{Z}_N \times \mathbb{Z}_N^{d+n+2}$ and, for the semi-functional parameters, $(\boldsymbol{\gamma}', \boldsymbol{\theta}') \xleftarrow{R} \mathbb{Z}_N^{n+2} \times \mathbb{Z}_N^d$. It then computes:

$$\text{SK} = \left(\left(g_1^{\boldsymbol{\rho}'}, g_1^r, g_1^{\tau+r(\varphi+\langle \mathbf{u}, \boldsymbol{\alpha} \rangle) - \langle \boldsymbol{\rho}', \boldsymbol{\sigma} \rangle} \right) \times (g_2^{y'} g_3^{z'})^{\boldsymbol{\gamma}'}, g_1^{rM^\top \boldsymbol{\alpha}} \times (g_2^{y'} g_3^{z'})^{\boldsymbol{\theta}'} \right) \times g_3^\mu$$

The \mathbb{G}_{p_1} and \mathbb{G}_{p_3} parthosos are properly distributed and the semi-functional parameters are $\boldsymbol{\gamma} = y' \boldsymbol{\gamma}'$ and $\boldsymbol{\theta} = y' \boldsymbol{\theta}'$. These are properly distributed as well.

- For the j key, \mathcal{B} will use the challenge term to generate the secret key. It picks $(\boldsymbol{\rho}', \boldsymbol{\mu}) \xleftarrow{R} \mathbb{Z}_N^n \times \mathbb{Z}_N^{d+n+2}$ and outputs

$$\text{SK} = \left(T^{\boldsymbol{\rho}'}, T, g_1^\tau T^{\varphi+\langle \mathbf{u}, \boldsymbol{\alpha} \rangle - \langle \boldsymbol{\rho}', \boldsymbol{\sigma} \rangle}, T^{M^\top \boldsymbol{\alpha}} \right) \times g_3^\mu$$

If $T = g_1^a g_2^b g_3^c$, then this key is semi-functional with

$$r = a \quad \boldsymbol{\rho} = a \boldsymbol{\rho}' \quad \boldsymbol{\gamma} = b(\boldsymbol{\rho}', 1, \varphi + \langle \mathbf{u}, \boldsymbol{\alpha} \rangle - \langle \boldsymbol{\rho}', \boldsymbol{\sigma} \rangle) \quad \boldsymbol{\theta} = b M^\top \boldsymbol{\alpha}.$$

Again, since the public parameters only determine φ and $\boldsymbol{\alpha}$ modulo p_1 , the semi-functional parameters are random modulo p_2 in \mathcal{A} 's view.

- From the $j + 1$ key, \mathcal{B} picks $(\boldsymbol{\rho}', r, \boldsymbol{\mu}) \xleftarrow{R} \mathbb{Z}_N^n \times \mathbb{Z}_N \times \mathbb{Z}_N^{d+n+2}$ and computes:

$$\text{SK} = \left(g_1^{\boldsymbol{\rho}'}, g_1^r, g_1^{\tau+r(\varphi+\langle \mathbf{u}, \boldsymbol{\alpha} \rangle) - \langle \boldsymbol{\rho}', \boldsymbol{\sigma} \rangle}, g_1^{rM^\top \boldsymbol{\alpha}} \right) \times g_3^\mu$$

Challenge. In this phase, \mathcal{B} has to create a semi-functional ciphertext with **EncryptSF**. \mathcal{A} gives \mathcal{B} two messages m_0 and m_1 and a challenge vector \mathbf{u}^* . \mathcal{B} chooses $\beta \xleftarrow{R} \{0, 1\}$ and outputs the ciphertext

$$\widetilde{\text{CT}} = (\mathbf{c}_\sigma, \mathbf{c}_{\mathbf{u}^*}, c_s, c_m) = \left((g_1^x g_2^y)^\sigma, (g_1^x g_2^y)^{-(\varphi+\langle \mathbf{u}^*, \boldsymbol{\alpha} \rangle)}, (g_1^x g_2^y), m_\beta \cdot \hat{e}((g_1^x g_2^y), g_1^\tau) \right).$$

Phase 2. \mathcal{B} works the same way as in **Phase 1**.

The ciphertext parameters are:

$$s = x \quad \text{and} \quad \delta = y(\sigma, -\varphi - \langle \mathbf{u}^*, \alpha \rangle, 1).$$

s is properly distributed, but the semi-functional parameters are not: if the space of the j key contains the challenge vector \mathbf{u}^* , the secret key is nominal with respect to the ciphertext. Indeed:

$$\begin{aligned} & (\gamma + (0, \dots, 0, \langle \mathbf{v}, \theta \rangle)) \cdot \delta \\ = & b \left(\rho', 1, \varphi + \langle \mathbf{u}, \alpha \rangle - \langle \rho', \sigma \rangle + \langle \mathbf{v}, M^\top \alpha \rangle \right) \cdot y(\sigma, -\varphi - \langle \mathbf{u}^*, \alpha \rangle, 1) \\ = & by \left(\langle \rho', \sigma \rangle - \varphi - \langle \mathbf{u}^*, \alpha \rangle + \varphi + \langle \mathbf{u}, \alpha \rangle - \langle \rho', \sigma \rangle + \langle M\mathbf{v}, \alpha \rangle \right) \\ = & by \langle \mathbf{u} + M\mathbf{v} - \mathbf{u}^*, \alpha \rangle \\ = & 0 \pmod{p_2} \end{aligned}$$

Notice that if the space of the j key contains a vector \mathbf{w} which is equal to \mathbf{u}^* modulo p_2 , we obtain the same result. With the help of two lemmas, we will show that the change in the adversary's advantage is negligible.

Lemma 2. *If assumption 2 holds, then for any PPT adversary \mathcal{A} , \mathcal{A} 's advantage in the MasterLeakC_j game, or in the MasterLeakC_{j+1} game, changes only by a negligible amount if we restrict it to make queries only on the challenge vector and on spaces such that no component of any vector of them is equal to a respective component from the challenge vector \mathbf{u}^* modulo p_2 .*

Proof. If there exists an adversary whose advantage changes by a non-negligible amount under this restriction, we can find a non-trivial factor of N with non-negligible probability. So we can break Assumption 2, using the same proof as [LW10]).

The simulator plays MasterLeakC_j or MasterLeakC_{j+1} , using the terms from Assumption 2 to create the semi-functional keys and ciphertext. It works in a way similar to the simulator in the reduction shown in Theorem 3. \square

Lemma 3. *We suppose the leakage is at most $\ell_{\text{SK}} = (n - 1 - 2c) \log p_2$, where $c > 0$ is a positive constant. Then, for any PPT adversary \mathcal{A} , \mathcal{A} 's advantage in the MasterLeakC_{j+1} game changes only by a negligible amount when the truly semi-functional j key is replaced by a nominal semi-functional key whenever \mathcal{A} declares the j key to be associated to a space that contains the challenge ciphertext vector.*

Proof. Let us suppose there exists a PPT algorithm \mathcal{A} whose advantage changes by a non-negligible amount when the MasterLeakC_{j+1} game changes as described above. Using \mathcal{A} , we will create a PPT algorithm \mathcal{B} which will distinguish between $(\Delta, f(\nu))$ and $(\Delta, f(\nu'))$

from Corollary 1 with non-negligible advantage (when $m = n + 1$ and $p = p_2$). This will yield a contradiction, since these distributions have a negligible statistical distance.

\mathcal{B} simulates MasterLeakC_{j+1} as follows. It runs **Setup** and gives \mathcal{A} the public parameters. Since \mathcal{B} knows α and generators of all the subgroups, it can respond to \mathcal{A} 's queries in **Phase 1**.

With non-negligible probability, the j key \mathcal{A} chooses in **Phase 1** must be associated to a space that contains \mathbf{u}^* . (If it only did this with negligible probability, then the difference in advantages whenever it happens would be negligible.)

\mathcal{B} will not create that key but instead will encode the leakage \mathcal{A} asks for, as a single polynomial-time computable function f with domain $\mathbb{Z}_{p_2}^{n+1}$ and with an image of size $2^{\ell_{\text{SK}}}$. \mathcal{B} receives a sample $(\Delta, f(\Gamma))$, where Γ is either distributed as ν or as ν' . \mathcal{B} will use $f(\Gamma)$ to answer all of \mathcal{A} 's leakage queries on the j key by implicitly defining this key as follows.

\mathcal{B} chooses $(r_1, r_2, \theta) \in \mathbb{Z}_{p_2}^{d+2}$. We let g_2 denote a generator of \mathbb{G}_{p_2} . \mathcal{B} implicitly sets the \mathbb{G}_{p_2} components of the key to be $g_2^{\Gamma'}$, where Γ' is defined to be:

$$\Gamma' = (\Gamma_1, \dots, \Gamma_n, \Gamma_{n+1} + r_1, r_2, \theta).$$

\mathcal{B} defines the other components of the key to fit their appropriate distribution.

At some point, \mathcal{A} declares the challenge vector \mathbf{u}^* . If the space of the j key does not contain \mathbf{u}^* , then \mathcal{B} aborts the simulation and guesses whether Γ is orthogonal to Δ randomly. However, the simulation continues with non-negligible probability. Suppose the space of the j key is $W = \text{Aff}(M, \mathbf{u})$ and let \mathbf{v} be the vector such that $\mathbf{u}^* = \mathbf{u} + M\mathbf{v}$.

\mathcal{B} chooses a random element $t_2 \in \mathbb{Z}_{p_2}$ such that $\Delta_{n+1}r_1 + (r_2 + \langle \mathbf{v}, \theta \rangle)t_2 = 0 \pmod{p_2}$. It then constructs the challenge ciphertext, using $(\Delta, t_2) \in \mathbb{Z}_{p_2}^{n+2}$ as parameter.

If Γ is orthogonal to Δ , then the j key is nominally semi-functional (and well distributed):

$$\begin{aligned} ((\Gamma_1, \dots, \Gamma_n, \Gamma_{n+1} + r_1, r_2) + (0, \dots, 0, \langle \mathbf{v}, \theta \rangle)) \cdot \Delta &= \Gamma \cdot \Delta + \Delta_{n+1}r_1 + (r_2 + \langle \mathbf{v}, \theta \rangle)t_2 \\ &= 0 \pmod{p_2}. \end{aligned}$$

If Γ is not orthogonal to Δ , then the challenge key is truly semi-functional (and also well distributed).

It is clear that \mathcal{B} can easily handle **Phase 2** queries, since the j key cannot be queried when its space contains \mathbf{u}^* . Therefore, \mathcal{B} can use the output of \mathcal{A} to gain a non-negligible advantage in distinguishing $(\Delta, f(\nu))$ and $(\Delta, f(\nu'))$, which violates Corollary 1. \square

The above lemmas conclude the theorem. \square

Theorem 5. *If assumption 3 holds, any polynomial-time attacker \mathcal{A} has only a negligible advantage in MasterLeakCK.*

Proof. We suppose there exists a PPT attacker \mathcal{A} which attains a non-negligible advantage in MasterLeakCK. We will build a PPT algorithm \mathcal{B} that breaks assumption 3 with non-negligible advantage.

\mathcal{B} receives $\mathcal{D}_3 = (N, \mathbb{G}, \mathbb{G}_T, \hat{e}, g_1, g_2, g_3, g_1^\tau g_2^y, g_1^{x'} g_2^{y'})$ and a challenge term T which is either $\hat{e}(g_1, g_1)^{\tau x'}$ or a random term of \mathbb{G}_T . Algorithm \mathcal{B} works as follows.

Setup. \mathcal{B} picks $(\sigma, \alpha, \varphi) \xleftarrow{R} \mathbb{Z}_N^n \times \mathbb{Z}_N^n \times \mathbb{Z}_N$. Notice that now τ is unknown. The term $\hat{e}(g_1, g_1)^\tau$ is computed as $\hat{e}(g_1^\tau g_2^y, g_1)$. It gives PP = $(N, g_1, g_3, g_1^\varphi, g_1^\alpha, \hat{e}(g_1, g_1)^\tau, g_1^\sigma)$ to \mathcal{A} , where N, g_1 and g_3 are given by the challenger.

Phase 1. For each secret key requested by the adversary, the simulator \mathcal{B} picks the random exponents $(r, \rho, \mu) \xleftarrow{R} \mathbb{Z}_N \times \mathbb{Z}_N^n \times \mathbb{Z}_N^{d+n+2}$ and for the semi-functional parameters, $\gamma' \xleftarrow{R} \mathbb{Z}_N^{n+2}$ and $\theta \xleftarrow{R} \mathbb{Z}_N^d$. It uses the secret key

$$\text{SK} = \left(\left(g_1^\rho, g_1^\tau, (g_1^\tau g_2^y) g_1^{r(\varphi + \langle \mathbf{u}, \alpha \rangle) - \langle \rho', \sigma \rangle} \right) \times g_2^{\gamma'}, g_1^{rM^\top \alpha} \times g_2^\theta \right) \times g_3^\mu$$

It is a properly distributed semi-functional key with parameters

$$\gamma = \gamma' + (0, \dots, 0, y) \quad \text{and} \quad \theta.$$

Challenge. The adversary \mathcal{A} gives \mathcal{B} two messages m_0, m_1 and a challenge vector \mathbf{u}^* . The simulator chooses $\beta \xleftarrow{R} \{0, 1\}$ and outputs the following ciphertext:

$$\text{CT} = \left((g_1^{x'} g_2^{y'})^\sigma, (g_1^{x'} g_2^{y'})^{-\varphi - \langle \mathbf{u}^*, \alpha \rangle}, (g_1^{x'} g_2^{y'}), m_\beta \cdot T \right).$$

Phase 2. \mathcal{B} works the same way as in **Phase 1**.

If $T = \hat{e}(g_1, g_1)^{\tau x'}$, then we get a semi-functional ciphertext of m_β with parameters

$$s = x' \quad \text{and} \quad \delta = y'(\sigma, -\varphi - \langle \mathbf{u}^*, \alpha \rangle, 1).$$

δ is properly distributed since all terms are random modulo p_2 . Therefore, \mathcal{B} has properly simulated game MasterLeakCK.

If $T \xleftarrow{R} \mathbb{G}_T$, then c_m is entirely random and we get a semi-functional ciphertext of a random message. Hence, the value of β is information-theoretically hidden and the probability of success of \mathcal{A} is exactly 1/2, since $\beta \xleftarrow{R} \{0, 1\}$. Therefore, \mathcal{B} can use the output of \mathcal{A} to break Assumption 3 with non-negligible advantage. \square

This concludes the proof of Theorem 1.