

Watermarking and Traitor Tracing for PRFs

David Wu


April 2020

based on joint work with Rishab Goyal, Sam Kim, and Brent Waters

Software Watermarking

[NSS99, BGIRSVY01, HMW07, CHNVW16]

```
static void AES_enc_blk(block *blk, const AES_KEY *key) {  
    unsigned j, rnds = ROUNDS(key);  
    const __m128i *sched = ((__m128i *) (key->rd_key));  
    *blk = _mm_xor_si128(*blk, sched[0]);  
    for (j = 1; j < rnds; ++j) {  
        *blk = _mm_aesenc_si128(*blk, sched[j]);  
    }  
    *blk = _mm_aesenc_si128(*blk, sched[j]);  
}
```



Embed a “mark” within a program



```
static void AES_enc_blk(block *blk, const AES_KEY *key) {  
    unsigned j, rnds = ROUNDS(key);  
    const __m128i *sched = ((__m128i *) (key->rd_key));  
    *blk = _mm_xor_si128(*blk, sched[0]);  
    for (j = 1; j < rnds; ++j) {  
        *blk = _mm_aesenc_si128(*blk, sched[j]);  
    }  
    *blk = _mm_aesenc_si128(*blk, sched[j]);  
}
```


If mark is removed, then program is destroyed

Applications: proving software ownership,
preventing unauthorized distribution of software

Software Watermarking

[NSS99, BGIRSVY01, HMW07, CHNVW16]

```
static void AES_enc_blk(block *blk, const AES_KEY *key) {  
    unsigned j, rnds = ROUNDS(key);  
    const __m128i *sched = ((__m128i *) (key->rd_key));  
    *blk = _mm_xor_si128(*blk, sched[0]);  
    for (j = 1; j < rnds; ++j) {  
        *blk = _mm_aesenc_si128(*blk, sched[j]);  
    }  
    *blk = _mm_aesencast_si128(*blk, sched[j]);  
}
```



Embed a “mark” within a program



```
static void AES_enc_blk(block *blk, const AES_KEY *key) {  
    unsigned j, rnds = ROUNDS(key);  
    const __m128i *sched = ((__m128i *) (key->rd_key));  
    *blk = _mm_xor_si128(*blk, sched[0]);  
    for (j = 1; j < rnds; ++j) {  
        *blk = _mm_aesenc_si128(*blk, sched[j]);  
    }  
    *blk = _mm_aesencast_si128(*blk, sched[j]);  
}
```

If mark is removed, then program is destroyed

Two main algorithms:

- $\text{Mark}(C, m) \rightarrow C'$: Takes circuit C and mark m and outputs a marked circuit C'
- $\text{Extract}(C') \rightarrow m/\perp$: Extracts the mark from a circuit C'

Software Watermarking

[NSS99, BGIRSVY01, HMW07, CHNVW16]

```
static void AES_enc_blk(block *blk, const AES_KEY *key) {  
    unsigned j, rnds = ROUNDS(key);  
    const __m128i *sched = ((__m128i *) (key->rd_key));  
    *blk = _mm_xor_si128(*blk, sched[0]);  
    for (j = 1; j < rnds; ++j) {  
        *blk = _mm_aesenc_si128(*blk, sched[j]);  
    }  
    *blk = _mm_aesenc_si128(*blk, sched[j]);  
}
```

Mark



```
static void AES_enc_blk(block *blk, const AES_KEY *key) {  
    unsigned j, rnds = ROUNDS(key);  
    const __m128i *sched = ((__m128i *) (key->rd_key));  
    *blk = _mm_xor_si128(*blk, sched[0]);  
    for (j = 1; j < rnds; ++j) {  
        *blk = _mm_aesenc_si128(*blk, sched[j]);  
    }  
    *blk = _mm_aesenc_si128(*blk, sched[j]);  
}
```



Functionality-preserving: On input a circuit C (and mark m), the Mark algorithm outputs a circuit C' where


$$C(x) = C'(x)$$

on almost all inputs x

Software Watermarking

[NSS99, BGIRSVY01, HMW07, CHNVW16]

```
static void AES_enc_blk(block *blk, const AES_KEY *key) {  
    unsigned j, rnds = ROUNDS(key);  
    const __m128i *sched = ((__m128i *) (key->rd_key));  
    *blk = _mm_xor_si128(*blk, sched[0]);  
    for (j = 1; j < rnds; ++j) {  
        *blk = _mm_aesenc_si128(*blk, sched[j]);  
    }  
    *blk = _mm_aesencast_si128(*blk, sched[j]);  
}
```



```
static void AES_enc_blk(block *blk, const AES_KEY *key) {  
    unsigned j, rnds = ROUNDS(key);  
    const __m128i *sched = ((__m128i *) (key->rd_key));  
    *blk = _mm_xor_si128(*blk, sched[0]);  
    for (j = 1; j < rnds; ++j) {  
        *blk = _mm_aesenc_si128(*blk, sched[j]);  
    }  
    *blk = _mm_aesencast_si128(*blk, sched[j]);  
}
```

- Unremovability:** Given a program C' with mark m , no efficient adversary can construct a circuit C^* where
- $C^*(x) = C'(x)$ on almost all inputs x
 - The circuit C^* does not preserve the mark: $\text{Extract}(C^*) \neq m$

Software Watermarking

[NSS99, BGIRSVY01, HMW07, CHNVW16]

```
static void AES_enc_blk(block *blk, const AES_KEY *key) {  
    unsigned j, rnds = ROUNDS(key);  
    const __m128i *sched = ((__m128i *) (key->rd_key));  
    *blk = _mm_xor_si128(*blk, sched[0]);  
    for (j = 1; j < rnds; ++j) {  
        *blk = _mm_aesenc_si128(*blk, sched[j]);  
    }  
    *blk =  
}
```



```
static void AES_enc_blk(block *blk, const AES_KEY *key) {  
    unsigned j, rnds = ROUNDS(key);  
    const __m128i *sched = ((__m128i *) (key->rd_key));  
    *blk = _mm_xor_si128(*blk, sched[0]);  
    for (j = 1; j < rnds; ++j) {  
        *blk = _mm_aesenc_si128(*blk, sched[j]);  
    }  
    *blk =  
}
```

Adversary is very powerful: sees the code of the marked program C' and has complete flexibility in crafting C^*


Unremovability: Given a program C' with mark m , no efficient adversary can construct a circuit C^* where

- $C^*(x) = C'(x)$ on almost all inputs x
- The circuit C^* does not preserve the mark: $\text{Extract}(C^*) \neq m$

Software Watermarking

[NSS99, BGIRSVY01, HMW07, CHNVW16]

```
static void AES_enc_blk(block *blk, const AES_KEY *key) {  
    unsigned j, rnds = ROUNDS(key);  
    const __m128i *sched = ((__m128i *) (key->rd_key));  
    *blk = _mm_xor_si128(*blk, sched[0]);  
    for (j = 1; j < rnds; ++j) {  
        *blk = _mm_aesenc_si128(*blk, sched[j]);  
    }  
    *blk = _mm_aesenc_si128(*blk, sched[j]);  
}
```




```
static void AES_enc_blk(block *blk, const AES_KEY *key) {  
    unsigned j, rnds = ROUNDS(key);  
    const __m128i *sched = ((__m128i *) (key->rd_key));  
    *blk = _mm_xor_si128(*blk, sched[0]);  
    for (j = 1; j < rnds; ++j) {  
        *blk = _mm_aesenc_si128(*blk, sched[j]);  
    }  
    *blk = _mm_aesenc_si128(*blk, sched[j]);  
}
```

Learning the original
(unmarked) function gives a
way to remove the watermark

- Notion only achievable for functions that are not learnable
- Focus has been on cryptographic functions


Watermarking Cryptographic Programs



Previous works: watermarking PRFs [CHNVW16, BLW17, KW17, QWZ18, KW19]

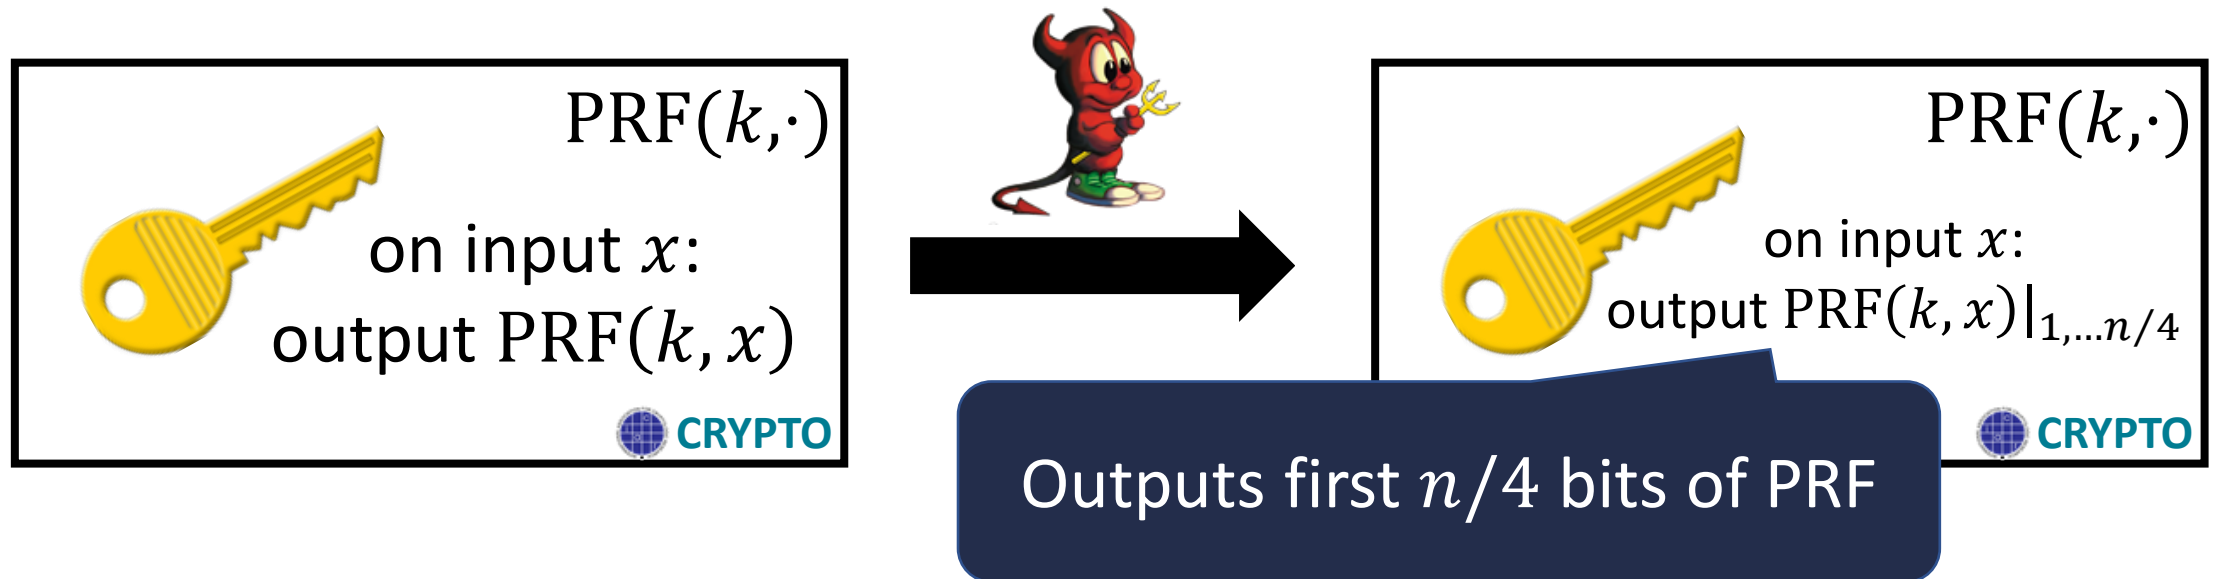
Suffices for watermarking other symmetric primitives:
(e.g., MAC signing key, symmetric decryption key)

A Closer Look at Watermarking Security



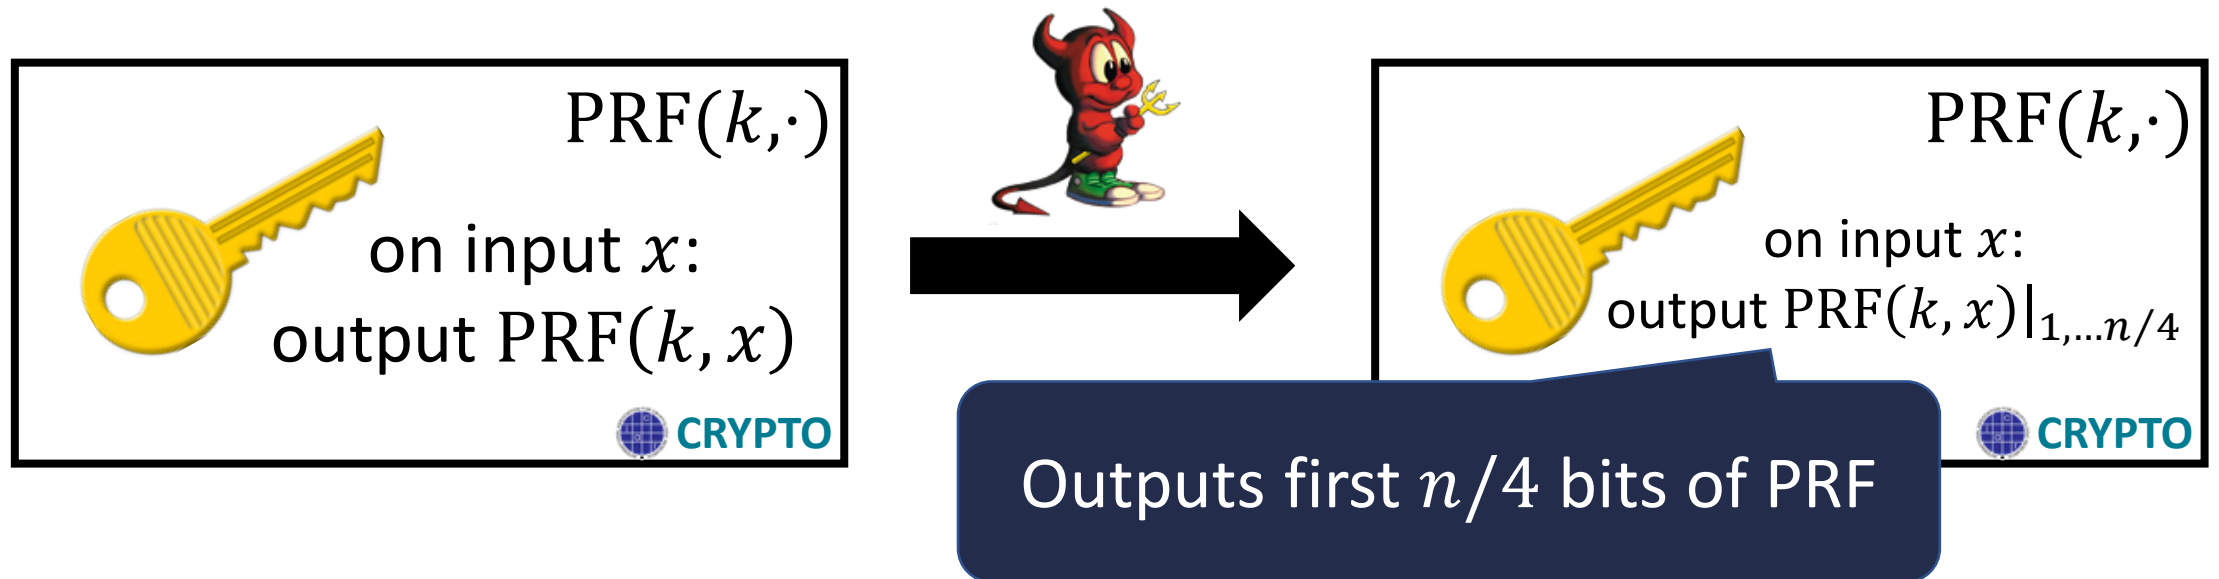
- Unremovability:** Given a program C' with mark m , no efficient adversary can construct a circuit C^* where
- $C^*(x) = C'(x)$ on almost all inputs x
 - The circuit C^* does not preserve the mark: $\text{Extract}(C^*) \neq m$

A Closer Look at Watermarking Security



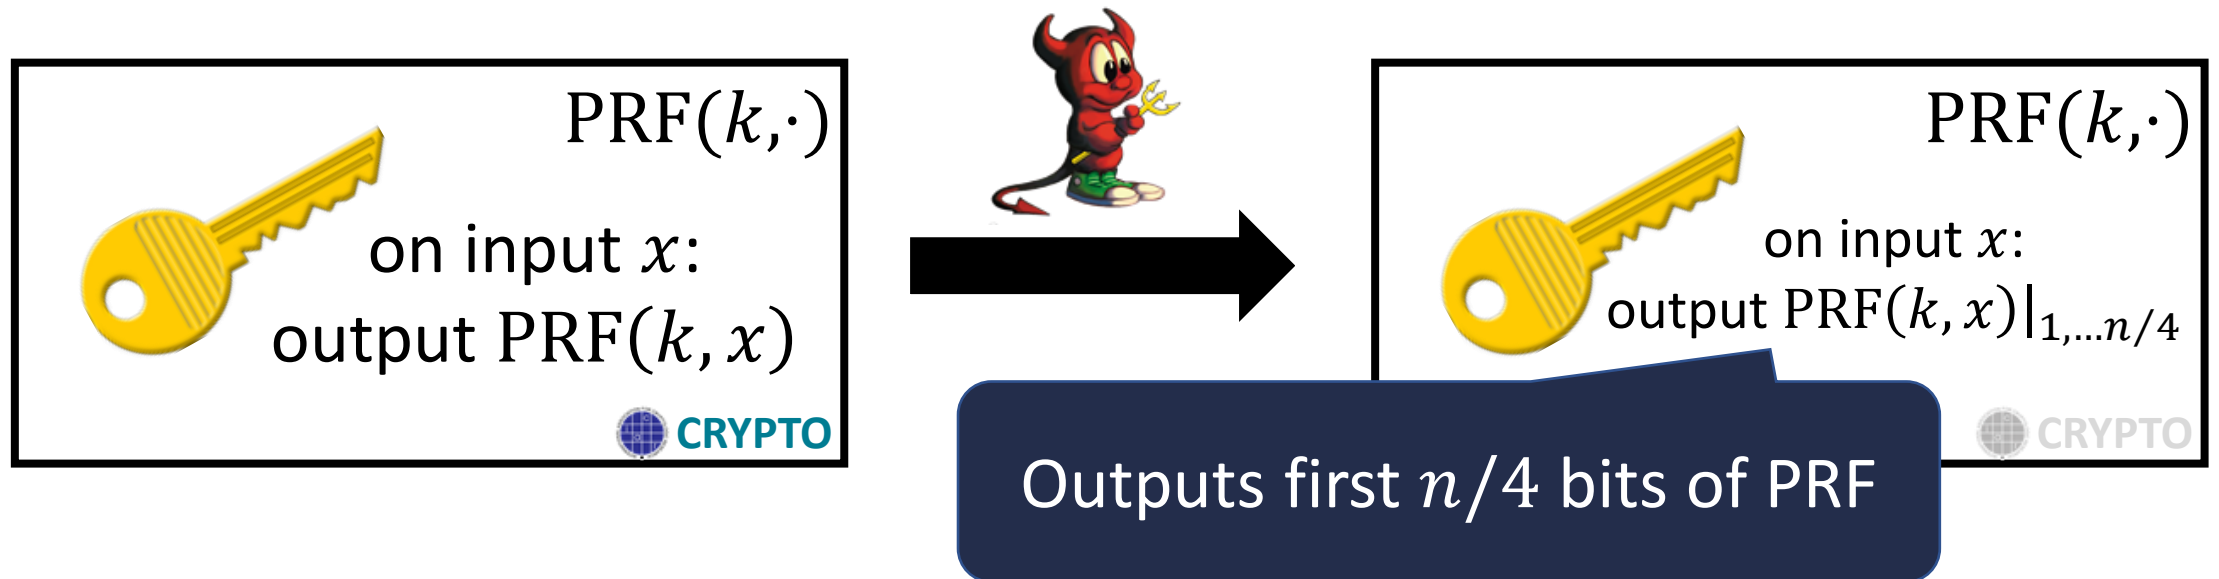
- Unremovability:** Given a program C' with mark m , no efficient adversary can construct a circuit C^* where
- $C^*(x) = C'(x)$ on almost all inputs x
 - The circuit C^* does not preserve the mark: $\text{Extract}(C^*) \neq m$

A Closer Look at Watermarking Security



- Unremovability:** Given a program C' with mark m , no efficient adversary can construct a circuit C^* where
- $C^*(x) = C'(x)$ on almost all inputs x
 - The circuit C^* does not preserve the mark: $\text{Extract}(C^*) \neq m$
- Adversary's circuit does not preserve functionality

A Closer Look at Watermarking Security




Unremovability: Given a program C' with mark m , no efficient adversary can construct a circuit C^* where


- $C^*(x) = C'(x)$ on almost all inputs x Adversary's circuit does not preserve functionality
- The circuit C^* does not preserve the mark: $\text{Extract}(C^*) \neq m$

No guarantees on whether the mark is preserved or not!

A Closer Look at Watermarking Security



$\text{PRF}(k, \cdot)$
on input x :
output $\text{PRF}(k, x)|_{1, \dots, n/4}$




Suppose circuit that only outputs leading $n/4$ bits does not contain the watermark

Is this a problem?

For building blocks like PRFs, we do not necessarily need to recover exact output to “break” functionality

Suppose watermarkable PRF used to protect against unauthorized distribution of decryption keys




Encrypted image
(PRF in counter mode)



Partial decryption
(using program on left)


Adversary’s program is “good enough” in most settings, but may not preserve watermark

A Closer Look at Watermarking Security




$\text{PRF}(k, \cdot)$


on input x :
output $\text{PRF}(k, x)|_{1, \dots, n/4}$



Suppose watermarkable PRF used to protect against unauthorized distribution of decryption keys



Encrypted image
(PRF in counter mode)




Partial decryption
(using program on left)

Adversary's program is "good enough" in most settings, but may not preserve watermark

Watermarking cryptographic programs:


- Exact functionality preserving does not seem like the right security notion
- If adversary's program can break the primitive, then watermark should be preserved

A Closer Look at Watermarking Security




$\text{PRF}(k, \cdot)$

on input x :
output $\text{PRF}(k, x)|_{1, \dots, n/4}$



Suppose watermarkable PRF used to protect against unauthorized distribution of decryption keys




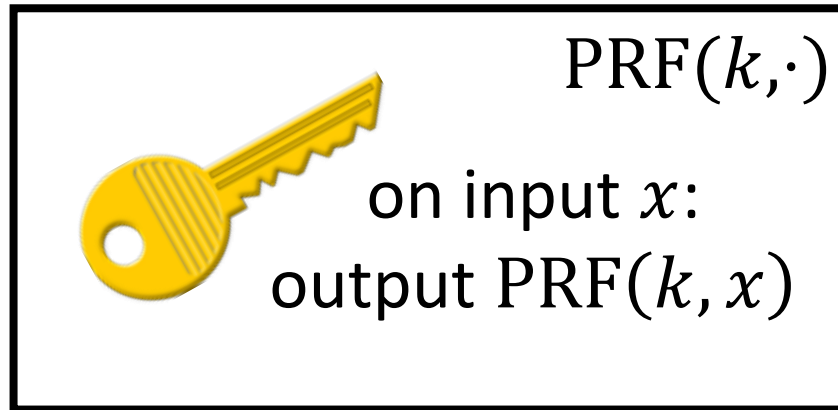
Encrypted image
(PRF in counter mode)

Partial decryption
(using program on left)


Existing watermarking constructions are unable to recover the watermark from this type of program

Adversary's program is "good enough" in most settings, but may not preserve watermark

Traceable PRFs



Mark




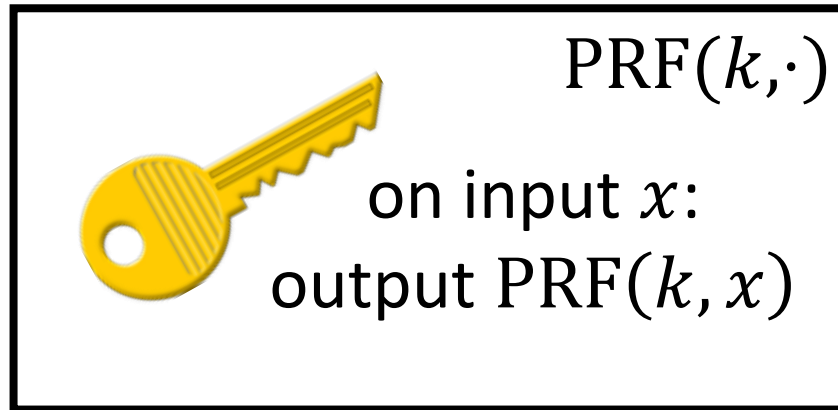
PRF security:

PRF(k, \cdot) indistinguishable
from random function


Marking security (informal):

if program C can distinguish
PRF(k, \cdot) from random, then mark
should be preserved

Traceable PRFs



Mark



Traitor tracing: if program can distinguish ciphertexts, then mark is preserved

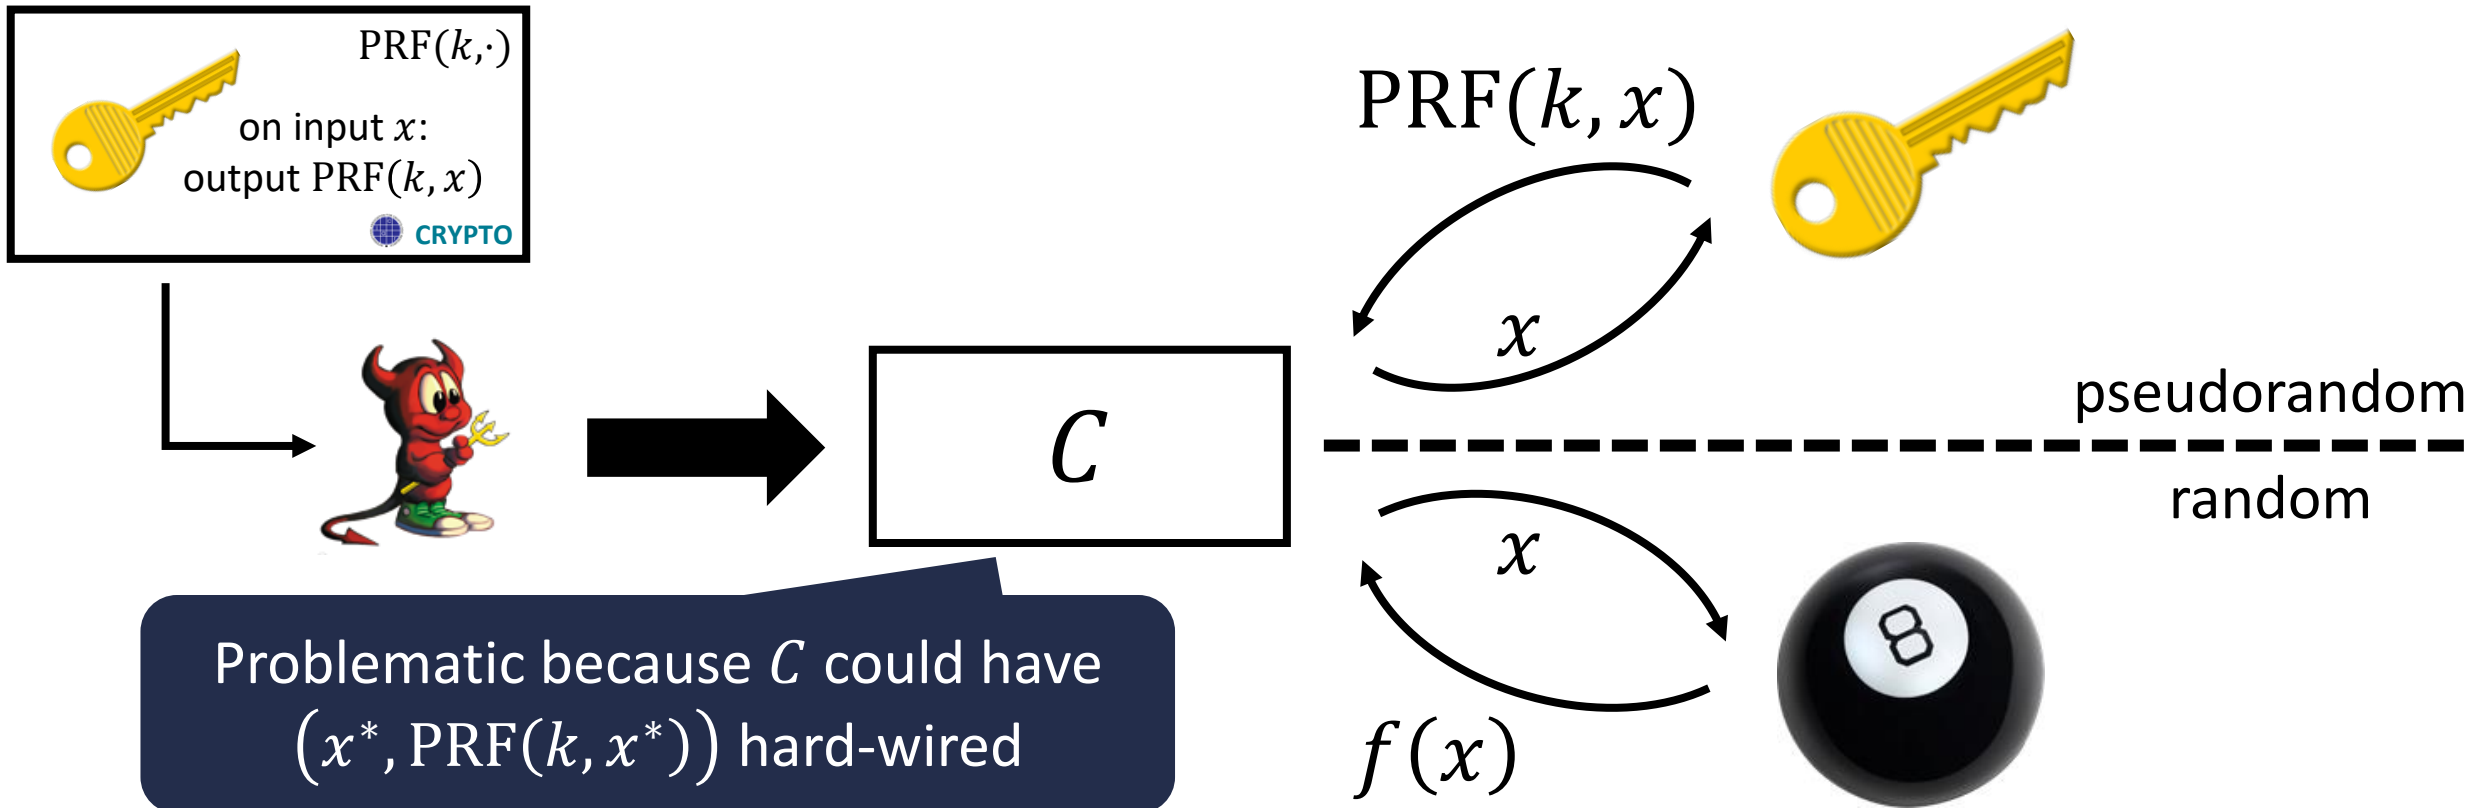
Traceable PRF: analog for PRFs

Marking security (informal):
if program C can distinguish PRF(k, \cdot) from random, then mark should be preserved

Traceable PRFs

Marking security (informal):

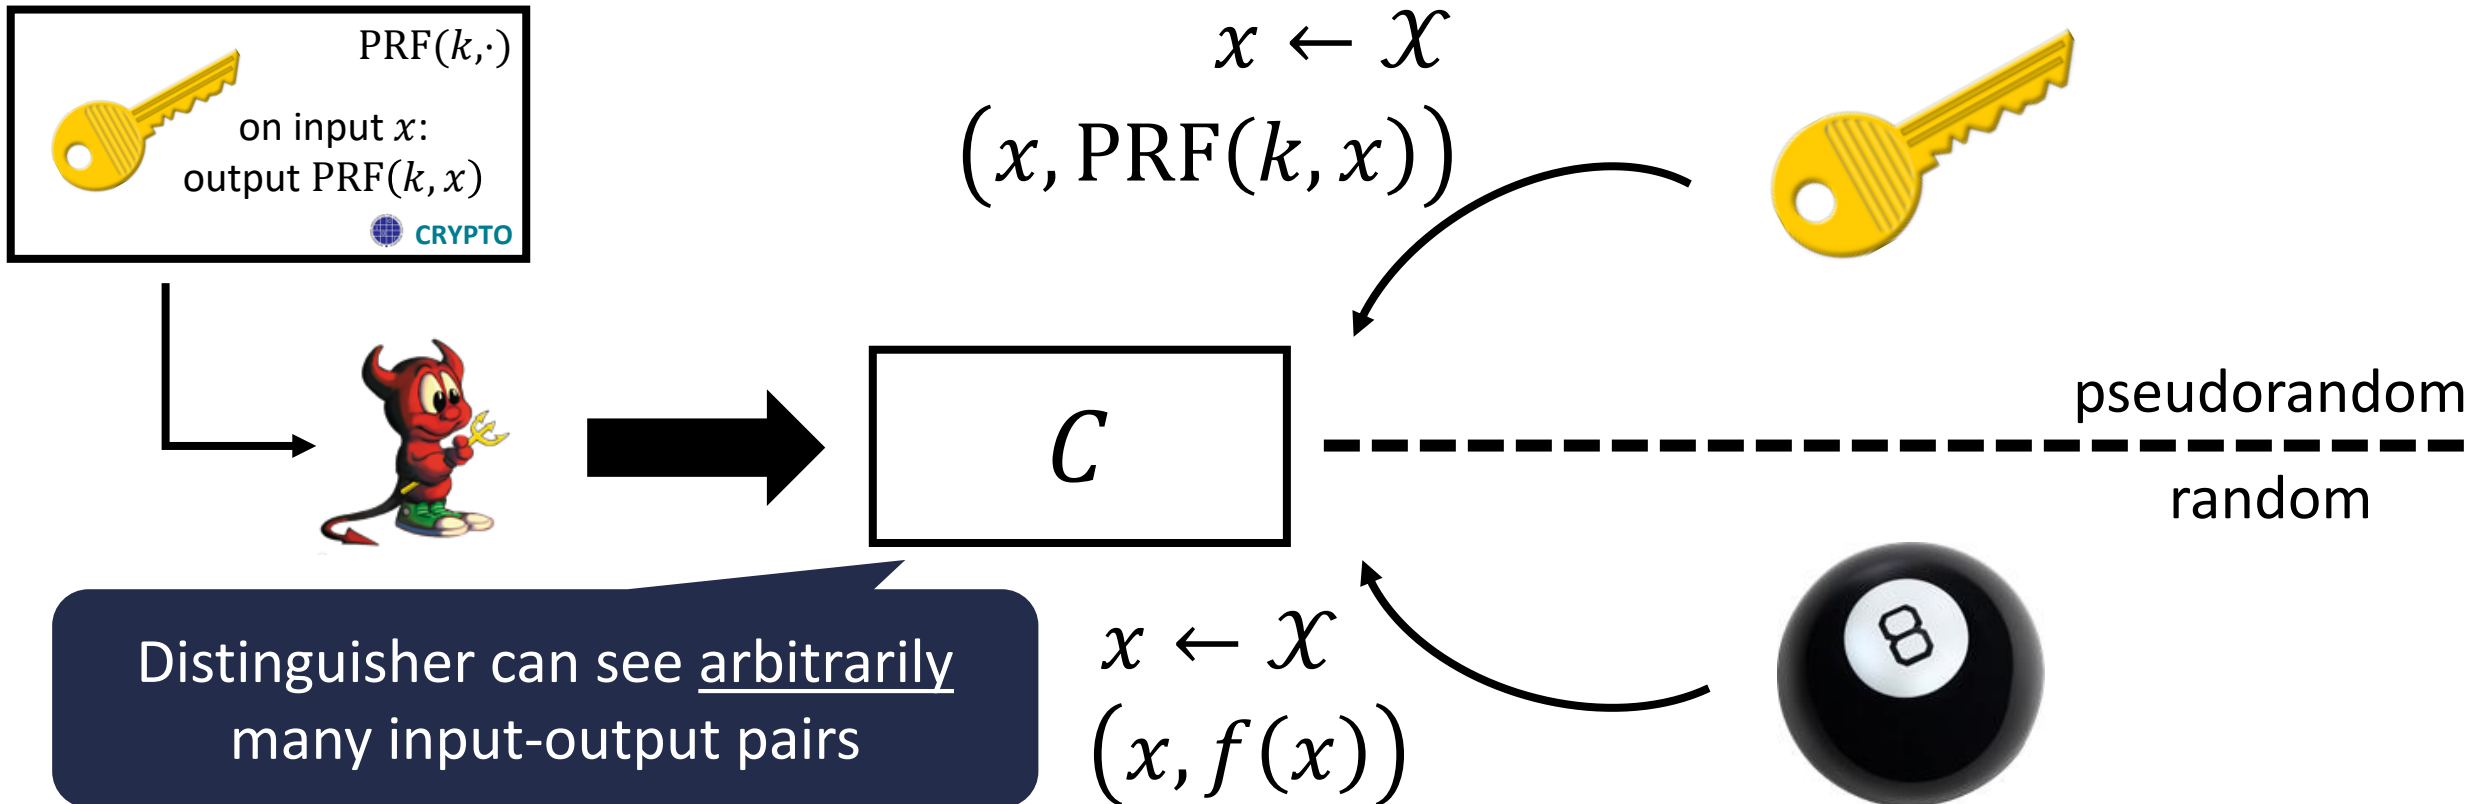
if program C can distinguish $\text{PRF}(k, \cdot)$ from random, then mark should be preserved



Traceable PRFs

Marking security (informal):


if program C can distinguish $\text{PRF}(k, \cdot)$ from random
on randomly sampled inputs, then mark should be preserved



Traceable PRFs

Marking security (informal):

if program C can **break weak pseudorandomness** of $\text{PRF}(k, \cdot)$, then mark should be preserved



Traceable PRFs

$\text{Setup}(1^\lambda) \rightarrow (\text{msk}, \text{tk})$

msk: master PRF key

tk: tracing key (can be public or secret)

$\text{KeyGen}(\text{msk}, \text{id}) \rightarrow \text{sk}_{\text{id}}$

embeds $\text{id} \in \{0,1\}^\ell$ into the key

$\text{Eval}(\text{sk}, x) \rightarrow y$

sk can be either msk or sk_{id}

$\text{Trace}^D(\text{tk}) \rightarrow T \subseteq \{0,1\}^\ell$

tracing algorithm given oracle access to weak PRF distinguisher

Traceable PRFs

$\text{Setup}(1^\lambda) \rightarrow (\text{msk}, \text{tk})$

msk: master PRF key

tk: tracing key (can be public or secret)

$\text{KeyGen}(\text{msk}, \text{id}) \rightarrow \text{sk}_{\text{id}}$

Tracing key is sampled with PRF key
(tracing algorithm needs to be able to
sample PRF evaluations)

$\text{Eval}(\text{sk}, x) \rightarrow y$

sk can be either msk or sk_{id}

$\text{Trace}^D(\text{tk}) \rightarrow T \subseteq \{0,1\}^\ell$

tracing algorithm given oracle access to
weak PRF distinguisher


Traceable PRFs

Correctness: marked and unmarked keys agree almost everywhere

$$\Pr_{x \leftarrow \mathcal{X}} [\text{Eval}(\text{msk}, x) = \text{Eval}(\text{sk}_{\text{id}}, x)] = 1 - \text{negl}(\lambda)$$

Pseudorandomness: $\text{Eval}(\text{msk}, \cdot)$ is pseudorandom

Tracing Security:



Traceable PRFs


Traceable PRF directly implies secret-key traitor tracing (via nonce-based encryption)

$$\text{Encrypt}(k, m) := (r, \text{PRF}(k, r) \oplus m)$$

Instantiate PRF with a traceable PRF

Not the case if we start with watermarkable PRF!

Tracing Security:



Traceable PRFs

Our results:

Assuming LWE, there exists a single-key traceable PRF with secret tracing

This talk

Assuming indistinguishability obfuscation and injective one-way functions, there exists a fully collusion-resistant traceable PRF with public tracing

Notably: assumptions are the same as those needed for watermarkable PRFs (and rely on similar building blocks)

Constructing Traceable PRFs

Rely on intermediate notion: **private linear constrained PRF**

(analog of private linear broadcast encryption from traitor tracing) [BSW06]



Constrained PRF key: can be used to evaluate at all points $x \in \mathcal{X}$ where $C(x) = 1$

Constructing Traceable PRFs

Rely on intermediate notion: **private linear constrained PRF**

(analog of private linear broadcast encryption from traitor tracing) [BSW06]



PRF key

Privacy: index associated with a domain element is hidden

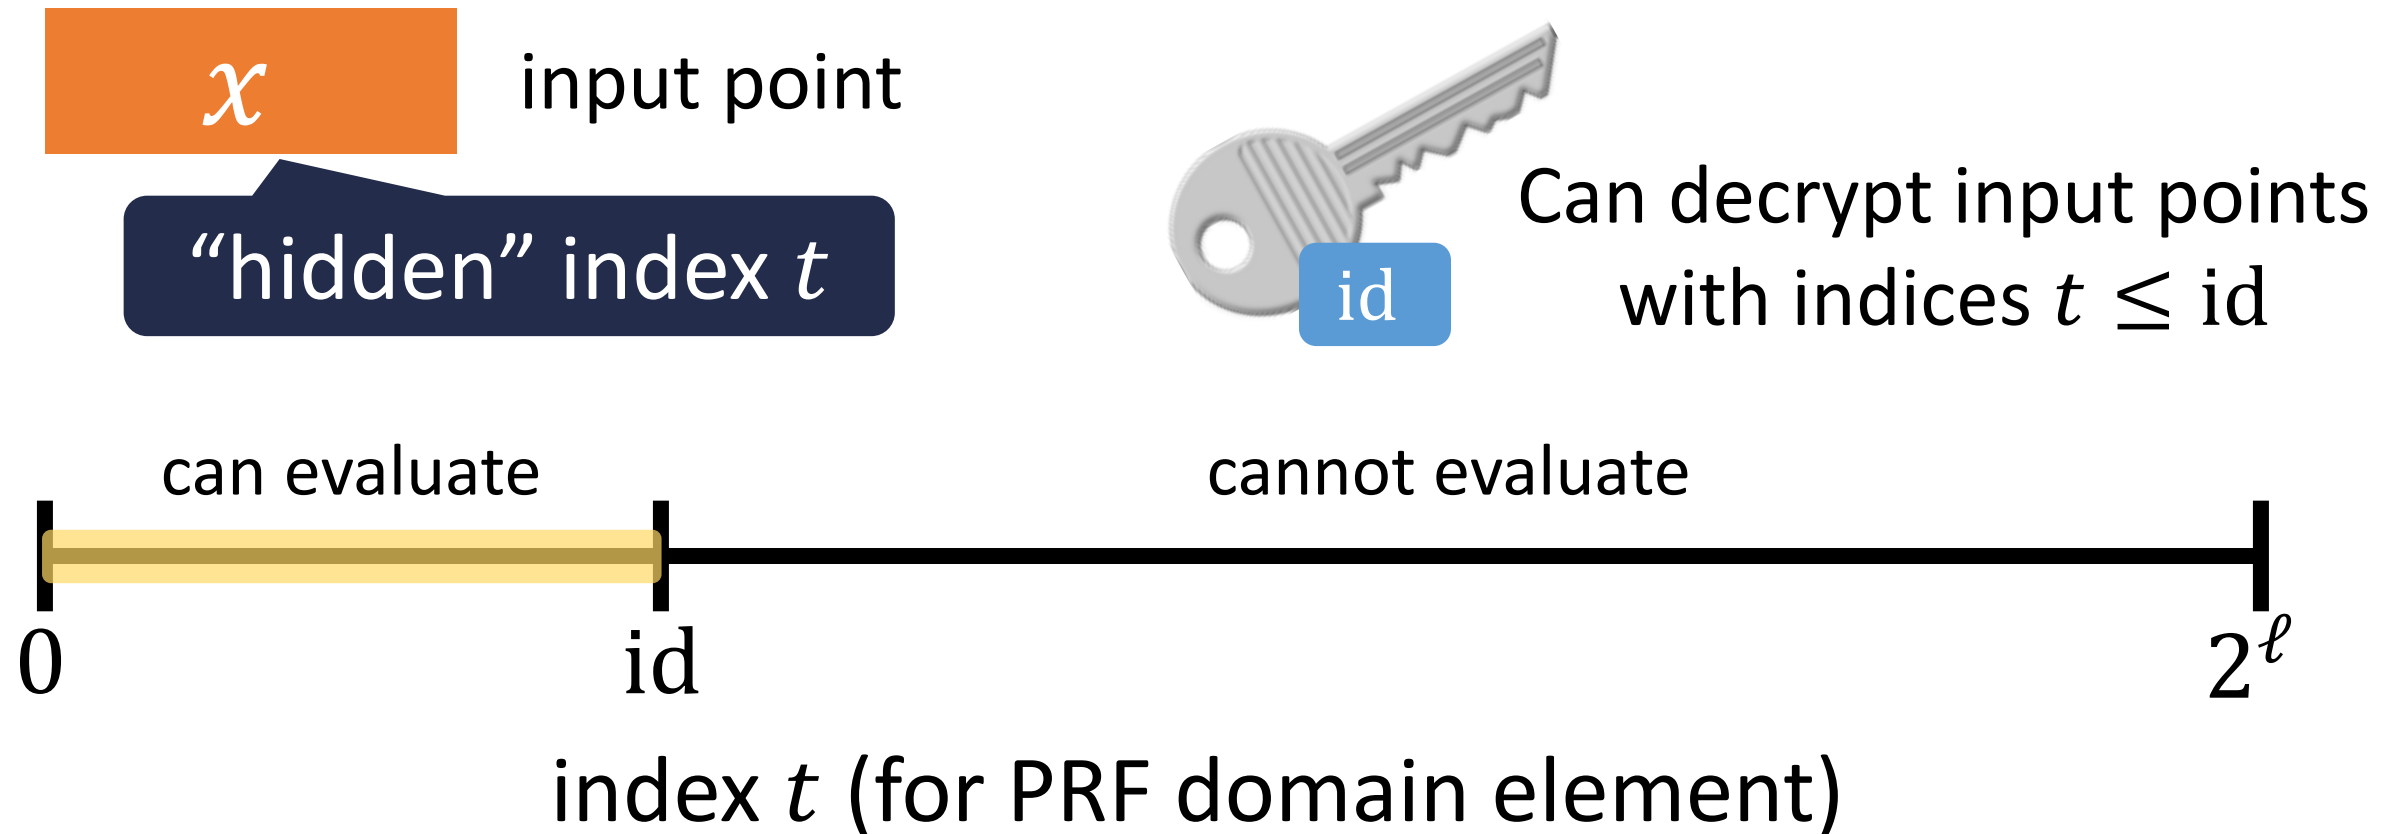
Linear constraint family:

- Some PRF inputs are associated with a (secret) index t between 0 and 2^ℓ
- Constrained key associated with $\text{id} \in [0, 2^\ell - 1]$ and can be used to evaluate on inputs whose index t satisfies $t \leq \text{id}$ (or no index)

Constructing Traceable PRFs

Rely on intermediate notion: **private linear constrained PRF**

(analog of private linear broadcast encryption from traitor tracing) [BSW06]



Constructing Traceable PRFs

Rely on intermediate notion: **private linear constrained PRF**


(analog of private linear broadcast encryption from traitor tracing) [BSW06]

Normal hiding: domain element with index 0 indistinguishable from random domain elements

Identity hiding: domain elements with index i and j are indistinguishable without key for $i \leq \text{id} < j$

Pseudorandomness: PRF outputs on inputs with index 2^ℓ are pseudorandom

(all of the properties should hold given constrained keys)




There exists a sampling algorithm to sample inputs with a specified index (could be secret-key algorithm)

Constructing Traceable PRFs

Tracing idea:

Assumption: Distinguisher D can break weak pseudorandomness with advantage ε

Implication: There must be a jump somewhere, and can only appear at id



Normal hiding

Inputs with index 0 are indistinguishable from random inputs, so decoder has advantage ε

Identity hiding

Distinguishing advantage changes negligibly when $\text{id} \notin [i, j - 1]$

Pseudorandomness

Inputs with index 2^ℓ are pseudorandom, so decoder has advantage 0

Constructing Private Linear Constrained PRF

Starting point: standard constrained PRF

Let domain $\mathcal{X} = \{0,1\}^\ell$

Problem: indices for domain element are public



$$C_{\text{id}}(t) = \begin{cases} 0, & t > \text{id} \\ 1, & t \leq \text{id} \end{cases}$$

Can decrypt input points with tags $t \leq \text{id}$

Constructing Private Linear Constrained PRF

Starting point: standard constrained PRF

Solution: Encrypt indices

Let domain $\mathcal{X} = \mathcal{CT}$ (ciphertext space for symmetric encryption scheme)



$$C_{k,\text{id}}(\text{ct}) = \begin{cases} 0, & \text{Decrypt}(k, \text{ct}) > \text{id} \\ 1, & \text{otherwise} \end{cases}$$

k : decryption key

Can decrypt input points corresponding to inputs that encrypt index greater than id

Constructing Private Linear Constrained PRF

Starting point: standard constrained PRF

Let domain $\mathcal{X} = \mathcal{CT}$

Problem: constrained key might leak k which leaks indices



$$C_{k, \text{id}}(\text{ct}) = \begin{cases} 0, & \text{Decrypt}(k, \text{ct}) > \text{id} \\ 1, & \text{otherwise} \end{cases}$$

k : decryption key

Can decrypt input points corresponding to inputs that encrypt index greater than id

Constructing Private Linear Constrained PRF

Starting point: standard constrained PRF

Let domain $\mathcal{X} = \mathcal{CT}$

Solution: use a private constrained PRF (constrained key hides constraint) [BLW17, CC17]



$$C_{k, \text{id}}(\text{ct}) = \begin{cases} 0, & \text{Decrypt}(k, \text{ct}) > \text{id} \\ 1, & \text{otherwise} \end{cases}$$

k : decryption key

Can decrypt input points corresponding to inputs that encrypt index greater than id


Constructing Traceable PRFs

Rely on intermediate notion: **private linear constrained PRF**


(analog of private linear broadcast encryption from traitor tracing) [BSW06]

Normal hiding: domain element with index 0 indistinguishable from random domain elements

Holds as long as encryption scheme has pseudorandom ciphertexts
(and constrained key hides secret key)



Constrain_c



$$C_{k,id}(ct) = \begin{cases} 0, & \text{Decrypt}(k, ct) > id \\ 1, & \text{otherwise} \end{cases}$$


Constructing Traceable PRFs

Rely on intermediate notion: **private linear constrained PRF**


(analog of private linear broadcast encryption from traitor tracing) [BSW06]

Identity hiding: domain elements with index i and j are indistinguishable without key for $i \leq \text{id} < j$

Holds as long as encryption scheme is semantically secure
(and constrained key hides secret key)



Constrain_c



$$C_{k,\text{id}}(\text{ct}) = \begin{cases} 0, & \text{Decrypt}(k, \text{ct}) > \text{id} \\ 1, & \text{otherwise} \end{cases}$$


Constructing Traceable PRFs

Rely on intermediate notion: **private linear constrained PRF**


(analog of private linear broadcast encryption from traitor tracing) [BSW06]

Pseudorandomness: PRF outputs on inputs with index 2^ℓ are pseudorandom

Holds by constrained security of constrained PRF
(constraint function always false if $\text{id} = 2^\ell$)



Constrain_c




$$C_{k,\text{id}}(\text{ct}) = \begin{cases} 0, & \text{Decrypt}(k, \text{ct}) > \text{id} \\ 1, & \text{otherwise} \end{cases}$$


Constructing Traceable PRFs

Rely on intermediate notion: **private linear constrained PRF**

(analog of private linear broadcast encryption from traitor tracing) [BSW06]




Constrain_c



$$C_{k,\text{id}}(\text{ct}) = \begin{cases} 0, & \text{Decrypt}(k, \text{ct}) > \text{id} \\ 1, & \text{otherwise} \end{cases}$$

LWE




single-key


private constrained PRF



symmetric encryption



single-key
private linear constrained PRF
(with secret sampling)




single-key
traceable PRF
(with secret tracing)

Public tracing: need a way to sample *PRF evaluations* (both inputs *and* outputs)

Unclear how to do so via private constrained PRFs, possible using indistinguishability obfuscation (with full collusion-resistance)

Traceable PRF Summary



Unremovability: Any program that can *distinguish* PRF outputs (on random inputs) must preserve the watermark

More generally: when considering software watermarking, should not always tie “functionality preserving” to “input-output preservation”


Traceable PRF Summary

Rely on intermediate notion: **private linear constrained PRF**

(analog of private linear broadcast encryption from traitor tracing) [BSW06]



$$C_{k,\text{id}}(\text{ct}) = \begin{cases} 0, & \text{Decrypt}(k, \text{ct}) > \text{id} \\ 1, & \text{otherwise} \end{cases}$$



<https://eprint.iacr.org/2020/316>

Private Constrained PRFs from Lattices

Overview of Brakerski-Vaikuntanathan and Brakerski-Tsabury-Vaikuntanathan-Wee constructions

Lattice-Based PRFs

[BPR12, BLMR13, BP14]

Learning with errors (LWE):

$$(A, s^T A + e^T) \approx (A, u^T)$$

$$A \leftarrow \mathbb{Z}_q^{n \times m}, s \leftarrow \mathbb{Z}_q^n, e \leftarrow \chi^m, u \leftarrow \mathbb{Z}_q^m$$

Learning with rounding (LWR) [BPR12]:

Replace error with deterministic rounding

$$(A, [s^T A]_p) \approx (A, u^T)$$

$$A \leftarrow \mathbb{Z}_q^{n \times m}, s \leftarrow \mathbb{Z}_q^n, u \leftarrow \mathbb{Z}_p^m$$

Lattice-Based PRFs

[BPR12, BLMR13, BP14]

Learning with rounding (LWR) [BPR12]:

replace error with deterministic rounding

$$\left(\mathbf{A}, \lfloor \mathbf{s}^T \mathbf{A} \rfloor_p \right) \approx (\mathbf{A}, \mathbf{u}^T)$$

General blueprint for lattice PRFs:

PRF family define by collection of public parameters: $\mathbf{A}_1, \dots, \mathbf{A}_\ell \in \mathbb{Z}_q^{n \times m}$

PRF key: $\mathbf{s} \leftarrow \mathbb{Z}_q^n$


PRF evaluation at x : $\mathbf{A}_1, \dots, \mathbf{A}_\ell, x \mapsto \mathbf{A}_x$

$$\text{PRF}(\mathbf{s}, x) := \lfloor \mathbf{s}^T \mathbf{A}_x \rfloor_p$$

} multiple ways to derive
 \mathbf{A}_x from $\mathbf{A}_1, \dots, \mathbf{A}_\ell$

The GSW FHE Scheme

[GSW13]



pk: $A \in \mathbb{Z}_q^{n \times m}$



sk: $s \in \mathbb{Z}_q^n$

Public key is an **LWE matrix**
(columns are LWE samples)

$$s^T A = e^T \approx \mathbf{0}^T$$

Ciphertext for $x \in \{0,1\}$:

$$A_x = AR + xG \text{ where } R \text{ is random short matrix}$$

Decryption:

$$s^T A_x = s^T AR + x \cdot s^T G \approx x \cdot s^T G$$

The GSW/BGG⁺ Homomorphisms

[GSW13, BGGHNSVV14]

$$A_1 = AR_1 + x_1G \quad \cdots \quad A_\ell = AR_\ell + x_\ell G$$

Input-independent evaluation:

$$A_1, \dots, A_\ell, f \mapsto A_f$$

Function of
 A_1, \dots, A_ℓ, f, x

$$A_f = AR_{f,x} + f(x)G \quad \text{where} \quad R_{f,x} = [R_1 \mid \cdots \mid R_\ell]H_{f,x}$$

and $H_{f,x}$ is short

Input-dependent evaluation:

$$[AR_1 \mid \cdots \mid AR_\ell]H_{f,x} = AR_{f,x}$$

$$[A_1 - x_1G \mid \cdots \mid A_\ell - x_\ell G]H_{f,x} = A_f - f(x)G$$

Lattice-Based Constrained PRFs

[BV15]

Domain: $\mathcal{X} = \{0,1\}^\rho$

Let $U_x(f) := f(x)$ be a universal circuit where $|f| = \ell$

can evaluate at x where $f(x) = 0$

Public parameters: $\mathbf{A}_1, \dots, \mathbf{A}_\ell \leftarrow \mathbb{Z}_q^{n \times m}$

PRF key: $\mathbf{s} \leftarrow \mathbb{Z}_q^n$

PRF evaluation at x :

$$\mathbf{A}_1, \dots, \mathbf{A}_\ell, x \mapsto \mathbf{A}_{U_x}$$

$$\text{PRF}(\mathbf{s}, x) := \left[\mathbf{s}^T \mathbf{A}_{U_x} \right]_p$$

to argue pseudorandomness, need to also multiply by $G^{-1}(\mathbf{D})$ where \mathbf{D} is part of public parameters

Constrained key for f :

$$\mathbf{s}^T [\mathbf{A}_1 - f_1 \cdot \mathbf{G} \mid \dots \mid \mathbf{A}_\ell - f_\ell \cdot \mathbf{G}] + \mathbf{e}^T$$

Constrained evaluation at x :

$$\begin{aligned} & \mathbf{s}^T [\mathbf{A}_1 - f_1 \cdot \mathbf{G} \mid \dots \mid \mathbf{A}_\ell - f_\ell \cdot \mathbf{G}] \mathbf{H}_{U_x, f} + \mathbf{e}^T \mathbf{H}_{U_x, f} \\ & \approx \mathbf{s}^T (\mathbf{A}_{U_x} - f(x) \cdot \mathbf{G}) \\ & = \mathbf{s}^T \mathbf{A}_{U_x} \quad \text{when } f(x) = 0 \end{aligned}$$

Lattice-Based Constrained PRFs

[BV15]

Domain: $\mathcal{X} = \{0,1\}^\rho$

Let $U_x(f) := f(x)$ be a universal circuit where $|f| = \ell$

Public parameters: $\mathbf{A}_1, \dots, \mathbf{A}_\ell \leftarrow \mathbb{Z}_q^{n \times m}$

PRF key: $\mathbf{s} \leftarrow \mathbb{Z}_q^n$

PRF evaluation at x :

$$\mathbf{A}_1, \dots, \mathbf{A}_\ell, x \mapsto \mathbf{A}_{U_x}$$

$$\text{PRF}(\mathbf{s}, x) := \left[\mathbf{s}^T \mathbf{A}_{U_x} \right]_p$$

Computing $\mathbf{H}_{U_x, f}$ requires knowledge of f
(construction does not hide the constraint)

Constrained evaluation at x :

$$\begin{aligned} & \mathbf{s}^T [\mathbf{A}_1 - f_1 \cdot \mathbf{G} \mid \dots \mid \mathbf{A}_\ell - f_\ell \cdot \mathbf{G}] \mathbf{H}_{U_x, f} + \mathbf{e}^T \mathbf{H}_{U_x, f} \\ & \approx \mathbf{s}^T (\mathbf{A}_{U_x} - f(x) \cdot \mathbf{G}) \\ & \approx \mathbf{s}^T \mathbf{A}_{U_x} \quad \text{when } f(x) = 0 \end{aligned}$$

Lattice-Based Private Constrained PRFs

[BTVW17]

Approach: encrypt the function f using an FHE scheme, and homomorphically evaluate U_x

$$\hat{f} := \text{Encrypt}(\text{pk}, f) \quad |\hat{f}| = L$$

$$\hat{U}_x(\hat{f}) := \text{FHE. Eval}(\text{pk}, U_x, \hat{f}) \quad \text{Homomorphic evaluation of } U_x \text{ on } f$$

Constrained key for f :

$$\mathbf{s}^T [\mathbf{A}_1 - \hat{f}_1 \cdot \mathbf{G} \mid \cdots \mid \mathbf{A}_\ell - \hat{f}_L \cdot \mathbf{G}] + \mathbf{e}^T$$

Constrained evaluation at x :

$$\begin{aligned} & \mathbf{s}^T [\mathbf{A}_1 - \hat{f}_1 \cdot \mathbf{G} \mid \cdots \mid \mathbf{A}_\ell - \hat{f}_L \cdot \mathbf{G}] \mathbf{H}_{\hat{U}_x, \hat{f}} + \mathbf{e}^T \mathbf{H}_{\hat{U}_x, \hat{f}} \\ & \approx \mathbf{s}^T (\mathbf{A}_{\hat{U}_x} - \hat{U}_x(\hat{f}) \cdot \mathbf{G}) \end{aligned}$$

Problem: $\hat{U}_x(\hat{f})$ is a bit of the encryption of $f(x)$, not $f(x)$

Lattice-Based Private Constrained PRFs

[BTVW17]

Straightforward to generalize homomorphic operations to matrix-valued functions:

$$\begin{aligned} A_1, \dots, A_\ell, f &\mapsto A_f \\ [A_1 - x_1 \mathbf{G} \mid \dots \mid A_\ell - x_\ell \mathbf{G}] \mathbf{H}_{f,x} &= A_f - f(x) \mathbf{G} \end{aligned}$$



$$\begin{aligned} A_1, \dots, A_\ell, f &\mapsto A_f \\ [A_1 - x_1 \mathbf{G} \mid \dots \mid A_\ell - x_\ell \mathbf{G}] \mathbf{H}_{f,x} &= A_f - \mathbf{X}_f \end{aligned}$$

$$f: \{0,1\}^\ell \mapsto b \in \{0,1\}$$



$$f: \{0,1\}^\ell \mapsto \mathbf{X}_f \in \mathbb{Z}_q^{n \times m}$$

Idea: compute \mathbf{X}_f bit-by-bit, and multiply encoding of the k^{th} bit of the j^{th} component of \mathbf{X}_f by $G^{-1}(2^k \mathbf{E}_j)$, where \mathbf{E}_j is 1 in the j^{th} component and 0 everywhere else

The GSW FHE Scheme

[GSW13]

Recall GSW decryption:

Ciphertext for $x \in \{0,1\}$: $\mathbf{A}_x = \mathbf{A}\mathbf{R} + x\mathbf{G}$

Decryption: $\mathbf{s}^T \mathbf{A}_x = \mathbf{s}^T \mathbf{A}\mathbf{R} + x \cdot \mathbf{s}^T \mathbf{G} = x \cdot \mathbf{s}^T \mathbf{G} + \text{error}$

Property: Multiplying secret key with ciphertext yields encoding of the plaintext message

Lattice-Based Private Constrained PRFs

[BTVW17]

Approach: encrypt the function f using an FHE scheme, and homomorphically evaluate U_x

$$\hat{f} := \text{Encrypt}(\text{pk}, f) \quad |\hat{f}| = L$$

$$\hat{U}_x(\hat{f}) := \text{FHE. Eval}(\text{pk}, U_x, \hat{f})$$

Homomorphic evaluation of U_x on f

Constrained key for f :

$$\mathbf{s}^T [\mathbf{A}_1 - \hat{f}_1 \cdot \mathbf{G} \mid \cdots \mid \mathbf{A}_\ell - \hat{f}_L \cdot \mathbf{G}] + \mathbf{e}^T$$

Define $\hat{U}_x(\hat{f})$ to output the GSW ciphertext (matrix-valued) obtained from homomorphic evaluation

Constrained evaluation at x :

$$\mathbf{s}^T [\mathbf{A}_1 - \hat{f}_1 \cdot \mathbf{G} \mid \cdots \mid \mathbf{A}_\ell - \hat{f}_L \cdot \mathbf{G}] \mathbf{H}_{\hat{U}_x, \hat{f}} + \mathbf{e}^T \mathbf{H}_{\hat{U}_x, \hat{f}}$$

$$\approx \mathbf{s}^T (\mathbf{A}_{\hat{U}_x} - \hat{U}_x(\hat{f}))$$

$$\approx \mathbf{s}^T (\mathbf{A}_{\hat{U}_x} - f(x) \cdot \mathbf{G})$$

Insight: If \mathbf{s} is also the secret key for the GSW encryption scheme, then

$$\mathbf{s}^T \hat{U}_x(\hat{f}) = f(x) \cdot \mathbf{s}^T \mathbf{G} + \text{error}$$

Lattice-Based Private Constrained PRFs

[BTVW17]

Approach: encrypt the function f using an FHE scheme, and homomorphically evaluate U_x

$$\hat{f} := \text{Encrypt}(\text{pk}, f) \quad |\hat{f}| = L$$

$$\hat{U}_x(\hat{f}) := \text{FHE. Eval}(\text{pk}, U_x, \hat{f})$$

Homomorphic evaluation of U_x on f

Constrained key for f :

$$\mathbf{s}^T [\mathbf{A}_1 - \hat{f}_1 \cdot \mathbf{G} \mid \cdots \mid \mathbf{A}_\ell - \hat{f}_L \cdot \mathbf{G}] + \mathbf{e}^T$$

Define $\hat{U}_x(\hat{f})$ to output the GSW ciphertext (matrix-valued) obtained from homomorphic evaluation

Constrained evaluation at x :

$$\mathbf{s}^T [\mathbf{A}_1 - \hat{f}_1 \cdot \mathbf{G} \mid \cdots \mid \mathbf{A}_\ell - \hat{f}_L \cdot \mathbf{G}] \mathbf{H}_{\hat{U}_x, \hat{f}} + \mathbf{e}^T \mathbf{H}_{\hat{U}_x, \hat{f}}$$

$$\approx \mathbf{s}^T (\mathbf{A}_{\hat{U}_x} - \hat{U}_x(\hat{f}))$$

$$\approx \mathbf{s}^T (\mathbf{A}_{\hat{U}_x} - f(x) \cdot \mathbf{G})$$

Some tweaks needed to argue security (see [BTVW17] for full details)

Summary

Input-independent evaluation:

$$\mathbf{A}_1, \dots, \mathbf{A}_\ell, f \mapsto \mathbf{A}_f \quad \text{PRF evaluation}$$

Input-dependent evaluation:

$$[\mathbf{A}_1 - x_1 \mathbf{G} \mid \dots \mid \mathbf{A}_\ell - x_\ell \mathbf{G}] \mathbf{H}_{f,x} = \mathbf{A}_f - f(x) \mathbf{G} \quad \text{constrained evaluation}$$

Constraint privacy:

- Encrypt constraint using GSW FHE scheme
- LWE secret reused for PRF secret key and FHE secret key

Thank you!