


Pattern Matching in Encrypted Stream from Inner Product Encryption

Elie Bouscaté^(1,2) Guilhem Castagnos⁽¹⁾ Olivier Sanders⁽²⁾

(1) Université de Bordeaux

(2) Orange Labs


End-to-End Encryption



More and more encrypted data

- about 90% of Internet traffic worldwide is encrypted
- increased use of encrypted messaging services (WhatsApp,...)


End-to-End Encryption



Standard encryption protocols **designed to prevent any processing**

- arbitration between privacy and functionalities

End-to-End Encryption



Current solutions imply **decryption by a gateway**

- the gateway can access all data exchanged through the channel
- **what is the point of end-to-end encryption?**

Pattern Matching

- Many applications perform pattern matching
 - Intrusion Detection System (IDS)
 - Content filtering
 - Searches on genomic data
 - ...
- Example of Snort rules:
 - alert tcp (msg:"MALWARE-BACKDOOR - Dagger_ 1.4.0";
content:"2| 00 00 00 06 00 00 00 | Drives | 24 00 |",depth 16;)
 - alert tcp (msg:"MALWARE-BACKDOOR QAZ Worm Client Login
access"; content:"qazwsx.hsq";)

Searchable Encryption

- SE is an encryption scheme with **additional features**:
 - given sk one can derive a trapdoor td_w on a pattern W
 - given td_w , the gateway can test whether $C = \text{Encrypt}(W)$

Searchable Encryption

- SE is an encryption scheme with **additional features**:
 - given sk one can derive a trapdoor td_w on a pattern W
 - given td_w , the gateway can test whether $C = \text{Encrypt}(W)$

SE does not address our problem

Searchable Encryption


- SE is an encryption scheme with **additional features**:
 - given sk one can derive a trapdoor td_w on a pattern W
 - given td_w , the gateway can test whether $C = \text{Encrypt}(W)$

SE does not address our problem

- We need:
 - given sk one can derive a trapdoor td_w on a pattern W
 - given td_w , the gateway can check whether W is contained in the **stream** encrypted in C

Dealing with Data Streams

SE-based solutions follow the **sliding window method**:



patterns

host

hostile

...

- Each C_i can be tested using td_W
- The process must be **repeated for each possible length** of keywords

Dealing with Data Streams


SE-based solutions follow the **sliding window method**:



- Each C_i can be tested using td_W
- The process must be **repeated for each possible length** of keywords

Dealing with Data Streams


SE-based solutions follow the **sliding window method**:



- Each C_i can be tested using td_W
- The process must be **repeated for each possible length** of keywords

Dealing with Data Streams


SE-based solutions follow the **sliding window method**:



- Each C_i can be tested using td_W
- The process must be **repeated for each possible length** of keywords

Dealing with Data Streams



SE-based solutions follow the **sliding window method**:



- Each C_i can be tested using td_W
- The process must be **repeated for each possible length** of keywords

Dealing with Data Streams

SE-based solutions follow the **sliding window method**:



- Each C_i can be tested using td_W
- Splitting keywords **harms privacy**

Hidden Vector Encryption

- n -HVE enables to encrypt vectors $\mathbf{x} = (x_1, \dots, x_n) \in \Sigma^n$
- Secret keys are associated with vectors $\mathbf{k} = (k_1, \dots, k_n)$ where
 - $k_i \in \Sigma$
 - or $k_i = *$ (wildcard)
- Given $sk_{\mathbf{k}}$ and an encryption of \mathbf{x} , Test returns
 - 1 if $\forall i \in [1, n] k_i = * \text{ or } k_i = x_i$
 - 0 otherwise
- No leakage beyond the output of $\text{Test}(sk_{\mathbf{k}}, \mathbf{x})$

Dealing with Data Streams

plaintext	\$	h	o	s	t	i	l	e
$k_{\text{host},0}$	h	o	s	t	*	*	*	*
$k_{\text{host},1}$	*	h	o	s	t	*	*	*
$k_{\text{host},2}$	*	*	h	o	s	t	*	*
$k_{\text{host},3}$	*	*	*	h	o	s	t	*
$k_{\text{host},4}$	*	*	*	*	h	o	s	t

pattern: host

- Wildcards enable to deal with offsets
- Secret keys **must be issued for each possible offset**
⇒ n keys per pattern for n bytes stream

SEPM

- Searchable Encryption supporting Pattern Matching (SEPM) is essentially a HVE with **key size independent of n**
- State-of-the-Art:

DFOS18 ¹: First construction, **$O(n)$ public key**, proof in the GGM

BCC20 ²: **public key size independent of n** , proof in the GGM

BCS21 ³: improved performance, **selective security under static assumption**

Ad-hoc constructions, **no adaptive security under standard assumptions**

¹Desmoulins *et al.* Pattern Matching on Encrypted Streams. Asiacrypt 2018.

²Bkakria *et al.* Privacy-Preserving Pattern Matching on Encrypted Data. Asiacrypt 2020.

³Bouscatié *et al.* Public Key Encryption with Flexible Pattern Matching. Asiacrypt 2021.

Our Contribution

Our Goal

We want to identify/improve relations between primitives



Very few constructions

Few constructions




Many constructions

- No known relation between SEPM and HVE

Our Goal

We want to identify/improve relations between primitives



- No known relation between SEPM and HVE
- Generic transformation (KSW^4) from IPE to HVE with two-fold ciphertext increase

⁴Katz *et al.* Predicate Encryption Supporting Disjunctions, Polynomial Equations, and Inner Products. Journal of Cryptology 2013.

First Step: HVE \rightarrow SEPM

[BCC20] introduced the **fragmentation** approach:

- streams are **split into overlapping fragments**
- d is a bound on the largest searchable pattern

$$\mathbf{x} = \overbrace{x_1, \dots, x_d}^{x_1}, \underbrace{x_{d+1}, \dots, x_{2d}}_{x_2}, \overbrace{x_{2d+1}, \dots, x_{3d}}^{x_3}, \underbrace{x_{3d+1}, \dots, x_{4d}}_{x_4}, \overbrace{x_{4d+1}, \dots, x_{5d}}^{x_5}, \dots$$

Fragmentation is much more powerful than initially thought

First Step: HVE \rightarrow SEPM

We use fragmentation to circumvent HVE limitations for streams

- streams are split into overlapping fragments
- d is a bound on the largest searchable pattern

$$\mathbf{x} = \overbrace{x_1, \dots, x_d}^{x_1} \underbrace{\overbrace{x_{d+1}, \dots, x_{2d}}^{x_2} \overbrace{x_{2d+1}, \dots, x_{3d}}^{x_3} \overbrace{x_{3d+1}, \dots, x_{4d}}^{x_4} \overbrace{x_{4d+1}, \dots, x_{5d}}^{x_5}}_{x_4}, \dots$$

- Use $2d$ -HVE to encrypt each fragment
- Use wildcards to deal with pattern offsets within a fragment
 $O(d)$ key size (independent of n)
- Run HVE.Test in each relevant fragment

Second Step: IPE \rightarrow HVE (KSW)

KSW proposed a generic conversion $2n$ -IPE \rightarrow n -HVE

$$\mathbf{x}' = (x_1 \cdot r_1, \dots, x_n \cdot r_n, -r_1, \dots, -r_n)$$

- HVE.Encrypt(x_1, \dots, x_n): $r \xleftarrow{\$} \mathbb{Z}_p$ and $C \leftarrow$ IPE.Encrypt(\mathbf{x}')

Second Step: IPE \rightarrow HVE (KSW)

KSW proposed a generic conversion $2n$ -IPE \rightarrow n -HVE

$$\begin{aligned} \mathbf{x}' &= (x_1 \cdot r_1, \dots, x_n \cdot r_n, -r_1, \dots, -r_n) \\ \mathbf{k}' &= (1, \dots, 1, k_1, \dots, k_n) \end{aligned}$$

- $\text{HVE.Keygen}(k_1, \dots, k_n): \text{sk}_k \leftarrow \text{IPE.Keygen}(\mathbf{k}')$

Second Step: IPE \rightarrow HVE (KSW)

KSW proposed a generic conversion $2n$ -IPE \rightarrow n -HVE

$$\begin{aligned} \mathbf{x}' &= (x_1 \cdot r_1, \dots, x_n \cdot r_n, -r_1, \dots, -r_n) \\ \mathbf{k}' &= (1, \dots, 1, k_1, \dots, k_n) \end{aligned}$$

- $\text{HVE.Test}(C, \text{sk}_{\mathbf{k}})$: return 1 iff $\langle \mathbf{x}', \mathbf{k}' \rangle = 0$ using $\text{IPE.Eval}(C, \text{sk}_{\mathbf{k}})$

Correctness:

- $(x_1, \dots, x_n) = (k_1, \dots, k_n) \Rightarrow \langle \mathbf{x}', \mathbf{k}' \rangle = 0$
- $\langle \mathbf{x}', \mathbf{k}' \rangle = 0 \Rightarrow (x_1, \dots, x_n) = (k_1, \dots, k_n)$ with overwhelming probability...

Second Step: IPE \rightarrow HVE (KSW)

KSW proposed a generic conversion $2n$ -IPE \rightarrow n -HVE

$$\mathbf{x}' = (x_1 \cdot r_1, \dots, x_n \cdot r_n, -r_1, \dots, -r_n)$$
$$\mathbf{k}' = (1, \dots, 1, k_1, \dots, k_n)$$

- $\text{HVE.Test}(C, \text{sk}_k)$: return 1 iff $\langle \mathbf{x}', \mathbf{k}' \rangle = 0$ using $\text{IPE.Eval}(C, \text{sk}_k)$

Correctness:

- $(x_1, \dots, x_n) = (k_1, \dots, k_n) \Rightarrow \langle \mathbf{x}', \mathbf{k}' \rangle = 0$
- $\langle \mathbf{x}', \mathbf{k}' \rangle = 0 \Rightarrow (x_1, \dots, x_n) = (k_1, \dots, k_n)$ with overwhelming probability...

for honest key queries!

Second Step: IPE \rightarrow HVE (KSW)

KSW proposed a generic conversion $2n$ -IPE \rightarrow n -HVE

$$\mathbf{x}' = (x_1 \cdot r_1, \dots, x_n \cdot r_n, -r_1, \dots, -r_n)$$
$$\mathbf{k}' = (1, \dots, 1, k_1, \dots, k_n)$$

- $\text{HVE.Test}(C, \text{sk}_{\mathbf{k}})$: return 1 iff $\langle \mathbf{x}', \mathbf{k}' \rangle = 0$ using $\text{IPE.Eval}(C, \text{sk}_{\mathbf{k}})$

Security:

- no formal security proof in KSW
- malicious queries create discrepancy between IPE/HVE experiments:
 - Given the challenge $\text{IPE.Encrypt}(\mathbf{x}')$, \mathcal{A} chooses \mathbf{k}' such that $\langle \mathbf{x}', \mathbf{k}' \rangle = 0$ but $(x_1, \dots, x_n) \neq (k_1, \dots, k_n)$
 - Valid query for the HVE experiment but not for the IPE one
- We show how to circumvent this issue in our paper

Second Step: IPE \rightarrow HVE (Ours)

We introduce a generic conversion $(n + 1)$ -IPE \rightarrow n -HVE

$$\begin{aligned} \mathbf{x}' &= (x_1, \dots, x_n, -1) \\ \mathbf{k}' &= (r_1, \dots, r_n, \langle \mathbf{k}, \mathbf{r} \rangle) \end{aligned}$$

- two-fold improvement compared to KSW
- $(x_1, \dots, x_n) = (k_1, \dots, k_n) \Rightarrow \langle \mathbf{x}', \mathbf{k}' \rangle = 0$
- $\langle \mathbf{x}', \mathbf{k}' \rangle = 0 \Rightarrow (x_1, \dots, x_n) = (k_1, \dots, k_n)$ with overwhelming probability... **still for honest key queries!**

Second Step: IPE \rightarrow HVE (Ours)

We introduce a generic conversion $(n + 1)$ -IPE \rightarrow n -HVE

$$\begin{aligned} \mathbf{x}' &= (x_1, \dots, x_n, -1) \\ \mathbf{k}' &= (r_1, \dots, r_n, \langle \mathbf{k}, \mathbf{r} \rangle) \end{aligned}$$


- two-fold improvement compared to KSW
- $(x_1, \dots, x_n) = (k_1, \dots, k_n) \Rightarrow \langle \mathbf{x}', \mathbf{k}' \rangle = 0$
- $\langle \mathbf{x}', \mathbf{k}' \rangle = 0 \Rightarrow (x_1, \dots, x_n) = (k_1, \dots, k_n)$ with overwhelming probability... **still for honest key queries!**

Security

- **Selective security holds** without additional assumptions
- **Adaptive security needs an additional assumption**


Second Step: IPE \rightarrow HVE (Ours)

Adaptive security experiment




Second Step: IPE \rightarrow HVE (Ours)

Adaptive security experiment



Second Step: IPE \rightarrow HVE (Ours)

Adaptive security experiment



Example of **malicious selection** of \mathbf{x}_0 and \mathbf{x}_1

- $\forall i, b, \mathbf{k}_i \neq \mathbf{x}_b$ (**valid HVE queries**)
- $\exists i^*, b^* : \langle \mathbf{k}'_{i^*}, \mathbf{x}_{b^*} \rangle = 0$ and $\langle \mathbf{k}'_{i^*}, \mathbf{x}_{1-b^*} \rangle \neq 0$ (**invalid IPE queries**)


Second Step: IPE \rightarrow HVE (Ours)

Main observations

- Knowledge of \mathbf{k}'_i enables to find such \mathbf{x}_b
 - $\Rightarrow \text{sk}_{\mathbf{k}'_i}$ must hide \mathbf{k}'_i
- Function-privacy **strongly depends on the context** (pattern entropy)
 - \Rightarrow we want a property independent of the context
- \mathcal{A} has some control over \mathbf{k}'_i as it selects \mathbf{k}_i
 - \Rightarrow must be modelled by our property

Second Step: IPE \rightarrow HVE (Ours)

We introduce **key privacy** for IPE

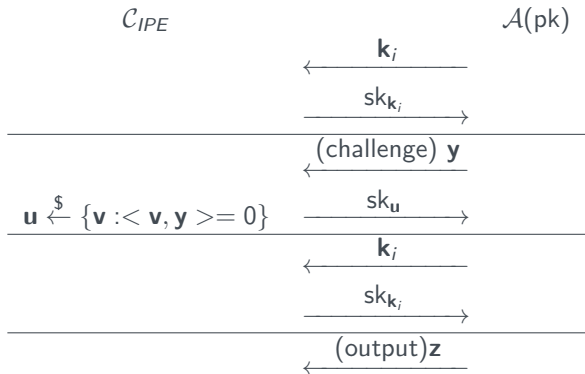


\mathcal{A} has some control on u through the choice of y

sufficient for our IPE \rightarrow HVE transformation

Second Step: IPE \rightarrow HVE (Ours)

We introduce **key privacy** for IPE



\mathcal{A} wins if:

- $\langle u, z \rangle = 0$
- $z \notin \mathbf{Vect}(y)$ (no trivial win)

Second Step: IPE \rightarrow HVE (Ours)

- We **prove adaptive security** of our conversion with **key private** IPE
- Key privacy can **easily** be **assessed**
 - independent of the context
 - achieved by KSW under Discrete Log assumption
 - achieved by OT⁵ under Discrete Log assumption
 - we provide examples of scheme that do not achieve it

\Rightarrow some IPE schemes yield better HVE than others

⁵Okamoto *et al.* Adaptively attribute-hiding inner product encryption. Eurocrypt 2012.

Conclusion

- We show how to generically build SEPM from HVE
- We revisit relations between HVE and IPE
 - We introduce a more efficient conversion IPE \rightarrow HVE than KSW
 - We prove selective security of our conversion without additional condition
 - We prove adaptive security assuming a mild condition we formalise
- Our work leads to many new SEPM constructions with new features