

LaPSuS: A Lattice-Based PSA Scheme under Scrutiny




SCN 2024, International Conference on Security and Cryptography for Networks

Johannes Ottenhues (Univ. St. Gallen), [Alexander Koch](#) (IRIF, Univ. Paris Cité, CNRS) | 2024-09-11



Private Stream Aggregation (PSA)

Setting: Each client has a sensor, producing a private value $x_i \in \mathbb{Z}_q$




What is $\sum_{i \in [n]} x_i$?

Private Stream Aggregation (PSA)

Setting: Each client has a sensor, producing a private value $x_{i,\ell} \in \mathbb{Z}_q$ at time step ℓ

$\ell =$ 




What is $\sum_{i \in [n]} x_{i,\ell}$ for a given time step ℓ ?

Private Stream Aggregation (PSA)

Setting: Each client has a sensor, producing a private value $x_{i,l} \in \mathbb{Z}_q$ at time step l

$l =$ 



What is $\sum_{i \in [n]} x_{i,l}$ for a given time step l ?

Private Stream Aggregation (PSA)

Setting: Each client has a sensor, producing a private value $x_{i,\ell} \in \mathbb{Z}_q$ at time step ℓ

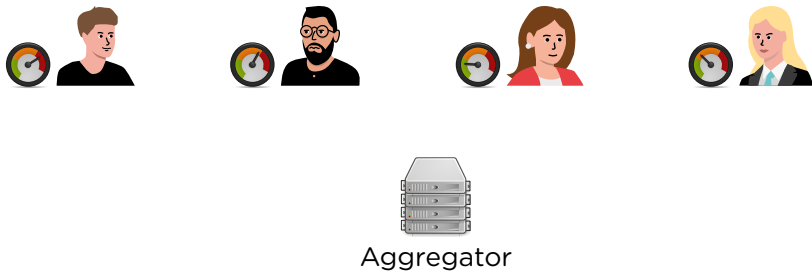


Image "server multiple" by RRZE, CC BY-SA 3.0

Private Stream Aggregation (PSA)

PSA: A 3-tuple of efficient algorithms (**Setup**, Enc, AggrDec)

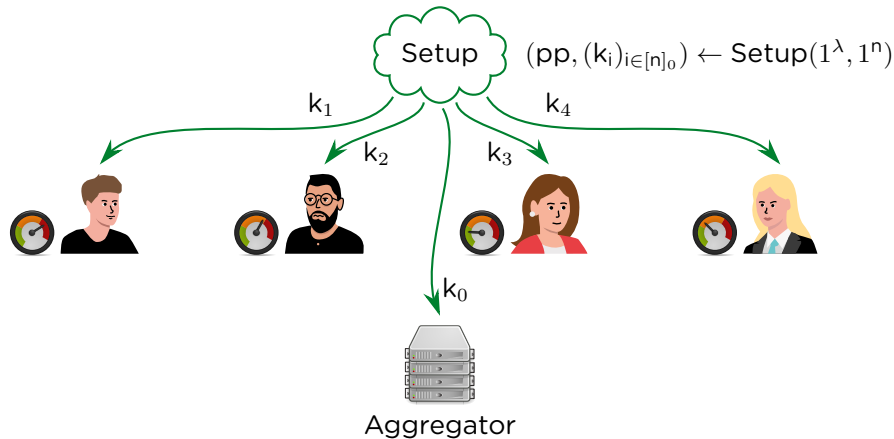


Image "server multiple" by RRZE, CC BY-SA 3.0

Private Stream Aggregation (PSA)

PSA: A 3-tuple of efficient algorithms (Setup, Enc, AggrDec)

$\ell =$ 




Image "server multiple" by RRZE, CC BY-SA 3.0

Private Stream Aggregation (PSA)

PSA: A 3-tuple of efficient algorithms (Setup, Enc, AggrDec)

$\ell =$ 

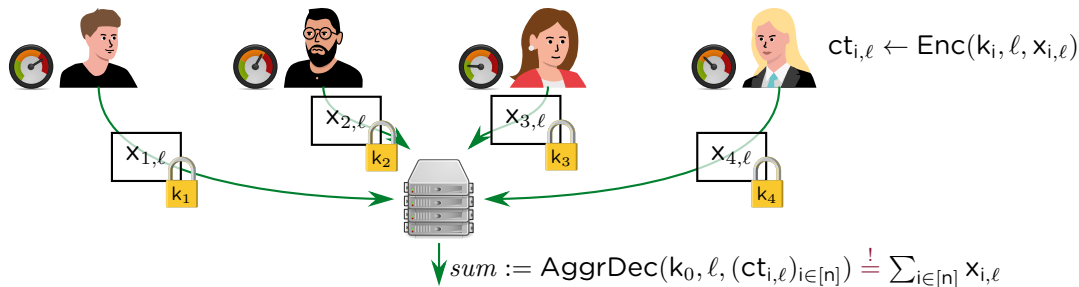


Image "server multiple" by RRZE, CC BY-SA 3.0

Private Stream Aggregation (PSA)

A note on **inherent leakage** if you get more than one ct per client and label:

$\ell =$ 

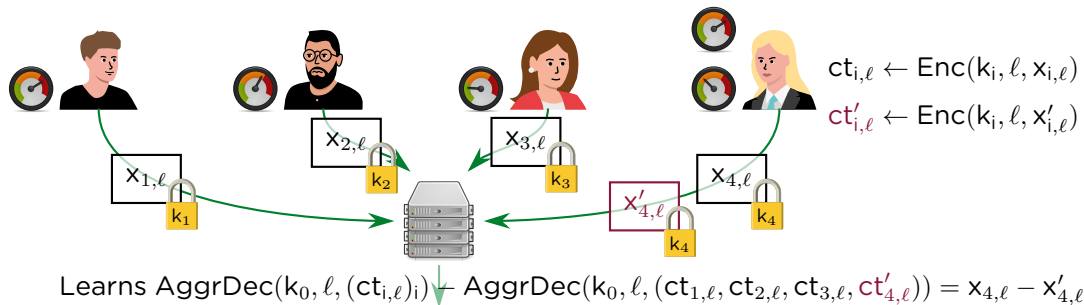
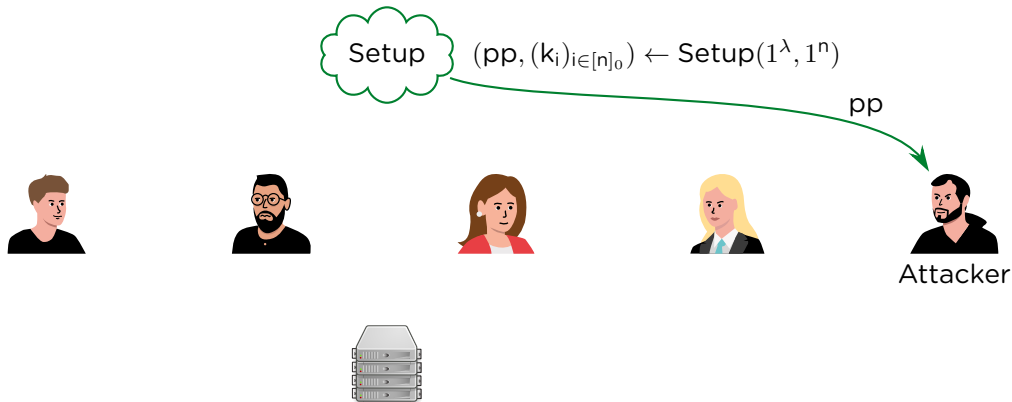


Image "server multiple" by RRZE, CC BY-SA 3.0

Security Notion: Aggregator Obliviousness (AO)

Attacker gets pp and oracles: $QEnc(i, x_i, \ell)$, $QCorrupt(i)$, $QChallenge(\mathcal{U}, (x_i^0)_{i \in \mathcal{U}}, (x_i^1)_{i \in \mathcal{U}}, \ell^*)$



Security Notion: Aggregator Obliviousness (AO)

Attacker gets pp and oracles: $\text{QEnc}(i, x_i, \ell)$, $\text{QCorrupt}(i)$, $\text{QChallenge}(\mathcal{U}, (x_i^0)_{i \in \mathcal{U}}, (x_i^1)_{i \in \mathcal{U}}, \ell^*)$

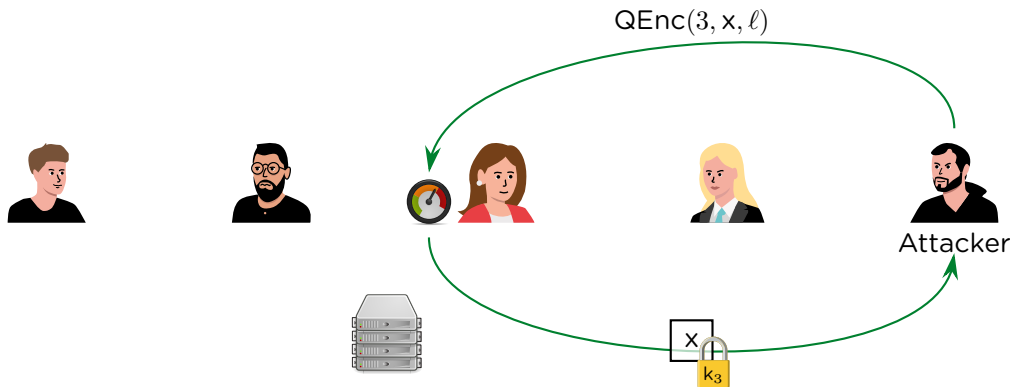


Image "server multiple" by RRZE, CC BY-SA 3.0

Security Notion: Aggregator Obliviousness (AO)

Attacker gets pp and oracles: $QEnc(i, x_i, \ell)$, $QCorrupt(i)$, $QChallenge(\mathcal{U}, (x_i^0)_{i \in \mathcal{U}}, (x_i^1)_{i \in \mathcal{U}}, \ell^*)$

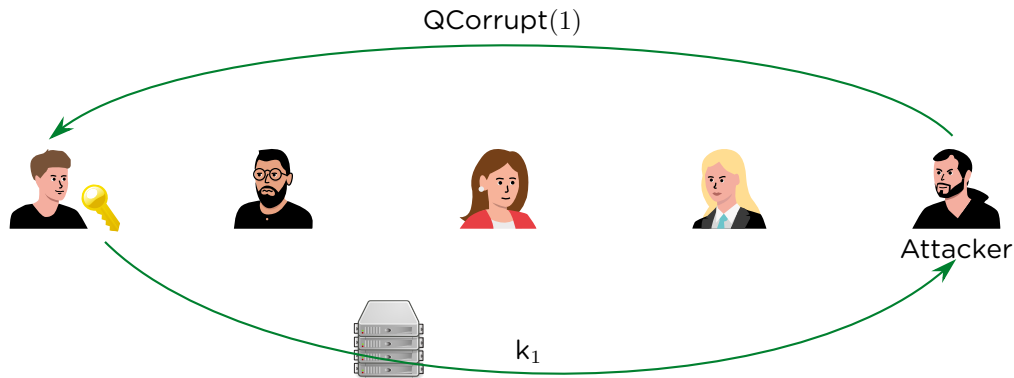


Image "server multiple" by RRZE, CC BY-SA 3.0

Security Notion: Aggregator Obliviousness (AO)

Attacker gets pp and oracles: $\text{QEnc}(i, x_i, \ell)$, $\text{QCorrupt}(i)$, $\text{QChallenge}(\mathcal{U}, (x_i^0)_{i \in \mathcal{U}}, (x_i^1)_{i \in \mathcal{U}}, \ell^*)$

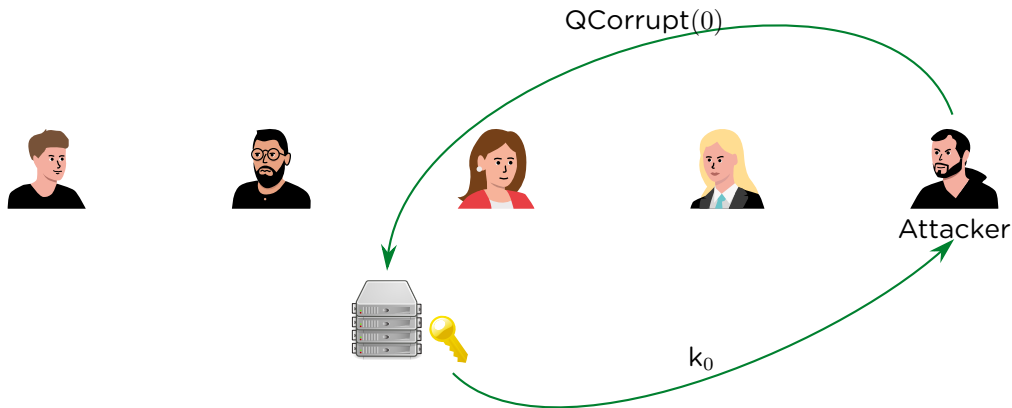


Image "server multiple" by RRZE, CC BY-SA 3.0

Security Notion: Aggregator Obliviousness (AO)

Attacker gets pp and oracles: $\text{QEnc}(i, x_i, \ell)$, $\text{QCorrupt}(i)$, $\text{QChallenge}(\mathcal{U}, (x_i^0)_{i \in \mathcal{U}}, (x_i^1)_{i \in \mathcal{U}}, \ell^*)$

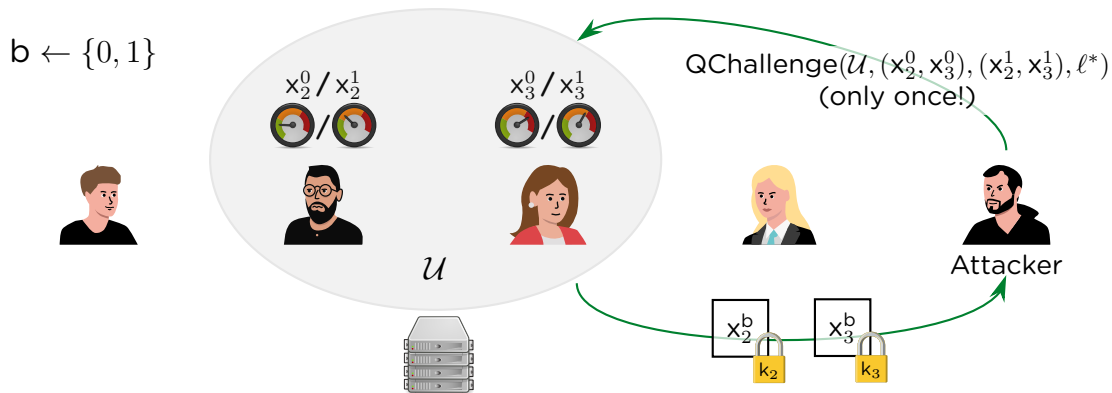


Image "server multiple" by RRZE, CC BY-SA 3.0

Security Notion: Aggregator Obliviousness (AO)

Aim of the attacker: Guess bit b .

$$b \leftarrow \{0, 1\}$$



Attacker Rules

- Keep \mathcal{U} disjoint from $\{\text{corrupted}\} =: \mathcal{CS}$!
- If you corrupted the aggregator **and** \mathcal{U} , \mathcal{CS} , and the set of users from which you got an encryption for ℓ^* covers everyone, **enforce** $\sum x_i^0 = \sum x_i^1$

... **or you loose!**



Attacker

Security Notion: Aggregator Obliviousness (AO)

Aim of the attacker: Guess bit b .

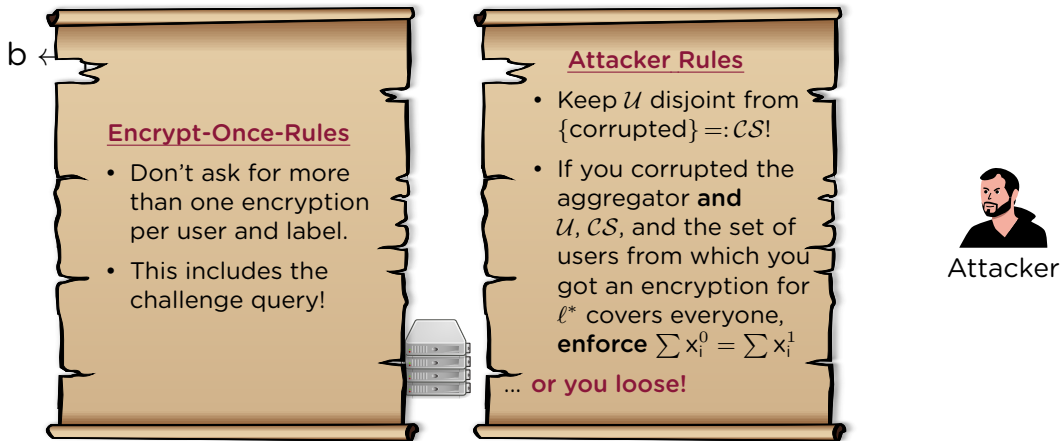


Image "server multiple" by RRZE, CC BY-SA 3.0

LaPS Scheme (Becker et al., NDSS 2018)

- $(\text{Gen}_{\text{ah}}, \text{Enc}_{\text{ah}}, \text{Dec}_{\text{ah}})$: **additively homomorphic** PKE with pseudorandom ciphertexts.
- \mathbf{G} : gadget matrix, i.e. $\mathbf{G} = \mathbf{I} \otimes (1 \ 2 \ \dots \ 2^{d-1})$, with $d = O(\lceil \log q \rceil)$
- $e \leftarrow \mathcal{D}_{\Lambda_{\mathbf{v}}^{\perp}(\mathbf{G}), \sigma}$: error distribution (with $\sigma > 2\sqrt{\lambda}$); defined so that $\mathbf{G} \cdot e \bmod q = \mathbf{v}$

Assumption: “Augmented LWE” (El Bansarkhani et al.): additionally hides message x inside error-term of LWE sample, via: encode x into pseudorandom \mathbf{v} , sample e from $\mathcal{D}_{\Lambda_{\mathbf{v}}^{\perp}(\mathbf{G}), \sigma}$

Setup($1^\lambda, 1^n$):

```
 $\mathbf{A} \leftarrow \mathbb{Z}_q^{\lambda \times d}$   
for  $i \in [n]$ :  $k_i \leftarrow \mathbb{Z}_q^\lambda$   
 $(\text{pk}_{\text{ah}}, \text{sk}_{\text{ah}}) \leftarrow \text{Gen}_{\text{ah}}(1^\lambda)$   
 $\text{pp} := (\mathbf{A}, \text{pk}_{\text{ah}})$   
 $k_0 := (\text{sk}_{\text{ah}}, \sum_{i \in [n]} k_i)$   
return  $(\text{pp}, (k_i)_{i \in [n]_0})$ 
```

Enc(pp, k_i, x_i):

```
parse  $\text{pp} = (\mathbf{A}, \text{pk}_{\text{ah}})$   
 $\mathbf{v}_i \leftarrow \text{Enc}_{\text{ah}}(\text{pk}_{\text{ah}}, x_i)$   
 $e_i \leftarrow \mathcal{D}_{\Lambda_{\mathbf{v}_i}^{\perp}(\mathbf{G}), \sigma}$   
 $\text{ct}_i = k_i^\top \mathbf{A} + e_i^\top \bmod q$   
return  $\text{ct}_i$ 
```

AggrDec($\text{pp}, k_0, (\text{ct}_i)_{i \in [n]}$):

```
parse  $k_0 = (\text{sk}_{\text{ah}}, k)$   
 $e = \sum_{i \in [n]} \text{ct}_i - k^\top \mathbf{A}$   
return  $\text{Dec}_{\text{ah}}(\text{sk}_{\text{ah}}, \mathbf{G} \cdot e \bmod q)$ 
```

First Attack on LaPS using two Ciphertexts of a Client

Given two encryptions ct_i and ct'_i of x_i and x'_i and $k_0 = (sk_{ah}, k)$.

$$ct_i - ct'_i = k_i^\top \mathbf{A} + \mathbf{e}_i^\top - k_i^\top \mathbf{A} + \mathbf{e}'_i{}^\top \pmod{q} = \mathbf{e}_i^\top - \mathbf{e}'_i{}^\top \pmod{q}.$$

Then:

$$\mathbf{G}(\mathbf{e}_i^\top - \mathbf{e}'_i{}^\top) \pmod{q} = \mathbf{v}_i - \mathbf{v}'_i \pmod{q}.$$

We can then decrypt to get

$$\begin{aligned} \text{Dec}_{ah}(sk_{ah}, \mathbf{v}_i - \mathbf{v}'_i \pmod{q}) &= \text{Dec}_{ah}(sk_{ah}, \mathbf{v}_i) - \text{Dec}_{ah}(sk_{ah}, \mathbf{v}'_i) \\ &= x_i - x'_i \pmod{q}. \end{aligned}$$

This is without knowing any other ciphertexts, i.e. goes beyond inherent leakage.

First Attack on LaPS using two Ciphertexts of a Client

Given two encryptions ct_i and ct'_i of x_i and x'_i and $k_0 = (sk_{ah}, k)$.

$$ct_i - ct'_i = k_i^\top \mathbf{A} + \mathbf{e}_i^\top - k_i^\top \mathbf{A} + \mathbf{e}'_i{}^\top \pmod{q} = \mathbf{e}_i^\top - \mathbf{e}'_i{}^\top \pmod{q}.$$

Then:

$$\mathbf{G}(\mathbf{e}_i^\top - \mathbf{e}'_i{}^\top) \pmod{q} = \mathbf{v}_i - \mathbf{v}'_i \pmod{q}.$$

We can then decrypt to get

$$\begin{aligned} \text{Dec}_{ah}(sk_{ah}, \mathbf{v}_i - \mathbf{v}'_i \pmod{q}) &= \text{Dec}_{ah}(sk_{ah}, \mathbf{v}_i) - \text{Dec}_{ah}(sk_{ah}, \mathbf{v}'_i) \\ &= x_i - x'_i \pmod{q}. \end{aligned}$$

This is without knowing any other ciphertexts, i.e. goes beyond inherent leakage.

A Problem with the Proof of LaPS: Removing QEnc

Shi et al. (p. 15):

Proof of Theorem 1: First, we will make a small modification to the aggregator oblivious security game. In the **Encrypt** queries, if the adversary submits a request for some tuple (i, x, t^*) where t^* is the time step specified in the **Challenge** phase, the challenger treats this as a **Compromise** query, and simply returns the sk_i to the adversary. Given sk_i , the adversary can compute the requested ciphertext herself. **Therefore, this modification actually gives more power to the adversary.** From now on, we will assume that the adversary does not make any **Encrypt** queries **for the time t^* .**

LaPS (Becker et al., p. 16):

We define the following intermediate game Game similar to [40] that is indistinguishable from the aggregator obliviousness security game according to Definition 6. **First, we treat any **Encrypt** query as a **Compromise** query from the adversary. Clearly, this turns the adversary actually more powerful.** Secondly, we change the **Challenge** phase to its real-

Problem Intuition: **Corruptions** lead to **stricter** conditions on what is allowed: corrupted users cannot be included in \mathcal{U} , but those from which we queried an encryption (for $\ell \neq \ell^*$ if encrypt-once) can be. Hence, LaPS is shown only to be “corruption-only-AO-secure”.

How Weak is corruption-only-AO Security?

The following “trivial” PSA scheme is (perfectly) corruption-only-AO secure:¹

<u>Setup($1^\lambda, 1^n$):</u> for $i \in [n]$: $k_i \leftarrow \mathbb{Z}_q$ $k_0 := -\sum_{i \in [n]} k_i$ $pp := q$ (the modulus) return $(pp, (k_i)_{i \in [n]_0})$	<u>Enc(pp, k_i, x_i, ℓ):</u> return $x_i + k_i$ <u>AggrDec($pp, k_0, \ell, (ct_i)_{i \in [n]}$):</u> return $\sum_{i \in [n]} ct_i + k_0$
--	---

Intuition: The adversary cannot get anyone to use their OTP key twice ($\mathcal{U} \cap \mathcal{CS} = \emptyset$)

Problem: In practice, a PSA setup is supposed to be used for more than one use per party. Hence, security guarantees are very weak.

Also: The scheme is not AO-secure.

¹This was already mentioned in a footnote in Waldner et al. (2021)

How Weak is corruption-only-AO Security?

The following “trivial” PSA scheme is (perfectly) corruption-only-AO secure:¹

<u>Setup($1^\lambda, 1^n$):</u> for $i \in [n]$: $k_i \leftarrow \mathbb{Z}_q$ $k_0 := -\sum_{i \in [n]} k_i$ $pp := q$ (the modulus) return $(pp, (k_i)_{i \in [n]_0})$	<u>Enc(pp, k_i, x_i, ℓ):</u> return $x_i + k_i$ <u>AggrDec($pp, k_0, \ell, (ct_i)_{i \in [n]}$):</u> return $\sum_{i \in [n]} ct_i + k_0$
--	---

Intuition: The adversary cannot get anyone to use their OTP key twice ($\mathcal{U} \cap \mathcal{CS} = \emptyset$)

Problem: In practice, a PSA setup is supposed to be used for more than one use per party. Hence, security guarantees are very weak.

Also: The scheme is not AO-secure.

Intuition: Using $\text{QEnc}(i^*, x = 0, \ell)$ you get $k_{i^*} = k_{i^*} + 0$ without corrupting.

¹This was already mentioned in a footnote in Waldner et al. (2021)

Key Recovery Attack on LaPS using Averaging

- $(\text{Gen}_{\text{ah}}, \text{Enc}_{\text{ah}}, \text{Dec}_{\text{ah}})$: additively homomorphic PKE with pseudorandom ciphertexts.
- \mathbf{G} : gadget matrix, i.e. $\mathbf{G} = \mathbf{I} \otimes (1 \ 2 \ \dots \ 2^{d-1})$, with $d = O(\lceil \log q \rceil)$
- $e \leftarrow \mathcal{D}_{\Lambda_{\mathbf{v}}^{\perp}(\mathbf{G}), \sigma}$: error distribution (with $\sigma > 2\sqrt{\lambda}$); defined so that $\mathbf{G} \cdot e \bmod q = \mathbf{v}$

Note: Each $\text{Enc}(pp, k_i, x_i)$ uses same \mathbf{A} and k_i .

Attack: Run QEnc many (poly(λ)) times to average out the error.

Setup($1^\lambda, 1^n$):

$\mathbf{A} \leftarrow \mathbb{Z}_q^{\lambda \times d}$

for $i \in [n]$: $k_i \leftarrow \mathbb{Z}_q^\lambda$

$(pk_{\text{ah}}, sk_{\text{ah}}) \leftarrow \text{Gen}_{\text{ah}}(1^\lambda)$

$pp := (\mathbf{A}, pk_{\text{ah}})$

$k_0 := (sk_{\text{ah}}, \sum_{i \in [n]} k_i)$

return $(pp, (k_i)_{i \in [n]_0})$

Enc(pp, k_i, x_i):

parse $pp = (\mathbf{A}, pk_{\text{ah}})$

$\mathbf{v}_i \leftarrow \text{Enc}_{\text{ah}}(pk_{\text{ah}}, x_i)$

$e_i \leftarrow \mathcal{D}_{\Lambda_{\mathbf{v}_i}^{\perp}(\mathbf{G}), \sigma}$

$ct_i = k_i^T \mathbf{A} + e_i^T \bmod q$

return ct_i

AggrDec($pp, k_0, (ct_i)_{i \in [n]}$):

parse $k_0 = (sk_{\text{ah}}, k)$

$e = \sum_{i \in [n]} c_i - k^T \mathbf{A}$

return $\text{Dec}_{\text{ah}}(sk_{\text{ah}}, \mathbf{G} \cdot e \bmod q)$

Another Proof Problem: Changing to Real-or-Random

The proof switches to a **Real-or-Random** (RoR) variant of the AO-game without really defining it. And, doing it naively leads to a potentially **weaker definition**.

Main Issue: How should the random x_i be chosen? In AO, the attacker can determine adaptively whether the balance condition needs to hold, here this is not possible.

Way out: Define RoR-AO by giving a choice b_{sum} in QChallenge on whether the randomness should add up to the same value as the provided x_i 's. We show this to be equivalent to AO.

LaPS (Becker et al., p. 16):

We define the following intermediate game Game similar to [40] that is indistinguishable from the aggregator obliviousness security game according to Definition [6]. First, we treat any **Encrypt** query as a **Compromise** query from the adversary. Clearly, this turns the adversary actually more powerful. **Secondly, we change the Challenge phase to its real-or-random version, i.e. instead of having the adversary specify two sets of plaintext-randomness pairs $\{(d_i, r_i)\}$ and $\{(d'_i, r'_i)\}$ and have her distinguish between encryptions of either one, we let the adversary pick one set $\{(d_i, r_i)\}$ and have her distinguish between a set of valid encryptions and a set of random values in \mathbb{Z}_q^λ .** It is straightforward that any adversary with more than negligible advantage in winning Game will also win the aggregator obliviousness security game with more than negligible advantage. Therefore, it suffices to show that

Summary and Discussion

Summary:

- 1 The LaPS scheme (NDSS 2018) was the first with claimed post-quantum and encrypt-multiple times security, and subsequently patented.
- 2 However, the proof contains at least two shaky/wrong steps
- 3 There are attacks that let you recover differences of their messages and the user's keys.
- 4 **Importantly:** First mention of problems was in a remark in Waldner et al. (ePrint 2021) and the proof mistakes have also spread to other papers.

Main Takeaway/Discussion:

- 1 Getting all the details of a security proof right is hard, but crucial.
- 2 Maybe: avoid putting the main proof in the appendix.

References

-  Elaine Shi, T.-H. Hubert Chan, Eleanor Gilbert Rieffel, Richard Chow, and Dawn Song. “Privacy-Preserving Aggregation of Time-Series Data”. In: NDSS 2011. The Internet Society, 2011. URL: <https://www.ndss-symposium.org/ndss2011/privacy-preserving-aggregation-of-time-series-data>.
-  Rachid El Bansarkhani, Özgür Dagdelen, and Johannes Buchmann. “Augmented learning with errors: The untapped potential of the error term”. In: FC 2015. Springer, 2015, pp. 333–352. DOI: 10.1007/978-3-662-47854-7_20.
-  Daniela Becker, Jorge Guajardo, and Karl-Heinz Zimmermann. “Revisiting Private Stream Aggregation: Lattice-Based PSA.”. In: NDSS 2018. The Internet Society, 2018. URL: https://www.ndss-symposium.org/wp-content/uploads/2018/02/ndss2018_02B-3_Becker_paper.pdf.
-  Hendrik Waldner, Tilen Marc, Miha Stopar, and Michel Abdalla. Private Stream Aggregation from Labeled Secret Sharing Schemes. Cryptology ePrint Archive, Report 2021/081.