


Closing the Gap: A Universal Privacy Framework for Outsourced Data

Dirk Achenbach, Matthias Huber, Jörn Müller-Quade, Jochen Rill | 3. September 2015

INSTITUT FÜR KRYPTOGRAPHIE UND SICHERHEIT





Imagine: Facebook wants to host its user database on Amazon's cloud service platform.

- The user database is huge and contains highly sensitive information, maybe even company secrets.
- Queries made to the outsourced database must be very efficient (therefore simple encryption is not an option).

We need to investigate the security (using formal models) of efficient solutions for data outsourcing.

A Model for Outsourced Data



\mathcal{S}	Server
\mathcal{C}	Client
\mathcal{Q}	Set of all allowed queries
$\pi \in \mathcal{Q}$	One query from the set
$\text{view}_{\mathcal{S}}^{\pi_j}(\text{Enc}(d, K))$	Server view for query π_j on encrypted database d
$(\text{Gen}, \text{Enc}, \text{Dec}, \mathcal{Q})$	A queryable outsourcing scheme

Existing Solutions and their Security:

PIR

Definition (Private Information Retrieval)

A queryable outsourcing scheme $(\text{Gen}, \text{Enc}, \text{Dec}, \{\pi\})$ exhibits *Computational Single-Server Private Information Retrieval* (PIR) when two conditions hold for any $n \in \mathbb{N}$, any security parameter $k \in \mathbb{N}$, and any data set d over $\Sigma = \{0, 1\}^n$:

- ① Correctness: $\forall i \in \{0, \dots, n-1\} : \pi_i^{\mathcal{C}}(d) = d[i]$.
- ② Privacy: $\forall c \in \mathbb{N}, i, j \in \{0, \dots, n-1\}, \forall \mathcal{A} \exists K \in \mathbb{N}$ such that $\forall k > K$

$$|\Pr[\mathcal{A}(\text{view}_{\mathcal{G}}^{\pi_i}(\text{Enc}(d, K))) = 1] - \Pr[\mathcal{A}(\text{view}_{\mathcal{G}}^{\pi_j}(\text{Enc}(d, K))) = 1]| < \frac{1}{\max(k, n)^c}.$$

Existing Solutions and their Security:

PIR

PIR is great, since it guarantees that queries are completely hidden . . .

Existing Solutions and their Security:

PIR

PIR is great, since it guarantees that queries are completely hidden ...

... but PIR can not be realised efficiently ...

... and gives no guarantees concerning the data.

Existing Solutions and their Security:

Goh's construction [3]

Security Game (IND-CKA_(Gen, Enc, Q)^A)

- 1 The experiment chooses a key $K \leftarrow \text{Gen}(1^k)$ and a random bit $b \leftarrow \{0, 1\}$.
- 2 The adversary \mathcal{A} is given input 1^k and access to an oracle $\text{view}_{\mathcal{G}}^{\pi_{\cdot}}(\text{Enc}(\cdot, K))$.
- 3 \mathcal{A} outputs two plaintexts m_0 and m_1 of the same length to the experiment. The adversary must not have queried the oracle for words that are only in one of the two plaintexts.
- 4 \mathcal{A} is given $\text{Enc}_K(m_b)$ and access to an oracle $\text{view}_{\mathcal{G}}^{\pi_{\cdot}}(\text{Enc}(m_b, K))$
- 5 \mathcal{A} submits a guess b' for b .

The result of the experiment is 1 if $b' = b$ and 0 else.

Existing Solutions and their Security:

Goh's construction [3]

IND-CKA restricts the adversary to allow for an efficient realisation . . .

Existing Solutions and their Security:

Goh's construction [3]

IND-CKA restricts the adversary to allow for an efficient realisation . . .

. . . but in a very specific way which may not be suited for other schemes . . .
. . . and it also gives no guarantees concerning queries.

There are multiple different and specifically tailored security notions for outsourcing schemes. They are hiding parts of the data, parts of the queries, or a combination of both.

This makes comparing different schemes nearly impossible!

... so we need a general framework to instantiate security notions from,
which fits more schemes.

What do we even want to hide?

- Queries: Query Privacy
- Results: Result Privacy
- The data itself: Data Privacy

Security Game ($\text{Q-IND}_{(\text{Gen}, \text{Enc}, \text{Q})}^{\mathcal{A}}(k)$)

- 1 *The experiment chooses a key $K \leftarrow \text{Gen}(1^k)$.*
- 2 *$\mathcal{A} \leftarrow \text{view}_{\mathcal{G}}^{\pi_{\cdot}}(\text{Enc}(\cdot, K))$.*
- 3 *$(q_0, q_1) \leftarrow \mathcal{A}$. q_0 and q_1 must yield protocols π_{q_0} and π_{q_1} with the same number of protocol messages.*
- 4 *The experiment draws a random bit $b \leftarrow \{0, 1\}$.*
- 5 *Challenge oracle: $\mathcal{A} \leftarrow \text{view}_{\mathcal{G}}^{\pi_{q_b}}(\text{Enc}(\cdot, K))$. The oracle takes any data set $d \in \Delta$ as input, internally runs the protocol π_{q_b} on $\text{Enc}(d, K)$, and outputs $\text{view}_{\mathcal{G}}^{\pi_{q_b}}(\text{Enc}(d, K))$ to the adversary.*
- 6 *\mathcal{A} outputs a guess b' for b .*

The result of the experiment is 1 if $b' = b$ and 0 else.

This is equivalent to PIR!

Security Game (Q-IND $^{\mathcal{A}, R_q, n_1, n_2, n_3}_{(\text{Gen}, \text{Enc}, \text{Q})}(k)$)

- ① The experiment chooses a key $K \leftarrow \text{Gen}(1^k)$.
- ② $\mathcal{A} \leftarrow \text{view}_{\mathcal{G}}^{\pi}(\text{Enc}(\cdot, K))$. \mathcal{A} is only allowed to query the oracle for a total number of n_1 times.
- ③ $(q_0, q_1) \leftarrow \mathcal{A}$. (q_0, q_1) is restricted with regard to equivalence relation $R_q \subseteq \Pi^2$, i.e. $(q_0, q_1) \in R_q$.
- ④ The experiment draws a random bit $b \leftarrow \{0, 1\}$.
- ⑤ Challenge oracle: $\mathcal{A} \leftarrow \text{view}_{\mathcal{G}}^{\pi_{q_b}}(\text{Enc}(\cdot, K))$. \mathcal{A} may call the challenge oracle for a total number of n_2 times.
- ⑥ Run oracle: $\mathcal{A} \leftarrow \text{run}_{\mathcal{G}}^{\pi_b}(\text{Enc}(\cdot, K))$. The run oracle executes queries just as the view oracle does, but has no output. \mathcal{A} is allowed to call the run oracle for a total number of n_3 times.
- ⑦ \mathcal{A} outputs a guess b' for b .

Security Game ($\text{D-IND}_{(\text{Gen}, \text{Enc}, \mathcal{Q})}^{\mathcal{A}}(k)$)

- ① *The experiment chooses a key $K \leftarrow \text{Gen}(1^k)$.*
- ② *$\mathcal{A} \leftarrow \text{view}_{\mathcal{E}}^{\pi_{\cdot}}(\text{Enc}(\cdot, K))$. The oracle takes a query q and a data set d as input and returns $\text{view}_{\mathcal{E}}^{\pi_q}(\text{Enc}(d, K))$.*
- ③ *$(d_0, d_1) \leftarrow \mathcal{A}$ (both of equal length).*
- ④ *The experiment draws a random bit $b \leftarrow \{0, 1\}$.*
- ⑤ *Challenge oracle: $\mathcal{A} \leftarrow \text{view}_{\mathcal{E}}^{\pi_{\cdot}}(\text{Enc}(d_b, K))$. That is, the oracle takes any query q such that $\pi_q \in \mathcal{Q}$ as input, internally runs the protocol π_q on $\text{Enc}(d_b, K)$, and outputs $\text{view}_{\mathcal{E}}^{\pi_q}(\text{Enc}(d_b, K))$ to the adversary.*
- ⑥ *\mathcal{A} outputs a guess b' for b .*

The result of the experiment is 1 if $b' = b$ and 0 else.

Security Game (D-IND $^{\mathcal{A}, R_d, n_1, n_2, n_3}_{(\text{Gen}, \text{Enc}, \text{Q})}(k)$)

- ① The experiment chooses a key $K \leftarrow \text{Gen}(1^k)$.
- ② $\mathcal{A} \leftarrow \text{view}_{\mathcal{G}}^{\pi}(\text{Enc}(\cdot, K))$. \mathcal{A} is only allowed to query the oracle for a total number of n_1 times.
- ③ $(d_0, d_1) \leftarrow \mathcal{A}$. (d_0, d_1) is restricted with regard to equivalence relation $R_d \subseteq \Delta^2$, i.e. $(d_0, d_1) \in R_d$
- ④ The experiment draws a random bit $b \leftarrow \{0, 1\}$.
- ⑤ Challenge oracle: $\mathcal{A} \leftarrow \text{view}_{\mathcal{G}}^{\pi}(\text{Enc}(d_b, K))$. \mathcal{A} may call the challenge oracle for a total number of n_2 times.
- ⑥ Run oracle: $\mathcal{A} \leftarrow \text{run}_{\mathcal{G}}^{\pi}(\text{Enc}(d_b, K))$. The run oracle executes queries just as the view oracle does, but has no output. \mathcal{A} is allowed to call the run oracle for a total number of n_3 times.
- ⑦ \mathcal{A} outputs a guess b' for b .

- $Q\text{-IND} \not\iff D\text{-IND}$
- $D\text{-IND} \wedge Q\text{-IND} \iff R\text{-IND}$

- $\text{PIR (Chor et. al [1])} \iff \text{Q-IND}$
- $\text{Ind}_{\text{SSE}} \text{ (Curtmola et al. [2])} \implies \text{D-IND}$
- $\text{D-IND} \wedge \text{Q-IND} \implies \text{Ind}_{\text{SSE}}$
- $\text{D-IND} \implies \text{IND-CKA}$
- $\text{PrivK}_{\mathcal{A}, \text{Enc}}^{\text{cpa}}(k) \text{ (Haynberg et al. [4])} \iff \text{D-IND}$
- $\text{IND-ICP (Huber et al. [5])} \iff \text{Static Security}$

- We strictly separated and formalised security goals for data outsourcing.
- We generalised our formalisation to capture mechanisms with weaker security properties.
- Our framework is applicable for all outsourcing schemes (i.e. searchable encryption or database outsourcing).
- Our framework makes the security of different schemes compareable.




- Investigate simulation based formulations and their relation to game based notions.
- Investigate relations among the generalised notions.
- Investigate limitations of “pure” Data Privacy for practical outsourcing schemes.



Benny Chor u. a. “Private Information Retrieval”. In: *J. ACM* 45.6 (Nov. 1998), S. 965–981. ISSN: 0004-5411. DOI: 10.1145/293347.293350. URL: <http://doi.acm.org/10.1145/293347.293350>.



Reza Curtmola u. a. “Searchable Symmetric Encryption: Improved Definitions and Efficient Constructions”. In: *Proceedings of the 13th ACM Conference on Computer and Communications Security. CCS '06*. Full version available at <https://eprint.iacr.org/2006/210>. Alexandria, Virginia, USA: ACM, 2006, S. 79–88. ISBN: 1-59593-518-5. DOI: 10.1145/1180405.1180417.

-  Eu-Jin Goh u. a. “Secure Indexes.” In: *IACR Cryptology ePrint Archive* 2003 (2003). <https://eprint.iacr.org/2003/216/>, S. 216.
-  Rolf Haynberg u. a. “Symmetric Searchable Encryption for Exact Pattern Matching using Directed Acyclic Word Graphs.” In: *SECRYPT*. 2013, S. 403–410.
-  Matthias Huber u. a. “Cumulus4j: A Provably Secure Database Abstraction Layer”. In: *CD-ARES Workshops*. 2013, S. 180–193.