


Oblivious Transfer

Sven Laur
swen@math.ut.ee

University of Tartu

Ideal implementation



The protocol is always carried out between a client \mathcal{P}_1 and a sender \mathcal{P}_2 .

- ▷ The server \mathcal{P}_2 has a database of two elements $x_0, x_1 \in \mathcal{M}$.
- ▷ The client \mathcal{P}_1 can fetch either x_0 or x_1 so that the server \mathcal{P}_2 cannot detect which element is fetched.
- ▷ The client should not learn anything more than x_b . Moreover, the client should be always aware of his or her choice b .

How to handle large databases?

Theorem. 1-out-of- 2^ℓ oblivious transfer protocol for k -bit strings can be implemented using 1-out-of-2 oblivious transfer protocol for $2^\ell \cdot k$ -bit strings.

Simplified proof

To encode x_{00}, \dots, x_{11} , generate uniformly matrices Y and Z such that

$$\begin{array}{|c|c|} \hline x_{00} & x_{01} \\ \hline x_{10} & x_{11} \\ \hline \end{array} = \begin{array}{|c|c|} \hline y_{00} & y_{01} \\ \hline y_{10} & y_{11} \\ \hline \end{array} \oplus \begin{array}{|c|c|} \hline z_{00} & z_{01} \\ \hline z_{10} & z_{11} \\ \hline \end{array}$$

Next the client uses 1-out-of-2 oblivious transfer twice.

- ▷ First, the client must fetch the correct column of Y .
- ▷ Second, the client must fetch the correct row of Z .

Even a malicious client can learn only a single entry x_{ab} and he or she must be aware of the location ab .

Solution to the millionaires problem

Let $w_1, w_2 \in \{1, \dots, n\}$ be the total wealth of two millionaires. Then one of them can find out who is richer and nothing more with the help of oblivious transfer protocol. The construction was first published by Yao (1982).

- ▷ The first millionaire creates an n -element table of possible answers

1	2	...	n
$w_1 > 1$	$w_1 > 2$...	$w_1 > n$

- ▷ The second millionaire fetches the w_2 th entry from the table and thus learns the value $w_1 > w_2$.
- ▷ The protocol is secure only if the first millionaire behaves semi-honestly.

This construction can be generalised for all functions with small input range.

Multiplication \Leftrightarrow Oblivious transfer

Theorem. Given an ideal multiplication protocol, we can implement 1-out-of-2 oblivious transfer. Given an ideal 1-out-of-2 oblivious transfer protocol we can implement multiplication over \mathbb{Z}_2 in the semihonest model.

Clarification

- ▷ Observe that $x_b = (1 - b)x_0 + bx_1$ and thus any multiplication protocol that provides shares is sufficient to implement oblivious transfer.
- ▷ Oblivious transfer is sufficient to implement multiplication, since the sender can use columns of the multiplication table as the input.

Kilian proved in 1988 that zero-knowledge proofs and commitments can be constructed using only oblivious transfer protocol. Hence, we can use commitment and zero knowledge proofs to eliminate malicious behaviour.

Homomorphic Oblivious Transfer

Homomorphic encryption

A public key cryptosystem $(\text{Gen}, \text{Enc}, \text{Dec})$ is an additively homomorphic cryptosystem if for any two message $m_1, m_2 \in \mathcal{M}$ the distributions

$$\text{Enc}_{\text{pk}}(m_1) \cdot \text{Enc}_{\text{pk}}(m_2) \equiv \text{Enc}_{\text{pk}}(m_1 + m_2)$$

coincide even if we fix a ciphertext $\text{Enc}_{\text{pk}}(m_1)$.

Multiplying a ciphertext $\text{Enc}_{\text{pk}}(m)$ with a newly generated $\text{Enc}_{\text{pk}}(0)$ completely destroys all extra information besides the value m .

We can compute also crypto-compute multiplication

$$\text{Enc}_{\text{pk}}(m_1)^{m_2} \cdot \text{Enc}_{\text{pk}}(0) \equiv \text{Enc}_{\text{pk}}(m_1 \cdot m_2) .$$

Famous examples

The Goldwasser-Micali cryptosystem is additively homomorphic over \mathbb{Z}_2 .

The lifted ElGamal cryptosystem is additively homomorphic over \mathbb{Z}_p

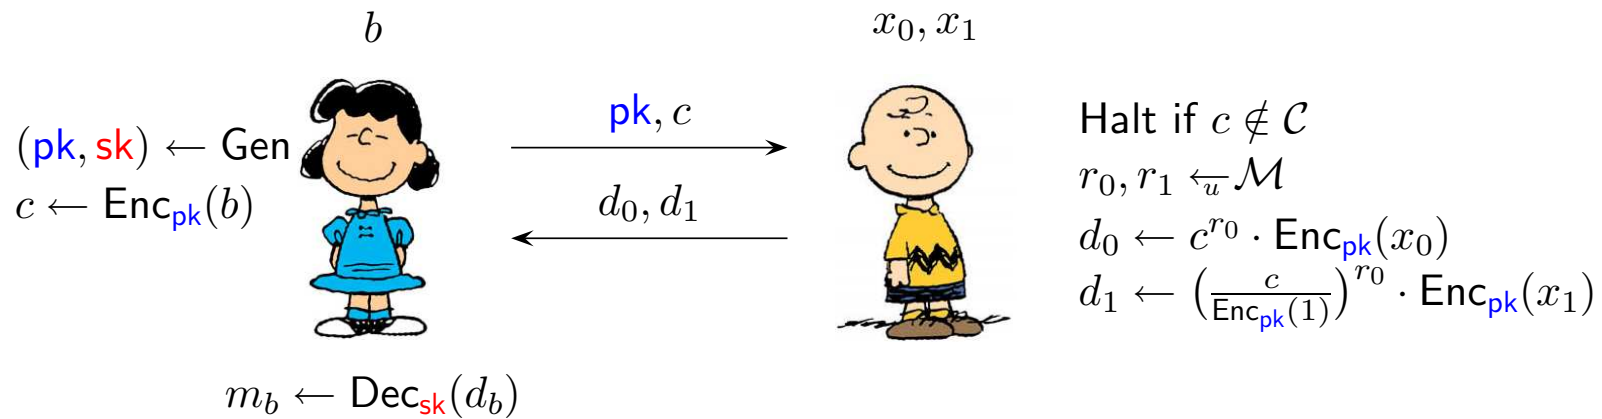
$$\overline{\text{Enc}}_{\text{pk}}(m) = \text{Enc}_{\text{pk}}(g^m) = (g^r, g^m y^r)$$

$$\overline{\text{Dec}}_{\text{sk}}(c_1, c_2) = \log_g[\text{Dec}(c_1, c_2)] = \log_g \left[\frac{c_2}{c_1^x} \right]$$

For obvious reason, the decryption rule $\overline{\text{Dec}}_{\text{sk}}(\cdot)$ can be efficiently computed for few ciphertexts or otherwise the cryptosystem would not be secure.

The Paillier cryptosystem uses lifting with together with a trapdoor that allows us to efficiently compute discrete logarithms. The corresponding message space is \mathbb{Z}_n where n is RSA modulus.

Aiello-Ishai-Reingold oblivious transfer



If $(\text{Gen}, \text{Enc}, \text{Dec})$ be an additively homomorphic cryptosystem then

$$d_0 \equiv \text{Enc}_{pk}(b)^{r_0} \cdot \text{Enc}_{pk}(x_0) \equiv \text{Enc}_{pk}(x_0 + br_0) ,$$

$$d_1 \equiv \text{Enc}_{pk}(b - 1)^{r_1} \cdot \text{Enc}_{pk}(x_1) \equiv \text{Enc}_{pk}(x_1 + (b - 1)r_1) .$$

If the message space has prime order then br_0 has uniform distribution if $b \neq 0$ and $(b - 1)r_1$ has uniform distribution if $b \neq 1$.

Security in the semi-honest model

Lemma 1. If the cryptosystem is additively homomorphic over \mathbb{Z}_p , then for any t -time **semi-honestly** corrupted receiver \mathcal{P}_1^* there exists $t + O(1)$ -time ideal world adversary \mathcal{P}_1° such that the joint output distributions are identical in the real and ideal world.

Proof

- ▷ For fixed value b , the messages received by \mathcal{P}_1^* have the following distribution: $d_b = \text{Enc}_{pk}(x_b)$ and $d_{b-1} = \text{Enc}_{pk}(m)$ where $m \leftarrow \mathcal{M}$.
- ▷ Given x_b from the trusted third party, we can perfectly simulate the reply in the real world.
- ▷ Since the output of \mathcal{P}_2 is \perp in both worlds the joint output distribution coincides in both worlds.

□

Security in the semi-honest model

Lemma 2. If the cryptosystem is (t, ε) -IND-CPA secure, then for any τ -time **semi-honestly** corrupted sender \mathcal{P}_2^* there exists $(\tau + O(1))$ -time ideal world adversary \mathcal{P}_2° such that the joint output distributions in the real and ideal world are $(t - \tau, \varepsilon)$ -indistinguishable.

Proof

- ▷ The sender receives an encryption of b that we cannot simulate, since the trusted third party sends only \emptyset to \mathcal{P}_2° .
- ▷ However, we can replace c with $\text{Enc}_{pk}(0)$. Since \mathcal{P}_1 outputs m_b in both worlds then the output distributions must be $(t - \tau, \varepsilon)$ -indistinguishable.
- ▷ Otherwise, we can construct a new adversary \mathcal{A} from the participant \mathcal{P}_2° and the output distinguisher \mathcal{B} that wins the IND-CPA game.

□

Interpretation of the results

Semi-honest receiver can carry out only the attacks that are possible against ideal implementation. The only benefit the receiver may gain in the real world is a marginal $O(1)$ speed-up compared to the ideal world.

Let us consider a specific security goal. Then any of those can be formalised as a predicate $\mathcal{B}(\cdot)$ that indicates whether \mathcal{P}_2^* was successful or not.

Lemma 2 indicates that is we consider specific $(t - \tau)$ -time security goals $\mathcal{B}(\cdot)$, then for any τ -time semi-honest sender \mathcal{P}_2^*

$$\Pr [\mathcal{P}_2^* \text{ wins}] \leq \Pr [\mathcal{P}_2^\circ \text{ wins}] + \varepsilon$$

where \mathcal{P}_2° is $(\tau + O(1))$ -time adversary. In other words, \mathcal{P}_2^* can achieve only marginal increase ε in success and a marginal $O(1)$ speed-up.

Security against malicious receivers

Lemma 3. If the cryptosystem is additively homomorphic over \mathbb{Z}_p and validity of the public key can be tested, then for any t -time maliciously corrupted receiver \mathcal{P}_1^* there exists unbounded ideal world adversary \mathcal{P}_1° such that the joint output distributions are identical in the real and ideal world.

Proof

- ▷ Given a valid public key pk we can always find the corresponding secret key by looking through all valid (pk, sk) pairs.
- ▷ Hence, we can decrypt c and find out the true input of \mathcal{P}_1^* .
- ▷ If $b \notin \{0, 1\}$ then the received messages d_0 and d_1 are both random encryptions. Thus we can always perfectly simulate the replies.
- ▷ Other steps in the proof are analogous.

Interpretation of the results

Lemma 3 indicates that for each real world attack there is a matching ideal world attack. Hence, the adversary can learn nothing that cannot be computed from the intended output m_b .

However, the participation in the real world protocol might give a huge computational speedup compared to the ideal world.

Hence, participation in the protocol might help \mathcal{P}_1^* to compute intractable functions from m_b . For example, if m_b is an encryption, then the protocol might reveal the underlying message.

How to achieve tight security guarantees?

If the receiver proves in zero-knowledge that he knows the secret key sk that corresponds to pk then the possible speedup becomes marginal.

- ▷ We can extract secret key by rewinding $ZKPOK_{sk} [(sk, pk) \in \text{Gen}]$.
- ▷ The simulation becomes efficient if we learn the secret key sk .

If the sender is assumed to be semi-honest and the protocol uses the ElGamal encryption, then we can use the Schnorr protocol.

To handle malicious senders, we must convert the corresponding sigma protocol $POK_x [g^x = y]$ to zero-knowledge proof of knowledge.

$$ZKPOK_x [g^x = y] = POK_x [g^x = y] + \text{coin flipping protocol}$$

Fortunately, we can reuse the same key in many protocol instances.

Security against malicious senders

To handle a malicious sender, we must extract x_0 and x_1 from \mathcal{P}_2^* .

- ▷ We can add zero-knowledge proofs of knowledge

$\text{ZKPOK}_{x_0, x_1} [d_0(x_0, x_1) \text{ and } d_1(x_0, x_1) \text{ are correctly formed}]$

and then we can construct the necessary simulator.

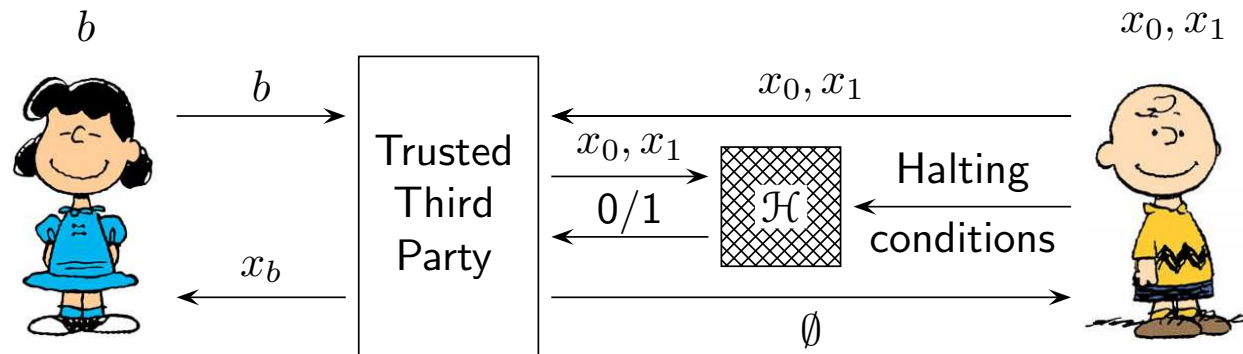
- ▷ One possibility is to commit x_0 and x_1 and then execute

$\text{ZKPOK}_{x_0, x_1} [\text{commitments are properly formed}]$

and then continue with the certified computation protocol.

As a result, we get a protocol with enormous computational overhead.

Output consistency



If the sender first commits pairs $(0, x_0)$ and $(1, x_1)$ and the oblivious transfer protocol is used to reveal the corresponding decommitment strings, then the malicious sender cannot alter the outputs without getting caught.

- ⇒ The sender can still cause selective halting.
- ⇒ Cheating behaviour is detectable with high probability.
- ⇒ Public complaints reveal information about receiver inputs.

Complete security vs output consistency

Both security levels reveal cheating with high probability:

- ▷ Complete security makes all deviations from the protocol that **could alter** the outcome for some receiver input detectable.
- ▷ Output consistency makes all deviations from the protocol that **alter** the output for this particular receiver input detectable.

Complete security has **large** computational overhead

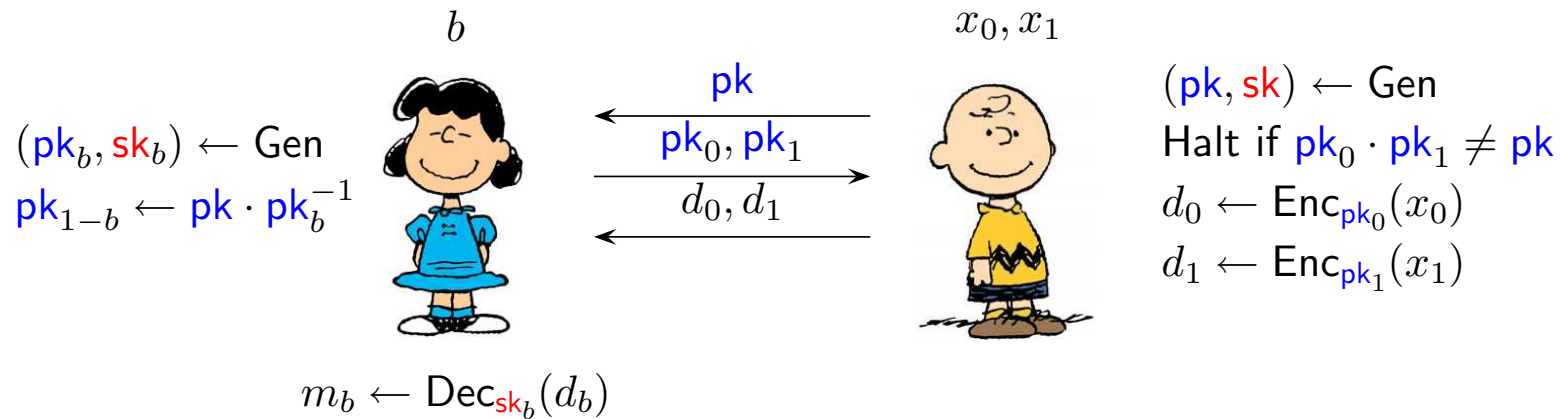
- ▷ Certified computations require extensive amount of extra steps

Output-consistent computations have **moderate** computational overhead

- ▷ Commitments are relatively easy to compute.
- ▷ Selective halting can cause privacy issues.

Bellare-Micali Protocol

Vanilla protocol



The protocol works under the assumption that all possible public keys form a group and the distribution of public keys is uniform.

- ▷ The ElGamal cryptosystem has such public key space.
- ▷ Since the public key pk_{1-b} is with correct distribution the corresponding ciphertext d_{1-b} is undecipherable.
- ▷ The protocol can tolerate unbounded senders.

Security in the semi-honest model

Lemma. If the public keys are uniformly distributed over some group, then for any t -time **semi-honestly** corrupted sender \mathcal{P}_2^* there exists $t + O(1)$ -time ideal world adversary \mathcal{P}_2° such that the joint output distributions are identical in the real and ideal world.

Proof

- ▷ In the simulator, we can first compute public keys $(pk_0, sk_0) \leftarrow \text{Gen}$ and $(pk_1, sk_1) \leftarrow \text{Gen}$ and then set the final key $pk \leftarrow pk_0 \cdot pk_1$.
- ▷ As a result, we can decrypt d_0 and d_1 and send the corresponding messages x_0 and x_1 to trusted third party.
- ▷ The simulation is perfect.

Security in the semi-honest model

Lemma. If the public keys are uniformly distributed over some group and the cryptosystem is (t, ε) -IND-CPA secure, then for any τ -time **semi-honestly** corrupted receiver \mathcal{P}_1^* there exists $\tau + O(1)$ -time ideal world adversary \mathcal{P}_1° such that the joint output distributions in the real and ideal world are $(t - \tau, \varepsilon)$ -indistinguishable.

Proof

- ▷ We can simulate the reply d_b with $\text{Enc}_{\text{pk}_b}(x_b)$ and $d_{1-b} \leftarrow \text{Enc}_{\text{pk}_{1-b}}(0)$.
- ▷ As pk_{1-b} can be taken from the IND-CPA game and the message pk defined as $\text{pk} \leftarrow \text{pk}_0 \cdot \text{pk}_1$, any distinguisher \mathcal{B} together with \mathcal{P}_1^* form a successful IND-CPA adversary.

□

How to strengthen the protocol?

Security against malicious receivers:

- ▷ Sigma protocol that proves that \mathcal{P}_1 knows one secret key sk_b is sufficient.
- ▷ If the protocol uses the ElGamal encryption, then we can use Schnorr protocol to prove $\text{POK}_{sk} [g^{sk} = pk_0 \vee g^{sk} = pk_1]$.

Security against malicious senders:

- ▷ If we do not care about efficiency, random generation of tuples (α, β, γ) until β coincides with the reply of \mathcal{P}_2^* provides a perfect simulation.
- ▷ Alternatively, we can use a sigma protocol with a high knowledge error and many rounds to get efficient simulator for \mathcal{P}_2^* .
- ▷ If we use the protocol as a sub-task in more complex protocol, then we must use certified computations to guarantee that inputs x_0 and x_1 are consistent with previous computations.