

Supersonic OT: Fast Unconditionally Secure Oblivious Transfer

Aydin Abadi^{*1} and Yvo Desmedt^{**2}

¹ Newcastle University

² The University of Texas at Dallas

Abstract. Oblivious Transfer (OT) is a fundamental cryptographic protocol with applications in secure Multi-Party Computation, Federated Learning, and Private Set Intersection. With the advent of quantum computing, it is crucial to develop unconditionally secure core primitives like OT to ensure their continued security in the post-quantum era. Despite over four decades since OT’s introduction, the literature has predominantly relied on computational assumptions, except in cases using unconventional methods like noisy channels or a fully trusted party. Introducing “Supersonic OT”, a highly efficient and unconditionally secure OT scheme that avoids public-key-based primitives, we offer an alternative to traditional approaches. Supersonic OT enables a receiver to obtain a response of size $O(1)$. Its simple (yet non-trivial) design facilitates easy security analysis and implementation. The protocol employs a basic secret-sharing scheme, controlled swaps, the one-time pad, and a third-party helper who may be corrupted by a semi-honest adversary. Our implementation and runtime analysis indicate that a single instance of Supersonic OT completes in 0.35 milliseconds, making it up to 2000 times faster than the state-of-the-art base OT.

1 Introduction

Oblivious Transfer (OT) is a vital cryptographic protocol that enables a user (called a receiver) interested in the s -th element of a database (m_0, m_1) (held by a sender) to learn only m_s while preserving the privacy of (i) index $s \in \{0, 1\}$ from the sender and (ii) the rest of the database’s elements from the receiver.

OT has found applications within numerous domains, such as generic secure Multi-Party Computation [46,3,24], Private Set Intersection [16], contract signing [18], Federated Learning [45,39,43], Zero-Knowledge proof systems [23], and accessing sensitive field elements of remote private databases while preserving privacy [9,2,31].

As quantum computing emerges, it is crucial to prioritize the development of fundamental security primitives, like OT, with unconditional security. This ensures their sustained resilience in the post-quantum era. However, despite

* aydin.abadi@newcastle.ac.uk

** y.desmedt@cs.ucl.ac.uk

over four decades having passed since Rabin introduced OT in 1981 [38], the security of OTs has predominantly leaned on computational assumptions that might not withstand the power of quantum computers. Efforts have been made to develop *unconditionally secure* OTs, ensuring security even against adversaries armed with quantum computers. Examples include the schemes proposed in [34,8,13,14,15,27,40].

However, the state-of-the-art unconditionally secure OTs either (i) depend on the multi-sender setting, where each sender possesses a database replica, (ii) utilize a specific communication channel (i.e., noisy channel), or (iii) require the presence of a fully trusted initializer. Nevertheless, distributing the same database across multiple servers, establishing a highly specific communication channel, or involving a fully trusted party would increase the overall deployment cost of these schemes.

In this work, we propose Supersonic OT, a novel unconditionally secure highly efficient 1-out-of-2 OT that does not need to rely on (i) multiple senders, (ii) noisy channel, or (iii) the involvement of a trusted initializer. Supersonic OT does not involve any public-key cryptography and allows the receiver to obtain a response of size $O(1)$. The scheme relies on four main components: a basic secret-sharing scheme [5], a controlled swap [19], the one-time pad [29], and a third-party helper. This helper could potentially be corrupted by a semi-honest adversary aiming to obtain the private information of the sender or receiver.

The design of Supersonic OT is simple (but elegant), facilitating straightforward security analysis and implementation. Indeed, we have implemented it, made its source code publicly available in [1], and evaluated its overhead. A single execution of Supersonic OT completes in approximately 0.35 milliseconds, making it about 1000 times faster than the base OT in [3] and up to 2000 times faster than the base OT in [35].

2 Preliminaries

2.1 Notations

By ϵ we mean an empty string. We denote a sender by S and a receiver by R . We assume parties interact with each other through a regular secure channel. U denotes a universe of messages m_1, \dots, m_t . We define σ as the maximum size of messages in U , i.e., $\sigma = \text{Max}(|m_1|, \dots, |m_t|)$.

2.2 Security Model

In this paper, we rely on the simulation-based model of secure multi-party computation [21] to define and prove the proposed protocols. Below, we restate the formal security definition within this model.

Two-party Computation. A two-party protocol Γ problem is captured by specifying a random process that maps pairs of inputs to pairs of outputs, one

for each party. Such process is referred to as a functionality denoted by $f : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}^* \times \{0, 1\}^*$, where $f := (f_1, f_2)$. For every input pair (x, y) , the output pair is a random variable $(f_1(x, y), f_2(x, y))$, such that the party with input x wishes to obtain $f_1(x, y)$ while the party with input y wishes to receive $f_2(x, y)$. In the setting where f is asymmetric and only one party (say the first one) receives the result, f is defined as $f := (f_1(x, y), \epsilon)$.

Security in the Presence of Passive Adversaries. In the passive adversarial model, the party corrupted by such an adversary correctly follows the protocol specification. Nonetheless, the adversary obtains the internal state of the corrupted party, including the transcript of all the messages received, and tries to use this to learn information that should remain private. Loosely speaking, a protocol is secure if whatever can be computed by a party in the protocol can be computed using its input and output only. In the simulation-based model, it is required that a party's view in a protocol's execution can be simulated given only its input and output.

This implies that the parties learn nothing from the protocol's execution. More formally, party i 's view (during the execution of Γ) on input pair (x, y) is denoted by $\text{View}_i^\Gamma(x, y)$ and equals $(w, r_i, m_1^i, \dots, m_i^i)$, where $w \in \{x, y\}$ is the input of i^{th} party, r_i is the outcome of this party's internal random coin tosses, and m_j^i represents the j^{th} message this party receives. The output of the i^{th} party during the execution of Γ on (x, y) is denoted by $\text{Output}_i^\Gamma(x, y)$ and can be generated from its own view of the execution.

Definition 1. *Let f be the deterministic functionality defined above. Protocol Γ securely computes f in the presence of a passive adversary if there exist polynomial-time algorithms $(\text{Sim}_1, \text{Sim}_2)$ such that:*

$$\begin{aligned} \{\text{Sim}_1(x, f_1(x, y))\}_{x, y} &\stackrel{c}{=} \{\text{View}_1^\Gamma(x, y)\}_{x, y} \\ \{\text{Sim}_2(y, f_2(x, y))\}_{x, y} &\stackrel{c}{=} \{\text{View}_2^\Gamma(x, y)\}_{x, y} \end{aligned}$$

2.3 Controlled Swap

The idea of the controlled swap was introduced by Fredkin and Toffoli [19]. It can be defined as function $\bar{\pi}(\cdot)$ which takes two inputs: a binary value s and a pair (c_0, c_1) . When $s = 0$, it returns the input pair (c_0, c_1) , i.e., it does not swap the elements. When $s = 1$, it returns (c_1, c_0) , effectively swapping the elements. It is evident that if s is uniformly chosen at random, then $\bar{\pi}(\cdot)$ represents a random permutation, implying that the probability of swapping or not swapping is $\frac{1}{2}$. We will use $\bar{\pi}(\cdot)$ in the protocol presented in Figure 1.

2.4 Secret Sharing

A (threshold) secret sharing $\text{SS}^{(t, n)}$ scheme is a cryptographic protocol that enables a dealer to distribute a string s , known as the secret, among n parties in

a way that the secret s can be recovered when at least a predefined number of shares, say t , are combined. If the number of shares in any subset is less than t , the secret remains unrecoverable and the shares divulge no information about s . This type of scheme is referred to as (n, t) -secret sharing or $\text{SS}^{(t,n)}$ for brevity.

In the case where $t = n$, there exists a highly efficient XOR-based secret sharing [5]. In this case, to share the secret s , the dealer first picks $n - 1$ random bit strings r_1, \dots, r_{n-1} of the same length as the secret. Then, it computes $r_n = r_1 \oplus \dots \oplus r_{n-1} \oplus s$. It considers each $r_i \in \{r_1, \dots, r_n\}$ as a share of the secret. To reconstruct the secret, one can easily compute $r_1 \oplus \dots \oplus r_n$. Any subset of less than n shares reveals no information about the secret. We will use this scheme in this paper. A secret sharing scheme involves two main algorithms; namely, $\text{SS}(1^\lambda, s, n, t) \rightarrow (r_1, \dots, r_n)$: to share a secret and $\text{RE}(r_1, \dots, r_t, n, t) \rightarrow s$ to reconstruct the secret.

3 Related Work

Oblivious Transfer (OT) is one of the important building blocks of cryptographic protocols and has been used in various mechanisms. The traditional 1-out-of-2 OT (\mathcal{OT}_1^2) is a protocol that involves two parties, a sender S and a receiver R . S has a pair of input messages (m_0, m_1) and R has an index s . The aim of \mathcal{OT}_1^2 is to enable R to obtain m_s , without revealing anything about s to S , and without allowing R to learn anything about m_{1-s} . The traditional \mathcal{OT}_1^2 functionality is defined as $\mathcal{F}_{\mathcal{OT}_1^2} : ((m_0, m_1), s) \rightarrow (\epsilon, m_s)$.

The notion of 1-out-of-2 OT was initially proposed by Rabin [38] which consequently was generalized by Even *et al.* [18]. Since then, numerous variants of OT have been proposed. For instance,

- 1-out-of- n OT, e.g., in [33,42,32]: which allows R to pick one entry out of n entries held by S ,
- k -out-of- n OT, e.g., in [12,28,11]: which allows R to select k entries out of n entries held by S ,
- OT extension, e.g., in [26,25,36,3]: that supports efficient executions of OT (that mainly relies on symmetric-key operations), in the case OT needs to be invoked many times
- distributed OT, e.g., in [34,13,48]: that allows the database to be distributed among m servers/senders.

In the remainder of this section, we discuss those variants of OT that are closer to our work. We refer readers to [44] for a recent survey of OT.

3.1 Unconditionally and Post-Quantum Secure OTs

There have been efforts to design (both-sided) unconditionally secure OTs. Some schemes, such as those proposed in [34,8,13], rely on multiple servers/senders that maintain an identical copy of the database. Other ones, like the one introduced

in [14,15,27], are based on a specific network structure, i.e., a noisy channel, to achieve unconditionally secure OT. There is also a scheme in [40] that achieves unconditionally secure OT using a fully trusted initializer.

Moreover, there exist OT schemes developed to maintain security in the presence of adversaries equipped with quantum computers. Examples include those proposed in [7,6,37,30,17,4]. However, these schemes are not unconditionally secure. Instead, they rely on various assumptions and problems (such as short integer solution, learning with errors, multivariate quadratic, decoding random linear codes, or computing isogenies between supersingular elliptic curves) as well as primitives (such as AES, pseudorandom generator, lattice-based Chameleon hash function, multivariate quadratic cryptography, McEliece cryptosystem, or supersingular isogeny Diffie-Hellman key exchange) that are deemed valid and secure in the era of quantum computing based on current knowledge and assessment. Their security could be compromised if any of the underlying assumptions or problems are proven to be solvable efficiently by future advancements in quantum algorithms or other unforeseen developments.

Hence, there exists no (efficient) unconditionally secure OT that does not use noisy channels, multi-server, and fully trusted initializer.

3.2 OT with Constant Response Size

Researchers have proposed several OTs, e.g., those proposed in [10,22,47], that enable a receiver to obtain a constant-size response to its query. To achieve this level of communication efficiency, these protocols require the receiver to locally store the encryption of the entire database, in the initialization phase. During the transfer phase, the sender assists the receiver with locally decrypting the message that the receiver is interested in.

The main limitation of these protocols is that a thin client with limited available storage space cannot use them, as it cannot locally store the encryption of the entire database.

4 Supersonic OT

In this section, we introduce a 1-out-of-2 OT, called “Supersonic OT”, which (i) operates at high speed by eliminating the need for public-key-based cryptography, (ii) delivers a response of size $O(1)$ to the recipient, R , and (iii) ensures information-theoretic security, making it post-quantum secure.

4.1 Security Definition

Supersonic OT involves three types of entities, a sender S , a receiver R , and a helper P . We assume each party can be corrupted by a passive non-colluding adversary. The functionality $\mathcal{F}_{\text{OT}_1^2}$ that Supersonic OT will compute is similar to that of conventional OT with the difference that now an additional party P is introduced, having no input and receiving no output. Thus, we define the

functionality as $\mathcal{F}_{\mathcal{OT}_1^2} : ((m_0, m_1), \epsilon, s) \rightarrow (\epsilon, \epsilon, m_s)$. Next, we present a formal definition of \mathcal{OT}_1^2 .

Definition 2 (\mathcal{OT}_1^2). *Let $\mathcal{F}_{\mathcal{OT}_1^2}$ be the OT functionality defined above. We assert that protocol Γ realizes $\mathcal{F}_{\mathcal{OT}_1^2}$ in the presence of a passive adversary, if for every non-uniform PPT adversary \mathcal{A} in the real model, there is a non-uniform PPT simulator Sim in the ideal model, where:*

$$\left\{ \text{Sim}_S((m_0, m_1), \epsilon) \right\}_{m_0, m_1, s} \stackrel{c}{\equiv} \left\{ \text{View}_S^\Gamma((m_0, m_1), \epsilon, s) \right\}_{m_0, m_1, s} \quad (1)$$

$$\left\{ \text{Sim}_P(\epsilon, \epsilon) \right\}_{m_0, m_1, s} \stackrel{c}{\equiv} \left\{ \text{View}_P^\Gamma((m_0, m_1), \epsilon, s) \right\}_{m_0, m_1, s} \quad (2)$$

$$\left\{ \text{Sim}_R\left(s, \mathcal{F}_{\mathcal{OT}_1^2}((m_0, m_1), \epsilon, s)\right) \right\}_{m_0, m_1, s} \stackrel{c}{\equiv} \left\{ \text{View}_R^\Gamma((m_0, m_1), \epsilon, s) \right\}_{m_0, m_1, s} \quad (3)$$

4.2 The Protocol

At a high level, the protocol operates as follows. Initially, R and S agree on a pair of keys. In the query generation phase, R splits its private index into two binary shares. It sends one share to S and the other to P . Given the share/query, S encrypts every message m_i (using a one-time pad) under one of the keys it agreed with R .

Accordingly, S permutes the encrypted messages using $\bar{\pi}$ and its share. It sends the resulting pair to P which permutes the received pair using $\bar{\pi}$ and its share. P sends only the first element of the resulting pair (which is a ciphertext) to R and discards the second element of the pair. Consequently, R decrypts the ciphertext and learns the message it was interested in. Figure 1 presents Supersonic OT in detail.

As it is evident, P plays a minimal role, involving permuting and obviously filtering out the message it receives from S . The size of the single message that R receives can be short (e.g., 64 or 128 bits) depending on the maximum bit size of the messages that S holds and the security parameter.

Theorem 1. *Let $\mathcal{F}_{\mathcal{OT}_1^2}$ be the functionality defined in Section 4.1. Then, Supersonic OT (presented in Figure 1) securely computes $\mathcal{F}_{\mathcal{OT}_1^2}$ in the presence of semi-honest adversaries, w.r.t. Definition 2.*

In the following subsections, we will prove the above theorem and then prove the correctness of Supersonic OT.

1. R-side Setup: $\text{Setup}(1^\lambda) \rightarrow (k_0, k_1)$
 - picks two random keys $(k_0, k_1) \xleftarrow{\$} \{0, 1\}^\sigma$ and sends them to S .
2. R-side Query Generation: $\text{GenQuery}(1^\lambda, s) \rightarrow q = (q_1, q_2)$
 - (a) splits the private index s into two shares (s_1, s_2) by calling:

$$\text{SS}(1^\lambda, s, 2, 2) \rightarrow (s_1, s_2)$$

(b) sends $q_1 = s_1$ to S and $q_2 = s_2$ to P .

3. S-side Response Generation: $\text{GenRes}(m_0, m_1, k_1, k_2, q_1) \rightarrow res$
 - (a) encrypts each message as follows.

$$\forall i, 0 \leq i \leq 1 : m'_i = m_i \oplus k_i$$

Let $e = (m'_0, m'_1)$ contain the encrypted messages.

(b) permutes the elements of e :

$$\bar{\pi}(s_1, e) \rightarrow e'$$

(c) sends $res = e'$ to P .

4. P-side Oblivious Filtering: $\text{OblFilter}(res, q_2) \rightarrow res'$
 - (a) permutes the elements of e' :

$$\bar{\pi}(s_2, e') \rightarrow e''$$

(b) sends (always) the first element in e'' , say $res' = e''_0$, to R and discards the second element in e'' .

5. R-side Message Extraction: $\text{Retrieve}(res', k_s) \rightarrow m_s$
 - retrieves the final related message m_s by decrypting e''_0 :

$$m_s = e''_0 \oplus k_s$$

Fig. 1: Supersonic OT.

4.3 Proof of Security

In this section, we prove the security of Supersonic OT, i.e., Theorem 1.

Proof. We consider the case where each party is corrupt at a time.

Corrupt Receiver R . In the real execution, R 's view is defined as:

$$\text{View}_R^{\text{Supersonic-OT}}((m_0, m_1), \epsilon, s) = \{r_R, e''_0, m_s\}$$

where r_R is the outcome of the internal random coin of R and is used to generate (s_1, s_2, k_0, k_1) . Below, we construct an idea-model simulator Sim_R which receives (s, m_s) from R .

1. initiates an empty view and appends a uniformly random coin r'_R to the view.
2. picks a random key $k \xleftarrow{\$} \{0, 1\}^\sigma$, using r'_R .
3. encrypts message m_s as follows $e = m_s \oplus k$.
4. appends e to the view and outputs the view.

Since we are in the passive adversarial model, the adversary picks its random coin r_R (in the real models) according to the protocol. Therefore, r_R and r'_R have identical distributions. Moreover, e''_0 in the real model and e in the ideal model have identical distributions as both are the result of XORing message m_s with a fresh uniformly random value. Also, m_s is the same in both models so it has identical distribution in the real and ideal models. We conclude that Relation 3 (in Section 4.1) holds.

Corrupt Sender S . In the real execution, S 's view can be defined as:

$$\text{View}_S^{\text{Supersonic-OT}}((m_0, m_1), \epsilon, s) = \{r_S, s_1, k_0, k_1\}$$

where r_S is the outcome of the internal random coin of S and is used to generate its random values. Next, we construct an idea-model simulator Sim_S which receives (m_0, m_1) from S .

1. initiates an empty view and appends a uniformly random coin r'_S to the view.
2. picks a binary random value $s' \xleftarrow{\$} \{0, 1\}$.
3. picks two uniformly random keys $(k'_0, k'_1) \xleftarrow{\$} \{0, 1\}^\sigma$.
4. appends s', k'_0, k'_1 to the view and outputs the view.

Next, we explain why the two views are indistinguishable. The random coins r_S and r'_S in the real and ideal models have identical distributions as they have been picked according to the protocol's description (as we consider the passive adversarial model). Moreover, s_1 in the real model and s' in the ideal model are indistinguishable, as due to the security of the secret sharing scheme, binary share s_1 is indistinguishable from a random binary value s' . Also, the elements of pair (k_0, k_1) in the real model and the elements of pair (k'_0, k'_1) in the ideal model have identical distributions as they have been picked uniformly at random from the same domain. Hence, Relation 1 (in Section 4.1) holds.

Corrupt Helper P . In the real execution, P 's view is defined as:

$$\text{View}_P^{\text{Supersonic-OT}}((m_0, m_1), \epsilon, s) = \{r_P, s_2, e'\}$$

where r_P is the outcome of the internal random coin of P and is used to generate its random values and e' is a pair (e'_0, e'_1) and is an output of $\bar{\pi}$. Next, we construct an idea-model simulator Sim_P .

1. initiates an empty view and appends a uniformly random coin r'_P to the view.
2. picks a binary random value $s' \xleftarrow{\$} \{0, 1\}$.
3. constructs a pair v of two uniformly random values $(v_0, v_1) \xleftarrow{\$} \{0, 1\}^\sigma$.
4. appends s', v to the view and outputs the view.

Since we consider the passive adversarial model, the adversary picks its random coins r_P and r'_P (in the real and ideal models respectively) according to the protocol. So, they have identical distributions. Moreover, s_2 in the real model and s' in the ideal model are indistinguishable, as due to the security of the secret sharing scheme, binary share s_2 is indistinguishable from a random binary value s' . In the real model, the elements of e' which are e'_0 and e'_1 have been encrypted/padded with two fresh uniformly random values. In the ideal model, the elements of v which are v_0 and v_1 have been picked uniformly at random. Due to the security of a one-time pad, e'_i ($\forall i, 0 \leq i \leq 1$) is indistinguishable from a uniformly random value, including v_0 and v_1 .

Also, in the real model, the pair e' that is given to P is always permuted based on the value of S 's share (i.e., $s_1 \in \{0, 1\}$) which is not known to P ; whereas, in the ideal model, the pair v is not permuted. However, given the permuted pair e' and not permuted pair v , a distinguisher cannot tell where each pair has been permuted with a probability greater than $\frac{1}{2}$. Therefore, Relation 2 (in Section 4.1) holds. \square

4.4 Proof of Correctness

In this section, we demonstrate that R always receives the message m_s corresponding to its query s . To accomplish this, we will show that (in step 4b) the first element of pair e'' always equals the encryption of m_s . This outcome is guaranteed by the following two facts: (a) $s = s_1 \oplus s_2$ and (b) S and T permute their pairs based on the value of their share, i.e., s_1 and s_2 respectively.

s	s_1	s_2
0	1	1
	0	0
1	1	0
	0	1

Table 1: Relation between query s and behaviour of permutation $\bar{\pi}$ from the perspective of S and P . When $s_i = 1$, $\bar{\pi}$ swaps the elements of its input pairs and when $s_i = 0$, $\bar{\pi}$ does not swap the elements of the input pairs.

As Table 1 indicates, when $s = 0$, then (i) either both S and P permute their pairs or (ii) neither does. In the former case, since both swap the elements of their pair, then the final permuted pair e'' will have the same order as the

original pair e (before it was permuted). In the latter case, again e'' will have the same order as the original pair e because neither party has permuted it. Thus, in both of the above cases (when $s = 0$), the first element of e'' will be the encryption of m_0 . Moreover, as Table 1 indicates, when $s = 1$, then only one of the parties S and P will permute their input pair. This means that the first element of the final permuted pair e'' will always equal the encryption of m_1 .

5 Performance Evaluation

We have implemented Supersonic OT in C++ and evaluated its concrete runtime. The source code for the implementation is publicly available in [1]. For the experiment, we utilized a MacBook Pro laptop equipped with a quad-core Intel Core i5, 2 GHz CPU, and 16 GB RAM. Given the low number and small size of exchanged messages in our protocol, implementing it in a network setting will yield negligible network overhead. We did not leverage parallelization or any other optimization. The experiment was run an average of 50 times. We utilized the GMP library [20] for big-integer arithmetic.

We analyzed the runtime of various phases of Supersonic OT across different invocation frequencies (1, 10, 10^3 , 10^5 , and 10^7 times). Table 2 shows the high-speed performance of Supersonic OT, requiring only 0.35 milliseconds (ms) for a single invocation. Notably, Phase 1 incurs the highest computation cost when the number of invocations is 1, 10, and 10^3 , while Phase 3 imposes the highest computation cost when the number of invocations is 10^5 and 10^7 . Overall, Phase 2 has the lowest computation cost.

Table 2: The table presents the runtime of Supersonic OT, categorized by various phases and measured in ms. The security parameter and message size are set at 128 bits.

Protocol	Phases	Number of OT Invocations				
		1	10	10^3	10^5	10^7
Supersonic OT	Phase 1	0.34	0.37	0.61	21.17	1990
	Phase 2	0.00069	0.00092	0.0071	0.7	63.86
	Phase 3	0.0011	0.0038	0.29	29.48	3058.09
	Phase 4	0.00065	0.0011	0.061	6.3	827.59
	Phase 5	0.00064	0.0012	0.063	6.016	675.06
	Total	0.35	0.38	1.05	63.7	6610

5.1 Runtime

Supersonic Versus Base OTs. Initially, we compare the runtime of Supersonic OT with that of *base* OTs proposed in [3,35]. These base OTs, known for their generality, efficiency, and widespread usage in literature, serve as efficient foundations in OT extensions. Table 3 and Figure 2a summarize this comparison. The runtime data for OTs in [3,35] is derived from the figures reported in [3], specifically from Table 3, where the GMP library was employed.³ Table 3 highlights that Supersonic OT demonstrates a speed advantage, being approximately 10^3 times faster than the OT in [3] and up to around 2×10^3 times faster than the OT in [35].

Table 3: The table compares the runtime of Supersonic OT with the following base OTs: standard OT (STD-OT) in [3], STD-OT in [35], and the random oracle OT (RO-OT) in [35]. The bit size of the security parameter is 128. The runtime has been measured in ms and is based on 128 invocations of each scheme. The enhancement ratio refers to the performance improvement that Supersonic OT offers in comparison to each specific scheme.

Scheme	Runtime	Enhancement Ratio
STD-OT in [3]	1,217	1,622
STD-OT in [35]	1,681	2,241
RO-OT in [35]	288	384
Supersonic OT	0.75	1

Note that Supersonic OT maintains a consistent runtime across different security parameters, whether lower (e.g., 80-bit) or higher (e.g., 256-bit) than 128-bit. In contrast, the runtime of the schemes in [3,35] would vary, fluctuating by at least a factor of 2.5 with changes in the security parameter.

Supersonic Versus OT Extensions. We proceed to compare the runtime of Supersonic OT with the runtime of efficient OT extensions presented in [3,41]. Given that OT extensions are designed for scenarios involving frequent executions, we evaluate the runtime of these three OTs when invoked 10^7 times. The runtime data for OTs in [3,41] is extracted from the figures reported in [3],

³ The implementation in [3] used a smaller RAM size (4 GB) compared to our usage of 16 GB. Their reported performance also accounts for LAN delay. It is important to note that doubling the RAM size doesn't necessarily equate to doubling the runtime. Moreover, the impact of LAN delay is negligible, particularly for our OT, which incurs a very low communication cost.

Table 4: The table compares the runtime of Supersonic OT with the following general OT (G-OT) extensions: G-OT in [3] and G-OT in [41]. The runtime has been measured in milliseconds and is for 10^7 invocations of 1-out-of-2 OT. The enhancement ratio refers to the performance improvement that Supersonic OT offers in comparison to each scheme.

Scheme	Sec. Param. Size	Runtime	Enhancement Ratio
G-OT in [3]	80-bit	14,272	2
G-OT in [41]	80-bit	20,717	3
Supersonic OT	128-bit	6,610	1

specifically from Tables 3 and 4 in [3], where the GMP library was utilized.⁴ Table 4 and Figure 2b present the outcome of this comparison.

Table 4 shows that invoking Supersonic OT 10^7 times takes approximately 6,610 ms with a 128-bit security parameter.

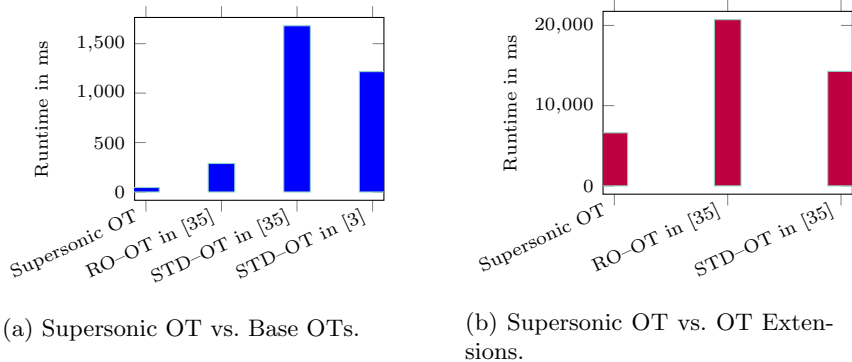


Fig. 2: Comparison of Runtimes for Different OTs.

Supersonic OT outperforms the OT in [3] by a factor of 2 and [41] by a factor of 3. Despite the higher 128-bit security parameter in Supersonic OT compared to the 80-bit parameter in the other two schemes, its runtime is still lower. We expect that increasing the security parameter in schemes in [3,41] would result in higher runtimes, given that the base OT's runtime increases accordingly.

⁴ Table 4 in [3] excludes the runtime of base OTs. Thus, the corresponding runtime of the base OT in Table 3 in [3] must be added to each figure reported in Table 4. For instance, for G-OT (in the LAN setting) we would have $13.92 \times 1000 + 352 = 14,272$.

5.2 Features

For the base OTs and OT extensions in [3,35,41] to achieve unconditional security, as discussed in Section 3.1, they typically require multiple replicas of the database, a noisy channel, or the involvement of a trusted initializer, all of which contribute to increased deployment costs. In contrast, Supersonic OT attains unconditional security without relying on database replications, noisy channels, or a fully trusted party. Although Supersonic OT involves an additional party, unlike base OTs or OT extensions that typically only involve the sender and receiver, it maintains its security even when this party is semi-honest.

6 Conclusion and Future Work

OT is a crucial privacy-preserving technology. OTs have found extensive applications in designing secure Multi-Party Computation protocols [46,3,24], Federated Learning [45,39,43], and accessing sensitive field elements of remote private databases while preserving privacy [9,2,31].

In this work, we presented Supersonic OT, an unconditionally secure very efficient 1-out-of-2 OT that will maintain security in the quantum computing era. Supersonic OT refrains from using any public-key cryptography. We have proved its security within the standard simulation-based paradigm.

Supersonic OT enables the receiver to obtain a response of size $O(1)$ whose bit size can be as small as 128 if the bit size of the secret messages maintained by the sender is at most 128. It relies on basic standard primitives. The simple design of Supersonic OT allows easy security analysis and implementation. We have studied the concrete performance of Supersonic OT. A single execution of Supersonic OT completes in approximately 0.35 milliseconds, making it up to 2000 times faster than the state-of-the-art base OT. Moreover, we studied the performance of Supersonic OT and the state-of-the-art OT extensions in the setting where they are invoked 10^7 times. Our analysis indicated that Supersonic OT takes approximately 6,610 ms, making it up to 3 times faster than current OT extensions.

As a future research direction, it would be interesting to see how the exceptional efficiency of Supersonic OT can improve the performance of those (generic MPC or Private Set Intersection) protocols that heavily rely on OT protocols.

References

1. Abadi, A.: Source code of Supersonic OT (2024), <https://github.com/AydinAbadi/Supersonic-OT>
2. Aiello, W., Ishai, Y., Reingold, O.: Priced oblivious transfer: How to sell digital goods. In: EUROCRYPT (2001)
3. Asharov, G., Lindell, Y., Schneider, T., Zohner, M.: More efficient oblivious transfer and extensions for faster secure computation. In: CCS'13 (2013)
4. Barreto, P., Oliveira, G., Benits, W.: Supersingular isogeny oblivious transfer. IACR Cryptol. ePrint Arch. (2018)

5. Blakley, G.R.: One time pads are key safeguarding schemes, not cryptosystems. fast key safeguarding schemes (threshold schemes) exist. In: IEEE S&P (1980)
6. Blazy, O., Chevalier, C.: Generic construction of uc-secure oblivious transfer. In: ACNS (2015)
7. Blazy, O., Chevalier, C., Vu, Q.: Post-quantum uc-secure oblivious transfer in the standard model with adaptive corruptions. In: ARES (2019)
8. Blundo, C., D'Arco, P., Santis, A.D., Stinson, D.R.: On unconditionally secure distributed oblivious transfer. *J. Cryptol.* (2007)
9. Camenisch, J., Dubovitskaya, M., Neven, G.: Oblivious transfer with access control. In: CCS (2009)
10. Camenisch, J., Neven, G., Shelat, A.: Simulatable adaptive oblivious transfer. In: Advances in Cryptology - EUROCRYPT (2007)
11. Chen, Y., Chou, J., Hou, X.: A novel k -out-of- n oblivious transfer protocols based on bilinear pairings. *IACR Cryptol. ePrint Arch.* (2010)
12. Chu, C., Tzeng, W.: Efficient k -out-of- n oblivious transfer schemes with adaptive and non-adaptive queries. In: PKC (2005)
13. Corniaux, C.L.F., Ghodosi, H.: A verifiable 1-out-of- n distributed oblivious transfer protocol. *IACR Cryptol. ePrint Arch.* (2013)
14. Crépeau, C., Kilian, J.: Achieving oblivious transfer using weakened security assumptions (extended abstract). In: FoCS (1988)
15. Crépeau, C., Morozov, K., Wolf, S.: Efficient unconditional oblivious transfer from almost any noisy channel. In: Security in Communication Networks, 4th International Conference, SCN (2004)
16. Dong, C., Chen, L., Wen, Z.: When private set intersection meets big data: an efficient and scalable protocol. In: CCS (2013)
17. Dowsley, R., van de Graaf, J., Müller-Quade, J., Nascimento, A.C.A.: Oblivious transfer based on the mceliece assumptions. *IEICE Trans. Fundam. Electron. Commun. Comput. Sci.* (2012)
18. Even, S., Goldreich, O., Lempel, A.: A randomized protocol for signing contracts. *Commun. ACM* (1985)
19. Fredkin, E., Toffoli, T.: Conservative logic. In: International Journal of Theoretical Physics (1982)
20. GNU Project: The gnu multiple precision arithmetic library (1991), <https://gmplib.org>
21. Goldreich, O.: The Foundations of Cryptography - Volume 2, Basic Applications. Cambridge University Press (2004)
22. Green, M., Hohenberger, S.: Universally composable adaptive oblivious transfer. In: ASIACRYPT (2008)
23. Gunupudi, V., Tate, S.R.: Generalized non-interactive oblivious transfer using count-limited objects with applications to secure mobile agents. In: FC (2008)
24. Harnik, D., Ishai, Y., Kushilevitz, E.: How many oblivious transfers are needed for secure multiparty computation? In: CRYPTO (2007)
25. Henecka, W., Schneider, T.: Faster secure two-party computation with less memory. In: CCS (2013)
26. Ishai, Y., Kilian, J., Nissim, K., Petrank, E.: Extending oblivious transfers efficiently. In: CRYPTO, (2003)
27. Ishai, Y., Kushilevitz, E., Ostrovsky, R., Prabhakaran, M., Sahai, A., Wullschlegel, J.: Constant-rate oblivious transfer from noisy channels. In: CRYPTO (2011)
28. Jarecki, S., Liu, X.: Efficient oblivious pseudorandom function with applications to adaptive OT and secure computation of set intersection. In: TCC (2009)

29. Katz, J., Lindell, Y.: Introduction to Modern Cryptography, Second Edition. CRC Press (2014)
30. Kundu, N., Debnath, S.K., Mishra, D.: 1-out-of-2: post-quantum oblivious transfer protocols based on multivariate public key cryptography. *Sādhanā* (2020)
31. Libert, B., Ling, S., Mouhartem, F., Nguyen, K., Wang, H.: Adaptive oblivious transfer with access control from lattice assumptions. *Theoretical Computer Science* (2021)
32. Liu, M., Hu, Y.: Universally composable oblivious transfer from ideal lattice. *Frontiers Comput. Sci.* (2019)
33. Naor, M., Pinkas, B.: Oblivious transfer and polynomial evaluation. In: *STOC* (1999)
34. Naor, M., Pinkas, B.: Distributed oblivious transfer. In: *ASIACRYPT*. Springer (2000)
35. Naor, M., Pinkas, B.: Efficient oblivious transfer protocols. *SODA* (2001)
36. Nielsen, J.B.: Extending oblivious transfers efficiently - how to get robustness almost for free. *IACR Cryptol. ePrint Arch.* (2007)
37. Peikert, C., Vaikuntanathan, V., Waters, B.: A framework for efficient and composable oblivious transfer. In: *CRYPTO* (2008)
38. Rabin, M.O.: How to exchange secrets with oblivious transfer (1981)
39. Ren, Z., Yang, L., Chen, K.: Improving availability of vertical federated learning: Relaxing inference on non-overlapping data. *ACM Trans. Intell. Syst.* (2022)
40. Rivest, R.: Unconditionally secure commitment and oblivious transfer schemes using private channels and a trusted initializer. technical report (1999)
41. Schneider, T., Zohner, M.: MW vs. Yao? efficient secure two-party computation with low depth circuits. In: *FC* (2013)
42. Tzeng, W.: Efficient 1-out-n oblivious transfer schemes. In: Naccache, D., Paillier, P. (eds.) *PKC* (2002)
43. Xu, G., Li, H., Zhang, Y., Xu, S., Ning, J., Deng, R.H.: Privacy-preserving federated deep learning with irregular users. *IEEE Trans. Dependable Secur. Comput.* (2022)
44. Yadav, V.K., Andola, N., Verma, S., Venkatesan, S.: A survey of oblivious transfer protocol. *ACM Comput. Surv.* (2022)
45. Yang, Q., Liu, Y., Chen, T., Tong, Y.: Federated machine learning: Concept and applications. *ACM Trans. Intell. Syst. Technol.* (2019)
46. Yao, A.C.: Protocols for secure computations (extended abstract). In: *23rd Annual Symposium on Foundations of Computer Science* (1982)
47. Zhang, B., Lipmaa, H., Wang, C., Ren, K.: Practical fully simulatable oblivious transfer with sublinear communication. In: *FC* (2013)
48. Zhao, S., Song, X., Jiang, H., Ma, M., Zheng, Z., Xu, Q.: An efficient outsourced oblivious transfer extension protocol and its applications. *Secur. Commun. Networks* (2020)