Actively Secure 1-out-of-N OT Extension with Application to Private Set Intersection

Michele Orrù[1]($\boxtimes$), Emmanuela Orsini[2], and Peter Scholl[2]

[1] CNRS, ENS Paris, Paris, France
michele.orru@ens.fr
[2] Department of Computer Science, University of Bristol, Bristol, UK
{Emmanuela.Orsini,Peter.Scholl}@bristol.ac.uk

Abstract. This paper describes a 1-out-of-N oblivious transfer (OT) extension protocol with active security, which achieves very low overhead on top of the passively secure protocol of Kolesnikov and Kumaresan (Crypto 2011). Our protocol obtains active security using a consistency check which requires only simple computation and has a communication overhead that is independent of the total number of OTs to be produced. We prove its security in both the random oracle model and the standard model, assuming a variant of correlation robustness. We describe an implementation, which demonstrates our protocol only costs around 5–30% more than the passively secure protocol.

Random 1-out-of-N OT is a key building block in recent, very efficient, passively secure private set intersection (PSI) protocols. Our random OT extension protocol has the interesting feature that it even works when N is exponentially large in the security parameter, provided that the sender only needs to obtain polynomially many outputs. We show that this can be directly applied to improve the performance of PSI, allowing the core private equality test and private set inclusion subprotocols to be carried out using just a single OT each. This leads to a reduction in communication of up to 3 times for the main component of PSI.

Keywords: Oblivious transfer · Private set intersection · Multi-party computation

1 Introduction

Oblivious transfer (OT) is a fundamental primitive in cryptography, first introduced by Rabin [Rab81] and now employed in a variety of protocols, ranging from contract signing [EGL85] to special-purpose tasks such as private set intersection [PSZ14]. It plays a decisive role in protocols for secure two-party and multi-party computation, including those based on Yao's garbled circuits [Yao82] and secret-sharing [NNOB12,LOS14,KOS16]. The most commonly studied form

---

Full version available at http://eprint.iacr.org/2016/933.pdf

M. Orrù–Work done while visiting University of Bristol.

of oblivious transfer is 1-out-of-2 OT, where a sender has two messages (x_0, x_1) as input, and a receiver chooses a bit b; the goal of the protocol is for the receiver to learn x_b, but no information on x_{1-b}, whilst the sender learns nothing about b. This can be generalized to 1-out-of-N OT and k-out-of-N OT, in which the receiver learns k of the sender's N messages.

Unfortunately, due to a result of Impagliazzo and Rudich [IR89], oblivious transfer is highly unlikely to be possible without the use of public-key cryptography; consequently, even the most efficient oblivious transfer constructions [PVW08, CO15] come with a relatively high cost.

OT Extensions. In 1996, Beaver [Bea96] first showed that it is possible to extend OT starting with a small number (say, security parameter κ) of "base" OTs, to create $\mathsf{poly}(\kappa)$ additional OTs using only symmetric primitives, with computational security κ. This construction is very impractical as it requires the evaluation of pseudorandom generators within Yao's garbled circuits.

Later, in 2003, Ishai et al. [IKNP03] proposed a protocol for extending oblivious transfers: the passively secure version of this protocol (hereafter IKNP) only requires black-box use of a correlation robust hash function, and is very efficient. Concretely, an optimized version of IKNP for OT on random strings (described in [ALSZ13, KK13]) requires sending κ bits and computing three hash function evaluations per OT, after a one-time cost of κ base OTs, for computational security κ. With a carefully optimized implementation, the dominant cost of this is communication [ALSZ16].

Kolesnikov and Kumaresan [KK13] showed how to modify the IKNP protocol using Walsh-Hadamard error-correcting codes and obtain a passively secure protocol for 1-out-of-N OT on random strings. The cost is only a small constant factor more than the 1-out-of-2 IKNP for values of N up 256.

Several recent works have proposed increasingly efficient protocols for 1-out-of-2 OT extension with active security [NNOB12, ALSZ15, KOS15]. The latter work of Keller et al. [KOS15], which is proven secure in the random oracle model, brings the cost of actively secure 1-out-of-2 OT to essentially the same as the passive IKNP protocol by adding a simple consistency check.

1.1   Contributions

Actively Secure 1-out-of-N OT Extension. Our main contribution is a practical, actively secure 1-out-of-N OT extension protocol with very low overhead on top of the passively secure protocol of Kolesnikov and Kumaresan [KK13]. For the case of random OT, where the sender's strings are sampled at random, our protocol (proven secure in the random oracle model) improves upon [KK13] by allowing for much larger values of N with a suitable choice of binary linear code. Our protocol even works when N is *exponential in the security parameter*, provided that the sender is only required to learn polynomially many output strings. The protocol requires only κ base OTs, and the extension phase has an amortized communication cost of $O(\kappa)$ bits per random OT.

At a high level, our protocol starts with the passively secure [KK13] protocol and adds a simple consistency check to obtain active security (similar to [KOS15] for 1-out-of-2 OT). However, there are several technical challenges to solve on the way. In [KOS15], a check is used to verify that pairs of strings are of the form $(\mathbf{x}_i, \mathbf{x}_i + \mathbf{b})$ for a fixed correlation $\mathbf{b}$ (with addition modulo 2), when the receiver only knows one string from each pair. In the [KK13] protocol, however, we must ensure that strings are of the form $\mathbf{x}_i + \mathbf{b} \odot \mathcal{C}(m_i)$, where $\mathcal{C}$ encodes a message m_i using an error-correcting code and $\odot$ denotes the component-wise product of bit vectors. The check of [KOS15] cannot be applied to this situation. We overcome this by adapting a check used previously in additively homomorphic UC commitments [FJNT16], which requires that $\mathcal{C}$ is a *linear* code with sufficiently large minimum distance.[1] The number of codewords in the binary linear code determines N in the 1-out-of-N OT, which gives a range of choices of N depending on the choice of code.

To be able to handle exponentially large N, it may seem that we just need to choose a suitable binary linear code of the right length. However, we need to take care that the security reduction does not contain any loss in security that scales with N: the reduction in [KK13] incurs a loss in $O(N^2)$, which would give a meaningless security result in this case. To ensure this, we modify the 1-out-of-N random OT functionality so that the sender can only obtain $N' = \mathsf{poly}(\kappa)$ of the output messages, and show that the loss in the resulting reduction is in $O(N')$.

Security in the Standard Model. For random OT extension, it is not known how to prove security without using a programmable random oracle as in [ALSZ13,KOS15]. However, for the case of non-random 1-out-of-N OT, we prove our protocol secure in the standard model, assuming a hash function that satisfies a variant of correlation robustness on high min-entropy secrets. This is a similar assumption to the protocol in [ALSZ15], but more general as we require the assumption to hold for a range of different parameters. This gives the first actively secure OT extension protocol needing only κ base OTs for security parameter κ and is proven secure without random oracles, even in the 1-out-of-2 case.[2]

Faster Private Set Intersection. We show that random 1-out-of-N OT with an exponentially large N can be directly applied to improve the efficiency of the previous fastest (semi-honest) private set intersection protocols. OT-based PSI protocols [PSZ14,PSSZ15] use random 1-out-of-N OT as a building block for a private equality test protocol, where two parties learn whether their inputs are

---

[1] We observe an interesting connection between our protocol and additively homomorphic UC commitment schemes [FJNT16,CDD+16]: our protocol essentially runs a homomorphic commitment protocol and hashes the resulting commitments to obtain random OTs. However, this mechanism seems very specific to the workings of these commitment schemes and appears unlikely to lead to a generic transformation.

[2] Note that our security reduction requires fixing the adversary's random coins, so is non-uniform. Obtaining a uniform reduction seems to need at least $\kappa + s$ base OTs, for statistical security parameter s.

equal (and nothing more). In that protocol, one random OT is used to perform an equality test on $\log N$-bit inputs. Since the random OT protocol of [KK13] only works for values of N up to 256 (due to the use of small Walsh-Hadamard codes) several OTs are XORed together to construct a protocol for comparing large (e.g. up to 128 bit) messages. Using our protocol with $N = 2^k$ gives a very simple private equality test on k-bit messages, for any $k = \mathsf{poly}(\kappa)$, using just a single 1-out-of-N random OT. This can be generalized to perform *private set inclusion*—where one party holds a single value and another party a set of m values—at the cost of one random OT and sending $m \cdot s$ bits, where s is the statistical security parameter. This results in a reduction in communication of around 2–5 times (depending on the bit-length of the input) for this component of the semi-honest PSI protocol in [PSSZ15].

Implementation. We have implemented and benchmarked our 1-out-of-N random OT extension protocol and compared its performance with the passive protocol of [KK13]. Although our implementation is not heavily optimized (it occupies around 800 lines of C in all), we show that the overhead of our consistency check for achieving active security is very low: the actively secure protocol takes only around 20% more time than the passive version, depending on parameters.

Towards Efficient Actively Secure PSI. Currently, the most efficient PSI protocols are the OT-based ones mentioned above, but these are only secure against a passive adversary. Since 1-out-of-N random OT is a key component in these protocols, our work can be seen as a step towards constructing more efficient PSI with active security. Actively secure PSI was recently studied by Lambæk [Lam16], who showed the protocol of [PSSZ15] can be modified to provide active security for one party, assuming the underlying OT protocol is actively secure; our protocol therefore provides an instantiation of this proposal.

Recent Work and Open Problems. In a very recent, independent work, Kolesnikov et al. [KKRT16] describe a batched oblivious PRF evaluation protocol with application to private set intersection. Although their protocol is phrased in the language of oblivious PRFs rather than 1-out-of-N OT, it is very similar to ours, only with passive security. Instead of using a traditional error-correcting code, they show that a random oracle has the necessary properties for passive security. In contrast, our protocol requires the linearity and erasure decoding properties of the binary code to achieve active security. They describe the same application to improved performance of PSI (with slightly better parameters than ours due to use of a random oracle) and give a thorough efficiency evaluation and implementation of the resulting PSI protocol. We note that it is still an interesting open problem to obtain a fully actively secure variant of the PSI protocol in [PSSZ15] with low overhead.

Regarding OT extension in general, there are still some interesting unsolved problems. Our 1-out-of-N OT extension cannot be used directly to improve performance of 1-out-of-2 OT on short secrets (as was done for the passive case in [KK13]), since the standard reduction from 1-out-of-N to 1-out-of-2 OT [NP99] is only passively secure. Therefore, it is still an open problem

to construct a practical 1-out-of-2 OT extension on short strings with communication sublinear in the security parameter. Also, the case of constructing k-out-of-N OT with active security using OT extensions is still open; there is an elegant passively secure protocol [SSR08], but it seems difficult to make this actively secure.

2   Preliminaries

Notation. We denote by κ and s the computational, resp. statistical, security parameters. We use bold lower case letters for vectors. Given a matrix A, we let $\mathbf{a}_i$ denote the i-th row of A, and $\mathbf{a}^j$ denote the j-th column of A. When referring to a vector $\mathbf{v} \in \mathbb{F}^n$, we write $\mathbf{v}[i]$, with $1 \leq i \leq n$, to mean the i-th component of $\mathbf{v}$. We identify bit strings as vectors over the finite field $\mathbb{F}_2$, and use "+" and "·" to mean addition and multiplication in this field. We use the notation $\mathbf{a} \odot \mathbf{b}$ to denote the component-wise product of vectors $\mathbf{a}, \mathbf{b} \in \mathbb{F}_2^n$. Given an integer N, we denote by $[N]$ the set of integers $\{1, \ldots, N\}$.

Error-Correcting Codes. Our protocol uses an $[n_\mathcal{C}, k_\mathcal{C}, d_\mathcal{C}]$ binary linear code $\mathcal{C}$, where $n_\mathcal{C}$ is the length, $k_\mathcal{C}$ the dimension and $d_\mathcal{C}$ the distance of $\mathcal{C}$. So, $\mathcal{C} : \mathbb{F}_2^{k_\mathcal{C}} \to \mathbb{F}_2^{n_\mathcal{C}}$ is a linear map such that for every pair of messages $\mathbf{m}_1, \mathbf{m}_2 \in \mathbb{F}_2^{k_\mathcal{C}}$, the Hamming weight of the sum of the encodings of the messages satisfies $\mathsf{wt}_\mathsf{H}(\mathcal{C}(\mathbf{m}_1) + \mathcal{C}(\mathbf{m}_2)) \geq d_\mathcal{C}$.

Oblivious Transfer Functionalities. We now recall some definitions of oblivious transfer. Following Even et al. [EGL85], 1-out-of-2 OT is a two-party protocol between a sender P_S, who inputs two messages v_0, v_1, and a receiver P_R who inputs a choice bit c and learns as output v_c and nothing about v_{1-c}, in such a way that P_S remains oblivious as what message was received by P_R. Formally, the general case of 1-out-of-N OT on κ-bit strings is defined as the functionality:

$$\mathcal{F}_{N\text{-OT}}((\mathbf{v}_0, \ldots, \mathbf{v}_{N-1}), c) = (\bot, \mathbf{v}_c),$$

where $\mathbf{x}_i \in \{0, 1\}^\kappa$ are the sender's inputs and $c \in \{0, \ldots, N-1\}$ is the receiver's input. We denote by $\mathcal{F}_{N\text{-OT}}^{\kappa, m}$ the functionality that runs $\mathcal{F}_{N\text{-OT}}$ m times on messages in $\{0, 1\}^\kappa$. For example, in $\mathcal{F}_{2\text{-OT}}^{\kappa, m}$, P_S inputs $(\mathbf{v}_{i,0}, \mathbf{v}_{i,1})$ and P_R inputs c_i for $i \in [m]$, and P_R receives the output $\mathbf{v}_{i,c_i}$.

Another important variant is the random OT functionality $\mathcal{F}_{N\text{-ROT}}^{\kappa, m}$, in which the sender provides no input, but receives random messages $(\mathbf{v}_0, \ldots, \mathbf{v}_{N-1})$ from the functionality as output.

2.1   Passively Secure OT Extension: The KK Protocol

We now recall the passively secure KK protocol for 1-out-of-N OT extension described in [KK13], which is a generalized version of the IKNP protocol for 1-out-of-2 OT [IKNP03].

Suppose the two parties wish to perform m sets of 1-out-of-N random OTs, where N is a power of two. There is a sender P_S with no input, and a receiver

P_R, who inputs the choices $w_1, \ldots, w_m \in \{0, \ldots, N-1\}$, which are represented as vectors $\mathbf{w}_i \in \mathbb{F}_2^{\log N}$. The two parties begin by performing n_C base 1-out-of-2 OTs on random inputs, with the roles of sender and receiver reversed. So, P_R obtains n_C pairs of random strings $(\mathbf{r}_0^j, \mathbf{r}_1^j)$ of length κ and P_S obtains $(b_j, \mathbf{r}_{b_j}^j)$, where $b_j \xleftarrow{\$} \{0,1\}$, for $j \in [n_C]$.

Next, both parties locally extend their base OT outputs to length m using a pseudorandom generator, where m is the final number of OTs desired. This results in κ sets of 1-out-of-2 OTs on m-bit strings, which we represent as matrices $T_0, T_1 \in \mathbb{F}_2^{m \times n_c}$, held by P_R, whilst P_S holds the vector $\mathbf{b} = (b_1, \ldots, b_{n_c}) \in \mathbb{F}_2^{n_c}$ and the matrix

$$T_{\mathbf{b}} := \left(\mathbf{t}_{b_1}^1 \ \ldots \ \mathbf{t}_{b_{n_C}}^{n_c} \right) \in \mathbb{F}_2^{m \times n_c},$$

where $\mathbf{t}_0^j, \mathbf{t}_1^j$ are the columns of T_0, T_1, for $j \in [n_C]$.

At this point P_R constructs a matrix $C \in \mathbb{F}_2^{m \times n_c}$, where each row $\mathbf{c}_i$ is the encoding $\mathcal{C}(\mathbf{w}_i)$ of the input $\mathbf{w}_i \in \mathbb{F}_2^{k_c}$, where $\mathcal{C}$ is a binary code of length n_C, dimension $k_C = \log_2 N$ and minimum distance $d_C \geq \kappa$. Then P_R sends to P_S the matrix

$$U = T_0 + T_1 + C.$$

Note that for each column of U, all information on the receiver's encoded input is masked by the value $\mathbf{t}_{1-b_j}^j$, which is unknown to P_S.

After this step P_S defines an $m \times n_C$ matrix Q with columns $\mathbf{q}^j = b_j \cdot \mathbf{u}^j + \mathbf{t}_{b_j}^j = b_j \cdot \mathbf{c}^j + \mathbf{t}_0^j$ (where $\mathbf{c}^j$ are the columns of C). Notice that the rows of Q are given by

$$\mathbf{q}_i = \mathbf{c}_i \odot \mathbf{b} + \mathbf{t}_i,$$

where $\mathbf{t}_i$ are the rows of T_0. Here, P_R holds $\mathbf{t}_i$ and P_S holds $(\mathbf{q}_i, \mathbf{b})$, for $i \in [m]$. The key observation to turn these values into OTs is that for each of the possible receiver choices $\mathbf{w} \in \mathbb{F}_2^{k_c}$, P_S can compute the value $\mathbf{q}_i + \mathcal{C}(\mathbf{w}) \odot \mathbf{b}$. If $\mathbf{w} = \mathbf{w}_i$ then this is equal to $\mathbf{t}_i$ so is known to P_R. Otherwise, for any $\mathbf{w} \neq \mathbf{w}_i$, P_R must guess κ bits of P_S's secret $\mathbf{b}$ to be able to compute $\mathbf{q}_i + \mathcal{C}(\mathbf{w}) \odot \mathbf{b}$, since the minimum distance of $\mathcal{C}$ guarantees that $\mathcal{C}(\mathbf{w})$ and $\mathcal{C}(\mathbf{w}_i)$ are at least Hamming distance κ apart.

Therefore, the parties can convert these values to random 1-out-of-N OTs by simply hashing their outputs with a random oracle, H. P_S outputs the values $\mathbf{v}_{w,i} = \mathsf{H}(i, \mathbf{q}_i + \mathcal{C}(\mathbf{w}) \odot \mathbf{b})$, for all $\mathbf{w} \in \mathbb{F}_2^{k_c}$, and P_R outputs $\mathbf{v}_{w_i,i} = \mathsf{H}(i, \mathbf{t}_i)$.

Instantiating the Code. As noticed in [KK13], if we instantiate the binary linear code $\mathcal{C}$ with the $[\kappa, 1, \kappa]$ binary repetition code, we obtain the 1-out-of-2 IKNP protocol [IKNP03]. In this case, each row of the matrix C constructed by the receiver is either 0^κ or 1^κ, depending on the receiver's choice bits. If instead $\mathcal{C}$ is chosen to be a Walsh-Hadamard code as in [KK13], then the result is a 1-out-of-2^{k_c} OT. This needs a code length of $N = 2^{k_c}$ with security parameter

$N/2$; this turns out to be more efficient than constructing 1-out-of-N OT from 1-out-of-2 OT for values of $N \leq 256$ with 128-bit security.

Security. The KK protocol (and hence IKNP) is actively secure against a corrupt sender, since after the base OTs, there is no opportunity for P_S to cheat. However, it only provides passive security against a corrupt receiver, since P_R may incorrect compute the encodings of their input in the matrix U. It was explained in [IKNP03, ALSZ15] that if P_R cheats in this way, and also learns (via a side-channel, for instance) the sender's outputs in just κ of the random OTs then P_R can compute the sender's secret $\mathbf{b}$, and thus learn all of the sender's outputs in every remaining OT.

3   Actively Secure Random 1-out-of-N OT Extension

In this section we present our actively secure OT extension protocol in the random oracle model. Since we want to construct 1-out-of-N random OT when N may be exponential in the security parameter, our protocol implements a modified random OT functionality $\mathcal{F}_{N\text{-ROT+}}$ (Fig. 1), which allows the sender to query the functionality to obtain their random OT outputs one at a time, so that all N need not be produced.

Functionality $\mathcal{F}_{N\text{-ROT+}}^{\kappa,m}$

Upon receiving (ROT) from P_S and $(\mathsf{ROT}, (w_1, \ldots, w_m))$, where $\forall i \in [m], 0 \leq w_i < N$, from P_R, do the following:

- Sample $\mathbf{v}_{j,i} \xleftarrow{\$} \mathbb{F}_2^\kappa$ for each $j \in \{0, \ldots, N-1\}$ and $i \in [m]$.
- Send the outputs $\mathbf{v}_{w_i,i}$ to P_R for each $i \in [m]$
- Upon receiving (SenderOutput, j, i) from P_S, where $i \in [m]$ and $j \in \{0, \ldots, N-1\}$, send $\mathbf{v}_{j,i}$ to P_S.
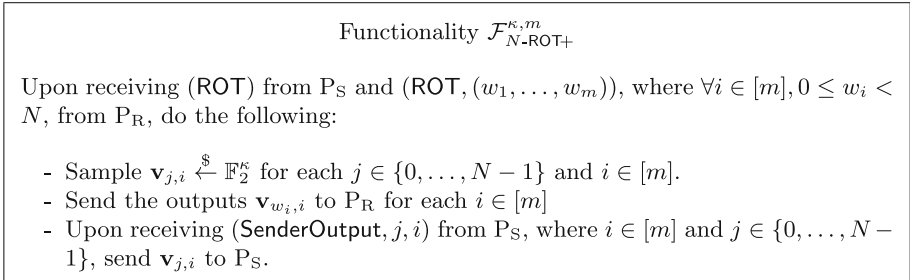
Fig. 1. Ideal functionality $\mathcal{F}_{N\text{-ROT+}}$ for m 1-out-of-$(\leq N)$ random OTs on κ-bit strings between a sender P_S and receiver P_R

The high-level idea of our protocol (Fig. 2) is that, to deal with a malicious receiver in the KK protocol, we add a consistency check that ensures P_R inputs codewords as rows of the matrix C when sending the matrix U in step 3. If the check passes then the protocol carries on and the correlated OTs are hashed to obtain random OTs. Otherwise, the protocol aborts.

The intuition behind security is that if not all the P_R's inputs $\mathbf{c}_i$ are codewords then to pass the check, the errors must 'cancel out' when taking the random linear combinations. However, the $x_i^{(\ell)}$ values used in the consistency check are unknown when P_R chooses $\mathbf{c}_i$ so this can only happen with negligible probability; since each $x_i^{(\ell)} \in \{0, 1\}$, there is a $1/2$ probability that $\mathbf{c}_i$ is not

---

<div style="text-align:center">Protocol N-$\mathsf{ROT}^{\kappa,m}$</div>

COMMON INPUT: κ and s are the computational and statistical security parameters, respectively; $\mathcal{C}$ is an 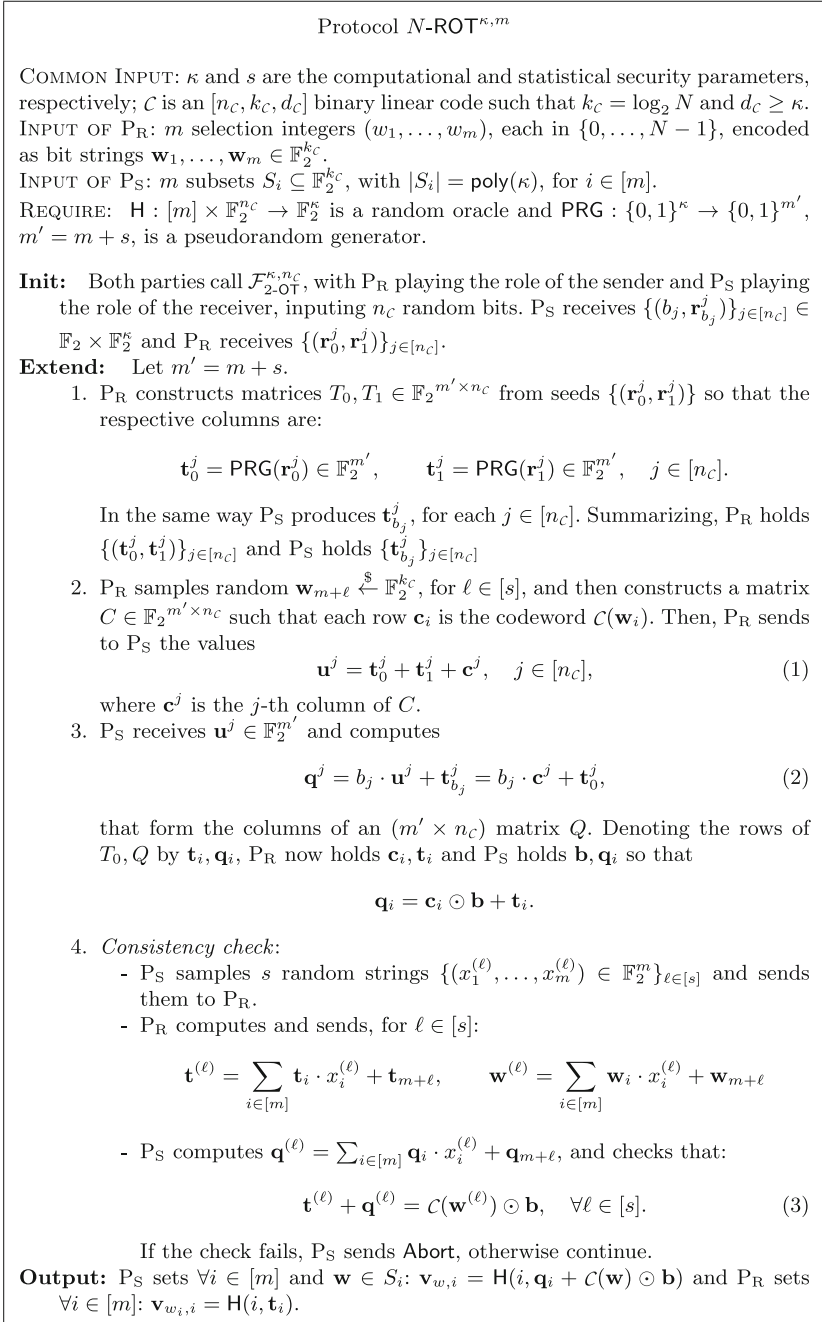$[n_\mathcal{C}, k_\mathcal{C}, d_\mathcal{C}]$ binary linear code such that $k_\mathcal{C} = \log_2 N$ and $d_\mathcal{C} \geq \kappa$.
INPUT OF $\mathrm{P_R}$: m selection integers $(w_1, \dots, w_m)$, each in $\{0, \dots, N-1\}$, encoded as bit strings $\mathbf{w}_1, \dots, \mathbf{w}_m \in \mathbb{F}_2^{k_\mathcal{C}}$.
INPUT OF $\mathrm{P_S}$: m subsets $S_i \subseteq \mathbb{F}_2^{k_\mathcal{C}}$, with $|S_i| = \mathsf{poly}(\kappa)$, for $i \in [m]$.
REQUIRE: $\mathsf{H} : [m] \times \mathbb{F}_2^{n_\mathcal{C}} \to \mathbb{F}_2^{\kappa}$ is a random oracle and $\mathsf{PRG} : \{0,1\}^\kappa \to \{0,1\}^{m'}$, $m' = m + s$, is a pseudorandom generator.

Init: Both parties call $\mathcal{F}_{2\text{-}\mathsf{OT}}^{\kappa, n_\mathcal{C}}$, with $\mathrm{P_R}$ playing the role of the sender and $\mathrm{P_S}$ playing the role of the receiver, inputing $n_\mathcal{C}$ random bits. $\mathrm{P_S}$ receives $\{(b_j, \mathbf{r}_{b_j}^j)\}_{j\in[n_\mathcal{C}]} \in \mathbb{F}_2 \times \mathbb{F}_2^{\kappa}$ and $\mathrm{P_R}$ receives $\{(\mathbf{r}_0^j, \mathbf{r}_1^j)\}_{j\in[n_\mathcal{C}]}$.

Extend: Let $m' = m + s$.
1. $\mathrm{P_R}$ constructs matrices $T_0, T_1 \in \mathbb{F}_2^{m' \times n_\mathcal{C}}$ from seeds $\{(\mathbf{r}_0^j, \mathbf{r}_1^j)\}$ so that the respective columns are:

$$\mathbf{t}_0^j = \mathsf{PRG}(\mathbf{r}_0^j) \in \mathbb{F}_2^{m'}, \qquad \mathbf{t}_1^j = \mathsf{PRG}(\mathbf{r}_1^j) \in \mathbb{F}_2^{m'}, \quad j \in [n_\mathcal{C}].$$

In the same way $\mathrm{P_S}$ produces $\mathbf{t}_{b_j}^j$, for each $j \in [n_\mathcal{C}]$. Summarizing, $\mathrm{P_R}$ holds $\{(\mathbf{t}_0^j, \mathbf{t}_1^j)\}_{j\in[n_\mathcal{C}]}$ and $\mathrm{P_S}$ holds $\{\mathbf{t}_{b_j}^j\}_{j\in[n_\mathcal{C}]}$

2. $\mathrm{P_R}$ samples random $\mathbf{w}_{m+\ell} \stackrel{\$}{\leftarrow} \mathbb{F}_2^{k_\mathcal{C}}$, for $\ell \in [s]$, and then constructs a matrix $C \in \mathbb{F}_2^{m' \times n_\mathcal{C}}$ such that each row $\mathbf{c}_i$ is the codeword $\mathcal{C}(\mathbf{w}_i)$. Then, $\mathrm{P_R}$ sends to $\mathrm{P_S}$ the values

$$\mathbf{u}^j = \mathbf{t}_0^j + \mathbf{t}_1^j + \mathbf{c}^j, \quad j \in [n_\mathcal{C}], \tag{1}$$

where $\mathbf{c}^j$ is the j-th column of C.

3. $\mathrm{P_S}$ receives $\mathbf{u}^j \in \mathbb{F}_2^{m'}$ and computes

$$\mathbf{q}^j = b_j \cdot \mathbf{u}^j + \mathbf{t}_{b_j}^j = b_j \cdot \mathbf{c}^j + \mathbf{t}_0^j, \tag{2}$$

that form the columns of an $(m' \times n_\mathcal{C})$ matrix Q. Denoting the rows of T_0, Q by $\mathbf{t}_i, \mathbf{q}_i$, $\mathrm{P_R}$ now holds $\mathbf{c}_i, \mathbf{t}_i$ and $\mathrm{P_S}$ holds $\mathbf{b}, \mathbf{q}_i$ so that

$$\mathbf{q}_i = \mathbf{c}_i \odot \mathbf{b} + \mathbf{t}_i.$$

4. *Consistency check*:
   - $\mathrm{P_S}$ samples s random strings $\{(x_1^{(\ell)}, \dots, x_m^{(\ell)}) \in \mathbb{F}_2^m\}_{\ell\in[s]}$ and sends them to $\mathrm{P_R}$.
   - $\mathrm{P_R}$ computes and sends, for $\ell \in [s]$:

$$\mathbf{t}^{(\ell)} = \sum_{i\in[m]} \mathbf{t}_i \cdot x_i^{(\ell)} + \mathbf{t}_{m+\ell}, \qquad \mathbf{w}^{(\ell)} = \sum_{i\in[m]} \mathbf{w}_i \cdot x_i^{(\ell)} + \mathbf{w}_{m+\ell}$$

   - $\mathrm{P_S}$ computes $\mathbf{q}^{(\ell)} = \sum_{i\in[m]} \mathbf{q}_i \cdot x_i^{(\ell)} + \mathbf{q}_{m+\ell}$, and checks that:

$$\mathbf{t}^{(\ell)} + \mathbf{q}^{(\ell)} = \mathcal{C}(\mathbf{w}^{(\ell)}) \odot \mathbf{b}, \quad \forall \ell \in [s]. \tag{3}$$

   If the check fails, $\mathrm{P_S}$ sends Abort, otherwise continue.

Output: $\mathrm{P_S}$ sets $\forall i \in [m]$ and $\mathbf{w} \in S_i$: $\mathbf{v}_{\mathbf{w},i} = \mathsf{H}(i, \mathbf{q}_i + \mathcal{C}(\mathbf{w}) \odot \mathbf{b})$ and $\mathrm{P_R}$ sets $\forall i \in [m]$: $\mathbf{v}_{w_i, i} = \mathsf{H}(i, \mathbf{t}_i)$.

Fig. 2. An actively secure protocol for $\mathcal{F}_{N\text{-}\mathsf{ROT}+}^{\kappa,m}$, extending $\mathcal{F}_{2\text{-}\mathsf{OT}}^{\kappa,n_\mathcal{C}}$.

included in the linear combination, so s sets of checks are needed to ensure a negligible cheating probability.

Compared with the consistency check of [KOS15] (for the 1-out-of-2 case), our check is simpler as we only require XOR operations instead of multiplications in the finite field $\mathbb{F}_{2^\kappa}$. However, being over $\mathbb{F}_2$ means that we must repeat the check s times, whereas [KOS15] only needs one check; in our case, working in $\mathbb{F}_{2^\kappa}$ would not allow the linear encoding relation to be verified, which is why we use $\mathbb{F}_2$.

We observe that our protocol, minus the final hashing step, is essentially the same as the additively homomorphic commitment protocol from [FJNT16] (which inspired our consistency check). Although our security proof requires quite some extra work to implement OT instead of commitments, it is interesting to see how the same construction can lead to two very different applications with just a small modification. More recently, another scheme [CDD+16] improved upon [FJNT16] by using a linear-time computable consistency check based on a special class of universal hash functions, and constructing a linear-time encodable error-correcting code. These changes can also be applied to our protocol, but it is not clear how efficient these would be in practice, and since we aim for practical (rather than asymptotic) efficiency we do not present this here.

Theorem 1. *Assuming that* H *is a random oracle and* PRG *a pseudo-random generator, the protocol* $N\text{-}\mathsf{ROT}^{\kappa,m}$ *in Fig. 2 securely implements* $\mathcal{F}_{N\text{-}\mathsf{ROT+}}^{\kappa,m}$ *(Fig. 1) in the* $\mathcal{F}_{2\text{-}\mathsf{OT}}$*-hybrid model with computational security parameter* κ *and statistical security parameter* s *against a static malicious adversary.*

Proof of this result can be found in the full version of this work.

The case of a corrupt sender is straightforward and reduces to the security of PRG, similar to previous works [ALSZ16, KOS15]. For a corrupt receiver, the first main challenge is for the simulator to extract the receiver's inputs, $\mathbf{w}_i$, to send to the functionality $\mathcal{F}_{N\text{-}\mathsf{ROT+}}$. This is done by using the values sent during the check to identify a set of positions where the receiver has attempted to 'guess' some bits of the sender's secret, $\mathbf{b}$. Removing these positions from the $\mathbf{c}_i$ values used by $\mathsf{P_R}$ leaves behind an incomplete codeword, which can be erasure decoded to recover the message.

After decoding the inputs, the simulator must then respond to random oracle queries made by the environment. We do this in an *optimistic* manner, meaning, we do not abort if conflicting queries are made, but answer at random in that case; the environment may not always notice this inaccurate behaviour if the sender did not learn all N outputs from the OTs. This allows us to obtain a security bound that depends on N', the *maximum* number of outputs learnt by $\mathsf{P_S}$ in any OT, rather than N, which may be exponential in κ.

Instantiating the Code. It remains to instantiate the binary linear code, $\mathcal{C}$, to obtain a 1-out-of-N random OT protocol for a desired power of two choice of N. As well as the repetition code (for 1-out-of-2 OT), we suggest a more efficient form of the Walsh-Hadamard code for $N \leq 512$; a binary Golay code

Table 1. Parameters for various choices of code

Code	N	Length	Distance/Security
Repetition [IKNP03]	2	128	128
Walsh-Hadamard [KK13]	≤256	256	128
Punctured Walsh-Hadamard	≤512	256	128
Binary Golay	2048	384	128
BCH-511	2^{76}	511	171
BCH-1023	2^{443}	1023	128

for $N = 2048$; and BCH codes for values of N that are exponential in the security parameter. The parameters of these codes are presented in Table 1; note that the code length determines exactly the amount of communication required per extended OT. We obtained the generator matrices for all of these codes using Sage[3]. For further details of the constructions, see the full version.

4   Security in the Standard Model

In this section we consider the case of non-random 1-out-of-N OT. In this protocol (Fig. 3), we remove the random oracle assumption and prove security in the standard model. Similarly to [ALSZ15], we need a stronger version of correlation robustness than that given in [IKNP03], and require that the secret correlation **b** comes from a distribution of min-entropy k and in addition is multiplied by a codeword in the binary linear code $\mathcal{C}$.

Protocol N-$\mathsf{OT}^{\kappa,m}$ (Standard Model)

INPUT: As in Protocol N-$\mathsf{ROT}^{\kappa,m}$.
REQUIRE: $\mathsf{H} : \mathbb{F}_2^{n_c} \to \mathbb{F}_2^\kappa$ a k-min-entropy strongly $\mathcal{C}$-correlation-robust and $\mathsf{PRG} : \{0,1\}^\kappa \to \{0,1\}^{m'}$.

Init: As in Protocol N-$\mathsf{ROT}^{\kappa,m}$.
Extend: As in Protocol N-$\mathsf{ROT}^{\kappa,m}$.
Output: P_S sends, $\forall i \in [m]$ and $0 \le w < N-1$: $\mathbf{y}_{w,i} = \mathbf{v}_{w,i} + \mathsf{H}(\mathbf{q}_i + \mathcal{C}(\mathbf{w}) \odot \mathbf{b})$.
     P_R recovers, $\forall i \in [m]$: $\mathbf{v}_{w_i,i} = \mathbf{y}_{w_i,i} + \mathsf{H}(\mathbf{t}_i)$.

Fig. 3. An implementation of $\mathcal{F}_{N\text{-}\mathsf{OT}}^{\kappa,m}$ extending $\mathcal{F}_{2\text{-}\mathsf{OT}}^{\kappa,n_c}$ in the Standard Model.

---

[3] http://www.sagemath.org.

Definition 1 (k-min-entropy code correlation robustness). *Let χ be a distribution on $\mathbb{F}_2^{n_C}$ with min-entropy k and $\mathcal{C}$ be an $[n_C, k_C, d_C]$ binary linear code. An efficiently computable function $\mathsf{H} : \mathbb{F}_2^{n_C} \to \mathbb{F}_2^{\kappa}$ is said to be k-min-entropy $\mathcal{C}$ -correlation robust if it holds that:*

$$\{\mathbf{t}_i, \mathsf{H}(\mathbf{t}_i + \mathbf{c} \odot \mathbf{b})\}_{i \in [m], \mathbf{c} \in C} \stackrel{c}{\equiv} \mathcal{U}^{m \cdot n_c + (m + |C|) \cdot \kappa},$$

where $\mathbf{b} \stackrel{\$}{\leftarrow} \chi$ and $\mathbf{t}_1, \ldots, \mathbf{t}_m \in \mathbb{F}_2^{n_C}$ are independent and uniformly distributed.

Similarly, $\mathsf{H}(\cdot)$ is said to be k-min-entropy strongly $\mathcal{C}$ -correlation robust if it holds that:

$$\{\mathsf{H}(\mathbf{t}_i + \mathbf{c} \odot \mathbf{b})\}_{i \in [m], \mathbf{c} \in C} \stackrel{c}{\equiv} \mathcal{U}^{(m + |C|) \cdot \kappa},$$

where $\mathbf{b} \stackrel{\$}{\leftarrow} \chi$, for any distribution of the $\{\mathbf{t}_i\}_{i \in [m]}$.

Notice that in the values used to mask P_S's inputs, it is the receiver that effectively chooses the $\mathbf{t}_j$'s, and they can not only choose these values non-uniformly, but even maliciously. This is the reason why we need a *strong* code-correlation robust hash function.

We claim that if H is a k-min-entropy strongly correlation robust function for all $n_C - s \leq k \leq n_C$, then the protocol is secure in the standard model. For further discussion on parameter choices regarding this assumption, see the full version.

Theorem 2. *Assuming that H is k-min-entropy strongly code-correlation robust for all $k \in \{n_C - s, \ldots, n_C\}$, and PRG is a pseudo-random generator, the protocol $N\text{-}\mathsf{OT}^{\kappa, m}$ in Fig. 3 securely implements $\mathcal{F}_{N\text{-}\mathsf{OT}}^{\kappa, m}$ in the $\mathcal{F}_{2\text{-}\mathsf{OT}}$-hybrid model against a static malicious adversary.*

Proof of this result can be found in the full version.

5   Application to Private Set Intersection

We now show how to apply the 1-out-of-N random OT extension protocol to increase the efficiency and obtain stronger security guarantees in existing private set intersection (PSI) protocols. We describe a simpler and more efficient private set inclusion protocol with active security, which is used as a key component of the most efficient passively secure PSI protocols.

5.1   Private Set Inclusion

A core building block of OT-based PSI protocols is a *private set inclusion* protocol, where party P_A has input $a \in \{0, 1\}^k$, party P_B has input a set $B \subset \{0, 1\}^k$ and the parties wish to learn whether $a \in B$. Note that the special case of $|B| = 1$ is a private equality test.

The previous most efficient protocol [PSSZ15, Sect. 6.1] requires $t = k/8$ executions of 1-out-of-256 random OT, and uses the KK protocol with length

256 Walsh-Hadamard codes. However, with the observation that our random OT protocol can be used for exponentially large values of N, we can in fact choose $N = 2^k$ and perform a private set inclusion with just a single 1-out-of-N random OT. This is possible because the OT sender is only required to learn one of the random OT outputs in order to run the set inclusion protocol.

The protocol, shown in Fig. 4, is very simple: P_A inputs their value a as the receiver's choice in a 1-out-of-N random OT, and P_B inputs each of their values $b \in B$ to obtain $|B|$ of the sender's random outputs. Thus, P_A learns a random value r_a and P_B learns a set of random values $R = \{r_b\}_{b \in B}$. P_B randomly permutes R and sends this to P_A, who checks whether $r_a \in R$ to determine the result (and can send this to P_B if desired).

Since P_A only learns one of the random OT outputs initially, all other possible elements of the set R are uniformly random so do not leak any information on P_B's input. Note that because our 1-out-of-N OT protocol is actively secure, we actually obtain an actively secure private set inclusion protocol (although this does not seem to suffice to make the PSI protocol of [PSSZ15] actively secure).

Protocol Π_{SetInc}

Inputs: P_A has input $a \in \{0,1\}^k$ and P_B has input $B \subset \{0,1\}^k$.

1. The parties run $\mathcal{F}^{s,1}_{N\text{-}\mathsf{ROT}+}$ with $N = 2^k$, where P_A plays receiver with input a.
2. P_A receives $r_a \in \{0,1\}^s$
3. For each $b \in B$, P_B calls $\mathcal{F}_{N\text{-}\mathsf{ROT}+}$ with $(\mathsf{SenderOutput}, 0, b)$ to obtain r_b.
4. P_B randomly permutes and sends $R = \{r_b\}_{b \in B}$ to P_A.
5. P_A outputs 1 if $r_a \in R$ and 0 otherwise.

Fig. 4. Private set inclusion protocol

The complete security proof and functionality that we implement is given in the full version.

Efficiency. The cost of the above protocol is precisely the cost of 1-out-of-N random OT, plus sending $s \cdot |B|$ bits. If the protocol is run in a large batch

Table 2. Comparing the communication cost of private set inclusion subprotocols on k-bit strings and size $|B|$ sets with statistical security s.

k	Cost with BCH (bit)	Cost with W-H (bit)
32	$467 + s \cdot |B|$	$1024 + s \cdot |B|$
64	$499 + s \cdot |B|$	$2048 + s \cdot |B|$
128	$708 + s \cdot |B|$	$4096 + s \cdot |B|$

using $\mathcal{F}_{N\text{-ROT+}}^{\kappa,m}$ for large m (which is possible for the application to private set intersection) then this gives an amortized cost of $n_\mathcal{C} + s \cdot |B|$ bits per execution, where $n_\mathcal{C}$ is the length of the code. The costs for this when instantiated with BCH codes (as described previously) are illustrated for various choices of k in Table 2, and compared with the Walsh-Hadamard code used in [PSSZ15]. In practice, the set size used in the set inclusion subprotocol for PSI in [PSSZ15] is around $|B| = 20$. For a large item length of $k = 128$ bits, and $s = 40$-bit statistical security, this gives a $3.3\times$ reduction in communication for the dominant component of PSI.

6   Implementation

We now evaluate the complexity of our random OT protocol, and compare its performance to a passively secure variant by analysing implementation results.

Complexity Analysis. The main overhead introduced by our protocol to produce m OTs, compared with the passively secure KK protocol, is the computation of $m \cdot s$ XORs (on $n_\mathcal{C}$-bit strings) by each party, and the communication of $s \cdot m$ random bits from the sender to receiver, followed by $s \cdot (n_\mathcal{C} + k_\mathcal{C})$ bits in the other direction, in the consistency check. However, the $s \cdot m$ bits can be reduced to κ by having P_S send only a single random seed for these values, and expanding the seed using a PRG.

Outside of the consistency check, the main protocol costs are the encodings, hash function evaluations and the $n_\mathcal{C}$ bits that are sent by P_R for each extended OT. Of course, the sender's computational cost also highly depends on the number of random OT outputs that are desired.

Implementation. We evaluated our protocol on two machines running over a 1 Gbps local network, and also simulated a WAN environment with 50 Mbps bandwidth and 100 ms round-trip latency to model a real-life scenario over the Internet. All benchmarks have been run on modern Core i7 machines at 2–3 GHz.

Our implementation is in plain C, and uses the SimpleOT [CO15] oblivious transfer software[4] to run the base OTs, and BLAKE2 [ANWOW13] for hashing, as this provides fast hashing on short inputs. Otherwise, it does not rely on any other software and is available in the public domain. The executable occupies $280K$.

The core protocol covers roughly 200 lines of C code. It mostly runs on single thread, except we use OpenMP to parallelize the encodings and hash function evaluations, which are the computational bottlenecks of the protocol. We fix the computational security parameter $\kappa = 128$ and statistical security parameter $s = 40$. We used Intel AVX instructions to efficiently implement vector addition, componentwise product, and matrix transposition. Encoding of the binary linear code is implemented with multiplication by the generator matrix.
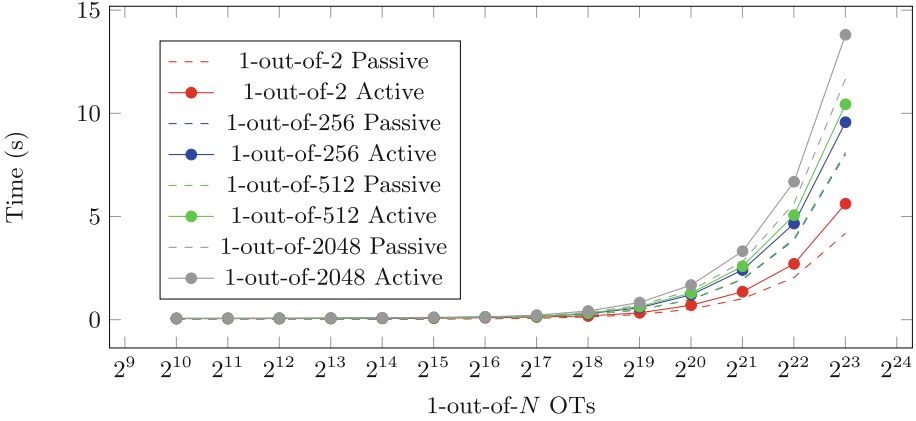
---

[4] http://users-cs.au.dk/orlandi/simpleOT/.

Fig. 5. Benchmarking different 1-out-of-N random OTs in LAN environment; average time for 20 runs.

Table 3. Data transmitted per OT and runtimes in seconds for 2^{23} OTs (LAN) or 2^{20} OTs (WAN), for several choices of N

Setting	$N = 2$	256	512	2048	2^{76}
Comms. (bit)	128	256	256	384	512
LAN, passive	4.1812	8.0260	8.1193	11.6642	23.4738
LAN, active	5.6191	9.5693	10.4379	13.8065	25.4001
WAN, passive	27.3982	54.2414	54.274	81.0548	108.89
WAN, active	27.882	54.7445	54.8189	81.6644	109.44

Our results for 1-out-of-N random OT for the small sized codes (repetition, Walsh-Hadamard, punctured Walsh-Hadamard and binary Golay) in the LAN setting can be seen in Fig. 5, for varying numbers of OTs. Table 3 compares the performance of the active and passively secure variants in both LAN and WAN settings, including the BCH-511 code, which could be used for the private set intersection application. We see that the overhead of active security is around 20–30% of the passive protocol over a LAN, and less than 5% in the WAN setting. This fits with the fact that the main cost of the check is computation, which is less significant in a WAN. The table also gives the amount of communication required for each choice of N, which shows that this reflects the main total cost of the protocol. Encoding of larger BCH codes (for $N = 2^{76}$) does have a noticeable effect in a LAN, though: here, BCH runtimes are 3 times higher than Walsh-Hadamard, but only have twice the communication cost. We expect that this could be improved by a more sophisticated encoding algorithm, rather than naive multiplication by the generator matrix.

References

[ALSZ13] Asharov, G., Lindell, Y., Schneider, T., Zohner, M.: More efficient oblivious transfer and extensions for faster secure computation. In: ACM Conference on Computer and Communications Security, CCS 2013, pp. 535–548 (2013)

[ALSZ15] Asharov, G., Lindell, Y., Schneider, T., Zohner, M.: More efficient oblivious transfer extensions with security for malicious adversaries. In: Proceedings of Advances in Cryptology - EUROCRYPT 2015, Sofia, Bulgaria, Part I, pp. 673–701, 26–30 April 2015

[ALSZ16] Asharov, G., Lindell, Y., Schneider, T., Zohner, M.: More efficient oblivious transfer extensions. J. Cryptol. 1–54 (2016)

[ANWOW13] Aumasson, J.-P., Neves, S., Wilcox-O'Hearn, Z., Winnerlein, C.: BLAKE2: simpler, smaller, fast as MD5. In: Jacobson, M., Locasto, M., Mohassel, P., Safavi-Naini, R. (eds.) ACNS 2013. LNCS, vol. 7954, pp. 119–135. Springer, Heidelberg (2013). doi:10.1007/978-3-642-38980-1_8

[Bea96] Beaver, D.: Correlated pseudorandomness and the complexity of private computations. In: Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing, pp. 479–488. ACM (1996)

[CDD+16] Cascudo, I., Damgård, I., David, B., Döttling, N., Nielsen, J.B.: Rate-1, linear time and additively homomorphic UC commitments. In: Proceedings of Advances in Cryptology - CRYPTO 2016, Santa Barbara, CA, USA, Part III, pp. 179–207, 14–18 August 2016

[CO15] Chou, T., Orlandi, C.: The simplest protocol for oblivious transfer. In: Progress in Cryptology - LATINCRYPT 2015, Guadalajara, Mexico, pp. 40–58, 23–26 August 2015

[EGL85] Even, S., Goldreich, O., Lempel, A.: A randomized protocol for signing contracts. Commun. ACM **28**(6), 637–647 (1985)

[FJNT16] Frederiksen, T.K., Jakobsen, T.P., Nielsen, J.B., Trifiletti, R.: On the complexity of additively homomorphic UC commitments. In: Proceedings of Theory of Cryptography - TCC 2016-A, Tel Aviv, Israel, 10–13 January 2016, Part I, pp. 542–565 (2016)

[IKNP03] Ishai, Y., Kilian, J., Nissim, K., Petrank, E.: Extending oblivious transfers efficiently. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 145–161. Springer, Heidelberg (2003). doi:10.1007/978-3-540-45146-4_9

[IR89] Impagliazzo, R., Rudich, S.: Limits on the provable consequences of one-way permutations. In: Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing, pp. 44–61. ACM (1989)

[KK13] Kolesnikov, V., Kumaresan, R.: Improved OT extension for transferring short secrets. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013. LNCS, vol. 8043, pp. 54–70. Springer, Heidelberg (2013). doi:10.1007/978-3-642-40084-1_4

[KKRT16] Kolesnikov, V., Kumaresan, R., Rosulek, M., Trieu, N.: Efficient batched oblivious PRF with applications to private set intersection. In: ACM Conference on Computer and Communications Security, CCS (2016)

[KOS15] Keller, M., Orsini, E., Scholl, P.: Actively secure OT extension with optimal overhead. In: Advances in Cryptology - CRYPTO Santa Barbara, CA, USA, pp. 724–741, 16–20 August 2015

[KOS16] Keller, M., Orsini, E., Scholl, P.: MASCOT: faster malicious arithmetic secure computation with oblivious transfer. In: ACM Conference on Computer and Communications Security, Vienna, Austria, pp. 830–842 (2016)

[Lam16] Lambæk, M.: Breaking and fixing private set intersection protocols. Cryptology ePrint Archive, Report 2016/665 (2016)

[LOS14] Larraia, E., Orsini, E., Smart, N.P.: Dishonest majority multi-party computation for binary circuits. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014. LNCS, vol. 8617, pp. 495–512. Springer, Heidelberg (2014). doi:10.1007/978-3-662-44381-1_28

[NNOB12] Nielsen, J.B., Nordholt, P.S., Orlandi, C., Burra, S.S.: A new approach to practical active-secure two-party computation. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 681–700. Springer, Heidelberg (2012). doi:10.1007/978-3-642-32009-5_40

[NP99] Naor, M., Pinkas, B.: Oblivious transfer and polynomial evaluation. In: Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing, Atlanta, GA, USA, pp. 245–254 (1999)

[PSSZ15] Pinkas, B., Schneider, T., Segev, G., Zohner, M.: Phasing: private set intersection using permutation-based hashing. In: 24th USENIX Security Symposium, Washington, D.C., USA, 12–14 August 2015, pp. 515–530 (2015)

[PSZ14] Pinkas, B., Schneider, T., Zohner, M.: Faster private set intersection based on OT extension. In: 23rd USENIX Security Symposium, San Diego, CA, pp. 797–812, August 2014

[PVW08] Peikert, C., Vaikuntanathan, V., Waters, B.: A framework for efficient and composable oblivious transfer. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 554–571. Springer, Heidelberg (2008). doi:10.1007/978-3-540-85174-5_31

[Rab81] Rabin, M.O.: How to exchange secrets with oblivious transfer (1981)

[SSR08] Shankar, B., Srinathan, K., Rangan, C.P.: Alternative protocols for generalized oblivious transfer. In: 9th International Conference on Distributed Computing and Networking, ICDCN (2008)

[Yao82] Yao, A.C.-C.: Protocols for secure computations (extended abstract). In: 23rd Annual Symposium on Foundations of Computer Science, pp. 160–164 (1982)