



Lower-Bounds on Public-Key Operations in PIR

Jesko Dujmovic^{1,2(✉)} and Mohammad Hajiabadi³

¹ Helmholtz Center for Information Security (CISPA), Saarbrücken, Germany

jesko.dujmovic@cispa.de

² Saarbrücken Graduate School of Computer Science, Saarbrücken, Germany

³ University of Waterloo, Waterloo, Canada

mdhajiabadi@uwaterloo.ca

Abstract. Private information retrieval (PIR) is a fundamental cryptographic primitive that allows a user to fetch a database entry without revealing to the server which database entry it learns. PIR becomes non-trivial if the server communication is less than the database size. We show that building (even) very weak forms of PIR protocols requires that the amount of public-key operations scale linearly in the database size.

We then use this bound to examine the related problem of communication efficient oblivious transfer (OT) extension.

Oblivious transfer is a crucial building block in secure multi-party computation (MPC). In most MPC protocols, OT invocations are the main bottleneck in terms of computation and communication. OT extension techniques allow one to minimize the number of public-key operations in MPC protocols. One drawback of all existing OT extension protocols is their communication overhead. In particular, the sender's communication is roughly double what is information-theoretically optimal.

We show that OT extension with close to optimal sender communication is impossible, illustrating that the communication overhead is inherent. Our techniques go much further; we can show many lower bounds on communication-efficient MPC. E.g. we prove that to build high-rate string OT with generic groups, the sender needs to do linearly many group operations.

Keywords: Private Information Retrieval · PIR · Lower Bounds · OT Extension · Generic group model · GGM · Communication Complexity

1 Introduction

Secure Multi-Party Computation (MPC), allows two parties to jointly evaluate a function f while leaking nothing about their input to the other party beyond the output of f . A central goal of modern cryptography is to construct efficient MPC protocols. This goal is important not only from a theoretical but also from a practical viewpoint. The computational efficiency of all MPC protocols

is typically bottlenecked by public-key operations (e.g., group operations, oblivious transfers (OT)). OT extension is a technique toward reducing public-key operations [3, 29], allowing one to get the results of many OTs at the cost of performing only a few OTs and some symmetric-key operations. This technique has revolutionized the practical development of MPC, leading to protocols which employ a small number of public-key operations for sophisticated tasks.

A significant limitation of existing OT extension techniques are their high communication cost: for performing ℓ 1-out-of-2 bit OTs, the sender communicates at least 2ℓ bits. This severely limits the use of OT extension in MPC settings where a low amount of communication is required ‘by design’ (e.g., Private-Information Retrieval (PIR)). The overarching goal of our paper is to understand communication-computation tradeoffs in such MPC settings. We show that in many such situations, performing many public-key operations is provably unavoidable. To put our results in context, let us illustrate how communication efficient OT would impact private information retrieval.

Private-Information Retrieval (PIR). Private information retrieval [15, 33] is a fundamental cryptographic primitive that allows a user to fetch a database entry without revealing to the server which database entry it learns. PIR becomes non-trivial if the server-to-user communication is strictly less than the database size (and ideally growing sub-linearly or even polylogarithmically in the database size). In some applications, one may need extra properties, such as an overall (as opposed to server-to-user) sub-linear communication or server privacy. Throughout the paper, we require neither of these unless otherwise stated. Since we prove lower bounds, this makes our results stronger. A truly efficient PIR protocol has significant real-world applications such as private certificate retrieval or private DNS lookups. By now, we know how to build PIR with communication complexity polylogarithmic in n from a wide range of assumptions [10, 12, 19, 27, 30]. While the amount of communication is attractively low, the computation overhead leaves much to be desired. [19] demonstrates that a communication efficient OT would imply PIR.

Computational Complexity of PIR. In (single-server) PIR protocols, the running time of the server cannot be sub-linear in n , the database size, without preprocessing [4]. If it was sub-linear the server could not read all the entries, leaking information about the user’s index i . Faced with this lower-bound, and the fact that PIR requires public-key assumptions [18], one may wish to settle for the next best thing: making the number of public-key operations independent of n . Somehow curiously, in all existing PIR protocols based on Diffie-Hellman or OT related assumptions [12, 19], not only the server’s running time, but the number of public-key operations performed by the server grows at least linearly with n . There is no evidence, however, if this is inherent, and in fact, it has remained an open problem whether one can build a PIR protocol where the number of public-key operations is sub-linear in n .

Is it possible construct PIR with a sub-linear amount of public-key operations and an arbitrarily large number of symmetric-key operations?

OT Extension. A major tool used for minimizing computation is OT extension. Existing OT extension techniques induce at least a linear amount of communication for the sender, making them unsuitable for PIR applications. Specifically, under existing constructions, an extended OT sender needs to communicate at least as many bits as its total input length. Beaver’s seminal construction [3] works by encoding all the sender messages into a garbled circuit, which the receiver can evaluate only on the labels that correspond to her choice bits; the IKNP protocol [29] establishes correlated randomness between the sender and the receiver, allowing the sender to XOR his messages with the corresponding masks such that the receiver can only de-mask the correct messages. In both these protocols, the sender’s outgoing protocol messages information-theoretically determine the entire sender’s input, causing the communication overhead. This state of affairs raises the following natural question.

Is it possible to construct OT extension where the sender communication is close to optimal?

In the above question, by ‘optimal sender communication’ we mean the best information-theoretically achievable communication: which is ℓ bits for the sender for performing ℓ 1-out-2 single-bit OTs. Since OT extension is crucially used in many MPC protocols, understanding its communication complexity is of both practical and theoretical value.

Having optimal sender communication for OT extension is reminiscent of rate-1 string OT: building 1-out-of-2 string OTs for a pair of ℓ -bit strings, where (roughly speaking) the sender communication grows as $\ell + \lambda$ (as opposed to $2\ell + \lambda$), where λ is the security parameter. Two-round rate-1 OT has found a number of applications, notably in the construction of PIR protocols with poly-logarithmic communication [1, 8, 12, 19, 21, 30]. We know how to build rate-1 OT from a wide variety of assumptions [1, 12, 19, 21, 30], but all these constructions make at least a linear number of public-key operations. In particular, computational efficiency (e.g., sub-linear number of public-key operations) and communication efficiency (e.g., sub-linear communication) seem to have largely been in conflict with each other — for reasons we have not been to justify so far. The goal of our paper is to elucidate this conflicting situation.

1.1 Our Results

We answer both of the above questions, and several other related ones, negatively. In particular, we give a lower-bound on the number of public-key operations that need to be performed by servers in PIR protocols, and use this lower-bound to derive similar results for related primitives. Our core idea is based on a compilation technique that allows one to remove public-key operation queries from a PIR protocol at the cost of proportionally increasing the communication complexity in the public-key operation free protocol. As applications of our main theorem, we obtain results that settle several open problems in MPC.

In the statement below, by an **SO** oracle we mean a *simulatable oracle*: roughly speaking, one that can be simulated via lazy sampling. Examples of

such oracles include generic-group oracles, public-key encryption oracles, etc. See Sect. 2 for more details. Also, we use the term a “party’s SO bit complexity” to indicate the total bit size of all SO queries made by the party.

Theorem 1 (Informal Main Theorem). *If there exists a PIR for n -bit databases with oracle access to simulatable oracle SO, arbitrary oracle O, server communication of $\eta < cn$ for $c < 1$, $r \in o(n)$ rounds of interaction with the user, and $q \in o(n)$ bits of communication with the SO oracle, then there exists a PIR with oracle access to O, server communication $\bar{\eta} \leq \bar{c}n$ for $\bar{c} < 1$, and no calls to SO.*

We derive the following corollary.

Corollary 1. *There exists no n -bit PIR protocol built solely¹ from a simulatable oracle SO and a random oracle O with $o(n)$ round complexity, with $o(n)$ server’s SO bit complexity and with $\eta \leq cn$ server’s communication for $c < 1$.*

For example, letting SO be a generic group oracle (GGM), we rule out all n -bit PIR protocols that have $o(n)$ rounds and where the server’s communication and the server’s total number of GGM queries are, respectively, cn and $o(n)$ for $c < 1$. This holds irrespective of the number of RO queries the protocol is allowed to make. This closely matches the known upper-bounds, as [19, 35] give n -bit PIR protocols based on the DDH assumption with server communication of $O(\lambda)$ and with the sever making $O(n)$ group operations.

The strength of the main theorem lies in its flexibility in instantiating the oracle SO: for example, one may let SO be an FHE oracle, and obtain similar results as long as the amount of server’s communication with the FHE oracle respects the bounds. The work of [35] shows how to obtain PIR generically from (additively) homomorphic encryption, where the sever performs $O(n)$ homomorphic additions. Our work shows that this is close to optimal.

We will show that our computational lower-bounds for PIRs give rise to communication lower-bounds for OT extension.

Corollary 2 (OT Extension: Communication Lower-Bounds). *There exists no ℓ -batch k -bit OT extension protocol with round complexity $r \in o(k\ell)$ and with server communication $\eta < c2k\ell$ for $c < 1$.*

In the above corollary, by ℓ -batch k -bit OT we mean performing ℓ OTs for pairs of k -bit strings. The IKNP protocol [29] in the ℓ single-bit OT case achieves sender communication of $> 2\ell$. Our result shows that the IKNP’s sender communication complexity is close to optimal.

Finally, we relate PIR to other MPC protocols such as rate-1 OT to arrive at the following corollary. For brevity, we describe the statements when O is the GGM oracle, and only for rate-1 OT. In fact, we can show that achieving any rate strictly greater than $1/2$ (measured as the information-theoretically optimal sender communication size divided by the sender’s communication size in the actual protocol) requires making an almost linear number of group operations.

¹ By “solely” we mean that the parties also have access to a PSPACE oracle. This stops the party from using any hardness assumptions other than the ones provided by the oracles.

Corollary 3 (String OT Corollary). *There exists no ℓ -bit string OT protocol in the GGM+RO model with sender communication of $\eta \leq c2\ell$ and $o(\ell)$ calls to the generic group for $c < 1$.*

Similar results can be proven about unbalanced PSI, see Sect. 5.2 for details.

2 Technical Overview

First, we will give a quick example of how to simulate a generic group efficiently to illustrate a simulatable oracle. Then, we sketch the proof of the main theorem. We proceed by showing why the main theorem is useful by demonstrating a few MPC protocols which imply non-trivial PIR, allowing us to apply our lower-bounds to them.

2.1 Generic Group Model

A generic group of order p is the group \mathbb{Z}_p together with the random injective encoding function $\sigma: \mathbb{Z}_p \rightarrow S$, where $S = \{0, \dots, p-1\}$. The algorithms can access this group via the oracle **Add** which decodes two encoded elements, computes a linear combination of them and gives back the encoded result. More formally, $\text{Add}: \mathbb{Z}_p^2 \times S^2 \rightarrow S$, $(a_1, a_2, \ell_1, \ell_2) \mapsto \sigma(a_1v_1 + a_2v_2)$, where $v_i = \sigma^{-1}(\ell_i)$ for $i \in \{1, 2\}$.² This way the algorithms interacting with the oracle can only access encodings instead of the real group.

To simulate a GGM oracle efficiently, the simulator dynamically generates the encoding function σ . More specifically, it maintains a partial set L of \mathbb{Z}_p -label pairs sampled by the simulator so far. (It is initially empty.) Whenever a query $(a_1, a_2, \ell_1, \ell_2)$ is made, the simulator checks if $(*, \ell_1) \in L$ (meaning that if for some v_1 , $(v_1, \ell_1) \in L$); if not, the simulator samples a random v_1 from \mathbb{Z}_p subject to $(v_1, *) \notin L$, and adds (v_1, ℓ_1) to L . The simulator does the same thing for ℓ_2 . Now assuming $(v_1, \ell_1) \in L$ and $(v_2, \ell_2) \in L$, letting $a_3 = a_1v_1 + a_2v_2$ if $(a_3, \ell_3) \in L$ for some ℓ_3 , the simulator responds to the query with ℓ_3 ; else, the simulator samples a random ℓ_3 subject to $(*, \ell_3) \notin L$, adds (a_3, ℓ_3) to L , and responds to the query with ℓ_3 .

Other simulatable oracles that are useful in our main theorem are black-box oblivious transfer [26], black-box public-key encryption [26], and ideal obfuscation [31].

2.2 Proof Sketch of Main Theorem

The observation that leads to the main theorem is that in PIR protocols the server does not need to be secure. This means the protocol would still be secure if the user learned all of the servers oracle queries. Also, the user knows all of

² One may set S to be a random subset of size p of a larger set $\{0, 1\}^u$ for $u > \log p$. Our analysis will remain unchanged, so we simply assume $S = \{0, \dots, p-1\}$. See [41] for differences between various models.

its own oracle queries. Therefore, the user can just simulate the oracle for both of the parties. For this the server just has to send all of its queries to the user. This modification increases the server communication roughly by the amount that the server would have communicated to the oracle. We will prove that this transformation preserves user security. Since we require the oracle communication to be in $o(n)$ and the server communication to be $< cn$ for some $c < 1$, the modified protocol will have no oracle queries and the server communication will be $< \bar{c}n$ for $\bar{c} < 1$. Moreover, in the actual compiled protocol, to enable the compiled user to distinguish query messages from normal protocol messages, we append a flag bit to the end of each server's protocol messages — causing a dependency on r , the number of rounds, in Theorem 1.

2.3 PIR Related Protocols

Now we exhibit a few protocols which imply non-trivial PIR, allowing us to apply our impossibility results for PIR to these protocols.

Low Sender-Communication OT. We show that an ℓ -batch k -bit OT protocol with sender communication $< c2k\ell$ for $c < 1$ implies a PIR. The transformation works as follows: suppose w.l.o.g the database size is $2k\ell$. The server runs the OT protocol and encodes the first half of the database into the messages $(m_i^{(0)})_{i \in [\ell]}$ and the second half into $(m_i^{(1)})_{i \in [\ell]}$. Now, if a user wants to look up the j -th element of the database, it acquires $(m_i^{(0)})_{i \in [\ell]}$ if $j \leq k\ell$ and $(m_i^{(1)})_{i \in [\ell]}$ otherwise. The database entry that the client wants to learn is contained in the OT output. The server communication is $< c2k\ell$ and the user's input j is hidden from the server by the OT's receiver security.

Low Total Communication OT. We next attempt to prove lower-bounds for the case where the OT protocol has low total communication (as opposed to low sender communication). For convenience we focus on ℓ -batch single-bit OT. By low total communication we mean an amount that is close the information-theoretically optimal communication, which is 2ℓ bits. The techniques for low sender-communication as above do not apply outright because the sender might cause the bulk of communication itself, an amount close to 2ℓ bits. We get around this issue via the following intuitive idea: When an OT protocol has low total communication it must have either low sender communication, from which we already showed how to obtain a non-trivial PIR, or it must have low receiver communication. In the latter case, we swap the roles of the two different parties with a role-flipping trick implicitly in [29] and explicitly in [39]. This role flipping trick turns the low receiver communication into low sender communication, which we can then turn into a non-trivial PIR. Figure 1 depicts the construction for $\ell = 1$. In essence, we show how to turn a communication-efficient OT into a sender-communication-efficient OT by introducing an additional round.

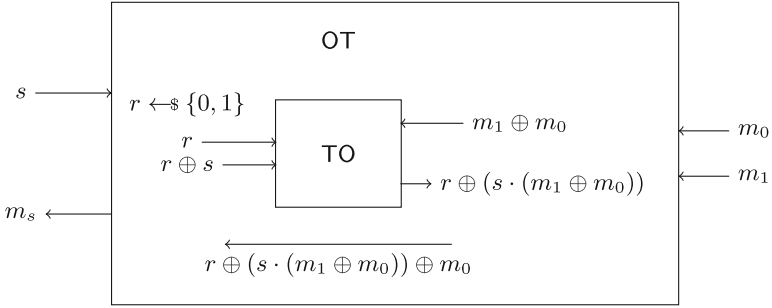


Fig. 1. Similar to [39] a visual representation of how to build an oblivious transfer OT from an oblivious transfer TO that goes in the opposite direction.

OT Extension. The above results immediately imply communication lower bounds for OT extension: showing that performing OT extension for ℓ -batch k -bit OTs with $c2k\ell$ bits of sender communication for $c < 1$, and with an $O(\lambda)$ (and even $o(k\ell)$) number of public-key operations is impossible.

Unbalanced PSI. Private set intersection (PSI) is an MPC protocol between two parties each holding a set and the party called the receiver learns the intersection of the two sets. No other information should be revealed to is to any of the parties. In Unbalanced PSI, a special case of PSI, the receiver set is much smaller than the sender set and the communication should only scale with the receiver set. To build a non-trivial PIR from such a protocol, for a client index i , the client sets $x := i$ (padding it out if necessary), and for a database DB, the server forms the set $\{i \mid \text{DB}[i] = 1\}$. An answer to $x \in^? S$ reveals $\text{DB}[i]$. This observation allows us to prove that in unbalanced PSI with sub-linear communication, the receiver should perform close to linear public-key operations. This shows that the large number of public-key operations used in unbalanced PSI protocols of [12, 19, 21] is inherent.

Non-Trivial PIR implies Oblivious Transfer. We know that non-trivial PIR implies oblivious transfer [18], and this is used to get our final impossibility results. The transformation utilizes the user security of the PIR protocol and deploys a compression argument to argue information loss. The entropy garnered from the information loss is then fed into a randomness extractor, the output of which can be used to guarantee sender security in the resulting oblivious transfer protocol.

2.4 Oracles

Notice that all these transformations only make black-box use of the protocols they transform. This means that if the starting protocol uses some oracle other than the one we want to remove, say the random oracle [6], then the resulting protocol will also use the random oracle even if we remove other oracle queries.

3 Related Work

Technique. Our ‘compilation-out’ techniques bear some similarities to ideas used by Gennaro and Trevisan [25] for giving lower-bounds on the query complexity of PRGs from OWPs. Essentially, they showed that if the number of queries is ‘small’, they can be encoded as part of the input, hence getting rid of OWP calls in the construction of a PRG. Gennaro et al. [24] built on that idea to give lower-bounds on the efficiency of various cryptographic primitives. These works mostly deal with non-interactive primitives. Our techniques are used in a different way in that we leverage the lack of security requirements for a party to get rid of oracle calls of an interactive protocol.

Private-Information Retrieval. In all but this section of the paper we talk about non-trivial single-server private information retrieval, which is why we will sometimes leave out the descriptor “single-server”. Traditionally, PIR [15, 33] is a protocol between one user and possibly multiple servers. Just like in the non-trivial single-server case the user with an index $i \in [n]$ learns the i -th element of a database $\text{DB} \in \{0, 1\}^n$ held by all the servers without disclosing i to the servers. In other words, if the caveat of non-triviality is not made, then not only the server communication needs to be sub-linear in n , but also the total communication. Thus, a non-trivial PIR is a weaker primitive than PIR. Multi-server PIR protocols assume some kind of non-collusion between the servers, which allows them to achieve statistical security as opposed to computational security in the single-server case.

By now PIR is a well studied primitive; here we focus on the single server setting. We know how to build PIR with communication complexity polylogarithmic in n from a wide range of assumptions [10, 12, 19, 30]. In the last few years, we have also made progress towards practically efficient PIR [16, 17, 28, 32, 37, 42, 43] and asymptotically efficient PIR [7, 11, 34] when the server and the client (or sometimes only the server) are allowed to preprocess the database. We even know some lower bounds for different preprocessing settings [5, 16, 17, 36, 40]. Preprocessing, however, only makes sense if one expects to use the PIR multiple times. We focus on a single query PIR and uniquely distinguish between public-key and symmetric key operations.

OT Extension. The intuition behind OT extension is that it only uses very few calls to an OT functionality to implement many more OTs. An equivalent description is that an OT extension protocol is an OT protocol that can make calls to an OT functionality. The protocol becomes valuable if the number of OT calls in the protocol is much less than the ‘size’ of the OT being implemented. OT calls are typically modelled as oracle calls to an OT functionality or the OT hybrid model.

Beaver [3] constructed the first OT extension protocol, which makes non-black-box use of pseudorandom generators and which has two rounds. Ishai et al. [29] give the first OT extension protocol only making black-box use of symmetric-key cryptography while increasing the rounds to three. Garg et al. [22]

show that three rounds are necessary in the OT hybrid model when only making black-box use of symmetric-key cryptography.

Rate-1 String OT. The notion of rate-1 OT has applications beyond the construction of PIR with polylog communication. In particular, a generalization of this notion, called trapdoor hash, has been used as a building block to build non-interactive zero knowledge for NP [9]. This has made the notion of rate-1 OT appealing from both a theoretical and practical points of view.

Unbalanced Private-Set Intersection (PSI). Private keyword search allows a receiver, with a single element x , to learn whether x is a member of a large set S held by a sender, or sometimes called PIR for keywords [14]. This is an instance of the so-called unbalanced PSI problem, defined earlier. A desirable feature of such unbalanced PSI protocols is sub-linear communication: the total amount of communication must be sub-linear the larger set size. We have protocols, from a wide variety of cryptographic assumptions, for unbalanced PSI whose communication complexity grows only polylogarithmically with the larger set size [12, 19, 21, 30].

Again, the Diffie-Hellman-based protocols come with a high sender computation cost: the number of group operations grows at least linearly in the bigger set size. As in PIR, one can prove that the strict running time of the sender in unbalanced PSI cannot be sub-linear in $|S_1|$, but that does not mean the number of public-key operations must also grow with n — especially, if the sender is allowed to make an arbitrarily-large number of symmetric-key operations. In fact, while the protocols in [12, 19, 21] induce little communication, the large number of public-key operations involved is a major bottleneck.

In the absence of the sub-linear communication requirement, one may use oblivious-transfer (OT) extension techniques [3, 29] to design unbalanced PSI protocols with a number of public-key operations independent of $|S_1|$. These protocols can be made concretely efficient as well (e.g., [13]). However, all these OT-extension-based protocols fail to achieve sub-linear communication.

4 Preliminaries

We denote the security parameter by λ . We say a function negl is negligible if for any polynomial poly we have $\text{negl}(\lambda) \in o(\frac{1}{\text{poly}(\lambda)})$. For two integers i and i' , we define $[i, i'] := \{i, i + 1, \dots, i'\}$. We let $[n] := \{1, \dots, n\}$.

For $i \in \{r, s\}$, denoting receiver (r) and sender (s), we let $\text{view}_i^\Pi(1^\lambda, x, y)$ denote the view of Party i in an honest execution of the protocol Π on 1^λ and on the parties' respective inputs, where the view contains the private input and the random coins of the respective party, the protocol's transcript, and the transcript of oracle queries and their responses. We may omit the security parameter 1^λ whenever it is clear from the context.

4.1 Oblivious Transfer

Definition 1 (Oblivious Transfer (OT)). An ℓ -batch k -bit string OT protocol OT is a protocol between two interactive PPT programs (OTR, OTS), where OTR and OTS denote, respectively, the receiver and the sender.

$\text{OTR}(1^\lambda, 1^\ell, 1^k, s)$: An interactive algorithm that takes in a security parameter 1^λ , batching parameter 1^ℓ , message parameter 1^k , and choice vector $s \in \{0, 1\}^\ell$, and outputs $m \in \{0, 1\}^\ell$.

$\text{OTS}(1^\lambda, 1^\ell, 1^k, m^{(0)}, m^{(1)})$: An interactive algorithm that takes in a security parameter 1^λ , batching parameter 1^ℓ , message parameter 1^k and two message vectors $m^{(0)}, m^{(1)} \in \mathbb{M}^\ell$, for $\mathbb{M} = \{0, 1\}^k$, and outputs \perp

We require the following.

Correctness. OT is $\alpha(\cdot)$ -correct if for any λ , $s \in \{0, 1\}^\ell$, $(m_i^{(0)}, m_i^{(1)})_{i \in [\ell]} \in (\mathbb{M} \times \mathbb{M})^\ell$, the probability over an honest interaction between $\text{OTR}(1^\lambda, 1^k, 1^\ell, s)$ and $\text{OTS}(1^\lambda, 1^k, 1^\ell, m^{(0)}, m^{(1)})$ that OTR outputs $(m_1^{(s_1)} \dots m_\ell^{(s_\ell)})$ is $\geq \alpha(\lambda)$. The protocol is perfectly correct if $\alpha = 1$. By default we require perfect correctness.

Semi-Honest Receiver Security. For any strings $s_0, s_1 \in \{0, 1\}^\ell$, $m^{(0)}, m^{(1)} \in \mathbb{M}^\ell$ we have that $\text{view}_s^{\text{OT}}(s_0, (m^{(0)}, m^{(1)}))$ and $\text{view}_s^{\text{OT}}(s_1, (m^{(0)}, m^{(1)}))$ are computationally indistinguishable.

Semi-Honest Sender Security. For any $s \in \{0, 1\}^\ell$ and $m^{(0)}, m^{(1)}, z^{(0)}, z^{(1)} \in \mathbb{M}^\ell$ such that $\{(m_i^{s_i})\} = \{(z_i^{s_i})\}$, we have that the two views $\text{view}_r^{\text{OT}}(s, (m^{(0)}, m^{(1)}))$ and $\text{view}_r^{\text{OT}}(s, (z^{(0)}, z^{(1)}))$ are computationally indistinguishable.

OT Terminologies. We may sometimes refer to an ℓ -batch single-bit OT as an ℓ -batch OT. Also, whenever we say a k -bit string OT we mean $\ell = 1$.

We define notions of *rate* as asymptotic ratios between the actual communication under a given protocol and the best achievable communication under a (possibly) insecure protocol; i.e., for ℓ -batch single-bit OT the sender must communicate at least ℓ bits to the receiver, if perfect correctness is required. Therefore, the optimal download communication is ℓ . Similarly, the optimal total communication is 2ℓ .

Expected Download Rate. An ℓ -batch single-bit OT protocol has expected download rate c if for all λ , s , $m^{(0)}$, $m^{(1)}$, and all but finitely many ℓ

$$\frac{\ell}{d(\lambda, \ell)} \geq c,$$

where $d(\lambda, \ell)$ is expected communication from $\text{OTS}(1^\lambda, 1^\ell, m^{(0)}, m^{(1)})$ to $\text{OTR}(1^\lambda, 1^\ell, s)$.

Expected (Overall) Rate. An ℓ -batch single-bit OT protocol has expected (overall) rate c if for all $\lambda, s, m^{(0)}, m^{(1)}$, and all but finitely many ℓ

$$\frac{2\ell}{t(\lambda, \ell)} \geq c,$$

where $t(\lambda, \ell)$ is the expected total communication.

We now define the notion of OT extension in the black-box OT model, which is stronger than the OT-hybrid model.

Definition 2 (Black-Box OT Extension). A black-box OT extension $\text{OTExt}^{\text{OT}} = (\text{OTRExt}^{\text{OT}}, \text{OTSExt}^{\text{OT}})$ is an ℓ -batch k -bit OT protocol that for a fixed polynomial poly , independent of ℓ , makes at most $\text{poly}(\lambda)$ calls to the base single-bit OT oracle $\text{OT} = (\text{OTR}, \text{OTS})$.

4.2 Private-Information Retrieval (PIR)

Definition 3 (Non-Trivial PIR). A non-trivial (single-server) private information retrieval ntPIR is an interactive protocol between two interactive PPT programs $(\text{PIRU}, \text{PIRS})$, where PIRU and PIRS denote, respectively, the client (user) and the server.

$\text{PIRU}(1^\lambda, 1^n, i)$: An interactive algorithm that takes in a security parameter 1^λ , the database size n , and a choice index $i \in [n]$, and at the end of the interaction outputs $y \in \{0, 1\}$.

$\text{PIRS}(1^\lambda, 1^n, \text{DB})$: An interactive algorithm that takes in a security parameter 1^λ , database size n and a database $\text{DB} \in \{0, 1\}^n$, and outputs \perp .

We require the following properties.

Correctness. The PIR protocol is $\alpha(\cdot)$ -correct if for any $\lambda, n, i \in [n]$ and $\text{DB} \in \{0, 1\}^n$, the probability over an honest interaction between $\text{PIRU}(1^\lambda, 1^n, i)$ and $\text{PIRS}(1^\lambda, 1^n, \text{DB})$ that PIRU outputs DB_i is $\geq \alpha(\lambda)$. The protocol is perfectly correct if $\alpha = 1$. By default we require perfect correctness.

Semi-Honest Client Security. For any $n, i, i' \in [n], \text{DB} \in \{0, 1\}^n$, $\text{view}_s^{\text{ntPIR}}(i, \text{DB})$ and $\text{view}_s^{\text{ntPIR}}(i', \text{DB})$ are computationally indistinguishable.

Non-Trivial Expected Download Communication. There exists a polynomial poly such that for all sufficiently large λ , for all $n \geq \text{poly}(\lambda)$, for all $i \in [n]$, and $\text{DB} \in \{0, 1\}^n$, the expected communication from $\text{PIRS}(1^\lambda, 1^n, i)$ to $\text{PIRU}(1^\lambda, 1^n, \text{DB})$ is $d(\lambda, n) < n$.

The following result shows that non-trivial PIR implies public-key cryptography in a black-box way. We use this theorem for our lower-bound results.

Theorem 2 ([18]). There exists a black-box construction of OT from a non-trivial PIR protocol.

5 Protocols that Imply Non-Trivial PIR

We substantiate the relevance of non-trivial PIR by showing that communication-efficient versions of some popular MPC protocols can be transformed into non-trivial PIR in a black-box manner. These transformations later let us transfer the lower-bounds regarding PIR to these protocols.

In the following, we focus on different variants of oblivious transfer and unbalanced private set intersection to demonstrate the concept. The same ideas apply to many other protocols such as vector oblivious linear evaluation and oblivious polynomial evaluation.

5.1 Oblivious Transfer

We show how to transform a protocol for k -bit string oblivious transfer $\text{OT} = (\text{OTR}, \text{OTS})$ that makes calls to an oracle \mathcal{O} into a PIR protocol $(\text{PIRU}, \text{PIRS})$ with database size $n = 2k$ that makes calls to the same oracle \mathcal{O} in a black-box manner. The construction is folklore and works by splitting the database in half, using each half as one of the two strings, and choosing the OT choice bit based on the PIR client's index accordingly.

- $\text{PIRU}^{\mathcal{O}}(1^\lambda, 1^n, i)$: For $n = 2k$, set the choice bit $b \leftarrow \lfloor (i-1)/k \rfloor$. Run the OT receiver $m_b \leftarrow \text{OTR}^{\mathcal{O}}(1^\lambda, b)$ to get the chosen string m_b . Return $m_b[i - kb]$
- $\text{PIRS}^{\mathcal{O}}(1^\lambda, 1^n, \text{DB})$: Let strings $m_0 \leftarrow \text{DB}[1, \dots, k]$ and $m_1 \leftarrow \text{DB}[k+1, \dots, 2k]$. Run the OT sender $\text{OTS}^{\mathcal{O}}(1^\lambda, m_0, m_1)$.

Lemma 1 (Folklore). *The PIR protocol $(\text{PIRU}, \text{PIRS})$ has the same correctness error, the same sender/receiver query complexity, and the same sender/receiver communication as those of the OT protocol OT .*

That means if the expected sender communication in OT is less than $n = 2k$, then $(\text{PIRU}, \text{PIRS})$ is a non-trivial PIR protocol.

Proof of Correctness. Correctness follows from the correctness of the OT protocol and for $b = \lfloor (i-1)/k \rfloor$ we have $m_b[i - kb] = \text{DB}[kb + 1, \dots, kb + k][i - kb] = \text{DB}[i]$. \square

Proof of Client Security. Suppose there exists $i, i' \in [N]$, $\text{DB} \in \{0, 1\}^N$ such that an adversary \mathcal{A} can distinguish $\text{view}_s^{\text{PIR}}(i, \text{DB})$ from $\text{view}_s^{\text{PIR}}(i', \text{DB})$ with non-negligible probability. Then $\lfloor (i-1)/k \rfloor \neq \lfloor (i'-1)/k \rfloor$ else $\text{view}_s^{\text{PIR}}(i, \text{DB})$ and $\text{view}_s^{\text{PIR}}(i', \text{DB})$ follow the exact same distribution. The same adversary \mathcal{A} distinguishes between $\text{view}_s^{\text{OT}}(\lfloor (i-1)/k \rfloor, (\text{DB}[1, \dots, k], \text{DB}[k+1, \dots, 2k]))$ and $\text{view}_s^{\text{OT}}(\lfloor (i'-1)/k \rfloor, (\text{DB}[1, \dots, k], \text{DB}[k+1, \dots, 2k]))$ with the same non-negligible probability since the views are exactly the same as $\text{view}_s^{\text{PIR}}(i, \text{DB})$ and $\text{view}_s^{\text{PIR}}(i', \text{DB})$ respectively. \square

Remark 1. One can transform any ℓ -batch k -bit OT protocol into a $k\ell$ -bit string OT protocol by reusing the same choice bit across all the ℓ batches. This works without any issues because we only talk about semi-honest security.

OT With Low Total Communications. Using the symmetric nature of OT we transform an OT protocol with low communication (not just low sender communication) into a low sender communication OT protocol in a black-box manner. This allows us to apply our PIR lower-bounds to OT with low expected communication. Our transformation works by noting that every communication efficient OT protocol has either low sender or low receiver communication; if the receiver communication is low, our transformation will swap the roles of the sender and receiver, to obtain an OT protocol with low sender communication, as desired.

The following transformation was implicitly used in [29] and explicitly in [39]. The transformation works as follows: Let $\text{OT} = (\text{OTR}, \text{OTS})$ be an ℓ -batch single-bit OT with expected total communication $t(\lambda, \ell)$, expected download communication $d(\lambda, \ell)$, expected upload communication $u(\lambda, \ell)$, and oracle accesses to \mathcal{O} . We define a $\text{OT}' = (\text{OTR}', \text{OTS}')$ as follows

$\text{OTR}'^{\mathcal{O}}(1^\lambda, 1^\ell, s)$:

1. If the expected download communication of OT is $d(\lambda, \ell) < u(\lambda, \ell) + \ell$:
 - (a) Run $(m'_1, \dots, m'_\ell) \leftarrow \text{OTR}'^{\mathcal{O}}(1^\lambda, 1^\ell, s)$, the ℓ -batch OT receiver on the choice string s .
 - (b) Return (m'_1, \dots, m'_ℓ)
2. Else:
 - (a) Sample $r \xleftarrow{\$} \{0, 1\}^\ell$ uniformly at random.
 - (b) Run $\text{OTS}'^{\mathcal{O}}(1^\lambda, 1^\ell, r, s \oplus r)$, the ℓ -batch OT sender on messages $m_0 = r$ and $m_1 = r \oplus s$.
 - (c) Receive v in the round after OTS is done.
 - (d) Return $v \oplus r$

$\text{OTS}'^{\mathcal{O}}(1^\lambda, 1^\ell, m^{(0)}, m^{(1)})$:

1. If the expected download communication of OT is $d(\lambda, \ell) < u(\lambda, \ell) + \ell$:
 - (a) Run $\text{OTS}'^{\mathcal{O}}(1^\lambda, 1^\ell, m^{(0)}, m^{(1)})$, the ℓ -batch OT sender on messages $m_0 = m^{(0)}$ and $m_1 = m^{(1)}$.
 - (b) Return
2. Else:
 - (a) Run $z \leftarrow \text{OTR}'^{\mathcal{O}}(1^\lambda, 1^\ell, m^{(1)} \oplus m^{(0)})$, the ℓ -batch OT receiver on the choice string $m^{(1)} \oplus m^{(0)}$ to receive the string z .
 - (b) Send $z \oplus m^{(0)}$ in the round after OTR is done
 - (c) Return

Lemma 2. *The constructed OT protocol $\text{OT}' = (\text{OTR}', \text{OTS}')$ has the same correctness as the base OT $\text{OT} = (\text{OTR}, \text{OTS})$. Moreover, $\text{OT}' = (\text{OTR}', \text{OTS}')$ is secure if $\text{OT} = (\text{OTR}, \text{OTS})$ is secure.*

Assuming $\text{OT} = (\text{OTR}, \text{OTS})$ has expected overall rate $r > 2/3$, the constructed OT has expected download rate $w > 1/2$.

Correctness. If $d(\lambda, \ell) < u(\lambda, \ell) + \ell$ then both parties behave exactly like OT and therefore correctness is inherited.

If $d(\lambda, \ell) \geq u(\lambda, \ell) + \ell$ the sender OTS' learns

$$\left(r_1 \oplus s_1 \cdot (m_1^{(1)} \oplus m_1^{(0)}), \dots, r_\ell \oplus s_\ell \cdot (m_\ell^{(1)} \oplus m_\ell^{(0)}) \right)$$

it then sends back

$$\left(r_1 \oplus s_1 \cdot (m_1^{(1)} \oplus m_1^{(0)}) \oplus m_1^{(0)}, \dots, r_\ell \oplus s_\ell \cdot (m_\ell^{(1)} \oplus m_\ell^{(0)}) \oplus m_\ell^{(0)} \right)$$

then the receiver OTR' computes

$$\begin{aligned} & \left(s_1 \cdot (m_1^{(1)} \oplus m_1^{(0)}) \oplus m_1^{(0)}, \dots, s_\ell \cdot (m_\ell^{(1)} \oplus m_\ell^{(0)}) \oplus m_\ell^{(0)} \right) \\ &= \left(m_1^{(s_1)}, \dots, m_\ell^{(s_\ell)} \right) \end{aligned}$$

□

Security. The security in the case that $d(\lambda, \ell) < u(\lambda, \ell) + \ell$ directly follows from the security of OT .

In the other case it follows from the security of OT and the work of [39] which proves that this exact construction is secure. For receiver security we have that the sender (according to sender security of OT) only learns z . Each bit z_i is either r_i or $s_i \oplus r_i$, in both cases it is uniformly random because r is uniformly random. Sender security of OT' follows because the execution of OT leaks nothing to the receiver (according to the receiver security of OT). That means all the receiver learns is $z \oplus m^{(0)}$. By correctness of OT' this is exactly $(m_1^{(s_1)} \oplus r_1, \dots, m_\ell^{(s_\ell)} \oplus r_\ell)$ and therefore contains no information about $(m_1^{(1-s_1)}, \dots, m_\ell^{(1-s_\ell)})$. □

Expected Download Communication. Let r be the rate, $t(\lambda, n)$ be the expected total communication which is the sum of the expected receiver-to-sender communication $u(\lambda, \ell)$ and the expected sender-to-receiver communication $d(\lambda, \ell)$. Then for all but finitely many ℓ we have $t(\lambda, \ell) < \frac{2\ell}{r}$. In the following, the expected sender-to-receiver communication of OT' will be called $d'(\lambda, \ell)$.

If $d(\lambda, \ell) < u(\lambda, \ell) + \ell$ then

$$\begin{aligned} d'(\lambda, \ell) &= d(\lambda, \ell) = t(\lambda, \ell) - u(\lambda, \ell) \leq t(\lambda, \ell) - d(\lambda, \ell) + \ell \quad \Leftrightarrow \\ & d'(\lambda, \ell) \leq \frac{t(\lambda, \ell) + \ell}{2} \end{aligned}$$

Else the new expected sender-to-receiver communication is

$$\begin{aligned} d'(\lambda, \ell) &= u(\lambda, \ell) + \ell = t(\lambda, \ell) - d(\lambda, \ell) + \ell \leq t(\lambda, \ell) - u(\lambda, \ell) - \ell + \ell \quad \Leftrightarrow \\ & d'(\lambda, \ell) \leq \frac{t(\lambda, \ell) + \ell}{2} \end{aligned}$$

Either way, $d'(\lambda, \ell) \leq \frac{t(\lambda, \ell) + \ell}{2}$ which means that for all but finitely many ℓ we have $d'(\lambda, \ell) < (\frac{1}{r} + \frac{1}{2})\ell$. Therefore, the expected download rate is $\frac{2}{(\frac{1}{r} + \frac{1}{2})}$ which is $> 1/2$ for $r > 2/3$. □

5.2 Unbalanced Private-Set Intersection

In unbalanced private set intersection we have a set A of n λ -bit messages, held by a sender $\text{PSIS}(1^\lambda, 1^n, A)$, and a singleton set B , held by a receiver $\text{PSIR}(1^\lambda, 1^n, B)$. The goal is for the receiver to learn $A \cap B$ while the sender should learn nothing. Semi-honest receiver security can be defined along the lines of receiver (client) security of PIR (Definition 3).

We show how to transform a protocol for unbalanced private-set intersection $\text{PSI} = (\text{PSIS}, \text{PSIR})$ that makes calls to oracle \mathcal{O} into a PIR protocol $(\text{PIRU}, \text{PIRS})$ that makes calls to the same oracle \mathcal{O} in a black-box manner. We do this by simply encoding the PIR-database and the PIR-query as sets.

- $\text{PIRU}^{\mathcal{O}}(1^\lambda, 1^n, i)$: Set $A := \{i\}$. Run the PSI receiver $I \leftarrow \text{PSIR}^{\mathcal{O}}(1^\lambda, A)$ to get the intersection I . Return 1 if $\{i\} = I$ and 0 otherwise.
- $\text{PIRS}^{\mathcal{O}}(1^\lambda, 1^n, \text{DB})$: Form the set $B := \{x \mid \text{DB}[x] = 1\}$. Run the PSI sender $\text{PSIS}^{\mathcal{O}}(1^\lambda, B)$.

Lemma 3 (Folklore). *The PIR protocol $(\text{PIRU}, \text{PIRS})$ has the same correctness error, sender and receiver communication and sender and receiver query complexity as those of PSI. Moreover, the resulting PIR protocol has client security if PSI provides receiver security.*

That means if the sender communication in PSI is less than n , then $(\text{PIRU}, \text{PIRS})$ is a non-trivial PIR protocol.

Proof of Correctness. Correctness follows from the correctness of the PSI protocol and the intersection of $\{i\}$ and B being $\{i\}$ if $i \in B \Leftrightarrow \text{DB}[i] = 1$ and \emptyset otherwise.

Proof of Client Security. Suppose there exists $i, i' \in [n]$, $\text{DB} \in \{0, 1\}^n$ such that an adversary \mathcal{A} can distinguish $\text{view}_s^{\text{PIR}}(i, \text{DB})$ from $\text{view}_s^{\text{PIR}}(i', \text{DB})$ with non-negligible probability. The same adversary \mathcal{A} distinguishes between $\text{view}_s^{\text{PSI}}(\{i\}, B)$ and $\text{view}_s^{\text{PSI}}(\{i'\}, B)$ with the same non-negligible probability since the views are exactly the same. \square

6 Lower-Bounds on the Number Oracle Queries in PIR

In this section, we show how to transform a private information retrieval (PIR) protocol with access to some *simulatable* oracle SO into one that does not query that oracle. To have something concrete in mind one may imagine SO being the generic group model, though the technique is much more general. We will later go into common instantiations of the oracle. This transformation allows us to transfer lower-bounds from PIR without oracle access to PIR with oracle access.

Simulatable Oracles. A simulatable oracle is an oracle SO which can efficiently be simulated by a stateful simulator Sim . More formally, a computationally unbounded adversary \mathcal{A} cannot win the following game with a non-negligible advantage in polynomially many rounds r , where Sim is a PPT algorithm:

1. Sample random bit $b \xleftarrow{\$} \{0, 1\}$.
2. Initialize the state of the oracle as $\text{st} \leftarrow \perp$.
3. Initialize the state of the adversary $\text{ast} \leftarrow \perp$.
4. The adversary produces a first query qu .
5. For $i \in [r]$:
 - (a) If $b = 0$:
 - Let the response be $\text{resp} \leftarrow \text{SO}(\text{qu})$
 - (b) Else:
 - Let response and new oracle state be $(\text{resp}, \text{st}) \leftarrow \text{Sim}(\text{qu}, \text{st})$
 - (c) Let new query and adversary state be $(\text{qu}, \text{ast}) \leftarrow \mathcal{A}(\text{ast}, \text{resp})$
6. Let the adversary output its guess $b' \leftarrow \mathcal{A}(\text{ast})$.
7. The adversary wins if $b = b'$.

Typical examples of simulatable oracles include the random oracle and the generic group oracle.

Construction 1. Let $\text{PIR} := (\text{PIRU}^{\text{SO}, \text{O}}, \text{PIRS}^{\text{SO}, \text{O}})$ be a bit PIR protocol that uses a simulatable oracle SO and another oracle O . We show how to compile out the SO -calls of $(\text{PIRU}^{\text{SO}, \text{O}}, \text{PIRS}^{\text{SO}, \text{O}})$, obtaining an SO -free PIR protocol $\overline{\text{PIR}} := (\overline{\text{PIRU}}^{\text{O}}, \overline{\text{PIRS}}^{\text{O}})$.

For notational convenience, in the following whenever calling PIRU or PIRS , we omit the private-state part of the input.

The protocol messages sent from $\overline{\text{PIRU}}$ to $\overline{\text{PIRS}}$ are tagged with either ‘protocol’ (or bit zero) signifying a normal protocol message, or with ‘query’ (or bit 1) signifying a query message.

$\overline{\text{PIRU}}^{\text{O}}(1^\lambda, 1^n, i)$:

- Initialize the state of the simulatable oracle $\text{st} \leftarrow \perp$.
- Run the interactive PPT $\text{PIRU}^{\text{SO}, \text{O}}$ with the following interactions:
 1. When $\text{PIRU}^{\text{SO}, \text{O}}$ calls O on a query qu forward the query to O and respond with the received response.
 2. When $\text{PIRU}^{\text{SO}, \text{O}}$ calls SO on a query qu , simulate the response and update the oracle simulators state $(\text{resp}, \text{st}) \leftarrow \text{Sim}(\text{qu}, \text{st})$.
 3. Upon $\overline{\text{PIRU}}$ receiving a message of the form $(\text{‘query’}, \text{msgs})$, interpret msgs as a query qu , simulate the oracle response and update the oracle simulators state as $(\text{resp}, \text{st}) \leftarrow \text{Sim}(\text{qu}, \text{st})$ and return resp to the sender $\overline{\text{PIRS}}$. If the message has the form $(\text{‘protocol’}, \text{msgs})$, run $\text{PIRU}^{\text{SO}, \text{O}}$ on the protocol message msgs until it produces the next message msg_r and send that to $\overline{\text{PIRS}}$. The oracle queries are handled as described above.

$\overline{\text{PIRS}}^{\text{O}}(1^\lambda, 1^n, \text{DB})$:

- Run the interactive PPT $\text{PIRS}^{\text{SO}, \text{O}}$ with the following interactions:

1. When $\text{PIRS}^{\text{SO},\text{O}}$ calls O on a query qu , forward the query to O and respond with the received response.
2. When $\text{PIRS}^{\text{SO},\text{O}}$ calls SO on a query qu , send a tagged query pair ('query', qu) to $\overline{\text{PIR}}_{\text{U}}$ and use the response msg_r as a query response for qu to PIRS .
3. Else, run $\text{PIRS}^{\text{SO},\text{O}}$ until it produces a message msgs and send the tagged message ('protocol', msgs) to $\overline{\text{PIR}}_{\text{U}}$, then wait for the response msg_r and continue.

Theorem 3. *If PIR is a non-trivial private information retrieval with server communication of $\eta < cn$ for $c < 1$, $r \in o(n)$ rounds of interaction with the user, and $q \in o(n)$ bits of communication with the SO oracle then $\overline{\text{PIR}}$ is a non-trivial private information retrieval with server communication $\bar{\eta} \leq \bar{c}n$ for $\bar{c} < 1$ and no calls to SO .*

Server Communication. The server's additional communication overhead includes 1 bit per round as well as a total of $O(q)$ bits. Since the number rounds is $o(n)$, the total server communication complexity becomes $cn + o(n)$, which is less than $\bar{c}n$ for some $\bar{c} < 1$. \square

Correctness. Notice that the above protocol will have different output from an execution of PIR either

1. if a with Sim simulated oracle behave differently from the real oracle behaviour or
2. if a message in the execution of PIR happens to start with t .

Both of these events happen with negligible probability. Therefore, if PIR has statistical correctness then so does $\overline{\text{PIR}}$. \square

Client Security. Suppose there exists $i, i' \in [n]$, $\text{DB} \in \{0,1\}^n$ such that an adversary $\overline{\mathcal{A}}$ can distinguish $\text{view}_s^{\overline{\text{PIR}}}(i, \text{DB})$ from $\text{view}_s^{\overline{\text{PIR}}}(i', \text{DB})$ with non-negligible probability. We construct a new adversary \mathcal{A} to distinguish between $\text{view}_s^{\text{PIR}}(i, \text{DB})$ and $\text{view}_s^{\text{PIR}}(i', \text{DB})$. The new adversary \mathcal{A} gets as input a view v either from $\text{view}_s^{\text{PIR}}(i, \text{DB})$ or $\text{view}_s^{\text{PIR}}(i', \text{DB})$ and does the following:

1. Generate an empty view \bar{v} .
2. Copy all O -oracle calls from v to \bar{v} .
3. Run PIRS on the randomness and DB as defined in the view v and simulate its interaction as follows:
 - (a) For PIRS's calls to the SO oracle with query qu and gets response resp enter ('query', qu) as a server message into \bar{v} and resp as a user message.
 - (b) For PIRS's messages msgs enter ('protocol', msgs) in the transcript \bar{v} as a server message and enter the users response msg_r as a users message.
4. Run $b \leftarrow \overline{\mathcal{A}}(\bar{v})$, on the view \bar{v} produced by PIRS
5. Return b

\mathcal{A} will distinguish $\text{view}_s^{\text{PIR}}(i, \text{DB})$ and $\text{view}_s^{\text{PIR}}(i', \text{DB})$ with negligibly close to the probability as $\bar{\mathcal{A}}$ can distinguish $\text{view}_s^{\text{PIR}}(i, \text{DB})$ from $\text{view}_s^{\text{PIR}}(i', \text{DB})$. This is because \bar{v} follows the same distribution as $\text{view}_s^{\text{PIR}}(i, \text{DB})$ (except that the SO queries are produced by the real oracle, not the simulator) if v was from $\text{view}_s^{\text{PIR}}(i, \text{DB})$ and \bar{v} follows the same distribution as $\text{view}_s^{\text{PIR}}(i', \text{DB})$ (same caveat here) if v was from $\text{view}_s^{\text{PIR}}(i', \text{DB})$. If the adversary could notice the simulation of SO then it would break its simulatability. \square

Remark 2. Theorem 3 is applicable to any two-PC protocol with one-sided receiver security. Of course, in the absence of further restrictions, such protocols are trivial to realize (e.g., by the sender sending its input in the clear to the receiver). One restriction that makes the problem non-trivial is to require the sender-to-receiver communication to be sub-linear in the sender's input size, as in PIR.

The utility of Theorem 3, beyond PIR itself, becomes apparent when one considers other protocols that imply non-trivial PIR while instantiating their underlying oracles via ideal forms of powerful primitives. We first discuss the implications of the theorem in terms of particular instantiations of the oracle, and in the next section we consider protocols that imply PIR.

Theorem 3 allows us to also rule out powerful non-black-box techniques for building PIR. We demonstrate this by letting SO include an OT oracle and an ideal obfuscation oracle that can obfuscate circuits with generic OT gates and random oracle gates. (See [2, 20] for capturing similar non-black-box techniques via oracle-aided circuits.)

Corollary 4. *For any constants $c < 1$, there exists no n -bit PIR protocol with server communication $\eta \leq cn$, round complexity $r \in o(n)$, and with oracle access to a PSPACE-complete oracle, a random oracle, a generic OT oracle, and an obfuscation oracle for circuits with OT and random oracle gates, and where the server only communicates $q \in o(n)$ bits to the ideal obfuscation and OT oracles.*

Proof. In Lemma 1 we show an OT protocol with the above mentioned characteristics implies a non-trivial PIR. Let SO consist of an OT oracle [26] and an ideal obfuscation oracle [2, 31] for obfuscating circuits with OT/RO gates. (Such an SO oracle is simulatable.) By invoking Theorem 3 one gets a non-trivial PIR with oracle access to the random oracle and an PSPACE-complete oracle. This in turn can be transformed into an OT protocol (while retaining the O oracles) via [18]. The existence of such an object however was ruled out by [26]. \square

Back to the black-box setting, other illustrative examples include the use of GGMs for building non-trivial PIR.

Corollary 5. *For any constants $c < 1$, a non-trivial n -bit PIR protocol with server communication of cn , round complexity $r \in o(n)$ and where the server makes sublinear in n many generic group queries requires MPC-hard assumptions, beyond the generic group.*

Proof. In Theorem 3, if one instantiates SO by a generic group [38] and let the \mathcal{O} oracle be empty, then one gets a non-trivial PIR without any oracle calls. This in turn can be transformed into an OT protocol without any oracles via [18]. OT is an MPC-complete protocol. \square

The above corollary is almost tight as there exists GGM-based PIR protocols with a linear number of GGM queries.

Lemma 4 ([19]). *Based on the DDH assumption, there exists a non-trivial n -bit PIR protocol with server communication of $O(\lambda)$ and with the server making $O(n)$ group operations.*

Finally, we may derive a statement for FHE oracles.

Corollary 6. *For any constants $c < 1$, a non-trivial n -bit PIR protocol with server communication of cn , round complexity $r \in o(n)$ where the server makes $q \in o(n)$ black-box use of fully homomorphic encryption³ requires MPC-hard assumptions, even beyond the fully homomorphic encryption.*

Proof. Let SO be an FHE oracle [23] (defined similarly to a generic PKE oracle of [26]). Let the \mathcal{O} oracle be empty. Invoking Theorem 3 one gets a non-trivial PIR without any oracle calls. This in turn can be transformed into an OT protocol without oracles via [18]. OT is an MPC-complete protocol. \square

7 Communication Lower-Bounds for OT Extension

Theorem 3 provides lower-bounds on the computational complexity of PIR protocols. In this section, we show that these computational lower-bounds give rise to communication lower-bounds for OT extension (i.e., the number of bits that an extended OT sender needs to communicate). The result of this section implies that the communication complexity of the sender in the IKNP OT extension protocol [29] is close to optimal.

Corollary 7 (OT Extension: Sender Communication Lower-Bound). *For any constants $c < 1$, there exist no ℓ -batch k -bit OT extension protocol with sender communication $\eta < c2k\ell$, round complexity $r \in o(k\ell)$, and with the sender making $q \in o(k\ell)$ OT calls.*

Proof. OT extension is just an OT protocol that makes use of only a black-box OT and a random oracle. An ℓ -batch k -bit OT naturally gives rise to a $k\ell$ -bit string OT. In Lemma 1 we show that such an OT protocol implies a non-trivial PIR for databases of $2k\ell$ bits. Under the resulting PIR protocol, the server communication is $\eta < c2k\ell$, round complexity $r \in o(k\ell)$, and the server communicates a total of $o(k\ell)$ bits with the OT oracle. Invoking Theorem 3, by instantiating SO with a generic OT oracle [26] and \mathcal{O} with the random oracle [6] and by also including a PSPACE-complete oracle, we get a non-trivial PIR with

³ This means the server communicates at most q bits to the FHE oracle.

oracle access to the random oracle and a PSPACE-complete oracle. This in turn can be transformed into an OT protocol (while retaining the \mathcal{O} oracles) via [18]. The existence of such an object however was ruled out by [26]. \square

Corollary 8 (OT Extension: Total Communication Lower-Bound). *For any constants $c < 1$, there exist no ℓ -batch k -bit OT extension protocol with total communication $\eta < \frac{3}{2}c(\ell + k\ell)$, round complexity $r \in o(k\ell)$, and with the sender making $q \in o(k\ell)$ OT calls.*

Proof. Follows from Lemma 2 and Corollary 7. \square

Acknowledgements. We thank the anonymous reviewers for their valuable comments and suggestions.

J.D.: Funded by the European Union (ERC, LACONIC, 101041207). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

M. Hajiabadi supported in part by an NSERC Discovery Grant 03270, and a Meta Research Award.

References

1. Aggarwal, D., Dottling, N., Dujmovic, J., Hajiabadi, M., Malavolta, G., Obremski, M.: Algebraic restriction codes and their applications. In: Braverman, M. (ed.), 13th Innovations in Theoretical Computer Science Conference, ITCS 2022, January 31–3 February 2022, Berkeley, CA, USA, volume 215 of LIPIcs, pp. 2:1–2:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2022)
2. Asharov, G., Segev, G.: Limits on the power of indistinguishability obfuscation and functional encryption. In: Guruswami, V. (ed.) 56th Annual Symposium on Foundations of Computer Science, pp. 191–209. IEEE Computer Society Press, October 2015
3. Beaver, D.: Correlated pseudorandomness and the complexity of private computations. In: 28th Annual ACM Symposium on Theory of Computing, pp. 479–488. ACM Press, May 1996
4. Beimel, A., Ishai, Y., Malkin, T.: Reducing the servers computation in private information retrieval: PIR with preprocessing. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, pp. 55–73. Springer, Heidelberg (2000). https://doi.org/10.1007/3-540-44598-6_4
5. Beimel, A., Ishai, Y., Malkin, T.: Reducing the servers’ computation in private information retrieval: PIR with preprocessing. *J. Cryptol.* **17**(2), 125–151 (2004)
6. Bellare, M., Rogaway, P.: Random oracles are practical: a paradigm for designing efficient protocols. In: Denning, D.E., Pyle, R., Ganesan, R., Sandhu, R.S., Ashby, V. (eds.), ACM CCS 93: 1st Conference on Computer and Communications Security, pp. 62–73. ACM Press, November 1993
7. Boyle, E., Ishai, Y., Pass, R., Wootters, M.: Can we access a database both locally and privately? In: Kalai, Y., Reyzin, L. (eds.) TCC 2017. LNCS, vol. 10678, pp. 662–693. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70503-3_22

8. Brakerski, Z., Branco, P., Dottling, N., Pu, S.: Batch-OT with optimal rate. In: Dunkelman, O., Dziembowski, S. (eds.) *Advances in Cryptology – EUROCRYPT 2022*. EUROCRYPT 2022. LNCS, vol. 13276, pp. 157–186. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-07085-3_6
9. Brakerski, Z., Koppula, V., Mour, T.: NIZK from LPN and trapdoor hash via correlation intractability for approximable relations. In: Micciancio, D., Ristenpart, T. (eds.) *CRYPTO 2020*. LNCS, vol. 12172, pp. 738–767. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-56877-1_26
10. Cachin, C., Micali, S., Stadler, M.: Computationally private information retrieval with polylogarithmic communication. In: Stern, J. (ed.) *EUROCRYPT 1999*. LNCS, vol. 1592, pp. 402–414. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-48910-X_28
11. Canetti, R., Holmgren, J., Richelson, S.: Towards doubly efficient private information retrieval. In: Kalai, Y., Reyzin, L. (eds.) *TCC 2017*. LNCS, vol. 10678, pp. 694–726. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70503-3_23
12. Chase, M., Garg, S., Hajiabadi, M., Li, J., Miao, P.: Amortizing rate-1 OT and applications to PIR and PSI. In: Nissim, K., Waters, B. (eds.) *TCC 2021*. LNCS, vol. 13044, pp. 126–156. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-90456-2_5
13. Chase, M., Miao, P.: Private set intersection in the internet setting from lightweight oblivious PRF. In: Micciancio, D., Ristenpart, T. (eds.) *CRYPTO 2020*. LNCS, vol. 12172, pp. 34–63. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-56877-1_2
14. Chor, B., Gilboa, N., Naor, M.: Private information retrieval by keywords. *Cryptology ePrint Archive*, Report 1998/003 (1998). <https://eprint.iacr.org/1998/003>
15. Chor, B., Kushilevitz, E., Goldreich, O., Sudan, M.: Private information retrieval. In: *36th Annual Symposium on Foundations of Computer Science*, pp. 41–50. IEEE Computer Society Press, October 1995
16. Corrigan-Gibbs, H., Henzinger, A., Kogan, D.: Single-server private information retrieval with sublinear amortized time. In: Dunkelman, O., Dziembowski, S. (eds.) *Advances in Cryptology – EUROCRYPT 2022*. EUROCRYPT 2022. LNCS, vol. 13276, pp. 3–33. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-07085-3_1
17. Corrigan-Gibbs, H., Kogan, D.: Private information retrieval with sublinear online time. In: Canteaut, A., Ishai, Y. (eds.) *EUROCRYPT 2020*. LNCS, vol. 12105, pp. 44–75. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45721-1_3
18. Di Crescenzo, G., Malkin, T., Ostrovsky, R.: Single database private information retrieval implies oblivious transfer. In: Preneel, B. (ed.) *EUROCRYPT 2000*. LNCS, vol. 1807, pp. 122–138. Springer, Heidelberg (2000). https://doi.org/10.1007/3-540-45539-6_10
19. Döttling, N., Garg, S., Ishai, Y., Malavolta, G., Mour, T., Ostrovsky, R.: Trapdoor hash functions and their applications. In: Boldyreva, A., Micciancio, D. (eds.) *CRYPTO 2019*. LNCS, vol. 11694, pp. 3–32. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-26954-8_1
20. Garg, S., Hajiabadi, M., Mahmoody, M., Mohammed, A.: Limits on the power of garbling techniques for public-key encryption. In: Shacham, H., Boldyreva, A. (eds.) *CRYPTO 2018*. LNCS, vol. 10993, pp. 335–364. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96878-0_12

21. Garg, S., Hajiabadi, M., Ostrovsky, R.: Efficient range-trapdoor functions and applications: rate-1 OT and more. In: Pass, R., Pietrzak, K. (eds.) TCC 2020. LNCS, vol. 12550, pp. 88–116. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-64375-1_4
22. Garg, S., Mahmoody, M., Masny, D., Meckler, I.: On the round complexity of OT extension. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018. LNCS, vol. 10993, pp. 545–574. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96878-0_19
23. Garg, S., Mahmoody, M., Mohammed, A.: Lower bounds on obfuscation from all-or-nothing encryption primitives. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017. LNCS, vol. 10401, pp. 661–695. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63688-7_22
24. Gennaro, R., Gertner, Y., Katz, J.: Lower bounds on the efficiency of encryption and digital signature schemes. In: 35th Annual ACM Symposium on Theory of Computing, pp. 417–425. ACM Press, June 2003
25. Gennaro, R., Trevisan, L.: Lower bounds on the efficiency of generic cryptographic constructions. In: 41st Annual Symposium on Foundations of Computer Science, pp. 305–313. IEEE Computer Society Press, November 2000
26. Gertner, Y., Kannan, S., Malkin, T., Reingold, O., Viswanathan, M.: The relationship between public key encryption and oblivious transfer. In: 41st Annual Symposium on Foundations of Computer Science, pp. 325–335. IEEE Computer Society Press, November 2000
27. Henzinger, A., Hong, M.M., Corrigan-Gibbs, H., Meiklejohn, S., Vaikuntanathan, V.: One server for the price of two: simple and fast single-server private information retrieval. Cryptology ePrint Archive, Report 2022/949 (2022). <https://eprint.iacr.org/2022/949>
28. Henzinger, A., Hong, M.M., Corrigan-Gibbs, H., Meiklejohn, S., Vaikuntanathan, V.: One server for the price of two: simple and fast single-server private information retrieval. In: Calandrino, J.A., Troncoso, C. (eds.), 32nd USENIX Security Symposium, USENIX Security 2023, Anaheim, CA, USA, 9–11 August 2023. USENIX Association (2023)
29. 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). https://doi.org/10.1007/978-3-540-45146-4_9
30. Ishai, Y., Paskin, A.: Evaluating branching programs on encrypted data. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 575–594. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-70936-7_31
31. Jain, A., Lin, H., Luo, J., Wichs, D.: The pseudorandom oracle model and ideal obfuscation. In: Handschuh, H., Lysyanskaya, A. (eds.) Advances in Cryptology – CRYPTO 2023. CRYPTO 2023. LNCS, vol. 14084, pp. 233–262. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-38551-3_8
32. Kogan, D., Corrigan-Gibbs, H.: Private blacklist lookups with checklist. In: Bailey, M., Greenstadt, R. (eds.) USENIX Security 2021: 30th USENIX Security Symposium, pp. 875–892. USENIX Association, August 2021
33. Kushilevitz, E., Ostrovsky, R.: One-way trapdoor permutations are sufficient for non-trivial single-server private information retrieval. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 104–121. Springer, Heidelberg (2000). https://doi.org/10.1007/3-540-45539-6_9

34. Lin, W.-K., Mook, E., Wichs, D.: Doubly efficient private information retrieval and fully homomorphic RAM computation from ring LWE. In: Saha, B., Servedio, R.A. (eds.) Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, 20–23 June 2023, pp. 595–608. ACM (2023)
35. Ostrovsky, R., Skeith, W.E.: A survey of single-database private information retrieval: techniques and applications. In: Okamoto, T., Wang, X. (eds.) PKC 2007. LNCS, vol. 4450, pp. 393–411. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-71677-8_26
36. Persiano, G., Yeo, K.: Limits of preprocessing for single-server PIR. In: Naor, J.S., Buchbinder, N. (eds.) Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference/Alexandria, VA, USA, 9–12 January 2022, pp. 2522–2548. SIAM (2022)
37. Shi, E., Aqeel, W., Chandrasekaran, B., Maggs, B.: Puncturable pseudorandom sets and private information retrieval with near-optimal online bandwidth and time. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021. LNCS, vol. 12828, pp. 641–669. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-84259-8_22
38. Shoup, V.: Lower bounds for discrete logarithms and related problems. In: Fumy, W. (ed.) EUROCRYPT 1997. LNCS, vol. 1233, pp. 256–266. Springer, Heidelberg (1997). https://doi.org/10.1007/3-540-69053-0_18
39. Wolf, S., Wullschlegel, J.: Oblivious transfer is symmetric. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 222–232. Springer, Heidelberg (2006). https://doi.org/10.1007/11761679_14
40. Yeo, K.: Lower bounds for (Batch) PIR with private preprocessing. In: Hazay, C., Stam, M. (eds.) Advances in Cryptology – EUROCRYPT 2023. EUROCRYPT 2023. LNCS, vol. 14004, pp. 518–550. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-30545-0_18
41. Zhandry, M.: To label, or not to label (in generic groups). In: Dodis, Y., Shrimpton, T. (eds.) Advances in Cryptology – CRYPTO 2022. CRYPTO 2022. LNCS, vol. 13509, pp. 66–96. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-15982-4_3
42. Zhou, M., Lin, W.K., Tselekounis, Y., Shi, E.: Optimal single-server private information retrieval. In: Hazay, C., Stam, M. (eds.) Advances in Cryptology – EUROCRYPT 2023. EUROCRYPT 2023. LNCS, vol. 14004, pp. 395–425. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-30545-0_14
43. Zhou, M., Park, A., Shi, E., Zheng, W.: PIANO: extremely simple, single-server PIR with sublinear server computation. IACR Cryptol. ePrint Arch., p. 452 (2023)