

On Quantum Advantage in Information Theoretic Single-Server PIR

Dorit Aharonov*

Zvika Brakerski[†]Kai-Min Chung[‡]

Ayal Green*

Ching-Yi Lai[§]Or Sattath[¶]

Abstract

In (single-server) Private Information Retrieval (PIR), a server holds a large database DB of size n , and a client holds an index $i \in [n]$ and wishes to retrieve $\text{DB}[i]$ without revealing i to the server. It is well known that information theoretic privacy even against an “honest but curious” server requires $\Omega(n)$ communication complexity. This is true even if quantum communication is allowed and is due to the ability of such an adversarial server to execute the protocol on a superposition of databases instead of on a specific database (“input purification attack”). Nevertheless, there have been some proposals of protocols that achieve sub-linear communication and appear to provide some notion of privacy. Most notably, a protocol due to Le Gall (ToC 2012) with communication complexity $O(\sqrt{n})$, and a protocol by Kerenidis et al. (QIC 2016) with communication complexity $O(\log(n))$, and $O(n)$ shared entanglement.

We show that, in a sense, input purification is the only potent adversarial strategy, and protocols such as the two protocols above are secure in a restricted variant of the quantum honest but curious (a.k.a specious) model. More explicitly, we propose a restricted privacy notion called *anchored privacy*, where the adversary is forced to execute on a classical database (i.e. the execution is anchored to a classical database). We show that for measurement-free protocols, anchored security against honest adversarial servers implies anchored privacy even against specious adversaries.

Finally, we prove that even with (unlimited) pre-shared entanglement it is impossible to achieve security in the standard specious model with sub-linear communication, thus further substantiating the necessity of our relaxation. This lower bound may be of independent interest (in particular recalling that PIR is a special case of Fully Homomorphic Encryption).

1 Introduction

Private Information Retrieval (PIR), introduced by Chor et al. [CGKS95], is perhaps the most basic form of joint computation with privacy guarantee. PIR is concerned with privately retrieving an entry from a database, without revealing which entry has been accessed. Formally, a PIR protocol is a communication protocol between two parties, a server holding a large database DB

*Hebrew University of Jerusalem.

[†]Weizmann Institute of Science.

[‡]Academia Sinica.

[§]National Chiao Tung University.

[¶]Ben-Gurion University.

containing n binary entries¹, and a client who wishes to retrieve the i th element of the database but without revealing the index i . Privacy can be defined using standard cryptographic notions such as indistinguishability or simulation (see [Gol04]). The simplicity of this primitive is since there is no privacy requirement for the database (i.e. we allow sending more information than necessary) and that the server is not required to produce any output in the end of the interaction, so functionality and privacy are *one sided*.

Clearly PIR is achievable by sending all of DB to the client. This will have communication complexity n and will be perfectly private under any plausible definition since the client sends no information. The absolute optimal result one could hope for is a protocol with logarithmic communication, matching the most communication efficient protocol without privacy constraints, in which the client sends the index i to the server and receives $\text{DB}[i]$ in response.

Alas, [CGKS95] proved that linear (in n) communication complexity is *necessary* for PIR, and that this is the case even in the presence of arbitrary setup information.² Despite its pessimistic outlook, this lower-bound served (already in [CGKS95] itself) as starting point to two extremely prolific and influential lines of research, showing that the communication complexity can be vastly improved if we place some restrictions on the server. The first considered *multiple non-interacting* servers (see, e.g., [Efr12, DG15] and references therein), instead of just a single server, and the second considered *computationally bounded* servers and relying on *cryptographic assumptions* (see, e.g., [CMS99, Gen09, BV11]).

While our discussion so far referred to protocols executed by classical parties over classical communication channels, the focus of this work is on the quantum setting, where there is a quantum communication channel between the client and server, and where the parties themselves are capable of performing quantum operations. Importantly, we still only require functionality for a classical database and a classical index.

One could hope that introducing quantum channels could allow an information theoretic solution to a problem that classically can only be solved using cryptographic assumptions, as has been the case for quantum key distribution [BB84], quantum money [Wie83], quantum digital signatures [GC01], quantum coin-flipping [Moc07, CK09, ACG⁺16] and more [BS16]. Indeed, the notion of Quantum PIR (or QPIR) is quite a natural extension of its classical counterpart and has also been extensively studied in the literature. Nayak’s famous result on the impossibility of random access codes [Nay99] implies a linear lower bound for non-interactive protocols (ones that consists of only a single message from the server to the client), and implicitly, via extension of the same methods, also for multi-round protocols. Formal variants of this lower bound were proven also by Jain, Radhakrishnan and Sen [JRS09] (in terms of quantum mutual information) and by Baumeler and Broadbent [BB15]. Indeed, one could trace back all of these results to the notion of *adversary purification* which was used to show the impossibility of various cryptographic tasks in the information-theoretic quantum model starting as early as [Lo97, LC97, May97]. In the context of QPIR, it can be shown that executing a QPIR protocol with sub-linear communication on a superposition of databases instead of on a single database, will leave the server at the end of the execution with a state that reveals some information about the index i . This is made explicit in [JRS09, Section 3.1] and is also implicit in the proof of [BB15].

¹Throughout this work we will focus on the setting of binary database. We do note that there is vast literature concerned with optimizations for the case of larger alphabet.

²Setup refers to any information that is provided to the parties prior to the execution of the protocol by a trusted entity, but crucially one that does not depend on the parties’ inputs. Shared randomness or shared entanglement are common examples.

Most relevant to our work is the aforementioned [BB15], which provides an analysis from a cryptographic perspective and considers a well defined adversarial model known as privacy against specious adversaries, or *the specious model* for short. This adversarial model was introduced by Dupuis, Nielsen and Salvail [DNS10] as a quantum counterpart to the classical notion of *honest but curious* (a.k.a semi-honest) adversaries.³ A specious adversary can be thought of as one that contains, as a part of its local state, a sub-state which is indistinguishable from that of the respective honest party, even when inspected jointly with the other party’s local state.⁴

Let us provide a high level description of the specious model. We provide a general outline for two-party protocols, and not one that is specific to QPIR. Consider a protocol executed between parties A, B on input registers X, Y respectively. Let A, B also denote the local state of the parties at a given point in time. Then the state of an honest execution of the protocol on inputs XY can be described by the joint density matrix of the registers $XABY$. A specious adversarial strategy for party A can be thought of as one where at any point in time, the local state of the adversary is of the form $A'XA$ (i.e. the adversary is allowed to maintain additional information, possibly in superposition with other parts of the system), such that the reduced density matrix of $XABY$ is still indistinguishable from the one obtained in an honest execution. This provides a potential advantage to a specious adversary (compared to an honest A) since it is quite possible that together with A' , the joint state is no longer honest. Thus the local view of the adversary, i.e. the registers $A'XA$, might in fact reveal information about B ’s input Y that was supposed to have been kept private.

In the QPIR setting, say taking A to be the server and B to be the client, the register X holds the database DB , and Y holds the index i . Indeed, [BB15] shows that it is sufficient that A' contains a purification of XA , where X is a uniform distribution over all databases. We call this the *purification attack*. Thus, while the adversary pretends to execute the protocol on a randomly sampled database, it is in fact executed on a superposition of all possible databases at the same time (indeed this is the case since A' contains a purification of X). As explained above, this methodology is not new, but [BB15] analyze and show that no meaningful notion of QPIR can be achieved against this class of adversaries.

While the negative results could leave us pessimistic as to the abilities of quantum techniques to improve the state of the art on single-server PIR, there is some optimism suggested by two works. Le Gall [LG12] proposed a protocol with sub-linear communication (specifically $O(\sqrt{n})$). Kerenidis et al. [KLGR16] proposed two protocols – an explicit one, with $O(\log n)$ communication, which requires linear pre-shared entanglement; and a second protocol, with poly-logarithmic communication (and does not require pre-shared entanglement). In terms of privacy, it is shown that in a perfectly honest execution of the protocol, client’s privacy is preserved. It might not be immediately clear how to translate this proof of privacy to the existing security models and reconcile it with the negative results. It is explained in [LG12] that the protocol is not actually secure if the server deviates from the protocol. However, as [BB15] observed, even a specious attacker that purifies the adversary can violate the security of the protocol, and the privacy proof strongly hinges on the honest execution using a classical database.

³As [DNS10] point out, their model is stronger, i.e. excludes a larger class of attacks, compared to the honest but curious model, even when restricted to a completely classical setting.

⁴More accurately, indistinguishability is required to hold even in the presence of an environment which can be arbitrary correlated (or entangled) with the parties’ inputs. In the quantum setting this usually corresponds to the environment.

Challenges. The state of affairs prior to this work, was that (non-trivial) QPIR was proven impossible even against fairly weak adversaries (namely, specious). Nevertheless, it appears that [LG12, KLGR16] achieve some non-trivial privacy guarantee using sub-linear communication. This privacy guarantee appears not to be captured by the existing security model. Lastly, we notice that all existing negative results are proven in a standalone model and did not consider protocols where the parties are allowed to share (honestly generated) setup information, such as the one by Kerenidis et al. [KLGR16]. In the quantum setting, a natural question is whether shared entanglement can help in achieving a stronger result.⁵ The goal of this work is to address these challenges.

1.1 Our Results

Anchored Privacy. We start by formalizing a refinement of the standard notion of quantum privacy - one where the adversary is not allowed to purify its input register. We show anchored privacy against specious adversaries follows from anchored privacy against an honest party, if the protocol itself does not require parties to perform measurements (i.e. is measurement-free). Formally, using our notation from above, privacy in our model is only required to hold if the reduced density matrix of the register X is a standard basis element, i.e. a fixed classical value. We call our model *anchored* privacy as we can view our adversary as anchored to a specific value for its input X .

We observe that Le Gall’s $O(\sqrt{n})$ protocol [LG12] and the two protocols mentioned above by Kerenidis et al. [KLGR16] are in fact private against honest servers. We prove that explicitly for the pre-shared entanglement protocol by Kerenidis et al. in Appendix B. Using our reduction we can deduce that these protocols are also anchored private against specious adversaries, namely that so long as the adversary does not attempt to execute the protocol on a superposition of databases (and is still specious in the manner explained above), privacy is guaranteed. In a sense, we formalize the folklore reliance on input purification to attack cryptographic schemes (and QPIR in particular), and show that in a model where input purification is impossible or prevented via some external restriction, it is possible to achieve security against specious adversaries.

We believe this model is interesting for three main reasons:

1. Conceptually, this model helps clarify the exact reason for the impossibility of QPIR - it is precisely because of the purification attack. Indeed, there is a formal sense in which some anchoring is necessary since we know that for any proposed protocol, allowing to execute on a superposition of inputs allows to violate security – see the preceding discussion in Section 1.
2. We view the anchored specious model as a stepping stone towards more robust notions. One intriguing future direction (mentioned briefly in our list of open problems) is to try to develop a malicious analog that still implements the ideology of “forbidden input purification, e.g. by forcing the adversary to “classically open the database before or after the execution in a manner that is consistent with the clients output. Another interesting direction is to try to enforce anchoring using a two-server setting, thus achieving logarithmic two-server QPIR (which is currently still beyond reach).

⁵We note that to the best of our understanding, even prior “entropic” results such as [JRS09] seem to fall short of capturing the potential additional power of shared entanglement. This is essentially due to the property that if AB are entangled, then it is possible that the reduced state of B will have (much) higher von Neumann entropy than the joint AB (whose entropy might even be 0).

3. We believe that our new model may be plausible in certain situations where one could certify that the server cannot employ a superposition on databases. We note that this model can be externally enforced, e.g. by conducting an inspection of the server’s local computation device (with a very low probability) and making sure that it complies, and otherwise apply a heavy penalty. One could imagine such an inspection verifying that a copy of the database is stored on a macroscopic device that cannot be placed in superposition using available technology. Another example of a setting where the anchored model could be applicable is when the database contains information with some semantic meaning, so that the client can easily notice when a nonsense value has been used (this is somewhat similar to the setting considered in [GLM08]). We recall that semi-honest protocols are often used as building blocks, with additional external mechanisms that are employed to validate the assumptions of the model, and hope that our model can also be used in this way. Lastly, from a purely scientific perspective, we believe that formalizing and pinpointing a non-trivial model where non-trivial QPIR is possible will allow to better understand this primitive and the relation between quantum privacy and its classical counterpart.

Improved Lower Bound. It would be instrumental to understand why the known QPIR lower bounds do not apply to our logarithmic protocol described above. Specifically, the protocol makes use of setup (pre-shared entanglement), and one could wonder whether this is the source of improvement, and perhaps with pre-shared entanglement it is possible to prove security even in the standard specious model. We show that this is not the case by providing a lower bound in the specious model even for the *one-sided communication* from the server to the client. Namely, we show that linear communication from the server to the client is necessary even if we allow arbitrary communication from the client to the server. In particular, this rules out the ability to use the setup to circumvent the lower bound, since the client (which is assumed to be honest) can generate the setup locally, and send the server’s share across the channel at the beginning of the protocol. This completes the picture in terms of the impossibility of QPIR in the specious model and further justifies our relaxation of the model in order to achieve meaningful results.

Noting that PIR can be thought of as a special case of Fully Homomorphic Encryption (FHE), our lower bound implies that even a Quantum Fully Homomorphic Encryption (QFHE) with (even approximate) information theoretic security cannot have non-trivial communication complexity, even if the QFHE protocol is allowed to make use of shared entanglement between the server and the client. We thus generalize (to allow shared prior entanglement) the impossibility results for (even imperfect) QPIR of [BB15] (as well as those of [YPF14] which explicitly referred to QFHE).

1.2 Overview of Our Techniques

Anchored-Specious Security. Recall the notation introduced above for two party protocol (A, B) on inputs (X, Y) , and recall that a specious adversary can be thought of as one where the local state of the adversary is of the form $A'XA$. Now let us consider the case of measurement-free protocols and also assume that the client’s input Y is a pure state (this can be justified since otherwise we can apply our argument on the joint state of Y and its purifying environment instead of Y itself). In such an execution, it holds that at any stage $XABY$ is a pure superposition (i.e. its density matrix is of rank 1). Now let us consider the joint state together with the specious adversary’s additional register, i.e. $A'XABY$. Since $(XABY)$ is pure, A' cannot be entangled with it, and therefore A' is in tensor product with the remainder of the state, namely $(XABY)$.

It follows that the status of the register A' can be simulated at any point in time without any knowledge of the other components of the protocol. There is a delicate point here, since A' may indeed be in tensor product, but we must also argue that it is independent of Y . Intuitively, to see why such dependence on Y cannot occur consider, e.g., $Y = |y_1\rangle + |y_2\rangle$. Then YA' is in the state $|y_1\rangle \otimes \rho_{A'} + |y_2\rangle \otimes \rho_{A'}$ (importantly the same $\rho_{A'}$ appears twice). However, this state is exactly the purification of executing the protocol either with $Y = |y_1\rangle$ or with $Y = |y_2\rangle$. We conclude that $\rho_{A'}$ must be the same in both settings, and by extension it can be shown to be the same for all Y .

After taking care of A' , we need to consider the other part of the adversary's state, namely the register (XA) . This register is, by definition, identical (or indistinguishable) from the state of an honest party during the execution. Recall that we assume our protocol is anchored private against honest servers. So the local honest state (XA) is guaranteed not to leak information about B 's input. Add to that the conclusion about A' being in tensor product and independent of B 's state, and we get that the entire local state of the specious adversary does not reveal any disallowed information.

As a conclusion, since we can show, e.g. in Le Gall's protocol or in our logarithmic protocol, that an honest execution with a classical X does not leak information about Y , this will also be the case in the anchored-specious setting.

Obviously many details are omitted from this high level overview. For example, a specious adversary is not required to make $(XABY)$ identical to an honest execution but rather only statistically close (in trace distance), which requires a more delicate analysis. Furthermore, the formal construction of a simulator for the adversary as required by the specious definition requires some care to detail. For the formal definitions and analysis see Section 3 below.

Our Lower Bound. We first note that previous lower bound proofs in [Nay99, BB15] bounded the *total* communication complexity by a reduction to quantum random access codes. It is not a-priori clear how to generalize this proof method to the presence of shared entanglement. To do so, we provide a new lower bound argument that establishes a linear lower bound on the *server's* communication complexity. Specifically, we show that the server needs to transmit at least roughly $n/2$ qubits to the client, no matter how many qubits is transmitted from the client to the server (assuming that the protocol has sufficiently small correctness and privacy error). As we mentioned above, such a lower bound trivially extends to hold with prior shared entanglement, since one can think of that the shared entanglement is established by the client sending messages to the server.

Our new lower bound argument is based on an interactive leakage chain rule in [LC18] and might even be considered conceptually simpler than previous methods. At a high level, we consider a server holding a uniformly random database $\mathbf{a} \in \{0, 1\}^n$ and running a QPIR protocol with a client. Initially, from the client's point of view, the database \mathbf{a} has n -bits of min-entropy, and the protocol execution can be viewed as an "interactive leakage" that leaks information about \mathbf{a} to the client. Let m_A and m_B denote the server and the client's communication complexity in the protocol. The interactive leakage chain rule in [LC18] states that the min-entropy of \mathbf{a} can only be decreased by at most $\min\{2m_A, m_A + m_B\}$. More precisely, let ρ_{AB} denote the states at the end of the protocol execution where the A register stores the (classical) random database \mathbf{a} and B denotes the client's local register. The interactive leakage chain rule states that $H_{\min}(A|B)_\rho \geq n - \min\{2m_A, m_A + m_B\}$. By the operational meaning of quantum min-entropy, given the client's state ρ_B , one cannot predict the database correctly with probability higher than $2^{-(n - \min\{2m_A, m_A + m_B\})}$. On the other hand, suppose the protocol is secure against specious servers with sufficiently small correctness and privacy

error. We can combine the by-now standard lower bound argument by Lo [Lo97] and gentle measurement [Win99, Aar04, ON07], we show that one can reconstruct the database \mathbf{a} from the client's state ρ_B with a constant probability. Combining both claims allows us to establish lower bounds on both the server's and the total communication complexity in a unified way.

1.3 Remaining Open Problems

We proposed a new model and a new protocol which, we believe, resurfaces the question of what can be achieved in the context of QPIR. We believe that a number of intriguing questions still remain for future work.

1. As discussed above, our model is a relaxation of the specious model, which is by itself a semi-honest model. Such models are fairly restrictive in the sense that they make structural assumptions on the adversary (i.e. that it follows the protocol, or contains a part that follows the protocol). Obviously, if we hope for non-trivial results, any model that we formalize must preclude purification of input. It is thus an intriguing question whether it is possible to formulate *malicious* adversarial models that are still purification-free, and what can be said about the plausibility of QPIR in such models. The current definition of anchored privacy will need to be amended, since a malicious server is allowed to just ignore its prescribed input, so a different method of anchoring needs to be devised.
2. Another natural question is whether setup is necessary to achieve logarithmic QPIR in the anchored specious model. We know from Kerenidis et al.'s result that polylogarithmic communication is achievable even without setup. Is there a reason can only improve it when assuming a setup? Another surprising aspect is that the shared entanglement created during the setup is not consumed during the protocol, and can be used for other needs after the execution of the protocol (e.g., running another execution of PIR, or teleportation). A similar phenomenon occurs in quantum information: catalyst quantum states are useful for mapping one bi-partite state to another using LOCC, without consuming the catalyst state [JP99, Kli07]. The related notion of quantum embezzlement [vdH03] has a similar property, but in this case, the original shared state changes slightly. The authors are not aware of any other cryptographic protocol with this non-consumption property.
3. Most state of the art classical PIR protocols (both in the multi-server setting and in the computational cryptographic setting) only require one round of communication. That is, one message (query) from the client to the server (or servers) and one response message. All the existing sublinear QPIR protocols have multiple rounds. Understanding the round complexity of QPIR in light of the classical state of the art is also an intriguing direction.
4. A main contribution of this work is to formalize the notion of anchored security and show it can be used to provide a non-trivial cryptographic primitive. It would be interesting to study the relevance of this notion (or adequately adapted versions) in the context of a variety of other cryptographic tasks. In particular, the question of whether it is possible to construct information theoretically secure fully homomorphic encryption (FHE) given quantum channels has received attention in recent years (see, e.g., [YPF14]). In homomorphic encryption, the server has a function f and the client has an input x , and the goal of the protocol is for the client to learn $f(x)$ without revealing any information about x . PIR and FHE functionalities

are intimately related (think about a function $f_{\text{DB}}(i) = \text{DB}[i]$ for FHE, and about executing PIR with database equal to the truth table of some function), and it is thus intriguing whether the anchored model is applicable in the context of FHE as well.

1.4 Paper Organization

General preliminaries are provided in Section 2. We present our new model, and the proof that for pure protocols honest security implies anchored specious security in Section 3. Our new lower bound is stated and proven in Section 4. In Appendix B, we show that the protocol by Kerenidis et al. is anchored private against specious adversaries.

2 Preliminaries

Standard preliminaries regarding Hilbert spaces and quantum states can be found in Appendix A. We provide below background and definitions concerning two-party quantum protocols, specious adversaries and quantum private information retrieval.

2.1 Two-Party Quantum Protocols

As in [BB15], we base our definitions on the works of [GW07] and [DNS10]. However, we make slight adaptations to allow for prior entanglement between the parties.

Definition 2.1 (Two-party quantum protocol). An s -round, two-party quantum protocol, denoted $\Pi = \{\mathcal{A}, \mathcal{B}, \rho_{\text{joint}}, s\}$ consists of:

1. input spaces \mathcal{A}_0 and \mathcal{B}_0 for parties \mathcal{A} and \mathcal{B} respectively,
2. initial spaces \mathcal{A}_p and \mathcal{B}_p (p for pre-shared state) for parties \mathcal{A} and \mathcal{B} respectively,
3. a joint initial state $\rho_{\text{joint}} \in \mathcal{A}_p \otimes \mathcal{B}_p$, split between the two parties,
4. memory spaces $\mathcal{A}_1, \dots, \mathcal{A}_s$ for \mathcal{A} and $\mathcal{B}_1, \dots, \mathcal{B}_s$ for \mathcal{B} , and communication spaces $\mathcal{X}_1, \dots, \mathcal{X}_s, \mathcal{Y}_1, \dots, \mathcal{Y}_{s-1}$,
5. an s -tuple of quantum operations $(\mathcal{A}_1, \dots, \mathcal{A}_s)$ for \mathcal{A} , where $\mathcal{A}_1 : L(\mathcal{A}_0 \otimes \mathcal{A}_p) \mapsto L(\mathcal{A}_1 \otimes \mathcal{X}_1)$, and $\mathcal{A}_t : L(\mathcal{A}_{t-1} \otimes \mathcal{Y}_{t-1}) \mapsto L(\mathcal{A}_t \otimes \mathcal{X}_t)$ ($2 \leq t \leq s$),
6. an s -tuple of quantum operations $(\mathcal{B}_1, \dots, \mathcal{B}_s)$ for \mathcal{B} , where $\mathcal{B}_1 : L(\mathcal{B}_0 \otimes \mathcal{B}_p \otimes \mathcal{X}_1) \mapsto L(\mathcal{B}_1 \otimes \mathcal{Y}_1)$, $\mathcal{B}_t : L(\mathcal{B}_{t-1} \otimes \mathcal{X}_t) \mapsto L(\mathcal{B}_t \otimes \mathcal{Y}_t)$ ($2 \leq t \leq s-1$), and $\mathcal{B}_s : L(\mathcal{B}_{s-1} \otimes \mathcal{X}_s) \mapsto L(\mathcal{B}_s)$.

Note that in order to execute a protocol as defined above, one has to specify the input, namely a quantum state $\rho_{\text{in}} \in S(\mathcal{A}_0 \otimes \mathcal{B}_0)$ from which the execution starts.

Definition 2.2 (Protocol Execution). If $\Pi = \{\mathcal{A}, \mathcal{B}, \rho_{\text{joint}}, s\}$ is an s -round two-party protocol, then the state after the t -th step ($1 \leq t \leq 2s$), and upon input state $\rho_{\text{in}} \in S(\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R})$, for any \mathcal{R} , is defined as

$$\rho_t(\rho_{\text{in}}) := (\mathcal{A}_{(t+1)/2} \otimes I_{\mathcal{B}_{(t-1)/2}}) \dots (\mathcal{B}_1 \otimes I_{\mathcal{A}_1})(\mathcal{A}_1 \otimes I_{\mathcal{B}_0, \mathcal{B}_p})(\rho_{\text{in}} \otimes \rho_{\text{joint}}),$$

for t odd, and

$$\rho_t(\rho_{in}) := (\mathcal{B}_{t/2} \otimes I_{\mathcal{A}_{t/2}}) \dots (\mathcal{B}_1 \otimes I_{\mathcal{A}_1})(\mathcal{A}_1 \otimes I_{\mathcal{B}_0, \mathcal{B}_p})(\rho_{in} \otimes \rho_{joint}),$$

for t even. We define the final state of protocol $\Pi = \{\mathcal{A}, \mathcal{B}, \rho_{joint}, s\}$ upon input state $\rho_{in} \in S(\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R})$ as: $[\mathcal{A} \otimes \mathcal{B}](\rho_{in}) := \rho_{2s}(\rho_{in})$.

The communication complexity of a protocol is the number of qubits that are exchanged between the parties. Slightly more generally, we can consider the logarithm of the dimension of the message registers $\mathcal{X}_t, \mathcal{Y}_t$. The formal definition thus follows.

Definition 2.3 (Communication Complexity). The *communication complexity* of a protocol as in Definition 2.1 is

$$\sum_{t=1}^s \log \dim(\mathcal{X}_t) + \sum_{t=1}^{s-1} \log \dim(\mathcal{Y}_t) .$$

We sometimes also refer to one-sided communication complexity, i.e. the total communication originating from one party to the other. The communication complexity of \mathcal{A} is defined to be the communication originating from \mathcal{A} , or formally $\sum_{t=1}^s \log \dim(\mathcal{X}_t)$. Symmetrically the communication complexity of \mathcal{B} is $\sum_{t=1}^{s-1} \log \dim(\mathcal{Y}_t)$.

2.2 Specious Adversary

Given a two-party quantum protocol $\Pi = \{\mathcal{A}, \mathcal{B}, \rho_{joint}, s\}$, an adversary $\tilde{\mathcal{A}}$ for \mathcal{A} is an s -tuple of quantum operations $(\mathcal{A}_1, \dots, \mathcal{A}_s)$, where $\mathcal{A}_1 : L(\mathcal{A}_0) \mapsto L(\mathcal{A}_1 \otimes \mathcal{X}_1)$ and $\mathcal{A}_t : L(\tilde{\mathcal{A}}_{t-1} \otimes \mathcal{Y}_{t-1}) \mapsto L(\tilde{\mathcal{A}}_t \otimes \mathcal{X}_t)$, $2 \leq t \leq s$, with $\tilde{\mathcal{A}}_1, \dots, \tilde{\mathcal{A}}_s$ being $\tilde{\mathcal{A}}$'s memory spaces. The global state after the t th step of a protocol run with $\tilde{\mathcal{A}}$ is $\tilde{\rho}_t(\tilde{\mathcal{A}}, \rho_{in})$. An adversary $\tilde{\mathcal{B}}$ for \mathcal{B} is similarly defined.

Definition 2.4 (Specious adversaries). Let $\Pi = \{\mathcal{A}, \mathcal{B}, \rho_{joint}, s\}$ be an s -round two-party protocol. An adversary $\tilde{\mathcal{A}}$ for \mathcal{A} is said to be γ -specious, if there exists a sequence of quantum operations (called recovery operators) $\mathcal{F}_1, \dots, \mathcal{F}_{2s}$, such that for $1 \leq t \leq 2s$ and for all $\rho_{in} \in S(\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R})$:

1. For all t even, $\mathcal{F}_t : L(\tilde{\mathcal{A}}_{t/2}) \mapsto L(\mathcal{A}_{t/2})$.
2. For all t odd, $\mathcal{F}_t : L(\tilde{\mathcal{A}}_{(t+1)/2} \otimes \mathcal{X}_{(t+1)/2}) \mapsto L(\mathcal{A}_{(t+1)/2} \otimes \mathcal{X}_{(t+1)/2})$.
3. For every input state $\rho_{in} \in S(\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R})$, for any \mathcal{R} ,

$$\Delta \left((\mathcal{F}_t \otimes I_{\mathcal{B}_t, \mathcal{R}}) \left(\tilde{\rho}_t(\tilde{\mathcal{A}}, \rho_{in}) \right), \rho_t(\rho_{in}) \right) \leq \gamma. \quad (1)$$

A γ -specious adversary $\tilde{\mathcal{B}}$ for \mathcal{B} is similarly defined.

2.3 Quantum Private Information Retrieval

We define QPIR similarly to [BB15].

Definition 2.5 (Quantum Private Information Retrieval). An s -round, n -bit Quantum Private Information Retrieval protocol (QPIR) is a two-party protocol $\Pi_{QPIR} = \{\mathcal{A}, \mathcal{B}, \rho_{joint}, s\}$, where \mathcal{A} is the server, \mathcal{B} is the client, and ρ_{joint} is an initial state shared between them prior to the protocol.

We call Π_{QPIR} $(1-\delta)$ -correct if, for all inputs $\rho_{in} = |x\rangle\langle x|_{\mathcal{A}_0} \otimes |i\rangle\langle i|_{\mathcal{B}_0}$, with $x = x_1, \dots, x_n \in \{0, 1\}^n$ and $i \in \{1, \dots, n\}$, there exists a measurement \mathcal{M} acting on \mathcal{B}_s with outcome 0 or 1, such that:

$$\Pr \{ \mathcal{M}(\text{tr}_{\mathcal{A}_s}[\mathcal{A} \otimes \mathcal{B}](\rho_{in})) = x_i \} \geq 1 - \delta .$$

If $\delta = 0$ we say that the protocol is perfectly correct.

We call Π_{QPIR} ϵ -private against a (possibly adversarial) server $\tilde{\mathcal{A}}$, if there exists a sequence of quantum operations (simulators) $\mathcal{J}_1, \dots, \mathcal{J}_{s-1}$, where $\mathcal{J}_t : L(\mathcal{A}_0 \otimes \mathcal{A}_p) \mapsto L(\tilde{\mathcal{A}}_t \otimes \mathcal{Y}_t)$, such that for all $1 \leq t \leq s-1$ and for all $\rho_{in} \in S(\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R})$,

$$\Delta \left(\text{tr}_{\mathcal{B}_0}(\mathcal{J}_t \otimes \mathbb{I}_{\mathcal{B}_0, \mathcal{R}}(\rho_{in})), \text{tr}_{\mathcal{B}_t}(\tilde{\rho}_{2t}(\tilde{\mathcal{A}}, \rho_{in})) \right) \leq \epsilon . \quad (2)$$

If $\epsilon = 0$ we say that the protocol is perfectly private.

We say that a QPIR protocol is ϵ -private against a class of servers if it is ϵ -private against any server from this class.

We note that in the above definition privacy is required to hold also for adversarial input states for the client and server, which also includes inputs in superposition, and even for the case where the client and server (and possibly a third party) are entangled. Nayak [Nay99, ANTsv02] showed that a perfectly private QPIR protocol, even only against 0-specious servers, must have communication complexity at least $(1 - H(1 - \delta))n$, where $H(p)$ is the binary entropy function. Baumeler and Broadbent [BB15] extended this lower bound to the case of $\epsilon > 0$ and presented a communication lower bound of

$$\left(1 - H \left(1 - \delta - 2\sqrt{\epsilon(2 - \epsilon)} \right) \right) n . \quad (3)$$

3 Anchored Privacy Against Specious Adversaries

We now present our new restricted notion of privacy, that we call *anchored privacy*. A protocol is anchored private if it satisfies the standard definition of privacy with respect to classical inputs on the adversary's side. There is no privacy requirement for superposition input states on the adversary's side (and therefore this notion of privacy is weaker, and hence, easier to achieve). A formal definition follows.

Definition 3.1 (Anchored Privacy). A QPIR protocol is anchored ϵ -private if Eq. (2) holds for all $\rho_{in} \in \mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R}$ (for any \mathcal{R}), for which $\rho_{in}|_{\mathcal{A}_0} = |x\rangle\langle x|$ for some $x \in \{0, 1\}^n$.

We note that prior intuitive notions of security such as that implied by the analysis of Le Gall [LG12] in fact correspond to anchored privacy against honest servers. Our main theorem below shows that this type of privacy extends to the specious setting as well.

Theorem 3.2. Let Π be a measurement-free QPIR protocol which is anchored ϵ -private against honest servers, then Π is anchored $(\epsilon + 3\sqrt{2\gamma})$ -private against γ -specious servers.

Critically, the theorem only holds for measurement-free QPIR protocols. To see this, consider the following protocol, which will be anchored-private against honest servers but not anchored-private against specious ones. Let Π be a QPIR protocol which is anchored-private against honest servers (e.g., Le-Gall's protocol [LG12]). Now consider the following protocol Π' which first generates a superposition over all possible databases, then measures this superposition to obtain a

classical value for the database. It then runs Π on this measured database (with the client using its real input index). Finally, both parties toss out the output of this first execution, and run Π again, now using the actual input database.

Let us first see that Π' is anchored-private against honest servers. This follows since Π is secure against honest adversaries when executed over input states in which the server's input is classical, and hence so is Π' which just consists of two sequential executions of Π over classical databases. However, a purification of an honest server allows to execute a purification attack on the first execution of Π in a way that allows to recover the client's input, even though the database used as input for Π' is completely classical.

Warm-up. We first give a proof under some simplifying assumptions: (i) $\gamma = \epsilon = 0$. (ii) the input is pure (iii) the purification space is trivial: $\mathcal{R} = \mathbb{C}$ and (iv) the specious server's quantum operations $\tilde{\mathcal{A}}_t$ are unitary. The main point that makes the analysis easier in this case is assumption (i).

Fix a step of the protocol t .

1. We claim that for every unitary γ -specious adversary, which is perfect (i.e. $\gamma = 0$) the entire state, (written in some *fixed* but maybe non standard basis), is of the form $|\eta\rangle_{\mathcal{S}'} \otimes |\psi_t\rangle_{\mathcal{S},\mathcal{C}}$ where $|\psi_t\rangle$ is the state that an honest server and client would have when running on the same input. Here, and later, we use the notation \mathcal{S} for all of the honest server registers at step t , \mathcal{C} for all of the client's registers at step t and \mathcal{S}' for the specious server's ancillary register at step t . Crucially, $|\eta\rangle$ is independent of the (server and client) input.

We now prove the above claim. By the specious property, we know that there exists a quantum operation \mathcal{F}_t which maps the global state at the t th stage to the state $|\psi_t\rangle$. We know that the state in step t in the honest run is necessarily pure since Π is measurement free. W.l.o.g. we can assume that the operation \mathcal{F}_t is a unitary U_t , followed by tracing out everything other than the \mathcal{S} and \mathcal{C} registers.

Let's assume towards contradiction that the state in the basis U_t^\dagger is of the form $|\eta(input)\rangle \otimes |\psi_t\rangle$, where $|\eta(input)\rangle$ depends on the input (where here we mean both the client and the server's input). There must be two different input states such that running them would give $|\eta(1)\rangle \otimes |\psi_t(1)\rangle$ and $|\eta(2)\rangle \otimes |\psi_t(2)\rangle$ for which $|\eta(1)\rangle \neq |\eta(2)\rangle$. Since the honest runs are entirely unitary (by the measurement-free property) and have different inputs, necessarily, $|\psi_t(1)\rangle \neq |\psi_t(2)\rangle$. By running the specious adversary on a superposition of these two inputs, we get that after applying \mathcal{F}_t , the state becomes a mixture of the two states, $|\psi_t(1)\rangle$ and $|\psi_t(2)\rangle$. This contradicts the perfect specious property (see Eq. (1)) – which requires the state to be the pure (since all the operations of the client and honest servers are unitary, and their input in this case is pure).

2. By the perfect anchored-privacy against the honest server, the state $\rho_t = \text{tr}_C(|\psi_t\rangle\langle\psi_t|_{\mathcal{S},\mathcal{C}})$ is independent of the client's input, and therefore, could only depend on x – the server's input. To emphasize that independence on the client's input (and possible dependence on the server's input), we denote the state ρ_t by $\rho_t(x)$.

Our goal is to show the anchored-privacy property for the specious server. Indeed, the two points above show that the specious server's state (in the fixed basis we choose to work in) is $|\eta\rangle\langle\eta| \otimes \rho_t(x)$, which is independent of the client's input. Therefore the simulator can generate that state exactly by using the server's classical input x , as required (see Eq. (2)).

Outline of the general proof. For each round t we construct a simulator for the server in the following way: we first construct a simulator $\tilde{\mathcal{S}}_t^{x_0,0}$ for input $|x_0\rangle \otimes |0\rangle$ where $|x_0\rangle$ is an input for the server and $|0\rangle$ is an input for the client. We construct this simulator using the simulator for the honest server along with the 'specious operator', and an ancillary state $|\sigma_{x_0,0}\rangle$. We then show that $|\sigma_{x_0,0}\rangle$ is also an appropriate ancillary state for any input $|x\rangle \otimes |\eta\rangle$. Using this, we show that $\tilde{\mathcal{S}}_t^{x,0}$ is indeed a simulator for any input $|x\rangle \otimes |\eta\rangle$, with slightly worse parameters.

We are now ready to give the proof in full generality:

Theorem 3.2 (Proof). Let Π be a purified QPIR protocol which is anchored ϵ -private against honest servers, and let $\tilde{\mathcal{S}}$ be a γ -specious server for Π . W.l.o.g we can assume that $\tilde{\mathcal{S}}$ is purified, namely, a unitary⁶. From now on, we will fix t . We can denote

$$|\psi_t^{\rho_{in}}\rangle\langle\psi_t^{\rho_{in}}| = \rho_t(\rho_{in}) \quad (4)$$

where $|\psi_t^{\rho_{in}}\rangle \in \mathcal{S} \otimes \mathcal{C} \otimes \mathcal{R}$ for some \mathcal{R} , and we use \mathcal{S} to represent the server's registers $\mathcal{S} = \mathcal{A}_t \otimes \mathcal{Y}_t \otimes \mathcal{A}_p$ (for t odd. otherwise $\mathcal{S} = \mathcal{A}_t \otimes \mathcal{A}_p$), and \mathcal{C} to represent the client's registers $\mathcal{C} = \mathcal{B}_t \otimes \mathcal{X}_t \otimes \mathcal{B}_p$ (for t even. otherwise $\mathcal{C} = \mathcal{B}_t \otimes \mathcal{B}_p$). Furthermore, w.l.o.g we assume the various recovery operators for $\tilde{\mathcal{S}}$ are purified. That is, there exist unitary operators $\hat{\mathcal{F}}_t$ such that $\mathcal{F}_t(\cdot) = \text{tr}_{\mathcal{S}'}(\hat{\mathcal{F}}_t(\cdot))$ for some purification space \mathcal{S}' which is at the hands of the server (from now on, for the sake of this proof, where we say "recovery operators" we regard these unitary $\hat{\mathcal{F}}_t$ operators). Therefore we can denote

$$|\tilde{\psi}_t^{\rho_{in}}\rangle\langle\tilde{\psi}_t^{\rho_{in}}| = \tilde{\rho}_t(\tilde{\mathcal{S}}, \rho_{in}) \quad (5)$$

where w.l.o.g $|\tilde{\psi}_t^{\rho_{in}}\rangle \in \mathcal{S}' \otimes \mathcal{S} \otimes \mathcal{C} \otimes \mathcal{R}$. We note that all of the unitary operators - $\mathcal{A}_t, \mathcal{B}_t$ which are used in the original protocol (by either the server or the client), $\tilde{\mathcal{A}}_t$ which are used by the specious server $\tilde{\mathcal{S}}$, and the recovery $\hat{\mathcal{F}}_t$ operators are independent of both the client's and the server's inputs

For each round t , we will start by constructing a simulator for $\tilde{\mathcal{S}}$ acting on $\rho_{in} = |x_0\rangle\langle x_0|_{\mathcal{A}_0} \otimes |0\rangle\langle 0|_{\mathcal{B}_0}$, where $x_0 \in \{0,1\}^n$ (in this specific input, \mathcal{R} is trivial and is thus omitted). By γ -speciousness of $\tilde{\mathcal{S}}$, along with our purification assumptions, there exists a unitary recovery operator $\hat{\mathcal{F}}_{2t} : L(\tilde{\mathcal{A}}_t) \mapsto L(\mathcal{S}' \otimes \mathcal{A}_t)$ such that

$$\Delta\left(\text{tr}_{\mathcal{S}'}\left(\left(\hat{\mathcal{F}}_{2t} \otimes \mathbb{I}_{\mathcal{C}}\right)|\tilde{\psi}_{2t}^{|x_0\rangle\otimes|0\rangle}\rangle\right), |\psi_{2t}^{|x_0\rangle\otimes|0\rangle}\rangle\right) \leq \gamma \quad (6)$$

By Lemma A.1, this means that there exists a state $|\sigma_{x_0,0}\rangle \in \mathcal{S}'$ such that:

$$\Delta\left(\left(\hat{\mathcal{F}}_{2t} \otimes \mathbb{I}\right)|\tilde{\psi}_{2t}^{|x_0\rangle\otimes|0\rangle}\rangle, |\sigma_{x_0,0}\rangle \otimes |\psi_{2t}^{|x_0\rangle\otimes|0\rangle}\rangle\right) \leq \sqrt{\gamma} \quad (7)$$

We can now operate on Eq. (7) with $\hat{\mathcal{F}}_{2t}^\dagger \otimes \mathbb{I}$ to get:

$$\Delta\left(|\tilde{\psi}_{2t}^{|x_0\rangle\otimes|0\rangle}\rangle, \left(\hat{\mathcal{F}}_{2t}^\dagger \otimes \mathbb{I}\right)\left(|\sigma_{x_0,0}\rangle \otimes |\psi_{2t}^{|x_0\rangle\otimes|0\rangle}\rangle\right)\right) \leq \sqrt{\gamma} \quad (8)$$

⁶This is because we can include the purification register at any point, as the server could have included himself rather than throwing it away

The above connects the states derived from the execution with the specious server to that with the honest server. By anchored ϵ -privacy of Π against honest servers, there exists a simulator $\mathcal{J}_t : L(\mathcal{A}_0 \otimes \mathcal{A}_p) \mapsto L(\mathcal{A}_t \otimes \mathcal{X}_t)$ such that for all $x \in \{0, 1\}^n$ and $|\alpha\rangle \in \mathcal{B}_0 \otimes \mathcal{R}$, for any \mathcal{R} ,

$$\Delta \left(\text{tr}_{\mathcal{B}_0, \mathcal{B}_p} \left((\mathcal{J}_t \otimes \mathbb{I}_{\mathcal{B}_0, \mathcal{B}_p}) \circ (|x\rangle\langle x|_{\mathcal{A}_0} \otimes |\alpha\rangle\langle \alpha|_{\mathcal{R}, \mathcal{B}_0} \otimes \rho_{\text{joint}}) \right), \text{tr}_{\mathcal{B}_t} \left(|\psi_{2t}^{[x] \otimes |\alpha\rangle}\rangle \right) \right) \leq \epsilon \quad (9)$$

(In fact, the above holds for any mixture over such α 's, by convexity). We can now define the simulator for ρ_{in} corresponding to input state $|x_0\rangle \otimes |0\rangle$ to be the following unitary embedding from $\mathcal{A}_0 \otimes \mathcal{A}_p$ to $\mathcal{S}' \otimes \mathcal{A}_0 \otimes \mathcal{A}_p$:

$$\tilde{\mathcal{J}}_t^{x_0, 0}(\cdot) = \hat{\mathcal{F}}_{2t}^\dagger \circ (|\sigma_{x_0, 0}\rangle\langle \sigma_{x_0, 0}| \otimes \mathcal{J}_t(\cdot)) \quad (10)$$

To show that it indeed satisfies the requirements from a simulator, we combine Eqs. (8), (10), and (9) for $x = x_0$, $|\alpha\rangle = |0\rangle$, to get that

$$\Delta \left(\text{tr}_{\mathcal{B}_0, \mathcal{B}_p} \left((\tilde{\mathcal{J}}_t^{x_0, 0} \otimes \mathbb{I}_{\mathcal{B}_0, \mathcal{B}_p}) \circ (|x_0\rangle\langle x_0|_{\mathcal{A}_0} \otimes |0\rangle\langle 0|_{\mathcal{B}_0} \otimes \rho_{\text{joint}}) \right), \text{tr}_{\mathcal{B}_t} \left(|\tilde{\psi}_{2t}^{[x_0] \otimes |0\rangle}\rangle \right) \right) \leq \epsilon + \sqrt{\gamma} \quad (11)$$

We now define the simulator for any input to be this exact simulator:

$$\tilde{\mathcal{J}}_t(\cdot) = \tilde{\mathcal{J}}_t^{x_0, 0}(\cdot); \quad (12)$$

In the remainder of the proof we show that $\tilde{\mathcal{J}}_t(\cdot)$ satisfies an inequality similar to Eq. (11) with respect to all classical server inputs $x \in \{0, 1\}^n$ (not necessarily x_0) and any input state $|\alpha\rangle \in \mathcal{B}_0 \otimes \mathcal{R}$ for any \mathcal{R} , as well as for a mixture of such α 's; this would imply anchored privacy for the specious server. To this end we show that also for this input, a similar inequality to Eq. (11) holds (with a slightly worse bound). Define

$$|x\alpha_+\rangle = \frac{1}{\sqrt{2}}|0\rangle_{\mathcal{R}'}|x_0\rangle_{\mathcal{A}_0}|0\rangle_{\mathcal{B}_0, \mathcal{R}} + \frac{1}{\sqrt{2}}|1\rangle_{\mathcal{R}'}|x\rangle_{\mathcal{A}_0}|\alpha\rangle_{\mathcal{B}_0, \mathcal{R}},$$

where we have added an additional (control) qubit in the space \mathcal{R}' . The specious adversary condition applies to this input state as well, and thus using the same derivation as for Eq. (8)) we get:

$$\Delta \left(|\tilde{\psi}_{2t}^{[x\alpha_+]}\rangle, \left(\hat{\mathcal{F}}_{2t}^\dagger \otimes \mathbb{I} \right) \left(|\sigma_{x\alpha_+}\rangle \otimes |\psi_{2t}^{[x\alpha_+]}\rangle \right) \right) \leq \sqrt{\gamma} \quad (13)$$

Using the fact that neither the server nor the client act on the \mathcal{R}' register, we get:

$$|\psi_{2t}^{[x\alpha_+]}\rangle = \frac{1}{\sqrt{2}}|0\rangle_{\mathcal{R}'} \otimes |\psi_{2t}^{[x_0] \otimes |0\rangle}\rangle_{\mathcal{S}, \mathcal{C}, \mathcal{R}} + \frac{1}{\sqrt{2}}|1\rangle_{\mathcal{R}'} \otimes |\psi_{2t}^{[x] \otimes |\alpha\rangle}\rangle_{\mathcal{S}, \mathcal{C}, \mathcal{R}} \quad (14)$$

Similarly, since the same is true for the adversarial run, we get:

$$|\tilde{\psi}_{2t}^{[x\alpha_+]}\rangle = \frac{1}{\sqrt{2}}|0\rangle_{\mathcal{R}'} \otimes |\tilde{\psi}_{2t}^{[x_0] \otimes |0\rangle}\rangle_{\mathcal{S}, \mathcal{C}, \mathcal{R}} + \frac{1}{\sqrt{2}}|1\rangle_{\mathcal{R}'} \otimes |\tilde{\psi}_{2t}^{[x] \otimes |\alpha\rangle}\rangle_{\mathcal{S}', \mathcal{S}, \mathcal{C}, \mathcal{R}} \quad (15)$$

We plug Eqs. (14) and (15) into Eq. (13), and project the register \mathcal{R}' in the resulting state onto $|1\rangle_{\mathcal{R}'}$ to get:

$$\Delta \left(\frac{1}{\sqrt{2}} |1\rangle_{\mathcal{R}'} \otimes |\tilde{\psi}_{2t}^{[x] \otimes |\alpha\rangle}\rangle_{\mathcal{S}, \mathcal{C}}, \left(\hat{\mathcal{F}}_{2t}^\dagger \otimes \mathbb{I}_{\mathcal{R}, \mathcal{C}} \right) \left(\frac{1}{\sqrt{2}} |1\rangle_{\mathcal{R}'} \otimes |\sigma_{x, \alpha_+}\rangle_{\mathcal{S}'} \otimes |\psi_{2t}^{[x] \otimes |\alpha\rangle}\rangle_{\mathcal{S}, \mathcal{C}} \right) \right) \leq \sqrt{\gamma} \quad (16)$$

Now we apply the fact that $\hat{\mathcal{F}}_{2t}^\dagger$ doesn't act on the client's input; the fact that a unitary operator doesn't change the distance between states; and the fact that tracing out doesn't increase that distance [AKN98], and Eq. (16) becomes:

$$\Delta \left(|\tilde{\psi}_{2t}^{[x] \otimes |\alpha\rangle}\rangle, \left(\hat{\mathcal{F}}_{2t}^\dagger \otimes \mathbb{I} \right) \left(|\sigma_{x\alpha_+}\rangle \otimes |\psi_{2t}^{[x] \otimes |\alpha\rangle}\rangle \right) \right) \leq \sqrt{2\gamma} \quad (17)$$

Similarly, by projecting onto $|0\rangle_{\mathcal{R}'}$ instead of $|1\rangle_{\mathcal{R}'}$ in the derivation of 16, we get

$$\Delta \left(|\tilde{\psi}_{2t}^{[x_0] \otimes |0\rangle}\rangle, \left(\hat{\mathcal{F}}_{2t}^\dagger \otimes \mathbb{I} \right) \left(|\sigma_{x\alpha_+}\rangle \otimes |\psi_{2t}^{[x_0] \otimes |0\rangle}\rangle \right) \right) \leq \sqrt{2\gamma} \quad (18)$$

We now want to apply the triangle inequality to (18), using Eqs. (8). Applying yet again the same sequence of simple argument, namely the fact that unitary transformations preserve the trace distance and tracing out can only decrease it, we get

$$\Delta \left(|\sigma_{x_0, 0}\rangle, |\sigma_{x\alpha_+}\rangle \right) \leq 2\sqrt{2\gamma} \quad (19)$$

And we can use Eq. (19) together with Eq. (17) to get:

$$\Delta \left(|\tilde{\psi}_{2t}^{[x] \otimes |\alpha\rangle}\rangle, \left(\hat{\mathcal{F}}_{2t}^\dagger \otimes \mathbb{I} \right) \left(|\sigma_{x_0, 0}\rangle \otimes |\psi_{2t}^{[x] \otimes |\alpha\rangle}\rangle \right) \right) \leq 3\sqrt{2\gamma} \quad (20)$$

And finally combine Eq. (20), (9) and (12) (in a similar way to how we derived Eq. (11)) to get:

$$\Delta \left(\text{tr}_{\mathcal{B}_0} \left(\left(\tilde{\mathcal{I}}_t \otimes \mathbb{I} \right) (|x\rangle\langle x| \otimes |\alpha\rangle\langle \alpha| \otimes \rho_{\text{joint}}) \right), \text{tr}_{\mathcal{B}_t} \left(|\tilde{\psi}_t^{[x] \otimes |\alpha\rangle}\rangle \right) \right) \leq \epsilon + 3\sqrt{2\gamma}. \quad (21)$$

This finishes our proof. \square

4 Linear Lower Bound in the Specious Model, Even with Prior Entanglement

In this section we show that in the standard specious model, even allowing arbitrarily long prior entanglement, it is still impossible to achieve QPIR with sublinear communication. We do so by presenting a new lower bound argument based on an interactive leakage chain rule in [LC18], which allows us to establish linear lower bounds on both the server's communication complexity and the total communication complexity in a unified way. Then we observe that the lower bound on the server's communication complexity extends trivially to the case with arbitrary prior entanglement. In the following, we state some useful preliminaries in Section 4.1 and present our lower bound in Section 4.2.

4.1 Quantum Information Theory Background

We first recall the notion of quantum min-entropy. Consider a bipartite quantum state ρ_{AB} . The quantum min-entropy of A conditioned on B is defined as

$$H_{\min}(A|B)_\rho = -\inf_{\sigma_B} \left\{ \inf \left\{ \lambda \in \mathbb{R} : \rho_{AB} \leq 2^\lambda I_A \otimes \sigma_B \right\} \right\}.$$

When ρ_{AB} is a cq-state (i.e., the A register is a classical state), the quantum min-entropy has a nice operational meaning in terms of guessing probability [KRS09]. Specifically, if $H_{\min}(A|B)_\rho = k$, then the optimal probability of predicting the value of A given ρ_B is exactly 2^{-k} .

In the following, we state the interactive leakage chain rule in [LC18]. Let $\rho = \rho_{AB}$ be a cq-state, that is, the system A is classical while B is quantum. The interactive leakage chain rule bounds how much the min-entropy $H_{\min}(A|B)_\rho$ can be decreased by an “interactive leakage” produced by applying a two-party protocol $\Pi = \{\mathcal{A}, \mathcal{B}, \rho_{\text{joint}}, s\}$ to ρ , where A is treated as a classical input to \mathcal{A} and B is given to \mathcal{B} as part of its initial state in ρ_{joint} .

Definition 4.1. Let $\rho = \rho_{AB}$ be a cq-state. Let $\Pi = \{\mathcal{A}, \mathcal{B}, \rho_{\text{joint}}, s\}$ be a two-party protocol where ρ_{joint} contains ρ_B in the \mathcal{B}_p system, and ρ_{in} be an input state where the classical state ρ_A is copied to \mathcal{A}_0 as the input for \mathcal{A} . (That is, \mathcal{A}_0 has an initial state $|0\rangle_{\mathcal{A}_0}$ and we do controlled NOT gates from ρ_A to $|0\rangle_{\mathcal{A}_0}$.) Consider the protocol execution $[\mathcal{A} \otimes \mathcal{B}](\rho_{\text{in}})$ and let σ_{AB_s} be the final state where A denotes the original classical state and B_s denotes the final state of \mathcal{B} . We say σ_{B_s} is an *interactive leakage* of A produced by Π .

Theorem 4.2. Let $\rho = \rho_{AB}$ be a cq-state. Let σ_{AB_s} be the final state of a two-party protocol $\Pi = \{\mathcal{A}, \mathcal{B}, \rho_{\text{joint}}, s\}$ with certain input state ρ_{in} . Let m_A and m_B be the communication complexity of \mathcal{A} and \mathcal{B} , respectively. We have

$$H_{\min}(A|B_s)_\sigma \geq H_{\min}(A|B)_\rho - \min\{m_A + m_B, 2m_A\}, \quad (22)$$

We will also use the following lemma about gentle measurement, which is first proved by Winter [Win99] and improved by Ogawa and Nagaoka [ON07], and is also referred to as the almost-as-good-as-new Lemma by Aaronson [Aar04]. It says that the post-measurement state of an almost-sure measurement will remain close to its original. The following version is taken from Wilde’s book [Wil13].

Lemma 4.3. Suppose $0 \leq \Lambda \leq I$ is a measurement operator such that for a mixed state ρ ,

$$\text{tr}(\Lambda\rho) \geq 1 - \epsilon.$$

Then the post-measurement state $\tilde{\rho}$ is $\sqrt{\epsilon}$ -close to the original state ρ :

$$\|\tilde{\rho} - \rho\|_{\text{tr}} \leq \sqrt{\epsilon}.$$

We will also need the following lemma, which can be proved by a standard argument using Uhlmann theorem and the Fuchs and van de Graaf inequality [FvdG99] (for a proof, see, e.g., [BB15]).

Lemma 4.4. Suppose $\rho_A, \sigma_A \in \mathcal{A}$ are two quantum states with purifications $|\phi\rangle_{AB}, |\psi\rangle_{AB} \in \mathcal{A} \otimes \mathcal{B}$, respectively, and $\|\rho_A - \sigma_A\|_{\text{tr}} \leq \epsilon$. Then there exists a unitary $U_B \in L(\mathcal{B})$ such that

$$\| |\phi\rangle_{AB} - I_A \otimes U_B |\psi\rangle_{AB} \|_{\text{tr}} \leq \sqrt{\epsilon(2 - \epsilon)}.$$

4.2 Our Lower Bound

Theorem 4.5. Let $\Pi = \{\mathcal{A}, \mathcal{B}, \rho_{\text{joint}} = |0\rangle\langle 0|, s\}$ be a QPIR protocol for the server's database of size n . Suppose Π is $(1 - \delta)$ -correct and ϵ -private against γ -specious servers with $\delta \leq n^{-4}/100, \epsilon \leq n^{-8}/100$. Then the server's communication complexity is at least $(n - 1)/2$ and the total communication complexity is at least $n - 1$.

In the above theorem, we consider protocols with no prior setup, i.e., $\rho_{\text{joint}} = |0\rangle\langle 0|$. We observe that the lower bound for the server's communication complexity extends for general ρ_{joint} , since one can think of ρ_{joint} as prepared by the client, who sends the server's initial state to the server at the beginning of the protocol. This simple reduction does not increase the server's communication complexity and extends the lower bound on the server's communication complexity for arbitrary ρ_{joint} .

Corollary 4.6. Let $\Pi = \{\mathcal{A}, \mathcal{B}, \rho_{\text{joint}}, s\}$ be a QPIR protocol for the server's database of size n with arbitrary ρ_{joint} . Suppose Π is $(1 - \delta)$ -correct and ϵ -private against γ -specious servers with $\delta \leq n^{-4}/100, \epsilon \leq n^{-8}/100$. Then the server's communication complexity is at least $(n - 1)/2$.

We now prove Theorem 4.5.

Proof. To establish communication complexity lower bound for Π , we consider a purified version $\bar{\Pi} = \{\bar{\mathcal{A}}, \bar{\mathcal{B}}, \rho_{\text{joint}}, s\}$ of Π , where both parties' operations are purified. Specifically, $\bar{\mathcal{A}}$ is modified from \mathcal{A} , where the sequence of quantum operations $\bar{\mathcal{A}}_1, \dots, \bar{\mathcal{A}}_s$ are unitaries

$$\begin{aligned}\bar{\mathcal{A}}_1 &: L(\mathcal{A}_0 \otimes \bar{\mathcal{A}}_0) \rightarrow L(\mathcal{A}_1 \otimes \bar{\mathcal{A}}_1 \otimes \mathcal{X}_1), \\ \bar{\mathcal{A}}_t &: L(\mathcal{A}_{t-1} \otimes \bar{\mathcal{A}}_{t-1} \otimes \mathcal{Y}_{t-1}) \rightarrow L(\mathcal{A}_t \otimes \bar{\mathcal{A}}_t \otimes \mathcal{X}_t), t = 2, \dots, s;\end{aligned}$$

$\bar{\mathcal{A}}_0$ is of sufficiently large dimension and initialized to $|0\rangle$; $\bar{\mathcal{A}}_t$ are called purifying spaces and

$$\text{tr}_{\bar{\mathcal{A}}_t}(\bar{\rho}_t(\rho_{\text{in}})) = \rho_t(\rho_{\text{in}})$$

for all $\rho \in \mathcal{A}_0 \otimes \mathcal{B}_0$. The purified $\bar{\mathcal{B}}$ for \mathcal{B} is similarly defined.

By inspection, it is easy to verify that $\bar{\Pi}$ preserves the properties of Π , i.e., $\bar{\Pi}$ is also $(1 - \delta)$ -correct, ϵ -private against γ -specious servers, and has the same communication complexity as Π . Thus, communication complexity lower bound for $\bar{\Pi}$ implies that for Π . Also, note that $\bar{\mathcal{A}}$ is a 0-specious adversary for $\bar{\Pi}$.

Now, let us consider an experiment that first samples a uniformly random database $\mathbf{a} \in \{0, 1\}^n$, and use \mathbf{a} as the server's database to run the protocol $\bar{\Pi}$ with an arbitrary fixed input of the client. Note that execution of the protocol can be viewed as producing an interactive leakage of \mathbf{a} . Let ρ_{AB} denote the final state where system A denotes the input \mathbf{a} and system B has the client's final local state. By Theorem 4.2, we have

$$H(A|B)_\rho \geq H(A)_\rho - \min\{2m_A, m_A + m_B\},$$

where m_A, m_B denote the server and the client's communication complexities, respectively. The operational meaning of min-entropy says that given the client's state ρ_B , one cannot guess the random database \mathbf{a} correctly with probability higher than $2^{-(H(A)_\rho - \min\{2m_A, m_A + m_B\})}$. To derive a lower bound on the communication complexity, we show a strategy to predict the database \mathbf{a} with probability at least $1 - n^2 \sqrt{\delta + 2\sqrt{\epsilon(1 - \epsilon)}} > 1/2$, which gives the desired lower bound.

Let $\sigma_B^i = \text{tr}_A[\bar{\mathcal{A}} \otimes \bar{\mathcal{B}}](|\mathbf{a}\rangle\langle\mathbf{a}|_{A_0} \otimes |i\rangle\langle i|_{B_0})$ and $\sigma_A^i = \text{tr}_B[\bar{\mathcal{A}} \otimes \bar{\mathcal{B}}](|\mathbf{a}\rangle\langle\mathbf{a}|_{A_0} \otimes |i\rangle\langle i|_{B_0})$.
By the definition of privacy, there exists a quantum operation \mathcal{F} such that

$$\Delta(\text{tr}_{B_0}\mathcal{F}_0 \otimes I_{\bar{B}_0}(\rho_{in}^1), \sigma_A^1) \leq \epsilon. \quad (23)$$

Since $\text{tr}_{B_0}\mathcal{F}_0 \otimes I_{\bar{B}_0}(\rho_{in}^1) = \text{tr}_{B_0}\mathcal{F}_0 \otimes I_{\bar{B}_0}(\rho_{in}^i)$ for all i ,

$$\Delta(\text{tr}_{B_0}\mathcal{F}_0 \otimes I_{\bar{B}_0}(\rho_{in}^1) - \sigma_A^i) \leq \epsilon \quad (24)$$

We have, by triangle inequality,

$$\Delta(\sigma_A^1 - \sigma_A^i) \leq 2\epsilon.$$

for all i .

By Lemma 4.4, we have

$$\Delta(I_A \otimes U_B^{1 \rightarrow i} |\psi^1\rangle_{\bar{A}\bar{B}}, |\psi^i\rangle_{\bar{A}\bar{B}}) \leq 2\sqrt{\epsilon(1-\epsilon)} \triangleq \epsilon', \quad (25)$$

where $|\phi\rangle_{\bar{A}\bar{B}}$ and $|\psi^i\rangle_{\bar{A}\bar{B}}$ are purifications of σ_A^1 and σ_A^i , respectively.

By the definition of correctness error, there exists measurement \mathcal{M}_i such that

$$\Pr\{\mathcal{M}_i(\sigma_B^i) = a_i\} \geq 1 - \delta.$$

Let

$$\mathcal{M}'_i = (U_B^{1 \rightarrow i})^\dagger \mathcal{M}_i U_B^{1 \rightarrow i}$$

for $i = 2, \dots, n$. Thus we have by Eq. (25)

$$\Pr\{\mathcal{M}'_i(\sigma_B^1) = a_i\} \geq 1 - \delta - \epsilon'. \quad (26)$$

By Lemma 4.3, the client can recover $\tilde{\sigma}_B^{(i)}$ such that

$$\Delta(\tilde{\sigma}_B^{(i)}, \sigma_B^1) \leq \sqrt{\delta + \epsilon'}. \quad (27)$$

Now we construct a protocol for the client to learn all the bits $\mathbf{a} = a_1, \dots, a_n$. First the client chooses input $|1\rangle\langle 1|$. Then he plays the protocol $\bar{\Pi}$ with Alice and obtains σ_B^1 . Measuring σ_B^1 by \mathcal{M}_1 , the client gets a_1 with probability at least $1 - \delta$. By Lemma 4.3, the client can recover $\tilde{\sigma}_B^1$ such that

$$\Delta(\tilde{\sigma}_B^1, \sigma_B^1) \leq \sqrt{\delta}.$$

Then the client measures \mathcal{M}'_2 on $\tilde{\sigma}_B^1$ and then recovers $\tilde{\sigma}_B^2$. Continue this process and $\tilde{\sigma}_B^k$ will be the state recovered from applying \mathcal{M}'_k to $\tilde{\sigma}_B^{k-1}$. We claim that

$$\Delta(\tilde{\sigma}_B^k, \sigma_B^1) \leq k\sqrt{\delta + \epsilon'}. \quad (28)$$

Suppose this is true for $i = 2, \dots, k$. If we measure \mathcal{M}'_{k+1} on $\tilde{\sigma}_B^{k+1}$ and on σ_B^1 , respectively, and recover $\tilde{\sigma}_B^{k+1}$ and $\tilde{\sigma}_B^{(k+1)}$, respectively, we have

$$\Delta(\tilde{\sigma}_B^{k+1}, \tilde{\sigma}_B^{(k+1)}) \leq \Delta(\tilde{\sigma}_B^k, \sigma_B^1) \leq k\sqrt{\delta + \epsilon'} \quad (29)$$

where the first inequality is because quantum operations do not increase trace distance. Now use the triangle inequality with Eqs. (27) and (29), and the claim follows by induction.

By Eqs. (26) and (28), the probability of recovering a_i by measuring \mathcal{M}'_i on $\tilde{\sigma}_B^{i-1}$ is at least $1 - i\sqrt{\delta + \epsilon'}$. Therefore, the client learns \mathbf{a} with probability at least

$$\prod_{i=1}^n \left(1 - i\sqrt{\delta + \epsilon'}\right) \geq 1 - n^2\sqrt{\delta + \epsilon'},$$

which is what we need to complete the proof. \square

Acknowledgments We thank the anonymous referees for presenting us with the work of Kerenidis et al. [KLGR16], and other valuable comments. ZB is supported by the Israel Science Foundation (Grant 468/14), Binational Science Foundation (Grants 2016726, 2014276), and by the European Union Horizon 2020 Research and Innovation Program via ERC Project REACT (Grant 756482) and via Project PROMETHEUS (Grant 780701). OS is supported by ERC Grant 280157, by the Israel Science Foundation (Grant 682/18), and by the Cyber Security Research Center at Ben-Gurion University. CYL is financially supported from the Young Scholar Fellowship Program by Ministry of Science and Technology (MOST) in Taiwan, under Grant MOST107-2636-E-009-005. KMC is partially supported by 2016 Academia Sinica Career Development Award under Grant No. 23-17 and the Ministry of Science and Technology, Taiwan under Grant No. MOST 103-2221-E-001-022-MY3. DA and AG were supported by ERC Grant 280157 for part of the work on this project, and are supported by the Israel Science Foundation (Grant 1721/17).

References

- [Aar04] S. Aaronson. **Limitations of quantum advice and one-way communication**. In *Proceedings. 19th IEEE Annual Conference on Computational Complexity, 2004.*, pages 320–332, June 2004.
- [ACG⁺16] D. Aharonov, A. Chailloux, M. Ganz, I. Kerenidis, and L. Magnin. **A Simpler Proof of the Existence of Quantum Weak Coin Flipping with Arbitrarily Small Bias**. *SIAM J. Comput.*, 45(3):633–679, 2016.
- [AKN98] D. Aharonov, A. Y. Kitaev, and N. Nisan. **Quantum Circuits with Mixed States**. In *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 20–30, 1998.
- [ANTSV02] A. Ambainis, A. Nayak, A. Ta-Shma, and U. Vazirani. **Dense Quantum Coding and Quantum Finite Automata**. *JACM*, 49(4):496–511, July 2002.
- [BB84] C. H. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing*, page 175, 1984.
- [BB15] Ä. Baumeler and A. Broadbent. **Quantum Private Information Retrieval has Linear Communication Complexity**. *J. Cryptology*, 28(1):161–175, 2015.

- [BS16] A. Broadbent and C. Schaffner. **Quantum cryptography beyond quantum key distribution**. *Des. Codes Cryptography*, 78(1):351–382, 2016.
- [BV11] Z. Brakerski and V. Vaikuntanathan. Efficient Fully Homomorphic Encryption from (Standard) LWE. In R. Ostrovsky, editor, *FOCS*, pages 97–106. IEEE, 2011. Full version in <https://eprint.iacr.org/2011/344.pdf>.
- [CGKS95] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan. **Private Information Retrieval**. In *36th Annual Symposium on Foundations of Computer Science, Milwaukee, Wisconsin, USA, 23-25 October 1995*, pages 41–50. IEEE Computer Society, 1995.
- [CK09] A. Chailloux and I. Kerenidis. **Optimal Quantum Strong Coin Flipping**. In *50th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2009, October 25-27, 2009, Atlanta, Georgia, USA*, pages 527–533. IEEE Computer Society, 2009.
- [CMS99] C. Cachin, S. Micali, and M. Stadler. **Computationally Private Information Retrieval with Polylogarithmic Communication**. In *EUROCRYPT*, pages 402–414, 1999.
- [DG15] Z. Dvir and S. Gopi. **2-Server PIR with Sub-Polynomial Communication**. In R. A. Servedio and R. Rubinfeld, editors, *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 577–584. ACM, 2015.
- [DNS10] F. Dupuis, J. B. Nielsen, and L. Salvail. **Secure two-party quantum evaluation of unitaries against specious adversaries**. In *Annual Cryptology Conference*, pages 685–706. Springer, 2010.
- [Efr12] K. Efremenko. **3-Query Locally Decodable Codes of Subexponential Length**. *SIAM J. Comput.*, 41(6):1694–1703, 2012.
- [FvdG99] C. A. Fuchs and J. van de Graaf. **Cryptographic distinguishability measures for quantum-mechanical states**. *IEEE Trans. Information Theory*, 45(4):1216–1227, May 1999.
- [GC01] D. Gottesman and I. Chuang. Quantum Digital Signatures, 2001, arXiv: [quant-ph/0105032](https://arxiv.org/abs/quant-ph/0105032).
- [Gen09] C. Gentry. *A fully homomorphic encryption scheme*. PhD thesis, Stanford University, 2009.
- [GLM08] V. Giovannetti, S. Lloyd, and L. Maccone. **Quantum Private Queries**. *Phys. Rev. Lett.*, 100:230502, Jun 2008.
- [Gol04] O. Goldreich. *The Foundations of Cryptography - Volume 2, Basic Applications*. Cambridge University Press, 2004.
- [GW07] G. Gutoski and J. Watrous. **Toward a general theory of quantum games**. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 565–574. ACM, 2007.

- [JP99] D. Jonathan and M. B. Plenio. **Entanglement-Assisted Local Manipulation of Pure Quantum States**. *Phys. Rev. Lett.*, 83:3566–3569, Oct 1999.
- [JRS09] R. Jain, J. Radhakrishnan, and P. Sen. **A property of quantum relative entropy with an application to privacy in quantum communication**. *J. ACM*, 56(6):33:1–33:32, 2009.
- [KLGR16] I. Kerenidis, M. Laurière, F. L. Gall, and M. Rennela. **Information cost of quantum communication protocols**. *Quantum Information & Computation*, 16(3&4):181–196, 2016.
- [Kli07] M. Klimesh. Inequalities that Collectively Completely Characterize the Catalytic Majorization Relation, 2007, arXiv: **0709.3680**.
- [KRS09] R. König, R. Renner, and C. Schaffner. **The Operational Meaning of Min- and Max-Entropy**. *IEEE Trans. Inf. Theory*, 55(9):4337–4347, Sept 2009.
- [LC97] H.-K. Lo and H. F. Chau. **Is Quantum Bit Commitment Really Possible?** *Phys. Rev. Lett.*, 78:3410–3413, Apr 1997.
- [LC18] C.-Y. Lai and K.-M. Chung. Interactive Leakage Chain Rule for Quantum Min-entropy, 2018, arXiv: **1809.10694**.
- [LG12] F. Le Gall. **Quantum Private Information Retrieval with Sublinear Communication Complexity**. *Theory of Computing*, 8(16):369–374, 2012.
- [Lo97] H.-K. Lo. **Insecurity of quantum secure computations**. *Physical Review A*, 56(2):1154, 1997.
- [May97] D. Mayers. **Unconditionally Secure Quantum Bit Commitment is Impossible**. *Phys. Rev. Lett.*, 78:3414–3417, Apr 1997.
- [Moc07] C. Mochon. Quantum weak coin flipping with arbitrarily small bias, 2007, arXiv: **0711.4114**.
- [Nay99] A. Nayak. **Optimal lower bounds for quantum automata and random access codes**. In *Foundations of Computer Science, 1999. 40th Annual Symposium on*, pages 369–376, 1999.
- [ON07] T. Ogawa and H. Nagaoka. **Making Good Codes for Classical-Quantum Channel Coding via Quantum Hypothesis Testing**. *IEEE Trans. Information Theory*, 53(6):2261–2266, June 2007.
- [vDH03] W. van Dam and P. Hayden. **Universal entanglement transformations without communication**. *Phys. Rev. A*, 67:060302, Jun 2003.
- [Wie83] S. Wiesner. **Conjugate Coding**. *SIGACT News*, 15(1):78–88, January 1983.
- [Wil13] M. M. Wilde. *Quantum Information Theory*. Cambridge University Press, 2013. Cambridge Books Online.
- [Win99] A. J. Winter. **Coding theorem and strong converse for quantum channels**. *IEEE Trans. Information Theory*, 45(7):2481–2485, 1999.

[YPF14] L. Yu, C. A. Pérez-Delgado, and J. F. Fitzsimons. [Limitations on information theoretically secure quantum homomorphic encryption](#), 2014, arXiv: [1406.2456](#).

A Hilbert Spaces and Quantum States

The Hilbert space of a quantum system A is denoted by the corresponding calligraphic letter \mathcal{A} and its dimension is denoted by $\dim(\mathcal{A})$. Let $L(\mathcal{A})$ be the space of linear operators on \mathcal{A} . A quantum state of system A is described by a *density operator* $\rho_A \in L(\mathcal{A})$ that is positive semidefinite and with unit trace ($\text{tr}(\rho_A) = 1$). Let $S(\mathcal{A}) = \{\rho_A \in L(\mathcal{A}) : \rho_A \geq 0, \text{tr}(\rho_A) = 1\}$ be the set of density operators on \mathcal{A} . When $\rho_A \in S(\mathcal{A})$ is of rank one, it is called a *pure* quantum state and we can write $\rho = |\psi\rangle\langle\psi|_A$ for some unit vector $|\psi\rangle_A \in \mathcal{A}$, where $\langle\psi| = |\psi\rangle^\dagger$ is the conjugate transpose of $|\psi\rangle$. If ρ_A is not pure, it is called a *mixed* state and can be expressed as a convex combination of pure quantum states.

The Hilbert space of a joint quantum system AB is the tensor product of the corresponding Hilbert spaces $\mathcal{A} \otimes \mathcal{B}$. For $\rho_{AB} \in S(\mathcal{A} \otimes \mathcal{B})$, its reduced density operator in system A is $\rho_A = \text{tr}_B(\rho_{AB})$, where

$$\text{tr}_B(\rho_{AB}) = \sum_i I_A \otimes \langle i|_B (\rho_{AB}) I_A \otimes |i\rangle_B$$

for an orthonormal basis $\{|i\rangle_B\}$ for \mathcal{B} . We sometimes use the equivalent notation,

$$\rho_{AB|A} := \text{tr}_B(\rho_{AB}).$$

Suppose $\rho_A \in S(\mathcal{A})$ of finite dimension $\dim(\mathcal{A})$. Then there exists \mathcal{B} of dimension $\dim(\mathcal{B}) \geq \dim(\mathcal{A})$ and $|\psi\rangle_{AB} \in \mathcal{A} \otimes \mathcal{B}$ such that

$$\text{tr}_B |\psi\rangle\langle\psi|_{AB} = \rho_A.$$

The state $|\psi\rangle_{AB}$ is called a purification of ρ_A .

The trace distance between two quantum states ρ and σ is

$$\Delta(\rho, \sigma) = \|\rho - \sigma\|_{\text{tr}},$$

where $\|X\|_{\text{tr}} = \frac{1}{2}\text{tr}\sqrt{X^\dagger X}$ is the trace norm of X . Hence the trace distance between two pure states $|\alpha\rangle, |\beta\rangle$ is

$$\Delta(|\alpha\rangle\langle\alpha|, |\beta\rangle\langle\beta|) = \sqrt{1 - |\langle\alpha|\beta\rangle|^2}. \quad (30)$$

Lemma A.1. Consider a quantum state ρ_{XY} over two registers X, Y , and denote $\rho_X = \text{tr}_Y(\rho_{XY})$. Then if there exists $\epsilon, |\varphi\rangle$ s.t. $\Delta(\rho_X, |\varphi\rangle\langle\varphi|) \leq \epsilon$, then there exists $\tilde{\rho}_Y$ s.t. $\Delta(\rho_{XY}, |\varphi\rangle\langle\varphi| \otimes \tilde{\rho}_Y) \leq \sqrt{\epsilon}$. Furthermore, if ρ_{XY} is pure then so is $\tilde{\rho}_Y$.

Proof. It is sufficient w.l.o.g to prove for a pure ρ_{XY} , since it is always possible to purify ρ_{XY} by adding an additional register Z , and consider the pure state ρ_{XYZ} . The transitivity of the partial trace operation implies that if the theorem is true for $X, (YZ)$, then it is also true for X, Y . Also assume w.l.o.g that $|\varphi\rangle = |0\rangle$ (this is just a matter of choosing a basis elements).

Thus we will provide a proof in the case where the joint state of X, Y can be written as a superposition $|\alpha\rangle = \sum_{x,y} w_{x,y} |x\rangle|y\rangle$. Define $P_0 = \Pr[X = 0] = \sum_y |w_{0,y}|^2$, and note that it must

be the case that $P_0 \geq 1 - \epsilon$. To see this, note that P_0 is the probability of measuring $X = 0$ in the experiment where we first trace out Y and then measuring X . Since $\Delta(\rho_X, |0\rangle\langle 0|) \leq \epsilon$, the probability of measuring $X = 0$ after tracing out Y is ϵ close to the probability of measuring $X = 0$ in $|0\rangle\langle 0|$, which is 1 (see, e.g., [AKN98]). The claim $P_0 \geq 1 - \epsilon$ follows.

Now define $|\beta\rangle = \frac{1}{\sqrt{P_0}} \sum_y w_{0,y} |y\rangle$, and let $\tilde{\rho}_Y = |\beta\rangle\langle\beta|$. Then

$$\Delta(\rho_{XY}, |0\rangle\langle 0| \otimes \tilde{\rho}_Y) = \Delta(|\alpha\rangle\langle\alpha|, |0\rangle\langle 0| \otimes |\beta\rangle\langle\beta|) = \sqrt{1 - |\langle\alpha|(0, \beta)\rangle|^2}. \quad (31)$$

We have

$$\langle\alpha|(0, \beta)\rangle = \frac{1}{\sqrt{P_0}} \sum_y |w_{0,y}|^2 = \sqrt{P_0}, \quad (32)$$

which implies that indeed $\Delta(\rho_{XY}, |0\rangle\langle 0| \otimes \tilde{\rho}_Y) = \sqrt{1 - P_0} \leq \sqrt{\epsilon}$. \square

B Security Analysis of Kereneidis et al.'s Protocol

For completeness, we restate⁷ the QPIR protocol with pre-shared entanglement by Kereneidis et al. [KLGR16, Section 6]. Given a database $\text{DB} \in \{0, 1\}^n$ for some $n = 2^\ell$ as input to the server, and index $i \in [n]$ as input to the client (If the client's input is a superposition, the algorithm is run in superposition), we denote the protocol Π_n as follows.

The protocol Π_n is recursive and calls $\Pi_{n/2}$ as a subroutine. For the execution of Π_n , the parties are required to pre-share a pair of entangled state registers $\frac{1}{2^{n/4}} \sum_{\mathbf{r} \in \{0,1\}^{n/2}} |\mathbf{r}\rangle_R |\mathbf{r}\rangle_{R'}$, where R is held by the server and R' is held by the client. They also share an entangled state needed for the recursive application of the protocol $\Pi_{n/2}$ (and the recursive calls it entails). Unfolding the recursion, this means that for all $n' = 2^{\ell'}$ with $\ell' \in [\ell - 1]$, there is an entangled register of length n' shared between the client and the server.

The protocol execution is described in shorthand Figure 1. In what follows we provide a detailed description and analyze the steps of the protocol to establish correctness and assert properties that will allow us to analyze privacy.

1. If $n = 1$ then the database contains a single value. In this case there is no need for shared entanglement, and the server sends a register F containing $|\text{DB}\rangle$ (the final response) to the client, and the protocol terminates. This is trivially secure and efficient. Otherwise proceed to the next steps.
2. The server denotes $\text{DB}_0, \text{DB}_1 \in \{0, 1\}^{n/2}$ s.t. $\text{DB} = [\text{DB}_0 || \text{DB}_1]$, i.e. the low-order and high-order bits of the database respectively. The server starts with two single-bit registers Q_0, Q_1 initialized to 0. The server CNOTs Q_b with the inner product of R and DB_b so that it contains $|\mathbf{r} \cdot \text{DB}_b\rangle_{Q_b}$, and sends Q_0, Q_1 to the client.

At this point, the joint state between the client and (an honest) server is

$$\sum_{\mathbf{r} \in \{0,1\}^{n/2}} |\mathbf{r}\rangle_R |\mathbf{r}\rangle_{R'} |\mathbf{r} \cdot \text{DB}_0\rangle_{Q_0} |\mathbf{r} \cdot \text{DB}_1\rangle_{Q_1}.$$

In particular the reduced density matrix of the server's state is independent of the index i .

⁷We make one minor adaptation – see Remark 1.

- Let $b^* = \lfloor \frac{i-1}{n} \rfloor$ denote the most significant bit of i . The client evaluates a Z gate on Q_{b^*} . It sends Q_0, Q_1 back to the server.

At this point, the joint state between the client and (an honest) server is

$$\sum_{\mathbf{r} \in \{0,1\}^{n/2}} (-1)^{\mathbf{r} \cdot \mathbf{DB}_{b^*}} |\mathbf{r}\rangle_R |\mathbf{r}\rangle_{R'} |\mathbf{r} \cdot \mathbf{DB}_0\rangle_{Q_0} |\mathbf{r} \cdot \mathbf{DB}_1\rangle_{Q_1} .$$

Importantly, the reduced density matrix of the server, which contains the registers R, Q_0, Q_1 , is the diagonal matrix that corresponds to the classical distribution of sampling a random \mathbf{r} in register R , and placing $\mathbf{r} \cdot \mathbf{DB}_0, \mathbf{r} \cdot \mathbf{DB}_1$ in Q_0, Q_1 . This density matrix is independent of b^* and therefore of i .

- The server again CNOTs Q_b with the inner product of R and \mathbf{DB}_b .

At this point, the joint state between the client and (an honest) server is

$$\sum_{\mathbf{r} \in \{0,1\}^{n/2}} (-1)^{\mathbf{r} \cdot \mathbf{DB}_{b^*}} |\mathbf{r}\rangle_R |\mathbf{r}\rangle_{R'} |0\rangle_{Q_0} |0\rangle_{Q_1} .$$

From this point on we disregard Q_0, Q_1 since they remain zero throughout. Since this step only involves a local unitary by the server, we are guaranteed that its reduced density matrix is still independent of i .

- The server performs QFT on R and the client performs QFT on R' . The resulting state is

$$\frac{1}{2^{3n/4}} \sum_{\mathbf{r}, \mathbf{y}, \mathbf{w} \in \{0,1\}^{n/2}} (-1)^{\mathbf{r} \cdot (\mathbf{DB}_{b^*} \oplus \mathbf{y} \oplus \mathbf{w})} |\mathbf{y}\rangle_R |\mathbf{w}\rangle_{R'} = \frac{1}{2^{n/4}} \sum_{\mathbf{y} \in \{0,1\}^{n/2}} |\mathbf{y}\rangle_R \underbrace{|\mathbf{y} \oplus \mathbf{DB}_{b^*}\rangle}_{\mathbf{w}}_{R'} .$$

Since we only performed local operations on the server and client side (without communication), the server's density matrix remains perfectly independent of b^* and thus of i .

- Note that at this point, the joint state of the client and server is a “shifted” entangled state where the shift corresponds to the half-database \mathbf{DB}_{b^*} that contains the element that the client wishes to retrieve. More explicitly, $\mathbf{DB}[i] = \mathbf{DB}_{b^*}[i^*]$ for $i^* = i \pmod{n/2}$ contains the $(\ell - 1)$ least significant bits of i . Therefore, for all \mathbf{y}, \mathbf{w} in the support of the joint state, it holds that $\mathbf{DB}[i] = \mathbf{w}[i^*] \oplus \mathbf{y}[i^*]$.

The client will now ignore (temporarily) the register R' and execute $\Pi_{n/2}$ recursively on index i^* . The (honest) server will carry out the protocol with the value \mathbf{y} from the register R serving as the server's database. Note that since the register R' is not touched, for the purposes of executing the protocol the value \mathbf{w} in R' is equivalent to have been measured, and the value \mathbf{y} in R is equivalent to the deterministic register $\mathbf{w} \oplus \mathbf{DB}_{b^*}$.

We are recursively guaranteed that in the end of the execution of $\Pi_{n/2}$, the client receives a register F containing the value $\mathbf{y}[i^*] = \mathbf{w}[i^*] \oplus \mathbf{DB}_{b^*}[i^*] = \mathbf{w}[i^*] \oplus \mathbf{DB}[i]$. Since the client still maintains the original register R' containing \mathbf{w} , it can CNOT the value $\mathbf{w}[i^*]$ from F and obtain $|\mathbf{DB}[i]\rangle_A$. Namely, in the end of the execution, the register F indeed contains the desired value $\mathbf{DB}[i]$.

7. Lastly, if the client and server desire to “clean up” and restore the shared entanglement so that it can be reused in consequent executions, the client can copy the contents of the register F to a fresh register (which is possible since this register contains a classical value). Since the client and server are pure (i.e. do not measure) throughout the protocol, they can rewind the execution of the protocol to restore their initial joint entanglement.

If the final cleanup step is not executed then the total number of rounds of Π_n is $2\ell + 1$, and the total communication complexity is $4\ell + 1$ (recall that $\ell = \log(n)$). If the cleanup step is executed, the round complexity and communication complexity are doubled due to rewinding the execution.

Remark 1. Note that in the original protocol by Kerenidis et al. step 7 does not appear, and it is not mentioned that the shared entanglement can be cleaned and reused.

We conclude that for classical inputs for both the client and the server, the honest server’s density matrix is independent of i . If the client first measures its input state, privacy holds. But since the very first operation is a CNOT operation (to determine the value of i) and since CNOT and measurements in the standard basis commute, we conclude that the server’s reduced density matrix is independent of the client’s input, even for inputs which are in superposition.

Lemma B.1. The protocol Π_n is a PIR protocol with perfect correctness and perfect anchored privacy against honest servers. It furthermore has communication complexity $O(\log n)$, and uses $O(n)$ bits of (reusable) shared entanglement.

Proof. The analysis in the body of the protocol establishes that the local view of the adversary is independent of i , when i is treated as a fixed parameter. For the sake of our privacy notion, we are required to establish that the server’s local state is independent of i even when i is an arbitrary quantum state. This follows since the client refers to the index i as constant, and therefore a superposition over i will translate to a superposition over classical executions of the protocol, each with a fixed i . Since the server’s local view is identical for any fixed i , it will also be in the same state for a superposition, and also for an arbitrary mixed state of i and some potential environment.

The communication complexity and the amount of reusable shared entanglement needed in this protocol follow directly from the protocol. \square

We can therefore apply Theorem 3.2 and conclude that Π is secure against anchored-specious adversaries.

Corollary B.2. There exists a PIR protocol Π with logarithmic communication complexity assuming linear shared entanglement, which is perfectly correct and anchored $O(\sqrt{\gamma})$ -private against γ -specious adversaries.

Recursive QPIR with Logarithmic Communication

Server input: Database $\text{DB} \in \{0, 1\}^n$.

Client input: Index $i \in [n]$.

Desired output: Value $\text{DB}[i]$ stored in register F on the client side.

Setup: Register R for server and R' for client in joint state $\frac{1}{2^{n/4}} \sum_{\mathbf{r} \in \{0,1\}^{n/2}} |\mathbf{r}\rangle_R |\mathbf{r}\rangle_{R'}$.

(This setup is for external recursion loop, internal loops require their own R, R' defined recursively.)

1. If $n = 1$, copy the (single-bit) database into a register and send to client, then terminate (or go to clean up step 7 below).
2. The server denotes $\text{DB}_0, \text{DB}_1 \in \{0, 1\}^{n/2}$ s.t. $\text{DB} = [\text{DB}_0 \parallel \text{DB}_1]$, i.e. the low-order and high-order bits of the database respectively. The server starts with two single-bit registers Q_0, Q_1 initialized to 0. The server CNOTs Q_b with the inner product of R and DB_b so that it contains $|\mathbf{r} \cdot \text{DB}_b\rangle_{Q_b}$. It sends Q_0, Q_1 to the client.
3. Let $b^* = \lfloor \frac{i-1}{n} \rfloor$ denote the most significant bit of i . The client evaluates a Z gate on Q_{b^*} . It sends Q_0, Q_1 back to the server.
4. The server again CNOTs Q_b with the inner product of R and DB_b .
5. The server performs QFT on R and the client performs QFT on R' .
6. Call $\Pi_{n/2}$ recursively (with fresh R, R' obtained from the setup). The server input is the contents of the register R (of length $n/2$). The client input is $i^* = i \pmod{n/2} \in [n/2]$. The client receives a response register F as the output of the recursive call. It then CNOTs $R'[i^*]$ into F . Finally, F contains the output of the recursive execution.
7. If it is desired to restore the shared entanglement, copy the (classical) output into a fresh register and rewind the execution of the protocol.

Figure 1: The QPIR protocol of Kerenidis et al.