

Quantum Symmetrically-Private Information Retrieval

Iordanis Kerenidis*
UC Berkeley
jkeren@cs.berkeley.edu

Ronald de Wolf†
CWI Amsterdam
rdewolf@cwi.nl

Abstract

Private information retrieval systems (PIRs) allow a user to extract an item from a database that is replicated over $k \geq 1$ servers, while satisfying various privacy constraints. We exhibit quantum k -server symmetrically-private information retrieval systems (QSPIRs) that use sub-linear communication, do not use shared randomness among the servers, and preserve privacy against honest users and dishonest servers. Classically, SPIRs without shared randomness do not exist at all.

Keywords: Private information retrieval. User privacy. Data privacy. Quantum computing.

1 Introduction

1.1 Setting

The *Private Information Retrieval* problem was introduced by Chor et al. [4]. A user wants to learn a bit x_i from an n -bit database $x = x_1 \dots x_n$, for some $i \in [n]$ of his choice. The privacy of the user requires that the database server learns nothing about i , in the information-theoretic sense, and general efficiency considerations require the communication between the user and the database to be limited. Clearly, PIR can be realized by making the server send the whole database to the user. This takes n bits of communication and can be shown to be optimal. Better protocols exist if the database is *replicated* among some $k \geq 2$ different servers, who cannot communicate [4, 1, 2]. Here we require that *individual* servers learn nothing about i . For $k = 2$, the best known scheme uses $O(n^{1/3})$ bits of communication [4], and asymptotically the best known k -server uses $n^{O(\log \log(k)/k \log(k))}$ bits [2]. For $k \geq 2$, no good lower bounds on the required communication are known for this setting.

In a recent paper, we showed how to obtain *quantum* PIR systems (QPIR, where the parties are quantum computers and the communication consists of qubits) that use slightly less communication than the best known classical schemes [9]. In Table 1 we list the best known bounds on the communication complexity for small numbers of servers, in the classical as well as quantum case.

In its standard form, PIR just protects the privacy of the *user*: the individual servers learn nothing about i . But now suppose we also want to protect the privacy of the *data*. That is, we don't want the user to learn anything about x beyond the x_i that he asks for. For example, because the user should pay a fee for every x_i that he learns (pay-per-view), or because the database contains

*Supported by DARPA under agreement number F 30602-10-2-0524. Part of this work was done while visiting CWI.

†Most of this work was done while a postdoc at UC Berkeley, supported by Talent grant S 62-565 from the Netherlands Organization for Scientific Research (NWO). Also partially funded by projects QAIP (IST-1999-11234) and RESQ (IST-2001-37559) of the IST-FET programme of the EC.

Servers	PIR complexity	QPIR complexity
$k = 1$	$\Theta(n)$	$\Theta(n)$
$k = 2$	$O(n^{1/3})$	$O(n^{3/10})$
$k = 3$	$O(n^{1/5.25})$	$O(n^{1/7})$
$k = 4$	$O(n^{1/7.87})$	$O(n^{1/11})$

Table 1: Best known bounds on the communication complexity of classical and quantum PIR

very sensitive information. This setting of *Symmetrically-Private* Information Retrieval (SPIR) was introduced by Gertner et al. [7] and is closely related to *oblivious transfer*. They showed that SPIR is *impossible* even if the user is honest—i.e., follows the protocol—and the servers can individually flip coins [7, Appendix A]. This no-go result holds no matter how many servers and how many bits and rounds of communication we allow. Therefore they extended the PIR model by allowing the servers to *share* a random string that is hidden from the user, and showed how to turn any PIR scheme into a SPIR scheme with shared randomness among the servers, at a small extra communication cost. The resulting schemes are information-theoretically secure even against *dishonest* users, and use a number of random bits that is of the same order as the communication.

The necessity of shared randomness for classical SPIR schemes is a significant drawback, since information-theoretic security requires *new* shared randomness for each application of the scheme. This either requires a lot of extra communication between the servers (if new shared randomness is generated for each new application) or much memory on the parts of the servers (if randomness is generated once for many applications, each server needs to store this).

1.2 Results

In this paper, we study the existence and efficiency of SPIR schemes in the *quantum* world, where user and servers have quantum computers and can communicate qubits. Here user privacy means that the states of individual servers should all be independent of i , and data privacy means that the concatenation of the various states of the user is independent of the bits x_j for all $j \neq i$. We can distinguish between honest-user and dishonest-user data privacy. In the first case, data privacy holds if the user is honest (follows the protocol). In the second case, data privacy should hold even if the user deviates from the protocol in any way.

Our main result is that honest-user quantum SPIR schemes exist even in the case where the servers do not share any randomness. As mentioned above, such honest-user SPIRs without shared randomness are impossible in the classical world. This gives another example of a cryptographic task that can be performed with information-theoretic security in the quantum world but that is impossible classically (key distribution [3] is the main example of this). The communication complexity of our k -server QSPIR schemes is of the same order as that of the best known classical k -server PIR schemes. At first sight, one might think this trivial: just take a classical scheme, ensure data privacy using shared randomness among the servers, and then get rid of the shared randomness by letting the user entangle the messages to the servers. However, this would violate data privacy, as the user would now have “access” to the servers’ shared randomness. In actuality we do something quite different, making use of the fact that the servers can add phases that multiply out to an overall phase. This phase allows the user to extract x_i , but nothing else.

The notion of an honest user is somewhat delicate, because clearly users cannot be trusted to follow the protocol in all cases. Still, there are scenarios where the assumption of a honest user is not unreasonable, for example in pay-per-view systems where the user accesses the system via some

box, attached to his TV, that is sealed or otherwise protected from tampering. In this case the user cannot deviate from the protocol, but he can still be curious, trying to observe what goes on inside of his box to try to extract more information about the database. Our honest-user QSPIRs are perfectly secure against such users.

It would be nice to have SPIR schemes that are secure even against *dishonest* users. However, we exhibit a large class of PIR schemes (quantum as well as classical) that can all be cheated by a dishonest quantum user. Our honest-user QSPIRs fall in this class and hence are not secure against dishonest users. Fortunately, if we are willing to allow shared randomness between the servers then the best classical SPIRs can easily be made secure against even dishonest quantum users: if the servers measure the communication in the computational basis, the scheme is equivalent to the classical scheme, even if the user is quantum.

Remarks:

(1) Often the PIR setting is generalized to *t-secure* PIR, where no colluding set of t servers together have any information about i . We focus on the $t = 1$ case here in order to simplify the presentation.

(2) Very efficient PIR and SPIR schemes are possible under *computational* assumptions, even for $k = 1$ servers (see e.g. the references at [10]). In this paper we focus on information-theoretic security.

2 Definitions

We assume familiarity with the quantum model [12]. The various variants of PIR are defined below.

Definition 1 *A one-round, k -server private information retrieval (PIR) scheme with recovery probability $1/2 + \varepsilon$, query size t , and answer size a , consists of a randomized algorithm (the user), and k randomized algorithms S_1, \dots, S_k (the servers), such that*

1. *On input $i \in [n]$, the user produces k t -bit queries q_1, \dots, q_k and sends these to the respective servers. The j th server sends back an a -bit string a_j . The user outputs a bit b depending on i, a_1, \dots, a_k , and his randomness.*
2. **Recovery:** *For all x and i , the probability (over the user's and servers' randomness) that $b = x_i$ is at least $1/2 + \varepsilon$.*
3. **User privacy:** *For all j , the distribution of q_j (over the user's randomness) is independent of i .*

The communication complexity of the scheme is the total length of the communication between the user and the servers, i.e. $k(t + a)$ bits.

All best known PIR schemes satisfy the above definitions with $\varepsilon = 1/2$ (i.e., no error probability), and we will hereafter take $\varepsilon = 1/2$ unless mentioned otherwise. It is open whether multiple-round schemes can be better than the one-round variety we defined here. For PIR schemes, we can assume without loss of generality that the servers are deterministic. However, we need randomized servers for the symmetrically-private variety:

Definition 2 *A symmetrically-private information retrieval (SPIR) scheme is a PIR scheme with the additional property of **data privacy**: the user's "view" (i.e. the concatenation of his various*

states during the protocol) does not depend on x_j , for all $j \neq i$. We distinguish between private-randomness and shared-randomness SPIR schemes, depending on whether the servers individually flip coins or have a shared random coin (hidden from the user). We also distinguish between honest-user and dishonest-user SPIR, depending on whether data privacy should hold even when the user deviates from the protocol.

Definition 3 We define quantum versions QPIR and QSPIR of PIR and SPIR, respectively, in the obvious way: the user and the servers are quantum computers, and the communication uses quantum bits; user privacy means that the density matrix of each server is independent of i at all points in the protocol; data privacy means that the concatenation of the density matrices that the user has at the various points of the protocol, is independent of x_j , for all $j \neq i$. For QSPIR, we still have the distinctions of private/public-randomness and honest/dishonest-user.

As mentioned in the introduction, Gertner et al. [7, Appendix A] exhibited quite efficient shared-randomness SPIR schemes. One might think that these can be turned into SPIR schemes with deterministic servers as follows: the user picks a random string, sends it to each of the servers (along with the queries) to establish shared randomness between them, and then erases (or “forgets”) his copy of the random string. However, this erasing of the random string by the user is ruled out by the definition, since the user’s view includes the random string he drew. In fact, Gertner et al. [7, Appendix A] showed that shared randomness between the servers is necessary for the existence of classical SPIR (even for multi-round protocols):

Fact 1 For every $k \geq 1$, there is no k -server private-randomness SPIR scheme.

Intuitively, the reason is that since the servers have no knowledge of i (by user privacy), their individual messages need to be independent of *all* bits of x , including x_i , to ensure data privacy. But since they cannot coordinate via shared randomness, their *joint* messages will be independent of the whole x as well, so the user cannot learn x_i .

Below we show that this negative result does not apply to the quantum world: using coordination via quantum entanglement, we can get honest-user QSPIRs without any communication or shared randomness between the servers at any stage of the protocol.

3 Honest-user quantum SPIR from PIR schemes

Our honest-user QSPIR schemes work on top of the PIR schemes recently developed by Beimel et al. [2]. These, as well as all others known, work as follows: the user picks a random string r , and depending on i and r , picks k queries $q_1, \dots, q_k \in \{0, 1\}^t$. He sends these to the respective servers, who respond with answers $a_1, \dots, a_k \in \{0, 1\}^a$. The user then outputs

$$\sum_{j=1}^k a_j \cdot b_j = x_i,$$

where $b_1, \dots, b_k \in \{0, 1\}^a$ are determined by i and r , and everything is modulo 2.

We will now describe the quantum SPIR scheme. As before, the user picks r, q_1, \dots, q_k . In addition, he picks k random strings $r_1, \dots, r_k \in \{0, 1\}^a$. He defines $r'_j = r_j + b_j$ and sets up the following $(k + 1)$ -register state

$$\frac{1}{\sqrt{2}}|0\rangle|q_1, r_1\rangle \cdots |q_k, r_k\rangle + \frac{1}{\sqrt{2}}|1\rangle|q_1, r'_1\rangle \cdots |q_k, r'_k\rangle.$$

The user keeps the first 1-qubit register to himself, and sends the other k registers to the respective servers. The j th server sees a random mixture of $|q_j, r_j\rangle$ and $|q_j, r'_j\rangle$. Since q_j gives no information about i (by the user privacy of the classical PIR scheme) and each of r_j and r'_j is individually random, the server learns nothing about i . The j th server performs the following unitary mapping

$$|q_j, r\rangle \mapsto (-1)^{a_j \cdot r} |q_j, r\rangle,$$

which he can do because a_j only depends on q_j and x . The servers then send everything back to the user; the overall communication is $2k(t+a)$ qubits, double that of the original scheme. The user now has the state

$$\frac{1}{\sqrt{2}}|0\rangle(-1)^{a_1 \cdot r_1} |q_1, r_1\rangle \cdots (-1)^{a_k \cdot r_k} |q_k, r_k\rangle + \frac{1}{\sqrt{2}}|1\rangle(-1)^{a_1 \cdot r'_1} |q_1, r'_1\rangle \cdots (-1)^{a_k \cdot r'_k} |q_k, r'_k\rangle.$$

Up to an insignificant global phase $(-1)^{\sum_j a_j \cdot r_j}$, this is equal to

$$\begin{aligned} \frac{1}{\sqrt{2}}|0\rangle|q_1, r_1\rangle \cdots |q_k, r_k\rangle + \frac{1}{\sqrt{2}}|1\rangle(-1)^{\sum_{j=1}^k a_j \cdot b_j} |q_1, r'_1\rangle \cdots |q_k, r'_k\rangle &= \\ \frac{1}{\sqrt{2}}|0\rangle|q_1, r_1\rangle \cdots |q_k, r_k\rangle + \frac{1}{\sqrt{2}}|1\rangle(-1)^{x_i} &|q_1, r'_1\rangle \cdots |q_k, r'_k\rangle. \end{aligned}$$

The user can learn x_i from this by returning everything except the first qubit to 0, and then applying the Hadamard transform to the first qubit, which maps $\frac{1}{\sqrt{2}}(|0\rangle + (-1)^{x_i}|1\rangle) \mapsto |x_i\rangle$. On the other hand, he can learn *nothing else*, since the various states of the user during the protocol never depend on any other x_j . Accordingly, we have an honest-user QSPIR scheme with recovery probability 1 and $2k(t+a)$ qubits of communication.

Note that nowhere in the protocol do the servers have shared randomness: they do not start with it, the random strings r_j, r'_j are not correlated between servers, and the servers do not end with any shared randomness (in fact they end with nothing). Moreover, there is hardly any entanglement in the state either: tracing out the one qubit that the user keeps to himself, the state becomes unentangled.

Plugging in the best known classical PIR schemes, due to [2], gives

Theorem 1 *For every $k \geq 2$, there exists a honest-user QSPIR (without shared randomness) with communication complexity $n^{O(\log \log(k)/k \log(k))}$.*

Slightly better complexities can be obtained for small k , as stated in the first column of Table 1 in the introduction. For $k = 1$ our scheme communicates $2n$ qubits (just start from a 1-server scheme with query length 0, $a_1 = x$ and $b_1 = e_i$), for $k = 2$ it uses $O(n^{1/3})$ qubits, for $k = 3$ it uses $O(n^{1/5.25})$ qubits etc. Notice that we cannot use the (slightly better) k -server QPIR schemes from the second column of Table 1, since these reveal more than 1 bit about x .

4 Dishonest-user quantum SPIR schemes

The assumption that the user is honest (i.e., follows the protocol) is somewhat painful, since the servers cannot rely on this. In particular, a dishonest quantum user can extract about $\log n$ bits of information about x of any honest-user QSPIR where the user's final state is pure, as follows. Consider such a pure QSPIR scheme, with as many servers and communication as you like. From the user's high level perspective, this can be viewed as a unitary that maps

$$|i\rangle|0\rangle \mapsto |i\rangle|x_i\rangle|\phi_{i,x_i}\rangle.$$

Because of data privacy, the state $|\phi_{i,x_i}\rangle$ only depends on i and x_i . Therefore by one application of the QSPIR and some unitary post-processing, the user can erase $|\phi_{i,x_i}\rangle$, mapping

$$|i\rangle|0\rangle \mapsto |i\rangle|x_i\rangle,$$

for any i or superposition of i s of his choice. That is, one run of the QSPIR can be used to make one *query* to x . Van Dam [6] has shown how one quantum query to x can be used to obtain $\Omega(\log n)$ bits of information about x (in the information-theoretic sense that is, not necessarily $\log n$ specific database-bits x_j). Accordingly, any pure QSPIR that is secure against an honest user will leak at least $\Omega(\log n)$ bits of information about x to a cheating user. This includes our schemes from the previous section. Even worse, the servers cannot even *detect* whether the user cheats, because they will have the same state in the honest scheme as well as in the cheating scheme.

By Holevo’s theorem [8], the information that a dishonest quantum user can obtain about x is upper bounded by the total communication of the scheme. This is sublinear for all $k \geq 2$, but still quite a lot. How to achieve perfect privacy against dishonest quantum users? In fact, for $k \geq 2$ servers we can just use a classical SPIR that is secure against dishonest users (of course, this will be a shared-randomness scheme again). If we require the servers to measure what they receive in the computational basis, then a dishonest quantum user cannot extract more information than a classical dishonest user—that is, nothing except one x_i .

The case of SPIR with a single server is different. This primitive is equivalent to 1-out-of- n Oblivious Transfer (OT) and, when we require perfect information-theoretic privacy against a dishonest server and user, it is impossible both in the classical and in the quantum world [11]. Crépeau [5] has exhibited a quantum scheme for so-called 1-out-of-2 OT, which suffices to construct 1-out-of- n OT and is perfectly secure against honest users and a dishonest server. However, a dishonest user can learn all n bits of the database. In fact, in any OT scheme which is perfectly secure against dishonest servers, a dishonest user can always learn all n bits of the database [11]. By replicating the database in more than one servers, we can overcome this difficulty.

5 Conclusion

We have shown that the best known PIR schemes can be turned into quantum PIR schemes that are symmetrically private with respect to a honest user, i.e., except for the bit x_i that he asks for, the honest user receives no information whatsoever about the database x . Shared randomness among the servers is necessary for achieving SPIR in the classical world. Our quantum SPIR schemes don’t need this.

Rather interestingly, the best known *quantum* PIR schemes use polynomially less communication than the best known classical schemes (Table 1), but our PIR-to-QSPIR reduction does not seem to work starting from a quantum PIR system. We leave it as an open question whether the communication complexity of QSPIR schemes can be significantly reduced, either based on the QPIR schemes of [9] or via some other method.

Acknowledgments

We thank the referees for some useful comments that improved the presentation of the paper.

References

- [1] A. Ambainis. Upper bound on communication complexity of private information retrieval. In *Proceedings of the 24th ICALP*, volume 1256 of *Lecture Notes in Computer Science*, pages 401–407, 1997.
- [2] A. Beimel, Y. Ishai, E. Kushilevitz, and J. Raymond. Breaking the $O(n^{1/(2k-1)})$ barrier for information-theoretic Private Information Retrieval. In *Proceedings of 43rd IEEE FOCS*, pages 261–270, 2002.
- [3] C. H. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of the IEEE International Conference on Computers, Systems and Signal Processing*, pages 175–179, 1984.
- [4] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan. Private information retrieval. *Journal of the ACM*, 45(6):965–981, 1998. Earlier version in FOCS'95.
- [5] C. Crépeau. Quantum oblivious transfer. *Journal of Modern Optics*, 41(12):2455–2466, 1994.
- [6] W. van Dam. Quantum oracle interrogation: Getting all information for almost half the price. In *Proceedings of 39th IEEE FOCS*, pages 362–367, 1998. quant-ph/9805006.
- [7] Y. Gertner, Y. Ishai, E. Kushilevitz, and T. Malkin. Protecting data privacy in private information retrieval schemes. *Journal of Computer and Systems Sciences*, 60(3):592–629, 2000. Earlier version in STOC 98.
- [8] A. S. Holevo. Bounds for the quantity of information transmitted by a quantum communication channel. *Problemy Peredachi Informatsii*, 9(3):3–11, 1973. English translation in *Problems of Information Transmission*, 9:177–183, 1973.
- [9] I. Kerenidis and R. de Wolf. Exponential lower bound for 2-query locally decodable codes via a quantum argument. In *Proceedings of 35th ACM STOC*, pages 106–115, 2003. quant-ph/0208062.
- [10] H. Lipmaa. <http://www.tcs.hut.fi/~helger/crypto/link/protocols/retrieval.html>.
- [11] H-K. Lo. Insecurity of quantum secure computations. *Physical Review A*, 56:1154, 1997. quant-ph/9611031.
- [12] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.