Quantum Garbled Circuits

Zvika Brakerski Henry Yuen

We introduce the notion of *quantum randomized encoding* (QRE), which is an analogue of the influential notion of randomized encodings (RE) in classical cryptography and complexity theory. Given the richness and utility of randomized encodings in theoretical computer science (see below), it is very natural to ask whether there it has a *quantum* analogue. Despite its appeal, this question has remained open, and as far as we know the notion was not even formally defined in the literature prior to this work.

In this work we provide a definition and proceed to present a QRE scheme called *Quantum Garbled Circuits*, which is a quantum analogue of the famous garbled circuits construction of Yao. We also present an illustration of the usefulness of QRE to designing zero-knowledge protocols for QMA. Just as RE has proved to be a fundamental object of study in classical cryptography, we expect that quantum RE will also be a source of interesting questions and applications in quantum cryptography and quantum computing in general.

First we present some background on REs in the classical setting. A randomized encoding of a function f is another function $\hat{f}$, computed probabilistically, such that on every input x, the output $f(x)$ can be recovered from $\hat{f}(x)$, and no other information about f or x is conveyed by $\hat{f}(x)$. A trivial example of a RE of a function f is f itself. Things become much more interesting when computing $\hat{f}(x)$ is simpler in some way than computing $f(x)$; for example, $\hat{f}(x)$ could be computed via a highly parallel process even if evaluating $f(x)$ itself requires a long sequential computation.

REs are central objects in cryptographic research and have applications to a variety of settings. A famous example of a RE is Yao's garbled circuits construction [Yao86], but it was only until the work of Applebaum, Ishai and Kushilevitz in [AIK04, AIK06] that the formal notion of randomized encodings was presented. Applications of RE range from secure multi-party computation, parallel cryptography, verifiable computation, software protection, functional encryption, key-dependent message security, program obfuscation and more (see [App17] for additional details and references). Interestingly, REs have also proved useful in recent circuit lower bounds [CR20].

A useful feature of many randomized encodings is *decomposability*: this is where a function f and a sequence of inputs $(x_1, \ldots, x_n)$ can be encoded in such a way that $\hat{f}(x_1, \ldots, x_n) = (\hat{f}_{\mathsf{off}}, \hat{f}_1, \ldots, \hat{f}_n)$ where $\hat{f}_{\mathsf{off}}$ (called the "offline" part of the encoding) depends only on f and the randomness r of the encoding, and $\hat{f}_i$ (the "online" part) only depends on x_i and the randomness r. Such randomized encodings are called *decomposable*.

A good illustration of the usefulness of decomposable REs (DREs) is the task of "private simultaneous messages" (PSM) introduced by Feige, Kilian and Naor [FKN94]. In a PSM protocol for computing a function f, a set of n separated players each have an input x_i and send a message m_i to a referee, who then computes the output value $y = f(x_1, \ldots, x_n)$. The messages m_i cannot reveal any information about the x_i's aside from the fact that $f(x_1, \ldots, x_n) = y$ (formally, the messages e_i can be simulated given y). The parties share a common random string r that is independent of their inputs and unknown to the referee, and the goal is to accomplish this task using minimal communication.

Using a DRE such as garbled circuits, the parties can simply send an encoding of the function f and their respective inputs respectively to the referee. In particular, using the point-and-permute garbled circuits scheme of Beaver, Micali and Rogaway [BMR90, Rog91], it is possible to construct DREs with perfect decoding correctness and perfect simulation security for any function f with complexity that scales with the formula size of f. Assuming the adversaries are computationally bounded, it is possible to reduce this complexity to scale polynomially with the circuit size of f. Thus, the PSM task can be performed efficiently using DREs.

Quantum Randomized Encodings

In this paper we introduce the notion of randomized encodings in the quantum setting, propose a construction, and analyze it. Our definition is an adaptation of the classical one: the *quantum randomized encoding (QRE)* of a quantum operation F (represented as a quantum circuit) and a quantum state $\boldsymbol{x}$ is another quantum state $\hat{F}(\boldsymbol{x})$ satisfying two properties:

1. (*Correctness*). The quantum state $F(\boldsymbol{x})$ can be *decoded* from $\hat{F}(\boldsymbol{x})$.
2. (*Privacy*). The encoding $\hat{F}(\boldsymbol{x})$ reveals no information about F or $\boldsymbol{x}$ apart from the output $F(\boldsymbol{x})$.

The privacy property is formalized by saying there is a simulator that, given $F(\boldsymbol{x})$, can compute the encoding $\hat{F}(\boldsymbol{x})$. We also refer to $\hat{F}$ as the encoding of F. Furthermore, we also define what it means for a QRE to be *decomposable*: the encoding $\hat{F}(\boldsymbol{x})$ can be computed in a way that each qubit of the input $\boldsymbol{x}$ is encoded independently, and the encoding takes in as input $\boldsymbol{x}$, a classical random string r, and a sequence of EPR pairs $\boldsymbol{e}$. [1] We note that in the formal definition, correctness and privacy are required to hold even in the presence of *quantum auxiliary input*.

We then present a construction of a decomposable QRE, which we call the *Quantum Garbled Circuits* scheme:

Theorem 1 (Main result, informal)**.** *Suppose* CRE *is a classical DRE scheme with perfect correctness, information-theoretic (resp. computational) privacy, and polynomial time decoding. Then there exists a decomposable QRE scheme* QGC *with the following properties:*

1. QGC *has perfect correctness and polynomial-time decoding.*
2. QGC *uses* CRE *as a black box, and has information-theoretic (resp. computational) privacy. (In fact, in the information-theoretic setting we get* perfect *privacy when properly instantiating* CRE*.)*
3. *If the encoding procedure of* CRE *can be computed in* $\mathbf{NC}^0$*, then the encoding procedure of* QGC *can be computed in* $\mathbf{QNC}^0_{\mathrm{f}}$ *(i.e. the class of constant-depth quantum circuits with unbounded fan-out gates).*
4. *If the input* $\boldsymbol{x}$ *is classical, then the encoding of* $\boldsymbol{x}$ *is also classical.*

The Quantum Garbled Circuits scheme assumes the existence of a classical DRE scheme CRE with specific correctness, privacy, and complexity properties; examples of such schemes can be found in [BMR90, Rog91]. In the case of computational privacy, we assume the existence of quantum-secure one-way functions. The properties of CRE will effect the *size* of the encoded function. In the computational case, the encoded $\hat{F}$ will have size (as a quantum circuit) that

[1] We recall that an EPR pair is the maximally entangled state $\frac{1}{\sqrt{2}}\big(|00\rangle + |11\rangle\big)$, the quantum analogue of a pair of classically correlated bits.

is proportional to that of F, times a polynomial in the so-called security parameter. This is asymptotically comparable to the classical setting. In the information theoretic setting, however, the size of $\hat{F}$ could grow even doubly-exponentially with the *depth* of the encoded F, this is in contrast to the classical setting where the growth is a single-exponent. Understanding the nature of this gap is an intriguing question for further investigation.

Application: A New Zero-Knowlege Σ-Protocol for QMA

To highlight the usefulness of the notion of QRE, we present an application to designing zero-knowledge (ZK) protocols for the complexity class **QMA**. Specifically we show how to easily obtain 3-round "sigma" (abbreviated by Σ) protocols for **QMA** using QRE as a black box, and in fact our construction achieves features that were not known before in the literature.

Zero-knowledge proof systems for **QMA** have only been studied fairly recently, and known results are still few [BJSW16, VZ20, CVZ20, BG19, BS20]. A first Σ-protocol was recently presented by Broadbent and Grilo [BG19].

In this work, we show another way to construct ZK Σ-protocols for **QMA** by using quantum randomized encodings. Like [BG19], our protocol also has constant soundness error. Our protocol does not require a reudction to a specific **QMA**-complete problem as is done in [BG19], but instead directly transforms arbitrary **QMA** verifiers into a zero-knowledge verifier via QRE. Our protocol also has the feature that it only requires a single-bit challenge from the verifier for achieving soundness 1/2, and supports *delayed inputs*: the prover can produce the first message without any knowledge of the instance or the witness. In the classical setting, protocols with single-bit challenges and the delayed-input property have been useful features for parallelizing the execution of ZK protocols (even though the statement being proved in one protocol may depend on the output of another protocol).

Future directions

We list several future directions and open questions.

1. **Applications of QRE**. We presented one application in the form of a simple zero-knowledge protocol for **QMA**. Given the variety of applications of RE in classical cryptography, we anticipate that there is similarly many analogous applications in the quantum setting. We elaborate on several potential applications in Appendix B of the full paper.

2. **Obtain statistically-private QRE for all log-depth circuits**. Our information-theoretic QRE has overhead that is, in the worst case, *doubly-exponential* in the depth of the circuit being encoded. Thus there is a gap between what is achievable with classical RE (where it is possible to encode all log-depth circuits with statistical privacy). Can information-theoretic QRE be achieved for all log-depth circuits, or is the gap inherent? We note that it is not known whether statistically secure RE can be performed for all polynomial-size *classical* circuits.

3. **Completely classical encoding for quantum circuits**. Can the encoding of a quantum circuit be made completely classical? This would be very useful for obtaining obfuscation for quantum circuits, assuming classical obfuscation.

References

[AIK04] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Cryptography in NC^0. In *45th Symposium on Foundations of Computer Science (FOCS 2004), 17-19 October 2004, Rome, Italy, Proceedings*, pages 166–175, 2004.

[AIK06] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Computationally private randomizing polynomials and their applications. *Computational Complexity*, 15(2):115–162, 2006.

[App17] Benny Applebaum. Garbled circuits as randomized encodings of functions: a primer. In Yehuda Lindell, editor, *Tutorials on the Foundations of Cryptography*, pages 1–44. Springer International Publishing, 2017.

[BG19] Anne Broadbent and Alex B Grilo. Zero-knowledge for qma from locally simulatable proofs. *arXiv preprint arXiv:1911.07782*, 2019.

[BJSW16] Anne Broadbent, Zhengfeng Ji, Fang Song, and John Watrous. Zero-knowledge proof systems for qma. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 31–40. IEEE, 2016.

[BMR90] Donald Beaver, Silvio Micali, and Phillip Rogaway. The round complexity of secure protocols (extended abstract). In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, May 13-17, 1990, Baltimore, Maryland, USA*, pages 503–513, 1990.

[BS20] Nir Bitansky and Omri Shmueli. Post-quantum zero knowledge in constant rounds. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 269–279, 2020.

[CR20] Lijie Chen and Hanlin Ren. Strong average-case circuit lower bounds from non-trivial derandomization. *Electronic Colloquium on Computational Complexity (ECCC)*, 27:10, 2020. To appear in STOC 2020.

[CVZ20] Andrea Coladangelo, Thomas Vidick, and Tina Zhang. Non-interactive zero-knowledge arguments for qma, with preprocessing. In *Annual International Cryptology Conference*, pages 799–828. Springer, 2020.

[FKN94] Uriel Feige, Joe Kilian, and Moni Naor. A minimal model for secure computation (extended abstract). In Frank Thomson Leighton and Michael T. Goodrich, editors, *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing, 23-25 May 1994, Montréal, Québec, Canada*, pages 554–563. ACM, 1994.

[Rog91] Philip Rogaway. *The Round-Complexity of Secure Protocols*. PhD thesis, MIT, 1991.

[VZ20] Thomas Vidick and Tina Zhang. Classical zero-knowledge arguments for quantum computations. *Quantum*, 4:266, 2020.

[Yao86] Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *FOCS*, pages 162–167, 1986.