Nine-Month Report

Yfke Dulek

March 2017

1 Introduction

Quantum computing on encrypted data (or QCED for short) is a family of techniques and protocols that allow the evaluation of some quantum function on a quantum input, while hiding that input from the party that performs the evaluation. This is useful in a setting involving two or more parties, some of which cannot be trusted: for example, a client holding some sensitive input may want to outsource a quantum computation to a computationally more powerful, but possibly malicious server. If the server can perform computations on *encrypted* data, the privacy of the client does not need to be compromised.

Different flavors of QCED are relevant for different applications: we can allow interaction between the different parties during the computation, we can consider a setting with two parties (a client-server setting) or one with multiple parties that each want to hide their own input, we can require the authentication of the data (QCAD: quantum computing on authenticated data), et cetera. There are currently still many open questions floating around: for certain requirements or combinations of requirements, no protocol is known yet. They might even be impossible under some circumstances.

During my master's thesis project and the first months of my PhD, I have been learning about this field of research by taking courses in quantum information (both online and traditional) to strengthen my theoretical understanding, attending a winter school and several conferences to extend my network and discuss ideas, and of course by reading and reviewing related papers, both old and new. Especially during reviews and discussions, it became apparent how subtle small differences between different results can be. Questions like "How *exactly* does this relate to existing work?" and "In what way does this protocol *really* provide an advantage over what we can already do?" kept coming up.

I am now starting to feel comfortable enough with my knowledge of the field to be able to answer these questions, given enough time to dig through the literature and understand the constructions of related protocols. However, since there are so many axes along which to improve QCED schemes, papers that focus on one aspect might not always be clear on how well their proposed algorithm performs with respect to the other axes. For example, minimizing the client-server communication may result in an increased demand on the client in terms of quantum memory. Such a scheme may be very suitable for some applications but not for others.

Some survey papers exist [Her12, Fit16] that aim to make a fair comparison between different schemes, but even these tend to focus on one specific dimension of QCED. Moreover, they don't always allow for a quick look-up, and run the risk of quickly becoming outdated in this fast-moving field.

2 Why a wiki?

A wiki can provide a comparison of the latest work in quantum computing on encrypted data that is objective, complete, and flexible. The wiki I am currently setting up[1] is centered around an overview page, which displays a list of works related to QCED and their most important properties (such as efficiency, security,

[1] A public, static copy is hosted at `http://homepages.cwi.nl/~schaffne/yfke/Overview.html`. An adaptive version is hosted at `http://quantumencryption.pbworks.com`, and is currently set to private.

and interaction) and features (such as verifiability, authentication, and blindness). From the overview page, the reader can navigate to pages that contain more detailed analysis of the listed schemes. These pages contain definitions, sketches of proofs, and justifications for the properties listed on the overview page.

The advantage of a wiki over a traditional survey paper in this case is clear: the reader can access different levels of details on the listed result, and can easily navigate to pages inside the wiki (on related works and shared definitions, constructions or proof techniques) as well as resources outside of it (the original article, a Complexity Zoo entry, or a recording of a conference talk). These features ensure that the wiki has the potential to become a useful and time-saving tool for research related to QCED.

The wiki is currently a work in progress, and is not intended to ever be completely finished. This ensures that the wiki can continue to be complete and relevant for new research. It can always be improved by adding the latest results into the comparison, without having to sacrifice details on earlier results because of some page limit, as is often the case in a traditional survey paper.

3 Open questions

Apart from acting as a resource for evaluating new contributions, the wiki can also play a role in establishing new research directions. It shows in what ways existing schemes could still be improved, exposes unexplored combinations of features, and helps develop intuitions about which improvements may not be possible. In this way, it can help close or tighten the gap between the achieved and the impossible. In this section, I list some concrete open questions that are candidates for future projects.

3.1 Non-interactive QCAD

So far, all QCAD schemes (i.e. QCED schemes that act on authenticated data, allowing the client to verify that the correct computation was performed) require the client and server to communicate during the computation[2] [Bro15, BGS13, ABOE10]. The required communication is only classical and in some protocols only necessary for performing T gates: Clifford circuits can be verifiably computed without any interaction using the trap code [BGS13]. However, in order to correctly apply logical T gates, the server needs help from the client in navigating the authenticated structure of the data. Checks and traps need to be left intact, and encrypted measurement outcomes need to be interpreted.

Given recent advances in quantum fully homomorphic encryption (QFHE) [DSS16], a form of quantum computing on encrypted that requires no interaction during the computation (but also does not in itself offer the option to verify the computation afterwards), it is natural to ask whether it is possible to design an encryption scheme that allows for *both* homomorphic computation *and* verification. This would require either upgrading an existing QCAD scheme to a non-interactive one by removing the communication described above, or to upgrade an existing QFHE scheme by adding verification. Since information-theoretically secure QFHE is not possible [YPDF14], not even in an imperfect setting [NS16], it is natural to consider this task in a computational setting. This is a currently active project, where we work together with researchers at QMATH in Copenhagen.

3.2 Obfuscation

Non-interactive QCAD could be a stepping stone toward a greater goal: quantum obfuscation [AF16]. Obfuscation can in some sense be viewed as an 'upgraded' version of non-interactive QCAD: the client (or obfuscator) provides an encrypted program description to the server (or user), who not only evaluates the encrypted program on its own input, but afterwards also decrypts the result without any help of the client. The challenge is that the server/user should not be able to decrypt anything other than the valid output of the computation: therefore, before decryption, it needs to be verified that he actually did what he was supposed to do. Little is currently known about the (im)possibility of various flavors of quantum obfuscation,

[2] See also the "Interaction" table in the overview page, and consider the schemes that are marked with authenitcation in the "Security" table)

but it is very plausible that non-interactive QCAD can lead to a construction of at least one of the weaker forms of obfuscation: indistinguishability obfuscation of quantum circuits with classical inputs and outputs.

3.3 QPIP systems with a classical verifier

It is conjectured that BQP, the class of problems that can be solved by bounded-error quantum polynomial-time algorithms, is not contained in NP. Therefore, a classical computer (acting as 'verifier') will likely not be able to efficiently check a claimed solution to a BQP problem, when presented that solution by some prover. A natural way in which the classical verifier can be granted extra power is to allow it to interact classically with the (quantum) prover during the verification process. This is known as a quantum prover interactive proof (QPIP) system. It is know known whether QPIP systems with a completely classical verifier can exist. There are partial results, approaching the desired goal of a classical verifier from different directions by either allowing the verifier to have a little bit of quantum power [Bro15], or by requiring two separate, non-communicating (but possibly entangled) provers [RUV12]. A third possible approach would be to slightly relax the definition to only require computational soundness, thus being able to embed techniques that rely on computational assumptions. Of course, such a result would be strengthened by a proof/argument of the impossibility of QPIP systems with classical verifier in the information-theoretic setting.

3.4 A link between verifiability and input privacy

Unlike in the classical world, quantum data that is authenticated is necessarily also encrypted [BCG$^+$02]. In particular, this means that if a quantum authentication protocol is extended to allow verified computations (resulting in a QCAD scheme), then input privacy is guaranteed as well. With simple and generic adaptations, the protocol can then always be made blind.

It seems that verification must also imply input privacy, and even blindness. This is also confirmed by the overview tables in the wiki. There is, however, a marked exception of a verification protocol which does not have the blindness property [FH15]. Understanding why the link is broken in this case, might give us more insight into the link between verifiability and blindness.

3.5 Generalization: more than two parties

The questions posed above, and also most existing work on QCED, involve two parties who do not trust each other. This setting can be easily extended to involve three or more parties, some of which might be untrustworthy and even collaborating with other untrustworthy parties. This related area of research is called (quantum) multiparty computation, where the multiple parties each hold a part of an input state. Some features of two-party quantum computation, such as verifiability or the amount/type of interaction required, are also interesting to study in a multi-party setting. Some open questions might be simply answered by generalizing existing two-party results, while others may require new ideas to answer. Again, a well-maintained overview in the wiki would help to identify the relevant variations to consider.

3.6 Continuous-variable quantum homomorphic encryption

Delegating quantum computations from a client to a computationally more powerful server is especially relevant in the onset period of quantum computers: a few companies and universities may have a universal quantum computer, while others (the clients) may not have physical access to such powerful machines yet. In this light, it is important to consider protocols for quantum information systems that can be practically implemented with reasonably few errors. One such practically realistic setting is quantum computing using continuous variables (CV) [BVL05]: compared to the discrete variable quantum computing models used by default, CV systems require less precise equipment for accurate results, are cheaper, and work at room temperature. A protocol for CV-QCED already exists [MJS$^+$16], but requires interaction during the computation. It seems plausible that the techniques from [DSS16] to eliminate interaction can be adapted to be applicable to the continuous-variable case.

4 Summary and outlook

Research on quantum computing on encrypted data develops quickly. I have spent the past months familiarizing myself with the field, and am working toward a contribution (see Section 3.1). As a concrete output, I have started a wiki to collect known results in the field and to allow for fast and accurate comparison of new results and techniques. I aim to keep expanding this wiki, adding old and new results to create a complete and up-to-date overview, and potentially making it publicly accessible at some point.

Several open questions, some of which emerged from the wiki overview pages, are listed in Section 3. These, together with new open questions that will undoubtedly come up over the years, can provide goals for me to work towards. The listed open questions all involve some QCED-related task for which no information-theoretically secure protocol is known. For some tasks, it makes sense to try to construct protocols based on computational assumptions. In these cases, the relevance of such computationally secure protocols could be emphasized by establishing impossibility in the information-theoretic setting. Although potentially very challenging, such an impossibility result would be very neat for e.g. the question posed in Section 3.3 (if it is true).

References

[ABOE10] Dorit Aharonov, Michael Ben-Or, and Elad Eban. Interactive proofs for quantum computations. *arXiv preprint arXiv:0810.5375*, 2010.

[AF16] Gorjan Alagic and Bill Fefferman. On quantum obfuscation. *arXiv preprint arXiv:1602.01771*, 2016.

[BCG+02] Howard Barnum, Claude Crépeau, Daniel Gottesman, Adam Smith, and Alain Tapp. Authentication of quantum messages. In *Foundations of Computer Science, 2002. Proceedings. The 43rd Annual IEEE Symposium on*, pages 449–458. IEEE, 2002.

[BGS13] Anne Broadbent, Gus Gutoski, and Douglas Stebila. Quantum one-time programs. In *Advances in Cryptology–CRYPTO 2013*, pages 344–360. Springer, 2013.

[Bro15] Anne Broadbent. How to verify a quantum computation. *arXiv preprint arXiv:1509.09180*, 2015.

[BVL05] Samuel L Braunstein and Peter Van Loock. Quantum information with continuous variables. *Reviews of Modern Physics*, 77(2):513, 2005.

[DSS16] Yfke Dulek, Christian Schaffner, and Florian Speelman. Quantum homomorphic encryption for polynomial-sized circuits. In *Annual Cryptology Conference*, pages 3–32. Springer, 2016.

[FH15] Joseph F Fitzsimons and Michal Hajdušek. Post hoc verification of quantum computation. *arXiv preprint arXiv:1512.04375*, 2015.

[Fit16] Joseph F. Fitzsimons. Private quantum computation: An introduction to blind quantum computing and related protocols. *arXiv:1611.10107*, 2016.

[Her12] Charles Herder. Blind quantum computation. `http://www.scottaaronson.com/showcase2/report/charles-herder.pdf`, 2012.

[MJS+16] Kevin Marshall, Christian S Jacobsen, Clemens Schäfermeier, Tobias Gehring, Christian Weedbrook, and Ulrik L Andersen. Continuous-variable quantum computing on encrypted data. *Nature Communications*, 7, 2016.

[NS16] Michael Newman and Yaoyun Shi. Quantum homomorphic encryption, transversal computation, and their limitations. draft, personal communication, 2016.

[RUV12] Ben W Reichardt, Falk Unger, and Umesh Vazirani. A classical leash for a quantum system: Command of quantum systems via rigidity of chsh games. *arXiv preprint arXiv:1209.0448*, 2012.

[YPDF14] Li Yu, Carlos A Pérez-Delgado, and Joseph F Fitzsimons. Limitations on information-theoretically-secure quantum homomorphic encryption. *Physical Review A*, 90(5):050303, 2014.