Quantum Authentication and Encryption with Key Recycling

Or: How to Re-use a One-Time Pad Even if $P = NP$ — Safely & Feasibly

Extended Abstract*

Serge Fehr[1] and Louis Salvail[2]

[1] Cryptology Group, Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands
[2] Department of Computer Science and Operations Research (DIRO), Université de Montréal, Canada

Abstract. We propose an information-theoretically secure encryption scheme for classical messages with quantum ciphertexts that offers *detection* of eavesdropping attacks, and *re-usability of the key* in case no eavesdropping took place: the entire key can be securely re-used for encrypting new messages as long as no attack is detected. This is known to be impossible for fully classical schemes, where there is no way to detect plain eavesdropping attacks.

This particular application of quantum techniques to cryptography was originally proposed by Bennett, Brassard and Breidbart in 1982, even before proposing quantum-key-distribution, and a simple candidate scheme was suggested but no rigorous security analysis was given. The idea was picked up again in 2005, when Damgård, Pedersen and Salvail suggested a new scheme for the same task, but now with a rigorous security analysis. However, their scheme is much more demanding in terms of quantum capabilities: it requires the users to have a quantum computer.

In contrast, and like the original scheme by Bennett *et al.*, our new scheme requires from the honest users merely to prepare and measure single BB84 qubits. As such, we not only show the first provably-secure scheme that is within reach of current technology, but we also confirm Bennett *et al.*'s original intuition that a scheme in the spirit of their original construction is indeed secure.

Background. Classical information-theoretic encryption (like the one-time pad) and authentication (like Carter-Wegman authentication) have the serious downside that the key can be re-used only a small number of times, e.g. only *once* in case of the one-time pad for encryption or a strongly universal$_2$ hash function for authentication. This is inherent since by simply *observing* the communication, an eavesdropper Eve inevitably learns a substantial amount of information on the key. Furthermore, there is no way for the communicating parties, Alice and Bob, to *know* whether Eve is present and has observed the communication or not, so they have to assume the worst.

This situation changes radically when we move to the quantum setting and let the ciphertext (or authentication tag) be a quantum state: then, by the fundamental properties of quantum mechanics, an Eve that *observes* the communicated state inevitably *changes* it, and so it is potentially possible for the receiver Bob to detect this, and, vice versa, to conclude that the key is still secure and thus can be safely re-used in case everything looks as it is supposed to be.

This idea of key re-usability by means of a quantum ciphertext goes back to a manuscript titled "*Quantum Cryptography II: How to re-use a one-time pad safely even if $P = NP$*" by Bennett, Brassard and Breidbart written in 1982. However, their paper was originally not published, and the idea was put aside after two of the authors discovered what then became known as BB84 quantum-key-distribution [1].[1] Only much later in 2005, this idea was picked up again by Damgård, Pedersen and Salvail in [3], where they proposed a new such encryption scheme and gave a rigorous security proof — in contrast, Bennett *et al.*'s original reasoning was very informal and hand-wavy.

The original scheme by Bennett *et al.* is simple and natural: you one-time-pad encrypt the message, add some redundancy by encoding the ciphertext using an error correction (or detection) code, and encode the result bit-wise into what we nowadays call BB84 qubits. The scheme by Damgård *et al.* is more involved; in particular, the actual quantum encoding is not done by means of single qubits, but by means of states that form a set of mutually unbiased bases in a Hilbert space of large dimension. This in particular means that their scheme requires a *quantum computer* to produce the quantum ciphertexts and to decrypt them.

* Full version published at *EUROCRYPT 2017* and available from `https://arxiv.org/abs/1610.05614`.

[1] A freshly typeset version of the original manuscript was then published more than 30 years later in [2].

Our Results. We are interested in the question of whether one can combine the simplicity of the originally proposed encryption scheme by Bennett *et al.* with a rigorous security analysis as offered by Damgård *et al.* for their scheme; in particular, whether there is a provably secure scheme that is within reach of being implementable with current technology — and we answer the question in the affirmative.

We start with the somewhat simpler problem of finding an *authentication* scheme that allows to re-use the key in case no attack is detected, and we show a surprisingly simple solution. In order to authenticate a (classical) message msg, the sender Alice encodes a random bit string $x \in \{0,1\}^n$ into BB84 qubits $H^\theta|x\rangle$, where $\theta \in \{0,1\}^n$ is part of the shared secret key, and she computes a tag $t = \mathsf{MAC}(k, msg\|x)$ of the message concatenated with x, where MAC is a classical information-theoretic one-time message authentication code, and its key k is the other part of the shared secret key. The pair consisting of

$$H^\theta|x\rangle \quad \text{and} \quad \mathsf{MAC}(k, msg\|x)$$

is then sent along with msg, and the receiver Bob verifies correctness of the received message in the obvious way by measuring the qubits to obtain x and checking t. If Bob accepts then Alice and Bob keep the key pair (k, θ); if he rejects then they keep k but replace θ by a fresh choice.

One-time security of the scheme is obvious, and the intuition for key-reusability is as follows. Since Eve does not know θ, she has a certain minimal amount of uncertainty in x, so that, if MAC has suitable extractor-like properties, the tag t gives away no information. Furthermore, if Eve tries to gain information on k and θ by measuring some qubits, she disturbs these qubits and is likely to be detected. A subtle issue is that if Eve measures only *very few* qubits then she has a good chance of not being detected, while still learning a little bit on θ by the fact that she has not been detected. However, as long as her uncertainty in θ is large enough this should not help her (much), and the more information on θ she tries to collect this way the more likely it is that she gets caught, and then θ gets refreshed.

We show that the above intuition is correct. Formally, we prove the following theorem. Let $\rho_{K\Theta E}$ be the state *before* the execution of the scheme, where (K, Θ) is the secret key pair, consisting of the MAC key and the secret choice of bases, and E collects all of Eve's (quantum) information, and let $\rho_{K\Theta' E'}$ be the state *after* the execution, where K is the same key as before, Θ' equals Θ *if* Bob accepted and otherwise it is freshly chosen, and E' is Eve's updated information. Then, for a suitable choice of parameters, the following holds, where $\lambda \in \mathbb{N}$ is the security parameter.

Theorem. *If $\rho_{K\Theta E} = \mu_{\mathcal{K}} \otimes \rho_{\Theta E}$, where $\mu_{\mathcal{K}}$ is the density matrix for a uniformly random K, then*

$$\mathrm{Guess}(\Theta'|E') \leq \mathrm{Guess}(\Theta|E) + 2^{-3\lambda} \quad \textit{and} \quad \delta(\rho_{K\Theta' E'}, \mu_{\mathcal{K}} \otimes \rho_{\Theta' E'}) \leq 2 \cdot 2^{-\lambda} + \frac{1}{2}\sqrt{\mathrm{Guess}(\Theta|E) \cdot 2^{\lambda}}\,.$$

Here, Guess stands for the guessing probability, naturally defined, and δ for the trace distance. The statement in particular implies that if Alice and Bob start off with a uniformly random K and a uniformly random Θ from a set of size at least $2^{3\lambda}$ (which they can do with the considered choice of parameters), then it is ensured that the guessing probability of (the possibly refreshed) Θ stays in the order of $O(2^{-3\lambda})$ and K stays $O(2^{-\lambda})$-close to uniformly random.[2] This implies the re-usability claim.

For the proof of the above theorem, we show that the typical choice of the classical MAC satisfies some extractor-like property, which ensures that the key k stays close to uniformly random if there is sufficient min-entropy in the message to which MAC is applied. And, we analyze variations and extensions of the techniques introduced in [7] for studying so-called monogamy-of-entanglement games, and we use those in order to bound the min-entropy of x, which is part of the message to which MAC is applied. We refer to the full version for the technical details.

Extending our authentication scheme to an encryption scheme is intuitively quite easy: we simply extract a one-time-pad key from x, using a strong extractor (with some additional properties) with a seed that is also part of the shared secret key. Similarly to above, we can prove that as long as the receiver Bob accepts, the key can be safely re-used, and if Bob rejects it is good enough to refresh θ.

Our schemes can be made *noise robust* so as to deal with a (slightly) noisy quantum communication; the generic solution proposed in [3] of using a quantum error correction code is not an option here as it would require a quantum computer for en- and decoding. It turns out that using standard error correction techniques, like sending along the syndrome of x with respect to a suitable error correcting code, renders our proof of the bound on $\mathrm{Guess}(\Theta'|E')$ invalid beyond an obvious fix, though we would expect the scheme to still be secure. Fortunately, we can deal with the issue by means of using error correction "without leaking partial information", as introduced by Dodis and Smith [4].

[2] It is easy to deal with the issue that the theorem assumes a *perfectly* random K but ensures "only" that K stays *almost-perfectly* random.

Encryption with Key Recycling vs QKD. A possible objection against the idea of encryption with key recycling is that one might just as well use QKD to produce a new key, rather than re-using the old one. However, there *are* subtle advantages of using encryption with key recycling instead. For instance, encryption with key recycling is (almost) non-interactive and requires only *1 bit* of authenticated feedback: "accept" or "reject", that can be provided *offline*, i.e., after the communication of the private message, as long as it is done before the scheme is re-used. This opens the possibility to provide the feedback by means of a different channel, like by confirming over the phone. In contrast, for QKD, a *large amount* of data needs to be authenticated *online* and in *both directions*. Furthermore, encryption with key recycling has the potential to be *more efficient* than QKD in terms of communication. Even though this is not the case for our scheme, there is certainly potential, because no sifting takes place and hence there is no need to throw out a fraction of the quantum communication. Altogether, on a stable quantum network for instance, encryption with key-recycling could well be the preferred choice over QKD. Last but not least, given that the re-usability of a one-time-pad-like encryption key was one of the very first proposed applications of quantum cryptography — even before QKD — we feel that giving a satisfactory answer should be of intellectual interest.

Related Work. Besides the work of Brassard *et al.* and of Damgård *et al.*, who focus on encrypting *classical* messages, there is a line of work (see the full version for references), that considers key recycling in the context of authentication and/or encryption of *quantum* messages. However, common to almost all this work is that only *part* of the key can be re-used if no attack is detected, or a new but *shorter* key can be extracted. The only exceptions we know of are the two recent works by Garg *et al.* [5] and by Portmann [6], which consider and analyze authentication schemes for quantum messages that do offer re-usability of the entire key in case no attack is detected. However, these schemes are based on techniques that require the honest users to perform quantum computations also when restricting to classical messages. Actually, [6] states it as an explicit open problem to "find a prepare-and-measure scheme to encrypt and authenticate a classical message in a quantum state, so that all of the key may be recycled if it is successfully authenticated". On the other hand, their schemes offer security against *superposition attacks*, where the adversary may trick the sender into authenticating a *superposition* of classical messages; this is something we do not consider here — as a matter of fact, it would be somewhat unnatural for us since such superposition attacks require the sender (wittingly or unwittingly) to hold a quantum computer, which is exactly what we want to avoid.

Open Problems. Our work gives rise to interesting open questions. An intriguing problem is whether in the noise-tolerant version we can do the error correction in a more straightforward way, just by sending the syndrome of x with respect to a *fixed* suitable code, rather than relying on the techniques from [4]. In return, the scheme would be simpler and it could tolerate more noise. From a practical perspective, it would be interesting to see to what extent it is possible to optimize the quantum communication rather than the amount of fresh bits needed for key refreshing, which was our focus (as discussed in the full version), and whether is it possible to beat QKD in terms of quantum communication.

References

1. C.H. Bennett, and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. In *IEEE International Conference on Computers, Systems & Signal Processing*, pp. 175–179, 1984.
2. C.H. Bennett, G. Brassard, and S. Breidbart. Quantum cryptography II: How to re-use a one-time pad safely even if $P = NP$. In *Natural Computing*, 13(4):453–458 (2014).
3. I. Damgård, T. Brochmann Pedersen, L. Salvail. A quantum cipher with near optimal key-recycling. In *CRYPTO 2005*, vol. 3621 of *LNCS*, pp. 494–510 (2005). Full version in *Natural Comp.*, 13(4):469–486 (2014).
4. Y. Dodis and A. Smith. Correcting errors without leaking partial information. In *37th ACM STOC*, pp. 654–663 (2005).
5. S. Garg, H. Yuen, and M. Zhandry. New security notions and feasibility results for authentication of quantum data. Manuscript, `arXiv:1607.07759v1` (2016)
6. C. Portmann. Quantum authentication with key recycling. Manuscript, `arxiv.org/abs/1610.03422v1` (2016), also to appear in these proceedings.
7. M. Tomamichel, S. Fehr, J. Kaniewski, and S. Wehner. One-sided device independent QKD and position-based cryptography from monogamy games. In *EUROCRYPT 2013*, vol. 7881 of *LNCS*, pp. 609–625 (2013).