



Price, A., Rarity, J., & Erven, C. (2017). A Quantum Key Distribution Protocol for Rapid Denial of Service Detection. *arXiv*.
<https://arxiv.org/abs/1707.03331>

Peer reviewed version

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the accepted author manuscript (AAM). It is available online via Arxiv at <https://arxiv.org/abs/1707.03331>. Please refer to any applicable terms of use of the author.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

A quantum key distribution protocol for rapid denial of service detection.

Alasdair B. Price

Centre for Quantum Photonics and Quantum Engineering Centre for Doctoral Training,
H. H. Wills Physics Laboratory & Department of Electrical and Electronic Engineering,
University of Bristol, Nanoscience and Quantum Information Building,
Tyndall Avenue, Bristol, BS8 1FD, United Kingdom
alsadair.price@bristol.ac.uk

John G. Rarity and Chris Erven

Quantum Engineering Technology Labs, H. H. Wills Physics Laboratory
& Department of Electrical and Electronic Engineering, University of Bristol,
Nanoscience and Quantum Information Building, Tyndall Avenue,
Bristol, BS8 1FD, United Kingdom.

November 15, 2017

Abstract

We introduce a quantum key distribution protocol designed to expose fake users that connect to Alice or Bob for the purpose of monopolising the link and denying service. It inherently resists attempts to exhaust Alice and Bob's initial shared secret, and is 100% efficient, regardless of the number of qubits exchanged above the finite key limit. Additionally, secure key can be generated from two-photon pulses, without having to make any extra modifications. This is made possible by relaxing the security of BB84 to that of the quantum-safe block cipher used for day-to-day encryption, meaning the overall security remains unaffected for useful real-world cryptosystems such as AES-GCM being keyed with quantum devices.

1 Introduction

Quantum key distribution (QKD) enables two remote parties (Alice and Bob) to generate a shared secret, using quantum mechanics to ensure security against all possible theoretical attacks [4, 25]. This means it is guaranteed quantum-safe, making BB84 (the first QKD protocol, of which there are now a number of different variants) a strong candidate for protecting future communications. Further advantage can be gained in the form of eavesdropper detection, by exploiting the disturbances introduced when an attacker measures the quantum states. Unfortunately, this opens up the potential for a denial of service (DoS) attack that can be carried out simply by increasing the error rate on the transmission line. While sometimes used as an argument against QKD [11], the risk of this happening is often overstated, as it requires an attacker to have physical access to the optical fibre, so the development of large-scale networks will mitigate any damage by enabling the quantum signal to be redirected. However, there is another way of performing DoS, that does not require an adversary to monitor all connections simultaneously, and to which all current QKD protocols are vulnerable.

To prevent man-in-the-middle attacks, it is required that the classical QKD channel be authenticated, and to retain information-theoretic security, this must be done using a Wegman-Carter message authentication code (MAC) [30] keyed with a pre-shared secret.




Figure 1: Time taken for a networked *ID Quantique Clavis²* to generate a $\sim 10^5$ -bit shared secret across a newly established connection. Each system is connected to an optical switch (introducing 1 dB of loss in each case, hence a total attenuation of 2 dB at 0 km), allowing different links to be selected. The lower bound (calculated from average secret key rates) skips device and fibre characterisation, and assumes this does not affect the performance of subsequent steps, to give the shortest possible denial of service attack duration.

The MAC has to be transmitted at the end of the QKD protocol, authenticating every message sent up to that point [28], as authenticating each message individually would prohibit net positive key generation. This means neither Alice or Bob will know whether the person they are communicating with is genuine until they have a secret key, so an imposter could deny service to other users simply by opening a connection and performing QKD. Figure 1 shows how long this could last for, assuming only one round of key generation is carried out by the attacker. For a 10km metropolitan-area network, an *ID Quantique Clavis²* will communicate with an illegitimate party for roughly 10 minutes before realising! We note that the *Clavis²* continues to work at attenuations above 9 dB, but key generation starts to become intermittent. The average time taken for a successful round of QKD at 10 dB is close to 20 minutes, however the DoS impact could be greater if other rounds fail, which happens in over 30% of cases. Ultimately, it makes sense for an attacker to maximise the attenuation on their link to keep the systems occupied for as long as possible.

In this paper, we discuss how QKD can be modified to eliminate the risk of DoS attacks that leverage provably fake users. In the real world, cryptosystems that use QKD are unlikely to employ the one-time pad in day-to-day communications. Instead, quantum-safe ciphers such as the Advanced Encryption Standard (AES) [20] feature heavily [22,28], as they utilise the key more efficiently. This means the overall system is not information-theoretically secure, so reducing the mathematical security of QKD in line with the encryption algorithm will not reduce the real-world security, and will actually increase it if DoS and side-channel attacks can be mitigated as a result. By making a few additional tweaks, we show that a computationally secure QKD protocol can securely generate key even from singly detected two-photon terms, and run at exactly 100% efficiency.

2 Preliminaries

Authentication in QKD is traditionally performed using a Wegman-Carter MAC [22, 28, 30]. This takes the form

$$\tau = h_{k_H}(m) \oplus k_M \quad (1)$$

where h is a universal hash function keyed with k_H , m is the message to be authenticated (in this case, a concatenation of every transmission made over the public channel) and k_M is the key used to mask the output of the hash. Alice calculates the tag τ for the information she publicly announced, and sends it to Bob. He then computes the tag for the information he received, and compares it with Alice's tag. So long as the two are the same, he can be confident that the information has not come from or been modified by a third party (Eve). The same can then be done for the messages sent from Bob to Alice.

As described in the introduction, this way of handling QKD authentication creates the opportunity for an attacker to carry out a DoS attack, which we now formalise.

Attack 1. *Eve establishes a high-loss connection with Alice and performs low bit rate QKD up to the point where she fails the authentication. During this period, Alice and Bob are unable to generate new shared keys, which may also lead to denial of service of their classical communications. The attack can be prolonged if agents of Eve are queued behind her, turning it into a distributed denial of service (DDoS) attack.*

After succumbing to attack 1, Alice and Bob may find that they have exhausted their supply of pre-shared secret. This, a well-established vulnerability that also has the potential to be exploited independently (see attack 2), can be counteracted by using a post-quantum public key algorithm to authenticate the next round of QKD [21]. So long as Eve cannot break said algorithm in the short amount of time for which it is useful to her, full security is retained for all keys thereafter. However, by taking this approach, we have introduced a primitive that was not already part of the system, assuming Alice and Bob's initial secret was shared without using post-quantum cryptography. The recovery mechanism can also be triggered relatively easily, allowing attack 2 to be used as a way of forcing public key algorithms to be used for every successful round of QKD. Therefore, from both simplicity and security perspectives, a reactive strategy is less than ideal.

Attack 2. *Eve establishes a low-loss connection with Alice and performs high bit rate QKD up to the point where she fails the authentication. She or her agents repeat this until Alice no longer has enough secret key with which to construct a MAC. At this point, Alice must switch to an alternative method of key distribution to avoid indefinite denial of service.*

We will later show that by careful construction of a computationally secure QKD protocol, it is possible to generate secret key from two-photon terms. This is not possible in canonical BB84 because of the following photon number splitting (PNS) attack.

Attack 3. *Eve performs a quantum nondemolition measurement on the number of photons in each pulse. She blocks all single photon terms, and splits those containing multiple photons. She retains at least one photon in a quantum memory, and allows the remainder to carry on towards Bob. When Alice announces her preparation bases, Eve measures the stored photons, returning the same raw key as Alice (assuming zero errors). This can be sifted correctly when Bob publicly responds to Alice's original announcement.*

Finally, for completeness, attack 4 demonstrates what would happen if a man in the middle were able to compromise the chosen authentication scheme. As Wegman-Carter MACs are unconditionally secure, a break of this nature is not considered possible for canonical BB84, assuming Eve does not have access to the initial secret key. However, the attack will be relevant later on, when discussing the use of AES in both the QKD authentication and data encryption.

Attack 4. *Eve intercepts the quantum bits (qubits), measures each one in a random basis and resends the results she observed in the bases she measured. She conceals this by modifying Alice's bases announcement and Bob's response, along with the authentication tags for each. Eve can now read all communications encrypted and/or authenticated using the key she*

shares with Alice, before forwarding them with or without modification, having re-encrypted or authenticated using the key she shares with Bob.

3 The Protocol

We begin by trying to fulfil the main objective of this paper; preventing attack 1. A trivial solution, which preserves the information-theoretic security of BB84, would be to implement some form of access control that requests Eve verify her identity before she is allowed to connect. However, if there are no further checks until the end of the protocol, this could easily be circumvented by Eve switching out Bob for herself once key generation begins. Therefore, the most sensible approach is to authenticate every message exchanged by Alice and Bob.

Ideally, this will mean modifying equation 1 such that the tags can be reused without increasing the risk of an attacker being able to decrypt messages that rely on quantum keys. Brassard proposed in [10] that k_M could be defined as the output of a random function. In practice, this can be the cipher used for the data encryption, independently keyed with k_C , so we rewrite equation 1 as

$$\tau_i = h_{k_H}(m_i) \oplus \text{AES}_{k_C}(s_i) \quad (2)$$

where s_i is a public one-time number, or “nonce”. A number of efficient authentication schemes such as poly1305-AES [7], UMAC [8] and VMAC [16] take this form (though their moduli for addition vary), and their security when accompanying a known message is well established. For a 128-bit tag, all forgeries will be rejected with probability close to 1, so long as AES cannot be distinguished from a uniform random one-to-one function, an attacker sees no more than 2^{64} messages and, as in conventional QKD, our hash function has small differential properties [6].

As a result, just under 2^{64} qubits can be individually accompanied by a MAC, assuming Bob uses a separate initial secret key with an independent nonce for sending authenticated replies to Alice. A small number of tags must also be retained for messages relating to other parts of the protocol, such as error correction.

For finite key security, $\Omega(10^5)$ raw bits must be exchanged and processed [24], meaning we can complete $\mathcal{O}(10^{14})$ rounds of QKD before the scheme *needs* to be rekeyed. The impact of this is two-fold. First, attack 2 is no longer viable, as an eavesdropper needs to establish more than eighteen billion billion connections before Alice and Bob will be prevented from constructing any more MACs of the form given by equation 2. Second, even if Eve were able to ensure key generation only failed at the very last moment, the number of times she would have to repeat her attack in order to exhaust Alice and Bob’s shared secret is still on the order of a hundred trillion, given the rekeying limit specified above, and assuming they only began with the minimum number of bits required to construct a secure MAC. For networks of sufficient size, we would expect them to find a link that she cannot influence long before reaching that limit.

We note that the choice to use AES-256 for both data encryption and QKD authentication is not just for the sake of simplicity, or so we can be confident our cryptosystem remains quantum-safe (although as this is our reason for using QKD in the first place, it is obviously important). Suppose that, despite all the analysis that has taken place up to this point, AES has an undisclosed flaw that allows attack 4 to be carried out by a select few. The result would be catastrophic. However, it would be no different compared to if the AES-based data encrypter had been paired with canonical BB84 instead, because the encryption can be broken directly in either case, meaning attack 4 offers no advantage. Of course, the chances of this happening are thought to be very low, and so even if the one-time pad were used for data encryption, the comparative reduction in mathematical security is outweighed by increased resilience against DoS attacks.

In a world where Eve cannot compromise AES, she may carry out an unsuccessful version of attack 4 on only some of the qubits. Although Alice and Bob will be aware of her presence, there would be no way of knowing which qubits had been targeted in standard BB84, so the entire protocol would have to be aborted. In our case, the individual authentication of

every basis would allow Alice and Bob to identify which qubits had been attacked in this way, giving them the option to keep those that were unaffected.

The above changes ensure that, if Eve tries to carry out attack 1, she will deny service for fractions of seconds rather than tens of minutes before her presence becomes obvious. This is achieved without a reduction in the mathematical security of real-world QKD-based cryptosystems. The next step is to look at whether we can gain any further benefits by capitalising on our use of a computationally secure MAC.

Now that every basis announcement is accompanied by an authentication tag, an interesting property emerges. There are only two possible tags for any given key/nonce pair, depending on whether the qubit was prepared in the X basis or the Z basis (though the exact values are unpredictable for anyone not in possession of the key). This means that if Alice decides to send the tags on their own, without the plaintext basis announcement that they authenticate, Bob can work out how he should have measured the qubit, by comparing the tags he would expect for each option. The bound for rejecting forgeries will remain the same, as the plaintext can always be ignored in the case where the bases are publicly announced.

Ideally, lack of knowledge about Alice and Bob's shared secret will prevent Eve from also identifying the correct bases using the authentication tags. That is, if they provide confidentiality, which is not a traditional requirement of a MAC, then she will no longer be able to carry out attack 3. This can easily be shown to be true for tags of the form given in equation 2.

AES-CTR (AES running in counter mode [12]) encrypts a plaintext message, p_j , as follows

$$c_j = p_j \oplus \text{AES}_{k_C}(s_j) \quad (3)$$

where c_j is the ciphertext, and s_j contains a counter that increments with each value of j , never repeating for any given k_C . So long as the counter is of length 64 bits or more (with the remainder of s_j comprising of random bits), 2^{64} messages can be sent with only a minimal chance of Eve being able to recover the plaintext [3].

We observe that if one were to set $p_j = h_{k_H}(m_i)$, then equations 2 and 3 become the same. This means an AES-based MAC of Wegman-Carter form can be viewed as an implementation of AES-CTR, and so Eve will be unable to work out which basis to measure in, given only a properly implemented 128-bit tag.

From the above, we have established that transmitting the basis information as proposed means two-photon pulses can contribute to the secure key rate. However, it is still possible to implement an alternative method for PNS, on higher-order multiphoton terms (see attack 5). All protocols are vulnerable to this unless, as in [18] and [27], additional eavesdropper detection mechanisms are in place.

Attack 5. *Eve performs a quantum nondemolition measurement on the number of photons in each pulse. She blocks all single and two-photon terms, but splits those containing three or more photons. She retains at least two photons in a quantum memory, and allows the remainder to carry on towards Bob. Eve then performs unambiguous state discrimination [29] on the qubits in her possession and returns a proportion of Alice's raw key dependent on the number of photons she split out of each pulse.*

Of course, if the tags provide a level of confidentiality sufficient to prevent attack 3, there is no longer any reason for them to be transmitted after Bob has measured the qubits, as Eve is unable to obtain the information required to perform a man-in-the-middle attack. If the tags are transmitted in advance, Bob can work out how he needs to measure before each qubit arrives, increasing the efficiency of the protocol from 50% to 100%.

As a result, our protocol does not need to be implemented using biased bases which, conditional on the number of photons transmitted, are used to asymptotically double the efficiency of BB84 [17]. In fact, given we have already waived our interest in information-theoretic security, transmitting the tags in advance of the qubits is a slightly preferable solution. This is partly because the efficiencies of real and simulated biased basis experiments are still noticeably lower than 100% [13,31], however assuming no additional countermeasures are employed, the protocol described in [17] is also vulnerable to a more simplistic PNS

attack than that which is applicable to vanilla BB84. This, attack 6, is possible due to the recommendation that key be generated from a single basis, with the other used only for eavesdropper detection. The fact a quantum memory is no longer required makes it a much more realistic exploit for modern-day implementations than attack 3, emphasising why it is imperative to use decoy states in any current system relying on biased bases. In contrast, the aforementioned *Clavis*² predominantly uses unbiased SARG04 [23], which has the same level of PNS-resistance as the protocol described herein, and falls back on unbiased BB84 for short distances, where SARG04 is not proven secure [9]. This may be considered acceptable so long as quantum memories remain in the early stages of development.

Attack 6. *Assume Eve does not possess a quantum memory, but is otherwise unchanged. She performs a quantum nondemolition measurement on the number of photons in each pulse and blocks all single photon terms. For the remainder, she splits off at least one photon from every pulse, and allows at least one photon to carry on towards Bob. Eve immediately measures her copy in the key generation basis. When Alice and Bob publicly sift their qubits, she can identify those used for eavesdropper detection, and discard any information she has on them. Every bit of her final key has now been correctly measured, without revealing her presence.*

Protocol 1 pulls together the methods we have developed for performing computationally secure, but still quantum-safe, QKD. A streamlined version is presented in figure 2, the details of which can be found in the next section. Up until now, we have focused solely on utilising AES, because of its ubiquity in modern communications, and position as the de facto quantum-safe alternative to the one-time pad. However, should AES ever become compromised in some way, it would be trivial to substitute in an alternative cipher (the post-quantum security of Serpent-256 is currently under evaluation [1], for example).

While we have assumed the quantum key will be used in computationally secure cryptosystems, it is still sensible to investigate the impact of a user who insists on encrypting their data with the one-time pad in a bespoke setting, despite its low efficiency and lack of authenticated encryption modes. In this scenario, we retain the advantages of our protocol, but also acquire everlasting security [19] (the plaintext cannot be recovered from the information available to Eve if she develops unlimited computational power after key exchange is complete). This, along with perfect forward secrecy (previously generated keys will be unaffected if the initial shared secret has not been refreshed and the current round of the protocol becomes compromised), cannot be achieved if the key is encrypted directly with AES. For such a scheme, perfect forward secrecy is unattainable because anyone in possession of the long-term secret can use it to extract past session keys from the ciphertexts, rather than returning a set of bases that are no longer of any use. Similarly, compromising a previous shared secret at a later date will expose all keys distributed thereafter, even if the secret is updated after every key exchange with material from that session. Therefore, one should take care not to be fooled into thinking direct encryption of the key is a valid simplification of our protocol. Of course, a system based on this would not provide eavesdropper detection either.

4 Optimisations

While it is perfectly feasible to implement protocol 1 as presented herein, there are a number of changes that can be made to reduce demand on the computational and/or communications resources. The first of these is summarised in protocol 2, where we allow Bob to check only whether the tag he receives is a match for that corresponding to a measurement in the X basis. This requires marginally less memory and processing time than individual basis authentication in otherwise-standard BB84. The trade-off is that if Eve measures in the Z basis, she no longer needs to be able to forge the corresponding authentication tag, ensuring only that the one she forwards, τ_i^E , is different to that sent by Alice. However, Eve still has not broken the authentication scheme (she cannot obtain any basis information or force Bob to measure in the X basis), and so this kind of interference will be exposed by the quantum bit error rate (QBER). Table 1 gives the outcomes for all of Eve’s possible strategies. It is clear that $\tau_i^E \equiv \tau_i^A$ remains optimal.




Figure 2: Block diagram showing the transmission of a single bit of key from Alice to Bob, as part of BB84-AES in its reduced processing form.

Next, we look at the effect of requiring the classical channel to transmit $128 \times$ the number of bits transferred over the quantum channel. Given the *Clavis*² emits laser pulses clocked at 5 MHz [14], the classical data rate needs to be 640 Mbit/s. For comparison, the Bristol and UK quantum networks on which the *Clavis*² systems are being deployed, both have SFP+ and QSFP+ channels with capacities of 10 Gbit/s and 40 Gbit/s respectively. While the gap appears large between what we need and what we can provide, pre-commercial quantum hardware has been shown to be capable of reaching super-GHz clock speeds [26]. Due to the way in which the states were encoded in this example, the actual clock rate of BB84 was only 560 MHz, however to avoid a potential future where our protocol necessitates two transceivers be multiplexed together, we can reduce our tag lengths as described in protocol 3. This remains secure for up to 2^{32} messages [6], allowing $\mathcal{O}(10^4)$ full rounds of QKD per initial key, and brings the classical communications requirements to within the capabilities of QSFP28 or CFP4 transceivers.

The final optimisation reduces demand on the classical channel by grouping multiple bases into a single authentication tag (protocol 4). The time taken to establish the presence of a fake user should not change significantly, because the tags are still transmitted ahead of the first qubit in every group. Of course, the processing at Bob's end will be expected to take slightly longer than before, as a MAC that represents ξ bases will have $\beta = 2^\xi$ possible values for each key/nonce pair. His method for identifying the correct set of measurements differs from protocol 1 in that he must compute all possible hashes and store them in a lookup table. He can then XOR the incoming tag with the AES-generated key, and compare. Combining protocol 3 with protocol 4 will speed up the hash function [16], thereby reducing the time taken to construct the table. The necessary calculations can be performed during downtime, or in parallel with device and fibre characterisation, or in parallel with a previous round of QKD provided each initial shared secret is used across multiple rounds. An important subtlety, that is also true for protocols 1, 2 and 3, is the hashes only need to be computed once so long as the initial secret remains unchanged, meaning that until this is refreshed, the lookup table does not need to be reconstructed.

To prevent a simple timing attack, Alice can never send the qubits until the worst-case lookup time has elapsed, so Bob must take care to select a search algorithm that is optimal in this regard, such as binary search [15] which makes no more than $\lfloor \log_2 \beta \rfloor + 1 = \xi + 1$ comparisons.

The exact value of ξ reflects a trade-off between computational and communications resources, and it is clear from figure 3 that the greatest benefits can be achieved when $1 < \xi \ll 32$, because of the exponential behaviour of both classical channel capacity and memory requirements. As a concrete example, we will consider the Bristol quantum network, which is hosted on pre-existing infrastructure, with each node's server containing 64 Intel




Figure 3: Illustrating how changing the number of bases represented by a single authentication tag affects both classical communication and computational resource requirements. To get a rough estimate for how our protocol will perform on a particular physical system, one can multiply the classical channel capacity by the QKD clock rate, and worst-case number of comparisons by the time taken to perform a single binary search comparison. The size of a 128-bit hash lookup table will always be double that of its 64-bit counterpart.

Xeon E5-2697A v4 processors. By implementing a binary search on a single CPU, without hardware-specific optimisation, we can estimate the performance of our protocol on a real system. If we assume a 64-bit tag and want to employ only a single SFP+ (QSFP+) channel, then $\xi = 8$ ($\xi = 2$) maximises the QKD clock rate while trying to use the least possible memory. In this case, it takes 6.940 ± 0.085 ns (2.085 ± 0.017 ns) to run the search, allowing for a 1.153 ± 0.014 GHz (0.959 ± 0.008 GHz) clock and consuming 2048 bytes (32 bytes) of memory, out of 87.7 GiB available and 131.7 GiB total RAM. To run a hypothetical 1.72 GHz-clock BB84 device based on the technology in [26] would require $\xi = 12$ ($\xi = 3$). In this instance, the search takes 9.692 ± 0.039 ns (2.881 ± 0.036 ns), and 32,768 bytes (64 bytes) of memory is required. However, it is important to note that while these parameters are sufficient to enable the use of presently-installed transceivers, the quantum clock is still capped at 1.238 ± 0.005 GHz (1.041 ± 0.013 GHz) because of the maximum search time. Hence, some parallelisation will also be required, in that each search must begin before the previous one is guaranteed to have finished.

Technically, the higher the value of ξ , the easier it is for Eve to guess one of the $2^\xi - 1$ other authentication tags that Bob will accept. A correct guess is still highly improbable, and so she will almost certainly be detected, however even if successful, Eve controls only whether or not Bob measures with the same bases as Alice. Hence, this is nothing more than a restricted version of the strategy she can employ in protocol 2 and, in the unlikely case of an odds-defying set of forgeries, Alice and Bob will be made aware of Eve’s presence by the QBER.

5 Conclusion

We have shown that, by reducing the mathematical security of BB84, it is possible to almost instantly detect denial of service that leverages fake users, something which no other quantum key distribution protocol has been shown to be capable of. Our design is inherently resilient against attacks that aim to exhaust Alice and Bob’s supply of initial secret key, but does not lead to large memory overheads because of this, nor does it operate reactively by falling back on public key cryptography. In changing how and when the bases are announced, we are able to achieve exactly 100% efficiency and, instead of posing a risk to security, two-photon terms now contribute positively to the final key rate, independently of the distance or number of bits exchanged, and without any further cost.

Such advantages are possible only so long as the output of the cipher used to construct our authenticators is indistinguishable from the output of a random permutation. This criterion is the same as that for ensuring the security of quantum-safe encryption schemes used in day-to-day communications, so having to sacrifice information-theoretic security is not overly concerning. At any rate, the chance that the above assumption will be violated is far lower than the likelihood of an attacker exploiting one of the weaknesses that our protocol defends against. If one were to insist on unconditional security, individual basis authentication could be performed using AES tags in standard BB84, reauthenticating everything at the end with a traditional Wegman-Carter MAC. However, attack vectors may still exist for exhausting the initial shared secret and, given the issues we have raised over implementing biased bases without the necessary hardware for decoy states, BB84-AES remains preferable, particularly for minimilistic implementations and retrofitting systems already in the field.

A final novelty of our protocol is that, by daisy-chaining multiple Alice/Bob pairs, it is possible to supply an arbitrary amount of quantum-safe quantum randomness with everlasting security to someone who cannot directly access a node containing a quantum random number generator (QRNG). Of course, the resource requirements scale badly (for a chain of d nodes, the QRNG would need to generate 2^{d-1} bit strings) and while the idea may be academically interesting, it is unclear whether such functionality is of any real-world use.

Adapting our work for the Six State Protocol [2] (which we call SSP-AES) and BBM92 [5] (likewise, BBM92-AES) is trivial, however the ease with which it can be applied to other forms of quantum key distribution is less well-defined. An advantage can certainly be gained by incorporating decoy states, although this is yet to be quantified. We have shown that the intersection between modern and quantum cryptography should be explored in more detail, with greater collaboration between researchers on both sides, as this area still seems largely untapped and ripe for real-world improvements in algorithms and implementations.

Acknowledgements

A. B. Price. was supported by the Bristol Quantum Engineering Centre for Doctoral Training, EPSRC grant EP/L015730/1. The authors acknowledge the UK Quantum Technology Hub for Quantum Communications Technologies, EPSRC grant EP/M013472/1. Thanks also go to K. G. Paterson and D. L. D. Lowndes for useful conversations.

References

- [1] D. Augot, L. Batina, D. J. Bernstein, J. Bos, J. Buchmann, W. Castryck, O. Dunkelmann, T. Güneysu, S. Gueron, A. Hülsing, T. Lange, M. S. E. Mohamed, C. Rechberger, P. Schwabe, N. Sendrier, F. Vercauteren, and B.-Y. Yang. Initial Recommendations of Long-Term Secure Post-Quantum Systems. *PQCRYPTO*, 2015.
- [2] H. Bechmann-Pasquinucci and N. Gisin. Incoherent and Coherent Eavesdropping in the Six-state Protocol of Quantum Cryptography. *Physical Review A*, 59(6):4238–4248, 1999.
- [3] M. Bellare, A. Desai, E. Jorjani, and P. Rogaway. A Concrete Security Treatment of Symmetric Encryption. In *Proceedings of the 38th Symposium on Foundations of Computer Science*, pages 394–403. IEEE, 1997. Full version available from <http://web.cs.ucdavis.edu/~rogaway/papers/sym-enc.pdf> [Last Accessed: 07/04/17].

- [4] 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*, volume 1, pages 175–179. 1984.
- [5] C. H. Bennett, G. Brassard, and N. D. Mermin. Quantum Cryptography Without Bell’s Theorem. *Physical Review Letters*, 68(5):557–559, 1992.
- [6] D. J. Bernstein. Stronger Security Bounds for Wegman–Carter–Shoup Authenticators. *Advances in Cryptology: EUROCRYPT 2005. Lecture Notes in Computer Science*, 3494:164–180, 2005.
- [7] D. J. Bernstein. The Poly1305-AES Message-Authentication Code. In H. Gilbert and H. Handschuh, editors, *Fast Software Encryption. FSE 2005. Lecture Notes in Computer Science*, volume 3557, pages 32–49. 2005.
- [8] J. Black, S. Halevi, A. Hevia, H. Krawczyk, and T. Krovetz (ed.). UMAC: Message Authentication Code Using Universal Hashing. *Network Working Group, The Internet Society*, 2006.
- [9] C. Branciard, N. Gisin, B. Kraus, and V. Scarani. Security of Two Quantum Cryptography Protocols Using the Same Four Qubit States. *Physical Review A*, 72(3):032301, 2005.
- [10] G. Brassard. On Computationally Secure Authentication Tags Requiring Short Secret Shared Keys. In D. Chaum, R. L. Rivest, and A. T. Sherman, editors, *Advances in Cryptology: Proceedings of Crypto 82*, pages 79–86. Springer, 1983.
- [11] CESG. Quantum Key Distribution. White Paper, 2016.
- [12] M. Dworkin. NIST SP 800-38A: Recommendation for Block Cipher Modes of Operation. *National Institute of Standards and Technology*, 2001.
- [13] C. Erven, X. Ma, R. Laflamme, and G. Weihs. Entangled Quantum Key Distribution with a Biased Basis Choice. *New Journal of Physics*, 11(4):045025, 2009.
- [14] ID Quantique SA. Quantum Key Distribution System Clavis2 User Guide (v 3.0). 2013.
- [15] D. E. Knuth. The Art of Computer Programming. 2nd Edition, *Addison-Wesley*, 3:414, 1998.
- [16] T. Krovetz. Message Authentication on 64-Bit Architectures. In E. Biham and A. M. Youssef, editors, *Selected Areas in Cryptography: 13th International Workshop. SAC 2006 Revised Selected Papers. Lecture Notes in Computer Science*, volume 4356, pages 327–341. Springer, 2007.
- [17] H.-K. Lo, H. F. Chau, and M. Ardehali. Efficient Quantum Key Distribution Scheme and a Proof of its Unconditional Security. *Journal of Cryptology*, 18(2):133–165, 2005.
- [18] H.-K. Lo, X. Ma, and K. Chen. Decoy State Quantum Key Distribution. *Physical Review Letters*, 94(23):230504, 2005.
- [19] M. Mosca, D. Stebila, and B. Ustaoglu. Quantum Key Distribution in the Classical Authenticated Key Exchange Framework. In P. Gaborit, editor, *Post-Quantum Cryptography. PQCrypto 2013. Lecture Notes in Computer Science*, volume 7932, pages 136–154. Springer, 2013.
- [20] NIST. Specification for the Advanced Encryption Standard (AES). *Federal Information Processing Standards Publication*, 2001.
- [21] R. Roscino, K. Layat, G. Ribordy, B. Huttner, and D. Caselunghe. Applicability of a Post-Quantum Signature in a QKD Public Channel (abstract). *6th International Conference on Quantum Cryptography (QCRYPT)*, 2016.
- [22] M. Sasaki, M. Fujiwara, H. Ishizuka, W. Klaus, K. Wakui, M. Takeoka, S. Miki, T. Yamashita, Z. Wang, A. Tanaka, K. Yoshino, Y. Nambu, S. Takahashi, A. Tajima, A. Tomita, T. Domeki, T. Hasegawa, Y. Sakai, H. Kobayashi, T. Asai, K. Shimizu, T. Tokura, T. Tsurumaru, M. Matsui, T. Honjo, K. Tamaki, H. Takesue, Y. Tokura, J. F. Dynes, A. R. Dixon, A. W. Sharpe, Z. L. Yuan, A. J. Shields, S. Uchikoga, M. Legré, S. Robyr, P. Trinkler, L. Monat, J.-B. Page, G. Ribordy, A. Poppe, A. Allacher, O. Maurhart, T. Länger, M. Peev, and A. Zeilinger. Field Test of Quantum Key Distribution in the Tokyo QKD Network. *Optics Express*, 19(11):10387–10409, 2011.
- [23] V. Scarani, A. Acín, G. Ribordy, and N. Gisin. Quantum Cryptography Protocols Robust against Photon Number Splitting Attacks for Weak Laser Pulse Implementations. *Physical Review Letters*, 92(5):057901, 2004.

- [24] V. Scarani and R. Renner. Quantum Cryptography with Finite Resources: Unconditional Security Bound for Discrete-Variable Protocols with One-Way Postprocessing. *Physical Review Letters*, 100(20):200501, 2008.
- [25] P. W. Shor and J. Preskill. Simple Proof of Security of the BB84 Quantum Key Distribution Protocol. *Physical Review Letters*, 85(2):441–444, 2000.
- [26] P. Sibson, C. Erven, M. Godfrey, S. Miki, T. Yamashita, M. Fujiwara, M. Sasaki, H. Terai, M. Tanner, C. Natarajan, R. Hadfield, J. O’Brien, and M. Thompson. Chip-based Quantum Key Distribution. *Nature Communications*, 8:13984, 2017.
- [27] D. Stucki, N. Brunner, N. Gisin, V. Scarani, and H. Zbinden. Fast and Simple One-way Quantum Key Distribution. *Applied Physics Letters*, 87(19):194108, 2005.
- [28] D. Stucki, M. Legré, F. Buntschu, B. Clausen, N. Felber, N. Gisin, L. Henzen, P. Junod, G. Litzistorf, P. Monbaron, L. Monat, J.-B. Page, D. Perroud, G. Ribordy, A. Rochas, S. Robyr, J. Tavares, R. Thew, P. Trinkler, S. Ventura, R. Voinil, N. Walenta, and H. Zbinden. Long-term Performance of the SwissQuantum Quantum Key Distribution Network in a Field Environment. *New Journal of Physics*, 13(12):123001, 2011.
- [29] S. J. van Enk. Unambiguous State Discrimination of Coherent States with Linear Optics: Application to Quantum Cryptography. *Physical Review A*, 66(4):042313, 2002.
- [30] M. N. Wegman and J. L. Carter. New Hash Functions and Their Use in Authentication and Set Equality. *Journal of Computer and System Sciences*, 22:265–279, 1981.
- [31] Z. Wei, W. Wang, Z. Zhang, M. Gao, Z. Ma, and X. Ma. Decoy-state Quantum Key Distribution with Biased Basis Choice. *Scientific Reports*, 3:2453, 2013.

Protocol 1. BB84-AES (basic version)

SUMMARY: Alice expands a shared secret with Bob, using computationally secure QKD and quantum-safe primitives.

1. *One-Time Setup.*
 - (a) An l_k -bit secret is shared between Alice and Bob using out-of-band communications, a trusted third party or a post-quantum public key algorithm.
 - (b) An l_v -bit initialisation vector is transmitted from Alice to Bob in the clear, where $l_v \leq 64$.
2. *Nonce Generation.* A single-use number s_i is constructed by appending a $(128 - l_v)$ -bit counter to the initialisation vector. The counter starts at 0 and increments after each call made to the generator. It must be maintained across all rounds of QKD that use the same initial shared secret, and is not to be confused with the index i used in the mathematics of this paper, where $1 \leq i \leq N$.
3. *Authentication Tags.*
 - (a) The shared secret is split into a 256-bit cipher key, k_C , and an $(l_k - 256)$ -bit hash key, k_H .
 - (b) Alice generates a cryptographically secure random number, which is used to select a basis $b_i \in \{X, Z\}$, and computes the tag $\tau_i^A = h_{k_H}(b_i) \oplus \text{AES}_{k_C}(s_i)$. h is a universal hash function, the output of which can be called from memory after it has been evaluated once for each basis, and AES is the Advanced Encryption Standard block cipher.
 - (c) Bob calculates $\tau_i^X = h_{k_H}(X) \oplus \text{AES}_{k_C}(s_i)$ and $\tau_i^Z = h_{k_H}(Z) \oplus \text{AES}_{k_C}(s_i)$.
4. *Key Exchange.*
 - (a) Alice prepares a qubit $|\psi\rangle_i$ by generating a cryptographically secure random number, $B_i \in \{0, 1\}$, and encoding it in the basis b_i .
 - (b) Alice sends τ_i^A to Bob, closely followed by $|\psi\rangle_i$.
 - (c) Bob compares τ_i^A with τ_i^X and τ_i^Z , to identify the basis in which he should measure. Upon receipt of $|\psi\rangle_i$, he will return B_i with probability $100\% - q$, where q is the quantum bit error rate.
 - (d) Bob announces whether or not the qubit arrived, by means of an authenticated response. He should maintain a separate nonce generator to Alice, paired with a different shared secret. As Bob's response need only be "Yes" or "No", he may choose to transmit it in the same way as Alice sends her bases.
5. *Loop.* Steps 3b, 3c and 4 are repeated for the remaining $N - i$ qubits sent from Alice to Bob. As multiple tags can be constructed in parallel, this may begin prior to completion of the previous iteration.
6. *Post Processing.*
 - (a) Error correction and privacy amplification are carried out as in BB84. The messages sent during this step can be authenticated in the same way as above.
 - (b) l_k bits are taken from the final key and stored for use as the initial secret in the next round of QKD, and a new initialisation vector is publicly agreed upon.

Protocol 2. BB84-AES (reduced processing)

SUMMARY: Replaces steps 3c and 4c in protocol 1, halving the number of XOR operations and tag comparisons that Bob has to carry out.

3. *Authentication Tags.*
 - (c) Bob calculates $\tau_i^X = h_{k_H}(X) \oplus \text{AES}_{k_C}(s_i)$.
4. *Key Exchange.*
 - (c) Bob compares τ_i^A with τ_i^X . If it matches, he will choose to measure in the X basis. Otherwise, he will choose to measure in the Z basis. Upon receipt of $|\psi\rangle_i$, he will return B_i with probability $100\% - q$, where q is the quantum bit error rate.

Can be combined with: BB84-AES (reduced bandwidth)

Table 1: Showing the probability of a bit-flip error occurring between Alice and Bob depending both on the bases chosen by each of the three parties, and whether or not Eve blindly modifies the authentication tag.

Alice's Basis	Eve's Basis	Forwarding Choice	Bob's Basis	Prob(error)
X	X	$\tau_i^E = \tau_i^A$	X	0
X	X	$\tau_i^E \neq \tau_i^A$	Z	0.5
X	Z	$\tau_i^E = \tau_i^A$	X	0.5
X	Z	$\tau_i^E \neq \tau_i^A$	Z	0.5
Z	X	$\tau_i^E = \tau_i^A$	Z	0.5
Z	X	$\tau_i^E \neq \tau_i^A$	Z	0.5
Z	Z	$\tau_i^E = \tau_i^A$	Z	0
Z	Z	$\tau_i^E \neq \tau_i^A$	Z	0

Protocol 3. BB84-AES (reduced bandwidth)

SUMMARY: Replaces the 128-bit tags in protocol 1 with 64-bit tags of the same form. UMAC [8] and VMAC [16] both provide such functionality, without dropping below the required security level.

Can be combined with: BB84-AES (reduced processing), BB84-AES (dense information transfer)

Protocol 4. BB84-AES (dense information transfer)

SUMMARY: Replaces steps 3b, 3c, 4a, 4b, 4c and 5 in protocol 1, grouping multiple bases into a single tag to reduce the necessary channel capacity by a factor of $l_\tau(\xi - 1)$. l_τ is the tag length in bits, and ξ is the number of bases per tag. We redefine the range of i such that $1 \leq i \leq \frac{N}{\xi}$.

3. *Authentication Tags.*

(b) Alice generates ξ cryptographically secure random numbers, which are used to select bases b_η through $b_{\eta+\xi-1}$, where $b_{\eta+\Xi} \in \{X, Z\}$, $\eta = 1+(i-1)\xi$ and $\Xi \in \{0, \dots, \xi-1\}$. It is required that $1 < \xi \ll N$. She computes the tag $\tau_i^A = h_{k_H}(b_\eta || \dots || b_{\eta+\xi-1}) \oplus \text{AES}_{k_C}(s_i)$. h is a universal hash function, AES is the Advanced Encryption Standard block cipher, and $||$ is used to indicate a concatenation.

(c) Bob calculates $h_{k_H}(b_\eta || \dots || b_{\eta+\xi-1})$ for all 2^ξ possible values of $b_\eta || \dots || b_{\eta+\xi-1}$, storing the results in ascending order. He also evaluates $\text{AES}_{k_C}(s_i)$ separately.

4. *Key Exchange.*

(a) Alice prepares the qubits $|\psi\rangle_\eta$ to $|\psi\rangle_{\eta+\xi-1}$. This is done by generating ξ cryptographically secure random numbers B_η through $B_{\eta+\xi-1}$, where $B_{\eta+\Xi} \in \{0, 1\}$, and encoding them in the bases b_η through $b_{\eta+\xi-1}$ respectively.

(b) Alice sends τ_i^A to Bob, closely followed by all $|\psi\rangle_{\eta+\Xi}$ for the corresponding value of i .

(c) Bob computes $\tau_i^A \oplus \text{AES}_{k_C}(s_i)$ and checks it against the lookup table he constructed in step 3c, to identify the bases in which he should measure. Upon receipt of $|\psi\rangle_{\eta+\Xi}$, he will return $B_{\eta+\Xi}$ with probability $100\% - q$, where q is the quantum bit error rate.

5. *Loop.* Steps 3b, 3c and 4 are repeated for the remaining $N - i\xi$ qubits sent from Alice to Bob. As multiple tags can be constructed in parallel, this may begin prior to completion of the previous iteration.

Can be combined with: BB84-AES (reduced bandwidth)
