Appl. Phys. B 69, 395–400 (1999) / Digital Object Identifier (DOI) 10.1007/s003409900167

Applied Physics B
Lasers
and Optics

Quantum key distribution: theory for application

N. Lütkenhaus

Helsinki Institute of Physics, PL 9, FIN-00014 Helsingin yliopisto, Finland

Received: 15 June 1999/Published online: 3 November 1999

Abstract. Quantum key distribution bears the promise to set new standards in secure communication. However, on the way from the theoretical principles to the practical implementation we find many obstacles that need to be taken care of. In this article I show how to obtain a key with a realistic setup such that the security of this key can be proven for an important restricted class of eavesdropping attacks, namely the individual attacks.

PACS: 03.67.Dd; 03.67.-a; 42.79.Sz

Quantum key distribution (QKD), often refered to as quantum cryptography, promises to show a way to provable secure communication. For an introduction and an overview to different protocols see for example [1, 2]. In an idealized setting, it allows us to establish a secret key between two distant parties who do not need to share any secret beforehand. This secret key can be used as a one-time pad, or Vernam cipher, to allow secure communication given that the key is secret and that all keys of the same length are equally probable. The difference from today's cryptographic schemes lies in the fact that in this scenario the security can be proven, rather than having to rely on computational difficulty of inverse operations, such as factorising, which are needed to break a code.

When we start to implement a QKD scheme experimentally, however, we are faced with several problems. One is that the idealized setting for QKD involves a public channel which is assumed to be faithful, meaning that an eavesdropper can listen to the conversation relayed by that channel, but cannot change the signals exchanged on that channel. We find that in an experimental setting it is rather a problem to implement such a channel. Other problems involve the signal preparation, the design of a quantum channel to transport the signals, and the state detection scheme. Unlike an eavesdropper, sender and receiver are quite limited by current technology. It is therefore important to investigate the influence on the security of the final key for all deviations made from the ideal protocol. These deviations are unavoidable, and a positive proof of security has to take account of the technological restrictions of todays realizations.

In this article I describe some of the constraints for secure key distribution as far as quantum optical implementations are concerned. It turns out that many experiments performed today are actually provable insecure in principle, although an eavesdropper still faces tough problems to break the code. On the other hand, under some restrictions it is possible to give a positive proof of security for experiments using today's technology. In Sect. 1 I outline the Bennett–Brassard protocol (BB84) in an idealized setting. This protocol is used throughout this paper. Section 2 introduces measures to quantify the eavesdropper's knowledge. In Sect. 3 I summarize tools of classical computer science needed in an adaption to noisy channels. The following section clarifies what we mean by 'proving security of QKD'. The technological restraints are taken account of in Sect. 5 leading to the main results of this paper regarding realistic experiments.

1 Bennett–Brassard protocol (BB84)

The goal of the Bennett–Brassard Protocol [3] is to establish a random secret key between two parties, which are conventionally called Alice and Bob. To reach this goal, Alice sends a sequence of signals, each chosen at random from a set of four signal states. These signal states comprise two sets of orthogonal states such that the overlap probability between signals from different sets is $1/2$. An example are single photons with the first set of signals given by horizontal and vertical linear polarization, while the second set is given by right and left hand circular polarization. The receiver, Bob, uses at random one of two measurement apparatus on the incoming signals. The first is a polarizer which distinguishes vertical and horizontal linear polarization, the second distinguishes right and left circular polarization. It is clear that a signal photon prepared in a linear polarization will give a deterministic measurement result when measured with the corresponding linear polarization analyzer, and the same is true for circular polarized photons measured with the circular polarization analyzer. The other combinations, for example of linear polarized photons measured with the circular polarization analyzer, give random measurement results.

It is an essential ingredient of quantum key distribution that Alice and Bob can identify a set of highly correlated events within all events generated by randomly selected signals and measurements. This assures their advantage over an eavesdropper Eve. In the BB84 protocol this selection works using a public channel. We assume for the moment that this channel transmits signals between Alice and Bob faithfully, that is, the signals are not affected by errors and they cannot be changed by the eavesdropper. Then Alice and Bob can exchange information about the signal set from which each individual signal was drawn (linear or circular polarized), and about the measurement apparatus used. They retain only those events where the signal polarization matches the measurement apparatus. If we now translate the signals within each set as zeros and ones, for example by calling horizontal linear and right circular polarization ‘zero’ and vertical linear and left circular polarization ‘one’, then we obtain what is called the *sifted key*.

In an ideal environment and in absence of an eavesdropper, this sifted key is shared by Alice and Bob, that is, Alice’s and Bob’s version of the key coincide completely. On the other hand, as soon as an eavesdropper interacts with the signals and draws some information about them, it becomes inevitable that, on average, some error rate is introduces into the sifted key. By comparing one part of the key, this error rate becomes observable within the statistical uncertainty. If no errors are observed, then we can assume the remaining part of the sifted key to be secure and we can use it as a one-time pad for cryptography.

2 Quantifying Eve’s knowledge

Before we head into details concerning implementations of QKD we introduce some measures of Eve’s knowledge on a key. These measures compare Eve’s a priori and a posteriori probability distributions $p(x)$ and $p(x|m)$ for the key $x \in X$ of length n. The variable $m \in M$ denotes Eve’s accumulated knowledge, which occurs with total probability $q(m)$, including the communication over the public channel and her measurement results on the signals. If we compare the expected difference of the Shannon entropy for both distributions, then we obtain the Shannon information

$$\begin{aligned} I_{\mathrm{S}} = &-\sum_{x\in X} p(x)\log_2 p(x) \\ &+\sum_{m\in M} q(m)\sum_{x\in X} p(x|m)\log_2 p(x|m)\,. \end{aligned} \tag{1}$$

For an equally distributed key, $p(x) = 2^{-n}$, we find

$$I_{\mathrm{S}} = n + \sum_{m\in M} q(m)\sum_{x\in X} p(x|m)\log_2 p(x|m)\,. \tag{2}$$

In that case we obtain $I_{\mathrm{S}} = 0$ iff $p(x|m) = 2^{-n}$, which means that Eve did not obtain any information on the key. On the other hand, complete knowledge of the key is characterized by $I_{\mathrm{S}} = n$.

Another measure is the expected collision probability of the a posteriori probability $p(x|m)$. It is defined by

$$p_{\mathrm{c}} = \sum_{m\in M} q(m)\sum_{x\in X} p^2(x|m)\,. \tag{3}$$

The equally distributed a priori probability distribution gives a collision probability of $p_{\mathrm{c}}^{\text{a priori}} = 2^{-n}$, which is the lowest obtainable value. If Eve does not know anything about the key, then $p_{\mathrm{c}} = 2^{-n}$, whereas complete knowledge is characterized by $p_{\mathrm{c}} = 1$.

It is instructive to study the trade-off between the amount of information Eve can gain on the signals in relation to the disturbance she causes in the form of the observable error rate in the sifted key. The simplest example of an eavesdropping strategy is the intercept–resend strategy in one of the signal bases. One implementation is as follows: Eve measures all signals in the horizontal/vertical polarization basis and forwards a single photon to Bob which corresponds to her measurement results. It is easy to check that this will cause an error fraction of $e = 0.25$ in the sifted key and gives her a Shannon information of $I_{\mathrm{S}} = n/2$ on the sifted key. Now, Eve can perform this attack on a fraction p of the signals only. Then the error fraction goes as $e = 0.25p$ and the information as $I_{\mathrm{S}} = np/2$. In an experiment we expect to find error rates $e = 0.01$. With the described attack this means that Eve could have gained an amount of Shannon information of $I_{\mathrm{S}} = \frac{1}{2}\frac{0.01}{0.25} = 0.02$. This value is too high to use the sifted key directly as a secret key. Even more, by a better choice of the eavesdropping basis [4] or by using more sophisticated measurements [5] Eve could get even more information on the key. This motivates the use of the tools presented in the following section to extract a key from the sifted key on which Eve has negligible amount of information.

3 Classical tools for realistic environment

In a realistic environment we encounter several problems. The first one is the problem of the public channel. The major point, however, is that in a realistic setup noise in the system is basically unavoidable, and there is no way to tell noise from eavesdropping activity apart. We have seen that Eve can be in possession of non-negligible information about the key even for small error rates in the range of 1%–3%, which are typical for present day setups. This problem can be overcome by the technique of *privacy amplification* which allows us to extract a shorter key from the partially compromised sifted key such that Eve’s information on the new key is negligible. But before we can apply this technique, we need to reconcile Alice’s and Bob’s versions of the sifted key by performing error correction.

3.1 Implementing a public channel

The usual example for a public channel is a radio transmitter, but that implementation is rather impractical, and, with some substantial effort, even this channel can be tampered with. At present, we know only one reliable way to implement a public channel. This methods is that of *authentication* [6]. The idea is to use any classical channel between Alice and Bob such that both parties keep a record of what has been sent and received on that channel. At the end of the protocol, they map their respective records of the communication with a secret function into a short sequence, called the message ‘tag’. They exchange this tag so that both parties can compare their own tag with that of the other partner. The security relies on the secrecy of the mapping into the tag, and the stability of tags. The last point means that it is highly unlikely that Eve could construct the correct tag to a message she altered, if she knows only the tag for the correct message. This statement can be made more precise [6, 7], but that is beyond the scope of this paper.

The important consequence of this implementation of a public channel is that Alice and Bob now have to share some secret before they can start QKD. Nevertheless, we are able to show that it is possible to create a long secret key out of a short secret key. In other words, we expand a secret key rather than create one.

3.2 *Privacy amplification*

We cannot use the sifted key directly for secure communication if it is affected even by a small error rate of about 1%–3%. However, the tool of *generalized privacy amplification* allows us to cut Eve's knowledge from the key at the cost of the key length. To convince ourselves that this might be possible let us have a look at the following example: assume that Eve knows each bit of the key with probability $p_1 = 1/2(1+\epsilon)$. Here $\epsilon = 0$ indicates that Eve does not have a clue about the value of the bit. Now let us define a new key by taking the parity bit of two subsequent bits, thereby halving the length of the key. We find that Eve will know each bit of the new key with probability $p_2 = 1/2(1+\epsilon^2)$. With other words, for small ϵ, Eve knows much less about the new shorter key than on the old one.

As shown in [8], one does not need necessarily to half the key length. There are more subtle methods to map the original sifted key into a new shorter key. A precondition for this to work is that Alice and Bob share the same key. We have to perform error correction before we can apply privacy amplification.

By what fraction do we need to shorten the key? As long as Eve does interact with each signal separately (individual attack, see below) the answer is clear. It depends on Eve's knowledge on the key measured in the collision probability. If we shorten the key by the fraction

$$\tau_1 = 1 + \frac{1}{n} \log_2 p_c \tag{4}$$

and then by additional n_S bits so that we obtain a new key of length $n_{\text{fin}} = (1-\tau_1)n - n_s$ then the Shannon information in Eve's hand on the key is bounded by $I_S \leq \frac{2^{-n_s}}{\ln 2}$.

3.3 *Error correction*

We need to correct the key prior to privacy amplification. Of course, error correction codes are a well-studied field in computer science. However, the situation here is a non-standard one for two reasons. First, error correction works with redundant information, for example in form of giving parity bits for subsets of the key. Obviously, we cannot exchange arbitrary amounts of such information over the public channel since all such information will become available to Eve. Second, the transmission over the public channel can be made error-free, thereby allowing for specialized codes.

Table 1. Performance of the bi-directional error reconciliation protocol by Brassard and Salvail. The values are taken from that paper. Here e is the observed error rate, while f is the ratio of actually needed redundant bits to the corresponding number of the Shannon limit. (I used the bounds for $I(4)$ provided in the reference.)

e	f
0.01	1.16
0.05	1.16
0.1	1.22
0.15	1.35

To avoid the flow of side information to Eve due to exchange of parity bits over the public channel, we can encode these parity bits with secret bits shared by Alice and Bob. We then need to check at the end that we actually gain more secure bits than we put in in authentication and error correction. Alternative equivalent methods are possible [9].

From the work of Shannon [10] in classical information theory we know that error correction codes exist which can correct a key affected by error rate e in the limit of long keys such that (a) the number of redundant bits N_{rec} is given by

$$N_{\text{rec}} = n \left[e \log_2 e + (1-e) \log_2 (1-e) \right] , \tag{5}$$

and (b) the errors are corrected with unit probability. It is not possible to use fewer redundant bits. I will refer to this situation as the Shannon limit of error correction. The theorem by Shannon states the existence of such codes in the limit of long keys, but it does not construct such codes. It is therefore important to investigate what kind of codes are known and practically available for the implementation in connection with quantum key distribution. It turns out that it is rather hard to find error correction codes that use uni-directional communication only, that is, Alice sends information to Bob while Bob is passive. On the other hand, we do know one effective error correction code which uses bi-directional classical communication [11] where Alice and Bob exchange information both ways. This protocol works close to the Shannon limit, as can be seen from Table 1.

4 Proving security

It is the goal of the security analysis to state a protocol which extracts a final key from the sifted key such that (a) the Shannon information of Eve, or some other similar quantity, on that key can be bounded to be exponentially small, (b) the key can be proven to be shared between Alice and Bob with the exception of some exponentially small probability. Exponentially small here means that any specified bound on the relevant quantity can be matched without changing the ratio of the length of the sifted and the final key in the limit of long keys.

We are still far from being able to present a protocol that guarantees security of the final key without making physical assumptions about the setup. For example, we will assume that Eve cannot interfere with Alice's and Bob's apparatus, for example by looking into their apparatus to measure the signal setting or the detection setting.

Other assumptions can be made to restrict the eavesdropping ability of Eve. The study of security of QKD started considering only individual attacks. Here Eve interacts with each signal independently. This strategy can always be described as follows: Eve attaches to each signal an auxiliary system and both systems interact with each other. Then Eve stores the auxiliary system until she learns the full public discussion on the public channel. Then she decides which measurement to perform on the auxiliary system and measures each auxiliary system individually. The next scenario is that of a *collective attack* [12] which differs from that of the individual attack in that Eve can perform coherent measurements on the auxiliary systems. For that she needs the ability to manipulate those systems coherently. The most general scenario [13, 14] is that of coherent attacks. Here we drop the assumption that the auxiliary systems attached to the signals are independent of each other and of the previously sent signals. Instead, those systems can depend on the whole history of previously sent signals.

In the limit of long keys each security protocol is now characterized by the ratio between the length of the final and the sifted key. For implementations we are interested in the ratio between the length of the final key per signal sent by Alice, shortly denoted by secret bits per time slot. This ratio can be calculated for the coherent attack following a security analysis by Mayers [13]. It should be noted that these results are valid for single photon signals and for the Shannon limit of error correction while using uni-directional protocols only. It would be desirable to extend these results to cope with realistic signals and to work with bi-directional protocols.

It turns out that it is rather hard to perform the security analysis to accommodate realistic experiments in the scenario of coherent attacks. Therefore I will restrict the following results to the scenario of individual attacks. The advantage is that now elements of the realistic implementation, for example use of coherent states as signal states, and bi-directional error correction protocols, can be handled. These results can therefore be used to explore the ground for the generalization to coherent attacks. Results in the scenario of individual attacks will bound those in the scenario of coherent attacks. Even apart from that, the study of individual attacks is justified in its own right. In contrast to classical encryption methods, quantum key distribution needs to be secure only against technology available today and at the place of the transmission. If tomorrow we will have tools available to perform unlimited coherent interactions and storage of quantum systems, then this does not help to eavesdrop on today's transmissions of quantum signals in QKD.

5 Adaption to realistic experiments

In this section I will show how we can obtain a key secure against individual attacks. I will be interested in the limit of long keys only and therefore concentrate only on the ratio between secure bits and number of time slots used to generate it. After a brief review of the ideal case using single photons I present the extension to arbitrary signal states in the four BB84 polarization states.

The common ground is that the total gain of secure bits is given by

$$\begin{aligned} N_{\text{gain}} &= n_{\text{sif}}\,[1-\tau_1(e)] \\ &\quad - n_{\text{sif}}\, f\, \left[e \log_2 e + (1-e)\log_2(1-e)\right] , \end{aligned} \tag{6}$$

with the first term describing the length of the final key after privacy amplification while the second term counts the cost of error correction taking into account the factor f (see Table 1) to accommodate the fact that even the bi-directional error correction protocol does not operate at the Shannon limit described by $f=1$.

The quantity we need to know is the value of τ_1 as a function of the observed error rate. There are certain subtleties in the calculation of the observed error rate coming from the fact that Eve is not restricted to send single photon states only to Bob. However, the possibility to send multi-photon states can be can be excluded by taking into account the (unavoidable) events that the polarization analyzer will show photons in *both* output modes, leading to a photon detection in two detectors. We are not allowed to discard these events. Instead, we should randomly assign one or the other outcome. Therefore the number of double clicks will lead to an additional error rate of half of the double click rate.

5.1 Single-photon signal states

For single-photon states the value of τ_1 has been calculated in [15]. It is given via the collision probability per single signal $p_c^{(1)} = p_c^{1/n}$),

$$p_c^{(1)}(e) \le \begin{cases} \frac{1}{2} + 2e - 2e^2 & \text{for } e \le 1/2 \\ 1 & \text{for } 1/2 \le e \end{cases} . \tag{7}$$

by

$$\tau_1(e) \le \begin{cases} \log_2\left(1+4e-4e^2\right) & \text{for } e \le 1/2 \\ 1 & \text{for } 1/2 \le e \end{cases} . \tag{8}$$

The resulting gain of secure bits per time slot is given by $N_{\text{gain}}/(2n_{\text{sif}})$, taking into account that half of the signals do not contribute to the sifted key. This assumes a loss-free quantum channel and ideal detectors. Therefore this curve (Fig. 1) represents an upper bound on the rate QKD can deliver.

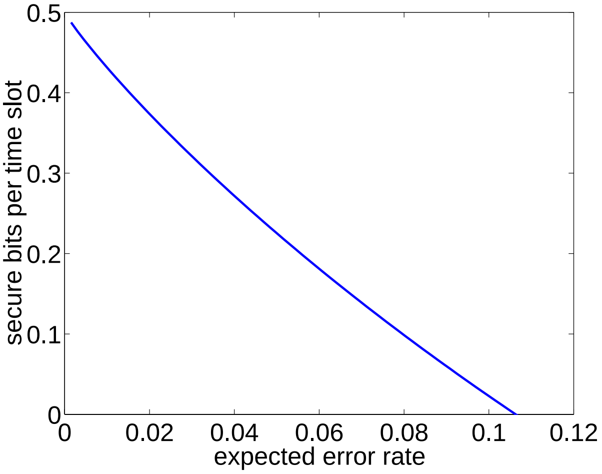

Fig. 1. Gain of secure bits per time slot as a function of the observed error rate e

5.2 Multi-photon signal states

To deal with multi-photon signals we use the observation that in a physical realization we can think of the signal states to be a mixture of Fock states in the four BB84 polarization modes. This is due to the fact that Eve does not have a phase reference for the pulses, therefore she sees a state averaged over all optical phases. This average leads to a density matrix which equals the non-averaged density matrix on the diagonal elements (in the Fock basis) while all off-diagonal elements vanish. In some setups such as the "plug and play" setup [16] Eve will have a phase reference available. However, using a phase randomizer for the signal states, we can return to the phase-averaged situation.

As a consequence of this observation, we can think of the signals to contain $0, 1, 2, \ldots$ photons in total following a classical probability distribution. Eve can perform a quantum-nondemolition measurement on the total photon number and therefore knows the photon number of each signal. Whenever she finds a multi-photon signal she can extract one photon from the signal such that the remaining signal and the extracted photon retain their original polarization. [17] This can be done using Jaynes–Cummings Hamiltonians. Since Eve can store the extracted photon until she learns the polarization basis in which the signal was prepared in (linear or circular), she will always recover the full information of signals encoded in multi-photon signals.

In the presence of loss in the quantum channel only a small fraction of signals will be successfully detected by Bob. This allows Eve to perform a powerful eavesdropping attack. She replaces the lossy quantum channel by a perfect channel. Then she attacks the multi-photon signals which allow her to get the complete signal information while at the same time Bob will successfully and error-free receive a bit as well. Eve might block all single-photon signals completely, or she might eavesdrop on a fraction on them using the optimal single-photon eavesdropping attack. This way, Eve can make sure that Bob finds precisely the number of successful detections (with or without error) he expects.

We can take care of this strategy, as the following calculation shows. It is important to observe that these results are rigorous, they are not only tailored to protect against the above attack but they give a security proof against all attacks within the scenario of individual attacks. If we can bound the number of multi-photon signals m which might contribute to the n received signals, then we can bound Eve's collision probability by

$$p_c \leq \left(p_{\mathrm{c}}^{(1)}\right)^{n-m} \quad . \tag{9}$$

since the collision probability for one multi-photon signal is given by $p_{\mathrm{c}}^{(\mathrm{multi})} = 1$. In the limit of large keys we can use the expected number of contributing multi-photon signals as the bound m. The resulting expression for τ_1 is given by

$$\tau_1^{(m)}(e^{(1)}) = 1 + \frac{n-m}{n} \log_2 p_{\mathrm{c}}^{(1)}(e^{(1)}) \; . \tag{10}$$

Here $e^{(1)}$ is not the observed error rate. As a matter of fact, we have to assume that all observed errors are due only to eavesdropping on single-photon signals. The corresponding error rate drawn on those signals then is derived from the observed error rate e as $e^{(1)} \leq e\, n/(n-m)$. It is important to point out that n is the length of the sifted key and not the number of signals sent by Alice to establish the sifted key. We therefore expect m to depend on n and of the number of signals sent by Alice.

5.3 Evaluation for experiments

To obtain a secure key in an experiment we establish a sifted key of length n using a total of $2n_{\mathrm{T}}$ signals sent by Alice. (The factor 2 takes care of those bits being discarded since Alice's signal did not match the basis of Bob's measurement.) The rate of errors in n is found to be e. In the limit of large keys ($n \to \infty$), we would like to know how long the resulting secure key could be. To calculate this, we need the source characteristics given as the probability to send no photon, one photon, or more than one photon S_0, S_1, S_m respectively. The expected number of multi-photon signals used to establish the sifted key is given by $m = S_m n_{\mathrm{T}}$, which should satisfy $m < n$ to allow secure communication [17, 18]. With that the number of secure bits extracted is given by

$$\begin{aligned} \frac{n_{\mathrm{fin}}}{n_T} = \frac{n}{n_T} & \\ & \times \left\{ -\frac{n-m}{n} \log_2 \left[\frac{1}{2} + 2e\frac{n}{n-m} - 2\left(e\frac{n}{n-m}\right)^2 \right] \right. \\ & \left. \times\; f(e)\, \left[e \log_2 e + (1-e)\log_2 (1-e) \right] \right\} \; . \end{aligned} \tag{11}$$

The observables in the experiment are S_m, n_{T}, n, e while $f(e)$ is a characteristic of the used error correction protocol. Therefore, with $m = S_m n_{\mathrm{T}}$, we can calculate directly the expected final length of a key for a given setup.

5.4 Prediction of secure key bit rates

To design a QKD experiment it is important to know the key rates we can expect to obtain. It is relatively easy to model the expected values of S_m, n_{T}, n, e for a given setup. The signal source determines S_m and we can evaluate, for example, the corresponding value for weak laser pulses. If we know the dark count rate of Bob's detection unit, the detection efficiency, the loss in the quantum channel and in Bob's detection unit, and the intrinsic error rate due to misalignment etc., then we can predict n/n_{T} and e. These three numbers allow us to predict $n_{\mathrm{fin}}/n_{\mathrm{T}}$ according to (11). In the example of weak coherent pulses the resulting fraction of secure bits is a function of the strength of the weak coherent pulse, i.e. the average photon number per signal. It turns out that there is an optimal photon number to choose which is, typically, in the order of the total transmission factor of the system (including channel loss, loss in Bob's detector and detection efficiency).

It is now interesting to predict this optimal key creation rate as a function of the distance for realistic parameters. For these parameters I have chosen two sets of numbers drawn from publication by British Telecom. The first data come from an experiment performed by Marand and Townsend [19] in the second telecommunication window at 1.3 μm. The reported values are a detection efficiency $\eta_{\mathrm{B}} = 0.11$, a dark-count rate $d_{\mathrm{B}} = 10^{-5}$ while the fiber shows a loss of $\alpha =$

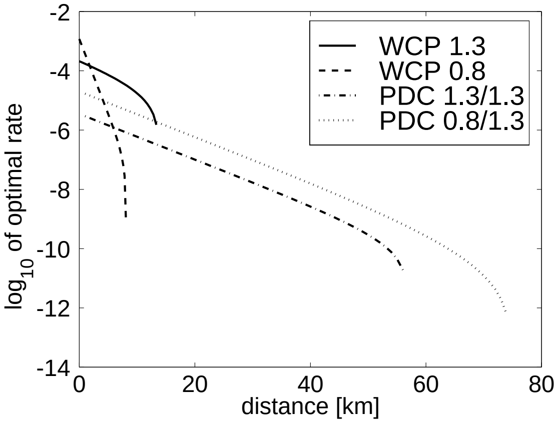

Fig. 2. Simulation of gain rate of secure bits per time slot as function of the covered distance

0.38 dB/km and the detection unit shows a loss of $c = 5$ dB. The alignment error is given by $e = 0.8\%$ of the signals. The second experiment performed by Townsend [20] at 0.8 μm is characterized by $\eta_B = 0.5$, $d_B = 5 \times 10^{-8}$, $\alpha = 2.5$ dB/km, and a loss in Bob's detection unit of $c = 8$ dB.

With these data I explore the optimal achievable secure key creation rate for four scenarios. The actual experiments were not performed at these rates. The first scenario uses weak coherent pulses at 1.3 μm (WCP 1.3), the second uses weak coherent pulses at 0.8 μm (WCP 0.8). The other two scenarios do not use weak coherent pulses as sources but they use a downconversion source with a gated output of the signal mode conditioned on the detection of a photon in the idler mode. The gating detector is modeled using the detector from the above experiments. The third scenario uses downconversion with the signal and the idler mode at 1.3 μm. The last scenario uses non-degenerate downconversion such that the idler mode at 0.8 μm is used to make use of the better detector, while the signal mode uses the lower absorption at 1.3 μm.

The simulation of these scenarios is shown in Fig. 2. We note that in all four cases there is a maximum distance for which secure communication is possible. The weak coherent schemes do not reach as far as the downconversion schemes. The rate per signal is for small distances by order of magnitudes higher for the WCP case than in the PDC case. For quite small distances, the 0.8-μm scheme gives a better rate than the 1.3-μm scheme. However, here we have to be careful with the interpretation. The four curves represent the rate of secure bits per transmitted signal. To obtain the real rate of secure bits per second we need to multiply these values with the repetition rate of the experiment. Here it turns out that the 0.8-μm scheme has been driven faster than the 1.3-μm scheme giving the 0.8 μm the leading edge up to distances higher than obvious from the graph. It is therefore important to keep an eye on the achievable repetition rate of the setup.

6 Conclusion

In this article I have analyzed the security of quantum key distribution under the restriction that the eavesdropper is restricted to individual attacks. As the analysis of the examples shows, this already restricts the range up to which QKD can be performed. This confirms that the restriction of individual attacks is a useful tool to explore the possibilities of QKD for realistic setups. For the future one can hopefully merge this approach with that of Mayers to be able to drop this restriction of Eve. Even then, however, we have not succeeded with the total proof of security of QKD with practical setups. This is due to other assumptions made on the way: we assume that Bob sends the right polarizations as signal states, and that the polarization is the only difference between the signals. Additionally, we assume that Eve cannot penetrate Alice's and Bob's setup to read off settings of phase-shifters etc. which would likewise reveal the whole key to her. In practice, one needs to protect the setup against these intrusive attacks in a convincing way. Hopefully, we will be able to drop these assumptions in the future.

Acknowledgements. I would like to thank Steve Barnett, Mohammed Bourenanne, Mila Dušek, Ondřej Haderka, Martin Hendrych, Hitoshi Inamori, Anders Karlsson, Dominic Mayers, Tal Mor, Eugene Polzik, Louis Salvail, Christian Schori, Jens-Lykke Sørensen, and Paul Townsend for all the discussions during the long-standing research into the security of practical quantum cryptography.

References

1. S.J.D. Phoenix, P.D. Townsend: Contemp. Phys. **36**, 165 (1995)
2. D. Bruß, N. Lütkenhaus: quant-ph/9901061
3. C.H. Bennett, G. Brassard: In Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing (Bangalore, India, IEEE New York 1984) pp. 175–179
4. B. Huttner, A.K. Ekert: J. Mod. Opt. **41**, 2455 (1994)
5. C.A. Fuchs, N. Gisin, R.B. Griffiths, C.-S. Niu, A. Peres: Phys. Rev. A **56**, 1163 (1997)
6. M.N. Wegman, J.L. Carter: J. Comp. Syst. Sci. **22**, 265 (1981)
7. J.L. Carter, M.N. Wegman: J. Comp. Syst. Sci. **18**, 143 (1979)
8. C.H. Bennett, G. Brassard, C. Crépeau, U.M. Maurer: IEEE Trans. Inf. Theory **41**, 1915 (1995)
9. C. Cachin, U.M. Maurer: J. Cryptology **10**, 97 (1997)
10. C. Shannon: Bell Syst. Tech. J. **27**, 379, 623 (1948)
11. G. Brassard, L. Salvail: In *Advances in Cryptology - EUROCRYPT '93* ed. by T. Helleseth, Vol. 765 of *Lecture Notes in Computer Science* (Springer, Berlin, Heidelberg 1994) pp. 410–423
12. E. Biham, M. Boyer, G. Brassard, J. van de Graaf, T. Mor: Report quant-ph/9801022 (1998)
13. D. Mayers: Report quant-ph/9802025v4 (1998)
14. H.-K. Lo, H.F. Chau: Science **283**, 2050 (1999)
15. N. Lütkenhaus: Phys. Rev. A **59**, 3301 (1999)
16. A. Muller, T. Herzog, B. Huttner, W. Tittel, H. Zbinden, N. Gisin: Applied Phys. Lett. **70**, 793 (1997)
17. G. Brassard, N. Lütkenhaus, T. Mor, B. Sanders: In preparation, 1999
18. B. Huttner, N. Imoto, N. Gisin, T. Mor: Phys. Rev. A **51**, 1863 (1995)
19. C. Marand, P.T. Townsend: Opt. Lett. **20**, 1695 (1995)
20. P.D. Townsend: IEEE Photonics Technology Letters **10**, 1048 (1998)