Journal of Modern Optics
Vol. 52, No. 5, 20 March 2005, 691–705

Taylor & Francis
Taylor & Francis Group

Faked states attack on quantum cryptosystems

VADIM MAKAROV* and DAG R. HJELME
Department of Electronics and Telecommunications, Norwegian University of Science and Technology (NTNU), NO-7491 Trondheim, Norway

(*Received 4 February 2004; in final form 3 June 2004*)

A new type of attack on quantum cryptography systems is proposed. In this attack, Eve utilizes various optical imperfections in Bob's scheme and constructs light pulses so that Bob does not distinguish his detection results from normal, whereas they give Bob the basis and bit value chosen at Eve's discretion. Applying this attack to systems with passive basis choice on Bob's side is considered. Also, a general workflow of breaking into a running quantum cryptolink using this or Trojan horse attack is discussed.

1. Introduction

Quantum cryptography was introduced as a perfectly secure way of communication based on the laws of physics. However, as the field matured and moved towards more and more practical implementations, it was slowly realized that their security consists of many components and that there are many fine points in the protocol and hardware.

A thorough discussion of quantum attacks gradually evolved to include such imperfections of physical apparatus as faint pulse sources (as opposed to true single-photon sources), loss in the transmission line and non-ideal detectors [1, 2]. Realistic key extraction protocols involving necessary authentication steps and probability estimates have been developed [3]. It has been realized that the equipment manufacturer must be trusted because there is no way for the user to verify the equipment [4]. Finally, the search came down to optical loopholes in particular implementations and classes of schemes, and eventually to electronic and software loopholes [5, 6]. It should also not be forgotten that 'classical' security at the end points of the communication link is just as important (even though this is not a task for the designer or manufacturer of communication equipment).

The whole security is only as strong as the weakest link in it. While it is still true that the laws of physics form the foundation of security in quantum cryptography, its real security will probably be determined by technological implementations, technical measures and unexpected loopholes in them [7].

So far, the search for such loopholes has attracted limited interest. One reason for the lack of interest is that these issues have little connection with fundamental physics, which most people working in the field have background in. Another reason is that quantum cryptosystems have not really taken off into widespread practical use

*Corresponding author. http://www.vad1.com

Journal of Modern Optics
ISSN 0950–0340 print/ISSN 1362–3044 online
http://www.tandf.co.uk/journals
DOI: 10.1080/09500340410001730986

(only a few devices have been sold), hence a lack of motivation to try and crack them, and to protect systems from the crackers. Nevertheless we think that a researcher has to pay attention to possible loopholes in implementations because it reflects on the technology he can use.

In this paper, we introduce a new class of practical attacks which we named *faked states attacks*. The point of this paper is not so much to estimate how easy or difficult it would be to carry out these attacks, but to raise awareness of their existence. Attacks that have not been discussed at all and have been blissfully ignored may end up as real security holes.

The structure of the paper is as follows. In section 2, we define the faked states attack, discuss it using an example of one particular implementation of quantum cryptosystem, and consider protective measures. In section 3, we give an idea of how Eve could proceed to break into a running cryptolink in practice. We discuss the general steps involved, what factors influence them and what present-day technology Eve can employ.

2. Faked states attack

Definition. Faked states attack on a quantum cryptosystem is an intercept-and-resend attack where Eve does not try to reconstruct the original states, but generates instead light pulses that get detected by the legitimate parties in a way controlled by her while not setting off any alarms.

It is well known that intercept-and-resend attack is a strategy doomed to fail if it attempts to regenerate the quantum states as close to the original as possible after detection. However, legitimate parties could sometimes be fooled, using imperfections of their set-ups, into thinking they are detecting original quantum states while they are in fact detecting light pulses generated by Eve. We call these light pulses *faked states*. Faked states are specific to each particular scheme or even particular sample of equipment being attacked.

A successful faked states attack gives Eve full knowledge of the key. (A partially successful faked states attack gives Eve partial information about the key.)

We have chosen to explain the attack using the example of the entanglement-based quantum key distribution (QKD) system developed by the Geneva group [8]. While we do not consider any other system in this paper, the above attack definition is applicable to any type of quantum cryptosystem. In particular, faked states attack can also be run against a system with *active* basis choice on Bob's side, which we would like to detail in our next paper.

The QKD system in [8] we consider here exploits photon pairs entangled in energy-time, where the sums of both the energy and the momenta of the down-converted photons equal those of the pump photon. We recap here how the system works. The photon pair source is located at Alice and is asymmetric, producing pairs where one photon in the pair has a wavelength optimized for detection (810 nm) and the other photon in the pair has a wavelength optimized for long-distance transmission (1550 nm); see figure 1. The 810 nm photon goes into Alice's interferometer, while the 1550 nm photon is sent to Bob over an optical fibre and goes into Bob's interferometer. Unbalanced Mach–Zehnder interferometers are used: an open-path, bulk optics interferometer at Alice and a fibre-optic interferometer at Bob. While the interferometers have different construction, the path difference between the short

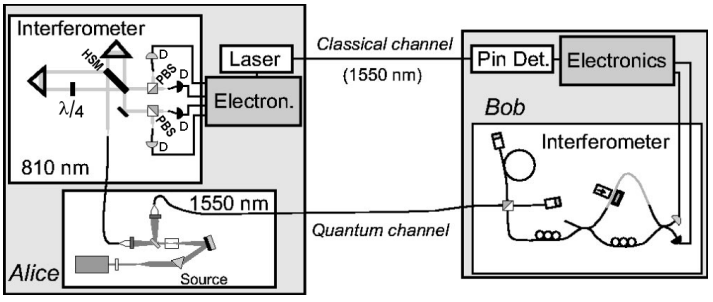

Figure 1. Reprinted from [8]: asymmetric system for quantum key distribution utilizing photon pairs (PBS, polarizing beam splitter; D, single-photon detector; $\lambda/4$, quarter-wave plate; HSM, half-silvered mirror).

and long arm is matched to a fraction of wavelength between them. Photons can propagate in four ways: both photons through the short arms at Alice and Bob, both photons through the long arms at Alice and Bob, one through the short arm at Alice and the other through the long arm at Bob, one through the long arm at Alice and the other through the short arm at Bob. The short–short and long–long processes are indistinguishable and yield two-photon interference, registered as coinciding counts at Alice's and Bob's photon detectors. Actually, whenever one of Alice's four detectors registers a count, it generates a pulse that is transmitted to Bob via a classical channel and gates his detectors in order to try to detect the other photon in the pair.

In order to do QKD, Alice and Bob must implement two incompatible measurement bases in their interferometers. The basis can be selected in each interferometer by randomly introducing either 0 or $\pi/2$ phase shift in one of the interferometer arms. In this system, the basis choice is passive on both Alice's and Bob's side. In Alice's interferometer, the $\pi/2$ phase shift is introduced at $\lambda/4$ plate (see figure 1) for one linear polarization of the beam only. The two linear polarizations get separated from one another at the polarizing beam splitters (PBSs). Photons are inserted into the interferometer and polarized such that they have about equal probability of experiencing 0 or $\pi/2$ phase shift and going either way at the PBSs; the basis for each photon is known by which pair of detectors registered it. In Bob's set-up, each photon chooses its detection basis at the PBS (see figure 2) and experiences a different delay for the two bases. Bob's detectors are gated twice and the detection basis is known by the gate that yielded a click. For additional details about this scheme, we refer the reader to [8].

We show that passive basis choice on Bob's side in this scheme in fact represents a vulnerability that can be exploited in a faked states attack.

Let us consider how faked states attack can be implemented (from Eve's standpoint) and how it can be thwarted (from Bob's standpoint) for several *attack–countermeasure* iterations.

In all implementations considered below, Eve's attack is dependent on forcing Bob to detect not in a randomly chosen basis, but in the basis chosen by Eve. Eve cuts into the line and connects the fibre running from Alice to an equivalent of Bob's set-up, noting the detection basis and bit value for every quantum state she detects. Then, she sends a faked state towards Bob for every quantum state detected, programming her detection basis into the faked state. Bob always detects it in the

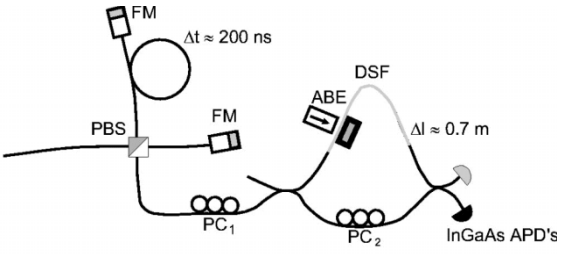

Figure 2. Reprinted from [8]: schematic diagram of Bob's interferometer (PBS, polarizing beam splitter; FM, Faraday mirror; PC, polarization controllers; ABE, adjustable birefringent element; DSF, dispersion-shifted fibre).

basis programmed by Eve. Eve's presence remains hidden, because after the sifting step all the bits in the raw key have been detected by her in the proper basis, and the subsequent check by Alice and Bob shows no increase in quantum bit error rate (QBER).

(1) Basis choice via polarization

In the original set-up used in the experiment, each photon chooses its detection basis in Bob's set-up randomly at the polarizing beam splitter (PBS), see figure 2 [8]. The choice is random because the photons in the set-up are depolarized after travelling the line that connects Alice and Bob.

In the original paper, it is noted that 'Eve could devise a strategy where she could benefit from forcing detection of a given qubit in a particular basis, [so] we must introduce a polarizer aligned at 45° or a polarization scrambler in front of the PBS'. Indeed, without this component Eve could launch polarized photons, so that they are directed into one or another PBS output port at Eve's discretion, thus allowing her to choose Bob's basis and run a successful faked states attack as described above.

(2) Basis choice via polarization using polarizer imperfections

Let us suppose Bob uses the first of the named defenses, a polarizer aligned at 45°.

Eve's task would be to make sure her photons have the desired polarization (0° or 90°) *after* the polarizer. No polarizer is perfect. We speculate in Appendix A that there always exist two input polarization states (close to the maximum extinction state of the polarizer), for which the output polarization states become the required 0° linear and 90° linear. The polarizer may have high attenuation for these polarization states (several tens of dB), but Eve can easily compensate for this by increasing the intensity of her pulses.

Thus, the polarizer alone is not a viable defense.

Let us now suppose Bob uses the second of the named defenses, a polarization scrambler. An active polarization scrambler driven from a random-number generator would transform the incoming polarization state in a way unpredictable to Eve at each moment in time. This would be a sufficient defense against this attack.

(3) Basis choice via timing using reflections off optical interfaces.

Unable to force basis choice via polarization, Eve can now exploit the fact that Bob's gated detectors are not sensitive to incoming light most of the time, and use parasitic reflections that always exist in the set-up.

The normal path for light pulses in Bob's set-up would be to reflect off the Faraday mirrors (FM) (figure 2). Let us suppose the timing of the light pulse reflected off the FM in the short arm is such that it strikes Bob's detectors during their detection window. Then, should this same pulse or some part thereof take the path in the long arm and reflect off the FM there, it would also strike the detectors during their other detection window. The timing of the two detection windows at Bob is chosen such that both parts of Alice's pulse arrive during the windows—the one that has travelled the short arm and the one that has travelled the long arm.

The two arms in Bob's interferometer would, however, likely contain other reflection points besides the FMs at the ends. There will be weaker reflections off splices, off connectors in the arms, and also off collimating optics at the PBS ports. These reflections will likely *not* be time-matched between the arms, i.e. a pulse reflected off such a parasitic reflection in one arm and hitting the detectors during their detection window will not be reflected at the corresponding point in the other arm and would not reach the detectors during the other detection window.

Thus Eve gets to choose the basis again by sending a pulse timed to reach the detectors during only one of the two detection windows, via a suitable loophole reflection path.

If a suitable single parasitic reflection does not exist for one or for both bases, Eve can search for more complex multiple-reflection paths.

Reflection levels from optical components and connectors vary widely depending on the nature, specifications and quality of the component, measuring -10 to -70 dB (see Appendix B for some examples). Using a weak parasitic reflection to route a light pulse into the detector during its detection window means a much stronger pulse that has travelled the normal path hits the detector *outside* the detection window.

If the residual sensitivity of the detector outside the detection window is high enough, this may cause an error count at Bob, which is no good for Eve. In APDs used as single-photon detectors, we envision two possible mechanisms of residual sensitivity.

(a) An APD reverse-biased below breakdown has a sensitivity to incoming light. Its A/W ratio is determined by how close to the breakdown voltage it is biased. If a light pulse causes a current through the APD comparable to the current during an avalanche, the electronics may react to it and register a 'photon count' (provided it is able to register a count timed off the normal avalanche).
(b) Current flowing through the APD caused by light outside the detection window may leave charges trapped in the junction and cause an equivalent of the afterpulsing effect at the next detection window. We shall note, however, that Bob's detection system must cope with afterpulses caused by normal avalanches. This ensures that there are at least some time zones during which a current flowing through the APD would

not have a substantial probability of causing an avalanche in the next detection window.

A possible countermeasure to this attack from Bob's side would be to eliminate the Michelson interferometer from his set-up and use two identical Mach–Zehnder interferometers and four detectors instead of one interferometer and two detectors.

(4) Basis and bit value choice via timing using non-overlapping parts of detection window

A perfect alignment of detection windows between the bases is not necessary for Bob's operation. It is enough for him if Alice's pulse arrives at the detector during the time when the detection windows for both bases overlap (figure 3); the set-up is probably adjusted to achieve this and nothing more. The detection windows, however, may and most likely will remain shifted relative to one another. By using a short pulse timed to the non-overlapping parts of the detection windows, Eve can choose the basis.

The detection windows for '0' and '1' detectors may likewise have some non-overlap between them, allowing to choose the bit value. This attack may be useful in combination with attack (3). Suppose Eve has found optical paths allowing her to choose the basis, but is having difficulty injecting light with properly aligned linear polarization into the Mach–Zehnder interferometer via these paths in Bob's set-up (or more generally, cannot align the

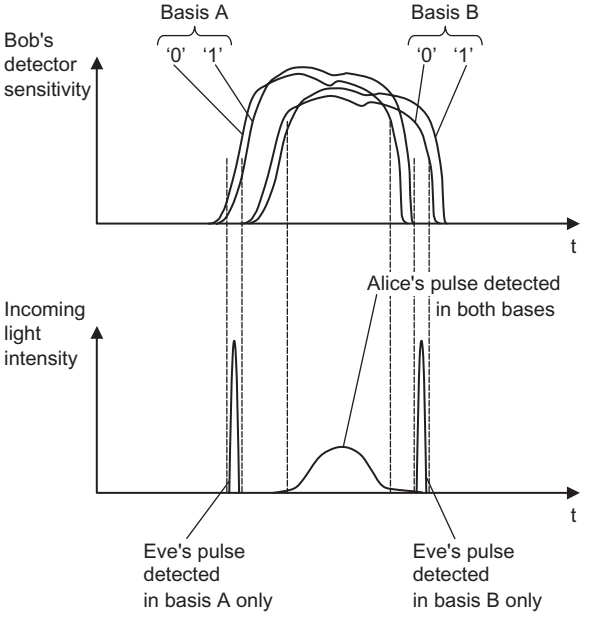

Figure 3. Bob's detection windows shown in the time frame of the incoming light pulses. Bob's detection windows may be shifted relative to each other between the bases and bit values. This does not affect normal operation: Alice's pulse is shown arriving during the time when all the windows overlap. However, if well-defined non-overlapping zones exist like on this drawing, Eve may exploit them to choose the basis (her pulses shown on the diagram), and also the bit value.

polarizations at the coupler where the pulses should interfere). Thus she is unable to obtain interference with good visibility to choose the bit value interferometrically. She may then try to choose the bit value via timing, using non-overlapping parts of '0' and '1' detection windows.

Bob's defense would be to check that his detection windows are aligned sufficiently well and do not have any non-overlap between them. This is an additional manufacturing step, or at least design step to take care of.

We would love to verify experimentally whether the vulnerabilities described in (3) and (4) exist, and if they can be practically exploited. However, this would require access to the Geneva group's experimental set-up, which we do not have. Besides, there is only a limited value in testing particular vulnerabilities of a laboratory set-up. If the set-up makes it to the production version, these vulnerabilities would not be exactly like those in the laboratory prototype and new vulnerabilities may be introduced in development.

Note that up to this point, we have implied Eve uses the same wavelength as Alice. However, Eve may use various wavelengths for her pulses, in order to exploit wavelength-dependent properties of Bob's set-up and have additional flexibility in constructing faked states. Wavelength-dependent properties useful for Eve would be different reflection and transmission coefficients, different detector sensitivity and different light speed in the fibre. In particular, antireflection coatings can have very large reflection coefficients outside the wavelength range they are designed for.

The difference in group speed can be useful in timing attack if the path length is different between the bases and between the bit values. For example, in the scheme that we are considering, the two arms of the Michelson interferometer have a path difference of 200 ns (figure 2). If we assume that the long arm is made of Corning SMF-28 fibre, the path difference would be smaller by about 80 ps when Eve uses 1310 nm wavelength pulses instead of 1550 nm used by Alice [9]. While this 80 ps imbalance is not sufficient alone for the timing attack, it may contribute towards the total detection window's misalignment required for a successful attack.

To protect from attacks (1)–(4), Bob can employ a combination of measures (figure 4): a sensitive monitoring detector, a narrowband filter that only passes Alice's wavelengths and a polarization scrambler. He must also make sure all detection windows are aligned.

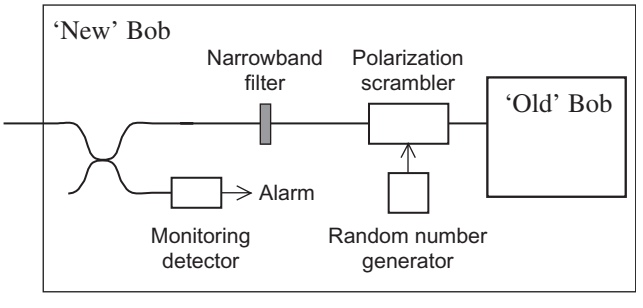

Figure 4. Additional security equipment needed to protect Bob (who is using passive basis choice) from the faked states attack described in this paper. In addition to what is shown on the diagram, Bob must match the timing of his detection windows to each other exactly.

In addition, the control software should perform all kinds of 'sanity checks' on the detector data. For example, Eve may have difficulty creating faked states for one particular basis, bit value, or basis/bit value combination, in which case she would send it less frequently or avoid it altogether. Bob should check that his detector data contains a proper mix of all possible detection outcomes (including 'double clicks' etc.), and preferably also check that the relative ratio of different outcomes does not fluctuate in time more than would be expected statistically.

(5) Incapacitation of monitoring detector

A possible strategy for Eve would be now to render the monitoring detector insensitive. It could be possibly done by damaging it with a very strong light pulse when the system is not operational (or if Bob would disregard a single alarm resulting from this action).

If the detector is employed, it is therefore advisable to power the detector and alarm trigger circuit from a battery-backed power supply. A single registered alarm, or a detection circuit downtime due to power failure should both be regarded as potential security breach events.

At this point we stop conjecturing ways to attack and note that Bob now has several extra optical components in his set-up while it is still possible for Eve to successfully compromise the link if she is lucky. Some readers will surely say that the above attacks are difficult and not very likely to succeed, or that more studies are needed to claim the attack will work. Note, however, that neither *not likely to succeed* nor *more studies are needed* is a definition that equals *perfect security*—which is what quantum cryptosystems have been supposed to be.

3. Workflow of breaking into a cryptolink

While breaking into a cryptolink, four factors influence Eve's workflow.

(1) The ability to stage the attack on samples or replicas of Alice's and Bob's apparatus. In a commercial environment, Eve would be able to buy the equipment for detailed study of its innards and for troubleshooting the attack sequence. If, however, neither the equipment nor information about its detailed construction is available to Eve, this does not make the attack impossible, just more difficult.

(2) The ability to install a tap on the optical quantum channel while it is not in use. If the channel is constantly running and an interruption of the connection for the time needed to connect to the line by conventional methods (cutting and splicing) would raise alarm, Eve must use a more elaborate technology for a non-interrupting tap.

(3) How much time elapses between quantum transmission and public discussion between Alice and Bob about this quantum transmission. By watching the public discussion, Eve can infer the QBER Alice and Bob perceive. Monitoring QBER, especially QBER for groups of bits or for particular bits, provides Eve important feedback to tune parameters of her attack (see section 3.3 below). The quicker she has the feedback, the faster she can optimize the parameters until the full eavesdropping can be run. We assume here that Alice and Bob use a bi-directional (interactive) error correction protocol, which is

more practical but leaks information about error positions to Eve [3]. If Alice and Bob, however, implement a uni-directional error correction, it will not provide Eve information about error positions.

(4) The presence of a continuous monitoring detector in the equipment being attacked. This one can make Eve's life considerably more difficult, if done properly.

The worst case Eve can face is that of breaking into a running cryptolink, not having detailed knowledge of the set-ups and the very long time between the quantum transmission and public discussion usable for inferring QBER for this transmission. This worst case combination is only a deterrent but no guarantee the attack is impossible. If any of the factors (1)–(3) turn to Eve's advantage, that makes the attack easier and more likely to succeed.

The presence of monitoring detector (4) can be a formidable deterrent in many cases. However, some types of faked states attack do not in principle require significant excess optical power at Bob. Eve can manage to keep below the detector's threshold, especially if the threshold is known to Eve from staging (1). For example, two of the attacks discussed in section 2 run with little optical power: *basis choice via polarization* and *basis and bit value choice via timing using non-overlapping parts of detection window*.

Let us now consider the general workflow, which can be divided into three stages: establishing an optical connection with the line, optical time domain reflectometry (OTDR) measurements, and testing/optimizing the attack parameters.

3.1 *Establishing optical connection with the line*

There are three possible cases.

(1) Installing a tap before installation of the quantum cryptolink. If Eve knows about installation plans in advance, this is the easiest case for her. Eve just installs her equipment onto a dark fibre and waits until Alice and Bob begin to use it for a quantum cryptolink.

The equipment Eve installs may consist of an electro-optical switch. She connects to the line with two standard fusion splices (figure 5). Depending on the timing requirements of the attack, Eve may need two switches, one for splitting off photons from the line and another for injecting pulses later down the line. The electrical signal speed in Eve's equipment can be made faster than the speed of light in the fibre line (which is about $2c/3$) through the use

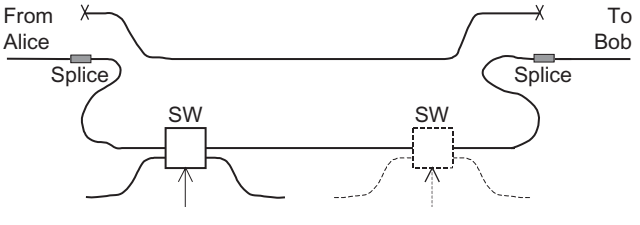

Figure 5. Tap on fibre channel that can be installed before installation of quantum cryptolink or when a quantum cryptolink is temporarily not in use (SW, electro-optical switch).

of, for example, free-space radio signals with travel speed close to c. She can compensate for her processing delay and even inject pulses at the second switch earlier than Alice's photon would have passed it.

The second switch can be substituted by a simple weak-coupling-ratio coupler. Electro-optical switches may be substituted by mechanical ones if their switching speed is sufficient to perform the attack.

(2) Installing a tap into a quantum cryptolink when it is not in operation. The times when the link is temporarily not in use may include maintenance, upgrade, equipment failure and power outage. Some of these events can be anticipated and even arranged by Eve.

If the link is not in use for the time sufficient to make two splices, Eve can employ the same approach as in (1). There are two possible differences, however: Eve may need to take care of the optical delay in the line (it should not change after installing the tap), and of the additional attenuation she introduces (it should be small enough so that Alice and Bob do not become aware of the tap).

(3) Installing a tap into a quantum cryptolink when it is constantly running. Conventional splicing cannot be used in this case. To do the tap, Eve needs a technology that neither interrupts the line nor introduces a noticeable attenuation for a significant time spell.

Although the authors admit that they are unaware of any commercially available technology that fits the requirements in the last case, there are some ideas on what it may look like. The existing evanescent-wave fibre technology would be a good place to start looking. In this technology, a part of the fibre cladding is polished away, allowing access to the mode field (figure 6(*a*)). A number of passive and active devices based on this technology are available: fixed and variable attenuators, shutters, polarizers and depolarizers, optical fibre taps, fixed- and variable-ratio couplers [10, 11]. A variable-ratio coupler consists of two fibres with parts of their claddings removed, placed in contact to one another for the length of several millimetres (figure 6 (*b*)). By laterally shifting the fibre, the coupling ratio can be varied from 0% to 100%. An added advantage of this device is virtually zero intrinsic insertion loss.

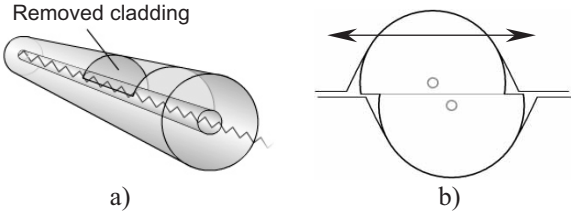

Figure 6. Evanescent-wave technology: (*a*) using side-polishing techniques, a small portion of the cladding is removed to access the evanescent tail of the propagating wave in the fibre. The removed cladding is replaced with a material to facilitate the function required for the component. The method is non-invasive to the optical core and the components modify the wave propagation by locally changing the guiding conditions, rather than impinging the propagation path (image courtesy Phoenix Photonics); (*b*) cross-section of a variable-ratio evanescent-wave coupler. The coupling ratio is varied from 0% to 100% by laterally shifting the fibres relative to each other (image courtesy Canadian Instrumentation & Research Ltd.).

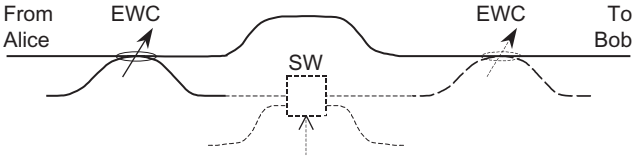

Figure 7. Tap on fibre channel that can be installed while the quantum cryptolink is in continuous operation (EWC, evanescent-wave coupler; SW, electro-optical switch).

Perhaps such couplers could be manufactured on a running quantum cryptolink (figure 7). Then, either both of them are quickly switched from 0% to 100% coupling ratio to connect Eve's devices into the line, or the coupling ratio of the first coupler is slowly varied from 0% to 100% as the attack progresses while the second one is used in weak-coupling mode for pulse injection.

Many QKD implementations also utilize a dedicated fibre-optic channel between Alice and Bob to carry synchronization signals and sometimes the public discussion. The same technology used to tap into the quantum channel can be used to tap into this classical channel.

3.2 *OTDR measurements on Bob's set-up*

After installing the tap, Eve might want to measure exact time delays to and between various reflections in Bob's set-up. She can use the standard OTDR technique to study reflections in Bob's set-up. This is especially needed if Eve does not have complete information about Bob's set-up.

There are two dangers for Eve to avoid at this stage.

(1) Getting OTDR probing pulses into Bob's detection windows, increasing QBER and setting off alarms. To avoid this, Eve can start probing with weak pulses and monitor public discussion between Alice and Bob for QBER. She can adjust the timing of her pulses, scan the whole time interval and gradually increase their brightness, until a slight raise in QBER is detected. Thus she can learn the positions of Bob's detection windows.
(2) Detection of the scanning activity by a continuous monitoring detector, if Bob has one. Perhaps, in this case Eve could use weak and infrequent pulses, performing OTDR over a long time.

3.3 *Optimizing the attack parameters*

Eve tries to proceed with the attack, substituting at first not all photons in the line, but only a few randomly chosen photons. To do this seamlessly, a high-speed switch at the tap is required. However, if the tap has been made with a single coupler, simply changing its coupling ratio and substituting those photons that get split off at the coupler and detected by Eve may be good enough.

Monitoring detection probability and QBER for those photons Eve substitutes with faked states is crucial at this stage. Eve listens to the public discussion and adjusts the parameters of the pulses she sends until they become indistinguishable from real quantum states for Bob.

After that, Eve switches to substituting every photon in the link. To say more accurately, she blocks all Alice's light from reaching Bob, diverts it to Eve's detectors, and substitutes enough detected photons with faked states so that Bob experiences the same detection rate as before.

The stages of sections 3.2 and 3.3 are optional. For a well-tested attack sequence, Eve may not need them.

The technology for breaking into a quantum cryptolink discussed above is also applicable to executing Trojan horse attacks, for example large pulse attack [6].

4. Conclusion

(1) Use of passive basis choice on Bob's side is risky from the security standpoint. Employ a random number generator and an optical modulator to set the detection basis explicitly. This way, Bob knows the basis for sure.

    To ensure true randomness of the basis, it is advisable to use a physical random number generator, for example the quantum random number generator described in [12].

(2) Installing a narrowband filter and sensitive, continuous monitoring detector on both Alice's and Bob's sides as a standard security equipment may be a good idea, whether with passive basis choice or not. Given that several attacks have been discovered that depend on shining light into legitimate parties' set-ups, this would be a justified precaution to hinder future exploits.

An interesting question arises if Eve does not find sufficiently strong vulnerabilities to run a successful faked states attack, but finds nevertheless *some* imperfections in Bob's set-up that would allow her to influence Bob's detection probabilities, for example, depending on the time delay or polarization of Alice's quantum states. This can ultimately contribute towards Eve's quantum attack and allow her to execute an attack that causes a smaller increase in the QBER level than the theoretically optimal quantum attack on a perfect set-up would cause [2]. This means the threshold QBER taking such inevitable set-up imperfections into account should be *lower* than follows from the pure theory (e.g. *less* than 11%), and Alice and Bob should compress the key more during the privacy amplification—but by how much? While this problem is addressed theoretically to some extent in [2], the task of finding the imperfections, quantifying them and refining the theoretical model remains.

The most general conclusion we have come to in our security study is that the perfect basic principles behind quantum cryptography are *not* a magic bullet that automatically provides perfect, unbreakable security. The real security story is still the perpetual cat-and-mouse game that has kept busy generations of codemakers and codebreakers over the centuries. The game continues at the next level of technology. Also, the widespread belief that the danger of quantum cryptography being broken down overnight is negligible, should be amended. This belief rests on two points [7].

(1) QKD is futureproof in the sense that any advances in technology in the future cannot be used to attack QKD keys that are created today, contrary to cryptosystems based on mathematical assumptions.

(2) Progress in technology is much easier to monitor than progress in mathematics.

While we do not dispute the first point, the second one does not appear to be always valid. An implementation loophole like one of those discussed in this paper may be discovered in secrecy and the first exploits can be made rather quickly. Depending on the diversity of the installed base of quantum cryptosystems, such an implementation loophole can temporarily compromise anything from a small percentage of systems to nearly all of them. The loophole will require a *hardware* fix once it becomes known to the public.

Is it really possible for a determined intruder to explore these attack possibilities, some of them seemingly extremely effort- and time-consuming? We find the historical example of cracking Enigma cipher inspiring in this regard. The task required increasingly larger effort, the use of an expertise unconventional to the codebreaking field at the time (mathematics), and new complex technology (electromechanical automated machines) [13]. Nevertheless, the German communications were routinely deciphered by Polish cryptanalysts and then by the Allies during 1933–1945. At the end of the war, several thousand cryptanalysts had to work to provide the British command with daily intelligence. Still, it was done. And of course, the German High Command never believed their unbreakable cipher was cracked (and so were unsuspecting former British colonies using the Enigmas in the years following the war).

In the present-day world, three-letter government agencies might be content with the fact that quantum cryptography equipment provides potential loopholes that may be exploited given a large budget and qualified dedicated staff, just like these agencies have, plus some motivation. So that they can listen to the quantum-encrypted communications while everybody else cannot.

Acknowledgments

We thank Dr Norbert Lutkenhaus for his valuable comments on the manuscript.

Appendix A: Preparing linear polarization states through an imperfect polarizer

Any linear optical component can be represented by its Jones matrix $\mathbf{T}$. The Jones matrix determines its effect on the polarization state and intensity of the incident wave $|\boldsymbol{s}\rangle$, i.e. it relates the output Jones vector $|\boldsymbol{t}\rangle$ to the input Jones vector $|\boldsymbol{s}\rangle$ via $|\boldsymbol{t}\rangle = \mathbf{T}|\boldsymbol{s}\rangle$. To generate a particular output state $|\boldsymbol{t}\rangle$ we need to apply an input state $|\boldsymbol{s}\rangle = \mathbf{T}^{-1}|\boldsymbol{t}\rangle$. Provided the Jones matrix is non-singular we can in principle generate an arbitrary output polarization state.

The Jones matrix of an ideal polarizer is singular, however, a real polarizer would have imperfections rendering the Jones matrix non-singular in general (albeit close to singular). The Jones matrix of a real polarizer rotated 45° can be written as $\mathbf{T} = \mathbf{T}_{\mathrm{LIN}} + \mathbf{T}_{\mathrm{SC}}$, where $\mathbf{T}_{\mathrm{LIN}}$ represents the ideal polarizer and $\mathbf{T}_{\mathrm{SC}}$ represents the small imperfections. The inverse of $\mathbf{T}$ will have large components, on the order of the inverse of the elements in $\mathbf{T}_{\mathrm{SC}}$. Thus to generate a linear polarized output state $|\boldsymbol{t}\rangle$

oriented at $\pm 45^\circ$ to the extinction axis will require high input power with polarization state $|\boldsymbol{s}\rangle$ close to the extinction axis.

Some of the best commercially available polarizers based on birefringent prisms have an extinction ratio of the order of 50 dB [14]. The attenuation of up to $\sim$50 dB such a polarizer would inflict on Eve's pulses can be easily compensated by their increased energy. Obtaining the proper output polarization states would require, however, a very precise setting of the input polarization states at the polarizer, which will require precise polarization control at Eve. Slight polarization instabilities in the path between Eve and Bob's polarizer may present an additional difficulty for Eve, and even effectively act as a random polarization scrambler.

Appendix B: Reflection coefficients for different fibre-optic components

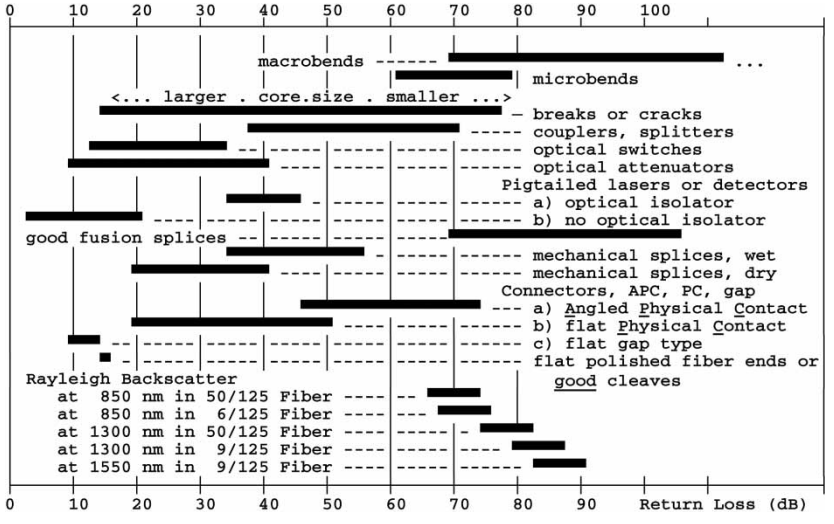

Reflection coefficient values of common fibre-optic features, components and faults. Note that the measured Rayleigh backscatter coefficients depend on the OTDR distance resolution; on this diagram, they are specific to the OFM130 system by Opto-Electronics, Inc., which has better than 1 m distance resolution in Rayleigh measurement mode. Courtesy of Opto-Electronics, Inc. (http://www.opto-electronics.com/).

References

[1] H. Inamori, N. Lutkenhaus and D. Mayers, quant-ph/0107017 (2001).
[2] D. Gottesman, H.-K. Lo, N. Lutkenhaus and J. Preskill, quant-ph/0212066 (2002).
[3] N. Lutkenhaus, Phys. Rev. A **59** 3301 (1999).
[4] J. Larsson, Quantum Inf. Comput. **2** 434 (2002).
[5] C. Kurtsiefer, P. Zarda, S. Mayer and H. Weinfurter, J. Mod. Optics **48** 2039 (2001).
[6] A. Vakhitov, V. Makarov and D.R. Hjelme, J. Mod. Optics **48** 2023 (2001).
[7] N. Gisin, G. Ribordy, W. Tittel and H. Zbinden, Rev. Mod. Phys. **74** 145 (2002).
[8] G. Ribordy, J. Brendel, J.-D. Gautier, *et al.*, Phys. Rev. A **63** 012309 (2001).

[9] According to *Corning® SMF-28™ CPC6 single-mode optical fibre*, a product information sheet PL1036 (Corning Incorporated, 1999), the effective group index of refraction (N_{eff}) is 1.4675 at 1310 nm and 1.4681 at 1550 nm, which makes for 0.04% difference in group speed between these wavelengths.
[10] Phoenix Photonics, http://www.phoenix-photonics.com/, see Products section.
[11] Canadian Instrumentation & Research Ltd, http://www.cirl.com/, see Products section.
[12] A. Stefanov, N. Gisin, O. Guinnard, L. Guinnard and H. Zbinden, quant-ph/9907006. The optical quantum random number generator described in the above paper is commercially available from id Quantique, http://www.idquantique.com/ See also more recent (introduced in March 2004) Quantis quantum random number generator (1999).
[13] S. Singh, *The Code Book* (Fourth Estate, London, 2000).
[14] Melles Griot Inc, *Melles Griot® Product Catalog* (Melles Griot, Carlsbad, California, 1999).