Telephones and faxes are not perfectly secure, but send a secret message made up of quantum bits, and you can know for sure if it was read before reaching its intended target

Quantum cryptography

Wolfgang Tittel, Grégoire Ribordy and Nicolas Gisin

WE LIVE in a quantum world – something that physicists have considered with amazement for more than seventy years. But we now realize that quantum physics is more than a radical departure from classical physics. It also offers many new possibilities for information processing.

Quantum cryptography is the most mature prospect of this fascinating new field. It is based on the fundamental postulate of quantum physics that "every measurement perturbs a system". Imagine sending a message carried by single quantum states, such as linearly polarized photons oriented at various angles. If the bits are not altered during transmission, you can be sure that no eavesdropper has measured the values of those bits. In other words, quantum cryptography turns an apparent limitation – namely that a measurement perturbs the system – into a potentially useful process, in which the perturbation uncovers the presence of an eavesdropper.

This idea of turning quantum conundrums into potentially useful processes is a characteristic of the whole field of "quantum information processing". For example, the famous Einstein Podolsky–Rosen paradox has lead to novel techniques such as "dense coding" and "quantum teleportation" (see "Fundamentals of quantum information" by Zeilinger on page 35). "Quantum entanglement", meanwhile, could make it possible to build quantum computers that could factorize large integers exponentially faster than the best-known algorithm for classical computers (see "Quantum computation" by Deutsch and Ekert on page 47).

1 Carry on spying

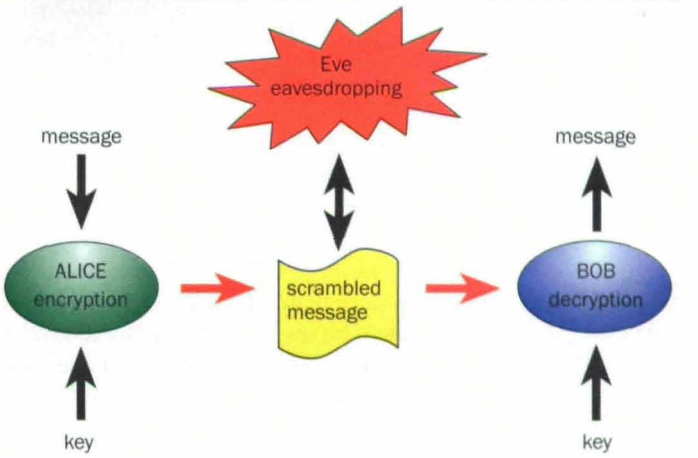

All forms of secret communication follow the following basic principles. Using an algorithm of some sort, the sender, Alice, combines a message with a key to create a scrambled message that she sends to Bob. He decrypts the scrambled message using his key to reveal the real message. Eve is an unauthorized eavesdropper. Quantum versions of this set-up enable Alice and Bob to exchange the key with absolute security and to find out whether or not Eve read their message.

Standard crypto-systems

Cryptography is the art of hiding information in a string of bits that are meaningless to any unauthorized party. To achieve this goal, an algorithm is used to combine a message with some additional information – known as the "key" – to produce a cryptogram. This technique is known as "encryption" (figure 1). The person who encrypts and transmits the message is traditionally known as Alice, while the person who receives it is called Bob. Eve is the unauthorized, malevolent eavesdropper. For a crypto-system to be secure, it should be impossible to unlock the cryptogram without Bob's key. In practice, this demand is often softened so that the system is just extremely difficult to crack. The idea is that the message should remain protected as long as the information it contains is valuable.

Crypto-systems come in two main classes – depending on whether the key is shared in secret or in public. The "one-time pad" system, which was proposed by Gilbert Vernam at AT&T in 1935, involves sharing a secret key and is the only crypto-system that provides proven, perfect secrecy. In this scheme, Alice encrypts a message using a randomly generated key and then simply adds each bit of the message to the corresponding bit of the key (figure 2). The scrambled text is then sent to Bob, who decrypts the message by subtracting the same key. Because the bits of the scrambled text are as random as those of the key, they do not contain any information.

Although perfectly secure, the problem with this system is that it is essential for Alice and Bob to share a common secret key, which must be at least as long as the message itself. They can also only use the key for a single encryption – hence the name "one-time pad". (If they used the key more than once, Eve could record all of the scrambled messages and start to build up a picture of the key.) Furthermore, the key has to be transmitted by some trusted means, such as a courier, or through a personal meeting between Alice and Bob. This procedure can be complex and expensive, and may even amount to a loophole in the system. (It is interesting to note that if Eve wanted to crack the one-time pad by trying out all possible keys one by one, she would obtain a message for each key and would then have to search through all of them. But she would have absolutely no way of knowing which was the right one!)

The other class of crypto-systems shares a public key. The first "public key crypto-systems" were proposed in 1976

by Whitfield Diffie and Martin Hellman, who were then at Stanford University in the US. These systems are based on so-called one-way functions, in which it is easy to compute the function $f(x)$ given the variable x, but difficult to go in the opposite direction and compute x from $f(x)$. In this context, the word "difficult" means that the time to do a task grows exponentially with the number of bits in the input. Factoring large integers is a candidate for such a one-way function. For example, it only takes a few seconds to work out that 107×53 is 5671, but it takes much longer to find the prime factors of 5671.

However, some of these one-way functions have a "trapdoor", which means that there is in fact an easy way of doing the computation in the difficult direction, provided that you have some additional information. For example, if you were told that 107 was one of the prime factors of 5671, the calculation would be relatively easy.

2 Perfect security

Alice		
Message		10110101
Add key	+	01101001
Scrambled text	=	11011100
	↓	
Transmit		
	↓	
Bob		
Scrambled text		11011100
Subtract key	−	01101001
Message	=	10110101

The "one-time pad" crypto-system allows messages to be sent with perfect security. Alice chooses a random number for the key and encrypts her message by adding the key to her message. She then transmits the scrambled message to Bob, who decrypts it by subtracting the key to reveal the real message. In this example, all of the calculations are performed in modulo 2 (i.e. $1+1=0+0=0$; $1+0=0+1=1$; $1-1=0-0=0$; and $1-0=0-1=1$). The problem with this system is that both the sender and recipient have to share the key.

For Alice to transmit a message with a public-key crypto-system, Bob first chooses a private key. He uses this key to compute a public key, which he discloses publicly. Alice then uses this public key to encrypt her message. She transmits the encrypted message to Bob, who decrypts it with his private key. The encryption–decryption process can be described mathematically as a one-way function with a trapdoor – namely, the private key. One therefore only needs to know this key to obtain the original message. In other words, if Bob knows what the trapdoor is, he can do the reverse calculation and reveal the message from the encrypted text.

Public-key crypto-systems are convenient and they have become very popular over the last 20 years. The security of the Internet, for example, is partially based on such systems. The most common example is the RSA crypto-system, which was developed by Ronald Rivest, Adi Shamir and Leonard Adleman of the Massachusetts Institute of Technology in 1977. Its secrecy is actually based on the fact that (as far as we know) the time needed to calculate the prime factors of an integer – and hence to work out the private key – increases exponentially with the number of input bits.

However, this system suffers from two potential major flaws. First, nobody knows for sure if factorization is actually as difficult as we currently think. Of course, one could easily improve the safety of the RSA by choosing a longer key, but if an algorithm were found that could factorize numbers quickly, it would immediately annihilate the security of the RSA system. Although such an algorithm has not yet been discovered – or if it has, it has not been published! – there is no guarantee that such an algorithm does not exist.

The second drawback to the RSA system is that problems that are difficult for a classical computer could become easy for a quantum computer (see box). With the recent developments in the theory of quantum computation, there are reasons to believe that it will eventually become possible to build these machines. If either of these possibilities were fulfilled, the RSA system would become obsolete. Meanwhile, other public-key crypto-systems also rely on unproven assumptions for their security, which could themselves be weakened or suppressed by theoretical or practical advances. One would then have no choice but to turn to secret-key crypto-systems.

Quantum cryptography on paper

The principles of cryptography that we have so far described have all been entirely general. Vernam's system, however, requires Bob and Alice to share a secret key, and it is here that quantum physics enters the scene. Quantum cryptography allows two physically separated parties to create a random secret key without resorting to the services of a courier. It also allows them to verify that the key has not been intercepted. ("Quantum key distribution" is therefore really a better name for quantum cryptography.) When used with Vernam's one-time pad scheme, the key allows the message to be transmitted with proven and absolute security. Quantum cryptography is not therefore a totally new crypto-system. But it does allow a key to be securely distributed and is consequently a natural complement to Vernam's cipher.

To understand how quantum cryptography works, consider the "BB84" communication protocol, which was introduced in 1984 by Charles Bennett of IBM in Yorktown Heights, US, and Gilles Brassard from the University of Montreal in Canada (figure 3). Alice and Bob are connected by a quantum channel and a classical public channel. If single photons are being used to carry the information, the quantum channel is usually an optical fibre. The public channel, however, can be any communication link, such as a phone line or an Internet connection. In practice, the public link is usually also an optical fibre, with both channels differing only in the intensity of the light pulses that code the bits: one photon per bit for the quantum channel, hundreds of photons per bit for the classical public channel. So how does it work?

First, Alice has four polarizers, which can transmit single photons that are linearly polarized either vertically, horizontally, at +45° or at –45°. She sends a series of photons down the quantum channel, having chosen at random one of the polarization states for each photon. She also records her choice.

Second, Bob has two analysers. One analyser allows him to distinguish between horizontally and vertically polarized photons. The other allows him to distinguish between photons polarized at +45° and –45°. Bob selects one analyser at random, and uses it to record each photon. He writes down which analyser he used and what it recorded. Note that every time Bob uses an analyser that is not compatible with Alice's choice of polarization, he will not be able to get any information about the state of the photon. For example, if Alice sent a vertically polarized photon and Bob chose the analyser designed to detect photons at ±45°, there is a 50% chance that he will find the photon in either the +45° channel or the –45° channel. And even if he finds out later that he chose the wrong analyser, he will have no way of finding out which polarization state Alice sent.

Third, after exchanging enough photons, Bob announces on the public channel the sequence of analysers he used, but

3 Encryption with polarized light

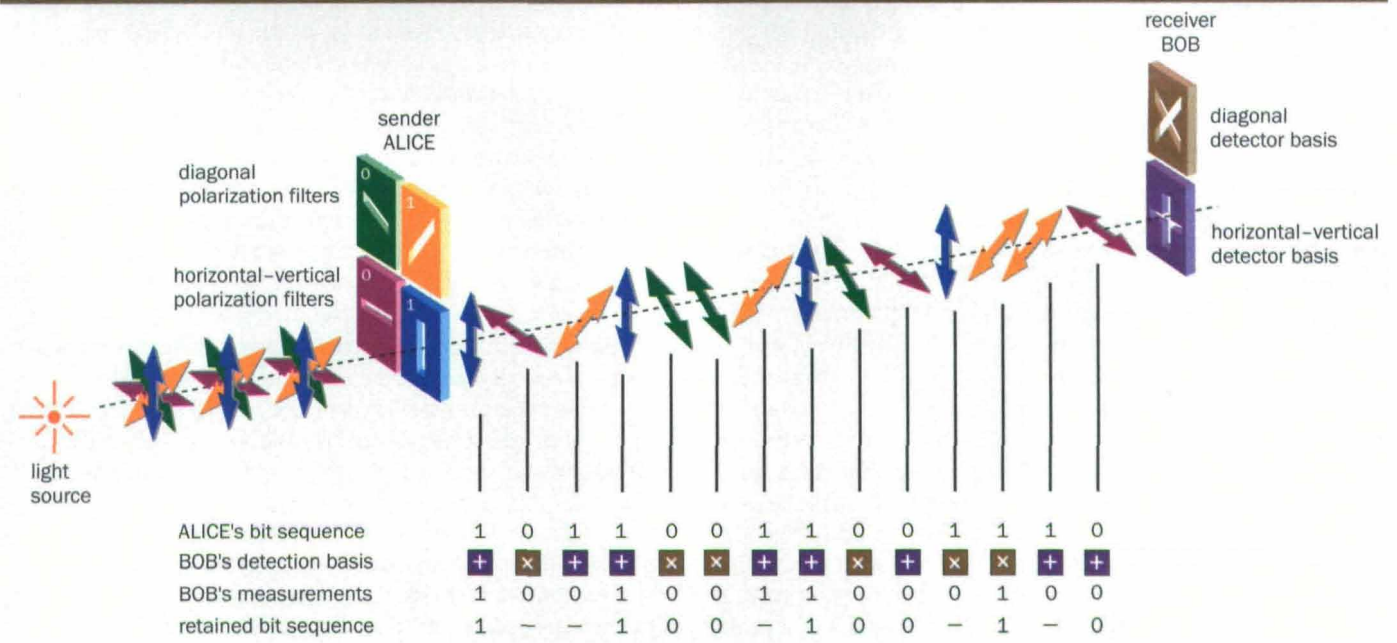

Quantum cryptography relies on creating keys that can be used with absolute secrecy to encrypt and decrypt a message. In the "BB84" communication protocol, Alice has four filters that can linearly polarize photons either vertically, horizontally, at +45° or at –45°. For each photon she sends down an optical fibre, she chooses one of these filters at random (row 1). Bob has an analyser that can distinguish between horizontally and vertically polarized photons, and another one that can distinguish between those polarized at ±45°. Every time he expects a photon to arrive, he chooses one analyser at random (row 2). He records whether or not he detects a signal, which analyser he used and which detector registered the count (row 3). He then tells Alice which photons he detected, the sequence of analysers he used, but not the results he obtained. Alice looks at her data and tells Bob when his analyser was compatible with the polarization of the photon she sent. If the analyser was incompatible, or if Bob did not detect the photon, the bit is discarded. For the bits that remain (row 4), Alice and Bob know for sure that they have the same values and the retained bits can now be used to generate a secret key.

not the results that he obtained.

Fourth, Alice compares this sequence with the list of bits that she originally sent, and tells Bob on the public channel on which occasions his analyser was compatible with the photon's polarization. She does not, however, tell him which polarization states she sent. If Bob used an analyser that was not compatible with Alice's photon, the bit is simply discarded. For the bits that remain, Alice and Bob know that they have the same values – provided that an eavesdropper did not perturb the transmission. They can now use these bits to generate a key, and send encrypted messages to one another.

To assess the secrecy of their communication, Alice and Bob select a random part of their key, and compare it over the public channel. Obviously, the disclosed bits cannot then be used for encryption any more. If their key had been intercepted by an eavesdropper, the correlation between the values of their bits will have been reduced. For example, if Eve has the same equipment as Bob and cuts the fibre and measures the signal, she will always get a random bit whenever she chooses the wrong analyser, i.e. in 50% of cases. But having intercepted the signal, Eve still has to send a photon to Bob, to cover her tracks. Therefore in half of the cases in which Alice's and Bob's analysers match, Eve will have sent a photon that is incorrectly polarized. However, in half of these cases, the photon will accidentally leave Bob's analyser through the correct channel – in which case, Eve's presence goes undetected. The point is that if Eve had been listening in, one in four of Alice's and Bob's bit values would disagree. In other words, her eavesdropping strategy could be easily detected.

There are other eavesdropping strategies that produce a lower disagreement rate. But since all measurements perturb either the vertical–horizontal polarization states or the diagonal states, or all four states, all eavesdropping strategies perturb the system to some extent. Hence, if Alice and Bob do not notice any discrepancies in the subset of their keys, they can be sure that their key has not been intercepted by Eve. They can then use their key with total confidence to encrypt a message.

Quantum cryptography in the real world

So how do people achieve quantum cryptography in practice? Photons are the best candidates for carrying the different quantum states. They are relatively easy to produce and can be transmitted using existing optical fibres. Over the last 25 years, the attenuation of light at a wavelength of 1300 nm has been reduced from several decibels per metre of fibre to just 0.35 decibels per kilometre. This means that photons can travel up to 10 km in a fibre before half of them are absorbed, which is sufficient to perform quantum cryptography in local networks. (Amplifiers cannot be used to transmit the photons further, because quantum states cannot be copied.) Although most quantum-key distribution prototypes use optical fibres, there are some projects aiming to establish quantum communication from a satellite down to earth or to another satellite.

As always happens in physics, however, there is a gap between theory and experiment. In practice, there will always be some errors in the transmission, usually up to a few per cent. The number of errors that are transmitted as a fraction of the total number of detected bits is called the quantum-bit error rate and is one of the parameters that characterizes how well a quantum-cryptography system works.

Uncorrelated bits may originate from several experimental imperfections. For example, Alice has to ensure that she creates photons that are in exactly the states she chose. If, for instance, a vertical photon is incorrectly polarized at an angle

of 84°, there is a 1% possibility that Bob will find it in the channel for horizontally polarized photons. A similar problem arises for Bob. If his polarizer cannot distinguish perfectly between two orthogonal states, he will detect photons in the wrong channel from time to time. Another difficulty is ensuring that the encoded bits are maintained during transmission. A vertically polarized photon, for example, should still be vertically polarized by the time it reaches Bob. But due to the birefringence of the fibre, the polarization states received by Bob will, in general, be different from those sent by Alice.

Even worse, changes to the mechanical or thermal environment can produce fluctuations on a time-scale of seconds or minutes, which means that the alignment of the two analysers has to be continuously monitored. This is possible in principle, but is not very convenient. In fact, the number of transmission errors – and hence the quantum-bit error rate – is dominated by the noise of the detector. In other words, most errors are not due to photons that have been incorrectly detected. The errors arise when a photon fails to reach a detector as expected and the wrong detector registers a dark count instead. Unfortunately, at the wavelengths where the fibre losses are low (i.e. 1310 nm), relatively noisy, low-efficiency home-made single-photon detectors have to be used.

To overcome these problems, Alice and Bob have to apply a classical error-correction algorithm to their data so that they can reduce the errors below an error rate of 10^{-9} – the industry standard for digital telecommunications. And since they cannot be sure if the presence of uncorrelated bits was due to the poor performance of their set-up or to an eavesdropper, they have to assume the worst-case scenario – namely that all of the errors were caused by Eve. To reduce the amount of information that Eve may have obtained, Alice and Bob therefore use a procedure known as "privacy amplification", in which several bits are combined into one. This procedure ensures that the combined bits only correlate if Alice and Bob's initial bits are the same. But Eve ends up with a totally different series of bits, because she only knows a fraction of the initial bits. The problem with privacy amplification is that it shortens the key length a lot and it is only possible up to certain error, which means that Alice and Bob have to be careful to introduce as few errors as possible when they initially send their quantum bits.

How fast are computers?

The continuous dialogue between basic quantum physics and fascinating potential applications leads to one basic question: are quantum computers really faster than classical ones? The consequences of solving this question will be dramatic whatever the answer. If quantum computers are indeed much faster, it would obviously be worth investing money in this field, although the very concept of information would then have to be changed. Instead of being part of mathematics, information would become part of physics! On the other hand, if classical computers can be as fast as quantum ones, then presumably the best classical algorithms have not yet been found. This finding could destroy all of the major security systems, which our IT-dependent society relies so heavily upon.

One of the fathers of quantum computing, David Deutsch of Oxford University, has recently argued that physics is more fundamental than mathematics, because answers to mathematical questions (like working out to which class of complexity a mathematical problem belongs) depend on physics. This claim has come as a shattering blow to mathematicians, who in an attempt to keep their science as the root of all others, are now trying to prove that classical computers are actually as efficient as quantum computers.

It is amusing to follow these debates that have been provoked by quantum physics, but it is important to realize that progress on these fundamental issues could happen soon, since some excellent theorists have recently joined the field. (Security managers, however, might be having nightmares!) But whatever the outcome of these debates, quantum cryptography and other applications of quantum communication are already proving that quantum mechanics can do useful things that are impossible with classical physics.

Cryptography experiments

Quantum cryptography moved from the realms of theory to experiment in 1989, when researchers at IBM built the first prototype that could securely distribute a key. They coded their message using polarized photons (figure 3), and managed to send it over a distance of 30 cm in air. Since then, the improvements have been immense, and several groups have shown that quantum cryptography works outside the lab as well. (We will only consider those systems that use 1310 nm photons, which could one day be used over long distances.) At Geneva University in 1995, the authors also demonstrated the feasibility of the polarization-encoding scheme with installed Swisscom fibres, and BT (formerly British Telecom) followed in 1997 with a similar system.

Another set-up, which encodes the message using the photons' phase rather than their polarization, was developed in 1993 by Paul Townsend and colleagues at BT. In this scheme, both Alice and Bob use identical unbalanced "Mach–Zehnder interferometers", in which one arm is longer than the other (figure 4). They are used to produce and detect photons with a particular phase shift. This scheme is also being used by Richard Hughes and his group at the Los Alamos National Laboratory in New Mexico.

Pulses that go down the short arm in Alice's interferometer and then the long arm in Bob's interferometer interfere with pulses that take the long first and then the short one. When Alice sends her message, she randomly applies phase shifts of $0, \pi/2, \pi$ or $3\pi/2$ to her photons. Bob, however, only has the option of applying a phase shift of $\pi/2$ or none at all. If Bob applies no phase shift, he can work out whether Alice's photon has a phase shift of 0 or π. On the other hand, if Bob applies a phase shift of $\pi/2$, he can distinguish between Alice's choice of $\pi/2$ and $3\pi/2$. After the message has been sent, Alice and Bob compare their settings using the public channel. If they chose compatible settings, Bob knows which phase Alice applied. A secret key can therefore be established by interpreting phase shifts of 0 and $\pi/2$ as "1", and π and $3\pi/2$ as "0". Incompatible measurements are disregarded.

As with polarization encoding, this scheme has to be actively controlled. For example, the arms in the two interferometers have to be adjusted so that the differences in the path length are the same. These differences also have to be kept stable. Another problem is that the two pulses at Bob's interferometer interfere perfectly only if they are in the same polarization state, which means that the scheme also requires an active polarization control.

In collaboration with Swisscom, we have recently proposed and tested a new type of interferometer that is self-balanced

4 Phase encryption

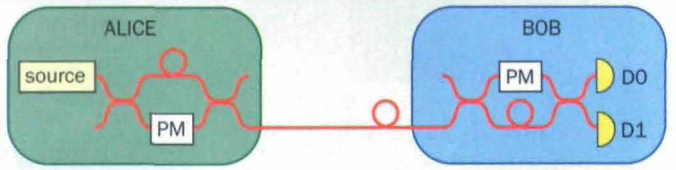

This set-up has been used by researchers at British Telecom and Los Alamos to encode and send a message using a photon's phase rather than polarization. Alice and Bob both have identical unbalanced Mach–Zehnder interferometers, each of which consists of one short arm containing a phase modulator (PM) and one long arm (denoted by the circle). Light entering Alice's interferometer is split in two and passes down the separate arms, before recombining where the arms join up. Alice's phase modulator is used to add a phase shift of either 0, $\pi/2$, π or $3\pi/2$. Bob can either apply a phase shift of 0 or $\pi/2$. Depending on whether the photon is detected by detector D0 or D1, this allows him to distinguish between Alice's phase shift of 0 or π, or between a phase shift of $\pi/2$ or $3\pi/2$.

and in which all birefringence fluctuations are automatically compensated. This set-up uses "time-multiplexed interferometry" – in other words, the pulses that interfere travel along precisely the same paths, but at different times. The advantage is that thermal drifts do not have to be controlled. Moreover, any fluctuations in the polarization of the interfering pulses are wiped out using "Faraday mirrors" at the end of the fibres – instruments that reflect light and transform the state of polarization to the orthogonal one.

There are also prototypes that work at other wavelengths. However, due to higher losses in the fibres, these systems cannot be used to transmit quantum bits any further than a few kilometres. For example, James Franson from Johns Hopkins University in Baltimore demonstrated polarization encoding in 1995 using 830 nm photons. Last year, BT tested a similar system, working at a frequency of 1.2 MHz, which is the highest transmission rate for quantum-key distribution to have so far been achieved.

Cryptography on noisy channels

Although quantum cryptography on noiseless channels has proved to be perfectly secure, noisy channels are much more difficult to handle. The problem with noisy channels is that if Eve intercepts and reads a message, she could then pass on a partially garbled message and get away with it. And if the quantum-bit error rate on her message is lower than the level of noise, Alice and Bob would never suspect anything.

Before they send any messages, Alice and Bob therefore have to evaluate how much information Eve could possibly obtain. They assume that Eve has unlimited technology, and that her eavesdropping strategy is only restricted by the laws of physics. Once Alice and Bob establish an upper limit on the amount of information that Eve knows, they can – provided this limit is not too high – use error-correction and privacy-amplification algorithms to reduce the information that she can get her hands on. Although this approach will produce a final key that is shorter than the raw data, Eve's information about the final key will then be arbitrarily small.

The drawback is that the complete solution to this problem is not yet known. However, if one assumes that Eve can only interact one by one with the quantum bits that Alice sent to Bob, it turns out that Eve will never know as much as Bob, provided that the quantum-bit error rate is less than 15%. Remarkably, this result establishes a connection with the famous "Bell inequality" – an inequality that is satisfied by all local hidden variable theories, but not by quantum mechanics. Eve's information is lower than Bob's if and only if Bob's results cannot be explained by any local hidden variable theory! This point nicely illustrates the fascinating dual nature of quantum information theory. It deals on the one hand with practical issues, such as the security of crypto-systems and fundamental questions about quantum physics – like non-locality – on the other.

The future starts here

Several groups have now shown that quantum cryptography is possible outside the laboratory. The error rates in sending quantum bits are now low enough to guarantee that the key can be securely distributed. Although the systems still suffer from low transmission rates – and messages can only be sent over a few tens of kilometres – they could, even today, provide a means of securely transmitting messages if the public-key systems that are used on the Internet were suddenly cracked. But, above all, quantum cryptography is fun. Not only does it naturally complement standard crypto-systems, it is also an excellent example of the interplay between fundamental and applied research.

Further reading

C A Fuchs 1997 Optimal eavesdropping in quantum cryptography. 1. Information bound and optimal strategy *Phys. Rev.* **A56** 1163–1172

R J Hughes 1995 Quantum cryptography *Contemp. Phys.* **36** 149–163

N D Mermin 1981 Bringing home the atomic world: quantum mysteries for anybody *Am. J. Phys.* **49** 940–943

S J D Phoenix and P D Townsend 1995 Quantum cryptography: how to beat the code breakers using quantum mechanics *Contemp. Phys.* **36** 165–195

J G Rarity June 1994 Dreams of a quiet light *Physics World* July pp46–51

T Spiller 1996 Q information processing: cryptography, computation and teleportation *Proc. IEEE* **84** 1719–46

Wolfgang Tittel, **Grégoire Ribordy** and **Nicolas Gisin** are in GAP-Optique, Université de Genève, 20 rue de l'Ecole de Medicine, CH-1211 Genève, Switzerland