

The battle between cryptographers, who encrypt messages, and cryptanalysts, who break those codes, has raged for centuries. As quantum computing promises to help cryptanalysts break many of the encryption methods used today, quantum cryptography promises to keep our secrets safe forever.





A New Face for Cryptography

Jane E. Nordholt and Richard J. Hughes

Cryptography, the mathematical science of secret communications, has had a long and distinguished history dating back to the time of the ancient Greeks. It is a subject noted for the never-ending struggle for one-upmanship between code makers and code breakers, a struggle in which the future of nations has literally been at stake. The code breakers' need to read another party's secret communications has been a tremendous force driving the development of new information-processing technologies. The code makers have responded by using those new technologies to develop more complex methods for ensuring the security of communications.

The latest round in this struggle seems set to be played out in the world's physics laboratories, with the combatants drawing upon fundamental principles of quantum physics, principles that were only of academic interest until about 15 years ago. The code breakers believe that a large-scale quantum computer—a device that uses the nonclassical aspects of quantum systems to manipulate information—could defeat the most widespread cryptosystems in use today. They are pushing the physics community to develop such a computer,

which necessarily involves controlling atoms and photons in ways that were barely dreamed of—until recently. Meanwhile, the code makers are ready for battle and are already exploiting quantum mechanics in a new code-making technology—quantum key distribution (QKD)—that could counter the quantum computing threat.

Classical Cryptography

The main goal of cryptography is to allow two parties (conventionally referred to as “Alice” and “Bob”) to communicate while simultaneously preventing a third party (“Eve”) from understanding those communications. Alice and Bob's messages should remain secret even when Eve is able to passively monitor the exchanges. (A more intrusive Eve might want to prevent Alice and Bob from communicating at all, but such a denial-of-service attack is a different type of communication problem that we will not consider here.) Cryptography provides Alice with the means to render her messages to Bob in a form that is indistinguishable from random noise but that, nevertheless, allows Bob to recover the original message.


Figure 1. A Symmetric-Key Cryptography System: The One-Time Pad

(a) Alice, the sender, first generates a string of binary bits (the key) that is as long as her binary message. Then she applies the XOR operation—bit by bit—to the key and her message, and sends the encrypted string to Bob over an open communications channel. (b) Bob, the receiver, uses the same key as Alice to decrypt the message by the same XOR operation, applied bit by bit. His decrypted message is identical to the original message sent by Alice. Because the value of each key bit is random, the message cannot be recovered without the key. As long as Alice and Bob use the key only once to encrypt and decrypt one message, this one-time pad system is absolutely secure, but distributing the secret keys remains a problem. (c)–(e) This series of photographs shows an aerial view of the St. Louis International Airport before encryption, as encrypted by Alice, and as decrypted by Bob. Whereas Alice’s encrypted photo is indistinguishable from random noise, Bob is able to reproduce the original faithfully.

(a) Encryption, One-Time Pad




(b) Decryption



(c) Original



(d) Encrypted




(e) Recovered Original



XOR operation, \oplus :
 $0 \oplus 0 = 0$: $0 \oplus 1 = 1$: $1 \oplus 0 = 1$: $1 \oplus 1 = 0$

| | | | | |
|-------------------|------|------|------|------|
| Alice's message | 1001 | 0000 | 0110 | 1001 |
| \oplus | | | | |
| Key | 1000 | 0100 | 0101 | 0001 |
| Encrypted message | 0001 | 0100 | 0011 | 1000 |



| | | | | |
|-------------------|------|------|------|------|
| Encrypted message | 0001 | 0100 | 0011 | 1000 |
| \oplus | | | | |
| Key | 1000 | 0100 | 0101 | 0001 |
| Original message | 1001 | 0000 | 0110 | 1001 |

This process of encryption (by Alice) and decryption (by Bob) can be accomplished if the two parties share a string of randomly generated binary bits known as a cryptographic key. In a system called the “one-time pad,” Alice and Bob must have identical copies of the key. (How they get the key will be discussed later). As seen in Figure 1, Alice adds the key to her message, bit by bit, using the binary exclusive OR- (XOR-, \oplus) operation, which is equivalent to addition modulo 2. Mathematically, the XOR operation is defined as

$$\begin{aligned}
 0 \oplus 0 &= 0 \quad , \\
 0 \oplus 1 &= 1 \quad , \\
 1 \oplus 0 &= 1 \quad , \text{ and} \\
 1 \oplus 1 &= 0 \quad .
 \end{aligned}
 \tag{1}$$

Alice’s encrypted communication at this point is indistinguishable from

random noise. Alice sends this message to Bob, who takes his copy of the key and subtracts it from the message, again using an XOR-operation. The original script is recovered. Provided a key is used to encipher only one message, the one-time pad encryption process is provably secure. In fact, it is the only completely secure cryptographic system.

The one-time pad is an example of a symmetric-key system (symmetric because Alice and Bob have the same key), and it requires a key that is as long as the message. In another type of symmetric key system, Alice and Bob use a short key to seed a high-quality random number generator of which they have identical copies. They then need to share fewer initial key bits in order to encrypt and decrypt large messages. In the Data Encryption Standard (DES)—a sym-

metric-key algorithm that was adopted as a United States government standard in 1977—the key length is 56 bits.

The security of all symmetric-key cryptographic systems rests entirely on the secrecy of the shared key because the structure of the cryptographic algorithm used by Alice and Bob is public knowledge. Certainly, the eavesdropper Eve understands and can implement the decryption algorithm. Should Eve obtain the key, she could immediately read Alice and Bob’s messages. Without the key, Eve must attempt a mathematical attack on the encrypted message (or parts thereof) in order to crack it. In a properly designed symmetric-key cryptosystem, no attack should be more efficient than an exhaustive search over all possible keys.

Consider, for example, the 56-bit

DES key. Because there is a choice of either 0 or 1 for every bit in a binary key, there are 2^{56} (or nearly 10^{18}) possible DES keys. A desktop computer testing a million keys a second would require more than two thousand years to search the entire key space. But the phenomenal increase in computational speed and capability has made the 56-bit key vulnerable. Today's supercomputers can search all possible keys in a matter of hours.

The simple solution is to use longer keys. Adding a bit to the key length doubles the search time, whereas doubling the key length makes the search problem exponentially harder. In the forthcoming Advanced Encryption Standard (AES), the key length will be up to 256 bits, in which case a search of the entire key space would be so computationally demanding that it would not be feasible on any computer system within the useful lifetime of the information.

The Key Distribution Problem.

A DES-type cryptographic system reduces the act of communicating a long secret (the message) to that of creating and sending a short secret (the key). But the central issue within any system is that any information about the key must remain out of the hands of unwanted parties. This latter requirement creates what is known as the key-distribution problem.

Traditionally, cryptographic keys were distributed by trusted couriers immortalized in spy movies as strangers in trench coats handcuffed to locked briefcases. But the infrastructure required to manage the key material makes this type of distribution impractical in our computer-driven, global community. Picture the logistics nightmare if a courier had to deliver a cryptographic key every time Alice wanted to use her credit card over the Internet—and imagine the added cost! In some cases, courier key

distribution is even impossible, such as when Bob is not a person but a satellite in Earth's orbit. Furthermore, the existence of the key material before delivery by courier introduces an insider threat, in that the key material could be copied and delivered surreptitiously to Eve.

About 30 years ago, researchers at Britain's Government Communications Headquarters (GCHQ), and later (independently) in the United States, found a new, more convenient way to securely distribute cryptographic keys. The system is known generically as public-key cryptography. One public-key protocol begins when Bob generates two very large prime numbers, p and q , which are multiplied to form the especially large number N . He then selects an integer g , and uses the numbers p , q , and g to generate a fourth number, d . The two numbers (N, g) constitute Bob's public key, which he makes widely available. The number d constitutes Bob's private key, which he keeps secret. (The protocol is discussed in greater detail in the box "Public-Key Cryptography: RSA" on the [next page](#).)

When Alice wants to send an encrypted message to Bob, she grabs a copy of his public key and uses it in an algorithm that mathematically scrambles her communication. The algorithm, however, is a clever one-way operation: Bob's public key (N, g) cannot be used to unscramble Alice's encrypted message. Instead, one needs the secret number d from Bob's private key to decrypt. Given only N , it is extremely difficult to find the prime factors p and q that are needed to generate d ; hence, the system is considered secure.

Because the public-key cryptography system is asymmetric—only Bob needs to have a secret key—it has become the enabling technology for electronic commerce. Alice can grab the public key from the Bob.com website and safely encrypt and send

her credit card number. In addition, public-key encryption also provides a means for Alice to authenticate her transaction.

But public-key cryptography has its downside. Because of the computational difficulty in calculating asymmetric keys, Alice and Bob use it only to produce and distribute a symmetric key that they then use for the bulk of their discussions. More disturbing is the lack of proof that the methodology is secure. A clever person could come up with a new factoring algorithm that allows finding the secret number d , thus making public-key cryptography obsolete.

In 1994, Peter Shor of AT&T did invent such an algorithm. If implemented, that algorithm would undermine the public-key cryptography in use today. Fortunately, Shor's algorithm must be run on a quantum computer, which is currently unavailable and will probably remain so for many years.

Public-key cryptography clearly has a place where security need not be guaranteed to last for years. Because it is not provably secure, however, and because a quantum computer may render it useless in the future, a better system is needed for highly valuable data such as government or trade secrets. That better system is quantum cryptography.

Quantum Cryptography

Quantum cryptography is a type of symmetric-key distribution that allows Alice and Bob to create and share a secret key, while Eve is prevented from obtaining any more than a tiny fraction of one bit of information about the final key's binary sequence. The secret key can actually be used in any symmetric encryption method desired. Because quantum cryptography is used to send these key bits, it is more correctly called

Public-Key Cryptography: RSA

Public-key cryptography is an asymmetric key-distribution system, wherein Bob generates two keys: a public key, which he makes available to anyone, and a private key, which he keeps secret. Alice uses the public key to encrypt her message, which she then sends to Bob, who uses his private key to decrypt that message. Perhaps the most widely used public-key cryptography algorithm is RSA, which was invented in 1978 by Ron Rivest, Adi Shamir, and Leonard Adleman and was named for its inventors. The RSA algorithm uses two keys that are constructed as follows:

- Bob generates two prime numbers, p and q , which are typically very large (several hundred bits in length).
- He calculates the product, $N = pq$, known as the modulus.
- He calculates Euler's quotient function $\Phi(N)$, which is simply the number of integers less than N that are coprime* to N . If p is a prime number, every number less than p is coprime to it, so $\Phi(p) = p - 1$. Since the modulus $N = pq$ is the product of prime numbers, $\Phi(N) = (p - 1)(q - 1)$. Let $\Phi(N)$ be designated by η .
- Bob chooses an integer g such that $g < N$, and g has no factors in common with η .
- Bob calculates $d = g^{\Phi(\eta)-1} \bmod \eta$, where $\bmod \eta$ is the modulus operation.†

Bob's public key is (N, g) . His private key is the number d .

* Two integers are coprime if they share no common divisors except 1.
 † For an introduction to modular arithmetic, see the article "From Factoring to Phase Estimation" on [page 38](#).

When Alice wants to send a message to Bob, she first represents her message as a series of numbers. To encrypt, she grabs Bob's public key (N, g) and uses it in the following mathematical transformation:

$$c = m^g \bmod N, \tag{1}$$

where m is a number representing a piece of her message. She sends the new number c off to Bob, who uses his private key (N, d) to perform the operation

$$m = c^d \bmod N, \tag{2}$$

thereby recovering Alice's number.

Public-key cryptography is based on a theorem by Euler, which states that $x^{\Phi(y)} = 1 \bmod y$, for any integer x that is coprime to the number y . The number d was chosen such that $d = g^{\Phi(\eta)-1} \bmod \eta$, or $dg = g^{\Phi(\eta)} \bmod \eta$, which by Euler's theorem becomes $dg = 1 \bmod \eta$. Subtracting 1 will result in $dg - 1 = 0 \bmod \eta$.

The last statement indicates that the number $dg - 1$ is evenly divisible by η , so that $dg - 1 = k\eta$, where k is an integer. In decrypting the message, Bob has

$$\begin{aligned} c^d \bmod N &= (m^g)^d \bmod N, \\ &= m (m^{dg-1} \bmod N), \text{ and} \\ &= m (m^{k\eta} \bmod N). \end{aligned} \tag{3}$$

But $\eta = \Phi(N)$. By Euler's theorem, $m^{\Phi(N)} = 1 \bmod N$. Thus,

$$\begin{aligned} c^d \bmod N &= m (1)^k \bmod N, \text{ and} \\ &= m \bmod N. \end{aligned} \tag{4}$$

In other words, $c^d \bmod N = m$, so that the decryption algorithm recovers Alice's message.

quantum key distribution (QKD). Adding to the security of a QKD system is the fact that any attempt to steal or copy a key can be detected, thus revealing information about the security environment.

The quantum part of quantum cryptography comes from the transmission and reception of single photons. In addition to keeping an eavesdropper at bay (primarily

because a photon cannot be split or copied reliably), quantum cryptographic systems exhibit strange quantum mechanical behaviors that are not normally observed in the classical world of everyday experience. The best example of such behavior occurs in our fiber-based quantum cryptographic system, in which we use the interference of single photons with themselves

to transmit information.

Before describing how a photon interfering with itself helps us encrypt messages, we will present an overview of the steps involved in executing a secure exchange of messages and then describe a simple protocol. Protocols are the rules used for the quantum mechanical and conventional transmissions at the heart of QKD.

A QKD Session. To perform QKD, Alice and Bob communicate in two different ways. The first is over a quantum channel, which allows Alice to reliably send single photons to Bob. While Eve may attempt to breach the quantum channel, her tampering can be detected. The second means of communication is an ordinary, public channel assumed to be monitored by Eve. Alice and Bob use this open channel to construct their secret key, implement any of several error-correction techniques, and coordinate a “privacy amplification” scheme that effectively prevents Eve from gaining any knowledge about the final key. In all, six steps are implemented in a QKD session. These are summarized in the box to the right.

As a first step, Alice and Bob authenticate their communications; that is, they verify each other’s identity. If this step is ignored, Eve can perform a “man-in-the-middle” attack and convince Alice that she is Bob, and Bob that she is Alice, in which case no form of key distribution or encryption can prevent Eve from reading all of Alice and Bob’s communications.

After authentication, Alice and Bob begin their QKD session. First, each generates a random bit stream. Alice then uses a QKD protocol, such as BB84 (discussed in the next section), that specifies how she is to encode each bit as the quantum state of a single particle. For example, she may use the specific polarization state of a single photon to encode for either a 0 or a 1. Then, Alice would send a stream of polarized photons to Bob, who follows the protocol in determining how to measure the polarization and hence deduce a bit sequence. Because of the way the protocol works, Alice and Bob can have a public conversation and select an overlapping subset of bits without revealing to each other the value of those bits.

Six Steps to a QKD Session

Authenticate. Over an open communication line, Alice confirms she is talking to Bob, and Bob confirms he is talking to Alice.

Use a quantum protocol. The protocol dictates how Alice is to encode her random bit stream as a quantum state of a single photon. Bob measures photons according to the protocol.

Construct the sifted key. Alice and Bob use an open line to discover which photons were sent and measured in the same basis. The bit values associated with that subset of photons form the sifted key.

Construct the reconciled key. Over the open line, Alice and Bob find and remove errors from the sifted key to make the reconciled key.

Construct the secret key. Alice and Bob use privacy amplification to construct a secret key from the reconciled key. An eavesdropper has essentially no information about the bits in the secret key.

Save some bits. A few secret bits are retained to enable authenticating future QKD sessions.

For example, if Alice’s random sequence is 0111 1010 1001 and as a result of his measurements Bob obtains the sequence 1001 1100 0100, then the protocol provides a means for Alice and Bob to know—without specifically telling each other—that the fourth, fifth, eighth, and eleventh bits form a common subsequence of 1100. This subsequence is called the “sifted” key.

In the real world, hardware is noisy, and transmission media are lossy, so the sifted key will contain some errors. Alice and Bob continue their public conversation and create a “reconciled” key, in which those errors are removed. During this process, some information about the sifted key becomes available to any potential listener (Eve). But Alice and Bob can calculate the maximum information Eve could have about their reconciled key, and using privacy amplification, reduce Eve’s information to substantially less than one bit. The result is a secret key known only to Alice and Bob. The one remaining step before closing the

session is to save a few key bits and thereby have a means to authenticate the next QKD session.

The BB84 Protocol. In 1984, Charles Bennett and Gilles Brassard published a paper describing how orthogonal and nonorthogonal quantum states could be used to construct a cryptographic key. Known today as BB84, the protocol is at the heart of our experimentally realized QKD systems. In the free-space version, Alice encodes random bit values in the polarization states of photons and then sends the single photons to Bob over the quantum channel. Bob’s measurement of the photon’s polarization and subsequent communication with Alice over a public channel allow the two parties to construct a sifted key.

A stylized version of the BB84 protocol is shown in Figure 2. (The box “Photons, Polarizers, and Projections” on [page 76](#) also provides some background material for this section.) Alice generates a random sequence of bits and then chooses—also at random—between one of two




Figure 2. The BB84 Protocol

The BB84 protocol works because Alice randomly chooses to encode the photons in two, nonorthogonal bases. (a) An attenuated laser produces close to single photons. (b) Alice uses a random number generator (RNG) to select a bit value: 0s are encoded as horizontally polarized photons and 1s as vertically polarized photons (c) A second RNG selects the basis. To choose the H/V basis, Alice does nothing, (the photons are already either $|H\rangle$ or $|V\rangle$). To choose the diagonal ($-45^\circ/+45^\circ$) basis, she inserts a half-wave plate (HWP) that rotates the polarization by -45° , so that $|H\rangle$ goes to $| -45^\circ\rangle$ and $|V\rangle$ to $| +45^\circ\rangle$. (d) Bob uses an RNG to select his measurement basis, choos-

ing either to do nothing (H/V) or to rotate the photon by $+45^\circ$ ($-45^\circ/+45^\circ$). He detects photons using an H/V oriented polarizing beam splitter (PBS), which transmits horizontally polarized photons but reflects vertically polarized ones (see inset). Photons polarized at $\pm 45^\circ$ have an equal probability to go to either detector. Table I shows that, when Alice and Bob choose the same basis, they know that their bit values coincide. When they choose different bases, their bit values are randomly correlated. At the end of the session, Bob and Alice openly compare their bases for each measurement. They keep only those bits that were sent and measured in the same basis.

Table I. Details of the BB84 Protocol

| Sender (Alice) | | | Receiver (Bob) | | | | | Joint Action |
|----------------|-----|--------------|----------------|------------------------|-----------------|--------|--------|--------------|
| Alice's Basis | Bit | Polarization | Bob's Basis | Resulting Polarization | Probability (%) | | Bit | |
| | | | | | H-Det. | V-Det. | | |
| H/V | 0 | H | H/V | H | 100 | 0 | 0 | Keep bit |
| H/V | 1 | V | H/V | V | 0 | 100 | 1 | Keep bit |
| H/V | 0 | H | Diag. | $+45^\circ$ | 50 | 50 | 0 or 1 | Discard bit |
| H/V | 1 | V | Diag. | -45° | 50 | 50 | 0 or 1 | Discard bit |
| Diag. | 0 | -45° | H/V | -45° | 50 | 50 | 0 or 1 | Discard bit |
| Diag. | 1 | $+45^\circ$ | H/V | $+45^\circ$ | 50 | 50 | 0 or 1 | Discard bit |
| Diag. | 0 | -45° | Diag. | H | 100 | 0 | 0 | Keep bit |
| Diag. | 1 | $+45^\circ$ | Diag. | V | 0 | 100 | 1 | Keep bit |

polarization bases, either the horizontal/vertical (H/V) basis, or the diagonal ($-45^\circ/+45^\circ$) basis. If she chooses the H/V basis, the bit values of 0 are encoded as horizontally polarized photons, and bit values of 1 are encoded as vertically polarized photons, that is, $0 = |H\rangle$ and $1 = |V\rangle$. Similarly, if she chooses the diagonal basis, 0 and 1 bit values are encoded as $0 = |-45\rangle$ and $1 = |+45\rangle$. She sends the stream of polarized photons off to Bob.

At his end, Bob chooses at random to measure polarizations in either the H/V or diagonal basis. As shown in Figure 2, he uses a special dual-detector system. If he chooses the H/V basis, then photons in the state $|H\rangle$ go through to his H-detector, while those in the state $|V\rangle$ are reflected to the V-detector. Photons in the $|-45\rangle$ or $|+45\rangle$ state go randomly to either detector. If Bob measures in the diagonal basis, then his setup directs $|-45\rangle$ photons to the H-detector, $|+45\rangle$ photons to the V-detector, and $|H\rangle$ or $|V\rangle$ photons to either detector with equal probability.

Table I shows how the results differ depending on which polarization states were sent and how they were detected. When Alice and Bob used the same basis, a photon hit on Bob's H-detector means that Alice had a bit value of 0; a hit on his V-detector, that she had a bit value of 1. If the bases differ, there is no such correspondence. Bob and Alice therefore use the public channel and simply compare the sequence of bases. They keep the corresponding bits when the bases agree and disregard the bits when they don't agree. In this way, they can build a sifted-key sequence over a public channel without ever revealing the value of the individual key bits.

Because Alice and Bob have a 50 percent chance of choosing the same basis, in an ideal implementation of BB84, half of the photons are used to create the sifted key. In practice, the efficiency is much less

because the real world unavoidably introduces errors into the sifted-key sequence—polarizers are not perfect, photons do not always reach Bob, and detectors do not always fire when hit with a photon and sometimes fire on their own. Alice and Bob must check and correct their sequence for errors.

Error Correction. One example of a simple error-correction scheme is illustrated in Figure 3. Alice tells Bob the parity of each of her bytes, that is, whether the sum of each 8 bits of the sifted key is even or odd. Bob then checks the parity of his bytes. They keep those bytes that have the same parity and initiate a 20-questions-type deductive process to find the problem bit when the parity differs.¹ Because parity checks can only find an odd number of errors in a bit sequence, in practice, sifted bits are shuffled and then checked for errors several times. All errors must be eliminated to a high degree of certainty. If Alice and Bob's keys differ by even a single bit, the keys will be unusable.

Alice and Bob make their byte comparisons over the open channel, so Eve now has—at a minimum—information about the parity of each retained byte. To eliminate even this limited knowledge on Eve's part, Alice and Bob can agree to drop the last bit of each byte. In addition, they have to sacrifice some key bits to find the errors in their sequences. The reconciled key is therefore shorter than the sifted key. While undertaking the error correction process, however, Alice and Bob obtain an estimate of the bit error rate (BER), which is the number of errors they had in their sifted sequences. Alice and Bob use the BER and knowledge of the quantum mechanical and physical principles of the QKD technique to put a rigorous upper bound on the possible information that Eve may have about their bit sequences.

Privacy Amplification. In this step, Alice and Bob do an XOR operation on sequences of bits from the reconciled key to produce fewer, but brand new, bits. The amount of compression required depends on their estimate of Eve's acquired knowledge.

For example, suppose Alice and Bob share a reconciled sequence consisting of six bits, a, b, c, d, e, and f, and they suspect that Eve knows three of the six bits. Alice and Bob make two new bits out of the original six by doing the following operation:

$$\begin{aligned} a \oplus b \oplus c \oplus d &= \text{Bit 1} \quad , \text{ and} \\ c \oplus d \oplus e \oplus f &= \text{Bit 2} \quad . \end{aligned} \quad (2)$$

Although Eve may have known three bits of the reconciled key sequence, she knows nothing about the new bits generated by privacy application. Alice and Bob can apply this procedure to reduce Eve's knowledge to less than one bit in a key that is several hundred bits long and thereby produce a completely secure key. In general, if the original sequence is n -bits long, privacy amplification will compress it to $R(n)$ bits, where

$$R(n) = -n \log_2[\zeta^2 + (1 - \zeta)^2] \quad (3)$$

and ζ is the BER.

Foiling Eve. We are now in a better position to discuss how the complete QKD session prevents Eve from gaining information about the secret key. First, Eve cannot get any information about the key over the open channel; although the BB84 protocol allows her to know which bits Alice and Bob had in common, she knows nothing about the values

¹ Bits that get transmitted correctly are valuable. Although Alice and Bob could drop all eight bits of a problem byte, it is usually worthwhile to winnow through the byte and retain as many bits as possible.

Photons, Polarizers, and Projection

Our realization of the BB84 protocol uses the polarization state of individual photons to encode bit values. But the key feature that prevents an eavesdropper from detecting the polarizations without being noticed is the use of two nonorthogonal linear polarizations to represent 0 and 1. Rather than preparing a random sequence of horizontally or vertically polarized photons in the quantum states $|H\rangle$ or $|V\rangle$, respectively, Alice (the sender) polarizes photons in the quantum states $|H\rangle$ or $| -45\rangle$ when she wants to send a 0 to Bob (the receiver) and $|V\rangle$ or $|+45\rangle$ when she wants to send a 1.

We can do a simple experiment to demonstrate the quantum mechanical properties of nonorthogonal photons. We need just 3 sheets of linearly polarizing filters, which are readily available from scientific education kits or suppliers. The filter is made from a material that has an intrinsic transmission axis for photons (the polarization axis). As shown in Figure A, if randomly polarized light (for example, sunlight), made up of a large number of photons goes through a linear polarizer with its axis aligned, say, horizontally, the photons that emerge are polarized in the state $|H\rangle$.

We perform our experiment by orienting the first polarizer filter horizontally and holding it up to sunlight. The light intensity decreases by about 50 percent, which indicates that about half the photons get through. We then place a second polarizer behind the first and rotate it until no light passes. At that point, the polarization axes of the two filters are orthogonal to each other, that is, the polarization axis of the second polarizer is in the vertical direction. If we place the third filter between the first two with its polarization axis at -45° to the others, we naively expect no change in the light transmission, but suddenly one eighth of the sunlight gets through the stack, even though the axes of the outer two polarizers are still perpendicular.

These spooky results are a direct consequence of the quantum properties of single photons. A linearly polarized photon is described by a quantum mechanical wave function. Mathematically, it is represented by a “ket” $|\psi\rangle$, which is analogous to an ordinary unit vector in 2-dimensions. Just as a plane vector can be written in terms of two orthogonal plane vectors, we can express $|\psi\rangle$ as a superposition of two orthogonal kets, $|\phi\rangle$ and $|\phi+90\rangle$, in a two-dimensional Hilbert space, with real (as opposed to complex) coefficients. The ket $|\phi\rangle$ represents a photon linearly polarized at the angle ϕ to the horizontal, while $|\phi+90\rangle$ represents a photon polarized at the angle $(\phi + 90^\circ)$. The orthogonal kets are a basis for the Hilbert space. We have

$$|\psi\rangle = \cos\theta |\phi\rangle + \sin\theta |\phi+90\rangle, \tag{1}$$

where θ is the angle between $|\phi\rangle$ and $|\psi\rangle$. The coefficients in front of the kets $-\cos\theta$ and $\sin\theta$ are probability amplitudes. Nature has dictated that the outcome of a measurement of the photon’s polarization state (for example, by transmission through a polarizing filter) is indeterminate—it depends on the basis (the orientation of the polarization axis) used to make the measurement. The probability p that a measurement of $|\psi\rangle$ yields the result $|\phi\rangle$ is given by the expression

$$p = \cos^2\theta, \tag{2}$$

that is, p is the square of the probability amplitude in front of the ket $|\phi\rangle$.




Figure A. Polarizing Filter
The filter projects photons into polarization states parallel to its polarization axis.

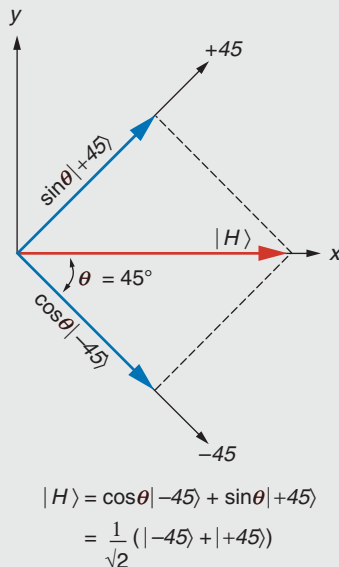


Figure B. Decomposition into Diagonal Basis
A horizontally polarized photon is expressed in terms of the $+45/-45$ basis.

We are now in a position to understand the simple experiment discussed earlier. The polarization axis of the first polarizing filter is set to be horizontal. Equation (1) tells us we can express an incoming photon as a superposition of a ket that is aligned parallel the polarization axis, that is, $|\phi\rangle = |0\rangle \equiv |H\rangle$, and a ket that is orthogonal to the axis, that is, $|\phi+90\rangle = |90\rangle \equiv |V\rangle$. We have

$$|\psi\rangle = \cos\theta |H\rangle + \sin\theta |V\rangle, \quad (3)$$

where the angle θ is now seen to describe the angle between the incoming photon's polarization and the filter's polarization axis. According to Equation (2), the probability that a linearly polarized photon passes through the horizontal polarizer is $p = \cos^2\theta$, that is, the square of the probability amplitude for the state $|H\rangle$. Because photons of all polarizations impinge on the first filter, the amount of light that gets through found by taking the average of p over all angles, that is, $\langle \cos^2\theta \rangle = 1/2$. Half the light makes it through the first filter.

Every photon that makes it through has been projected into the state $|H\rangle$. These photons then interact with the second filter in the stack with polarization axis aligned at $\phi = -45^\circ$. We express the horizontal photon in the diagonal ($-45^\circ/+45^\circ$) basis as (see Figure B):

$$|H\rangle = \cos(45) |-45\rangle + \sin(45) |+45\rangle = 1/\sqrt{2} (|-45\rangle + |+45\rangle). \quad (4)$$

The probability that a photon passes through the second filter is $\cos^2(45) = 1/2$, so 1/4 of the sunlight makes it through the two filters. The photons that emerge are polarized at -45° . The third filter is aligned vertically ($\phi = 90^\circ$), so we rewrite the ket $|-45\rangle$ in the horizontal/vertical (H/V) basis:

$$|-45\rangle = \cos(-135) |V\rangle + \sin(-135) |-H\rangle = 1/\sqrt{2} (-|V\rangle + |-H\rangle). \quad (5)$$

The probability that a photon passes through the vertical filter is $\cos^2(-135) = 1/2$. Again, half the photons make it through the last filter, so in total one eighth of the sunlight makes it through the stack.

This demonstration of nonorthogonal photon polarizations and polarizers reveals another important property of photons: All information about the initial polarization state is lost as a result of the photon-polarizer interaction. For cryptography, that has an unfortunate implication for someone (Eve) who is trying intercept the encrypted bit stream. Eve can intercept the photons going to Bob, but unless she measures the polarization of those photons in the correct basis, she cannot correlate the results of her measurements with a bit value. With her polarizer set to -45° , she has a probability to detect photons in the state $|-45\rangle$, $|H\rangle$, or $|V\rangle$, corresponding to bit values of 0, 0, and 1. Her measurement does not reveal Alice's bit value, nor does it reveal the original polarization state of the photon. A certain fraction of the photons she sends to Bob (which she must do to cover her tracks) will be in error. Thus, by choosing to send a random sequence of nonorthogonally polarized photons, Alice and Bob assure that Eve cannot attempt to measure the sequence without introducing detectable errors in their QKD protocol.

of those bits. If Eve is to get bit information, she is forced to breach the quantum channel by intercepting the photons and measuring their polarizations. She must then send new photons on to Bob in order to cover her tracks.

But Eve must know the exact state of a photon if she is to send a new one correctly. She cannot, however, make a deterministic measurement of the photon's polarization state because Alice sends photons in two nonorthogonal bases. For example, suppose Eve has a detection apparatus identical to Bob's and she detects a photon in her first detector (bit value of 0) when she measures in the diagonal basis. Did Alice send a photon in the $|H\rangle$, $|V\rangle$, or $|+45\rangle$ state? Eve has no idea because, given her measurement basis, she can detect each of those states. A hit on Eve's detector does not reveal whether Alice sent a 0 or a 1; that information "materializes" only after Alice and Bob compare bases. In fact, Eve can choose any type of detection system or measurement strategy and still be uncertain about the original state of Alice's photon.

One might ask whether Eve can make copies of Alice's photon before making a measurement. Then she could send the original off to Bob, save her string of photons (somehow), and make deterministic polarization measurements after listening to Alice and Bob compare bases. But quantum mechanics prevents Eve from accurately copying an unknown photon. (See the box "The No-Cloning Theorem" on page 79.) She would have to make a deterministic measurement, but that action would inevitably reveal her presence to Alice and Bob.

If she were to guess the polarization state, Eve would have, at best, a 50 percent chance of forwarding the correct one to Bob. But in making her guess, she will necessarily introduce



Figure 3. A Simple Error-Correction Scheme

Error correction removes single-bit errors from the sifted key. A simple scheme involves checking the parity of each byte (8-bit) sequence. The parity of a byte is 0 if the number of 1s in the byte is even or 1 if the number of 1s in the byte is odd. In this case, Alice and Bob start a public conversation to compare the parity of each of their three bytes. Because there is a mismatch, caused by the seventh bit (indicated in red) in the third byte, they try to locate the problem. They must eliminate all errors, or else their keys are unusable. Because the conversation takes place over an open communication line, Eve initially gains information about the parity of the sifted key. That information, however, can be eliminated if Alice and Bob drop some bits from their sequence. Relying on her old information, Eve will not understand anything about the new bit sequence.

errors into Alice and Bob's sifted-key sequence and, hence, increase the BER. When Alice and Bob check their sifted-key sequences for mismatches, they conservatively assume that Eve caused all the errors. They make corrections to those sequences, compute the maximum information Eve could have about the reconciled key, and then use privacy amplification to compress out Eve's possible knowledge about their shared secret strings to substantially

less than one bit. The secret key is truly secret.

Experiments

To date, the three major experiments performed at Los Alamos National Laboratory are free-space, fiber, and entangled-state QKD systems. All of the systems were constructed from readily available pieces of equipment, and we were able to

show that a complete QKD session could be communicated over long distances and still produce a useful secret-bit yield. All three systems use the BB84 protocol.

Here, we describe the free-space and fiber-based experiments. Entangled-state QKD is described on page 58 in the article "Quantum State Entanglement."

Free-Space QKD. In free-space QKD, photons are transmitted through

The No-Cloning Theorem

In 1982, Bill Wootters and Wojciech Zurek applied the linear properties of quantum mechanics to prove that an arbitrary quantum state cannot be cloned. Although their argument is entirely general, we will illustrate the theorem with polarized photons. Suppose we have a perfect cloning device in the initial state $|A_0\rangle$ and an incoming photon in an arbitrary polarization state $|s\rangle$. The device duplicates the photon as follows:

$$|A_0\rangle|s\rangle \rightarrow |A_s\rangle|ss\rangle, \quad (1)$$

where $|A_s\rangle$ is the device final state, which may or may not depend on the polarization of the original photon, and $|ss\rangle$ refers to the state of the electromagnetic field in which there are two photons, each with polarization $|s\rangle$. Suppose that the device can duplicate both the vertical $|V\rangle$ and the horizontal $|H\rangle$ polarization, that is,

$$|A_0\rangle|V\rangle \rightarrow |A_V\rangle|VV\rangle, \quad \text{and} \quad (2)$$

$$|A_0\rangle|H\rangle \rightarrow |A_H\rangle|HH\rangle. \quad (3)$$

According to quantum mechanics, this transformation should be representable by a linear operator, which means the operator acts independently on each orthogonal state in the Hilbert space. Therefore, if the incoming photon has some arbitrary polarization given by the linear superposition $|s\rangle = \alpha|V\rangle + \beta|H\rangle$, the result of its interaction with the apparatus will be a superposition of Equations (2) and (3):

$$\begin{aligned} |A_0\rangle|s\rangle &= |A_0\rangle (\alpha|V\rangle + \beta|H\rangle) \\ &= \alpha|A_V\rangle|VV\rangle + \beta|A_H\rangle|HH\rangle. \end{aligned} \quad (4)$$

If the apparatus states $|A_V\rangle$ and $|A_H\rangle$ are not identical, the two photons emerging from the apparatus are in a mixed state of polarization; if they are identical, the emerging two photons are in a pure entangled state, $\alpha|VV\rangle + \beta|HH\rangle$. In neither case does the apparatus produce a final state $|ss\rangle$ consisting of two completely independent photons, each in the polarization state $\alpha|V\rangle + \beta|H\rangle$:

$$\begin{aligned} |ss\rangle &= (\alpha|V\rangle + \beta|H\rangle) (\alpha|V\rangle + \beta|H\rangle) \\ &= \alpha^2|VV\rangle + \alpha\beta|VH\rangle + \beta\alpha|HV\rangle + \beta^2|HH\rangle. \end{aligned} \quad (5)$$

Linearity, therefore, rules out the existence of a device that could faithfully clone a photon in an arbitrary polarization state.

open air. The protocol uses polarization states, as previously described, because the atmosphere preserves polarization over a wide range of photon wavelengths (including the full range of visible and infrared light). The major difficulty is detecting the

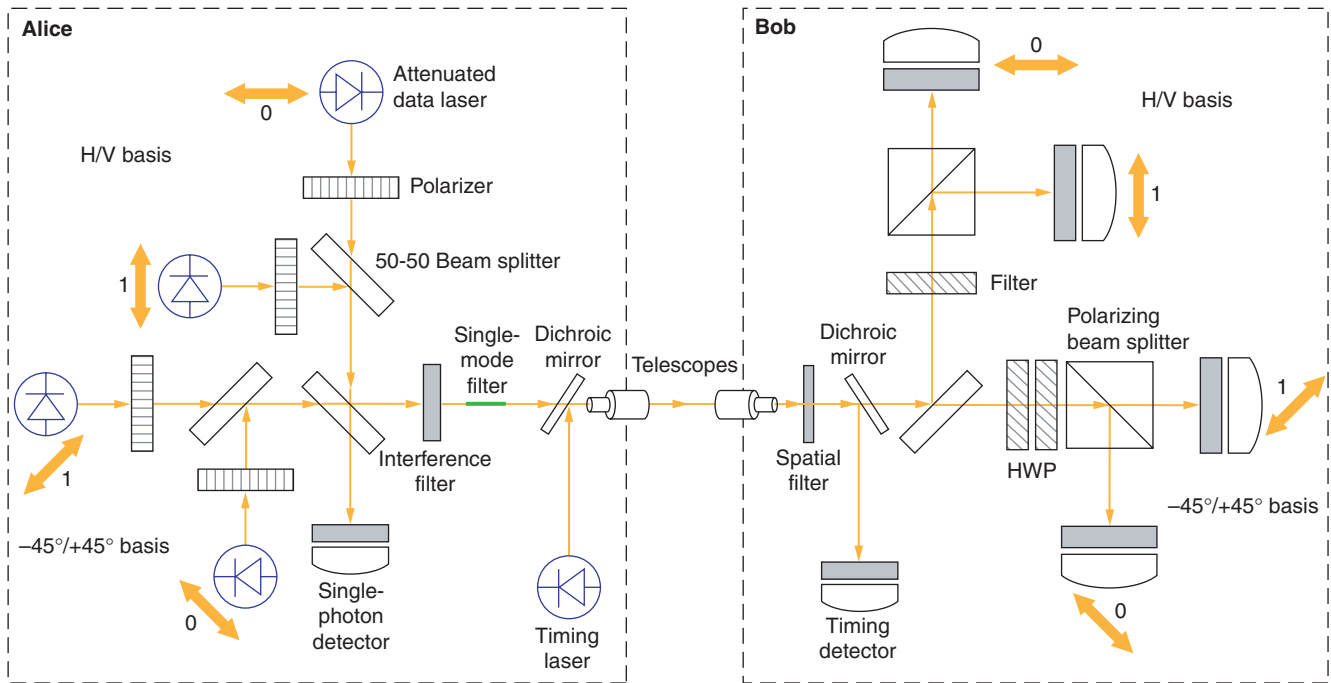
single QKD photons from within the enormous background of daytime photons, namely, $\geq 10^{10}$ background photons per centimeter squared, per second, per angstrom, per steradian ($\gamma/\text{cm}^2/\text{s}/\text{\AA}/\text{sr}$). This problem exists even at night because the background

from, say, moonlight or the light of urban areas is still much larger than the QKD signal. A second difficulty is dealing with losses due to atmospheric distortions. We are able to overcome both of these problems and can distinguish the QKD photons from background photons by using interference filters that transmit only photons of a specific wavelength, by carefully limiting the field of view, and by using a clever trick. The free-space QKD system is shown in Figure 4.


Alice and Bob have identical copies of the interference filters, which allow Alice to send photons at a selected wavelength and Bob to receive photons only at that wavelength. The preferred wavelength is about 772 nanometers, which is in the infrared and just outside the normal range of vision. The atmosphere is highly transmitting for light of this color, and single-photon detectors with good quantum efficiency at this wavelength are readily available. Furthermore, polarization selection and control components and diode lasers that produce the desired wavelength are all easily obtained.

A receiver telescope with a narrow field of view helps limit unwanted photons. Behind the telescope is a spatial filter that passes photons coming from a precise location (Alice's) while excluding all the others. The telescope must be employed with care, however. As anyone who has ever looked at the twinkling stars knows, the atmosphere can make a source of light appear to move. The magnitude of the movement varies considerably with the time of day, the weather, and the local terrain. If not accounted for, the atmosphere could cause Alice to shift rapidly in and out of Bob's field of view. Over short distances, these atmospheric distortions are not a serious problem. Over long distances, Alice

(a) Conceptual Diagram



(b) Alice's Optics Table



(c) Alice's Electronics



(d) Bob's System



Figure 4. Free-Space QKD

(a) In the BB84 protocol, Alice (the sender) encodes bits in the polarization states of single photons either as $0 = |H\rangle$ and $1 = |V\rangle$ or as $0 = |-45\rangle$ and $1 = |+45\rangle$. The data stream begins with a bright output pulse from the timing laser, which sets the timing of the pulse. A few nanoseconds later, one of the four data lasers ($\lambda = 772 \text{ nm}$) fires. Each data laser has its own attenuator, focusing optics, and polarizer. Each laser outputs a uniform pulse of the desired brightness in one of the four polarization states. The output of all four data lasers is combined by a series of beam splitters, which have been carefully arranged so that the distances between the lasers and output optics are the same (therefore eliminating any timing differences between the pulses). The final beam

splitter either directs the photons to a detector that monitors the average number of photons per laser pulse or sends the polarized photons through a narrow-pass interference filter (to remove any frequency differences) and a single mode fiber (to eliminate any spatial mode differences). The photons that pass through Alice's telescope are identical in every respect except for polarization. Bob (the receiver) uses spatial filtering, time-domain filtering, and wavelength selection to pick out Alice's photons from background. His telescope, with a field of view that is nominally 45 arc seconds (or 220 microradians), acts as a spatial filter that allows only photons from Alice's location to pass. The photons then pass through an interference filter (wavelength selection)

that is matched to the one in Alice's transmitter. Photons are sent to a 50-50 beam splitter, which acts as a basis selector by randomly directing a photon to one of the two measurement stations. Each station consists of a polarizing beam splitter and two single-photon detectors. A half-wave plate (HWP) rotates the photon's polarization before the $-45^\circ/+45^\circ$ station. A detector must fire within a set period following detection of the bright timing pulse (time-domain filtering). (b) Alice's compact optics table and (c) electronics are shown here. (d) Bob's telescope peers out from the door of the mobile trailer containing all his electronics and optical systems. Bob (and Alice) can be easily transported to different sites. Moreover, one person can operate the system.

corrects for atmospheric variations by observing Bob's beacon laser and is thus able to rapidly vary the point to which she sends the photons.

Finally, the clever trick is to send a bright laser pulse from Alice to Bob just before a single photon is sent so there is a known delay between the photon and the bright pulse. Bob accepts only photons that enter the system approximately 1 nanosecond after the bright pulse. This time-domain filtering greatly limits the possibility of a background photon being detected instead of a QKD photon. This system of multiple filtering techniques works so well that single QKD photons can be distinguished from background even in daylight.

One issue complicating the free-space system (as well as the other systems described below) is that the photon sources are actually attenuated laser diodes that produce weak laser pulses instead of true single photons. (Single-photon sources are currently too large and exotic for systems intended for use in the field.) The number of photons in a weak laser pulse is governed by Poisson statistics, and the number of photons in each pulse varies. The probability $P(n)$ that a pulse will contain n photons is,

$$P(n) = \frac{e^{-\mu} \mu^n}{n!}, \quad (4)$$

where μ is the average number of photons per pulse. If $\mu = 1$, there is roughly a 37 percent chance that a pulse will contain no photons, 37 percent that it will contain one photon, and 26 percent that the pulse will contain more than one photon.

By adjusting the attenuation, Alice can choose a specific value of μ . If she chooses a relatively high μ , say, above 1 photon per pulse, each time more than one photon is sent, it must be assumed that a clever eaves-

dropper would be able to detect and measure the extra photons. A great deal of privacy amplification—comcomitant with a large consumption of reconciled bits—is needed to keep the system secure, so overall, the secret bit yield decreases. If μ is too small, say, 0.05, then most of the time Alice is sending nothing over the quantum channel and experimental errors (such as background light getting into the receiver, dark counts in detectors, or even the actions of an eavesdropper) may dominate. Again, the secret-bit yield decreases. The choice of μ is therefore an important free parameter at Alice's disposal.

Our experiments have shown that the secret-bit yield depends strongly on atmospheric conditions. Turbulence along the optical path between Alice and Bob, for example, affects the transmission efficiency. To help show trends in the data, we construct a pseudo signal-to-noise ratio, η/C , where η is the transmission efficiency (obtained by dividing the number of sifted bits by μ) and C is the number of background photons detected by Bob.

Figure 5 shows data from a free-space QKD experiment that ran successfully at a 10-kilometer separation in daylight. The open communication channel was a wireless Ethernet. During the numerous experimental runs, Alice would send 10^6 laser pulses over a 1-second period. The value of μ was typically set between 0.1 and 0.8.

The experimental run labeled "Sample" in Figure 5 is a typical example. Approximately 22 percent of the pulses had a single photon ($\mu = 0.29$). After comparing Alice and Bob's bases, we constructed a sifted key of 651 bits. Following error correction, calculation of the BER, and privacy amplification, we obtained a secret key consisting of 264 bits, which is sufficient for the

new AES. Note that the secret-bit yield can be substantially higher at night (high η/C), because the background is reduced.

Our free-space system is a preliminary prototype for a system that could be flown on a spacecraft. Because the atmosphere has an effective thickness of only a few kilometers if one were to look straight up, our results are a good indicator of the feasibility of ground-to-satellite free-space QKD.

Fiber-Based QKD. The polarization state of a photon is not preserved in conventional optical fibers. That is why another physical property that could express the desired quantum mechanical properties for QKD had to be found in order to implement a fiber-based system. The solution was to have a photon interfere with itself after it travels down two paths of a twin Mach-Zehnder interferometer setup.

The concepts underlying the fiber-based QKD scheme are illustrated in Figure 6. Briefly, quantum mechanics tells us that a single photon entering a Mach-Zehnder interferometer behaves as if it has taken both paths through the instrument. The entrance beam splitter places the photon in a quantum mechanical superposition, with a component that describes a photon traversing the upper path and a component that describes the photon traversing the lower path. The two components have a definite phase relationship and can interfere with each other.

As seen in the figure, Alice can introduce a phase shift ϕ_A to the photon on one arm of the interferometer, while Bob can introduce a phase shift ϕ_B on the other. Depending on the phases set by both Alice and Bob, the interference at the exit beam splitter is such that the photon has a definite probability

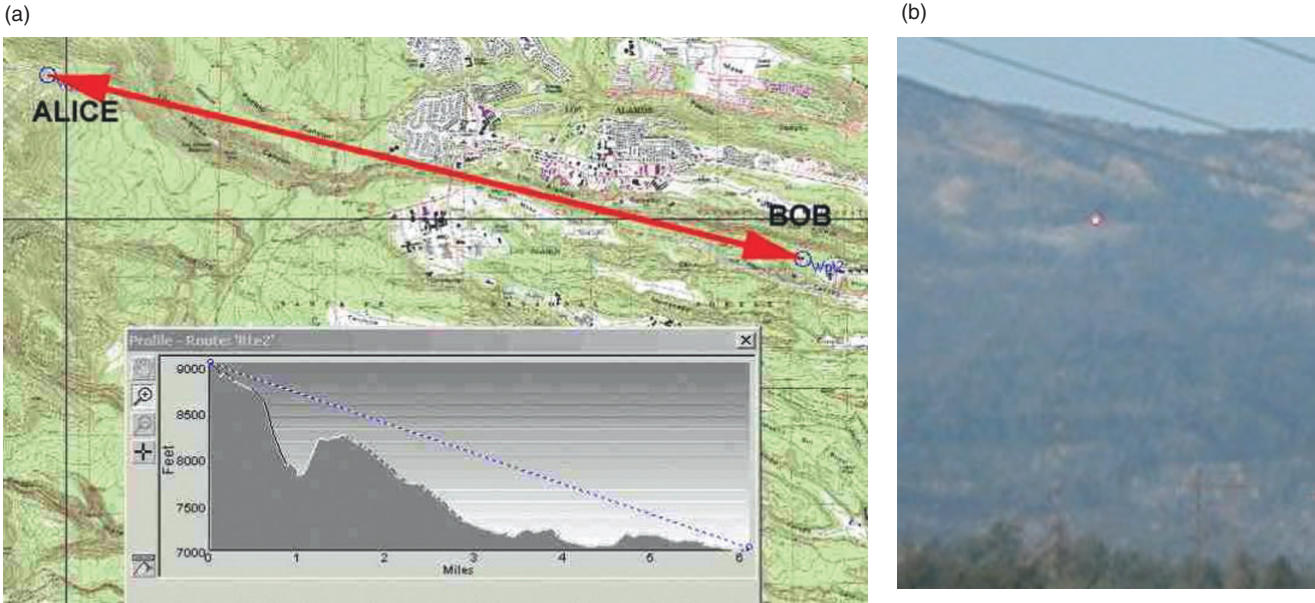
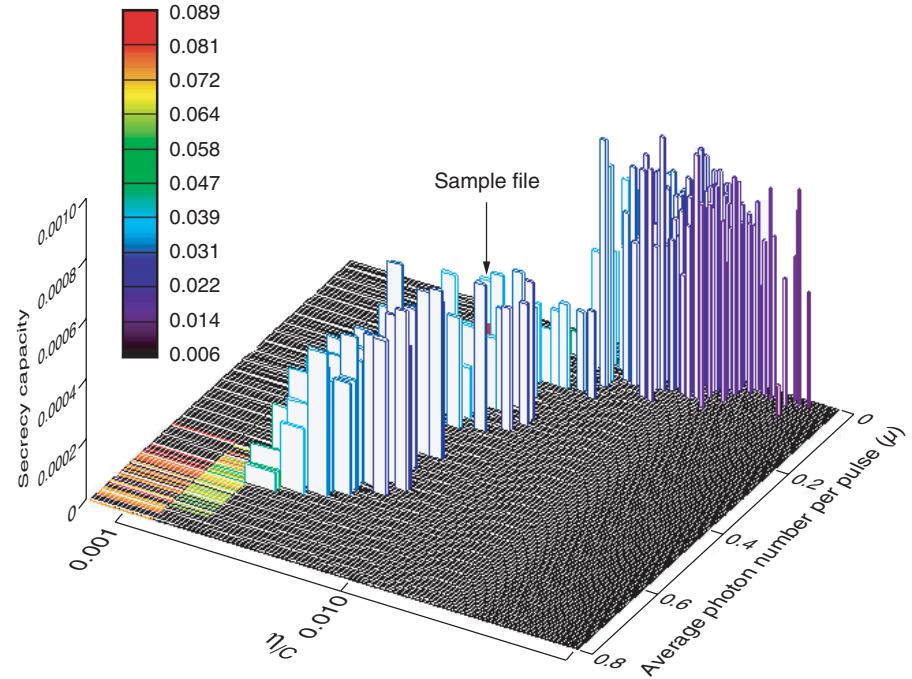


Figure 5. Data from a 10-km Free-Space QKD Experiment

(a) Alice was located halfway up Pajarito Mountain, in New Mexico, while Bob was 10 km away, at a Los Alamos lab site. (b) The bright red dot near the center of the picture is a spotting laser sent through Alice’s telescope. It was used to optically align the transmitter and receiver for the quantum channel. (c) Data from the experiment show the dependence of the secret-bit yield (normalized to the number of sifted bits) on the average number of photons per pulse μ and on the pseudo signal-to-noise ratio η/C (discussed in the text). Each vertical column corresponds to an experimental run in which Alice sent 10^6 polarized photons in 1 s. The flat, black regions of the graph are areas for which no data are available. With favorable atmospheric conditions or low background (high η/C), we can run at lower μ values and still obtain a high bit yield. Poorer conditions (low η/C) require higher μ values and result in a lesser yield.

(c) Sifted-Bit Error Rate (r)



to hit either of two detectors. The probability P_U that the photon hits the upper detector is given by

$$P_U = \sin^2\left(\frac{\phi_A - \phi_B}{2}\right), \quad (5)$$

whereas the probability P_L that the photon hits the lower detector

is given by

$$P_L = \cos^2\left(\frac{\phi_A - \phi_B}{2}\right) \quad (6)$$

We make use of these relations to implement the BB84 protocol. Alice chooses at random between two bases, X and Y. If she chooses the X-basis, then

for a bit value of 0 or 1, she sets $\phi_A = 0^\circ$ or 180° , respectively. If Alice chooses the Y-basis, then she chooses $\phi_A = 90^\circ$ or 270° for bit values of 0 or 1, respectively. At his end, Bob sets his phase angle ϕ_B to 0° if he is in the X-basis and to 90° if he is in the Y-basis.

Table II summarizes Alice and Bob’s choices and shows the value of the probabilities P_U and P_L , given the

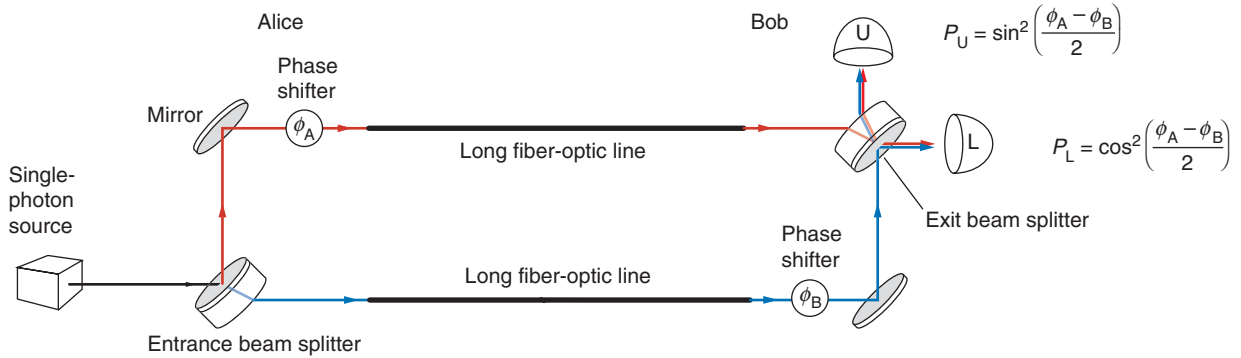


Figure 6. Mach-Zehnder Interferometer and Fiber-Based QKD Concept

In a Mach-Zehnder interferometer, a photon is placed in a superposition of two states by the entrance beam splitter. It travels down both arms simultaneously, and interferes with itself at the exit beam splitter. In the conceptual fiber-based QKD system illustrated here, a phase shifter is placed in each arm of the interferometer. Alice randomly chooses a bit value and a basis and sets the angle of her phase shifter according to her choices (see Table II below). Bob sets the

angle of his phase shifter according to his basis choice. The table shows the probability that Bob detects a photon in a given detector. When Alice and Bob use the same basis for sending and measuring, a hit in Bob’s lower detector means that Alice sent a bit value of 0, whereas a hit on the upper detector means she sent a 1. Because there is no such correlation when Alice and Bob use different bases, those bit values are discarded.

Table II. Fiber-Based QKD

| Sender (Alice) | | | Receiver (Bob) | | | | Action | |
|----------------|-----|-----------------------------|----------------|-----------------------------|-----------------|-------|--------|-------------|
| Basis | Bit | Phase ϕ_A ($^\circ$) | Basis | Phase ϕ_B ($^\circ$) | Probability (%) | | | Bit |
| | | | | | P_L | P_U | | |
| X | 0 | 0 | X | 0 | 100 | 0 | 0 | Keep bit |
| X | 1 | 180 | X | 0 | 0 | 100 | 1 | Keep bit |
| X | 0 | 0 | Y | 90 | 50 | 50 | 0 or 1 | Discard bit |
| X | 1 | 180 | Y | 90 | 50 | 50 | 0 or 1 | Discard bit |
| Y | 0 | 90 | X | 0 | 50 | 50 | 0 or 1 | Discard bit |
| Y | 1 | 270 | X | 0 | 50 | 50 | 0 or 1 | Discard bit |
| Y | 0 | 90 | Y | 90 | 100 | 0 | 0 | Keep bit |
| Y | 1 | 270 | Y | 90 | 0 | 100 | 1 | Keep bit |

various combinations of ϕ_A and ϕ_B . Because we are implementing BB84, Table II is essentially the same as Table I. When Alice and Bob choose the same basis, a photon representing Alice’s 1 always goes to the upper detector, and a photon representing her 0 always goes to the lower. If Alice and Bob use different bases, the photon has equal probability to emerge from either port, and Bob has no information about what bit value Alice has sent. At the end of the session, Bob calls Alice on the open

communications line, and the two compare which bases they used for each photon. They keep the bit values when the bases agree and discard the other bits.

In the scheme discussed above, a single Mach-Zehnder interferometer stretches between Alice and Bob. In practice, that is a bad idea. The photon needs to maintain phase coherence as it propagates down the two optical fibers that make up the long arms of the interferometer. Photons often experience random phase shifts as

they go through long fiber-optic cables, and because the shifts in one arm are independent of those in the other, the interference condition at the exit beam splitter changes in a random fashion. Furthermore, having two dedicated fibers would be expensive to operate in the real world.

A better idea is for Alice and Bob each to have a Mach-Zehnder interferometer, with the two connected by a single long fiber—see Figure 7.

Each interferometer is modified to have a long arm and a short arm, and




Figure 7. Implementation of Fiber-Based QKD

Our fiber-based QKD system uses two modified interferometers connected by a single, long optical fiber. Each interferometer has a long (L) arm and a short (S) arm. In going from Alice’s entrance beam splitter to Bob’s exit beam splitter, the photon can take paths S_1S_2 , L_1L_2 , S_1L_2 , and L_1S_2 . The latter two paths have the same length, and the photon traveling them can maintain phase coherence and interfere with itself. The protocol then works as described in Figure 6.

the path length differences between the two arms are greater than the coherence length of the photon. There is no interference as the photon leaves Alice’s instrument. But of the four possible paths through the entire system (refer to the figure), the two designated as S_1L_2 and L_1S_2 are of equal length (to within the phase coherence length of the photon). A photon that travels down those two paths interferes with itself at Bob’s exit beam splitter. The system therefore behaves as if it were a single instrument. Alice and Bob are still free to vary the phase on one arm of their interferometers, as needed, to carry out the protocol.

Our system transmits bits through 48 kilometers of fiber. As in the free-space experiments, Alice first sends a bright pulse to trigger the detectors and to limit background interference. Single photons are sent at 1310 nanometers, and the bright timing pulse is at 1550 nanometers. The secret-bit yield is lower than that obtained in the free-space experiment.

Summary

Quantum cryptography can enable secure transmission of sensitive, proprietary, or national security information across a metropolitan area or

corporate campus and provide the long-term security guarantees such data require. It is the only technology that will be secure no matter what technology an adversary develops in the future. Furthermore, it raises the stakes for eavesdroppers because they must perform risky, active attacks against a system. Currently, a public-key encrypted system can be attacked through passive, standoff monitoring.

Because of the inherent advantages of quantum cryptography, we can envision a future in which a QKD system provides secure communications in metropolitan areas between banks, between off-site stock-trading centers and central stock exchanges, between corporate offices, and between offices and broadband data networks. Money transfers between banks now amount to over \$2 trillion per day worldwide and well justify the expense of implementing QKD systems. Optical wireless “last-mile” communications systems could even provide broadband access to most homes.

By combining theoretical analyses with innovative experimental advances, the Los Alamos quantum cryptography team has already demonstrated the practicality of free-space quantum cryptography in a series of record-setting experiments. In 1996, the team demonstrated

atmospheric quantum-key transmission at night, quickly followed by a record-setting 0.5-kilometer point-to-point transmission in full daylight, then a 1.6-, and finally a 10-kilometer transmission. The world record for the longest QKD distribution in fiber—48 kilometers—was also held by the Los Alamos team for many years. Several of the first demonstrations of entanglement-based QKD have also been performed at the Laboratory.

In the near future, the free-space quantum cryptography system could provide secure satellite communications—using a low-orbit satellite—between cities anywhere in the world. When deployed on a spacecraft, our system can be used to generate cryptographic keys between any two users who are anywhere on the planet and can view that spacecraft. Each user would individually generate a key with the spacecraft. The second user would then be instructed to change specific bits so that the two users’ keys would match. Because the spacecraft only needs to instruct the user which bits to change, and can do so without revealing any bit values, this is a secure key-generation methodology.

On a more philosophical note, the challenging demands of cryptography have already produced a huge growth in research into the foundations of

quantum mechanics. Fundamental concepts that were previously thought to be testable only in thought experiments have been subjected to experimental verification. Many concepts, such as entanglement, that have been almost completely neglected since the early days of quantum physics have been explored and realized. This trend will continue, and we will find out to what extent the creation and control of “mesoscopic” quantum systems, that is, the netherworld between single-particle behavior and collective-particle behavior, can be performed. This research may help elucidate the puzzling transition between the quantum and classical regime. The development of quantum technology will open up other applications of quantum physics, such as quantum-enhanced sensors and improvements to atomic clocks and satellite navigation systems. Whether or not quantum cryptography becomes a widely adopted technology, we are in for an interesting next decade. ■

Acknowledgments

The Quantum Cryptography team combines the talents of numerous scientists and engineers, including those of Kevin P. McCabe, George L. Morgan, Michael J. Pigue, Steven A. Storms, Paul A. Montano, James T. Thrasher, and especially Charles G. Peterson. The authors wish to thank Derek Derkacs for technical support. We gratefully acknowledge support for the 10-kilometer free-space experiment from the National Reconnaissance Office Director’s Innovation Initiative program, administered by Col. John Comtois and Peter Hendrickson.

Further Reading

- Bennett, C. H. 1992. Quantum Cryptography: Uncertainty in the Service of Privacy. *Science* **257** (5071): 752.
- Bennett, C. H., and G. Brassard. 1984. Quantum Cryptography: Public-Key Distribution and Coin Tossing. In *Proceedings of the IEEE International Conference on Computers, Systems, and Signal Processing, Bangalore, India, 1984*. p. 175. New York: IEEE.
- Bennett, C. H., G. Brassard, C. Crépeau, and U. M. Maurer. 1995. Generalized Privacy Amplification. *IEEE Trans. Inf. Theory* **41** (6): 1915.
- Bennett, C. H., G. Brassard, and A. K. Ekert. 1992. Quantum Cryptography. *Sci. Am.* **267** (4): 50.
- Hughes, R. J., D. G.L. Morgan, and C. G. Peterson. 2000. Quantum Key Distribution over a 48-km Optical Fiber Network. *J. Mod. Opt.* **47**: 533.
- Hughes, R. J., J. E. Nordholt, D. Derkacs, and C. G. Peterson. 2002. Practical Free-Space Quantum Key distribution over 10 km in Daylight and at Night. *New J. Phys.* **4**: 43. [Online]: <http://www.njp.org>
- Hughes, R. J., W. T. Buttler, P. G. Kwiat, S. K. Lamoreaux, G. L. Morgan, J. E. Nordholt, and C. G. Peterson. 2000. Free-Space Quantum Key Distribution in Daylight. *J. Mod. Opt.* **47**: 549.
- Hughes, R. J., W. T. Buttler, P. G. Kwiat, S. K. Lamoreaux, G. L. Morgan, J. E. Nordholt, and C. G. Peterson. 2000. Quantum Cryptography for Secure Satellite Communications. In *2000 IEEE Aerospace Conference Proceedings*. p. 191. New York: IEEE.
- Hughes, R., and J. Nordholt. 1999. Quantum Cryptography Takes to the Air. *Phys. World* **12** (5): 31.
- Hughes, R. J., W. T. Buttler, P. G. Kwiat, S. K. Lamoreaux, G. G. Luther, G. L. Morgan, J. E. Nordholt, and C. G. Peterson. 1999. Quantum Cryptography for Secure Free-Space Communications. *Proc. SPIE-Int. Soc. Opt. Eng.* **3615**: 98.
- Nordholt, J. E., R. J. Hughes, G. L. Morgan, C. G. Peterson, and C. C. Wipf. 2002. Present and Future Free-Space Quantum Key Distribution. *Proc. SPIE-Int. Soc. Opt. Eng.* **4635**: 116.
- Schneier, B. 1995. *Applied Cryptography: Protocols, Algorithms Source Code in C*. New York: John Wiley & Sons.
- Singh, S. 1999. *The Code Book: The Evolution of Secrecy from Mary, Queen of Scots to Quantum Cryptography*. New York: Doubleday.
- Wootters, W. K., and W. H. Zurek. 1982. A Single Quantum Cannot be Cloned. *Nature* **299**: 802.


Jane E. (Beth) Nordholt has broad experience in quantum-key distribution, experimental astrophysics, high-energy physics, computing, and space plasma physics. Currently a technical



project leader at Los Alamos National Laboratory, she is the coinventor for the free-space quantum key distribution project and holds several patents on quantum-key distribution and spacecraft instrumentation. Beth received

four NASA group achievement awards and two Los Alamos Distinguished Performance Awards. In 2001, she received an R&D 100 Award for her work on free-space quantum cryptography from the *Research and Development* magazine. Her interests include quantum cryptography, quantum communications, quantum metrology, the composition of planetary magnetospheres, planetary science, and advanced instrumentation.

Richard J. Hughes is a Laboratory Fellow and Quantum Information Science team leader in the Neutron Science and Technology Group of the Physics Division at Los Alamos National Laboratory. He is the principal investigator for several projects in quantum computation and quantum cryptography. Richard obtained his Ph.D. in theoretical elementary particle physics from the University of Liverpool and held research positions at Oxford



University and The Queen’s College, Oxford; California Institute of Technology; and CERN, the European Center for Nuclear Research. He was a distinguished visiting scientist at Oxford University and the University of Oslo. Richard was awarded the Los Alamos Fellows Prize for his work on quantum information science; he was twice awarded Los Alamos Distinguished Performance Awards for his quantum cryptography research; and he was cowinner of an R&D 100 Award for the entry “Free-Space Quantum Cryptography.” He became a Fellow of the American Physical Society in 1999. He has authored over 100 scientific papers on quantum field theory, the foundations of quantum mechanics, quantum cryptography, and quantum computation. In his spare time, Richard enjoys ultramarathon trail running over distances of up to 100 miles.