QUANTUM CRYPTOGRAPHY DEFIES EAVESDROPPING

Students of physics can be excused for finding their quantum mechanics courses a little cryptic from time to time. Now researchers are using the properties of quantum mechanics to encrypt information for secure transmission.

Artur Ekert (Oxford University) has shown how to use Bell's inequalities and the Einstein–Podolsky–Rosen effect to guarantee secure distribution of a cryptographic key.[1] Ekert's scheme can be simplified to one that is equivalent in many ways to a scheme proposed in 1984 by Charles H. Bennett (IBM) and Gilles Brassard (University of Montreal) that uses only one-particle states and not the paired states of EPR.[2,3] Recently Bennett showed that any two nonorthogonal states will suffice for transmitting a key.[4] Don't imagine that quantum cryptography is only a field of theoretical musings. Bennett, Brassard and their collaborators built a working quantum cryptography device in 1989, and other groups are working on similar devices.[5,6]

The key to secrecy

Quantum cryptography began in the late 1960s with work by Stephen J. Wiesner (then at Columbia University). In a paper that was written in about 1970 but remained unpublished until 1983, he showed how quantum effects in principle could be used to manufacture "bank notes" immune to counterfeiting. Wiesner also proposed a quantum multiplexing channel, in which two or three messages are combined and the reading of one message will destroy the others. In the 1980s quantum versions of various cryptographic techniques were suggested, including quantum key distribution.

The central idea behind many of the quantum techniques is that an eavesdropper can't monitor transmissions based on quantum mechanics without being noticed by the participants. Quantum mechanics, however, is more suited to transmitting random data than a predetermined message. Even if some bits are lost because of imperfect detectors, the recipient still gets a random bit stream. The random message can subsequently be used as the key for a classical transmission encoded using the one-time pad technique.

The one-time pad, devised in about 1918, is one of the simplest and most secure encryption schemes. The message is converted to binary, and both sender and receiver have a copy of a secret key—a random sequence of 1's and 0's. The sender combines the key and the message using the exclusive OR (XOR) operation (equivalent to addition modulo 2), and the receiver decodes the message with a similar application of XOR to the encrypted message and the key.

Although unbreakable, the one-time pad has the drawback of using large amounts of key, which can only be used once: If two messages are sent using the same key, an eavesdropper can take the XOR of the two encrypted messages, eliminating the random key and leaving something that is relatively easy to crack.

More efficient classical means of transmitting information, such as public-key cryptography, rely on the difficulty of solving certain "hard" problems such as the factoring of large numbers. However, these techniques can be defeated by exhaustive computer analysis or by the discovery of better algorithms for solving the problems on which they are based. By contrast, information theory and the laws of physics guarantee the security of messages sent by one-time pad using a key distributed by quantum mechanics.

Bennett and Brassard[3] developed a quantum scheme for transmitting a random key in 1984. (See the figure on page 22.) In that scheme, the sender of the key (traditionally code-named Alice) prepares states randomly out of a selection of four states. For example, she could use photons with vertical, horizontal, left circular and right circular polarizations. (Recently Bennett showed that any two nonorthogonal states suffice to distribute a key in much the same way as the scheme described here.[4]) Alice sends the photons to her ally, Bob, who randomly chooses to measure either the rectilinear or the circular polarization of each photon. After completing all the measurements, Alice and Bob discuss their data using a classical communications channel, which may be bugged by their antagonist, Eve. In fact, this discussion can be completely public. The only requirement is that the classical messages can't be altered or suppressed by Eve. Alice and Bob discard results in which Bob failed to detect a photon and those for which he made measurements using a polarization basis different from the one prepared by Alice. They then compare a large subset of the remaining results. If there has been no tampering, Bob's measurements of these will perfectly match what Alice prepared, and they can deduce that the unchecked ones are almost certainly also untampered with and can be used as bits of a secret random key. Without knowing Alice and Bob's polarization choices in advance, Eve cannot determine the photon states without introducing errors in Bob's measurements.

Einstein–Podolsky–Rosen

Ekert's EPR-based system[1] could be built by modifying the apparatus that Alain Aspect and coworkers (Institute of Theoretical and Applied Optics, Orsay) used to test Bell's inequalities.[7] In Ekert's scheme, a source emits pairs of spin-$\frac{1}{2}$ particles in a singlet state (or, equivalently, pairs of correlated photons). Alice receives one of each pair; Bob the other. Alice and Bob each have a set of three axes (such as horizontal, vertical and diagonal), and for each pair of particles they each independently choose one axis at random and measure their particle's spin along that axis. After a series of particle pairs have been transmitted, they announce which axes they used for each measurement, and they discard results in which either of them failed to detect a particle.

They compare the results of the measurements that were taken with different axes, and check a correlation function of those results to ensure that the particles were not disturbed by Eve. The correlation function is related to Bell's inequalities and shows whether or not the pairs that they received were still correctly correlated. If the function checks out, Alice and Bob are assured that their remaining results—measurement pairs taken using identical axes—are almost certainly precisely correlated and usable as the bits of a secret key.

It might seem that Ekert's EPR-based scheme is automatically secure: While the particles are in transit the information that Eve wishes to purloin does not yet exist. However, Eve might try to circumvent the system by substituting her own choice of particle states in place of those from Alice and Bob's source. But even this will fail, because she will not know the orientations that Alice and Bob choose for their analyzers. "Her in-

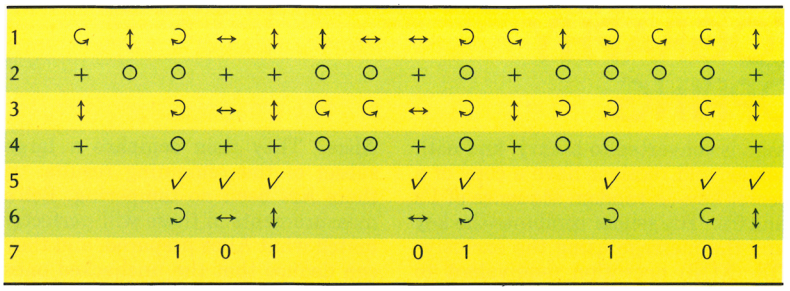

Quantum key distribution via photon polarization states. Alice sends a random sequence of polarized photons (1) to Bob, who uses a random sequence of polarization bases (2) to measure them (3). Bob tells Alice which bases he used (4) and she tells him which were correct (5). They keep the correct data (6) and interpret it as a binary sequence (7). This raw key now undergoes error correction and "distillation" as described in the text. (Adapted from ref. 4.)

tervention would be equivalent to the introduction of elements of physical reality, and the correlation function won't have the correct quantum mechanical value," Ekert says.

David Mermin (Cornell) realized that Bell's theorem isn't needed, and he proposed a simplification of Ekert's scheme that uses only two orientations of the polarizers.[2] Alice and Bob perform essentially the same procedure as in Ekert's scheme, but instead of computing a correlation function they discard the mismatched measurements and check a subset of the matched measurements, as in Bennett and Brassard's 1984 scheme. Mermin proved that Eve couldn't fake out Alice and Bob by generating pairs correlated to a third system that she would later measure. Any correlation that avoids detection by Alice and Bob necessarily yields no useful information to Eve. (Ekert had conjectured, but did not prove, that this held for his original scheme.)

Bennett pointed out[2] that the simplified scheme was in many ways equivalent to the scheme he and Brassard had proposed in 1984. The equivalence can be seen in this way: If Alice produces EPR pairs, measures one particle of each pair as in Ekert's scheme and sends the other to Bob, then one has essentially the scheme of Bennett and Brassard. Alice is merely using EPR as a device to generate the random states.

For both schemes Mermin, Bennett and Brassard proved that attacks that can avoid detection will yield no information to Eve. Even if Eve could measure the entire sequence of transmitted states as a single coherent entity—a possibility Brassard calls science fiction—she could not secretly extract information about the key.

Ekert points out that his EPR scheme is not entirely equivalent to the one-particle scheme and does have an advantage. In the one-particle scheme, once Alice and Bob create their key it exists in classical form and must be stored securely until it is used. In the EPR scheme, however, Alice and Bob could in principle store their particle pairs without measuring them until just before they were ready to use them. This would greatly reduce the time during which the key is in classical form and vulnerable to Eve's snooping. In practice, however, no one knows how to store EPR-correlated particles without destroying the correlations.

Privacy amplification

Serious practical problems challenge the implementation of any of these schemes, because they assume perfect one- or two-particle states and perfect detectors. In reality, erroneous counts and noise can conceal Eve's efforts to learn or predetermine the key, or can cause Alice and Bob to discard their key, believing it to be compromised when it is not. A realistic scheme needs error correction and privacy amplification—the distillation of a perfectly secret key out of a sequence of raw bits that Eve may have partial knowledge of.

Bennett, Brassard and Jean-Marc Robert (University of Quebec at Rimouski) have devised just such a scheme.[5,8] Alice and Bob take blocks of bits of their raw key and add each block's bits modulo 2 to obtain the parity (0 or 1) of the block. Bob's parity will differ from Alice's for any block that suffered an odd number of errors in transmission. By systematically comparing the parities of many overlapping blocks and subdividing those that reveal errors, Alice and Bob can almost certainly locate and eliminate all the errors introduced by either equipment problems or Eve's interference. For each bit of parity that is compared, one bit must be discarded from the raw key to ensure that Eve does not learn anything about the final key by listening to the error-correcting discussion.

The number of errors found lets Alice and Bob estimate an upper bound on how many bits of the key Eve is likely to know. Finally, Alice chooses random groups of bits in what remains of the key, and she and Bob use the parities of these groups of bits as their secret key. Even if Eve knows some of the remaining bits of the key, this hashing procedure can, with probability extremely close to 1, extract a shorter key about which Eve has not a single bit of information.

Bennett and Brassard put their amplification scheme to work in an apparatus that they built in 1989 with the help of their students John Smolin (UCLA) and François Bessette and Louis Salvail (University of Montreal).[5] In place of single photons, their device used faint incoherent pulses of light produced by an LED and subsequently linearly or circularly polarized. Each transmitted pulse had an intensity of about 0.1 photons. Software avatars of Alice and Bob ran on a PC, and eavesdropping was simulated by a software Eve with idealized abilities.

With faint incoherent pulses, Eve has two strategies: beam splitting and beam interception. The use of faint pulses minimizes the expected number of photons per pulse, but even if the expected number of photons per pulse is much less than 1, the pulses will have multiphoton components to their state. Hence Eve can split the beam and occasionally learn the polarization that Alice sent by detecting one photon while an identical photon travels on to Bob. In the interception strategy, Eve intercepts some pulses, reads their polarizations in randomly chosen bases and sends Bob pulses whose polarizations correspond to her measurement results. Assuming Eve has perfect detectors, each intercept has a 50% chance of giving her the correct bit and a 50% chance of yielding a random bit. 25% of the intercepts will introduce an error detectable by Bob. After monitoring Alice and Bob's discussion, Eve will know which 50% of her intercepts were correct. The possibility of better strategies is still being investigated.

A major drawback of the prototype used by Bennett, Brassard and their coworkers is the short distance over which the key can be transmitted: The photons traveled in a channel only 32 cm long. Bennett and Brassard point out that their device is hardly state of the art. It was, after all, built out of off-the-shelf parts by

theorists. Nevertheless, extending polarization-based schemes to longer distances could be problematical: Light-carrying fibers have not yet been developed that will preserve the required range of polarizations over long distances.

Alternative schemes that are more amenable to long distances use interferometry, relying on differences in phase instead of differences in polarization.

One such scheme is EPR based. It uses parametric down-conversion to produce pairs of photons whose phases are correlated in much the same way that spins are correlated in Ekert's spin-based scheme. The photons travel along fibers to Alice and Bob, who each have an identical Mach–Zehnder interferometer. The setting of a phase shift in one arm of each of their interferometers corresponds to the choice of a polarization axis in the spin system. John Rarity and Paul Tapster of the British Defence Research Agency (formerly the Royal Signals and Radar Establishment) are collaborating with Ekert and G. Massimo Palma (University of Palermo) to develop such a device.[6] Already they have a device that can distribute keys over a few meters, and they expect to increase that to kilometers with new fibers and photodetectors.

Bennett has proposed a non-EPR scheme using interferometry.[4] (See the figure below.) Alice prepares a dim pulse and a following brighter reference pulse by splitting an initial pulse with a beam splitter. The reference pulse is delayed in the long arm of her half of the interferometer, while the signal pulse undergoes a phase shift of 0 or π in the short arm. These pulses are then sent through an optical fiber to Bob, who splits the beam with his half of the interferometer. If he selects the same phase shift as Alice, he will detect a signal pulse; if he selects the wrong phase shift, the signal will be destroyed by interference. These results are interpreted as bit streams and processed as in the polarization-based schemes.

Post-cold-war cryptography

Even if the technological challenges of transmitting over long distances are overcome, quantum distribution of keys may prove too expensive or troublesome to be competitive with classical encryption schemes. However, there are other "post-cold-war" roles for cryptography that do not require transmission over distances and that could be of commercial interest in the near future. "In this scenario there are no enemies, exactly," says Bennett, "but you must negotiate with everyone and you don't entirely trust them."

A typical problem is two-party secure computation, in which both parties wish to know the result of a computation and wish to be sure that the result is computed correctly from true data, but neither side wishes to reveal all of its own data. This goal can be achieved through protocols based on public-key cryptography methods or through the use of trusted intermediaries. There are also quantum protocols, which ensure security by the laws of physics. For these protocols the ability to transmit over a distance is not important; one can imagine the participants sitting around the same table.

Many of these protocols can be built out of one protocol, oblivious transfer, which is almost identical to the quantum multiplexing channel that Wiesner proposed in 1970. Michael O. Rabin (Harvard) independently formulated oblivious transfer in 1981. The essence of oblivious transfer sounds quite pointless: One sends a stream of bits and arranges that each bit has only a fixed probability of arriving. The receiver knows which bits arrived, while the sender does not. However, out of this random ignorance can be built up a sophisticated partial ignorance where both parties, with extremely high probability, know the value of the function in question, but neither party knows the value of all the data that went into computing the function.

Many researchers have participated in reducing complicated cryptographic schemes to simpler protocols. In 1988 Claude Crépeau (Ecole Normale Supérieure) and Joe Kilian (NEC Research, Princeton, New Jersey) showed how the ideal quantum channel could be used to implement secure versions of oblivious transfer and how these could in turn be built into more complicated protocols such as multiparty secure computations. More recently Crépeau, Bennett and Brassard have developed somewhat more practical schemes for quantum oblivious transfer and related protocols. At present these schemes take millions of oblivious transfers to build up the two-party protocols. Crépeau, Brassard and coworkers are working at developing more efficient schemes.

Thus far a quantum device has not been built for performing secure oblivious transfer. The known protocols will tolerate error rates of about 1%, but the existing devices have error rates of 4–5%. Bennett points out that even when the remaining practical problems of quantum cryptography are overcome, one will need a high level of distrust to favor the quantum schemes over the less secure classical schemes. But as he says, "Quantum cryptography is another interesting way in which the quantum world differs qualitatively from the classical world."

—Graham P. Collins

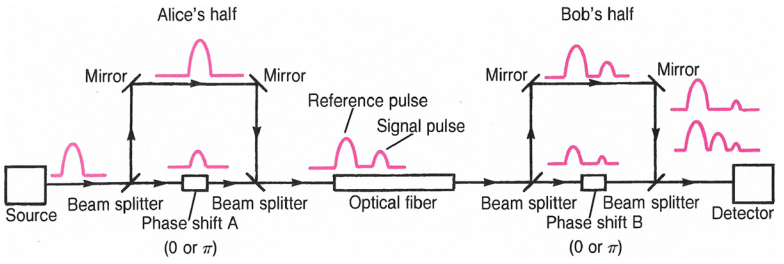

Interferometric quantum channel can use an optical fiber to distribute a key over long distances. The unsymmetric beam splitters delay the bulk of the pulse in the long arms, while a small fraction of the pulse undergoes a phase shift in the short arm. Bob detects a signal pulse only if he chooses the same phase shift as Alice. (Adapted from ref. 4.)

References

1. A. K. Ekert, Phys. Rev. Lett. **67**, 661 (1991).
2. C. H. Bennett, G. Brassard, N. D. Mermin, Phys. Rev. Lett. **68**, 557 (1992).
3. C. H. Bennett, G. Brassard, in *Proc. IEEE Int. Conf. on Computers, Systems, and Signal Processing, Bangalore, India*, IEEE, New York (1984), p. 175.
4. C. H. Bennett, Phys. Rev. Lett. **68**, 3121 (1992).
5. C. H. Bennett, F. Bessette, G. Brassard, L. Salvail, J. Smolin, J. Cryptol. **5**, 3 (1992).
6. A. K. Ekert, J. G. Rarity, P. R. Tapster, G. M. Palma, Phys. Rev. Lett. **69**, 1293 (1992).
7. A. Aspect, P. Grangier, G. Roger, Phys. Rev. Lett. **49**, 91 (1982). A. Aspect, J. Dalibard, G. Roger, Phys. Rev. Lett. **49**, 1804 (1982).
8. C. H. Bennett, G. Brassard, J.-M. Robert, SIAM J. Comput. **17**, 210 (1988). ■