

Experimental Quantum Cryptography

Charles H. Bennett
IBM Research *

François Bessette[†]
Université de Montréal[‡]

Gilles Brassard[§]
Université de Montréal[‡]

Louis Salvail
Université de Montréal[‡]

John Smolin[¶]
UCLA **

Abstract

We describe initial results from an apparatus and protocol designed to implement *quantum public key distribution*, by which two users, who share no secret information initially: 1) exchange a random quantum transmission, consisting of very faint flashes of polarized light; 2) by subsequent public discussion of the sent and received versions of this transmission estimate the extent of eavesdropping that might have taken place on it, and finally 3) if this estimate is small enough, can distill from the sent and received versions a smaller body of shared random information (key), which is certifiably secret in the sense that any third party's expected information on it is an exponentially small fraction of one bit. Because the system depends on the uncertainty principle of quantum physics, instead of usual mathematical assumptions such as the difficulty of factoring, it remains secure against an adversary with unlimited computing power.

* Yorktown Heights, New York, NY 10598, USA.

[†] Supported in part by an NSERC Postgraduate Scholarship.

[‡] Département IRO, Université de Montréal, C.P. 6128, succursale "A", Montréal (Québec), Canada H3C 3J7.

[§] Supported in part by NSERC under grant A4107.

[¶] This work was performed while this author was visiting IBM Research.

** Physics Department, University of California at Los Angeles, Los Angeles, CA 90024, USA.

1 Introduction and History

Quantum cryptography has recently entered the experimental era. The first convincingly successful quantum exchange took place in October 1989. After a short historical review of quantum cryptography, we report on the new apparatus and the results obtained with it. These results extend the first report on the apparatus, which appeared in *SIGACT News* [4].

Quantum cryptography was born in the late sixties when Stephen Wiesner wrote “Conjugate Coding”. Unfortunately, this highly innovative paper was unpublished at the time and it went mostly unnoticed. There Wiesner explained how quantum physics could be used in principle to produce bank notes that would be impossible to counterfeit and how to implement what he called a “multiplexing channel”, a notion strikingly similar to what Rabin was to put forward more than ten years later under the name of “oblivious transfer” (in our opinion, it would be fair to give at least equal credit to Wiesner for the concept of oblivious transfer).

Fortunately, Charles H. Bennett knew Wiesner quite well and heard about his idea from the horse’s mouth. Nevertheless, it was only when he met Gilles Brassard that quantum cryptography was revived. This happened on the occasion of the 20th IEEE Symposium on the Foundations of Computer Science, held in Puerto Rico in October 1979. Following our discussion of Wiesner’s idea, we discovered how to incorporate the (almost new at the time) notion of public key cryptography, resulting in a CRYPTO ’82 paper [6]. This brought Wiesner’s paper back to life, and it was subsequently published in *SIGACT News* [17], together with a selection of papers from the earlier CRYPTO ’81 workshop (for which “real” proceedings were not published).

Initially, quantum cryptography was thought of by everyone (including ourselves) mostly as a work of science-fiction because the technology required to implement it was out of reach (for instance, quantum bank notes [6] require the ability to store a single polarized photon or spin-1/2 particle for days without significant absorption or loss of polarization). Unfortunately, the impact of the CRYPTO ’82 conference had left most people under the impression that everything having to do with quantum cryptography was doomed from the start to being unrealistic.

The main breakthrough came when Bennett and Brassard realized that photons were never meant to *store* information, but rather to *transmit* it (although it should be said that half of Wiesner’s original paper dealt precisely with the use of quantum physics for the transmission of information). This lead initially to the *self-winding reusable one-time pad* [5] which was still not very practical. Later, Bennett thought of the *quantum public key distribution channel* and Brassard designed the somewhat less realistic *quantum coin-tossing protocol* [1, 2]. Quantum cryptography was also picked up by other researchers. For instance, Crépeau and Kilian showed how the quantum channel could be used to implement oblivious transfer in a strong way (Wiesner’s original multiplexing channel could leak information on both channels), zero-knowledge protocols, and secure two-party computation [12, 11].

The principle of quantum cryptography has been described in major popular magazines such as *Scientific American* [15], *The Economist* [14], and *New Scientist* [13]. Also, Brickell and Odlyzko close their very thorough survey of recent (1988) results in cryptanalysis with these words: "If such systems [quantum cryptography] become feasible, the cryptanalytic tools discussed here [in their paper] will be of no use" [10].

In this paper, we report on the first experimental quantum public key distribution channel ever designed and actually put together. Although we assume that the reader is already familiar with the principles of quantum cryptography, the following section should provide sufficient background. (A good description of the quantum channel itself can be found in chapter 6 of [9].)

2 Quantum Public Key Distribution

The purpose of public key distribution is for two users "Alice" and "Bob", who share no secret information initially, to agree on a random key, which remains secret from an adversary "Eve", who eavesdrops on their communications. In conventional cryptography and information theory it is taken for granted that digital communications can always be passively monitored, so that the eavesdropper learns their entire contents, without the sender or receiver being aware that any eavesdropping has taken place. By contrast, when digital information is encoded in non-orthogonal states of an elementary quantum system, such as single photons with polarization directions 0, 45, 90 and 135 degrees, one obtains a communications channel with the property that its transmissions cannot in principle be reliably read or copied by an eavesdropper ignorant of certain key information used in forming the transmission. The eavesdropper cannot even gain partial information about such a transmission without disturbing it in a random and uncontrollable way likely to be detected by the channel's legitimate users.

The protocol we describe here is secure even against an enemy possessing unlimited computing power (even if $\mathcal{P} = \mathcal{NP}$!), under any attack in which she is limited to measuring photons (or in the subsequent generalization, light pulses) one at a time, and combining the classical results of these measurements with information subsequently overheard during the public discussion. The formalism of quantum mechanics allows a more general kind of measurement, completely infeasible at present or in the foreseeable future. Such a measurement would treat the entire sequence of n photons sent during a key-distribution session as a single 2^n -state quantum system, cause it to interact coherently with an intermediate quantum system of comparable complexity, maintain the phase coherence of the intermediate system for an arbitrarily long time, then finally measure the intermediate system in a way depending on the information overheard during the public discussion. It is not known whether the protocol is secure against such an attack.

We first review the original quantum public key distribution (QPKD) protocol of [2], which illustrates the method most plainly. Then, we describe subsequent mod-

ifications of the protocol [7, 8, 3], which give it the ability, necessary in practice, to function despite partial information leakage to the eavesdropper and partial corruption of the quantum transmissions by noise. Finally, we describe the physical apparatus by which QPKD has actually been carried out. The essential quantum property involved, a manifestation of the uncertainty principle, is the fact that any measurement of a single photon's rectilinear (0 vs 90 degree) polarization randomizes its diagonal (45 vs 135 degree) polarization, and vice versa.

The basic QPKD protocol begins with Alice sending a random sequence of the four kinds of polarized photons to Bob. Bob then chooses randomly and independently for each photon (and independently of the choices made by Alice, of course, since these choices are unknown to him at this point) whether to measure the photon's rectilinear or diagonal polarization. Bob then announces publicly which kind of measurement he made (but not the result of the measurement), and Alice tells him, again publicly, whether he made the correct measurement (i.e. rectilinear for a 0 or 90 degree photon, diagonal for a 45 or 135 degree photon). Alice and Bob then agree publicly to discard all bit positions for which Bob performed the wrong measurement. Similarly, they agree to discard bit positions where Bob's detectors failed to detect the photon at all—a fairly common event with existing detectors at optical wavelengths. The polarizations of the remaining photons should be shared secret information between Alice and Bob, provided that no eavesdropping on the quantum channel has taken place. In the basic protocol, Alice and Bob next test for eavesdropping by publicly comparing polarizations of a random subset of the photons on which they think they should agree. In [2] it is shown that any measurement the eavesdropper makes on one of these photons while it is in transit from Alice to Bob has a $1/4$ chance of inducing a discrepancy when the data of Bob and Alice are compared, assuming that this photon is detected in the correct basis by Bob (otherwise, this photon is lost to all parties). If Alice and Bob find no discrepancies, they may safely conclude that there are few or no errors in the remaining uncomparing data, and that little or none of it is known to any eavesdropper.

The elementary protocol described above is inadequate in practice for two reasons:

1. Realistic detectors have some noise; therefore, Alice's and Bob's data will differ even in the absence of eavesdropping. Accordingly, they must be able to recover from a reasonably small error frequency.
2. It is technically difficult to produce a light pulse containing exactly one photon. It is much easier to produce a coherent pulse, which may be regarded as a superposition of quantum states with 0, 1, 2... photons; or an incoherent pulse, which may be regarded as a statistical mixture of coherent states. In either case, let μ be the expected number of photons per pulse. If μ is small (i.e. significantly less than 1), there is a probability approximately $\mu^2/2$ that an eavesdropper will be able to split a pulse into two or more photons, reading one and allowing the other(s) to go to Bob.

Below we briefly describe a practical protocol that allows these difficulties to be overcome. Further details may be found in [3, 7, 8].

The first task is for Alice and Bob to exchange public messages enabling them to reconcile the differences between their data, while revealing to Eve as little information as possible. We assume throughout that Eve listens to all the public messages between Bob and Alice.

An effective way for Alice and Bob to do this is for them first to agree on a random permutation of the bit positions in their strings (to randomize the locations of errors), then partition the permuted strings into blocks of size k such that single blocks are believed to be unlikely to contain more than one error. For each such block, Alice and Bob compare the block's parity. Blocks with matching parity are tentatively accepted as correct, while those of discordant parity are subject to a bisection search, disclosing $\log(k)$ further parities of sub-blocks, until the error is found and corrected. If the initial block size was much too large or too small, due to a bad a priori guess of the error rate, that fact will become apparent, and the procedure can be repeated with a more suitable block size. In order to avoid leaking information to Eve during the reconciliation process, Alice and Bob agree to discard the last bit of each block or sub-block whose parity they have disclosed. It is easy to see that Eve cannot know more information about the remaining truncated block than she did about the whole block before disclosure of its parity. However, because the string gets shorter, Eve's *proportion* of known information increases.

Of course, even with an appropriate block size, some errors will typically remain undetected, having occurred in blocks or sub-blocks with an even number of errors. To remove additional errors, the random permutation and block parity disclosure is repeated several more times, with increasing block sizes, until Alice and Bob estimate that very few errors remain in the data as a whole. At this point a different strategy is adopted to eliminate any errors that may remain and to verify, with high probability, that they have in fact been eliminated.

In each iteration of this strategy, Alice and Bob compare parities of a publicly chosen random subset of the bit positions in their entire respective data strings. If the data strings are not identical, then the random-subset parities will disagree with probability $1/2$. If a disagreement is found, Alice and Bob undertake a bisection search, similar to that described above, to find and remove the error. As in the preceding block-parity stage of the reconciliation, the last bit of each compared subset is discarded to avoid leaking any information to Eve. Each subsequent random subset parity is, of course, computed with a new independent random subset of bit positions in the remaining string.

At some point, all errors will have been removed, but Alice and Bob will not yet be aware of their success. When this occurs, subsequent random subset parities will of course always agree. After the last detected error, Alice and Bob continue comparing random subset parities until sufficiently many consecutive agreements (say 20) have been found to assure them that their strings are indeed identical, with a negligible probability of not detecting the existence of remaining errors.

Alice and Bob are now in the possession of a string that is almost certainly shared, but only partly secret. As described in the next section, they can find a conservative upper bound on Eve's partial information on their string from the detected error frequency, the optical pulse intensity, and the 0/1 ratio of the received string. If their reconciled string x has length n , and if they estimate that Eve knows at most k bits about it (these need not be k physical bits, but any k bits of information about the string), it is shown in [8] that for any security parameter $s > 0$, a hash function h randomly and publicly chosen from an appropriate class of functions $\{0,1\}^n \rightarrow \{0,1\}^{n-k-s}$ will map their string into a value $h(x)$ about which Eve's expected information is less than $2^{-s}/\ln 2$ bit. An adequate hash function for this purpose can be obtained by continuing to compute $n - k - s$ additional independent random subset parities, but now keeping their values secret instead of revealing them. The class of hash functions thus realized is essentially the strongly-universal₂ class H3 discussed by Wegman and Carter [16].

3 Physical Apparatus

The apparatus occupies an optical bench approximately one meter long inside a light-tight box measuring approximately $1.5 \times .5 \times .5$ meters. It is controlled by a program running on an IBM PC computer, which contains separate software representations of the sender Alice, who controls the sending apparatus, the receiver Bob, who controls the receiving apparatus, and optionally an eavesdropper Eve. The program can also run in simulation mode, without the attached experimental apparatus. Even though they reside in the same computer, no direct communication is allowed between the software Alice and the software Bob, except the public channel communication called for by the protocol.

Alice's light source, at the left end of the optical bench, consists of a green light-emitting diode (Stanley type HBG5566X) as the source of incoherent light, a 50 micron pinhole and 25 mm focal length lens to form a collimated beam, a 500 ± 20 nm interference filter (Ealing type 45-5040) to reduce the intensity and spectral width of the light and select a portion of the spectrum at which the photomultipliers have relatively high quantum efficiency, and finally a Polaroid filter (i.e. a dichroic sheet polarizer) to polarize the beam horizontally. The LED is driven by current pulses (about 10^{-7} coulombs in 50 nanoseconds) yielding, after collimation, filtration and polarization, an intensity of about 0.1–0.2 photon per pulse. The low intensity serves to minimize the chance that an eavesdropper will be able to split any one pulse into two or more photons.

Alice modulates the polarization of the beam by means of two Pockels cells (IN-RAD type 102-020), operated at + or – the quarter-wave voltage (about 800 volts), so as to be able to choose among the four polarization states {horizontal, vertical, left-circular, or right-circular} (circular polarizations are used instead of diagonal because they require only half the Pockels cell voltage). High voltage NPN transistors (type BU-205), in series with 200K ohm pull-up resistors, are used to switch the high

voltage for the Pockels cells under control of low voltage TTL signals on output lines of the PC's parallel port (5.1 volt Zener diodes protect the computer from exposure to high voltage in case of transistor failure).

The quantum channel itself is a free air optical path of approximately 32 centimeters.

Bob's receiving apparatus, at the right end of the optical bench, consists of another Pockels cell and a calcite Wollaston prism (Melles-Griot type 03PPW001/C), oriented so as to split the beam into vertically and horizontally polarized beams, which are directed into two photomultiplier tubes (Hamamatsu type R1463-01) with integral preamplifiers and voltage dividers in the sockets (Hamamatsu type C716-05). Bob's Pockels cell is also operated at quarter wave voltage, allowing him to use the same Wollaston prism to make a measurement of either rectilinear or circular polarization, depending on whether the voltage is off or on.

The timing for each experiment is controlled by a timing and detection unit, which also contains the hardware for handling asynchronous communication with the PC's parallel port, and two potentiometers for setting the discrimination levels for rejecting small pulses from each photomultiplier preamplifier (no rejection of large pulses is necessary, owing to their infrequency). The pulse-height discrimination is carried out by fast ECL voltage comparators (Plessey type SP9687).

Upon receiving a "start" signal on one of the PC parallel port's output lines, the timing unit waits 60 μ sec for the Pockels cell voltages to settle, turns the LED on for about 75 ns, gates the photomultiplier detection logic on for about 100 ns, and sets two input lines of the parallel port according to the result (for each photomultiplier, whether a count was detected during the gate interval). When it has done all this, the timing unit turns on another of the parallel port's input lines to signify "done", and begins waiting for the next start signal. When the computer sees the done signal it knows it can read the results of the present experiment and thereafter safely start the next experiment.

Alice's choice of polarization and Bob's choice of reading basis are made randomly (not pseudorandomly) using a large file of random bits supplied to the computer on a diskette. Of course, Alice and Bob feed on different bits from this diskette (recall that although they live on the same computer, they do not communicate or otherwise share information that is not called for by the public channel discussion). These random bits had been previously generated using the same experimental apparatus, by taking the physically random output of one of the photomultipliers, illuminated by an auxiliary nearby LED of intensity such as to yield a count in about 1/2 the time windows, removing the 0/1 bias by von Neumann's trick (i.e. in each consecutive pair of tosses taking $HT=1$, $TH=0$, ignoring HH and TT), and XORing the resulting bits with pseudorandom bits from the computer to hide any residual deviations from randomness caused by time-variation of the photomultiplier and pulse-detection circuit. The same file is used to supply additional random bits as needed by Alice and Bob during the data reconciliation and privacy amplification protocols described in the previous section.

The photomultipliers had quantum efficiency approximately 9%, with dark count rates of about 1500 per second, or about 0.00015 per 100 ns time window. When using pulses of 0.15 expected photons per pulse, this dark count rate would yield a bit error rate of approximately 1%; the actual error rate, about 5%, was due primarily to imperfect alignment of the Pockels cells.

The driver program on the PC provides the ability to simulate two principal kinds of eavesdropping: intercept/resend and beamsplitting, by a hypothetical adversary "Eve" who has detectors of 100% quantum efficiency.

Recall that μ is the expected number of photons per light pulse. If μ is sufficiently smaller than 1, it is approximately also the probability that a pulse would be detected by a perfectly efficient detector. In intercept/resend, Eve intercepts a light pulse and reads it in a basis of her choosing (she cannot be sure of choosing the correct basis, which has not yet been announced). If, with probability approximately μ , she is successful in detecting a photon, she fabricates and sends to Bob a pulse of the same polarization as she detected. It can be shown that the canonical bases, rectilinear or circular, are optimal for Eve to use in this attack, yielding an expected information 1/2 bit per intercepted photon, and inducing an error with probability 1/4 if the fabricated pulse is later detected in the correct basis by Bob. Other bases for Eve yield less information, induce more errors, or both. To avoid suspicion, Eve's fabricated pulses should be of such intensity (slightly higher than one expected photon per pulse) as to yield the same net rate of pulse detection by Bob as if no eavesdropping were taking place.

No additional hardware is needed to simulate this attack: when the software Eve wishes to intercept a pulse, she borrows the real receiving apparatus from Bob; when she wishes to resend to Bob, she borrows the sending apparatus from Alice. While Eve is borrowing the receiving apparatus, Alice obliges her by repeating the same transmission $1/q$ times, where q is the quantum efficiency of the actual detectors. This allows the software Eve to obtain a count with the same probability μ as a physical eavesdropper with perfectly efficient detectors.

The other attack, beamsplitting, would be technically easy for a real Eve, and depends on the fact that the transmitted light pulses are not pure single-photon states. To carry out this attack Eve would use a partly-silvered mirror or equivalent device to divert a fraction f of the original beam's intensity to detectors of her own, letting the remainder pass undisturbed to Bob. With probability approximately $f\mu/2$, Eve will succeed in detecting a photon, and will have by good luck measured it in the correct basis. This attack induces no errors, but does attenuate the intensity reaching Bob by a factor $1 - f$. If Eve is in control of the channel between Alice and Bob, and if this channel has significant attenuation, she can conceal the attenuation due to her beamsplitting by substituting a more transparent channel. Assuming conservatively that she can do this, she will divert most of the beam to herself, and learn a fraction roughly $\mu/2$ of the polarizations later correctly measured by Bob. An Eve with superior technology might be able to store her portion of the split beam and delay measuring it until after the correct bases were announced, thereby doubling her

information yield. On the other hand, if Alice and Bob suspected Eve of having this capability, they could send and receive all the pulses first, wait an arbitrarily long time for Eve's stored beam to decay, and only then announce all the bases.

A dramatic but harmless variant on the beamsplitting attack would be for Eve to attempt to detect enough photons in the incoming pulse to determine its polarization uniquely, even without knowing the correct basis. An example of such a measurement would be for Eve to further split the intercepted portion of beam into two beams of intensity $f\mu/2$, and measure the rectilinear polarization of one and the circular polarization of the other. If, by extreme good luck, this measurement yielded three photons with polarizations vertical, horizontal, and right circular, Eve would know that the original pulse's polarization was definitely right-circular, and she could capitalize on this knowledge by sending Bob such a bright pulse of right-circular light that he would be sure to detect it. Fortunately this attack succeeds so rarely (roughly with probability $\mu^3/32$) that it is a less serious threat than simple 2-photon beamsplitting.

The driver program simulates beamsplitting simply by having the software Alice disclose directly to the software Eve the correct polarizations of a fraction $\mu/2$ of the pulses.

The expected information leaked to Eve through both kinds of eavesdropping is bounded above by

$$k = N(\mu/2 + 2p) \text{ bits}, \quad (1)$$

where N is the number of Alice's pulses received by Bob in the correct basis, μ is the pulse intensity at the upstream end of the channel, and p is the bit error rate. This estimate assumes that Eve has been able to manipulate the channel attenuation as described above to maximize her share f , that she does not have the superior technology required to delay measurement until after announcement of the correct bases (if she did, the first term above would be increased from $\mu/2$ to μ), and that intercept/resend eavesdropping is the only cause of transmission errors. These assumptions will in most cases be excessively conservative: e.g. in our case, the channel has negligible attenuation and many of the bit errors can be confidently attributed to causes other than eavesdropping. In our experiments, Eve has a small additional source of information: the excess 0:1 ratio (about 62/38) in the received data resulting from an inequality of quantum efficiency between Bob's two photomultipliers. This imbalance gives Eve a small additional amount of information (about 0.03N bits) on Alice's and Bob's string. This leakage could have been prevented by using photomultipliers of equal sensitivity, or by having Bob randomly permute the photomultipliers between measurements.

The present apparatus is only an experimental prototype. In a more realistic demonstration, the error rate could be reduced several orders of magnitude by better optical alignment and cooling the photomultipliers to reduce dark current, the quantum channel could be made much longer (e.g. a few km of optical fiber), and the protagonists Alice, Bob, and Eve could reside in separate buildings [3]. The feasible

distance over which a QPKD system can operate depends on the noise and quantum efficiency of the detectors and especially on the attenuation of the optical channel: the weak signal entering the channel must still be recognizable above background upon leaving the channel.

4 Sample Data from the Apparatus

Here we give examples of data actually transmitted through the quantum channel, the subsequent public discussion, and the shared secret key ultimately distilled. The first batch of data is from a run in which there was in fact no eavesdropping, but the eavesdropper's potential information was nevertheless conservatively estimated as described above, from the known pulse intensity and error rate. The second batch of data illustrates the ability to distill a small amount of shared secret key from a run with significant amounts of both kinds of eavesdropping.

Here is raw data obtained from the quantum channel on Friday, April 13, 1990.

Alice 0101010010 1010001011 1100010100 0110001100 0001111001 1001001100 0001100000 1101000100 0001000100 0010100000
 1011011001 0001000100 0000000100 1010001011 0011010101 0101010010 0000100011 1110000001 0101000001 0011010000
 0111000100 0000011100 1100110100 0000101011 0000100001 1110000001 1100010000 0000100010 0010010110 0110000101
 0111000000 1010011110 1101100111 0000000000 0010010100 0000000001 0000110110 0010010001 0001011110 1101100101
 0101011000 1100100001 0000000100 1111111110 0010011010 0011000010 0111110000 0011000000 0000001010 1100010110
 1010001010 0101000010 1110011000 0011111000 1100100011 1000100000 0001000101 1000101100 1101010111 0111011010
 1001100101 0010000010 1000001100 0001110100

Bob 0101000010 1011001011 1100010100 0110001100 0001111001 1001000100 0001100000 1101000100 0001000100 0010100000
 1011011001 0001000000 0000000100 1010001010 1011000101 0101010010 1000100011 1110000001 0101000011 0011000010
 0111000100 0000111100 1100110100 0000101011 0000000000 1110000001 1100010000 0000110010 0010010100 0110000101
 0111010000 1000011110 1101100111 0000000000 0010010100 0000000001 0000110110 0010010101 0001011110 1101100101
 0101011000 1100100001 0000000100 1011011110 0010011010 0011000010 0111110000 0011000000 0100001010 0100010110
 1010001010 0101000010 1110011000 0011111000 1100100011 1000100000 0001000101 1000101100 1100010101 0101011010
 1001100101 0010100010 1000001000 0001110100

In this first example, out of about 85,000 light pulses of intensity 0.17 sent by Alice, 640 were received in the correct basis by Bob. Alice's corresponding string contained 242 ones out of 640 bits. Bob's string contained 28 errors, an error frequency of 4.375 %.

A random permutation and block parity comparison was performed with block size 10, reducing the string length to 509 bits with 8 undiscovered errors.

A second random permutation and block parity comparison was performed with block size 20, reducing the string length to 457 bits with 2 undiscovered errors.

Random subset parity comparison was then begun, revealing an error on the first attempt. Removal of this error by bisective search reduced the string length to 448 bits.

Another random subset parity was computed, revealing another error. Removal of this error reduced the string length to 439 bits with no undiscovered errors.

Twenty more random subset parities were compared and found to agree, confirming to Bob and Alice that with high probability their remaining strings, now 419 bits long, were identical.

From the 28 errors corrected during reconciliation, Bob and Alice estimated that the original error rate was 4.50 %. Estimated potential information leakage to Eve was 140 bits, including

- 58 bits from intercept/resend,
- 54 bits from beamsplitting, based on pulse intensity $\mu = 0.17$, and
- 28 bits from redundancy due to 0:1 imbalance (398/242) of initial string.

Therefore, allowing 60 bits excess compression for safety (e.g. in case Eve was especially lucky in her eavesdropping, and obtained several standard deviations more information than expected), it was decided to compress the string 200 bits by random subset hashing, leaving 219 bits of shared secret key distilled from 640 original bits.

The resulting secret key, the same for Alice and Bob, was

```
0000101010 1101100010 0100101100 0110010010 1000010100
1110011101 1001000011 1111101010 0000111010 0011111100
1100101000 1101011111 1110001101 0001100100 1000110011
0101110110 0011110110 1010100100 1011111010 0101111101
0110000000 010000101.
```

In the second example, out of another approximately 85,000 light pulses of intensity $\mu = 0.17$ sent by Alice, 640 were received in the correct basis by Bob. Alice's corresponding string contained 239 ones out of 640 bits. Bob's string contained 59 errors, an error frequency of 9.219 %. Through attempting to beamsplit all the pulses, and intercept/resending one sixth of them, the simulated Eve learned 100 individual bits of Alice's data as well as knowing 30 bits of distributed information about the string as a whole due to its 0/1 imbalance. (Her total information was actually slightly less than 130 bits, because of the correlation between these two kinds of information. As remarked earlier, Eve's absolute amount of information does not increase during reconciliation).

A random permutation and block parity comparison was performed with block size 5, reducing the string length to 399 bits with 13 undiscovered errors. (They start with a block size smaller than in the previous example because of reason (1) given at the end of this section.) Eve's information about the remaining string was still less than 130 bits, and included knowledge of 61 individual bits.

A second random permutation and block parity comparison was performed with block size 10, reducing the string length to 337 bits with 6 undiscovered errors. Eve's information about the remaining string was still less than 130 bits, and included knowledge of 51 individual bits.

A third random permutation and block parity comparison was performed with block size 20, reducing the string length to 292 bits and leaving no undiscovered errors. Eve's information about the remaining string was still less than 130 bits, and included knowledge of 43 individual bits.

Random subset parity comparison was then begun: Twenty consecutive successful comparisons with no failures convinced Alice and Bob that their strings, now consisting of 272 bits, were very probably identical. Eve's knowledge about the remaining string was still less than 130 bits, and included knowledge of 38 individual bits.

From the errors found and corrected, Alice and Bob estimated the error probability had been 9.59 % in the original string. From this error rate and from the known pulse intensity $\mu = 0.17$, Alice and Bob computed an upper bound of 207 bits on Eve's probable information, including:

- 123 bits from intercept/resend,
- 54 bits from beamsplitting,
- 30 bits from redundancy due to 0:1 imbalance.

Therefore, allowing 20 bits excess compression (about all we can afford), it was decided to compress the string 227 bits by random subset hashing, leaving 45 bits of shared secret key, distilled from 640 original bits. Since Eve's actual information was less than 130 bits, this amount of compression left Eve with an utterly negligible (less than 10^{-29} bit) expected information about the output of the hash function.

The resulting secret key, the same for Alice and Bob, was

0001110110 0011101001 1000100011 1111000010 10010.

The 640-bit batch size used above for illustrative purposes is far from optimal. In production use, a larger batch size (at least 10,000 bits) should be used for two reasons: 1) It would allow the users, by preliminary sampling, to get a good estimate of the bit error rate and so optimize the choice of block sizes used in the reconciliation stage; and 2) by reducing the statistical uncertainty in estimating Eve's possible information, it would reduce the proportional amount of compression needed in the privacy amplification stage to assure a given level of security.

Acknowledgements

We wish to thank Manuel Blum, Claude Crépeau, David Deutsch, Myron Mandel, and Stephen Wiesner for many helpful discussions.

References

- [1] Bennett, C.H. and G. Brassard, "An update on quantum cryptography", *Advances in Cryptology: Proceedings of Crypto '84*, August 1984, Springer-Verlag, pp. 475-480.
- [2] Bennett, C.H. and G. Brassard, "Quantum cryptography: Public key distribution and coin tossing", *Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing*, Bangalore, India, December 1984, pp. 175-179.
- [3] Bennett, C.H. and G. Brassard, "Quantum public key distribution system", *IBM Technical Disclosure Bulletin*, Vol. 28, 1985, pp. 3153-3163.
- [4] Bennett, C.H. and G. Brassard, "The dawn of a new era for quantum cryptography: The experimental prototype is working!", *SIGACT News*, Vol. 20, no. 4, Fall 1989, pp. 78-82.
- [5] Bennett, C.H., G. Brassard and S. Breidbart, "Quantum cryptography II: How to re-use a one-time pad safely even if $P = NP$ ", unpublished manuscript available from the authors, November 1982.
- [6] Bennett, C.H., G. Brassard, S. Breidbart and S. Wiesner, "Quantum cryptography, or unforgeable subway tokens", *Advances in Cryptology: Proceedings of Crypto '82*, August 1982, Plenum Press, pp. 267-275.
- [7] Bennett, C.H., G. Brassard and J.-M. Robert, "How to reduce your enemy's information", *Advances in Cryptology — Crypto '85 Proceedings*, August 1985, Springer-Verlag, pp. 468-476.
- [8] Bennett, C.H., G. Brassard and J.-M. Robert, "Privacy amplification by public discussion", *SIAM Journal on Computing*, Vol. 17, no. 2, April 1988, pp. 210-229.
- [9] Brassard, G., *Modern Cryptology: A Tutorial*, Lecture Notes in Computer Science, Vol. 325, Springer-Verlag, Heidelberg, 1988.
- [10] Brickell, E.F. and A.M. Odlyzko, "Cryptanalysis: A survey of recent results", *Proceedings of the IEEE*, Vol. 76, no. 5, May 1988, pp. 578-593.
- [11] Crépeau, C., "Correct and private reductions among oblivious transfers", PhD Thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, February 1990.
- [12] Crépeau, C. and J. Kilian, "Achieving oblivious transfer using weakened security assumptions", *Proceedings of 29th IEEE Symposium on the Foundations of Computer Science*, White Plains, New York, October 1988, pp. 42-52.
- [13] Deutsch, D., *New Scientist*, December 9, 1989, pp. 25-26.
- [14] Gottlieb, A., "Conjugal secrets — The untappable quantum telephone", *The Economist*, Vol. 311, no. 7599, 22 April 1989, p. 81.
- [15] Wallich, P., "Quantum cryptography", *Scientific American*, Vol. 260, no. 5, May 1989, pp. 28-30.
- [16] Wegman, M.N. and J.L. Carter, "New hash functions and their use in authentication and set equality", *Journal of Computer and System Sciences*, Vol. 22, 1981, pp. 265-279.
- [17] Wiesner, S., "Conjugate Coding", manuscript written circa 1970, unpublished until it appeared in *SIGACT News*, Vol. 15, no. 1, 1983, pp. 78-88.