Defense frontier analysis of quantum cryptographic systems

Boris Slutsky, Ramesh Rao, Pan-Cheng Sun, Ljubiša Tancevski, and Shaya Fainman

When a quantum cryptographic system operates in the presence of background noise, security of the key can be recovered by a procedure called key distillation. A key-distillation scheme effective against so-called individual (bitwise-independent) eavesdropping attacks involves sacrifice of some of the data through privacy amplification. We derive the amount of data sacrifice sufficient to defend against individual eavesdropping attacks in both BB84 and B92 protocols and show in what sense the communication becomes secure as a result. We also compare the secrecy capacity of various quantum cryptosystems, taking into account data sacrifice during key distillation, and conclude that the BB84 protocol may offer better performance characteristics than the B92.

OCIS codes: 060.4570, 270.0270, 060.4510.

1. Introduction

Quantum cryptography is a technique that permits two parties, Alice and Bob, to establish between themselves a shared secret key that cannot be compromised by an eavesdropper, Eve. This security is derived from encoding the data on nonorthogonal quantum states of a physical carrier particle. Since such quantum states cannot be duplicated or analyzed in transit without disturbing them, any attempt to interfere with the particle introduces transmission errors and thereby reveals itself to Alice and Bob.[1–5]

Unfortunately, in practice a physical connection always has losses and errors of its own, and therefore Alice and Bob can never be certain that Eve is not tampering with their communication. In the presence of noise, the sequence leading to a secure key is more complicated. Alice and Bob must first locate and discard the errors, for example, by publicly exchanging a series of block checksums. They then assume that all errors are eavesdropping induced and estimate Eve's potential knowledge of their data in this worst-case situation. Last, Alice and Bob apply a scrambling algorithm to extract from the data a shorter but secure key. The three-step procedure just described, which is referred to as key distillation,[4,6] requires, particularly in its second step, a careful investigation of possible eavesdropping attacks.

A necessary part of any eavesdropping attack is an interaction between Alice's particle and Eve's own quantum system, called a probe. The particle and the probe emerge from this interaction in an entangled quantum state, which is jointly measured by Bob and Eve. A strategy in which the eavesdropper performs a separate, independent measurement for each intercepted particle is called an individual attack. Alternatively, Eve can use a separate probe for each particle and later measure all probes together as a single quantum system, taking advantage of checksums and other group information Alice and Bob would have revealed in the course of key distillation. Even more generally, a single probe can be entangled with the entire set of particles. Attacks of the last two kinds are known respectively as collective and joint. Although quantum cryptography is believed to be secure against such attacks, there appears to be disagreement among researchers whether published results amount to a formal proof of this claim.[7–9] To our knowledge, no specific joint attacks have yet been proposed.[7]

In this paper we consider only individual attacks, of which the theory is more advanced. A key-distillation technique that can, at least in principle, defend against individual attacks was first outlined in Ref. 4 and is summarized in Fig. 1.[6] Starting from raw data obtained in the course of quantum trans-

The authors are with the Department of Electrical and Computer Engineering, University of California, San Diego, La Jolla, California 92093.

Received 11 August 1997; revised manuscript received 5 January 1998.

0003-6935/98/142869-10$15.00/0

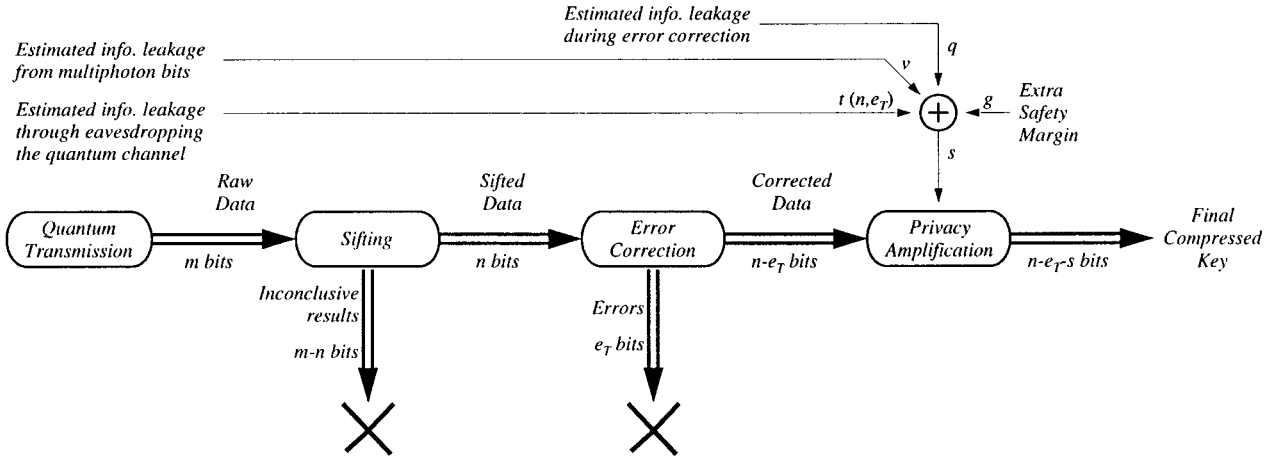

Fig. 1. Distillation of a secret key from a quantum transmission. Alice and Bob arrive at privacy-amplification compression level s by summing estimates of possible information leakage at various stages of the protocol, together with an arbitrary safety margin.

mission, Alice and Bob first discard so-called inconclusive bits.[10] They then exchange a series of block checksums, and where the checksums do not match, use a bisective search within the block to identify and discard the error. The resulting corrected data are input into the (so-called classical) privacy-amplification algorithm,[11] which produces a shorter but more secure key. Specifically, the privacy-amplification theorem (Ref. 11, corollary 4) asserts that if Eve's Renyi information on an l-bit data string does *not* exceed some quantity r, then her Shannon information (averaged over the choice of the privacy-amplification hash function) on the reduced $(l - s)$-bit string does not exceed $2^{r-s}/\ln 2$. Here the term "Renyi information on a w-bit string X" refers to the quantity

$$I^R = w + \log_2 E[P_X(X)] = w + \log_2 \sum_X P_X^2(X), \quad (1a)$$

where $P_X(X)$ is the probability distribution of X (as seen by Eve) and $E[\cdots]$ denotes the expected value. Similarly, "Shannon information on a w-bit string X" refers to

$$I^H = w + E[\log_2 P_X(X)] = w + \sum_X P_X(X) \log_2 P_X(X). \quad (1b)$$

Thus the secrecy of the final key (in the Shannon information sense) can be recovered, albeit at the expense of reducing its size, if Alice and Bob can upper-bound Eve's Renyi information I^R on the corrected data.

Eve's information can come from a number of sources, three of which are indicated in Fig. 1. (A particular technical implementation may have additional vulnerabilities.) Information leakage from multiphoton bits reflects the anticipated beam-splitting attacks on bit cells in which Alice's device emitted multiple photons[4,6,12–14]; this term is the easiest to bound because it cannot exceed the total number v of affected bit cells, which is determined by the parameters of Alice's transmitter.[15]

Information leakage during error correction is caused by the disclosure of block checksums. Using the technique from Ref. 4, Alice and Bob divide their data string into blocks of length l and reveal the parity of each block. A disagreement in a particular block triggers the bisective-search procedure illustrated by example in Fig. 2, in which a single error (bit 8) is located in a data block of $l = 16$ bits by the disclosure, in succession, of the parities of bits 1–16, 1–8, 1–4, 5–6, and 7. We note in passing that, if a block contains an odd number of errors greater than one, only one error would be found by the bisective search, and any blocks with an even number of errors would be entirely overlooked. Nonetheless, Alice and Bob can ultimately detect all errors with high probability if they eliminate any errors they do find, randomly permute the remaining bits, and repeat the same procedure, perhaps with a different block size l.[4]

It is easily verified that, regardless of the position of the error within the block, the kind of bisective search illustrated in Fig. 2 entails disclosure of a total of $\log_2 l$ parity bits plus the value of the error bit itself. Because of this last circumstance, Eve derives no benefit from any knowledge of error bits that she may have obtained previously through quantum eavesdropping, and Alice and Bob have no interest in maintaining these bits in their data string. It is

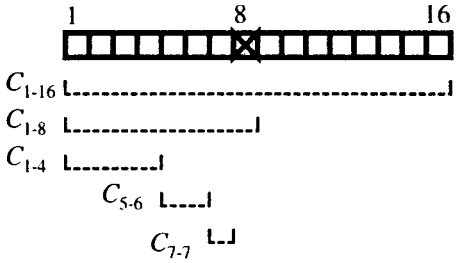

Fig. 2. Error correction by bisective search. The error bit (marked by a cross) is identified by public comparison of parities of the subblocks indicated by the dashed lines.

henceforth assumed that error bits are discarded before privacy amplification, as indicated in Fig. 1.

Note that Eve's Renyi information gain from a public disclosure may in principle exceed the number of bits disclosed. Fortunately, it has been shown to be close to that number in the special case in which the data consist of a lengthy sequence of independent and independently attacked bits and the disclosure is in the form of block parities.[16]

In the sequel, we accept as given the upper bounds ν and q on the information leakage from multiphoton bits and from error correction without questioning how Alice and Bob might arrive at these bounds. Instead, the focus is on the third information leakage source, marked in Fig. 1 as "eavesdropping the quantum channel," which reflects Eve's gains from measurements of her probe after the probe has interacted with Alice's particle. As explained above, such gains are inevitably accompanied by a disturbance of the particle state and result in transmission errors. The upper-bound estimate of information leakage through eavesdropping on the quantum channel is denoted by $t(n, e_T)$ in Fig. 1. As this notation implies, Alice and Bob determine the estimate on the basis of the size of the sifted data n and the number of errors e_T. (In principle, the number of inconclusive results $m - n$ could also play a role in the decision, but in the particular key-distillation scheme considered here it does not.) We call $t(n, e_T)$ the defense function because it is chosen by Alice and Bob to defend against an eavesdropping attack. It can be supposed without loss of generality that the defense function is agreed on in advance of the transmission. Furthermore, since Eve can always modify her strategy to induce additional errors deliberately, it would be illogical for Alice and Bob to ever reduce their estimate t of Eve's knowledge as the number of errors e_T increases. Discussion is therefore restricted to defense functions $t(n, e_T)$ that are monotonically nondecreasing with e_T for any fixed n.

The design of the defense function $t(n, e_T)$ is based on the trade-off relation between the average amount of information that Eve can extract from Alice's particle and the average degree of disturbance induced in the particle state. This trade-off has recently become the subject of intensive investigation.[17–19] Published results differ in their definitions of disturbance (raw-data error rate e_T/m in Ref. 17 versus the sifted-data error rate e_T/n in Refs. 18 and 19), as well as in their chosen measure of Eve's advantage (Shannon information on the raw data in Ref. 17 versus Shannon information on sifted data in Ref. 18 versus Renyi information on corrected data in Ref. 19).

The purpose of this paper is to supply the missing link between the average gain versus the average disturbance relation and the defense function $t(n, e_T)$. Specifically, we derive the defense function that adequately protects the final key against any given eavesdropping strategy for which the expected sifted-data error rate e_T/n and Eve's expected Renyi information on corrected data are known. We also formulate a precise statistical definition of the "adequate protection." Additionally, knowing both the raw throughput of a quantum cryptographic system and the defense function permits us to make a preliminary estimate of its secrecy capacity and compare such systems with one another. We find in particular that the four-state protocol BB84 promises better performance than does the two-state B92, while the optimal choice among different two-state systems depends on the channel-error rate.

This paper is organized as follows. In Section 2 a description of the eavesdropping process, which is needed to make the necessary definitions, is given and the underlying assumptions are discussed. In Section 3 we design and prove the safety of a defense function for a single eavesdropping strategy. Section 4 extends this treatment to multiple eavesdropping strategies, introducing the concept of defense frontier. It also uses the results from Ref. 19 to construct explicitly the defense frontier for both two-state and four-state systems and presents a comparison of the secrecy capacities of these systems. In Section 5 a brief discussion is given of an alternative definition of security in a quantum cryptographic channel, which is more difficult to satisfy than the definition used in Sections 3 and 4.

2. Procedure and Definitions

The two principal classes of quantum cryptographic schemes are known as BB84, or four state, and B92, or two state. In the four-state scheme,[1] Alice and Bob agree on two mutually conjugate orthogonal bases $\mathcal{B}_1$ and $\mathcal{B}_2$ in a two-dimensional plane in the Hilbert state space of a quantum particle. Alice transmits each data bit by first selecting at random one of the bases, $\mathcal{B}_1$ or $\mathcal{B}_2$, and then preparing the particle in one of the two basis states of her chosen basis, depending on the value of the bit (0 or 1). Bob, on his part, also selects at random one of the bases, $\mathcal{B}_1$ or $\mathcal{B}_2$, independently of Alice and performs a measurement on Alice's particle in his chosen basis. At the end of the transmission both Alice and Bob publicly announce the bases they used. In those cases in which Bob's choice of basis matches Alice's, Bob will have correctly received Alice's emitted data, absent eavesdropping and noise in the channel.

Similarly, in two-state quantum cryptography,[5] Alice and Bob agree on two nonorthogonal particle states $\mathbf{u}$ and $\mathbf{v}$. Alice transmits each bit of data by preparing the particle in one of the two states, depending on the value of the bit (0 or 1). Although two nonorthogonal quantum states $\mathbf{u}$ and $\mathbf{v}$ cannot be reliably distinguished, Bob can perform a measurement that, at least in a fraction of cases, conclusively indicates state $\mathbf{u}$ or state $\mathbf{v}$. At the end of the transmission Bob publicly announces which bits have yielded such conclusive results in his apparatus. In these cases, Bob will have correctly received Alice's emitted data, absent eavesdropping and noise in the channel.

Our third actor, the eavesdropper Eve, has access to the channel and is using an apparatus that implements some specific eavesdropping strategy.

"Throw a die: if the die shows 1, measure and retransmit in basis $\mathcal{B}_1$; if the die shows 2, measure and retransmit in basis $\mathcal{B}_2$; if the die shows 3, measure and retransmit in the basis halfway between; otherwise do nothing and let the particle proceed" is an example of an eavesdropping strategy. One terminal of the apparatus accepts quantum particles from Alice, while the other terminal sends particles toward Bob.

As stated above, our analysis is limited to individual eavesdropping strategies, wherein Eve attacks all Alice's particles independently of one another. We further assume that the attacks on all particles are identical. The latter assumption is made without apparent loss of generality because both the preparation of a particular particle by Alice and its measurement by Bob are random and statistically independent of their preparation and measurement of other particles. Whether Eve decides to subject each passing particle to the same measurement or, as in the example strategy above, chooses to treat particles differently, she has nothing to gain by discriminating between them on a systematic as opposed to a random basis. Moreover, any discernible pattern in Eve's activities may result in a statistical asymmetry that could be noticed by Alice and Bob.

Let K_i represent the joint preparation of Alice's, Bob's, and Eve's apparata in advance of bit i, where $i = 1, 2, \ldots, m$. K_i can take values k that index all combinations of preparation options available to Alice, Bob, and Eve. For example, in BB84 Alice and Bob have eight preparation options, corresponding to bit values 0 and 1 encoded by Alice in each of the two bases $\mathcal{B}_1$ and $\mathcal{B}_2$ and measured by Bob in $\mathcal{B}_1$ or $\mathcal{B}_2$. In B92 Alice and Bob have four options, consisting of two possible measurements by Bob of each of the two nonorthogonal states **u** and **v** that might be transmitted by Alice. In a somewhat more advanced version of B92,[20] Bob's receiver need not be preset differently for different bits, so that the number of Alice–Bob configurations is reduced to two. Eve's options depend on the capabilities of her probe.

Given the initial configuration $K_i = k$ and assuming that the design of each party's apparatus is publicly known, the physical system becomes fully determined, and quantum mechanics can predict the result in terms of a joint probability distribution $P_{YZ}^{(k)}(y, z)$ of Bob's and Eve's measurement outcomes Y_i and Z_i, respectively. The assumption of independent bit-by-bit eavesdropping can now be restated as follows: It is assumed that the preparation-outcome random vectors $\{K_i, Y_i, Z_i\}$ are independent identically distributed (i.i.d.) as a function of bit position i.

For each bit the result Y_i in Bob's laboratory either matches or does not match the bit value transmitted by Alice or the result is inconclusive. Let C be the set of bits for which the result is not inconclusive, and for all $i \in C$ let $E_i = 0$ if Bob receives the bit correctly and $E_i = 1$ if an error occurs. Both inconclusive and erroneous results are eventually identified publicly and removed from the transmission. For each bit that produces $E_i = 0$, Eve analyzes her outcome Z_i (in the case of BB84, together with information regarding the encoding basis of bit i, which is eventually disclosed) and arrives at the quantity P_i, which represents, from her point of view, the probability of the bit value's being 1.[21] Eve's Renyi and Shannon information on the bit, in accord with Eqs. (1a) and (1b), are given by

$$I_i^R = 1 + \log_2[P_i^2 + (1 - P_i)^2],$$
$$I_i^H = 1 + P_i \log_2 P_i + (1 - P_i)\log_2(1 - P_i). \quad (2)$$

Random vectors $\{I_i^H, I_i^R, E_i\}$, $i \in C$, are i.i.d. from bit to bit, being a deterministic function of i.i.d. random vectors $\{K_i, Y_i, Z_i\}$. Because bit values are independent of one another from Eve's point of view (as well as in general), both information measures I_i^H and I_i^R are additive over the transmission. Hence Eve's Renyi and Shannon information on the corrected data as a whole and the total number of errors between Alice and Bob are given respectively by the random variables

$$I_T^R = \sum_{i \in C} (1 - E_i) I_i^R, \quad I_T^H = \sum_{i \in C} (1 - E_i) I_i^H, \quad E_T = \sum_{i \in C} E_i. \quad (3)$$

The quantity I_T^R represents the information leakage through eavesdropping on the quantum channel referred to in Section 1 and Fig. 1. The defense function $t(n, e_T)$ must serve as an upper bound on I_T^R. Unfortunately, I_T^R is a random variable for which the only deterministic upper bound is the total length $n - e_T$ of corrected data. Since this estimate is unacceptable to Alice and Bob (for it would make them throw away the entire transmission), it must be abandoned in favor of some lower, probabilistic bound, so that, with some small but nonzero probability, $I_T^R > t(n, e_T)$. Alice and Bob eventually combine $t(n, e_T)$ with upper bounds of information leakage from other sources and add an arbitrary safety margin g to arrive at the compression level s (see Fig. 1), which they believe would render the output of the privacy-amplification algorithm sufficiently secret. In the few cases in which $I_T^R > t(n, e_T)$, however, the final key may not be as secret as Alice and Bob believe it is. This naturally leads to

Definition 1. An eavesdropping attack is defined as successful if it introduces some number of errors e_T into an n-bit sifted-data string resulting from an m-bit quantum transmission, while yielding the attacker $I_T^R > t(n, e_T)$ total Renyi information on the $(n - e_T)$-bit corrected data.

Note that Alice and Bob do not seek to detect eavesdropping activity but rather to satisfy themselves that their shared data are sufficiently confidential. Indeed, the former objective is unattainable in the presence of noise and losses in the channel. An adversary in possession of superior technology can secretly improve channel quality or frustrate Alice and Bob's pretransmission estimation of channel quality and thereby camouflage eavesdropping-induced er-

rors as natural channel errors. Nor does the key-distillation framework offer protection against the kind of hostile activity that might make any transmission impossible. What it does offer is an assurance that, unless an eavesdropping attack is successful in the meaning of definition 1, any key ultimately produced is at least as secure (in terms of Shannon information) as Alice and Bob believe it to be.

Let S denote the event of a successful eavesdropping attack in the sense of definition 1. The task facing Alice and Bob is to construct a defense function $t(n, e_T)$ to minimize the chance of S. We consider three distinct mathematical interpretations of this goal:

(1) In terms of the *a priori* probability Prob[S] over all quantum transmissions of length m: This measure might be of most interest to Eve when deciding whether her monitoring effort is likely to pay off.

(2) In terms of the probability Prob[$S|N = n, E_T = e_T$] over all quantum transmissions of length m that result in n bits of sifted data with e_T errors, which can be viewed as *a posteriori* in the sense that it includes all relevant information known to Alice and Bob at the end of the transmission: This measure might be most valuable to Alice and Bob because it most closely relates to confidentiality of a specific transmission.

(3) In terms of the "intermediate" probability Prob[$S|N = n$] over all quantum transmissions of length m that result in n bits of sifted data.

We show that Alice and Bob can construct a defense function $t(n, e_T)$ such that both the *a priori* probability (1) and the "intermediate" probability (3) do not exceed an arbitrary small value $\alpha > 0$. This defense function is the subject of Section 3. The *a posteriori* probability (2), however, cannot be reduced to an arbitrary low level by any reasonable choice of defense function. This circumstance, which is discussed briefly in Section 5, does not imply that Alice and Bob cannot achieve security in the usual informal sense. Roughly speaking, as long as the *a priori* probability (1) is kept low, Alice and Bob have as much reason to believe that a particular transmission is secure as an observer who records an equal number of heads and tails after multiple tosses of a particular coin has reason to believe that the coin is fair.

3. *A Priori* Defense Function

Theorem 1. Suppose Alice and Bob conduct and Eve monitors an m-bit quantum cryptotransmission that yields Alice and Bob N bits of raw data, including E_T errors. The defense function used by Alice and Bob is given by $t = t_1(n, e_T)$ and is a monotonically nondecreasing function of e_T for any fixed n. Let S denote the event of a successful eavesdropping in the meaning of definition 1, and let

$$\overline{E} \triangleq E[E_i|i \in C], \qquad \overline{I^R} \triangleq E[I_i^R|i \in C, E_i = 0], \tag{4}$$

where $E[\cdots|\cdots]$ denotes the conditional expected value and E_i, I_i^R, and C are defined by Eqs. (2) and in the text immediately preceding it. Then, for any $x > 0$, provided only that $x < n\overline{E}$ and $t_1(n, x)/(n - x) > \overline{I^R}$,

$$\text{Prob}[S|N = n] \le \frac{1}{2}\left\{1 - \text{erf}\left[\sqrt{2n}\left(\overline{E} - \frac{x}{n}\right)\right]\right\} + \frac{1}{2}\left(1 - \text{erf}\left\{[2(n - x)]^{1/2}\left[\frac{t_1(n, x)}{n - x} - \overline{I^R}\right]\right\}\right), \tag{5}$$

where the standard error function is

$$\text{erf}(z) \triangleq 2\frac{1}{\sqrt{\pi}}\int_0^z \exp(-\zeta^2)\text{d}\zeta.$$

Proof. With E_T and I_T^R given by Eqs. (3),

$$\begin{aligned}
\text{Prob}[S|N = n] &= \text{Prob}[I_T^R > t_1(N, E_T)|N = n] \\
&= \sum_{e_T} \text{Prob}[E_T = e_T \cap I_T^R > t_1(n, e_T)|N = n] \\
&= \sum_{e_T \le x} \text{Prob}[E_T = e_T \cap I_T{}^R > t_1(n, e_T)|N = n] \\
&\quad + \sum_{e_T > x} \text{Prob}[E_T = e_T \cap I_T^R > t_1(n, e_T)|N = n] \\
&\le \sum_{e_T \le x} \text{Prob}[E_T = e_T|N = n] \\
&\quad + \sum_{e_T > x} \text{Prob}[E_T = e_T \cap I_T^R > t_1(n, x)|N = n] \\
&= \text{Prob}[E_T \le x|N = n] \\
&\quad + \sum_{e_T > x} \text{Prob}[E_T = e_T|N = n] \\
&\quad \times \text{Prob}[I_T^R > t_1(n, x)|E_T = e_T, N = n],
\end{aligned} \tag{6}$$

where the inequality holds for any $x > 0$.

Let $E = \{E_i\}$ be the random vector describing the locations of all errors in the n-bit raw data and $e = \{e_i\}$ be a particular realization of $\{E_i\}$. By assumption $\{I_i^H, I_i^R, E_i\}$, $i \in C$, are i.i.d. random vectors, I_i^R does not depend on the outcomes of bits other than bit i itself:

$$\begin{aligned}
&\text{Prob}[I_i^R = r|i \in C, E = e, N = n] \\
&\qquad = \text{Prob}[I_i^R = r|i \in C, E_i = e_i].
\end{aligned}$$

For every realization e,

$$I_T^R = \sum_{i \in C}(1 - e_i)I_i^R = \sum_{i \in C, e_i = 0} I_i^R$$

is a sum of a large number $n - \Sigma\, e_i = n - e_T$ of i.i.d. terms relating to bits for which $e_i = 0$, and the mean of each term $\overline{I^R}$ is given by Eqs. (4). By the central limit theorem, given $E = e$ and $N = n$, I_T^R is Gaussian distributed with a mean $(n - e_T)\overline{I^R}$ and a variance $(n - e_T)\sigma_{I^R}^2 \le \frac{1}{4}(n - e_T)$, where the variance is capped because I_i^R is itself bounded between 0 and 1. Note that the distribution function of I_T^R depends on

only the total e_T and not on further details of a particular realization $e = \{e_i\}$, so that

$$\begin{aligned}\text{Prob}[I_T^R > t_1(n,x)|E_T = e_T, N = n] &= \text{Prob}[I_T^R > t_1(n,x)|E = e, N = n] \\ &= \frac{1}{2}\left(1 - \text{erf}\left\{\frac{1}{\sqrt{2}}\frac{t_1(n,x) - (n - e_T)\overline{I^R}}{[\frac{1}{4}(n - e_T)]^{1/2}}\right\}\right),\end{aligned}$$

where, within the theorem's premise of $t_1(n, x)/(n - x) > \overline{I^R}$, the argument of the error function is positive for $e_T > x$. With this,

$$\begin{aligned}&\sum_{e_T > x} \text{Prob}[E_T = e_T|N = n] \\ &\quad \times \text{Prob}[I_T^R > t_1(n,x)|E_T = e_T, N = n] \\ &= \sum_{e_T > x} \text{Prob}[E_T = e_T|N = n] \\ &\quad \times \frac{1}{2}\left(1 - \text{erf}\left\{\frac{1}{\sqrt{2}}\frac{t_1(n,x) - (n - e_T)\overline{I^R}}{[\frac{1}{4}(n - e_T)]^{1/2}}\right\}\right) \\ &\le \frac{1}{2}\left(1 - \text{erf}\left\{\frac{1}{\sqrt{2}}\frac{t_1(n,x) - (n - x)\overline{I^R}}{[\frac{1}{4}(n - x)]^{1/2}}\right\}\right) \\ &\quad \times \sum_{e_T > x} \text{Prob}[E_T = e_T|N = n] \\ &\le \frac{1}{2}\left(1 - \text{erf}\left\{[2(n - x)]^{1/2}\left[\frac{t_1(n,x)}{n - x} - \overline{I^R}\right]\right\}\right). \quad (7)\end{aligned}$$

Similarly, the total number of errors $E_T = \Sigma E_i$, $i \in C$ is the sum of a large number n of i.i.d. random variables E_i, the mean $\overline{E}$ of each being given by Eqs. (4); it follows from the central limit theorem that, given $N = n$, the sum has a Gaussian distribution with a mean $n\overline{E}$ and a variance $n\sigma_E^2 \le \frac{1}{4}n$, so that, in its left tail,

$$\text{Prob}[E_T \le x|N = n] \le \frac{1}{2}\left\{1 - \text{erf}\left[\sqrt{2n}\left(\overline{E} - \frac{x}{n}\right)\right]\right\}, \quad (8)$$

where the argument of the error function is again positive within the theorem's premise of $x < n\overline{E}$. Direct substitution of expressions (7) and (8) into Eq. (6) completes the proof. □

Corollary 1. For any given eavesdropping strategy, Alice and Bob can reduce Eve's chance of success in the "intermediate" sense $\text{Prob}[S|N = n]$ to an arbitrary low level α by selecting as their defense function

$$t_1(n, e_T) = n\begin{cases}0 & \dfrac{e_T}{n} < \overline{E} - \xi \\ [1 - (\overline{E} - \xi)]\overline{I^R} + \xi[1 - (\overline{E} - \xi)]^{1/2}, & \dfrac{e_T}{n} \ge \overline{E} - \xi\end{cases},$$

$$\xi = \frac{1}{\sqrt{2n}}\,\text{erf}^{-1}(1 - \alpha), \quad (9)$$

where the averages $\overline{E}$ and $\overline{I^R}$ are given by Eqs. (4). The same defense function also serves to limit the *a priori* probability $\text{Prob}[S]$ through the identity $\text{Prob}[S] \le \max_n\{\text{Prob}[S|N = n]\} < \alpha$.

Proof. The step function (9) is a monotonically nondecreasing function of e_T, as required by theorem 1. Application of theorem 1 with $x = n(\overline{E} - \xi)$ yields the desired bound for $\text{Prob}[S|N = n]$ because in this case each of the two terms on the right-hand side of relation (5) is less than $\alpha/2$. □

It must be noted that we do not claim that the step function of Eq. (9), which is depicted in Fig. 3(a), is the most economical way to limit Eve's probability of success. A better defense function might exist that accomplishes the same goal with less data sacrifice, particularly if it makes use of the number $m - n$ of inconclusive results.

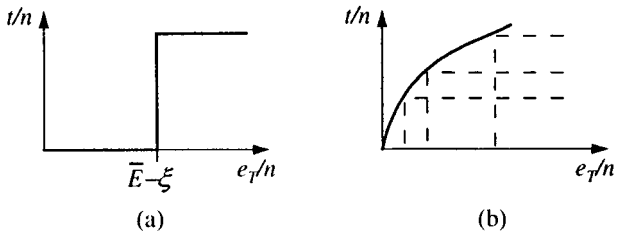

Fig. 3. (a) Defense function $t_1(n, e_T)$ for a single eavesdropping strategy. (b) Defense frontier $t_F(n, e_T)$ (solid curve) constructed to lie above and to the left of multiple defense functions (dashed lines).

4. Defense Frontier and Secrecy Capacity

In Section 3 it was assumed that Eve has at her disposal only one eavesdropping strategy, characterized by the conditional averages $\overline{E}$ and $\overline{I^R}$ defined by Eqs. (4). Let us now remove this limitation.

If Alice and Bob expect multiple eavesdropping strategies to be deployed, they must determine the quantities $\overline{E}$ and $\overline{I^R}$ and plot the defense function from Eq. (9) for each strategy so that a defense frontier $t_F(n, e_T)$ can be constructed above and to the left of all the individual defense functions $t_1(n, e_T)$ [Fig. 3(b)]. In general, no one eavesdropping strategy dominates the rest across the entire range of error rates e_T. However, it is clear both intuitively and from Eq. (9) that, of the two strategies characterized by equal values of $\overline{E}$, the strategy with the greater $\overline{I^R}$ is always superior. We thus restrict our attention to

the set of strategies $\{G(\bar{E})\}$ that yield the greatest attainable values of $\bar{I}^R$ for a given $\bar{E}$. For each such strategy $G(\bar{E})$, denote the corresponding value of $\bar{I}^R$ by $\bar{I}^R_{\max}(\bar{E})$ and the corresponding defense function by $t_{\bar{E}}(n, e_T)$. It is seen from Eq. (9) that the corner points of the step functions $t_{\bar{E}}(n, e_T)$ are given by

$$\begin{cases} \dfrac{e_T}{n} = \bar{E} - \xi; \\ \dfrac{t}{n} = [1 - (\bar{E} - \xi)]\bar{I}^R_{\max}(\bar{E}) + \xi[1 - (\bar{E} - \xi)^{1/2}], \end{cases}$$

which places them on the curve

$$t = (n - e_T)\bar{I}^R_{\max}\left(\frac{e_T}{n} + \xi\right) + \xi[n(n - e_T)]^{1/2}.$$

The defense frontier positioned above and to the left of the entire set of step functions $t_{\bar{E}}(n, e_T)$ is therefore given by

$$t_F(n, e_T) = \max_{e \le e_T}\left\{(n - e)\bar{I}^R_{\max}\left(\frac{e}{n} + \xi\right) + \xi[n(n - e)]^{1/2}\right\}, \tag{10}$$

with ξ given by Eq. (9).

Equation (10) is our main result. It can be seen that the only required input for the construction of the defense frontier is the relation between the expected sifted-data error rate $\bar{E}$ and the maximum attainable expected information on the corrected data $\bar{I}^R_{\max}$. This relation has recently been found for both B92 and BB84 quantum cryptographic protocols.[19] Specifically, for a B92 system with carrier-state overlap $\langle \mathbf{u}|\mathbf{v}\rangle = \sin 2\alpha$, the solution is defined parametrically on an auxiliary variable γ, for $-\delta < \gamma < \delta$, by

$$\bar{I}^R_{\max} = 1 + \log_2\left[1 - \frac{1}{2}\left(\frac{Q' \cos^2 2\alpha - E'}{\cos^2 2\alpha + E'}\right)^2\right],$$
$$\bar{E} = \frac{1}{2}\left(1 - \frac{\cos^2 2\alpha}{E'}\right), \tag{11a}$$

where

$$Q' \triangleq \frac{\sin(\gamma + \delta) + \frac{1}{2}\sin 2\gamma}{\sin\delta\cos\delta},$$
$$E' \triangleq \frac{\cos(\gamma + \delta) - \frac{1}{2}\cos 2\gamma}{\cos^2\delta} + \frac{1}{2}\cos^2 2\alpha,$$
$$\delta \triangleq \arccos\frac{1}{(1 + \sin^2 2\alpha)^{1/2}}, \tag{11b}$$

and for BB84,

$$\bar{I}^R_{\max}(\bar{E}) = 1 + \log_2\left[1 - \frac{1}{2}\left(\frac{1 - 3\bar{E}}{1 - \bar{E}}\right)^2\right]. \tag{12}$$

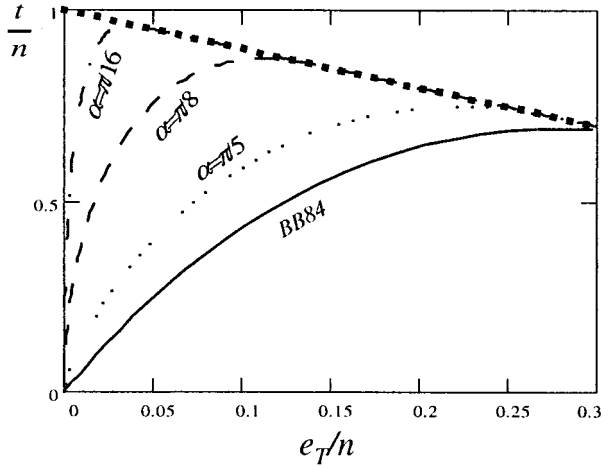

Fig. 4. Defense frontier curves for a BB84 system (solid curve) and for B92 systems with half-separations $\alpha = \pi/5$ (dotted curve), $\alpha = \pi/8$ (dashed curve), and $\alpha = \pi/16$ (dashed–dotted curve) between carrier states **u** and **v**. These are based on the optimal eavesdropping strategies[19] and on Eq. (10) in the limit of long transmission, $m, n \to \infty$, $\xi \to 0$. The boxes indicate the positions where the whole of the corrected data would be sacrificed.

Figure 4 shows the defense frontier curves for BB84 (solid curve) and for three B92 systems with different overlaps $\langle \mathbf{u}|\mathbf{v}\rangle = \sin 2\alpha$. The plots are based on Eqs. (11) and (12) and on Eq. (10), with the latter taken in the limit of a very long transmission so that $\xi \to 0$. All curves terminate on the straight line $t = n - e_T$ (indicated by boxes), where Alice and Bob would be forced to sacrifice the whole of their corrected data. As Fig. 4 suggests, a BB84 communication can be protected at a lower cost than can B92 and remains feasible at higher error rates.

It is interesting to note that the convex shape of the defense frontier curves in Fig. 4 is not coincidental. For any two eavesdropping strategies G_1 and G_2 represented by points D_1 and D_2 along the curve, there exists a family of mixed strategies wherein G_1 is applied with some probability $\lambda < 1$ and G_2 with the complementary probability $1 - \lambda$. In Fig. 4 such strategies would fall on a straight line connecting D_1 and D_2 and, just as all other strategies, must necessarily lie below the defense frontier curve.

For a better comparison of quantum cryptographic systems, let us consider their average secrecy capacity, defined as the number of secret bits produced per bit originally transmitted by Alice[20] (in terms of Fig. 1, this is the ratio of the size $n - e_T - s$ of the final key to the size m of the raw data), in the limit of a very long transmission:

$$C_s' = \lim_{m\to\infty} E\left[\frac{n - e_T - s}{m}\right]$$
$$= \lim_{m\to\infty} E\left[\frac{n}{m}\left\{1 - \frac{e_T}{n} - \frac{t_F(n, e_T)}{n} - \frac{q}{n}\right\} - \frac{g}{m}\right]. \tag{13}$$

As noted above, Alice and Bob do not pursue the impossible goal of communicating regardless of inter-

ference but seek only to satisfy themselves that their exchange is confidential. Consequently, when judging the quality of the channel, the expected value in Eq. (13) is taken with the assumption that no eavesdropping is actually present but that Alice and Bob nevertheless follow the key-distillation procedures illustrated in Fig. 1.

The quantities n/m and e_T/n in Eq. (13) each converge in distribution to their respective means $\chi = E[n/m]$ and $\epsilon = E[e_T/n]$, because they represent cumulative totals generated by an i.i.d. process and obey the central limit theorem. The term $t_F(n, e_T)/n$ likewise converges in distribution to its mean because it is a determinisitic monotonic function of e_T/n. If similar convergence in distribution to a constant can be assumed for the remaining term q/n, it can be shown that the product also converges in distribution:

$$\frac{n}{m}\left[1 - \frac{e_T}{n} - \frac{t_F(n, e_T)}{n} - \frac{q}{n}\right] \xrightarrow{\text{ind}} \chi\left(1 - \epsilon - \left.\frac{t_F}{n}\right|_{e_T/n=\epsilon} - \lim_{m\to\infty} E\left[\frac{q}{n}\right]\right).$$

Furthermore, convergence in distribution to a constant implies convergence in probability, which, for uniformly bounded variables, in turn implies convergence in the mean. Neglecting the term g/m, which is not a random variable and tends to zero as $m \to \infty$, we thus finally obtain

$$C_s' = \chi\left(1 - \epsilon - \left.\frac{t_F}{n}\right|_{e_T/n=\epsilon} - \lim_{m\to\infty} E\left[\frac{q}{n}\right]\right),$$

where

$$\chi \triangleq E\left[\frac{n}{m}\right], \qquad \epsilon \triangleq E\left[\frac{e_T}{n}\right], \tag{14}$$

where, as noted above, the expected values are taken under the assumption that no eavesdropping is present.

The secrecy capacity from Eqs. (14) is plotted in Fig. 5 versus the intrinsic (i.e., eavesdropping-free) channel-error rate ϵ, assuming for the purposes of illustration that $q = 0$. It is apparent that in B92 the optimum choice of α for Alice and Bob generally depends on the error rate in their channel. This is because greater values of α make the carrier states $\mathbf{u}$ and $\mathbf{v}$ difficult to distinguish for Bob as well as for Eve; hence they reduce the yield of conclusive results $\chi = 1 - \sin 2\alpha$.[22] The solid curve in Fig. 5 shows the secrecy-capacity function for the BB84 protocol, for which the rate of conclusive results is always $\chi = 1/2$. The BB84 system is seen to be more efficient in terms of secrecy capacity than any of the B92 systems depicted. Accounting for information leakage q during error correction will not alter this conclusion because q depends on only the number of errors in the raw data and not on the system in which the raw data are generated.[23]

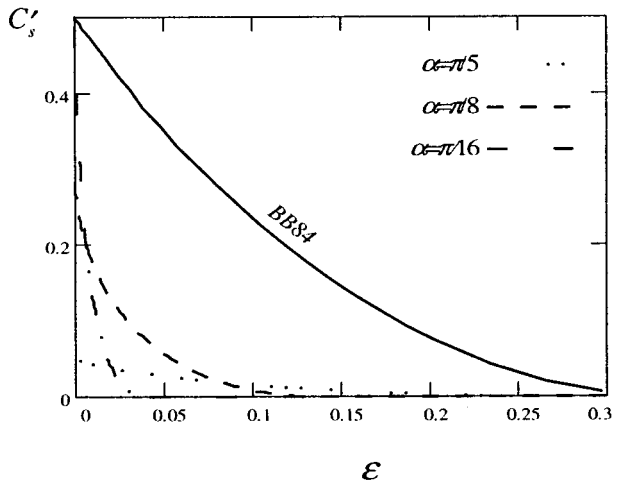

Fig. 5. Secrecy capacities of a BB84 system (solid curve) and B92 systems with $\alpha = \pi/5$ (dotted curve), $\alpha = \pi/8$ (dashed curve), and $\alpha = \pi/16$ (dashed–dotted curve) versus the intrinsic error rate, not including information leakage q during error correction. These are based on the optimal eavesdropping strategies[19] and on Eqs. (10) and (14) in the limit of long transmission ($m \to \infty$).

5. *A Posteriori* Defense Function

In this section we turn briefly to the second interpretation of transmission security in terms of the successful eavesdropping probability Prob[$S|N = n$, $E_T = e_T$]. It was observed in Section 3 that this would be the preferred measure for Alice and Bob since it incorporates all relevant information available to them at the end of the transmission and thus comes closest to an estimate of security for this particular transmission. In contrast, the *a priori* measure Prob[S] reflects only the average security of a class of similar m-bit transmissions. Unfortunately, an attempt to protect the conversation in the sense of Prob[$S|N = n$, $E_T = e_T$] results in impossibly high data sacrifice requirements.

As argued immediately above Eq. (7), for fixed values of $N = n$ and $E_T = e_T$ Eve's Renyi information I_T^R is Gaussian distributed with a mean of $(n - e_T)\overline{I^R}$, where $\overline{I^R}$ must be evaluated for the anticipated eavesdropping strategy by use of Eqs. (4). Just to maintain the *a posteriori* successful attack probability,

$$\text{Prob}[S|N = n, E_T = e_T] = \text{Prob}[I_T^R > t(n, e_T)|N = n, E_T = e_T],$$

below 50% (a rather modest goal), the defense function $t_1(n, e_T)$ must satisfy $t_1(n, e_T) > (n - e_T)\overline{I^R}$. In particular, $t_1(n, e_T = 0) > n\overline{I^R}$. The corresponding defense frontier $t_F(n, e_T)$ against an entire class $\mathcal{G}$ of eavesdropping strategies G must therefore satisfy

$$t_F(n, e_T = 0) = \max_{G\in\mathcal{G}}[t_1(n, e_T = 0)] > \max_{G\in\mathcal{G}}(n\overline{I^R}).$$

Given that some of the strategies represented in Eqs. (11) and (12) produce an $\overline{I^R}$ as high as 1 (albeit at the expense of high error rates $\overline{E}$), Alice and Bob would

have to sacrifice the entire transmission even when the data are error free!

This seemingly paradoxical result can be understood with the aid of the following observation. Quantum cryptographic security relies on the premise that the more productive (from the attacker's point of view) eavesdropping strategies also tend to be more intrusive in terms of the errors introduced. For a *given* strategy, however, it is often the case that the attacker obtains the most information precisely on those occasions when she causes the least disturbance. Alice and Bob would be mistaken if they were to take a low error rate as evidence that a particular transmission is secure. They would also be wrong if they were to conclude that the transmission was "unlikely to have been attacked" or that the attacker was "unlikely to have used an intrusive strategy": The term "likely" is inapplicable to the attacker's activities, for these activities are not random. Their *result*, however, which includes both the number of errors introduced and the amount of information leaked, *is* random. The central argument of Section 3 is that a *combination* of a low error rate and high information leakage is unlikely no matter what strategy the eavesdropper uses—as distinct from the (false) assertion that high information leakage is unlikely *given* a low error rate.

6. Conclusion

Quantum cryptography offers an assurance of communication confidentiality that is unique among cryptographic systems. Even an adversary Eve in possession of superior technology is prevented by the fundamental physical principle from monitoring the transmission unknowingly to the legitimate users, Alice and Bob, although such an adversary can disrupt the transmission. In this paper we have considered only so-called individual eavesdropping attacks, wherein Eve treats each bit of the transmission independently of other bits. At least within this class of attacks, an appropriate key-distillation procedure permits secure communication, even in the presence of noise and losses in the channel, albeit at a reduced throughput. The security that can be achieved, however, is not absolute but is probabilistic in nature, for which a precise definition has been given in Section 2.

To conduct key distillation properly, Alice and Bob must construct a defense function for each anticipated eavesdropping strategy. Combined together these functions lead to the defense frontier, which can protect the communication against the entire set of strategies. Sections 3 and 4 have derived defense frontiers for both the BB84 and the B92 protocols. The proposed defense frontiers are safe against all individual eavesdropping attacks but are not necessarily the most economical in terms of the channel throughput sacrificed for security.

An important figure of merit of a quantum cryptochannel is its secrecy capacity, defined as the number of secret bits obtained per one bit initially transmitted. Secrecy capacity is computed under the assumption that no hostile interference is in fact present, but that Alice and Bob nevertheless implement secure key distillation as a precaution. Secrecy capacity depends on the proportion of bit outcomes that are conclusive, on the defense function adopted by Alice and Bob, and on the amount of information leaked to Eve during the error-correction stage of key distillation. Since the last component is independent of the quantum cryptosystem, for the purposes of comparing the systems it can be omitted. Such a comparison of systems in terms of their throughput versus the intrinsic channel error-rate has been presented in Fig. 5 and has indicated that a BB84 system is more efficient than B92.

This research is supported in part by the Focused Research Initiative program of the Ballistic Missile Defense Organization, the U.S. Air Force Office of Scientific Research, and the U.S. National Science Foundation.

References and Notes

1. C. H. Bennett and G. Brassard, "Quantum cryptography: public key distribution and coin tossing," in *Proceedings of the IEEE International Conference on Computers, Systems, and Signal Processing* (IEEE, New York, 1984), pp. 175–179.
2. A. K. Ekert, "Quantum cryptography based on Bell's theorem," Phys. Rev. Lett. **67,** 661–663 (1991).
3. C. H. Bennett, G. Brassard, and N. D. Mermin, "Quantum cryptography without Bell's theorem," Phys. Rev. Lett. **68,** 557–559 (1992).
4. C. H. Bennett, F. Bessette, G. Brassard, L. Salvail, and J. Smolin, "Experimental quantum cryptography," J. Cryptol. **5,** 3–28 (1992).
5. C. H. Bennett, "Quantum cryptography using any two nonorthogonal states," Phys. Rev. Lett. **68,** 3121–3124 (1992).
6. B. Slutsky, P. C. Sun, Y. Mazurenko, R. Rao, and Y. Fainaman, "Effect of channel imperfection on the secrecy capacity of a quantum cryptographic system," J. Mod. Opt. **44,** 953–961 (1997).
7. E. Biham and T. Mor, "Security of quantum cryptography against collective attacks," Phys. Rev. Lett. **78,** 2256–2259 (1997).
8. E. Biham and T. Mor, "Bounds on information and the security of quantum cryptography," Phys. Rev. Lett **79,** 4034–4037 (1997).
9. D. Mayers, "Quantum key distribution and string oblivious transfer in noisy channels," in *Advances in Cryptology, CRYPTO'96*, N. Kobitz, ed., Vol. 1109 of Springer Lecture Notes in Computer Science Series (Springer, New York, 1996), pp. 343–357.
10. Inconclusive bits are those whose value is not revealed with certainty by Bob's measurement, for example, those measured in the wrong BB84 basis by Bob.[1] Inconclusive bits are an integral feature of quantum cryptographic protocols, even in the absence of channel and detector imperfections.
11. C. H. Bennett, G. Brassard, C. Crepeau, and U. M. Maurer, "Generalized privacy amplification," IEEE Trans. Inf. Theory **41,** 1915–1923 (1995).
12. B. Huttner, N. Imoto, N. Gisin, and T. Mor, "Quantum cryptography with coherent states," Phys. Rev. A **51,** 1863–1869 (1995).
13. H. Yuen, "Quantum amplifiers, quantum duplicators, and quantum cryptography," Quantum Semiclass. Opt. **8,** 939–949 (1996).
14. This condition is unavoidable because a perfect single-photon state is fundamentally impossible to prepare (although a good

approximation can be produced with phenomena such as parametric downconversion).
15. Strictly speaking, the total number of multiphoton bit cells is a Gaussian random variable, and only its average and variance are determined. Still, based on these parameters, it can be bounded from above with any desired confidence level.
16. C. Cachin and U. M. Maurer, "Linking information reconciliation and privacy amplification," J. Cryptol. **10,** 97–110 (1997).
17. C. A. Fuchs and A. Peres, "Quantum-state disturbance versus information gain: uncertainty relations for quantum information," Phys. Rev. A **53,** 2038–2045 (1996).
18. C. A. Fuchs, N. Gisin, R. B. Griffiths, C.-S. Niu, and A. Peres, "Optimal eavesdropping in quantum cryptography. I. Information bound and optimal strategy," Phys. Rev. A **56,** 1163–1172 (1997).
19. B. Slutsky, R. Rao, P.-C. Sun, and Y. Fainman, "Security of quantum cryptography against individual attacks," Phys. Rev. A (to be published).
20. A. K. Ekert, B. Huttner, G. M. Palma, and A. Peres, "Eavesdropping on quantum cryptographical systems," Phys. Rev. A **50,** 1047–1056 (1994).
21. Eve cannot use group information such as block checksums, revealed later in the protocol, because, by assumption, she must attack each bit independently of other bits.
22. The B92 curves in Fig. 5 are qualitatively similar to those in Fig. 4 of Ref. 20, although the latter are computed based on a suboptimal family of eavesdropping strategies and with Shannon rather than Renyi entropy.
23. Because individual bits are transmitted and received independently of one another, errors are distributed uniformly throughout raw data, regardless of the quantum cryptosystem used.