How to Reduce your Enemy's Information (extended abstract)†

Charles H. Bennett

IBM T. J. Watson Research Laboratory [1]
Yorktown Heights
NY 10598

Gilles Brassard [2]

Dépt. IRO, Université de Montréal [3]
C.P. 6128, Succ. "A", Montréal
Québec, H3C 3J7

Jean-Marc Robert [4]

Génie Electrique, Ecole Polytechnique
C.P. 6128, Succ. "A", Montréal
Québec, H3C 3A7

1. INTRODUCTION

In this paper, we investigate how a channel with perfect authenticity but no privacy can be used to repair the defects of a channel with imperfect privacy but no authenticity. More precisely, let us assume that Alice and Bob wish to agree on a secret random bit string. In order to achieve this goal, they have at their disposal an imperfect private channel and an authenticated public channel. The private channel is imperfect in various ways: transmission errors can occur, and partial information can leak to Eve, the eavesdropper, who also can modify the transmissions arbitrarily, as explained below. The only thing Eve cannot do is learn the *entire* contents of the original message sent by Alice. An interesting example of imperfect private "channel", used to exchange (not so random) strings, is Diffie and Hellman's public key distribution scheme [DH], which leaks partial information, even if the discrete logarithm is indeed hard to compute, because it is always feasible for an eavesdropper to determine whether the resulting secret is a quadratic residue or not. The quantum channel [BB1,BB2] is also susceptible to a limited amount of information leakage.

We allow Eve to toggle bits of her choice on the private channel transmissions, or jumble them around, even if she cannot actually read them. This could occur, for instance, if privacy were attempted by enciphering the individual bits with a one-time pad or with a probabilistic encryption scheme [GM] (to toggle an encoded bit, it suffices to multiply its code by the public quadratic non-residue), or alternatively, if a quantum channel were used (by passing selected photons through an appropriate sugar solution). Eve can also suppress the transmission of selected bits and replace them by bits of her choice.

† A full paper was submitted for publication in SIAM J. Comp. as *Privacy Amplification Through Public Discussion.*

1. Present address: Boston University, 111 Cummington Street, Boston, MA 02215.

2. Partially supported by NSERC grant A4107 and by NSF grant MCS-8204506.
3. Part of this research was conducted while this author was visiting the University of California, Berkeley.
4. Partially supported by NSERC grant A4107.

H.C. Williams (Ed.): Advances in Cryptology - CRYPTO '85, LNCS 218, pp. 468-476, 1986.

On the other hand, the public channel transmits information accurately (possibly because it is supplemented by a classical error-correcting code [MS]), and these transmissions cannot be modified or suppressed by Eve, but their entire contents becomes known to her. The authentication capability can either be enforced by physical properties of the channel or through the use of a universal hashing based authentication scheme [WC]. In the latter case, a small number of random secret bits must be shared initially between Alice and Bob, and some of them can be used only once, so that the net effect of the protocol can be viewed as key expansion rather than key distribution. Computationally secure authentication [Br,GGM] can also be used if protection against unlimited computing power is not sought. We shall assume throughout that Alice and Bob did not share initially any secret information, except perhaps for this public channel authentication feature.

It is instructive to compare our setting with the problem solved by the wire-tap channel of Wyner [W], which achieves similar results in a more classically information-theoretic setting. In Wyner's setting, Alice encodes information by a channel code of her choice. The output of her encoder is fed into two classic (discrete, memoryless) communications channels: the *main channel*, leading to the intended receiver Bob, and the *wire-tap channel*, of lesser capacity than the main channel, leading to the eavesdropper. All participants know the channel code and the statistical properties of both channels. Under these conditions, Wyner showed that by appropriate choice of the channel code, Alice can exploit the difference in capacity between the two channels to communicate reliably with Bob while maintaining almost perfect secrecy from the eavesdropper. In our setting, the users have an additional resource: the authenticated public channel. This allows them to cope with a more powerful eavesdropper. Our eavesdropper is more powerful in two ways, either of which would be fatal in Wyner's setting: she can tamper with Alice's communications as well as listen to them, and she eavesdrops by evaluating an N-bit to K-bit function *of her choice*, unknown to Alice and Bob, as we shall see in Section 4.2.

In this paper, we assume that some random bit string has already been transmitted from Alice to Bob over the private channel. We investigate authenticated public channel protocols that, with high probability, detect tampering and transmission errors. Subsequent protocols transform both strings in such a way as to eliminate most, and in some cases all, of Eve's information on the resulting string, except for its length. These public channel protocols remain secure against unlimited computing power. Although excessive tampering on the private channel can result in suppressing communications between Alice and Bob, it cannot fool them into thinking that they share a secret random string when in fact their strings are different or otherwise compromised.

This extended abstract contains no proofs and only a selection of the results found in the complete paper [BBR]. For easier reference, we retain here the full paper's numbering for sections, theorems, etc.. In Section 2, we explain why classical error-correcting codes are inappropriate in this context. In Section 3, we investigate how transmission errors and tampering can be detected with high probability, and sometimes corrected, at the cost of leaking some information to Eve. In Section 4, we investigate how Alice and Bob can subsequently reduce arbitrarily Eve's information at the cost of reducing slightly the length of their shared random string, assuming they have an *a priori* upper bound on the amount of information collected by Eve on the private channel. In Section 5, we investigate the possibility of depriving Eve entirely from any information on the final shared random string, at the cost of reducing its length more substantially.

Before we get started, let us give the following definition and some notation: if $i < j$, a function $f : \{0,1\}^j \to \{0,1\}^i$ is *equitable* if $\#\{x \mid f(x) = a\} = 2^{j-i}$ for every binary string a of length i. If x and

y are equal length bit strings, $x \oplus y$ denotes their bit-by-bit exclusive-or. Finally, if x is a length N bit string and if $0 \le K \le N$, x **mod** 2^K denotes the length K bit string consisting of the rightmost K bits of x, and x **div** 2^K denotes the length $N-K$ bit string obtained from x by deleting its rightmost K bits. We shall herein assume that the reader is familiar with the classical notions of error-correcting codes [MS], information theory [G], universal hashing [CW,WC], and the theory of finite fields [Be].

2. THE INADEQUACY OF CLASSICAL ERROR-CORRECTING CODES

Let us recall that the imperfect private channel considered here is susceptible, not only to random transmission errors, but also to any amount of controlled tampering. The classical theory of error-correcting codes [MS], on the other hand, is based on the assumptions that few errors are more likely to occur than many, and that errors are not maliciously set by an opponent. It is therefore not quite adequate for our purpose.

For instance, let x and y be Alice and Bob's strings, respectively, and let N be their length. Eve's ability to toggle bits of her choice enables her to actually select $x \oplus y$, barring actual transmission errors. This is clearly intolerable if error detection is attempted through a linear error-correcting code [MS]. Indeed, let x be the private channel transmitted codeword corresponding to Alice's chosen random string. Let z be any codeword chosen by Eve. If she perturbs the private channel transmission so that Bob receives $y = x \oplus z$, it will not be possible for him to detect tampering. Notice that Eve can achieve this without gaining any knowledge on the contents of the original transmission x.

Should Alice randomly shuffle the codeword bits, in an attempt to preclude this threat, and publicly tell Bob how to unscramble them only after the private channel transmission is completed, it would no longer be possible for Eve to toggle selected bits and be certain to escape detection. However, if a Hamming code of dimension $[N,K]$ is used, for instance, Eve can toggle 3 random bits and escape detection with probability $1/(N-2)$. Using such a protocol, Alice and Bob could only achieve a high probability of not being fooled, say $1 - 2^{-50}$, at the cost of exchanging unreasonably long strings. In Section 3.1, we describe error detection schemes such that the probability of undetected tampering and transmission errors is independent of the number and position of altered bits. Moreover, this probability can be exponentially small in the length of the strings transmitted.

3. DETECTION AND CORRECTION OF TRANSMISSION ERRORS AND TAMPERING

Let x be some random bit string selected by Alice. Assume she transmits it directly through the imperfect private channel, and let y be the string thus received by Bob. Let N be the length of both strings. We investigate public channel protocols that allow Alice and Bob to detect whenever $x \neq y$ with an arbitrarily small error probability, independently of how y differs for x. The fact that these protocols leak information to Eve about x is considered in Section 4.

3.1. Error detection

A very simple but impractical way of testing whether $x = y$ is for Alice to choose a random function $f : \{0,1\}^N \to \{0,1\}^K$, where K is a security parameter. After the private channel transmission is completed, she sends $f(x)$ to Bob over the public channel, together with a complete description of the function f. Should Bob find out that $f(y) = f(x)$, this would be considered as strong evidence that $y = x$, the error probability being 2^{-K}. On the other hand, should $f(y)$ be different from $f(x)$, Bob

could report to Alice with certainty that he did not receive the correct string. The amount of information on x leaking to Eve from this protocol depends only on the security parameter K, and not on the length N of the strings (except of course for the fact that $K < N$). This would not be the case if a classical error-detecting code had been used. Unfortunately, this scheme cannot be used in practice because there are too many such functions, so that $K2^N$ bits are typically needed to merely transmit a description of the randomly chosen function.

Universal hashing [CW] provides an efficient way to achieve the same goal. After the private channel transmission is completed, Alice randomly chooses a function $f: \{0,1\}^N \rightarrow \{0,1\}^K$ among some standard universal$_2$ class of functions. She then sends both $f(x)$ and a description of f to Bob. Thanks to universal hashing, the description of f can be transmitted efficiently. After computing $f(y)$, Bob checks whether it agrees with $f(x)$. If it does, a basic property of universal hashing allows them to assume that $x = y$, their probability of error being bounded by 2^{-K}.

3.2. Reconciliation of the strings

Whether $f: \{0,1\}^N \rightarrow \{0,1\}^K$ is chosen as a completely random function or within some universal$_2$ class of functions, what should Alice and Bob do whenever $f(x)$ differs from $f(y)$? If the private channel is reliable enough that only one or perhaps two errors are to be expected at most, it may be worthwhile for Bob to try computing $f(z)$ on all strings z differing from y by only one bit or two, in the hope of finding a match with $f(x)$ and thus a likely candidate z for x.

If many transmission errors are to be expected, this would be much too time consuming. In the full paper, we offer two different solutions to this problem, one based on the *post-facto* application of a convolutional code and one based on a blockwise exclusive-or strategy. The effect of the convolutional code protocol is to allow Bob to transform y into x with high probability, at the cost of disclosing to Eve some information about x. Protocols from Section 4 can subsequently be applied to reduce that information. On the other hand, the effect of the exclusive-or strategy is to transform both x and y into a probably common shorter string z on which Eve has no more information than she initially had on x from eavesdropping over the private channel.

4. REDUCTION OF THE EAVESDROPPER'S INFORMATION

Assuming that Alice and Bob agree on their strings as a result of one of the protocols discussed above, Eve has two different sources of information on that string: deterministic information obtained from eavesdropping on the private channel, as the original random bit string was being transmitted, and stochastic information resulting from eavesdropping on the public channel, as the agreement protocol was being carried out.

In this section, we investigate how to reduce Eve's information arbitrarily close to zero, at the cost of slightly shrinking the random bit string shared between Alice and Bob at the end of the protocol. In a first step, we assume that no eavesdropping on the private channel has occurred, but that tampering and transmission errors were possible. In a second step, we assume to the contrary that a limited amount of eavesdropping on the private channel is susceptible of having occurred, but that it is not necessary to carry out an agreement protocol from Section 3, thus depriving Eve from this potential stochastic information. Finally, the full paper considers the case where both sources of information are simultaneously available to her. All these protocols are secure against an eavesdropper with unlimited computing power.

4.1. Reducing the public channel eavesdropper's information

Let us assume for the moment that Eve did not attempt eavesdropping on the private channel, but that she has complete information on the error detection protocol carried out between Alice and Bob over the public channel. Let x be the random string of length N on which Alice and Bob have just agreed, and let $f : \{0,1\}^N \to \{0,1\}^K$ be their error detection function. Eve knows the K-bit value of $f(x)$, together with the function f itself. Her information can be characterized by the set $C = \{z \in \{0,1\}^N \mid f(z) = f(x)\}$ of possible candidates for x. From Eve's point of view, each element of C is equally likely to be the string x currently shared between Alice and Bob. Notice that Alice and Bob also have complete knowledge on the set C.

In order to reduce Eve's information, Alice and Bob publicly agree on a function $g : \{0,1\}^N \to \{0,1\}^R$, for some integer $R \leq N-K$, such that knowledge of the set C gives arbitrarily little information on $g(x)$, or perhaps even none at all. The final string on which Alice and Bob agree is thus $g(x)$. In other words, the purpose of this function g is to shrink the string x by at least K bits, in order to compensate for the K bits of information that knowledge of C gives Eve.

4.1.1. The case of truly random functions

Assume the error detection function f was chosen randomly among all N-bit to K-bit functions. Let $g : \{0,1\}^N \to \{0,1\}^R$ be the function $g(x) = x \bmod 2^R$. Let $S = N-K-R$, then

Theorem 8. The expected amount of information known by Eve on $g(x)$ from knowledge of f, g and $f(x)$ is less than $2^{-S}/\ln 2$ bit.

Here, S should be thought of as the number of additional bits sacrificed to privacy. Sacrificing one more bit in the final string chops in half Eve's information about it. This holds even if Eve knows in advance which information reduction function g is to used. Any other equitable N-bit to R-bit function would have performed just as well.

4.1.3. The case of universal hashing

Let us now assume that a practical error detection protocol was used: the function $f : \{0,1\}^N \to \{0,1\}^K$ was randomly chosen among some universal$_2$ class of hash functions. Rather than developing a general theory of information reduction in this context, let us design an *ad hoc* technique for a given universal$_2$ class.

Let a and b be elements of GF(2^N) [Be] such that $a \neq 0$. The degree one polynomial $q_{a,b}(x) = ax + b$, arithmetic being done in GF(2^N), defines a permutation of GF(2^N). If we let $\sigma : \{0,1\}^N \to \mathrm{GF}(2^N)$ stand for the natural one-one correspondance, this induces a permutation $\pi_{a,b} : \{0,1\}^N \to \{0,1\}^N$ defined by $\pi_{a,b}(x) = \sigma^{-1}(q_{a,b}(\sigma(x)))$. Therefore, for any fixed $K \leq N$, the function $h_{a,b} : \{0,1\}^N \to \{0,1\}^K$ defined by $h_{a,b}(x) = \pi_{a,b}(x) \bmod 2^K$ is equitable. Futhermore, the class of all such functions $h_{a,b}$, for every $a, b \in \mathrm{GF}(2^N)$, $a \neq 0$, forms a universal$_2$ class of hash functions, so that it can be used for the error detection protocol.

Theorem 15. Let a and b be any elements of GF(2^N) such that $a \neq 0$. Let x be a random string of length N. Then knowledge of a, b and $h_{a,b}(x)$ gives no information on the length $N-K$ string defined as $\pi_{a,b}(x) \operatorname{div} 2^K$.

Use of this universal$_2$ class allows Alice and Bob to verify whether their strings are identical, with a probability of error at most 2^{-K}. If they turn out to be the same, they can be transformed into a new string that is only K bits shorter, on which Eve has no information at all. This is optimal.

4.2. Reducing the private channel eavesdropper's information

Let us now assume that partial eavesdropping has occurred on the private channel. Let K be an upper bound on the number of bits of information thus obtained by Eve, where $K < N$. This can be formalized as follows in general: Eve chooses any function $e : \{0,1\}^N \rightarrow \{0,1\}^K$, and she obtains the value of $e(x)$ after x has been transmitted over the private channel. Of course, Alice and Bob have no information on which function e was chosen by Alice, except for an upper bound on K.

The effect of eavesdropping over the private channel is very similar to that of eavesdropping over the public channel, as described in Section 3, in that the information gained by Eve can be characterized by a set $E = \{z \in \{0,1\}^N \mid e(z) = e(x)\}$ of possible candidates for x. However, there is a fundamental difference: it is no longer true that Alice and Bob have complete knowledge on E. For this reason, it is not possible for them, in general, to eliminate Eve's information with certainty.

Theorem 17. No matter how Alice and Bob choose their function $g : \{0,1\}^N \rightarrow \{0,1\}^R$, for any $R > 0$, there always is an equitable function $e : \{0,1\}^N \rightarrow \{0,1\}^K$, for any $K > 0$, such that knowledge of e, g and $e(x)$ yields information on $g(x)$.

Therefore, the best Alice and Bob can hope for is to reduce arbitrarily Eve's information. There can be no analogue to Theorem 15. Nonetheless, if we restrict even further Eve's choice of e, so that she can only read a selection of K physical bits of x. it becomes possible again for Alice and Bob to eliminate her information entirely, as discussed in Section 5.

For simplicity, let us assume that transmission errors and tampering are not a worry for Alice and Bob, so that an error detection protocol is not carried out. This assumption is removed in Section 4.3 of the full paper. Let x be the length N bit string common to Alice and Bob, and let $e(x)$ be the K-bit information known by Eve about x. Alice and Bob wish to publicly agree on some function $g : \{0,1\}^N \rightarrow \{0,1\}^R$, for some $R \leq N-K$, such that knowledge of e, $e(x)$ and g leaves Eve with an arbitrarily small fraction of one bit of information about $g(x)$.

Here again, we consider two approaches for the reduction of Eve's information: one based on truly random functions and one based on universal hashing techniques. The first approach is only of theoretical interest, but the second one is efficient in practice.

4.2.1. The case of truly random functions

Theorem 19. Let $e : \{0,1\}^N \rightarrow \{0,1\}^K$ be any function, let $S < N-K$ be a security parameter, and let $R = N-K-S$. If $g : \{0,1\}^N \rightarrow \{0,1\}^R$ is chosen randomly, the expected amount of information on $g(x)$ given by knowledge of e, g and $e(x)$ is at most $2^{-S}/\ln 2$ bit.

4.2.2. The case of universal hashing

Contrary to the error detection protocols of Section 3, it is no longer sufficient to consider universal$_2$ classes: here, we use *strongly* universal$_2$ classes [WC].

Theorem 21. Let e, S and R be as in Theorem 19, let H be a publicly known strongly universal$_2$ class of hash function from $\{0,1\}^N$ to $\{0,1\}^R$ and and let g be a function chosen randomly within H. The expected amount of information on $g(x)$ given by knowledge of e, g and $e(x)$ is at most $2^{-S}/\ln 2$ bit.

The above theorem is true despite the fact that Eve already knows the class H, but of course not the specific function g, when she gets to choose her function e.

5. ELIMINATION OF THE EAVESDROPPER'S INFORMATION

The protocols of Section 4.2 should be sufficient for most applications, despite the fact that Eve still has an arbitrarily small fraction of one bit of information on the resulting shared random string. Although we were able to eliminate her information entirely in Theorem 15, the techniques used could only be applied because Alice and Bob had complete knowledge of Eve's information. As shown in Theorem 17, this cannot be extended whenever Eve is allowed to access information *of her choice* from the private channel transmission.

In this section, we investigate a protocol by which Alice and Bob can nonetheless wipe out Eve's information, assuming that she obtained a maximum of K *physical* bits of her choice from the private channel transmission. Although the value of K is known to Alice and Bob, they do not know, of course, which particular bits of their string are compromised. This protocol is expensive in the sense that the resulting string is generally substantially shorter than those resulting from the protocols of Section 4.2; however, this is the unavoidable price to pay in order to make sure that Eve is left with no information at all.

5.1. The notion of (N, J, K)-functions

For any integers N, J and K such that $N \geq J+K$, $J > 0$ and $K > 0$, a function $f : \{0,1\}^N \rightarrow \{0,1\}^J$ is said to be (N, J, K) if, no matter how one fixes any K of its input bits, each of the 2^J output bits can be produced in exactly 2^{N-J-K} different ways by varying the remaining $N-K$ input bits. Intuitively, an (N, J, K)-function compresses an N bit string into a J bit string in such a way that knowledge of any K of the input bits gives no information on the output. This is equivalent to the notion of t-resilient functions independently introduced by [CGHFRS].

Given such a function, Alice and Bob can apply it to their respective strings, thus producing a new (shorter) string on which Eve has no information. Notice that this still holds even if she already knows which function will be used by Alice and Bob in advance of her deciding which K bits to read from the private channel. Therefore, the subsequent public transmission between Alice and Bob is not necessary in this case, as it can be replaced by a standard protocol.

The case $J = N-K$ is the best possible because there is no hope to produce a completely secret string of length $N-K+1$ if Eve knows K of the original N bits. A function f that is $(N, N-K, K)$ is said to be (N, K). The following theorem shows how to build (N, K)-functions whenever they exist.

Theorem 23.

1) For any $N > 1$, there are $(N,1)$ and $(N, N-1)$-functions.

2) For any $N > 3$, there are no (N, K)-functions whenever $1 < K < N-1$.

5.2. How to build (N, J, K)-functions

We wish to answer the following question: given N and K, what is the maximum value for J such that an (N, J, K)-function exists? In other words, what is the longest secret random string on which Alice and Bob can agree if they start from a random string of length N, of which K bits are compromised. Theorem 23 shows that J must be strictly smaller than $N-K$ unless $K = 1$ or $K = N-1$.

We were unable to answer the above question in its full generality. For this reason, we restrict our attention to the special class of (N, J, K)-functions for which every output bit is produced as the exclusive-or of some of the input bits. Such functions are referred to as xor-(N, J, K)-functions. We conjecture that these functions are as efficient as possible, in the sense that if no xor-(N, J, K)-functions exist for given values of N, J and K, then no general (N, J, K)-functions exist either. This *Xor-Conjecture* is proved in [CGHFRS] for the case $J = 2$, but it is not believed in general by all members of [CGHFRS].

The following characterization, known as the *Xor-Lemma*, allows to establish an equivalence between xor-(N, J, K)-functions and binary linear codes [MS].

Lemma 25 (independently discovered by [CGHFRS]). Let M be a JxN Boolean matrix. Let $f: \{0,1\}^N \rightarrow \{0,1\}^J$ be the function represented by M in the natural way (i.e. $f(x)^t = Mx^t$, all operations being performed modulo 2). The function f is (N, J, K) if and only if the exclusive-or of any non-empty set of rows of M contains at least $K+1$ ones.

The equivalence is now stated:

Theorem 26 (independently discovered by [CGHFRS]). For given values of N, J and K, there exists an xor-(N, J, K)-function if and only if there exists an $[N,J]$ binary linear code with minimum distance at least $K+1$ between any two codewords.

Consequently, our problem is equivalent to a classical problem of algebraic coding theory. Unfortunately, no efficient algorithms are known, much less closed formed formulae, to determine the largest possible minimum codeword distance among all $[N, J]$ binary linear codes. There are, however, several classical lower and upper bounds on this value [MS], and these bounds apply just as well to our problem.

For instance, Hamming codes tell us that xor-$(2^L-1, 2^L-L-1, 2)$-functions exist for every $L \geq 2$. Conversely, Hamming's upper bound show that no xor-$(2^L-1, 2^L-L, 2)$-functions can exist. Elimination of Eve's information in this case (K=2) costs $L-2-S$ more bits than if we had been satisfied to reduce her information below $2^{-S}/\ln 2$ bit, as in Section 4.2. Similarly, Griesmer's upper bound and the simplex code allow to build xor-$(2^L-1, L, 2^{L-1}-1)$-functions for any $L \geq 2$, whereas neither xor-$(2^L-1, L, 2^{L-1})$-functions nor xor-$(2^L-1, L+1, 2^{L-1}-1)$-functions can exist. Finally, Varsharmov-Gilbert's lower bound together with McEliece's upper bound allow to construct xor-(N, J, K)-functions such that J is at least half the optimal (xor) value, as long as $K/N < 0.3$ and N is large enough. We encourage the reader to consult [CGHFRS] for additional results on (N, J, K) (*alias* t–resilient) functions.

6. CONCLUSIONS

If no eavesdropping occurred over the private channel, it is possible for Alice and Bob to publicly verify that no transmission errors nor tampering occurred either, with a 2^{-K} error probability,

nd end up with an entirely secret final string that is only K bits shorter than the original private ransmission. This is optimal. A somewhat shorter common string, on which Eve still has no informa- ion, can also be obtained with high probability despite transmission errors over the private channel.

If partial eavesdropping occurred over the private channel, leaking up to K bits of information to Eve, in Shannon's sense, it is still possible for Alice and Bob to publicly verify that no transmission errors nor tampering occurred, with a 2^{-L} error probability, and end up with a final string that is $K+L+S$ bits shorter than the original private transmission, on which Eve has less than $2^{-S}/\ln 2$ bit of information. Here again, transmission errors can be handled at the cost of reducing some more the length of the final common string.

Finally, if partial eavesdropping over the private channel is restricted to K physical bits secretly chosen by Eve, it becomes possible again for Alice and Bob to verify with high probability that no errors nor tampering occurred, and end up with a new string on which Eve has no information whatsoever. However, the new string is substantially shorter than if Alice and Bob had tolerated knowledge by Eve of an arbitrarily small fraction of one bit of information.

7. REFERENCES

[Be] E. R. Berlekamp, *Algebraic Coding Theory*, McGraw-Hill, New York, 1968.

[Br] G. Brassard, "On Computationally Secure Authentication Tags Requiring Short Secret Shared Keys", in *Advances in Cryptology: Proc. of Crypto 82*, D. Chaum, R. L. Rivest and A. T. Sherman, eds., Plenum, New York, 1983, pp. 267-275.

[BB1] C. H. Bennett and G. Brassard, "Quantum Cryptography and its Application to Provably Secure Key Expansion, Public-Key Distribution and Coin-Tossing", in *IEEE International Conference on Computers, Systems and Signal Processing*, Bangalore, India, December 1984, pp. 175-179.

[BB2] C. H. Bennett and G. Brassard, "An Update on Quantum Cryptography", in *Advances in Cryptology: Proc. of Crypto 84*, G. R. Blakley and D. Chaum, eds., Lecture Notes in Computer Science 196, Springer-Verlag, Berlin, 1985, pp. 475-480.

[BBR] C. H. Bennett, G. Brassard and J.-M. Robert, "Privacy Amplification through Public Discussion", submitted to *SIAM J. Comput.*, 1985.

[CW] J. L. Carter, and M. N. Wegman, "Universal Classes of Hash Functions", *J. Comput. System Sci.*, 18 (1979), pp. 143-154.

[CGHFRS] B. Chor, O. Goldreich, J. Hastad, J. Freidmann, S. Rudich and R. Smolensky, "The Bit Extraction Problem or t-Resilient Functions", in *Proc. 26th IEEE Symposium on Foundations of Computer Science*, IEEE Computer Society Press, 1985, pp. 396-407.

[DH] W. Diffie and M. Hellman, "New Directions in Cryptography", *IEEE Trans. Information Theory*, IT-22 (1976), pp. 644-654.

[G] R. G. Gallager, *Information Theory and Reliable Communication*, John Wiley and Sons, New York, 1968.

[GGM] O. Goldreich, S. Goldwasser and S. Micali, "How to Construct Random Functions", in *Proc. 25th IEEE Symposium on Foundations of Computer Science*, IEEE Computer Society Press, 1984, pp. 464-479.

[GM] S. Goldwasser and S. Micali, "Probabilistic Encryption", *J. Comput. System Sci.*, 28 (1984), pp. 270-299.

[MS] F. J. MacWilliams and N. J. A. Sloane, *The Theory of Error-Correcting Codes*, North-Holland, New York, 1977.

[WC] M. N. Wegman and J. L. Carter, "New Hash Functions and Their Use in Authentication and Set Equality", *J. Comput. System Sci.*, 22 (1981), pp. 265-279.

[W] A. D. Wyner, "The Wire-Tap Channel", *Bell System Journal*, 54 (1975), pp. 1355-1387.