The Exact Price for Unconditionally Secure Asymmetric Cryptography

Renato Renner[1] Stefan Wolf[2]

[1] Department of Computer Science, ETH Zürich, Switzerland.
renner@inf.ethz.ch
[2] Département d'Informatique et R.O., Université de Montréal, Canada.
wolf@iro.umontreal.ca

Abstract. A completely insecure communication channel can only be transformed into an unconditionally secure channel if some information-theoretic primitive is given to start from. All previous approaches to realizing such authenticity and privacy from weak primitives were symmetric in the sense that security for both parties was achieved. We show that asymmetric information-theoretic security can, however, be obtained at a substantially lower price than two-way security—like in the computational-security setting, as the example of public-key cryptography demonstrates. In addition to this, we show that also an unconditionally secure bidirectional channel can be obtained under weaker conditions than previously known. One consequence of these results is that the assumption usually made in the context of quantum key distribution that the two parties share a short key initially is unnecessarily strong.

Keywords. Information-theoretic security, authentication, information reconciliation, privacy amplification, quantum key agreement, reductions of information-theoretic primitives.

1 Motivation and Main Results

1.1 Realizing Unconditional Security from Other Primitives

There are mainly two types of cryptographic security, namely *computational* and *information-theoretic* security. Systems of the first type can in principle be broken by adversaries with sufficient computing power; their security is based on the hardness of certain computational tasks—such as factoring large integers or computing discrete logarithms. However, no proofs can be given up to date for the security of such schemes. To make things even worse, the realization of a *quantum computer* would allow for breaking many presently-used systems efficiently. These facts serve as a strong motivation for the study of *information-theoretically secure cryptography.* Systems of this type are provably unbreakable even by computationally unlimited adversaries. Clearly, this is the most desirable type of security—but it has its price [21], the exact determination of which has been an open problem and subject to intensive study. Most generally speaking, this price is some *information-theoretic primitive* [15] I, such as shared keys that are fully [20], [8], [11] or partially random [9], [18] and secret [19], authenticated and/or noisy classical [22], [6] or quantum [1] communication channels, or correlated pieces of information [13].

In order to describe these previous—and our new—results on a conceptual level, we use the following “channel calculus” introduced in [16]. Here, $A \longrightarrow B$ denotes an insecure communication channel from Alice to Bob, $A \bullet\!\!\longrightarrow B$ is an authentic channel from Alice to Bob (i.e., the *exclusivity*—represented by “$\bullet$”—sits on the sender's side, whereas the actual *security* is on the receiver's side: according to *his* view and knowledge, the message comes indeed from the legitimate sender), $A \longrightarrow\!\!\bullet B$ is a confidential channel (in the sender's view, the channel's output is accessible exclusively by the legitimate receiver), and the channel $A \bullet\!\!\longrightarrow\!\!\bullet B$ offering both authenticity *and* confidentiality is called a secure channel. The bidirectional channel $A \bullet\!\!\longleftrightarrow B$, for instance, is authentic from Alice to Bob and confidential in the opposite direction.

A number of previous results showed when and how an unconditionally secure channel can be obtained from completely insecure and from authentic but public channels, respectively. In [22], [6], [17], [19], examples of information-theoretic primitives I are given that allow for obtaining an unconditionally secure channel from completely insecure communication, i.e., for realizing the transformation

$$\left.\begin{array}{c} A \longleftrightarrow B \\ I \end{array}\right\} \; A \bullet\!\!\longleftrightarrow\!\!\bullet B \ ,$$

whereas it was shown in [13], for instance, that the required primitive I' can generally be much weaker if the communication channel is assumed to be *authentic* initially:

$$\left.\begin{array}{c} A \bullet\!\!\longrightarrow B \\ A \longleftarrow\!\!\bullet B \\ I' \end{array}\right\} \; A \bullet\!\!\longleftrightarrow\!\!\bullet B \ .$$

Note that in the context of *computational* security, this latter channel transformation is possible *without* any additional primitive I' (e.g., by using the Diffie-Hellman protocol [7]). In sharp contrast to this, unconditional authenticity *alone* is not sufficient for realizing unconditional confidentiality [13], [14], [17].

Clearly, a typical example of a primitive I which works in both of the above cases is a shared secret key of sufficient length. The question is whether much weaker primitives can be sufficient as well. More specifically, some of the open questions are the following.

- All known examples of protocols achieving the above transformations do this via the generation of a shared private key. The generated (unconditional) security then sits on both sides of the channel (as shown in the diagrams above). Is it possible to realize unconditional security *on only one end of the channel* under weaker assumptions? In other words, what is the price for realizing *asymmetric*[1] *unconditional security*? What is the minimal price for an unconditional “$\bullet$”?

- Unconditional secret-key agreement protocols consist of different phases (such as interactive error correction, called information reconciliation, or privacy amplification). The assumption is normally made that the public communication channel over which these protocols are carried out is authentic. Which of these protocol steps *do* require authentic channels, and which do not?
- If authentic channels *are* indeed necessary (such as in quantum key distribution), what is the minimal price (i.e., the weakest possible primitive) for obtaining them?

In the present paper, we give answers to all three questions. First, we describe a class of information-theoretic primitives I'' that allow for obtaining unconditional asymmetric security, i.e., for realizing the transformation

$$\left.\begin{array}{c} A \longleftrightarrow B \\ I'' \end{array}\right\} \; A \bullet\!\!\!\longleftrightarrow B \ .$$

We show that such a primitive I'' is generally *not* sufficient for obtaining a *two-way* secure channel, and that our class of primitives is optimal in the sense that weaker primitives do normally not allow for obtaining *any information-theoretic security at all* in the setting of completely insecure communication. Because of these two optimality results, one can say that we give the exact price for unconditional security, i.e., for realizing an unconditional “$\bullet$”, which can be seen as an “atom of information-theoretic security”.

Among the protocols used to achieve these results are methods for so-called *information reconciliation* (i.e., interactive error correction) not requiring authentic channels. Together with a similar result for privacy amplification [19], this implies that in many cases, information-theoretically secure key agreement protocols exist which do not require authentic channels *at all.*

If, on the other hand, such authenticity *is* required for a protocol, it can be achieved under much weaker assumptions than previously believed. For instance, it has been a standard assumption in quantum key distribution that the processing of the key requires a short secret key to start with—therefore, quantum key agreement is sometimes said to be key *expansion.* We show that neither a short secret key [11] nor a partially secret *common* string [19] are required for quantum key distribution, but that much weaker assumptions are in fact sufficient.

[1] Note that the term *asymmetric* is used here with respect to the high-level functionality and not—as usual—with respect to the keys held by the parties. In spite of this difference, it is fair to say that we try to realize the functionality of public-key authentication and encryption, but in the setting of unconditional security.

1.2 Main Results

We now give the main results of this paper. We first introduce the entropy measures required to formulate them. (For an introduction to information theory, see, for instance, [5].)

Definition 1. Let X and Y be two random variables (with ranges $\mathcal{X}$ and $\mathcal{Y}$). The *min-entropy* $H_\infty(Y)$ of Y is[2] $H_\infty(Y) := -\log(\max_{y \in \mathcal{Y}}(P_Y(y)))$. The 0-*entropy* $H_0(Y)$ is defined as $H_0(Y) := \log|\{y \in \mathcal{Y} \mid P_Y(y) > 0\}|$, and let

$$H_0^{\max}(Y|X) := \max_{x \in \mathcal{X}}(H_0(Y|X = x)) \ .$$

It has been shown in [19] that a common key S an arbitrarily large fraction of which (in terms of min-entropy) is known to the adversary is sufficient for obtaining two-way unconditional security.

Previous Result. [19] Let Alice and Bob be connected by a completely insecure bidirectional channel and share a binary string S, whereas an adversary Eve knows a random variable U such that[3]

$$H_\infty(S|U = u) = \Omega(\text{len}(S))$$

holds (where $u \in \mathcal{U}$ is the particular value known to Eve). Then Alice and Bob can obtain an unconditionally authentic and confidential bidirectional channel between each other.[4]

In this paper, we show that unconditional security on only one side of the channel can be achieved at a substantially lower price; in particular, the parties are not required to share any common string initially. The following result and its tightness are shown in Sections 2 and 3.

Asymmetric Result. Assume that Alice and Bob—who are connected by a completely insecure bidirectional channel—, and an adversary Eve know random variables X, Y, and U, respectively, such that

$$H_\infty(Y|U = u) - H_0^{\max}(Y|X) = \Omega(\log|\mathcal{Y}|) \tag{1}$$

holds. Then Alice and Bob can obtain an unconditionally authentic channel from Alice to Bob and an unconditionally confidential channel from Bob to Alice.

[2] All logarithms in this paper are with respect to the base 2.

[3] It is only for simplicity that we give asymptotic formulations of the previous and new results here. The involved hidden constants are small, and the protocols are useful already for relatively small values of n.

[4] More precisely, the length of a message that can be sent in an almost-perfectly secret way, for instance, is $(1 - o(1))H_\infty(S|U = u)$.

The length of the message which can be sent in a confidential way is (asymptotically) equal to the expression on the left hand side of (1). It is shown in Section 3.2 that this is optimal.

We also give a symmetric result which improves on the previous result above: Even a completely secure bidirectional channel can be obtained by parties not sharing a common string to start with. This is shown in Section 4.

Symmetric Result. Assume that Alice and Bob—who are connected by a completely insecure bidirectional channel—, and an adversary Eve know random variables X, Y, and U, respectively, such that

$$\begin{aligned}\max(H_\infty(X|U=u), H_\infty(Y|U=u)) - H_0^{\max}(Y|X) - H_0^{\max}(X|Y) \\ = \Omega(\max(\log|\mathcal{X}|, \log|\mathcal{Y}|))\end{aligned}$$

holds. Then Alice and Bob can obtain an unconditionally authentic and confidential bidirectional channel between each other.

In contrast to many previous secret-key agreement protocols, our protocols are not restricted to specific probability distributions but are universal in the sense that they work for *any* element in the class of distributions characterized by the given entropy conditions, where Alice and Bob do not have to know what the actual distribution is. Of course, such a condition is just *one* possible way of defining classes of distributions; it is a natural one, however, since a direct connection can be made to, for instance, an adversary's memory space. In Section 3 it is shown that our protocols are—in their universality—optimal.

Note that we have conditioned the involved random variables on an adversary's knowledge $U = u$. Alternatively, our results can be interpreted as to concern the model of unconditional security from keys generated by *correlated weak random sources* (other examples of such results are given in [9] and [18]).

If, on the other hand, Y is a *a priori* uniformly distributed key and U is Eve's information, then inequality (1) can be replaced by the—somewhat stronger—assumption

$$H_0(U) + H_0^{\max}(Y|X) = (1 - \Omega(1)) \log|\mathcal{Y}| \tag{2}$$

because of Lemma 2 below. Condition (2) is directly comparable to related bounds and results in *quantum* cryptography since all the involved quantities now have natural "translations" on the quantum side: The entropy of the involved random variables can simply be replaced by the entropy of the corresponding quantum states. Bounds on these quantities naturally arise from bounds on the size of an adversary's (quantum) memory [12], for instance.

2 Asymmetric Unconditional Security from Minimal Primitives

2.1 Authentication Between Parties NOT Sharing a Common String

The first ingredient for our protocols is an unconditional authentication method that is secure even between parties not sharing the same string; furthermore, none of the two parties' initial strings has to be secret, the only condition being that a non-vanishing fraction of the receiver's string is unknown to the adversary (in terms of min-entropy). More precisely, we show that the interactive authentication method presented in [19]—there in the context of parties sharing a partially secret key—has the following property: Under the *sole condition* that an adversary Eve is not fully aware of the receiver Bob's string, the latter can receive authenticated messages from Alice: He will (almost) never accept if the message was *not* the one sent by Alice (whatever her string and Eve's knowledge about it is). In other words, the protocol is secure also if Alice and Bob do not share the same key. More precisely, whereas they will only accept if their initial strings *are* identical—a fact that they enforce by interactive error correction—, Eve is unable to mount a successful active attack *even if they are not.*

We review Protocol AUTH of [19]—using identical keys s there; here, we will later replace s by two not necessarily equal strings y and y'. For parameters $k \cdot l = n$, let $s = s_0||s_1|| \cdots ||s_{k-1}$ be the decomposition of the n-bit string s into l-bit substrings, interpreted as elements of $GF(2^l)$, and let, for $x \in GF(2^l)$,

$$p_s(x) := \sum_{i=0}^{k-1} s_i \cdot x^i \tag{3}$$

be the evaluation in x of the polynomial represented by s. Then the protocol consists of repeating the following three rounds: First, Alice—the sender of the message to be authenticated—sends a random challenge $c' \in \{0,1\}^l$ to Bob which he replies to by sending back the pair $(p_s(c'), c)$, where $c \in \{0,1\}^l$ is another random challenge. Alice (after having checked the correctness of Bob's message—if it is incorrect, she rejects and aborts the protocol) then sends a message bit and, if this bit is 1, the value $p_s(c)$ to confirm. Under the assumption that an encoding of messages is used such that any insertion of a 0-bit (something Eve obviously can do) as well as any bit flip from 1 to 0 can be detected—because the resulting string is not a valid codeword—, this protocol was proven secure in [19]; more precisely, it was shown to be hard for Eve (having non-vanishing uncertainty in terms of min-entropy about $S = s$) to respond to a challenge, made by one party, without being able to use the other as an oracle, and that this fact implies the security of the protocol. Furthermore, it was shown that an encoding of m-bit messages with the mentioned properties exists with code word length $M = (1 + o(1))m$.

Below, we will show the security of this protocol—from the receiver's point of view (like in one-way authentication)—even when the parties do *not* share the same string and under the only assumption that Eve has some uncertainty about Bob's string (y). The main technical ingredient of this is Lemma 1, which implies, roughly speaking, that under the given conditions, Eve can, with overwhelming probability, either not respond to Alice's challenges (c') or not to Bob's (c)—even when given Alice's string (y'). The intuitive reason for this is that it is *either useless or impossible* for Eve to (impersonate Bob and) talk to Alice—depending on whether her uncertainty about Alice's string is small or not. Without loss of generality, we state and prove Lemma 1 with respect to *deterministic* adversarial strategies (given by the functions f and g).

Lemma 1. *Let Y' and Y be two random variables with joint distribution $P_{Y'Y}$ and ranges $\mathcal{Y}' = \mathcal{Y} = \{0,1\}^l$. Let $f : \{0,1\}^l \to \{0,1\}^l$ and $g : \{0,1\}^l \times \{0,1\}^n \to \{0,1\}^l$ be two functions and, for uniformly—and independently of $Y'Y$—distributed random variables C' and C with ranges $\{0,1\}^l$, let*

$$\alpha := \text{Prob}_{\,Y'YC'C}[p_{Y'}(C') = f(C') \ \text{ and } \ p_Y(C) = g(C, Y')] \ ,$$

where $p_{\cdot}(\,\cdot\,)$ is the polynomial function (3). *Then there exists $y \in \mathcal{Y}$ with*

$$P_Y(y) \geq \left(\alpha - \frac{2k}{2^l}\right)^k \ .$$

Proof. Let for every particular value $y' \in \mathcal{Y}'$

$$r_{y'} := \text{Prob}_{\,C'}[p_{y'}(C') = f(C')] = \frac{|\{c' \in \{0,1\}^l \,|\, p_{y'}(c') = f(c')\}|}{2^l} \ ,$$

and for every pair $(y, y') \in \mathcal{Y} \times \mathcal{Y}'$

$$r_{y|y'} := \text{Prob}_{\,C}[p_y(C) = g(C, y')] = \frac{|\{c \in \{0,1\}^l \,|\, p_y(c) = g(c, y')\}|}{2^l} \ .$$

Then we have

$$\alpha = \text{E}_{Y'Y}[r_{Y'} \cdot r_{Y|Y'}] \ . \tag{4}$$

Let us now consider the random experiment defined by

$$P_{Y'YC'_1 \cdots C'_k C_1 \cdots C_k} := P_{Y'Y} \cdot P_{C'_1 \cdots C'_k C_1 \cdots C_k} \ ,$$

where $P_{C'_1 \cdots C'_k C_1 \cdots C_k}$ is the uniform distribution over the subset of $(\{0,1\}^l)^{2k}$ satisfying that all the C'_i and all the C_i are distinct among each other. We then have

$$\begin{aligned}
&\text{Prob}\,[p_{Y'}(C'_i) = f(C'_i) \text{ for } i = 1, \ldots, k \ \text{ and } \ p_Y(C_i) = g(C_i, Y') \text{ for } i = 1, \ldots, k] \\
&\quad \geq \ \text{E}_{Y'Y}\left[r_{Y'} \cdot \left(r_{Y'} - \frac{1}{2^l}\right) \cdots \left(r_{Y'} - \frac{k-1}{2^l}\right) \cdot r_{Y|Y'} \cdots \left(r_{Y|Y'} - \frac{k-1}{2^l}\right)\right] \\
&\quad \geq \ \text{E}_{Y'Y}\left[\left(r_{Y'} - \frac{k-1}{2^l}\right)^k \left(r_{Y|Y'} - \frac{k-1}{2^l}\right)^k\right] \\
&\quad \geq \ \text{E}_{Y'Y}\left[\left(r_{Y'} \cdot r_{Y|Y'} - (r_{Y'} + r_{Y|Y'}) \cdot \frac{k-1}{2^l}\right)^k\right] \ \geq \ \left(\alpha - \frac{2k}{2^l}\right)^k \ .
\end{aligned} \tag{5}$$

The last inequality in (5) follows from the fact that $x \mapsto x^k$ is a convex function and Jensen's inequality [5], from (4), and from $r_{Y'}, r_{Y|Y'} \leq 1$.

Let $\mathcal{A}_k$ be the event the probability of which is bounded in (5). Since, for $x \in \{0,1\}^n$, k values $p_x(c)$ (for k distinct c's) uniquely determine x, we have, given that $\mathcal{A}_k$ occurs, that Y' is uniquely determined and Y is uniquely determined given Y'; together, we get that there exist $y' \in \mathcal{Y}'$ and $y \in \mathcal{Y}$ such that $P_{Y'Y|\mathcal{A}_k}(y', y) = 1$, hence $P_Y(y) \geq \text{Prob}\,[\mathcal{A}_k]$ for this particular value y. □

We will now state and prove the described property of the interactive authentication protocol AUTH (Theorem 3). This and other proofs in the paper make use of Lemma 2 (see [4], [17], [19]), which implies that when d (physical) bits of side information about a random variable are leaked, then its conditional min-entropy is not reduced by much more than d except with small probability.

Lemma 2. [4], [17], [19] *Let S, V, and W be random variables such that S and V are independent, and let $b \geq 0$. Then*

$$\text{Prob}_{\,VW}[H_\infty(S|V = v, W = w) \geq H_\infty(S) - \log|\mathcal{W}| - b] \geq 1 - 2^{-b}\ .$$

Theorem 3. *Assume that two parties Alice and Bob know n-bit strings Y' and Y, respectively. Given that $H_\infty(Y|U = u) \geq tn$ holds for some constant $0 < t \leq 1$, where $U = u$ summarizes an adversary Eve's entire knowledge, Alice can use Protocol AUTH to send authenticated messages of length m of order at most $O(tn/(\log n)^2)$ to Bob by communication over a completely insecure channel. The probability of a successful active attack, which is the event that Bob accepts although the message he received is not the correct one (or although Alice rejects) is of order $2^{-\Omega(tn/m)}$. If, on the other hand, Eve is passive and $Y' = Y$ holds, then Alice and Bob accept with certainty and Bob receives the correct message.*

Proof. Let m be the length of the message Alice wants to send to Bob; the number of executions of the three-round step in Protocol AUTH is then $M = (1+o(1))m$.

Since each party responds to at most M challenges during the protocol execution (and would then reject and abort), the min-entropy of Y, from Eve's viewpoint, at *any* point of the protocol, given all the communication $C = c$ she has seen, is, according to Lemma 2 (applied $2M$ times), at least

$$H_\infty(Y|U = u, C = c) \geq tn - 2Ml - 2Ma$$

with probability at least $1 - 2M2^{-a}$. We conclude that there exist choices of the protocol parameters of order $l = \Theta(n/M)$ and $k = \Theta(M)$—and a suitable choice of the auxiliary parameter a—such that we get the following:

$$\text{There exists } f(n) = \Omega(n) \text{ with Prob}\,[H_\infty(Y|U = u, C = c) \leq f(n)] \leq 2^{-f(n)}\ . \tag{6}$$

As described above, a successful attack of the protocol implies that Eve has been able to answer a challenge generated by one of the parties without help from the other party (i.e., without receiving any message from the other party between receiving the challenge and sending the corresponding response). The first possibility is that a challenge of *Alice* is responded without Bob's help; here, it is necessary for Eve to also answer at least one of Bob's challenges successfully (an attack is successful only if Bob is fooled)—possibly with Alice's "help", however. Let therefore $\mathcal{A}$ be the event that Eve correctly responds to one of the at most M challenges by Alice, and to one of Bob's at most M challenges *given Alice's string* Y'. According to Lemma 1, and because of the union bound, we have

$$\mathrm{Prob}\,[\mathcal{A}] \leq M^2 \cdot \left(2^{-H_\infty(Y|U=u,C=c)/k} + 2k/2^l\right) \ .$$

Hence, because of (6), the success probability of this attack is at most

$$M^2 \cdot \left(2^{-\Omega(n/M)} + \frac{\Theta(M)}{2^l}\right) + 2^{-\Omega(n)} = 2^{-\Omega(n/M)}$$

(note that $M/2^l = 2^{-\Omega(n/M)}$ and $M^2 2^{-\Omega(n/M)} = 2^{-\Omega(n/M)}$ hold since $M = O(tn/(\log n)^2)$). The second possibility of an attack is that a challenge of *Bob* is responded without Alice's help. The probability of this is, because of (6) and by a similar but simpler reasoning as the one used above, of order $2^{-\Omega(n/M)}$. The application of the union bound concludes the proof. □

2.2 Information Reconciliation over Unauthenticated Channels

We will now use the described authentication protocol, and its new property established in the previous section, for the construction of a protocol for *information reconciliation* by completely insecure communication. Information reconciliation is interactive error correction: Two parties, knowing strings X and Y, respectively, should share a common string at the end (e.g., one of the initial strings). The idea is to use Protocol AUTH in such a way that the parties can detect active attacks at any point in the protocol.

According to Lemma 4, the error correction itself can be done by exchanging redundancy, where the latter is generated by applying universal hashing[5] to the input strings; this is efficient with respect to the required communication, but computationally inefficient for one of the parties (Alice in our case). In the special but typical scenario where X and Y are bitstrings which differ in a certain limited number of positions, more efficient methods, based on *concatenated codes* [10], can be used instead in Protocol IR below.

Lemma 4. *Let X and Y be distributed according to P_{XY} such that $H_0^{\max}(Y|X) \leq r$ holds. Let, for some integer $s \geq 0$, $\mathcal{H}$ be a universal class of functions $h : \mathcal{Y} \to \{0,1\}^{r+s}$, and let H be the random variable corresponding to the random choice, independently of X and Y, of a function in $\mathcal{H}$ according to the uniform distribution. Then*

$$\mathrm{Prob}\left[\text{there exists } \overline{Y} \neq Y \text{ with } H(\overline{Y}) = H(Y) \text{ and } P_{Y|X}(Y,X) > 0\right] \leq 2^{-s} \ .$$

[5] A class $\mathcal{H}$ of functions $h : \mathcal{A} \to \mathcal{B}$ is *2-universal*—or *universal* for short—if, for all $a, a' \in \mathcal{A}$, $a \neq a'$, we have $|\{h \,|\, h(a) = h(a')\}| = |\mathcal{H}|/|\mathcal{B}|$.

Proof. For $x \in \mathcal{X}$, let $\mathcal{Y}_x := \{y \in \mathcal{Y} \,|\, P_{Y|X}(y, x) > 0\}$. We have $|\mathcal{Y}_x| \leq 2^r$. Since for any $y, \overline{y} \in \mathcal{Y}_x$, $\overline{y} \neq y$, and random $H \in \mathcal{H}$ the probability that $H(y) = H(\overline{y})$ holds is at most $1/2^{r+s}$, we have

$$\begin{aligned} &\mathrm{Prob}_{YH}[\text{there exists } \overline{Y} \in \mathcal{Y}_x,\ \overline{Y} \neq Y,\ \text{such that } H(\overline{Y}) = H(Y)] \\ &\quad \leq\ |\mathcal{Y}_x| \cdot \mathrm{Prob}\,[H(\overline{y}) = H(y) \text{ for some } \overline{y} \neq y]\ \leq\ 2^r \cdot \frac{1}{2^{r+s}}\ =\ \frac{1}{2^s} \end{aligned}$$

by the union bound. The statement then follows when the expectation over X is taken. □

In Protocol IR, D and T are parameters to be determined below, and $\mathcal{H}$ is a universal class of functions from $\{0,1\}^n$ to $\{0,1\}^D$. Furthermore, $\mathrm{AUTH}_{Y',Y}(M)$ means that the message M is sent using Protocol AUTH, where the "keys" used by the sender (Alice) and the receiver (Bob) are Y' and Y, respectively.

Protocol IR (Information Reconciliation)

Alice		**Bob**
$X \in \{0,1\}^n$		$Y \in \{0,1\}^n$
		$H \in_r \mathcal{H}$, $H\,:\,\{0,1\}^n \to \{0,1\}^D$
	$\xleftarrow{H,\ H(Y)}$	
$Y' \in \mathcal{Y}_X$ with $H(Y') = H(Y)$ $R \in_r \{0,1\}^T$ compute $p_{Y'}(R)$		
	$\xrightarrow{\mathrm{AUTH}_{Y',Y}((R, p_{Y'}(R)))}$	
accept, Y'		if $p_{Y'}(R) = p_Y(R)$: accept, Y otherwise: reject.

The content of the second message serves as a verification of whether the string Y' computed by Alice is correct. Clearly, it has to be authenticated because of possible substitution attacks. It is an interesting point that because of this authentication, Alice can choose the "challenge" string R herself: If the authentication is successful, Bob knows that R is indeed the challenge generated by Alice, and hence random.

Note that although applied in a—symmetric—context where two parties want to generate a common secret key, Protocol IR is secure (for Bob) in the same—asymmetric—sense as the authentication protocol: Either everything goes well or Bob will know it did not (with high probability).

Theorem 5. *Assume that two parties Alice and Bob know the value of a random variable X and an n-bit string Y, respectively, and that*

$$H_\infty(Y|U = u) - H_0^{\max}(Y|X) \geq tn \tag{7}$$

holds for some constant $0 < t \leq 1$, where $U = u$ summarizes an adversary's entire knowledge. Then Protocol IR (with suitable parameter choices)—carried out over a completely insecure channel—achieves the following. If Eve is passive, then Alice and Bob both accept and the string Y' computed by Alice is equal to Y except with probability $2^{-\Omega(n/\log n)}$. In general, it is true except with probability $2^{-\Omega(\sqrt{n}/\log n)}$ that either Bob rejects or both accept and $Y' = Y$ holds. Furthermore, the remaining conditional min-entropy of Y given Eve's initial information and the protocol communication is of order $(1-o(1))tn$ with probability $1 - 2^{-\Omega(n/\log n)}$.

Proof. Let us assume that Eve is passive. Let the parameter D be of order $D = H_0^{\max}(Y|X) + \Theta(n/\log n)$. Then we have, according to Lemma 4, that Alice's guess Y'—from X and $H(Y)$—is uniquely determined and hence correct except with probability $2^{-\Omega(n/\log n)}$.

Let us now consider the general case where Eve is possibly an active adversary. We first analyze the properties of the authentication of the confirmation message sent from Alice to Bob. Let the parameter T be of order $T = \Theta(\sqrt{n})$. We will argue that with high probability, either Bob rejects or Alice and Bob both accept and the values $(R, p_{Y'}(R))$ as received by Bob are the ones sent by Alice and, finally, that this implies that $Y' = Y$ holds, i.e., that Alice and Bob share the same string, the min-entropy of which, from Eve's viewpoint, is still $(1 - o(1))tn$.

First, we get, using Lemma 2 with the parameter choice $b = \Theta(n/\log n)$, that there exist functions $f(n) = (1 - o(1))tn$ and $g(n) = \Omega(n/\log n)$ such that

$$\text{Prob}\left[H_\infty(Y|U = u, H = h, H(Y) = h(y)) \geq f(n)\right] \geq 1 - 2^{-g(n)} \; .$$

Because of this, Theorem 3 implies that the authentication works—even if, for instance, Eve had modified the error-correction information sent in the first message and knows Y' perfectly. The length of the message to be authenticated with Protocol AUTH is of order $\Theta(\sqrt{n})$, and we choose the protocol parameter l to be $l = \Theta(\sqrt{n}/\log n)$ to make sure that the remaining min-entropy, given all the communication, is still an arbitrarily large fraction of tn. The success probability of the protocol is then, according to the proof of Theorem 3, $1 - 2^{-\Omega(\sqrt{n}/\log n)}$.

Let us hence assume now that Bob actually received the correct message $(R, p_{Y'}(R))$ as sent by Alice. Since R are the truly random bits (in particular, independent of Y') chosen by Alice, and since $p_y(r) = p_{y'}(r)$ can hold for at most $\deg(p_y) = n/T - 1 = \Theta(\sqrt{n})$ different values of r for any $y' \neq y$, we have that with probability $1 - 2^{-\Omega(\sqrt{n})}$ either Alice has the correct string, or Bob realizes that she does not.

Finally, the remaining min-entropy is still roughly the same with high probability since the total number of bits sent is of order $\Theta(n/\log n) = o(n)$. From Lemma 2, we get that there exist $f(n) = (1 - o(1))tn$ and $g(n) = \Omega(n/\log n)$ such that we have $\text{Prob}\left[H_\infty(Y|U = u, C = c) \geq f(n)\right] \geq 1 - 2^{-g(n)}$, where $C = c$ is the entire protocol communication. This concludes the proof. □

Remark. In Theorem 5—as well as in Theorems 6, 7, and 8 and Corollary 9 below—the assumed entropy bounds can be conditioned on an event $\mathcal{A}$ if at the same time the claimed protocol failure probabilities are increased by $1-\text{Prob}[\mathcal{A}]$. An example for which this can lead to substantially stronger statements is when the random variables $X = (X_1, \ldots, X_n)$, $Y = (Y_1, \ldots, Y_n)$, and $U = (U_1, \ldots, U_n)$ arise from n independent repetitions of a certain random experiment $P_{X_i Y_i U_i}$. In this case, $\mathcal{A}$ can be the event that the actual outcome sequences are *typical* (see [5]). This is a good choice because $\mathcal{A}$ occurs except with exponentially (in n) small probability, and because

$$H_\infty(Y|U=u,\mathcal{A}) \approx H(Y_i|U_i)\cdot n \gg H_\infty(Y|U=u)$$

and

$$H_0^{\max}(Y|X,\mathcal{A}) \approx H(Y_i|X_i)\cdot n \ll H_0^{\max}(Y|X)$$

can hold. (See also Example 1 below.)

2.3 The Price for One-Sided Authenticity and Confidentiality

In [19], Protocol PA, allowing for *privacy amplification* over a completely insecure channel, was presented. Privacy amplification [3], [2] means to generate, from an only weakly secret shared string, a shorter but highly secret key. Protocol PA—which uses Protocol AUTH as well as *extractors* as its main ingredients—has been shown to extract virtually all the min-entropy of an arbitrarily weakly secret string.

Theorem 6. [19] *Assume that Alice and Bob both know the same n-bit string Y satisfying $H_\infty(Y|U=u) \geq tn$ for some constant $0 < t \leq 1$, where $U = u$ summarizes Eve's entire information about Y. Then Protocol PA, using two-way communication over a completely insecure channel, has the following properties. Both Alice and Bob either reject or accept and compute strings S_A and S_B, respectively, such that if Eve is passive, then Alice and Bob accept and there exists a $(1-o(1))tn$-bit string S that is uniformly distributed from Eve's viewpoint and such that $S_A = S_B = S$ holds except with probability $2^{-\Omega(n/(\log n)^2)}$. In general (i.e., if Eve is possibly active), either both parties reject or there exists a string S with the above properties, except with probability $2^{-\Omega(\sqrt{n}/\log n)}$.*

Putting everything together, we can now conclude that the combination of Protocols IR and PA achieves what we had stated initially, namely asymmetric unconditional security for Bob from a very weak initial primitive. Given that Bob accepts at the end of the protocol, he shares a secret key with Alice. He can then send unconditionally confidential messages *to* her and receive authenticated messages *from* her.

Theorem 7. *Assume that two parties Alice and Bob know a random variable X and an n-bit string Y, respectively, and that $H_\infty(Y|U=u) - H_0^{\max}(Y|X) \geq tn$ holds for some constant $0 < t \leq 1$, where $U = u$ summarizes an adversary's entire knowledge. Then the combination of Protocols IR and PA, carried out* *over a completely insecure channel, achieves the following. Alice and Bob both either reject or accept and compute strings S_A and S_B, respectively, such that if Eve is passive, then Alice and Bob accept and there exists a $(1-o(1))tn$-bit string S that is uniformly distributed from Eve's viewpoint and such that $S_A = S_B = S$ holds except with probability $2^{-\Omega(n/(\log n)^2)}$. In general, either Bob rejects or Alice and Bob accept, and the above holds, except with probability $2^{-\Omega(\sqrt{n}/\log n)}$.*

Proof. Follows from Theorems 5 and 6. □

3 Impossibility Results and Lower Bounds

3.1 Two-Sided Security Requires Stronger Conditions

All protocols presented in Section 2 are asymmetric in the sense that the generated security is on Bob's side only. (Alice, for instance, could be talking to Eve instead of Bob without realizing this.) Example 1 shows that security for Alice simply *cannot* be achieved under assumptions as weak as that. This implies that the price for unconditional security on one side is strictly lower than for such security on both sides. The same is already well-known in the computational-security model, as the example of public-key cryptography demonstrates.

Example 1. Let $X = (X_1, \ldots, X_n)$ be a uniformly distributed n-bit string, and let $Y = (Y_1, \ldots, Y_n)$ and $U = (U_1, \ldots, U_n)$ be n-bit strings jointly distributed with X according to[6]

$$P_{YU|X}((y_1, \ldots, y_n), (u_1, \ldots, u_n), (x_1, \ldots, x_n)) = \prod_{i=1}^{n} |\delta_{y_i x_i} - \varepsilon| \cdot |\delta_{u_i x_i} - \varepsilon| \quad (8)$$

for some $0 < \varepsilon < 1/2$. Equation (8) means that the i-th bits of Y and U are generated by sending X_i over two independent binary symmetric channels with error probability ε.

Let now $\mathcal{A}$ be the event—which occurs except with exponentially (in n) small probability—that all the involved strings are typical sequences. Then we have, roughly,[7] $H_\infty(Y|U = u, \mathcal{A}) \approx h(2\varepsilon - 2\varepsilon^2)n$ and $H_0^{\max}(Y|X, \mathcal{A}) \approx h(\varepsilon)n$. Because of $2\varepsilon - 2\varepsilon^2 > \varepsilon$, the condition of Theorem 7 is satisfied. On the other hand, Bob has no advantage over Eve from Alice's viewpoint since Eve is able to *simulate* [17] Bob towards Alice: She can generate a random variable from U—in fact, she can use U itself—which has the same joint distribution with X as Y does—$P_{XU} = P_{XY}$. Hence Alice will never be able to tell Bob and Eve apart.

[6] Here, δ_{ij} is the Kronecker symbol, i.e., $\delta_{ij} = 1$ if $i = j$ and $\delta_{ij} = 0$ otherwise.

[7] We denote by $h(p)$ the binary entropy function $h(p) = -(p \log p + (1-p) \log(1-p))$.

3.2 Optimality of the Achieved Secret-Key Length

The protocols we have presented in Section 2 are universal and work for all possible specific probability distributions under the only assumption that the entropy condition (7) is satisfied. In other words, our protocols work for large *classes* of probability distributions, where Alice and Bob do not have to know the nature of Eve's information, i.e., the particular distribution, but only that the corresponding entropy bound is satisfied. In this sense, our protocols are optimal: In many situations, no protocol can extract a longer secret key—*even when the communication channel is assumed authentic*. (It should be noted, however, that there are *specific* settings in which key agreement by *authenticated* public communication is possible even though the expression in (7) is negative [13].)

This can be illustrated with the setting where Bob's random variable Y is uniformly distributed (also from Eve's viewpoint) and Alice's X provides her uniformly with *deterministic* information about Y: For every value x it can take, $P_{Y|X=x}$ is the uniform distribution over the set $|\mathcal{Y}_x|$ of size $|\mathcal{Y}|/|\mathcal{X}|$ (and these sets are disjoint for different values of X). After the execution of a key-agreement protocol, Alice has to know (with overwhelming probability) the key S generated by Bob. Eve, on the other hand, should be (almost) completely ignorant about it. Clearly, this can be satisfied only if there are at least as many possible values Alice can initially have as possible keys. Therefore, we always have, roughly, $|\mathcal{S}| \leq |\mathcal{X}| = |\mathcal{Y}|/|\mathcal{Y}_x|$, and hence

$$\mathrm{len}(S) \approx \log|\mathcal{S}| \leq \log|\mathcal{Y}| - \log|\mathcal{Y}_x| = H_\infty(Y) - H_0^{\max}(Y|X) \ .$$

4 Two-Way Security Under New and Weaker Assumptions

In this section we determine the price for achieving unconditional security for *both* Alice and Bob. The conditions we will find are weaker than the ones known previously (such as, for instance, a highly insecure but *common* string [19]).

We first give Protocol IR+, an extension of Protocol IR offering security also for Alice. After the first two protocol steps, which are the same as in Protocol IR, Alice sends error correction information $H'(X)$ about her initial string X (here, H' is from a universal class $\mathcal{H}'$ with suitable parameters) to Bob, who then uses his "estimate" X' of X as the authentication key for sending a challenge-response pair for Y. If Alice receives this correctly, and if it corresponds to the value $p_{Y'}(R')$ she can compute herself, she can be convinced that $Y' = Y$ holds. The crucial observation for proving Theorem 8 is that the given entropy condition on Y *also* implies that Eve, having seen all the error-correction information and other messages, still has $\Omega(n)$ of min-entropy about X—because the same holds for Y. The reason is that given all the protocol communication, Y can—with overwhelming probability—be computed from X, and vice versa.

Protocol IR+ (Two-Secure Information Reconciliation)

Alice		**Bob**
$X \in \{0,1\}^n$		$Y \in \{0,1\}^n$
		$H \in_r \mathcal{H}$, $H : \{0,1\}^n \to \{0,1\}^D$
	$\xleftarrow{H,\ H(Y)}$	
$Y' \in \mathcal{Y}_X$ with $H(Y') = H(Y)$ $R \in_r \{0,1\}^T$ compute $p_{Y'}(R)$		
	$\xrightarrow{\mathrm{AUTH}_{Y',Y}((R, p_{Y'}(R)))}$	
$H' \in_r \mathcal{H}'$, $H' : \{0,1\}^n \to \{0,1\}^{D'}$		if $p_{Y'}(R) \neq p_Y(R)$: reject
	$\xrightarrow{H',\ H'(X)}$	
		$X' \in \mathcal{X}_Y$ with $H'(X') = H'(X)$ $R' \in_r \{0,1\}^T$
	$\xleftarrow{\mathrm{AUTH}_{X',X}((R', p_Y(R')))}$	
if $p_Y(R') \neq p_{Y'}(R')$: reject otherwise: accept, Y'		accept, Y.

Theorem 8. *Assume that two parties Alice and Bob know n-bit strings X and Y, respectively, and that*

$$H_\infty(Y|U = u) - H_0^{\max}(Y|X) - H_0^{\max}(X|Y) \geq tn$$

holds for some constant $0 < t \leq 1$, where $U = u$ summarizes an adversary's entire knowledge. Then Protocol IR+ (for suitable parameter choices)—carried out over a completely insecure channel—achieves the following. If Eve is passive, then Alice and Bob both accept and the string Y' computed by Alice is equal to Y except with probability $2^{-\Omega(n/\log n)}$. In general, it is true except with probability $2^{-\Omega(\sqrt{n}/\log n)}$ that either both parties reject or $Y' = Y$ holds. Furthermore, the remaining min-entropy of Y given Eve's initial information and the protocol communication is of order $(1 - o(1))tn$ with probability $1 - 2^{-\Omega(n/\log n)}$.

Proof. Follows from Theorem 5, Lemma 4, and Theorem 3. □

Corollary 9. *Assume that two parties Alice and Bob know n-bit strings X and Y, respectively, and that*

$$H_\infty(Y|U = u) - H_0^{\max}(Y|X) - H_0^{\max}(X|Y) \geq tn$$

holds for some constant $0 < t \leq 1$, where $U = u$ summarizes an adversary's entire knowledge. Then the combination of Protocols IR+ and PA, carried out over a completely insecure channel, achieves the following. Alice and Bob both either reject or accept and compute strings S_A and S_B, respectively, such that if Eve is passive, then Alice and Bob both accept and there exists a $(1 - o(1))tn$-bit string S that is uniformly distributed from Eve's viewpoint and such that $S_A = S_B = S$ holds except with probability $2^{-\Omega(n/(\log n)^2)}$. In general, either both parties reject or there exists a string S with the above properties, except with probability $2^{-\Omega(\sqrt{n}/\log n)}$.

Proof. Follows from Theorems 8 and 6. □

5 Concluding Remarks

In this paper we have determined, so to speak, a minimal price for unconditional security. For two parties connected by a completely insecure bidirectional communication channel, we have described the weakest possible information-theoretic primitive necessary for obtaining security on *one* end of the channel—i.e., guaranteed *exclusivity* of read and write access to the channel on its other end. Roughly speaking, we found that whenever Eve's uncertainty about the information of the party at one end of the channel exceeds the uncertainty about the same information as seen by the party at the channel's other end, then the entire entropy difference can be transformed into a key which is secret for the former party. This asymmetric notion of security for one party means that *either* the two parties share a secret key, *or* this—designated—party knows that they do not.

One of the consequences of our protocols is that the required conditions for the possibility of secret-key agreement in general, and quantum key distribution in particular, can be relaxed substantially: Quantum key agreement has sometimes been perceived to be rather key *extension* than actual key *generation* in view of the usually-made assumption that the two parties share a short unconditionally secret key already initially, from which they can then produce a longer key (where the initial key is required for authenticating the public communication exchanged for processing the raw key). Our results show that this condition is unnecessary and can be replaced by a much weaker assumption.

References

1. C. H. Bennett and G. Brassard, Quantum cryptography: public key distribution and coin tossing, *Proceedings of the IEEE International Conference on Computers, Systems and Signal Processing*, pp. 175–179, 1984.
2. C. H. Bennett, G. Brassard, C. Crépeau, and U. M. Maurer, Generalized privacy amplification, *IEEE Trans. on Information Theory*, Vol. 41, No. 6, pp. 1915–1923, 1995.
3. C. H. Bennett, G. Brassard, and J.-M. Robert, Privacy amplification by public discussion, *SIAM Journal on Computing*, Vol. 17, pp. 210–229, 1988.

4. C. Cachin, *Entropy measures and unconditional security in cryptography*, Ph. D. Thesis, ETH Zürich, Hartung-Gorre Verlag, Konstanz, 1997.
5. T. M. Cover and J. A. Thomas, *Elements of information theory*, Wiley Series in Telecommunications, 1992.
6. I. Csiszár and J. Körner, Broadcast channels with confidential messages, *IEEE Trans. on Information Theory*, Vol. 24, pp. 339–348, 1978.
7. W. Diffie and M. E. Hellman, New directions in cryptography, *IEEE Trans. on Information Theory*, Vol. 22, No. 6, pp. 644–654, 1976.
8. Y. Dodis and A. Smith, Fooling an unbounded adversary with a short key: a notion of indistinguishability, relations to extractors, and lower bounds, manuscript, 2003.
9. Y. Dodis and J. Spencer, On the (non)universality of the one-time pad, *Proceedings of FOCS 2002*, pp. 376–385, 2002.
10. G. D. Forney Jr., *Concatenated codes*, Massachusetts Institute of Technology, Cambridge, Massachusetts, 1966.
11. P. Gemmell and M. Naor, Codes for interactive authentication, *Advances in Cryptology - CRYPTO '93*, LNCS, Vol. 773, pp. 355–367, Springer-Verlag, 1993.
12. R. König, U. M. Maurer, and R. Renner, On the power of quantum memory, available on www.arxiv.org, quant-ph/0305154, 2003.
13. U. M. Maurer, Secret key agreement by public discussion from common information, *IEEE Trans. on Information Theory*, Vol. 39, No. 3, pp. 733–742, 1993.
14. U. M. Maurer, Information-theoretically secure secret-key agreement by NOT authenticated public discussion, *Advances in Cryptology - EUROCRYPT '97*, LNCS, Vol. 1233, pp. 209–225, Springer-Verlag, 1997.
15. U. M. Maurer, Information-theoretic cryptography, *Advances in Cryptology - CRYPTO '99*, LNCS, Vol. 1666, pp. 47–64, Springer-Verlag, 1999.
16. U. M. Maurer and P. Schmid, A calculus for security bootstrapping in distributed systems, *Journal of Computer Security*, Vol. 4, No. 1, pp. 55–80, 1996.
17. U. M. Maurer and S. Wolf, Secret-key agreement over unauthenticated public channels – Parts I–III, *IEEE Trans. on Information Theory*, Vol. 49, No. 4, pp. 822–851, 2003.
18. J. L. McInnes and B. Pinkas, On the impossibility of private key cryptography with weakly random keys, *Advances in Cryptology - CRYPTO '90*, LNCS, Vol. 537, pp. 421–436, Springer-Verlag, 1990.
19. R. Renner and S. Wolf, Unconditional authenticity and privacy from an arbitrarily weak secret and completely insecure communication, *Advances in Cryptology - CRYPTO 2003*, LNCS, Vol. 2729, pp. 78–95, Springer-Verlag, 2003.
20. A. Russell and H. Wang, How to fool an unbounded adversary with a short key, *Advances in Cryptology - EUROCRYPT 2002*, LNCS, Vol. 2332, pp. 133–148, Springer-Verlag, 2002.
21. C. E. Shannon, Communication theory of secrecy systems, *Bell System Technical Journal*, Vol. 28, pp. 656–715, 1949.
22. A. D. Wyner, The wire-tap channel, *Bell System Technical Journal*, Vol. 54, No. 8, pp. 1355–1387, 1975.