Long Term Confidentiality: a Survey*

Johannes Braun[1], Johannes Buchmann[1], Ciaran Mullan[1], and Alex Wiesmaier[2]

[1] Technische Universität Darmstadt
Hochschulstraße 10, 64283 Darmstadt, Germany
{jbraun,buchmann,cmullan}@cdc.informatik.tu-darmstadt.de
[2] AGT Group (R&D) GmbH
Hilpertstraße 20a, 64295 Darmstadt, Germany
awiesmaier@agtgermany.com

Abstract. Sensitive electronic data may be required to remain confidential for long periods of time. Yet encryption under a computationally secure cryptosystem cannot provide a guarantee of long term confidentiality, due to potential advances in computing power or cryptanalysis. Long term confidentiality is ensured by information theoretically secure ciphers, but at the expense of impractical key agreement and key management. We overview known methods to alleviate these problems, whilst retaining some form of information theoretic security relevant for long term confidentiality.

Keywords: Long term confidentiality, Information theoretic security, Perfect secrecy, Everlasting security

1 Introduction

Consider the following scenario: Alice encrypts her data today using the most trusted computationally secure cryptosystem. An adversary Eve gains access to the encrypted data and stores it in a safe place. In 20 or 30 years time the cryptosystem is broken, due to a novel cryptanalysis or feasibility of a brute force attack on the key. A patient, passive Eve finally learns Alice's secrets.

Long term data security is a challenging goal, and from a cryptographic point of view may be divided into long term authenticity and long term confidentiality. While much work has been done on long term authenticity, which can be accomplished with encapsulation and re-signing [40, 88], achieving practical long term confidentiality is a major open problem [19].

This article provides a survey of cryptographic schemes relevant for long term confidentiality. We make an albeit hazy distinction between long term confidentiality of *stored data*, which requires a method to distribute and store keys, and long term confidentiality of *transmitted data*, which may only require secure key distribution.

The phrase 'long term' varies in the literature from 'longer periods of time' [22,57] to indefinitely [80, 91, 105, 106]. Of course, in theory one can store electronic data indefinitely at a physically secure, isolated location. But this imposes all sorts of impracticalities, so is not of major interest. It is also clear that in certain scenarios the requirement of confidentiality diminishes or vanishes over time, for example patents after acceptance and subsequent publication. Yet some forms of data do require long

* Accepted for publication in Designs, Codes and Cryptography, Springer. The final publication is available at www.springerlink.com (DOI : 10.1007/s10623-012-9747-6)

term confidentiality, often by law. For example, German law stipulates that medical and legal data remain confidential from third parties even after death of a patient or client [50]. As more and more sensitive medical and legal data is processed and stored electronically, ensuring its long term confidentiality is highly desirable. Non-legal scenarios where long term confidentiality is desirable include industrial, and governmental or military secrets (e.g. defence strategies, construction plans).

The current trend towards cloud computing means that more and more data is being processed and stored by online resources beyond physical and logical control of the owner. It is a simple task for an adversary to intercept, copy, and store any data sent across a public network and from this point on, confidentiality is determined solely by the original encryption scheme. Hence re-encryption is useless, and even deletion of data cannot be guaranteed.

Current cryptographic schemes in wide deployment today, such as RSA, Diffie–Hellman, and AES, do not offer long term confidentiality guarantees. This is because concrete security is based on the current infeasibility of a specific computational problem, such as factoring a 1024-bit RSA modulus or computing a 128-bit AES key, and there are no known techniques to prove the hardness of such problems.

Indeed, if quantum computers are fully realised then factoring and computing discrete logarithms will become feasible, and hence most public key cryptosystems in use will become insecure [101]. Even post-quantum schemes [19], which provide resistance against quantum adversaries, offer no long term guarantees of confidentiality, as they are still based in the computational model of security.

The effect of quantum computers on the security of symmetric key ciphers (e.g. due to Grover's algorithm [51]) is considerably less severe, and may be compensated by e.g. doubling key sizes [13]. Nevertheless, data encrypted with key sizes considered secure today (e.g. 128 bit) are still at risk of retroactive decryption.

At best, future security is estimated heuristically, with variable key length recommendations [15,65]. The block cipher DES provides a good illustration. Since its FIPS standardisation in 1976 [82], DES enjoyed widespread deployment for over 20 years. But by the late 1990s computers were so cheap and powerful that a 2^{56} brute force search for the key became a feasible task [38]. Although this development was anticipated by security experts and DES was replaced by Triple-DES, the confidentiality protection of prior ciphertexts was lost.

Yet cryptosystems *do* exist that provide long term confidentiality, and were known long before the modern era of public key cryptography. Such schemes are said to be *information theoretically secure*, *unconditionally secure*, or *perfectly secret*. Since every public key cryptosystem is susceptible to a brute force search on the key, unconditionally secure ciphers (such as the famous One Time Pad) necessarily belong in the realm of symmetric key cryptography.

However, such ciphers are not used in practice (at least on any moderate scale) because of the practical problems of key agreement and key management: how do Alice and Bob securely and practically agree on large amounts of secret key material, and how do they securely and practically store such material? Attempts to solve these problems whilst retaining some form of information theoretic security has led to a closer scrutiny of the assumptions underlying perfect secrecy, which may be stated as:

1. *An adversary has complete and perfect access to the channel.*
2. *Given any ciphertext, a computationally unbounded adversary learns no information about the corresponding plaintext.*

Here the *channel* refers to any medium whereby communicating parties Alice and Bob exchange information. This channel may be *authenticated*, in which case Alice and Bob have a guarantee that they are talking to each other and not an imposter Eve, or the channel may be *unauthenticated*, and there is no such guarantee.

We will see how by relaxing one or other of the above two assumptions, it is possible to make progress on the problems of key agreement and key management. However, all of the methods we discuss currently have limitations, either in terms of practicality or by imposing unrealistic assumptions on an adversary. Nevertheless, the ideas discussed motivate very challenging research problems.

We mention several related survey articles which overlap to some extent with our exposition [2,7,48,72,97,103,116]. This article differs through having a wider scope by bringing together all the models relevant for long term confidentiality. Moreover, we provide a discussion on current limitations and practicalities of each model considered.

Outline. Section 2 introduces the basic notions of entropy and perfect secrecy relevant for later discussions. In Section 3 we survey information theoretic key agreement protocols, relevant for long term confidentiality of transmitted data, and in Section 4 we discuss information theoretic secret sharing schemes, relevant for long term confidentiality of stored data. We close with a brief summary in Section 5.

2 Shannon, entropy, and perfect secrecy

Claude Shannon's *A Mathematical Theory of Communication* [100] is a masterpiece. His theorems tell us precisely how reliably we can communicate and how much meaningful information we can convey over a given channel. All of this is achieved in the natural language of *entropy*, a fundamental notion throughout the sciences, and of particular importance for cryptography.

2.1 Entropy

Let X be a discrete random variable over an alphabet χ. The *entropy* of X is defined as

$$H(X) := -\sum_{x \in \chi} P(X = x) \log_2 P(X = x).$$

Intuitively, $H(X)$ is a measure of the uncertainty of an outcome prior to an observation of X, or equivalently, information gained by an observation of X. For example, let X be the event that a coin is tossed n times, and suppose the coin is unbiased. There are 2^n possible outcomes and it is easily seen that $H(X) = n$. (It is convention that logarithms are in base 2 so that entropy is measured in *bits*.) Now suppose we repeat the experiment with a completely biased coin that turns up heads every time. Then the outcome is certain and $H(X) = 0$. In general, entropy is bounded by $0 \leq H(X) \leq log_2|\chi|$ with the upper bound obtained when X is uniformly distributed.

To facilitate our later discussion on information theoretic cryptographic protocols we need a few more concepts.

The *joint entropy* of two random variables X and Y is given by

$$H(XY) := -\sum_x \sum_y P(X = x, Y = y) \log_2 P(X = x, Y = y),$$

and is bounded by $H(XY) \leq H(X) + H(Y)$ with equality when X and Y are statistically independent. The *conditional entropy* of X given Y may be defined as

$$H(X|Y) := H(XY) - H(Y),$$

and the *mutual information* of X and Y is given by

$$I(X;Y) := H(X) - H(X|Y) = H(X) + H(Y) - H(XY).$$

Mutual information is a symmetric measure of the mutual dependence of X and Y, and the conditional entropy $H(X|Y)$ is a measure of how much uncertainty remains about X after observing Y. For example, if X and Y are independent, then $H(X|Y) = H(X)$ and $I(X;Y) = 0$. In general, $H(X|Y) \leq H(X)$, saying that information cannot increase uncertainty. Finally, the *conditional mutual information* of X and Y given Z can be defined as

$$I(X;Y|Z) := H(X|Z) - H(X|YZ).$$

A detailed treatment of the rich subject of information theory is given in [23].

2.2 Perfect secrecy

A *cryptosystem* consists of a set of possible plaintexts, a set of possible ciphertexts, and a set of possible keys, together with some method of encryption and decryption. Let M, C and K denote random variables for plaintext, ciphertext, and key, respectively. A cryptosystem is said to have *perfect secrecy* if $H(M|C) = H(M)$, i.e. that any given ciphertext yields no information about the underlying plaintext.

Perfect secrecy is the strongest possible security measure for a cryptosystem. However, Shannon proved that any cryptosystem satisfying perfect secrecy must also satisfy $H(K) \geq H(M)$. Put plainly, the key must be at least as long as the plaintext. This limit imposes a severe restriction on the practicality of perfectly secure cryptosystems, as distribution and management of such large keys becomes a major issue.

The most famous example of a perfectly secure cryptosystem is the One Time Pad, attributed to Vernam [113]. Encryption of a binary message m with a random key k is simply componentwise XOR: $c = m \oplus k$. Decryption is likewise given by $c \oplus k = (m \oplus k) \oplus k = m$. The One Time Pad is clearly optimal in the sense that the bound $H(K) \geq H(M)$ is met, but still impractical for the reasons just mentioned.

3 Key agreement from information theory

In this section we survey known solutions for information theoretic key agreement. The main models in question are noisy channel models, quantum key distribution, the Bounded Storage Model and the Limited Access Model. We describe the basics of each model and discuss limitations in terms of practicalities and adversarial assumptions. We mention two relevant survey articles of Maurer [72] and Wolf [116] which overlap with parts of this section.

3.1 Noisy channel models

Noise is a naturally occurring phenomenon of every physical communication channel. Remarkably, this simple fact alone allows for key agreement over a public channel. With his wire-tap channel, which is still the focus of ongoing research [8, 9],

Wyner [120] was the first to demonstrate the possibility of key agreement utilising noise. In a one way communication setting where two parties are connected via a noisy channel, Wyner showed that key agreement is possible under the assumption that an adversary Eve receives only a *degraded version* of the received signal. Wyner's model was enhanced by Csiszár and Körner [24] to a more realistic broadcast setting, whereby Eve receives a noisy version of the transmitted signal via a different channel. Key agreement is possible in this setting only if Eve's channel is noisier than the receiver's, again a rather unwanted premise. Allowing for interactive two way communication, Maurer [77] demonstrated the possibility of secret key agreement *even if Eve's channel is less noisy than Alice and Bob's*. We sketch the details of this impressive result.

Maurer's model. Suppose a random bitstring R is publicly broadcast. Alice, Bob, and Eve receive R via different noisy channels with respective error probabilities $\alpha, \beta, \epsilon \neq 0$. Suppose Alice receives the string X, Bob receives the string Y, and Eve receives the string Z, where we need to assume $Z \neq X$ and $Z \neq Y$. Crucially, we suppose $0 < \epsilon < \alpha, \beta < 1/2$, i.e. that Eve has the most reliable channel and

$$I(X;R) < I(Z;R), \text{ and}$$

$$I(Y;R) < I(Z;R).$$

To derive a secure key from the strings X and Y, Alice and Bob first engage in an *advantage distillation phase*. Over a secondary (noiseless, authenticated) public channel, advantage distillation enables Alice and Bob to exchange information about X and Y in such a way that they gain some informational advantage over Eve. More precisely, after an exchange of messages Alice computes a string S_A and Bob computes a string S_B such that

$$I(S_A;T) < I(S_A;S_B), \text{ and}$$
$$I(S_B;T) < I(S_A;S_B),$$

where T summarises Eve's total knowledge. Next, in an *information reconciliation phase*, Alice and Bob use their extra information to agree on a common string S with high probability. Finally, in a *privacy amplification phase*, again by a public exchange of messages Alice and Bob shrink S to form a secure key about which Eve has only negligible information.

Each of these phases is a fairly involved process. Advantage distillation and information reconciliation are achieved using error-correcting codes; details and specific suggestions are given in [18,70,77,116]. Privacy amplification may be achieved via universal hashing or randomness extraction techniques, and may be considered a subject in its own right [11, 12, 93, 95, 109, 111, 122].

Secret key rates. An interesting notion in Maurer's model is that of a *secret key rate*, which generalises Csiszár and Körner's earlier notion of *secrecy capacity* to a two way communication setting. The secret key rate $S(X;Y||Z)$ of the joint probability distribution P_{XYZ} is defined as the maximum rate R as follows [116]: For every $\epsilon > 0$ there exists N_0 such that for all $N \geq N_0$ Alice and Bob can agree on k-bit strings S and S' based on N independent realisations of X, Y (X^N denotes the block of the first N realizations of X) with the properties:

$$k > (R - \epsilon)N$$

$$H(S|X^N U) = 0$$
$$H(S'|Y^N U) = 0$$
$$P(S \neq S') < \epsilon$$
$$I(S; Z^N U) < \epsilon$$
$$H(S) > k - \epsilon.$$

Herein U represents the entire communication held over the public channel. Informally speaking, $S(X;Y||Z)$ is the maximum rate at which Alice and Bob can generate an almost uniformly distributed string about which Eve has virtually no information and which hence can be used securely as a cryptographic key. Intuitive bounds (as shown by Maurer [77]) on the secret key rate are given by

$$\max\{I(X;Y) - I(X;Z), I(Y;X) - I(Y;Z)\} < S(X;Y||Z),$$

$$S(X;Y||Z) < \min\{I(X;Y), I(X;Y|Z)\}.$$

Further theoretical results on secret key rates are given in [1,74,75,77,92,115] including some interesting open problems.

Limitations of noisy channel models. The models of Wyner and Csiszár and Körner clearly place unrealistic assumptions on the quality of an adversary's channel. An objection to Maurer's model, and all noisy channel models, is that in practice one does not know the error probability ϵ of Eve's noisy channel. (Indeed, Eve may be listening via several different channels.) Even if one fixes a bound for ϵ, one must also agree on an acceptable amount of information that an adversary can learn. Also, the complexities of error correction and privacy amplification make it difficult to judge practical secret key rates. The number of messages exchanged must also be considered. All of this makes it difficult to set a concrete security level at which to confidently engage in key agreement (despite experimental work [62]). Moreover it is not clear how to physically realise noisy channels with the required properties. While noisy channel models seem reasonable for wireless settings, there is not much hope for networked settings such as the Internet, as one cannot easily build a noisy channel on top of the TCP/IP stack.

We remark that the case of active adversaries (i.e. unauthenticated channels) has been extensively studied in [71,73,93,94,121,122]. However, even in this scenario it is generally assumed that Eve is passive for the initial noisy broadcast, which is again unrealistic.

3.2 Quantum key distribution

Quantum key distribution offers a quite different method of key agreement, yet has many features in common with noisy channel models.

In quantum computing, information is transmitted at the quantum level, rather than the classical level. The analog of the classical bit is the quantum bit, or *qubit*. The qubit may be measured as being in one of two states (0 or 1 say), but until measurement exists in a combination, or *superposition*, of the two states. A qubit can be physically realised as a polarised photon. A fibre optic cable may thus act as a quantum channel, and measurements taken with a beamsplitter. Peculiarities at the

quantum level are such that measurement influences the state of a qubit, and in general it is impossible to copy a qubit. This fact, the so-called *no-cloning theorem* [119], may be utilised by Alice and Bob to generate a secret key, since in theory eavesdropping is detectable over a quantum channel.

The BB84 Protocol. Bennett and Brassard [10] were the first to propose a quantum key agreement protocol, named BB84. (Although Wiesner [114] was the first to recognise the significance of quantum theory for cryptography.) The BB84 Protocol involves two public channels, one quantum and one classical, and essentially works as follows.

First Alice and Bob agree on two different bases. Each of the used bases is a pair of orthogonal quantum states, and the states of one basis are conjugate and non-orthogonal to the states of the other basis. Preparation and measurement of a qubit is performed with respect to one of these bases. In each basis, one of the states is used to encode 0 and the other state to encode 1. To agree on a secret key, Alice generates a random bitstring and for each bit randomly chooses one of the two bases to encode the bit. She then prepares each qubit depending on the respective bit value and basis and sends it to Bob over the quantum channel. Bob measures each qubit in one of the two bases (which he randomly selects) and stores the pairs (bit, basis). Then Bob and Alice communicate over the classical channel and compare their choices of the basis for each pair. The pairs where the basis is different are discarded by both parties. In the end, both hold a list of bits where Alice's encoding matches Bob's measurement. Bob and Alice now reveal a random sample of the bits and determine the error rate. If no errors occurred they can directly use the remaining bits as a secure key. However, if there is noise or an eavesdropper (who disturbs the original states by his measurements) on the quantum channel this introduces errors in Bob's measurements. In that case, information reconciliation and privacy amplification (over the classical channel) are employed to agree on a highly secure key.

Since Bennett and Brassard's original paper, there has been an explosion of interest and theoretical and practical progress in quantum key distribution. We merely refer the reader to several excellent survey articles [2, 48, 97], and [16] for important early references.

Limitations of quantum key distribution. Quantum key distribution is currently expensive, requiring dedicated hardware and networks to prepare, transmit and measure photons. There have been several attacks on specific implementations, using for example bright illumination [68] and indeed eavesdropping [17, 47]. Keys can nevertheless be exchanged [49] but at the cost of additional privacy amplification and reduction of the secret key rate. See [47, 89] for further references on quantum hacking.

Practicality of quantum key distribution is currently limited by low distances and key rates, although steady progress is being made. Recent experiments [33, 56, 66, 98, 110] achieve distances of 140 to 250 kilometres with key rates of several bits per second. Other experiments reach significantly higher key rates of 1 Mbits/s at a distance of 20 kilometres and 24 kbits/s at a distance of 100 kilometres [31, 78, 81]. On the theoretical side, studies have shown that over free space inter-satellite and satellite-to-ground stations, quantum communication links are possible due to lower atmospheric density and hence lower link attenuation [4, 87]. Thus large distances might be bridged in the future.

Another limitation of quantum key distribution is the inherent point to point infrastructure. Two paradigms have emerged to extend capabilities to networks and achieve scalability and robustness against denial of service attacks and technical breakdowns. The first is the quantum channel switching paradigm realised by the DARPA Quantum Network [39] and the Tokyo QKD Network [96]. An end-to-end quantum channel is established on demand between two participants by optical switching at the network nodes. However, this does not allow to extend the distance for key transmission and it requires a secret pre-shared key between any two participants. The second paradigm is that of trusted repeaters, which is followed by the SECOQC implementation [86]. In this setting independent quantum channels exist between network nodes. A key is then routed via intermediary nodes from the sender to the receiver. Based on this topology any distance can be bridged, but intermediary nodes learn the key and thus must be trustworthy in order to guarantee security.

3.3 The Bounded Storage Model

The Bounded Storage Model proposed by Maurer [76] provides an alternative method for key agreement. In this model, Alice and Bob are assumed to share a short secret key k which they expand into a much longer key x suitable for use, say, with the One Time Pad. An adversary Eve is assumed to have a limited storage capacity, but is otherwise computationally unbounded; in particular Eve has unlimited memory available to her during any function evaluation.

The model is conceptually simple, and works as follows. Suppose Eve's storage capacity is s, and suppose there exists a very large public source of randomness R satisfying $|R| > s$. The source of R may be, for example, a satellite broadcast. Alice and Bob use k to agree on, and access, a small portion r of R. They then apply a known *randomiser*, or *key expansion function* f to create the key $x = f(r, k)$, where now $|x| \gg |k|$. Eve meanwhile is allowed to read the entirety of R, compute an arbitrary function h depending on R, and store the output $h(R)$. The only restriction on h is that its output satisfies $|h(R)| \leq s$. After Eve stores $h(R)$, it is assumed that the signal R is lost forevermore. For security, the function f must satisfy the property that with high probability the probability distribution of x conditioned on $h(R)$ and k is close to uniform. (An explicit choice of f is given in [76].)

Significant improvements have been made to the Bounded Storage Model since its inception. In Maurer's original work [76], Eve was not assumed to have such a powerful function h; she could merely store an arbitrary s *actual* bits of R. The relaxation to the function h was made by Cachin and Maurer [21] using privacy amplification techniques, but came with considerable storage costs to Alice and Bob as well as non-negligible probability that Eve learns a non-negligible amount of information about x. These problems were removed by Aumann and Rabin [6], but with the undesirable requirement that $|R| \gg s$. Dziembowski and Maurer [34, 36] then showed that secure key agreement is possible for $s/|R|$ arbitrarily close to 1. Everlasting security is proved in [5, 29], meaning that after the signal R is lost, Eve learns no new information about x from learning k, even if she now has an unbounded amount of storage.

There is a small mountain of literature devoted to the Bounded Storage Model. In several works [5,30,36,67], reuse of the initial key k with new randomisers is considered and shown to be secure. Lu [67] and Vadhan [112] reduced the initially required size of k and used randomness extractors for the construction of their schemes. In [21, 35,

37] it is shown that information theoretic key agreement is possible without a pre-shared secret key k, but at the impractical expense of enormous storage requirements $\Omega(\sqrt{|R|})$ for honest parties. The so-called *Hybrid Bounded Storage Model* has been considered in [5, 29]. This model suggests to use a computationally secure initial key k. Dziembowski and Maurer [35] showed that this approach is in general not secure, but there may exist natural initial key agreement protocols providing everlasting security. The formalisation and security proof of such a scheme is still open. Harnik and Naor [54] showed that black box proofs for everlasting security in the hybrid model cannot exist. Ding [28] and Dodis and Smith [32] consider the practical problem of transmission errors, and the case of quantum adversaries has been considered in [25, 61].

Limitations of the Bounded Storage Model. In spite of all this theoretical work, there still remain two fundamental questions regarding practicality:

- how is the source of randomness R physically realised?
- to what extent is storage a limiting factor?

To realise R, it has been suggested to use a broadcast satellite or signal of a deep space radio source at an estimated rate of 100 gigabit/sec [21, 29, 30]. Besides the expense of this approach, such transmission rates are beyond current capabilities. Even with this transmission rate, it takes considerable time to listen to the signal. Given the example from Dziembrowski and Maurer [36] and assuming Eve's storage capacity to be at most one petabit ($s = 10^{15}$), Alice and Bob must listen to R for a day and a half to reach the required randomiser size of 12.5 petabits.

The second issue is of an even more fundamental character. Unless R exceeds all current storage facilities combined, a storage capacity is really a financial cap on the adversary's budget. And storage is cheap. Amazon Simple Storage Service allows to store 1GB at a monthly cost of \$0.037 in its cheapest variant [3]. In the above example this would imply a total cost of about \$60,000 to store R entirely.

3.4 The Limited Access Model

The Limited Access Model proposed by Rabin [90] and implemented in [45, 59] is another recent conceptual proposal for information theoretic key agreement. The model requires a large network of public servers called Page Server Nodes (PSNs). Each PSN generates random, fixed length bitstrings and stores each string in a *page*. Each page has the property that after it is accessed twice its content is overwritten by the PSN with a fresh random string. As in the Bounded Storage Model, the Limited Access Model supposes Alice and Bob share a short secret key k which they expand into a much longer key x,

The model essentially works as follows. A portion of the key k is used to agree on a selection of pages from a selection of PSNs. Alice and Bob access the relevant pages and download the associated bitstrings. The bitstrings are XORed together to form x. Security is based on the assumptions that:

- an adversary cannot monitor all PSNs.
- an adversary cannot monitor all strings downloaded by any one user.

Thus if Alice and Bob download n of N total pages and Eve monitors $u < N$ random pages, then the probability that Eve is able to derive x is at most $(u/N)^n$. Under

these assumptions, everlasting security is ensured. For suppose the initial key k is later compromised. Then an adversary Eve learns nothing of x since the constituent bitstrings have all been overwritten by the PSN, having been accessed twice, by Alice and Bob. And under the two assumptions, Eve has not stored all the strings downloaded by either Alice or Bob, and thus she gains no extra information about x from discovering k.

But what if Eve becomes active and accesses a relevant page of a PSN after Alice yet before Bob, say? When Bob accesses the page, he will download a different string from Alice and a common key x will not be established. This is where the remaining part of the key k comes into play.

After downloading, Alice and Bob perform a so-called *page reconciliation protocol*, to identify and discard any bitstrings that do not match. Encrypting under the remaining part of k, Alice sends to Bob the page locations and hashes (or MACs, as in [45]) of the associated bitstrings. Bob compares hash (or MAC) values and sends back information on any discrepancies found. The key x is formed from XORing matching bitstrings, part of which is reserved to update the initial key k to generate further key material.

Limitations of the Limited Access Model. There are clear limitations of the Limited Access Model. The first assumption may be seen as a distributed bounded storage assumption. Whilst it is a nice idea that the task of generating randomness is distributed across a network, one may question to what extent monitoring (distributed) storage locations acts as a limiting factor on an adversary. The second assumption is certainly not standard (the exact opposite is usually assumed), but clearly needed for security. Rabin also suggests that Alice evades monitoring by visiting Internet cafes or friends and download pages anonymously. While this may work on a small scale, it is clearly not practical at any interesting scale. Moreover, there is nothing to prevent an adversary monitoring Alice's local Internet cafes.

As in the Bounded Storage Model, an initial key k is required. The authors of [46, 59] suggest to use a common computationally secure protocol. But in this case the page reconciliation protocol may affect everlasting security: suppose an all-powerful Eve computes k and then finds all preimages of the hash (or MAC) values. It may then be possible to mount a better-than-brute-force attack on the key x. Thus security depends on the particulars of the hash or MAC function.

A less severe practical issue concerns the necessarily large number of page server nodes. Rabin presents a search engine based construction not requiring PSNs. Using k as search engine input, common web pages are randomly selected as sources of randomness. However, it is an open question as to how much randomness such a system provides. Furthermore the web pages are not destroyed after downloading twice, hence everlasting security might be at risk.

4 Long term confidential data storage by proactive secret sharing

For long term confidentiality of stored data, different solutions are needed than those for data transmission. This is because encryption with the One Time Pad is not suitable since the problem of secure long term storage of data is replaced by secure long term storage of the key.

The challenge of long term confidentiality of stored data is met by *proactive perfect secret sharing*. *Proactive security* means to periodically update or reestablish the security of a system, even if no threats to the system are found. The rationale behind proactive security is that intrusion detection is very difficult due to slow attacks as considered by Storer et al. [103]. Because of the long lifetime of the data, an attacker has a large time window within which to mount an attack. Occasionally entering a system making only small changes at any given time may go undetected. Thus over long periods of time it seems reasonable to assume that intrusion will occur at some point.

Perfect secret sharing was introduced independently by Shamir [99] and Blakley [14] in 1979. This technique allows for information theoretically secure storage of data in a distributed environment; a secret is divided into shares that are distributed to different participants (e.g. storage servers) in such a way that any nonqualified subset of participants learns no information about the secret. Hence, information theoretic security is preserved as long as an adversary does not compromise a qualified subset of participants.

To protect against this kind of attack, Herzberg et al. [55] introduced the process of *share renewal* to secret sharing schemes. This proactive measure allows for shares (compromised or otherwise) to be updated at periodic intervals, so that in order to learn the secret an adversary must now compromise a qualified subset of participants within a given time period. Thus in theory an adversary can compromise every participant over different time periods and still learn no information about the secret. This so-called *mobile adversary model* has been widely adopted for long term storage security [22, 52, 53, 55, 117, 118].

Security of (proactive) secret sharing schemes relies on several assumptions. First, it is assumed that an adversary does not eavesdrop on communications during share distribution and share renewal. Thus shares must be distributed using secure channels and long term secure data transmission (as considered in Section 3) is a preliminary to long term security of storage. Second, it is clearly necessary to assume that uncompromised participants securely erase old shares. Third, it must be possible to end an adversary's access to a system by e.g. a reboot, thus a system is not compromised once and for all. Furthermore, in case an adversary has access to a system during the share renewal phase, the renewed shares of the considered system are still compromised, following Herzberg et al. [55] and approved by Nikov and Nikova [83].

Note that an active adversary can easily destroy a secret by providing inconsistent share updates during share renewal. In order to provide security against such active adversaries, the technique is extended by applying verifiable secret sharing protocols [22, 55]. In summary, in proactive verifiable secret sharing schemes, the lifetime of a system is divided into fixed time periods, each ending with an update phase. The update phase consists of the steps *share renewal*, *detection of corrupted shares* and *share recovery*. Verifiability is only required to protect the secret from destruction, and not mandatory for confidentiality.

4.1 Functionality

In the general model of secret sharing schemes, there are n participants (or shareholders) $P_1, \ldots, P_n$, and a special entity called the Dealer. The Dealer distributes the secret to the n participants in such a way that only *qualified subsets* can jointly reconstruct the secret. The set of all qualified subsets Γ is called the *access structure*

of the scheme. Thus a subset P of participants is able to jointly reconstruct the secret if and only if $P \in \Gamma$. A secret sharing scheme whereby any k out of n participants can compute the secret is called a (k, n)-*threshold scheme*. To achieve information theoretic security for any secret sharing scheme, each share must be of length at least as the secret itself [63]. Schemes that achieve the lower bound on share sizes are called *optimal*.

In the following we describe Shamir's secret sharing scheme [99] as a basis for proactive systems. It is the best known perfect secret sharing scheme and considered in a great variety of works, including practical applications [20,55,58,105]. For a more general view on secret sharing schemes, see the recent survey by Beimel [7].

Shamir's secret sharing scheme. Shamir's secret sharing scheme is a perfect (k, n)-threshold scheme that makes use of polynomials over a finite field $\mathbb{F}$. It is optimal in the share size, and $n < |\mathbb{F}|$ is required.

To share a secret $s \in \mathbb{F}$, $k-1$ secret coefficients $a_i \in \mathbb{F},\ i = 1, 2, \ldots, k-1$ are chosen uniformly at random by the Dealer. The secret s forms the constant term of the polynomial

$$f(x) = s + \sum_{i=1}^{k-1} a_i x^i.$$

For n mutually different $x_j \in \mathbb{F},\ j = 1, 2, \ldots, n$, the evaluations of the polynomial $y_j = f(x_j)$ form the shares. From each subset of k shares and the corresponding x_j the secret can unambiguously be reconstructed by interpolation using Lagrange's formula

$$s = f(0) = \sum_{i=1}^{k} y_i \prod_{l=1, l \neq i}^{k} \frac{x_l}{x_l - x_i}.$$

The application of Shamir's scheme to arbitrarily large secrets requires arbitrarily large fields, leading to inefficiencies. Miyamoto et al. [79] provide an approach avoiding large fields and admitting an efficient implementation. The main idea is to first split the secret into blocks of several bits say, such that each block can be seen as an element of a given, relatively small field $\mathbb{F}$. An example would be to split the secret into bytes and use $\mathbb{F} = GF(2^8)$. Then each block is shared individually using Shamir's scheme while reusing the x_j. After sharing all blocks, all shares belonging to the same x_j are respectively assembled into a vector of shares.

Share renewal. There are essentially three approaches for share renewal with Shamir's scheme. The first and most general way to renew the shares works with any secret sharing scheme. First the share is reconstructed by a central instance, e.g. the Dealer. New shares are then created and distributed to the participants; old shares are securely deleted. However, this method exposes the secret to the central instance [53]. Hence during share renewal, intrusion into a single system suffices to compromise the secret.

One method to renew shares without reconstruction uses the fact that Shamir's scheme is additively homomorphic. To renew the shares, one shares the zero element of F, and provides the shares to the participants. They add these shares locally to their prior shares of the secret. This method is applied by Herzberg at al. [55], where each of the participants independently share the zero element and distribute the shares to the other participants.

While this approach keeps the current access structure, a third approach proposed by Frankel at al. [42] and Desmedt and Jojoda [27] uses homomorphic properties of exponentiation, and optionally allows change of the access structure during share renewal. As described by Gupta and Gopinath [53] for Shamir's scheme, the participants share their initial shares according to the intended access structure and distribute the obtained temporary shares to the other participants. Given a (k, n)-threshold secret sharing scheme, a participant can then compute its new share from k of these temporary shares.

Verifiability. The problem with share renewal without reconstruction is that an active adversary may destroy the secret such that it cannot be reconstructed anymore. To provide security against destruction, these schemes are combined with share verification schemes obtaining verifiable proactive secret sharing. Most of these schemes use zero knowledge proofs [102] or commitment schemes, where some information is broadcasted. Feldman [41] proposed a conditionally secure verification scheme based on the intractability of solving discrete logarithms. However, the confidentiality of the secret is only computationally secure and hence not applicable in the context of long term confidentiality.

Pederson [85] provides an information theoretically secure verification scheme similar to Feldman. He achieves information theoretic security by involving a random value into his commitment scheme to mask the actual value of the shares. However, this approach doubles the share size and security against an alteration of the secret is still computational.

The scheme of Herzberg et al. [55] works with either Feldman's or Pederson's approach.

Wong et al. [117] combine Desmedt and Jojoda's share renewal protocol with Feldman's verification scheme, which was extended by Gupta and Gopinath [52] to be robust against participants that behave maliciously when receiving share updates. To achieve information theoretic security Gupta and Gopinath [53] apply Pedersons's verification scheme to their approach.

Stinson and Wei [102] propose a completely unconditionally secure proactive verifiable secret sharing scheme, which is improved in [26]. In [83] an attack utilising the broadcast information against the schemes from [26, 102] is presented. Yet, they consider a changed adversarial model. This attack is generalised to general access structures in [84].

4.2 Practicality and limitations

Information theoretically secure secret sharing schemes suffer from large computational overheads during share distribution and renewal. This limits their practical applicability as discussed by Subbiah and Blough [108]. Verifiability to resist malicious dealers makes the schemes even more complex and might threaten confidentiality by broadcasting verification messages. Considering the total amount of data to be stored, given the secret s, an overall amount of data of at least $k \cdot |s|$ must be stored to achieve security against a passive adversary that can compromise $(k-1)$ participants. If we consider active adversaries that aim to destroy the secret, this amount is further increased: in the Herzberg et al. [55] mobile adversary model, for a (k, n)-threshold secret sharing scheme, in order to guarantee security even if up to $k-1$ servers are dishonest, one requires $2(k-1) < n$. This results in a data storage of $n \cdot |s|$.

Proactivity results in a lot of data traffic during share renewal. Each participant sends and receives an amount of data at least $(n-1)\cdot|s|$, thus periodically producing a massive network load. This is accompanied by the computational overhead of sharing and reconstruction.

However, the main hurdle to overcome is that secure channels between participants must be assumed since in general, confidentiality is lost when an adversary is able to eavesdrop during share distribution. Because of these problems, proactive, or even verifiable proactive secret sharing schemes are rarely used in practice. However, secret sharing is applied in several archival systems such as POTSHARDS [104, 105], Pasis [44], GridSharing [108], and the approaches of Hühnlein et al. [58] and Masinter and Welch [69].

But none of the above systems apply proactive secret sharing. While this is mentioned to be investigated further for POTSHARDS, Masinter and Welch apply periodic availability and integrity checks, and only apply reconstruction if defects are recognised. We refer to the survey on existing archival systems by Storer et al. [103] for further reading.

Schemes proposed to reduce the overhead come with new disadvantages. Krawczyk [63] proposed a secret sharing scheme with short shares using a combination of encryption, information dispersal and perfect secret sharing. However, the result is a computationally secure scheme. Subbiah [107] proposed a scheme that uses an additional system secret, which is shared among the participants and is claimed to provide proactive security against passive adversaries. Shares of the system secret are periodically updated and used to refresh the shares of all the stored secrets. However, the scheme has a flaw: the use of one and the same set of shares for share renewal of all stored secrets is comparable to One Time Pad encryption always reusing one key, which is clearly insecure.

5 Summary

We have reviewed known information theoretic methods for key agreement and key management relevant for long term confidentiality. For key agreement, noisy channel models are problematic as it is hard to see how weakening the assumption that an adversary has complete and perfect access to the channel can lead to concrete security guarantees. Quantum key distribution is a promising approach, but currently has engineering and scalability issues. The Bounded Storage Model is impractical to physically realise, and the Limited Access Model needs to overcome simple eavesdropping attacks. Moreover, each model requires an enormous amount of randomness: achievable bit rates of physical random number generators range from several kbit/s up to 400 Mbits/s [43, 64], and recent experiments even achieve 300 Gbit/s [60], but are rather costly. Concerning long term storage of confidential data, proactive secret sharing presents a solution. However, information theoretic key agreement is a prerequisite.

Despite a mountain of theoretical work, a practical solution to long term confidentiality remains an elusive goal.

References

1. H. Ahmadi and R. Safavi-Naini. Secret keys from channel noise. In *Proceedings of the 30th Annual international conference on Theory and applications of cryptographic tech-*

niques: advances in cryptology, EUROCRYPT'11, pages 266–283, Berlin, Heidelberg, 2011. Springer-Verlag.

2. R. Alléaume, N. Lütkenhaus, R. Renner, P. Grangier, T. Debuisschert, G. Ribordy, N. Gisin, P. Painchault, T. Pornin, L. Slavail, M. Riguidel, A. Shilds, T. Länger, M. Peev, M. Dianati, A. Leverrier, A. Poppe, J. Bouda, C. Branciard, M. Godfrey, J. Rarity, H. Weinfurter, A. Zeilinger, and C. Monyk. Quantum key distribution and cryptography: a survey. In S. L. Braunstein, H.-K. Lo, K. Paterson, and P. Ryan, editors, *Classical and Quantum Information Assurance Foundations and Practice*, number 09311 in Dagstuhl Seminar Proceedings, Dagstuhl, Germany, 2010. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany.
3. Amazon web services. Amazon simple storage service (amazon s3), 2012. `http://aws.amazon.com/s3`.
4. J. M. P. Armengol, B. Furch, C. J. de Matos, O. Minster, L. Cacciapuoti, M. Pfennigbauer, M. Aspelmeyer, T. Jennewein, R. Ursin, T. Schmitt-Manderbach, G. Baister, J. Rarity, W. Leeb, C. Barbieri, H. Weinfurter, and A. Zeilinger. Quantum communications at esa: Towards a space experiment on the iss. *Acta Astronautica*, 63(1-4):165 – 178, 2008.
5. Y. Aumann, Y. Z. Ding, and M. O. Rabin. Everlasting security in the bounded storage model. *IEEE Transactions on Information Theory*, 48(6):1668 –1680, 2002.
6. Y. Aumann and M. O. Rabin. Information theoretically secure communication in the limited storage space model. In *Proceedings of the 19th Annual International Cryptology Conference on Advances in Cryptology*, CRYPTO '99, pages 65–79, London, UK, 1999. Springer-Verlag.
7. A. Beimel. Secret-sharing schemes: a survey. In *Proceedings of the Third international conference on Coding and cryptology*, IWCC'11, pages 11–46, Berlin, Heidelberg, 2011. Springer-Verlag.
8. M. Bellare and S. Tessaro. Polynomial-time, semantically-secure encryption achieving the secrecy capacity. Cryptology ePrint Archive, Report 2012/022, 2012. `http://eprint.iacr.org/`.
9. M. Bellare, S. Tessaro, and A. Vardy. A cryptographic treatment of the wiretap channel. Cryptology ePrint Archive, Report 2012/015, 2012. `http://eprint.iacr.org/`.
10. C. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. In IEEE, editor, *Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing*, pages 175–179, 1984.
11. C. Bennett, G. Brassard, C. Crépeau, and U. M. Maurer. Generalized privacy amplification. *IEEE Transactions on Information Theory*, 41:1915–1923, 1995.
12. C. H. Bennett, G. Brassard, and J.-M. Robert. Privacy amplification by public discussion. *SIAM J. Comput.*, 17:210–229, April 1988.
13. D. J. Bernstein. Introduction to post-quantum cryptography. In D. J. Bernstein, J. Buchmann, and E. Dahmen, editors, *Post-Quantum Cryptography*, pages 1–14. Springer Berlin Heidelberg, 2009.
14. G. R. Blakley. Safeguarding cryptographic keys. In *Proceedings of the 1979 AFIPS National Computer Conference*, pages 313–317, Monval, NJ, USA, 1979. AFIPS Press.
15. BlueKrypt. Cryptographic key length recommendation. `http://www.keylength.com`.
16. G. Brassard. A bibliography of quantum cryptography. *Journal of Modern Optics*, December 1993. http://www.cs.mcgill.ca/ crepeau/CRYPTO/Biblio-QC.html.
17. G. Brassard, N. Lütkenhaus, T. Mor, and B. C. Sanders. Limitations on practical quantum cryptography. *Physical Review Letters*, 85(6):1330–1333, Aug. 2000.
18. G. Brassard and L. Salvail. Secret-key reconciliation by public discussion. In T. Helleseth, editor, *Advances in Cryptology – EUROCRYPT 93*, volume 765 of *Lecture Notes in Computer Science*, pages 410–423. Springer Berlin / Heidelberg, 1994. 10.1007/3-540-48285-7_35.
19. J. Buchmann, A. May, and U. Vollmer. Perspectives for cryptographic long-term security. *Commun. ACM*, 49:50–55, September 2006.

20. C. Cachin, R. Haas, and M.Vukolić. Dependable storage in the intercloud. Technical report, IBM Research, October 2010. RZ 3783.
21. C. Cachin and U. M. Maurer. Unconditional security against memory-bounded adversaries. In *Proceedings of the 17th Annual International Cryptology Conference on Advances in Cryptology*, pages 292–306, London, UK, 1997. Springer-Verlag.
22. R. Canetti, R. Gennaro, A. Herzberg, and D. Naor. Proactive security: Long-term protection against break-ins. *CryptoBytes*, 3:1–8, 1997.
23. T. M. Cover and J. A. Thomas. *Elements of information theory (2. ed.)*. Wiley, 2006.
24. I. Csiszar and J. Korner. Broadcast channels with confidential messages. *Information Theory, IEEE Transactions on*, 24(3):339 – 348, may 1978.
25. I. B. Damgård, S. Fehr, L. Salvail, and C. Schaffner. Cryptography in the bounded quantum-storage model. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '05, pages 449–458, Washington, DC, USA, 2005. IEEE Computer Society.
26. P. D'Arco and D. R. Stinson. On unconditionally secure robust distributed key distribution centers. In *Proceedings of the 8th International Conference on the Theory and Application of Cryptology and Information Security: Advances in Cryptology*, ASIACRYPT '02, pages 346–363, London, UK, UK, 2002. Springer-Verlag.
27. Y. Desmedt and S. Jajodia. Redistributing secret shares to new access structures and its applications. Technical report, George Mason University, July 1997. ISSE TR-97-1.
28. Y. Ding. Error correction in the bounded storage model. In J. Kilian, editor, *Theory of Cryptography*, volume 3378 of *Lecture Notes in Computer Science*, pages 578–599. Springer Berlin / Heidelberg, 2005.
29. Y. Ding and M. Rabin. Hyper-encryption and everlasting security. In H. Alt and A. Ferreira, editors, *STACS 2002*, volume 2285 of *Lecture Notes in Computer Science*, pages 731–731. Springer Berlin / Heidelberg, 2002. 10.1007/3-540-45841-7_1.
30. Y. Z. Ding. *Provable everlasting security in the bounded storage model*. PhD thesis, Harvard University, Cambridge, MA, USA, 2001. AAI3011357.
31. A. R. Dixon, Z. L. Yuan, J. F. Dynes, A. W. Sharpe, and A. J. Shields. Gigahertz decoy quantum key distribution with 1 mbit/s secure key rate. *Opt. Express*, 16(23):18790–18979, Nov 2008.
32. Y. Dodis and A. Smith. Correcting errors without leaking partial information. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, STOC '05, pages 654–663, New York, NY, USA, 2005. ACM.
33. J. F. Dynes, H. Takesue, Z. L. Yuan, A. W. Sharpe, K. Harada, T. Honjo, H. Kamada, O. Tadanaga, Y. Nishida, M. Asobe, and A. J. Shields. Efficient entanglement distribution over 200 kilometers. *Opt. Express*, 17(14):11440–11449, Jul 2009.
34. S. Dziembowski and U. Maurer. Tight security proofs for the bounded-storage model. In *Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*, STOC '02, pages 341–350, New York, NY, USA, 2002. ACM.
35. S. Dziembowski and U. Maurer. On generating the initial key in the bounded-storage model. In C. Cachin and J. Camenisch, editors, *Advances in Cryptology - EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 126–137. Springer Berlin / Heidelberg, 2004. 10.1007/978-3-540-24676-3_8.
36. S. Dziembowski and U. Maurer. Optimal randomizer efficiency in the bounded-storage model. *Journal of Cryptology*, 17:5–26, 2004. 10.1007/s00145-003-0309-y.
37. S. Dziembowski and U. Maurer. The bare bounded-storage model: The tight bound on the storage requirement for key agreement. *Information Theory, IEEE Transactions on*, 54(6):2790 –2792, june 2008.
38. Electronic Frontier Foundation. *Cracking DES - Secrets of Encryption Research, Wiretap Politics & Chip Design*. O'Reilly Media, July 1998.
39. C. Elliott, A. Colvin, D. Pearson, O. Pikalo, J. Schlafer, and H. Yeh. Current status of the darpa quantum network, 2005.

40. European Telecommunications Standards Institute (ETSI). Electronic Signatures and Infrastructures (ESI) – Electronic Signatures and Infrastructures (ESI); CMS Advanced Electronic Signatures (CAdES). ETSI Technical Specification TS 101 733, Version 1.7.4, Juli 2008. `http://www.etsi.org/`.
41. P. Feldman. A practical scheme for non-interactive verifiable secret sharing. In *Proceedings of the 28th Annual Symposium on Foundations of Computer Science*, SFCS '87, pages 427–438, Washington, DC, USA, 1987. IEEE Computer Society.
42. Y. Frankel, P. Gemmell, P. D. MacKenzie, and M. Yung. Optimal-resilience proactive public-key cryptosystems. In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science*, pages 384–393, Washington, DC, USA, 1997. IEEE Computer Society.
43. M. Fürst, H. Weier, S. Nauerth, D. G. Marangon, C. Kurtsiefer, and H. Weinfurter. High speed optical quantum random number generation. *Opt. Express*, 18(12):13029–13037, Jun 2010.
44. G. R. Ganger, P. K. Khosla, M. Bakkaloglu, M. W. Bigrigg, R. Garth, S. Oguz, P. Vijay, C. A. N. Soules, J. D. Strunk, and J. J. Wylie. Survivable storage systems. In *In DARPA Information Survivability Conference and Exposition, IEEE*, volume 2, pages 184–195, 2001.
45. R. E. H. García. The analysis and implementation of a practical crypto-system in the limited access model. Master's thesis, Centro de Investigación y de Estudios Avanzados del Instituto Politécnico Nacional, Departamento de Computación, 2010.
46. R. E. H. García, I. Cabrera, and D. Chakraborty. On implementation of a practical crypto-system in the limited access model. In *Electrical Engineering Computing Science and Automatic Control (CCE), 2010 7th International Conference on*, pages 418–423, sept. 2010.
47. I. Gerhardt, Q. Liu, A. Lamas-Linares, J. Skaar, C. Kurtsiefer, and V. Makarov. Full-field implementation of a perfect eavesdropper on a quantum cryptography system. *Nature Communications*, 2:349+, June 2011.
48. N. Gisin, G. Ribordy, W. Tittel, and H. Zbinden. Quantum cryptography. *Rev. Mod. Phys*, 74:145–195, 2002.
49. D. Gottesman, H.-K. Lo, N. Lütkenhaus, and J. Preskill. Security of quantum key distribution with imperfect devices. *Quantum Info. Comput.*, 4:325–360, September 2004.
50. G. F. Government. §203 StGB Violation of private secrets (Verletzung von Privatgeheimnissen), 2012. German Criminal Code (Strafgestzbuch StGB).
51. L. K. Grover. Quantum mechanics helps in searching for a needle in a haystack. *PHYS.REV.LETT.*, 79:325, 1997.
52. V. Gupta and K. Gopinath. An extended verifiable secret redistribution protocol for archival systems. In *Availability, Reliability and Security, 2006. ARES 2006. The First International Conference on*, page 8 pp., april 2006.
53. V. H. Gupta and K. Gopinath. g^2_{its} vsr: An information theoretical secure verifiable secret redistribution protocol for long-term archival storage. *Security in Storage Workshop, International IEEE*, 0:22–33, 2007.
54. D. Harnik and M. Naor. On everlasting security in the *Hybrid* bounded storage model. In M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, editors, *Automata, Languages and Programming*, volume 4052 of *Lecture Notes in Computer Science*, pages 192–203. Springer Berlin / Heidelberg, 2006. 10.1007/11787006_17.
55. A. Herzberg, S. Jarecki, H. Krawczyk, and M. Yung. Proactive secret sharing or: How to cope with perpetual leakage. In *Lecture Notes in Computer Science*, pages 339–352. Springer-Verlag, 1995.
56. P. A. Hiskett, D. Rosenberg, C. G. Peterson, R. J. Hughes, S. Nam, A. E. Lita, A. J. Miller, and J. E. Nordholt. Long-distance quantum key distribution in optical fibre. *New Journal of Physics*, 8(9):193, 2006.

57. J. Hughes and J. N. Roge. Long-term security vulnerabilities of encrypted data. *Issues in Information Systems*, 8:522–528, 2007.
58. D. Hühnlein, U. Korte, L. Langer, and A. Wiesmaier. A comprehensive reference architecture for trustworthy long-term archiving of sensitive data. In I. Press, editor, *3rd International Conference on New Technologies, Mobility and Security (NTMS)*, pages 1–5, Dez 2009.
59. J. K. Juang. Practical implementation and analysis of hyper-encryption. Master's thesis, Massachusetts Institute of Technology. Dept. of Electrical Engineering and Computer Science., 2009.
60. I. Kanter, Y. Aviad, I. Reidler, E. Cohen, and M. Rosenbluh. An optical ultrafast random bit generator. *Nature Photonics*, 4(1):58–61, Dec. 2009.
61. R. Konig and B. Terhal. The bounded-storage model in the presence of a quantum adversary. *Information Theory, IEEE Transactions on*, 54(2):749 –762, feb. 2008.
62. V. Korzhik, V. Yakovlev, and A. Sinuk. Achieveability of the key-capacity in a scenario of key sharing by public discussion and in the presence of passive eavesdropper. In V. Gorodetsky, L. Popyack, and V. Skormin, editors, *Computer Network Security*, volume 2776 of *Lecture Notes in Computer Science*, pages 308–315. Springer Berlin / Heidelberg, 2003.
63. H. Krawczyk. Secret sharing made short. In *Proceedings of the 13th annual international cryptology conference on Advances in cryptology*, pages 136–146, New York, NY, USA, 1994. Springer-Verlag New York, Inc.
64. LE Tech Co.,Ltd. Genuine random number generator, 2012.
65. A. K. Lenstra and E. R. Verheul. Selecting cryptographic key sizes. *Journal of Cryptology*, 14:255–293, 2001. 10.1007/s00145-001-0009-4.
66. Y. Liu, T.-Y. Chen, J. Wang, W.-Q. Cai, X. Wan, L.-K. Chen, J.-H. Wang, S.-B. Liu, H. Liang, L. Yang, C.-Z. Peng, K. Chen, Z.-B. Chen, and J.-W. Pan. Decoy-state quantum key distribution with polarized photons over 200 km. *Opt. Express*, 18(8):8587–8594, Apr 2010.
67. C.-J. Lu. Encryption against storage-bounded adversaries from on-line strong extractors. *Journal of Cryptology*, 17:27–42, 2004. 10.1007/s00145-003-0217-1.
68. L. Lydersen, C. Wiechers, C. Wittmann, D. Elser, J. Skaar, and V. Makarov. Hacking commercial quantum cryptography systems by tailored bright illumination. *Nature Photonics*, 4(10):686–689, Aug. 2010.
69. L. Masinter and M. Welch. A system for long-term document preservation. In *IS&T Archiving 2006*, volume 3, pages 61–68, Ottawa, Canada: Society For Imaging Science And Technology, 2006.
70. U. Maurer. Protocols for secret key agreement by public discussion based on common information. In E. Brickell, editor, *Advances in Cryptology – CRYPTO 92*, volume 740 of *Lecture Notes in Computer Science*, pages 461–470. Springer Berlin / Heidelberg, 1993. 10.1007/3-540-48071-4_32.
71. U. Maurer. Information-theoretically secure secret-key agreement by not authenticated public discussion. In *Advances in Cryptology - EUROCRYPT '97, Lecture*, pages 209–225. Springer-Verlag, 1997.
72. U. Maurer. Information-theoretic cryptography. In M. Wiener, editor, *Advances in Cryptology — CRYPTO '99*, volume 1666 of *Lecture Notes in Computer Science*, pages 47–64. Springer-Verlag, Aug. 1999.
73. U. Maurer, R. Renner, and S. Wolf. Unbreakable keys from random noise. In P. Tuyls, B. Skoric, and T. Kevenaar, editors, *Security with Noisy Data*, pages 21–44. Springer London, 2007.
74. U. Maurer and S. Wolf. Towards characterizing when information-theoretic secret key agreement is possible. In K. Kim and T. Matsumoto, editors, *Advances in Cryptology ASIACRYPT '96*, volume 1163 of *Lecture Notes in Computer Science*, pages 196–209. Springer Berlin / Heidelberg, 1996. 10.1007/BFb0034847.

75. U. Maurer and S. Wolf. Unconditionally secure key agreement and the intrinsic conditional information. *Information Theory, IEEE Transactions on*, 45(2):499 –514, mar 1999.
76. U. M. Maurer. Conditionally-perfect secrecy and a provably-secure randomized cipher. *Journal of Cryptology*, 5:53–66, 1992. 10.1007/BF00191321.
77. U. M. Maurer. Secret key agreement by public discussion from common information. *IEEE Transactions on Information Theory*, 39(3):733–742, may 1993.
78. A. Mink, X. Tang, L. Ma, T. Nakassis, B. Hershman, J. C. Bienfang, D. Su, R. Boisvert, C. W. Clark, and C. J. Williams. High speed quantum key distribution system supports one-time pad encryption of real-time video. In *Proceedings of SPIE*, volume 6244, pages 62440M–1–7. SPIE, 2006.
79. T. Miyamoto, S. Doi, H. Nogawa, and S. Kumagai. Autonomous distributed secret sharing storage system. *Syst. Comput. Japan*, 37(6):55–63, 2006.
80. J. Müller-Quade and D. Unruh. Long-term security and universal composability. *Journal of Cryptology*, 23:594–671, 2010. 10.1007/s00145-010-9068-8.
81. N. Namekata, H. Takesue, T. Honjo, Y. Tokura, and S. Inoue. High-rate quantum key distribution over 100 km using ultra-low-noise, 2-ghz sinusoidally gated ingaas/inp avalanche photodiodes. *Opt. Express*, 19(11):10632–10639, May 2011.
82. National Institute of Standards and Technology. Data encryption standard (DES). FIPS Publication 46-3, October 1999.
83. V. Nikov and S. Nikova. On proactive secret sharing schemes. In *Selected Areas in Cryptography*, pages 308–325, 2004.
84. V. Nikov, S. Nikova, B. Preneel, and J. Vandewalle. Applying general access structure to proactive secret sharing schemes. Cryptology ePrint Archive, Report 2002/141, 2002. `http://eprint.iacr.org/`.
85. T. P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In *Proceedings of the 11th Annual International Cryptology Conference on Advances in Cryptology*, CRYPTO '91, pages 129–140, London, UK, 1992. Springer-Verlag.
86. M. Peev, C. Pacher, R. Allaume, C. Barreiro, J. Bouda, W. Boxleitner, T. Debuisschert, E. Diamanti, M. Dianati, J. F. Dynes, S. Fasel, S. Fossier, M. Frst, J.-D. Gautier, O. Gay, N. Gisin, P. Grangier, A. Happe, Y. Hasani, M. Hentschel, H. Hbel, G. Humer, T. Lnger, M. Legr, R. Lieger, J. Lodewyck, T. Lornser, N. Ltkenhaus, A. Marhold, T. Matyus, O. Maurhart, L. Monat, S. Nauerth, J.-B. Page, A. Poppe, E. Querasser, G. Ribordy, S. Robyr, L. Salvail, A. W. Sharpe, A. J. Shields, D. Stucki, M. Suda, C. Tamas, T. Themel, R. T. Thew, Y. Thoma, A. Treiber, P. Trinkler, R. Tualle-Brouri, F. Vannel, N. Walenta, H. Weier, H. Weinfurter, I. Wimberger, Z. L. Yuan, H. Zbinden, and A. Zeilinger. The secoqc quantum key distribution network in vienna. *New Journal of Physics*, 11(7):075001, 2009.
87. M. Pfennigbauer, W. R. Leeb, M. Aspelmeyer, T. Jennewein, and A. Zeilinger. Free-space optical quantum key distribution using intersatellite. In *Links, Proceedings of the CNES - Intersatellite Link Workshop*, 2003.
88. D. Pinkas, J. Ross, and N. Pope. Cms advanced electronic signatures (cades). Request For Comments – RFC 5126, Februar 2008. `http://www.ietf.org/rfc/rfc5126.txt`.
89. Quantum Hacking. Papers and preprints. `http://www.iet.ntnu.no/groups/optics/qcr/publications.html`.
90. M. O. Rabin. Provably unbreakable hyper-encryption in the limited access model. In *Theory and Practice in Information-Theoretic Security, 2005. IEEE Information Theory Workshop on*, pages 34 – 37, oct. 2005.
91. T. A. Ramos, N. da Silva, L. C. Lung, J. G. Kohler, and R. F. Custódio. An infrastructure for long-term archiving of authenticated and sensitive electronic documents. In *Proceedings of the 7th European conference on Public key infrastructures, services and applications*, EuroPKI'10, pages 193–207, Berlin, Heidelberg, 2011. Springer-Verlag.
92. R. Renner and S. Wolf. New bounds in secret-key agreement: The gap between formation and secrecy extraction. In E. Biham, editor, *Advances in Cryptology EUROCRYPT*

2003, volume 2656 of *Lecture Notes in Computer Science*, pages 643–643. Springer Berlin / Heidelberg, 2003.
93. R. Renner and S. Wolf. Unconditional authenticity and privacy from an arbitrarily weak secret. In D. Boneh, editor, *Advances in Cryptology - CRYPTO 2003*, volume 2729 of *Lecture Notes in Computer Science*, pages 78–95. Springer Berlin / Heidelberg, 2003.
94. R. Renner and S. Wolf. The exact price for unconditionally secure asymmetric cryptography. In *Advances in Cryptology - EUROCRYPT 2004, International Conference on the Theory and Applications of Cryptographic Techniques, Interlaken, Switzerland, May 2-6, 2004, Proceedings*, pages 109–125, 2004.
95. R. Renner and S. Wolf. Simple and Tight Bounds for Information Reconciliation and Privacy Amplification. In B. Roy, editor, *Advances in Cryptology - ASIACRYPT 2005*, volume 3788 of *Lecture Notes in Computer Science*, chapter 11, pages 199–216. Springer Berlin / Heidelberg, Berlin, Heidelberg, 2005.
96. M. Sasaki, M. Fujiwara, H. Ishizuka, W. Klaus, K. Wakui, M. Takeoka, S. Miki, T. Yamashita, Z. Wang, A. Tanaka, K. Yoshino, Y. Nambu, S. Takahashi, A. Tajima, A. Tomita, T. Domeki, T. Hasegawa, Y. Sakai, H. Kobayashi, T. Asai, K. Shimizu, T. Tokura, T. Tsurumaru, M. Matsui, T. Honjo, K. Tamaki, H. Takesue, Y. Tokura, J. F. Dynes, A. R. Dixon, A. W. Sharpe, Z. L. Yuan, A. J. Shields, S. Uchikoga, M. Legré, S. Robyr, P. Trinkler, L. Monat, J.-B. Page, G. Ribordy, A. Poppe, A. Allacher, O. Maurhart, T. Länger, M. Peev, and A. Zeilinger. Field test of quantum key distribution in the tokyo qkd network. *Opt. Express*, 19(11):10387–10409, May 2011.
97. V. Scarani, H. Bechmann-Pasquinucci, N. J. Cerf, M. Dušek, N. Lütkenhaus, and M. Peev. The security of practical quantum key distribution. *Rev. Mod. Phys.*, 81:1301–1350, Sep 2009.
98. T. Schmitt-Manderbach, H. Weier, M. Fürst, R. Ursin, F. Tiefenbacher, T. Scheidl, J. Perdigues, Z. Sodnik, C. Kurtsiefer, J. G. Rarity, A. Zeilinger, and H. Weinfurter. Experimental demonstration of free-space decoy-state quantum key distribution over 144 km. *Phys. Rev. Lett.*, 98(1):010504, 2007.
99. A. Shamir. How to share a secret. *Commun. ACM*, 22(11):612–613, 1979.
100. C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423, July, October 1948.
101. P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26:1484–1509, October 1997.
102. D. R. Stinson and R. Wei. Unconditionally secure proactive secret sharing scheme with combinatorial structures. In *Proceedings of the 6th Annual International Workshop on Selected Areas in Cryptography*, SAC '99, pages 200–214, London, UK, 2000. Springer-Verlag.
103. M. W. Storer, K. Greenan, and E. L. Miller. Long-term threats to secure archives. In *Proceedings of the second ACM workshop on Storage security and survivability*, StorageSS '06, pages 9–16, New York, NY, USA, 2006. ACM.
104. M. W. Storer, K. M. Greenan, E. L. Miller, and K. Voruganti. Potshards: secure long-term storage without encryption. In *2007 USENIX Annual Technical Conference on Proceedings of the USENIX Annual Technical Conference*, pages 11:1–11:14, Berkeley, CA, USA, 2007. USENIX Association.
105. M. W. Storer, K. M. Greenan, E. L. Miller, and K. Voruganti. Potshards -a secure, recoverable, long-term archival storage system. *Trans. Storage*, 5:5:1–5:35, June 2009.
106. A. Subbiah, M. Ahamad, and D. M. Blough. Using byzantine quorum systems to manage confidential data. Technical report, Georgia Institute of Technology, 2004.
107. A. Subbiah and D. Blough. Practical share renewal for large amounts of data. Technical report, School of Electrical and Computer Engineering - Georgia Institute of Technology, 2005.
108. A. Subbiah and D. M. Blough. An approach for fault tolerant and secure data storage in collaborative work environments. In *In Proceedings of the 2005 ACM Workshop on Storage Security and Survivability*, pages 84–93. ACM Press, 2005.

109. L. Trevisan. Construction of extractors using pseudo-random generators (extended abstract). In *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, STOC '99, pages 141–148, New York, NY, USA, 1999. ACM.
110. R. Ursin, F. Tiefenbacher, T. Schmitt-Manderbach, H. Weier, T. Scheidl, M. Lindenthal, B. Blauensteiner, T. Jennewein, J. Perdigues, P. Trojek, B. Ömer, M. Fürst, M. Meyenburg, J. Rarity, Z. Sodnik, C. Barbieri, H. Weinfurter, and A. Zeilinger. Entanglement-based quantum communication over 144km. *Nature Physics*, 3(7):481–486, June 2007.
111. S. Vadhan. Extracting all the randomness from a weakly random source. Technical report, Electronic Colloquium on Computational Complexity, 1998.
112. S. P. Vadhan. Constructing locally computable extractors and cryptosystems in the bounded-storage model. *Journal of Cryptology*, 17:43–77, 2004. 10.1007/s00145-003-0237-x.
113. G. S. Vernam. Cipher printing telegraph systems for secret wire and radio telegraphic communications. *American Institute of Electrical Engineers*, XLV:109–115, 1926.
114. S. Wiesner. Conjugate coding. *SIGACT News*, 15:78–88, January 1983.
115. S. Wolf. *Information-theoretically and computationally secure key agreement in cryptography*. PhD thesis, ETH Zurich, Zurich, 1999.
116. S. Wolf. Unconditional security in cryptography. In I. Damgård, editor, *Lectures on Data Security*, volume 1561 of *Lecture Notes in Computer Science*, pages 217–250. Springer Berlin / Heidelberg, 1999.
117. T. Wong, C. Wang, and J. Wing. Verifiable secret redistribution for archive systems. In *Security in Storage Workshop, 2002. Proceedings. First International IEEE*, pages 94 – 105, dec. 2002.
118. T. M. Wong, C. Wang, and J. M. Wing. Verifiable secret redistribution for threshold sharing schemes. Technical report, School of Computer Science Carnegie Mellon University Pittsburgh, 2002.
119. W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, Oct. 1982.
120. A. D. Wyner. The wire-tap channel. *Bell System Technical Journal*, 54(8):1355–1387, 1975.
121. V. Yakovlev, V. Korzhik, and G. Morales-Luna. Key distribution protocols based on noisy channels in presence of an active adversary: Conventional and new versions with parameter optimization. *Information Theory, IEEE Transactions on*, 54(6):2535 –2549, june 2008.
122. V. Yakovlev, V. I. Korzhik, G. Morales-Luna, and M. Bakaev. Key distribution protocols based on extractors under the condition of noisy channels in the presence of an active adversary. *CoRR*, abs/1005.3184, 2010.