Quantum information theory (20110401)

Lecturer: Jens Eisert

Chapter 7: Quantum key distribution

Contents

Chapter 7

Quantum key distribution

7.1 Elements of cryptography

7.1.1 A short history of cryptography

Ideas of cryptography and secret communication are presumably about as old as mankind is. There are many reasons why one would like to communicate with a legitimate recipient while making sure at the same time that nobody else listens in to the conversation. This feature of communication is intricately intertwined with rather obvious features of human behaviour. For this reason it may not be a huge surprise that the history of cryptography reads like a crime story in the first place.

Examples of applications of cryptography from the more recent past (viewed from the perspective of the history of mankind, that is) include the cryptographic encoding of messages by a scytale, a device used as a cipher by the ancient Greeks and Spartans during military campaigns, first mentioned by the Greek poet Archilochus, who lived in the 7th century BC. It already features many aspects of a modern cryptographic scheme. It consists of a cylinder with a strip of parchment wound around it on which a message is written. The encryption arises from the fact that both the sender and the legitimate receiver share the cylinder. Once this is available, one can wind the parchment around it to generate a perfectly readable message. Without it, the message seems scrambled. The key point is that while two legitimate parties share the same object (a cylinder in this case, so a key in more modern terms), illegitimate users would not have access to this object. While this idea gives rise to a code that can obviously be broken, it has a security level that is presumably sufficient to reflect combat situations in the ancient world.

Turning to more recent events, it is well known that the fates of history in times of the second world war have been deeply intertwined with the history of secure communication. For example, Admiral Isoroku Yamamoto, the leading military commander of the Japanese Navy during World War II and the architect for the attack on Pearl Harbor, announced his advent to the front line base on the island of Bougainville to boost morale – of course strictly encrypted, that is. Only that it was not sufficiently encrypted after all. The encryption system used – the Japanese Naval Cipher JN-25D

in this case – was intercepted and with some effort successfully decrypted by US naval intelligence units. By the time, Yamamoto was arriving, the US was already there.

Maybe even more prominently, during World War II, efforts of encryption and efforts of deciphering messages had a decisive impact. Submarines obviously make sense only if their precise location can be concealed. The *Enigma machine* was the machine in the focus of a number of pivotal events. It was an electro-mechanical rotor cipher machine, invented by the German engineer Arthur Scherbius at the end of World War I and later developed into various variants, that were developed in the early 20th century to protect commercial, diplomatic and military communication. What the Enigma does, basically, is to transform each letter into a product of permutations. Unlike the previously mentioned cryptosystems, it required serious effort to break the code. An early version of the Enigma was broken by the Polish General Staff's Cipher Bureau in December 1932. Later versions used by Nazi Germany could initially not be deciphered; early on in WW II, the British Government Code and Cypher School at Bletchley Park built up an extensive cryptanalytic capability to break later versions of the machine. Alan Turing, a Cambridge University mathematician and logician and the inventor of the famous paradigmatic Turing machine, provided much of the key insights that eventually led to the breaking the naval Enigma, which had a major influence on the naval war. Once messages sent by submarines could be deciphered, the advantage of submarines was gone, with significant implications on how the war developed. That is to say, Alan Turing and his team at Bletchley Park had a major contribution to allied victory (a state of affairs that was later less appreciated when he was very badly treated, but this is a different matter, intertwined with another historical development).

These examples are mentioned only to highlight how the history of cryptography – as a history of code making and code breaking – is intimately intertwined with important events in history. By no means is this supposed to mean that the use of cryptography is confined to the military realm. Quite to the contrary, the use of cryptography is permeant to many aspects of our modern lives, in fact, it is ubiquitous. Whenever one uses WhatsApp, `https`, or any instance of internet banking, one resorts to a cryptographic scheme. Secure communication has become a pillar of how we communicate.

7.1.2 One-time pads

Getting more concrete: One can communicate securely if two parties share the same key. In retrospect this may seem obvious, at the same time it used to be far from clear. The one-time pad was developed by Gilbert Vernam in 1917, proving that there is an absolutely secure coding scheme which is secure against eavesdroppers with unlimited computational power. In the one-time pad, a plaintext is encoded making use of a secret key (a pad, for that matter) that has the same length as the plaintext itself. The very same key is also employed by the legitimate receiver to decode the message. Given that one makes use of the key only once, the encryption scheme is absolutely secure, a statement that has later been proven by Claude Shannon. In modern cryptographic systems (such as the *Data Encryption Standard (DES)* and the *Advanced Encryption Standard (AES)*) now used widely, shorter keys are being made use of to encrypt longer messages, for obvious pragmatic reasons. Such an approach uses fewer resources, but at the same time is not to the same extent provably secure as the one-time pad is. In

any case, ultimately, at the heart of the matter is how to establish a secure key in the first place.

7.1.3 Public key distribution schemes

The most commonly used scheme is based on so-called *public key cryptographic protocols*, prominently the famous *RSA scheme* named after Ron Rivest, Adi Shamir, and Leonard Adleman. This ingenious idea has actually been invented twice, once by RSA and once by James H. Ellis. Ellis was a British engineer and cryptographer who in 1970 also invented a public key distribution scheme while working at the Government Communications Headquarters (GCHQ) in Cheltenham. At the time his results were kept secret; they became available only later after the embargo had been lifted. In public key distribution schemes, a message receiver, now and later on referred to as Bob, prepares two different cryptographic keys. One that is public and one that is private. Subsequently, Bob broadcasts the public key through an authenticated channel so that everyone who listens to this channel can acquire a copy of the public key. There is no requirement whatsoever to keep this public key secret. The original sender of the message, referred to as Alice, encodes her message with the public key from Bob and then sends out the encrypted message through a public insecure channel. The algorithm is set up in such a fashion so that the message encrypted with the public key can only be decrypted in conjunction with the private key.

Public key systems are widely used, basically any cryptographic scheme one encounters in electronic communication is based on a public key cryptographic scheme. The RSA scheme is practically secure, with a security level depending on the key length. Unfortunately, its security has not been proven. It rests on the existence of *one-way functions*: The multiplication is in $\mathtt{P}$, while factoring is contained in $\mathtt{NP}$.

The *RSA algorithm* involves basically four steps: key generation, key distribution, encryption and decryption. The core idea is the observation that it is practically possible to identify three very large positive integers e, d and n with the property that the modular exponentiation for all integers m with $0 \leq m < n$ satisfies

$$(m^e)^d = m(\mathrm{mod}\, n) \tag{7.1}$$

and that even knowing e and n or even m it can be extremely difficult to find d. RSA involves a public key and a private key. e basically takes the role of the public key, d is kept as the private key exponent. Primality test, the decision problem that asks whether a given number is a prime number or not, used to be in $\mathtt{NP}$, until a probabilistic algorithm in $\mathtt{BPP}$ became known, and later the algorithm was de-randomized to an algorithm in $\mathtt{P}$ (look for the Miller-Rabin primality test and Solovay-Strassen primality test). A proof of $\mathtt{P} = \mathtt{NP}$ would indeed prove that one-way functions do not exist, shaking the basis on which RSA rests. This would imply that there cannot be proven security in public key distribution schemes. However, the precise practical implications would depend on the specifics of the argument. For example, if the proof of $\mathtt{P} = \mathtt{NP}$ was not constructive, then this proof would not give advice on how to actually break the key.

7.1.4 Quantum computers potentially breaking public key schemes

In any case, there is no denying that the lack of provable security poses a significant security risk. RSA itself was a highly unexpected discovery, and one should hence not rule out the possibility that someone could find an efficient factoring algorithm and thus compromise most public cryptographic systems. What is more, a *quantum computer* can solve factoring in polynomial time (*Shor's algorithm* provides a quantum algorithm for factoring the runtime of which scales polynomially in the length of the input - it is in BQP in the language of computational complexity). Large-scale quantum computers do not exist yet, but the development is fast. In 2016, IBM made a 16 qubit cloud quantum computer publicly available as a cloud service based on superconducting circuits, which has been characterized using randomized benchmarking and developed into a 50 qubit machine in 2018. More recently still, Google announced the 128 qubit Brizzlecone chip, based on a similar architecture. These devices are still way too small (and too noisy) to pose a security risk. But their development is fast and the case for quantum computing is open. And indeed, large-scale quantum computers could break essentially all RSA based cryptographic schemes used today over night.

7.1.5 What quantum key distribution can deliver

Quantum key distribution is different. Its security on the level of the scheme is mathematically proven. Its security is based on very fundamental physical laws of nature. These are the laws of quantum mechanics. *Quantum mechanics* is the theory of the world at the small scale: That of atoms, ions and light quanta. But since the macroscopic world is ultimately built from such building blocks, it equally applies to the macroscopic world: It is the best physical theory of nature that we have today. In quantum key distribution, one envisions to make use of constituents in which the quantum features are most manifest. Practically speaking, one sends single photons (excitations of light modes), weak pulses or Gaussian light through fibres (the same kind of fibres that are used by the Telekom) or free space, even via satellites.

Ultimately, the security is rooted in structure elements of quantum mechanics: One cannot learn about the unknown quantum state of a quantum systems without disturbing the state. There are trade-offs: One can perform a gentle measurement, learn very little and at the same time disturb very little. And one can do hard projective measurements. But there is no way one can obtain some information about an unknown quantum state without changing the same state to some extent. An implication of this feature is that quantum information cannot be *copied* or *cloned*, as one commonly says in this context. It is impossible to build a machine that takes a physical system in an unknown quantum state and produces two quantum systems in the very same state. If one could do disturbance-free measurements, that would be possible, but the *no cloning* feature of quantum mechanics forbids that. We will see that this is a simple consequence of the linearity of quantum mechanical laws. Quantum key distribution is no far-fetched dream: It is already reality. One can commercially buy quantum cryptographic devices: The company IDQuantique is only one out of many offering such products. It has been one of the early successes of the field of quantum cryptography to implement a BB84 scheme (the simplest and most used scheme for quantum key distribution that we will

discuss soon) making use of an installed optical fibre cable linking Geneva and Nyon over 23 km through Lake Geneva in 1995, at remarkably low quantum bit error rates. This effort basically started the development of long-distance quantum key distribution. In the meantime, satellite-based quantum key distribution is being pursued.

Why is not all modern cryptography done via quantum key distribution and it is still a market niche? This has various reasons. The core reason is that reliable quantum key distribution over arbitrary distances is still hindered by serious technological obstacles. One needs to build so-called *quantum repeaters* to compensate for losses, in order to maintain security in the presence of realistically high noise levels. There is significant progress in this direction, but fully fledged quantum repeaters have not been implemented yet. This means that quantum key distribution is still confined to relatively short distances. Then, it is a marketing issue: The market may well grow a lot if people realize that the security claim in quantum key distribution is very different from that in public key distribution. Such processes take time. The BMBF (Bundesministerium für Bildung und Forschung) has "bug-proof communication" on its web page as one of the strategic aims, and indeed, there is a large scale project on realizing quantum repeaters by the BMBF, called Q.Link-X (which we are part of). It should be clear that quantum key distribution is no science fiction, but an important technology of tomorrow.

7.2 Quantum key distribution

7.2.1 BB84 quantum key distribution scheme as an example

We now turn to our first cryptographic application. This scheme, the famous BB84 scheme for *quantum key distribution*, was the historically first scheme for quantum key distribution, featuring in its name the initial letters of Bennett and Brassard, the names of its inventors. It is built on earlier work by Wiesner on *quantum money* and *conjugate coding*. It is both an ingenious scheme that lives up to the expectations of a modern quantum key distribution scheme and it also serves as a nice example of the quantum formalism laid out above.[1] It is based on the iterated use of single qubits only, so the state spaces we need to consider are merely $\mathcal{S}(\mathbb{C}^2)$. Interestingly, rigorous security proofs were only found more than a decade later.

In the BB84 protocol, Alice sends a string of qubits to Bob, prepared one by one, and hence in a product state. She prepares either states that are eigenstates of Pauli Z or she prepares eigenstates of Pauli X. That is to say, she prepares orthogonal states taken from two non-orthogonal bases. Specifically, the protocol proceeds as follows.

- Alice picks an i.i.d. random bit string $a \in \{0,1\}^n$.
- Alice picks a second i.i.d. random bit string $b \in \{0,1\}^n$. At this point, she does not reveal either of the two bit strings.
- Alice now prepares quantum states of single qubits that she sends to Bob. The basis picked will depend on b: If $b_i = 0$, she prepares the i-th state in the Z

[1] The recollection of how the BB84 scheme came about is an interesting story in its own right. Note also that Bennett, a theoretical physicist, actually implemented a first experimental demonstration of the scheme himself, using a setup that is still standing on his desk in his office at the IBM Watson Center.

basis with eigenvectors $\{|0\rangle, |1\rangle\}$, if $b_i = 1$, she prepares states in the X basis with eigenvectors $\{|+\rangle, |-\rangle\}$. That is to say, for the following values (b_j, a_j) she prepares

$$(0,0): \quad |0\rangle, \tag{7.2}$$
$$(0,1): \quad |1\rangle, \tag{7.3}$$
$$(1,0): \quad |+\rangle, \tag{7.4}$$
$$(1,1): \quad |-\rangle. \tag{7.5}$$

Note that $|0\rangle, |1\rangle \in \mathbb{C}^2$ are orthogonal, and so are $|+\rangle, |-\rangle \in \mathbb{C}^2$, but the respective bases are not orthogonal to each other.

- These states are sent to Bob via a quantum channel.
- Bob picks an i.i.d. random bit string $c \in \{0,1\}^n$.
- Depending on the value c_j, Bob measures in the Z basis ($c_j = 0$) or the X basis ($c_j = 1$).
- If the basis picked by Bob is the same one as the one Alice picked, so if for a value j one has $c_j = b_j$, then the outcome of the measurement will be deterministic: Bob will receive the measurement outcome d_j. If no eavesdropper is present, $d_j = a_j$ with certainty, following the above rule for quantum measurements.
- If in contrast the basis picked is different, so if for a value j one has $c_j \neq b_j$, then he will receive an i.i.d. random number, not correlated with a_j. At this point, Bob cannot judge, however, which is the case, since at this point, Bob has not received any classical information from Alice yet.
- Alice and Bob communicate classically over the bases used, so they reveal the bit strings b and c. The string a is not revealed at any time, however.
- Alice and Bob discard all cases j for which $c_j \neq b_j$. This will happen in expectation in half the cases. They end up with a bit string I of expected length $n/2$.
- They take the measurement outcomes and values a_j, $j \in I$.
- In order to determine the presence of an eavesdropper, Alice and Bob now compare a predetermined subset $J \subset I$ of the bit string I established. According to the quantum mechanical rules, the values $d_j = a_j$ should follow, if no third party (an "eavesdropper", commonly referred to as Eve) was present. If an eavesdropper has gained any information about the quantum states sent, this must introduce errors in Bob's measurements. Other environmental conditions can give rise to errors of the same type. If the rate of bits differing in J is $p > p_0$, they will abort the key and try again, possibly with a different quantum channel, as the security of the key can not be guaranteed under these circumstances. The threshold value p_0 is chosen so that the number of bits available to the eavesdropper Eve is less

than this number, privacy amplification can be used to reduce Eve's knowledge of the key to an arbitrarily small amount at the cost of reducing the length of the key.

Alice's random bit	0	1	1	0	1	0	0	1
Alice's random sending basis	+	+	×	+	×	×	×	+
Photon polarization Alice sends	↑	→	↘	↑	↘	↗	↗	→
Bob's random measurement basis	+	×	×	×	+	×	+	+
Photon polarization Bob measures	↑	↗	↘	↗	→	↗	→	→
Public discussion of basis								
Shared secret key	0		1			0		1

The remaining bit string $I \backslash J$ is the raw key. Why does this give rise to a secure key? We will look at security proofs later. The point is that by the time the quantum systems are being sent, Eve has no chance to guess the correct basis any better than making random choices. In this way, she has to introduce errors to the quantum state with high probability. In case she guessed right, she will get the right outcome, as $|\langle 0|0\rangle| = |\langle 1|1\rangle| = 1$. In the other cases, however, she will get uniformly random outcomes, as $|\langle +|0\rangle| = |\langle -|1\rangle| = 1/2$. Of course, she is not forced to precisely do such measurements, as she is free in her choices. However, this will not help her. This is a consequence of a very basic theorem we will encounter later. By the time the measurement bases are revealed, it is too late, and she cannot make use of that information any more. It is the key point of quantum key distribution that this idea does not only work if Eve sticks to performing von-Neumann measurements in the given basis. It works if Eve performs arbitrary measurements, even ones that are entangled over all invocations of the preparations, and making unrealistic assumptions about her. She might even have a quantum computer at her disposal, allowing for arbitrary coherent manipulation of all qubits sent. Still, asymptotically, she will not gain information about the key.

7.2.2 Security proofs

Quantum key distribution offers the claim of secure key distribution in the presence of an eavesdropped that is attributed unlimited, even unrealistic, resources. Historically, however, before full security proofs were available, particular kinds of attacks were considered.

- *Intercept-resend attack:* The simplest type of a possible attack is the intercept-resend attack, in which Eve performs projective measurements, makes use of the

measurement outcome, and prepares a new quantum system in a suitable state. It is easy to see that the BB84 scheme is secure against such intercept-resend attacks.

- *Individual attacks:* In this type of attack, Eve performs generalized measurements as laid out above, but interacts with each qubit (or other quantum system) in the channel separately and independently. Invoking the above Stinespring dilation, physically, this means Eve lets the quantum system transmitted interact with an auxiliary system each which is subsequently measured in a von-Neumann measurement. The intercept-resend attack is an instance of such an attack. Generally, individual attacks are the most realistic ones given present technology. *Photon number splitting attacks* in quantum optical schemes in which weak pulses are being sent are specifically important instances of such individual attacks.

- *Collective attacks:* This is a yet more general kind of attack. Here, Eve again performs generalized measurement. Again, she prepares independent auxiliary systems which interact with the quantum systems transmitted. But now she can perform a joint measurement on the collection of auxiliary systems.

- *Coherent attacks:* This is an attack that is in no way limited in what Even is allowed to do.

Any attack will give rise to errors in the transmission. This the *quantum bit error rate* $Q \geq 0$ captures the rate of errors in transmission. A key quantity used in security proofs is that of the *quantum mutual information* between Alice and Bob, as well as Alice and Eve and Bob and Eve. The quantum mutual information of a bipartite state defined on $\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_B$ is given by

$$I(A:B) = S(\rho_A \otimes \rho_B) - S(\rho). \tag{7.6}$$

This quantity captures correlations in quantum states (both classical and quantum correlations, i.e., entanglement), taking a zero value for product states.

7.2.3 General strategies of security proofs

Historically, the discussion of security of quantum key distribution protocols centred around the discussion of specific attacks. While this is instructive, it falls short of the actual promise of quantum key distribution. The first security proofs that considered an unbounded adversary (and hence coherent attacks) were given more than a decade after the introduction of the first schemes. Work by Preskill and Shor is noteworthy in this respect, in that it takes a very physical approach and links the theory of security of quantum key distribution to that of quantum error correction and entanglement distillation explained above.

Only again much later, it was noticed that the security criterion used so far may well be insufficient: It does guarantees that an eavesdropper cannot guess the key, so in this sense the scheme is secure. But this is only true of the key is not used subsequently. If part of the key is ultimately to an eavesdropper (e.g., when it is used

to encrypt a message that is known to her), the rest may become insecure. Based on these insights, a more stringent security criterion for quantum key distribution has been introduced, concomitant with new security proofs. If $\rho_{K,E}$ is he joint state of the final key generated and the quantum information gathered by an eavesdropper Eve, then this state must be close to an ideal key τ_K which is perfectly uniform and is independent from the adversary's information ρ_E, as

$$(1 - p_{\text{abort}})D(\rho_{K,E}, \tau_K \otimes \rho_E) \leq \varepsilon \tag{7.7}$$

where p_{abort} is the probability that the protocol aborts, $D(.,.)$ is again the trace-norm distance and $\varepsilon \in [0, 1]$ is a small real number.

7.3 Quantum repeaters for secure long-distance quantum key distribution

7.3.1 Entanglement based key distribution schemes

In Section 7.2.1 we have encountered the BB84 scheme as a scheme for quantum key distribution in which quantum systems are being prepared and then sent through a quantum channel. It is the most important and still most practical scheme for quantum key distribution. Having said that, there are many other schemes for quantum key distribution. One way of categorizing them is into *prepare-and-measure* schemes (such as the original BB84 scheme) and into *entanglement-based schemes*. The latter type of scheme at first sight seems quite distinctly different: One first prepares an entangled state, to then – once distributed – performs local measurements to establish a key. However, a moment of thought reveals that this is something very similar: In fact, every prepare-and-measure scheme can be seen as an equivalent to an entanglement based scheme. Take the maximally entangled state vector of two qubits

$$|\Omega\rangle = \frac{1}{\sqrt{2}}(|0, 1\rangle - |1, 0\rangle). \tag{7.8}$$

If Alice on one side performs a Z measurement and obtains 0 as her outcome, Bob's system will be in $|1\rangle$. Similarly, upon a 1 outcome, Bob's system will projected to be in $|0\rangle$. That is to say, the measurement on Alice's side is effectively like a (non-deterministic) preparation of a quantum state on Bob's side. Since the state vector is UU-invariant, i.e., $(U \otimes U)|\Omega\rangle = |\Omega\rangle$ for all $U \in U(2)$, the same holds true for X measurements, and in fact any other measurement in the same basis on both sides. If Alice projects her system into $|+\rangle$, Bob will see $|-\rangle$, and if she encounters $|-\rangle$, Bob will have $|+\rangle$. Again, this can be seen as a probabilistic preparation. This connection between prepare-and-measure schemes has long been observed. In fact, a precondition for security in any scheme, including prepare-and-measure schemes, is the presence of entanglement in the equivalent entanglement (or correlation) based scheme.

7.3.2 Entanglement swapping and distillation

However, there is an important conceptual difference: In prepare-and-measure schemes, there is little one can do about losses when sending quantum systems through quantum channels. In entanglement based schemes, one can do something about it. Accepting this, the key question is: How can one establish a maximally entangled state between arbitrary distances in the first place? This is possible by means of *quantum repeaters*. They consist of two steps:

- First, they involve *entanglement distillation*. We have discussed this above: This is an LOCC protocol aimed at extracting maximally entangled states. In effect, both parties will end up with fewer, almost maximally entangled states. Entanglement has been "distilled", in a similar way as one can extract high percentage alcohol from a liquid in which alcohol is only present in a dilute form. These maximally entangled states can be used in subsequent steps. This is a highly interesting procedure: Entanglement, so intrinsic quantum correlations, are here manipulated like an interconvertible resource.

- Then there are steps of *entanglement swapping*. This step is maybe even more intricate and interesting. Think of two maximally entangled states shared by Alice and Bob on the one hand and Bob and Charlie on the other hand. So let us start from

$$|\psi\rangle = |\Omega\rangle_{A,B_1} \otimes |\Omega\rangle_{B_2,C}, \tag{7.9}$$

with again

$$|\Omega\rangle = (|0,0\rangle + |1,1\rangle)/\sqrt{2}. \tag{7.10}$$

That is to say, Bob holds two halves of maximally entangled states. The two copies will not have any shared history. Now Bob can perform a projective measurement, projecting the state vector into

$$(\mathbb{1}_A \otimes \langle\Omega|_{B_1,B_2} \otimes \mathbb{1}_C)|\psi\rangle = \frac{1}{\sqrt{2}}|\Omega\rangle_{A,C}. \tag{7.11}$$

That is to say, after the projective measurement, A and C are in a maximally entangled state, even though these particles have no joint history whatsoever. The entanglement has been "swapped". Of course, in a projective measurement, this would only work in a probabilistic fashion. However, a moment of thought reveals that this can be made deterministic, in that for each outcome of a joint measurement on B_1 and B_2, one can find a Pauli correction on A and C so that deterministically, $|\Omega\rangle_{A,C}$ is reached. The reason for this is ultimately that

$$\{(\mathbb{1}\otimes\mathbb{1})|\Omega\rangle,\ (X\otimes\mathbb{1})|\Omega\rangle,\ (Y\otimes\mathbb{1})|\Omega\rangle,\ (Z\otimes\mathbb{1})|\Omega\rangle\} \tag{7.12}$$

for Pauli operators $X, Y, Z, \mathbb{1}$ constitute a basis of the maximally entangled states on $\mathbb{C}^2 \otimes \mathbb{C}^2$.

7.3.3 Full quantum repeater schemes

A *quantum repeater* now makes use of such such steps in a hierarchical, tree-like fashion. It involves steps of entanglement distillation between neighbours, followed by entanglement swapping steps. There are many variants of such quantum repeaters, as well as numerous suggestions for experimental realizations thereof. In fact, the reliable realization of quantum repeaters is the key obstacle on the path to secure quantum key distribution over arbitrary distances – while near-distance quantum key distribution is perfectly feasible. The trouble is that notions of entanglement distillation and entanglement swapping are generally assumed to rely on *quantum memories* that store quantum information reliably. Since quantum information has to be transmitted via light (other quantum systems are hardly feasible for this task), and quantum information is stored in matter qubits, one needs coherent frequency converters (to align the respective frequencies) and needs to map quantum states of light onto atoms, ions, or atomic ensembles. Then, it has to be read out, but needless to say, all coherently again at negligible losses. This still constitutes a technological road block, even though progress is fast. In fact, surprising as this may sound, each of the above mentioned components has already been achieved in experiments. Once this road block is overcome, secure quantum communication over arbitrary distances is feasible.