The Universal Composable Security of Quantum Key Distribution

Michael Ben-Or[1,4,6], Michał Horodecki[2,6], Debbie W. Leung[3,4,6],
Dominic Mayers[3,4], and Jonathan Oppenheim[1,5,6]

[1] Institute of Computer Science, The Hebrew University, Jerusalem, Israel
[2] Institute of Theoretical Physics and Astrophysics, University of Gdańsk, Poland
[3] Institute of Quantum Information, California Institute of Technology, Pasadena, USA
[4] Mathematical Science Research Institute, Berkeley, USA
[5] DAMTP, University of Cambridge, Cambridge, UK
[6] Isaac Newton Institute, University of Cambridge, Cambridge, UK

benor@cs.huji.ac.il, fizmh@univ.gda.pl, wcleung@cs.caltech.edu,
dmayers@cs.caltech.edu, and J.Oppenheim@damtp.cam.ac.uk

Abstract. The existing unconditional security definitions of quantum key distribution (QKD) do not apply to joint attacks over QKD and the subsequent use of the resulting key. In this paper, we close this potential security gap by using a universal composability theorem for the quantum setting. We first derive a composable security definition for QKD. We then prove that the usual security definition of QKD still implies the composable security definition. Thus, a key produced in any QKD protocol that is unconditionally secure in the usual definition can indeed be safely used, a property of QKD that is hitherto unproven. We propose two other useful sufficient conditions for composability. As a simple application of our result, we show that keys generated by repeated runs of QKD degrade slowly.

1 Introduction

Quantum cryptography differs strikingly from its classical counterpart. On one hand, quantum effects are useful in the construction of many cryptographic schemes. On the other hand, dishonest parties can also employ more powerful quantum strategies when attacking cryptographic schemes.

The security of quantum key distribution. One of the most important quantum cryptographic applications is quantum key distribution (QKD) [1–3]. The goal of key distribution (KD) is to allow two *remote* parties, Alice and Bob, to share a *secret* bit string. Classically, KD cannot be unconditionally secure (i.e. secure against all possible classical attacks) (see Sect. 2). Furthermore, the security of existing KD schemes is based on assumptions in computation complexity or limitations of the memory space of the adversary, Eve. In contrast, QKD is based on an intrinsic property of quantum mechanics, "extracting information about an unknown quantum state inevitably disturbs it," [4] which allows eavesdropping activities to be detected in principle. Indeed, QKD can be *unconditionally secure*, i.e., against Eve whose capability is only limited by quantum mechanics [5–11]. Furthermore, QKD remains secure even if the quantum states are sent through a noisy quantum channel, as long as the observed error rates are below certain threshold values.

In what sense is QKD secure? We will describe the assumptions and security definitions more formally in Sect. 2. In QKD, Alice and Bob are assumed to start with a small initial key K_i (for authentication purposes). They have access to uncorrelated randomness that is not controlled by Eve. They may exchange quantum and classical messages in both directions via channels that are completely under the control of Eve, and may perform local quantum operations and measurements. Based on their measurement outcomes, Alice and Bob either abort QKD or generate their respective keys $K_{\rm A}, K_{\rm B}$. Correspondingly, we say that the QKD test is failed or passed, and the events can be described as $M{=}0$ or $M{>}0$, where M is the length of the key generated. Eve also obtains quantum and classical data (her "view" or "transcript") from which she extracts classical data $K_{\rm E}$ via a measurement. What happens during a specific run of QKD depends on Eve's strategy as well as the particular outcomes of the coins and quantum measurements of all the parties. However, the security of QKD can still be captured by requiring that (1) the conditional mutual information $I(K_{\rm E}\,{:}\,K_{\rm A}, K_{\rm B}\,|M)$ is negligible and (2) for all eavesdropping strategies with nonnegligible $\Pr(M{>}0)$, $K_{\rm A}$, $K_{\rm B}$ are near-uniform and $\Pr(k_{\rm A} \neq k_{\rm B})$ is negligible. Throughout the paper, we use capitalized letters $K_{\rm A}$, $K_{\rm B}$, $K_{\rm E}$, and M to denote the random variables, and uncapitalized letters to denote specific outcomes.

The security problem of using QKD. Proofs of security of QKD (in the sense described above) address all attacks on the QKD scheme allowed by quantum mechanics. The problem is that QKD is *not* the only occasion for attack — further attack may occur when Alice and Bob use the keys generated. In particular, Eve may never have made a measurement during QKD to obtain any $K_{\rm E}$. Eve's transcript is a quantum state. She could have delayed measurements until after more attack during the application, a strategy with power that has no classical counterpart. In other words, security statements in QKD that revolve around bounding $I(K_{\rm E}\,{:}\,K_{\rm A}, K_{\rm B}\,|M)$ is *not applicable* if the key is to be used!

The limitations of mutual-information-based security statements were known as a folklore for some time (for example, see Sec. 4.2 in [11]). One of the earliest known security problems in QKD is the "key degradation problem" [12]: QKD requires a key for authentication, which in turns may come from previous runs of QKD. Since each run of QKD is slightly imperfect, repeated runs of QKD produce less and less secure keys. A conclusive analysis on the degradation has been elusive, since joint attacks over all runs of QKD have to be considered.

As it turns out, joint attacks on QKD and the subsequent use of the generated key have to be considered in many other occasions. For example, suppose Alice and Bob perform QKD to obtain a key, and then use the key to encrypt quantum states [13, 14]. Eve eavesdrops during both QKD and encryption and performs a collective measurement on the two eavesdropped states. It is well-known that such a collective measurement may yield more *accessible information* than the sum of information obtained in two separate measurements [15].

Our current study is further motivated by the results in [16, 17], which show that there are ensembles of quantum states that provide little accessible information on their own, but can provide *much more* information when a little more *classical* data is available. The extra information can be arbitrarily large compared to both the initial information and the amount of extra classical data. Such strange property reveals a new, unexpected, inadequacy of mutual-information-based statements. In particular, in the context of QKD, the usefulness of bounding the initial accessible information of Eve becomes very questionable, if Eve delays her measurement until further data is available during the application of the key — the security of the key is questionable even in *classical* applications!

The goal of the current paper is to study the security of using a key generated by QKD, i.e., the composability of QKD.

The universal composability approach. Composability is an active area of research that is concerned with the security of composing cryptographic primitives in a possibly complex manner. The simplest example is the security of using a cryptographic primitive as a subroutine in another application. We will follow the *universal composability* approach. For a specific task (functionality), a primitive that realizes the task is said to be universal composable if any application using the primitive is about as secure as one using the ideal functionality. A security definition that ensures *universal* composability was recently proposed by Canetti [18]. A simpler model in the quantum setting and a corresponding universal composable security definition were reported by by some of us [19, 20]. Universal composable security definitions are useful because they are in terms of the ideal functionality only, without reference to the potential application. The security of a complex protocol can then be analyzed in terms of the security of each individual component in a systematic and error-proof manner. In the quantum setting, universal composability provides the only existing systematic technique for analyzing security in the presence of subtleties including entanglement and collective attacks. In this paper, universal composability provides the precise framework for proving the security of using the keys generated from QKD, a problem that appears intractable at first sight.

An alternative approach to composability in the classical setting was obtained in [21], with a generalization to the quantum setting studied in [22, 23].

Main Results. We have pointed out a serious potential security problem in using the keys generated from QKD. We will address the problem in the rest of the paper. We derive a new security definition for QKD that is universal composable. The essence is that QKD and certain ideal KD should be indistinguishable from the point of view of potential adversaries. Then, we prove that the original mutual-information-based security definition implies the new composable definition. Other simple sufficient conditions for the composable security of QKD will be discussed. One of these conditions, high singlet-fidelity, has always been an intermediate step in the widely-used "entanglement-based" security proofs of QKD. We show that high singlet-fidelity is much more closely related to composable security than the usual security definition, and we obtain much better security bounds for known QKD schemes. We thus prove the security of using a key generated by QKD in various ways, and provide simple criteria for future schemes. As a corollary, we analyze the extent of key-degradation in repeated use of QKD [12].

Our work also has non-cryptographic applications in the study of correlations in quantum systems. The various security conditions are tied to correlation measures in quantum systems. Each derivation for the composable security for QKD is based on relating a pair of correlation measures.

Related work. Since the current result was initially presented [24, 25], various related results were reported. The composable security of generic classes of QKD schemes were proved in [26, 27], following a different approach of showing the composable security of certain privacy amplification procedures against quantum adversaries [26]. These related works share the concerns raised in this paper, with results complementary to ours.

Organization of the paper. We end this section by introducing some basic elements in the quantum setting. We review QKD in Sect. 2, stating our definitions and assumptions more formally. In Sect. 3, we review the quantum universal composability theorem. We will restrict ourselves to the simpler case concerning unconditional security. We start describing our main results in Sect. 4, which contains a derivation of a simple criteria for the universal composable security for QKD. In Sect. 5, we prove that the usual security definition for QKD implies the universal composable security. In addition, we demonstrate two other sufficient conditions for composable security. One is based on bounding the Holevo information of Eve on the key. The other is based on bounding the singlet-fidelity in security proofs using entanglement-purification. The latter implies much better security of existing QKD protocols than is generically implied by the usual security definition. We conclude with lessons learnt from the current results. Frequently used notations and functions are listed in the appendix.

Basic elements of quantum mechanics. A quantum system or register is associated with a Hilbert space $\mathbb{H}$. We only consider finite dimensional Hilbert spaces. Let $\mathcal{B}(\mathbb{H})$ and $\mathbb{U}(\mathbb{H})$ denote, respectively, the set of bounded operators and the unitary group acting on $\mathbb{H}$. We loosely refer to the system as $\mathbb{H}$ also. A composite quantum system is associated with the tensor product of the Hilbert spaces associated with the constituent systems.

The state of $\mathbb{H}$ is specified by a positive semidefinite *density matrix* $\rho \in \mathcal{B}(\mathbb{H})$ of unit trace. A density matrix is a convex combination of rank-1 projectors (commonly called *pure states*) and represents a probabilistic mixture of pure states. Pure states can be represented as vectors in $\mathbb{H}$, up to a physically unobservable phase. $|\psi\rangle$ and $|\psi\rangle\langle\psi|$ denote the vector and rank-1 projector respectively.

A measurement $\mathcal{M}$ on $\mathbb{H}$ is specified by a POVM — a set of positive semidefinite operators $\{O_k\}$ such that $\sum_k O_k = I$. If the state is initially ρ, the measurement $\mathcal{M}$ yields the outcome k with probability $\mathrm{Tr}(O_k \rho)$ and changes the state to $\sqrt{O_k}\rho\sqrt{O_k}/\mathrm{Tr}(O_k\rho)$. $\mathcal{M}$ is said to be along a basis $\{|k\rangle\}$ if $\{O_k\} = \{|k\rangle\langle k|\}$. Measuring an unknown state generally disturbs it.

The most general evolution of the state is given by a trace-preserving completely-positive (TCP) linear map $\mathcal{E}$ acting on $\mathcal{B}(\mathbb{H})$. Any such $\mathcal{E}$ can be implemented by preparing a pure state in some ancillary system $\mathbb{H}'$, applying a joint unitary operator $U \in \mathbb{U}(\mathbb{H} \otimes \mathbb{H}')$, and discarding $\mathbb{H}'$ (i.e., a partial trace over $\mathbb{H}'$).

We mention two distance measures for quantum states. First, the trace distance $\|\rho_1 - \rho_2\|_1$ between two density matrices is twice the maximum probability of distinguishing between the two states. Second, the fidelity is $F(\rho_1, \rho_2) = \max_{|\psi_1\rangle, |\psi_2\rangle} |\langle\psi_1|\psi_2\rangle|^2$, where $\rho_{1,2} \in \mathcal{B}(\mathbb{H})$, $|\psi_{1,2}\rangle \in \mathbb{H}\otimes\mathbb{H}'$ are "purifications" of $\rho_{1,2}$ (i.e., $\mathrm{Tr}_{\mathbb{H}'}|\psi_{1,2}\rangle\langle\psi_{1,2}| = \rho_{1,2}$), and $\langle\cdot|\cdot\rangle$ is the inner product in $\mathbb{H}$.

We refer our readers to the excellent textbook by Nielsen and Chuang [28] for a more comprehensive review of the quantum model of information processing.

2 Quantum Key Distribution

The goal of key distribution (KD) is to allow two *remote* parties, Alice and Bob, to share a *secret* bitstring such that no third party, Eve, will have much information about the bitstring. KD is impossible unless Alice and Bob can identify one another and detect alterations of their communication. In other words, the task of *message authentication* is necessary for KD. There are unconditionally secure methods for authenticating a classical message with a much shorter key [29]. Thus, KD uses authentication as a subroutine, and achieves key expansion (producing a key using a much shorter initial key).

Classically, unconditionally secure KD between two remote parties is impossible. Classical physics permits an eavesdropper to have exact duplicates of all communications in any KD procedure without being detected. In contrast, while quantum key distribution (QKD) cannot prevent eavesdropping, it can detect eavesdropping. This allows Alice and Bob to avoid generating compromised keys with high probability. The usefulness of QKD is to avoid Alice and Bob being fooled into having a false sense of security. It is worth emphasizing what QKD does not offer. First, QKD does not promise to always produce a key, since Eve can cause QKD to be aborted with high probability by intense eavesdropping. Second, there is a vanishing but non-zero chance that Eve is undetected, so that one cannot make simple security statements conditioned on not aborting QKD.

How and why QKD works, through an example. Various QKD schemes have been proposed and we only name a few here: BB84 [1], E91 [2], B92 [3], and the six-state scheme [30, 31]. We illustrate the general features and principles behind QKD by describing the class of "prepare-&-measure schemes." Recall that Alice and Bob are given secure local coin tosses. Step 1: Alice first generates a random bitstring, encodes it in some quantum state ρ_{A}, and sends ρ_{A} to Bob through an insecure quantum channel controlled by Eve. During this time, Eve can manipulate the message (system $\mathbb{A}$) in any way allowed by quantum mechanics. Eventually, she will have to give some quantum message ρ_{B}

to Bob for QKD to proceed. Mathematically, Eve's most general operation can be described as attaching a private system $\mathbb{E}$ in the state $|0\rangle\langle 0|_{\mathrm{E}}$, applying a joint unitary operation U to produce a joint state $\rho = U\,(\rho_{\mathrm{A}} \otimes |0\rangle\langle 0|_{\mathrm{E}})\,U^{\dagger}$, and passing system $\mathbb{A}$ to Bob (relabeled as system $\mathbb{B}$). Thus, Bob and Eve share the joint state ρ, and $\rho_{\mathrm{B}} := \mathrm{Tr}_{\mathbb{E}}\,\rho$, $\rho_{\mathrm{E}} := \mathrm{Tr}_{\mathbb{B}}\rho$ are their respective reduced density matrices. Meanwhile, Bob measures ρ_{B} (according to his coin tosses). Step 2: Bob acknowledges to Alice receipt of the quantum message. Step 3: Only *after* Alice hears from Bob will further classical discussion be conducted over a public but authenticated channel. Step 4: At the end, based on their measurement outcomes and discussions, Alice and Bob either abort QKD ($m = 0$), or generate keys K_{A} and K_{B} ($m > 0$), and they announce m. Eve will have access to all the classical communication between Alice and Bob, besides the state ρ_{E}. She can measure ρ_{E} at any time to obtain a classical string K_{E}, though it is to her advantage to wait until after she receive the classical communication. See Fig. 1 for a schematic diagram for the class of prepare-&-measure QKD schemes.

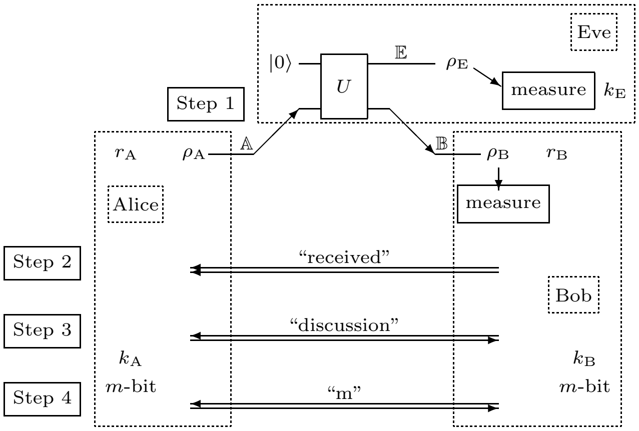

Fig. 1. Schematic diagram for the class of prepare-&-measure QKD schemes. The classical messages, represented by double lines, are available to Eve. Eve can make her measurement any time after step 1. Dashed boxes represent private laboratory spaces. Outcomes of Alice and Bob's local coins are represented by $r_{\mathrm{A}}, r_{\mathrm{B}}$.

The principle behind QKD is that, in quantum mechanics, one can only reversibly extract information from an unknown quantum state if the state is drawn from an orthogonal set [4]. Thus in the prepare-&-measure scheme described above, if Alice encodes her message using a random basis chosen from several nonorthogonal possibilities, and Eve is to obtain any information on the outcomes of K_{A}, K_{B}, then $\rho_{\mathrm{B}} \neq \rho_{\mathrm{A}}$. To detect the disparity, Bob measures some of the received qubits (the "test-qubits" chosen randomly to avoid Eve tailoring her attack) and discusses with Alice to check if his measurement outcomes are consistent with what Alice has sent. This intuition can be turned into a provably secure procedure. Alice and Bob estimate various error rates on the test-qubits. If the observed error rates are above certain thresholds, Alice and Bob abort

QKD. Otherwise, error reconciliation and privacy amplification are applied to the untested qubits to extract bitstrings k_A and k_B for Alice and Bob respectively. It is unlikely that the observed error rates are low while untested qubits have higher error rates. QKD remains secure whether the observed noise is due to natural channel noise or due to eavesdropping.

General features of any QKD scheme. There are other QKD schemes besides prepare-&-measure schemes, for example, the entanglement-based QKD schemes (see [2, 7, 32]). Unless otherwise stated, our discussion applies to all QKD schemes. The basic ingredients are still secure local coins, completely insecure quantum communication, and authenticated public classical communication between Alice and Bob. In the most general QKD scheme, the ingredients may be used in any possible way. Alice and Bob still obtain some bitstrings as the output keys, k_A and k_B, of certain length m. Eve's view is still given by some quantum and classical data, denoted collectively by $\rho_{\mathrm{E},k_\mathrm{A},k_\mathrm{B}}$, with explicit dependence on k_A, k_B. (Her view is a draw from an ensemble.)

We emphasize a limitation in QKD. Eve can be "lucky." For example, she may have attacked only the untested qubits, or may have attacked every qubit without causing inconsistency in Alice and Bob's measurements. Thus, it is unlikely, but still possible, for Eve to have a lot of information on the generated key without being detected. QKD does not promise that "*conditioned* on passing the test, the keys K_A, K_B will be so-and-so." With the above limitation in mind, there are several approaches to a proper security statement. The approach that is most commonly used in existing security proofs is to bound the probability that Alice and Bob generate bitstrings that are not equal, uniform, or private. We will use a more compact statement in the following.

Let n be a security parameter in QKD (for example, the number of qubits transmitted from Alice to Bob). Fix an arbitrary eavesdropping strategy. The attack induces a distribution $\Pr(M{=}m)$ on the key length M. The average value of M is typically a small fraction of n. The outcome m in a particular run of QKD depends on the outcome of the coins and measurements by Alice and Bob. We can assume that m is made *public* at the end of QKD. Recall $m > 0$ if the QKD test is passed and $m = 0$ if QKD is aborted.

Let $p_\mathrm{qkd}^{(m)}$ denote the distribution of $K_\mathrm{A}, K_\mathrm{B}$ generated in QKD conditioned on $|K_\mathrm{A}| = |K_\mathrm{B}| = m$, i.e.,

$$p_\mathrm{qkd}^{(m)}(k_\mathrm{A}, k_\mathrm{B}) = \Pr(K_\mathrm{A} = k_\mathrm{A}, K_\mathrm{B} = k_\mathrm{B} | M = m) \ . \tag{1}$$

Let $p_\mathrm{ideal}^{(m)}$ be the following distribution over two m-bit strings,

$$\begin{cases} p_\mathrm{ideal}^{(m)}(l, l) \ \ = 2^{-m} & \\ p_\mathrm{ideal}^{(m)}(l, l') = 0 & \text{if } l \neq l' \,. \end{cases} \tag{2}$$

Let $\mathcal{V}$ denote the set of exponentially decaying functions of n. With these notations, a simple statement for the security condition can be made.

Usual security definition for QKD. A QKD scheme is said to be secure if the following properties hold for all eavesdropping strategies.

- *Equality-and-uniformity:* $\exists \mu_1 \in \mathcal{V}$ s.t.

$$\sum_{m=0} \Pr(m) \, \| \, p^{(m)}_{\text{ideal}} - p^{(m)}_{\text{qkd}} \, \|_1 \ \leq \ \mu_1 \tag{3}$$

- *Privacy:* $\exists \mu_2 \in \mathcal{V}$ s.t.

$$\sum_{m=0} \Pr(m) \times I(K_{\text{E}} \!:\! K_{\text{A}}, K_{\text{B}} \,|\, M = m \,) \ \leq \ \mu_2 \tag{4}$$

where I denotes the mutual information [33] between K_{E} and $K_{\text{A}}, K_{\text{B}}$ conditioned on $M = m$. Using the equality condition, we only need to focus on $k_{\text{A}} =: k$ in (4). In particular,

- *Privacy:* $\exists \mu_2' \in \mathcal{V}$ s.t.

$$\sum_{m=0} \Pr(m) \times I(K_{\text{E}} \!:\! K \,|\, M = m \,) \ \leq \ \mu_2' \tag{5}$$

The above security conditions revolve around bounding expressions that can be interpreted as deviations from the desired properties, averaged over m. The product in each summand is bounded, precisely capturing the security requirement that an undesired event occurs with low probability. Note that the $m = 0$ terms do not contribute, as $\| \, p^{(m)}_{\text{ideal}} - p^{(m)}_{\text{qkd}} \, \|_1 \ = \ 0$ and $I(K_{\text{E}} \!:\! K_{\text{A}}, K_{\text{B}} \,|\, M = 0 \,) = 0$.

3 Quantum Universal Composability Theorem

Cryptographic protocols often consist of a number of simpler components. A single primitive is rarely used alone. A strong security definition for the primitive should thus reflect the security of using it within a larger application. This allows the security of a complex protocol to be based only on the security of the components and how they are put together, but not in terms of the details of the implementation.

A useful approach is to consider the *universal composability* of cryptographic primitives [18–20]. The first ingredient is to ensure the security of a *basic composition.* We need a security definition stated for a single execution of the primitive that still guarantees security of composition with other systems. This definition involves a description of some ideal functionality of the primitive (i.e. the ideal task the primitive should achieve). More concretely, we want a security definition such that, if σ is a secure realization of an ideal subroutine σ_{I}, and a protocol $\mathcal{P}$ using σ_{I}, written as $\mathcal{P}{+}\sigma_{\text{I}}$, is a secure realization of $\mathcal{P}_{\text{I}}$ (the ideal functionality of $\mathcal{P}$), $\mathcal{P}{+}\sigma$ is also a secure realization of $\mathcal{P}_{\text{I}}$. Throughout the paper, we denote the associated ideal functionality of a protocol by adding a subscript I, and we denote a protocol $\mathcal{P}$ calling a subprotocol σ as $\mathcal{P}{+}\sigma$ (this last expression stretches the meaning of $\mathcal{P}$ a little bit to refer to the module of $\mathcal{P}$ calling σ). The second ingredient is a universal composability theorem stating how a complex protocol can be built out of secure components. It is simply a recipe on how to securely perform basic composition recursively.

The simplifications in analyzing the composable security of QKD. Our goal is to analyze the unconditional security of QKD using known results in quantum universal composability [19, 20]. The setting for QKD is simpler than that considered in [19, 20] in two important aspects. First, we are only concerned with unconditional security. Second, in QKD, Alice and Bob are known to be honest, and Eve is known to be adversarial, and no party is corrupted unpredictedly. The formal corruption rules are not used in our derivation of a composable security definition for QKD. We will describe a simplified model that is sufficient for our derivation of a universal composable security definition for QKD. This definition is applicable in the general setting considered in [19, 20] – so long as an appropriate model is used for analyzing the rest of the application when applying Theorems 1 and 2.

The simplified model. We first describe the model for quantum protocols and other concepts involved in the quantum composable security definition. We base our discussion on the (acyclic) quantum circuit model (see, for example, [34, 35]), with an important extension [20] (see also the endnotes [36]). Throughout the paper, we only consider circuits in the extended model.

1. Structure of a protocol. A (cryptographic) protocol $\mathcal{P}$ can be viewed as a quantum circuit in the extended model [20, 36], consisting of inputs, outputs, a set of registers, and some partially ordered operations. A protocol may consist of a number of subprotocols and parties. Each subprotocol consists of smaller units called "unit-roles," within each the operations are considered "local." For example, the operations and registers of each party in each subprotocol form a unit-role. Communications between unit-roles within a subprotocol represent *internal communications*; those between unit-roles in different subprotocols represent input/output of data to the subprotocols. A channel is modeled by an ordered pair of operations by the sender and receiver on a shared register. The channel available for the communication determines its security features.

2. The game: security in terms of indistinguishability from the ideal functionality. Let $\mathcal{P}_{\rm I}$ denote the ideal functionality of $\mathcal{P}$. Intuitively, $\mathcal{P}$ is secure (in a sense defined by $\mathcal{P}_{\rm I}$) if $\mathcal{P}$ and $\mathcal{P}_{\rm I}$ behave similarly under any adversarial attack. "Similarity" between $\mathcal{P}$ and $\mathcal{P}_{\rm I}$ is modeled by a game between *an environment* $\mathcal{E}$ and *a simulator* $\mathcal{S}$. These are sets of registers and operations to be defined, and they are sometimes personified in our discussion. In general, $\mathcal{P}$ and $\mathcal{P}_{\rm I}$ have very different internal structures and are very distinguishable, and the simulator $\mathcal{S}$ is added to $\mathcal{P}_{\rm I}$ to make an extended ideal protocol $\mathcal{P}_{\rm I}+\mathcal{S}$ that is less distinguishable from $\mathcal{P}$. $\mathcal{E}$ consists of the adversaries that act against $\mathcal{P}$ and an application protocol that calls $\mathcal{P}$ as a subprotocol. At the beginning of the game, $\mathcal{P}$ or $\mathcal{P}_{\rm I}+\mathcal{S}$ are picked at random. $\mathcal{E}$ will call and act against the chosen protocol, and will output a bit Γ at the end of the game. The similarity between $\mathcal{P}$ and $\mathcal{P}_{\rm I}+\mathcal{S}$ (or the lack of it) is captured in the statistical difference in the output bit Γ.

3. Valid $\mathcal{E}$. The application and adversarial strategy of $\mathcal{E}$ are first chosen (the same whether it is interacting with $\mathcal{P}$ or $\mathcal{P}_{\rm I}$+$\mathcal{S}$). $\mathcal{E}$ has to obey quantum mechanics, but is otherwise unlimited in computation power. If $\mathcal{P}$ is chosen in the game, $\mathcal{E}$ can (I) control the input/output of $\mathcal{P}$, (II) attack insecure internal communication as allowed by the channel type, (III) direct the adversarial parties to interact with the honest parties in $\mathcal{P}$. $\mathcal{E}$+$\mathcal{P}$ has to be an acyclic circuit in the extended model [20, 36].

4. Valid $\mathcal{P}_{\rm I}$ *and* $\mathcal{S}$*.* If $\mathcal{P}_{\rm I}$+$\mathcal{S}$ is chosen in the game, $\mathcal{E}$ (I) controls the input/output of $\mathcal{P}_{\rm I}$ as before. However, the interaction given by (II) and (III) above will now occur between $\mathcal{E}$ and $\mathcal{S}$ instead. ($\mathcal{S}$ is impersonating or simulating $\mathcal{P}$.) The strategy of $\mathcal{S}$ can depend on the strategy of $\mathcal{E}$. $\mathcal{P}_{\rm I}$ should have the same input/output structure as $\mathcal{P}$, but is otherwise arbitrary. (Of course, the security definition is only useful if $\mathcal{P}_{\rm I}$ carries the security features we want to prove for $\mathcal{P}$.) In particular, $\mathcal{P}_{\rm I}$ may be defined with internal channels and adversaries different from those of $\mathcal{P}$. $\mathcal{S}$ can (II′) attack insecure internal communication of $\mathcal{P}_{\rm I}$ and (III′) direct the adversarial parties to interact with the honest parties in $\mathcal{P}_{\rm I}$. Thus, $\mathcal{P}_{\rm I}$ exchanges information with $\mathcal{S}$, and this can modified the security features of $\mathcal{P}_{\rm I}$. To $\mathcal{E}$, $\mathcal{S}$ acts like part of $\mathcal{P}_{\rm I}$, "padding" it to look like $\mathcal{P}$, while to $\mathcal{P}_{\rm I}$, $\mathcal{S}$ acts like part of $\mathcal{E}$. It is amusing to think of $\mathcal{S}$ as making a "man-in-the-middle" attack between $\mathcal{E}$ and $\mathcal{P}_{\rm I}$. Finally, $\mathcal{E}$+$\mathcal{P}_{\rm I}$+$\mathcal{S}$ has to be an acyclic circuit in the extended circuit model [20, 36]. See Fig. 2 for a summary of the game and the rules.

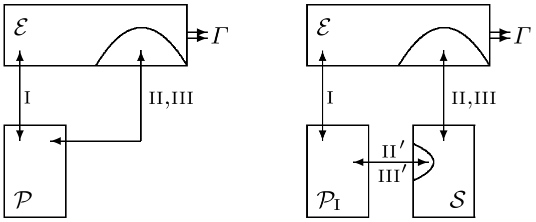

Fig. 2. The game defining the composable security definition. The curved region in $\mathcal{E}$ represents the adversaries against $\mathcal{P}$, and the curved region in $\mathcal{S}$ represents the adversaries against $\mathcal{P}_{\rm I}$. We label the types of interactions as described in the text.

With a slight abuse of language, the symbols $\mathcal{P}$ and $\mathcal{P}_{\rm I}$+$\mathcal{S}$ are also used to denote the respective events of their being chosen at the beginning of the game. We can now state the universal composable security definition.

Definition 1: $\mathcal{P}$ is said to ϵ-securely realize $\mathcal{P}_{\rm I}$ (shorthand $\mathcal{P}$ ϵ-s.r. $\mathcal{P}_{\rm I}$) if

$$\forall \mathcal{E} \;\; \exists \mathcal{S} \;\; \text{s.t.} \;\; \big| \Pr(\Gamma{=}0|\mathcal{P}) - \Pr(\Gamma{=}0|\mathcal{P}_{\rm I}{+}\mathcal{S}) \big| \leq \epsilon \, . \tag{6}$$

We call ϵ in (6) the *distinguishability-advantage* between $\mathcal{P}$ and $\mathcal{P}_{\rm I}$. This security definition (in the model described) is useful because security of basic composition follows "by definition" [19, 20]. We have the following simple version of a universal composability theorem.

Theorem 1. *Suppose a protocol* $\mathcal{P}$ *calls a subroutine* σ*. If* σ ϵ_σ*-s.r.* $\sigma_{\rm I}$ *and* $\mathcal{P}$+$\sigma_{\rm I}$ $e_{\mathcal{P}}$*-s.r.* $\mathcal{P}_{\rm I}$*, then* $\mathcal{P}$+σ ϵ*-s.r.* $\mathcal{P}_{\rm I}$ *for* $\epsilon \leq \epsilon_{\mathcal{P}}$+$\epsilon_\sigma$*.*

Theorem 1 can be generalized to any arbitrary protocol with a proper *modular structure*. An example of an improper modular structure is one with a security deadlock, in which the securities of two components are interdependent.

Proper modular structures can be characterized as follows. Let $\mathcal{P}+\sigma_1+\sigma_2+\cdots$ be any arbitrary protocol using a number of subprotocols. This can be represented by a 1-level tree, with $\mathcal{P}$ being the parent and $\{\sigma_i\}$ the children. For each σ_i that uses other subprotocols, replace the corresponding node by an appropriate 1-level subtree. This is done recursively, until the highest-level subprotocols (the leaves) call no other subprotocols. These are the primitives. It was proved in [20] that more general modular structures, represented by an acyclic directed graph, can be transformed to a tree. The following composability theorem relates the security of a protocol $\mathcal{P}$ to the security of all the components in the tree.

Theorem 2. *Let $\mathcal{P}$ be a protocol and $T_{\mathcal{P}}$ its associated tree. Let $\mathcal{M}$ be a subprotocol corresponding to any node in $T_{\mathcal{P}}$ with subprotocols $\{\mathcal{N}_i\}_{i=1,\cdots,l}$. Then, if $\mathcal{M}+\mathcal{N}_{1\mathrm{I}}+\cdots+\mathcal{N}_{l\mathrm{I}}$ $\epsilon_{\mathcal{M}}$-s.r. $\mathcal{M}_{\mathrm{I}}$, we have $\mathcal{P}$ ϵ-s.r. $\mathcal{P}_{\mathrm{I}}$ for $\epsilon \leq \sum_{\mathcal{M}} \epsilon_{\mathcal{M}}$.*

Theorem 2 is obtained by recursive use of Theorem 1 and the triangle inequality. The idea is to recursively replace each subprotocol by its ideal functionality, from the highest to the lowest level toward the root. The distinguishability-advantage between $\mathcal{P}$ and $\mathcal{P}_{\mathrm{I}}$ is upper bounded by the sum of all the individual distinguishability-advantages between pairs of protocols before and after each replacement. See Fig. 4 for an example of $T_{\mathcal{P}}$ that describes repeated QKD.

In the next section, we analyze QKD in the composability framework. This is part of our main result and it also illustrates the composability framework.

4 Universal Composable Security Definition of QKD

We first describe a general QKD scheme in the composability framework. Then, we tailor an ideal functionality for KD that resembles QKD. Finally, the universal composable security definition of QKD is restated as a distinguishability criteria.

4.1 QKD in the Game Defining Security

Our discussion relies on the existence of authentication schemes that are universal composable in the quantum setting. Furthermore, the authentication scheme should use a key much shorter than the message to be authenticated (so that QKD indeed expands a key). For example, the scheme in [29] satisfies such conditions (composability is proved in [37]). Let α denote any such authentication scheme and let α_{I} denote ideal authentication. Let $\kappa+\alpha$ denote QKD using authentication scheme α and let κ_{I} denote an ideal KD protocol to be defined. By Theorem 1, we can focus on the security of $\kappa+\alpha_{\mathrm{I}}$, i.e., QKD using perfectly authenticated classical channels. The initial key requirement is embedded in the subroutine α_{I}. In this case, QKD has no input and outputs some bitstrings k_{A}, k_{B} of certain length m to Alice and Bob, with $m = 0$ if and only if QKD is aborted. (We can assume that m is a publicly announced output of QKD.) Eve's view (including both quantum and classical data) is given by the state $\rho_{\mathrm{E},k_\mathrm{A},k_\mathrm{B}}$.

We now turn to the game defining the composable security definition of QKD. Eve is an adversary that is part of the environment $\mathcal{E}$. Following the discussion in Sect. 3, $\mathcal{E}$ will fix an arbitrary strategy. Since there is no input to QKD, the optimal application in $\mathcal{E}$ is simply to receive the output keys from $\kappa{+}\alpha_\mathrm{I}$ or κ_I. $\mathcal{E}$ will also consist of the action of Eve and other circuits that compute Γ. A schematic diagram is given in Fig. 3.

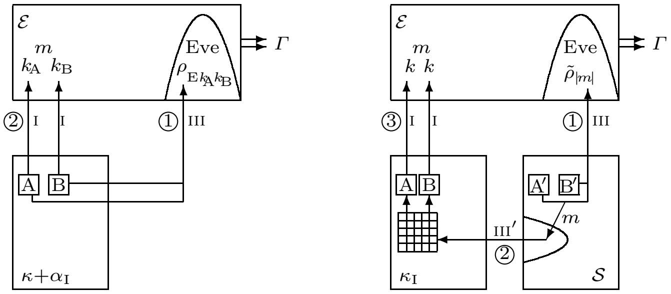

Fig. 3. The game defining the composable security definition of QKD, with our choice of ideal KD and simulator. An ordering of the interactions is given in circles. We also label the types of interactions (see rules 3 and 4 in Sect. 3) explicitly. Upon an input m, the checkered box generates a perfect key of length m to Alice and Bob.

If $\mathcal{E}$ is interacting with $\kappa{+}\alpha_\mathrm{I}$, $\mathcal{E}$ will: (I) receive the output bitstrings k_A, k_B, and $m = |k_\mathrm{A}| = |k_\mathrm{B}|$, and (III) obtain $\rho_{\mathrm{E},k_\mathrm{A},k_\mathrm{B}}$ which depends on Eve's strategy and k_A, k_B. Altogether, $\mathcal{E}$ will be in possession of the state

$$\rho_\mathrm{qkd} = \sum_{k_\mathrm{A},k_\mathrm{B}} \Pr(k_\mathrm{A},k_\mathrm{B})\ |k_\mathrm{A},k_\mathrm{B}\rangle\langle k_\mathrm{A},k_\mathrm{B}| \otimes \rho_{\mathrm{E},k_\mathrm{A},k_\mathrm{B}} \tag{7}$$

in which $\rho_{\mathrm{E},k_\mathrm{A},k_\mathrm{B}}$ and $k_\mathrm{A}, k_\mathrm{B}$ can be correlated. We have omitted an explicit register for m, because the information is redundant given $k_\mathrm{A}, k_\mathrm{B}$.

4.2 Ideal KD and the Simulator

We now define the ideal functionality for QKD. In general, when formulating an ideal functionality, one need not be concerned with how the functionality is realized. What is important is to impose the essential security features while mimicking the analyzed protocol from the point of view of $\mathcal{E}$.

Our ideal KD functionality κ_I has to model both the possibility to generate a perfect key, and the possibility for Eve to cause QKD to be aborted. Besides Alice and Bob, κ_I has a box that accepts a value m from an adversary "Devil" and outputs a perfect m-bit key K to Alice and Bob ($m = 0$ means abort). When κ_I is run, Devil sends m to the box, which sends K to Alice and Bob.

This formulation of $\kappa_{\rm I}$ satisfies the security conditions (3) and (5) perfectly ($\mu_1, \mu_2 = 0$). See Fig. 3 for a schematic diagram.

Consider the following simulator $\mathcal{S}$. $\mathcal{S}$ runs a "fake QKD" with fake Alice$'$ and Bob$'$. They interact with Eve (in $\mathcal{E}$) and run verification procedure as in QKD. A value m is announced for the fake QKD, but the fake output keys are unused and kept secret in $\mathcal{S}$. The Devil in $\mathcal{S}$ then sends m to the box in $\kappa_{\rm I}$, which generates a perfect m-bit key string k to Alice and Bob in $\kappa_{\rm I}$, who forward their outputs to $\mathcal{E}$. Let

$$\tilde{\rho}_m = \sum_{k_{\rm A},k_{\rm B}:|k_{\rm A}|=|k_{\rm B}|=m} \Pr(k_{\rm A}, k_{\rm B}|M{=}m)\ \rho_{{\rm E},k_{\rm A},k_{\rm B}} \ \ . \tag{8}$$

Then, at the end of the game, $\mathcal{E}$ will be in possession of the state

$$\rho_{\rm ideal} = \sum_k \Pr(M{=}|k|)\ 2^{-|k|}\ |k,k\rangle\langle k,k| \otimes \tilde{\rho}_{|k|} \ \ . \tag{9}$$

How $\kappa_{\rm I} + \mathcal{S}$ interacts with $\mathcal{E}$ is summarized in Fig. 3.

4.3 Universal Composable Security Definition and Simple Privacy Condition

Recall that at the beginning of the game, one of κ and $\kappa_{\rm I}{+}\mathcal{S}$ is chosen at random to interact with $\mathcal{E}$. The distinguishability-advantage is upper bounded by the trace distance of the two possible final states of $\mathcal{E}$ right before $\mathit{\Gamma}$ is computed,

$$\big|\Pr(\mathit{\Gamma}{=}0\,|\,\kappa) - \Pr(\mathit{\Gamma}{=}0\,|\,\kappa_{\rm I}{+}\mathcal{S})\,\big| \le \tfrac{1}{2}\, \big\|\, \rho_{\rm qkd} - \rho_{\rm ideal} \big\|_1 \le \tag{10}$$

$$\le \tfrac{1}{2}\, \big\|\, \rho_{\rm qkd} - \rho_{\rm qi1} \big\|_1 + \tfrac{1}{2}\, \big\|\, \rho_{\rm qi1} - \rho_{\rm qi2} \big\|_1 + \tfrac{1}{2}\, \big\|\, \rho_{\rm qi2} - \rho_{\rm ideal} \big\|_1 \ \ , \tag{11}$$

where $\rho_{\rm qi1}$ and $\rho_{\rm qi2}$ are hybrid, intermediate, states between $\rho_{\rm qkd}$ and $\rho_{\rm ideal}$ defined as

$$\rho_{\rm qi1} = \sum_k \Pr(M{=}|k|)\ 2^{-|k|} |k,k\rangle\langle k,k| \otimes \rho_{{\rm E},k,k} \ \ , \tag{12}$$

$$\rho_{\rm qi2} = \sum_k \Pr(M{=}|k|)\ 2^{-|k|}\ |k,k\rangle\langle k,k| \otimes \bar{\rho}_{|k|} \ \ , \tag{13}$$

with $\bar{\rho}_m = \frac{1}{2^m} \sum_{k:|k|=m} \rho_{{\rm E},k,k}$. The sum of the first and the last terms in (11) can be bounded by μ_1 in the equality-and-uniformity condition ((3) in Sect. 2) as follows. Using (7) and (12),

$$\big\|\, \rho_{\rm qkd} - \rho_{\rm qi1} \big\|_1 = \Bigg\|\ \sum_{k_{\rm A}\neq k_{\rm B}} \Pr(k_{\rm A}{,}k_{\rm B})\, |k_{\rm A}{,}k_{\rm B}\rangle\langle k_{\rm A}{,}k_{\rm B}| \otimes \rho_{{\rm E},k_{\rm A},k_{\rm B}} \ + $$

$$+ \ \sum_k \Big[\Pr(k{,}k){-}\Pr(|k|)2^{-|k|}\Big]\, |k{,}k\rangle\langle k{,}k| \otimes \rho_{{\rm E},k,k} \Bigg\|_1 \le \mu_1 \ \ .$$

Using (9) and (13),

$$\left\| \, \rho_{\rm qi2} - \rho_{\rm ideal} \, \right\|_1 \ \leq \ \sum_m \Pr(M{=}m) \, \left\| \, \bar{\rho}_m - \tilde{\rho}_m \right\|_1 \ \leq \ \mu_1$$

where we have used

$$\bar{\rho}_m = \sum_{k_{\rm A},k_{\rm B}} p^{(m)}_{\rm ideal}(k_{\rm A},k_{\rm B}) \, \rho_{{\rm E},k_{\rm A},k_{\rm B}} \ , \quad \tilde{\rho}_m = \sum_{k_{\rm A},k_{\rm B}} p^{(m)}_{\rm qkd}(k_{\rm A},k_{\rm B}) \, \rho_{{\rm E},k_{\rm A},k_{\rm B}} \ , \quad (14)$$

and the equality-and-uniformity condition (3) for the last inequality. The remaining term in the composable security condition (11) is given by

$$\begin{aligned} \frac{1}{2} \left\| \, \rho_{\rm qi1} - \rho_{\rm qi2} \, \right\|_1 &= \frac{1}{2} \left\| \sum_k \Pr(M{=}|k|) \; 2^{-|k|} \; |k{,}k\rangle\langle k{,}k| \otimes \left[\bar{\rho}_{|k|} - \rho_{{\rm E},k,k} \right] \right\|_1 \\ &\leq \frac{1}{2} \sum_k \Pr(M{=}|k|) \; 2^{-|k|} \; \left\| \, \bar{\rho}_{|k|} - \rho_{{\rm E},k,k} \right\|_1 \ , \qquad (15) \end{aligned}$$

which can be interpreted as a new privacy condition.

We have thus compartmentalized the quantity governing the composable security definition for QKD, (10) or (11), into two parts: a term governed by the equality-and-uniformity condition (3) and a new term (15) related to privacy, a bound of which will be called a "composable privacy condition" for QKD. Once (15) is bounded by some μ_2^*, QKD using ideal authentication $\kappa{+}\alpha_{\rm I}$ ϵ_κ-securely realizes the ideal KD $\kappa_{\rm I}$, if $\mu_1 + \mu_2^* \leq \epsilon_\kappa$. Following Theorems 1 and 2, one can use the key "as if it were perfect." Proving such a bound on (15) is relatively straightforward, as compared to a direct proof of the security of using a slightly imperfect key from QKD (without the composability theorem).

In the following section, we prove several bounds for (15). First, we show that for any QKD scheme satisfying the usual privacy condition (5), (15) can be bounded as well, albeit with a potentially large but manageable degradation. Second, we prove a tighter bound on (15) assuming a privacy condition in terms of Eve's Holevo information on the key. Finally, we propose a new, tight, sufficient condition for bounding (10) (the full composable security condition) based on the singlet-fidelity considered in most existing security proofs for QKD. This bypasses (5) and incorporates all of equality, uniformity, and privacy. As an application, we obtain sharp upper bounds for (10) for existing QKD schemes.

5 Universal Composability of QKD

We state some composable security results of QKD; proofs can be found in [38].

Usual privacy condition implies composable privacy condition. Given the usual privacy condition (5), $\sum_{m=0} \Pr(m) \times I(K_{\rm E} \!:\! K \,|\, M = m\,) \ \leq \ \mu_2$, the following bound for (15) holds, ensuring composable privacy:

$$\left\| \rho_{\rm qi1} - \rho_{\rm qi2} \, \right\|_1 \leq 2^{\max(m)/2+1} \sqrt{\mu_2} \ . \qquad (16)$$

Typically, $\max(m)$ is a small fraction of n, the security parameter such as the number of qubits communicated. Since $\mu_2 \in \mathcal{V}$, the set of exponentially decaying functions of n, bounding the key rate m/n ensures the above is in $\mathcal{V}$ also.

Small Holevo information implies composable privacy. Suppose, instead of the usual privacy condition (5) in terms of the accessible information, we have

- *Privacy:* $\exists \mu_2' \in \mathcal{V}$ s.t.

$$\sum_m \Pr(M{=}m) \times \chi(\mathcal{F}_m) \ \leq \ \mu_2' \tag{17}$$

where χ is the Holevo information [39], and $\mathcal{F}_m$ is the ensemble $\{2^{-m}, \rho_{\mathrm{E},k,k}\}_{|k|=m}$. Equation (17) is more stringent than (5) since the Holevo information is an upper bound for the accessible information. In fact, (17) implies

$$\left\|\rho_{\mathrm{qi1}} - \rho_{\mathrm{qi2}}\right\|_1 \leq \sqrt{2\,(\ln 2)\;\mu_2'} \tag{18}$$

which does not have an overhead exponential in the length of the key generated.

A new sufficient condition for composable security. We can easily analyze the composable security of any QKD scheme that has a security proof based on entanglement purification protocol. All existing QKD schemes have such security proofs. The final keys K_A, K_B are outcomes of Alice and Bob's measurements on a shared state ρ_{AB}^m for some m, and ρ_{AB}^m is supposed to be $\varPhi^{\otimes m}$ in the absence of eavesdropping. Here, m is again the key length and $\varPhi = \frac{1}{2}(|00\rangle + |11\rangle)(\langle 00| + \langle 11|)$. The usual privacy condition (5) is often obtained by showing the following.

- *High fidelity:* $\exists \mu_2'' \in \mathcal{V}$ s.t.

$$\sum_m \Pr(m) \left[\, 1 - F(\rho_{\mathrm{AB}}^m, \varPhi^{\otimes m}) \,\right] \leq \mu_2'' \tag{19}$$

(See Sect. 1 for the definition of F.) The above turns out implying a sharp bound on (15):

$$\frac{1}{2} \left\|\, \rho_{\mathrm{qkd}} - \rho_{\mathrm{ideal}} \,\right\|_1 \leq \sqrt{\mu_2''}\ . \tag{20}$$

Equation (19) is thus a good new sufficient condition for *composable security*, being part of the standard QKD proof and implying a tight bound on (10) simultaneously. It also implies *both* equality-and-uniformity and privacy (unlike a bound on Holevo information or mutual information which only implies the composable privacy condition).

6 Discussions and Applications

We have motivated this work with a discussion of the potential gap between the desired security of using a key generated by QKD and the security promised by the privacy condition (5) used in many previous studies of "unconditional security" of QKD. Then, we apply the universal composability theorem to obtain a new security condition that will guarantee the security of using a key generated from QKD. We propose a new composable privacy condition based on bounding (15), and we propose useful sufficient conditions such as bounds on (17) or (19). Most interesting of all, we show that a bound on the singlet-fidelity (19) directly implies the composable security condition (a bound on (10)). These are our main contributions (in the context of cryptography).

We also provide a proof that the existing privacy condition (5) does imply composable security (a bound on (15)) though the bound degrades exponentially in the key size. Despite the existence of such connections, we emphasize that future security proofs should bound (10), (15), (17), or (19) directly. We also provide a sharp bound on (15) based on Holevo's information (17) or singlet-fidelity (19). We show that most existing security proofs for QKD imply sharp bounds on (10), when bypassing the usual privacy condition (5). Outside the context of cryptography, these connections between various privacy conditions can be useful for the study of correlations in quantum systems.

The pathologies of the accessible information exhibited recently [16, 17] suggest a conjecture that, when going from (5) to (15), the degradation of the security parameter exponential in the key size is necessary.

As a final application, we analyze the security of repeating QKD t times, without assuming the availability of an authenticated classical channel. (Note that t is a fixed parameter that does not grow with the problem size.) Each run of QKD κ calls a composable authentication scheme α as a subroutine, and each run of α requires a composably secure key, which is provided by the previous round of κ (as a subroutine to α). Call the t rounds of QKD our protocol $\mathcal{P}$. The associated tree for $\mathcal{P}$, and the ideal realization $\mathcal{P}_{\mathrm{I}}$ are given in the far left and right of Fig. 4.

If $\kappa{+}\alpha_{\mathrm{I}}$ ϵ_κ-s.r. κ_{I} (as in (10)) and if $\alpha{+}\kappa_{\mathrm{I}}$ ϵ_α-s.r. α_{I}, $\mathcal{P}$ $t(\epsilon_\kappa{+}\epsilon_\alpha)$-s.r. $\mathcal{P}_{\mathrm{I}}$. In other words, each extra around of QKD degrades the overall distinguishability-advantage by an additive constant $(\epsilon_\kappa + \epsilon_\alpha)$. The same result can be obtained by using Theorem 2, or conversely, this simple exercise illustrates the idea behind Theorem 2.

Acknowledgement We thank Charles Bennett, Daniel Gottesman, Aram Harrow, and John Smolin for interesting discussions on the security concerns of using a key obtained from QKD. We also thank Dominique Unruh and Jörn Müller-Quade for interesting discussions on their alternative framework of composability.

Part of this work was completed while MH and JO were visiting the MSRI program on quantum information, Berkeley, 2002. MB acknowledges the support of the Israel Science Foundation and a research grant from the Israeli Ministry of Defense. MH is supported by EU grants RESQ (IST-2001-37559) and QUPRODIS (IST-2001-38877). DL acknowledges the support from the Tolman Foundation and the Croucher Foundation. DL and DM acknowledge support

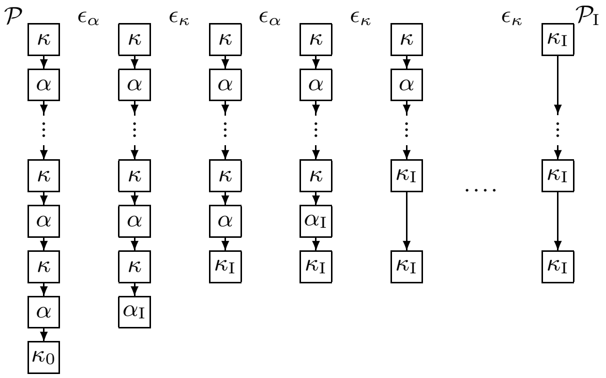

Fig. 4. Associated tree for t rounds of κ in the left. κ_0 represents some initially shared key. The arrows point from parents to children. Each tree to the right is obtained by replacing one node by its ideal functionality. The distinguishability-advantage of each pair of consecutive schemes is marked between their trees near the roots. Authentication is omitted in the ideal functionality $\mathcal{P}_{\rm I}$.

from the US NSF under grant no. EIA-0086038. JO is supported by an EU grant PROSECCO (IST-2001-39227) and a grant from the Cambridge-MIT Institute.

A Notations

We gather notations frequently used in the paper, roughly in order of first appearance:

- KD: key distribution
- QKD: quantum key distribution
- Alice and Bob: two honest parties trying to establish a common key
- Eve: an active adversary
- A, B, E: subscripts labelling objects related to Alice, Bob, and Eve
  $\mathbb{A}$, $\mathbb{B}$, $\mathbb{E}$: labels of their respective quantum systems
- **Capitalized letters denote random variables and the corresponding uncapitalized letters denote particular outcomes**
- $K_{\rm A}$, $k_{\rm A}$, $K_{\rm B}$, $k_{\rm B}$: output keys for Alice and Bob
- K, k: $k := k_{\rm A}$ when $k_{\rm A} = k_{\rm B}$
- M, m: length of key generated by QKD, with $M = 0$ iff QKD is aborted
- $K_{\rm E}$, $k_{\rm E}$: classical data possibly extracted by Eve at the end of QKD by measuring her quantum state
- $\Pr(\cdot)$: probability of the event "$\cdot$"
- log: logarithm in base 2
- For random variables X, Y, Z:
  $H(X) := -\sum_x \Pr(x) \log \Pr(x)$ is the entropy of X
  $I(X{:}Y) := H(X) + H(Y) - H(XY)$ is the mutual information between X, Y

$I(X{:}Y|Z{=}z)$ is the mutual information between X, Y conditioned on $Z{=}z$
$I(X{:}Y|Z) := \sum_z \Pr(z) I(X{:}Y|Z{=}z)$ is the conditional mutual information

- ρ: generic symbol for a density matrix
- $|\cdot\rangle$: a vector in a Hilbert space, with label "$\cdot$"
  $|\cdot\rangle\langle\cdot|$: the projector onto the subspace spanned by $|\cdot\rangle$, also known as "outer-product" of $|\cdot\rangle$ and $\langle\cdot|$
- $\mathrm{Tr}(\cdot)$: the trace
- $\mathrm{Tr}_{\mathbb{H}_1}(\cdot)$: the partial trace over the system $\mathbb{H}_1$. Let ρ_{12} be the density matrix for a joint state on $\mathbb{H}_1$ and $\mathbb{H}_2$. $\mathrm{Tr}_{\mathbb{H}_1}(\rho_{12})$ is the state after $\mathbb{H}_1$ is discarded.
- $\|\cdot\|_1$: the trace distance, which can be taken as the sum of the singular values
- F: the fidelity. For two states ρ_1, ρ_2 in H, $F(\rho_1, \rho_2) = \max_{|\psi_1\rangle, |\psi_2\rangle} |\langle\psi_1|\psi_2\rangle|^2$ where $|\psi_{1,2}\rangle \in \mathbb{H}{\otimes}\mathbb{H}'$ are "purifications" of $\rho_{1,2}$ (i.e., $\mathrm{Tr}_{\mathbb{H}'}|\psi_{1,2}\rangle\langle\psi_{1,2}| = \rho_{1,2}$), and $\langle\cdot|\cdot\rangle$ is the inner product. Here, we can take $\dim(\mathbb{H}') = \dim(\mathbb{H})$.
- $\rho_{\mathrm{E},k_\mathrm{A},k_\mathrm{B}}$: Eve's view (both quantum and classical data) when the key outputs to Alice and Bob are $k_\mathrm{A}, k_\mathrm{B}$.
- n: security parameter such as the number of qubits communicated in QKD
- $p^{(m)}_{\mathrm{qkd}}$: $p^{(m)}_{\mathrm{qkd}}(k_\mathrm{A}, k_\mathrm{B}) = \Pr(K_\mathrm{A}{=}k_\mathrm{A}, K_\mathrm{B}{=}k_\mathrm{B}|M{=}m)$, i.e., the distribution of $K_\mathrm{A}, K_\mathrm{B}$ generated in QKD conditioned on $|K_\mathrm{A}| = |K_\mathrm{B}| = m$
- $p^{(m)}_{\mathrm{ideal}}$: the distribution over two m-bit strings l, l' defined as
  $p^{(m)}_{\mathrm{ideal}}(l, l') = 0$ if $l \neq l'$, $p^{(m)}_{\mathrm{ideal}}(l, l) = 2^{-m}$.
- $\mathcal{V}$: the set of exponentially decaying functions of n
- σ, $\mathcal{P}$, σ_I, $\mathcal{P}_\mathrm{I}$: σ and $\mathcal{P}$ are generic labels for protocols, with σ possibly used as a subroutine. The symbol of a protocol with a subscript I denotes the ideal functionality of the protocol. $\mathcal{P}{+}\sigma$: a protocol $\mathcal{P}$ calling a subroutine σ.
- $\mathcal{E}$, $\mathcal{S}$: the environment and the simulator. These are sets of registers and operations and they are sometimes personified in our discussion.
- $\mathit{\Gamma}$: output bit of $\mathcal{E}$
- ϵ-s.r. : $\mathcal{P}$ ϵ-s.r. $\mathcal{P}_\mathrm{I}$ is a shorthand for $\mathcal{P}$ ϵ-securely realizes $\mathcal{P}_\mathrm{I}$ (see mathematical definition in (6)). ϵ is called the *distinguishability-advantage* between $\mathcal{P}$ and $\mathcal{P}_\mathrm{I}$.
- $T_\mathcal{P}$: the associated tree for a protocol $\mathcal{P}$
- α, α_I: universal composable authentication with negligible key requirement and its ideal functionality
- $\kappa{+}\alpha$, $\kappa{+}\alpha_\mathrm{I}$, κ_I: QKD using authentication α, QKD using ideal authentication α_I, and ideal KD defined in Sect. 4.2
- Devil: an adversary that determines the key length m generated by κ_I
- ρ_{qkd}: state possessed by $\mathcal{E}$ after interacting with $\kappa{+}\alpha_\mathrm{I}$, see (7)
- ρ_{ideal}: state possessed by $\mathcal{E}$ after interacting with κ_I, see (9)
- ρ_{qi1}, ρ_{qi2}: hybrid, intermediate, states between ρ_{qkd} and ρ_{ideal}, see (12), (13)
- $\tilde{\rho}_m$: Eve's state when $M = m$, averaged over K_A, K_B. See (8)
- $\bar{\rho}_m$: uniform average of $\rho_{\mathrm{E},k,k}$ for $|k| = m$. See the line right after (13)
- **Ensemble $\{q_x, \varrho_x\}_x$: a distribution $\{q_x\}_x$ of quantum states ϱ_x**
- I_{acc}: accessible information of an ensemble $\{q_x, \varrho_x\}_x$, i.e., the maximum mutual information between X and outcome Y obtained from measuring a specimen ϱ_x
- $\mathcal{F}_m$: the ensemble $\{2^{-m}, \rho_{\mathrm{E},k,k}\}_{|k|=m}$

– $\chi(\{q_x, \varrho_x\})$: Holevo information of the ensemble $\{q_x, \varrho_x\}$, given by $S(\sum_x q_x \varrho_x) - \sum_x q_x S(\varrho_x)$ where $S(\cdot)$=$\mathrm{Tr}(\cdot \log(\cdot))$ is the von Neumann entropy
– ρ^m_{AB}: state on which measurements by Alice and Bob output K_{A}, K_{B} in QKD-security-proofs based on entanglement purification
– Φ: a perfect EPR pair $\frac{1}{2}(|00\rangle + |11\rangle)(\langle 00| + \langle 11|)$
– Singlet fidelity: $F(\rho^m_{\mathrm{AB}}, \Phi^{\otimes m})$. Note that "singlet" usually refers to a state that is only unitarily equivalent to Φ, but we borrow the term in this paper.

References

1. C. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of IEEE International Conference on Computers, Systems and Signal Processing*, pages 175–179, New York, 1984. IEEE. Bangalore, India, December 1984.
2. A. Ekert. Quantum cryptography based on Bell's theorem. *Phys. Rev. Lett.*, 67(6):661–663, 1991.
3. C. Bennett. Quantum cryptography using any two nonorthogonal states. *Phys. Rev. Lett.*, 68(21):3121–3124, 1992.
4. C. Bennett, G. Brassard, R. Jozsa, D. Mayers, A. Peres, B. Schumacher, and W. Wootters. Reduction of quantum entropy by reversible extraction of classical information. *Journal of Modern Optics*, 41(12):2307–2314, 1994.
5. D. Mayers. Quantum key distribution and string oblivious transfer in noisy channels. In *Advances in Cryptography–Proceedings of Crypto'96*, pages 343–357, New York, 1996. Springer-Verlag.
6. D. Mayers. Unconditional security in quantum cryptography. *J. Assoc. Comp. Mach*, 48:351, 2001. quant-ph/9802025.
7. H.-K. Lo and H. F. Chau. Unconditional security of quantum key distribution over arbitrarily long distances. *Science*, 283:2050–2056, 1999. quant-ph/9803006.
8. E. Biham, M. Boyer, P. Boykin, T. Mor, and V. Roychowdhury. A proof of the security of quantum key distribution. In *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 715–724, New York, 2000. ACM. quant-ph/9912053.
9. P. Shor and J. Preskill. Simple proof of security of the bb84 quantum key distribution protocol. *Phys. Rev. Lett.*, 85:441–444, 2000. quant-ph/0003004.
10. K. Tamaki, M. Koashi, and N. Imoto. Unconditionally secure key distribution based on two nonorthogonal states. *Phys. Rev. Lett.*, 90:167904, 2003. quant-ph/0212162.
11. G. Gottesman and H.-K. Lo. Proof of security of quantum key distribution with two-way classical communications. *IEEE Transactions on Information Theory*, 49(2):457–475, 2003. quant-ph/0105121.
12. C. Bennett and J. Smolin first suggested the key degradation problem to one of us, and A. Harrow has obtained partial results.
13. A. Ambainis, M. Mosca, A. Tapp, and R. de Wolf. Private quantum channels. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 547–553, 2000. quant-ph/0003101.
14. P. Boykin and V. Roychowdhury. Optimal encryption of quantum bits. quant-ph/0003059.
15. A. Peres and W. Wootters. Optimal detection of quantum information. *Phys. Rev. Lett.*, 66:1119–1122, 1991.

16. D. DiVincenzo, M. Horodecki, D. Leung, J. Smolin, and B. Terhal. Locking classical correlation in quantum states. *Phys. Rev. Lett.*, 92:067902, 2004. quant-ph/0303088.
17. P. Hayden, D. Leung, P. Shor, and A. Winter. Randomizing quantum state: constructions and applications. quant-ph/0307104.
18. R. Canetti. Universal composable security: A new paradigm for cryptographic protocols. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 136–145. IEEE, 2001.
19. M. Ben-Or and D. Mayers. Composability theorem. Part I of presentation by D. Mayers, QIP 2003, MSRI, Berkeley. See http://www.msri.org/publications/ln/msri/2002/qip/mayers/1/index.html .
20. M. Ben-Or and D. Mayers. Composing quantum and classical protocols. quant-ph/0409062.
21. M. Backes, B. Pfitzmann, and M. Waidner. A general composition theorem for secure reactive systems. In *First Theory of Cryptography Conference (TCC)*, pages 336–354, 2004.
22. D. Unruh. Relating formal security for classical and quantum protocols. Presentation at the Special week on Quantum crytography, Isaac Newton Institute for Mathematical Sciecnes, September 2004. Available at http://www.unruh.de/DniQ/publications.
23. D. Unruh. Simulation security for quantum protocols. quant-ph/0409125.
24. M. Ben-Or, M. Horodecki, D. Leung, D. Mayers, and J. Oppenheim. Composability of QKD. Part II of presentation by D. Mayers, QIP 2003, MSRI, Berkeley. See http://www.msri.org/publi-cations/ln/msri/2002/qip/mayers/1/index.html .
25. M. Ben-Or, M. Horodecki, D. Leung, D. Mayers, and J. Oppenheim. Composability of quantum proocols - applications to quantum key distribution and quantum authentication. Part II of presentation by D. Leung, QIP 2004, IQC, University of Waterloo. See http://www.iqc.ca/conferences/qip/presentations/leung-.pdf.
26. R. Renner and Konig. Universally composable privacy amplification against quantum adversaries. quant-ph/0403133.
27. M. Christandl, R. Renner, and A. Ekert. A generic security proof for quantum key distribution. quant-ph/0402131.
28. M. Nielsen and I. Chuang. *Quantum computation and quantum information.* Cambridge University Press, Cambridge, U.K., 2000.
29. M. Wegman and J. Carter. New hash functions and their use in authentication and set equality. *Journal of Computer and System Sciences*, 22:265–279, 1981.
30. D. Bruss. Optimal eavesdropping in quantum cryptography with six states. *Phys. Rev. Lett.*, 81:3018–3021, 1998.
31. H. Bechmann-Pasquinucci and N. Gisin. Incoherent and coherent eavesdropping in the 6-state protocol of quantum cryptography. quant-ph/9807041.
32. D. Deutsch, A. Ekert, R. Jozsa, C. Macchiavello, S. Popescu, and A. Sanpera. Quantum privacy amplification and the security of quantum cryptography over noisy channels. *Phys. Rev. Lett.*, 77:2818, 1996. quant-ph/9604039.
33. T. Cover and J. Thomas. *Elements of Information Theory.* John Wiley and Sons, New York, 1991.
34. A. Yao. Quantum circuit complexity. *Proc. of the 34th Ann. IEEE Symp. on Foundations of Computer Science*, pages 352–361, 1993.
35. D. Aharonov, A. Kitaev, and N. Nisan. Quantum circuits with mixed states. quant-ph/9806029.

36. An acyclic circuit is a partially ordered set of gates. However, associating the circuit with constraints on the timing of the adversarial attack is a delicate issue. Suppose the circuit contains conditional gates controlled by random public classical registers. The gates on the target may or may not be applied depending on the values of the control registers. When the gates are not applied, the associated time-constraints of the adversarial attack disappear. In the extension to the usual acyclic circuit model, we consider all possible values of the control registers and the resulting sets of *nontrivial* partially ordered operations, and the corresponding constraints on the adversarial attack.
37. P. Hayden, D. Leung, and D. Mayers. On the composability of quantum message authentication and key recycling.
38. M. Ben-Or, M. Horodecki, D. Leung, D. Mayers, and J. Oppenheim. The universal composable security of quantum key distribution. quant-ph/0409078.
39. A. Holevo. Information-theoretical aspects of quantum measurement. *Problemy Peredachi Informatsii*, 9(2):31–42, 1973. [A. S. Kholevo, *Problems of Information Transmission*, vol. 9, pp. 110-118 (1973)].