On Deniability in Quantum Key Exchange

Donald Beaver

Syntechnica, LLC

Abstract. We show that claims of "perfect security" for keys produced by quantum key exchange (QKE) are limited to "privacy" and "integrity." Unlike a one-time pad, QKE does not necessarily enable Sender and Receiver to pretend later to have established a different key. This result is puzzling in light of Mayers' "No-Go" theorem showing the impossibility of quantum bit commitment. But even though a simple and intuitive application of Mayers' protocol transformation appears sufficient to provide deniability (else QBC would be possible), we show several reasons why such conclusions are ill-founded. Mayers' transformation arguments, while sound for QBC, are insufficient to establish deniability in QKE.

Having shed light on several unadvertised pitfalls, we then provide a candidate deniable QKE protocol. This itself indicates further shortfalls in current proof techniques, including reductions that preserve privacy but fail to preserve deniability. In sum, purchasing undeniability with an off-the-shelf QKE protocol is significantly more expensive and dangerous than the mere optic fiber for which "perfect security" is advertised.

1 Introduction

Privacy and integrity are the cornerstones of security. But a third, more subtle property is often overlooked: deniability, or the ability to pretend, after sending a message, that a different message was sent (perhaps by pretending a different key was used). The ability to deny a message is important in settings such as voting (to inhibit selling or coercion) and free, private speech.

A one-time pad is private, supports integrity, and easily provides deniability: after actually sending $c = m_0 \oplus k$, one can pretend that the cleartext was m_1 by pretending the key was $k' = k \oplus m_0 \oplus m_1$. But OTP's cannot be generated from scratch, and the length of an OTP must be at least as long as the total cleartext. Otherwise, the key equivocation – Shannon's pioneering measure of information-theoretic security – will not be sufficient, and information will be leaked, ultimately limiting the range of alternate, fake keys.

Private and public key cryptography address these issues by ensuring that finding the key or cleartext is computationally difficult, under widely-accepted complexity assumptions. But it is amply clear that they provide no key equivocation whatsoever. Even though it is difficult to find m from $m^e \bmod n$, or $g^{ab} \bmod p$ from g^a and g^b, it is obvious that there is a unique solution in each case. (Pretending that $m' \neq m$ was sent is impossible, since $(m')^e \not\equiv m^e$.) Moreover, the mere fact that the keys are used to encrypt long messages makes it immediately obvious that, from an information-theoretic viewpoint, equivocation is limited.

With limited equivocation, and with the obviously unique mathematical solutions behind m^e or g^{ab}, it is unsurprising that RSA and Diffie-Hellmann key exchange are *undeniable*. There is no false message or key that can possibly match the public record, even though finding the real message or key might be difficult.

1.1 Perfect, Unconditional Security

Along comes quantum key exchange (QKE), offering to establish a key with "perfect, unconditional security." How is this possible in a world where each party has unlimited computing power? Unlike the classical information-theoretic world, determining the precise state of a particle prepared in an unknown fashion is generally impossible. Thus, one can obtain asymmetries in knowledge that are not achievable in classical settings.

Wiesner pioneered the cryptographic application of this principle in a proposal to authenticate money, and Bennett and Brassard showed the first key exchange protocol based on it [Wi83,BB84]. The canonical example is conjugate coding: a polarized photon represents a bit in one of two ways: either using a $+$ basis where 0° indicates 0 and 1° indicates 1, or using a $\times$ basis where 45° indicates 0 and 135° indicates 1. Knowing the right basis, it is easy to discover the bit. Without knowing the basis, any attempted observation is likely to extract possibly-incorrect information at the cost of irreversibly leaving the photon in a changed state. (This is not an *engineering* issue but a fundamental corollary of physics.)

Thus it is possible to detect attempted eavesdropping, unlike the classical world, where complete information about a given transmission can be obtained by direct inspection (in ideal principle). Moreover, if the level of eavesdropping is low, then the extracted information is low, ultimately enabling successful key exchange (and quantum money). Best of all, this can be done "from scratch," assuming that a Sender and Receiver also have a classical, authenticated public line.

A series of papers show that QKE is "perfectly, unconditionally secure." Want to avoid Shannon's key equivocation bounds? Just generate more key. Since there are no computational issues, this looks like a convenient way (engineering aside) to create a OTP of unlimited length.

The current work offers a strong note of caution: advertised "perfect, unconditional security" is not the same as equivalence to a OTP. In particular, while privacy and integrity have been provably established for QKE, deniability is not covered, nor is it implied.

1.2 How to Bind the Message?

Imagine that Eve measures only one photon. With probability $1/2$, she chooses the correct basis, obtains complete information on that photon, and transmits an unchanged photon to R. (Or, she uses an incorrect basis, thereby "disturbing" the photon, but R's later measurement coincidentally restores it.) No secrecy is compromised; no theorems are violated.

As for *deniability*, however, things are different. If S and R later try to open up their accounting records to show that a different key was established, then they must change *something* in their actual records. There is some nonzero chance that they decide to pretend a different bit was used for the one photon Eve measured. With significantly nonzero probability, the false record they provide will not match Eve's observation.

1.3 Quantum Subtleties

The issue is clouded by related results in quantum cryptography. Although it was thought that the asymmetries in knowledge might also enable bit commitment,[1] Mayers showed that quantum bit commitment (QBC) is in fact impossible. One subtle insight is that the programming command "(step n) Party P measures particle X privately" is not enforceable against an adversary. Intuitively speaking, a cheater can postpone certain required measurements (or more general "collapses" of quantum systems), thereby keeping her options open.

At face value, this serves only to add optimism to the QKE setting, where one might now happily conclude that QKE is fully deniable, since otherwise it would enable QBC. Although we will extract inspiration from Mayers' insightful work, we also show that such off-the-shelf conclusions are logically unfounded.

It is easy to invoke Mayers' no-commit theorem blindly to (1) dismiss the significance of a positive result ("it follows easily") or (2) dismiss correctness of negative results ("it contradicts no-commitment"). Ordinarily, one would imagine that a demonstration that BB84 is binding would be sufficient *prima facie* to show that the no-commit theorem does not apply.

But nothing quantum is *prima facie*. Instead, without deeper investigation, many have been tempted to challenge the "counterexample" (namely, the assertion of undeniability). Our deeper investigation displays why Mayers' result is true against QBC but insufficient against QKE. In fact, it is somewhat surprising that deniability might in fact be achievable through LOCC (local operations and classical communication), since the no-commitment proof demands potentially nonlocal operations on S and R.

1.4 Contributions

Our work is directed at (1) making the properties of QKE fully apparent for existing and newly-designed protocols, and (2) analyzing the extent to which existing techniques (such as Mayers' methods for QBC) apply to QKE, and extending them where needed. In sum:

[1] In bit commitment, a bit to be committed must be kept secret from the "receiving" party, at least until much later, when it is to be "decommitted/unveiled.".

- The merely "perfectly, unconditionally secure" key established through quantum key exchange is not completely equivalent to a one-time-pad.
- Perfect privacy and integrity do not imply deniability.
- Quantum protocols are not necessarily deniable, even if "perfectly secure."
- The BB84 quantum protocol is binding.
- Mayers' "no-commitment" theorem is not sufficient to imply deniability.
- Deniability can be achieved through extensions that require a quantum computer.

Our positive result balances the need for purification with the need to use the public, authenticated, classical channel of QKE. It represents the first *deniable* quantum key exchange protocol.

2 Background, Notation, Definitions

We employ standard terminology from quantum computing and cryptography. Let $|0\rangle$ and $|1\rangle$ denote an orthonormal basis for a two-dimensional complex Hilbert space $\mathcal{H}_2$. Using the Dirac bra-ket notation, $\langle\phi| \equiv_{def} |\phi\rangle^{\dagger}$. The Pauli matrices are

$$X = \tfrac{1}{2}\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \; Y = \tfrac{1}{2}\begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \; Z = \tfrac{1}{2}\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

Conjugate coding uses two bases: $B_+ = \{|0\rangle, |1\rangle\}$ and $B_\times = \{(1/\sqrt{2})(|0\rangle + |1\rangle), (1/\sqrt{2})(|0\rangle - |1\rangle)\}$. (N.b.: subscript "+" is interchangeable with "0", and subscript "×" with "1.") The Bell basis describes entanglement:

$$\beta_{00} = \tfrac{1}{\sqrt{2}}(|00\rangle + |11\rangle)\ , \ \ \beta_{10} = \tfrac{1}{\sqrt{2}}(|00\rangle - |11\rangle)$$
$$\beta_{01} = \tfrac{1}{\sqrt{2}}(|01\rangle + |10\rangle)\ , \ \ \beta_{11} = \tfrac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$$

A state can be represented as a *density matrix* over $\mathcal{H}_n$ for some n. A density matrix ρ is a weighted sum of projectors, with $\mathrm{Tr}(\rho) = 1$. A density matrix can represent a mixed state or, in the case that $\rho = |\phi\rangle\langle\phi|$, a pure state ϕ. We typically consider "binary" Hilbert spaces, expressible as $(\mathcal{H}_2)^{\otimes n}$.

Let A be Hermitian. Traditionally, a measurement of A is seen as "collapsing" the state ρ to an eigenvector of A. The expected value of A will be $\langle A\rangle = \mathrm{Tr}(\rho A)$. We allow parties to perform generalized measurements through the standard toolkit: (1) appending an unentangled ancillary subsystem; (2) applying a unitary transformation; (3) making an orthogonal measurement; (4) tracing out a local part of the system.[2]

Let C_1 and C_2 be $[n, k_1]$ and $[n, k_2]$ binary codes, respectively, with

$$\{0\} \subset C_2 \subset C_1 \subset \mathrm{GF}(2)^n.$$

[2] This is tantamount to discarding part of the system. In this paper, we accept this approach at face value.

The $\mathrm{CSS}_{x,z}$ quantum encoding [CSS96] maps $v \in C_1$ to the following codeword:

$$v \rightarrow \frac{1}{\sqrt{|C_2|}} \sum_{w \in C_2} (-1)^{z \cdot w} |x + v + w\rangle.$$

In the context of BB84, phase errors turn out to be irrelevant and the straightforward protocol purification of BB84 (see *infra*) is more like $\mathrm{CSS}_{x,z}$ with z omitted or averaged out.

2.1 Protocol Execution

We use a circuit model for protocol execution. A global state Φ, described over a basis B_+^k for some k, is advanced through applying each party P (*i.e.* each circuit) to a collection of registers, where a register is a local subset of the "wires" of the overall circuit. We use superscripts to indicate location of given registers; thus, *e.g.*, $\Phi = \sum_{v_1..v_5} \alpha_{v_1..v_5} |v_1 v_2 v_3\rangle^A \otimes |v_4 v_5\rangle^B$ describes a state with Alice holding the first three registers/wires and Bob holding the rest. The tensor product sign $\otimes$ is omitted when clear from context. We also use the shorthand $\mathcal{H}_A \otimes \mathcal{H}_B$ to express the state space.

A transition of the system thus consists of applying a local unitary transformation $U_P \otimes \mathbf{1}$ to the registers held by party P, along with a possible orthogonal measurement. More precisely, U_P applies to the Hilbert subspace $\mathcal{H}_2^m$ at indices $k_1, \ldots, k_m$ that are labelled as under P's control. Subsequent communication is modelled by reassignments of those labels.

Initialization. We say that a protocol is *properly initialized* if each party starts with quantum registers in unentangled states $|0\rangle$. In particular, there is no entanglement with other parties, nor with the auxiliary-parties that comprise the "environment."

Communication. As noted, communication is generally just a reassignment of the register labels. When an eavesdropper is present, the register is assigned to the eavesdropper first before being reassigned to the destination party. Generally, the eavesdropper can forward any transformation or substitution she pleases.

This suffices to model QBC, but in QKE there is an additional channel: the reliable public classical channel. One way to regard this channel is as a separate party who first measures the input, then broadcasts the result to the source, destination, and eavesdropper (who is not allowed to alter the result). This irrevocable measurement induces a mixture among several outcomes of the overall protocol.

The collection of such auxiliary parties consists the "environment" and is described by an environment space $\mathcal{H}_{env}$. Thus a two-party protocol is executed over $\mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_{env}$.

Views and Outputs. The distribution produced by running a quantum protocol is obtained by tracing out each party. The *view* of party P includes her state in $\mathcal{H}_P$ along with the classical strings in any environmental/classical channels she used or saw. In particular, direct *erasure* of classical information is not allowed. For some purposes, we may more generally regard the *view* as the collection of registers along with the measurement history of a party, rather than tracing out the final state.

Parameters. We generally use κ to denote a security parameter, k to denote a key (generated or exchanged), and we sometimes overload k when describing error correction: $[n, k]$ connotes an error correcting code mapping k logical bits to n representation bits.

2.2 Deniability

There are several variants on the meaning of *deniability* [CDNO97,Be96] and *binding*, depending on the parties attempting to equivocate and what their success rates may be.

Let m_1 and m_2 be arbitrary messages. Run $S(m_1)$, R and E, obtaining global state $\rho(m_1)$ whose registers are $|\phi_S\rangle$, $|\phi_R\rangle$, $|\phi_E\rangle$, $|\phi_{env}\rangle$. Let D_S and D_R be local computations (not necessarily unitary operations). Let $\rho(m_1, m_2) = \sum |D_S(m_1, m_2, \phi_S)\rangle |D_R(m_1, m_2, \phi_R)\rangle |\phi_E\rangle |\phi_{env}\rangle$, representing an attempt at denial: pretending that m_2 was really sent.

Let J be a judge, who has inputs for registers ϕ_S, ϕ_R, ϕ_E, and ϕ_{env}. J's final state is described in registers d and J', where d is a single-bit, "decision" register. Flip a coin c in the environment to determine whether denial will be attempted. If $c = 0$, run J on $\rho(m_1)$; if $c = 1$, run J on $\rho(m_1, m_2)$ (namely apply D_R and D_S before submitting to the judge). The final result is of the form

$$\rho(m_1, m_2, c) = \sum |cd\rangle |\phi_{J'}\rangle |\phi_{env}\rangle.$$

Tracing out J' and the environment gives a mixture over $|cd\rangle$'s.

A judge is *safe* if $|01\rangle$ has zero probability, namely the judge makes no false accusations.

Definition 1. *Let (S, R) be a quantum key exchange protocol with denial programs (D_S, D_R). For eavesdropper E and judge J, let $P_{J,E}(m_1, m_2, \kappa)$ be the probability J gives $|11\rangle$, on security parameter κ. We say the protocol is* deniable *if, for any E, any safe J, and for any m_1, m_2: $P_{J,E}(m_1, m_2, \kappa) = \kappa^{-\omega(1)}$.*

Simplifying, let $P(\kappa)$ be the maximal probability of $|11\rangle$ over all messages of size $O(\kappa)$. A protocol family indexed by integers C is *perilously deniable* if $P(\kappa, C) = O(\kappa^{-C})$, namely S and R can reduce their vulnerability to a small (but non-negligible) polynomial fraction. A protocol is *weakly binding* if $P(\kappa, c) = \Omega(\kappa^{-c})$ for some $c > 0$. (Thus a protocol family can be perilously deniable while each given value of C produces a weakly binding protocol.) A protocol is *binding* if $P(\kappa) = 1 - \kappa^{-\omega(1)}$.

(Variants on this approach include sender-only and receiver-only deniability, single-bit messages, unsafe judges, and many others. Note that when an exchanged key is used like a one-time pad, "key" can often be interchanged with "message" to simplify the discussion.)

BB84

1. S selects $2(4+\delta)n$ random bits $\{b[i,0], b[i,1]\}$.
2. S encodes $\{b[i,0]\}$ as qubits $\{p[i]\}$, each in basis $+$ or $\times$ depending on $\{b[i,1]\}$.
3. S sends $\{p[i]\}$.
4. S chooses random $v_k \in C_1$.
5. E forwards $\{q[i]\}$ to R (possibly unchanged).
6. R measures each $\{q[i]\}$ in random basis $\{c[i]\}$.
7. S announces $\{b[i,1]\}$ on the classical channel.
8. S and R discard indices wherever $b[i,0] \neq c[i]$. S selects and announces a random remaining subset of $2n$ bits, along with a random n-subset π of check indices. (Abort if impossible.)
9. S and R reveal $p[i]$ and $q[i]$ classically for $i \in \pi$ and abort if any (resp. more than t) disagree.
10. S announces $x \oplus v_k$, where x is the n-bit remaining string in $\{b[i,0]\}$.
11. R computes $y \oplus x \oplus v_k$, where y is the n-bit remaining string in $\{q[i]\}$, and applies C_1 to correct it to v_k (presumably).
12. S and R calculate k from the coset $v_k + C_2$.

Fig. 1. BB84 protocol for n-bit key k, in modern conventions, with C_1 used for reconciliation and cosets of C_2 used for privacy amplification.

3 Quantum Key Exchange

The relevant steps of the BB84 protocol, with eavesdropper, are sketched in Fig. 1. While arbitrary hashing and privacy amplification techniques can be used with the basic BB84 approach, we have illustrated the typical approach of employing binary codes. More particularly, we follow the conventions of [ShPr01] so that we can connect to related work.

It can be shown that an eavesdropper gains information only with $O(2^{-k})$ probability, or gains at most $O(2^{-k})$ as measured by entropy, where $k = k_1 - k_2$.

3.1 Variants of BB84

In the original [BB84] protocol, R measured the photons before knowing the proper bases; the mistaken bases were discarded. Since photons were difficult to store without measuring, this made the theoretical protocol more feasible. As discussed in [BBM92], however, if S and R simply try to establish EPR pairs, the result is similar to having R postpone her measurements until S announces the bases on the classical channel. Since S will wait until he has confirmed that R has the particles, namely that Eve no longer has a chance to change what she has forwarded, this is "okay." [BBM92] suggest that such a protocol has security equivalent to the original [BB84].

Further degrees of purification are possible. Note that the specific resulting protocol depends on the particular reconciliation and amplification routines.

- (BB84) No entanglement or purification.

- (BB84-EPR/Ekert) Use entangled qubits, then measure; no other purification.
- (PQECC) Use entangled qubits; measure check subset; measure key; leave other registers in superposition; (more details later).
- (BB84-Key) Purify completely according to [Ma96]; measure key.
- (BB84-Pur) Purify completely according to [Ma96].

Apart from PQECC, these variants turn out to be either binding, vacuous (non-transmitting), or or unimplementable within the communication model.

The common instantiations of universal hashing and privacy amplification [BBCM95,Ma93,BBR88] correspond to error correcting codes over GF(2). Although any number of reconciliation protocols are available, let us expand and simplify a natural path. S and R will statistically sample half the particles to put a cap on how much interfering Eve did. As long as it is sub-threshold (say, less than 1%), they continue. This means that (with exponentially-high probability) Eve has only measured a small fraction (say << 1/100) of the remaining indices. (More precisely, she has applied a measurement/alteration with "small" quantum entropy; it may affect any number of particles and in superposition.)

For single-bit messages, there are two canonical superpositions that S will have sent, depending on whether $k = 0$ or $k = 1$:

$$\phi_k = \frac{1}{\sqrt{|C_2|}} \sum_{z \in P(k)} |z\rangle$$

where $P(k)$ denotes all strings in the coset $v_k + C_2$ (with v_k being a fixed member corresponding to k).

For example, let BB-EPR-PAR(k) be the variant of BB84 in which $P(m)$ contains all k-bit strings whose parity is m. To alter a given word requires reversing at least one bit.

4 Eavesdropping and Binding BB84

We now consider what happens when an eavesdropper listens in on BB84. Taking a cue from standards bodies, we let "must" indicate an apparent intuitive requirement (not necessarily necessary), and let *MUST* indicate a requirement that we do assert as factual.

The intuitive observation is that S and R "must" change a bit somewhere in order to pretend to have sent the opposite cleartext. While this classical reasoning is not justifiable in the quantum setting, it approaches the same final conclusions.

Even though (S, R) face an unknown adversary, they *MUST* select a strategy (D_S, D_R) in advance. Clearly, the actual computation can depend dynamically on the results of the transmission. But there can be no argument: the encryption/denial programs (S, R, D_S, D_R) *MUST* be written down by a designer.

Let eavesdropper $\mathrm{M}(a, k)$ trade a qubit at position a:

$$|z\rangle^S |00\rangle^E |0^k\rangle^R \rightarrow |z\rangle^S |z(a)0\rangle^E |z(a{:}0)\rangle^R$$

where $z(a)$ indicates the a^{th} bit in z, and $z(a{:}0)$ indicates replacing bit a in z by 0. Like measuring a photon, this action disturbs the stream of bits, although it does not give M information about the overall "key." (Recall that the notation was simplified by adjusting it after S announces the bases. When M acts, she does not know the correct basis, and simply uses $+$. We now consider only those paths in which S announced $+$ at index a.)

Consider BB-EPR-PAR (parity-based) with odd $k \geq 3$. Let $\Phi(E, m)$ be the state obtained after sending bit m. There are a variety of equivocation transforms U available to patch $\Phi(\emptyset, 0)$ to $\Phi(\emptyset, 1)$ against a passive eavesdropper. Some simple ones are (*cf.* one-time pad) to negate all qubits; a given qubit; a random qubit. Luckily, these are all local transforms, too. We take U_{neg} which negates all qubits.

We now show this leads to catastrophe, as do other similar choices. Let $P(m; a, b) = \{z \in P(m) \mid z(a) = b\}$, and let $\rho(E, m) = \mathrm{Tr}_{\mathcal{H}_{E,env}} \Phi(E, m)$.

$$\rho(\mathrm{M}(a,k),0) = \sum_{P(0;a,0)} |z\rangle^S |z\rangle^R \langle Z|^R \langle Z|^S + \sum_{P(0;a,1)} |z\rangle^S |z(a{:}0)\rangle^R \langle Z(a{:}0)|^R \langle Z|^S$$

$$\rho(\mathrm{M}(a,k),1) = \sum_{P(1;a,0)} |z\rangle^S |z\rangle^R \langle Z|^R \langle Z|^S + \sum_{P(1;a,1)} |z\rangle^S |z(a{:}0)\rangle^R \langle Z(a{:}0)|^R \langle Z|^S$$

$$U_{neg}\rho(\mathrm{M}(a,k),0)U_{neg}^{\dagger} = \sum_{P(1;a,0)} |\bar{z}\rangle^S |\bar{z}\rangle^R \langle \bar{Z}|^R \langle \bar{Z}|^S + \sum_{P(1;a,1)} |\bar{z}\rangle^S |\bar{z}(a{:}1)\rangle^R \langle \bar{Z}(a{:}1)|^R \langle \bar{Z}|^S$$

$$= \sum_{P(1;a,0)} |z\rangle^S |z(a{:}1)\rangle^R \langle Z(a{:}1)|^R \langle Z|^S + \sum_{P(1;a,1)} |z\rangle^S |z\rangle^R \langle Z|^R \langle Z|^S$$

There is clearly 0 fidelity between $\rho(\mathrm{M}(a,k),1)$ and $U_{neg}\rho(\mathrm{M}(a,k),0)$. A judge simply calculates $d = (\oplus_i z^S(i)) \oplus (z^R(k))$. Even if the protocol designer opts to use U_{neg} only some of the time, then against a passive adversary or against $\mathrm{M}(a,k)$, equivocation is detected with probability at least $1/2$ (again, in at least $1/2$ of all paths). By inspection, the judge is both safe and binding.

If the protocol designer uses "$U_{neg,p}$", in which qubit at location p is negated, then $\mathrm{M}(p,k)$ (versus passive) presents a problem. If "$U_{neg,r}$" is used, in which a randomly selected qubit is negated, then $\mathrm{M}(1,k)$ leads to $(1 - 1/k)$ fidelity. This is better for S and R, but still binding.

Observations. We avoided constructing a single eavesdropper who randomly chooses to be passive or to invade. Other objections aside, this could pose the risk of enabling an equivocation strategy (insofar as E's program were to become public), as mentioned earlier (§6.1). This is a subtle but important issue that differentiates classical cryptography from quantum.

It is straightforward to imagine extensions to BB84 that apparently provide *perilous deniability.* For instance, by lengthening the "ciphertext" to contain κ^{2C} bits, arranging equivocation could be as simple as changing a "small" $(1/\kappa^C)$

fraction of bits. For each fixed C, there remains a non-negligible chance to get caught (hence perilous and still weakly binding). But even this weak goal is much harder to prove in the quantum setting than initial intuition suggests, since Eve can spread her eavesdropping among many bits, not just focus on a particular small subset. Space does not permit a full analysis here.

5 Mayers' Theorem

Mayers takes great care to state the model precisely, and we follow his descriptions [Ma96].

5.1 Defining Bit Commitment

In a bit commitment protocol, Alice encodes her input bit b into a state $|\psi_b\rangle$ of $\mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_{env}$, using an initial protocol, *commit*(b). A second protocol, *unveil*($|\psi_b\rangle$), is used to give Bob b or a "refusal" string $\perp$.

Alice has the power to choose $p(b \mid \text{not } \perp)$ by behaving honestly. The goal of the *commit* protocol is to prevent her from subsequently changing $p(b \mid \text{not } \perp)$ even if she was dishonest. Let *unveil'* denote running *unveil* with possibly dishonest A. A state $|\psi\rangle$ is *perfectly committing* if every attack *unveil'* either returns $\perp$ with probability 1 or returns b with probability $p'(b \mid \text{not } \perp) = p(b \mid \text{not } \perp)$.

Bob may attempt to gain b prematurely. Let η_B be Bob's classical information from $\mathcal{H}_{env \cap B}$, as the state is collapsed into $|\psi_{b,\eta}\rangle$ of $\mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_{env}$. The reduced density matrix of Bob given η is:

$$\rho_B(|\psi_{b,\eta}\rangle) = \mathrm{Tr}_{\mathcal{H}_{env \cap A}}(|\psi_{b,\eta}\rangle\langle\psi_{b,\eta}|).$$

If b is fixed by η, we let $F(\eta) = 0$. Otherwise, we let $F(\eta)$ be the *fidelity* between $\rho_B(|\psi_{0,\eta}\rangle)$ and $\rho_B(|\psi_{1,\eta}\rangle)$. (The fidelity $F(\alpha, \beta)$ is the supremum of $\mid \psi_\alpha^\dagger \psi_\beta \mid$ over all purifications ψ_α, ψ_β of α, β. Note that $F(\alpha, \beta) = 1$ only when $\alpha = \beta$.) A state is *perfectly concealing* if η is independent of b and the expected value of $F'(\eta)$ is 1, where F' refers to executions with a possibly cheating Bob.

A commitment protocol is *perfectly secure* if it is both perfectly committing and perfectly concealing. One can replace "perfect" by a tolerance of ϵ (or $\epsilon(k)$ for some security parameter k) in the preceding definitions.

5.2 Previous Work: No Bit Commitment

Mayers' Theorem states that quantum bit commitment is impossible:

Theorem 1. **[Mayers' Theorem, or No-Commitment]** *No properly initialized quantum bit commitment protocol is unconditionally secure.*

The proof is based on what we call an *equivocation strategy* that allows A to change the bit even after protocol *commit*. We sketch some of the ideas here.

Consider a properly initialized quantum bit commitment protocol. Let dishonest A' refrain from making measurements. Mayers shows that there is a purification $|\psi_{01}\rangle$ of $\rho_B(|\psi'_{0,\eta}\rangle)$ such that $\langle\psi_{01} \mid \psi'_{1,\eta}\rangle \geq F'(\eta)$. Moreover, there is a unitary transformation $U = U_A \otimes \mathbf{1}$ mapping $|\psi'_{0,\eta}\rangle \to |\psi_{01}\rangle$.

Because an ϵ-concealing protocol must have fidelity $F(\eta) \geq 1-\epsilon$, this implies an equivocation strategy for A'. If she wishes to set $b=0$, she makes the measurements required of A and continues with *unveil*. To set $b=1$, she applies U, performs the measurements required of A in *commit*, and continues with *unveil*. The result is that Bob accepts this unveiling with probability approaching 1.

We refer to this strategy as *Mayers Equivocation*. Unfortunately, there is an intuitive but incorrect way to paraphrase the theorem, which goes something like this:

Claim 2 [NoGo Folk "Theorem"] *In any quantum protocol, whenever $F(\rho_B(|\psi_{0,\eta}\rangle), \rho_B(|\psi_{1,\eta}\rangle)) \approx 1$, then Alice can equivocate successfully with probability ≈ 1. (to be disproved)*

6 Limitations on Applying No-Commitment

The No-Commitment theorem is obtained through two methods: (1) abstaining from private measurement, followed by (2) applying a unitary transformation to change $|\psi_0\rangle$ to $|\psi_1\rangle$.

There are several aspects of the model and the result that make it insufficient to apply automatically to QKE. These include the quantifiers, colocation, and the impact of generic abstinence from measurement.

6.1 Quantifier Problems

Mayers' result essentially says, $(\forall B)(\exists U_A(B))$ such that a cheating committer A can employ $U_A(B)$ to equivocate. The strategy can depend on B's program, which is acceptable because the honest programs for A and B must be declared. Naturally, the cheating programs need not be disclosed, but to disprove security, it suffices to show that no honest program is protected.

Likewise, it is hard to imagine an encryption protocol in which S and R are allowed to know what E's program is. Their attempts to communicate, and to deny, must be successful even without being given details of E's program. (There may be some dynamic deductions to make about E based on her behavior, but this is vastly different than knowing her full program.) A quantification, "$(\forall E)(\exists S, R)$ such that S and R successfully communicate," is insufficient.

Yet the natural generalization of Mayers' result to QKE is precisely backwards. S and R are not given E's program and then allowed to equivocate.

The "folk no-commit theorem" fails to hold: an arbitrary E does indeed gain no information, and indeed there mathematically exists a Mayers equivocating transform on the joint (S, R) state, but S and R have no way to determine what it is, since E is arbitrary and inaccessible. Further problems occur; see below.

6.2 Colocation

To equivocate a one-time pad, one merely needs to reverse a bit. This can be done individually and locally by S and R, without communication (apart from knowing they must equivocate).

The direct application of No-Commitment treats (S, R) as the committer, who, for any trustee/eavesdropper, has an equivocation transformation U. (Forget about the order of quantifiers.) There is no mathematical guarantee that U can be factored into local transforms U_S and U_R. Therefore, the (nevertheless correct) proof given in [Ma96] does not provide sufficient grounds to apply to QKE. There may indeed be local strategies for S and R for particular QKE protocols, but they are not implied directly by [Ma96].

If S and R need to communicate or be co-located in order to equivocate, the deniability property is weakened. A vote-coercing Mafioso simply interrogates them separately. While deniability that requires colocation may be better than nothing, it is not always sufficient.

6.3 Abstinence Makes the Bit Grow Weaker

An extremely critical (and clever) aspect of Mayers' approach is the demand that parties refrain from making internal measurements. This gives them the flexibility to equivocate later.

In bit commitment, abstaining seems to have no obvious impact. There is no particular reason why A should do any measurements. She already knows her bit and doesn't stand to discover anything much.

But [BB84] specifically requires measurements to be peformed, both to check how invasive E is and to discern what the message is. In the full protocol purification, BB84-Pur, R cannot receive the cleartext privately. Depending on interpretation, either no measurement is made at all, or the only actionable information is whatever was transmitted over the clear classic channel.

In the BB84-Key variant, close inspection reveals that the previously classical channel is now used to send check-information as qubits to R. But this presupposes the end result: a secure quantum channel. Thus BB84-Key is not implementable within the rules of the model.

Even the limited purification to BB84-EPR/Ekert does not buy anything. An attack similar to the one against BB84 in §4 will work against entanglement purification of a tainted EPR source.

7 PQECC: Deniable QKE

We now propose a protocol to achieve deniability, although it requires a quantum computer. Normally, quantum error-correcting codes are useful to protect against decoherence. Typically, a syndrome is measured and then used to restore the original state. We applied QECC in an unusual twist: we avoid measuring the relative syndrome. Instead, a postulated quantum computer applies the QECC

PQECC

1. S constructs duplicate registers $(K, B, \mathit{\Pi}, Y, V, W)$ with respective lengths k, $2n$, $2n \log 2n$, n, n, n, in superposition $\sum |K', B', \mathit{\Pi}', Y', V', W'\rangle |K, B, \mathit{\Pi}, Y, V, W\rangle$, (except $W' = W = |0^n\rangle$). All computations are quantum (except explicit measurements later), including error correction.
2. S quantum-computes the Y-selected ECC representation of K and places it in W, leaving $K = |0^n\rangle$.
3. S applies interleaving $\mathit{\Pi}$ against the 2n qubits in (V, W), leaving $\mathit{\Pi} = |0^{2n \log 2n}\rangle$.
4. Let T be the joint $2n$-qubit register (V, W). S applies Hadamards based on B to corresponding qubits in T.
5. S sends T to E.
6. E attacks T via operator U_E.
7. R receives the manipulated T register and reports receipt.
8. S measures $(b, \pi, v, y) \leftarrow (B, \mathit{\Pi}, V', Y')$ and announces them.
9. R applies Hadamards to T according to b, then applies π^{-1} to T. Consider T again as registers (V, W).
10. R measures register V and aborts if there are any (or more than a threshold t) mismatches with v.
11. R sets new register E to $|0^n\rangle$ and applies error correction (using y) to (W, E).
12. S measures key $k_S \leftarrow K'$ and R measures key $k_R \leftarrow W$. Encryption is as a OTP: $ciph \leftarrow k_S \oplus m$, $m_{rcv} \leftarrow ciph \oplus k_R$.

Fig. 2. PQECC protocol for n-bit key k, rewritten to be compatible with modern conventions. The ECC steps are quantum-level calculations for the random hashing and error correction in BB84 *et seq.*

without measuring the relative syndrome at all. The goal is to purify the BB84 protocol as much and as carefully as possible. Only the most necessary measurements are actually performed. (Note that the key/message is emphatically not the only register that is measured.)

This *purified* QECC, or PQECC protocol, is the basis for establishing deniability. Although it was investigated several years ago [Be96d], it is closely related by convergent evolution to the "modified Lo-Chau" and "QKD-CSS" protocols [ShPr01], discussed below. We have tried to present it in a form that illustrates its connections to those protocols.

First, we comment on the registers and computations. The B and $\mathit{\Pi}$ registers correspond to random axis choices and random interleaving of key bits with check bits. ($\mathit{\Pi}$ can certainly be a permutation, as suggested by [ShPr01]. The random hashing and amplification in BB84 is purified to a random ECC selected by Y (corresponding to the (x, z) choice in [ShPr01]). The decoding procedure decouples the signal (W) from the noise (E).

7.1 Related Protocols

Two related protocols, called QKD-CSS and "modified Lo-Chau," appeared subsequent to the first consideration of PQECC as the maximal effective protocol purification of BB84 [Be96d]. If one imagines a hierarchy of protocol classes depending on purification and/or specific error-correction codes, PQECC is more or less subsumed by “modified Lo-Chau” while it more or less subsumes QKD-CSS.

The motivation for PQECC was to obtain deniability, while the motivation for Lo-Chau and QKD-CSS was to find a proof of privacy for any kind of QKE, and especially QKE without quantum computation. Hindsight shows that this independent convergence to similar protocols is natural given the drive to use protocol purification as (1) a tool for proving privacy and (2) a tool for achieving deniability.

7.2 Privacy and Deniability

Recently, simplified proofs of privacy have appeared for BB84 [LC99,GoPr00,ShPr01]. There are two important ingredients that concern us.

It is useful to try to establish the secret key as a sequence of k ebits shared by S and R, namely $\Phi = \beta_{00}^{\otimes k}$, starting with a noisy pair of n-bit registers in state ρ. Let ρ' describe the state of the k-bit key registers after the full protocol. Lo and Chau’s approach employs a result shown by Gottesman and Preskill:

$$\langle \beta_{00}^{\otimes k} | \rho' | \beta_{00}^{\otimes k} \rangle \geq \mathrm{Tr}(\Pi \rho),$$

where Π projects onto Bell states differing from $\beta_{00}^{\otimes n}$ by at most t bit-flip errors (applications of Pauli X) and at most t phase-flip errors (applications of Pauli Z). Indirectly argued, the sampling test gives an accurate bound on $\mathrm{Tr}(\Pi \rho)$, ultimately implying that the fidelity between ρ' and $\beta_{00}^{\otimes k}$ is exponentially close to 1. This demonstrates privacy for the Lo-Chau protocol, although a quantum computer is necessary.

In a second stage, Shor and Preskill apply this to a protocol (QKD-CSS) that uses CSS codes. By then instructing S and R to perform measurements sooner rather than later, the QKD-CSS protocol “reduces to” BB84, and the requisite quantum computer can be avoided without introducing any information leak to Eve.

For deniability, the first ingredient is essential. It ultimately allows us to conclude that Eve is entangled negligibly with K' (and instead overwhelmingly with register E) in protocol PQECC. This means that the simple OTP denial strategy is overwhelmingly effective: pretend $k' = k \oplus m_0 \oplus m_1$. Without sufficient space for proof, we merely assert:

Claim. PQECC is a deniable quantum cryptosystem.

This analysis should also extend to other cryptosystems such as the “modified Lo-Chau” protocol.

8 Conclusions

The BB84 protocol is weakly binding on S and R. Despite the reaction of some, Mayers’ no-commit theorem [Ma96] does not suffice to turn BB84 into a deniable QKE. There are formal improprieties with quantifiers, insufficient support for equivocation without co-location, conflicts between measurement abstinence and correct or allowable transmission, and counterintuitive adversarial binding arguments. (None of this impugns the correct work on bit commitment.)

This paper seeks to make practitioners aware of the incomplete analysis of QKE, despite claims of "perfect, unconditional security." An off-the-shelf optic-fiber-based QKE might enable a private electronic vote but it will support coercion and vote-selling.

Using protocol purification in a refined manner can provide a deniable QKE, in the form of the PQECC protocol, but a quantum computer is required. But there remains a great deal of turbidity in current quantum "security reductions."

Acknowledgements I thank Claude Crépeau and Jeroen van de Graaf for helpful conversations (long ago!), and Joern Müller-Quade for more recent ones. Several referees made extensive remarks (over many years) and I thank them for their efforts. All mistakes remain mine.

References

[ADH97] L. Adleman, J. Demarrais, M. Huang. "Quantum Computability." SIAM J. Comput., **26**:5, 1997, 1524–1540.

[Ba95] A. Barenco. "A Universal Two-Bit Gate for Quantum Computation." Proc. Royal Society of London, **449**, 1995, 679–683.

[BDM+95] A. Barenco, C. Bennett, R. Cleve, D. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. Smolin, H. Weinfurter. "Elementary Gates for Quantum Computation." Phys. Rev. Letters A, **52**, 1995, 3457–3467.

[Be96] D. Beaver. "Plausible Deniability." Proc. of PragoCrypt 1996, J. Prybl, Ed., CTU Publishing House, Prague, 1996, 272–288.

[Be96d] D. Beaver. Unpublished manuscript, 1996.

[Be99] D. Beaver. "Imperfections in Perfectly Secure Key Exchange." IEEE Information Theory and Networking Workshop, Metsovo, 1999.

[Be92] C. Bennett. "Quantum Cryptography Using Any Two Orthogonal States." Phys. Rev. Letters, **67**:21, 1992, 2121–2124.

[BBBSS92] C.H. Bennett, F. Bessette, G. Brassard, L. Salvail, J. Smolin. "Experimental Quantum Cryptography." *Journal of Cryptography*, **5**:1, 1992, 3–28.

[BB84] C. Bennett, G. Brassard. "Quantum Cryptography: Public-Key Distribution and Coin-Tossing." Proceedings of IEEE CSSP, Bangalore, India, 1984, 175–179.

[BBCM95] C. Bennett, G. Brassard, C. Crépeau, U. Maurer. "Generalized Privacy Amplification." IEEE Trans. Information Theory, **41**:6, 1995.

[BBM92] C. Bennett, G. Brassard, D. Mermin. "Quantum Cryptography Without Bell's Theorem." Phys. Rev. Letters, **68**:5, 1992, 557–559. Also see Manuscript, March 6, 1995.

[BBR88] C. Bennett, G. Brassard, J.M. Robert. "Privacy Amplification by Public Discussion." SIAM J. Computing, **16**:2, 1988, 210–229.

[BCMS97] G. Brassard, C. Crépeau, D. Mayers, L. Salvail. "A Brief Review on the Impossibility of Quantum Bit Commitment." Los Alamos Preprint Archive quant-ph/9712023, 1997.

[BCMS98] G. Brassard, C. Crépeau, D. Mayers, L. Salvail. "Defeating Classical Bit Commitments with a Quantum Computer." Los Alamos Preprint Archive quant-ph/9806031, 1998.

[BCJL93] G. Brassard, C. Crépeau, R. Josza, D. Langlois. "A Quantum Bit Commitment Scheme Provably Unbreakable by Both Parties." *Proc. of* 34th *FOCS*, IEEE, 1993, 362–371.

[BS93] G. Brassard, L. Salvail. "Secret-Key Reconciliation by Public Discussion." *Advances in Cryptology – EuroCrypt '93*, Springer Verlag LNCS **765**, 1993, 410–423.

[CSS96] A. R. Calderbank, P. Shor, "Good Quantum Error Correcting Codes Exist." Phys. Rev. A **54**, 1996, 1098–1105. A. M. Steane, "Multiple Particle Interference and Error Correction." Proc. R. Soc. London A **452**, 1996, 2551–2577.

[CDNO97] R. Canetti, C. Dwork, M. Naor, R. Ostrovsky. "Deniable Encryption." *Advances in Cryptology – Crypto '97*, Springer-Verlag LNCS **1294**, 1997, 90–104.

[De89] D. Deutsch. "Quantum Computational Networks." Proc. Royal Society of London, **425**, 1989, 73–90.

[Di95] D. DiVincenzo. "Two-Bit Gates are Universal for Quantum Computation." Phys. Rev. A, **50**, 1995, 1015-1022.

[Ek91] A. Ekert. "Quantum Cryptography Based on Bell's Theorem." Phys. Rev. Letters, **67**:6, 1991, 661-663.

[ERTP92] "Practical Quantum Cryptography Based on Two-Photon Interferometry." Phys. Rev. A, **48**:1, 1993, R5–R8.

[Fe86] R. Feynman. "Quantum Mechanical Computers." Found. Phys. **16**, 1986, 507-531.

[GoPr00] D. Gottesman, J. Preskill. "Secure Quantum Key Distribution using Squeezed States." Los Alamos Preprint Archive quant-ph/0008046, 2000.

[HJW93] L. Hughson, R. Josza, W. Wooters. "A Complete Classification of Quantum Ensembles Having a Given Density Matrix." *Phys Letters A*, **183**, 1993, 14–18.

[LC96] H.K. Lo, H.F. Chau. "Is Quantum Bit Commitment Really Possible?" Los Alamos Preprint Archive quant-ph/9603004, 1996.

[LC97] H.K. Lo, H.F. Chau. "Why Quantum Bit Commitment and Ideal Quantum Coin Tossing are Impossible." Los Alamos Preprint Archive quant-ph/9711065, 1997.

[LC99] H.-K. Lo, H. F. Chau, "Unconditional Security of Quantum Key Distribution over Arbitrarily Long Distances." Science **283**, 1999, 2050–2056.

[Ma93] U. Maurer. "Secret Key Agreement by Public Discussion from Common Information." IEEE Trans. Information Theory, **39**:3, 1993, 733–742.

[Ma96t] D. Mayers. "The Trouble with Quantum Bit Commitment." Los Alamos Preprint Archive quant-ph/9603015, 1996.

[Ma96] D. Mayers. "Unconditionally Secure Quantum Bit Commitment is Impossible." *PhysComp '96*, Boston, November 1996.

[Ma97] D. Mayers. "Unconditionally Secure Quantum Bit Commitment is Impossible." *Phys. Rev. Letters*, **78**, 1997, 3414–3417.

[ShPr01] P. Shor, J. Preskill. "Simple Proof of Security of the BB84 Quantum Key Distribution Protocol." Los Alamos Preprint Archive quant-ph/0003004, 2000.

[Wi83] S. Wiesner. "Conjugate Coding." SIGACT News, **15**:1, 1983, 78–88; orig. manuscript circa 1970.