



Public-key cryptography based on bounded quantum reference frames



Lawrence M. Ioannou^{a,*}, Michele Mosca^{a,b}

^a Institute for Quantum Computing and Department of Combinatorics and Optimization, University of Waterloo, 200 University Avenue, Waterloo, Ontario, N2L 3G1, Canada

^b Perimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo, Ontario, N2L 2Y5, Canada

ARTICLE INFO

Article history:

Received 2 March 2009

Accepted 31 May 2011

Available online 17 September 2014

Keywords:

Public-key authentication

Identification

Quantum-public-key cryptography

Reference frames

ABSTRACT

We demonstrate that the framework of bounded quantum reference frames has application to building quantum-public-key cryptographic protocols and proving their security. Thus, the framework we introduce can be seen as a public-key analogue of the framework of Bartlett et al. [1], where a private shared reference frame is shown to have cryptographic application. The protocol we present in this paper is an identification scheme, which, like a digital signature scheme, is a type of authentication scheme. We prove that our protocol is both reusable and secure under the honest-verifier assumption. Thus, we also demonstrate that secure reusable quantum-public-key authentication is possible to some extent.

© 2014 Published by Elsevier B.V.

1. Introduction

Since its inception, the focus of quantum cryptography has been on the symmetric-key model, where Alice and Bob attempt to generate private shared correlations (as in quantum key distribution [2,3]) or are assumed to hold them (as in quantum authentication [4]). Such correlations can usually be defined or encoded by a string of bits—the secret key—but Bartlett et al. [1] showed that they may also take the form of a private shared reference frame. Symmetric-key quantum protocols are usually unconditionally secure, meaning that the sole assumption is that (some part of) quantum theory is correct; however, Damgaard et al. [5,6] have investigated information-theoretically secure protocols, such as password-based identification and bit commitment, in the bounded quantum storage model, where an extra assumption is that the size or quality of the adversary's quantum memory is limited (see also Refs. [7–9]).¹

Going beyond the symmetric-key model, but retaining unconditional security, Gottesman and Chuang [11] introduced quantum-public-key cryptography—where the public keys are quantum systems, each of whose state encodes the (same) classical private key—by giving a secure one-time (digital) signature scheme for signing classical messages.

A public-key framework eliminates the need for Alice and Bob to establish private shared correlations, which has practical advantages in large networks of users (where there may be many “Alices” or “Bobs”). Alice chooses a random private key, creates copies of the corresponding public key, and distributes the copies in an authenticated fashion to all potential “Bobs”. In principle, this asymmetric setup allows, e.g., any Bob to send encrypted messages to Alice or to verify any signature for a message that Alice digitally signed, thus significantly reducing the number of secret/private keys involved as compared to

* Corresponding author.

E-mail addresses: lmioannou@gmail.com (L.M. Ioannou), mmosca@iqc.ca (M. Mosca).

¹ The bounded storage model for classical protocols (e.g. Ref. [10]), where the adversary's classical memory is assumed to be bounded, also gives information-theoretic security.

the case where each Alice–Bob pair shares a secret key and uses symmetric-key protocols. Thus the public-key framework vastly simplifies key distribution, which is often the most costly part of any cryptosystem. Note that the security of classical public-key protocols is necessarily based on computational assumptions [12].

The mapping that takes a private key to the state of the corresponding quantum public key is always assumed to be publicly known. Furthermore, in any reasonable quantum-public-key system, the states of two quantum public keys corresponding to two different private keys always have overlap less than $(1 - \delta)$, for some positive and publicly known δ . Thus, a striking aspect of quantum-public-key cryptography that sets it apart from its classical counterpart is that the number of copies of the (quantum) public key in circulation must be limited. If this were not the case, then an adversary could collect an arbitrarily large number of copies, measure them all, and determine the private key.

The limit on the number of copies of the quantum public key implies that not everyone can use the protocol; however, in practice, the maximum number of users (or uses) of any particular protocol can be estimated, and thus the parameters of the protocol can be adjusted so that the limit allows for this maximum. Increasing this limit would presumably result in a less efficient instance of the protocol, and this is one kind of tradeoff between efficiency and usability in the quantum-public-key setting. Another kind concerns reusability. For instance, the abovementioned signature scheme is “one-time” because only one message may be signed under a particular key-value, even though many different users can verify that one signature. If a second message needs to be signed, the signer must choose a new private key and then distribute corresponding new public keys. One open problem is thus whether there exist reusable signature schemes, where either the same copy of the public key can be used to verify many different message-signature pairs securely, or where just the same key-values can be used to verify many different message-signature pairs securely (but a fresh copy of the public key is needed for each verification). The latter notion of “reusability” is what we adopt here.

What makes a key *public*? In principle, Alice’s public-key-generation algorithm, which takes as input the private key and outputs one copy of the quantum public key, may output a system in a pure state or a mixed state, from Alice’s point of view (a mixed state is a fixed probabilistic distribution of pure states). In the original framework of Gottesman and Chuang, the algorithm is assumed to produce a system in a pure state. For some applications, like digital signature schemes, this purity is crucial; for, otherwise, Alice could cheat by sending different public keys to different “Bobs”. Purity prevents Alice’s cheating here because different “Bobs” can compare their copies of the public key via a “distributed SWAP-test” [11] to see if they are the same (with high probability), much like can be done in the case of classical public keys. But the ability to do an equality test benefits any scheme, since an adversary who tries to substitute bad keys for legitimate ones could thus be caught. Indeed, if the public-key-generation algorithm produces a mixed state, since there is no equality test guaranteed to recognize when two mixed states are equal, then no such test for equality of public keys may be possible—this is at odds with what it means to be “public”, i.e., publicly verifiable.² While the scheme we present in this paper does not explicitly make use of the “distributed SWAP-test” (since we assume the public keys have been distributed securely), it can do so in principle. We view this as analogous to how modern public-key protocols do not explicitly specify an equality test among unsure “Bobs”, but how the framework naturally allows such a test which would thwart attempts to distribute fake public keys.

Our work appears to be of a dramatically different character when compared to other explorations of quantum-public-key protocols [13–17]: we demonstrate that the framework of bounded quantum reference frames [18] has application to building such protocols and proving their security. Thus, the framework we introduce can be seen as a public-key analogue of the framework of Bartlett et al. [1].

We stress that our work in public-key quantum cryptography strives for unconditional security, as opposed to security based on computational assumptions [12]. In particular, our work is unrelated to the work in Ref. [19], where classical public-key systems (whose security must be based on computational assumptions) are constructed that require a quantum computer for the generation of the public keys.

The protocol we present in this paper is an identification scheme, which, like a digital signature scheme, is a type of authentication scheme. Authentication schemes are not concerned with ensuring the *privacy* of information, but rather seek to ensure its *integrity*. For example, digital signature schemes (and message authentication codes) ensure the integrity of origin of messages, whereas identification schemes ensure the integrity of origin of communication in real time [12]. Identification protocols are said to ensure “aliveness”—that the entity proving its identity is active at the time the protocol is executed.

We prove that our identification protocol is both reusable and secure under the honest-verifier assumption (defined in the next section). Thus, we also demonstrate that secure reusable quantum-public-key authentication is possible to some extent.

We now proceed with a description of our protocol (Section 2) and the honest-verifier security proof (Section 3).

² Other authors have defined the framework to include mixed public keys, and Ref. [13] proposes an encryption scheme with mixed public keys that is reusable and unconditionally secure [14].

2. An identification scheme

In the following, Alice and Bob are always assumed to be honest players and Eve is always assumed to be the adversary. Suppose Alice generates a private key and authentically distributes copies of the corresponding public key to any potential users of the scheme, including Bob.

The following is a description (adapted from Section 4.7.5.1 in Goldreich's book [20]) of how a secure public-key identification scheme works. If Alice wants to identify herself to Bob (i.e. prove that it is she with whom he is communicating), she invokes the identification protocol by first telling Bob that she is Alice, so that Bob knows he should use the public key corresponding to Alice (assuming Bob possesses public keys from many different people). The ensuing protocol (whatever it is) has the property that the *prover* Alice can convince the *verifier* Bob (except, possibly, with negligible probability) that she is indeed Alice, but an adversary Eve cannot fool Bob (except with negligible probability) into thinking that she is Alice, even after having listened in on the protocol between Alice and Bob or having participated as a (devious) verifier in the protocol with Alice several times. An *honest-verifier identification protocol* is only intended to be secure under the extra assumption that, whenever Eve engages the prover Alice in the protocol, Eve follows the verification protocol as if she were honest. Note that no identification protocol is secure against an attack where Eve concurrently acts as a verifier with Alice and as a prover with Bob (but note also that, in such a case, the “aliveness” property is still guaranteed). Note also that, by our definition of “reusable,” an identification scheme is considered reusable if Alice can prove her identity many times using the same key-values but the verifier needs a fresh copy of the public key for each instance of the protocol.

A couple of remarks are in order:

- In practice, public-key identification schemes are implemented in smart-card systems (e.g., inside an automated teller machine (ATM) for access to a bank account, or beside a doorway for access to a building), so that the smart card “proves” to the card reader that it is authorized.³ In such situations, it may be relatively difficult for an adversary to tamper with the verification procedure that is encoded in the card reader, in which case an honest-verifier identification protocol may suffice. For example, an honest-verifier identification protocol is secure against the class of attacks whereby an adversary collects the maximum number of legitimate copies of the public key and uses these in conjunction with a phony smart card to act as a dishonest prover.
- Note that public-key identification can be trivially achieved via a digital signature scheme (Alice signs a random message presented by Bob), but since we do not know of an unconditionally secure and reusable digital signature scheme, our scheme is noteworthy. Similarly, public-key identification can be achieved with a public-key encryption scheme (Bob sends an encrypted random challenge to Alice, who returns it decrypted), but we do not know of an unconditionally secure and reusable public-key encryption scheme (that uses pure public keys; though, see Ref. [15] for a promising candidate).

A summary of our protocol is as follows. Alice chooses a private phase reference and distributes a limited number of samples of her reference frame as quantum public keys. The samples are used by Bob to verify that the prover is actually Alice. Because Alice has a perfect phase reference, she can carry out the identification protocol with no error (assuming perfect quantum channels). But, because Eve only has a bounded quantum phase reference frame (in the form of a limited number of copies of the public key), she inevitably incurs an error that Bob can detect with sufficiently high probability (we discuss quantum phase reference frames in more detail in Section 3.4.1).

2.1. Protocol specification

Our identification protocol takes the form of a typical “challenge-response” interactive proof system, consisting of a kernel (or subprotocol) that is repeated several times in order to amplify the security, i.e., reduce the probability that an adversary can break the protocol. We assume all quantum channels are perfect.

Parameters

- The *security* parameter $s \in \mathbb{Z}^+$
 - ◊ equals the number of kernel iterations.
 - ◊ The probability that Eve can break the protocol (in an honest-verifier setting) is exponentially small in s .
- The *reusability* parameter $r \in \mathbb{Z}^+$
 - ◊ equals the maximum number of copies of the quantum public key in circulation and
 - ◊ equals the maximum number of times the protocol may be executed by Alice, before she needs to pick a new private key.

³ Note that it is not a user's personal identification number (PIN) that functions as the prover's private key; the PIN only serves to authenticate the user to the smart card (not the smart card to the card reader).

Keys

- The *private key* is

$$(x_1, x_2, \dots, x_s), \quad (1)$$

where Alice chooses each x_j , $j = 1, 2, \dots, s$, independently and uniformly randomly from $\{1, 2, \dots, 2r + 1\}$.

◊ The value x_j is used only in the j th kernel-iteration.

- One copy of the *public key* is an s -partite system in the state

$$\bigotimes_{j=1}^s |\psi_{x_j}\rangle, \quad (2)$$

where (omitting normalization factors)

$$|\psi_{x_j}\rangle := |0\rangle + e^{2\pi i x_j / (2r+1)} |1\rangle. \quad (3)$$

◊ Alice authentically distributes (e.g. via trusted courier) at most r copies of the public key.

◊ The j th subsystem of the public key (which is in the state $|\psi_{x_j}\rangle$) is only used in the j th kernel-iteration.

Actions

- The *kernel* $\mathcal{K}(x)$ of the protocol is the following three steps, where we use the shorthand

$$\phi_x := 2\pi x / (r + 1), \quad (4)$$

and where we have dropped the subscript “ j ” from “ x_j ”.

(1) Bob creates $|0\rangle|1\rangle + |1\rangle|0\rangle$, and sends one register of this system to Alice.

(2) Alice measures the received register in the basis $\{|0\rangle \pm e^{i\phi_x}|1\rangle\}$. If the state of the register immediately after the measurement is $|0\rangle + e^{i\phi_x}|1\rangle$, then Alice sends “0” to Bob; otherwise, Alice sends “1”.

(3) If Bob receives “1”, then he applies the Pauli-Z gate

$$Z := \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \quad (5)$$

to the register that he kept in Step 1. Finally, Bob SWAP-tests⁴ this register with his authentic copy of $|\psi_x\rangle$.

- When Alice wants to identify herself to Bob, they take the following actions:

(i) Alice checks that she has not yet engaged in the protocol r times before with the current value of the private key; if she has, she aborts (and refreshes the private and public keys).

(ii) Alice sends Bob her purported identity (“Alice”), so that Bob may retrieve the public keys corresponding to Alice.

(iii) The kernel $\mathcal{K}(x)$ is repeated s times, for $x = x_1, x_2, \dots, x_s$. Bob “accepts” if all the SWAP-tests passed; otherwise, Bob “rejects”.

2.2. Completeness of the protocol

It is clear that the protocol is correct for honest players, that is, Bob always “accepts” when Alice is the prover. To see this, note that, up to global phase, the state $|0\rangle|1\rangle + |1\rangle|0\rangle$ equals

$$(|0\rangle + e^{i\phi_x}|1\rangle)(|0\rangle + e^{i\phi_x}|1\rangle) - (|0\rangle - e^{i\phi_x}|1\rangle)(|0\rangle - e^{i\phi_x}|1\rangle). \quad (7)$$

In the next section, we will prove that the protocol is also secure, under the honest-verifier assumption.

3. Honest-verifier security

Let us clearly define what Eve is allowed to do in our honest-verifier model. Eve can

- passively monitor Alice’s and Bob’s interactions (which means that Eve can read the classical bits sent by Alice, and read the bit that indicates whether Bob “accepts” or “rejects”), and

⁴ The SWAP-test of two registers (labeled 2 and 3) in the states $|\xi\rangle_2$ and $|\chi\rangle_3$ is a measurement (with respect to the computational basis $\{|0\rangle_1, |1\rangle_1\}$) of the control register (labeled 1) of the state

$$(H_1 \otimes I_2 \otimes I_3)(c - \text{SWAP}_{2,3})(|0\rangle_1 + |1\rangle_1)|\xi\rangle_2|\chi\rangle_3/\sqrt{2}, \quad (6)$$

where H_1 is the usual Hadamard gate (applied to register 1) and $c - \text{SWAP}_{2,3}$ is the controlled-SWAP gate. The probability that the state is $|0\rangle_1$ immediately after the measurement—which corresponds to a *pass*—is $(1 + |\langle \xi | \chi \rangle|^2)/2$. When the registers 2 and 3 are in the mixed states ρ and ρ' , this probability is $(1 + \text{tr}(\rho\rho'))/2$.

- participate as the verifier, but only performing the actions as if she were honest, and
- participate as the prover in one or more instances of the protocol.

Note that Eve is assumed not to be able to actively interfere with Alice’s and Bob’s communications when Alice and Bob are participating in the protocol, as this would allow Eve to be a dishonest verifier by stealing the qubits Bob sends to Alice and replacing them with her own qubits.

Immediately, we see that Eve’s passive monitoring only gives her independent and random bits (plus the bit corresponding to “accept”), and thus gives her no useful information (in the sense that she may as well generate random bits herself). We can therefore ignore the effects of her passive monitoring.

With regard to Eve acting as verifier, note that, in the kernel iteration $\mathcal{K}(x)$, Eve can at best extract one extra copy of $|\psi_x\rangle$ from Alice when Eve follows the verifier protocol honestly in Step 1, by not bothering to do the SWAP-test in Step 3. Eve is technically not allowed not to do the SWAP-test in our honest-verifier setting, but we show that—even if she is allowed—then the protocol is secure (as long as Eve is honest in Step 1). This allows Eve to obtain a maximum of r extra copies (recall Alice only participates in the protocol r times before refreshing her keys) of the public key in addition to any copies she obtains legitimately. Let t be the total number of copies of the public key that Eve has in her possession. Note that $t \leq 2r - 1$, since we always assume that at least one copy is left for Bob, so that Eve can carry out the protocol with him.

Therefore, to prove security in our setting, it suffices to consider attacks where Eve is armed with her t copies of the public key and she participates as a prover in order to try to cause Bob to “accept”. We use the following definition of “security”.

Definition 1 (Security). An honest-verifier identification protocol (for honest prover Alice and honest verifier Bob) is *secure with error ϵ* if the probability that Bob “accepts” when any adversary Eve participates in the protocol as a prover is less than ϵ (assuming that, whenever Eve engages Alice in the protocol, Eve follows the verification protocol honestly).

We will assume that Eve has always extracted the r illegitimate copies of the public key from Alice, and we define t' to be the number of copies that Eve obtained legitimately:

$$t = r + t'. \quad (8)$$

Note that Eve can make at most $(r - t')$ attempts at fooling Bob, i.e., causing Bob to “accept”. Most of the argument, beginning in Section 3.1, is devoted to showing that

$$\begin{aligned} \Pr[\text{Eve fools Bob on first attempt, using } t \text{ copies}] & \quad (9) \\ & \leq (1 - 1/8(t + 1))^5. \quad (10) \end{aligned}$$

In general, Eve learns something from one attempt to the next; however, because Eve can simulate her interaction with Bob at the cost of using one copy of $|\psi_x\rangle$ per simulated iteration of $\mathcal{K}(x)$, we have, for $\ell = 2, 3, \dots, (r - t')$,

$$\begin{aligned} \Pr[\text{Eve fools Bob on } \ell\text{th attempt, using } t \text{ copies}] \\ \leq \Pr[\text{Eve fools Bob on first attempt, using } (t + \ell - 1) \text{ copies}]. \end{aligned}$$

Given this, we use the union bound:

$$\begin{aligned} \Pr[\text{Eve fools Bob at least once, using } t \text{ copies}] \\ & \leq \sum_{\ell=1}^{r-t'} \Pr[\text{Eve fools Bob on } \ell\text{th attempt, using } t \text{ copies}] \\ & \leq \sum_{\ell=1}^{r-t'} \Pr[\text{Eve fools Bob on first attempt, using } (t + \ell - 1) \text{ copies}] \\ & \leq \sum_{\ell=1}^{r-t'} (1 - 1/8(t + \ell))^5 \\ & \leq (r - t')(1 - 1/16r)^5. \end{aligned}$$

It follows that the probability that Eve can fool Bob at least once, that is, break the protocol, is

$$P_{\text{break}} \leq r(1 - 1/16r)^5, \quad (11)$$

which, for fixed r , is exponentially small in s . Note that this bound is likely not tight, since it ultimately assumes that all of Eve's attempts are equally as powerful. In particular, this bound assumes that Eve's copies do not degrade with each use. A more detailed analysis using results about degradation of quantum reference frames [18] may be possible.

From Eq. (11) follows our main theorem:

Theorem 2 (Honest-verifier-security of the protocol). *For any $\epsilon > 0$ and any $r \in \mathbb{Z}^+$, the identification protocol specified in Section 2.1 is secure with error ϵ according to Definition 1 if*

$$s \in \Omega(r \log(r/\epsilon)). \quad (12)$$

The theorem shows how the efficiency of the protocol scales with its reusability. The remainder of the paper establishes the bound in Lines (9) and (10).

3.1. Sufficiency of individual attacks

At each iteration, Eve should take some action, which we may assume takes the form of a measurement, in order to get an answer to send back to Bob. In general, Eve can mount a coherent attack, whereby her actions during iteration j may involve systems that she used or will use in previous or future iterations as well as systems consisting of copies of the state $|\psi_{x_k}\rangle$ for any k —not just for $k = j$. Intuition suggests that, since each x_j is *independently* selected from the set $\{1, 2, \dots, 2r + 1\}$, Eve's measurement at iteration j may be assumed to be independent of her measurement at any other iteration and in particular does not need to act on any components of copies of the public key other than those corresponding to the copies of $|\psi_{x_j}\rangle$. In other words, it seems plausible that Eve's optimal strategy can without loss of generality consist of the “product” of identical optimal strategies for each iteration individually. Indeed, this intuition can be shown to be correct by combining a technique from Ref. [21], for expressing the maximum probability of acceptance in a two-message quantum interactive proof system as a semidefinite program, with a result in Ref. [22], which implies that the semidefinite program satisfies the product rule that we need; see Appendix A for a proof. The remainder of Section 3 establishes that the probability of passing the SWAP-test for any particular iteration is at most $1 - 1/8(t + 1)$, and thus, by the result proved in Appendix A, the probability of passing all s SWAP-tests is at most $(1 - 1/8(t + 1))^s$.

3.2. Equivalence of discrete and continuous private phases

Now, we show that, from Bob's and Eve's points of view, Alice's choosing the private phase angle ϕ_x from the discrete set $\{2\pi x/(2r + 1) : x = 1, 2, \dots, 2r + 1\}$ is equivalent to her choosing the phase angle from the continuous interval $[0, 2\pi)$. We have argued that the only information that Eve or Bob (or anyone but Alice) has about ϕ_x may be assumed to come from a number of copies of $|\psi_x\rangle$ that can be no greater than $2r$ (there are r legitimate copies of the public key, and one can extract r more copies from Alice); let this number be c , where $1 \leq c \leq 2r$. We may describe the state of these c systems by the density operator

$$\frac{1}{2r + 1} \sum_{x=1}^{2r+1} \frac{1}{2^c} ((|0\rangle + e^{2\pi ix/(2r+1)}|1\rangle)(\langle 0| + e^{-2\pi ix/(2r+1)}\langle 1|))^{\otimes c}. \quad (13)$$

Had ϕ_x been chosen uniformly from $\{2\pi x/(2r + 1) : x \in [0, 2r + 1)\} = [0, 2\pi)$, they would describe the state by

$$\frac{1}{2\pi} \int_0^{2\pi} \frac{1}{2^c} ((|0\rangle + e^{i\phi}|1\rangle)(\langle 0| + e^{-i\phi}\langle 1|))^{\otimes c} d\phi. \quad (14)$$

It is straightforward to show that the above two density operators are both equal to

$$\frac{1}{2^c} \sum_{w=0}^c \binom{c}{w} |S_w^c\rangle\langle S_w^c|, \quad (15)$$

where $|S_w^c\rangle$ is the normalized symmetric sum of all $\binom{c}{w}$ states in $\{|0\rangle, |1\rangle\}^{\otimes c}$ whose binary labels have Hamming weight w .⁵ Thus, without loss, we may drop the subscript “ x ” on “ ϕ_x ”, write “ ϕ ” for Alice’s private phase angle, and assume she did (somehow) choose ϕ uniformly randomly from $[0, 2\pi)$.⁶

3.3. Sufficiency of maximizing successful guessing probability

The security of our protocol follows from a result of Bartlett et al. [18], which concerns a slightly different problem for Eve than the problem of her trying to cheat (fool) Bob. This different problem is for Eve, using her t copies of $|0\rangle + e^{i\phi}|1\rangle$, to guess whether she has been given the system $|0\rangle + e^{i\phi}|1\rangle$ or $|0\rangle - e^{i\phi}|1\rangle$, where each case occurs with equal probability and ϕ is unknown and uniformly randomly chosen from $[0, 2\pi)$. The purpose of this section is to show that any good cheating strategy gives a good guessing strategy; we will show that an upper bound on the average successful guessing probability gives an upper bound on the cheating probability, so that, in order to prove security, it suffices to show that the maximum successful guessing probability is sufficiently small (which we will do in the next section).

Any cheating strategy of Eve can be modeled as follows. Let

$$|\pm\rangle := |0\rangle \pm e^{i\phi}|1\rangle. \quad (18)$$

Recall that Bob creates a system in the state $|0\rangle|1\rangle + |1\rangle|0\rangle$, which equals $|+\rangle|+\rangle - |-\rangle|-\rangle$ up to global phase. Eve’s system before Bob sends one of his registers can be represented by $|\mathcal{E}\rangle$, which consists of the t copies of $|+\rangle$ as well as any ancillary registers (which we can assume are in a pure state). Eve’s (optimal) measurement can thus be modeled by a unitary operation U_E acting on her system (labeled E , which now includes the qubit Bob sends), which transforms the state of the total system as follows:

$$\frac{1}{\sqrt{2}}(|+\rangle_B|+\rangle_E - |-\rangle_B|-\rangle_E)|\mathcal{E}\rangle_E \quad (19)$$

$$\mapsto U_E \frac{1}{\sqrt{2}}(|+\rangle_B(\alpha|0\rangle_E|\psi_0^+\rangle_E + \beta|1\rangle_E|\psi_1^+\rangle_E) \quad (20)$$

$$- |-\rangle_B(\gamma|0\rangle_E|\psi_0^-\rangle_E + \delta|1\rangle_E|\psi_1^-\rangle_E)), \quad (21)$$

so that the leftmost register of Eve’s system encodes her measurement outcome. Bob’s application of the Z gate conditioned on the value of the measurement outcome can be modeled by a controlled- Z gate (where Bob’s kept qubit, labeled B , is the target-qubit and the leftmost qubit of Eve’s system is the control-qubit), which will take the state of the total system to

$$\frac{1}{\sqrt{2}}(|+\rangle_B(\alpha|0\rangle_E|\psi_0^+\rangle_E - \delta|1\rangle_E|\psi_1^-\rangle_E) \quad (22)$$

$$+ |-\rangle_B(\beta|1\rangle_E|\psi_1^+\rangle_E - \gamma|0\rangle_E|\psi_0^-\rangle_E)). \quad (23)$$

Let τ represent the density operator for this state after Eve’s system has been traced out. The probability that Bob’s SWAP-test passes is easily calculated to be

$$P_{\text{pass}} = \frac{1 + \langle +|\tau|+\rangle}{2} \quad (24)$$

$$= \frac{1 + (|\alpha|^2 + |\delta|^2)/2}{2}. \quad (25)$$

Now, suppose Eve is faced with the different problem of guessing whether Bob gave her $|+\rangle$ or $|-\rangle$, where each case occurs with probability $1/2$ (and where ϕ is unknown and uniformly random in $[0, 2\pi)$). Since, as can be seen from the mapping in Line (19), U_E maps

⁵ This requires the following two facts: (1) for any integer a ,

$$\frac{1}{2\pi} \int_0^{2\pi} e^{ia\theta} d\theta = \begin{cases} 0 & \text{if } a \neq 0, \\ 1 & \text{otherwise;} \end{cases} \quad (16)$$

and (2) for any integer $p \geq 2$ and integer a :

$$\frac{1}{p} \sum_{k=1}^p e^{2\pi iak/p} = \begin{cases} 0 & \text{if } a \text{ is not a multiple of } p, \\ 1 & \text{otherwise,} \end{cases} \quad (17)$$

where the second fact is applied at $p = 2r + 1$.

⁶ One way to interpret this result is that even if Alice encodes infinitely many bits into ϕ , it is no better than if she encoded $\lceil \log_2(2r + 1) \rceil$ bits. Note that if Eve performs an optimal phase estimation [23] in order to learn ϕ and then cheat Bob, she can only learn at most $\lfloor \log_2(2r - 1) \rfloor$ bits of ϕ (here, we assume Eve has $2r - 1$ copies of the public key, having left Bob one copy), whereas Alice actually encoded $\lceil \log_2(2r + 1) \rceil$ bits into ϕ .

$$|+\rangle_E |\mathcal{E}\rangle_E \mapsto U_E \alpha |0\rangle_E |\psi_0^+\rangle_E + \beta |1\rangle_E |\psi_1^+\rangle_E \quad (26)$$

$$|-\rangle_E |\mathcal{E}\rangle_E \mapsto U_E \gamma |0\rangle_E |\psi_0^-\rangle_E + \delta |1\rangle_E |\psi_1^-\rangle_E. \quad (27)$$

Eve can use the same procedure she used for her attack in order to guess which state Bob prepared: upon measuring her leftmost register, she guesses “|+⟩” if she gets outcome “0”, and otherwise she guesses “|−⟩”. The probability that she guesses successfully on average using this strategy is clearly

$$P_{\text{succ}} = \frac{1}{2} \times \Pr(\text{outcome} = \text{“0”} | \text{Bob prepared } |+\rangle) \quad (28)$$

$$+ \frac{1}{2} \times \Pr(\text{outcome} = \text{“1”} | \text{Bob prepared } |-\rangle) \quad (29)$$

$$= (|\alpha|^2 + |\delta|^2)/2. \quad (30)$$

Thus, any upper bound on P_{succ} gives an upper bound on P_{pass} .

3.4. Bounding the successful guessing probability

Bartlett et al. [18] give an expression for the average successful guessing probability in terms of the state of a bounded phase reference frame (which, for us, takes the form of copies of the public key). In Section 3.4.1, we give some background on (bounded) phase references so that we can better understand the result in Ref. [18] in order to apply it. In Section 3.4.2, we derive the bound on Eve’s successful guessing probability, establishing what was claimed in Lines (9) and (10).

3.4.1. Phase reference frames

Consider the two qubit states “|0⟩ + |1⟩” and “|0⟩ − |1⟩” (the use of quotation marks will become clear below). Given a qubit promised to be prepared in one of those two states, how could you decide which state the given qubit is in? In general, without any other system to help you, you cannot, because the question is not well defined: the states |0⟩ and |1⟩ are only defined up to global phase, so, e.g., replacing |1⟩ by −|1⟩ changes your answer. A *phase reference (frame)* is a quantum-mechanical system that, when taken together with the given qubit, fixes the relative phase between |0⟩ and |1⟩ in the state “|0⟩ + |1⟩”. This intuitive definition suffices for our purposes.

Note that usually in quantum information processing, it is assumed that one has a phase reference, which ascribes definite meaning to the state “|0⟩ + |1⟩”. Similarly, in multiparty quantum communication/cryptography protocols, it is usually assumed that all players involved have access to a common phase reference, so that “|0⟩ + |1⟩” means the same thing to each party. This is in fact a reasonable assumption, since it has been shown that, with modest overhead, any set of players can simulate having a common (phase) reference frame, by using the symmetric and anti-symmetric subspaces of entangled states [24].

The most popular phase reference frame occurs in optics, and is known as a “coherent state”, which is usually defined as the state

$$|C_\theta\rangle := e^{-\frac{\alpha^2}{2}} \sum_{w=0}^{\infty} (\alpha^w / \sqrt{w!}) e^{i w \theta} |w\rangle, \quad (31)$$

of a single (optical) mode, where α is a real number; this state *encodes* the (relative) phase $e^{i\theta}$. The number w in $|w\rangle$ is the *photon number*. Employing a system in this state to first prepare the given qubit and then to measure it, in order to solve the decision problem posed above, fixes the relative phase of |0⟩ and |1⟩ in the superposition “|0⟩ + |1⟩” to be $e^{i\theta}$, i.e., this superposition is now more correctly written $|0\rangle + e^{i\theta}|1\rangle$.⁷ Normally, one would redefine |1⟩ as $|1\rangle := e^{i\theta}|1\rangle$, and write the superposition as $|0\rangle + |1\rangle$.

Note that there is still a slight problem in that θ is not really well defined, since it could be replaced by any $\theta' \neq \theta$ and the physics of the problem would not change. We say that, in the definition of the coherent state, the phase $e^{i\theta}$ is defined relative to a *hidden absolute phase reference*, which in practice means that the actual value of θ need not be known, but what is important is that the relative phase $e^{i\theta}$ stays consistent throughout all the quantum operations.

⁷ The precise way in which this preparation and measurement works in practice is beyond the scope of this paper, but we give an intuitive explanation. First note that it is sufficient to implement the Hadamard gate

$$|0\rangle \mapsto |0\rangle + e^{i\phi}|1\rangle \quad (32)$$

$$e^{i\theta}|1\rangle \mapsto |0\rangle - e^{i\phi}|1\rangle. \quad (33)$$

Since the coherent state is unchanged (up to global phase) under the operation $|w\rangle \mapsto |w-1\rangle$ (i.e. the annihilation of a photon, which we note is not a unitary operation), it is possible to approximate the Hadamard gate by approximately mapping $|C_\theta\rangle|0\rangle \mapsto |C_\theta\rangle(|0\rangle + e^{i\theta}|1\rangle)$ by taking a photon from the coherent state to use in the right-hand qubit; similarly, the operation $|C_\theta\rangle|1\rangle \mapsto |C_\theta\rangle(e^{-i\theta}|0\rangle - |1\rangle)$ may be approximated. The quality of the approximation depends on the total energy of the coherent state. See Ref. [25] for a complete analysis of a similar task.

A single mode is mathematically modeled by \mathbb{C}^N , and a basis for this space is $\{|n\rangle : n = 0, 1, \dots, N-1\}$; in general, as for the coherent state, N can equal ∞ . Similarly, a k -mode (multimode) phase reference is modeled by $(\mathbb{C}^N)^{\otimes k}$. It is convenient to adopt this optics-based nomenclature (modes, photons) when discussing reference frames, though the following results are completely general and do not rely on optical implementations.

In practice, any phase reference frame is *bounded*, meaning that the total energy, or average total photon number, of the state is upper-bounded. For a multimode phase reference state, this bound may take the form of an upper bound on the number k of modes and perhaps an upper bound on the photon number of each mode. Generally, the higher the total energy of the phase reference frame, the better it performs in practice, i.e., the better it maintains consistent relative phase throughout a quantum computation.

An example of a k -mode bounded phase reference frame encoding the phase $e^{i\theta}$ is a system of k qubits in the state

$$(|0\rangle + e^{i\theta}|1\rangle)^{\otimes k} = \sum_{w=0}^k \sqrt{\binom{k}{w}} e^{iw\theta} |S_w^k\rangle, \quad (34)$$

where $|S_w^k\rangle$ is the k -mode state defined just after Eq. (15) (see Ref. [26] for a detailed discussion of such “refbits”). For this multimode phase reference frame, the maximum photon number is 1 for each mode.

Define the unitary re-phasing map $U(\theta)$ on $(\mathbb{C}^{k+1})^{\otimes k}$ as the mapping

$$|w_1\rangle|w_2\rangle \cdots |w_k\rangle \mapsto e^{i(w_1+w_2+\cdots+w_k)\theta} |w_1\rangle|w_2\rangle \cdots |w_k\rangle \quad (35)$$

for any $\theta \in [0, 2\pi]$ and all $w_l = 0, 1, \dots, k$, for $l = 1, 2, \dots, k$. A unitary operation V on $(\mathbb{C}^k)^{\otimes k}$ is said to be *phase invariant* if $U(\theta)VU(\theta)^\dagger = V$ for all $\theta \in [0, 2\pi]$. If V is phase invariant, one does not need any phase reference to perform V ; e.g., Eve could use her own phase reference (say, a coherent state encoding the phase $e^{i\theta_E}$) to carry out V on some register, and the result would be the same (up to global phase) as if Alice performed V on the same register using her own phase reference (a coherent state encoding the phase $e^{i\theta_A}$, $\theta_A \neq \theta_E$); for simplicity, we have assumed that Eve’s and Alice’s phase references are perfect (see Section II.B of Bartlett et al. [27] for more details). We will use the fact that $U(\theta)$ is phase invariant if and only if it is block-diagonal with respect to subspaces of constant total photon number (see e.g. Ref. [25] for a proof).

Definition 3 (*Equivalence of phase reference frames*). Suppose ρ_1 and ρ_2 are two states of a multimode phase reference frame. Then a phase reference in state ρ_1 and a phase reference in state ρ_2 are *equivalent* if there exists a phase-invariant unitary operation V such that

$$\rho_1 = V\rho_2V^\dagger. \quad (36)$$

Note that there is a phase-invariant unitary transformation on $(\mathbb{C}^{k+1})^{\otimes k}$ that maps

$$|S_w^k\rangle \mapsto |0\rangle^{\otimes(k-1)}|w\rangle, \quad \text{for all } w = 0, 1, \dots, k, \quad (37)$$

because this mapping may be completed on $(\mathbb{C}^{k+1})^{\otimes k}$ to a unitary operator that is block-diagonal with respect to subspaces of constant total photon number. Therefore, a multimode phase reference in the state in Line (34) is equivalent to a single-mode phase reference in the state

$$\sum_{w=0}^k \sqrt{\binom{k}{w}} e^{iw\theta} |w\rangle \quad (38)$$

(where we have omitted the ancilla in the state $|0\rangle^{\otimes(k-1)}$), which we note looks like the coherent state but for the moduli of the coefficients.

Finally, we define a special class of phase references.

Definition 4 (*Covariant family of phase reference states*). Let $U(\theta)$ be the unitary rephasing map on \mathbb{C}^N such that $U(\theta)|w\rangle = e^{iw\theta}|w\rangle$ for all $\theta \in [0, 2\pi]$. Suppose that $\{\rho(\phi)\}_{\phi \in [0, 2\pi]}$ is a family of (single-mode) phase reference states on \mathbb{C}^N , where $\rho(\phi)$ encodes the relative phase $e^{i\phi}$. Then $\{\rho(\phi)\}_{\phi \in [0, 2\pi]}$ is *covariant* if

$$\rho(\phi) = U(\phi)\rho(0)U(\phi)^\dagger \quad (39)$$

for all $\phi \in [0, 2\pi]$.

3.4.2. The bound

First we note that the effect of Alice’s selection of ϕ serves to completely randomize the relative phase between $|0\rangle$ and $|1\rangle$ in any superposition of the two states, from the point of view of anyone other than Alice. Thus, even though we make the usual assumption that all players (Alice, Bob, Eve) share a common phase reference, the protocol effectively forces Alice to have a private phase reference, leaving the other players with maximal ignorance (but for the information contained in the copies of the public key) of what the “correct” relative phase is in each iteration of the protocol. Therefore, Eve’s t copies of $|0\rangle + e^{i\phi}|1\rangle$ may be seen as a bounded multimode phase reference encoding the phase $e^{i\phi}$ relative to a *known* common phase reference—but for *unknown* and uniformly random $\phi \in [0, 2\pi)$.⁸

Bartlett et al. [18] prove the following theorem (rephrased for our purposes).

Theorem 5 (Optimal probability of successful guessing). (See [18].) Suppose $\{\rho(\phi)\}_{\phi \in [0, 2\pi]}$ is a covariant family of single-mode phase reference states, where $\rho(\phi)$ encodes the phase $e^{i\phi}$ (relative to a known common phase reference). Given a single mode in the state $\rho(\phi)$, for ϕ unknown and uniformly random in $[0, 2\pi)$, and a qubit in one of the two states $|0\rangle \pm e^{i\phi}|1\rangle$, where each state occurs with probability $1/2$, the optimal probability of successfully guessing which state the given qubit is in is

$$P_{\text{succ}} = \frac{1}{2} + \frac{1}{2} \sum_{m=0}^{\infty} \Re(\langle m+1 | \rho(0) | m \rangle). \quad (40)$$

We showed in the previous section that Eve’s multimode phase reference is equivalent to the single-mode phase reference in the state

$$\rho(\phi) := \frac{1}{2^t} \sum_{w=0}^t \sum_{w'=0}^t \sqrt{\binom{t}{w} \binom{t}{w'}} e^{i(w-w')\phi} |w\rangle \langle w'|. \quad (41)$$

Substituting this value of $\rho(\phi)$ into the above theorem gives

$$P_{\text{succ}} = \frac{1}{2} + \frac{1}{2} \frac{1}{2^t} \sum_{m=0}^{t-1} \sqrt{\binom{t}{m} \binom{t}{m+1}}, \quad (42)$$

which we can show to be in $1 - \Omega(1/t)$ (up to logarithmic factors) using some simple approximations. Cheung [28] has improved our asymptotic bound on this quantity by showing that

$$\frac{1}{2^t} \sum_{m=0}^{t-1} \sqrt{\binom{t}{m} \binom{t}{m+1}} \leq 1 - \frac{1}{2(t+1)} - \frac{1}{2^{t+1}}, \quad (43)$$

which implies

$$P_{\text{succ}} \leq 1 - 1/4(t+1). \quad (44)$$

It follows that

$$P_{\text{pass}} \leq 1 - 1/8(t+1), \quad (45)$$

which we recall is an upper bound on the probability that Bob’s SWAP-test passes in any particular kernel-iteration, when Eve is acting as a dishonest prover and using t copies of the public key. Thus, as we argued in Section 3.1 (and Appendix A), the total probability that Eve causes all s of Bob’s SWAP-tests to pass is

$$\begin{aligned} & \Pr[\text{Eve fools Bob on first attempt, using } t \text{ copies}] \\ & \leq (1 - 1/8(t+1))^s, \end{aligned}$$

as claimed in Lines (9) and (10). This completes the proof of security of the protocol.

Acknowledgements

We acknowledge useful discussions with Giulio Chiribella, Daniel Gottesman, Rob Spekkens, and, in particular, John Watrous, who gave us the argument in Appendix A that proves that the protocol is secure against coherent attacks. L.M. Ioannou was supported by EPSRC, SCALA, QuantumWorks, Mitacs, and IQC. M. Mosca was partially supported by NSERC, DTO-ARO, CFI, CIFAR, Ontario-MRI, CRC, OCE, QuantumWorks and Mitacs.

⁸ Note that our assumption that Alice, Bob, and Eve all share a perfect common phase reference implies that Alice can make r samples of her phase reference with no degradation of the original phase reference. Thus, while we are using the theoretical framework of (bounded) quantum reference frames from Refs. [1,18] in our analysis, our initial assumptions are different than in those works (where the standard assumption that everyone shares a perfect common phase reference is usually not used).

Appendix A

Consider the following non-cryptographic, two-message interactive protocol (or game) between Evelyn and Bobby (neither of whom is considered adversarial, hence we distinguish these two players from Eve and Bob), denoted $\mathcal{L} = \mathcal{L}(\Phi)$, where Φ is a quantum operation (super-operator) that specifies Evelyn's action in Step 2' below (the quantities r and t are as defined previously):

- (1') Bobby chooses a uniformly random $x \in \{1, 2, \dots, 2r + 1\}$ and creates a $(t + 3)$ -qubit system in the state $|\psi_x\rangle^{\otimes(t+1)}(|0\rangle|1\rangle + |1\rangle|0\rangle)$; Bobby sends to Evelyn t copies of $|\psi_x\rangle$ as well as one qubit of the system in the state $|0\rangle|1\rangle + |1\rangle|0\rangle$.
- (2') Evelyn carries out the quantum operation Φ on the received qubits which outputs one qubit, which Evelyn sends to Bobby.
- (3') Bobby measures the received qubit in the computational basis $\{|0\rangle, |1\rangle\}$; if the measurement outcome is "1", then he applies the Pauli-Z gate

$$Z := \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \quad (\text{A.1})$$

to the qubit of the system in the state $|0\rangle|1\rangle + |1\rangle|0\rangle$ that he kept in Step 1'. Finally, Bobby SWAP-tests this qubit with his remaining copy of $|\psi_x\rangle$.

The following proposition is immediate:

Proposition 6. *The probability that Eve causes Bob's SWAP-test to pass in any particular iteration of the protocol in Section 2.1 is at most*

$$\alpha := \max_{\Phi} \Pr[\text{Bobby's SWAP-test passes in } \mathcal{L}(\Phi)], \quad (\text{A.2})$$

where Φ ranges over all admissible quantum operations that Evelyn can apply in Step 2'.

Now consider the parallel s -fold repetition of \mathcal{L} , which we denote $\mathcal{L}^{\parallel s} = \mathcal{L}^{\parallel s}(\Phi')$, where now Φ' is Evelyn's quantum operation in the second step of $\mathcal{L}^{\parallel s}$. The following proposition is also immediate:

Proposition 7. *The probability that Eve fools Bob on the first attempt using t copies in the protocol in Section 2.1 is at most*

$$\alpha' := \max_{\Phi'} \Pr[\text{all of Bobby's SWAP-tests pass in } \mathcal{L}^{\parallel s}(\Phi')], \quad (\text{A.3})$$

where Φ' ranges over all admissible quantum operations that Evelyn can apply in the second step of $\mathcal{L}^{\parallel s}$.

Therefore, in order to prove that it is sufficient to consider individual (as opposed to coherent) attacks by Eve, it suffices to show that $\alpha' = \alpha^s$.

In Ref. [21], it is shown that the maximum acceptance probability of any two-message interactive proof system can be expressed as a semidefinite (optimization) program (see Ref. [29] for a relevant review of semidefinite programming). Before we apply this fact, we need to make some definitions. Let \mathcal{X} and \mathcal{Y} be the input and output spaces, respectively, of Evelyn's quantum operation Φ in \mathcal{L} , i.e. $\Phi : L(\mathcal{X}) \rightarrow L(\mathcal{Y})$, where $L(\mathcal{X})$ is the space of all linear operators from the complex Euclidean space \mathcal{X} to itself (and likewise for $L(\mathcal{Y})$). Let $\text{Pos}(\mathcal{Y} \otimes \mathcal{X})$ denote the set of all positive semidefinite operators in $L(\mathcal{Y} \otimes \mathcal{X})$. Similarly, for $\mathcal{L}^{\parallel s}$, we have that $\Phi' : L(\mathcal{X}^{\otimes s}) \rightarrow L(\mathcal{Y}^{\otimes s})$. Viewing Bobby's SWAP-test passing as "acceptance" in an interactive proof system, we thus have, according to Ref. [21], that α and α' can be expressed, respectively, as solutions to the following semidefinite programs π_{α} and $\pi_{\alpha'}$:

$$\begin{array}{ll} \text{maximize:} & \text{Tr}(B^{\dagger}X) \\ \text{subject to:} & \text{Tr}_{\mathcal{Y}}(X) = \mathbb{I}_{\mathcal{X}}, \\ & X \in \text{Pos}(\mathcal{Y} \otimes \mathcal{X}) \end{array} \quad \begin{array}{ll} \text{maximize:} & \text{Tr}((B^{\otimes s})^{\dagger}X') \\ \text{subject to:} & \text{Tr}_{\mathcal{Y}^{\otimes s}}(X') = I_{\mathcal{X}^{\otimes s}}, \\ & X' \in \text{Pos}((\mathcal{Y} \otimes \mathcal{X})^{\otimes s}), \end{array}$$

where X and X' are the Choi–Jamiołkowski representations of Φ and Φ' , and B is a positive semidefinite operator representing Bobby's actions, i.e., $B \in \text{Pos}(\mathcal{Y} \otimes \mathcal{X})$. Furthermore, it is shown in Ref. [21] that such semidefinite programs (arising from two-message interactive proof systems) satisfy the condition of strong duality, which means that the solution to each semidefinite program above coincides with that of its dual.

In Ref. [22], the following theorem is proven:

Theorem 8. (See [22].) Suppose that the following two semidefinite programs π_1 and π_2 satisfy strong duality:

$$\begin{array}{ll} \underline{\pi_1} & \underline{\pi_2} \\ \text{maximize: } & \text{Tr}(J_1^\dagger X) & \text{maximize: } & \text{Tr}(J_2^\dagger X) \\ \text{subject to: } & \Phi_1(X) = C_1, & \text{subject to: } & \Phi_2(X) = C_2, \\ & X \in \text{Pos}(\mathcal{X}_1) & & X \in \text{Pos}(\mathcal{X}_2), \end{array}$$

where $\Phi_1 : L(\mathcal{X}_1) \rightarrow L(\mathcal{Y}_1)$ and $\Phi_2 : L(\mathcal{X}_2) \rightarrow L(\mathcal{Y}_2)$, for complex Euclidean spaces $\mathcal{X}_1, \mathcal{Y}_1, \mathcal{X}_2, \mathcal{Y}_2$, and $J_1 \in L(\mathcal{X}_1)$ and $J_2 \in L(\mathcal{X}_2)$ are Hermitian. Let $\alpha(\pi_1)$ and $\alpha(\pi_2)$ denote the semidefinite programs' solutions. If J_1 and J_2 are positive semidefinite, then the solution to the following semidefinite program, denoted $\pi_1 \otimes \pi_2$, is $\alpha(\pi_1 \otimes \pi_2) = \alpha(\pi_1)\alpha(\pi_2)$:

$$\begin{array}{l} \underline{\pi_1 \otimes \pi_2} \\ \text{maximize: } \text{Tr}((J_1 \otimes J_2)^\dagger X) \\ \text{subject to: } \Phi_1 \otimes \Phi_2(X) = C_1 \otimes C_2, \\ X \in \text{Pos}(\mathcal{X}_1 \otimes \mathcal{X}_2). \end{array}$$

Since B is positive semidefinite and $\pi_{\alpha'} = \pi_{\alpha}^{\otimes s}$ (using the associativity of \otimes), Theorem 8 can be applied $s - 1$ times in order to prove that $\alpha' = \alpha^s$ as required. See Ref. [30] for a similar approach, based on ideas in Ref. [31].

Note that this argument, combined with the arguments in the main body of the paper, shows that both the serial and parallel versions of our identification protocol are secure.

References

- [1] S.D. Bartlett, T. Rudolph, R.W. Spekkens, Decoherence-full subsystems and the cryptographic power of a private shared reference frame, *Phys. Rev. A* 70 (2004) 032307.
- [2] C.H. Bennett, G. Brassard, Quantum cryptography: public key distribution and coin tossing, in: *Proceedings of IEEE International Conference on Computers, Systems and Signal Processing*, IEEE Press, New York, 1984, pp. 175–179.
- [3] A.K. Ekert, Quantum cryptography based on Bell's theorem, *Phys. Rev. Lett.* 67 (6) (1991) 661–663.
- [4] H. Barnum, C. Crépeau, D. Gottesman, A. Smith, A. Tapp, Authentication of quantum messages, in: I. Press (Ed.), *Proc. 43rd Annual IEEE Symposium on the Foundations of Computer Science (FOCS'02)*, 2002, pp. 449–458.
- [5] I. Damgaard, S. Fehr, L. Salvail, C. Schaffner, Cryptography in the bounded quantum-storage model, in: I. Press (Ed.), *Proceedings of the 46th IEEE Symposium on Foundations of Computer Science – FOCS 2005*, 2005, pp. 449–458.
- [6] I. Damgaard, S. Fehr, L. Salvail, C. Schaffner, Secure identification and QKD in the bounded-quantum-storage model, in: *CRYPTO 2007*, in: *Lecture Notes in Computer Science*, vol. 4622, 2007, pp. 342–359.
- [7] C. Schaffner, Cryptography in the bounded-quantum-storage model, Ph.D. thesis, University of Aarhus, 2007.
- [8] R. Koenig, S. Wehner, J. Wullschlegler, Unconditional security from noisy quantum storage, arXiv:0906.1030, 2009.
- [9] C. Schaffner, Simple protocols for oblivious transfer and secure identification in the noisy-quantum-storage model, arXiv:1002.1495, 2010.
- [10] C. Cachin, U.M. Maurer, Unconditional security against memory-bounded adversaries, in: *CRYPTO'97: Proceedings of the 17th Annual International Cryptology Conference on Advances in Cryptology*, 1997.
- [11] D. Gottesman, I.L. Chuang, Quantum digital signatures, arXiv:quant-ph/0105032, 2001.
- [12] A.J. Menezes, P. van Oorschot, S. Vanstone, *Handbook of Applied Cryptography*, CRC Press LLC, Boca Raton, 1996.
- [13] A. Kawachi, T. Koshihara, H. Nishimura, T. Yamakami, Computational indistinguishability between quantum states and its cryptographic application, in: *Advances in Cryptology EUROCRYPT 2005*, in: *Lecture Notes in Computer Science*, vol. 3494, Springer, 2005, pp. 268–284, full version at <http://arxiv.org/abs/quant-ph/0403069>.
- [14] M. Hayashi, A. Kawachi, H. Kobayashi, Quantum measurements for hidden subgroup problems with optimal sample complexity, *Quantum Inf. Comput.* 8 (2008) 0345–0358.
- [15] D. Gottesman, Quantum public key cryptography with information-theoretic security, in: *Workshop on Classical and Quantum Information Security*, Caltech, 15–18 December, 2005, <http://www.cpi.caltech.edu/quantum-security/program.html>, see also <http://www.perimeterinstitute.ca/personal/dgottesman>.
- [16] S. Kak, A three-stage quantum cryptography protocol, *Found. Phys. Lett.* 19 (2006) 293–296.
- [17] G.M. Nikolopoulos, Applications of single-qubit rotations in quantum public-key cryptography, *Phys. Rev. A* 77 (2008) 032348, see also *Phys. Rev. A* 78, 019903.
- [18] S.D. Bartlett, T. Rudolph, R.W. Spekkens, P.S. Turner, Degradation of a quantum reference frame, *New J. Phys.* 8 (2006) 58.
- [19] T. Okamoto, K. Tanaka, S. Uchiyama, Quantum public-key cryptosystems, in: *Proc. of CRYPTO 2000*, in: *Lecture Notes in Computer Science*, vol. 1880, Springer-Verlag, 2000, pp. 147–165.
- [20] O. Goldreich, *Foundations of Cryptography (Volume 1): Basic Tools*, Cambridge University Press, Cambridge, 2001.
- [21] R. Jain, S. Upadhyay, J. Watrous, Two-message quantum interactive proofs are in PSPACE, in: *Annual IEEE Symposium on Foundations of Computer Science*, 2009, pp. 534–543.
- [22] R. Mittal, M. Szegedy, Product rules in semidefinite programming, in: E. Csuhaj-Varjú, Z. Ésik (Eds.), *FCT*, in: *Lecture Notes in Computer Science*, vol. 4639, Springer, 2007, pp. 435–445.
- [23] W. van Dam, G.M. D'Ariano, A. Ekert, C. Macchiavello, M. Mosca, Optimal phase estimation in quantum networks, *J. Phys. A, Math. Theoret.* 40 (2007) 7971–7984.
- [24] S.D. Bartlett, T. Rudolph, R.W. Spekkens, Classical and quantum communication without a shared reference frame, *Phys. Rev. Lett.* 91 (2003) 027901.
- [25] L.M. Ioannou, M. Mosca, Universal quantum computation in a hidden basis, *Quantum Inf. Comput.* (2014), in press, <http://arxiv.org/abs/0810.2780>.

- [26] S.J. van Enk, Quantifying the resource of sharing a reference frame, *Phys. Rev. A* 71 (2005) 032339.
- [27] S.D. Bartlett, T. Rudolph, R.W. Spekkens, Reference frames, superselection rules, and quantum information, *Rev. Modern Phys.* 79 (2007) 555.
- [28] D. Cheung, 2009, Unpublished notes.
- [29] J. Watrous, Theory of quantum information, Lecture notes for course CS 789, University of Waterloo, available at <http://www.cs.uwaterloo.ca/~watrous/>, 2008.
- [30] G. Gutoski, Quantum strategies and local operations, Ph.D. thesis, University of Waterloo, 2009.
- [31] R. Cleve, W. Slofstra, F. Unger, S. Upadhyay, Strong parallel repetition theorem for quantum XOR proof systems, arXiv:quant-ph/0608146v1, 2006.