

Secure Two-Party Quantum Evaluation of Unitaries against Specious Adversaries

Frédéric Dupuis^{1,*}, Jesper Buus Nielsen², and Louis Salvail^{3,**}

¹ Institute for Theoretical Physics, ETH Zurich, Switzerland
dupuis@phys.ethz.ch

² DAIMI, Aarhus University, Denmark
jbn@cs.au.dk

³ Université de Montréal (DIRO), QC, Canada
salvail@iro.umontreal.ca

Abstract. We describe how any two-party quantum computation, specified by a unitary which simultaneously acts on the registers of both parties, can be privately implemented against a quantum version of classical semi-honest adversaries that we call specious. Our construction requires two ideal functionalities to guarantee privacy: a private SWAP between registers held by the two parties and a classical private AND-box equivalent to oblivious transfer. If the unitary to be evaluated is in the Clifford group then only one call to SWAP is required for privacy. On the other hand, any unitary not in the Clifford requires one call to an AND-box per R-gate in the circuit. Since SWAP is itself in the Clifford group, this functionality is universal for the private evaluation of any unitary in that group. SWAP can be built from a classical bit commitment scheme or an AND-box but an AND-box cannot be constructed from SWAP. It follows that unitaries in the Clifford group are to some extent the easy ones. We also show that SWAP cannot be implemented privately in the bare model.

1 Introduction

In this paper, we address the problem of privately evaluating some unitary transform U upon a joint quantum input state held by two parties. Since unitaries model what quantum algorithms are implementing, we can see this problem as a natural extension of secure two-party evaluation of functions to the quantum realm. Suppose that a state $|\phi_{\text{in}}\rangle \in \mathcal{A} \otimes \mathcal{B}$ is the initial shared state where Alice holds register \mathcal{A} and Bob holds register \mathcal{B} . Let $U \in U(\mathcal{A} \otimes \mathcal{B})$ be some unitary transform acting upon \mathcal{A} and \mathcal{B} . What cryptographic assumptions are needed for a private evaluation of $|\phi_{\text{out}}\rangle = U|\phi_{\text{in}}\rangle$ where *private* means that each player learns no more than in the ideal situation depicted in Fig. 1? Of course, answers to this question depend upon the adversary we are willing to tolerate.

* Supported by Canada's NSERC Postdoctoral Fellowship Program.

** Supported by Canada's NSERC discovery grant, MITACS, and the QuantumWorks networks(NSERC).

In [17], it was shown that unitaries cannot be used to implement classical cryptographic primitives. Any non-trivial primitive implemented by unitaries will necessarily leak information toward one party. Moreover, this leakage is available to a weak class of adversaries that can be interpreted as the quantum version of classical semi-honest adversaries. It follows that quantum two-party computation of unitaries cannot be used to implement classical cryptographic primitives. This opens the possibility that the cryptographic assumptions needed for private evaluations of unitaries are weaker than for their classical counterpart. So, what classical cryptographic assumptions, if any, are required to achieve privacy in our setting? Are there unitaries more difficult to evaluate privately than others?




Fig. 1. Ideal Functionality for unitary U

In this work, we answer these questions against a class of weak quantum adversaries, called specious, related to classical semi-honest adversaries. We say that a quantum adversary is specious if at any step during the execution of a protocol, it can provide a judge with some state that, when joined with the state held by the honest player, will be indistinguishable from a honest interaction. In other words, an adversary is specious if it can pass an audit with success at any step. Most known impossibility proofs in quantum cryptography apply when the adversary is restricted to be specious. Definitions similar to ours have been proposed for the quantum setting and usually named semi-honest. However, translating our definition to the classical setting produces a strictly stronger class of adversaries than semi-honest¹ which justifies not adopting the term *semi-honest*. We propose the name *specious* as the core of the definition is that the adversary must appear to act honestly.

Contributions. First, we define two-party protocols for the evaluation of unitaries having access to oracle calls. This allows us to consider protocols with security relying on some ideal functionalities in order to be private. We then say that a protocol is in the *bare model* if it does not involve any call to an ideal functionality. We then formally define what we mean by specious adversaries. Privacy is then defined via simulation. We say that a protocol for the two-party evaluation of unitary U is private against specious adversaries if, for any joint input state and at any step of the protocol, there exists a simulator that can reproduce the adversary’s view having only access to its own part of the joint input state. Quantum simulation must rely on a family of simulators for the view of the adversary rather than one because quantum information does not accumulate but can vanish as the protocol evolves. For instance, consider the

¹ As an example, assume there exist public key cryptosystems where you can sample a public key without learning the secret key. Then this is a semi-honest oblivious transform: The receiver, with choice bit c , samples pk_c in the normal way and learns its corresponding secret key and samples pk_{1-c} without learning its secret key. He sends (pk_0, pk_1) . Then the sender sends $(E_{pk_0}(m_0), E_{pk_1}(m_1))$ and the receiver decrypts $E_{pk_c}(m_c)$. This is not secure against a specious adversary who can sample pk_{1-c} along with its secret key sk_{1-c} and then delete sk_{1-c} before the audit.

trivial protocol that let Alice send her input register to Bob so that he can apply locally $|\phi_{\text{out}}\rangle = U|\phi_{\text{in}}\rangle$ before returning her register. The final state of such a protocol is certainly private, as Bob cannot clone Alice's input and keep a copy, yet at some point Bob had access to Alice's input thus violating privacy. No simulator can possibly reproduce Bob's state after he received Alice's register without having access to her input state.

Second, we show that no protocol can be shown statistically private against specious adversaries in the bare model for a very simple unitary: the swap gate. As the name suggests, the swap gate simply permutes Alice's and Bob's input states. Intuitively, the reason why this gate is impossible is that at some point during the execution of such protocol, one party that still has almost all its own input state receives a non-negligible amount of information (in the quantum sense) about the other party's input state. At this point, no simulator can possibly re-produce the complete state held by the receiving party since a call to the ideal functionality only provides access to the other party's state while no call to the ideal functionality only provides information about that party's own input. Therefore, any simulator cannot re-produce a state that contains information about the input states of both parties. It follows that cryptographic assumptions are needed for the private evaluation of unitaries against specious adversaries. On the other hand, a classical bit commitment is sufficient to implement the swap privately in our model.

Finally, we give a very simple protocol for the private evaluation of any unitary based on ideas introduced by [7,6] in the context of fault tolerant quantum computation. Our construction is similar to Yao's original construction in the classical world[22,9]. We represent any unitary U by a quantum circuit made out of gates taken from the universal set $\mathcal{UG} = \{X, Y, Z, \text{CNOT}, \text{H}, \text{P}, \text{R}\}$ [13]. The protocol evaluates each gate of the circuit upon shared encrypted input where the encryption uses the Pauli operators $\{X, Y, Z\}$ together with the identity. In addition to the Pauli gates X, Y , and Z , gates CNOT, H, and P can easily be performed over encrypted states without losing the ability to decrypt. Gates of that kind belong to what is called the *Clifford group*. The CNOT gate is the only gate in \mathcal{UG} acting upon more than one qubit while the R-gate is the only one that does not belong to the Clifford group. In order to evaluate it over an encrypted state while preserving the ability to decrypt, we need to rely upon a classical ideal functionality computing securely an additive sharing for the AND of Alice's and Bob's input bits. We call this ideal functionality an AND-box. Upon input $x \in \{0, 1\}$ for Alice and $y \in \{0, 1\}$ for Bob, it produces $a \in_R \{0, 1\}$ and $b \in \{0, 1\}$ to Alice and Bob respectively such that $a \oplus b = x \wedge y$. An AND-box can be obtained from any flavor of oblivious transfer and is defined the same way than an NL-box[14,15] without the property that its output can be obtained before the input of the other player has been provided to the box (i.e., NL-boxes are non-signaling). The *equivalence* between AND-boxes, NL-boxes, and oblivious transfer is discussed in [21]. At the end of the protocol, each part of the shared key allowing to decrypt the output must be exchanged in a fair way. For this task, Alice and Bob rely upon an ideal swap functionality called

SWAP. The result is that any U can be evaluated privately upon any input provided Alice and Bob have access to one AND-box per R-gate and one call to the an ideal swap. If the circuit happens to have only gates in the Clifford group then only one call to an ideal swap is required for privacy. In other words, SWAP is universal for the private evaluation of circuits in the Clifford group (i.e., those circuits having no R-gate) and itself belongs to that group (SWAP is not a classical primitive). To some extent, circuits in the Clifford group are the *easy* ones. Privacy for circuits containing R-gates however needs a classical cryptographic primitive to be evaluated privately by our protocol. It means that AND-boxes are universal for the private evaluation of any circuit against specious adversaries. We don't know whether there exist some unitary transforms that are universal for the private evaluation of any unitary against specious adversaries.

Previous works. All impossibility results in quantum cryptography we are aware of apply to classical primitives. In fact, the impossibility proofs usually rely upon the fact that an adversary with a seemingly honest behavior can force the implementation of classical primitives to behave quantumly. The result being that implemented that way, the primitive must leak information to the adversary. This is the spirit behind the impossibility of implementing oblivious transfer securely using quantum communication[10]. In that same paper the impossibility of any one-sided private evaluation of non-trivial primitives was shown. All these results can be seen as generalizations of the impossibility of bit commitment schemes based on quantum communication[11,12]. The most general impossibility result we are aware of applies to any *non-trivial* two-party classical function[17]. It states that it suffices for the adversary to *purify* its actions in order for the quantum primitive to leak information. An adversary purifying its actions is specious as defined above. None of these impossibility proofs apply to quantum primitives characterized by some unitary transform applied to joint quantum inputs. Blind quantum computation is a primitive that shows similarities to ours. In [4], a protocol allowing a client to get its input to a quantum circuit evaluated blindly has been proposed. The security of their scheme is unconditional while in our setting almost no unitary allows for unconditional privacy.

An unpublished work of Smith[19] shows how one can devise a private protocol for the evaluation of any unitary that seems to remain private against all quantum adversaries. However, the techniques used require strong cryptographic assumptions like homomorphic encryption schemes, zero-knowledge and witness indistinguishable proof systems. The construction is in the spirit of protocols for multiparty quantum computation[3,5] and fault tolerant quantum circuits[18,1]. Although our protocol only guarantees privacy against specious adversaries, it is obtained using much weaker cryptographic assumptions.

2 Preliminaries

The N -dimensional complex Euclidean space (i.e., Hilbert space) will be denoted by \mathcal{H}_N . We denote quantum registers using calligraphic typeset \mathcal{A} . As usual, $\mathcal{A} \otimes \mathcal{B}$

denotes the space of two such quantum registers. We write $\mathcal{A} \approx \mathcal{B}$ when \mathcal{A} and \mathcal{B} are such that $\dim(\mathcal{A}) = \dim(\mathcal{B})$. A register \mathcal{A} can undergo transformations as a function of time; we denote by \mathcal{A}_i the state of space \mathcal{A} at time i . When a quantum computation is viewed as a circuit accepting input in \mathcal{A} , we denote all wires in the circuit by $\mathbf{w} \in \mathcal{A}$. If the circuit accepts input in $\mathcal{A} \otimes \mathcal{B}$ then the set of all wires is denoted $\mathbf{w} \in \mathcal{A} \cup \mathcal{B}$.

The set of all linear mappings from \mathcal{A} to \mathcal{B} is denoted by $L(\mathcal{A}, \mathcal{B})$ while $L(\mathcal{A})$ stands for $L(\mathcal{A}, \mathcal{A})$. To simplify notation, for $\rho \in L(\mathcal{A})$ and $M \in L(\mathcal{A}, \mathcal{B})$ we write $M \cdot \rho$ for $M\rho M^\dagger$. We denote by $\text{Pos}(\mathcal{A})$ the set of positive semi-definite operators in \mathcal{A} . The set of positive semi-definite operators with trace 1 acting on \mathcal{A} is denoted $D(\mathcal{A})$; $D(\mathcal{A})$ is the set of all possible quantum states for register \mathcal{A} . An operator $A \in L(\mathcal{A}, \mathcal{B})$ is called a *linear isometry* if $A^\dagger A = \mathbb{1}_{\mathcal{A}}$. The set of unitary operators (i.e., linear isometries with $\mathcal{B} = \mathcal{A}$) acting in \mathcal{A} is denoted by $U(\mathcal{A})$. The identity operator in \mathcal{A} is denoted $\mathbb{1}_{\mathcal{A}}$ and the completely mixed state in $D(\mathcal{A})$ is denoted by $\mathbb{I}_{\mathcal{A}}$. For any positive integer $N > 0$, $\mathbb{1}_N$ and \mathbb{I}_N denote the identity operator respectively the completely mixed state in \mathcal{H}_N . When the context requires, a pure state $|\psi\rangle \in \mathcal{AB}$ will be written $|\psi\rangle^{\mathcal{AB}}$ to make explicit the registers in which it is stored.

A linear mapping $\Phi : L(\mathcal{A}) \mapsto L(\mathcal{B})$ is called a *super-operator* since it belongs to $L(L(\mathcal{A}), L(\mathcal{B}))$. Φ is said to be *positive* if $\Phi(A) \in \text{Pos}(\mathcal{B})$ for all $A \in \text{Pos}(\mathcal{A})$. The super-operator Φ is said to be *completely positive* if $\Phi \otimes \mathbb{1}_{L(\mathcal{Z})}$ is positive for every choice of the Hilbert space \mathcal{Z} . A super-operator Φ can be physically realized or is *admissible* if it is completely positive and preserves the trace: $\text{tr}(\Phi(A)) = \text{tr}(A)$ for all $A \in L(\mathcal{A})$. We call such a super-operator a *quantum operation*. Another way to represent any quantum operation is through a linear isometry $W \in L(\mathcal{A}, \mathcal{B} \otimes \mathcal{Z})$ such that $\Phi(\rho) = \text{tr}_{\mathcal{Z}}(W \cdot \rho)$, for some extra space \mathcal{Z} . Any such isometry W can be implemented by a physical process as long as the resource to implement \mathcal{Z} is available. This is just a unitary transform in $U(\mathcal{A} \otimes \mathcal{Z})$ where the system in \mathcal{Z} is initially in known state $|0_{\mathcal{Z}}\rangle$.

For two states $\rho_0, \rho_1 \in D(\mathcal{A})$, we denote by $\Delta(\rho_0, \rho_1)$ the trace norm distance between ρ_0 and ρ_1 : $\Delta(\rho_0, \rho_1) := \frac{1}{2} \|\rho_0 - \rho_1\|$. If $\Delta(\rho_0, \rho_1) \leq \varepsilon$ then any quantum process applied to ρ_0 behaves exactly as for ρ_1 except with probability at most ε [16].

Let X, Y , and Z be the three non-trivial one-qubit Pauli operators. The Bell measurement is a complete orthogonal measurement on two qubits made out of the measurement operators $\{|\Psi_{x,y}\rangle\langle\Psi_{x,y}|\}_{x,y \in \{0,1\}}$ where $|\Psi_{x,y}\rangle := \frac{1}{\sqrt{2}}(|0, x\rangle + (-1)^y|1, \bar{x}\rangle)$. We say that the outcome of a Bell measurement is $(x, y) \in \{0, 1\}^2$ if $|\Psi_{x,y}\rangle\langle\Psi_{x,y}|$ has been observed. The quantum one-time-pad is a perfectly secure encryption of quantum states[2]. It encrypts a qubit $|\psi\rangle$ as $X^x Z^z |\psi\rangle$, where the key is two classical bits, $(x, z) \in \{0, 1\}^2$ and $X^0 Z^0 = \mathbb{1}$, $X^0 Z^1 = Z$, $X^1 Z^0 = X$ and $X^1 Z^1 = Y$ are the Pauli operators.

2.1 Modeling Two-Party Strategies

Consider an interactive two-party strategy Π^θ between parties \mathcal{A} and \mathcal{B} and oracle calls θ . Π^θ can be modeled by a sequence of quantum operations for

each player together with some oracle calls also modeled by quantum operations. Each quantum operation in the sequence corresponds to the action of one party at a certain step of the strategy. The following definition is a straightforward adaptation of n -turn interactive quantum strategies as described in [8]. The main difference is that here, we provide a joint input state to both parties and that quantum transmissions taking place during the execution is modeled by a quantum operation; one that is moving a state on one party’s side to the other party.

Definition 2.1. *A n -step two party strategy with oracle calls denoted $\Pi^\mathcal{O} = (\mathcal{A}, \mathcal{B}, \mathcal{O}, n)$ consists of:*

1. *input spaces \mathcal{A}_0 and \mathcal{B}_0 for parties \mathcal{A} and \mathcal{B} respectively,*
2. *memory spaces $\mathcal{A}_1, \dots, \mathcal{A}_n$ and $\mathcal{B}_1, \dots, \mathcal{B}_n$ for \mathcal{A} and \mathcal{B} respectively,*
3. *an n -tuple of quantum operations $(\mathcal{A}_1, \dots, \mathcal{A}_n)$ for \mathcal{A} , $\mathcal{A}_i : L(\mathcal{A}_{i-1}) \mapsto L(\mathcal{A}_i)$, $(1 \leq i \leq n)$,*
4. *an n -tuple of quantum operations $(\mathcal{B}_1, \dots, \mathcal{B}_n)$ for \mathcal{B} , $\mathcal{B}_i : L(\mathcal{B}_{i-1}) \mapsto L(\mathcal{B}_i)$, $(1 \leq i \leq n)$,*
5. *memory spaces $\mathcal{A}_1, \dots, \mathcal{A}_n$ and $\mathcal{B}_1, \dots, \mathcal{B}_n$ can be written as $\mathcal{A}_i = \mathcal{A}_i^\mathcal{O} \otimes \mathcal{A}'_i$ and $\mathcal{B}_i = \mathcal{B}_i^\mathcal{O} \otimes \mathcal{B}'_i$, $(1 \leq i \leq n)$, and $\mathcal{O} = (\mathcal{O}_1, \mathcal{O}_2, \dots, \mathcal{O}_n)$ is an n -tuple of quantum operations: $\mathcal{O}_i : L(\mathcal{A}_i^\mathcal{O} \otimes \mathcal{B}_i^\mathcal{O}) \mapsto L(\mathcal{A}_i^\mathcal{O} \otimes \mathcal{B}_i^\mathcal{O})$, $(1 \leq i \leq n)$.*

If $\Pi = (\mathcal{A}, \mathcal{B}, n)$ is a n -turn two-party protocol then the final state of the interaction upon input state $\rho_{\text{in}} \in D(\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R})$, where \mathcal{R} is a system of dimension $\dim \mathcal{R} = \dim \mathcal{A}_0 \dim \mathcal{B}_0$, is:

$$\begin{aligned}
 [\mathcal{A} \otimes \mathcal{B}](\rho_{\text{in}}) := & (\mathbb{1}_{L(\mathcal{A}'_n \otimes \mathcal{B}'_n \otimes \mathcal{R})} \otimes \mathcal{O}_n)(\mathcal{A}_n \otimes \mathcal{B}_n \otimes \mathbb{1}_{\mathcal{R}}) \\
 & \dots (\mathbb{1}_{L(\mathcal{A}'_1 \otimes \mathcal{B}'_1 \otimes \mathcal{R})} \otimes \mathcal{O}_1)(\mathcal{A}_1 \otimes \mathcal{B}_1 \otimes \mathbb{1}_{\mathcal{R}})(\rho_{\text{in}}) .
 \end{aligned}$$

Step i of the strategy corresponds to the actions of \mathcal{A}_i and \mathcal{B}_i followed by the oracle call \mathcal{O}_i .

Note that we consider input states defined on the input systems together with a reference system \mathcal{R} ; this allows us to show the correctness and privacy of the protocol not only for pure inputs, but also for inputs that are entangled with a third party. This is the most general case allowed by quantum mechanics.

A two-party strategy is therefore defined by quantum operation tuples $(\mathcal{A}_1, \dots, \mathcal{A}_n)$, $(\mathcal{B}_1, \dots, \mathcal{B}_n)$, and $(\mathcal{O}_1, \dots, \mathcal{O}_n)$. These operations also define working spaces $\mathcal{A}_0, \dots, \mathcal{A}_n, \mathcal{B}_0, \dots, \mathcal{B}_n$ together with the input-output spaces to the oracle calls $\mathcal{A}_i^\mathcal{O}$ and $\mathcal{B}_i^\mathcal{O}$ for $1 \leq i \leq n$.

A *communication oracle* from Alice to Bob is modeled by having $\mathcal{A}_i^\mathcal{O} \approx \mathcal{B}_i^\mathcal{O}$ and letting \mathcal{O}_i move the state in $\mathcal{A}_i^\mathcal{O}$ to $\mathcal{B}_i^\mathcal{O}$ and erase $\mathcal{A}_i^\mathcal{O}$. Similarly for communication in the other direction. We define a *bare model* protocol to be one which only uses communication oracles.

3 Specious Quantum Adversaries

3.1 Protocols for Two-Party Evaluation

Let us consider two-party protocols for the quantum evaluation of unitary transform $U \in U(\mathcal{A}_0 \otimes \mathcal{B}_0)$ between parties \mathcal{A} and \mathcal{B} upon joint input state $\rho_{\text{in}} \in D(\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R})$:

Definition 3.1. A two-party protocol $\Pi_U^\mathcal{O} = (\mathcal{A}, \mathcal{B}, \mathcal{O}, n)$ for $U \in U(\mathcal{A}_0 \otimes \mathcal{B}_0)$ is an n -step two-party strategy with oracle calls, where $\mathcal{A}_n \approx \mathcal{A}_0$ and $\mathcal{B}_n \approx \mathcal{B}_0$. It is said to be ε -correct if

$$\Delta([\mathcal{A} \otimes \mathcal{B}](\rho_{\text{in}}), (U \otimes \mathbb{1}_{\mathcal{R}}) \cdot \rho_{\text{in}}) \leq \varepsilon \quad \text{for all } \rho_{\text{in}} \in D(\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R}) .$$

We denote by Π_U a two-party protocol in the bare model where, without loss of generality, we assume that \mathcal{O}_{2i+1} ($0 \leq i \leq \lfloor \frac{n}{2} \rfloor$) implements a communication channel from \mathcal{A} to \mathcal{B} and \mathcal{O}_{2i} ($1 \leq i \leq \lfloor \frac{n}{2} \rfloor$) implements a communication channel from \mathcal{B} to \mathcal{A} . Communication oracles are said to be trivial.

In other words, a two-party protocol $\Pi_U^\mathcal{O}$ for unitary U is a two-party interactive strategy where, at the end, the output of the computation is stored in the memory of the players. $\Pi_U^\mathcal{O}$ is correct if, when restricted to the output registers (and \mathcal{R}), the final quantum state shared by \mathcal{A} and \mathcal{B} is $(U \otimes \mathbb{1}_{\mathcal{R}}) \cdot \rho_{\text{in}}$.

As it will become clear when we discuss privacy in Sect. 3.3, we need to consider the joint state at any step during the evolution of the protocol:

$$\begin{aligned} \rho_1(\rho_{\text{in}}) &:= (\mathbb{1}_{L(\mathcal{A}'_1 \otimes \mathcal{B}'_1 \otimes \mathcal{R})} \otimes \mathcal{O}_1)(\mathcal{A}_1 \otimes \mathcal{B}_1 \otimes \mathbb{1}_{L(\mathcal{R})})(\rho_{\text{in}}), \\ \rho_{i+1}(\rho_{\text{in}}) &:= (\mathbb{1}_{L(\mathcal{B}'_{i+1} \otimes \mathcal{A}'_{i+1} \otimes \mathcal{R})} \otimes \mathcal{O}_{i+1})(\mathcal{A}_{i+1} \otimes \mathcal{B}_{i+1} \otimes \mathbb{1}_{L(\mathcal{R})})(\rho_i(\rho_{\text{in}})) , \end{aligned} \quad (1)$$

for $1 \leq i < n$. We also write the final state of $\Pi_U^\mathcal{O}$ upon input state ρ_{in} as $\rho_n(\rho_{\text{in}}) = [\mathcal{A} \otimes \mathcal{B}](\rho_{\text{in}})$.

3.2 Modeling Specious Adversaries

Intuitively, a specious adversary acts in any way apparently indistinguishable from the honest behavior, in the sense that no audit can distinguish the behavior of the adversary from the honest one.

More formally, a specious adversary in $\Pi_U^\mathcal{O} = (\mathcal{A}, \mathcal{B}, \mathcal{O}, n)$ may use an arbitrary large quantum memory space. However, at any step $1 \leq i \leq n$, the adversary can transform its own current state to one that is indistinguishable from the honest joint state. These transforms are modeled by quantum operations, one for each step of the adversary in $\Pi_U^\mathcal{O}$, and are part of the adversary's specification. We denote by $(\mathcal{T}_1, \dots, \mathcal{T}_n)$ these quantum operations where \mathcal{T}_i produces a valid transcript at the end of the i -th step.

Let $\tilde{\mathcal{A}}$ and $\tilde{\mathcal{B}}$ be adversaries in $\Pi_U^\mathcal{O}$. We denote by $\Pi_U^\mathcal{O}(\tilde{\mathcal{A}}) = (\tilde{\mathcal{A}}, \mathcal{B}, \mathcal{O}, n)$ and $\Pi_U^\mathcal{O}(\tilde{\mathcal{B}}) = (\mathcal{A}, \tilde{\mathcal{B}}, \mathcal{O}, n)$ the resulting n -step two-party strategies. We denote by $\tilde{\rho}_i(\tilde{\mathcal{A}}, \rho_{\text{in}})$ the state defined in (1) for protocol $\Pi_U^\mathcal{O}(\tilde{\mathcal{A}})$ and similarly by $\tilde{\rho}_i(\tilde{\mathcal{B}}, \rho_{\text{in}})$ that state for protocol $\Pi_U^\mathcal{O}(\tilde{\mathcal{B}})$.

Adding the possibility for the adversary to be ε -close to honest, we get the following definition:

Definition 3.2. Let $\Pi_U^\mathcal{O} = (\mathcal{A}, \mathcal{B}, \mathcal{O}, n)$ be an n -step two-party protocol with oracle calls for $U \in \mathbf{U}(\mathcal{A}_0 \otimes \mathcal{B}_0)$. We say that:

- $\tilde{\mathcal{A}}$ is ε -specious if $\Pi_U^\mathcal{O}(\tilde{\mathcal{A}}) = (\tilde{\mathcal{A}}, \mathcal{B}, \mathcal{O}, n)$ is an n -step two-party strategy with $\tilde{\mathcal{A}}_0 = \mathcal{A}_0$ and there exists a sequence of quantum operations $(\mathcal{T}_1, \dots, \mathcal{T}_n)$ such that:
 1. for every $1 \leq i \leq n$, $\mathcal{T}_i : \mathbf{L}(\tilde{\mathcal{A}}_i) \mapsto \mathbf{L}(\mathcal{A}_i)$,
 2. for every input state $\rho_{\text{in}} \in \mathbf{D}(\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R})$, and for all $1 \leq i \leq n$,

$$\Delta \left((\mathcal{T}_i \otimes \mathbb{1}_{\mathbf{L}(\mathcal{B}_i \otimes \mathcal{R})}) \left(\tilde{\rho}_i(\tilde{\mathcal{A}}, \rho_{\text{in}}) \right), \rho_i(\rho_{\text{in}}) \right) \leq \varepsilon .$$

- $\tilde{\mathcal{B}}$ is ε -specious if $\Pi_U^\mathcal{O}(\tilde{\mathcal{B}}) = (\mathcal{A}, \tilde{\mathcal{B}}, \mathcal{O}, n)$ is a n -step two-party strategy with $\tilde{\mathcal{B}}_0 = \mathcal{B}_0$ and there exists a sequence of quantum operations $(\mathcal{T}_1, \dots, \mathcal{T}_n)$ defined as before with $\mathcal{B}_i, \tilde{\mathcal{B}}_i$, and $\tilde{\rho}_i(\tilde{\mathcal{B}}, \rho_{\text{in}})$ replacing $\mathcal{A}_i, \tilde{\mathcal{A}}_i$, and $\tilde{\rho}_i(\tilde{\mathcal{A}}, \rho_{\text{in}})$ respectively.

If a party is $\varepsilon(m)$ -specious with $\varepsilon(m)$ negligible for m a security parameter then we say that this party is statistically specious.

3.3 Privacy

Privacy for $\Pi_U^\mathcal{O}$ is defined as the ability for a simulator, having only access to the adversary’s input and the ideal functionality U , to reproduce the state of the adversary at any step in the execution of $\Pi_U^\mathcal{O}$. Our definition is similar to the one introduced in [20] for statistical zero-knowledge proof systems.

A simulator for an adversary in $\Pi_U^\mathcal{O}$ is represented by a sequence of quantum operations $(\mathcal{S}_i)_{i=1}^n$, where \mathcal{S}_i re-produces the view of the adversary after step i . \mathcal{S}_i initially receives the adversary’s input and has access to the ideal functionality for U evaluated upon the joint input of the adversary and the honest player. Because of no-cloning, a simulator calling U loses its input, and the input might be required to simulate e.g. early steps in the protocol, so we have to allow that \mathcal{S}_i does not call U . For this purpose we introduce a bit $q_i \in \{0, 1\}$. When $q_i = 0$, \mathcal{S}_i does not call U and when $q_i = 1$, \mathcal{S}_i must first call the ideal functionality U before performing some post-processing. More precisely,

Definition 3.3. Let $\Pi_U^\mathcal{O} = (\mathcal{A}, \mathcal{B}, \mathcal{O}, n)$ be an n -step two-party protocol for $U \in \mathbf{D}(\mathcal{A}_0 \otimes \mathcal{B}_0)$. Then,

- $\mathcal{S}(\tilde{\mathcal{A}}) = ((\mathcal{S}_1, \dots, \mathcal{S}_n), q)$ is a simulator for adversary $\tilde{\mathcal{A}}$ in $\Pi_U^\mathcal{O}$ if it consists of:
 1. a sequence of quantum operations $(\mathcal{S}_1, \dots, \mathcal{S}_n)$ where for $1 \leq i \leq n$, $\mathcal{S}_i : \mathbf{L}(\mathcal{A}_0) \mapsto \mathbf{L}(\tilde{\mathcal{A}}_i)$,
 2. a sequence of bits $q \in \{0, 1\}^n$ determining if the simulator calls the ideal functionality at step i : $q_i = 1$ iff the simulator calls the ideal functionality.

- Similarly, $\mathcal{S}(\tilde{\mathcal{B}}) = \langle (\mathcal{S}_1, \dots, \mathcal{S}_n), q' \rangle$ is a simulator for adversary $\tilde{\mathcal{B}}$ in $\Pi_U^\mathcal{O}$ if it satisfies conditions 1 and 2 above with $q', \mathcal{B}_0, \mathcal{B}_i$, and $\tilde{\mathcal{B}}_i$ replacing $q, \mathcal{A}_0, \mathcal{A}_i$, and $\tilde{\mathcal{A}}_i$ respectively.

Given an input state $\rho_{\text{in}} \in D(\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R})$, we define the $\tilde{\mathcal{A}}_i$'s respectively $\tilde{\mathcal{B}}_i$'s simulated views as:

$$\begin{aligned} \nu_i(\tilde{\mathcal{A}}_i, \rho_{\text{in}}) &:= \text{tr}_{\mathcal{B}_0} \left((\mathcal{S}_i \otimes \mathbb{1}_{L(\mathcal{B}_0 \otimes \mathcal{R})}) ((U^{q_i} \otimes \mathbb{1}_{\mathcal{R}}) \cdot \rho_{\text{in}}) \right) , \\ \nu_i(\tilde{\mathcal{B}}_i, \rho_{\text{in}}) &:= \text{tr}_{\mathcal{A}_0} \left((\mathbb{1}_{L(\mathcal{A}_0 \otimes \mathcal{R})} \otimes \mathcal{S}_i) \left((U^{q'_i} \otimes \mathbb{1}_{\mathcal{R}}) \cdot \rho_{\text{in}} \right) \right) . \end{aligned}$$

We say that protocol $\Pi_U^\mathcal{O}$ is private against specious adversaries if there exists a simulator for the view at any step of any such adversary. In more details,

Definition 3.4. Let $\Pi_U^\mathcal{O} = (\mathcal{A}, \mathcal{B}, \mathcal{O}, n)$ be a protocol for $U \in U(\mathcal{A}_0 \otimes \mathcal{B}_0)$ and let $0 \leq \delta \leq 1$. We say that $\Pi_U^\mathcal{O}$ is δ -private against ε -specious $\tilde{\mathcal{A}}$ if there exists a simulator $\mathcal{S}(\tilde{\mathcal{A}})$ such that for all input states $\rho_{\text{in}} \in D(\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R})$ and for all $1 \leq i \leq n$, $\Delta \left(\nu_i(\tilde{\mathcal{A}}_i, \rho_{\text{in}}), \text{tr}_{\mathcal{B}_i}(\tilde{\rho}_i(\tilde{\mathcal{A}}_i, \rho_{\text{in}})) \right) \leq \delta$. Similarly, we say that Π_U is δ -private against ε -specious $\tilde{\mathcal{B}}$ if there exists a simulator $\mathcal{S}(\tilde{\mathcal{B}})$ such that for all input states $\rho_{\text{in}} \in D(\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R})$ and for all $1 \leq i \leq n$, $\Delta \left(\nu_i(\tilde{\mathcal{B}}_i, \rho_{\text{in}}), \text{tr}_{\mathcal{A}_i}(\tilde{\rho}_i(\tilde{\mathcal{B}}_i, \rho_{\text{in}})) \right) \leq \delta$. Protocol $\Pi_U^\mathcal{O}$ is δ -private against ε -specious adversaries if it is δ -private against both $\tilde{\mathcal{A}}$ and $\tilde{\mathcal{B}}$. For $\gamma > 0$, if $\Pi_U^\mathcal{O}$ is $2^{-\gamma m}$ -private for $m \in \mathbb{N}^+$ a security parameter then we say that $\Pi_U^\mathcal{O}$ is statistically private.

We show next that for some unitary, statistical privacy cannot be satisfied by any protocol in the bare model.

4 Unitaries with No Private Protocols

In this section, we show that no statistically private protocol for the swap gate exists in the bare model. The swap gate, denoted SWAP, is the following unitary transform:

$$\text{SWAP} : |\phi_A\rangle^{\mathcal{A}_0} |\phi_B\rangle^{\mathcal{B}_0} \mapsto |\phi_B\rangle^{\mathcal{A}_0} |\phi_A\rangle^{\mathcal{B}_0} ,$$

for any one qubit states $|\phi_A\rangle \in \mathcal{A}_0$ and $|\phi_B\rangle \in \mathcal{B}_0$ (i.e., $\dim(\mathcal{A}_0) = \dim(\mathcal{B}_0) = 2$). Notice that SWAP is in the Clifford group since it can be implemented with three CNOT gates. It means that universality is not required (gates in the Clifford groups are not universal for quantum computation) for a unitary to be impossible to evaluate privately. The impossibility of SWAP essentially follows from no cloning.

Theorem 4.1 (Impossibility of swapping). *There is no correct and statistically private two-party protocol $\Pi_{\text{SWAP}} = (\mathcal{A}, \mathcal{B}, \mathcal{O}, n(m))$ in the bare model.*

Proof. Suppose that there exists an ε -correct, ε -private protocol in the bare model for SWAP for sufficiently small ε ; we will show that this implies that one of the two players must *lose* information upon receiving a message, which is clearly impossible.

We will consider the following particular pure input state: $|\varphi\rangle := |\Psi_{0,0}\rangle^{\mathcal{A}_0\mathcal{R}_A} \otimes |\Psi_{0,0}\rangle^{\mathcal{B}_0\mathcal{R}_B}$, a maximally entangled state between $\mathcal{A}_0 \otimes \mathcal{B}_0$ and the reference system $\mathcal{R}_A \otimes \mathcal{R}_B$ that is broken down into two subsystems for convenience. Furthermore, we will consider the “purified” versions of the honest players for this protocol; in other words, we will assume that the super-operators $\mathcal{A}_1, \dots, \mathcal{A}_n$ and $\mathcal{B}_1, \dots, \mathcal{B}_n$ are in fact linear isometries and that therefore the players never discard any information unless they have to send it to the other party. The global state $\rho_i(\varphi)$ after step i is therefore a pure state on $\mathcal{A}_i \otimes \mathcal{B}_i \otimes \mathcal{R}_A \otimes \mathcal{R}_B$.

After step i of the protocol (i.e., after the i th message has been sent), Alice’s state must either depend only on her own original input (if $q_i = 0$ for her simulator), or on Bob’s original input (if $q_i = 1$). More precisely, by the definition of privacy (Definition 3.4), we have that

$$\Delta(\nu_i(\mathcal{A}, \varphi), \text{tr}_{\mathcal{B}_i}[\rho_i(\varphi)]) \leq \varepsilon ,$$

where $\nu_i(\mathcal{A}, \varphi)$ is \mathcal{A} ’s simulated view after step i and $\rho_i(\varphi)$ is the global state in the real protocol after step i . Now, suppose that $q_i = 0$, and let $|\xi\rangle \in \mathcal{A}_i \otimes \mathcal{R}_A \otimes \mathcal{R}'_B \otimes \mathcal{Z}$ be a purification of $\nu_i(\mathcal{A}, \varphi)$ with \mathcal{Z} being the purifying system, and \mathcal{R}_B renamed for upcoming technical reasons. The pure state $|\xi\rangle \otimes |\Psi_{0,0}\rangle^{\mathcal{R}_B\mathcal{B}_0}$ has the same reduced density matrix as $\nu_i(\mathcal{A}, \varphi)$ on $\mathcal{A}_i \otimes \mathcal{R}_A \otimes \mathcal{R}_B$. Hence, by Uhlmann’s theorem, there exists a linear isometry $V : \mathcal{B}_i \rightarrow \mathcal{B}_0 \otimes \mathcal{Z} \otimes \mathcal{R}'_B$ such that

$$V\nu_i(\mathcal{A}, \varphi)V^\dagger = |\xi\rangle\langle\xi| \otimes |\Psi_{0,0}\rangle\langle\Psi_{0,0}|^{\mathcal{B}_0\mathcal{R}_B}$$

and hence

$$\Delta\left(V\rho_i(\varphi)V^\dagger, |\xi\rangle\langle\xi| \otimes |\Psi_{0,0}\rangle\langle\Psi_{0,0}|^{\mathcal{B}_0\mathcal{R}_B}\right) \leq \sqrt{2\varepsilon} .$$

This means that if $q_i = 0$, then Bob is still capable of reconstructing his own input state after step i by applying V to his working register. Clearly, this means that $q'_i = 0$ (i.e., Bob’s simulator must also not call SWAP), and therefore, by the same argument, Alice must also be able to reconstruct her own input with an isometry $V_A : \mathcal{A}_i \rightarrow \mathcal{B}_0 \otimes \mathcal{Z} \otimes \mathcal{R}'_A$. The same argument also holds if $q_i = 1$: we then conclude that $q'_i = 1$ and that Alice and Bob must have each other’s inputs; no intermediate situation is possible. We conclude that, at every step i of the protocol, $q_i = q'_i$.

Now, before the protocol starts, Alice must have her input, and Bob must have his, hence, $q_0 = q'_0 = 0$. At the end, the two inputs must have been swapped, which means that $q_n = q'_n = 1$; there must therefore be a step k in the protocol after which the two inputs are swapped but not before, meaning that $q_k = 1$ and $q_{k-1} = 0$. But at each step, only one player receives information, which means that at this step k , the player who received the message must lose the ability to reconstruct his own input, which is clearly impossible. \square

Using this line of reasoning, Theorem 4.1 can be extended to apply to any protocol for almost any unitary preventing both parties to recover their input states from its output.

Sufficient Assumptions for Private SWAP. A private protocol for SWAP in the bare model would exist if the players could rely on special relativity and a lower bound on their separation in space: they simply send their messages simultaneously. The fact that messages cannot travel faster than the speed of light ensures that the messages are independent of each other. It is also straightforward to devise a private protocol for SWAP based on commitment schemes. \mathcal{A} sends one half EPR-pair to \mathcal{B} while keeping the other half. \mathcal{A} then teleports (without announcing the outcome of the measurement) her register and commits on the outcome of the Bell measurement. \mathcal{B} sends his register to \mathcal{A} before she opens her commitment.

5 The Protocol

We now describe a private protocol for the two-party evaluation of any unitary $U \in \mathcal{U}(\mathcal{A}_0 \otimes \mathcal{B}_0)$ denoted by $P_U^\mathcal{O} = (\mathcal{A}^*, \mathcal{B}^*, \mathcal{O}, n_U + 1)$ where U is represented by a circuit C_U with u gates in \mathcal{UG} . We slightly abuse the notation with respect to the parameter $n_U + 1$. Given circuit C_U , we let n_U be the number of oracle calls (including calls to communication oracles). Setting the last parameter to $n_U + 1$ instead of n_U comes from the fact that in our protocol, \mathcal{A}^* and \mathcal{B}^* have to perform a last operation each in order to get their outcome. These last operations do not involve a call to any oracle. Let G_j be the j -th gate in $C_U = G_u G_{u-1} \dots G_1$. The protocol is obtained by composing sub-protocols for each gate similarly to well-known classical constructions[22,9]. Notice that $P_U^\mathcal{O}$ will not be presented in the form of Definition 3.1. \mathcal{A}^* is not necessarily sending the first and the last messages. This can be done without consequences since we provide a simulation for each step where a message from the honest party is received or the output of a call to an ideal functionality is available. Putting $P_U^\mathcal{O}$ in the standard form of Definition 3.1 is straightforward and changes nothing to the proof of privacy.

The evaluation of each gate is performed over shared encrypted states. Each wire in C_U will be updated from initially holding the input $\rho_{\text{in}} \in \mathcal{D}(\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R})$ to finally holding the output $(U \otimes \mathbb{I}_{\mathcal{R}}) \cdot \rho_{\text{in}} \in \mathcal{D}(\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R})$. The state of wires $\mathbf{w} \in \mathcal{A}_0 \cup \mathcal{B}_0$ after the evaluation of G_j are stored at \mathcal{A}^* 's or \mathcal{B}^* 's according if $\mathbf{w} \in \mathcal{A}_0$ or $\mathbf{w} \in \mathcal{B}_0$. The shared encryption keys for wire $\mathbf{w} \in \mathcal{A}_0 \cup \mathcal{B}_0$ updated after the evaluation of G_j are denoted by $K_{\mathcal{A}^*}^j(\mathbf{w}) = (X_{\mathcal{A}^*}^j(\mathbf{w}), Z_{\mathcal{A}^*}^j(\mathbf{w})) \in \{0, 1\}^2$ and $K_{\mathcal{B}^*}^j(\mathbf{w}) = (X_{\mathcal{B}^*}^j(\mathbf{w}), Z_{\mathcal{B}^*}^j(\mathbf{w})) \in \{0, 1\}^2$ for \mathcal{A}^* and \mathcal{B}^* respectively and are held privately in internal registers of each party.

The final phase of the protocol is where a call to an ideal functionality is required. \mathcal{A}^* and \mathcal{B}^* exchange their own part of each encryption key for the other party's wires. In order to do this, the *key-releasing phase* invokes an ideal SWAP-gate as functionality: $\mathcal{O}_{n_U} : \mathcal{L}(\mathcal{A}_{n_U}^\mathcal{O} \otimes \mathcal{B}_{n_U}^\mathcal{O}) \mapsto \mathcal{L}(\mathcal{A}_{n_U}^\mathcal{O} \otimes \mathcal{B}_{n_U}^\mathcal{O})$, where

$\mathcal{O}_{n_U}(\rho) := \text{SWAP} \cdot \rho$. Upon joint input state $\rho_{\text{in}} \in \mathcal{D}(\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R})$, protocol $P_U^{\mathcal{O}(U)}$ runs the following phases:

Initialization: We assume that \mathcal{A}^* and \mathcal{B}^* have agreed upon a description of U by a circuit C_U made out of u gates (G_1, \dots, G_u) in \mathcal{UG} . For all wires $\mathbf{w} \in \mathcal{A}_0 \cup \mathcal{B}_0$, \mathcal{A}^* and \mathcal{B}^* set their initial encryption keys as $K_{\mathcal{A}^*}^0(\mathbf{w}) = (X_{\mathcal{A}^*}^0(\mathbf{w}), Z_{\mathcal{A}^*}^0(\mathbf{w})) := (0, 0)$ and $K_{\mathcal{B}^*}^0(\mathbf{w}) = (X_{\mathcal{B}^*}^0(\mathbf{w}), Z_{\mathcal{B}^*}^0(\mathbf{w})) := (0, 0)$ respectively.

Evaluation: For each gate number $1 \leq j \leq u$, \mathcal{A}^* and \mathcal{B}^* evaluate G_j as described in details below. This evaluation results in shared encryption under keys $K_{\mathcal{A}^*}^j(\mathbf{w}) = (X_{\mathcal{A}^*}^j(\mathbf{w}), Z_{\mathcal{A}^*}^j(\mathbf{w}))$ and $K_{\mathcal{B}^*}^j(\mathbf{w}) = (X_{\mathcal{B}^*}^j(\mathbf{w}), Z_{\mathcal{B}^*}^j(\mathbf{w}))$ for all wires $\mathbf{w} \in \mathcal{A}_0 \cup \mathcal{B}_0$, which at that point hold a shared encryption of $((G_j G_{j-1} \dots G_1) \otimes \mathbb{1}_{\mathcal{R}}) \cdot \rho_{\text{in}}$. Only the evaluation of the R-gate requires a call to an ideal functionality (i.e., an AND-BOX).

Key-Releasing: Let $\mathcal{A}_{n_U}^{\mathcal{O}}$ and $\mathcal{B}_{n_U}^{\mathcal{O}}$ be the set of registers holding respectively $K_{\mathcal{A}^*}^u(\mathbf{w}) = (X_{\mathcal{A}^*}^u(\mathbf{w}), Z_{\mathcal{A}^*}^u(\mathbf{w}))$ for $\mathbf{w} \in \mathcal{B}_0$ and $K_{\mathcal{B}^*}^u(\mathbf{w}) = (X_{\mathcal{B}^*}^u(\mathbf{w}), Z_{\mathcal{B}^*}^u(\mathbf{w}))$ for $\mathbf{w} \in \mathcal{A}_0$. We assume w.l.g that dimensions of both sets of registers are identical²:

1. \mathcal{A}^* and \mathcal{B}^* run the ideal functionality for the SWAP-gate upon registers $\mathcal{A}_{n_U}^{\mathcal{O}}$ and $\mathcal{B}_{n_U}^{\mathcal{O}}$.
2. \mathcal{A}^* applies the decryption operator $K_{\mathcal{A}^*}^u(\mathbf{w}) = (X_{\mathcal{A}^*}^u(\mathbf{w}) \oplus X_{\mathcal{B}^*}^u(\mathbf{w}), Z_{\mathcal{A}^*}^u(\mathbf{w}) \oplus Z_{\mathcal{B}^*}^u(\mathbf{w}))$ to each of her wires $\mathbf{w} \in \mathcal{A}_0$.
3. \mathcal{B}^* applies the decryption operator for key $K_{\mathcal{B}^*}^u(\mathbf{w}) = (X_{\mathcal{A}^*}^u(\mathbf{w}) \oplus X_{\mathcal{B}^*}^u(\mathbf{w}), Z_{\mathcal{A}^*}^u(\mathbf{w}) \oplus Z_{\mathcal{B}^*}^u(\mathbf{w}))$ to each of his wires $\mathbf{w} \in \mathcal{B}_0$.

Swapping for Key-Releasing. Notice that the key-releasing phase only uses the SWAP-gate with classical input states. The reader might therefore wonder why this functionality is defined quantumly when a classical swap would work equally well. The reason is that, perhaps somewhat surprisingly, a classical swap is a potentially stronger primitive than a quantum swap. From a classical swap one can build a quantum swap by encrypting the quantum states with classical keys, exchange the encrypted states using quantum communication, and then using the classical swap to exchange the keys. Obtaining a classical swap from a quantum one, however, is not obvious. Suppose that registers \mathcal{A} and \mathcal{B} should be swapped classically while holding quantum states beforehand. These registers could be entangled with some purification registers before being swapped. Using a quantum swap between \mathcal{A} and \mathcal{B} will always leave these registers entangled with the purification registers until they become measured while a classical swap will ensure that \mathcal{A} and \mathcal{B} become unentangled with the purification registers after its invocation. In other words, a classical swap could prevent an adversary from exploiting entanglement in his attack.

² Otherwise, add enough registers initially in state $|0\rangle$ to the smaller set.

The Ideal AND-Box Functionality. As we are going to see next, a call to an ideal AND-box is required during the evaluation of the R-gate. Unlike the ideal SWAP used for key-releasing, the AND-box will be modeled by a purely classical primitive denoted AND-BOX. This is required for privacy of our protocol since any implementation of it by some unitary will necessarily leak[17]. The quantum operation implementing it will first measure the two one-qubit input registers in the computational basis in order to get classical inputs $x, y \in \{0, 1\}$ for \mathcal{A}^* and \mathcal{B}^* respectively. The classical output bits are then set to $a \in_R \{0, 1\}$ for \mathcal{A}^* and $b = a \oplus xy$ for \mathcal{B}^* .

5.1 Computing over Encrypted States

Before the execution of G_{j+1} in C_U , \mathcal{A}^* and \mathcal{B}^* share an encryption of $\rho_j = ((G_j \cdot G_{j-1} \cdot \dots \cdot G_1) \otimes \mathbb{1}_{\mathcal{R}}) \cdot \rho_{\text{in}}$ in registers³ holding wires $\mathbf{w} \in \mathcal{A}_0 \cup \mathcal{B}_0$. Each wire $\mathbf{w} \in \mathcal{A}_0 \cup \mathcal{B}_0$ is encrypted by a shared quantum one-time pad as

$$\left(\left(\bigotimes_{\mathbf{w} \in \mathcal{A}_0 \cup \mathcal{B}_0} X^{X_{\mathcal{A}^*}^j(\mathbf{w}) \oplus X_{\mathcal{B}^*}^j(\mathbf{w})} Z^{Z_{\mathcal{A}^*}^j(\mathbf{w}) \oplus Z_{\mathcal{B}^*}^j(\mathbf{w})} \right) \otimes \mathbb{1}_{\mathcal{R}} \right) \cdot \rho_j, \tag{2}$$

where $K_{\mathcal{A}^*}^j(\mathbf{w}) := (X_{\mathcal{A}^*}^j(\mathbf{w}), Z_{\mathcal{A}^*}^j(\mathbf{w})) \in \{0, 1\}^2$ and $K_{\mathcal{B}^*}^j(\mathbf{w}) := (X_{\mathcal{B}^*}^j(\mathbf{w}), Z_{\mathcal{B}^*}^j(\mathbf{w})) \in \{0, 1\}^2$ are two bits of secret keys for \mathcal{A}^* and \mathcal{B}^* respectively. In other words, wires $\mathbf{w} \in \mathcal{A}_0 \cup \mathcal{B}_0$ are encrypted by $X^x Z^z$ where $x = X_{\mathcal{A}^*}^j(\mathbf{w}) \oplus X_{\mathcal{B}^*}^j(\mathbf{w})$ and $z = Z_{\mathcal{A}^*}^j(\mathbf{w}) \oplus Z_{\mathcal{B}^*}^j(\mathbf{w})$ are additive sharings for the encryption of \mathbf{w} . Then, evaluating G_{j+1} upon state (2) will produce a new sharing $K_A^{j+1}(\mathbf{w}) := (X_A^{j+1}(\mathbf{w}), Z_A^{j+1}(\mathbf{w}))$ and $K_B^{j+1}(\mathbf{w}) := (X_B^{j+1}(\mathbf{w}), Z_B^{j+1}(\mathbf{w}))$ for the encryption of state $\rho_{j+1} = (G_{j+1} \otimes \mathbb{1}_{\mathcal{R}}) \cdot \rho_j$. In the following, we describe how to update the keys for the wires involved in the current gate to be evaluated—all other wires retain their previous values.

5.2 Evaluation of Gates in the Pauli and Clifford Groups

Non-trivial Pauli gates (i.e., X, Y , and Z) can easily be computed on encrypted quantum states since they commute or anti-commute pairwise. Let $G_{j+1} \in \{X, Y, Z\}$ be the Pauli gate to be executed on wire \mathbf{w} . It means that up to an irrelevant phase factor, it suffices for the owner of \mathbf{w} to apply G_{j+1} without the need for neither party to update their shared keys, i.e., $K_{\mathcal{A}^*}^{j+1}(\mathbf{w}) := K_{\mathcal{A}^*}^j(\mathbf{w})$ and $K_{\mathcal{B}^*}^{j+1}(\mathbf{w}) := K_{\mathcal{B}^*}^j(\mathbf{w})$.

Now, suppose that $G_{j+1} \in \{\mathbf{H}, \mathbf{P}\}$. Each of these one-qubit gates applied upon wire \mathbf{w} will be computed by simply letting the party owning \mathbf{w} apply G_{j+1} . Encryption keys are updated locally as:

$$\begin{aligned} \mathbf{H} : K_{\mathcal{A}^*}^{j+1} &= (X_{\mathcal{A}^*}^{j+1}(\mathbf{w}), Z_{\mathcal{A}^*}^{j+1}(\mathbf{w})) := (Z_{\mathcal{A}^*}^j(\mathbf{w}), X_{\mathcal{A}^*}^j(\mathbf{w})), \\ K_{\mathcal{B}^*}^{j+1} &= (X_{\mathcal{B}^*}^{j+1}(\mathbf{w}), Z_{\mathcal{B}^*}^{j+1}(\mathbf{w})) := (Z_{\mathcal{B}^*}^j(\mathbf{w}), X_{\mathcal{B}^*}^j(\mathbf{w})), \end{aligned}$$

³ To ease the notation in the following, we assume $\rho_j \in \mathcal{D}(\mathcal{A}_0 \otimes \mathcal{B}_0)$ rather than in $\mathcal{D}(\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R})$. It is easy to see that this can be done without loss of generality.

$$\begin{aligned}
 P : K_{\mathcal{A}^*}^{j+1} &= (X_{\mathcal{A}^*}^{j+1}(\mathbf{w}), Z_{\mathcal{A}^*}^{j+1}(\mathbf{w})) := (X_{\mathcal{A}^*}^j(\mathbf{w}), X_{\mathcal{A}^*}^j(\mathbf{w}) \oplus Z_{\mathcal{A}^*}^j(\mathbf{w})) , \\
 K_{\mathcal{B}^*}^{j+1} &= (X_{\mathcal{B}^*}^{j+1}(\mathbf{w}), Z_{\mathcal{B}^*}^{j+1}(\mathbf{w})) := (X_{\mathcal{B}^*}^j(\mathbf{w}), X_{\mathcal{B}^*}^j(\mathbf{w}) \oplus Z_{\mathcal{B}^*}^j(\mathbf{w})) .
 \end{aligned}$$

Any one-qubit gate in the Clifford group can be implemented the same way using their own commutation relations with the Pauli operators used for encryption. A CNOT-gate on local wires can be evaluated in a similar way. That is, whenever both wires \mathbf{w} and \mathbf{w}' feeding the CNOT belong to the same party. Assume that \mathbf{w} is the control wire while \mathbf{w}' is the target and that \mathcal{A}^* holds them both. Then, \mathcal{A}^* simply applies CNOT on wires \mathbf{w} and \mathbf{w}' . Encryption keys are updated as:

$$\begin{aligned}
 \text{CNOT} : K_{\mathcal{A}^*}^{j+1}(\mathbf{w}) &= (X_{\mathcal{A}^*}^{j+1}(\mathbf{w}), Z_{\mathcal{A}^*}^{j+1}(\mathbf{w})) := (X_{\mathcal{A}^*}^j(\mathbf{w}), Z_{\mathcal{A}^*}^j(\mathbf{w}) \oplus Z_{\mathcal{A}^*}^j(\mathbf{w}')) , \\
 K_{\mathcal{A}^*}^{j+1}(\mathbf{w}') &= (X_{\mathcal{A}^*}^{j+1}(\mathbf{w}'), Z_{\mathcal{A}^*}^{j+1}(\mathbf{w}')) := (X_{\mathcal{A}^*}^j(\mathbf{w}') \oplus X_{\mathcal{A}^*}^j(\mathbf{w}), Z_{\mathcal{A}^*}^j(\mathbf{w}')) , \\
 K_{\mathcal{B}^*}^{j+1}(\mathbf{w}) &:= K_{\mathcal{B}^*}^j(\mathbf{w}) \text{ and } K_{\mathcal{B}^*}^{j+1}(\mathbf{w}') := K_{\mathcal{B}^*}^j(\mathbf{w}') .
 \end{aligned}$$

When \mathcal{B}^* holds both wires, the procedure is simply performed with the roles of \mathcal{A}^* and \mathcal{B}^* reversed.

Nonlocal CNOT. We now look at the case where $G_{j+1} = \text{CNOT}$ upon wires \mathbf{w} and \mathbf{w}' , one of which is owned by \mathcal{A}^* while the other is owned by \mathcal{B}^* . In this case, interaction is unavoidable for the evaluation of the gate. Let us assume w.l.g that \mathcal{A}^* holds the control wire \mathbf{w} while \mathcal{B}^* holds the target wire \mathbf{w}' (i.e., $\mathbf{w} \in \mathcal{A}_0$ and $\mathbf{w}' \in \mathcal{B}_0$). We start from a construction introduced in [7] in the context of fault tolerant quantum computation.

The idea behind the sub-protocol is depicted in Fig. 2. The effect of the Bell measurement is to *teleport* the input state of wires \mathbf{w} and \mathbf{w}' through the CNOT-gate[7]. The input to the CNOT appearing in the circuit of Fig. 2 is independent of both input wires \mathbf{w} and \mathbf{w}' (they are just two half EPR-pairs).

The sub-protocol for the evaluation of CNOT simply consists in executing the circuit of Fig. 2 without the decryption part (i.e., the part inside the dotted rectangle). The state $|\xi\rangle := (\mathbb{1}_A \otimes \text{CNOT} \otimes \mathbb{1}_B)|\Psi_{0,0}\rangle|\Psi_{0,0}\rangle$ can be prepared by one party. We let the holder of the *control wire* (i.e., \mathcal{A}^* in Fig. 2) prepare $|\xi\rangle$ before sending its two rightmost registers to the other party. The decryption in the dotted-rectangle is used to update the encryption keys according to the measurement outcomes (a_x, a_z, b_x, b_z) :




Fig. 2. Evaluation of CNOT

$$\begin{aligned}
 \text{CNOT} : K_{\mathcal{A}^*}^{j+1}(\mathbf{w}) &:= (X_{\mathcal{A}^*}^j(\mathbf{w}) \oplus a_x, Z_{\mathcal{A}^*}^j(\mathbf{w}) \oplus a_z) , \\
 K_{\mathcal{B}^*}^{j+1}(\mathbf{w}) &:= (X_{\mathcal{B}^*}^j(\mathbf{w}), Z_{\mathcal{B}^*}^j(\mathbf{w}) \oplus b_z) , \\
 K_{\mathcal{A}^*}^{j+1}(\mathbf{w}') &:= (X_{\mathcal{A}^*}^j(\mathbf{w}') \oplus a_x, Z_{\mathcal{A}^*}^j(\mathbf{w}')) , \\
 K_{\mathcal{B}^*}^{j+1}(\mathbf{w}') &:= (X_{\mathcal{B}^*}^j(\mathbf{w}') \oplus b_x, Z_{\mathcal{B}^*}^j(\mathbf{w}') \oplus b_z) .
 \end{aligned}$$

As for all previous gates, the key updating phase is performed locally without the need for communication.

5.3 Evaluation of the R-Gate

The only gate left in \mathcal{UG} is $G_{j+1} := R$. We assume without loss of generality that \mathcal{A}^* owns wire w upon which R is applied (i.e., $w \in \mathcal{A}_0$). The subprotocol needs a call to an ideal AND-BOX in order to guarantee privacy during the key updating process. Observe first that the R -gate commutes with Pauli encryption operator Z . It means that applying the R -gate upon a state encrypted with Z produces the correct output state still encrypted with Z . However, the equality $R \cdot X = e^{-i\pi/4} Y P \cdot R$ tells us that a P -gate should be applied for the decryption of the output when the input has been encrypted using X . This breaks the invariant that wires after each gate are all encrypted by Pauli operators. We remove the P -gate by converting it into a sequence of Pauli operators.

Ignoring an irrelevant global phase, the result of applying R on wire w is

$$RZ^{X_{\mathcal{A}^*}^j(w) \oplus Z_{\mathcal{B}^*}^i(w)} X^{X_{\mathcal{A}^*}^j(w) \oplus X_{\mathcal{B}^*}^j(w)} = Z^{X_{\mathcal{A}^*}^j(w) \oplus Z_{\mathcal{B}^*}^j(w) \oplus X_{\mathcal{A}^*}^j(w) \oplus X_{\mathcal{B}^*}^j(w)} X^{X_{\mathcal{A}^*}^j(w) \oplus X_{\mathcal{B}^*}^j(w)} P^{X_{\mathcal{A}^*}^j(w) \oplus X_{\mathcal{B}^*}^j(w)} R, \tag{3}$$

To remove the P -gate, we let each party remove his part of $P^{X_{\mathcal{A}^*}^j(w) \oplus X_{\mathcal{B}^*}^j(w)}$ in a private interactive process. To do this, \mathcal{A}^* picks random bits r and r' , and \mathcal{B}^* picks random bits s and s' . \mathcal{A}^* applies the operator $X^r Z^{r'} P^{X_{\mathcal{A}^*}^j(w)}$ and sends the resulting quantum state to \mathcal{B}^* . \mathcal{B}^*




Fig. 3. Implementation of the R-gate

applies the operator $X^s Z^{s'} P^{X_{\mathcal{B}^*}^j(w)}$ and sends the result back to \mathcal{A}^* . The resulting protocol is shown in Fig. 3. It starts with \mathcal{A}^* applying R upon the encrypted state before the one-round interactive process described above starts.

After \mathcal{A}^* 's application of R , the resulting state is as described on the right-hand side of (3). At the end of the process (i.e., circuit of Fig. 3), the encryption becomes:

$$Z^{s'} X^s P^{X_{\mathcal{B}^*}^j(w)} Z^{r'} X^r P^{X_{\mathcal{A}^*}^j(w)} Z^{X_{\mathcal{A}^*}^j(w) \oplus Z_{\mathcal{B}^*}^j(w) \oplus X_{\mathcal{A}^*}^j(w) \oplus X_{\mathcal{B}^*}^j(w)} X^{X_{\mathcal{A}^*}^j(w) \oplus X_{\mathcal{B}^*}^j(w)} P^{X_{\mathcal{A}^*}^j(w) \oplus X_{\mathcal{B}^*}^j(w)}. \tag{4}$$

Now, we use the fact that Z commutes with P and $P \cdot X = XZ \cdot P$. In addition, since for $a, b \in \{0, 1\}$, $P^{a+b} = Z^{ab} P^{a \oplus b}$ we re-write (4) as

$$Z^{s' \oplus r' \oplus X_{\mathcal{A}^*}^j(w) \oplus X_{\mathcal{B}^*}^j(w) \oplus Z_{\mathcal{A}^*}^j(w) \oplus Z_{\mathcal{B}^*}^j(w) \oplus (r \oplus X_{\mathcal{A}^*}^j(w)) \cdot X_{\mathcal{B}^*}^j(w)} X^{s \oplus r \oplus X_{\mathcal{A}^*}^j(w) \oplus X_{\mathcal{B}^*}^j(w)}. \tag{5}$$

Encryption (5) is not a proper additive sharing since the Z -operator depends on $(r \oplus X_{\mathcal{A}^*}^j(w)) \cdot X_{\mathcal{B}^*}^j(w)$; the logical AND between a value known only by \mathcal{A}^* (i.e., $r \oplus X_{\mathcal{A}^*}^j(w)$) and a value known only by \mathcal{B}^* (i.e., $X_{\mathcal{B}^*}^j(w)$).

To get back to an additive sharing, \mathcal{A}^* and \mathcal{B}^* can simply call the AND-BOX once with inputs $r \oplus X_{\mathcal{A}^*}^j(\mathbf{w})$ and $X_{\mathcal{B}^*}^j(\mathbf{w})$ respectively as depicted in Fig. 4. After this, \mathcal{A}^* and \mathcal{B}^* share a proper encryption of the resulting state. The new encryption key for \mathcal{A}^* 's wire \mathbf{w} becomes:




Fig. 4. $\alpha \oplus \beta = (r \oplus X_{\mathcal{A}^*}^j(\mathbf{w})) \cdot X_{\mathcal{B}^*}^j(\mathbf{w})$ from an AND-BOX.

$$\begin{aligned}
 R : K_{\mathcal{A}^*}^{j+1}(\mathbf{w}) &:= (r \oplus X_{\mathcal{A}^*}^j(\mathbf{w}), r' \oplus \alpha \oplus Z_{\mathcal{A}^*}^j(\mathbf{w}) \oplus X_{\mathcal{A}^*}^j(\mathbf{w})), \\
 K_{\mathcal{B}^*}^{j+1}(\mathbf{w}) &:= (s \oplus X_{\mathcal{B}^*}^j(\mathbf{w}), s' \oplus \beta \oplus Z_{\mathcal{B}^*}^j(\mathbf{w}) \oplus X_{\mathcal{B}^*}^j(\mathbf{w})).
 \end{aligned}$$

5.4 On the Necessity of Swapping Privately

One may ask whether relying upon SWAP is necessary for the protocol to be private against specious adversaries. For instance, what would happen if one party announces the encryption keys before the other party? We now show that as soon as one party gets the other party's decryption key before having announced its own, a specious adversary can break privacy.

Consider the protocol for a quantum circuit made out of one single CNOT-gate. Suppose that \mathcal{A}^* holds the control wire \mathbf{w} while \mathcal{B}^* holds the target wire \mathbf{w}' . Suppose also the key-releasing phase first asks \mathcal{B}^* to announce the encryption keys $K_{\mathcal{B}^*}(\mathbf{w})$ before \mathcal{A}^* announces $K_{\mathcal{A}^*}(\mathbf{w}')$. Suppose \mathcal{A}^* 's input state is $|0\rangle$.

The adversary $\tilde{\mathcal{A}}$ can now act as follows. $\tilde{\mathcal{A}}$ runs the protocol for CNOT without performing the Bell measurement until she receives the encryption key b_z from \mathcal{B}^* . Clearly, $\tilde{\mathcal{A}}$'s behavior is specious up to that point since she could re-produce the honest state by just applying the Bell measurement on her input state. However, given b_z she could also in principle compute the CNOT upon any input state of her choice. This means that the state she holds after b_z has been announced and before applying her Bell measurement contains information about \mathcal{B}^* 's input. On the one hand, when $\tilde{\mathcal{A}}$'s input state is $|0\rangle$ no information whatsoever on \mathcal{B}^* 's input state should be available to her (i.e., in this case CNOT behaves like the identity). On the other hand, had her input state been $|-\rangle$, information about \mathcal{B}^* 's state would have become available since the control and target wires exchange their roles when the input states are in the Hadamard basis. However, when $\tilde{\mathcal{A}}$'s input state is $|0\rangle$, any simulation of her view can only call the ideal functionality with input state $|0\rangle$. It follows that no simulator can reproduce $\tilde{\mathcal{A}}$'s state right after the announcement of b_z .

6 Proof of Privacy

Privacy of the Evaluation Phase. We start by showing privacy of protocol $P_U^{\mathcal{G}} = (\mathcal{A}^*, \mathcal{B}^*, n_U + 1)$ at all steps $1 \leq i \leq n_U - 1$ occurring during the *evaluation phase* of quantum circuit C_U implementing U with u gates in \mathcal{UG} . The last step of the evaluation phase is $n_U - 1$ since only one oracle call is left to complete the execution. This phase is the easy part of the simulation since

all transmissions are independent of the joint input state $\rho_{\text{in}} \in \mathcal{D}(\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R})$. The lemma below can easily be proven and provides a perfect simulation of any adversary's view generated during the evaluation of any gate in C_U . No call to the ideal functionality for U is required.

Lemma 6.1. $P_U^\mathcal{O} = (\mathcal{A}^*, \mathcal{B}^*, n_U + 1)$ admits a simulator $\mathcal{S}(\tilde{\mathcal{A}})$ for any adversary $\tilde{\mathcal{A}}$ (not necessarily specious) that does not call the ideal functionality for $U \in \mathcal{U}(\mathcal{A}_0 \otimes \mathcal{B}_0)$ such that for any joint input state $\rho_{\text{in}} \in \mathcal{D}(\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R})$, every $1 \leq i \leq n_U - 1$:

$$\Delta \left(\nu_i(\tilde{\mathcal{A}}, \rho_{\text{in}}), \text{tr}_{\mathcal{B}_i} \left(\tilde{\rho}_i(\tilde{\mathcal{A}}, \rho_{\text{in}}) \right) \right) = 0 .$$

The same holds against any adversary $\tilde{\mathcal{B}}$.

Privacy of the Key-Releasing Phase. Before proving privacy of the key-releasing phase, we need the following lemma establishing that at the end of the protocol, specious adversaries must leave their extra working registers (used to implement the attack) independent of the joint input state. In other words, no extra information is available to the adversary at the very end of any correct protocol. Hence, if the adversary can break the privacy of a protocol, then he must “rush” to do so before the last step.

Lemma 6.2 (Rushing Lemma). Let $\Pi_U^\mathcal{O} = (\mathcal{A}, \mathcal{B}, n)$ be a correct protocol for the two party evaluation of U . Let $\tilde{\mathcal{A}}$ be any ε -specious adversary in $\Pi_U^\mathcal{O}$. Then, there exists an isometry $T : \tilde{\mathcal{A}}_n \rightarrow \mathcal{A}_n \otimes \hat{\mathcal{A}}$ and a mixed state $\tilde{\varrho} \in \mathcal{D}(\hat{\mathcal{A}})$ such that for all joint input states $\rho_{\text{in}} \in \mathcal{D}(\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R})$,

$$\Delta \left((T \otimes \mathbb{1}_{\mathcal{B}_n \otimes \mathcal{R}}) \cdot \left([\tilde{\mathcal{A}} \otimes \mathcal{B}](\rho_{\text{in}}) \right), \tilde{\varrho} \otimes (U \otimes \mathbb{1}_{\mathcal{R}}) \cdot \rho_{\text{in}} \right) \leq 12\sqrt{2\varepsilon} .$$

The same also applies to any ε -specious adversary $\tilde{\mathcal{B}}$.

Proof. We shall only prove the statement for an ε -specious $\tilde{\mathcal{A}}$; the statement for an ε -specious $\tilde{\mathcal{B}}$ is identical. Furthermore, by convexity, it is sufficient to prove the theorem for pure ρ_{in} .

Consider any pair of pure input states $|\psi_1\rangle$ and $|\psi_2\rangle$ in $\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R}$. Now, let $\mathcal{R}' := \mathcal{R} \otimes \mathcal{R}_2$, where $\mathcal{R}_2 = \text{span}\{|1\rangle, |2\rangle\}$ represents a single qubit, and define the state $|\psi\rangle := \frac{1}{\sqrt{2}}(|\psi_1\rangle|1\rangle + |\psi_2\rangle|2\rangle) \in \mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R}'$. Note that $\text{tr}_{\mathcal{R}_2}(|\psi\rangle\langle\psi|) = \frac{1}{2}|\psi_1\rangle\langle\psi_1| + \frac{1}{2}|\psi_2\rangle\langle\psi_2|$. Due to the correctness of the protocol and to the speciousness of $\tilde{\mathcal{A}}$, there exists a quantum operation $\mathcal{T}_n : \mathcal{L}(\tilde{\mathcal{A}}_n) \rightarrow \mathcal{L}(\mathcal{A}_n)$ such that

$$\Delta \left((\mathcal{T}_n \otimes \mathbb{1}_{\mathcal{L}(\mathcal{B}_n \otimes \mathcal{R}')})([\tilde{\mathcal{A}} \otimes \mathcal{B}](|\psi\rangle\langle\psi|)), (U \otimes \mathbb{1}_{\mathcal{R}'} \cdot |\psi\rangle\langle\psi|) \right) \leq 2\varepsilon .$$

Now, consider any isometry $T : \tilde{\mathcal{A}}_n \rightarrow \mathcal{A}_n \otimes \hat{\mathcal{A}}$ such that $\mathcal{T}_n(\sigma) = \text{tr}_{\hat{\mathcal{A}}}(T\sigma T^\dagger)$ for every $\sigma \in \mathcal{L}(\tilde{\mathcal{A}}_n)$ — in other words, any operation that implements \mathcal{T}_n while keeping any information that would otherwise be destroyed in $\hat{\mathcal{A}}$. By Uhlmann's theorem, there must exist a state $\tilde{\varrho} \in \mathcal{D}(\hat{\mathcal{A}})$ such that

$$\Delta \left((T \otimes \mathbb{1}_{\mathcal{B}_n \otimes \mathcal{R}'}) \cdot \left([\tilde{\mathcal{A}} \otimes \mathcal{B}](|\psi\rangle\langle\psi|) \right), \tilde{\varrho} \otimes ((U \otimes \mathbb{1}_{\mathcal{R}'}) \cdot |\psi\rangle\langle\psi|) \right) \leq 2\sqrt{2\varepsilon} .$$

Now, the trace distance is monotonous under completely positive, trace non-increasing maps. In particular, we can apply the projector $P_1 = \mathbb{1}_{\mathcal{L}(\mathcal{A}_n \otimes \mathcal{B}_n \otimes \mathcal{R})} \otimes |1\rangle\langle 1|$ to both states in the above trace distance and the inequality will still hold. In other words, we project both states onto $|1\rangle$ on \mathcal{R}_2 , thereby turning $|\psi\rangle\langle\psi|$ into $\frac{1}{2}|\psi_1\rangle\langle\psi_1|$. Factoring out the $\frac{1}{2}$, we get that

$$\Delta\left((T \otimes \mathbb{1}_{\mathcal{B}_n \otimes \mathcal{R}}) \cdot \left([\tilde{\mathcal{A}} \otimes \tilde{\mathcal{B}}](|\psi_1\rangle\langle\psi_1|)\right), \tilde{\varrho} \otimes ((U \otimes \mathbb{1}_{\mathcal{R}}) \cdot |\psi_1\rangle\langle\psi_1|)\right) \leq 4\sqrt{2\varepsilon} .$$

Likewise, projecting onto $|2\rangle$ yields

$$\Delta\left((T \otimes \mathbb{1}_{\mathcal{B}_n \otimes \mathcal{R}}) \cdot \left([\tilde{\mathcal{A}} \otimes \tilde{\mathcal{B}}](|\psi_2\rangle\langle\psi_2|)\right), \tilde{\varrho} \otimes ((U \otimes \mathbb{1}_{\mathcal{R}}) \cdot |\psi_2\rangle\langle\psi_2|)\right) \leq 4\sqrt{2\varepsilon} .$$

Our only problem at this point is that $\tilde{\varrho}$ in principle depends on $|\psi_1\rangle$ and $|\psi_2\rangle$. However, repeating the above argument with $|\psi_1\rangle$ and $|\psi_3\rangle$ for any $|\psi_3\rangle$ will yield a $\tilde{\varrho}'$ with

$$\Delta\left((T \otimes \mathbb{1}_{\mathcal{B}_n \otimes \mathcal{R}}) \cdot \left([\tilde{\mathcal{A}} \otimes \tilde{\mathcal{B}}](|\psi_1\rangle\langle\psi_1|)\right), \tilde{\varrho}' \otimes ((U \otimes \mathbb{1}_{\mathcal{R}}) \cdot |\psi_1\rangle\langle\psi_1|)\right) \leq 4\sqrt{2\varepsilon}$$

and hence, by the triangle inequality, $\Delta(\tilde{\varrho}, \tilde{\varrho}') \leq 8\sqrt{2\varepsilon}$. Therefore, for any state $|\varphi\rangle \in \mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R}$, there exists a state $\tilde{\rho} \in \tilde{\mathcal{A}}$ with $\Delta(\tilde{\rho}, \tilde{\varrho}) \leq 8\sqrt{2\varepsilon}$ such that

$$\Delta\left((T \otimes \mathbb{1}_{\mathcal{B}_n \otimes \mathcal{R}}) \cdot \left([\tilde{\mathcal{A}} \otimes \tilde{\mathcal{B}}](|\varphi\rangle\langle\varphi|)\right), \tilde{\rho} \otimes ((U \otimes \mathbb{1}_{\mathcal{R}}) \cdot |\varphi\rangle\langle\varphi|)\right) \leq 4\sqrt{2\varepsilon} .$$

The lemma then follows by the triangle inequality:

$$\begin{aligned} & \Delta\left((T \otimes \mathbb{1}_{\mathcal{B}_n \otimes \mathcal{R}}) \cdot \left([\tilde{\mathcal{A}} \otimes \tilde{\mathcal{B}}](|\varphi\rangle\langle\varphi|)\right), \tilde{\varrho} \otimes ((U \otimes \mathbb{1}_{\mathcal{R}}) \cdot |\varphi\rangle\langle\varphi|)\right) \\ & \leq \Delta\left((T \otimes \mathbb{1}_{\mathcal{B}_n \otimes \mathcal{R}}) \cdot \left([\tilde{\mathcal{A}} \otimes \tilde{\mathcal{B}}](|\varphi\rangle\langle\varphi|)\right), \tilde{\rho} \otimes ((U \otimes \mathbb{1}_{\mathcal{R}}) \cdot |\varphi\rangle\langle\varphi|)\right) + \Delta(\tilde{\rho}, \tilde{\varrho}) \\ & \leq 4\sqrt{2\varepsilon} + 8\sqrt{2\varepsilon} = 12\sqrt{2\varepsilon} . \quad \square \end{aligned}$$

In order to conclude the privacy of $P_U^\mathcal{O}$, families $\mathcal{S}(\tilde{\mathcal{A}})$ and $\mathcal{S}(\tilde{\mathcal{B}})$ need one more simulator each: $\mathcal{S}_{n_U} \in \mathcal{S}(\tilde{\mathcal{A}})$ and $\mathcal{S}'_{n_U} \in \mathcal{S}(\tilde{\mathcal{B}})$ corresponding to the simulation of the key-releasing phase. This time, these simulators need to query the ideal functionality for U and also need the adversary to be specious. We show that privacy of the key-releasing phase follows from the ‘‘Rushing Lemma’’ (Lemma 6.2). This is the role of the ideal SWAP to make sure that before the adversary gets the output of the computation, the information needed by the honest player to recover its own output has been given away by the adversary.

It should be mentioned that we’re not explicitly simulating the final state of the adversary since simulating the SWAP allows also to get $\tilde{\mathcal{A}}$ ’s final state by simply adding $\tilde{\mathcal{A}}$ ’s last quantum operation to the simulated view. We therefore set step n_U in $P_U^\mathcal{O}$ to be the step reached after the call to SWAP. This abuses the notation a bit since after SWAP, $\tilde{\mathcal{A}}$ and $\tilde{\mathcal{B}}^*$ must each apply a final quantum operation with no more oracle call. We’ll denote by $\tilde{\mathcal{A}}_{n_U+1}$ and $\tilde{\mathcal{B}}^*_{n_U+1}$ these last operations allowing to reconstruct the output of the computation (no communication).

Lemma 6.3. *For any ε -specious quantum adversary $\tilde{\mathcal{A}}$ against $P_U^\mathcal{O} = (\mathcal{A}^*, \mathcal{B}^*, n_U + 1)$, there exist simulators $\mathcal{S}_{n_U} \in \mathcal{S}(\tilde{\mathcal{A}})$ such that for all $\rho_{\text{in}} \in \mathcal{D}(\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R})$,*

$$\Delta\left(\nu_{n_U}(\tilde{\mathcal{A}}, \rho_{\text{in}}), \text{tr}_{\mathcal{B}_{n_U}}\left(\tilde{\rho}_{n_U}(\tilde{\mathcal{A}}, \rho_{\text{in}})\right)\right) \leq 24\sqrt{2\varepsilon} .$$

Simulator \mathcal{S}_{n_U} calls the ideal functionality for U and can be used directly to simulate step $n_U + 1$ as well. The same holds for adversary $\tilde{\mathcal{B}}$.

Proof (sketch). We only prove privacy against adversary $\tilde{\mathcal{A}}$, privacy against $\tilde{\mathcal{B}}$ follows directly since the key-releasing phase is symmetric. The idea behind the proof is to run $\tilde{\mathcal{A}}$ and \mathcal{B}^* upon a dummy joint input state until the end of the protocol. Since the adversary is specious, it can re-produce the honest state at the end. The Rushing Lemma tells us that at this point, the output of the computation is essentially in tensor product with all the other registers. Moreover, the state of all other registers is independent of the input state upon which the protocol is executed. The *dummy output* can then be replaced by the output of the ideal functionality for U before $\tilde{\mathcal{A}}$ goes back to the stage reached just after SWAP.

More formally, we define a simulator $\mathcal{S}_{n_U} \in \mathcal{S}(\tilde{\mathcal{A}})$ producing $\tilde{\mathcal{A}}$'s view just after the call to SWAP. Let $\tilde{\mathcal{A}}_{\text{SWAP}} \in \mathcal{L}(\mathcal{A}_0, \tilde{\mathcal{A}}_{n_U})$ and $\mathcal{B}_{\text{SWAP}}^* \in \mathcal{L}(\mathcal{B}_0, \tilde{\mathcal{B}}_{n_U})$ be the quantum operations run by $\tilde{\mathcal{A}}$ and \mathcal{B}^* respectively until after SWAP is executed. Notice that at this point, $\tilde{\mathcal{A}}$'s and \mathcal{B}^* 's registers do not have any further oracle registers since no more communication or oracle call will take place. Let $\tilde{\mathcal{A}}_{n_U} \in \mathcal{L}(\tilde{\mathcal{A}}_{n_U}, \tilde{\mathcal{A}}_{n_U+1} \otimes \mathcal{Z})$ be the isometry implementing $\tilde{\mathcal{A}}$'s last quantum operation taking place after the call to SWAP (and producing her final state) and let $B_{n_U} \in \mathcal{L}(\mathcal{B}_{n_U}, \mathcal{B}_{n_U+1} \otimes \mathcal{W})$ be the isometry implementing \mathcal{B}^* 's last quantum operation. Finally, let $T \in \mathcal{L}(\tilde{\mathcal{A}}_{n_U+1}, \mathcal{A}_{n_U+1} \otimes \hat{\mathcal{A}})$ be the isometry implementing \mathcal{T}_{n_U+1} as defined in Lemma 6.2 (i.e., the transcript produced at the very end of the protocol). As usual, let $\rho_{\text{in}} \in \mathcal{D}(\mathcal{A}_0 \otimes \mathcal{B}_0 \otimes \mathcal{R})$ be the joint input state. The simulator \mathcal{S}_{n_U} performs the following operations:

1. \mathcal{S}_{n_U} generates the quantum state $\sigma(\phi^*) = [\tilde{\mathcal{A}}_{\text{SWAP}} \otimes \mathcal{B}_{\text{SWAP}}^*](|\phi^*\rangle\langle\phi^*|) \in \mathcal{D}(\tilde{\mathcal{A}}_{n_U} \otimes \mathcal{B}_{n_U})$ implementing $\tilde{\mathcal{A}}$ interacting with \mathcal{B}^* until SWAP is applied. The execution is performed upon a predetermined (dummy) arbitrary input state $|\phi^*\rangle \in \mathcal{A}_0 \otimes \mathcal{B}_0$.
2. \mathcal{S}_{n_U} sets $\sigma'(\phi^*) = (T \tilde{\mathcal{A}}_{n_U} \otimes B_{n_U}) \cdot \sigma(\phi^*) \in \mathcal{D}(\mathcal{A}_{n_U+1} \otimes \mathcal{B}_{n_U+1} \otimes \mathcal{Z} \otimes \hat{\mathcal{A}} \otimes \mathcal{W})$.
3. \mathcal{S}_{n_U} replaces register $\mathcal{A}_{n_U+1} \approx \mathcal{A}_0$ by \mathcal{A}^* 's output of the ideal functionality for U evaluated upon ρ_{in} . That is, \mathcal{S}_{n_U} generates the state $\sigma'(\rho_{\text{in}}) = (U \otimes \mathbb{1}_{\mathcal{R}}) \cdot \rho_{\text{in}} \otimes \text{tr}_{\mathcal{A}_{n_U+1}\mathcal{B}_{n_U+1}}(\sigma'(\phi^*)) \in \mathcal{D}(\mathcal{A}_{n_U+1} \otimes \mathcal{B}_{n_U+1} \otimes \mathcal{R} \otimes \mathcal{Z} \otimes \hat{\mathcal{A}} \otimes \mathcal{W})$.
4. \mathcal{S}_{n_U} finally sets $\nu_{n_U}(\tilde{\mathcal{A}}, \rho_{\text{in}}) = \text{tr}_{\mathcal{B}_{n_U+1}\mathcal{W}}((T \tilde{\mathcal{A}}_{n_U} \otimes \mathbb{1}_{\mathcal{B}_{n_U+1}\mathcal{R}})^\dagger \cdot \sigma'(\rho_{\text{in}})) \in \mathcal{D}(\tilde{\mathcal{A}}_{n_U} \otimes \mathcal{R})$.

Notice that execution of the ideal SWAP ensures that the keys swapped are independent of each other and of the joint input state ρ_{in} . This is because for any input state, all these keys are uniformly distributed bits if they are outcomes

of Bell measurements and otherwise are set to 0. By Lemma 6.2 and the fact that $\tilde{\mathcal{A}}$ is ε -specious, we have:

$$\Delta(\text{tr}_{\mathcal{Z}\hat{\mathcal{A}}\mathcal{W}}(\sigma'(\phi^*)), \tilde{\varrho} \otimes U \cdot |\phi^*\rangle\langle\phi^*|) \leq 12\sqrt{2\varepsilon} \text{ and}$$

$$\Delta\left((\mathcal{T}_{n_U+1} \otimes \mathbb{1}_{L(\mathcal{B}_{n_U+1})})\left([\tilde{\mathcal{A}} \otimes \mathcal{B}^*\right](\rho_{\text{in}})\right), \tilde{\varrho} \otimes U \cdot \rho_{\text{in}}\right) \leq 12\sqrt{2\varepsilon} .$$

It follows using the triangle inequality that,

$$\Delta\left((\mathcal{T}_{n_U+1} \otimes \mathbb{1}_{L(\mathcal{B}_{n_U+1})})\left([\tilde{\mathcal{A}} \otimes \mathcal{B}^*\right](\rho_{\text{in}})\right), \text{tr}_{\mathcal{Z}\hat{\mathcal{A}}\mathcal{W}}(\sigma'(\rho_{\text{in}}))\right) \leq 24\sqrt{2\varepsilon} . \quad (6)$$

Using the fact that isometries cannot increase the trace-norm distance and that $(T\hat{A}_{n_U})^\dagger$ allows $\tilde{\mathcal{A}}$ to go back from the end of the protocol to the step reached after SWAP, we get from (6) that

$$\Delta\left(\nu_{n_U}(\tilde{\mathcal{A}}, \rho_{\text{in}}), \text{tr}_{\mathcal{B}_{n_U}}\left(\tilde{\rho}_{n_U}(\tilde{\mathcal{A}}, \rho_{\text{in}})\right)\right) =$$

$$\Delta\left((\mathcal{T}_{n_U+1} \otimes \mathbb{1}_{L(\mathcal{B}_{n_U+1})})\left([\tilde{\mathcal{A}} \otimes \mathcal{B}^*\right](\rho_{\text{in}})\right), \text{tr}_{\mathcal{Z}\hat{\mathcal{A}}\mathcal{W}}(\sigma'(\rho_{\text{in}}))\right) \leq 24\sqrt{2\varepsilon} .$$

The proof of the statement follows. □

7 Main Result and Open Questions

Putting Lemma 6.1 and Lemma 6.3 together gives the desired result:

Theorem 7.1 (Main Result). *Protocol $P_U^\mathcal{O}$ is statistically private against any statistically specious quantum adversary and for any $U \in \text{U}(\mathcal{A}_0 \otimes \mathcal{B}_0)$. If U is in the Clifford group then the only non-trivial oracle call in \mathcal{O} is one call to an ideal SWAP. If U is not in the Clifford group then \mathcal{O} contains an additional oracle call to AND-BOX for each R-gate in the circuit for U .*

It should be mentioned that it is not too difficult to modify our protocol in order to privately evaluate quantum operations rather than only unitary transforms. Classical two party computation together with the fact that quantum operations can be viewed as unitaries acting in larger spaces can be used to achieve this extra functionality. Privacy can be preserved by keeping these extra registers encrypted after the execution of the protocol. We leave this discussion to the full version of the paper.

A few interesting questions remain open:

- It would be interesting to know whether there exists a unitary transform that can act as a universal primitive for private two-party evaluation of unitaries. This would allow to determine whether classical cryptographic assumptions are required for this task.
- Finally, is there a way to compile quantum protocols secure against specious adversaries into protocols secure against arbitrary quantum adversaries? An affirmative answer would allow to simplify greatly the design of quantum protocols. Are extra assumptions needed to preserve privacy against any adversary?

Acknowledgements

The authors would like to thank the referees for their comments and suggestions. We would also like to thank Thomas Pedersen for numerous helpful discussions in the early stage of this work.

References

1. Aharonov, D., Ben-Or, M.: Fault-tolerant quantum computation with constant error. In: 29th Annual ACM Symposium on Theory of Computing (STOC), pp. 176–188 (1997)
2. Ambainis, A., Mosca, M., Tapp, A., de Wolf, R.: Private quantum channels. In: 41st Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 547–553 (2000)
3. Ben-Or, M., Crépeau, C., Gottesman, D., Hassidim, A., Smith, A.: Secure multiparty quantum computation with (only) a strict honest majority. In: 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 249–260 (2006)
4. Broadbent, A., Fitzsimons, J., Kashefi, E.: Universal blind quantum computation (December 2009), <http://arxiv.org/abs/0807.4154>
5. Crépeau, C., Gottesman, D., Smith, A.: Secure multi-party quantum computation. In: 34th Annual ACM Symposium on Theory of Computing (STOC), pp. 643–652 (2002)
6. Gottesman, D., Chuang, I.L.: Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations. *Nature* 402, 390–393 (1999)
7. Gottesman, D., Chuang, I.L.: Quantum teleportation is a universal computational primitive (August 1999), <http://arxiv.org/abs/quant-ph/9908010>
8. Gutoski, G., Watrous, J.: Quantum interactive proofs with competing provers. In: Diekert, V., Durand, B. (eds.) STACS 2005. LNCS, vol. 3404, pp. 605–616. Springer, Heidelberg (2005)
9. Kilian, J.: Founding cryptography on oblivious transfer. In: 20th Annual ACM Symposium on Theory of Computing (STOC), pp. 20–31 (1988)
10. Lo, H.-K.: Insecurity of quantum secure computations. *Physical Review A* 56(2), 1154–1162 (1997)
11. Lo, H.-K., Chau, H.F.: Is quantum bit commitment really possible? *Physical Review Letters* 78, 3410–3413 (1997)
12. Mayers, D.: Unconditionally secure quantum bit commitment is impossible. *Physical Review Letters* 78, 3414–3417 (1997)
13. Nielsen, M.A., Chuang, I.L.: *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge (2000)
14. Popescu, S., Rohrlich, D.: Quantum nonlocality as an axiom. *Foundations of Physics* 24(3), 379–385 (1994)
15. Popescu, S., Rohrlich, D.: Causality and nonlocality as axioms for quantum mechanics. In: *Symposium on Causality and Locality in Modern Physics and Astronomy: Open Questions and Possible Solutions* (1997), <http://arxiv.org/abs/quant-ph/9709026>
16. Renner, R., König, R.: Universally composable privacy amplification against quantum adversaries. In: Kilian, J. (ed.) TCC 2005. LNCS, vol. 3378, pp. 407–425. Springer, Heidelberg (2005)

17. Salvail, L., Sotáková, M., Schaffner, C.: On the power of two-party quantum cryptography. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 70–87. Springer, Heidelberg (2009)
18. Shor, P.W.: Fault-tolerant quantum computation. In: 37th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 56–65 (1996)
19. Smith, A.: Techniques for secure distributed computing with quantum data. Presented at the Field's institute Quantum Cryptography and Computing Workshop (October 2006)
20. Watrous, J.: Limits on the power of quantum statistical zero-knowledge. In: 43rd Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 459–468 (2002)
21. Wolf, S., Wullschleger, J.: Oblivious transfer and quantum non-locality. In: International Symposium on Information Theory (ISIT 2005), pp. 1745–1748 (2005)
22. Yao, A.: How to generate and exchange secrets. In: 27th Annual IEEE Symposium on Foundations of Computer Science (FOCS) (1986)