On the Power of Oblivious State Preparation

James Bartusek[*] Dakshita Khurana[†]

Abstract

We put forth Oblivious State Preparation (OSP) as a cryptographic primitive that unifies techniques developed in the context of a quantum server interacting with a classical client. OSP allows a *classical* polynomial-time sender to input a choice of one out of two public observables, and a quantum polynomial-time receiver to recover an eigenstate of the corresponding observable – while keeping the sender's choice hidden from any malicious receiver.

We obtain the following results:

- The existence of (plain) trapdoor claw-free functions implies OSP, and the existence of dual-mode trapdoor claw-free functions implies round-optimal (two-round) OSP.
- OSP implies the existence of proofs of quantumness, test of a qubit, blind classical delegation of quantum computation, and classical verification of quantum computation.
- Two-round OSP implies quantum money with classical communication, classically-verifiable position verification, and (additionally assuming classical FHE with log-depth decryption) quantum FHE.

Thus, the OSP abstraction helps separate the cryptographic layer from the information-theoretic layer when building cryptosystems across classical and quantum participants. Indeed, several of the aforementioned applications were previously only known via tailored LWE-based constructions, whereas our OSP-based constructions yield new results from a wider variety of assumptions, including hard problems on cryptographic group actions.

Finally, towards understanding the minimal hardness assumptions required to realize OSP, we prove the following:

- OSP implies oblivious transfer between one classical and one quantum party.
- Two-round OSP implies public-key encryption with classical keys and ciphertexts.

In particular, these results help to ”explain” the use of public-key cryptography in the known approaches to establishing a ”classical leash” on a quantum server. For example, combined with a result of Austrin et al. (CRYPTO 22), we conclude that *perfectly-correct* OSP cannot exist unconditionally in the (quantum) random oracle model.

[*] NYU. Email: `bartusek.james@gmail.com`

[†] UIUC and NTT Research. Email: `dakshita@illinois.edu`

Contents

1 Introduction

One of the central concepts driving research in quantum cryptography over the past decade has been that of a "classical leash" on quantum systems [RUV13]. In other words, how can we enable a classical device, using just classical communication, to exert some element of *control* over a quantum mechanical system?

In a major conceptual advance from 2018 [BCM+18], the use of (public-key) cryptography was identified as a useful tool for establishing this desired control over a quantum server. There has since been an explosion of results on classical-client quantum-server protocols, ranging from proofs of quantumness under quantum-hard assumptions [BCM+18], certifiable randomness generation [BCM+18], quantum homomorphic encryption with classical ciphertexts [Mah18a], classical verification of quantum computation [Mah18b], self-testing a single quantum device [MV21], position verification [LLQ22], secure quantum computation [Bar21], and quantum money [RS19, Shm22] with classical communication, and proofs of contextuality [ABCC24], among others.

The resounding success of this line of work begs a deeper understanding of the basic principles underlying the paradigm introduced in [BCM+18]. For example, [BCM+18] based their results on the existence of a fairly ad-hoc and unwieldy cryptographic primitive: a noisy trapdoor claw-free function (TCF) with an adaptive hardcore bit. This primitive has only been shown to exist from the learning with errors (LWE) assumption,[1] and several followups, including many of the aforementioned results, inherited the use of this primitive. This raises the following (informal) question.

Is there a conceptually-simple and easy-to-instantiate primitive that suffices for building powerful classical-client quantum-server applications?

We note that some partial progress has been made towards a more "generic" approach to constructing some of these end applications. For example, a recent line of work [KMCVY21, KLVY23, NZ23] has yielded classical verification of quantum computation from the assumption of quantum fully-homomorphic encryption (QFHE), and [GV24] has shown that QFHE follows from any classical FHE (with decryption in NC_1) plus an appropriate notion of "dual-mode" trapdoor claw-free functions (with no need for an adaptive hardcore bit). However, as we show, these approaches can still be generalized much further.

Another aspect of [BCM+18]'s approach that demands further investigation is their use of public-key (i.e. trapdoor-based) cryptography. This has been justified informally by observing that, because the classical client is computationally weaker than the quantum server, we need to introduce some mechanism for the client to gain the "upper hand" on the server. One way to do this is to introduce asymmetric cryptography, allowing the client to send the server a public key while keeping the secret key to themselves. However, as far as we are aware, nothing more than this informal intuition has been proposed in an attempt to address the following fundamental (and again, informal) question.

Is public-key cryptography necessary to establish a classical leash on a quantum server?

As an example meant to further illustrate the importance of this question, we note that progress in this direction may shed more light on the recent breakthrough techniques of Yamakawa-Zhandry

[1] [AMR22] has shown how to obtain a weaker variant of the adaptive hardcore bit property from hard problems on cryptographic group actions.

[YZ24] for establishing quantum advantage. Indeed, while they show that proofs of quantumness exist in the random oracle model, which can be heuristically instantiated using a cryptographic hash function (i.e. *symmetric* cryptography), their techniques have so far resisted attempts at constructing, say, a test of a qubit, verifiable delegation of quantum computation, or other "classical leash"-style primitives mentioned above. Establishing the necessity of public-key cryptography for these primitives would explain this gap. While we do not completely close this question, our results do show that the principles underlying current approaches to classical-leash primitives (inspired by [BCM+18, KMCVY21], etc.) also yield public-key style primitives. We provide further details on these results later in the introduction.

1.1 Oblivious state preparation

Aiming to make progress on these two motivating questions, we put forth the idea of *Oblivious State Preparation* (OSP) as a unifying cryptographic primitive in the realm of classical-client quantum-server protocols.

OSP is simple to describe. It is a protocol that takes place between a classical sender and a quantum receiver. The classical sender has as input a bit $b \in \{0, 1\}$ which specifies a choice of one of two public observables. We usually take one to be Z and the other to be X, but in principle they could be arbitrary. At the end of the protocol, the receiver outputs a quantum state. We require two properties.

- **Correctness:** If the receiver is honest, their output is an eigenstate of the observable chosen by the sender, and the sender receives a description of this state. In the usual "standard" case, this means that when $b = 0$, the receiver outputs either $|0\rangle$ or $|1\rangle$, and when $b = 1$, the receiver outputs either $|+\rangle$ or $|-\rangle$.

- **Security:** Any quantum polynomial-time (QPT) malicious receiver has $\mathrm{negl}(\lambda)$ advantage in guessing the sender's input bit b.

OSP highlights an inherent cryptographic property of quantum information arising from the uncertainty principle. That is, it is not necessarily possible to determine the basis of a given state, even though this basis information is well-defined and fixed by the description of the state. Indeed, given the resource of quantum communication, information-theoretically secure OSP is trivial. In this work, however, OSP always refers to the classical-communication case. At a high level, we ask, (1) what cryptography is necessary to obtain OSP, i.e. the ability for a classical client to, roughly speaking, set up an instance of the uncertainty principle on a quantum server, and (2) what are the applications of this ability.

In fact, the idea of OSP for the Z and X observables as described above has been previously proposed by [CCKW19] under the name "malicious 4-states QFactory with basis-blindness." They showed that such a protocol can be obtained from a TCF with a particular type of "homomorphic, hardcore predicate". However, in order to derive most applications, they required an additional and conjectured *verifiability* property (see also [CCKW21]). In this work, we give OSP a more general treatment as a primitive, show that it is possible to relax the assumptions under which it can be built, and vastly expand its set of applications.

1.2 Results

1.2.1 Constructions

As mentioned above, researchers beginning with [BCM+18] identified the usefulness of (variants of) trapdoor claw-free functions (TCFs) in realizing applications such as a test of a qubit, quantum fully-homomorphic encryption, and classical verification of quantum computation. While several variants of TCFs have appeared over the years, e.g. extended, dual-mode, with adaptive hardcore bit, etc., we show that perhaps the most stripped down notion of a TCF suffices to build OSP.

We define a (plain) TCF as a family of functions f that can be sampled along with a trapdoor td. The guarantee, roughly, is that there is a (QPT preparable) distribution $\mathcal{D}$ over inputs such that with some inverse polynomial probability over $x \leftarrow \mathcal{D}$, x has exactly one sibling x' (similarly weighted by $\mathcal{D}$) such that $f(x) = f(x')$, and moreover, both x and x' can be recovered given $f(x)$ and td. Finally, claw-freeness demands that no QPT adversary can recover any such "claw" x and x' given only the description of f. Building on techniques from [BGKM+23], we show the following.

Theorem 1.1 (Informal)**.** *(Plain) TCFs imply OSP.*

Our construction of OSP from plain TCFs requires multiple of rounds of interaction. However, a desirable feature for some applications is limited interaction. The best we can hope for is *two-round* OSP, i.e. one message from the sender followed by one from the receiver.[2] Adapting techniques from [GV24], we show that TCFs with an additional *dual-mode* property imply such two-round OSP.

Theorem 1.2 (Informal)**.** *Dual-mode TCFs imply two-round OSP.*

Briefly, a dual-mode TCF is similar to a plain TCF except that the function may be sampled in an "injective" mode, where there are no collisions, and it is computationally difficult to distinguish this from the normal "lossy" mode. We note that dual-mode (and thus plain) TCFs are known from LWE [BCM+18] and from the "extended linear hidden shift" assumption on cryptographic group actions [AMR22, GV24].

In what follows, we will highlight the power of OSP by establishing several classical-client quantum-server applications, as well as cryptographic implications. We will build everything from "standard" OSP, where the two observables are Z and X. However, it is meaningful to consider OSP for *any* pair of non-commuting observables, in particular pairs that may not be maximally anti-commuting. We begin to explore the landscape of OSP as a more general primitive, and in particular establish the following result.

Theorem 1.3 (Informal)**.** *OSP for any pair of two-outcome observables that are at a* $1/\mathrm{poly}$ *angle implies standard (*Z *and* X*) OSP.*

1.2.2 Applications

From OSP. We show that OSP is sufficient to obtain several classical-client quantum-server protocols of interest.

[2] No one-message OSP can be secure. Indeed, the message from sender to receiver would have to fix the description of the desired state $H^b |s\rangle$, and correctness would imply that the receiver could then generate multiple copies of this state, eventually enough to determine the basis with near certainty.

Theorem 1.4 (Informal)**.** *OSP implies proofs of quantumness, a test of a qubit, blind classical delegation of quantum computation, and classical verification of quantum computation.*

A few remarks on these results are in order. The proof of quantumness from OSP can be seen as a modular instantiation of the template proposed by [KMCVY21] and later tweaked by [BGKM$^+$23] and [ABCC24]. In particular, we show that a *single* instance of OSP suffices to implement a "computational Bell test" between the client and server based on the CHSH game.

Blind classical delegation of quantum computation allows a classical client to outsource a quantum computation of its choice to a quantum server without leaking anything about the actual description of the computation. In a major breakthrough, [Mah18a] gave the first construction of this primitive by building *quantum fully-homomorphic encryption* (QFHE) from LWE. This approach of course relies on techniques that at least imply classical fully-homomorphic encryption, and, to the best of our knowledge, fully-homomorphic encryption has remained the only solution to blind classical delegation of quantum computation that has been made explicit in the literature.[3] While perhaps folklore, we formalize the fact that classical FHE is not required for blind classical delegation of quantum computation, showing that OSP suffices.

In another major breakthrough, [Mah18b] designed a protocol that allows a classical client to verifiably delegate a BQP computation to a potentially cheating quantum server. Since then, classical verification of quantum computation (CVQC) has remained a central primitive of study in quantum cryptography [Zha22, BKL$^+$22, CLLW22, NZ23, GKNV24, MNZ24]. However, until now, all known constructions relied on the hardness of LWE. As a corollary of our result, we show that CVQC follows from any (plain) TCF, and thus from hard problems on cryptographic group actions. In a nutshell, we formalize the fact that the recent approach of [KLVY23, NZ23] establishing CVQC from QFHE can in fact be instantiated from any (potentially interactive, non-compact) blind classical delegation of quantum computation protocol, and thus, from any OSP.

From two-round OSP. Our next batch of results makes use of *two-round* OSP.

Theorem 1.5 (Informal)**.** *Two-round OSP implies (privately-verifiable) quantum money with classical communication, position verification with classical communication, and (assuming classical FHE with decryption in NC_1) QFHE.*

Briefly, the first two results go via the intermediate primitive of a "1-of-2 puzzle" [RS19], which we build from OSP using similar techniques to our CHSH-based proof of quantumness. Then, we appeal to [RS19], who showed that 1-of-2 puzzles imply privately-verifiable quantum money with classical communication, and [LLQ22], who showed that 1-of-2 puzzles imply position verification with classical communication. Prior to our work, the only construction of 1-of-2 puzzles, due to [RS19], relied specifically on LWE (via TCFs with the adaptive hardcore bit property).[4]

The third result on QFHE follows by adapting the recent techniques of [GV24], who showed how to construct QFHE from any classical FHE (with decryption in NC_1) and dual-mode TCFs. We observe that, in fact, *any* two-round OSP suffices in place of the dual-mode TCF.

[3] [CCKW21] show how to realize blind classical delegation of quantum computation but only against an *honest-but-curious* server, using their protocol for "pseudo-secret random qubit generator."

[4] We note that a recent concurrent and independent work has shown how to construct classically-verifiable position verification from certified randomness protocols [ACC$^+$24].

Discussion. Before moving on, it is worth pointing out that each of these results individually do not require deep new techniques. Rather, they mostly follow by adapting, modularizing, and generalizing existing approaches in the literature. However, in our view, the identification of OSP as a simple primitive that yields all of these applications is useful, both as a pedagogical tool and to enhance future research. OSP abstracts away the cryptographic essence of major protocols in the area, allowing for easy, information-theoretic design of protocols that call an underlying OSP functionality. This separates the "cryptographic layer" from the "information-theoretic layer" in the design of these protocols. In some cases, this modular approach also allows us to broaden the set of assumptions under which these applications are known to exist.

1.2.3 Implications

Finally, given the broad reach of OSP, we seek to understand the cryptography necessary in order to realize it. Our results on this are summarized as follows.

Theorem 1.6 (Informal)**.** *OSP implies commitments and oblivious transfer (OT) with classical communication (where one party is completely classical), while two-round OSP implies public-key encryption.*

So what can we conclude about OSP from these results? As mentioned earlier, one of the original motivations for our work was to "justify" the use of public-key cryptography in the recent line of work aimed at establishing a classical leash on quantum systems. Progress towards this goal can be appreciated by noting that the techniques introduced in [BCM+18, Mah18a, Mah18b, KMCVY21] all at the very least provide some way to perform an OSP between the classical client and quantum server, and thus, by our result, also provide a way to build OT with classical communication.

So far, the community does not have any approach for building OT with classical communication from minicrypt assumptions, or even from arbitrary trapdoor functions. Thus, our result helps explain why the known constructions of OSP require TCFs, which are more structured than even injective trapdoor functions. In fact, in the classical setting, we actually have an oracle separation between OT and injective trapdoor functions [GKM+00].

However, it remains an open question to give such strong oracle separations in the quantum setting. Progress came when [ACC+22] showed that *perfectly correct* key agreement between one classical and one quantum party does not exist in the quantum random oracle model. This notion of key agreement is implied by the perfectly correct variant of our notion of OT between one classical and one quantum party, and thus, we obtain the following corollary.

Corollary 1.7 (Informal)**.** *Perfectly correct OSP does not exist in the quantum random oracle model.*

This leaves a sliver of possibility that *non-perfectly-correct* OSP can yet be constructed without public-key assumptions. However, we have established that if one can build OSP from minicrypt primitives (say, by adapting the techniques of [YZ24]), or even from arbitrary trapdoor functions, then this would also represent a major breakthrough in cryptography more generally - a construction of classical-communication OT from new assumptions.

2 Technical Overview

2.1 Realizing oblivious state preparation

We show that OSP follows from any trapdoor claw-free function (TCF), i.e. we don't require an additional adaptive hardcore bit or dual-mode property. Our presentation abstracts out the key role of the TCF, which is to generate a claw-state correlation. That is, we define a "claw-state generator" (CSG) as any protocol between a classical sender and quantum receiver that outputs $\frac{1}{\sqrt{2}}(|0,x_0\rangle + |1,x_1\rangle)$ to the receiver and $x_0, x_1 \in \{0,1\}^n$ to the sender.[5] We say that the protocol has *search security* if no QPT receiver can output both (x_0, x_1) except with negligible probability.

Now, we show that CSG with search security implies OSP by relying on a sub-protocol from [BGKM+23].[6] After the CSG is performed, the sender chooses two random strings $r_0, r_1 \leftarrow \{0,1\}^n$ and sends them to the receiver. The receiver then maps

$$\frac{1}{\sqrt{2}}\left(|x_0\rangle + |x_1\rangle\right) \rightarrow \frac{1}{\sqrt{2}}\left(|x_0\rangle\,|x_0 \cdot r_0\rangle + |x_1\rangle\,|x_1 \cdot r_1\rangle\right),$$

and measures all but the last qubit in the Hadamard basis to obtain a string d, which it returns to the sender.

It is easy to check that if $x_0 \cdot r_0 = x_1 \cdot r_1$, then the receiver obtain a standard basis eigenstate, while if $x_0 \cdot r_0 \neq x_1 \cdot r_1$, then the receiver obtains a Hadamard basis eigenstate. Moreover, which eigenstate obtained can be computed by the sender, who knows x_0, x_1, and d. To see why this is secure, note that the bit that determines the basis is equal to $x_0 \cdot r_0 \oplus x_1 \cdot r_1 = (x_0, x_1) \cdot (r_0, r_1)$. Thus, by Goldreich-Levin, any adversarial receiver that can predict the basis of their received state can be used to extract an entire claw (x_0, x_1), which breaks the search security of the CSG.

Now, while the protocol above implements *random-input* OSP where the sender ends up with a random choice of basis, it is generically possible to reorient this into a chosen-input OSP, which we show in Lemma 4.3. Finally, we note that in the body, we also construct OSP with the optimal round-complexity of two (i.e. one message from the sender followed by one from the receiver), by assuming a slightly stronger variant of TCFs, namely *dual-mode* TCFs. This construction adapts the recent techniques of [GV24], and we refer the reader to Section 5.2 for details.

2.2 Applications

2.2.1 Proofs of quantumness

Our first use case for OSP is to instantiate a "computational Bell test", first introduced by [KMCVY21]. The resulting protocol is essentially a generalized presentation of [ABCC24]'s recent proof of quantumness protocol, in which they instantiated the OSP using an "encrypted CNOT" operation based on a structured type of TCF.

Consider the server's state at the end of an OSP protocol: If the client's input was $a = 0$, they obtain $|x\rangle$ for some bit x, and if the client's input was $a = 1$, they obtain $H\,|x\rangle$ for some bit x. This is *exactly* the same as Bob's state in the CHSH game once (honest) Alice provides an answer x on

[5] Technically, we call this a *differentiated-bit* CSG, since the first qubit holds a bit that differentiates the two members of the claw state. It is easy to show that one can generically add the differentiated-bit property to any CSG, while maintaining search security (see Lemma 4.6).

[6] Later, we will actually show that OSP implies CSG (with an even stronger security property called *indistinguishability* security), meaning that these notions are in fact equivalent.

input question a. Indeed, in the CHSH game, Alice and Bob initially share an EPR pair, and Alice measures her half in the standard basis if $a = 0$ and in the Hadamard basis in $a = 1$. This suggests the following proof of quantumness protocol.

- The verifier samples $a \leftarrow \{0,1\}$ and performs an OSP with the prover in order to deliver $H^a |x\rangle$.
- The parties "complete" the CHSH game as follows. The verifier samples $b \leftarrow \{0,1\}$ and sends it to the prover. If $b = 0$, the prover measures their state in the $X + Z$ basis to obtain y, and if $b = 1$, the prover measures their state in the $X - Z$ basis to obtain y, and sends y back to the verifier.[7]
- The verifier runs the CHSH verification predicate, accepting if $x \oplus y = a \cdot b$.

By a standard analysis of the CHSH game, an honest QPT prover following the strategy outlined above makes the verifier accept with probability $\cos^2(\pi/8) > 0.85$. Now, consider any classical polynomial-time prover. Note that an equivalent way to write the verification predicate is to accept if $y = x \oplus a \cdot b$. Given any classical prover that wins with probability $3/4 + 1/\text{poly}$, we can rewind them to extract answers on both $b = 0$ and $b = 1$ that simultaneously accept with probability at least $1/2 + 1/\text{poly}$. That is, we can obtain $y_0 = x$ and $y_1 = x \oplus a$ with probability $1/2 + 1/\text{poly}$. However, this contradicts the security of the OSP, since $y_0 \oplus y_1 = a$, and OSP demands that no polynomial-time adversary has noticeable advantage in guessing the basis choice a.

2.2.2 1-of-2 puzzles

Next, we extend the above ideas to realize more applications, via the intermediate primitive of 1-of-2 puzzles. Introduced by [RS19], a 1-of-2 puzzle consists of four algorithms defined as follows.

- $\mathsf{KeyGen}(1^\lambda) \rightarrow (\mathsf{pk}, \mathsf{vk})$. The PPT key generation algorithm takes as input the security parameter 1^λ and outputs a public key pk and a secret verification key vk.
- $\mathsf{Obligate}(\mathsf{pk}) \rightarrow (|\psi\rangle, y)$: The QPT obligate algorithm takes as input the public key, and outputs a classical obligation string y and a quantum state $|\psi\rangle$.
- $\mathsf{Solve}(|\psi\rangle, b) \rightarrow a$: The QPT solve algorithm takes as input a state $|\psi\rangle$ and a bit $b \in \{0,1\}$ and outputs a string a.
- $\mathsf{Ver}(\mathsf{vk}, y, b, a) \rightarrow \{\top, \bot\}$: The PPT verify algorithm takes as input the verification key vk, a string y, a bit $b \in \{0,1\}$, and a string a, and either accepts or rejects.

Correctness requires that on either challenge $b \in \{0,1\}$, the Solve algorithm produces an accepting answer a, while security stipulates that no QPT adversary can *simultaneously* produce an accepting answer a_0 on challenge $b = 0$ and an accepting answer a_1 on challenge $b = 1$. This is clearly an inherently quantum primitive that has been shown by prior work to imply both (privately-verifiable) quantum money with classical communication [RS19], and position verification with classical communication [LLQ22].

[7] To be concrete, these bases are defined as $X + Z = \{\cos(\pi/8)|0\rangle + \sin(\pi/8)|1\rangle, -\sin(\pi/8)|0\rangle + \cos(\pi/8)|1\rangle\}$, and $X - Z = \{\cos(-\pi/8)|0\rangle + \sin(-\pi/8)|1\rangle, -\sin(-\pi/8)|0\rangle + \cos(-\pi/8)|1\rangle\}$.

Our idea to obtain a 1-of-2 puzzle is as follows. The KeyGen and Obligate algorithms will run λ parallel instances of two-round OSP on a uniformly random sender's bit $r \leftarrow \{0,1\}$.[8] This results in a state $(H^r)^{\otimes\lambda}|s\rangle$, where s is a λ-bit string. Now, challenge $b=0$ asks for a string that matches s on at least 0.85 fraction of indices, while challenge $b=1$ asks for a string that matches $s \oplus (r,\ldots,r)$ on at least 0.85 fraction of indices. Roughly, this is a λ-parallel repetition of the above proof of quantumness, all using the same verifier bit r. Thus, correctness follows from measuring all states in the $X+Z$ basis when $b=0$, and in the $X-Z$ basis when $b=1$ (and a tail bound). We can also establish a weak form of security. Suppose a (quantum) adversary has a $1/2+1/\mathrm{poly}$ probability of passing both challenges simultaneously. Then, by XORing its answers and taking the majority bit, we have that with $1/2+1/\mathrm{poly}$, this adversary can be used to predict the bit r, a contradiction to the security of the two-round OSP. Finally, to obtain a full-fledged 1-of-2 puzzle (i.e. with negligible security), we appeal to an amplification lemma of [RS19] that is itself based on the parallel repetition for weakly verifiable puzzles of [CHS05].

2.2.3 Blind delegation

Next, we show that OSP is sufficient to obtain blind classical delegation of any quantum computation. While there are likely several routes to showing this, our approach essentially instantiates the protocol of [Bro15] using OSP in place of quantum communication from the client to the server.

In full generality, our definition of blind delegation allows the classical client to delegate the computation of some publicly-known quantum operation Q that takes as input a private classical string x from the client and a quantum state on register $\mathcal{V}$ from the server. At the end of the protocol, the prover recovers the output $Q(x,\mathcal{V})$ up to a quantum one-time pad X^rZ^s with keys (r,s) known to the client. Note that this implies the ability to deliver to the client a classical output, by having the prover measure the output register and deliver the result to the client, which will be correct up to a *classical* one-time pad defined by r.

We show that OSP implies this notion as follows.[9] First, we write the circuit Q as a sequence of alternating Clifford operations and $T^\dagger$ gates $Q = C_{\ell+1}T^\dagger C_\ell \ldots C_2T^\dagger C_1$, where $T^\dagger$ is the $\pi/4$ rotation clockwise around the XY plane. To begin the protocol, the client sends $x \oplus r_{\mathsf{inp}}$, where x is their input and r_{inp} is a classical one-time pad. As is typically the case, Clifford operations are straightforward: the server can apply them directly to the current state of the system, and the client can perform a corresponding update to their one-time pad keys. On the other hand, each time we come to a $T^\dagger$ gate, we will use one instance of OSP (and thus some interaction).

The idea is that $T^\dagger X^rZ^s|\psi\rangle = (P^\dagger)^rX^rZ^sT^\dagger|\psi\rangle$, and so applying the $T^\dagger$ reduces to what we call an *encrypted phase gate*. That is, the client holds a private bit r, the server holds a single-qubit quantum state $|\psi\rangle$,[10] and we want the server to obtain $P^r|\psi\rangle$ after interaction, potentially up to some Pauli error (concretely, our protocol will result in $Z^mP^r|\psi\rangle$, where m is known to the client).

To enable this, the client first uses OSP to transmit a fresh state $Z^sP^r|+\rangle$ to the server, where s is known to the client. That is, the client and server engage in an OSP for the X and Y observables, which follows by using "standard" OSP for the Z and X observables and then having the server

[8] We note that it is possible to define a more interactive version of 1-of-2 puzzles and write down a candidate from OSP rather than two-round OSP. However, for security we rely on an amplification lemma from [RS19] that crucially uses the two-round setup, so we leave an exploration of this generalization to future work.

[9] It is also easy to see that this definition implies OSP as well, meaning OSP is both necessary and sufficient.

[10] In general, this could be entangled with the rest of their system, but we suppress this in the overview to avoid clutter.

apply the appropriate (public) rotation. Now, a straightforward calculation confirms that if the server applies CNOT from their state $|\psi\rangle$ onto $Z^s P^r |+\rangle$ and then measures the second register in the standard basis, the remaining state on the first register will be exactly $P^r |\psi\rangle$ up to some Pauli Z error.

Security of the delegation protocol is immediate from security of the OSP. One by one, we can switch the client's input to each of the OSP instances to 0. Once this is done, the only relevant information the server obtains from running the protocol is $x \oplus r_{\mathsf{inp}}$, which perfectly hides the client's private input x.

2.2.4 Verifiable delegation

A protocol for classical verification of arbitrary BQP computation (CVQC) was first shown by [Mah18b], albeit from a specific type of TCF known only from LWE. Recent work [KLVY23, NZ23, MNZ24] has explored a different approach that has given us CVQC from the more generic assumption of QFHE. However, QFHE is itself a strong primitive, and is only known from lattices. In this work, we further generalize the approach, showing that it can be instantiated with any blind classical delegation of quantum computation protocol in place of the QFHE, and thus from any OSP.

The starting point for this approach is the "KLVY compiler", which uses QFHE to compile any classical-verifier game sound against two non-communicating (but potentially entangled) provers into a single prover game, with the hope that the semantic security of the QFHE will give soundness against any QPT prover. Given any two-prover game, the compiled protocol is defined as follows.

- The verifier sends a QFHE encryption $\mathsf{Enc}(x)$ of Alice's question to the prover.
- The prover runs Alice's strategy under the QFHE to obtain an encryption of her answer $\mathsf{Enc}(a)$, along with some auxiliary quantum state. The prover sends $\mathsf{Enc}(a)$ to the verifier.
- The verifier sends Bob's question y in the clear.
- The prover uses its auxiliary state and y to obtain Bob's answer b, which it sends to the verifier.
- The verifier decrypts $\mathsf{Enc}(a)$ to obtain a and applies its verification predicate to (a, b).

Intuitively, QFHE is used to enforce the non-communicating assumption *computationally*. That is, the semantic security of QFHE implies that the first prover operation (Alice) cannot transmit any information about her question a to the second prover operation (Bob) that can be efficiently recovered. While we don't have a general theorem establishing optimal soundness of the KLVY compiler[11] for *any* game (see [BGKM$^+$23, CMM$^+$24, KMP$^+$24] for progress in this direction), [NZ23] presented a two-player game for arbitrary BQP computation and showed its soundness under the KLVY compiler, giving CVQC from QFHE as a corollary.

In this work, we begin by defining what we call the *generalized* KLVY compiler, which replaces the first round above with any (potentially interactive, non-compact) blind classical delegation of quantum computation protocol. Then, in order to simplify the task of proving soundness of the

[11] Here, we mean soundness against *quantum* polynomial-time provers. [KLVY23] showed a general theorem establishing soundness against *classical* polynomial-time provers.

generalized KLVY compiler, we define a clean class of strategies for two-prover games, which we call *computationally non-local strategies* (Definition 6.19). While the standard notion of a non-local strategy requires that Alice's operation $\{A^x\}_x$ and Bob's operation $\{B^y\}_y$ (where both are written as sets of strategies parameterized by their question) must be applied to disjoint Hilbert spaces, say $\mathcal{H}_\mathcal{A}$ and $\mathcal{H}_\mathcal{B}$, our notion relaxes this requirement as follows. It includes any strategy $\{A^x\}_x$ on $\mathcal{H}_\mathcal{A} \otimes \mathcal{H}_\mathcal{B}$ followed by $\{B^y\}_y$ on $\mathcal{H}_\mathcal{B}$ such that no QPT distinguisher given the state on register $\mathcal{B}$ output by A^x can guess x with noticeable advantage.

We next prove a theorem (Theorem 6.23) showing that any upper bound on the value of a two-prover game against computationally non-local strategies is also an upper bound on the soundness of the generalized KLVY-compiled game. While straightforward to show, this theorem is quite useful. It allows one to forget all of the underlying details of the cryptographic component when attempting to prove the soundness of a (generalized) KLVY-compiled protocol. Indeed, prior work (e.g. [KLVY23, NZ23, MNZ24]) carried around clunky notation specific to QFHE including public / secret key pairs, ciphertexts, etc., when upper-bounding the soundness of their compiled non-local games.

Here, we re-visit [NZ23]'s proof strategy, showing that it in fact establishes an upper bound on *any* computationally non-local strategy (i.e. not only ones that arise from the use of QFHE). In particular, there are only a handful of places where QFHE is used in their proof, and each time it is only used to show that Bob cannot distinguish between two different Alice questions with better than negligble advantage. Thus, we conclude that CVQC follows generically from any blind classical delegation of quantum computation protocol, and thus from any OSP.

2.2.5 Encrypted CNOT and applications

Next, we re-visit a notion that was informally introduced in the influential work of [Mah18a], called "encrypted CNOT". In [Mah18a], encrypted CNOT was built assuming TCFs with a particular structural requirement on the claws, and it was incorporated into their construction of quantum FHE (in a non-black-box way). Here, we define encrypted CNOT formally as a special case of blind classical delegation of quantum computation, which in particular implies that it follows from any OSP. However, we go a step further, and provide a simple and direct construction of encrypted CNOT, which in particular shows that if we start with a *two-round* OSP, then we obtain a two-round encrypted CNOT.

To be precise, we define encrypted CNOT as a protocol that takes a private bit b from the client and a two-qubit state from the server, and, if $b = 0$, does nothing to the server's state, while if $b = 1$, applies a CNOT to the server's state. The output state will only be correct up to a quantum one-time pad that must be known to the client. To implement this from OSP, we operate in two steps. Suppose for simplicity that the server initially holds a two qubit state of the form

$$(\alpha_0 |0\rangle + \alpha_1 |1\rangle) \otimes (\beta_0 |0\rangle + \beta_1 |1\rangle).$$

1. First, depending on the bit b, either entangle the first qubit with a fresh register, or not. This can be accomplished using OSP as follows. Execute two instances of OSP, where if $b = 0$, the client inputs are $(0, 1)$ while if $b = 1$, the client inputs are $(1, 0)$. Thus, up to a one-time pad, the server's state can now be written as

$$\begin{aligned} &(\alpha_0 |0\rangle + \alpha_1 |1\rangle) \otimes |0\rangle \otimes |+\rangle \otimes (\beta_0 |0\rangle + \beta_1 |1\rangle) \quad \text{if } b = 0 \\ &(\alpha_0 |0\rangle + \alpha_1 |1\rangle) \otimes |+\rangle \otimes |0\rangle \otimes (\beta_0 |0\rangle + \beta_1 |1\rangle) \quad \text{if } b = 1. \end{aligned}$$

The server then applies a CNOT from the 1st to the 3rd qubit and a CNOT from the 2nd to the 3rd qubit. In the first case, where the 3rd qubit is $|+\rangle$, this has no effect, while in the second case, this entangles the 1st and 2nd qubit. Then, after measuring the 3rd qubit in the standard basis, the server's state becomes (up to a one-time pad)

$$\begin{aligned}&(\alpha_0 |0\rangle + \alpha_1 |1\rangle) \otimes |0\rangle \otimes (\beta_0 |0\rangle + \beta_1 |1\rangle) \quad \text{if } b = 0\\ &(\alpha_0 |00\rangle + \alpha_1 |11\rangle) \otimes (\beta_0 |0\rangle + \beta_1 |1\rangle) \quad \text{if } b = 1.\end{aligned}$$

2. Next, apply CNOT from the 2nd qubit to the 3rd qubit, then "delete" the 2nd qubit by measuring it in the Hadamard basis. Clearly, in the $b = 0$ case this again has no effect, while in the $b = 1$ case, this accomplishes a CNOT between the two input qubits. In particular, it can be confirmed that after this step, the server's state becomes (up to a one-time pad)

$$\begin{aligned}&(\alpha_0 |0\rangle + \alpha_1 |1\rangle) \otimes (\beta_0 |0\rangle + \beta_1 |1\rangle) \quad \text{if } b = 0\\ &(\alpha_0\beta_0 |00\rangle + \alpha_0\beta_1 |1\rangle + \alpha_1\beta_0 |11\rangle + \alpha_1\beta_1 |10\rangle) \quad \text{if } b = 1.\end{aligned}$$

For details (in particular, how the client recovers the one-time pad keys from the OSP information and the server's measurement results), refer to Section 6.5. Here, we mention the applications we obtain from our encrypted CNOT protocol.

Quantum fully-homomorphic encryption. The first quantum fully-homomorphic encryption (QFHE) scheme was constructed by [Mah18a]. As alluded to above, the construction combines a particular encrypted CNOT protocol with a particular classical FHE protocol in a non-black-box manner in order to achieve QFHE. A recent work of [GV24] pioneered a *generic* approach to QFHE from *any* classical FHE (with log-depth decryption) and *any* dual-mode TCF. In Section 6.5, we observe that their work can in fact be seen as constructing QFHE from any (two-round) *encrypted CNOT* protocol plus classical FHE (with log-depth decryption), and thus we establish that QFHE follows from two-round OSP plus classical FHE (with log-depth decryption).

Claw-state generators with indistinguishability security. Recall the notion of a claw-state generator (CSG) introduced above, which is a protocol that delivers a state $\frac{1}{\sqrt{2}}(|0, x_0\rangle + (-1)^z |1, x_1\rangle)$ to a quantum receiver and strings (x_0, x_1, z) to a classical sender.[12] We say that such a protocol has *indistinguishability* security if for all $i \in [n]$, no QPT server can predict the bit $x_{0,i} \oplus x_{1,i}$ with better than $\mathrm{negl}(\lambda)$ advantage.

In Section 6.5, we show that (two-round) encrypted CNOT can be used to obtain a (two-round) CSG with indistinguishability security. The construction is straightforward: the sender begins by sampling a string $\Delta \leftarrow \{0,1\}^n$, and the receiver initializes the state $|+\rangle_{\mathcal{B}} \otimes |0\rangle_{\mathcal{C}_1} \otimes \cdots \otimes |0\rangle_{\mathcal{C}_n}$. Then, the parties engage in n encrypted CNOTs, where the i'th protocol takes input Δ_i from the sender, and applies a CNOT from register $\mathcal{B}$ to register $\mathcal{C}_i$. It can be confirmed that the receiver ends up with a state of the form

$$\frac{1}{\sqrt{2}} (|0, x\rangle + (-1)^z |0, x + \Delta\rangle),$$

[12] In full generality, we allow the claw-state to have a phase specified by the bit z, as long as this bit is known to the sender.

where $x \in \{0,1\}^n$ and $z \in \{0,1\}$ can be recovered by the sender. Crucially, note that XOR of the two members of the claw is equal to Δ, and thus, breaking indistinguishability security of the CSG yields an attack on the encrypted CNOT protocol.

Combined with our construction of OSP from Section 2.1, this shows that OSP and CSG (with either search or indistinguishability security) are *equivalent*. Moreover, the notion of a CSG with indistinguishability security will be central to our implications in the next section establishing cryptographic lower bounds for constructing OSP.

2.3 Implications

Our next goal is to understand what cryptographic hardness is necessary for OSP. Towards addressing this, we show that OSP implies commitments and oblivious transfer (where one participant only requires classical capabilities), and that *two-round* OSP implies public-key encryption.

2.3.1 Commitments

Our commitment scheme proceeds as follows: the classical committer acts as the sender in λ executions of an OSP with chosen input basis b in all executions. The committer thus obtains bits $s_1, \ldots, s_\lambda$, while an (honest) receiver ends up with $H^b \left|s_1\right\rangle, \ldots, H^b \left|s_\lambda\right\rangle$.

In the decommit phase, the committer reveals $s_1, \ldots, s_\lambda$ along with its comitted bit b, and the receiver accepts iff for every $i \in [\lambda]$, the projection of its i^{th} qubit onto $H^b \left|s_i\right\rangle\!\left\langle s_i\right| H^b$ accepts.

The hiding of this commitment against a malicious receiver follows from the fact that OSP hides the sender input b from an arbitrary (malicious) receiver, together with a straightforward hybrid argument.

To see why this satisfies statistical (sum) binding, consider the state $\left|\psi\right\rangle$ that an honest receiver ends up with after interacting with an arbitrary malicious committer. For any fixing of this state $\left|\psi\right\rangle$, the probability that a decommitment to $(0, s_0)$ is accepted is $\mathsf{pr}_{0,s_0} = \|\left\langle s_0 | \psi\right\rangle\|^2$, and a decommitment to $(1, s_1)$ is accepted is $\mathsf{pr}_{1,s_1} = \|\left\langle s_1\right| H^{\otimes\lambda} \left|\psi\right\rangle\|^2$. Let pr_0 denote $\max_{s_0}(\mathsf{pr}_{0,s_0})$ and pr_1 denote $\max_{s_1}(\mathsf{pr}_{1,s_1})$. Since for every s_0, s_1,

$$\|\left\langle s_0\right| H^{\otimes\lambda} \left|s_1\right\rangle\|^2 = \frac{1}{2^\lambda},$$

we conclude that $\mathsf{pr}_0 + \mathsf{pr}_1 \leq 1 + \mathrm{negl}(n)$, as desired.

2.3.2 Oblivious transfer

We obtain oblivious transfer (OT) by building on the notion of a claw-state generator (CSG) with indistinguishability security, introduced above.[13] The basic idea is as follows. The OT receiver will delegate the preparation of a state

$$\frac{1}{\sqrt{2}}\left(\left|0, x_0\right\rangle + (-1)^z \left|1, x_1\right\rangle\right)$$

[13] Our basic construction achieves a somewhat non-standard definition that we call search security (against a malicious receiver). We also show that, by additionally assuming one-way functions, we can achieve a more standard indistinguishability-based definition. We refer the reader to Section 7.2 for details.

to the OT sender, where x_0 and x_1 are single bits. We will take $b = x_0 \oplus x_1$ to be the receiver's choice bit, which, by the indistinguishability security of the CSG, is computationally unpredictable to any QPT sender. Then, the sender measures their state in the standard basis to obtain two bits (c, y), and defines their OT bits to be $r_0 = y, r_1 = y \oplus c$.

Note that if $b = 0$, then $y = x_0 = x_1$ is known to the receiver, while the bit c is uniformly random, meaning r_1 is unpredictable. On the other hand, if $b = 1$, then $x_0 = 1 \oplus x_1$, so the bit $r_1 = c \oplus y$ is known to the receiver, while the bit y is uniformly random, meaning r_0 is unpredictable. However, this analysis relies on the fact that the state measured by the sender is indeed the desired state $\frac{1}{\sqrt{2}}(|0, x_0\rangle + (-1)^z |1, x_1\rangle)$. Unfortunately, the notion of CSG (and also the underlying notion of OSP) does not guarantee any *verifiability* property, meaning that we have no guarantees on what the OT sender's state might look like if the OT receiver is acting maliciously in the CSG protocol. To remedy this, we repeat the CSG protocol several times and use a cut-and-choose protocol to allow the sender to check that the receiver is behaving (close to) honestly. In the end, after combining the several protocols, we arrive at a (game-based) notion of oblivious transfer between a classical (unbounded) receiver and a quantum (polynomial-time) sender. For full details, please refer to Section 7.2.

2.3.3 Public-key encryption

Finally, we show that two-round OSP implies (CPA-secure) public-key encryption. Our protocol uses the same building block as the OT protocol from above: a (two-round) CSG for generating the state

$$\frac{1}{\sqrt{2}}(|0, x_0\rangle + (-1)^z |1, x_1\rangle)$$

with indistinguishability security.

The public key in our scheme is the first-round (classical) message msg_1 of the CSG protocol, while the secret key is the secret state of the classical sender who generated this message. To encrypt a bit m, use msg_1 to generate the above state along with a second-round message msg_2, measure the state in the standard basis to obtain (b, x_b), and output $(\mathsf{msg}_2, b, m \oplus x_b)$ as the ciphertext.

Given the sender's state, the key generator can recover the values of x_0, x_1 from msg_2 and thus decrypt the message. However, breaking the CPA security of this scheme reduces to being able to predict x_b given msg_2 for a random bit b. In turn, this yields an adversary attacking the indistinguishability security of the CSG: Given msg_1, it prepares a state $\frac{1}{\sqrt{2}}(|0, x_0\rangle + (-1)^z |1, x_1\rangle)$ honestly along with msg_2, samples a random b, and runs the PKE adversary to obtain a guess for x_b. Then, it measures its state to obtain $(b', x_{b'})$. If $b \neq b'$ (which occurs with probability 1/2), this adversary then knows $x_0 \oplus x_1$, breaking the indistinguishability security of the CSG. Again, we refer to the body, in particular Section 7.3, for the full details.

3 Preliminaries

Let λ denote the security parameter. We write $\operatorname{negl}(\cdot)$ to denote any negligible function, which is a function f such that for every constant $c \in \mathbb{N}$ there exists $N \in \mathbb{N}$ such that for all $n > N$, $f(n) < n^{-c}$. We write $\mathsf{non\text{-}negl}(\cdot)$ to denote any function f that is not negligible, that is, there exists a constant c such that for infinitely many n, $f(n) \geq n^{-c}$. Finally, we write $\mathrm{poly}(\cdot)$ to denote any polynomial function f, that is, there exist constants c and N such that for all $n > N$, $f(n) < n^c$.

A probabilistic polynomial-time (PPT) family of circuits $\{C_\lambda\}_{\lambda \in \mathbb{N}}$ is a family of randomized classical circuits with $|C_\lambda| \leq \mathrm{poly}(\lambda)$, and a quantum polynomial-time (QPT) family of circuits $\{Q_\lambda\}_{\lambda \in \mathbb{N}}$ is a family of quantum circuits with $|Q_\lambda| \leq \mathrm{poly}(\lambda)$.

Let Tr denote the trace operator. The *trace distance* between two quantum (mixed) states ρ_0, ρ_1, denoted $\mathsf{TD}(\rho_0, \rho_1)$ is defined as

$$\frac{1}{2}\|\rho_0 - \rho_1\|_1,$$

where $\|\cdot\|_1$ is the *trace norm*, defined by

$$\|\rho\|_1 \coloneqq \mathrm{Tr}\sqrt{\rho^\dagger \rho}.$$

The trace distance between two states ρ_0 and ρ_1 is an upper bound on the probability that any (unbounded) algorithm can distinguish ρ_0 and ρ_1.

Given two quantum operations Q_0, Q_1 that take as input a state on register $\mathcal{A}$, their diamond distance is defined as

$$D_\diamond(Q_0, Q_1) \coloneqq \sup_{\mathcal{B}} \max_{\rho_{\mathcal{A},\mathcal{B}}} \|(Q_0 \otimes \mathcal{I}_{\mathcal{B}})\rho_{\mathcal{A},\mathcal{B}} - (Q_1 \otimes \mathcal{I}_{\mathcal{B}})\rho_{\mathcal{A},\mathcal{B}}\|_1,$$

where $\mathcal{I}_{\mathcal{B}}$ is the identity matrix on register $\mathcal{B}$. In words, the diamond distance upper bounds the trace distance between the outputs of Q_0 and Q_1 on any input (which could be entangled with an arbitrary auxiliary register $\mathcal{B}$).

We will use the usual convention that Z refers to the basis $\{|0\rangle, |1\rangle\}$, X refers to the basis $\{|+\rangle, |-\rangle\}$, and Y refers to the basis $\{P|+\rangle, P|-\rangle\}$, where P is the phase gate. We will often refer to the $X+Z$ and $X-Z$ bases, defined as

$$X + Z = \{\cos(\pi/8)|0\rangle + \sin(\pi/8)|1\rangle, -\sin(\pi/8)|0\rangle + \cos(\pi/8)|1\rangle\},$$

$$X - Z = \{\cos(-\pi/8)|0\rangle + \sin(-\pi/8)|1\rangle, -\sin(-\pi/8)|0\rangle + \cos(-\pi/8)|1\rangle\}.$$

Lemma 3.1 (Gentle measurement [Win99]). *Let ρ be a quantum state and let $(\Pi, \mathcal{I} - \Pi)$ be a projective measurement such that $\mathrm{Tr}(\Pi\rho) \geq 1 - \delta$. Let*

$$\rho' = \frac{\Pi\rho\Pi}{\mathrm{Tr}(\Pi\rho)}$$

be the state after applying $(\Pi, \mathcal{I} - \Pi)$ to ρ and post-selecting on obtaining the first outcome. Then, $\mathsf{TD}(\rho, \rho') \leq 2\sqrt{\delta}$.

Lemma 3.2 (Quantum Goldreich-Levin [AC02]). *Suppose there exists a string $s \in \{0,1\}^n$, a state $|\psi\rangle$ on m qubits, a unitary U on $n + m + t$ qubits that is classically controlled on its first n qubits, and an $\epsilon \in (0,1)$ such that for uniformly random $r \leftarrow \{0,1\}^n$, measuring the last qubit of $U|x\rangle|\psi\rangle|0^t\rangle$ yields $r \cdot s$ with probability at least $1/2 + \epsilon$. Then given $|\psi\rangle$, there exists a quantum algorithm that outputs s with probability at least $4\epsilon^2$ using a single invocation of U and $U^\dagger$.*

4 Oblivious State Preparation

In this section, we first define a "standard" notion of oblivious state preparation (OSP), and then investigate variants of the definition. The standard notion we propose enables a quantum server communicating with a classical client to prepare a single-qubit state in either the standard or Hadamard basis, without actually learning the basis. This corresponds exactly to the functionality of "Malicious 4-states QFactory with basis-blindness" proposed by [CCKW19].

However, the concept of OSP is not fundamentally tied to the standard and Hadamard bases. Conceptually, it captures the ability for a client to enable the preparation of a state in one of two arbitrary bases on the server's system. Thus, later in the section we define a generalized notion of OSP, which enables the angle between the bases to be arbitrary, and we initiate the study of this generalized notion.

4.1 Basic definitions

Definition 4.1 (Oblivious State Preparation). *Oblivious state preparation (OSP) is a protocol that takes place between a PPT sender S with input $b \in \{0,1\}$ and a QPT receiver R:*

$$(s,\, |\psi\rangle) \leftarrow \langle S(1^\lambda, b), R(1^\lambda)\rangle,$$

where $s \in \{0,1\}$ is the sender's output and $|\psi\rangle$ is the receiver's output. It should satisfy the following properties.

- ***Correctness.*** *For any $b \in \{0,1\}$, let*

  $$\Pi_{\mathsf{OSP},b} \coloneqq \sum_{s\in\{0,1\}} |s\rangle\langle s| \otimes H^b\, |s\rangle\langle s|\, H^b.$$

  Then for any $b \in \{0,1\}$,

  $$\mathbb{E}\left[\|\Pi_{\mathsf{OSP},b}\,|s\rangle\ |\psi\rangle\,\| : (s,\, |\psi\rangle) \leftarrow \langle S(1^\lambda, b), R(1^\lambda)\rangle\right] = 1 - \mathrm{negl}(\lambda).$$

  We say that the protocol satisfies perfect *correctness if the expectation above is equal to 1.*

- ***Security.*** *For any QPT adversary $\{\mathsf{Adv}_\lambda\}_{\lambda\in\mathbb{N}}$,*

  $$\Big|\Pr\left[b_{\mathsf{Adv}} = 0 : (s, b_{\mathsf{Adv}}) \leftarrow \langle S(1^\lambda, 0), \mathsf{Adv}_\lambda\rangle\right] - \Pr\left[b_{\mathsf{Adv}} = 0 : (s, b_{\mathsf{Adv}}) \leftarrow \langle S(1^\lambda, 1), \mathsf{Adv}_\lambda\rangle\right]\Big| = \mathrm{negl}(\lambda).$$

We say that the protocol is a two-round OSP *if it consists of just two messages: one from the sender followed by one from the receiver. In this case, we use the following notation to describe the algorithms of the protocol.*

- $\mathsf{OSP.Sen}(1^\lambda, b) \to (\mathsf{msg}_S, \mathsf{st}_S)$. *The PPT sender takes as input the security parameter 1^λ and a bit b, and outputs a message msg_S and state st_S.*

- $\mathsf{OSP.Rec}(\mathsf{msg}_S) \to (\,|\psi\rangle\,, \mathsf{msg}_R)$. *The QPT receiver takes as input the sender's message msg_S and outputs its final state $|\psi\rangle$ and a message msg_R.*

- $\mathsf{OSP.Dec}(\mathsf{st}_S, \mathsf{msg}_R) \to s$. *The PPT sender takes as input its state* st_S *and the receiver's message* msg_R*, and produces its output bit* s*.*

Sometimes, we will refer to the above definition as a *chosen-input* OSP, in order to distinguish it from a random-input variant defined below, where the sender does not fix a choice of b at the beginning of the protocol.

Definition 4.2 (Random-Input Oblivious State Preparation). *Random-input OSP is a protocol that takes place between a PPT sender S and a QPT receiver R:*

$$((s,b),\ |\psi\rangle) \leftarrow \langle S(1^\lambda), R(1^\lambda)\rangle,$$

where (s,b) is the sender's output and $|\psi\rangle$ is the receiver's output. It should satisfy the following properties.

- ***Correctness.*** *Let*
$$\Pi_{\mathsf{OSP}} \coloneqq \sum_{s,b\in\{0,1\}} |s,b\rangle\langle s,b| \otimes H^b\, |s\rangle\langle s|\, H^b.$$
Then
$$\mathbb{E}\left[\|\Pi_{\mathsf{OSP}}\, |s,b\rangle\ |\psi\rangle\ \| : ((s,b),\ |\psi\rangle) \leftarrow \langle S(1^\lambda), R(1^\lambda)\rangle\right] = 1 - \mathrm{negl}(\lambda).$$

- ***Security.*** *For any QPT adversary $\{\mathsf{Adv}_\lambda\}_{\lambda\in\mathbb{N}}$,*
$$\left|\Pr\left[b_{\mathsf{Adv}} = b : ((s,b), b_{\mathsf{Adv}}) \leftarrow \langle S(1^\lambda), \mathsf{Adv}_\lambda\rangle\right] - \frac{1}{2}\right| = \mathrm{negl}(\lambda).$$

Lemma 4.3. *Random-input OSP implies (chosen-input) OSP.*

Proof. To obtain chosen-input OSP with sender's choice bit b, the parties begin by running a random-input OSP, which (up to negligible trace distance) delivers output $H^{b'}|s\rangle$ to the receiver and (s,b') to the sender. The sender then sends the bit $c = b \oplus b'$ to the receiver, and the receiver applies H^c to its state to obtain $H^b|s\rangle$. Security follows from the security of the random-input OSP, which guarantees that the bit b' is unpredictable and thus that the bit $c = b \oplus b'$ masks the sender's choice of b. □

4.2 Claw-state generators

Next, we define (variants of) a claw-state generation (CSG) protocol, and show that CSG implies OSP. Later, in section Section 6.5, we will show that in fact OSP implies CSG, meaning that these notions are equivalent.

Definition 4.4 (Claw-State Generator). *A claw-state generator (CSG) is a protocol that takes places between a PPT sender S and a QPT receiver R:*

$$((x_0, x_1, z), |\psi\rangle) \leftarrow \langle S(1^\lambda, n), R(1^\lambda, n)\rangle,$$

where the sender's output consists of $x_0, x_1 \in \{0,1\}^n$ and $z \in \{0,1\}$, and $|\psi\rangle$ is the receiver's output. It should satisfy the following notion of correctness, and, depending on the setting, it should also satisfy either search security or indistinguishability security.

- ***Correctness.*** *Let*

$$\Pi_{\mathsf{CSG}} \coloneqq \sum_{x_0 \neq x_1 \in \{0,1\}^n, z \in \{0,1\}} |x_0, x_1, z\rangle\langle x_0, x_1, z| \otimes \frac{1}{2} \left(|x_0\rangle + (-1)^z |x_1\rangle\right)\left(\langle x_0| + (-1)^z \langle x_1|\right).$$

  Then

$$\mathbb{E}\left[\|\Pi_{\mathsf{CSG}} |x_0, x_1, z\rangle |\psi\rangle \| : ((x_0, x_1, z), |\psi\rangle) \leftarrow \langle S(1^\lambda, n), R(1^\lambda, n)\rangle\right] = 1 - \mathrm{negl}(\lambda).$$

  We say that the protocol has perfect *correctness if the above probability is equal to 1.*

- ***Search security.*** *For any QPT adversary $\{\mathsf{Adv}_\lambda\}_{\lambda \in \mathbb{N}}$,*

$$\Pr\left[x_{\mathsf{Adv}} = (x_0, x_1) : ((x_0, x_1, z), x_{\mathsf{Adv}}) \leftarrow \langle S(1^\lambda, n), \mathsf{Adv}_\lambda\rangle\right] = \mathrm{negl}(\lambda).$$

- ***Indistinguishability security.*** *For any QPT adversary $\{\mathsf{Adv}_\lambda\}_{\lambda \in \mathbb{N}}$ and any $i \in [n]$,*

$$\left|\Pr\left[b_{\mathsf{Adv}} = x_{0,i} \oplus x_{1,i} : ((x_0, x_1, z), b_{\mathsf{Adv}}) \leftarrow \langle S(1^\lambda, n), \mathsf{Adv}_\lambda\rangle\right] - \frac{1}{2}\right| = \mathrm{negl}(\lambda).$$

We say that the protocol is a two-round CSG *if it consists of just two messages: one from the sender followed by one from the receiver. In this case, we use the following notation to describe the algorithms of the protocol.*

- $\mathsf{CSG.Sen}(1^\lambda, n) \rightarrow (\mathsf{msg}_S, \mathsf{st}_S)$. *The PPT sender takes as input the security parameter 1^λ and outputs a message msg_S and state st_S.*
- $\mathsf{CSG.Rec}(\mathsf{msg}_S) \rightarrow (|\psi\rangle, \mathsf{msg}_R)$. *The QPT receiver takes as input the sender's message msg_S and outputs its final state $|\psi\rangle$ and a message msg_R.*
- $\mathsf{CSG.Dec}(\mathsf{st}_S, \mathsf{msg}_R) \rightarrow (x_0, x_1, z)$. *The PPT sender takes as input its state st_S and the receiver's message msg_R, and produces its output (x_0, x_1, z).*

We also define a version of a claw-state generator where the honest receiver obtains $\frac{1}{\sqrt{2}}(|0, x_0\rangle + (-1)^z |1, x_1\rangle)$, that is, where the two members of the claw are differentiated by the first bit.

Definition 4.5 (Differentiated-Bit Claw-State Generator)**.** *A differentiated-bit claw-state generator is defined exactly like a claw-state generator except that the correctness property is stated as follows. Let*

$$\Pi_{\mathsf{DBCSG}} \coloneqq \sum_{x_0 \neq x_1 \in \{0,1\}^n, z \in \{0,1\}} |x_0, x_1, z\rangle\langle x_0, x_1, z| \otimes \frac{1}{2} \left(|0, x_0\rangle + (-1)^z |1, x_1\rangle\right)\left(\langle 0, x_0| + (-1)^z \langle 1, x_1|\right).$$

Then

$$\mathbb{E}\left[\|\Pi_{\mathsf{DBCSG}} |x_0, x_1, z\rangle |\psi\rangle \| : ((x_0, x_1, z), |\psi\rangle) \leftarrow \langle S(1^\lambda, n), R(1^\lambda, n)\rangle\right] = 1 - \mathrm{negl}(\lambda).$$

It is straightforward to obtain a differentiated-bit CSG with search security from a (plain) CSG with search security.

Lemma 4.6. *CSG with search security (Definition 4.4) implies differentiated-bit CSG with search security (Definition 4.5).*

Proof. The protocol for differentiated-bit CSG goes as follows. First, run a plain CSG. Then, the sender samples a uniformly random y conditioned on $y \cdot x_0 = 0$ and $y \cdot x_1 = 1$, and sends y to the receiver. Finally, the receiver applies the map

$$\frac{1}{\sqrt{2}}(|x_0\rangle + (-1)^z |x_1\rangle) \to \frac{1}{\sqrt{2}}(|y \cdot x_0, x_0\rangle + (-1)^z |y \cdot x_1, x_1\rangle = \frac{1}{\sqrt{2}}(|0, x_0\rangle + (-1)^z |1, x_1\rangle).$$

Correctness is immediate, and security follows by reduction. In particular, the reduction to the security of the CSG will run the adversary for differentiated-bit CSG, sample a truly uniform y to feed to the adversary in the last round, and return the adversary's guess x_{Adv}. The y will be properly distributed with probability 1/2, and thus the reduction succeeds with probability at least half that of the differentiated-bit CSG adversary. □

Next, we prove that OSP follows from any CSG with search security.

Theorem 4.7. *CSG with search security implies OSP.*

Proof. We show how to use the differentiated-preimage variant of CSG to build a random-input OSP, and then appeal to Lemma 4.6 to obtain differentiated-preimage CSG from CSG, and Lemma 4.3 to obtain (chosen-input) OSP from random-input OSP. The main idea is to use Goldreich-Levin, similar to how it is used in [BGKM+23], in order to use a distinguisher for the OSP basis to obtain a predictor for the claw-state. The protocol is given in Fig. 1.

OSP from CSG

- The sender S and receiver R begin by running a differentiated-bit CSG, which delivers (x_0, x_1, z) to S and (up to negligible trace distance) $\frac{1}{\sqrt{2}}(|0, x_0\rangle + (-1)^z |1, x_1\rangle)$ to R, where $x_0, x_1 \in \{0,1\}^n$ and $z \in \{0,1\}$.
- Next, S samples $r_0, r_1 \leftarrow \{0,1\}^n$ and sends them to R.
- Using r_0, r_1, the receiver R applies the operation that maps

  $$\frac{1}{\sqrt{2}}(|0, x_1\rangle + |1, x_1\rangle) \to \frac{1}{\sqrt{2}}(|0, x_0\rangle |r_0 \cdot x_0\rangle + (-1)^z |1, x_1\rangle |r_1 \cdot x_1\rangle),$$

  and then measures all but the last qubit in the Hadamard basis to obtain a string $d \in \{0,1\}^{n+1}$, which it returns to S. R outputs its remaining qubit.
- S sets $b := (x_0, x_1) \cdot (r_0, r_1)$. If $b = 0$, then S sets $s := x_0 \cdot r_0 = x_1 \cdot r_1$. If $b = 1$, then S sets $s := z \oplus d \cdot (1, x_0 \oplus x_1)$. S outputs (s, b).

Figure 1: Random-input OSP from any differentiated-bit claw-state generator.

First, we argue correctness. If the sender's output is $b = (x_0, x_1) \cdot (r_0, r_1) = 0$ and $s = x_0 \cdot r_0 = x_1 \cdot r_1$, then the receiver's state before their final measurement is (negligibly close to)

$$\frac{1}{\sqrt{2}}\left(|0, x_0\rangle + (-1)^z |1, x_1\rangle \right) \otimes |s\rangle .$$

So their Hadamard basis measurement has no effect on the last qubit, and their output state will be $|s\rangle = H^b |s\rangle$.

Next, if the sender computes $b = (x_0, x_1) \cdot (r_0, r_1) = 1$, then the receiver's state before their final measurement is either (negligibly close to)

$$\frac{1}{\sqrt{2}} \left(|0, x_0\rangle \ |0\rangle + (-1)^z \ |1, x_1\rangle \ |1\rangle \right) \text{ or } \frac{1}{\sqrt{2}} \left(|0, x_0\rangle \ |1\rangle + (-1)^z \ |1, x_1\rangle \ |0\rangle \right).$$

Either way, a standard calculation shows that if they measure all but their last qubit in the Hadamard basis to obtain d, the last qubit becomes $Z^{z \oplus d \cdot (1, x_0 \oplus x_1)} |+\rangle = H |s\rangle$, for $s = z \oplus d \cdot (1, x_0 \oplus x_1)$.

Now, we establish security. Suppose there exists $\mathsf{Adv} = \{\mathsf{Adv}_\lambda\}_{\lambda \in \mathbb{N}}$ that has $\mathsf{non\text{-}negl}(\lambda)$ advantage in the OSP security game. Then there must be some $\mathsf{non\text{-}negl}(\lambda)$ probability that, after Adv and S interact in the CSG protocol, Adv still has $\mathsf{non\text{-}negl}(\lambda)$ advantage *conditioned on the interaction so far*.

That is, let $|\psi\rangle$ be the state of Adv_λ right after the conclusion of the CSG protocol, and define B to be the routine that takes $r_0, r_1 \in \{0,1\}^n$ as input, runs the remainder of Adv's strategy using state $|\psi\rangle$ and strings (r_0, r_1), and outputs Adv_λ's guess for b. Then we have that

$$\Pr_{|\psi\rangle} \left[\mathop{\mathbb{E}}_{r_0, r_1 \leftarrow \{0,1\}^n} [B(|\psi\rangle, (r_0, r_1)) = (x_0, x_1) \cdot (r_0, r_1)] = \frac{1}{2} + \mathsf{non\text{-}negl}(\lambda) \right] = \mathsf{non\text{-}negl}(\lambda).$$

Now, we appeal to Lemma 3.2, which implies that there exists a B' such that when B has $\mathsf{non\text{-}negl}(\lambda)$ advantage given advice state $|\psi\rangle$, $B'(|\psi\rangle)$ has $\mathsf{non\text{-}negl}(\lambda)$ probability of outputting (x_0, x_1). But this yields an adversary that breaks the security of the CSG with $\mathsf{non\text{-}negl}(\lambda)$ probability, completing the proof.

□

4.3 OSP with generalized angle

Next, we consider a generalized notion of OSP, where the protocol is defined by *any* choice of two (not necessarily mutually unbiased) single-qubit bases. By post-processing with an appropriate rotation, we can without loss of generality consider one basis to be $\{|+\rangle, |-\rangle\}$ and the other to be a rotated basis on the XY plane of the Bloch sphere. For example, by having the receiver apply a Hadamard gate followed by a $\sqrt{X}$ rotation to their state received as output from the protocol described in Definition 4.1, we have that if $b = 0$, the receiver obtains either $|+\rangle$ or $|-\rangle$ and if $b = 1$, the receiver obtains either $\frac{1}{\sqrt{2}}(|0\rangle + i|1\rangle)$ or $\frac{1}{\sqrt{2}}(|0\rangle - i|1\rangle)$.

While this is an example of OSP with mutually unbiased bases (two bases at a maximum angle), one can consider OSP with arbitrary angle between the chosen bases. For any angle $\theta \in [2\pi]$, we define $|+_\theta\rangle := \frac{1}{\sqrt{2}}(|0\rangle + e^{i\theta}|1\rangle)$. Then, OSP with mutually unbiased bases corresponds to OSP with basis $\{|+\rangle, Z|+\rangle\}$ or basis $\{|+_{\pi/2}\rangle, Z|+_{\pi/2}\rangle\}$. For any $\epsilon \in (0,1]$, we define ϵ-OSP to be an OSP with bases $\{|+\rangle, Z|+\rangle\}$ and $\{|+_{\epsilon\pi/2}\rangle, Z|+_{\epsilon\pi/2}\rangle\}$, defined formally as follows.

Definition 4.8 (ϵ-OSP). *An OSP with generalized angle, or ϵ-OSP, is a protocol that takes place between a PPT sender S with input $b \in \{0,1\}$ and a QPT receiver R:*

$$(s, |\psi\rangle) \leftarrow \langle S(1^\lambda, b), R(1^\lambda) \rangle,$$

where $s \in \{0,1\}$ is the sender's output and $|\psi\rangle$ is the receiver's output. It should satisfy the following properties.

- ***Correctness.*** *For any $b \in \{0,1\}$, let*

$$\Pi_{\epsilon\text{-}\mathsf{OSP},b} \coloneqq \sum_{s\in\{0,1\}} |s\rangle\langle s| \otimes Z^s \,|+_{b\epsilon\pi/2}\rangle\langle +_{b\epsilon\pi/2}|\, Z^s.$$

Then for any $b \in \{0,1\}$,

$$\mathbb{E}\left[\|\Pi_{\epsilon\text{-}\mathsf{OSP},b}\,|s\rangle\,|\psi\rangle\,\| : (s,\,|\psi\rangle) \leftarrow \langle S(1^\lambda,b), R(1^\lambda)\rangle\right] = 1 - \mathrm{negl}(\lambda).$$

- ***Security.*** *For any QPT adversary $\{\mathsf{Adv}_\lambda\}_{\lambda\in\mathbb{N}}$,*

$$\Big|\Pr\left[b_{\mathsf{Adv}} = 0 : (s, b_{\mathsf{Adv}}) \leftarrow \langle S(1^\lambda,0), \mathsf{Adv}_\lambda\rangle\right] \\ - \Pr\left[b_{\mathsf{Adv}} = 0 : (s, b_{\mathsf{Adv}}) \leftarrow \langle S(1^\lambda,1), \mathsf{Adv}_\lambda\rangle\right]\Big| = \mathrm{negl}(\lambda).$$

In this work, we give constructions of "standard" OSP (with $\epsilon = 1$), and focus on deriving implications of this notion. However, it would be useful to know if this flavor of OSP is in some sense the "minimal" OSP assumption, and more generally, how ϵ-OSP relates for various choices of ϵ. While we do not fully resolve these questions in this work, we do show that ϵ-OSP implies OSP for many choices of ϵ. In particular, we show the following claim.[14]

Claim 4.9. *For any rational constant $\epsilon \in (0,1]$ that can be written as $\epsilon = c/d$ where c is an odd integer, it holds that ϵ-OSP implies OSP. In particular, for any constant n, $1/n$-OSP implies OSP.*[15]

Proof. The main idea is to use the fact that given two states $Z^{s_1}|+_{\phi_1}\rangle$ and $Z^{s_2}|+_{\phi_2}\rangle$, it is possible to produce the state $Z^{s_1\oplus s_2}|+_{\phi_1+\phi_2}\rangle$ with probability 1/2. That is, we can sum the angles of the states using a procedure that succeeds with probability 1/2 (this procedure is used, for example, in Kuperberg's algorithm [Kup05]). This follows by simply applying a CNOT from the first to the second state, and then measuring the second state in the standard basis. If the result is 0, the first state is now $Z^{s_1\oplus s_2}|+_{\phi_1+\phi_2}\rangle$. To confirm this, we have

$$\begin{aligned}
&\mathrm{CNOT}\left(Z^{s_1}|+_{\phi_1}\rangle \otimes Z^{s_1}|+_{\phi_2}\rangle\right)\\
&\quad= \frac{1}{2}\mathrm{CNOT}\left(Z^{s_1}\otimes Z^{s_2}\right)\Big(|0\rangle + e^{i\phi_1}|1\rangle\Big)\Big(|0\rangle + e^{i\phi_2}|1\rangle\Big)\\
&\quad= \frac{1}{2}\left(Z^{s_1\oplus s_2}\otimes Z^{s_2}\right)\Big(|00\rangle + e^{i\phi_2}|01\rangle + e^{i\phi_1}|11\rangle + e^{i(\phi_1+\phi_2)}|10\rangle\Big)\\
&\quad= \frac{1}{2}\left(Z^{s_1\oplus s_2}\otimes Z^{s_2}\right)\Big(\Big(|0\rangle + e^{i(\phi_1+\phi_2)}|1\rangle\Big)|0\rangle + e^{i\phi_2}\Big(|0\rangle + e^{i(\phi_1-\phi_2)}|1\rangle\Big)|1\rangle\Big)\\
&\quad= \frac{1}{\sqrt{2}}Z^{s_1\oplus s_2}|+_{\phi_1+\phi_2}\rangle|0\rangle + \frac{1}{\sqrt{2}}Z^{s_2}e^{i\phi_1}|+_{\phi_1-\phi_2}\rangle|1\rangle.
\end{aligned}$$

[14] Ideally, we would like to show the implication for any rational ϵ, but the current approach does not appear to work for $\epsilon =$ even/odd.

[15] In fact, our proof shows that ϵ-OSP implies OSP even for any inverse-polynomial $\epsilon = 1/\mathrm{poly}(\lambda)$. Thus, OSP with any small enough but still non-trivial angle implies standard OSP. We note that [DK16] has established similar results in the information-theoretic setting.

Now, for simplicity suppose that d is a power of 2 (a similar procedure works for arbitrary d), and let λ be the security parameter. Given input $b \in \{0, 1\}$, the parties will run $8^{\log d}\lambda = d^3\lambda = \mathrm{poly}(\lambda)$ many ϵ-OSP protocols with sender input b. The receiver is now in possession of $8^{\log d}\lambda$ many states (negligibly close to) $Z^s |+_{b\epsilon\pi/2}\rangle$, each with potentially different s. The receiver now runs a procedure to "sum" d of them together. If $b = 0$, this results in a state

$$Z^{s'} |+\rangle \in \{ |+\rangle , |-\rangle \} ,$$

whereas if $b = 1$, this results in a state

$$Z^{s'} |+_{d\epsilon\pi/2}\rangle = Z^{s'} |+_{c\pi/2}\rangle \in \left\{ |+_{\pi/2}\rangle , Z |+_{\pi/2}\rangle \right\} ,$$

where the final inclusion uses the fact that c is an odd integer, and s' is the XOR of all the s corresponding to states involved in the sum.

To obtain the desired sum, the receiver operates in layers. Given $8k\lambda$ pairs of states at angle ϕ, the receiver splits them into $4k\lambda$ pairs of states, and applies the above CNOT-and-measure procedure to each pair. The expected number of "successes" is $2k\lambda$, and by Chernoff, there will be at least $k\lambda$ successes with all but $\mathrm{negl}(\lambda)$ probability. Thus, with all but $\mathrm{negl}(\lambda)$ probability, the receiver obtains $k\lambda$ states at angle 2ϕ.

Now, starting with $8^{\log d}\lambda$ many states at angle ϕ and operating for $\log d$ layers, the receiver will end up with $\lambda \geq 1$ state at angle $2^{\log d}\phi = d\phi$. This establishes correctness of the OSP protocol.

Finally, security of the OSP protocol follows by a standard hybrid argument from the security of the ϵ-OSP protocol, which we have simply repeated $d^3\lambda$ times.

□

We conclude this section by proving a natural structural lemma about OSP, and more generally ϵ-OSP. It shows that in the honest case, the sender's output s must be (negligibly) close to uniformly random for either choice of $b \in \{0, 1\}$. That is, an honest run of ϵ-OSP produces a uniformly random eigenstate of the observable specified by b (though we caution that this is no longer necessarily true when the receiver is adversarial).

Lemma 4.10. *For any $\epsilon \in (0, 1]$ such that $\epsilon \geq 1/\mathrm{poly}(\lambda)$, any secure ϵ-OSP protocol, and any $b \in \{0, 1\}$, it holds that*

$$\left| \Pr\left[s = 0 : (s, |\psi\rangle) \leftarrow \langle S(1^\lambda, b), R(1^\lambda) \rangle \right] - \frac{1}{2} \right| = \mathrm{negl}(\lambda).$$

Proof. Suppose otherwise, and without loss of generality suppose that when $b = 1$, the probability that $s = 0$ is equal to $1/2 + \delta$ for some $\delta = \textsf{non-negl}(\lambda)$ (the other cases are symmetric). We show that this would contradict the security of the ϵ-OSP. Consider an adversary Adv that participates as an honest receiver, measures their final state in the Y basis $\{ |+_{\pi/2}\rangle , Z |+_{\pi/2}\rangle\}$, and guesses $b = 1$ if they obtain outcome $|+_{\pi/2}\rangle$. First note that, no matter what the distribution on s is when $b = 0$, we have that

$$\left| \Pr[\mathsf{Adv} = 1 : b = 0] - \frac{1}{2} \right| = \mathrm{negl}(\lambda),$$

since measuring either the $|+\rangle$ or $|-\rangle$ state in the Y basis yields a uniformly random outcome. Next, let

$$\epsilon' := \cos^2 \left((1 - \epsilon) \frac{\pi}{4} \right)$$

be the probability of obtaining outcome $|+_{\pi/2}\rangle$ when measuring the state $|+_{\epsilon\pi/2}\rangle$ in the Y basis, and note that $\epsilon' \geq 1/2 + 1/\mathrm{poly}(\lambda)$ whenever $\epsilon \geq 1/\mathrm{poly}(\lambda)$. Then to complete the proof, we have that

$$\begin{aligned}\Pr\left[\mathsf{Adv} = 1 : b = 1\right] &\geq \left(\frac{1}{2} + \delta\right)\epsilon' + \left(\frac{1}{2} - \delta\right)(1 - \epsilon') - \mathrm{negl}(\lambda) \\ &= \frac{1}{2} - \delta + 2\delta\epsilon' - \mathrm{negl}(\lambda) \\ &\geq \frac{1}{2} + \frac{2\delta}{\mathrm{poly}(\lambda)} - \mathrm{negl}(\lambda) \\ &= \frac{1}{2} + \mathsf{non\text{-}negl}(\lambda).\end{aligned}$$

□

5 Constructions

In this section, we provide constructions of OSP from trapdoor claw-free functions (TCFs).

In Section 5.1, we show how to construct OSP from any (plain) TCF, meaning we assume no extra properties such as dual-mode or adaptive hardcore bit. Moreover, all that we require from the TCF is that there is some inverse-polynomial probability that an image has exactly two preimages (which can be obtained efficiently using the trapdoor), and otherwise it can have 1 or 3 or more (as long as the trapdoor correctly identifies these as "bad" images).

In Section 5.2, we show how to obtain a *two-round* OSP by assuming an extra *dual-mode* property of the TCF. A dual-mode TCF (dTCF) can be sampled in either a disjoint mode or lossy mode. Again, we only assume that there is some inverse polynomial probability that an image has two preimages in lossy mode, and we use the amplification lemma for dTCFs recently established by [GV24] in order to show that such dTCFs still imply two-round OSP.

Before coming to the constructions, we provide definitions of TCFs.

Definition 5.1 (Trapdoor claw-free function)**.** *A trapdoor claw-free function (TCF) consists of a PPT parameter generation algorithm* $\mathsf{Gen}(1^\lambda) \to \mathsf{pp}, \mathsf{sp}$ *and a keyed family of PPT computable functions*

$$\left\{\left\{F_{\mathsf{pp}} : \{0,1\}^{n(\lambda)} \to \{0,1\}^{m(\lambda)}\right\}_{(\mathsf{pp},\cdot)\in\mathsf{Gen}(1^\lambda)}\right\}_{\lambda\in\mathbb{N}}.$$

There exists a family of distributions

$$\left\{\{\mathcal{D}_{\mathsf{pp}}\}_{(\mathsf{pp},\cdot)\in\mathsf{Gen}(1^\lambda)}\right\}_{\lambda\in\mathbb{N}}$$

over $\{0,1\}^{n(\lambda)}$ *and a PPT algorithm* $\mathsf{Invert}(\mathsf{sp}, y)$ *such that the following properties are satisfied.*

- ***Efficient state preparation.*** *There is a QPT algorithm that, given any* $(\mathsf{pp}, \cdot) \in \mathsf{Gen}(1^\lambda)$, *outputs a state within negligible trace distance of the state*

$$|\psi_{\mathsf{pp}}\rangle := \sum_{x\in\{0,1\}^{n(\lambda)}} \sqrt{\mathcal{D}_{\mathsf{pp}}(x)}\,|x\rangle\,.$$

- ***Efficient inversion.*** *For any* $(\mathsf{pp}, \cdot) \in \mathsf{Gen}(1^\lambda)$*, let* $\mathsf{Claw}_\lambda \subseteq \{0,1\}^{m(\lambda)}$ *be the set of* $y \in \{0,1\}^{m(\lambda)}$ *such that there exists exactly two* $x_0, x_1 \in \{0,1\}^{n(\lambda)}$ *such that* $F_{\mathsf{pp}}(x_0) = F_{\mathsf{pp}}(x_1) = y$*, and*

$$\left|\frac{\mathcal{D}_{\mathsf{pp}}(x_0)}{\mathcal{D}_{\mathsf{pp}}(x_0) + \mathcal{D}_{\mathsf{pp}}(x_1)} - \frac{\mathcal{D}_{\mathsf{pp}}(x_1)}{\mathcal{D}_{\mathsf{pp}}(x_0) + \mathcal{D}_{\mathsf{pp}}(x_1)}\right| = \mathrm{negl}(\lambda).$$

Then there exists $\delta(\lambda) = 1/\mathrm{poly}(\lambda)$ *such that*

$$\Pr\left[F_{\mathsf{pp}}(x_0) = F_{\mathsf{pp}}(x_1) = y : \begin{array}{r} (\mathsf{pp}, \mathsf{sp}) \leftarrow \mathsf{Gen}(1^\lambda) \\ x \leftarrow \mathcal{D}_{\mathsf{pp}} \\ y := F_{\mathsf{pp}}(x) \\ \{x_0, x_1\} \leftarrow \mathsf{Invert}(\mathsf{sp}, y) \end{array}\right] \geq \delta(\lambda),$$

and for all $y \notin \mathsf{Claw}_\lambda$*,* $\mathsf{Invert}(\mathsf{sp}, y) = \bot$.

- ***Claw-free.*** *For any QPT adversary* $\{\mathsf{Adv}_\lambda\}_{\lambda \in \mathbb{N}}$*,*

$$\Pr\left[F_{\mathsf{pp}}(x_0) = F_{\mathsf{pp}}(x_1) : \begin{array}{r} (\mathsf{pp}, \mathsf{sp}) \leftarrow \mathsf{Gen}(1^\lambda) \\ \{x_0, x_1\} \leftarrow \mathsf{Adv}_\lambda(\mathsf{pp}) \end{array}\right] = \mathrm{negl}(\lambda).$$

Next, we define a *dual-mode* variant of TCFs.

Definition 5.2 (Dual-mode trapdoor claw-free function)**.** *A dual-mode trapdoor claw-free function (dTCF) consists of a PPT parameter generation algorithm* $\mathsf{Gen}(1^\lambda, \mu) \to \mathsf{pp}, \mathsf{sp}$ *that takes as input a "mode" bit* $\mu \in \{0,1\}$*, and a keyed family of PPT computable functions*[16]

$$\left\{\left\{F_{\mathsf{pp}} : \{0,1\} \times \{0,1\}^{n(\lambda)} \to \{0,1\}^{m(\lambda)}\right\}_{(\mathsf{pp},\cdot) \in \mathsf{Gen}(1^\lambda, \mu)}\right\}_{\mu \in \{0,1\}, \lambda \in \mathbb{N}}.$$

There exists a family of distributions

$$\left\{\{\mathcal{D}_{\mathsf{pp}}\}_{(\mathsf{pp},\cdot) \in \mathsf{Gen}(1^\lambda, \mu)}\right\}_{\mu \in \{0,1\}, \lambda \in \mathbb{N}}$$

over $\{0,1\}^{n(\lambda)}$ *and a PPT algorithm* $\mathsf{Invert}(\mathsf{sp}, b, y)$ *such that the following properties are satisfied.*

- ***Efficient state preparation.*** *There is a QPT algorithm that, given any* $(\mathsf{pp}, \cdot) \in \mathsf{Gen}(1^\lambda, \mu)$*, outputs a state within negligible trace distance of the state*

$$|\psi_{\mathsf{pp}}\rangle := \sum_{x \in \{0,1\}^{n(\lambda)}} \sqrt{\mathcal{D}_{\mathsf{pp}}(x)}\, |x\rangle\,.$$

- ***Efficient inversion.*** *For any* $\mu \in \{0,1\}$ *and* $b \in \{0,1\}$*,*[17]

$$\Pr\left[\mathsf{Invert}(\mathsf{sp}, b, y) = x : \begin{array}{r} (\mathsf{pp}, \mathsf{sp}) \leftarrow \mathsf{Gen}(1^\lambda, \mu) \\ x \leftarrow \mathcal{D}_{\mathsf{pp}} \\ y := F_{\mathsf{pp}}(b, x) \end{array}\right] = 1 - \mathrm{negl}(\lambda).$$

[16] Notice that, as compared to plain TCFs, dual-mode TCFs take an extra bit of input. We will require that each claw has one preimage that starts with 0 and one that starts with 1, which is important for the amplification lemma of [GV24].

[17] Note that this property implies that with overwhelming probability over $(\mathsf{pp}, \mathsf{sp}) \leftarrow \mathsf{Gen}(1^\lambda, \mu)$ and $x \leftarrow \mathcal{D}_{\mathsf{pp}}$, x has no siblings x' such that $F_{\mathsf{pp}}(b, x) = F_{\mathsf{pp}}(b, x')$. That is, $F_{\mathsf{pp}}(0, \cdot)$ and $F_{\mathsf{pp}}(1, \cdot)$ are effectively injective.

- ***Dual-mode.***
  - *Disjoint mode ($\mu = 0$): For any $b \in \{0,1\}$,*

$$\Pr\left[\exists x' \text{ s.t. } F_{\mathsf{pp}}(1-b,x') = y : \begin{array}{r} (\mathsf{pp},\mathsf{sp}) \leftarrow \mathsf{Gen}(1^\lambda, 0) \\ x \leftarrow \mathcal{D}_{\mathsf{pp}} \\ y := F_{\mathsf{pp}}(b,x) \end{array}\right] = \mathrm{negl}(\lambda).$$

  - *Lossy mode ($\mu = 1$): There exists $\delta(\lambda) = 1/\mathrm{poly}(\lambda)$ such that for any $b \in \{0,1\}$,*

$$\Pr\left[\exists x' \text{ s.t. } F_{\mathsf{pp}}(1-b,x') = y \;:\; \begin{array}{r} (\mathsf{pp},\mathsf{sp}) \leftarrow \mathsf{Gen}(1^\lambda, 1) \\ x \leftarrow \mathcal{D}_{\mathsf{pp}} \\ y := F_{\mathsf{pp}}(b,x) \end{array}\right] \geq \delta(\lambda),$$

    and for any $\nu(\lambda) = \mathsf{non\text{-}negl}(\lambda)$,

$$\Pr\left[\begin{array}{l} \exists x' \text{ s.t. } F_{\mathsf{pp}}(1-b,x') = y \\ \wedge \left|\frac{\mathcal{D}_{\mathsf{pp}}(x)}{\mathcal{D}_{\mathsf{pp}}(x)+\mathcal{D}_{\mathsf{pp}}(x')} - \frac{\mathcal{D}_{\mathsf{pp}}(x')}{\mathcal{D}_{\mathsf{pp}}(x)+\mathcal{D}_{\mathsf{pp}}(x')}\right| \geq \nu(\lambda) \end{array} : \begin{array}{r} (\mathsf{pp},\mathsf{sp}) \leftarrow \mathsf{Gen}(1^\lambda, 1) \\ x \leftarrow \mathcal{D}_{\mathsf{pp}} \\ y := F_{\mathsf{pp}}(b,x) \end{array}\right] = \mathrm{negl}(\lambda),$$

    where this last requirement enforces that there are (effectively) no "unbalanced" claws.

- ***Mode indistinguishability.*** *For any QPT adversary $\{\mathsf{Adv}_\lambda\}_{\lambda\in\mathbb{N}}$,*

$$\left|\Pr\left[\mathsf{Adv}_\lambda(\mathsf{pp}) = 1 : (\mathsf{pp},\mathsf{sp}) \leftarrow \mathsf{Gen}(1^\lambda, 0)\right] - \Pr\left[\mathsf{Adv}_\lambda(\mathsf{pp}) = 1 : (\mathsf{pp},\mathsf{sp}) \leftarrow \mathsf{Gen}(1^\lambda, 1)\right]\right| = \mathrm{negl}(\lambda).$$

Remark 5.3. *Dual-mode (and thus plain) TCFs are known from LWE [BCM+18] (even with polynomial modulus-to-noise ratio, since we can take $\delta = 1 - 1/\mathrm{poly}$) and from the "extended linear hidden shift" assumption on cryptographic group actions [AMR22, GV24].*

5.1 OSP from plain TCFs

Theorem 5.4. *Any TCF satisfying Definition 5.1 implies OSP (Definition 4.1).*

Proof. This follows fairly immediately by using the TCF to construct a claw-state generator (Definition 4.4), and then appealing to Theorem 4.7, which shows how to construct OSP from any CSG. Let $\delta = \delta(\lambda)$ be the parameter from the efficient inversion property of the TCF. Then the CSG is constructed as follows.

- S samples $(\mathsf{sp},\mathsf{pp}) \leftarrow \mathsf{Gen}(1^\lambda)$ and sends pp to R.
- Run the following at most λ/δ times. If the protocol has not terminated at that point, R outputs $\bot$, and S samples $x_0, x_1 \leftarrow \{0,1\}^n$ and outputs $(x_0, x_1, 0)$.
  - R prepares a state within negligible trace distance of $|\psi_{\mathsf{pp}}\rangle$, applies F_{pp} in superposition to a fresh register, and measures that register to obtain y.
  - R sends y to S. If $\mathsf{Invert}(\mathsf{sp}, y) = \{x_0, x_1\}$, S outputs $(x_0, x_1, 0)$, instructs R to terminate, and R outputs its state. Otherwise, if $\mathsf{Invert}(\mathsf{sp}, y) = \bot$, S instructs R to repeat.

Correctness of the CSG follows first by noting that the probability that the parties never continue from the loop is at most $(1-\delta)^{\lambda/\delta} \leq e^{-\lambda} = \text{negl}(\lambda)$. Then, when $\mathsf{Invert}(\mathsf{sp}, y) = \{x_0, x_1\}$, we know that R's state is negligibly close to $\frac{1}{\sqrt{2}}(|x_0\rangle + |x_1\rangle)$, by the efficient state preparation and efficient inversion properties of the TCF.

To show security of the CSG, first note that any adversary with noticeable probability of guessing S's output (x_0, x_1) must cause S to terminate at some point during the loop, since otherwise (x_0, x_1) is uniformly random and independent of their view. Then, consider the following reduction to the claw-freeness of the TCF. The reduction receives pp from its challenger, samples a random round $i \leftarrow [\lambda/\delta]$, runs the adversary and instructs them to terminate on the i'th invocation of the loop, and returns the adversary's guess for (x_0, x_1). With probability at least $1/(\lambda/\delta) = 1/\text{poly}(\lambda)$, this matches the adversary's view in the real protocol, which means that the reduction has a noticeable probability of outputting a claw (x_0, x_1).

□

5.2 Two-round OSP from dual-mode TCFs

In [GV24], it is shown that any dTCF satisfying Definition 5.2 (i.e. with any $\delta = 1/\text{poly}(\lambda)$) implies the following variant of dTCF with a *phase computation* property that succeeds with all but negligible probability. Note that we also relax the inversion property to a *partial inversion* property, which only requires that the first bit b of the preimage be recovered. This allows for the possibility that the functions $F_{\mathsf{pp}}(0, \cdot)$ and $F_{\mathsf{pp}}(1, \cdot)$ are non-injective.

Definition 5.5 (dTCF with efficient phase computation). *A dTCF with* efficient phase computation *consists of a PPT parameter generation algorithm* $\mathsf{Gen}(1^\lambda, \mu) \rightarrow \mathsf{pp}, \mathsf{sp}$ *that takes as input a "mode" bit* $\mu \in \{0,1\}$, *and a keyed family of PPT computable functions*

$$\left\{ \left\{ F_{\mathsf{pp}} : \{0,1\} \times \{0,1\}^{n(\lambda)} \rightarrow \{0,1\}^{m(\lambda)} \right\}_{(\mathsf{pp}, \cdot) \in \mathsf{Gen}(1^\lambda, \mu)} \right\}_{\mu \in \{0,1\}, \lambda \in \mathbb{N}}.$$

There exists a family of distributions

$$\left\{ \{\mathcal{D}_{\mathsf{pp}}\}_{(\mathsf{pp}, \cdot) \in \mathsf{Gen}(1^\lambda, \mu)} \right\}_{\mu \in \{0,1\}, \lambda \in \mathbb{N}}$$

over $\{0,1\}^{n(\lambda)}$ *and PPT algorithms* $\mathsf{PartialInvert}(\mathsf{sp}, y)$ *and* $\mathsf{PhaseInvert}(\mathsf{sp}, y, d)$ *such that the following properties are satisfied.*

- ***Efficient state preparation.*** *There is a QPT algorithm that, given any* $(\mathsf{pp}, \cdot) \in \mathsf{Gen}(1^\lambda, \mu)$, *outputs a state within negligible trace distance of the state*

$$|\psi_{\mathsf{pp}}\rangle := \sum_{x \in \{0,1\}^{n(\lambda)}} \sqrt{\mathcal{D}_{\mathsf{pp}}(x)}\, |x\rangle .$$

- ***Efficient partial inversion.*** *For any* $\mu \in \{0,1\}$,

$$\Pr\left[B = \{b : \exists x \text{ s.t. } F_{\mathsf{pp}}(b, x) = y\} : \begin{array}{r} (\mathsf{pp}, \mathsf{sp}) \leftarrow \mathsf{Gen}(1^\lambda, \mu) \\ x \leftarrow \mathcal{D}_{\mathsf{pp}} \\ y := F_{\mathsf{pp}}(b, x) \\ B \leftarrow \mathsf{PartialInvert}(\mathsf{sp}, y) \end{array} \right] = 1 - \text{negl}(\lambda).$$

- ***Dual-mode.***
  - *Disjoint mode ($\mu = 0$): For any $b \in \{0,1\}$,*

$$\Pr\left[\exists x' \text{ s.t. } F_{\mathsf{pp}}(1-b,x') = y : \begin{array}{r}(\mathsf{pp},\mathsf{sp}) \leftarrow \mathsf{Gen}(1^\lambda, 0)\\ x \leftarrow \mathcal{D}_{\mathsf{pp}}\\ y := F_{\mathsf{pp}}(b,x)\end{array}\right] = \mathrm{negl}(\lambda).$$

  - *Lossy mode ($\mu = 1$): For any $(\mathsf{pp},\cdot) \in \mathsf{Gen}(1^\lambda,1)$, $y \in \{0,1\}^{m(\lambda)}$, and $d \in \{0,1\}^{n(\lambda)}$, define*

$$w_{\mathsf{pp},y,d,0} := \sum_{x:F_{\mathsf{pp}}(0,x)=y} (-1)^{d\cdot x}\sqrt{\mathcal{D}_{\mathsf{pp}}(x)}, \quad w_{\mathsf{pp},y,d,1} := \sum_{x:F_{\mathsf{pp}}(1,x)=y} (-1)^{d\cdot x}\sqrt{\mathcal{D}_{\mathsf{pp}}(x)},$$

and re-normalize

$$\widetilde{w}_{\mathsf{pp},y,d,0} := \frac{w_{\mathsf{pp},y,d,0}}{\sqrt{w^2_{\mathsf{pp},y,d,0} + w^2_{\mathsf{pp},y,d,1}}}, \quad \widetilde{w}_{\mathsf{pp},y,d,1} := \frac{w_{\mathsf{pp},y,d,1}}{\sqrt{w^2_{\mathsf{pp},y,d,0} + w^2_{\mathsf{pp},y,d,1}}}.$$

Then for any $d \in \{0,1\}^{n(\lambda)}$, there exists $\nu(\lambda) = \mathrm{negl}(\lambda)$ such that

$$\Pr\left[\left|\widetilde{w}_{\mathsf{pp},y,d,0} - (-1)^s \widetilde{w}_{\mathsf{pp},y,d,1}\right| \le \nu(\lambda) : \begin{array}{r}(\mathsf{pp},\mathsf{sp}) \leftarrow \mathsf{Gen}(1^\lambda, 1)\\ b \leftarrow \{0,1\}\\ x \leftarrow \mathcal{D}_{\mathsf{pp}}\\ y := F_{\mathsf{pp}}(b,x)\\ s \leftarrow \mathsf{PhaseInvert}(\mathsf{sp},y,d)\end{array}\right] = 1 - \mathrm{negl}(\lambda).$$

- ***Mode indistinguishability.*** *For any QPT adversary $\{\mathsf{Adv}_\lambda\}_{\lambda\in\mathbb{N}}$,*

$$\left|\Pr\left[\mathsf{Adv}_\lambda(\mathsf{pp}) = 1 : (\mathsf{pp},\mathsf{sp}) \leftarrow \mathsf{Gen}(1^\lambda,0)\right] - \Pr\left[\mathsf{Adv}_\lambda(\mathsf{pp}) = 1 : (\mathsf{pp},\mathsf{sp}) \leftarrow \mathsf{Gen}(1^\lambda,1)\right]\right| = \mathrm{negl}(\lambda).$$

Lemma 5.6 ([GV24], Amplication for dTCFs)**.** *Any dTCF satisfying Definition 5.2 implies a dTCF with efficient phase computation (Definition 5.5).*[18]

Now, we show how to construct OSP from any dTCF satisfying Definition 5.5.

Theorem 5.7. *Any dTCF with efficient phase computation (Definition 5.5) implies two-round OSP.*

Proof. The construction is given in Fig. 2. Security follows immediately from the mode indistinguishability property of the dTCF, so it remains to argue correctness.

In the case that $b = 0$, the efficient partial inversion and disjoint mode properties of the dTCF imply that with all but negligible probability, (1) all of y's preimages begin with the same bit s, and thus the receiver's state on $\mathcal{B}$ collapses to a standard basis state $|s\rangle$, and (2) the $\mathsf{PartialInvert}$ algorithm will return this s given (sp, y), so the sender will obtain the correct description of the receiver's state.

[18] Technically, [GV24] assume that the dTCF input to their amplification lemma has claws that are perfectly balanced, as opposed to almost perfectly balanced, but it can be confirmed that their amplification lemma holds even when the claws are $1 - \mathrm{negl}(\lambda)$ balanced.

In the case that $b = 1$, the receiver's state after measuring y is

$$\propto \sum_{x : F_{\mathsf{pp}}(0,x)=y} \sqrt{\mathcal{D}_{\mathsf{pp}}(x)} \left|0\right\rangle_{\mathcal{B}} \left|x\right\rangle_{\mathcal{X}} + \sum_{x : F_{\mathsf{pp}}(1,x)=y} \sqrt{\mathcal{D}_{\mathsf{pp}}(x)} \left|1\right\rangle_{\mathcal{B}} \left|x\right\rangle_{\mathcal{X}} .$$

Thus, after measuring $\mathcal{X}$ in the Hadamard basis to obtain d, the state collapses to a state

$$\propto \sum_{x : F_{\mathsf{pp}}(0,x)=y} (-1)^{d \cdot x} \sqrt{\mathcal{D}_{\mathsf{pp}}(x)} \left|0\right\rangle + \sum_{x : F_{\mathsf{pp}}(1,x)=y} (-1)^{d \cdot x} \sqrt{\mathcal{D}_{\mathsf{pp}}(x)} \left|1\right\rangle = w_{\mathsf{pp},y,d,0} \left|0\right\rangle + w_{\mathsf{pp},y,d,1} \left|1\right\rangle ,$$

which after normalizing, is equal to

$$\widetilde{w}_{\mathsf{pp},y,d,0} \left|0\right\rangle + \widetilde{w}_{\mathsf{pp},y,d,1} \left|1\right\rangle .$$

The efficient phase computation property implies that with all but negligible probability, the sender outputs s such that $|\widetilde{w}_{\mathsf{pp},y,d,0} = (-1)^s \widetilde{w}_{\mathsf{pp},y,d,1}| = \text{negl}(\lambda)$, in which case the receiver's state is negligibly close to $Z^s \left|+\right\rangle$. This completes the proof.

Two-round OSP from dTCF

- $\mathsf{OSP.Sen}(1^\lambda, b)$: Sample $(\mathsf{pp}, \mathsf{sp}) \leftarrow \mathsf{Gen}(1^\lambda, b)$, and define $\mathsf{msg}_S := \mathsf{pp}$, and $\mathsf{st}_S := \mathsf{sp}$.
- $\mathsf{OSP.Rec}(\mathsf{msg}_S)$: Prepare a state within negligible trace distance of $\left|+\right\rangle_{\mathcal{B}} \left|\psi_{\mathsf{pp}}\right\rangle_{\mathcal{X}}$, apply F_{pp} in superposition to a fresh register, and measure that register to obtain y. Then, measure register $\mathcal{X}$ in the Hadamard basis to obtain $d \in \{0,1\}^n$. Finally, output the remaining qubit on register $\mathcal{B}$ and set $\mathsf{msg}_R := (y, d)$.
- $\mathsf{OSP.Dec}(\mathsf{st}_S, \mathsf{msg}_R)$:
  - If $b = 0$, compute $s \leftarrow \mathsf{PartialInvert}(\mathsf{sp}, y)$, and output s.
  - If $b = 1$, compute $s \leftarrow \mathsf{PhaseInvert}(\mathsf{sp}, y, d)$, and output s.

Figure 2: A consruction of two-round OSP from any dTCF with efficient phase computation (Definition 5.5), which is known from any dTCF (Definition 5.2).

□

6 Applications

In this section, we move to the applications of OSP, establishing the following results.

- Section 6.1: OSP implies proofs of quantumness (as well as a "test of a qubit" and certifiable randomness).
- Section 6.2: Two-round OSP implies 1-of-2 puzzles, which previous work has shown implies (privately-verifiable) quantum money with classical communication, and position verification with classical communication.
- Section 6.3: OSP implies blind classical delegation of quantum computation.

- Section 6.4: Blind classical delegation generically implies verifiable classical delegation, so we can conclude that OSP implies verifiable classical delegation.
- Section 6.5: (Two-round) OSP implies (two-round) encrypted CNOT, which yields claw-state generators with indistinguishability security and (additionally assuming classical FHE with log-depth decryption) quantum FHE.

6.1 Proofs of quantumness

Definition 6.1 (Proof of quantumness)**.** *A proof of quantumness protocol is an interaction between a QPT prover and a PPT verifier*

$$\{\top, \bot\} \leftarrow \langle P(1^\lambda), V(1^\lambda)\rangle,$$

where $\{\top, \bot\}$ is the output of the verifier. There exists $\epsilon(\lambda), \delta(\lambda)$ with $\epsilon(\lambda) - \delta(\lambda) = 1/\mathrm{poly}(\lambda)$ such that the following properties hold.

- ***Completeness.***
$$\Pr\left[\top \leftarrow \langle P(1^\lambda), V(1^\lambda)\rangle\right] \geq \epsilon(\lambda).$$
- ***Soundness.*** *For any PPT adversary $\{\mathsf{Adv}_\lambda\}_{\lambda\in\mathbb{N}}$,*
$$\Pr\left[\top \leftarrow \langle \mathsf{Adv}_\lambda, V(1^\lambda)\rangle\right] \leq \delta(\lambda) + \mathrm{negl}(\lambda).$$

Theorem 6.2. *OSP (Definition 4.1) implies a proof of quantumness (Definition 6.1).*

Proof. We describe the protocol in Fig. 3. The construction and proof follow the presentation in [ABCC24], which is based on ideas originated in [KMCVY21].

Proof of quantumness from OSP

- The verifier samples $r \leftarrow \{0,1\}$ and acts as the sender in an OSP with the prover. The verifier receives a bit s and the prover receives a state (negligibly close to) $H^r\left|s\right\rangle$.
- The verifier samples $a \leftarrow \{0,1\}$ and sends a to the prover.
- If $a = 0$, the prover measures their state in the $X + Z$ basis, and if $a = 1$, the prover measures their state in the $X - Z$ basis, to obtain a bit b. The prover sends b to the verifier.
- The verifier accepts if $b = s \oplus r \cdot a$.

Figure 3: Proof of quantumness from OSP.

First we show that the protocol has completeness $\epsilon > 0.85$. By applying gentle measurement (Lemma 3.1), we can take the prover's state at the conclusion of the OSP protocol to be exactly $H^r\left|s\right\rangle$, and only lose a $\mathrm{negl}(\lambda)$ factor in the final bound. By a standard calculation, measuring $H^r\left|s\right\rangle$ in the X+Z basis yields s with probability $\cos^2(\pi/8)$, and measuring $H^r\left|s\right\rangle$ in the X-Z basis yields $s \oplus r$ with probability $\cos^2(\pi/8)$. Thus, the verifier accepts with probabilty at least $\cos^2(\pi/8) - \mathrm{negl}(\lambda) > 0.85$.

Now, we show that the protocol has soundness $\delta = 0.75$. To see this, suppose a PPT prover $\mathsf{Adv} = \{\mathsf{Adv}_\lambda\}_{\lambda\in\mathbb{N}}$ has advantage $0.75 + \mathsf{non\text{-}negl}(\lambda)$ in the protocol. Now consider the following procedure Adv' attacking the security of the OSP protocol.

- Interact as Adv in the OSP protocol and let $\mathsf{st}_{\mathsf{Adv}}$ be the state of Adv at the conclusion of the OSP protocol.
- Run $\mathsf{Adv}(\mathsf{st}_{\mathsf{Adv}}, 0) \to b_0$ and $\mathsf{Adv}(\mathsf{st}_{\mathsf{Adv}}, 1) \to b_1$ to obtain a final-round answer on each possible question $a \in \{0,1\}$ from the verifier.
- Output $b_0 \oplus b_1$ as the guess for r.

Let D be the distribution over $(\mathsf{st}_{\mathsf{Adv}}, r, s)$ that results from running the OSP protocol with Adv on a random input r, and defining s to be the sender's output. For any $(\mathsf{st}_{\mathsf{Adv}}, r, s)$ in the support of D and any $a \in \{0,1\}$, define $p_{\mathsf{st}_{\mathsf{Adv}},r,s}[a] := \Pr[\mathsf{Adv}(\mathsf{st}_{\mathsf{Adv}}, a) = s \oplus r \cdot a]$.

Then we have that

$$\begin{aligned}
\Pr\left[\mathsf{Adv}' = r\right] &\geq \mathop{\mathbb{E}}_{(\mathsf{st}_{\mathsf{Adv}},r,s)\leftarrow D}\left[p_{\mathsf{st}_{\mathsf{Adv}},r,s}[0] \cdot p_{\mathsf{st}_{\mathsf{Adv}},r,s}[1]\right] \\
&\geq \mathop{\mathbb{E}}_{(\mathsf{st}_{\mathsf{Adv}},r,s)\leftarrow D}\left[p_{\mathsf{st}_{\mathsf{Adv}},r,s}[0] + p_{\mathsf{st}_{\mathsf{Adv}},r,s}[1] - 1\right] \\
&= 2 \mathop{\mathbb{E}}_{(\mathsf{st}_{\mathsf{Adv}},r,s)\leftarrow D}\left[\frac{1}{2}\left(p_{\mathsf{st}_{\mathsf{Adv}},r,s}[0] + p_{\mathsf{st}_{\mathsf{Adv}},r,s}[1]\right)\right] - 1 \\
&= 2(0.75 + \mathsf{non\text{-}negl}(\lambda)) - 1 \\
&= 0.5 + \mathsf{non\text{-}negl}(\lambda),
\end{aligned}$$

where the second inequality follows from the fact that $x \cdot y \geq x + y - 1$ for any $x, y \in [0,1]$, and the second equality follows from the fact that

$$\mathop{\mathbb{E}}_{(\mathsf{st}_{\mathsf{Adv}},r,s)\leftarrow D}\left[\frac{1}{2}\left(p_{\mathsf{st}_{\mathsf{Adv}},r,s}[0] + p_{\mathsf{st}_{\mathsf{Adv}},r,s}[1]\right)\right]$$

is exactly Adv's advantage in the proof of quantumness. This yields a contradiction to the security of OSP, completing the proof.

□

To conclude this section, we note that Fig. 3 fits the protocol template from [BGKM+23, Figure 1] (unsurprisingly, since it is essentially a more modular presentation of the protocol from [BGKM+23, Section 5.3]), and thus inherits the results established by [BGKM+23] about this class of proof of quantumness protocols. In particular it implies (1) a "test of a qubit" ([BGKM+23, Theorem 4.7]), meaning that any *quantum* prover with advantage close to $\cos^2(\pi/8)$ must be using (close to) anti-commuting operators in the final round, and (2) the ability to generate certifiable randomness from a quantum prover. We refer the reader to [BCM+18, BGKM+23, Vid20, MAF23] for formal definitions and more discussion on the notions of test of a qubit and certifiable randomness. We stress that, in general, proofs of quantumness (e.g. Shor's algorithm [Sho97] and Yamakawa-Zhandry [YZ24]) do not always imply a test of a qubit, but our OSP-based proof of quantumness does since it fits the template described in [BGKM+23].

6.2 1-of-2 puzzles

Next, we show that two-round OSP implies a *1-of-2 puzzle*, which was originally defined by [RS19]. In words, a 1-of-2 puzzle is a task with two challenges such that a prover can answer either but not both simultaneously. It is a useful abstraction, as it has been shown to imply both privately-verifiable quantum money [RS19] and position verification with classical communication [LLQ22]. We begin by providing the definition.

Definition 6.3 (1-of-2 puzzle [RS19])**.** *A 1-of-2 puzzle consists of four algorithms with the following syntax.*

- $\mathsf{KeyGen}(1^\lambda) \to (\mathsf{pk}, \mathsf{vk})$*. The PPT key generation algorithm takes as input the security parameter* 1^λ *and outputs a public key* pk *and a secret verification key* vk*.*
- $\mathsf{Obligate}(\mathsf{pk}) \to (|\psi\rangle, y)$*: The QPT obligate algorithm takes as input the public key, and outputs a classical obligation string y and a quantum state $|\psi\rangle$.*
- $\mathsf{Solve}(|\psi\rangle, b) \to a$*: The QPT solve algorithm takes as input a state $|\psi\rangle$ and a bit $b \in \{0,1\}$ and outputs a string a.*
- $\mathsf{Ver}(\mathsf{vk}, y, b, a) \to \{\top, \bot\}$*: The PPT verify algorithm takes as input the verification key* vk*, a string y, a bit $b \in \{0,1\}$, and a string a, and either accepts or rejects.*

We say the puzzle is an $\epsilon(\lambda)$-1-of-2 puzzle if it satisfies the following properties.

- ***Completeness.***

$$\Pr\left[\top \leftarrow \mathsf{Ver}(\mathsf{vk}, y, b, a) : \begin{array}{r} (\mathsf{pk}, \mathsf{vk}) \leftarrow \mathsf{KeyGen}(1^\lambda) \\ (|\psi\rangle, y) \leftarrow \mathsf{Obligate}(\mathsf{pk}) \\ b \leftarrow \{0,1\} \\ a \leftarrow \mathsf{Solve}(|\psi\rangle, b) \end{array}\right] = 1 - \mathrm{negl}(\lambda).$$

- ***Soundness.*** *For any QPT adversary* $\{\mathsf{Adv}_\lambda\}_{\lambda\in\mathbb{N}}$*,*

$$\Pr\left[\top \leftarrow \mathsf{Ver}(\mathsf{vk}, y, 0, a_0) \wedge \top \leftarrow \mathsf{Ver}(\mathsf{vk}, y, 1, a_1) \ : \ \begin{array}{r} (\mathsf{pk}, \mathsf{vk}) \leftarrow \mathsf{KeyGen}(1^\lambda) \\ (y, a_0, a_1) \leftarrow \mathsf{Adv}_\lambda(\mathsf{pk}) \end{array}\right] \leq \epsilon(\lambda) + \mathrm{negl}(\lambda).$$

We call a 1-of-2 puzzle *strong* if $\epsilon = 0$, and note the following amplification theorem due to [RS19]. We will then construct an ϵ-1-of-2 puzzle for $\epsilon = 0.5$ from OSP, which gives a strong 1-of-2 puzzle as a corollary.

Theorem 6.4 ([RS19])**.** *For any $\epsilon = 1 - 1/\mathrm{poly}(\lambda)$, an ϵ-1-of-2 puzzle implies a strong 1-of-2 puzzle.*

Theorem 6.5. *Two-round OSP (Definition 4.1) implies an ϵ-1-of-2 puzzle (Definition 6.3) for $\epsilon = 0.5$.*

Proof. Let $(\mathsf{OSP.Sen}, \mathsf{OSP.Rec}, \mathsf{OSP.Dec})$ be any two-round OSP protocol (see Definition 4.1). We define the 1-of-2 puzzle as follows.

- $\mathsf{KeyGen}(1^\lambda)$: Sample $r \leftarrow \{0,1\}$ and for $i \in [\lambda]$, sample $(\mathsf{msg}_{S,i}, \mathsf{st}_{S,i}) \leftarrow \mathsf{OSP.Sen}(1^\lambda, r)$. Define $\mathsf{pk} := (\mathsf{msg}_{S,1}, \dots, \mathsf{msg}_{S,\lambda})$ and $\mathsf{vk} := (\mathsf{st}_{S,1}, \dots, \mathsf{st}_{S,\lambda})$.

- $\mathsf{Obligate}(\mathsf{pk})$: For each $i \in [\lambda]$, run $(|\psi_i\rangle, \mathsf{msg}_{R,i}) \leftarrow \mathsf{OSP.Rec}(\mathsf{msg}_{S,i})$. Define $|\psi\rangle := (|\psi_1\rangle, \dots, |\psi_\lambda\rangle)$ and $y := (\mathsf{msg}_{R,1}, \dots, \mathsf{msg}_{R,\lambda})$.
- $\mathsf{Solve}(|\psi\rangle, b)$: If $b = 0$, measure each $|\psi_{R,i}\rangle$ in the $X + Z$ basis and if $b = 1$, measure each $|\psi_{R,i}\rangle$ in the $X - Z$ basis. This results in a string $a \in \{0,1\}^\lambda$.
- $\mathsf{Ver}(\mathsf{vk}, y, b, a)$: For each $i \in [\lambda]$, set $s_i := \mathsf{OSP.Dec}(\mathsf{st}_{S_i}, \mathsf{msg}_{R,i})$, define $s := (s_1, \dots, s_\lambda)$, and define $s \oplus r := (s_1 \oplus r, \dots, s_\lambda \oplus r)$. If $b = 0$, accept iff $\Delta(a, s) \geq 0.85$ and if $b = 1$, accept iff $\Delta(a, s \oplus r) \geq 0.85$.

A standard calculation shows that for each $i \in [\lambda]$, $\Pr[a_i = s_i] = \cos^2(\pi/8) > 0.85$ if $b = 0$ and $\Pr[a_i = s_i \oplus r] = \cos^2(\pi/8) > 0.85$ if $b = 1$. Moreover, for each $b \in \{0,1\}$, these λ events are independent. Thus, by a standard tail bound, for each $b \in \{0,1\}$ we have that $\Pr[\mathsf{Ver}(\mathsf{vk}, y, b, a) = \top] = 1 - \mathrm{negl}(\lambda)$.

Now, suppose towards contradiction that there exists a QPT $\mathsf{Adv} = \{\mathsf{Adv}_\lambda\}_{\lambda \in \mathbb{N}}$ such that

$$\Pr\left[\begin{array}{c} \top \leftarrow \mathsf{Ver}(\mathsf{vk}, y, 0, a_0) \\ \wedge \top \leftarrow \mathsf{Ver}(\mathsf{vk}, y, 1, a_1) \end{array} : \begin{array}{c} (\mathsf{pk}, \mathsf{vk}) \leftarrow \mathsf{KeyGen}(1^\lambda) \\ (y, a_0, a_1) \leftarrow \mathsf{Adv}_\lambda(\mathsf{pk}) \end{array}\right] = \frac{1}{2} + \mathsf{non\text{-}negl}(\lambda).$$

This implies that with probability $1/2 + \mathsf{non\text{-}negl}(\lambda)$, the majority bit in $a_0 \oplus a_1$ is equal to r. However, by a standard hybrid argument, the security of the two-round OSP implies that r cannot be predicted with better than $\mathrm{negl}(\lambda)$ advantage, a contradiction.

□

Corollary 6.6. *Two-round OSP implies a strong 1-of-2 puzzle.*

In turn, we obtain the following results as corollaries from prior work.

- [RS19]: Two-round OSP implies privately-verifiable quantum money with classical communication.
- [LLQ22]: Two-round OSP implies position verification with classical communication.

6.3 Blind delegation

We first give a very general definition for blind delegation of quantum computation with a classical client. It allows the client to delegate the computation of an arbitrary publicly-known quantum operation that takes a quantum input (provided by the server, and potentially entangled with an auxiliary system held by the server) and a private classical input (chosen by the client). After interaction, the server is able to obtain the (potentially quantum) output up to a one-time pad with keys known to the client.

Definition 6.7 (Blind Classical Delegation of Quantum Computation)**.** *Let $\mathcal{H}_\mathcal{V}, \mathcal{H}_\mathcal{W}$ be Hilbert spaces of arbitrary dimension, and let $Q : \{0,1\}^* \times \mathcal{H}_\mathcal{V} \to \mathcal{H}_\mathcal{W}$ be a polynomial-size quantum circuit that takes as input a classical string x and a state on register $\mathcal{V}$, and outputs a state on register $\mathcal{W}$. A protocol for blind classical delegation of quantum computation consists of an interaction*

$$(\mathcal{W}, (r, s)) \leftarrow \langle S(1^\lambda, Q, \mathcal{V}), C(1^\lambda, Q, x)\rangle$$

between

- *a QPT server $S(1^\lambda, Q, \mathcal{V})$ with input the security parameter 1^λ, circuit Q, and state on register $\mathcal{V}$, and*
- *a PPT client $C(1^\lambda, Q, x)$ with input the security parameter 1^λ, circuit Q, and classical string x.*

At the end of the interaction, the server outputs a state on register $\mathcal{W}$ and the client outputs classical strings (r, s). The protocol must satisfy the following properties.

- ***Correctness.** For any circuit Q and input x, let $\mathsf{IDEAL}[Q, x]$ be the map from $\mathcal{V} \to \mathcal{W}$ defined by $Q(x, \cdot)$, and let $\mathsf{REAL}[Q, x]_\lambda$ be the map from $\mathcal{V} \to \mathcal{W}$ induced by running the protocol*

$$(\mathcal{W}, (r, s)) \leftarrow \langle S(1^\lambda, Q, \mathcal{V}), C(1^\lambda, Q, x)\rangle$$

and then applying $X^r Z^s$ to $\mathcal{W}$. Then for any Q and x,

$$D_\diamond \left(\mathsf{REAL}[Q, x]_\lambda, \mathsf{IDEAL}[Q, x]\right) = \text{negl}(\lambda).$$

- ***Security.** For any circuit Q, QPT adversary $\{\mathsf{Adv}_\lambda\}_{\lambda\in\mathbb{N}}$, and two strings x_0, x_1, it holds that*

$$\Bigg| \Pr\left[b_{\mathsf{Adv}} = 0 : (b_{\mathsf{Adv}}, (r, s)) \leftarrow \langle \mathsf{Adv}_\lambda, C(1^\lambda, Q, x_0)\rangle\right] - \Pr\left[b_{\mathsf{Adv}} = 0 : (b_{\mathsf{Adv}}, (r, s)) \leftarrow \langle \mathsf{Adv}_\lambda, C(1^\lambda, Q, x_1)\rangle\right] \Bigg| = \text{negl}(\lambda),$$

where b_{Adv} denotes a single bit output by Adv_λ after the interaction (which could result from an arbitrary QPT operation applied to its state after the interaction).

Remark 6.8. *Note that the above definition implies that if Q has a* classical *output, then the client can obtain this output with one extra message from the server. That is, suppose $Q : \{0,1\}^* \times \mathcal{H}_\mathcal{V} \to \{0,1\}^*$. Then at the conclusion of the protocol defined above, the server has a classical output $y \oplus r$, and the client has the one-time pad key r (note that s is irrelevant once the output has been measured in the standard basis). Then, the server returns $y \oplus r$ to the client, who recovers the output y. In this case, we denote the protocol*

$$y \leftarrow \langle S(1^\lambda, Q, \mathcal{V}), C(1^\lambda, Q, x)\rangle,$$

where y is the client's output.

Towards showing that OSP implies blind classical delegation of arbitrary quantum computation, we first show that it implies what we call an "encypted phase" protocol.

Definition 6.9 (Encrypted phase). *A protocol for encrypted phase is the special case of blind classical delegation of quantum computation where $\mathcal{H}_\mathcal{V} = \mathcal{H}_\mathcal{W}$ is a single-qubit register, the client's private input is a bit $b \in \{0,1\}$, and $Q(b, \mathcal{V})$ is the identity if $b = 0$ and applies a phase gate P to $\mathcal{V}$ if $b = 1$.*

Lemma 6.10. *OSP (Definition 4.1) implies encrypted phase.*

Proof. We first describe the protocol.

- The server and client begin by running an OSP protocol, where the client plays the role of the sender with bit b equal to the client's input bit b, and the server plays the role of the receiver. Then, the server applies a Hadamard gate followed by a $\sqrt{X}$ gate to their output state. Up to a negligible trace distance, this results in the state

  - $Z^s\,|+\rangle$ if $b=0$,
  - $Z^sP\,|+\rangle$ if $b=1$,

  where s is the sender's output bit. Let $\mathcal{M}$ be the name of the output state's register.

- Next, the server applies a CNOT gate from register $\mathcal{V}$ to register $\mathcal{M}$, and then measures $\mathcal{M}$ in the standard basis to obtain a bit m. The server sends m to the client, and outputs $\mathcal{V}$.

- If $b=0$, the client outputs $(0,s)$ and if $b=1$, the client outputs $(0,s\oplus m)$.

Security follows immediately from the security of OSP, so it remains to check correctness. Let $\alpha_0\,|0\rangle_{\mathcal{V}}+\alpha_1\,|1\rangle_{\mathcal{V}}$ be the input state on register $\mathcal{V}$ (note that it could be entangled with some auxiliary register, i.e. $\alpha_0\,|0\rangle_{\mathcal{V}}\,|\psi_0\rangle+\alpha_1\,|1\rangle_{\mathcal{V}}\,|\psi_1\rangle$, but we suppress writing the auxiliary register to avoid clutter).

In the case $b=0$, we have that

$$\begin{aligned}&\mathsf{CNOT}\,(\alpha_0\,|0\rangle_{\mathcal{V}}+\alpha_1\,|1\rangle_{\mathcal{V}})\otimes Z^s\,|+\rangle_{\mathcal{M}}\\&\quad=(Z^s_{\mathcal{V}}\otimes Z^s_{\mathcal{M}})\mathsf{CNOT}\,(\alpha_0\,|0\rangle_{\mathcal{V}}+\alpha_1\,|1\rangle_{\mathcal{V}})\otimes|+\rangle_{\mathcal{M}}\\&\quad=(Z^s_{\mathcal{V}}\otimes Z^s_{\mathcal{M}})\,(\alpha_0\,|0\rangle_{\mathcal{V}}+\alpha_1\,|1\rangle_{\mathcal{V}})\otimes|+\rangle_{\mathcal{M}}\,,\end{aligned}$$

so measuring $\mathcal{M}$ does not affect the state on $\mathcal{V}$. Thus, the server is left with their original state up to a Z^s error, which is the desired outcome.

In the case $b=1$, we have that

$$\begin{aligned}&\mathsf{CNOT}\,(\alpha_0\,|0\rangle_{\mathcal{V}}+\alpha_1\,|1\rangle_{\mathcal{V}})\otimes PZ^s\,|+\rangle_{\mathcal{M}}\\&\quad=\frac{1}{\sqrt{2}}(Z^s_{\mathcal{V}}\otimes Z^s_{\mathcal{M}})\mathsf{CNOT}\,(\alpha_0\,|0\rangle_{\mathcal{V}}+\alpha_1\,|1\rangle_{\mathcal{V}})\otimes(|0\rangle_{\mathcal{M}}+i\,|1\rangle_{\mathcal{M}})\\&\quad=\frac{1}{\sqrt{2}}(Z^s_{\mathcal{V}}\otimes Z^s_{\mathcal{M}})\,(\alpha_0\,|0\rangle_{\mathcal{V}}\,|0\rangle_{\mathcal{M}}+i\alpha_0\,|0\rangle_{\mathcal{V}}\,|1\rangle_{\mathcal{M}}+\alpha_1\,|1\rangle_{\mathcal{V}}\,|1\rangle_{\mathcal{M}}+i\alpha_1\,|1\rangle_{\mathcal{V}}\,|0\rangle_{\mathcal{M}})\\&\quad=\frac{1}{\sqrt{2}}(Z^s_{\mathcal{V}}\otimes Z^s_{\mathcal{M}})\,((\alpha_0\,|0\rangle_{\mathcal{V}}+i\alpha_1\,|1\rangle_{\mathcal{V}})\,|0\rangle_{\mathcal{M}}+(i\alpha_0\,|0\rangle_{\mathcal{V}}+\alpha_1\,|1\rangle_{\mathcal{V}})\,|1\rangle_{\mathcal{M}})\,.\end{aligned}$$

So if the measured bit $m=0$, the remaining state is

$$Z^s_{\mathcal{V}}\,(\alpha_0\,|0\rangle_{\mathcal{V}}+i\alpha_1\,|1\rangle_{\mathcal{V}})\,,$$

which is the desired outcome, and if the measured bit $m=1$, the remaining state is

$$Z^s_{\mathcal{V}}\,(i\alpha_0\,|0\rangle_{\mathcal{V}}+\alpha_1\,|1\rangle_{\mathcal{V}})=Z^s_{\mathcal{V}}\,(\alpha_0\,|0\rangle_{\mathcal{V}}-i\alpha_1\,|1\rangle_{\mathcal{V}})=Z^{s\oplus 1}_{\mathcal{V}}\,(\alpha_0\,|0\rangle_{\mathcal{V}}+i\alpha_1\,|1\rangle_{\mathcal{V}})\,,$$

which is again the desired outcome.

□

Next, we show that the ability to perform an encrypted phase implies blind classical delegation of quantum computation for arbitrary quantum operations.

Theorem 6.11. *Any protocol for encrypted phase (Definition 6.9) implies blind classical delegation of quantum computation (Definition 6.7).*

Proof. Given any circuit Q, write it using Clifford and $T^\dagger$ gates, where $T^\dagger = \sqrt{P^\dagger}$. That is, Q can be written as $C_{\ell+1}T^\dagger C_\ell T^\dagger \dots C_2 T^\dagger C_1$, where each C_i is Clifford. We will use the fact that for any polynomial-size Clifford C and one-time padded state $X^r Z^s |\psi\rangle$, it holds that $CX^rZ^s|\psi\rangle = X^{r'}Z^{s'}C|\psi\rangle$, where r' and s' are efficiently computable from C, r, and s. Moreover, for a single qubit state $|\psi\rangle$, it holds that $T^\dagger X^r Z^s |\psi\rangle = (P^\dagger)^r X^r Z^s T^\dagger |\psi\rangle$.

Given a protocol for encrypted phase (Definition 6.9), we describe a protocol for blind classical delegation of quantum computation:

- Given inputs $(1^\lambda, Q, x)$, the client writes $Q = C_{\ell+1}T^\dagger C_\ell T^\dagger \dots C_2 T^\dagger C_1$, where each C_i is Clifford and the i'th $T^\dagger$ gate is applied to qubit t_i of the server's register $\mathcal{V}$.
- The client samples a classical one-time pad r_{inp} and sends $x \oplus r_{\mathsf{inp}}$ to the server. Throughout the protocol, the client will hold the classical keys for a quantum one-time pad applied to the server's register $\mathcal{V}$ (which we consider to now include the encrypted input $x \oplus r_{\mathsf{inp}}$). The client initializes these keys to $(r, s) \coloneqq ((r_{\mathsf{inp}}, 0, \dots, 0), (0, \dots, 0))$.
- For $i \in [\ell]$, perform the following steps.
  - The server applies C_i to register $\mathcal{V}$, and the client updates the quantum one-time pad keys (r, s) according to C_i.
  - The server applies a $T^\dagger$ gate to qubit t_i.
  - Let b_i be the X-bit of the one-time pad key on qubit t_i. The server and client apply an encrypted phase (Definition 6.9) to qubit t_i with client input equal to b_i.
  - The client uses their output (r_i, s_i) from the encrypted phase to update the one-time pad key on qubit t_i.
- The server applies the final Clifford $C_{\ell+1}$ to register $\mathcal{V}$ and outputs $\mathcal{V}$, and the client does the final one-time pad update according to $C_{\ell+1}$ and outputs their final one-time pad keys $(r_{\mathsf{out}}, s_{\mathsf{out}})$.

Correctness is straightforward using the properties listed prior to the description of the protocol. Security follows from a standard hybrid argument: Starting from the final $T^\dagger$ gate, we switch the client's input to the encrypted phase to 0. Once this is completed, the protocol no longer depends on the client's initial classical one-time pad r_{inp}, and thus we can switch between client inputs x_0 and x_1 without affecting the view of the server. □

Thus, we obtain the following corollary.

Corollary 6.12. *OSP implies blind classical delegation of quantum computation (Definition 6.7).*

In fact, we also observe the following, which shows that OSP is *both* necessary and sufficient for blind classical delegation of quantum computation.

Claim 6.13. *Blind classical delegation of quantum computation according to Definition 6.7 implies OSP.*

Proof. Let Q be the circuit that takes as input a bit b (and no quantum input) and outputs the state $H^b|0\rangle$. To perform OSP, the sender inputs their bit b to a classical delegation protocol for Q, with the receiver acting as the server. The receiver ends up with a state (negligibly close to) $X^rZ^sH^b|0\rangle = H^b|x\rangle$, where $x = r$ if $b = 0$ and $x = s$ if $b = 1$, and the sender ends up with (r, s). Thus, the sender outputs $x = r$ if $b = 0$ and $x = s$ if $b = 1$ to complete the description of the OSP protocol. □

6.4 Verifiable delegation

In this section, we show that *blind* classical delegation of quantum computation (Definition 6.7) generically implies *verifiable* classical delegation of quantum computation, which then gives verifiable delegation from OSP as a corollary. Previously, it was shown by [NZ23] that QFHE, which is a special case of blind delegation, implies verifiable delegation. Here, we observe that their approach, which builds on the "KLVY compiler" of [KLVY23] can be generalized to establish the result from any blind delegation. We begin by defining verifiable delegation, which we call classical verification of quantum computation (CVQC).

Definition 6.14 (Classical Verification of Quantum Computation)**.** *Classical verification of quantum computation (CVQC) is an interaction between a QPT prover and PPT verifier on input an instance x*

$$\{\top, \bot\} \leftarrow \langle P(1^\lambda, x), V(1^\lambda, x)\rangle,$$

where $\{\top, \bot\}$ is the verifier's output. For any language $\mathcal{L}$ in BQP, there exists some $\epsilon(\lambda), \delta(\lambda)$ with $\epsilon(\lambda) - \delta(\lambda) = 1/\mathrm{poly}(\lambda)$ such that the following properties hold.

- ***Completeness.*** *For any $x \in \mathcal{L}$,*

$$\Pr\Big[\top \leftarrow \langle P(1^\lambda, x), V(1^\lambda, x)\rangle\Big] \geq \epsilon(\lambda).$$

- ***Soundness.*** *For any $x \notin \mathcal{L}$ and QPT adversary $\{\mathsf{Adv}_\lambda\}_{\lambda\in\mathbb{N}}$,*

$$\Pr\left[\top \leftarrow \langle \mathsf{Adv}_\lambda, V(1^\lambda, x)\right] \leq \delta(\lambda) + \mathrm{negl}(\lambda).$$

This section is dedicated to proving the following theorem.

Theorem 6.15. *Blind classical delegation of quantum computation (Definition 6.7) implies classical verification of quantum computation (Definition 6.14).*

Corollary 6.16. *OSP implies classical verification of quantum computation.*

6.4.1 Nonlocal games and the KLVY compiler

Towards proving this theorem, we recall the KLVY compiler, which uses QFHE to compile any nonlocal game into a single-prover protocol. In this subsection, we observe that the KLVY compiler can be instantiated with any (potentially interactive, non-compact) blind delegation protocol, which yields what we call the *generalized KLVY compiler*. We also define a set of nonlocal game strategies that we call *computationally nonlocal strategies*. This definition provides a clean way to establish soundness of the generalized KLVY compiler - soundness of any compiled game can be upper-bounded by the value of the best computationally nonlocal strategy for that game.

First, we present standard definitions of (families of) nonlocal games, as well as nonlocal strategies for these games.

Definition 6.17 (Nonlocal game)**.** *A* nonlocal game *$G = (Q, V)$ is specified by a distribution Q over pairs $(x, y) \in \{0,1\}^{n_1} \times \{0,1\}^{n_2}$ and a verification predicate $V(x, y, a, b) \in \{0,1\}$, where $a \in \{0,1\}^{m_1}$ and $b \in \{0,1\}^{m_2}$. A* family *of nonlocal games $\mathcal{G} = \{\mathcal{G}_\lambda\}_{\lambda\in\mathbb{N}}$ is a set of games parameterized by λ, where* *each $\mathcal{G}_\lambda$ is itself a set of games $G \in \mathcal{G}_\lambda$. Each game $G \in \mathcal{G}$ is defined by a distribution Q_G over pairs $(x, y) \in \{0,1\}^{n_{1,G}} \times \{0,1\}^{n_{2,G}}$ and a verification predicate $V_G(x, y, a, b) \in \{0,1\}$, where $a \in \{0,1\}^{m_{1,G}}$ and $b \in \{0,1\}^{m_{2,G}}$. For any game $G \in \mathcal{G}$, we define λ_G to be the choice of λ such that $G \in \mathcal{G}_\lambda$. We say that the family of games is* efficient *if there exists a polynomial $p(\cdot)$ and a procedure that, for any $G \in \mathcal{G}$, samples from Q_G and computes V_G in time $p(\lambda_G)$.*

Definition 6.18 (Nonlocal strategy)**.** *A* nonlocal strategy *$\mathscr{S}$ for game G consists of the following.*

- *A state $|\psi\rangle \in \mathcal{H}_\mathcal{A} \otimes \mathcal{H}_\mathcal{B}$.*
- *For every $x \in \{0,1\}^{n_1}$, a projective measurement $\{A^x_a\}_a$ acting on $\mathcal{H}_\mathcal{A}$ with outcomes $a \in \{0,1\}^{m_1}$.*
- *For every $y \in \{0,1\}^{n_2}$, a projective measurement $\{B^y_b\}_b$ acting on $\mathcal{H}_\mathcal{B}$ with outcomes $b \in \{0,1\}^{m_2}$.*

The value *of this strategy is given by*

$$\omega(G, \mathscr{S}) \coloneqq \mathop{\mathbb{E}}_{(x,y)\leftarrow Q} \sum_{a,b} V(x, y, a, b) \langle\psi| A^x_a \otimes B^y_b |\psi\rangle .$$

A strategy $\mathscr{S}$ for a family *of games $\mathcal{G}$ consists of a strategy*

$$\mathscr{S}_G = \left(|\psi_G\rangle , \{A^x_{a,G}\}_a, \{B^y_{b,G}\}_b \right)$$

for each $G \in \mathcal{G}$. We say that $\mathscr{S}$ is efficient *if $|\psi\rangle$ is QPT-preparable and $\{A^x_{a,G}\}_a$ and $\{B^y_{b,G}\}_b$ are QPT-implementable.*

Next, we present our new definition of a *computationally nonlocal strategy*.

Definition 6.19 (Computationally nonlocal strategy)**.** *A* computationally nonlocal strategy *$\mathscr{C}$ for a family of games $\mathcal{G} = \{\mathcal{G}_\lambda\}_\lambda$ consists of the following for each $G \in \mathcal{G}$.*

- *A QPT-preparable state $|\psi_G\rangle \in \mathcal{H}_{\mathcal{A},G} \otimes \mathcal{H}_{\mathcal{B},G}$.*
- *For every $x \in \{0,1\}^{n_{1,G}}$, a QPT-implementable unitary U^x_G acting on $\mathcal{H}_{\mathcal{A},G} \otimes \mathcal{H}_{\mathcal{B},G}$. For each $a \in \{0,1\}^{m_{1,G}}$, define*
$$A^x_{a,G} \coloneqq |a\rangle\langle a| \, U^x_G,$$
where the projection $|a\rangle\langle a|$ is applied to some specified sub-register of $\mathcal{H}_{\mathcal{A},G}$.
- *For every $y \in \{0,1\}^{n_{2,G}}$, a QPT-implementable projective measurement $\{B^y_{a,G}\}_b$ acting on $\mathcal{H}_{\mathcal{B},G}$ with outcomes $b \in \{0,1\}^{m_{2,G}}$.*

For this to be a valid computationally nonlocal strategy, the "Alice" operations must satisfy the following property. There exists a negligible function $\mu(\lambda)$ such that for any sequence of games $\{G_\lambda\}_{\lambda\in\mathbb{N}}$ where each $G_\lambda \in \mathcal{G}_\lambda$, QPT distinguisher $\{M_\lambda, I - M_\lambda\}_\lambda$ ***acting only on*** *$\mathcal{H}_{\mathcal{B},G_\lambda}$, and sequence of pairs of questions $\{x_{0,\lambda}, x_{1,\lambda} \in \{0,1\}^{n_{1,G_\lambda}}\}_\lambda$,*

$$\left| \sum_a \langle\psi_{G_\lambda}| A^{x_{0,\lambda},\dagger}_{a,G_\lambda} M_\lambda A^{x_{0,\lambda}}_{a,G_\lambda} |\psi_{G_\lambda}\rangle - \sum_a \langle\psi_{G_\lambda}| A^{x_{1,\lambda},\dagger}_{a,G_\lambda} M_\lambda A^{x_{1,\lambda}}_{a,G_\lambda} |\psi_{G_\lambda}\rangle \right| \le \mu(\lambda).$$

That is, Alice's questions must be computationally hidden by the state passed from the Alice operation to the Bob operation. The value *of this strategy is a function $\omega(G, \mathscr{C})$ of the game $G \in \mathcal{G}$, defined by*

$$\omega(G, \mathscr{C}) := \mathop{\mathbb{E}}_{(x,y) \leftarrow Q_G} \sum_{a,b} V_G(x, y, a, b) \langle \psi_G | A^{x,\dagger}_{a,G} B^y_{b,G} A^x_{a,G} | \psi_G \rangle .$$

We say that the computationally nonlocal value of $\mathcal{G}$ is upper-bounded by a function $\omega_C(G)$ if for all computationally nonlocal strategies $\mathscr{C}$, there exists a negligible function $\nu(\lambda)$ such that for any sequence of games $\{G_\lambda\}_{\lambda \in \mathbb{N}}$ where each $G_\lambda \in \mathcal{G}_\lambda$, it holds that

$$\omega(G_\lambda, \mathscr{C}) \leq \omega_C(G_\lambda) + \nu(\lambda).$$

Remark 6.20. *We will sometimes refer to a computationally nonlocal strategy $\mathscr{C}$ for some* fixed *game G (such as the CHSH game). In this case, we view G as a family of games $\mathcal{G} = \{\mathcal{G}_\lambda\}_\lambda = \{G\}_\lambda$. That is, each $\mathcal{G}_\lambda$ just consists of G itself, and $\mathscr{C}$ consists of one strategy for each $\lambda \in \mathbb{N}$.*

The following remark formalizes a simple claim about computational indistinguishability.

Remark 6.21. *By following the arguments in [NZ23, Lemma 7, Lemma 8], it is straightforward to see that the Alice operations for any computationally nonlocal strategy $\mathscr{C}$ must also satisfy the following. There exists a negligible function $\mu(\lambda)$ such that for any sequence of games $\{G_\lambda\}_{\lambda \in \mathbb{N}}$ where each $G_\lambda \in \mathcal{G}_\lambda$, QPT-implementable POVM $\{\{M_{\gamma,\lambda}\}_\gamma\}_\lambda$ with outcomes in $\gamma \in [0, 1]$ acting only on $\mathcal{H}_{\mathcal{B},G_\lambda}$, and sequence of pairs of questions $\{x_{0,\lambda}, x_{1,\lambda} \in \{0, 1\}^{n_{1,G_\lambda}}\}_\lambda$,*

$$\left| \sum_a \sum_\gamma \gamma \langle \psi_{G_\lambda} | A^{x_{0,\lambda},\dagger}_{a,G_\lambda} M_{\gamma,\lambda} A^{x_{0,\lambda}}_{a,G_\lambda} | \psi_{G_\lambda} \rangle - \sum_a \sum_\gamma \gamma \langle \psi_{G_\lambda} | A^{x_{1,\lambda},\dagger}_{a,G_\lambda} M_{\gamma,\lambda} A^{x_{1,\lambda}}_{a,G_\lambda} | \psi_{G_\lambda} \rangle \right| \leq \mu(\lambda).$$

Next, we state the "generalized" version of the KLVY compiler, where in place of a QFHE protocol, we use any blind classical delegation of quantum computation protocol (Definition 6.7).

Definition 6.22 (Generalized KLVY Compiler)**.** *Let $\mathcal{G}$ be a family of nonlocal games, let $\Pi = \langle S, C \rangle$ be a blind classical delegation of quantum computation protocol, and let $\mathscr{S}$ be an efficient nonlocal strategy for $\mathcal{G}$. For each $G \in \mathcal{G}$, let $A_G : \{0, 1\}^{n_{1,G}} \times \mathcal{H}_\mathcal{A} \to \{0, 1\}^{m_{1,G}}$ be the QPT circuit that performs the Alice measurement of $\mathscr{S}_G$ and $B_G : \{0, 1\}^{n_{2,G}} \times \mathcal{H}_\mathcal{B} \to \{0, 1\}^{m_{2,G}}$ be the QPT circuit that performs the Bob measurement of $\mathscr{S}_G$. The* KLVY-compiled *protocol* $\mathsf{KLVY}[\mathcal{G}, \Pi, \mathscr{S}]$ *is an interaction between a QPT prover* Prove *and a PPT verifier* Ver *that is parameterized by a game $G \in \mathcal{G}_\lambda \subset \mathcal{G}$, and operates as follows.*

1. *The verifier samples $(x, y) \leftarrow Q_G$.*

2. *Let $|\psi_G\rangle \in \mathcal{H}_\mathcal{A} \otimes \mathcal{H}_\mathcal{B}$ be the initial state used in $\mathscr{S}_G$. The prover and verifier engage in a blind classical delegation of quantum computation protocol (with classical output, see Remark 6.8)*

$$a \leftarrow \langle S(1^\lambda, A_G, |\psi_G\rangle), C(1^\lambda, A_G, x) \rangle,$$

   with the prover playing the role of the server S and the verifier playing the role of the client C.

3. *The verifier sends y to the prover.*

4. *The prover runs $b \leftarrow B_G(y, \mathcal{B})$ and sends b to the verifier.*

5. *The verifier outputs* $V_G(x, y, a, b)$.

This interaction is denoted

$$b_{\mathsf{Ver}} \leftarrow \langle \mathsf{Prove}(1^\lambda, G), \mathsf{Ver}(1^\lambda, G)\rangle,$$

where b_{Ver} *denotes the bit output by* Ver.

The completeness *of* $\mathsf{KLVY}[\mathcal{G}, \Pi, \mathscr{S}]$ *is defined by a function*

$$c(G) \coloneqq \Pr\left[b_{\mathsf{Ver}} = 1 : b_{\mathsf{Ver}} \leftarrow \langle \mathsf{Prove}(1^\lambda, G), \mathsf{Ver}(1^\lambda, G)\rangle\right],$$

and $\mathsf{KLVY}[\mathcal{G}, \Pi, \mathscr{S}]$ *has soundness* $s(G)$ *if for any QPT adversary* $\{\mathsf{Adv}_\lambda\}_{\lambda\in\mathbb{N}}$

$$\Pr\left[b_{\mathsf{Ver}} = 1 : b_{\mathsf{Ver}} \leftarrow \langle \mathsf{Adv}_{\lambda_G}(1^\lambda, G), \mathsf{Ver}(1^\lambda, G)\rangle\right] \leq s(G).$$

Finally, we prove the main theorem of this subsection. Essentially, we show that the computational nonlocal value of any game G upper bounds the soundness of the (generalized) KLVY compiled version of G.

Theorem 6.23. *Let* $\mathcal{G}$ *be a family of nonlocal games, let* Π *be a blind classical delegation of quantum computation protocol, and let* $\mathscr{S}$ *be an efficient nonlocal strategy for* $\mathcal{G}$. *Then the completeness of* $\mathsf{KLVY}[\mathcal{G}, \Pi, \mathscr{S}]$ *satisfies*

$$c(G) \geq \omega(G, \mathscr{S}) - \mathrm{negl}(\lambda_G),$$

and, for any upper bound $\omega_C(G)$ *on the computationally nonlocal value of* $\mathcal{G}$, $\mathsf{KLVY}[\mathcal{G}, \Pi, \mathscr{S}]$ *has soundness*

$$s(G) \leq \omega_C(G) + \mathrm{negl}(\lambda_G),$$

where λ_G *is defined to be the* λ *such that* $G \in \mathcal{G}_\lambda$.

Proof. The completeness claim follows directly from the correctness of the blind delegation protocol Π. The soundness claim follows by observing that any Adv interacting in $\mathsf{KLVY}[\mathcal{G}, \Pi, \mathscr{S}]$ defines a computationally nonlocal strategy for $\mathcal{G}$. This can be argued as follows, where for convenience we will drop parameterizations by G and λ. Consider purifying the interaction between Adv and Ver during the second step of the protocol (the blind delegation step), and then measuring the verifier's output a. For each verifier input x, this defines a unitary U^x that is applied to initial state $|0\rangle_{\mathcal{A}} |\psi\rangle_{\mathcal{B}}$, where $|\psi\rangle$ is the initial state of Adv, $\mathcal{H}_{\mathcal{B}}$ includes the working register of Adv along with the register that holds the transcript of communication between Adv and Ver, and $\mathcal{H}_{\mathcal{A}}$ is the register required to implement the operations of Ver. Then A^x_a is defined as $|a\rangle\langle a| U^x$, i.e. the unitary U^x followed by a standard basis projection onto the verifier's output a on a sub-register of $\mathcal{H}_{\mathcal{A}}$. It follows immediately from the security of Π that there exists a negligible function $\mu(\lambda)$ such that for any QPT distinguisher $\{M, I - M\}$ acting only on $\mathcal{H}_{\mathcal{B}}$, and $x_0, x_1 \in \{0,1\}^{n_1}$,

$$\left|\sum_a \langle\psi| \langle 0| A^{x_0,\dagger}_a M A^{x_0}_a |\psi\rangle |0\rangle - \sum_a \langle\psi| \langle 0| A^{x_1,\dagger}_a M A^{x_1}_a |\psi\rangle |0\rangle\right| = \mathrm{negl}(\lambda),$$

which shows that this is a valid computationally nonlocal strategy.

□

6.4.2 The CHSH game

Next, we recall the CHSH game, which is an important building block in the CVQC protocol of [NZ23].

Definition 6.24 (CHSH)**.** *The CHSH game is defined by question distribution $(x, y) \leftarrow \{0, 1\} \times \{0, 1\}$ and predicate $V(x, y, a, b)$ that accepts iff $x \cdot y = a \oplus b$, where answers $a, b \in \{0, 1\}$.*

An important lemma from [NZ23] establishes a *rigidity* property of the KLVY-compiled CHSH game. That is, any strategy in the (QFHE-based) KLVY-compiled game that approaches the optimal value of $\cos^2(\pi/8)$ for nonlocal strategies must be such that the Bob operations nearly anticommute. It is not hard to see that the same holds for *any* computationally nonlocal strategy,[19] and in fact this claim was already essentially shown by [BGKM$^+$23, Theorem 4.7].

Theorem 6.25. *Let $\omega_{\mathsf{CHSH}} = \cos^2(\pi/8)$. Fix any computationally nonlocal strategy for the CHSH game with value $\omega_{\mathsf{CHSH}} - \delta(\lambda)$, and, dropping the parameterization by λ, let $B^0 := B^0_0 - B^0_1, B^1 := B^1_0 - B^1_1$ be binary observables defined by the "Bob" measurements. Then for any $x \in \{0, 1\}$, it holds that*

$$\sum_{a \in \{0,1\}} \langle\psi| A^{x,\dagger}_a \{B^0, B^1\}^2 A^x_a |\psi\rangle \leq O(\delta(\lambda)) + \mathrm{negl}(\lambda).$$

Proof. This follows readily from [BGKM$^+$23, Theorem 4.7], following their arguments in Section 5.4.[20] It can also be verified that a proper adaptation of Lemma 34 from [NZ23] yields this claim using a different proof technique. □

6.4.3 The [NZ23] BQP verification game

Now, we define the [NZ23] family of nonlocal games for verifying arbitrary BQP computation.[21] The family $\mathcal{G}$ consists of a game $G[H]$ for each XX/ZZ local Hamiltonian

$$H = \sum_{W \in \{X,Z\}, i \neq j \in [\lambda]} p_{W,i,j} W(e_i + e_j),$$

where $\sum_{W,i,j} p_{W,i,j} = 1$. Before coming to the formal specification of the nonlocal game $G[H]$ associated with H, we provide a high-level overview.

To begin with, two provers Alice and Bob share several halves of EPR pairs, and Alice prepares a minimum eigenvalue state $|\psi\rangle$ for H. Alice's question instructs her to either (1) **Teleport** the state to Bob, (2) participate in an anti-commutation (**CHSH**) game, or (3) participate in a **Commutation** game. Bob's question is always a single bit that instructs him to either measures all of his halves of EPR pairs in the standard basis or the Hadamard basis, and return the results.

In the teleportation case, an honest Alice simply teleports $|\psi\rangle$ to Bob and returns the teleportation errors to the verifier. The verifier now either samples an XX or a ZZ term of the Hamiltonian

[19] Note that strategies in the QFHE-based KLVY-compiled game are a special case of computationally nonlocal strategies.

[20] Our formalization of computationally nonlocal strategies is slightly more general than [BGKM$^+$23]'s quantum device formalization (since a computationally nonlocal strategy does not necessarily have to be defined by some prover interacting in a protocol), but it is straightforward to verify that their techniques apply to our more general setting.

[21] We note that the same set of games should suffice for verifying arbitrary QMA languages as well, by providing sufficient copies of the witness state to the prover, but we follow [NZ23] and stick to BQP for this exposition.

to match Bob's question, and can then recover the outcome of measuring $|\psi\rangle$ with that term by combining Bob's answers with Alice's teleportation errors. In the case that Alice and Bob are honest, repeating this round multiple times suffices to estimate the minimum eigenvalue of H. In fact, as long as Bob's operations are (nearly) isometric to a tensor of Z observables when his question asks for standard basis measurements and a tensor of X observables when his question asks for Hadamard basis measurements, we can conclude that the parties cannot on average convince the verifier that H has a significantly lower eigenvalue that it really has. The purpose of the CHSH and Commutation questions is to enforce this structure on Bob's operations. Combined, they yield a variant of the *Pauli braiding test* [NV17, Gri20].

Now, for the purpose of CVQC, we care about the value of any *computationally* nonlocal strategy for this family of games. First, it is imperative that if the strategy passes the CHSH test, then we can conclude that Bob's operations anti-commute. This indeed follows in the computationally nonlocal setting, as discussed in Section 6.4.2. It turns out that the only other crucial property is that Bob cannot change his strategy based on whether Alice received a Teleport, CHSH, or Commutation question. This clearly holds for any computationally nonlocal strategy, not just those that arise from the use of QFHE. While these are the main ideas, we now give a precise description of the [NZ23] family of games that can be run through the generalized KLVY compiler to yield CVQC from any blind classical delegation of quantum computation protocol.

The game $H[G]$. Define $\mathcal{G}_\lambda$ to be the set of games $G[H]$ where H is a λ-qubit Hamiltonian. Let $\beta = \beta(\lambda), \alpha = \alpha(\lambda) \in [-1, 1]$ be functions of the security parameter, and let $\kappa = \Theta((\beta - \alpha)^2)$ be a parameter to be set later. The game $G[H]$ will allow us to decide whether the minimum eigenvalue of H is smaller than α or larger than β.

Given H, define D_X to be the renormalized distribution over $X(e_i + e_j)$ operators induced by H, and define D_Z analogously. Define D_Q to be $U_\lambda \times D_X$, where U_λ is the uniform distribution over $\{0, 1\}^\lambda$. Then for any H, the game $G[H]$ is defined as follows.

- $Q_{G[H]}$: Sample the Alice question q_A and Bob question q_B by choosing one of the following options, where the first two are chosen with probability $(1 - \kappa)/2$ and the last one with probability κ.
  - **CHSH.** Sample $(a, b) \leftarrow U_\lambda \times D_X$ conditioned on $a \cdot b = 1$ (later, we will refer to this distribution as D_Q^1). Sample $x, y \leftarrow \{0, 1\}$, and set $q_A = (\text{CHSH}, (a, b, x))$ and $q_B = y$.
  - **Commutation.** Sample $(a, b) \leftarrow U_\lambda \times D_X$ conditioned on $a \cdot b = 0$ (later, we will refer to this distribution as D_Q^0). Sample $y \leftarrow \{0, 1\}$, and set $q_A = (\text{Commutation}, (a, b))$ and $q_B = y$.
  - **Teleport.** Sample $y \leftarrow \{0, 1\}$ and set $q_A = \text{Teleport}$ and $q_B = y$.
- $V_{G[H]}$: Receive Alice answer s_A (of varying size) and Bob answer $s_B \in \{0, 1\}^\lambda$. Compute the following, determined by the type of Alice question.
  - **CHSH.** Let $z := (1 - y)(a \cdot s_B) + y(b \cdot s_B)$, and output 1 iff $s_A + z = x \cdot y$.
  - **Commutation.** Let $z := (1 - y)(a \cdot s_B) + y(b \cdot s_B)$, and output 1 iff $(s_A)_y = z$.
  - **Teleport.** Sample w such that $w = 0$ with probability $\sum_{i,j} p_{X,i,j}$ and $w = 1$ with probability $\sum_{i,j} p_{Z,i,j}$. If $w \neq q_B$ then output 1. Otherwise, sample a term $W(e_i + e_j)$ from

the distribution induced by $p_{W,i,j}$, where $W = X$ if $w = 0$ and $W = Z$ if $w = 1$. Parse $s_A = (z, x)$, where $z, x \in \{0,1\}^\lambda$, and compute the following.

* If $W = Z$, output 1 iff $(-1)^{(s_B)_i + (s_B)_j + (s_A)_i + (s_A)_j} = -1$.
* If $W = X$, output 1 iff $(-1)^{(s_B)_i + (s_B)_j + (s_A)_{\lambda+i} + (s_A)_{\lambda+j}} = -1$.

Finally, we obtain a classical verification protocol for BQP (Definition 6.14) from blind classical delegation of quantum computation by combining the following theorem with Theorem 6.23, thus proving Theorem 6.15.

Theorem 6.26 (Adaptation of Theorem 46 from [NZ23])**.** *Let $\mathcal{G}[\mathsf{YES}]$ be the family of games $\mathcal{G}$ restricted to those defined by a Hamiltonian H with lowest eigenvalue at most α. There exists an efficient nonlocal strategy $\mathscr{S}$ with value*

$$\omega(G, \mathscr{S}) \geq \frac{1}{2}(1-\kappa)(1+\omega_{\mathsf{CHSH}}) + \kappa(1 - \frac{1}{4}\alpha)$$

for all $G \in \mathcal{G}[\mathsf{YES}]$. Next, let $\mathcal{G}[\mathsf{NO}]$ be the family of games $\mathcal{G}$ restricted to those defined by a Hamiltonian H with lowest eigenvalue at least β. Then there exists a choice of $\kappa = \Theta((\beta-\alpha)^2)$ such that the computationally nonlocal value of $\mathcal{G}[\mathsf{NO}]$ is upper-bounded by

$$\omega_C(G) = \frac{1}{2}(1-\kappa)(1+\omega_{\mathsf{CHSH}}) + \kappa(1 - \frac{1}{4}\alpha) - \frac{1}{8}\kappa(\beta - \alpha).$$

The proof of this theorem follows readily from arguments made in [NZ23]. For completeness, we give an overview of their main claims written in our notation in Appendix A.

6.5 Encrypted CNOT and applications

[Mah18a] informally introduced the notion of a (two-round) “encrypted CNOT” protocol as a sub-routine in her construction of quantum fully-homomorphic encryption. Here, we define encrypted CNOT formally as a special case of blind classical delegation of quantum computation. In particular, this means that it follows generically from any OSP, due to the results from Section 6.3.

However, we also provide a simple and direct protocol for encrypted CNOT which in particular implies that *two-round* encrypted CNOT follows from any two-round OSP. Finally, we discuss applications of encrypted CNOT to quantum fully-hommorphic encryption and claw-state generators with *indistinguishability* security (Definition 4.4).

Definition 6.27 (Encrypted CNOT)**.** *An encrypted CNOT protocol is the special case of blind classical delegation of quantum computation (Definition 6.7) where $\mathcal{H}_\mathcal{V} = \mathcal{H}_\mathcal{W}$ is a two-qubit register, the client's private input is a bit $b \in \{0,1\}$, and $Q(b, (\mathcal{V}_0, \mathcal{V}_1))$ is the identity if $b = 0$ and applies a CNOT gate from $\mathcal{V}_0$ to $\mathcal{V}_1$ if $b = 1$. We say that the protocol is* two-round *if it just consists of one message from the client followed by one message from the server:*

- $\mathsf{E\text{-}CNOT.Gen}(1^\lambda, b) \to (\mathsf{msg}_C, \mathsf{st}_C)$: *The PPT client takes as input the security parameter 1^λ and a bit b and outputs a message msg_C and state st_C.*
- $\mathsf{E\text{-}CNOT.Apply}((\mathcal{V}_0, \mathcal{V}_1), \mathsf{msg}_C) \to ((\mathcal{V}_0, \mathcal{V}_1), \mathsf{msg}_S)$: *The QPT server takes as input the client's message msg_C, performs an operation on the registers $\mathcal{V}_0, \mathcal{V}_1$, and produces a message msg_S.*
- $\mathsf{E\text{-}CNOT.Dec}(\mathsf{st}_C, \mathsf{msg}_S) \to (r, s)$: *The PPT decoding algorithm takes as input the client's state st_C and the server's message msg_S and outputs one-time pad keys (r, s).*

Theorem 6.28. *OSP (resp. two-round OSP) implies encrypted CNOT (resp. two-round encrypted CNOT).*

Proof. Encrypted CNOT from OSP follows generically from Corollary 6.12, so we focus on the two-round case. We first describe the protocol, where $(\mathsf{OSP.Sen}, \mathsf{OSP.Rec}, \mathsf{OSP.Dec})$ is any two-round OSP.

- $\mathsf{E\text{-}CNOT.Gen}(1^\lambda, b)$: Sample OSP first-round messages

  $$(\mathsf{OSP.msg}_{S,0}, \mathsf{OSP.st}_{S,0}) \leftarrow \mathsf{OSP.Sen}(1^\lambda, b), \quad (\mathsf{OSP.msg}_{S,1}, \mathsf{OSP.st}_{S,1}) \leftarrow \mathsf{OSP.Sen}(1^\lambda, 1-b),$$

  and define

  $$\mathsf{msg}_C \coloneqq (\mathsf{OSP.msg}_{S,0}, \mathsf{OSP.msg}_{S,1}), \quad \mathsf{st}_C \coloneqq (\mathsf{OSP.st}_{S,0}, \mathsf{OSP.st}_{S,1}).$$

- $\mathsf{E\text{-}CNOT.Apply}((\mathcal{V}_0, \mathcal{V}_1), \mathsf{msg}_C)$: Generate OSP states and second-round messages

  $$\left(|\psi_0\rangle_{\mathcal{O}_0}, \mathsf{OSP.msg}_{R,0}\right) \leftarrow \mathsf{OSP.Rec}(\mathsf{OSP.msg}_{S,0}), \quad \left(|\psi_1\rangle_{\mathcal{O}_1}, \mathsf{OSP.msg}_{R,1}\right) \leftarrow \mathsf{OSP.Rec}(\mathsf{OSP.msg}_{S,1}).$$

  Apply $\mathsf{CNOT}_{\mathcal{V}_0 \to \mathcal{O}_1}, \mathsf{CNOT}_{\mathcal{O}_0 \to \mathcal{O}_1}, \mathsf{CNOT}_{\mathcal{O}_0 \to \mathcal{V}_1}$, measure $\mathcal{O}_0$ in the Hadamard basis to obtain bit m_0, and measure $\mathcal{O}_1$ in the standard basis to obtain bit m_1. Define

  $$\mathsf{msg}_S \coloneqq (\mathsf{OSP.msg}_{R,0}, \mathsf{OSP.msg}_{R,1}, m_0, m_1),$$

  and output $(\mathcal{V}_0, \mathcal{V}_1)$.

- $\mathsf{E\text{-}CNOT.Dec}(\mathsf{st}_C, \mathsf{msg}_S)$: Decode OSP output bits

  $$t_0 \leftarrow \mathsf{OSP.Dec}(\mathsf{OSP.st}_{S,0}, \mathsf{OSP.msg}_{R,0}), \quad t_1 \leftarrow \mathsf{OSP.Dec}(\mathsf{OSP.st}_{S,1}, \mathsf{OSP.msg}_{R,1}),$$

  and compute the one-time pad keys r, s as follows.

  - If $b = 0$: $r = (0, t_0)$, $s = (t_1, 0)$.
  - If $b = 1$: $r = (0, m_1 \oplus t_1)$, $s = (m_0 \oplus t_0, 0)$.

Security is immediate from the security of the OSP. Thus, it remains to check correctness, which we check separately for $b = 0$ and $b = 1$. To avoid clutter, we will assume that registers $\mathcal{V}_0, \mathcal{V}_1$ hold pure states $\alpha_0 |0\rangle + \alpha_1 |1\rangle$ and $\beta_0 |0\rangle + \beta_1 |1\rangle$, but note that correctness extends readily to the setting where $\mathcal{V}_0, \mathcal{V}_1$ may be entangled with each other and with auxiliary systems.

First, suppose that $b = 0$. We will break the server's actions into two stages: (1) apply $\mathsf{CNOT}_{\mathcal{V}_0 \to \mathcal{O}_1}, \mathsf{CNOT}_{\mathcal{O}_0 \to \mathcal{O}_1}$ and measure $\mathcal{O}_1$, and (2) apply $\mathsf{CNOT}_{\mathcal{O}_0 \to \mathcal{V}_1}$ and measure $\mathcal{O}_0$. By the correctness of OSP, we have that right before the measurement in the first stage, the state of the system on registers $(\mathcal{V}_0, \mathcal{O}_0, \mathcal{O}_1)$ is (negligibly close to)

$$\begin{aligned}
&\mathsf{CNOT}_{\mathcal{V}_0 \to \mathcal{O}_1} \mathsf{CNOT}_{\mathcal{O}_0 \to \mathcal{O}_1} X^{0,t_0,0} Z^{0,0,t_1} \left((\alpha_0 |0\rangle + \alpha_1 |1\rangle)_{\mathcal{V}_0} \otimes |0\rangle_{\mathcal{O}_0} \otimes |+\rangle_{\mathcal{O}_1}\right) \\
&= X^{0,t_0,t_0} Z^{t_1,t_1,t_1} \mathsf{CNOT}_{\mathcal{V}_0 \to \mathcal{O}_1} \mathsf{CNOT}_{\mathcal{O}_0 \to \mathcal{O}_1} \left((\alpha_0 |0\rangle + \alpha_1 |1\rangle)_{\mathcal{V}_0} \otimes |0\rangle_{\mathcal{O}_0} \otimes |+\rangle_{\mathcal{O}_1}\right) \\
&= X^{0,t_0,t_0} Z^{t_1,t_1,t_1} \left((\alpha_0 |0\rangle + \alpha_1 |1\rangle)_{\mathcal{V}_0} \otimes |0\rangle_{\mathcal{O}_0} \otimes |+\rangle_{\mathcal{O}_1}\right).
\end{aligned}$$

Thus, measuring $\mathcal{O}_1$ in the standard basis has no affect on the remaining system. Next, we write the state of the system on registers $(\mathcal{V}_0, \mathcal{V}_1, \mathcal{O}_0)$ right before the measurement in the second stage, which we imagine implementing by applying a Hadamard gate and then measuring in the standard basis.

$$\begin{aligned}
&H_{\mathcal{O}_0}\mathsf{CNOT}_{\mathcal{O}_0\to\mathcal{V}_1}X^{0,0,t_0}Z^{t_1,0,t_1}\left((\alpha_0\left|0\right\rangle+\alpha_1\left|1\right\rangle)_{\mathcal{V}_0}\otimes(\beta_0\left|0\right\rangle+\beta_1\left|1\right\rangle)_{\mathcal{V}_1}\otimes\left|0\right\rangle_{\mathcal{O}_0}\right)\\
&=X^{0,t_0,t_1}Z^{t_1,0,t_0}H_{\mathcal{O}_0}\mathsf{CNOT}_{\mathcal{O}_0\to\mathcal{V}_1}\left((\alpha_0\left|0\right\rangle+\alpha_1\left|1\right\rangle)_{\mathcal{V}_0}\otimes(\beta_0\left|0\right\rangle+\beta_1\left|1\right\rangle)_{\mathcal{V}_1}\otimes\left|0\right\rangle_{\mathcal{O}_0}\right)\\
&=X^{0,t_0,t_1}Z^{t_1,0,t_0}\left((\alpha_0\left|0\right\rangle+\alpha_1\left|1\right\rangle)_{\mathcal{V}_0}\otimes(\beta_0\left|0\right\rangle+\beta_1\left|1\right\rangle)_{\mathcal{V}_1}\otimes\left|+\right\rangle_{\mathcal{O}_0}\right)
\end{aligned}$$

Thus, measuring $\mathcal{O}_0$ in the standard basis has no affect on the system $(\mathcal{V}_0,\mathcal{V}_1)$, which is

$$X^{0,t_0}Z^{t_1,0}\left((\alpha_0\left|0\right\rangle+\alpha_1\left|1\right\rangle)_{\mathcal{V}_0}\otimes(\beta_0\left|0\right\rangle+\beta_1\left|1\right\rangle)_{\mathcal{V}_1}\right),$$

as desired.

Next, suppose $b=1$. We break the server's actions into two stages in the same manner. By the correctness of OSP, we have that right before the measurement in the first stage, the state of the system on registers $(\mathcal{V}_0,\mathcal{O}_0,\mathcal{O}_1)$ is (negligibly close to)

$$\begin{aligned}
&\mathsf{CNOT}_{\mathcal{V}_0\to\mathcal{O}_1}\mathsf{CNOT}_{\mathcal{O}_0\to\mathcal{O}_1}X^{0,0,t_1}Z^{0,t_0,0}\left((\alpha_0\left|0\right\rangle+\alpha_1\left|1\right\rangle)_{\mathcal{V}_0}\otimes\left|+\right\rangle_{\mathcal{O}_0}\otimes\left|0\right\rangle_{\mathcal{O}_1}\right)\\
&=X^{0,0,t_1}Z^{0,t_0,0}\mathsf{CNOT}_{\mathcal{V}_0\to\mathcal{O}_1}\mathsf{CNOT}_{\mathcal{O}_0\to\mathcal{O}_1}\left((\alpha_0\left|0\right\rangle+\alpha_1\left|1\right\rangle)_{\mathcal{V}_0}\otimes\left|+\right\rangle_{\mathcal{O}_0}\otimes\left|0\right\rangle_{\mathcal{O}_1}\right)\\
&=X^{0,0,t_1}Z^{0,t_0,0}\left(\frac{1}{\sqrt{2}}\alpha_0\left|000\right\rangle+\frac{1}{\sqrt{2}}\alpha_0\left|011\right\rangle+\frac{1}{\sqrt{2}}\alpha_1\left|101\right\rangle+\frac{1}{\sqrt{2}}\alpha_1\left|110\right\rangle\right)\\
&=X^{0,0,t_1}Z^{0,t_0,0}\left(\frac{1}{\sqrt{2}}\left(\alpha_0\left|00\right\rangle+\alpha_1\left|11\right\rangle\right)\left|0\right\rangle+\frac{1}{\sqrt{2}}\left(\alpha_0\left|01\right\rangle+\alpha_1\left|10\right\rangle\right)\left|1\right\rangle\right)\\
&=\sum_{c\in\{0,1\}}\frac{1}{\sqrt{2}}X^{0,c,t_1}Z^{0,t_0,0}\Big(\left(\alpha_0\left|00\right\rangle+\alpha_1\left|11\right\rangle\right)\otimes\left|c\right\rangle\Big)\\
&=\sum_{m_1\in\{0,1\}}\frac{1}{\sqrt{2}}\Big(X^{0,m_1\oplus t_1}Z^{0,t_0}\left(\alpha_0\left|00\right\rangle+\alpha_1\left|11\right\rangle\right)\Big)\otimes\left|m_1\right\rangle,
\end{aligned}$$

where the last equality follows by a change of variables $m_1=c\oplus t_1$. Now, the server measures the last register $(\mathcal{O}_1)$ in the standard basis to obtain m_1, and then applies the second stage. We write the state of the system on registers $(\mathcal{V}_0,\mathcal{V}_1,\mathcal{O}_0)$ right before the measurement in the second stage, which again we implement by applying a Hadamard gate and then measuring in the standard basis. Below, the one-time pad keys are always written in the order $(\mathcal{V}_0,\mathcal{V}_1,\mathcal{O}_0)$.

$$\begin{aligned}
&H_{\mathcal{O}_0}\mathsf{CNOT}_{\mathcal{O}_0\to\mathcal{V}_1}X^{0,0,m_1\oplus t_1}Z^{0,0,t_0}\Big(\left(\alpha_0\left|00\right\rangle+\alpha_1\left|11\right\rangle\right)_{\mathcal{V}_0,\mathcal{O}_0}\otimes\left(\beta_0\left|0\right\rangle+\beta_1\left|1\right\rangle\right)_{\mathcal{V}_1}\Big)\\
&=X^{0,m_1\oplus t_1,t_0}Z^{0,0,m_1\oplus t_1}H_{\mathcal{O}_0}\mathsf{CNOT}_{\mathcal{O}_0\to\mathcal{V}_1}\Big(\left(\alpha_0\left|00\right\rangle+\alpha_1\left|11\right\rangle\right)_{\mathcal{V}_0,\mathcal{O}_0}\otimes\left(\beta_0\left|0\right\rangle+\beta_1\left|1\right\rangle\right)_{\mathcal{V}_1}\Big)\\
&=X^{0,m_1\oplus t_1,t_0}Z^{0,0,m_1\oplus t_1}H_{\mathcal{O}_0}\left(\alpha_0\beta_0\left|000\right\rangle+\alpha_0\beta_1\left|010\right\rangle+\alpha_1\beta_0\left|111\right\rangle+\alpha_1\beta_1\left|101\right\rangle\right)_{\mathcal{V}_0,\mathcal{V}_1,\mathcal{O}_0}\\
&=X^{0,m_1\oplus t_1,t_0}Z^{0,0,m_1\oplus t_1}\frac{1}{\sqrt{2}}\Big(\big(\alpha_0\beta_0\left|00\right\rangle+\alpha_0\beta_1\left|01\right\rangle+\alpha_1\beta_0\left|11\right\rangle+\alpha_1\beta_1\left|10\right\rangle\big)\left|0\right\rangle\\
&\qquad\qquad+\big(\alpha_0\beta_0\left|00\right\rangle+\alpha_0\beta_1\left|01\right\rangle-\alpha_1\beta_0\left|11\right\rangle-\alpha_1\beta_1\left|10\right\rangle\big)\left|1\right\rangle\Big)\\
&=\sum_{c\in\{0,1\}}\frac{1}{\sqrt{2}}X^{0,m_1\oplus t_1,t_0}Z^{c,0,m_1\oplus t_1}\Big(\alpha_0\beta_0\left|00\right\rangle+\alpha_0\beta_1\left|01\right\rangle+\alpha_1\beta_0\left|11\right\rangle+\alpha_1\beta_1\left|10\right\rangle\Big)\otimes\left|c\right\rangle
\end{aligned}$$

$$= \sum_{m_0 \in \{0,1\}} \frac{1}{\sqrt{2}} X^{0, m_1 \oplus t_1} Z^{m_0 \oplus t_0, 0} \Big(\alpha_0\beta_0 |00\rangle + \alpha_0\beta_1 |01\rangle + \alpha_1\beta_0 |11\rangle + \alpha_1\beta_1 |10\rangle \Big)_{\mathcal{V}_0, \mathcal{V}_1} \otimes Z^{m_1 \oplus t_1} |m_0\rangle_{\mathcal{O}_0} ,$$

where the equality follows by a change of variables $m_0 = c \oplus t_0$. Thus, measuring $\mathcal{O}_0$ in the standard basis to obtain m_0 yields the final state

$$X^{0, m_1 \oplus t_1} Z^{m_0 \oplus t_0, 0} \left(\alpha_0\beta_0 |00\rangle + \alpha_0\beta_1 |01\rangle + \alpha_1\beta_0 |11\rangle + \alpha_1\beta_1 |10\rangle \right)_{\mathcal{V}_0, \mathcal{V}_1} ,$$

as desired.

□

Application: QFHE. [Mah18a] combined a particular classical FHE protocol with a particular two-round encrypted CNOT protocol in a non-black-box way in order to achieve quantum FHE. However, [GV24] recently pioneered a *generic* approach to constructing quantum FHE from *any* classical FHE (with log-depth decryption) and *any* dual-mode TCF. Implicit in their work is that, in fact, one can use any two-round encrypted CNOT protocol, and thus (due to our work) any two-round OSP. We confirm this in the following theorem.

Theorem 6.29 (Adapted from [GV24]). *Given any FHE with decryption in NC$_1$ and any two-round encrypted CNOT, there exists QFHE.*

Corollary 6.30. *Given any FHE with decryption in NC$_1$ and any two-round OSP, there exists QFHE.*

Proof. In [GV24, Section 5], it is shown that classical FHE (with decryption in NC$_1$) plus a procedure for generating a "hidden Bell pair" is sufficient to construct quantum FHE. We now describe the requirements for the hidden Bell pair procedure, where Enc refers to a classical FHE encryption.

- Before the procedure begins, the client holds a bit μ and the server holds a state

$$X^r Z^s \left(\frac{1}{\sqrt{2}} \left(|0\rangle_{\mathcal{W}_0} |\phi_0^0\rangle + |1\rangle_{\mathcal{W}_0} |\phi_0^1\rangle \right) \otimes \frac{1}{\sqrt{2}} \left(|0\rangle_{\mathcal{W}_1} |\phi_1^0\rangle + |1\rangle_{\mathcal{W}_1} |\phi_1^1\rangle \right) \right) ,$$

  where the one-time pad (r, s) is on registers $\mathcal{W}_0, \mathcal{W}_1$, along with encryptions $\mathsf{Enc}(s, r)$ of the one-time pad keys.

- The client runs a parameter generation algorithm $(\mathsf{pp}, \mathsf{sp}) \leftarrow \mathsf{Gen}(1^\lambda, \mu)$ and publishes $\mathsf{pp}, \mathsf{Enc}(\mathsf{sp})$. We requires that pp is a semantically secure encryption of μ.

- The server runs a procedure to obtain a state that is negligibly close in trace distance to

$$X^{r'} Z^{s'} \left(\frac{1}{\sqrt{2}} \left(|00\rangle_{\mathcal{W}_\mu, \mathcal{W}_2} |\phi_\mu^0\rangle + |11\rangle_{\mathcal{W}_\mu, \mathcal{W}_2} |\phi_\mu^1\rangle \right) \otimes \frac{1}{\sqrt{2}} \left(|0\rangle_{\mathcal{W}_{1-\mu}} |\phi_{1-\mu}^0\rangle + |1\rangle_{\mathcal{W}_{1-\mu}} |\phi_{1-\mu}^1\rangle \right) \right) ,$$

  where the one-time pad (s', r') is applied to registers $\mathcal{W}_0, \mathcal{W}_1, \mathcal{W}_2$. The server also obtains encryptions $\mathsf{Enc}(r', s')$ of the one-time pad keys.

It is straightforward to implement such a procedure given a two-round encrypted CNOT:

- The client's parameter generation algorithm runs $(\mathsf{msg}_{C,0}, \mathsf{st}_{C,0}) \leftarrow \textsf{E-CNOT.Gen}(1^\lambda, \mu)$ and $(\mathsf{msg}_{C,1}, \mathsf{st}_{C,1}) \leftarrow \textsf{E-CNOT.Gen}(1^\lambda, 1-\mu)$, and sets $\mathsf{pp} := (\mathsf{msg}_{C,0}, \mathsf{msg}_{C,1})$ and $\mathsf{sp} := (\mathsf{st}_{C,0}, \mathsf{st}_{C,1})$.
- The server initializes register $\mathcal{W}_2$ to $|+\rangle_{\mathcal{W}_2}$, applies

  $$((\mathcal{W}_0, \mathcal{W}_2), \mathsf{msg}_{S,0}) \leftarrow \textsf{E-CNOT.Apply}((\mathcal{W}_0, \mathcal{W}_2), \mathsf{msg}_{C,0}),$$

  applies

  $$((\mathcal{W}_1, \mathcal{W}_2), \mathsf{msg}_{S,1}) \leftarrow \textsf{E-CNOT.Apply}((\mathcal{W}_1, \mathcal{W}_2), \mathsf{msg}_{C,1}),$$

  and finally uses $\mathsf{Enc}(r, s), \mathsf{Enc}(\mathsf{sp}) = \mathsf{Enc}(\mathsf{st}_{C,0}, \mathsf{st}_{C,1})$ and $\mathsf{msg}_{S,0}, \mathsf{msg}_{S,1}$ to homomorphically obtain $\mathsf{Enc}(r', s')$ under the FHE.

Correctness is straightforward, as μ determines whether the CNOT is applied from $\mathcal{W}_0$ to $\mathcal{W}_2$ or from $\mathcal{W}_1$ to $\mathcal{W}_2$. Security follows immediately from the security of the encrypted CNOT.

□

Application: CSG. Finally, we return to the notion of a claw-state generator (CSG), as defined in Section 4.2. There, it was shown that any CSG with search security implies OSP. Here, we show (a strengthening of) the reverse implication: Any OSP implies CSG with *indistinguishability* security. This implication goes via the intermediate primitive of encrypted CNOT.

Theorem 6.31. *Encrypted CNOT (resp. two-round encrypted CNOT) implies differentiated-bit* CSG *(resp. two-round* CSG*) with indistinguishability security, for any $n = \mathrm{poly}(\lambda)$.*

Corollary 6.32. *OSP (resp. two-round OSP) implies differentiated-bit* CSG *(resp. two-round* CSG*) with indistinguishability security, for any $n = \mathrm{poly}(\lambda)$.*

Proof. We give the protocol in the two-round case, which can also be instantiated with any (not necessarily two-round) encrypted CNOT to yield a (not necessarily two-round) CSG. Let

$$(\textsf{E-CNOT.Gen}, \textsf{E-CNOT.Apply}, \textsf{E-CNOT.Dec})$$

be any two-round encrypted CNOT.

- $\mathsf{CSG.Sen}(1^\lambda, n)$: Sample a uniformly random string $\Delta \leftarrow \{0,1\}^n$, and for each $i \in [n]$, sample

  $$(\textsf{E-CNOT.msg}_{C,i}, \textsf{E-CNOT.st}_{C,i}) \leftarrow \textsf{E-CNOT}(1^\lambda, \Delta_i).$$

  Define

  $$\mathsf{msg}_S := (\textsf{E-CNOT.msg}_{C,1}, \ldots, \textsf{E-CNOT.msg}_{C,n}), \qquad \mathsf{st}_S := (\textsf{E-CNOT.st}_{C,1}, \ldots, \textsf{E-CNOT.st}_{C,n}).$$

- $\mathsf{CSG.Rec}(\mathsf{msg}_S)$: Initialize a register $\mathcal{B}$ to $|+\rangle_{\mathcal{B}}$. For each $i \in [n]$, initialize a register $\mathcal{C}_i$ to $|0\rangle_{\mathcal{C}_i}$ and apply

  $$((\mathcal{B}, \mathcal{C}_i), \textsf{E-CNOT.msg}_{S,i}) \leftarrow \textsf{E-CNOT.Apply}((\mathcal{B}, \mathcal{C}_i), \textsf{E-CNOT.msg}_{C,i}).$$

  Output the state on registers $\mathcal{B}, \mathcal{C}_1, \ldots, \mathcal{C}_n$, and define

  $$\mathsf{msg}_R := (\textsf{E-CNOT.msg}_{S,1}, \ldots, \textsf{E-CNOT.msg}_{S,n}).$$

- $\mathsf{CSG.Dec}(\mathsf{st}_S, \mathsf{msg}_R)$: For each $i \in [n]$, compute

  $$((r_{i,0}, r_{i,1}), (s_{i,0}, s_{i,1})) \leftarrow \mathsf{E\text{-}CNOT.Dec}(\mathsf{E\text{-}CNOT.st}_{C,i}, \mathsf{E\text{-}CNOT.msg}_{S,i}),$$

  define

  $$r_0 \coloneqq \bigoplus_{i\in[n]} r_{i,0} \in \{0,1\},\ \ r \coloneqq (r_{1,1}, \ldots, r_{n,1}) \in \{0,1\}^n,\ \ s \coloneqq \left(\bigoplus_{i\in[n]} s_{i,0}, s_{1,1}, \ldots, s_{n,1}\right) \in \{0,1\}^{n+1},$$

  and output

  $$x_0 = r \oplus r_0 \cdot \Delta,\ \ x_1 = r \oplus (1 - r_0) \cdot \Delta,\ \ z = s \cdot (1, \Delta).$$

Observe that, by the correctness of the encrypted CNOT protocol, the state on registers $(\mathcal{B}, \mathcal{C}_1, \ldots, \mathcal{C}_n)$ output by $\mathsf{CSG.Rec}$ is (negligibly close to)

$$\begin{aligned}
&X^{\bigoplus_{i\in[n]} r_{i,0}, r_{1,1}, \ldots, r_{n,1}} Z^{\bigoplus_{i\in[n]} s_{i,0}, s_{1,1}, \ldots, s_{n,1}} \left(\frac{1}{\sqrt{2}} |0, 0^n\rangle + \frac{1}{\sqrt{2}} |1, \Delta\rangle\right) \\
&= X^{r_0, r} Z^s \left(\frac{1}{\sqrt{2}} |0, 0^n\rangle + \frac{1}{\sqrt{2}} |1, \Delta\rangle\right) \\
&= \frac{1}{\sqrt{2}} |0, r \oplus r_0 \cdot \Delta\rangle + (-1)^{s\cdot(1,\Delta)} \frac{1}{\sqrt{2}} |1, r \oplus (1 - r_0) \cdot \Delta\rangle,
\end{aligned}$$

and thus correctness holds. Indistinguishability security follows immediately from the security of the encrypted CNOT, since for each $i \in [n]$, $x_{0,i} \oplus x_{1,i} = \Delta_i$.

□

7 Implications

In this section, we show that OSP implies both commitments and oblivious transfer with one classical party, while *two-round* OSP implies public-key encryption with classical public keys and ciphertexts. The purpose of exploring this direction is to place lower bounds on the cryptography necessary to build OSP.

7.1 Commitments from OSP

Definition 7.1 (Commitment). *A (statistically-binding, computationally-hiding) commitment between a classical committer and quantum receiver consists of an interaction*

$$(\mathsf{st}_{\mathsf{Com}}, |\psi\rangle_{\mathsf{Rec}}) \leftarrow \langle \mathsf{Com}(1^\lambda, b), \mathsf{Rec}(1^\lambda)\rangle,$$

where Com *is a PPT machine with input bit* $b \in \{0,1\}$ *and* Rec *is a QPT machine, along with algorithms* $(\mathsf{Open}, \mathsf{Ver})$ *with the following syntax.*

- $\mathsf{Open}(\mathsf{st}_{\mathsf{Com}}) \to (b, s)$ *is a PPT algorithm that takes as input the committer's state* $\mathsf{st}_{\mathsf{Com}}$ *and produces a bit* b *and opening information* s.
- $\mathsf{Ver}(|\psi\rangle_{\mathsf{Rec}}, b, s) \to \{\top, \bot\}$ *is a QPT algorithm that takes as input the receiver's state* $|\psi\rangle_{\mathsf{Rec}}$*, a bit* b*, and opening information* s*, and either accepts or rejects.*

Correctness *requires that for any* $b \in \{0,1\}$,

$$\Pr\left[\mathsf{Ver}(|\psi\rangle_{\mathsf{Rec}}, b, s) = \top : \begin{array}{r} (\mathsf{st}_{\mathsf{Com}}, |\psi\rangle_{\mathsf{Rec}}) \leftarrow \langle \mathsf{Com}(1^\lambda, b), \mathsf{Rec}(1^\lambda)\rangle \\ (b, s) \leftarrow \mathsf{Open}(\mathsf{st}_{\mathsf{Com}}) \end{array}\right] = 1 - \mathrm{negl}(\lambda).$$

The commitment satisfies computational hiding *if for any QPT adversarial receiver* $\{\mathsf{Adv}_\lambda\}_{\lambda \in \mathbb{N}}$,

$$\begin{aligned} \Bigg| &\Pr\left[b_{\mathsf{Adv}} = 0 : (\mathsf{st}_{\mathsf{Com}}, b_{\mathsf{Adv}}) \leftarrow \langle \mathsf{Com}(1^\lambda, 0), \mathsf{Adv}_\lambda\rangle\right] \\ &- \Pr\left[b_{\mathsf{Adv}} = 0 : (\mathsf{st}_{\mathsf{Com}}, b_{\mathsf{Adv}}) \leftarrow \langle \mathsf{Com}(1^\lambda, 1), \mathsf{Adv}_\lambda\rangle\right] \Bigg| = \mathrm{negl}(\lambda). \end{aligned}$$

The commitment satisfies statistical binding *if for any unbounded adversarial committer* $\{\mathsf{Adv}_\lambda\}_{\lambda \in \mathbb{N}}$,

$$\Pr\left[\mathsf{Ver}(|\psi\rangle_{\mathsf{Rec}}, b, s) = \top : \begin{array}{r} (\mathsf{st}_{\mathsf{Adv}}, |\psi\rangle_{\mathsf{Rec}}) \leftarrow \langle \mathsf{Adv}_\lambda, \mathsf{Rec}(1^\lambda)\rangle \\ b \leftarrow \{0,1\} \\ s \leftarrow \mathsf{Adv}_\lambda(\mathsf{st}_{\mathsf{Adv}}, b) \end{array}\right] \leq \frac{1}{2} + \mathrm{negl}(\lambda).$$

Theorem 7.2. *OSP implies a commitment between a classical committer and quantum receiver.*

Proof. Given any OSP, consider the following commitment scheme.

- For all $i \in [\lambda]$, execute an OSP between the committer playing the role of the sender with input b and the receiver playing the role of the OSP receiver. Each execution results in an output s_i for the sender a state $|\psi_i\rangle$ for the receiver. Define $\mathsf{st}_{\mathsf{Com}} := (b, s_1, \ldots, s_\lambda)$ and $|\psi\rangle_{\mathsf{Rec}} := |\psi_1\rangle \otimes \cdots \otimes |\psi_\lambda\rangle$.
- $\mathsf{Open}(\mathsf{st}_{\mathsf{Com}})$ outputs $b, s = (s_1, \ldots, s_\lambda)$.
- $\mathsf{Ver}(|\psi\rangle_{\mathsf{Rec}}, b, s)$ measures $|\psi\rangle_{\mathsf{Rec}}$ in the standard basis if $b = 0$ or the Hadamard basis if $b = 1$, and accepts if the outcome is s.

Correctness is straightforward from the correctness of OSP, and hiding follows from the security of OSP via a standard hybrid argument.

To show binding, consider any state $|\psi\rangle_{\mathsf{Rec}}$ that the receiver could have after interacting with Adv_λ in the commit phase. Note that for any s_0, s_1,

$$\Pr[\mathsf{Ver}(|\psi\rangle_{\mathsf{Rec}}, 0, s_0) = \top] = \|\langle s_0 | \psi\rangle_{\mathsf{Rec}}\|^2 \text{ and } \Pr[\mathsf{Ver}(|\psi\rangle_{\mathsf{Rec}}, 1, s_1) = \top] = \|\langle s_1 | H^{\otimes\lambda} |\psi\rangle_{\mathsf{Rec}}\|^2.$$

Let

$$s_0 := \max_s \{\Pr[\mathsf{Ver}(|\psi\rangle_{\mathsf{Rec}}, 0, s) = \top]\} \text{ and } s_1 := \max_s \{\Pr[\mathsf{Ver}(|\psi\rangle_{\mathsf{Rec}}, 1, s) = \top]\},$$

meaning that in the binding game, the adversary's optimal strategy given b is to output s_b.

Then since

$$\|\langle s_0 | H^{\otimes\lambda} |s_1\rangle\|^2 = \frac{1}{2^\lambda} = \mathrm{negl}(\lambda),$$

we conclude that

$$\Pr[\mathsf{Ver}(|\psi\rangle_{\mathsf{Rec}}, 0, s_0) = \top] + \Pr[\mathsf{Ver}(|\psi\rangle_{\mathsf{Rec}}, 1, s_1) = \top] \leq 1 + \mathrm{negl}(\lambda),$$

which completes the proof.

□

Remark 7.3. *Even though the above commitment satisfies a standard notion of statistical binding (sum binding), it is unclear if the effective committed bit can actually be extracted (even inefficiently) from an adversarial committer. Note that this is always possible in both the fully-classical and fully-quantum setting. Intuitively, the difficulty here arises from the fact that the span of receiver states that can be (perfectly) opened to 0 and the span of receiver states that be (perfectly) opened to 1 are equivalent: the span is the entire Hilbert space. Thus, we cannot define an inefficient projective measurement that successfully distinguishes these spaces. We mention this both out of curiosity, and also because in the next section we will actually use an inefficiently-extractable commitment scheme. Since we are not able to show that our OSP-based commitment is inefficiently-extractable, we rely on a one-way-function based commitment for this purpose.*

7.2 OT from OSP

First, we define a notion of game-based oblivious transfer (OT) with one-sided statistical security, where one party (the receiver) is classical and the other party (the sender) is quantum. For security, we allow the receiver to be an unbounded algorithm, but require the adversarial sender to be QPT.

Definition 7.4 (OT with one-sided statistical security between classical receiver and quantum sender)**.** *Oblivious transfer (OT) with one-sided statistical security is a protocol that takes place between a PPT receiver R with input $b \in \{0,1\}$ and a QPT sender S:*

$$(r, (r_0, r_1), \tau) \leftarrow \langle R(1^\lambda, b), S(1^\lambda) \rangle,$$

where $r_0, r_1 \in \{0,1\}^n$ is the sender's output, $r \in \{0,1\}^n$ is the receiver's output, and τ is the (classical) transcript of communication produced during the protocol. We require the following properties.

- ***Correctness.*** *For any $b \in \{0,1\}$,*

  $$\Pr\left[r = r_b : (r, (r_0, r_1), \tau) \leftarrow \langle R(1^\lambda, b), S(1^\lambda) \rangle\right] = 1 - \mathrm{negl}(\lambda).$$

  We say the protocol satisfies perfect *correctness if the above probability is equal to 1.*

- ***Security against a QPT sender.*** *For any QPT adversarial sender $\{\mathsf{Adv}_\lambda\}_{\lambda \in \mathbb{N}}$,*

  $$\Big| \Pr\left[b_{\mathsf{Adv}} = 0 : (r, b_{\mathsf{Adv}}, \tau) \leftarrow \langle R(1^\lambda, 0), \mathsf{Adv}_\lambda \rangle\right] \\ - \Pr\left[b_{\mathsf{Adv}} = 0 : (r, b_{\mathsf{Adv}}, \tau) \leftarrow \langle R(1^\lambda, 1), \mathsf{Adv}_\lambda \rangle\right] \Big| = \mathrm{negl}(\lambda).$$

We consider two different security properties that may hold against an unbounded receiver.

- ***Search security against an unbounded receiver.*** *For any unbounded adversarial receiver* Adv,

  $$\Pr\Big[r_{\mathsf{Adv}} = (r_0, r_1) : (r_{\mathsf{Adv}}, (r_0, r_1), \tau) \leftarrow \langle \mathsf{Adv}, S(1^\lambda) \rangle\Big] = \mathrm{negl}(\lambda).$$

- ***Indistinguishability security against an unbounded receiver.*** *For this definition, we restrict our attention to the case where $n = 1$ (i.e. r_0, r_1 are single bits). There exists an unbounded extractor* Ext *such that for any unbounded adversarial receiver* Adv,

  $$\left| \Pr\left[r_{1-b,\mathsf{Adv}} = r_{1-b} : \begin{array}{r} ((r_{0,\mathsf{Adv}}, r_{1,\mathsf{Adv}}), (r_0, r_1), \tau) \leftarrow \langle \mathsf{Adv}, S(1^\lambda) \rangle \\ b \leftarrow \mathsf{Ext}(\tau) \end{array}\right] - \frac{1}{2} \right| = \mathrm{negl}(\lambda).$$

We first provide a construction of OT with search security against an unbounded receiver, assuming only a differentiated-bit CSG with indistinguishability security (defined in Definition 4.4 and constructed from OSP in Section 6.5).

Since search security is a somewhat non-standard notion of OT security, we also show how to tweak the construction to obtain indistinguishability security by additionally using an *inefficiently-extractable* commitment (defined in Appendix B). We use a fully-classical (post-quantum) inefficiently-extractable commitment, which is known from any (post-quantum) one-way function [Nao91]. Our protocols are given in Fig. 4 and Fig. 5.

OT with search security from OSP

- For each $i \in [2\lambda]$:
    - The classical receiver R and quantum sender S engage in a differentiated-bit CSG with indistinguishability security where R plays the role of the sender $\mathsf{CSG.Sen}$ and S plays the role of the receiver $\mathsf{CSG.Rec}$:

      $$((x_{0,i}, x_{1,i}, z_i), |\psi_i\rangle) \leftarrow \langle \mathsf{CSG.Sen}(1^\lambda, 1), \mathsf{CSG.Rec}(1^\lambda, 1)\rangle.$$

      In the case that both parties are honest, the state $|\psi_i\rangle$ obtained by the sender is (negligibly close to)

      $$|\psi_i\rangle = \frac{1}{\sqrt{2}} \left(|0, x_{0,i}\rangle + (-1)^{z_i} |1, x_{1,i}\rangle\right).$$
- S samples a uniformly random subset $T \subset [2\lambda]$ of size λ and sends T to R. Define $\overline{T} \coloneqq [2\lambda] \setminus T$.
- Given input bit b, R defines $b_i \coloneqq b \oplus x_{0,i} \oplus x_{1,i}$ for all $i \in \overline{T}$, and defines r to be the concatenation of the bits $\{x_{0,i}\}_{i \in \overline{T}}$. R sends $\{x_{0,i}, x_{1,i}, z_i\}_{i \in T}, \{b_i\}_{i \in \overline{T}}$ to S, and outputs r.
- For each $i \in T$, S attempts to project $|\psi_i\rangle$ onto the state $\frac{1}{\sqrt{2}}(|0, x_{0,i}\rangle + (-1)^{z_i} |1, x_{1,i}\rangle)$. If any measurement fails, output $r_0, r_1 \leftarrow \{0,1\}^n$ sampled uniformly at random. Otherwise, for all $i \in \overline{T}$, measure $|\psi_i\rangle$ in the standard basis to obtain bits (c_i, y_i), and define $r_{0,i} \coloneqq y_i \oplus b_i \cdot c_i$ and $r_{1,i} \coloneqq y_i \oplus (1 \oplus b_i) \cdot c_i$. Finally, define r_0 to be the concatenation of all bits $\{r_{0,i}\}_{i \in \overline{T}}$ and r_1 to be the concatenation of all bits $\{r_{1,i}\}_{i \in \overline{T}}$, and output (r_0, r_1).

Figure 4: A protocol for OT with search security against an unbounded receiver, assuming only claw-state generators with indistinguishability security (which follow from OSP).

Theorem 7.5. *The protocol in Fig. 4 (resp. Fig. 5) satisfies OT with search (resp. indistinguishability) security against an unbounded receiver (Definition 7.4). That is, OSP implies OT with search security, while OT plus one-way functions implies OT with indistinguishability security.*

Proof. We write out the proofs for the protocol in Fig. 5, and note that essentially the same proof strategy works to show security of the protocol in Fig. 4. In what follows, we prove correctness, security against a QPT sender, and indistinguishability security against an unbounded receiver.

Correctness. We need to show that for any $b \in \{0,1\}$, the receiver's output r is equal to the sender's output r_b with overwhelming probability. To see this, it suffices to show that for all $i \in \overline{T}$,

OT with indistinguishability security from OSP plus OWF

- For each $i \in [2\lambda]$:
  - The classical receiver R and quantum sender S engage in a differentiated-bit CSG with indistinguishability security where R plays the role of the sender CSG.Sen and S plays the role of the receiver CSG.Rec:

    $$((x_{0,i}, x_{1,i}, z_i), |\psi_i\rangle) \leftarrow \langle \mathsf{CSG.Sen}(1^\lambda, 1), \mathsf{CSG.Rec}(1^\lambda, 1)\rangle.$$

    In the case that both parties are honest, the state $|\psi_i\rangle$ obtained by the sender is (negligibly close to)

    $$|\psi_i\rangle = \frac{1}{\sqrt{2}}\left(|0, x_{0,i}\rangle + (-1)^{z_i}\,|1, x_{1,i}\rangle\right).$$
  - R and S engage in an inefficiently-extractable commitment, with R playing the role of Com and S playing the role of Rec:

    $$(\mathsf{st}_{\mathsf{Com},i}, \mathsf{st}_{\mathsf{Rec},i}, \tau_i) \leftarrow \langle \mathsf{Com}(1^\lambda, (x_{0,i}, x_{1,i}, z_i)), \mathsf{Rec}(1^\lambda)\rangle.$$
- S samples a uniformly random subset $T \subset [2\lambda]$ of size λ and sends T to R. Define $\overline{T} := [2\lambda] \setminus T$.
- For each $i \in T$, R computes the opening to its commitment $((x_{0,i}, x_{1,i}, z_i), w_i) \leftarrow \mathsf{Open}(\mathsf{st}_{\mathsf{Com},i})$. Then, given input bit b, R defines $b_i := b \oplus x_{0,i} \oplus x_{1,i}$ for all $i \in \overline{T}$, and defines $r := \bigoplus_{i\in\overline{T}} x_{0,i}$. Finally, R sends $\{x_{0,i}, x_{1,i}, z_i, w_i\}_{i\in T}, \{b_i\}_{i\in\overline{T}}$ to S, and outputs r.
- For each $i \in T$, S checks that $\mathsf{Ver}(\mathsf{st}_{\mathsf{Rec}}, (x_{0,i}, x_{1,i}, z_i), w_i) = \top$, and attempts to project $|\psi_i\rangle$ onto the state $\frac{1}{\sqrt{2}}(|0, x_{0,i}\rangle + (-1)^{z_i}\,|1, x_{1,i}\rangle)$. If any verification check or measurements fails, output $r_0, r_1 \leftarrow \{0, 1\}$ sampled uniformly at random. Otherwise, for all $i \in \overline{T}$, measure $|\psi_i\rangle$ in the standard basis to obtain bits (c_i, y_i), and define $r_{0,i} := y_i \oplus b_i \cdot c_i$ and $r_{1,i} := y_i \oplus (1 \oplus b_i) \cdot c_i$. Finally, define

  $$r_0 := \bigoplus_{i\in\overline{T}} r_{0,i}, \quad r_1 := \bigoplus_{i\in\overline{T}} r_{1,i},$$

  and output (r_0, r_1).

Figure 5: A protocol for OT with indistinguishability security against an unbounded receiver, assuming CSG plus one-way functions. The differences from the OT protocol in Fig. 4 are highlighted in red.

it holds that (with overwhelming probability) $x_{0,i} = y_i \oplus (b \oplus b_i) \cdot c_i$, where $b_i = b \oplus x_{0,i} \oplus x_{1,i}$, and (c_i, y_i) are obtained by measuring a state that is (negligibly close to)

$$\frac{1}{\sqrt{2}}\left(|0, x_{0,i}\rangle + |1, x_{1,i}\rangle\right)$$

in the standard basis. This is easy to check by plugging in the two choices of $(c_i, y_i) = (0, x_{0,i})$ and $(c_i, y_i) = (1, x_{1,i})$.

Security against a QPT sender. This follows directly from the indistinguishability security of the CSG. Indeed, the only information that the sender obtains about b are the bits $b_i = b \oplus x_{0,i} \oplus x_{1,i}$, and the security of CSG implies that each bit $x_{0,i} \oplus x_{1,i}$ is computationally unpredictable to the QPT sender.

Security against an unbounded receiver. First, we define the unbounded-time extractor Ext. Given the classical transcript τ of the OT protocol, Ext identifies the transcripts $\tau_{\mathsf{Com},1}, \dots, \tau_{\mathsf{Com},2\lambda}$ of the commitment protocols, along with the set T and bits $\{b_i\}_{i \in \overline{T}}$. It runs the unbounded-time extractor for the commitment scheme on each $\{\tau_{\mathsf{Com},i}\}_{i \in \overline{T}}$ to obtain $\{(x_{0,i}, x_{1,i}, z_i)\}_{i \in \overline{T}}$, and outputs $b = \mathsf{maj}\{b_i \oplus x_{0,i} \oplus x_{1,i}\}_{i \in \overline{T}}$.

For any $\{0,1\}$-valued random variable s sampled during the protocol, we define

$$\mathsf{Bias}(s) := \left| \Pr[s = 0 \mid \mathsf{st}_R] - \frac{1}{2} \right|,$$

where st_R is the final state of the receiver at the conclusion of the protocol. That is, $\mathsf{Bias}(s)$ captures that advantage that any unbounded receiver has in guessing the bit s at the conclusion of the protocol. Our goal is to show that $\mathsf{Bias}(r_{1 \oplus b}) = \mathrm{negl}(\lambda)$. Towards showing this, we prove the following claim.

Claim 7.6. *Consider running the commitment scheme extractor on all commitment transcripts* $\{\tau_{\mathsf{Com},i}\}_{i \in [2\lambda]}$ *to obtain* $\{(x_{0,i}, x_{1,i}, z_i)\}_{i \in [2\lambda]}$. *Define*

$$|\psi_i\rangle := \frac{1}{\sqrt{2}} \left(|0, x_{0,i}\rangle + |1, x_{1,i}\rangle \right)$$

and let $|\psi_i'\rangle$ *be the state obtained by the sender after the i'th CSG protocol (note that these states are unentangled with each other). Then with probability* $1 - \mathrm{negl}(\lambda)$, *either (1) one of the sender's measurements fails, or (2) for at least* $7\lambda/4$ *indices* $i \in [2\lambda]$, *it holds that* $|\langle \psi_i | \psi_i' \rangle|^2 \geq 1 - 1/25$.

Proof. Suppose that condition (2) does not hold. We will show that this implies that with probability $1 - \mathrm{negl}(\lambda)$, one of the sender's measurements fails, which suffices to show the claim.

First note that due to the extractability property of the commitment, there is only $\mathrm{negl}(\lambda)$ probability that the receiver can open any commitment i to a value other than the value $(x_{0,i}, x_{1,i}, z_i)$ that was extracted. Thus, with all but $\mathrm{negl}(\lambda)$ probability, the sender attempts to project each $|\psi_i'\rangle$ for $i \in T$ onto the state $|\psi_i\rangle$ defined above. Since T is sampled as a uniformly random subset of $[2\lambda]$ of size λ, a straightforward tail bound shows that, except with $\mathrm{negl}(\lambda)$ probability, there will be $\Omega(\lambda)$ many indices i such that $i \in T$ and $|\langle \psi_i | \psi_i' \rangle|^2 < 1 - 1/25$. For each such i, the probability that the sender's measurments fails is at least $1/25$, and these probabilities are independent. Thus the overall probability of *not* failing is at most $(1 - 1/25)^{\Omega(\lambda)} \leq e^{-\Omega(\lambda)} = \mathrm{negl}(\lambda)$. □

Now, in the case that one of the sender's measurements fails, they output uniformly random bits r_0, r_1, meaning $\mathsf{Bias}(r_{1 \oplus b}) = 0$. If not, the above claim shows that it suffices to consider the scenario where condition (2) holds, with only a $\mathrm{negl}(\lambda)$ difference in the final bound on $\mathsf{Bias}(r_{1 \oplus b})$.

In this case, let $S \subseteq \overline{T}$ be the set of indices such that (1) $|\langle \psi_i | \psi_i' \rangle|^2 \geq 1 - 1/25$ and (2) $b = b_i \oplus x_{0,i} \oplus x_{1,i}$. By the claim above and the definition of b, we know that $|S| \geq \lambda/4$. For each $i \in S$, we have that $r_{1 \oplus b, i} = y_i \oplus (1 \oplus b \oplus b_i) \cdot c_i = y_i \oplus (1 \oplus x_{0,i} \oplus x_{1,i}) \cdot c_i$, where (c_i, y_i) are obtained by measuring $|\psi_i'\rangle$ in the standard basis. Note that if instead, (c_i, y_i) were obtained by measuring

$$|\psi_i\rangle = \frac{1}{\sqrt{2}}\left(|0, x_{0,i}\rangle + |1, x_{1,i}\rangle\right)$$

in the standard basis, then $r_{1\oplus b,i}$ would be a uniformly random bit (even conditioned on st_R). To see this, note that in the case that $x_{0,i} = x_{1,i}$ we have that $r_{1\oplus b,i} = y_i \oplus c_i$ where y_i is fixed and c_i is a uniformly random bit, while in the case that $x_{0,i} \neq x_{1,i}$ we have that $r_{1\oplus b,i} = y_i$ where y_i is a uniformly random bit. Thus, by Gentle Measurement (Lemma 3.1), we have that the total variation distance between $r_{1\oplus b,i}$ and a uniformly random bit is $\leq 2\sqrt{1/25} = 2/5$, which implies that $\mathsf{Bias}(r_{1\oplus b,i}) \leq 4/5$. Finally, noting that each $r_{1\oplus b,i}$ is an independent random variable (since the states $|\psi_i'\rangle$ are all unentangled), we have that

$$\mathsf{Bias}(r_{1\oplus b}) \leq \prod_{i\in S} \mathsf{Bias}(r_{1\oplus b,i}) \leq (4/5)^{\lambda/4} = \mathrm{negl}(\lambda),$$

which completes the proof.

□

Next, we argue that combining this result with [ACC+22, Theorem 3.1] allows us to conclude the following.

Corollary 7.7. *Perfectly correct OSP does not exist in the quantum random oracle model.*

Proof. First, it is straightforward to verify that when the protocol in Fig. 4 is instantiated with a perfectly correct claw-state generator (which can be constructed from a perfectly-correct OSP), then the resulting protocol is a perfectly-correct OT with search security against an unbounded receiver.

Next, we confirm that our notion of perfectly-correct OT (with a classical receiver) implies perfectly-correct key agreement (with one classical party). The key agreement protocol will have Alice sample a random $b \leftarrow \{0,1\}$ and act as the OT receiver in a protocol with Bob, obtaining string r_b. Then Alice sends her bit b to Bob, and their shared key is defined to be r_b. Any QPT eavesdropper that can output r_b with noticeable probability can be used to break the search security of the OT. Indeed, by security against a QPT sender, such an eavesdropper will *also* be able to output r_{1-b} with noticeable probability if the final bit of the transcript they are given is flipped to $1-b$. Thus, an adversarial receiver can first run the OT protocol honestly with the sender to obtain (b, r_b), and then run the eavesdropper on $1-b$ to obtain r_{1-b} with noticeable probability.

Finally, we appeal to [ACC+22, Theorem 3.1], which establishes that perfectly-correct key agreement between one classical and one quantum party does not exist in the quantum random oracle model. Thus, we conclude that perfectly-correct OSP does not exist in the quantum random oracle model. □

7.3 PKE from two-round OSP

In this section, we show that two-round OSP implies CPA-secure public-key encryption with a classical key generator / decryptor and a quantum encryptor. In particular, the scheme has classical keys and ciphertexts. First, we define CPA-secure public-key encryption.

Definition 7.8 (CPA-secure PKE)**.** *A CPA-secure public-key encryption scheme with classical key generator and quantum encryptor consists of the following algorithms.*

- $\mathsf{KeyGen}(1^\lambda) \to (\mathsf{pk}, \mathsf{sk})$: *The PPT key generation algorithm takes as input the security parameter and outputs a public key* pk *and secret key* sk*.*
- $\mathsf{Enc}(\mathsf{pk}, m) \to \mathsf{ct}$: *The QPT encryption algorithm takes as input the public key and a plaintext bit* $m \in \{0,1\}$*, and outputs a (classical) ciphertext* ct*.*
- $\mathsf{Dec}(\mathsf{sk}, \mathsf{ct}) \to m$: *The PPT decryption algorithm takes as input the secret key and a ciphertext, and outputs a plaintext* m*.*

It should satisfy the following properties.

- ***Correctness**: For any* $m \in \{0,1\}$*,*

$$\Pr\left[\mathsf{Dec}(\mathsf{sk}, \mathsf{ct}) = m : (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda), \mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{pk}, m)\right] = 1 - \mathrm{negl}(\lambda).$$

- ***Security**: For any QPT adversary* $\{\mathsf{Adv}_\lambda\}_{\lambda \in \mathbb{N}}$*,*

$$\Big| \Pr\left[\mathsf{Adv}_\lambda(\mathsf{pk}, \mathsf{ct}) = 0 : (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda), \mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{pk}, 0)\right] - \Pr\left[\mathsf{Adv}_\lambda(\mathsf{pk}, \mathsf{ct}) = 0 : (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda), \mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{pk}, 1)\right] \Big| = \mathrm{negl}(\lambda).$$

We build our scheme from any two-round differentiated-bit CSG with indistinguishability security (defined in Definition 4.4 and constructed from two-round OSP in Section 6.5). Letting $(\mathsf{CSG.Sen}, \mathsf{CSG.Rec}, \mathsf{CSG.Dec})$ be the CSG algorithms, the scheme is constructed as follows.

- $\mathsf{KeyGen}(1^\lambda)$: Sample $(\mathsf{msg}_S, \mathsf{st}_S) \leftarrow \mathsf{CSG.Sen}(1^\lambda, 1)$ and output $\mathsf{pk} = \mathsf{msg}_S$ and $\mathsf{sk} = \mathsf{st}_S$.
- $\mathsf{Enc}(\mathsf{pk}, m)$: Given a message bit $m \in \{0,1\}$, sample $(|\psi\rangle, \mathsf{msg}_R) \leftarrow \mathsf{CSG.Rec}(\mathsf{msg}_S)$ and measure $|\psi\rangle$ in the standard basis to obtain bits (b, x_b). Output $\mathsf{ct} := (\mathsf{msg}_R, b, m \oplus x_b)$.
- $\mathsf{Dec}(\mathsf{sk}, \mathsf{ct})$: Parse $\mathsf{ct} = (\mathsf{msg}_R, b, m')$, run $(x_0, x_1, z) \leftarrow \mathsf{CSG.Dec}(\mathsf{st}_S, \mathsf{msg}_R)$, and output $m = m' \oplus x_b$.

Theorem 7.9. *The scheme described above satisfies Definition 7.8.*

Proof. First, we show correctness. By the correctness of CSG, the state $|\psi\rangle$ sampled by Enc is (negligibly close to) $\frac{1}{\sqrt{2}}(|0, x_0\rangle + (-1)^z |1, x_1\rangle)$, where the bits (x_0, x_1, z) can be recovered by running $(x_0, x_1, z) \leftarrow \mathsf{CSG.Dec}(\mathsf{st}_S, \mathsf{msg}_R)$. Thus, given b, the decryptor can determine x_b, and unmask $m \oplus x_b$ to recover the message.

Next, we show security. It suffices to show that x_b is unpredictable, that is, for any QPT adversary $\{\mathsf{Adv}_\lambda\}_{\lambda \in \mathbb{N}}$, it holds that

$$\left| \Pr\left[\mathsf{Adv}_\lambda(\mathsf{msg}_S, \mathsf{msg}_R, b) = x_b : \begin{array}{r} (\mathsf{msg}_S, \mathsf{st}_S) \leftarrow \mathsf{CSG.Sen}(1^\lambda, 1) \\ (|\psi\rangle, \mathsf{msg}_R) \leftarrow \mathsf{CSG.Rec}(\mathsf{msg}_S) \\ (x_0, x_1, z) \leftarrow \mathsf{CSG.Dec}(\mathsf{st}_S, \mathsf{msg}_R) \\ b \leftarrow \{0,1\} \end{array}\right] - \frac{1}{2} \right| = \mathrm{negl}(\lambda).$$

Indeed, note that in the real scheme, b is obtained by measuring $|\psi\rangle$ in the standard basis, but since $|\psi\rangle$ is (negligibly close to) a uniform superposition over $(0, x_0)$ and $(1, x_1)$, we can imagine just sampling a random bit $b \leftarrow \{0,1\}$, and Adv_λ will have negligibly close to the same advantage.

Now, suppose there exists Adv_λ that has noticeable advantage in the above game. We use such an adversary to break the indistinguishability security of the CSG. The CSG adversary will do the following:

- Receive msg_S from its challenger.
- Run $(|\psi\rangle, \mathsf{msg}_R) \leftarrow \mathsf{CSG.Rec}(\mathsf{msg}_S)$.
- Sample $b \leftarrow \{0,1\}$ and run $\mathsf{Adv}_\lambda(\mathsf{msg}_S, \mathsf{msg}_R, b)$ to obtain a guess for x_b.
- Measure $|\psi\rangle$ in the standard basis to obtain $(b', x_{b'})$.
- If $b = b'$, output a random bit, and otherwise, if $b \neq b'$, output $x_b \oplus x_{b'}$.

Note that with probability 1/2, the CSG adversary makes a uniformly random guess, and otherwise, the CSG makes a guess for $x_0 \oplus x_1$ with noticeable advantage, due to the guarantee on Adv_λ. Thus the CSG adversary has a noticeable advantage in breaking the indistinguishability security of the CSG.

□

References

[ABCC24] Atul Singh Arora, Kishor Bharti, Alexandru Cojocaru, and Andrea Coladangelo. A computational test of quantum contextuality, and even simpler proofs of quantumness, 2024.

[AC02] Mark Adcock and Richard Cleve. A quantum goldreich-levin theorem with cryptographic applications. In *Proceedings of the 19th Annual Symposium on Theoretical Aspects of Computer Science*, STACS '02, page 323–334, Berlin, Heidelberg, 2002. Springer-Verlag.

[ACC$^+$22] Per Austrin, Hao Chung, Kai-Min Chung, Shiuan Fu, Yao-Ting Lin, and Mohammad Mahmoody. On the impossibility of key agreements from quantum random oracles. In *Advances in Cryptology – CRYPTO 2022: 42nd Annual International Cryptology Conference, CRYPTO 2022, Santa Barbara, CA, USA, August 15–18, 2022, Proceedings, Part II*, page 165–194, Berlin, Heidelberg, 2022. Springer-Verlag.

[ACC$^+$24] Omar Amer, Kaushik Chakraborty, David Cui, Fatih Kaleoglu, Charles Lim, Minzhao Liu, and Marco Pistoia. Certified randomness implies secure classical position-verification. Cryptology ePrint Archive, Paper 2024/1726, 2024.

[AMR22] Navid Alamati, Giulio Malavolta, and Ahmadreza Rahimi. Candidate trapdoor claw-free functions from group actions with applications to quantum protocols. In Eike Kiltz and Vinod Vaikuntanathan, editors, *Theory of Cryptography*, pages 266–293, Cham, 2022. Springer Nature Switzerland.

[Bar21] James Bartusek. Secure quantum computation with classical communication. In Kobbi Nissim and Brent Waters, editors, *Theory of Cryptography*, pages 1–30, Cham, 2021. Springer International Publishing.

[BCM$^+$18] Zvika Brakerski, Paul Christiano, Urmila Mahadev, Umesh V. Vazirani, and Thomas Vidick. A cryptographic test of quantumness and certifiable randomness from a single quantum device. *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 320–331, 2018.

[BGKM$^+$23] Zvika Brakerski, Alexandru Gheorghiu, Gregory D. Kahanamoku-Meyer, Eitan Porat, and Thomas Vidick. Simple tests of quantumness also certify qubits. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023*, pages 162–191, Cham, 2023. Springer Nature Switzerland.

[BKL$^+$22] James Bartusek, Yael Tauman Kalai, Alex Lombardi, Fermi Ma, Giulio Malavolta, Vinod Vaikuntanathan, Thomas Vidick, and Lisa Yang. Succinct classical verification of quantum computation. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022*, pages 195–211, Cham, 2022. Springer Nature Switzerland.

[Bro15] Anne Broadbent. Delegating private quantum computations. *Canadian Journal of Physics*, 93(9):941–946, 2015.

[CCKW19] Alexandru Cojocaru, Léo Colisson, Elham Kashefi, and Petros Wallden. Qfactory: Classically-instructed remote secret qubits preparation. In Steven D. Galbraith and Shiho Moriai, editors, *Advances in Cryptology – ASIACRYPT 2019*, pages 615–645, Cham, 2019. Springer International Publishing.

[CCKW21] Alexandru Cojocaru, Léo Colisson, Elham Kashefi, and Petros Wallden. On the possibility of classical client blind quantum computing. *Cryptography*, 5(1):3, January 2021.

[CHS05] Ran Canetti, Shai Halevi, and Michael Steiner. Hardness amplification of weakly verifiable puzzles. In Joe Kilian, editor, *TCC 2005*, volume 3378 of *LNCS*, pages 17–33. Springer, Heidelberg, February 2005.

[CLLW22] Kai-Min Chung, Yi Lee, Han-Hsuan Lin, and Xiaodi Wu. Constant-round blind classical verification of quantum sampling. In *Advances in Cryptology – EUROCRYPT 2022: 41st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Trondheim, Norway, May 30 – June 3, 2022, Proceedings, Part III*, page 707–736, Berlin, Heidelberg, 2022. Springer-Verlag.

[CMM$^+$24] David Cui, Giulio Malavolta, Arthur Mehta, Anand Natarajan, Connor Paddock, Simon Schmidt, Michael Walter, and Tina Zhang. A computational tsirelson's theorem for the value of compiled xor games, 2024.

[DK16] Vedran Dunjko and Elham Kashefi. Blind quantum computing with two almost identical states, 2016.

[GKM$^+$00] Y. Gertner, S. Kannan, T. Malkin, O. Reingold, and M. Viswanathan. The relationship between public key encryption and oblivious transfer. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, FOCS '00, page 325, USA, 2000. IEEE Computer Society.

[GKNV24] Sam Gunn, Yael Tauman Kalai, Anand Natarajan, and Agi Villanyi. Classical commitments to quantum states, 2024.

[Gri20] Alex B. Grilo. A simple protocol for verifiable delegation of quantum computation in one round, 2020.

[GV24] Aparna Gupte and Vinod Vaikuntanathan. How to construct quantum fhe, generically. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology – CRYPTO 2024*, pages 246–279, Cham, 2024. Springer Nature Switzerland.

[KLVY23] Yael Kalai, Alex Lombardi, Vinod Vaikuntanathan, and Lisa Yang. Quantum advantage from any non-local game. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, STOC 2023, page 1617–1628, New York, NY, USA, 2023. Association for Computing Machinery.

[KMCVY21] Gregory D. Kahanamoku-Meyer, Soonwon Choi, Umesh V. Vazirani, and Norman Y. Yao. Classically verifiable quantum advantage from a computational bell test. *Nature Physics*, 18:918 – 924, 2021.

[KMP$^+$24] Alexander Kulpe, Giulio Malavolta, Connor Paddock, Simon Schmidt, and Michael Walter. A bound on the quantum value of all compiled nonlocal games, 2024.

[Kup05] Greg Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. *SIAM Journal on Computing*, 35(1):170–188, 2005.

[LLQ22] Jiahui Liu, Qipeng Liu, and Luowen Qian. Beating Classical Impossibility of Position Verification. In Mark Braverman, editor, *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, volume 215 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 100:1–100:11, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[MAF23] Ilya Merkulov and Rotem Arnon-Friedman. Entropy accumulation under post-quantum cryptographic assumptions, 2023.

[Mah18a] Urmila Mahadev. Classical homomorphic encryption for quantum circuits. *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 332–338, 2018.

[Mah18b] Urmila Mahadev. Classical verification of quantum computations. *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 259–267, 2018.

[MNZ24] Tony Metger, Anand Natarajan, and Tina Zhang. Succinct arguments for qma from standard assumptions via compiled nonlocal games, 2024.

[MV21] Tony Metger and Thomas Vidick. Self-testing of a single quantum device under computational assumptions. *Quantum*, 5:544, September 2021.

[Nao91] Moni Naor. Bit commitment using pseudorandomness. *J. Cryptol.*, 4(2):151–158, January 1991.

[NV17] Anand Natarajan and Thomas Vidick. A quantum linearity test for robustly verifying entanglement. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2017, page 1003–1015, New York, NY, USA, 2017. Association for Computing Machinery.

[NZ23] Anand Natarajan and Tina Zhang. Bounding the quantum value of compiled nonlocal games: From chsh to bqp verification. *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1342–1348, 2023.

[RS19] Roy Radian and Or Sattath. Semi-quantum money. In *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*, AFT '19, page 132–146, New York, NY, USA, 2019. Association for Computing Machinery.

[RUV13] Ben W. Reichardt, Falk Unger, and Umesh Vazirani. A classical leash for a quantum system: command of quantum systems via rigidity of chsh games. In *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science*, ITCS '13, page 321–322, New York, NY, USA, 2013. Association for Computing Machinery.

[Shm22] Omri Shmueli. Public-key quantum money with a classical bank. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2022, page 790–803, New York, NY, USA, 2022. Association for Computing Machinery.

[Sho97] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, oct 1997.

[Vid20] Thomas Vidick. Interactions with quantum devices (course), 2020. http://users.cms.caltech.edu/~vidick/teaching/fsmp/fsmp.pdf.

[Win99] Andreas J. Winter. Coding theorem and strong converse for quantum channels. *IEEE Trans. Inf. Theory*, 45(7):2481–2485, 1999.

[YZ24] Takashi Yamakawa and Mark Zhandry. Verifiable quantum advantage without structure. *J. ACM*, 71(3), June 2024.

[Zha22] Jiayu Zhang. Classical verification of quantum computations in linear time. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 46–57, 2022.

A Adaptation of the proof from [NZ23]

In this section, we overview how the proof of Theorem 6.26 goes, following arguments made in [NZ23]. Fix some computationally nonlocal strategy $\mathscr{C}$ that associates to every $G \in \mathcal{G}[\mathsf{NO}]$ a set

$$\left(|\psi_G\rangle , \left\{ A^{q_A}_{s_A,G} \right\}_{s_A} , \left\{ B^{q_B}_{s_B,G} \right\}_{s_B} \right) ,$$

and fix any sequence $\{G[H]_\lambda\}_\lambda$ where $G[H]_\lambda \in \mathcal{G}[\mathsf{NO}]_\lambda$ for each $\lambda \in \mathbb{N}$. From now on, we will drop the parameterization by λ and refer to a single game $G[H]$ parameterized by a Hamiltonian H, as well as a fixed strategy $(|\psi\rangle, \{A^{q_A}_{s_A}\}_{s_A}, \{B^{q_B}_{s_B}\}_{s_B})$, when we really mean an infinite sequence of games, Hamiltonians and strategies. To prove the theorem, we must show that

$$\mathop{\mathbb{E}}_{(q_A,q_B)\leftarrow Q} \sum_{s_A,s_B} V(q_A,q_B,s_A,s_B) \langle\psi| A^{q_A,\dagger}_{s_A} B^{q_B}_{s_B} A^{q_A}_{s_A} |\psi\rangle \leq \omega_C + \mathrm{negl}(\lambda).$$

Following [NZ23], we now introduce some notation. First, for any observable O, we define $\langle O\rangle \coloneqq \langle\psi| O |\psi\rangle$, where $|\psi\rangle$ is the state defined by the strategy we fixed above. Next, since $q_B \in \{0,1\}$ is just a single bit, and $s_B \in \{0,1\}^\lambda$ is a λ-bit string, we will define

$$\{Z_\gamma\}_{\gamma\in\{0,1\}^\lambda} \coloneqq \{B^0_{s_B}\}_{s_B}, \quad ,\{X_\gamma\}_{\gamma\in\{0,1\}^\lambda} \coloneqq \{B^1_{s_B}\}_{s_B},$$

and assume wlog that there exist unitaries U_Z, U_X such that

$$\{Z_\gamma\}_\gamma = \{U_Z^\dagger(|\gamma\rangle\langle\gamma| \otimes I)U_Z\}_\gamma, \quad \{X_\gamma\}_\gamma = \{U_X^\dagger(|\gamma\rangle\langle\gamma| \otimes I)U_X\}_\gamma.$$

Finally, we define sets of binary observables $\{Z(a)\}_{a\in\{0,1\}^\lambda}, \{X(a)\}_{a\in\{0,1\}^\lambda}$ as follows:

$$\begin{aligned} Z(a) &\coloneqq \sum_\gamma (-1)^{a\cdot\gamma} U_Z^\dagger(|\gamma\rangle\langle\gamma| \otimes I)U_Z, \\ X(a) &\coloneqq \sum_\gamma (-1)^{a\cdot\gamma} U_X^\dagger(|\gamma\rangle\langle\gamma| \otimes I)U_X. \end{aligned}$$

Now, we adapt several lemmas from [NZ23].

Lemma A.1 (Adaptation of Lemma 36 from [NZ23]). *Suppose the strategy succeeds in the CHSH subtest with probability at least $\omega_{\mathsf{CHSH}} - \epsilon$. Then*

$$\mathop{\mathbb{E}}_{\substack{(a,b)\leftarrow D^1_Q,\\ q_A \coloneqq (\mathsf{CHSH},(a,b,0))}} \sum_{s_A} \langle A^{q_A,\dagger}_{s_A} \cdot |\{Z(a), X(b)\}|^2 \cdot A^{q_A}_{s_A}\rangle \leq O(\epsilon).$$

Proof. Following [NZ23], for any fixed a, b, this can be seen as an instance of a computationally nonlocal strategy applied to the CHSH game. Thus, the claim follows from Theorem 6.25. □

Lemma A.2 (Adaptation of Lemma 37 from [NZ23]). *Suppose the strategy succeeds in the commutation subtest with probability at least $1 - \epsilon$. Then*

$$\mathop{\mathbb{E}}_{\substack{(a,b)\leftarrow D^0_Q,\\ q_A \coloneqq (\mathsf{Commutation},(a,b))}} \sum_{s_A} \langle A^{q_A,\dagger}_{s_A} \cdot |[Z(a), X(b)]|^2 \cdot A^{q_A}_{s_A}\rangle \leq O(\epsilon).$$

Proof. Again, for any fixed a, b, this can be seen as an instance of a computationally nonlocal strategy applied to the commutation game described in [NZ23, Section 3]. Thus, this follows from [NZ23, Lemma 23], which analyzes the commutation game. Since this analysis does not use the blindness of QFHE at all (which is not required because the commutation game has no Alice question), there is no change to the proof in our setting. □

Lemma A.3 (Adaptation of Lemma 38 from [NZ23]). *Suppose the strategy succeeds in the CHSH subtest with probability at least $\omega_{\mathsf{CHSH}} = \epsilon$, and in the commutation subtest with probability at least $1 - \epsilon$. Then*

$$\mathop{\mathbb{E}}_{\substack{(a,b)\leftarrow D_Q,\\ q_A:=\mathsf{Tel}}} \sum_{s_A} \langle A_{s_A}^{q_A,\dagger} \cdot |(-1)^{a\cdot b} Z(a)X(b) - X(b)Z(a)|^2 \cdot A_{s_A}^{q_A} \rangle \leq O(\epsilon) + \mathrm{negl}(\lambda).$$

Proof. Here, we crucially use the fact that the strategy is computationally nonlocal to switch the Alice questions in the above lemmas to Tel. Indeed, this lemma is implied by Lemma A.1 and Lemma A.2 by following the proof of [NZ23, Lemma 38], where equations (192) and (195) follow in our setting from the fact that the strategy is computationally nonlocal. □

Lemma A.4 (Adaptation of Lemma 39 from [NZ23]). *For any $u_1, u_2 \in \{0,1\}$, it holds that*

$$\mathop{\mathbb{E}}_{\substack{(a,b=(e_i+e_j))\leftarrow D_Q,\\ q_A:=\mathsf{Tel}}} \sum_{\substack{s_A:(s_A)_i=u_1,\\ (s_A)_j=u_2}} \langle A_{s_A}^{q_A,\dagger} \cdot |(-1)^{a\cdot b} Z(a)X(b)Z(a) - X(b)| \cdot A_{s_A}^{q_A} \rangle \leq O(\epsilon^{1/2}) + \mathrm{negl}(\lambda).$$

Proof. This is implied by Lemma A.3 by following the proof of [NZ23, Lemma 39] with no changes. □

Next, we import the following definitions.

- **SWAP isometry.** Let $\mathcal{H}_{\mathcal{Y}}$ and $\mathcal{H}_{\mathcal{Z}}$ be two copies of $(\mathbb{C}^2)^{\otimes\lambda}$. The λ-qubit SWAP isometry $V : \mathcal{H}_{\mathcal{B}} \to \mathcal{H}_{\mathcal{B}} \otimes \mathcal{H}_{\mathcal{Y}} \otimes \mathcal{H}_{\mathcal{Z}}$ is defined by the following expression:

$$V |\phi\rangle = \left(\frac{1}{2^\lambda} \sum_{u,v\in\{0,1\}^\lambda} Z(u)X(v) \otimes I \otimes \sigma_Z(u)\sigma_X(v) \right) |\phi\rangle |+\rangle^{\otimes\lambda}.$$

- **Expected verifier outcomes.** Let H_X be H restricted to XX terms. Let $\hat{\mathbb{E}}[H_X]$ be the expected value of the meausurement outcome computed by the verifier in a teleport round, conditioned on 1) $w = q_B$, so the verifier perform an energy check instead of accepting automatically, and 2) the verifier choosing an XX term to check. Then

$$\hat{\mathbb{E}}[H_X] = \sum_{u_1,u_2\in\{0,1\}} (-1)^{u_1+u_2} \mathop{\mathbb{E}}_{\substack{(b=e_i+e_j)\leftarrow D_X,\\ q_A:=\mathsf{Tel}}} \sum_{\substack{s_A:(s_A)_i=u_1,\\ (s_A)_j=u_2}} \langle A_{s_A}^{q_A,\dagger} X(b) A_{s_A}^{q_A} \rangle.$$

  Define $\hat{\mathbb{E}}[H_Z]$ analogously, which gives

$$\hat{\mathbb{E}}[H_Z] = \sum_{v_1,v_2\in\{0,1\}} (-1)^{v_1+v_2} \mathop{\mathbb{E}}_{\substack{(a=e_i+e_j)\leftarrow D_Z,\\ q_A:=\mathsf{Tel}}} \sum_{\substack{s_A:(s_A)_{\lambda+i}=v_1,\\ (s_A)_{\lambda+j}=v_2}} \langle A_{s_A}^{q_A,\dagger} Z(a) A_{s_A}^{q_A} \rangle.$$

Lemma A.5 (Adaptation of Lemmas 43 and 44 from [NZ23]). *Define*

$$\rho_{s_A} := \mathrm{Tr}_{\mathcal{B},\mathcal{Z}}[V A_{s_A}^{\mathsf{Tel}} |\psi\rangle\langle\psi| A_{s_A}^{\mathsf{Tel},\dagger} V^\dagger].$$

Then, assuming the strategy succeeds with probability $\omega_{\mathsf{CHSH}} - \epsilon$ in the CHSH subtest and with probability $1-\epsilon$ in the commutation subtest,

$$\sum_{u_1,u_2}(-1)^{u_1+u_2}\sum_{\substack{s_A:(s_A)_i=u_1,\\(s_A)_j=u_2}}\mathop{\mathbb{E}}_{b\leftarrow D_X}\mathrm{Tr}[\sigma_X(b)\rho_{s_A}]\approx_{O(\epsilon^{1/2})+\mathrm{negl}(\lambda)}\hat{\mathbb{E}}[H_X].$$

Moreover,

$$\sum_{v_1,v_2}(-1)^{v_1+v_2}\sum_{\substack{s_A:(s_A)_{\lambda+i}=v_1,\\(s_A)_{\lambda+j}=v_2}}\mathop{\mathbb{E}}_{a\leftarrow D_Z}\mathrm{Tr}[\sigma_Z(a)\rho_{s_A}]=\hat{\mathbb{E}}[H_Z].$$

Proof. This is implied by Lemma A.4 by following the proofs of [NZ23, Lemma 43] and [NZ23, Lemma 44] with no changes. □

Lemma A.6 (Adaptation of Lemma 45 from [NZ23])**.** *Assuming the strategy succeeds with probability $\omega_{\mathsf{CHSH}} - \epsilon$ in the CHSH subtest and with probability $1-\epsilon$ in the commutation subtest, there exists a state ρ such that*

$$\mathop{\mathbb{E}}_{a\leftarrow D_Z}\mathrm{Tr}[\rho_Z(a)\rho]=\hat{\mathbb{E}}[H_Z],$$
$$\mathop{\mathbb{E}}_{b\leftarrow D_X}\mathrm{Tr}[\rho_X(b)\rho]\approx_{O(\epsilon^{1/2})+\mathrm{negl}(\lambda)}\hat{\mathbb{E}}[H_X].$$

Proof. This is implied by Lemma A.5 by following the proof of [NZ23, Lemma 45] with no changes. □

Now, following analysis in the proof of [NZ23, Theorem 46] and assuming for contradiction that the strategy succeeds with probability greater than

$$\begin{aligned}&\frac{1}{2}(1-\kappa)(1+\omega_{\mathsf{CHSH}})+\kappa(1-\frac{1}{4}\alpha)-\frac{1}{8}\kappa(\beta-\alpha)\\=&\frac{1}{2}(1-\kappa)(1+\omega_{\mathsf{CHSH}})+\kappa(1-\frac{1}{4}\beta)+\frac{1}{8}\kappa(\beta-\alpha),\end{aligned}$$

we can conclude that for an appropriate choice of $\kappa=\Theta((\beta-\alpha)^2)$, Lemma A.6 implies that there exists a state ρ such that $\mathrm{Tr}[H\rho]>\beta$, which gives a contradiction.

B Inefficiently extractable commitments

In this section, we define (post-quantum) classical inefficiently-extractable commitments, which can be constructed from any post-quantum one-way function [Nao91].

Definition B.1 (Inefficiently-extractable commitment)**.** *An inefficiently-extractable commitment between a classical committer and classical receiver consists of an interaction*

$$(\mathsf{st}_{\mathsf{Com}},\mathsf{st}_{\mathsf{Rec}},\tau)\leftarrow\langle\mathsf{Com}(1^\lambda,b),\mathsf{Rec}(1^\lambda)\rangle,$$

where τ is the (classical) transcript of interaction produced by the protocol, along with algorithms (Open, Ver) *with the following syntax.*

- $\mathsf{Open}(\mathsf{st}_{\mathsf{Com}}) \to (b, w)$ *is a PPT algorithm that takes as input the committer's state* $\mathsf{st}_{\mathsf{Com}}$ *and produces a bit* b *and opening information* w*.*
- $\mathsf{Ver}(\mathsf{st}_{\mathsf{Rec}}, b, w) \to \{\top, \bot\}$ *is a PPT algorithm that takes as input the receiver's state* $\mathsf{st}_{\mathsf{Rec}}$*, a bit* b*, and opening information* w*, and either accepts or rejects.*

Correctness *requires that for any* $b \in \{0, 1\}$*,*

$$\Pr\left[\mathsf{Ver}(\mathsf{st}_{\mathsf{Rec}}, b, w) = \top : \begin{array}{r} (\mathsf{st}_{\mathsf{Com}}, \mathsf{st}_{\mathsf{Rec}}, \tau) \leftarrow \langle \mathsf{Com}(1^\lambda, b), \mathsf{Rec}(1^\lambda) \rangle \\ (b, w) \leftarrow \mathsf{Open}(\mathsf{st}_{\mathsf{Com}}) \end{array}\right] = 1 - \mathrm{negl}(\lambda).$$

The commitment satisfies (post-quantum) computational hiding *if for any QPT adversarial receiver* $\{\mathsf{Adv}_\lambda\}_{\lambda \in \mathbb{N}}$*,*

$$\begin{vmatrix} \Pr\left[b_{\mathsf{Adv}} = 0 : (\mathsf{st}_{\mathsf{Com}}, b_{\mathsf{Adv}}) \leftarrow \langle \mathsf{Com}(1^\lambda, 0), \mathsf{Adv}_\lambda \rangle\right] \\ - \Pr\left[b_{\mathsf{Adv}} = 0 : (\mathsf{st}_{\mathsf{Com}}, b_{\mathsf{Adv}}) \leftarrow \langle \mathsf{Com}(1^\lambda, 1), \mathsf{Adv}_\lambda \rangle\right] \end{vmatrix} = \mathrm{negl}(\lambda).$$

The commitment is inefficiently extractable *if there exists an unbounded extractor* Ext *such that for any unbounded adversarial committer* Adv*,*

$$\Pr\left[\mathsf{Ver}(\mathsf{st}_{\mathsf{Rec}}, 1 - b, w) = \top : \begin{array}{r} (\mathsf{st}_{\mathsf{Adv}}, \mathsf{st}_{\mathsf{Rec}}, \tau) \leftarrow \langle \mathsf{Adv}, \mathsf{Rec}(1^\lambda) \rangle \\ b \leftarrow \mathsf{Ext}(\tau) \\ w \leftarrow \mathsf{Adv}(\mathsf{st}_{\mathsf{Adv}}, 1 - b) \end{array}\right] = \mathrm{negl}(\lambda).$$