

Classical Commitments to Quantum States

Sam Gunn^{*1}, Yael Tauman Kalai^{†2}, Anand Natarajan^{‡2}, and Ági Villányi^{§2}

¹UC Berkeley

²MIT

April 19, 2024

Abstract

We define the notion of a classical commitment scheme to quantum states, which allows a quantum prover to compute a classical commitment to a quantum state, and later open each qubit of the state in either the standard or the Hadamard basis. Our notion is a strengthening of the measurement protocol from Mahadev (STOC 2018). We construct such a commitment scheme from the post-quantum Learning With Errors (LWE) assumption, and more generally from any noisy trapdoor claw-free function family that has the distributional strong adaptive hardcore bit property (a property that we define in this work).

Our scheme is *succinct* in the sense that the running time of the verifier in the commitment phase depends only on the security parameter (independent of the size of the committed state), and its running time in the opening phase grows only with the number of qubits that are being opened (and the security parameter). As a corollary we obtain a classical succinct argument system for **QMA** under the post-quantum LWE assumption. Previously, this was only known assuming post-quantum secure indistinguishability obfuscation. As an additional corollary we obtain a generic way of converting any X/Z quantum PCP into a succinct argument system under the quantum hardness of LWE.

*gunn@berkeley.edu

†tauman@mit.edu

‡anandn@mit.edu

§agivilla@mit.edu

Contents

1	Introduction	3
1.1	The Definition	6
1.2	The Construction	7
1.3	Applications	10
1.4	Related Works	10
2	Technical Overview	11
2.1	Mahadev’s measurement protocol	15
2.2	Our Single-Qubit Commitment Scheme	16
2.3	Succinct commitments	18
2.4	Applications	19
3	Preliminaries	20
3.1	Quantum information facts	22
3.2	Hash Family with Local Opening	22
3.3	Noisy Trapdoor Claw-Free Functions	24
4	The Distributional Strong Adaptive Hardcore Bit Property	26
5	Classical Commitments to Quantum States	31
5.1	Syntax	31
5.1.1	Syntax for Succinct Commitments	33
5.2	Properties	35
5.2.1	Correctness	35
5.2.2	Binding	36
6	Constructions	39
6.1	Construction for Single Qubit States	39
6.2	Construction of Commitments for Multi-Qubit States	41
6.3	Construction of Succinct Multi-Qubit Commitments	43
6.3.1	Construction	44
7	Analysis of the Multi-Qubit Commitment Schemes from Section 6	46
7.1	Correctness	46
7.2	Binding	49
7.3	Binding for the Succinct Commitment Scheme	72
8	Applications	79
8.1	Succinct Interactive Arguments for QMA	79
8.2	Succinct Interactive Arguments from X/Z Quantum PCPs	82
9	Acknowledgements	84
A	Weak commitments to Quantum States (WCQ)	87
B	Proof of Lemma 4.5	89

1 Introduction

A commitment scheme is one of the most basic primitives in classical cryptography, with far reaching applications ranging from zero-knowledge proofs [GMW86, BCC88], identification schemes and signature schemes [FS87], secure multi-party computation protocols [GMW87, CDv88], and succinct arguments [Mic94]. There is a long history of studying commitments to *classical* information, both in the classical and post-quantum worlds. In this work, we focus on the less-studied question of committing to *quantum states*. This notion was first systematically explored in a recent work by Gunn et al. [GJMZ22], who defined and constructed a commitment scheme for quantum states using quantum messages. In this work, we study the existence of *classical* commitments to quantum states, where all messages (the commitment and the opening) are classical, and the receiver is a classical machine. Our major contributions are a definition of a classical commitment to quantum states, a construction based on the post-quantum Learning With Errors (LWE) assumption, and a construction of a *succinct* commitment to quantum states (analogous to Merkle hashing in the classical setting [Mer87]), also under post-quantum LWE.¹ As an immediate application, we obtain a succinct classical argument system for **QMA** based only on post-quantum hardness of LWE, improving on previous work which required indistinguishability obfuscation [BKL⁺22]. To our knowledge, our work constitutes the first work to define a binding classical commitment to quantum states, and to give a construction that achieves this.

Our construction builds directly on the seminal *measurement protocol* of Mahadev [Mah18], which was used by her to construct the first classical argument system for **QMA**. Loosely speaking, a measurement protocol is a way for a classical verifier to request a quantum prover to measure each qubit of a quantum state (of the prover’s choice) in the X or Z basis with the guarantee that the prover’s opening must be “consistent with a quantum state.” This motivates our definition of a classical *opening* of a quantum state: the receiver should be able to request the sender to open each qubit of the committed state in either the X or Z basis. (One could imagine asking for openings in more general bases, but these two seem to be a desirable minimum.) However, a measurement protocol does not automatically give rise to a commitment, for several reasons. First, there is a major structural difference: in a measurement protocol, all phases of the protocol—even the keys chosen in the initial setup—may depend on the choice of opening basis! (Indeed, in Mahadev’s protocol, the keys consist of either “2-to-1” or “injective” claw-free functions depending on the basis to be measured.) This is far from what we would like in a commitment: the initial “commitment” phase should be *completely* independent of the basis in which the receiver ultimately chooses to request an opening.

Thus, the first step to building our construction is to convert Mahadev’s measurement protocol into something having the syntax of a commitment², and henceforth we refer to this modified protocol as Mahadev’s “weak” commitment³. In the most basic version of this protocol, a quantum sender holding a qubit in state $|\psi\rangle$ interacts with a classical receiver, sending a classical message that commits to $|\psi\rangle$. Later, the sender is requested by the receiver to “open” the committed qubit in either

¹More generally, our constructions are based on the existence of a (noisy) claw-free trapdoor function family with a distributional strong adaptive hard-core bit property, which in particular can be instantiated under the LWE assumption.

²Technically, we do this by always using the “2-to-1” mode of the claw-free function. Moreover, we do not even rely on the existence of a dual-mode (as was done by Mahadev [Mah18]), and simply use a “2-to-1” claw-free family.

³We refer to it as a weak commitment since (as we elaborate on below) it does not have the desired binding property.

the standard or the Hadamard basis. To open, the sender performs an appropriate measurement and returns the outcome, which can be *decoded* by the receiver (using a cryptographic trapdoor), to obtain an outcome from measuring $|\psi\rangle$ in the appropriate basis.

A commitment scheme must be *binding*, meaning that the sender cannot change their mind about the committed state once the commitment has been sent. It turns out that the modified Mahadev scheme is a “weak” commitment because it partially satisfies the binding property: it is binding in the standard basis, but *not at all* binding in the Hadamard basis. In fact, the sender, after committing to $|+\rangle$, can always freely change the committed state to $|-\rangle$ without ever being detected! Relatedly, in the modified Mahadev scheme, the receiver performs a test on the opening in the standard basis case, and only accepts the opening if it is valid, but performs no test in the Hadamard case.

Motivated by this observation, we show that a simple twist on Mahadev’s weak commitment is truly binding (in a rigorous sense which we define) in both bases. We elaborate on our binding definition in Sections 1.1 and 2, and on our construction in Sections 1.2 and 2, and below only give a teaser of it. In our construction, the sender first commits to $|\psi\rangle$ under Mahadev’s weak commitment, generating a commitment string y_0 and a (multi-qubit) post-commitment state $|\psi_1\rangle$. It then coherently *opens* this state in the Hadamard basis—that is, it executes a unitary version of the opening algorithm, but does not perform the final measurement, instead producing a quantum state $|\psi_1\rangle$. Finally, the sender applies Mahadev’s weak commitment *again* to the state $|\psi_1\rangle$, qubit-by-qubit, obtaining a vector of commitment strings \vec{y} and a post-commitment state $|\psi_2\rangle$. The strings (y_0, \vec{y}) now constitute a classical commitment to the state $|\psi\rangle$. To open this commitment in the Hadamard basis, the sender simply applies the standard basis opening procedure for the second Mahadev commitment, yielding a string z which the receiver will test and decode using the commitment vector \vec{y} . By the standard-basis binding of Mahadev’s commitment, we are guaranteed that the decoded outcome from z —assuming the test passes—yields the same result as measuring $|\psi_1\rangle$ in the standard basis, and by construction, this gives a Hadamard-basis opening of $|\psi\rangle$, which it can then decode using the commitment string y_0 . But how do we open the commitment in the standard basis? It is far from obvious that this is even possible! For this we exploit specific features of the Mahadev scheme—in particular, the fact that the opening procedure is “native”: Opening in the standard basis constitutes measuring the registers in the standard basis, and opening in the Hadamard basis constitutes measuring the registers in the Hadamard basis. This fact is useful both to argue that the opening is correct and to prove that the binding property is achieved. We note that in our new scheme the verifier tests the validity of both the standard basis opening and the Hadamard basis opening, and decodes both opening using the cryptographic trapdoor.⁴

Our basic construction for a single qubit can be extended to states with any number of qubits to get a *non-succinct* commitment to a quantum state. We next ask whether our commitment scheme can be made *succinct*: can the sender commit to an ℓ qubit state, and open to a small number of these qubits, by exchanging much fewer than ℓ bits with the receiver? Here, already in the case of “weak” commitments, there is a significant technical obstacle with just the *very first message* from the receiver to the sender: openings in Mahadev’s scheme can leak information about the secret key, so each committed qubit must use a fresh secret key to maintain any security at all. This means that already in the initial key-exchange phase, the receiver must send the sender $\geq \ell$

⁴We mention that in Mahadev’s scheme, the verifier only tests the validity of the standard basis opening, and this test, as well as the decoding, is done publicly (without the trapdoor). The verifier uses the trapdoor only to decode the Hadamard basis opening, which it did not test.

bits. We show that, surprisingly, the “strong” binding property of our commitment, together with specific properties of the underlying (noisy) trapdoor claw-free family, allows us to overcome this barrier. Namely, we show that strong binding, together with specific properties of the underlying (noisy) trapdoor claw-free family, implies that the openings do not leak information about the key in our scheme, allowing us to use the same key for all committed qubits. We emphasize that, even to obtain a succinct “weak” commitment, or a succinct measurement protocol, the only route we know of using standard (post-quantum) cryptographic assumptions is through our strongly binding commitments! We view this as an interesting indication of the possible usefulness of our strong binding property in further applications.

As a teaser for how exactly the leakage occurs, and how we avoid it, for now we remark that in the Mahadev weak commitment, the adversary can cause the receiver to generate outputs of the form $d' \cdot s$, for known vectors d' of its choice, where s is the secret. This means that the output for sufficiently many qubits, may leak the secret s . For an honest sender, this would not be an issue because the vectors d' would be obtained by a quantum measurement with unpredictable answers, and thus have high min-entropy. We show that in our scheme, even *dishonest* senders are forced to produce d' with (sufficient) min-entropy, because of the additional tests done in our opening procedure. This is what prevents the outcomes from leaking information about s .

Reusing the key directly only gives us a short first message, which yields a “semi-succinct” commitment, in which messages from the receiver are short, but messages from the sender are long. In fact, this already yields an application of our results: a *fully-succinct* classical argument system for **QMA** which is secure assuming post quantum security of **LWE**. We obtain this by following the template of Bartusek et al. [BKL⁺22], but replacing their use of Mahadev’s measurement protocol with our succinct commitment.

Theorem 1.1 (Informal). *There exists a (classical) succinct interactive argument for **QMA** under the post-quantum Learning With Errors (LWE) assumption.*⁵

This improves on the result of [BKL⁺22] in terms of cryptographic assumptions: they required the assumption of post-quantum indistinguishability obfuscation (iO) to succinctly generate ℓ keys for Mahadev’s protocol, whereas our protocol only requires the post-quantum security of **LWE**. It is currently not known how to deduce post-quantum iO from *any* standard cryptographic assumptions, whereas **LWE** is the “paradigmatic” post-quantum cryptographic assumption.

To construct a succinct argument system for **QMA**, the approach we and [BKL⁺22] both follow is to construct a semi-succinct argument system, and then make it fully succinct by composing with (state-preserving) post-quantum interactive arguments of knowledge [CMSZ21, LMS22]. It turns out that the same tools let us construct outright a fully succinct commitment scheme: for this to be meaningful, we imagine that the sender only opens to a small number of qubits chosen by the receiver, rather than to all of the qubits. In classical cryptography, succinct commitments are natural partners of PCPs, as they enable a verifier to delegate the task of checking a PCP to the prover. While quantum PCPs do not currently exist, we hope that our succinct commitment can be paired with a suitable future PCP to design interesting protocols.

⁵More generally, assuming the existence of a (noisy) trapdoor claw free function family with a distributional strong adaptive hard-core bit property, which we elaborate on later on.

1.1 The Definition

Defining a non-succinct commitment scheme Our definition of a (non-succinct) commitment scheme is a natural extension of the classical counterpart. It consists of a key generation algorithm **Gen** that takes as input the security parameter 1^λ and a length parameter 1^ℓ and outputs a pair of public and secret keys $(\mathbf{pk}, \mathbf{sk})$; a commit algorithm **Commit** that takes as input a public key \mathbf{pk} and an ℓ -qubit quantum state σ and outputs a classical string \mathbf{y} and a post-commitment state ρ , where \mathbf{y} is the commitment to the quantum state σ ;⁶ an open algorithm **Open** that takes as input the post-commitment state ρ and a basis choice $\mathbf{b} = (b_1, \dots, b_\ell) \in \{0, 1\}^\ell$, where $b_i = 0$ corresponds to opening the i 'th qubit in the standard basis and $b_i = 1$ corresponds to opening the i 'th qubit in the Hadamard basis, and outputs an opening $\mathbf{z} \in \{0, 1\}^{\ell \cdot \text{poly}(\lambda)}$; and the final algorithm **Out** that takes as input a secret key \mathbf{sk} , a commitment string \mathbf{y} , a basis choice $\mathbf{b} \in \{0, 1\}^\ell$ and an opening \mathbf{z} , and outputs the measurement result $\mathbf{m} \in \{0, 1\}^\ell$ or \perp if the opening is rejected.⁷

We mention that the above syntax yields a commitment scheme that is *privately verifiable* in the sense that \mathbf{sk} is needed to decode the measurement value \mathbf{m} from the opening value \mathbf{z} . While it would be desirable to construct a commitment scheme that is publicly verifiable, where **Gen** only generates a public key \mathbf{pk} , and this public key is used by the opening algorithm to generate the output \mathbf{m} along with an opening \mathbf{z} which can be verified given \mathbf{pk} , we believe that this public key variant is impossible to achieve. This impossibility was formalized on the quantum setting (i.e., where the commitment is a quantum state) by [GJMZ22], and we leave it as an open problem to prove the impossibility in the classical setting.

We require two properties from our commitment scheme: completeness and binding. We note that for commitments to classical strings it is common to require a *hiding* property. We do not require it since one can easily obtain hiding by committing to the commitment string \mathbf{y} using a classical commitment scheme (that is binding and hiding).

- **Correctness.** The correctness property asserts that if an honest committer commits to an ℓ -qubit state σ then for any basis choice $\mathbf{b} \in \{0, 1\}^\ell$, the algorithm **Out**, applied to the opening string \mathbf{z} generated by **Open**, yields an output \mathbf{m} whose distribution is statistically close to the distribution obtained by simply measuring σ in the basis \mathbf{b} .
- **Binding.** Loosely speaking, the binding property asserts that for any (possibly malicious) QPT algorithm **Commit*** that commits to an ℓ -qubit quantum state, there is a *single* extracted quantum state τ such that for *any* QPT algorithm **Open*** and *any* basis (b_1, \dots, b_ℓ) , where $b_i = 0$ corresponds to measuring the i 'th qubit in the standard basis and $b_i = 1$ corresponds to measuring it in the Hadamard basis, the output obtained by **Open*** (b_1, \dots, b_ℓ) is computationally indistinguishable from measuring τ in basis (b_1, \dots, b_ℓ) , assuming **Open*** is always accepted. We relax the requirement that **Open*** is always accepted, and allow **Open*** to be rejected with probability δ at the price of the two distributions being $O(\sqrt{\delta})$ -computationally indistinguishable. We elaborate on the binding property in Section 2.

⁶We note that both the length of \mathbf{pk} and the length of the commitment string \mathbf{y} may grow polynomially with the length ℓ of the committed state σ .

⁷We note that in the actual definition we partition this algorithm into two parts: **Ver** and **Out** where the former only outputs a bit indicating if the opening is valid or not and the latter outputs the actual opening if valid. This partition is only for convenience.

Comparison with Mahadev’s measurement protocol. Our commitment scheme is stronger than that of a *measurement protocol*, originally considered in [Mah18] and formally defined in [BKL⁺22]. Beyond the syntactic difference, where in a measurement protocol the opening basis must be determined during the key generation phase (and the key generation algorithm takes as input the basis $\mathbf{b} \in \{0, 1\}^\ell$), our binding property is significantly stronger. A measurement protocol guarantees that any (possibly malicious) QPT algorithm Open^* must be consistent with an ℓ -qubit state, but different opening algorithms can be consistent with different quantum states.

Defining a succinct commitment scheme. The syntax for a succinct commitment differs quite substantially from the syntax of a non-succinct commitment described above. First, Gen only takes as input the security parameter 1^λ (and does not take as input the length parameter 1^ℓ); in addition, Commit is required to output a succinct commitment of size $\text{poly}(\lambda)$. However, there is a more substantial difference which stems from the fact that, similarly to the non-succinct variant, we require a succinct commitment to have an extraction property that asserts that one can extract an ℓ -qubit quantum state τ such that the output distribution of any successful opening is indistinguishable from measuring τ . Since in this setting we consider opening algorithms that only open a few of the qubits, there is no way we can extract an ℓ -qubit state from such algorithms. As a remedy, we add an *interactive test phase*. This test phase is executed with probability $1/2$, and if executed then at the end of it the verifier outputs 0 or 1, indicating accept or reject, and the protocol terminates without further executing the opening phase, since the test protocol destroys the state. We note that Mahadev’s measurement protocol has a non-interactive test phase which is executed with probability $1/2$. In our setting this test phase is *interactive*. It is this interactive nature that allows us to extract a large state from a succinct protocol.

1.2 The Construction

Our construction: the single qubit case We construct the commitment scheme in stages. We first construct a *single-qubit* commitment scheme; this scheme is inspired by the construction from Mahadev [Mah18]. We elaborate on it in Section 2, but give a very high-level description here. First, let us recall Mahadev’s weak commitment for a single qubit. In this scheme, the sender receives a public key that enable it to evaluate a *two-to-one trapdoor claw-free (TCF) function* $f : \{0, 1\} \times \mathcal{X} \rightarrow \mathcal{Y}$.⁸ For every image $y \in \mathcal{Y}$, there are exactly two preimages, which have the form $(0, x_0)$ and $(1, x_1)$, where $x_0, x_1 \in \{0, 1\}^n$, but any such pair (called a “claw”) is cryptographically hard to find. In Mahadev’s scheme, to commit to a qubit in state $|\psi\rangle = \sum_{b \in \{0, 1\}} \alpha_b |b\rangle$, the sender first prepares

$$\sum_{b \in \{0, 1\}} \sum_{x \in \mathcal{X}} \alpha_b |b\rangle |x\rangle |f(b, x)\rangle,$$

and then measures the last register to obtain a random outcome y . The resulting state is the $(n + 1)$ -qubit state

$$\sum_{b \in \{0, 1\}} \alpha_b |b\rangle |x_b\rangle.$$

To open this in the standard basis, the honest sender measures in the standard basis and returns (b, x_b) ; the receiver checks that $f(b, x_b) = y$, and if so, records a measurement outcome of b .

⁸We mention that under the LWE assumption we only have a “noisy” TCF function family, which was constructed in [BCM⁺18]. We do not go into this technicality in the introduction and overview sections.

Intuitively, this constitutes a “binding” commitment in the standard basis because it is impossible for the sender to know both x_0 and x_1 , and thus impossible to flip between them. To open in the Hadamard basis, the honest sender measures in the *Hadamard* basis; a short calculation shows that the outcome is a random string $d \in \{0, 1\}^{n+1}$, where the probability that $d \cdot (1, x_0 \oplus x_1) \equiv 0 \pmod{2}$ is exactly equal to $|\alpha_0 + \alpha_1|^2/2$, the probability that a Hadamard basis measurement on the *original* state $|\psi\rangle$ would have yielded $+$. The receiver uses the cryptographic trapdoor to compute $d \cdot (1, x_0 \oplus x_1) \pmod{2}$ as the measurement outcome of the opening, and performs *no* test. This is not at all a binding commitment: indeed, the “commitments” to a Hadamard basis states $|\pm\rangle$ look like

$$|\pm\rangle \mapsto \frac{1}{\sqrt{2}}(|0\rangle |x_0\rangle \pm |1\rangle |x_1\rangle),$$

and one can easily map from one state to the other by applying a Pauli Z operator to the first qubit.

We now describe our modification to convert this weak commitment (denoted commit_W) into a binding commitment: simply apply a Hadamard transform to the post-commitment state, and then weakly commit again to the resulting n -qubit state, applying the Mahadev scheme qubit by qubit, with a new TCF function f_i for each qubit.

$$\begin{aligned} \sum_b \alpha_b |b\rangle &\mapsto^{\text{commit}_W \rightarrow y_0} \sum_b \alpha_b |b, x_b\rangle \\ &\mapsto^{H^{\otimes(n+1)}} \sum_{d \in \{0,1\}^{n+1}} \beta_d |d\rangle \\ &\mapsto^{\text{commit}_W \rightarrow y_1, \dots, y_{n+1}} \sum_d \beta_d |d_1, x'_{1,d_1}\rangle \dots |d_{n+1}, x'_{n+1,d_{n+1}}\rangle. \end{aligned}$$

Here, d_j denotes the j th bit of d , and $x'_{j,b}$ denotes the corresponding preimage of y_j under the TCF function f_j (so $f_j(b, x'_{j,b}) = y_j$).

Let us see how to open this commitment. It will be easier to start with the Hadamard basis: to open in this basis, the sender measures their state in the *standard* basis, and returns the string $(d_1, x_1, \dots, d_{n+1}, x_{n+1})$. The receiver checks that each (d_i, x_i) is a preimage of the corresponding y_i , and the records the measurement outcome as $(d_1, \dots, d_{n+1}) \cdot (1, x_0 \oplus x_1)$. To open in the standard basis, the sender measures their state in the *Hadamard* basis, obtaining a (long) string z , and the receiver converts this into a measurement outcome by applying the Mahadev procedure for the *Hadamard* basis. Specifically, it first splits z into equal blocks of size $n+1$, and applies the Mahadev Hadamard procedure on each block, to get $n+1$ bits m_1, \dots, m_{n+1} .

$$\begin{aligned} z &= (z_1, \dots, z_{n+1}) \\ &\mapsto (m_1 = z_1 \cdot (1, x'_{1,0} \oplus x'_{1,1}), \dots, m_{n+1} = z_{n+1} \cdot (1, x'_{n+1,0} \oplus x'_{n+1,1})) \end{aligned}$$

Now, this corresponds to the outcome of opening the weak commitment of $\sum_d \beta_d |d\rangle$ in the Hadamard basis. But this state in turn was equal to the Hadamard transform of $\sum_b \alpha_b |b, x_b\rangle$. Thus, the outcomes m_1, \dots, m_{n+1} should look like the outcome of measuring $\sum_b \alpha_b |b, x_b\rangle$ in the standard basis: that is, like a preimage of y_0 under the TCF function f ! Thus, the receiver tests the outcomes by checking that

$$f(m_1, \dots, m_{n+1}) = y_0,$$

and if this passes, it records m_1 as the measurement outcome.

At an intuitive level, what makes this commitment scheme binding is that the receiver performs a test in *both* bases. More formally, we show binding in two parts: (1) there exists a qubit state consistent with the openings reported by the sender, and (2) for any two opening algorithms, the openings they generate are statistically indistinguishable. The proof of (1) uses standard techniques from the analysis of Mahadev’s protocol—in particular, the “swap isometry” as presented in [Vid20], but the proof of (2) is new to our work. Our arguments are based on the *collapsing* property of the TCF functions used to generate y_1, \dots, y_n (in the Hadamard basis case), and y_0 (in the standard basis case). Jumping ahead, we note in the succinct setting the situation is reversed. We can obtain (2) basically “for free” from the non-succinct setting, whereas the proof of (1) incurs most of the technical burden in this work.

Our construction: multiple qubits, and succinctness From the single-qubit scheme described above, we construct a *non-succinct multi-qubit* commitment scheme, by committing qubit-by-qubit, and thus repeating the single-qubit construction ℓ -times, where ℓ is the number of qubits we wish to commit to. This transformation is generic and can be used to convert *any* single-qubit commitment scheme into a *non-succinct* multi-qubit one. We emphasize that in the resulting ℓ -qubit scheme, both the public-key and the commitment string grow with ℓ , since the former consists of ℓ public-keys and the latter consists of ℓ commitment strings, where each corresponds to the underlying single-qubit scheme. We then convert this scheme into a succinct commitment scheme. This is done in two stages:

1. **Stage 1:** Reuse the same public key, as opposed to choosing ℓ independent ones. Namely, the public key consists of a single public key \mathbf{pk} corresponding the underlying single-qubit commitment scheme. To commit to an ℓ -qubit state, commit qubit-by-qubit while using the same public key \mathbf{pk} . We refer to such a commitment scheme as *semi-succinct* since the public key is succinct but the commitment is not.

We note that while this construction is generic, the analysis is not. In general, reusing the same public-key may break the binding property. We prove that if we start with our specific single-qubit commitment scheme then the resulting semi-succinct multi-qubit scheme remains sound. We recall, that as mentioned above, if we start with Mahadev’s single qubit weak commitment protocol and convert it into a multi-qubit weak commitment while reusing the same public key, then the resulting measurement protocol becomes insecure. The reason is that a malicious sender may generate openings d in the Hadamard basis that cause the receiver’s “decoding” outcomes $d \cdot (1, x_0 \oplus x_1)$ to leak bits of \mathbf{sk} —recall that the receiver must use the secret key to decode, as x_0 and x_1 cannot be computed efficiently without it. Indeed, the TCF function family that we (and Mahadev) use is the LWE based construction due to [BCM⁺18], which has the property that $d \cdot (1, x_0 \oplus x_1) = d' \cdot s$, where s is a secret key⁹ and d' can be efficiently computed from d and x_0 . Once enough information about the secret s has been revealed, the scheme is no longer a secure measurement protocol, let alone a secure commitment: with knowledge of s , it becomes easy to distinguish the outcomes of the commitment from outcomes of measuring a true quantum state! Thus, to argue the security of our semi-succinct scheme, we must exploit specific properties of our single-qubit scheme. Indeed, we crucially use the

⁹In their construction the public key is an LWE tuple $(A, As + e)$. The secret key is actually a trapdoor of the matrix A but revealing the secret s is sufficient to break security.

binding property of our scheme which implies that the openings z reported by a successful sender must always have high min-entropy, which in our construction implies that d' has min-entropy. We then use a specific property of the underlying TCF function family from [BCM⁺18], which we call the “distributional strong adaptive hardcore bit” property. Roughly, this property ensures that if the opening d has min-entropy then $d \cdot (1, x_0 \oplus x_1)$ (which in their construction is equal to $d' \cdot s$) does not reveal information about sk .

2. **Stage 2:** Convert any semi-succinct commitment scheme into a succinct one. This part is generic and shows how to convert *any* semi-succinct commitment scheme into a succinct one. Our transformation is almost identical to that from [BKL⁺22], who showed how to convert any semi-succinct interactive argument (which is one where only the verifier’s communication is succinct, and where the prover’s communication can be long) into a fully succinct one. We elaborate on the high-level idea behind this transformation in Section 2.

1.3 Applications

We show how to use our succinct commitment scheme to construct succinct interactive argument for **QMA**. As a simpler bonus, we also use it show how to compile a hypothetical quantum PCP in “ X/Z form” into a succinct interactive argument. For the X/Z PCP compiler the idea is simple: In the succinct interactive argument the prover first succinctly commits to the X/Z PCP, then the verifier sends its X/Z queries and finally the prover opens the relevant qubits in the desired basis. The succinct interactive argument for **QMA** is more complicated, and follows the blueprint from [Mah18, BKL⁺22]. We elaborate on this in Section 2.1.

1.4 Related Works

Our work is inspired by the measurement protocol of Mahadev [Mah18], which has the same correctness guarantee as our commitment scheme. However, a measurement protocol (as was formally defined in [BKL⁺22]) does not require binding to hold; rather it only requires that an opening is consistent with a qubit. This qubit may be different for different opening algorithms. Indeed, the measurement protocol of Mahadev, as well as the ones from followup works, are not binding in the Hadamard basis. Mahadev uses this measurement protocol to construct classical interactive arguments for **QMA**. Mahadev’s measurement protocol, which was proven to be secure under the post-quantum LWE assumption, is a key ingredient in our construction.

Mahadev’s measurement protocol is not succinct. In a followup work, Bartusek et al. [BKL⁺22] constructed a succinct measurement protocol, by using Mahadev’s measurement protocol as a key ingredient, and thus obtaining a succinct classical interactive arguments for **QMA**. However the security of their protocol, and thus the soundness of the resulting **QMA** argument, relies on the existence of a post-quantum secure indistinguishable obfuscation scheme (in addition the post-quantum LWE assumption). We mention that Chia, Chung and Yamakawa [CCY20] also construct a succinct measurement protocol, which they use to obtain a succinct 2-message argument for **QMA**. However, in their scheme the prover and verifier share a polynomial-sized structured reference string (which requires a trusted setup to instantiate), and their security is heuristic.¹⁰

¹⁰More specifically, their scheme uses a hash function h , and it is proved to be secure when h is modeled as a random oracle, but the *protocol description itself* explicitly requires the code of h (i.e. uses h in a non-black-box way).

We improve upon these works by constructing a succinct classical commitment scheme for quantum states that guarantees binding (which is a stronger security condition than the one offered by a measurement protocol), based only on the post-quantum LWE assumption. As a result, we obtain a succinct classical interactive arguments for **QMA**, under the post-quantum LWE assumption. Our analysis makes use of techniques developed in [Mah18, Vid20, BKL⁺22], in addition to several new ideas that are needed to obtain our results.

We mention that our work, as well as all prior works mentioned above, require the receiver (a.k.a the verifier) to hold a secret key sk which is needed to decode the prover’s message and obtain the measurement output. We mention that the recent work of Bartusek et al. [BKNY23] considers the public-verifiable setting, where decoding can be done publicly. They construct a publicly verifiable measurement protocol in an oracle model, which is used as a building block in their obfuscation of pseudo-deterministic quantum circuits.

So far we only focused on prior work where the verifier (and hence the communication) is classical. We mention that recently Gunn et al. [GJMZ22] defined and constructed a *quantum* commitment scheme to quantum states, where *both* parties are quantum. In their setting, the quantum committer sends a quantum commitment to the receiver, and later opens by sending a quantum opening. The receiver then applies some unitary operation to recover the committed quantum state. This is in contrast to the classical setting where the receiver is classical and cannot hope to recover the committed quantum state, and instead only obtains an opening in a particular basis (standard or Hadamard). We mention that the quantum commitment scheme from [GJMZ22] relies on very weak cryptographic assumptions, and in particular, ones that are implied by the existence of one-way functions.

Finally, simultaneously and using different techniques from this work, a succinct argument system for **QMA** based on the assumption of quantum Fully Homomorphic Encryption (qFHE) was achieved by [MNZ24]. While both papers use common techniques from [BKL⁺22] to go from semi-succinctness to full succinctness, the core techniques are essentially disjoint. In particular, [MNZ24] does not use commitments to quantum states, but instead directly analyzes the soundness of the KLVY [KLVY22] compilation of a particular semi-succinct two-prover interactive proof for **QMA**. We leave it as an interesting open question for future work whether their result can yield an alternate construction of our primitive of quantum commitments.

Roadmap We refer the reader to Section 2 for the high-level overview of our techniques, to Section 3 for all the necessary preliminaries, to Section 5 for the formal definition of a succinct and non-succinct commitment scheme, to Section 6 for the constructions, to Section 7 for the analysis, and to Section 8 for the applications.

2 Technical Overview

In this section we describe the ideas behind our commitment schemes and their applications in more depth yet still informally. Our first contribution is defining the notion of a classical commitment scheme to quantum states. Let us start with the non-succinct version, and in particular the single-qubit case. As mentioned in the introduction, such a commitment scheme consists of algorithms

(Gen, Commit, Open, Out)

where **Gen** is a PPT algorithm that takes as input the security parameter 1^λ and outputs a pair of keys (pk, sk) ; **Commit** is a QPT algorithm that takes as input a public key pk and a single-qubit quantum state σ and outputs a classical commitment string \mathbf{y} and a post-commitment state ρ ; **Open** is a QPT algorithm that takes as input the post-commitment state ρ and a bit $b \in \{0, 1\}$, where $b = 0$ corresponds to a standard basis opening and $b = 1$ corresponds to a Hadamard basis opening, and outputs a classical opening \mathbf{z} ; and **Out** is a polynomial-time algorithm that takes as input the secret key sk , a commitment string \mathbf{y} , a basis $b \in \{0, 1\}$ and an opening \mathbf{z} and it outputs an element in $\{0, 1, \perp\}$.

We require the scheme to satisfy a correctness and a binding property. The correctness property is straightforward and was formalized in prior work [BKL⁺22]. It is the binding property that is tricky to formulate and achieve.

Defining Binding: the single qubit setting In the classical setting, the binding condition asserts that for any poly-size algorithm **Commit**^{*} that generates a commitment \mathbf{y} (to some classical string), and for every poly-size algorithms **Open**₁^{*} and **Open**₂^{*}, the probability that they successfully open to different strings is negligible. In the quantum setting the analogous property is the following: For any QPT algorithm **Commit**^{*} that generates a commitment \mathbf{y} (to a quantum state), and for QPT algorithms **Open**₁^{*} and **Open**₂^{*} (that are accepted with probability 1) and every basis choice $b \in \{0, 1\}$, the output distributions of **Open**₁^{*} and **Open**₂^{*} are statistically close or computationally indistinguishable. This is indeed one of the properties we require.¹¹ But this property on its own is not enough. We also need to ensure that the opening is consistent with some qubit. Namely, we require that there exists a QPT extractor **Ext** such that for every QPT algorithm **Open**^{*} (that is accepted with probability 1), **Ext** given black-box access to **Open**^{*} can extract from **Open**^{*} a quantum state τ such that for every basis $b \in \{0, 1\}$ the output of **Open**^{*} is computationally indistinguishable from measuring τ in basis b . We mention that this latter condition was formalized in [BKL⁺22] as a security property from a measurement protocol.

We construct a commitment scheme that achieves the above two properties. However, to make this definition meaningful we must consider opening algorithms that are accepted with probability smaller than 1. Indeed, we consider opening algorithms that are accepted with probability $1 - \delta$ and obtain $O(\sqrt{\delta})$ -indistinguishability in both the requirements above. We note that we can assume that **Open**^{*} is accepted with probability $1 - \delta$ by repeating the commitment protocol $\Omega(1/\delta)$ times (assuming the committer has many copies of the state they wish to commit to).

The multi-qubit setting So far we focused on the single-qubit setting. When generalizing the definitions to the multi-qubit setting we distinguish between the non-succinct setting and the succinct setting, starting with the former. The syntax can be generalized to the non-succinct multi-qubit setting in a straightforward way by committing and opening qubit-by-qubit. Generalizing the binding definition to the multi-qubit setting is a bit tricky. In particular, recall that we assumed that **Open**^{*} is accepted with high probability when opening in both bases. As mentioned, this is a reasonable assumption since we can require the committer to commit to its state many ($\Omega(1/\delta)$) times, then open half of the commitments in the standard basis and half of them in the Hadamard

¹¹Jumping ahead, we note that our non-succinct commitment scheme achieves statistical closeness and our succinct commitment scheme achieves computational indistinguishability. We mention that Mahadev’s scheme [Mah18], as well as its successors [BKL⁺22], do not satisfy this property since these schemes offer no binding on the Hadamard basis.

basis. If any of them are rejected then output \perp and otherwise, choose a random one that was opened in the desired basis b and use that as the opening. Generalizing this to the ℓ -qubit setting must be done with care to avoid an exponential blowup in ℓ . Clearly, we do not want to assume that for every basis choice $(b_1, \dots, b_\ell) \in \{0, 1\}^\ell$, Open^* successfully opens in this basis with high probability, since we cannot enforce this without incurring an exponential blowup. Yet, in order for our extractor to be successful, we need to ensure that Open^* succeeds in opening each qubit in each basis with high probability. To achieve this, without incurring an exponential blowup, we require that Open^* succeeds with high probability to open all the qubits in the standard basis (i.e., succeeds with $(b_1, \dots, b_\ell) = (0, \dots, 0)$) and succeeds with high probability to open all the qubits in the Hadamard basis (i.e., succeeds $(b_1, \dots, b_\ell) = (1, \dots, 1)$). This can be achieved via repetitions, as in the single qubit setting. Specifically, in this setting we ask 1/3 of the repetitions to be opened in the 0^ℓ basis, 1/3 to be opened in the 1^ℓ basis, and the remaining 1/3 to be opened in the desired (b_1, \dots, b_ℓ) basis. Jumping ahead, we note that the extractor Ext uses Open^* with basis (b, \dots, b) to extract the state τ . We refer the reader to Definition 5.6 for the formal definition.

Our construction for the single qubit case We start by describing our commitment scheme in the single-qubit case. Our starting point is Mahadev’s [Mah18] measurement protocol. Her protocol is binding in the standard basis but offers no binding guarantees, and in fact fails to provide any form of binding, when opening in the Hadamard basis. Moreover, in her protocol the opening basis must be determined ahead of time and the public key pk used to compute the commitment string depends on this basis. Specifically, her protocol uses a family of (noisy) trapdoor claw-free functions, where functions can be generated either in an *injective* mode or in a *two-to-one* mode. The public key of the commitment scheme consists of a public key corresponding to an injective function if the verifier wishes to open in the standard basis, and corresponds to a two-to-one function if the verifier wishes to open in the Hadamard basis.

We first notice that it is not necessary to determine the opening basis in the key generation phase. In fact, we show that one can always use the two-to-one mode, irrespective of the basis we wish to open in. Moreover, we show that this “dual mode” property is not needed altogether. This observation is quite straightforward and was implicitly used in the analysis in prior work [Vid20, BKL⁺22].

Our first instrumental idea is that we can obtain binding in both bases if we compose Mahadev’s weak commitment twice! Namely, to commit to a state σ , we first apply Mahadev’s measurement protocol, denoted by Commit_W , to obtain

$$(\mathbf{y}, \rho) \leftarrow \text{Commit}_W(\text{pk}, \sigma).$$

As mentioned, this already guarantees binding when opening in the standard basis, but fails to provide binding when opening in the Hadamard basis. To fix this we make use of the fact that Mahadev’s measurement protocol has the property that the Open algorithm always measures the post-commitment state in either the standard basis or the Hadamard basis. We apply to the post-commitment state ρ the unitary that computes Hadamard opening $\text{Open}(\cdot, 1)$, which is simply the Hadamard unitary $H^{\otimes(n+1)}$, where $n+1$ is the number of qubits in ρ (n being the security parameter associated with the underlying NTCF family), and we commit to the resulting state. Namely, we compute

$$\rho' \leftarrow H^{\otimes(n+1)}[\rho] \quad \text{and} \quad (\mathbf{y}', \rho'') \leftarrow \text{Commit}_W(\text{pk}', \rho'),$$

where \mathbf{pk} and \mathbf{pk}' are independent keys,¹² and where throughout our paper we use the shorthand

$$U[\rho] = U\rho U^\dagger$$

to denote the application of a unitary U to a mixed state ρ .

To open the commitment in the Hadamard basis, we just need to measure ρ' in the *standard* basis. Binding in the Hadamard basis follows from the fact that ρ' was committed to via the classical string \mathbf{y}' , and from the fact that Mahadev's measurement protocol provides binding in the standard basis. However, it is no longer clear how to open in the standard basis, since the original post-commitment state ρ is no longer available, and has been replaced with ρ'' . Here we use the desired property mentioned above, specifically, that algorithm `Open` generates a standard basis opening by measuring the state in the standard basis, and generates a Hadamard basis opening by measuring the state in the Hadamard basis. This implies that measuring ρ in the standard basis is equivalent to measuring ρ' in the Hadamard basis.

The reader may be concerned that we may have lost the binding in the standard basis, since opening in the Hadamard basis is not protected. But this is not the case, since it is the commitment string \mathbf{y} that binds the standard basis measurement, and the commitment string \mathbf{y}' that binds the Hadamard basis measurement.

Multi-qubit commitments One can use this single qubit commitment scheme to commit to an ℓ -qubit state, by committing qubit-by-qubit. This results with a long commitment string of size $\ell \cdot \text{poly}(\lambda)$ and with a long public key, since the public key consists of ℓ public keys $(\mathbf{pk}_1, \dots, \mathbf{pk}_\ell)$, where each \mathbf{pk}_i is generated according to the single qubit scheme. As mentioned in the introduction, our main goal is to construct a succinct commitment scheme. Following the blueprint of [BKL⁺22], we do this in two steps. We first construct a *semi-succinct* commitment scheme where the commitment string is long, but the public-key is succinct. We then show how to convert the semi-succinct scheme into a fully succinct one.

Semi-succinct commitments In our semi-succinct commitment scheme we generate a single key pair $(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Gen}(1^\lambda)$ corresponding to the single-qubit scheme, and simply use \mathbf{pk} to commit to each and every one of the qubits. The question is whether this is sound. Let us first describe the main issue that comes up when trying to prove soundness, and then we will show how we overcome it. The issue is that our commitment scheme is privately verifiable, and thus a QPT algorithm `Open*`, which produces an opening \mathbf{z} , does not know the corresponding output bit $m = \text{Out}(\mathbf{sk}, \mathbf{y}, b, \mathbf{z})$ since \mathbf{sk} is needed to compute m . Therefore, perhaps a malicious QPT algorithm `Open*` can generate \mathbf{z} in a way such that m leaks information about \mathbf{sk} . In particular, perhaps `Open*` can generate ℓ openings $\mathbf{z}_1, \dots, \mathbf{z}_\ell$ such that their corresponding outputs m_1, \dots, m_ℓ completely leak \mathbf{sk} .

Recall that our binding property consists of two parts: The first asserts that for any QPT algorithm `Commit*` that commits to an ℓ -qubit state via a classical commitment string \mathbf{y} , it holds that for any two QPT opening algorithms `Open*_1` and `Open*_2` and any basis choice (b_1, \dots, b_ℓ) , the output distributions produced by these two opening algorithms are (computationally or statistically) close. In our construction we get *statistical closeness*, and hence the closeness holds even if \mathbf{sk} is leaked. Indeed, the proof of this property in the semi-succinct setting is the same as the proof in the non-succinct setting. The issue is with the second part: Given \mathbf{sk} , the distributions generated

¹²Using different and independent public keys \mathbf{pk} and \mathbf{pk}' is important in our analysis.

by $\text{Ext}^{\text{Open}^*}$ and Open^* are no longer computationally indistinguishable. Diving deeper into our scheme and its analysis, we note that the standard basis outputs produced by $\text{Ext}^{\text{Open}^*}$ and Open^* are actually statistically close, and it is the Hadamard basis outputs that are only computationally indistinguishable.

We next examine the leakage that the decoded messages m_1, \dots, m_ℓ may contain about the secret key, and argue that even given this leakage, the Hadamard basis outputs produced by $\text{Ext}^{\text{Open}^*}$ and Open^* remain computationally indistinguishable. To this end, we will need to use additional properties about Mahadev’s measurement protocol, and thus recall it in Section 2.1 below. Jumping ahead, we mention that one property that we rely on is the fact that in Mahadev’s protocol, Out does not use the secret key when generating standard basis outputs (and the secret key is only used to generate Hadamard basis outputs).

Recall that in our commitment scheme, the secret key consists of two parts, (sk, sk') , since we apply Mahadev’s protocol twice (where sk is for a single qubit state and sk' is for an $(n + 1)$ -qubit state). We mention that when opening in the standard basis, the output m can only leak information about sk' . This is the case since to open in the standard basis, we first use sk' to generate a standard basis opening \mathbf{z} for Mahadev’s protocol, and then use Mahadev’s Out algorithm to decode \mathbf{z} , which as mentioned above, can be done publicly without the secret key sk (since it is a standard basis opening). Importantly, we show that the computational indistinguishability of the Hadamard basis opening only relies on the fact that sk is secret, and does not rely on the secrecy of sk' . Thus, the remaining problem, which is at the heart of the technical complication, is the leakage of the Hadamard basis openings on sk . We note that in Mahadev’s protocol, the Hadamard basis openings may leak the entire sk . What saves us in our setting is the fact that we tie the hands of the adversary when opening in the Hadamard basis. To explain this in more detail we need to recall Mahadev’s measurement protocol.

2.1 Mahadev’s measurement protocol

As mentioned, Mahadev’s measurement protocol [Mah18] uses a noisy TCF family.¹³ In this overview, for the sake of simplicity, we describe her scheme assuming we have a noiseless TCF family, which is a function family associated with algorithms

$$(\text{Gen}_{\text{TCF}}, \text{Eval}_{\text{TCF}}, \text{Invert}_{\text{TCF}})$$

where Gen_{TCF} is a PPT algorithm that takes as input the security parameter 1^λ and outputs a key pair (pk, sk) ; Eval is a poly-time deterministic algorithm that takes as input the public key pk , and a pair (b, \mathbf{x}) where $b \in \{0, 1\}$ is a bit and $\mathbf{x} \in \{0, 1\}^n$ (where $n = \text{poly}(\lambda)$), and outputs a value \mathbf{y} , and $\text{Eval}(\text{pk}, \cdot)$ is a two-to-one function where every \mathbf{y} in the image has exactly two preimages of the form $(0, \mathbf{x}_0)$ and $(1, \mathbf{x}_1)$; $\text{Invert}_{\text{TCF}}$ takes as input the secret key sk and an element \mathbf{y} in the image and it outputs the two preimages $((0, \mathbf{x}_0), (1, \mathbf{x}_1))$.

In what follows we show how Mahadev uses a TCF family to construct a measurement protocol. The following protocol slightly differs from Mahadev’s scheme, and in particular the basis choice is not determined during the key generation algorithm. The measurement protocol consists of algorithms $(\text{Gen}, \text{Commit}, \text{Open}, \text{Out})$ defined as follows:

- Gen is identical to Gen_{TCF} ; it takes as input the security parameter 1^λ and outputs a key pair (pk, sk) .

¹³As mentioned above, her work, as well as followup works, use a dual-mode TCF family; we avoid this technicality.

- **Commit** takes as input \mathbf{pk} and a single-qubit pure state $|\psi\rangle = \alpha_0 |0\rangle + \alpha_1 |1\rangle$ and generates

$$|\psi'\rangle = \alpha_0 |0, x_0\rangle + \alpha_1 |1, x_1\rangle$$

such that $\text{Eval}(\mathbf{pk}, (0, \mathbf{x}_0)) = \text{Eval}(\mathbf{pk}, (1, \mathbf{x}_1)) = \mathbf{y}$, and outputs \mathbf{y} as the commitment string.

- **Open** takes as input the post-committed state $|\psi'\rangle$ and a basis $b \in \{0, 1\}$; if $b = 0$ it returns the outcome \mathbf{z} of measuring $|\psi'\rangle$ in the standard basis, which is of the form (b, x_b) , and if $b = 1$ it returns the outcome \mathbf{z} of measuring $|\psi'\rangle$ in the Hadamard basis.
- **Out** takes as input $(\mathbf{sk}, \mathbf{y}, b, \mathbf{z})$, and if $b = 0$ it checks that $\text{Eval}(\mathbf{pk}, \mathbf{z}) = \mathbf{y}$ and if this is the case it outputs the first bit of \mathbf{z} , and otherwise it outputs \perp . If $b = 1$ it outputs $\mathbf{z} \cdot (1, \mathbf{x}_0 \oplus \mathbf{x}_1)$ where $((0, \mathbf{x}_0), (1, \mathbf{x}_1)) = \text{Invert}_{\text{TCF}}(\mathbf{y})$.

Note that Mahadev's measurement protocol is not fully binding. The issue is that a cheating prover can produce any opening in the Hadamard basis, and will never be rejected. For instance, a cheating prover could commit to $|+\rangle$ honestly, apply a Z to the first qubit of the post-commitment state, and then open to $|-\rangle$.

2.2 Our Single-Qubit Commitment Scheme

We convert Mahadev's protocol into a binding commitment scheme by adding another step to the commitment algorithm, as described in the beginning of Section 2. More specifically, our commitment scheme consists of algorithms $(\text{Gen}, \text{Commit}, \text{Open}, \text{Out})$ defined as follows:

- $\text{Gen}(1^\lambda)$ generates $n + 2$ TCF keys $(\mathbf{pk}_i, \mathbf{sk}_i)_{i \in \{0, 1, \dots, n+1\}}$, where each $(\mathbf{pk}_i, \mathbf{sk}_i) \leftarrow \text{Gen}_{\text{TCF}}(1^\lambda)$, and outputs $\mathbf{pk} = (\mathbf{pk}_0, \mathbf{pk}_1, \dots, \mathbf{pk}_{n+1})$ and $\mathbf{sk} = (\mathbf{sk}_0, \mathbf{pk}_1, \dots, \mathbf{sk}_{n+1})$.
- $\text{Commit}(\mathbf{pk}, |\psi\rangle)$ operates as follows:

1. Parse $\mathbf{pk} = (\mathbf{pk}_0, \mathbf{pk}_1, \dots, \mathbf{pk}_{n+1})$.

2. Apply Mahadev's measurement protocol to commit to $|\psi\rangle = \alpha_0 |0\rangle + \alpha_1 |1\rangle$ w.r.t. \mathbf{pk}_0 ; i.e., generate

$$|\psi'\rangle = \alpha_0 |0, x_0\rangle + \alpha_1 |1, x_1\rangle$$

such that $\text{Eval}(\mathbf{pk}_0, (0, \mathbf{x}_0)) = \text{Eval}(\mathbf{pk}_0, (1, \mathbf{x}_1)) = \mathbf{y}_0$.

3. Compute

$$H^{\otimes(n+1)} |\psi'\rangle = \sum_{\mathbf{d} \in \{0, 1\}^{n+1}} \beta_{\mathbf{d}} |\mathbf{d}\rangle,$$

4. Use Mahadev's measurement protocol to commit qubit-by-qubit to the above $(n + 1)$ -qubit state, w.r.t. public keys $\mathbf{pk}_1, \dots, \mathbf{pk}_{n+1}$ to obtain the state

$$\sum_{\mathbf{d} \in \{0, 1\}^{n+1}} \beta_{\mathbf{d}} |\mathbf{d}\rangle |\mathbf{x}'_{1, \mathbf{d}_1}\rangle \cdots |\mathbf{x}'_{n+1, \mathbf{d}_{n+1}}\rangle$$

and strings $\mathbf{y}_1, \dots, \mathbf{y}_{n+1}$ such that for every $i \in [n + 1]$,

$$\text{Eval}(\mathbf{pk}_i, (0, \mathbf{x}'_{i, 0})) = \text{Eval}(\mathbf{pk}_i, (1, \mathbf{x}'_{i, 1})) = \mathbf{y}_i.$$

5. Output $(\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_{n+1})$, and (for simplicity) rearrange the post-commitment state to be

$$\sum_{\mathbf{d} \in \{0,1\}^{n+1}} \beta_{\mathbf{d}} |\mathbf{d}_1, \mathbf{x}'_{1,\mathbf{d}_1}\rangle \dots |\mathbf{d}_{n+1}, \mathbf{x}'_{n+1,\mathbf{d}_{n+1}}\rangle$$

- **Open** takes as input the post-commitment state ρ and a basis $b \in \{0,1\}$. If $b = 1$ (corresponding to opening in the Hadarmard basis) then it outputs the measurement of the state ρ in the standard. If $b = 0$ (corresponding to opening in the stanadard basis) then it outputs the measurement of the state ρ in the Hadamard basis.
- **Out** takes as input the secret key $\mathbf{sk} = (\mathbf{sk}_0, \mathbf{sk}_1, \dots, \mathbf{sk}_{n+1})$, a commitment string $\mathbf{y} = (\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_{n+1})$, a basis $b \in \{0,1\}$ and an opening string $\mathbf{z} \in \{0,1\}^{(n+1)^2}$ and does the following:

1. If $b = 1$ then parse

$$\mathbf{z} = (\mathbf{d}_1, \mathbf{x}'_{1,\mathbf{d}_1}, \dots, \mathbf{d}_{n+1}, \mathbf{x}'_{n+1,\mathbf{d}_{n+1}})$$

and check that for every $i \in [n+1]$ it holds that

$$\mathbf{y}_i = \text{Eval}(\mathbf{pk}_i, (\mathbf{d}_i, \mathbf{x}'_{i,\mathbf{d}_i})).$$

If all these checks pass then output $\mathbf{d} \cdot (1, \mathbf{x}_0 \oplus \mathbf{x}_1)$, where $((0, \mathbf{x}), (1, \mathbf{x}_1)) = \text{Invert}(\mathbf{sk}_0 \mathbf{y}_0)$. Otherwise, output \perp .

2. If $b = 0$ then parse $\mathbf{z} = (\mathbf{z}_1, \dots, \mathbf{z}_{n+1})$, and for every $i \in [n+1]$ compute

$$((0, \mathbf{x}'_{i,0}), (1, \mathbf{x}'_{i,1})) = \text{Invert}(\mathbf{sk}_i, \mathbf{y}_i) \quad \text{and} \quad m_i = \mathbf{z}_i \cdot (1, \mathbf{x}'_{i,0} \oplus \mathbf{x}'_{i,1})$$

If $\text{Eval}(\mathbf{pk}_0, (m_1, \dots, m_{i+1})) = \mathbf{y}_0$ then output m_1 , and otherwise output \perp .

Analyzing the leakage. We next analyze the leakage that a cheating QPT algorithm **Commit*** and a cheating QPT algorithm **Open*** obtain by, given $\mathbf{pk} = (\mathbf{pk}_0, \mathbf{pk}_1, \dots, \mathbf{pk}_{n+1})$, generating a commitment string $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_\ell)$, where each $\mathbf{y}_i = (\mathbf{y}_{i,0}, \mathbf{y}_{i,1}, \dots, \mathbf{y}_{i,n+1})$, a basis (b_1, \dots, b_ℓ) and an opening $\mathbf{z} = (\mathbf{z}_1, \dots, \mathbf{z}_\ell)$, and obtaining outputs $m_i = \text{Out}(\mathbf{sk}, \mathbf{y}_i, b_i, \mathbf{z}_i)$ for every $i \in [\ell]$. Denote by

$$I = \{i : b_i = 0\} \quad \text{and} \quad J = \{i : b_i = 1\}.$$

We distinguish between the leakage obtained from $\{m_i\}_{i \in I}$ and that obtained from $\{m_i\}_{i \in J}$. As mentioned above, $\{m_i\}_{i \in I}$ only leaks information about $\mathbf{sk}_1, \dots, \mathbf{sk}_{n+1}$, since \mathbf{sk}_0 is not used when computing $\{m_i\}_{i: b_i=0}$. For $i \in J$, it holds that

$$m_i = \mathbf{d}_i \cdot (1, \mathbf{x}_{i,0} \oplus \mathbf{x}_{i,1}) \quad \text{where} \quad ((0, \mathbf{x}_{i,0}), (1, \mathbf{x}_{i,1})) = \text{Invert}_{\text{TCF}}(\mathbf{sk}_0, \mathbf{y}_{i,0}),$$

where $\mathbf{z}_i = (\mathbf{d}_{i,1}, \mathbf{x}'_{i,1,\mathbf{d}_{i,1}}, \dots, \mathbf{d}_{i,n+1}, \mathbf{x}'_{i,n+1,\mathbf{d}_{i,n+1}})$. This may leak information about \mathbf{sk}_0 . In particular, if we use the underlying (noisy) TCF family from [BCM⁺18], along with an adversarially chosen $\mathbf{d} = (\mathbf{d}_1, \dots, \mathbf{d}_\ell)$ then $\{m_i\}_{i \in J}$ may leak part of the secret key which breaks the indistinguishability between the output produced by **Open*** and $\text{Ext}^{\text{Open}*}$.

We get around this problem by arguing that in our scheme if \mathbf{z} is accepted then it must be the case that “the important” bits of \mathbf{d} have min-entropy $\omega(\log \lambda)$.¹⁴ For this we rely on the fact that the

¹⁴We emphasize that this is not the case for Mahadev’s scheme, since in her scheme every Hadamard opening \mathbf{d} is accepted.

underlying TCF family has the adaptive hardcore bit property, which the (noisy) TCF family from [BCM⁺18] was proven to have under the LWE assumption. We actually need the stronger condition that “the important” bits of \mathbf{d} have min-entropy $\omega(\log \lambda)$ even given some auxiliary input (which comes into play due to the fact that we are opening many qubits). We prove this for the specific NTCF family from [BCM⁺18]. Specifically, we prove that under the LWE assumption, the NTCF family from [BCM⁺18] has a property which we refer to as the *distributional strong adaptive hardcore bit property*. We argue that this property, together with the min-entropy property of \mathbf{d} , implies that the leakage obtained from $\mathbf{d}_i \cdot (\mathbf{x}_{i,0} \oplus \mathbf{x}_{i,1})$ is benign and does not break the indistinguishability between the output produced by Open^* and $\text{Ext}^{\text{Open}^*}$.

In more detail, for Mahadev’s measurement protocol, the proof that the Hadamard outputs of $\text{Ext}^{\text{Open}^*}$ and Open^* are computationally indistinguishable relies on the adaptive hardcore bit property, which states that for every QPT adversary A ,

$$\Pr[A(\text{pk}) = (b, \mathbf{x}_b, \mathbf{d}, \mathbf{d} \cdot (1, \mathbf{x}_0 \oplus \mathbf{x}_1))] \leq \frac{1}{2} + \text{negl}(\lambda)$$

where $((0, \mathbf{x}_0), (1, \mathbf{x}_1)) = \text{Invert}(\text{sk}, \text{Eval}(\text{pk}, b, \mathbf{x}_b))$. We need to argue that this holds even if A gets as auxiliary input a bunch of elements of the form

$$(b_i, \mathbf{x}_{i,b_i}, \mathbf{d}_i \cdot (1, \mathbf{x}_{i,0} \oplus \mathbf{x}_{i,1})).$$

While in general this is not true, we prove that it is true for the (noisy) TCF family from [BBCM93], if each \mathbf{d}_i has $\omega(\log \lambda)$ min-entropy (even conditioned on $(\mathbf{d}_1, \dots, \mathbf{d}_{i-1})$), under the LWE assumption.

2.3 Succinct commitments

As mentioned in the introduction, our main result is a *succinct* commitment scheme, where Commit commits to an ℓ -qubit state by generating a *succinct* classical commitment string that consists of only $\text{poly}(\lambda)$ many bits, and Open generates an opening to any qubit $i \in [\ell]$ in any basis $b \in \{0, 1\}$, where the opening consists of only $\text{poly}(\lambda)$ many bits. Importantly, the guarantee we provide is that even if Open^* only opens to a few qubits, we should still be able to extract the entire ℓ -qubit quantum state from Open^* .¹⁵ This seems impossible to do, since how can we extract information about qubits that were never opened? Indeed, to achieve this we need to change the syntax.

We add to the syntax an *interactive test phase*. Similar to the test round in Mahadev’s protocol, our test phase is executed with probability $1/2$, and if it is executed then after the test phase the protocol is terminated and the opening phase is never run. This is the case since the test phase destroys the quantum state. Importantly, we allow the test phase to be *interactive*. It is this interaction that allows us to extract a long ℓ -qubit state from Open^* . Loosely speaking, in this test phase, we choose at random $b \leftarrow \{0, 1\}$ and ask the prover to provide an opening to all the ℓ -qubits in basis b^ℓ . To ensure that the protocol remains succinct, we ask for the openings to be sent in a succinct manner, using a Merkle hash. Then the prover and verifier engage in a succinct interactive argument where the prover proves knowledge of the committed openings. For this we use Kilian’s protocol and the fact that it is a proof-of-knowledge even in the post-quantum setting [CMSZ21, VZ21]. Then the verifier sends the prover the secret key sk and the prover and verifier engage in a succinct interactive argument where the prover proves that the committed openings are accepted (w.r.t. sk). This is also done using the Kilian protocol.

¹⁵This guarantee is important for our applications, as we will see in Section 2.4.

In addition, we allow the commit phase to be interactive. This allows `Commit` to first generate a non-succinct commitment \mathbf{y} , and send its Merkle hash, denoted by \mathbf{rt} . Then the committer can run a succinct proof-of-knowledge interactive argument, to prove knowledge of a preimage of \mathbf{y} . Importantly, the proof-of-knowledge must be *state-preserving*, which means that we can extract \mathbf{y} without destroying the state. Such a state-preserving proof-of-knowledge protocol was recently constructed in [LMS22]. This interactive commitment phase allows us to reduce the binding of the succinct commitment scheme to that of the semi-succinct one. This part of the analysis is similar to [BKL⁺22].

2.4 Applications

We construct a succinct interactive argument for **QMA** and a compiler that converts any X/Z PCP into a succinct interactive argument, both under the **LWE** assumption. For simplicity we do not use our succinct commitment scheme to construct these succinct interactive arguments. Rather we use our *semi-succinct* commitment scheme to construct a *semi-succinct* interactive argument. We then rely on a black-box transformation from [BKL⁺22] which shows a generic transformation for converting any semi-succinct interactive argument for **QMA** into a fully succinct one.¹⁶

Compiling an X/Z PCP into a semi-succinct interactive argument Our compiler uses a succinct commitment in a straightforward way. The succinct interactive argument proceeds as follows:

1. The verifier generates a key pair $(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Gen}(1^\lambda)$ corresponding to the underlying semi-succinct commitment scheme.
2. The prover commits to the X/Z PCP $|\pi\rangle$ by generating a classical commitment string $\mathbf{y} \leftarrow \text{Commit}(\mathbf{pk}, |\pi\rangle)$. It sends \mathbf{y} to the verifier.
3. With probability $1/2$ the verifier behaves as the PCP verifier and chooses small set of indices (i_1, \dots, i_k) along with basis choices (b_1, \dots, b_k) ; with probability $1/2$ the verifier chooses a random $b \leftarrow \{0, 1\}$ and sends b to the prover.
4. If the prover receives a bit b then it opens the entire PCP in the standard basis if $b = 0$ and in the Hadamard basis if $b = 1$. Otherwise, if the prover receives a set of indices (i_1, \dots, i_k) along with basis choices (b_1, \dots, b_k) then the prover opens these locations in the desired basis.

Completeness follows immediately from the completeness of the underlying semi-succinct commitment scheme. To argue soundness, fix a cheating prover P^* that is accepted with high probability. We rely on the soundness property of the underlying commitment scheme to argue that there exists a QPT extractor that extracts a state $|\pi^*\rangle$ from P^* , such that on a random challenge produced by the PCP verifier (for which P^* succeeds in opening with high probability), the output of P^* is close to the the outcome obtained by measuring $|\pi^*\rangle$ directly, which implies that $|\pi^*\rangle$ is an X/Z PCP that is accepted with high probability, implying that the soundness property holds.

¹⁶We mention that this transformation was used (in a non-black-box way to convert our semi-succinct commitment scheme into a succinct one.

Semi-succinct interactive argument for QMA To obtain a semi-succinct argument, we follow the blueprint of Mahadev [Mah18]. Namely, we first convert the **QMA** witness into one that can be verified by measuring only in the X/Z basis. For this we rely on a result due to Fitzsimons, Hajdušek, and Morimae [FHM18] which shows how to convert multiple copies of the **QMA** witness into an ℓ -qubit state $|\pi\rangle$ that can be verified by measuring it only in the X/Z basis. Importantly, this state can be verified by measuring it in a random basis $(b_1, \dots, b_\ell) \leftarrow \{0, 1\}^\ell$. Armed with this tool, the semi-succinct interactive argument proceeds as follows:

1. The verifier generates a key pair $(\mathbf{pk}, \mathbf{sk})$ corresponding to the underlying semi-succinct commitment scheme.
2. The prover converts its (multiple copies) of the **QMA** witness into a state $|\pi\rangle$ by relying on the [FHM18] result, and computes $\mathbf{y} \leftarrow \text{Commit}(\mathbf{pk}, |\pi\rangle)$.
3. With probability $1/2$ the verifier chooses at random a seed $s \in \{0, 1\}^\lambda$ and sends s to the prover, and with probability $1/2$ the verifier chooses a random $b \leftarrow \{0, 1\}$ and sends b to the prover.
4. If the prover receives a bit b then it sends the opening of the commitment in the basis b^ℓ . If it receives a seed s then it uses a pseudorandom generator to deterministically expand s to a pseudorandom string (b_1, \dots, b_ℓ) and sends an opening of the commitment in basis (b_1, \dots, b_ℓ) .
5. The verifier uses its secret key to compute the output corresponding to this opening. If any of the openings are rejected it rejects. Otherwise, in the case that it sent a seed, it accepts if the verifier from [FHM18] would have accepted.

To argue soundness, fix a cheating prover P^* that is accepted with high probability. We first rely on the soundness property of the underlying commitment scheme to argue that the QPT extractor extracts a state $|\pi^*\rangle$ from P^* , such that for any choice of basis $\mathbf{b} = (b_1, \dots, b_\ell)$ for which P^* succeeds in opening with high probability, the output corresponding to these openings are computationally indistinguishable from measuring $|\pi^*\rangle$ in basis \mathbf{b} . By the soundness of the underlying scheme [FHM18] we note that for a random basis (b_1, \dots, b_ℓ) , the state would be rejected with high probability. Hence it must also be the case if the basis is pseudorandom, as otherwise one can distinguish a pseudorandom string from a truly random one.

3 Preliminaries

Notations. For any random variables A and B (classical variables or quantum states), we use the notation $A \equiv B$ to denote that A and B are identically distributed, and use $A \stackrel{\epsilon}{\equiv} B$ to denote that A and B are ϵ -close, where closeness is measured with respect to total variation distance for classical variables, trace distance for mixed quantum states, and $\|\cdot\|_2$ distance for pure quantum states. For every two ensemble of distributions $A = \{A_\lambda\}_{\lambda \in \mathbb{N}}$ and $B = \{B_\lambda\}_{\lambda \in \mathbb{N}}$ we use the notation $A \approx B$ to denote that A and B are computationally indistinguishable, i.e., for every polynomial size distinguisher D there exists a negligible function $\mu = \mu(\lambda)$ such that for every $\lambda \in \mathbb{N}$,

$$|\Pr[D(a) = 1] - \Pr[D(b) = 1]| \leq \mu(\lambda)$$

where the probabilities are over $a \leftarrow A_\lambda$ and $b \leftarrow B_\lambda$. For every $\epsilon = \epsilon(\lambda) \in [0, 1]$, we use the notation $A \stackrel{\epsilon}{\approx} B$ to denote that for every polynomial size distinguisher D and for every $\lambda \in \mathbb{N}$,

$$|\Pr[D(a) = 1] - \Pr[D(b) = 1]| \leq \epsilon(\lambda)$$

where the probabilities are over $a \leftarrow A_\lambda$ and $b \leftarrow B_\lambda$.

For any random variable A , we denote by $\text{Supp}(A)$ the support of A ; i.e.,

$$\text{Supp}(A) = \{a : \Pr[A = a] > 0\}.$$

We denote strings in $\{0, 1\}^*$ by bold lower case letters, such as \mathbf{x} . We let PPT denote probabilistic polynomial time, QPT denote probabilistic quantum polynomial time, and QPT denote quantum polynomial time.

Let \mathcal{H} be a complex Hilbert space of finite dimension 2^n . Thus, $\mathcal{H} \simeq \mathbb{C}^{2^n}$ where \mathbb{C} denotes the complex numbers. A pure n -qubit quantum state is a unit vector $|\Psi\rangle \in \mathcal{H}$. Namely, it can be written as

$$|\Psi\rangle = \sum_{b_1, \dots, b_n \in \{0, 1\}} \alpha_{b_1, \dots, b_n} |b_1, \dots, b_n\rangle$$

where $\{|b_1, \dots, b_n\rangle\}_{b_1, \dots, b_n \in \{0, 1\}}$ forms an orthonormal basis of \mathcal{H} , and where $\alpha_{b_1, \dots, b_n} \in \mathbb{C}$ satisfy

$$\sum_{b_1, \dots, b_n \in \{0, 1\}} |\alpha_{b_1, \dots, b_n}|^2 = 1.$$

We refer to n as the number of qubits in $|\Psi\rangle$. We sometimes divide the registers of $|\Psi\rangle$ into named registers. We often denote these registers by calligraphic upper-case letters, such as \mathcal{A} and \mathcal{B} , in which case we also divide the Hilbert space into $\mathcal{H} = \mathcal{H}_{\mathcal{A}} \otimes \mathcal{H}_{\mathcal{B}}$, so that each quantum state $|\Psi\rangle$ is a linear combination of quantum states $|\Psi_{\mathcal{A}}\rangle \otimes |\Psi_{\mathcal{B}}\rangle \in \mathcal{H}_{\mathcal{A}} \otimes \mathcal{H}_{\mathcal{B}}$.¹⁷ We denote by

$$|\Psi\rangle_{\mathcal{A}} = \text{Tr}_{\mathcal{B}}(|\Psi\rangle \langle \Psi|) \in \mathcal{H}_{\mathcal{A}},$$

where $\text{Tr}_{\mathcal{B}}$ is the linear operator defined by

$$\text{Tr}_{\mathcal{B}}(|\Psi_{\mathcal{A}}\rangle \langle \Psi_{\mathcal{A}}| \otimes |\Psi_{\mathcal{B}}\rangle \langle \Psi_{\mathcal{B}}|) = |\Psi_{\mathcal{A}}\rangle \langle \Psi_{\mathcal{A}}| \cdot \text{Tr}(|\Psi_{\mathcal{B}}\rangle \langle \Psi_{\mathcal{B}}|),$$

where Tr is the trace operator.

Let $D(\mathcal{H})$ denote the set of all positive semidefinite operators on \mathcal{H} with trace 1. A mixed state is an operator $\rho \in D(\mathcal{H})$, and is often called a *density matrix*. We denote by

$$U[\sigma] = U\sigma U^\dagger.$$

For any binary observable O and bit $b \in \{0, 1\}$ we let $\Pi_{O,b}[\sigma]$ denote the *unnormalized* projection of σ to the state that has value b when measured in the O -basis. Namely,

$$\Pi_{O,b}[\sigma] = \Pi_{O,b}\sigma\Pi_{O,b}^\dagger.$$

We let X and Z denote the Pauli matrices:

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

For any single qubit register \mathcal{A} , we denote by $X_{\mathcal{A}}$ (respectively, $Z_{\mathcal{A}}$) the unitary that applies the Pauli X (respectively, Z) unitary to the \mathcal{A} register of a given quantum state and applies the identity unitary to all other registers.

¹⁷We sometimes give registers names that correspond to their purpose, such as a *coin* register or an *open* register.

3.1 Quantum information facts

We use the following infant version of the gentle measurement lemma.

Lemma 3.1. *Let $|\psi\rangle$ be a pure state and Π be a projector such that $\langle\psi|\Pi|\psi\rangle = 1 - \varepsilon$. Then $\|\Pi|\psi\rangle - |\psi\rangle\|_2 = \sqrt{\varepsilon}$.*

Proof. Calculate:

$$\|\Pi|\psi\rangle - |\psi\rangle\|_2^2 = 1 - \langle\psi|\Pi|\psi\rangle = \varepsilon.$$

□

We also use the following version for mixed states.

Lemma 3.2. *Let ρ be a mixed state and Π be a projector such that $\text{Tr}[\Pi\rho] = 1 - \varepsilon$. Then $\frac{1}{2}\|\rho - \Pi[\rho]\|_1 \leq \sqrt{\varepsilon}$.*

Proof. This is Lemma 9.4.2 of [Wil11].

□

Lemma 3.3. *Suppose $|\psi_1\rangle_{ABC}$ and $|\psi_2\rangle_{ABC}$ are pure states with A being a single-qubit register, and U is a unitary acting on register B . Let CU_{AB} denote the controlled version of U , with A being the control system and B the target. Then*

$$\|CU_{AB} \otimes I_C(|\psi_1\rangle - |\psi_2\rangle)\|_2 = \|Z_A \otimes I_{BC} (|\psi_1\rangle - |\psi_2\rangle)\|_2,$$

where Z is the Pauli Z operator.

Proof. In the following calculation we omit factors of identity that are clear from context.

$$\begin{aligned} \|CU_{AB}(|\psi_1\rangle - |\psi_2\rangle)\|_2^2 &= \||0\rangle\langle 0|_A (|\psi_1\rangle - |\psi_2\rangle) + |1\rangle\langle 1|_A \otimes U_B(|\psi_1\rangle - |\psi_2\rangle)\|_2^2 \\ &= \||0\rangle\langle 0|_A (|\psi_1\rangle - |\psi_2\rangle)\|_2^2 + \||1\rangle\langle 1|_A \otimes U_B(|\psi_1\rangle - |\psi_2\rangle)\|_2^2 \\ &= \||0\rangle\langle 0|_A (|\psi_1\rangle - |\psi_2\rangle)\|_2^2 + \||1\rangle\langle 1|_A \otimes U_B(|\psi_1\rangle - |\psi_2\rangle)\|_2^2 \\ &= \||0\rangle\langle 0|_A (|\psi_1\rangle - |\psi_2\rangle)\|_2^2 + \||1\rangle\langle 1|_A (|\psi_1\rangle - |\psi_2\rangle)\|_2^2 \\ &= \||0\rangle\langle 0|_A (|\psi_1\rangle - |\psi_2\rangle) - |1\rangle\langle 1|_A (|\psi_1\rangle - |\psi_2\rangle)\|_2^2 \\ &= \|Z_A(|\psi_1\rangle - |\psi_2\rangle)\|_2^2. \end{aligned}$$

□

3.2 Hash Family with Local Opening

A hash family with local opening consists of the following algorithms:

$\text{Gen}(1^\lambda) \rightarrow \text{hk}$. This PPT algorithm takes as input a security parameter λ (in unary) and outputs a hash key hk .

$\text{Eval}(\text{hk}, \mathbf{x}) \rightarrow \text{rt}$. This deterministic poly-time algorithm takes as input a hash key hk and a string $\mathbf{x} \in \{0, 1\}^N$ and outputs a hash value (often referred to as hash root) $\text{rt} \in \{0, 1\}^{\text{poly}(\lambda, \log N)}$.

$\text{Open}(\text{hk}, \mathbf{x}, i) \rightarrow (b, \mathbf{o})$. This deterministic poly-time algorithm takes as input a hash key hk , a string $\mathbf{x} \in \{0, 1\}^N$ and an index $i \in [N]$. It outputs a bit $b \in \{0, 1\}$ and an opening $\mathbf{o} \in \{0, 1\}^{\text{poly}(\lambda, \log N)}$.

$\text{Ver}(\text{hk}, \text{rt}, i, b, \mathbf{o}) \rightarrow 0/1$. This deterministic poly-time algorithm takes as input a hash key hk , a hash root rt , an index $i \in [N]$, a bit $b \in \{0, 1\}$ and an opening $\mathbf{o} \in \{0, 1\}^{\text{poly}(\lambda, \log N)}$. It outputs a bit indicating whether or not the opening is valid.

Definition 3.4. A hash family with local opening $(\text{Gen}, \text{Eval}, \text{Open}, \text{Ver})$ is required to satisfy the following properties.

Opening completeness. For any $\lambda \in \mathbb{N}$, any $N = N(\lambda) \leq 2^\lambda$, any $\mathbf{x} = (x_1, \dots, x_n) \in \{0, 1\}^N$, and any index $i \in [N]$,

$$\Pr \left[\begin{array}{l} b = x_i \\ \wedge \text{Ver}(\text{hk}, \text{rt}, i, b, \mathbf{o}) = 1 \end{array} : \begin{array}{l} \text{hk} \leftarrow \text{Gen}(1^\lambda), \\ \text{rt} = \text{Eval}(\text{hk}, \mathbf{x}), \\ (b, \mathbf{o}) = \text{Open}(\text{hk}, \mathbf{x}, i) \end{array} \right] = 1 - \text{negl}(\lambda).$$

Computational binding w.r.t. opening. For any poly-size adversary A , there exists a negligible function $\text{negl}(\cdot)$ such that for every $\lambda \in \mathbb{N}$,

$$\Pr \left[\begin{array}{l} \text{Ver}(\text{hk}, \text{rt}, i, 0, \mathbf{o}_0) = 1 \\ \wedge \text{Ver}(\text{hk}, \text{rt}, i, 1, \mathbf{o}_1) = 1 \end{array} : \begin{array}{l} \text{hk} \leftarrow \text{Gen}(\lambda), \\ (1^N, \text{rt}, i, \mathbf{o}_0, \mathbf{o}_1) \leftarrow A(\text{hk}) \end{array} \right] = \text{negl}(\lambda).$$

Theorem 3.5 ([Mer88]). Assuming the existence of a collision resistant hash family there exists a hash family with local opening (according to Definition 3.4).

Definition 3.6. A hash family with local opening $(\text{Gen}, \text{Eval}, \text{Open}, \text{Ver})$ is said to be collapsing if any QPT adversary A wins in the following game with probability $\frac{1}{2} + \text{negl}(\lambda)$:

1. The challenger generates $\text{hk} \leftarrow \text{Gen}(1^\lambda)$ and sends pk to A .
2. $A(\text{hk})$ generates a classical value (rt, j) and a quantum state σ .
 A sends (rt, j, σ) to the challenger.
3. The challenger does the following:
 - (a) Apply in superposition the algorithm $\text{Ver}(\text{hk}, \text{rt}, j, \cdot, \cdot)$ to σ , and measure the output. If the output is 0 then send \perp to A . Otherwise, denote the resulting state by σ'
 - (b) Choose a random bit $b \leftarrow \{0, 1\}$.
 - (c) If $b = 0$ then send σ' to A .
 - (d) If $b = 1$ then measure σ' in the standard basis and send the resulting state to A .
4. Upon receiving the quantum state (or the symbol \perp), A outputs a bit b' .
5. A wins if $b' = b$

Theorem 3.7 ([Unr16, CMSZ21]). There exists a hash family with local opening that is collapsing assuming the post-quantum hardness of LWE.

3.3 Noisy Trapdoor Claw-Free Functions

In what follows we define the notion of a *noisy trapdoor claw-free function family*. This notion is simpler than the notion of a *dual-mode noisy trapdoor claw-free function family* which was used for certifiable randomness generation in [BCM⁺18] and by Mahadev [Mah18] in her classical verification protocol for QMA. This simpler notion suffices for our work.¹⁸

Definition 3.8. *A noisy trapdoor claw-free function (NTCF) family is described by PPT algorithms (Gen, Eval, Invert, Check, Good) with the following syntax:*

$\text{Gen}(1^\lambda) \rightarrow (\text{pk}, \text{sk})$. *This PPT key generation algorithm takes as input a security parameter λ (in unary) and outputs a public key pk and a secret key sk .*

We denote by D_{pk} the domain of the (randomized) function defined by pk , and assume for simplicity that D_{pk} is an efficiently verifiable and samplable subset of $\{0, 1\}^{n(\lambda)}$. We denote by R_{pk} the range of this (randomized) function.

$\text{Eval}(\text{pk}, b, \mathbf{x}) \rightarrow \mathbf{y}$. *This PPT algorithm takes as input a public key pk , a bit $b \in \{0, 1\}$ and an element $\mathbf{x} \in D_{\text{pk}}$, and outputs a string \mathbf{y} distributed according to some distribution $\chi = \chi_{\text{pk}, b, \mathbf{x}}$.*

$\text{Invert}(\text{sk}, \mathbf{y}) \rightarrow ((0, \mathbf{x}_0), (1, \mathbf{x}_1))$. *This deterministic polynomial time algorithm takes as input a secret key sk , and an element \mathbf{y} in the range R_{pk} and outputs two pairs $(0, \mathbf{x}_0)$ and $(1, \mathbf{x}_1)$ with $\mathbf{x}_0, \mathbf{x}_1 \in D_{\text{pk}}$, or \perp .*

$\text{Check}(\text{pk}, b, \mathbf{x}, \mathbf{y}) \rightarrow 0/1$. *This deterministic poly-time algorithm takes as input a public key pk , a bit $b \in \{0, 1\}$, an element $\mathbf{x} \in D_{\text{pk}}$ and an element $\mathbf{y} \in R_{\text{pk}}$ and outputs a bit.*

$\text{Good}(\mathbf{x}_0, \mathbf{x}_1, \mathbf{d}) \rightarrow 0/1$. *This deterministic poly-time algorithm takes as input two domain elements $\mathbf{x}_0, \mathbf{x}_1 \in D_{\text{pk}}$ and a string $\mathbf{d} \in \{0, 1\}^{n+1}$. It outputs a bit that characterizes membership in the set:*

$$\text{Good}_{\mathbf{x}_0, \mathbf{x}_1} := \{\mathbf{d} : \text{Good}(\mathbf{x}_0, \mathbf{x}_1, \mathbf{d}) = 1\} \quad (1)$$

We specify that $\text{Good}(\mathbf{x}_0, \mathbf{x}_1, \mathbf{d})$ ignores the first bit of \mathbf{d} .

For the purpose of this work, we allow Good to output a vector (as opposed to a single bit).¹⁹ Specifically, Good may output a vector in $\{0, 1\}^k$ (for some $k = k(\lambda) \in \mathbb{N}$) in which case we define

$$\text{Good}_{\mathbf{x}_0, \mathbf{x}_1} := \{\mathbf{d} : \text{Good}(\mathbf{x}_0, \mathbf{x}_1, \mathbf{d}) \neq 0^k\}.$$

We require that the following properties are satisfied.

1. Completeness:

(a) *For all $(\text{pk}, \text{sk}) \in \text{Supp}(\text{Gen}(1^\lambda))$, every $b \in \{0, 1\}$, every $\mathbf{x} \in D_{\text{pk}}$, and $\mathbf{y} \in \text{Supp}(\text{Eval}(\text{pk}, b, \mathbf{x}))$,*

$$\text{Invert}(\text{sk}, \mathbf{y}) = ((0, \mathbf{x}_0), (1, \mathbf{x}_1))$$

such that $\mathbf{x}_b = \mathbf{x}$ and $\mathbf{y} \in \text{Supp}(\text{Eval}(\text{pk}, \beta, \mathbf{x}_\beta))$ for every $\beta \in \{0, 1\}$.

¹⁸Our formulation is from [BKL⁺22] (without the dual-mode requirement).

¹⁹We use this extension to analyze our succinct commitment scheme.

- (b) For all $(\text{pk}, \text{sk}) \in \text{Supp}(\text{Gen}(1^\lambda))$, there exists a perfect matching $M_{\text{pk}} \subseteq D_{\text{pk}} \times D_{\text{pk}}$ such that $\text{Eval}(\text{pk}, 0, \mathbf{x}_0) \equiv \text{Eval}(\text{pk}, 1, \mathbf{x}_1)$ if and only if $(\mathbf{x}_0, \mathbf{x}_1) \in M_{\text{pk}}$.
- (c) For all $(\text{pk}, \text{sk}) \in \text{Supp}(\text{Gen}(1^\lambda))$, every $b \in \{0, 1\}$ and every $\mathbf{x} \in D_{\text{pk}}$,

$$\Pr[\text{Check}(\text{pk}, b, \mathbf{x}, \mathbf{y}) = 1] = 1 \quad (2)$$

if and only if $\mathbf{y} \in \text{Supp}(\text{Eval}(\text{pk}, b, \mathbf{x}))$.

- (d) For all $(\text{pk}, \text{sk}) \in \text{Supp}(\text{Gen}(1^\lambda))$ and every pair of distinct domain elements $\mathbf{x}_0, \mathbf{x}_1$, the density of $\text{Good}_{\mathbf{x}_0, \mathbf{x}_1}$ is $1 - \text{negl}(\lambda)$.

2. **Efficient Range Superposition:** For every $(\text{pk}, \text{sk}) \in \text{Supp}(\text{Gen}(1^\lambda))$ and every $b \in \{0, 1\}$, there exists an efficient QPT algorithm to prepare a state $|\varphi_b\rangle$ such that:

$$|\varphi_b\rangle \equiv \frac{1}{\sqrt{|D_{\text{pk}}|}} \sum_{\substack{\mathbf{x} \in D_{\text{pk}} \\ \mathbf{y} \in R_{\text{pk}}}} \sqrt{p_{\text{pk}}(b, \mathbf{x}, \mathbf{y})} |\mathbf{x}\rangle |\mathbf{y}\rangle \quad (3)$$

for some negligible function $\mu(\cdot)$. Here, $p_{\text{pk}}(b, \mathbf{x}, \mathbf{y})$ denotes the probability density of \mathbf{y} in the distribution $\text{Eval}(\text{pk}, b, \mathbf{x})$.

3. **Adaptive Hardcore Bit:** For every QPT adversary A there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$,

$$\Pr[A(\text{pk}) = (\mathbf{y}, b, \mathbf{x}, \mathbf{d}, v) : \text{Check}(\text{pk}, b, \mathbf{x}, \mathbf{y}) = 1 \wedge \mathbf{d} \in \text{Good}_{\mathbf{x}_0, \mathbf{x}_1} \wedge \mathbf{d} \cdot (1, \mathbf{x}_0 \oplus \mathbf{x}_1) = v] \leq \frac{1}{2} + \mu(\lambda),$$

where the probability is over $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$, and where $((0, \mathbf{x}_0), (1, \mathbf{x}_1)) = \text{Invert}(\text{sk}, \mathbf{y})$.

Claim 3.9. [BCM⁺18] There exists a NTCF family assuming the post-quantum hardness of LWE.

In this work we rely on the fact that every NTCF family is *collapsing*, as defined below.²⁰

Definition 3.10. A NTCF family $(\text{Gen}, \text{Eval}, \text{Invert}, \text{Check}, \text{Good})$ is said to be *collapsing* if any QPT adversary A wins in the following game with probability $\frac{1}{2} + \text{negl}(\lambda)$:

1. The challenger generates $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$ and sends pk to A .
2. $A(\text{pk})$ generates a classical value $\mathbf{y} \in R_{\text{pk}}$ and an $(n(\lambda) + 1)$ -qubit quantum state $\sigma = \sigma_{\mathcal{S}, \mathcal{Z}}$, where the \mathcal{S} register contains a single qubit and the \mathcal{Z} register contains $n(\lambda)$ many qubits. A sends (\mathbf{y}, σ) to the challenger.
3. The challenger does the following:
 - (a) Apply in superposition the algorithm Check to σ , w.r.t. public key pk and the image string \mathbf{y} , and measure the bit indicating whether the output of Check is 1. If the output does not equal 1, send \perp to A . Otherwise, denote the resulting state by σ' .
 - (b) Choose a random bit $b \leftarrow \{0, 1\}$.
 - (c) If $b = 0$ then it send σ' to A .

²⁰This definition is similar to Definition 3.6 adapted to a NTCF family.

(d) If $b = 1$ then measure the \mathcal{S} register of σ' in the standard basis and send the resulting state to \mathbf{A} .

4. Upon receiving the quantum state (or the symbol \perp), \mathbf{A} outputs a bit b' .

5. \mathbf{A} wins if $b' = b$

Remark 3.11. An equivalent definition of collapsing is obtained by replacing Item 3d with the following: If $b = 1$ then send to \mathbf{A} the state $Z_{\mathcal{S}}[\sigma']$. In this work we use both of these formulations, since it is sometimes easier to work with one and other times with the other.

Claim 3.12. [Unr16] Every NTCF family is collapsing.

In what follows we define an extension of the collapsing property and argue that any NTCF family satisfies it. This extension may appear to be unnatural, but we make use of it when proving the binding property of our commitment schemes.

Claim 3.13. For every polynomial $\ell = \ell(\lambda)$, every NTCF family $(\text{Gen}, \text{Eval}, \text{Invert}, \text{Check}, \text{Good})$ is ℓ -extended collapsing, where the ℓ -extended collapsing definition asserts that every QPT adversary \mathbf{A} wins in the following extended collapsing game with probability $\frac{1}{2} + \text{negl}(\lambda)$:

1. The challenger generates ℓ independent public keys $\text{pk}_1, \dots, \text{pk}_\ell \leftarrow \text{Gen}(1^\lambda)$ and sends $(\text{pk}_1, \dots, \text{pk}_\ell)$ to \mathbf{A} .

2. $\mathbf{A}(\text{pk}_1, \dots, \text{pk}_\ell)$ generates a subset $J \subseteq [\ell]$, classical values $\{\mathbf{y}_j\}_{j \in J}$ where each $\mathbf{y}_j \in \mathbb{R}_{\text{pk}_j}$, and a $|J| \cdot (n(\lambda) + 1)$ -qubit quantum state $\sigma = \sigma_{\{\mathcal{S}_j, \mathcal{Z}_j\}_{j \in J}}$, where each register \mathcal{S}_j consists of a single qubit and each register \mathcal{Z}_j consists of $n(\lambda)$ -qubits.

\mathbf{A} sends $(J, \{\mathbf{y}_j\}_{j \in J}, \sigma)$ to the challenger.

3. The challenger does the following:

(a) For every $j \in J$ apply in superposition the algorithm Check to the $(\mathcal{S}_j, \mathcal{Z}_j)$ registers of σ w.r.t. pk_j and \mathbf{y}_j , and check that the output is 1. If this is not the case send \perp to \mathbf{A} .

(b) Otherwise, choose a random bit $b \leftarrow \{0, 1\}$ and measure the registers $\{\mathcal{S}_j\}_{j \in J}$ in the standard basis if and only if $b = 1$.

(c) Send the resulting state to \mathbf{A} .

4. Upon receiving a quantum state (or the symbol \perp), \mathbf{A} outputs a bit b' .

5. \mathbf{A} wins if $b' = b$

Claim 3.13 follows from Claim 3.12 together with a straightforward hybrid argument.

4 The Distributional Strong Adaptive Hardcore Bit Property

The binding property of our succinct commitment scheme relies on a variant of the adaptive hardcore bit property, which we define next.

Definition 4.1. A NTCF family $(\text{Gen}, \text{Eval}, \text{Invert}, \text{Check}, \text{Good})$ is said to have the **distributional strong adaptive hardcore bit** property if there exists a QPT algorithms \mathbf{A} and C such that the following holds: \mathbf{A} takes as input pk and a quantum state σ , and outputs a tuple $(\mathbf{y}, b, \mathbf{x}, \rho)$ such that $\text{Check}(\text{pk}, b, \mathbf{x}, \mathbf{y}) = 1$ and ρ is a state containing at least $n + 1$ qubits. Denote by \mathcal{O}_1 the registers containing the first $n + 1$ qubits, and denote by \mathcal{O}_2 all other registers. C takes as input the state $\rho_{\mathcal{O}_2}$ and outputs $\text{aux} \leftarrow C(\rho_{\mathcal{O}_2})$. Denote by ρ_{aux} the post measurement state, and assume that ρ satisfies that with overwhelming probability over aux , for every $\mathbf{d}' \in \{0, 1\}^k$ the probability of measuring registers \mathcal{O}_1 of ρ_{aux} in the standard basis and obtaining $\mathbf{d} \in \{0, 1\}^{n+1}$ such that $\text{Good}(\mathbf{d}, \mathbf{x}_0, \mathbf{x}_1) = \mathbf{d}'$ is negligible in λ . Then

$$(\text{pk}, \mathbf{y}, \mathbf{x}_0 \oplus \mathbf{x}_1, \mathbf{d} \cdot (1, \mathbf{x}_0 \oplus \mathbf{x}_1), \text{aux}) \approx (\text{pk}, \mathbf{y}, \mathbf{x}_0 \oplus \mathbf{x}_1, U, \text{aux}) \quad (4)$$

where $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$, $(\mathbf{y}, b, \mathbf{x}, \rho) \leftarrow \mathbf{A}(\text{pk}, \sigma)$, $((0, \mathbf{x}_0), (1, \mathbf{x}_1)) = \text{Invert}(\text{sk}, \mathbf{y})$, $(\text{aux}, \rho_{\text{aux}}) \leftarrow C(\rho_{\mathcal{O}_2})$, \mathbf{d} is obtained by measuring registers \mathcal{O}_1 of ρ_{aux} in the standard basis, and U is uniformly distributed in $\{0, 1\}$.

Moreover, there exists sk_{pre} , which is efficiently computable from sk , such that given $(\text{pk}, \mathbf{y}, b, \mathbf{x}_b)$ and sk_{pre} one can efficiently compute \mathbf{x}_{1-b} such that if $\text{Check}(\text{pk}, b, \mathbf{x}_b, \mathbf{y}) = 1$ then $\text{Check}(\text{pk}, 1 - b, \mathbf{x}_{1-b}, \mathbf{y}) = 1$, and Equation (4) holds even if C takes as input sk_{pre} in addition to the state $\rho_{\mathcal{O}_2}$.

We argue that the NTCF from [BCM⁺18], defined below, satisfies the distributional strong adaptive hardcore bit property (under LWE).

The NTCF family from [BCM⁺18] The NTCF family from [BCM⁺18] is a lattice based construction and makes use of the following theorem from [MP11].

Theorem 4.2 (Theorem 5.1 in [MP11]). *Let $n, m \geq 1$ and $q \geq 2$ be such that $m = \Omega(n \log q)$. There is an efficient randomized algorithm $\text{TrapGen}_{\text{MP}}(1^n, 1^m, q)$ that returns a matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ together with a trapdoor $\mathbf{t}_{\mathbf{A}} \in \mathbb{Z}_q^m$ such that the distribution of \mathbf{A} is negligibly (in n) close to the uniform distribution. Moreover, there is an efficient algorithm $\text{Invert}_{\text{MP}}$ that, on input $(\mathbf{A}, \mathbf{t}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e})$, where $\|\mathbf{e}\| \leq \frac{q}{C\sqrt{n \log q}}$ and where C is a universal constant, returns \mathbf{s} and \mathbf{e} with overwhelming probability over $(\mathbf{A}, \mathbf{t}_{\mathbf{A}}) \leftarrow \text{TrapGen}_{\text{MP}}(1^n, 1^m, q)$.*

The NTCF family $(\text{Gen}, \text{Eval}, \text{Invert}, \text{Check}, \text{Good})$ from [BCM⁺18] is defined as follows:

- $\text{Gen}(1^\lambda)$ is associated with the following:
 - Prime $q \leq 2^\lambda$ of size super-polynomial in λ .
 - $n = n(\lambda)$ and $m = m(\lambda)$, both polynomially bounded functions of λ , such that $m \geq n \log q$ and $n \geq \lambda$.
 - Two error distributions χ, χ' over \mathbb{Z}_q , that are associated with bounds $B, B' \in \mathbb{N}$ such that:
 1. $\frac{B}{B'} = \text{negl}(\lambda)$.
 2. $B' \leq \frac{q}{2C\sqrt{n \cdot m \cdot \log q}}$, where C is the universal constant from Theorem 4.2.
 3. $\Pr_{\mathbf{e} \leftarrow \chi^m} [\|\mathbf{e}\| > B] = \text{negl}(\lambda)$.
 4. $\Pr_{\mathbf{e}' \leftarrow (\chi')^m} [\|\mathbf{e}'\| > B'] = \text{negl}(\lambda)$.
 5. $\mathbf{e}' \equiv \mathbf{e}' + \mathbf{e}$, for $\mathbf{e} \leftarrow \chi^m$ and $\mathbf{e}' \leftarrow (\chi')^m$.

It does the following:

1. Generate $(\mathbf{A}, \mathbf{t}_\mathbf{A}) \leftarrow \text{TrapGen}_{\text{MP}}(1^n, 1^m, q)$.
 2. Choose a random bit string $\mathbf{s} \leftarrow \{0, 1\}^n$ and a random error vector $\mathbf{e} \leftarrow \chi^m$.
 3. Let $\mathbf{u} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e}$.
 4. Output $\text{pk} = (\mathbf{A}, \mathbf{u})$ and $\text{sk} = (\mathbf{A}, \mathbf{u}, \mathbf{t}_\mathbf{A})$.
- $\text{Eval}(\text{pk}, b, \cdot)$ is a function with domain \mathbb{Z}_q^n and range \mathbb{Z}_q^m . $\text{Eval}(\text{pk}, b, \mathbf{x})$ parses $\text{pk} = (\mathbf{A}, \mathbf{u})$, samples $\mathbf{e}' \leftarrow (\chi')^m$, and outputs $\mathbf{y} = \mathbf{A}\mathbf{x} + b\mathbf{u} + \mathbf{e}'$ (where all the operations are done modulo q).

Equivalently, we think of $\text{Eval}(\text{pk}, b, \cdot)$ as a function with domain $\{0, 1\}^w$ for $w \triangleq n \cdot \lceil \log q \rceil$, where each element $\mathbf{x} \in \mathbb{Z}_q^n$ is matched to its bit decomposition. Namely, denote by

$$J : \mathbb{Z}_q^n \rightarrow \{0, 1\}^w$$

the bit decomposition function where each element in \mathbb{Z}_q is converted to its bit decomposition in $\{0, 1\}^{\lceil \log q \rceil}$. We think Eval as taking as input an element $\mathbf{z} \in \{0, 1\}^w$, computing $\mathbf{x} = J^{-1}(\mathbf{z}) \in \mathbb{Z}_q^n$,²¹ and then applying $\text{Eval}(\text{pk}, b, \mathbf{x})$.

Remark 4.3. *We note the change in notation: In the definition of a NTCF family, we denoted the input length by n , and here we denote it by w .*

- $\text{Invert}(\text{sk}, \mathbf{y})$ does the following:
 1. Parse $\text{sk} = (\mathbf{A}, \mathbf{u}, \mathbf{t}_\mathbf{A})$.
 2. Compute $\text{Invert}_{\text{MP}}(\mathbf{A}, \mathbf{t}_\mathbf{A}, \mathbf{y}) = \mathbf{x}$
 3. If $\|\mathbf{y} - \mathbf{A} \cdot \mathbf{x}\| \leq 2\sqrt{m} \cdot B'$ then output $((0, J(\mathbf{x})), (1, J(\mathbf{x} - \mathbf{s})))$.
 4. Otherwise, output $((0, J(\mathbf{x} + \mathbf{s})), (1, J(\mathbf{x})))$.
- $\text{Check}(\text{pk}, b, J(\mathbf{x}), \mathbf{y})$ outputs 1 if and only if $\|\mathbf{y} - \mathbf{A}\mathbf{x} - b\mathbf{u}\| \leq 2\sqrt{m} \cdot B'$.
- $\text{Good}(J(\mathbf{x}_0), J(\mathbf{x}_1), \mathbf{d})$ outputs $\mathbf{d}' \in \{0, 1\}^n$ such that

$$\mathbf{d} \cdot (1, J(\mathbf{x}_0) \oplus J(\mathbf{x}_1)) = \mathbf{d}' \cdot \mathbf{s}$$

where the inner product in both sides of the equation above is done modulo 2. The vector \mathbf{d}' is computed as follows, using the fact that $\mathbf{x}_1 = \mathbf{x}_0 - \mathbf{s}$ (where subtraction is modulo q) and the fact that $\mathbf{s} \in \{0, 1\}^n$.

1. Partition $\mathbf{d} \in \{0, 1\}^{w+1}$ into its first bit, denoted by d_0 , and the following n blocks, each of size $\lceil \log q \rceil$, denoted by $\mathbf{d}[1], \dots, \mathbf{d}[n] \in \{0, 1\}^{\lceil \log q \rceil}$.
2. For every $b \in \{0, 1\}$, partition $J(\mathbf{x}_b)$ into blocks $J(\mathbf{x}_{b,1}), \dots, J(\mathbf{x}_{b,n})$, each of size $\lceil \log q \rceil$.
3. For every $i \in [n]$ let

$$\mathbf{d}'_i = \mathbf{d}[i] \cdot (J(\mathbf{x}_{0,i}) \oplus J(\mathbf{x}_{0,i} - 1)),$$

where \cdot denotes inner product mod 2 and where $\mathbf{x}_{0,i} - 1$ is done mod q

²¹ $J^{-1} : \{0, 1\}^w \rightarrow \mathbb{Z}_q^n$ is the function that breaks its input into n blocks of length $\lceil \log q \rceil$ each, and replaces each such block $(b_1, \dots, b_{\lceil \log q \rceil})$ with the element $(\sum_{i=1}^{\lceil \log q \rceil} b_i \cdot 2^{i-1}) \bmod q$, which is an element in \mathbb{Z}_q .

4. Output the string $\mathbf{d}' \in \{0, 1\}^n$.

Note that

$$\mathbf{d} \cdot (1, J(\mathbf{x}_0) \oplus J(\mathbf{x}_1)) = d_0 + \mathbf{d}' \cdot \mathbf{s} \pmod 2.$$

Claim 4.4. *The NTCF family from [BCM⁺18] satisfies the distributional strong adaptive hardcore bit property assuming the post-quantum hardness of LWE.*

The proof of Claim 4.4 makes use of the following lemma from [BCM⁺18].

Lemma 4.5. [BCM⁺18] *Let q be a prime, $k, n \geq 1$ integers, and $\mathbf{C} \leftarrow \mathbb{Z}_q^{k \times n}$ a uniformly random matrix. With probability at least $1 - q^{-k} \cdot 2^{-\frac{n}{8}}$ over the choice of \mathbf{C} the following holds for the fixed \mathbf{C} . For all $\mathbf{v} \in \mathbb{Z}_q^k$ and any distinct vectors $\mathbf{d}'_1, \mathbf{d}'_2 \in \{0, 1\}^n \setminus \{0^n\}$, the distribution of $(\mathbf{d}'_1 \cdot \mathbf{s}, \mathbf{d}'_2 \cdot \mathbf{s})$, where both inner products are done mod 2 and where $\mathbf{s} \leftarrow \{0, 1\}^n$ is uniform conditioned on $\mathbf{C}\mathbf{s} = \mathbf{v}$, is within statistical distance $O(q^{\frac{3k}{2}} \cdot 2^{-\frac{n}{40}})$ of the uniform distribution over $\{0, 1\}^2$.*

Remark 4.6. *We note that [BCM⁺18] proved this lemma for a single vector \mathbf{d}' (as opposed to two distinct ones). Their proof carries over to this setting as well, and we include it in Appendix B for completeness.*

Proof of Claim 4.4. We define sk_{pre} , corresponding to $\text{sk} = (\mathbf{A}, \mathbf{u}, \mathbf{t}_{\mathbf{A}})$, to be $\text{sk}_{\text{pre}} = \text{Invert}_{\text{MP}}(\mathbf{A}, \mathbf{t}_{\mathbf{A}}, \mathbf{u})$. Thus $\text{sk}_{\text{pre}} = \mathbf{s}$, where $\mathbf{u} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e}$ and \mathbf{e} is a low norm vector.

Fix any QPT circuit C and any QPT algorithm A that takes as input $\text{pk} = (\mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e})$ and a quantum state σ , and outputs a tuple $(\mathbf{y}, b, \mathbf{x}, \rho)$, such that $\text{Check}(\text{pk}, b, \mathbf{x}, \mathbf{y}) = 1$ and ρ is a state that has registers \mathcal{O}_1 and \mathcal{O}_2 , where \mathcal{O}_1 contains $w + 1$ qubits, and with overwhelming probability over $\text{aux} \leftarrow C(\rho_{\mathcal{O}_2}, \mathbf{s})$ it holds that the residual state ρ_{aux} satisfies that for every $\mathbf{d}' \in \{0, 1\}^n$ the probability of measuring registers \mathcal{O}_1 of ρ_{aux} in the standard basis and obtaining $\mathbf{d} \in \{0, 1\}^{w+1}$ such that $\text{Good}(\mathbf{d}, \mathbf{x}_0, \mathbf{x}_1) = \mathbf{d}'$ is negligible in λ .

We need to prove that

$$(\text{pk}, \mathbf{y}, J(\mathbf{x}_0) \oplus J(\mathbf{x}_1), \mathbf{d} \cdot (1, J(\mathbf{x}_0) \oplus J(\mathbf{x}_1)), \text{aux}) \approx (\text{pk}, \mathbf{y}, J(\mathbf{x}_0) \oplus J(\mathbf{x}_1), U, \text{aux}) \quad (5)$$

where $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$, $(\mathbf{y}, b, \mathbf{x}, \rho, C) \leftarrow A(\text{pk}, \sigma)$, $((0, \mathbf{x}_0), (1, \mathbf{x}_1)) = \text{Invert}(\text{sk}, \mathbf{y})$, (\mathbf{d}, aux) is obtained by letting $\text{aux} \leftarrow C(\rho_{\mathcal{O}_2}, \mathbf{s})$ and \mathbf{d} is the outcome of measuring the \mathcal{O}_1 registers of the residual state ρ_{aux} in the standard basis, and U is uniformly distributed in $\{0, 1\}$.

To this end, we define an alternative algorithm $\widehat{\text{Gen}}(1^\lambda)$, which is the same as $\text{Gen}(1^\lambda)$ with the only difference being that rather than choosing $(\mathbf{A}, \mathbf{t}_{\mathbf{A}})$ via the TrapGen algorithm, it chooses \mathbf{A} to be close to a low rank matrix. Specifically, $\widehat{\text{Gen}}(1^\lambda)$ does the following:

1. Let $\delta = \epsilon/2$ and let $k = n^\delta$.
2. Sample $(\mathbf{B}, \mathbf{t}_{\mathbf{B}}) \leftarrow \text{TrapGen}_{\text{MP}}(1^k, 1^m, q)$.
3. Sample $\mathbf{C} \leftarrow \{0, 1\}^{k \times n}$ and $\mathbf{N} \leftarrow \chi^{m \times n}$.
4. Let $\widehat{\mathbf{A}} = \mathbf{B} \cdot \mathbf{C} + \mathbf{N}$.
5. Sample $\mathbf{s} \leftarrow \{0, 1\}^n$ and $\mathbf{e} \leftarrow \chi^m$.
6. Let $\widehat{\mathbf{u}} = \widehat{\mathbf{A}} \cdot \mathbf{s} + \mathbf{e}$.

7. Output $\widehat{\mathbf{pk}} = (\widehat{\mathbf{A}}, \widehat{\mathbf{u}})$ and $\widehat{\mathbf{sk}} = (\widehat{\mathbf{A}}, \mathbf{B}, \mathbf{t}_B, \mathbf{s})$.

The LWE implies that $\mathbf{A} \approx \widehat{\mathbf{A}}$ which in turn implies that

$$(\mathbf{A}, \mathbf{u}, \mathbf{s}, \mathbf{y}, b, \mathbf{x}, \boldsymbol{\rho}) \approx (\widehat{\mathbf{A}}, \widehat{\mathbf{u}}, \mathbf{s}, \widehat{\mathbf{y}}, \widehat{b}, \widehat{\mathbf{x}}, \widehat{\boldsymbol{\rho}}) \quad (6)$$

where

- $\mathbf{s} \leftarrow \{0, 1\}^n$ and $\mathbf{e} \leftarrow \chi^m$.
- $\mathbf{u} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e}$ and $\widehat{\mathbf{u}} = \widehat{\mathbf{A}} \cdot \mathbf{s} + \mathbf{e}$.
- $(\mathbf{y}, b, \mathbf{x}, \boldsymbol{\rho}) = A(\mathbf{pk}, \boldsymbol{\sigma})$ for $\mathbf{pk} = (\mathbf{A}, \mathbf{u})$.
- $(\widehat{\mathbf{y}}, \widehat{b}, \widehat{\mathbf{x}}, \widehat{\boldsymbol{\rho}}) = A(\widehat{\mathbf{pk}}, \boldsymbol{\sigma})$ for $\widehat{\mathbf{pk}} = (\widehat{\mathbf{A}}, \widehat{\mathbf{u}})$.

We next argue that

$$(\widehat{\mathbf{A}}, \mathbf{u}, \mathbf{y}, J(\mathbf{x}_0) \oplus J(\mathbf{x}_1), \mathbf{d} \cdot (1, J(\mathbf{x}_0) \oplus J(\mathbf{x}_1)), \mathbf{aux}) \approx (\widehat{\mathbf{A}}, \mathbf{u}, \mathbf{y}, J(\mathbf{x}_0) \oplus J(\mathbf{x}_1), U, \mathbf{aux}) \quad (7)$$

where in the above equation, to avoid cluttering of notation, we omit some of the “hat” notation, and denote by $\widehat{\mathbf{pk}} = (\widehat{\mathbf{A}}, \mathbf{u})$ distributed according to $\widehat{\text{Gen}}(1^\lambda)$, $(\mathbf{y}, b, \mathbf{x}, \boldsymbol{\rho}) \leftarrow A(\widehat{\mathbf{pk}}, \boldsymbol{\sigma})$ for $\widehat{\mathbf{pk}} = (\widehat{\mathbf{A}}, \mathbf{u})$, $\mathbf{x}_b = \mathbf{x}$, $\mathbf{x}_{1-b} = \mathbf{x}_b - (-1)^b \mathbf{s}$, and $(\mathbf{d}, \mathbf{aux})$ is obtained by measuring computing $\mathbf{aux} \leftarrow C(\boldsymbol{\rho}_{\mathcal{O}_2}, \mathbf{s})$ and \mathbf{d} is the outcome of measuring the \mathcal{O}_1 registers of the residual state $\boldsymbol{\rho}_{\mathbf{aux}}$ in the standard basis. Equation (5) follows immediately from Equation (7), together with the fact that $\mathbf{A} \approx \widehat{\mathbf{A}}$ and the fact that the distributions in Equations (5) and (7) can be generated efficiently from \mathbf{A} and $\widehat{\mathbf{A}}$, respectively.

Let

$$\mathbf{d}' = \text{Good}(\mathbf{x}_0, \mathbf{x}_1, \mathbf{d}) \in \{0, 1\}^n,$$

then by the definition of `Good`,

$$\mathbf{d} \cdot (1, J(\mathbf{x}_0) \oplus J(\mathbf{x}_1)) = d_0 \oplus \mathbf{d}' \cdot \mathbf{s}.$$

Therefore, to prove Equation (7) it suffices to prove that

$$(\widehat{\mathbf{A}}, \mathbf{u}, \mathbf{s}, \mathbf{y}, b, \mathbf{x}, d_0, \mathbf{d}' \cdot \mathbf{s}, \mathbf{aux}) \approx (\widehat{\mathbf{A}}, \mathbf{u}, \mathbf{s}, \mathbf{y}, b, \mathbf{x}, d_0, U, \mathbf{aux}).$$

To prove the above equation it suffices to prove that with overwhelming probability over $\mathbf{C} \leftarrow \mathbb{Z}_q^{k \times n}$ it holds that for every $\mathbf{v} \in \mathbb{Z}_q^k$, for every distribution $D_{\mathbf{v}}$ (that depends on \mathbf{v}) that outputs $(\mathbf{d}', \boldsymbol{\rho}_{\mathcal{O}_2})$ such that with overwhelming probability \mathbf{d}' has min-entropy $\omega(\log \lambda)$ given $\mathbf{aux} \leftarrow C(\boldsymbol{\rho}_{\mathcal{O}_2}, \mathbf{s})$,

$$(\mathbf{s}, \mathbf{aux}, \mathbf{d}' \cdot \mathbf{s}) \equiv (\mathbf{s}, \mathbf{aux}, U), \quad (8)$$

where \mathbf{s} is sampled randomly from $\{0, 1\}^n$ conditioned on $\mathbf{C}\mathbf{s} = \mathbf{v}$. We note that the distribution of \mathbf{d}' may not have min-entropy $\omega(\log \lambda)$ (conditioned on \mathbf{aux}) since it is generated w.r.t. $\widehat{\mathbf{A}}$ and not w.r.t. \mathbf{A} , and while $\mathbf{A} \approx \widehat{\mathbf{A}}$, checking if a distribution has min-entropy cannot be done efficiently. Nevertheless, the distribution $D_{\mathbf{v}}$ is indistinguishable from having min-entropy $\omega(\log \lambda)$, and thus it suffices to prove Equation (8). To this end, for every $\mathbf{C} \in \mathbb{Z}_q^{k \times n}$ and $\mathbf{v} \in \mathbb{Z}_q^k$ consider the sets

$$\mathcal{S} = \{\mathbf{x} \in \{0, 1\}^n : \mathbf{C}(\mathbf{x}) = \mathbf{v}\} \quad \text{and} \quad \mathcal{X} = \{0, 1\}^n \setminus \{0^n\}$$

and the hash function

$$h : \mathcal{S} \times \mathcal{X} \rightarrow \{0, 1\},$$

defined by

$$h(\mathbf{x}, \mathbf{d}') = \mathbf{x} \cdot \mathbf{d}' \pmod{2}.$$

Lemma 4.5 implies that for all but negligible fraction of \mathbf{C} and \mathbf{v} it holds that h is 2-universal, which together with the leftover hash lemma, implies that Equation (8) holds, as desired. \square

5 Classical Commitments to Quantum States

In this section we define the notion of a classical commitment to quantum states. Our definition is stronger than the notion of a *measurement protocol*, originally considered in [Mah18] and formally defined in [BKL⁺22],²² in several ways. First, the opening basis is not determined during the key generation phase. Namely, the key generation algorithm, Gen , takes as input only the security parameter (in unary), as opposed to taking both the security parameter and the opening basis. In particular, the opening basis can be determined after the commitment phase, and can be chosen adaptively based on any information that the parties have access to. Importantly, our binding property is significantly stronger. It guarantees that for any QPT cheating committer $\text{C}^*.\text{Commit}$ that commits to an ℓ -qubit quantum state, there is a *single* extracted quantum state τ such that for *any* QPT algorithm $\text{C}^*.\text{Open}$ and *any* basis opening (b_1, \dots, b_ℓ) , where $b_i = 0$ corresponds to measuring the i 'th qubit in the standard basis and $b_i = 1$ corresponds to measuring it in the Hadamard basis, the opening obtained by $\text{C}^*.\text{Open}(b_1, \dots, b_\ell)$ is computationally indistinguishable from measuring τ in basis (b_1, \dots, b_ℓ) , assuming the opening of $\text{C}^*.\text{Open}$ is accepted. In contrast, in the soundness guarantee of a weak commitment scheme, the extracted state τ may depend on $\text{C}^*.\text{Open}$, which can be chosen adaptively after the commitment phase and hence is not truly binding.²³

5.1 Syntax

In what follows we define the syntax of a commitment scheme. We present two definitions. The first is the syntax for a *non-succinct* commitment scheme, where the length of the commitment string that commits to an ℓ -qubit quantum state, grows polynomially with ℓ . More specifically, in this definition the length ℓ is determined in the key generation algorithm, and the run-time of all the algorithms grow polynomially with ℓ . In Section 5.1.1 we define the syntax for a *succinct* commitment scheme, where the verifier's run-time grows poly-logarithmically with ℓ .

Definition 5.1. *A (non-succinct) classical commitment scheme for quantum states is associated with algorithms $(\text{Gen}, \text{Commit}, \text{Open}, \text{Ver}, \text{Out})$ and has the following syntax:*

1. Gen is a PPT algorithm that takes as input a security parameter λ and a length parameter ℓ (both in unary), and outputs a pair $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda, 1^\ell)$, where pk is referred to as the public key and sk is referred to as the secret key.

²²We refer to this weaker notion as a “weak classical commitment,” and recall its definition in Appendix A (for completeness).

²³Indeed, the weak classical commitment scheme from [Mah18] is only binding in the standard basis, and offers no binding guarantees when opening in the Hadamard basis.

2. **Commit** is a QPT algorithm that takes as input a public key \mathbf{pk} and an ℓ -qubit quantum state σ and outputs a pair $(\mathbf{y}, \rho) \leftarrow \text{Commit}(\mathbf{pk}, \sigma)$, where \mathbf{y} is a classical string referred to as the commitment string and ρ is a quantum state.
3. **Open** is a QPT algorithm that takes as input a quantum state ρ and a basis $(b_1, \dots, b_\ell) \in \{0, 1\}^\ell$ (where $b_j = 0$ corresponds to opening the j 'th bit in the standard basis and $b_j = 1$ corresponds to opening it in the Hadamard basis). It outputs a pair $(\mathbf{z}, \rho') \leftarrow \text{Open}(\rho, (b_1, \dots, b_\ell))$, where \mathbf{z} is a classical string, referred to as the opening string, and ρ' is the residual state (which is sometimes omitted).
4. **Ver** is a polynomial time algorithm that takes a tuple $(\mathbf{sk}, \mathbf{y}, (b_1, \dots, b_\ell), \mathbf{z})$, where \mathbf{sk} is a secret key, \mathbf{y} is a commitment string (to the quantum state), $(b_1, \dots, b_\ell) \in \{0, 1\}^\ell$ is a string specifying the opening basis, and \mathbf{z} is an opening string. It outputs 0 (if \mathbf{z} is not a valid opening) and outputs 1 otherwise.
5. **Out** is a polynomial time algorithm that takes a tuple $(\mathbf{sk}, \mathbf{y}, (b_1, \dots, b_\ell), \mathbf{z})$ (as above), and outputs an ℓ -bit string $\mathbf{m} \leftarrow \text{Out}(\mathbf{sk}, \mathbf{y}, (b_1, \dots, b_\ell), \mathbf{z})$.

The protocol associated with the tuple $(\text{Gen}, \text{Commit}, \text{Open}, \text{Ver}, \text{Out})$ is a two party protocol between a QPT committer \mathbf{C} and a PPT verifier \mathbf{V} and consists of two phases, COMMIT and OPEN. During the COMMIT phase, \mathbf{V} takes as input security parameter λ and a length parameter ℓ and \mathbf{C} takes in an arbitrary quantum state σ . During the OPEN phase, \mathbf{V} takes as input a basis bit $(b_1, \dots, b_\ell) \in \{0, 1\}^\ell$. The protocol proceeds as follows:

- COMMIT phase:
 1. $[\mathbf{C} \leftarrow \mathbf{V}]$: \mathbf{V} samples $(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Gen}(1^\lambda, 1^\ell)$ and sends the public key \mathbf{pk} to \mathbf{C} .
 2. $[\mathbf{C} \rightarrow \mathbf{V}]$: \mathbf{C} computes $(\mathbf{y}, \rho) \leftarrow \text{Commit}(\mathbf{pk}, \sigma)$ and sends the commitment string \mathbf{y} to \mathbf{V} .
- OPEN phase:
 1. $[\mathbf{C} \leftarrow \mathbf{V}]$: \mathbf{V} sends an opening basis (b_1, \dots, b_ℓ) to \mathbf{C} .
 2. $[\mathbf{C} \rightarrow \mathbf{V}]$: \mathbf{C} computes $(\mathbf{z}, \rho') \leftarrow \text{Open}(\rho, (b_1, \dots, b_\ell))$ and sends \mathbf{z} to \mathbf{V} .
 3. $[\mathbf{V}]$: \mathbf{V} checks that $\text{Ver}(\mathbf{sk}, \mathbf{y}, (b_1, \dots, b_\ell), \mathbf{z}) = 1$, and if so it outputs $\mathbf{m} \leftarrow \text{Out}(\mathbf{sk}, \mathbf{y}, (b_1, \dots, b_\ell), \mathbf{z})$ as the decommitment. Otherwise, it outputs \perp .

Remark 5.2. One could define **Open**, **Ver**, **Out** to operate on one qubit at a time. Namely, one could define **Open** to take as input a quantum state ρ an index $j \in [\ell]$ and a basis $b \in \{0, 1\}$, and output a pair $(\mathbf{z}, \rho') \leftarrow \text{Open}(\rho, (j, b))$, and define **Ver** and **Out** to take as input $(\mathbf{sk}, \mathbf{y}, (j, b), \mathbf{z})$ and output a bit (indicating accept/reject for **Ver** and indicating an output bit for **Out**). Indeed, in our definition of a succinct classical commitment to quantum state, stated in Section 5.1.1 below, **Open**, **Ver** and **Out** operate on one qubit at a time. In addition, our constructions in Section 6 are defined where **Open**, **Ver** and **Out** operate on one qubit at a time.

Note that in the syntax above the length of the public key \mathbf{pk} as well as the length of the commitment \mathbf{y} grows with the number of qubits in the committed state (denoted by ℓ). In this work we also construct *succinct* commitments where the length of \mathbf{pk} and the commitment \mathbf{y} grow only with the security parameter (and grow only poly-logarithmically with ℓ). In what follows we define the syntax of a *succinct* classical commitment scheme for multi-qubit quantum states.

5.1.1 Syntax for Succinct Commitments

The syntax of a succinct commitment is similar to that of a non-succinct commitment scheme (defined above), with the following main differences:

1. The key generation algorithm (**Gen**) takes as input only the security parameter λ and does not depend on the size ℓ of the committed quantum state.

This change ensures that the runtime of **Gen** does not grow with ℓ .²⁴

2. The opening algorithm (**Open**) opens one qubit at a time. Namely, it takes as input the post-commitment quantum state ρ , a single index $j \in [\ell]$ and a basis $b \in \{0, 1\}$, and it outputs an opening to the j 'th qubit.

The reason for this change is that in some of our applications we commit to a long quantum state but open only a small portion of it. For example, this is the case in our compilation of a X/Z quantum PCP into a succinct argument (see Section 8).

3. The succinct commitment has two additional components. The first is an interactive protocol that verifies that the prover “knows” a non-succinct commitment string corresponding to this succinct commitment. This protocol is referred to as **Ver.Commit**, and is a protocol between a poly-time (classical) prover P and a PPT verifier V .²⁵ The second is a **Test** protocol that tests that the committer can open all the qubits in a valid manner. We note that **Test** is executed with probability $1/2$, and if it is executed then **Open** is not executed (since **Test** destroys the quantum state needed for the **Open** algorithm).²⁶

Formally a succinct commitment scheme for quantum states consists of

$$(\text{Gen}, \text{Commit}, \text{Ver.Commit}, \text{Test}, \text{Open}, \text{Ver}, \text{Out})$$

such that

1. **Gen** is a PPT algorithm that takes as the security parameter λ (in unary) and outputs a pair $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$.
2. **Commit** is a QPT algorithm that takes as input a public key pk and an ℓ -qubit quantum state σ and outputs a tuple $(\text{rt}, \mathbf{y}, \rho) \leftarrow \text{Commit}(\text{pk}, \sigma)$, where rt is a succinct classical commitment to σ (of size $\text{poly}(\lambda, \log \ell)$), \mathbf{y} is its non-succinct counterpart, and ρ is the residual quantum state.
3. **Ver.Commit** is an interactive protocol between a poly-time prover P with input $(\text{pk}, \text{rt}, \mathbf{y})$ and a PPT verifier V with input (sk, rt) . At the end of the protocol, V outputs a verdict bit in $\{0, 1\}$, corresponding to accept or reject (1 corresponding to accept and 0 corresponding to reject). The communication complexity is $\text{poly}(\lambda, \log \ell)$.
4. **Test** is an interactive protocol between a QPT prover P_{Test} with input $(\text{pk}, \text{rt}, \mathbf{y}, \rho)$ and a BPP verifier V_{Test} with input (sk, rt) . At the end of the protocol, V outputs a verdict bit in $\{0, 1\}$, corresponding to accept or reject.

²⁴One could give **Gen** the parameter ℓ in binary, but our scheme does not require it.

²⁵We note that since P is a classical algorithm the quantum state remains unchanged.

²⁶This is also the case in the **Test** phase in Mahadev’s measurement protocol [Mah18].

5. **Open** is a QPT algorithm that takes as input a quantum state ρ , an index $j \in [\ell]$, and a basis $b_j \in \{0, 1\}$ (where $b_j = 0$ corresponds to measuring the j 'th qubit in the standard basis and $b_j = 1$ corresponds to measuring it in the Hadamard basis). It outputs a pair $(\mathbf{z}, \rho') \leftarrow \text{Open}(\rho, (j, b_j))$, where \mathbf{z} is a classical string of length $\text{poly}(\lambda, \log \ell)$, referred to as the *opening string*, and ρ' is the residual state (which is sometimes omitted).
6. **Ver** is a polynomial time algorithm that takes a tuple $(\text{sk}, \text{rt}, (j, b_j), \mathbf{z})$, where sk is a secret key, rt is a succinct classical commitment string to an ℓ -qubit quantum state, $j \in [\ell]$, $b_j \in \{0, 1\}$ is a bit specifying the opening basis, and \mathbf{z} is an opening string. It outputs 0 (if \mathbf{z} is not a valid opening) and outputs 1 otherwise.
7. **Out** is a polynomial time algorithm that takes a tuple $(\text{sk}, \text{rt}, (j, b_j), \mathbf{z})$, and outputs a bit $m \leftarrow \text{Out}(\text{sk}, \text{rt}, (j, b_j), \mathbf{z})$.

Remark 5.3. We extend **Ver** and **Out** to take as input $(\text{sk}, \text{rt}, (J, \mathbf{b}_J), \mathbf{z})$ instead of $(\text{sk}, \text{rt}, (j, b_j), \mathbf{z})$, where $J \subseteq [\ell]$ and $\mathbf{b}_J \in \{0, 1\}^{|J|}$, in which case the algorithms run with input $(\text{sk}, \text{rt}, (j, b_j), \mathbf{z})$ for every $j \in J$. We extend **Open** in a similar manner.

The succinct commitment protocol associated with the tuple

$$(\text{Gen}, \text{Commit}, \text{Ver.Commit}, \text{Test}, \text{Open}, \text{Ver}, \text{Out})$$

is a two party protocol that consists of three phases, COMMIT, CHECK and OPEN, as follows:

- **COMMIT** phase:
 1. $[\text{C} \leftarrow \text{V}]$: V samples $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$ and sends the public key pk to C .
 2. $[\text{C} \rightarrow \text{V}]$: C computes $(\text{rt}, \mathbf{y}, \rho) \leftarrow \text{Commit}(\text{pk}, \sigma)$ and sends the succinct commitment string rt to V .
- **CHECK** phase:
 1. Run **Ver.Commit** protocol between the prover P with input $(\text{pk}, \text{rt}, \mathbf{y})$ and the verifier V with input (sk, rt) . If V rejects then the commitment rt is rejected and the protocol ends.
 2. Otherwise, choose at random $c \leftarrow \{0, 1\}$.
 3. If $c = 0$ then go to the **OPEN** phase.
 4. If $c = 1$ then run the **Test** protocol, where the prover P_{Test} takes as input $(\text{pk}, \text{rt}, \mathbf{y}, \rho)$ and the verifier V_{Test} takes as input (sk, rt) . If V_{Test} rejects then the commitment rt is rejected and otherwise it is accepted. At the end of the **Test** protocol the commitment protocol ends (the **Open** phase is not executed).
- **OPEN** phase:
 1. $[\text{C} \leftarrow \text{V}]$: V sends a subset $J \subseteq [\ell]$ and an opening basis $\mathbf{b}_J \in \{0, 1\}^{|J|}$ to C .
 2. $[\text{C} \rightarrow \text{V}]$: C computes $(\mathbf{z}, \rho') \leftarrow \text{Open}(\rho, (J, \mathbf{b}_J))$ and sends \mathbf{z} to V .
 3. $[\text{V}]$: V checks that $\text{Ver}(\text{sk}, \text{rt}, (J, \mathbf{b}_J), \mathbf{z}) = 1$, and if so outputs $\mathbf{m} \leftarrow \text{Out}(\text{sk}, \text{rt}, (J, \mathbf{b}_J), \mathbf{z})$ as the decommitment bit. Otherwise, it outputs \perp .

5.2 Properties

We require that a commitment scheme satisfies two properties, *correctness* and *binding*, defined below.

5.2.1 Correctness

We define the correctness guarantee separately for the non-succinct and the succinct setting, starting with the former.

Definition 5.4 (Correctness). *A (non-succinct) classical commitment scheme is correct if for any ℓ -qubit quantum state σ , and any basis $\mathbf{b} = (b_1, \dots, b_\ell) \in \{0, 1\}^\ell$,*

$$\text{Real}(1^\lambda, \sigma, \mathbf{b}) \equiv \sigma(\mathbf{b}), \quad (9)$$

where $\sigma(\mathbf{b})$ is the distribution obtained by measuring each qubit j of σ in the basis specified by b_j (standard if $b_j = 0$, Hadamard if $b_j = 1$), and $\text{Real}(1^\lambda, \sigma, \mathbf{b})$ is the distribution resulting from the following experiment:

1. Generate $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda, 1^\ell)$.
2. Generate $(\mathbf{y}, \rho) \leftarrow \text{Commit}(\text{pk}, \sigma)$.
3. Compute $(\mathbf{z}, \rho') \leftarrow \text{Open}(\rho, \mathbf{b})$.
4. If $\text{Ver}(\text{sk}, \mathbf{y}, \mathbf{b}, \mathbf{z}) = 0$ then output \perp .
5. Otherwise, output $\text{Out}(\text{sk}, \mathbf{y}, \mathbf{b}, \mathbf{z})$.

Definition 5.5 (Succinct Correctness). *A succinct classical commitment scheme is correct if for any ℓ -qubit quantum state σ , any basis $\mathbf{b} = (b_1, \dots, b_\ell) \in \{0, 1\}^\ell$, and any subset $J \subseteq [\ell]$, the following two conditions holds:*

$$\text{Real}_{c=0}(1^\lambda, \sigma, J, \mathbf{b}_J) \equiv \sigma(J, \mathbf{b}_J) \quad \text{and} \quad \text{Real}_{c=1}(1^\lambda, \sigma, J, \mathbf{b}_J) \equiv 1 \quad (10)$$

where:

- $\sigma(J, \mathbf{b}_J)$ is the distribution obtained by measuring each qubit $j \in J$ of σ in the basis specified by b_j (standard if $b_j = 0$ and Hadamard if $b_j = 1$).
- $\text{Real}_{c=0}(1^\lambda, \sigma, J, \mathbf{b}_J)$ is the distribution resulting from the following experiment:
 1. Generate $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$.
 2. Generate $(\text{rt}, \mathbf{y}, \rho) \leftarrow \text{Commit}(\text{pk}, \sigma)$.
 3. Run the protocol Ver.Commit between P with input $(\text{pk}, \text{rt}, \mathbf{y})$ and V with input (sk, rt) . If V rejects then then output \perp .
 4. Otherwise, compute $(\mathbf{z}, \rho') \leftarrow \text{Open}(\rho, (J, \mathbf{b}_J))$.
 5. If $\text{Ver}(\text{sk}, \text{rt}, (J, \mathbf{b}_J), \mathbf{z}) = 0$ then output \perp .
 6. Otherwise, output $\text{Out}(\text{sk}, \text{rt}, (J, \mathbf{b}_J), \mathbf{z})$.

- $\text{Real}_{c=1}(1^\lambda, \sigma, J, \mathbf{b}_J)$ is the distribution resulting from the following experiment:

1. Generate $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$.
2. Generate $(\text{rt}, \mathbf{y}, \rho) \leftarrow \text{Commit}(\text{pk}, \sigma)$.
3. Run the protocol Ver.Commit between P with input $(\text{pk}, \text{rt}, \mathbf{y})$ and V with input (sk, rt) . If V rejects then then output \perp .
4. Execute Test where the prover P_{Test} takes as input $(\text{pk}, \text{rt}, \mathbf{y}, \rho)$ and the verifier V_{Test} takes as input (sk, rt) . If V_{Test} rejects then output \perp and if it accepts then output 1.

5.2.2 Binding

In what follows we define the binding condition. Intuitively, our binding guarantee is that a cheating committer cannot change the way they open based on *any* information they learn after the commitment phase, and that the opening distribution is consistent with the distribution of a qubit.

For simplicity, we consider only cheating algorithms that are accepted with high probability. This can be ensured by repetition. Namely, for every $\epsilon, \delta > 0$ by repeating the commitment and opening protocol $O\left(\frac{\log(1/\epsilon)}{\delta}\right)$ times, if a cheating C^* is accepted in all of executions with probability at least ϵ then a random execution is accepted with probability at least $1 - \delta$.

We first define the binding property for the (non-succinct) commitment and then define it for the succinct commitment.

Definition 5.6 (Binding). *A classical (non-succinct) commitment scheme to a multi-qubit quantum state is said to be computationally binding if there exists a QPT oracle machine Ext such that for any QPT algorithm $C^*.\text{Commit}$, any $\text{poly}(\lambda)$ -size quantum state σ , any polynomial $\ell = \ell(\lambda)$, any basis $\mathbf{b} = (b_1, \dots, b_\ell)$, and any QPT algorithms $C_1^*.\text{Open}$ and $C_2^*.\text{Open}$, for every $i \in \{1, 2\}$*

$$\text{Real}^{C^*.\text{Commit}, C_i^*.\text{Open}}(\lambda, \mathbf{b}, \sigma) \stackrel{\eta}{\approx} \text{Ideal}^{\text{Ext}, C^*.\text{Commit}, C_i^*.\text{Open}}(\lambda, \mathbf{b}, \sigma) \quad (11)$$

and

$$\text{Real}^{C^*.\text{Commit}, C_1^*.\text{Open}}(\lambda, \mathbf{b}, \sigma) \stackrel{\eta}{\approx} \text{Real}^{C^*.\text{Commit}, C_2^*.\text{Open}}(\lambda, \mathbf{b}, \sigma) \quad (12)$$

where $\eta = O\left(\sqrt{\delta}\right)$ and

$$\delta = \mathbb{E}_{\substack{(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda, 1^\ell) \\ (\mathbf{y}, \rho) \leftarrow C^*.\text{Commit}(\text{pk}, \sigma)}} \max_{\substack{i \in \{1, 2\}, \\ \mathbf{b}' \in \{\mathbf{b}, \mathbf{0}, \mathbf{1}\}}} \Pr[\text{Ver}(\text{sk}, \mathbf{y}, \mathbf{b}', C_i^*.\text{Open}(\rho, \mathbf{b}')) = 0]. \quad (13)$$

and where $\text{Real}^{C^*.\text{Commit}, C^*.\text{Open}}(\lambda, \mathbf{b}, \sigma)$ is defined as follows:

- $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda, 1^\ell)$.
 - $(\mathbf{y}, \rho) \leftarrow C^*.\text{Commit}(\text{pk}, \sigma)$.
1. Compute $(\mathbf{z}, \rho') \leftarrow C^*.\text{Open}(\rho, \mathbf{b})$.
 2. If $\text{Ver}(\text{sk}, \mathbf{y}, \mathbf{b}, \mathbf{z}) = 0$ then output \perp .
 3. Otherwise, let $\mathbf{m} = \text{Out}(\text{sk}, \mathbf{y}, \mathbf{b}, \mathbf{z})$.

4. Output $(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m})$.

$\text{Ideal}^{\text{Ext}, \text{C}^*. \text{Commit}, \text{C}^*. \text{Open}}(\lambda, \mathbf{b}, \boldsymbol{\sigma})$ is defined as follows:

1. $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda, 1^\ell)$.
2. $(\mathbf{y}, \boldsymbol{\rho}) \leftarrow \text{C}^*. \text{Commit}(\text{pk}, \boldsymbol{\sigma})$.
3. Let $\tau_{\mathcal{A}, \mathcal{B}} = \text{Ext}^{\text{C}^*. \text{Open}}(\text{sk}, \mathbf{y}, \boldsymbol{\rho})$.
4. Measure $\tau_{\mathcal{A}}$ in the basis $\mathbf{b} = (b_1, \dots, b_\ell)$ to obtain $\mathbf{m} \in \{0, 1\}^\ell$.
5. Output $(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m})$.

Remark 5.7. Throughout this write-up to avoid cluttering of notation we omit the superscript $\text{C}^*. \text{Commit}$ from

$$\text{Real}^{\text{C}^*. \text{Commit}, \text{C}^*. \text{Open}}(\lambda, \mathbf{b}, \boldsymbol{\sigma}) \quad \text{and} \quad \text{Ideal}^{\text{Ext}, \text{C}^*. \text{Commit}, \text{C}^*. \text{Open}}(\lambda, \mathbf{b}, \boldsymbol{\sigma}),$$

and denote these by

$$\text{Real}^{\text{C}^*. \text{Open}}(\lambda, \mathbf{b}, \boldsymbol{\sigma}) \quad \text{and} \quad \text{Ideal}^{\text{Ext}, \text{C}^*. \text{Open}}(\lambda, \mathbf{b}, \boldsymbol{\sigma}),$$

respectively.

Remark 5.8. We prove that our commitment scheme is sound with $\eta \leq 10\sqrt{\delta}$. Note that δ is a bound on the probability that the openings of $\text{C}_i^*. \text{Open}$ are rejected not only on basis \mathbf{b} , but also on basis $\mathbf{0}$ and $\mathbf{1}$. We note that for Equation (12) we do not need to bound the probability that $\text{C}_i^*. \text{Open}$ is rejected on basis $\mathbf{0}$ and basis $\mathbf{1}$, and indeed we do not bound these probabilities in the proof (see Lemma 7.10). The reason we need to bound these probabilities to prove Equation (11) is that our extractor uses the openings of $\text{C}_i^*. \text{Open}$ on basis $\mathbf{0}$ and $\mathbf{1}$ to extract the quantum state.

Remark 5.9. We mention that we prove a stronger condition than the one given in Equation (12). This is done in Lemma 7.10. The strengthening is due to two reasons. First, we prove Equation (12) by induction on ℓ , and for the induction step to go through we need to strengthen the induction hypothesis, and as a result we prove a stronger guarantee. Second, we allow the cheating algorithm $\text{C}^*. \text{Open}$ to depend on a part of the secret key sk . This is needed to obtain our succinct interactive argument for **QMA** in Section 8.1 and is needed for our applications in Section 8.

Definition 5.6 assumes that $\text{C}^*. \text{Open}$ opens all the qubits of the committed state. Indeed, we use $\text{C}^*. \text{Open}$ to extract an ℓ -qubit quantum state. In what follows we define the notion of binding for a succinct commitment, where $\text{C}^*. \text{Open}$ may only open to a subset $J \subseteq [\ell]$ of the qubits, and hence cannot be used to extract the entire state (as was done in Definition 5.6). While we can extract a state consisting of $|J|$ qubits from $\text{C}^*. \text{Open}$, for our applications we will need to extract the entire ℓ -qubit state, even if $\text{C}^*. \text{Open}$ only opens to the qubits in J (without blowing up the communication). This is precisely the purpose of the **Ver.Commit** protocol; instead of extracting from $\text{C}^*. \text{Open}$ we extract from the (cheating) prover P^* of the **Ver.Commit** protocol. We note that even though **Ver.Commit** is a succinct protocol, its interactive nature will allow us to extract the (non-succinct) ℓ -qubit state from P^* .

Definition 5.10 (Succinct Binding). *A succinct classical commitment scheme to a multi-qubit quantum state is said to be computationally binding if there exists a QPT oracle machine Ext such that for any QPT algorithm $\text{C}^*. \text{Commit}$, any $\text{poly}(\lambda)$ -size quantum state σ , any polynomial $\ell = \ell(\lambda)$, any QPT prover P^* for the Ver. Commit protocol, any QPT prover P_{Test}^* for the Test protocol, any $J \subseteq [\ell]$ and $\mathbf{b}_J = (b_j)_{j \in J} \in \{0, 1\}^{|J|}$, any $\epsilon > 0$, and any QPT algorithms $\text{C}_1^*. \text{Open}$ and $\text{C}_2^*. \text{Open}$, for every $i \in \{1, 2\}$*

$$\text{Real}^{\text{C}^*. \text{Commit}, P^*, \text{C}_i^*. \text{Open}}(\lambda, (J, \mathbf{b}_J), \sigma) \stackrel{\zeta}{\approx} \text{Ideal}^{\text{Ext}, \text{C}^*. \text{Commit}, P^*, P_{\text{Test}}^*}(\lambda, (J, \mathbf{b}_J), \sigma, \epsilon) \quad (14)$$

and

$$\text{Real}^{\text{C}^*. \text{Commit}, P^*, \text{C}_1^*. \text{Open}}(\lambda, (J, \mathbf{b}_J), \sigma) \stackrel{\eta}{\approx} \text{Real}^{\text{C}^*. \text{Commit}, P^*, \text{C}_2^*. \text{Open}}(\lambda, (J, \mathbf{b}_J), \sigma) \quad (15)$$

where $\text{Real}^{\text{C}^*. \text{Commit}, P^*, \text{C}^*. \text{Open}}(\lambda, (J, \mathbf{b}_J), \sigma)$ is defined as follows:

1. Generate $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$.
2. Compute $(\text{rt}, \rho) \leftarrow \text{C}^*. \text{Commit}(\text{pk}, \sigma)$.²⁷
3. Compute the Ver. Commit protocol between $P^*(\text{pk}, \text{rt}, \rho)$ and $V(\text{sk}, \text{rt})$. If V rejects then output \perp . Denote the resulting quantum state of P^* at the end of this protocol by ρ_{post} .²⁸
4. Compute $(\mathbf{z}_J, \rho') \leftarrow \text{C}^*. \text{Open}(\rho_{\text{post}}, (J, \mathbf{b}_J))$.
5. If $\text{Ver}(\text{sk}, \text{rt}, (J, \mathbf{b}_J), \mathbf{z}_J) = 0$ then output \perp .
6. Otherwise, let $\mathbf{m}_J = \text{Out}(\text{sk}, \text{rt}, (J, \mathbf{b}_J), \mathbf{z}_J)$.
7. Output $(\text{pk}, \text{rt}, (J, \mathbf{b}_J), \mathbf{m}_J)$.

Let δ_0 be the probability that at the end of the protocol Ver. Commit the verifier rejects. Namely, δ_0 is the probability that Item 3 above outputs \perp . Denote by ρ_{post} the state of P^* after the Ver. Commit protocol. Let δ'_0 be the probability that the verifier $V_{\text{Test}}(\text{sk}, \text{rt})$ outputs \perp in the Test protocol when interacting with $P_{\text{Test}}^*(\text{pk}, \text{rt}, \rho_{\text{post}})$. Let

$$\delta = \max_{i \in \{1, 2\}} \Pr[\text{Ver}(\text{sk}, \text{rt}, \mathbf{b}_J, \mathbf{z}_J) = 0] \quad \text{and} \quad \mathbf{z} = \text{C}_i^*. \text{Open}(\rho_{\text{post}}, \mathbf{b}_J), \quad (16)$$

and let

$$\eta = O\left(\sqrt{\delta_0 + \delta}\right) \quad \text{and} \quad \zeta = O\left(\sqrt{\delta_0 + \delta'_0 + \delta}\right) + \epsilon. \quad (17)$$

$\text{Ideal}^{\text{Ext}, \text{C}^*. \text{Commit}, P^*, P_{\text{Test}}^*}(\lambda, (J, \mathbf{b}_J), \sigma, \epsilon)$ is defined as follows:

1. Generate $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$.
2. Compute $(\text{rt}, \rho) \leftarrow \text{C}^*. \text{Commit}(\text{pk}, \sigma)$.
3. Let $\tau_{\mathcal{A}, \mathcal{B}} = \text{Ext}^{P^*, P_{\text{Test}}^*}(\text{sk}, \text{rt}, \rho, 1^{[1/\epsilon]})$.

²⁷Note that a malicious $\text{C}^*. \text{Commit}$ may choose rt maliciously without a corresponding non-succinct commitment string \mathbf{y} . We assume without loss of generality that all the auxiliary information it has about rt is encoded in ρ .

²⁸If P^* was honest then it would have been classical and hence ρ would have remained unchanged. But since we are considering a malicious P^* it may alter its quantum state during the Ver. Commit protocol.

4. Measure the J qubits of τ_A in the basis \mathbf{b}_J to obtain $\mathbf{m}_J \in \{0, 1\}^{|J|}$.
5. Output $(\text{pk}, \text{rt}, (J, \mathbf{b}_J), \mathbf{m}_J)$.

Remark 5.11. In the succinct soundness definition above we assume that $\ell = \ell(\lambda)$ is polynomial in the security parameter. We could also consider ℓ that is super-polynomial in λ , in which case we will obtain binding assuming the post-quantum ℓ -security of the LWE assumption; i.e., assuming that a $\text{poly}(\ell)$ size quantum circuit cannot break the LWE assumption. The proof for a general (super-polynomial) ℓ is exactly the same as the one where $\ell = \text{poly}(\lambda)$, the only difference is that now we consider adversaries that run in $\text{poly}(\ell)$ time.

6 Constructions

In this section we present our constructions. We first construct a classical commitment scheme for committing to a *single* qubit state. This can be found in Section 6.1. Then, we show a generic transformation that converts any single-qubit commitment scheme into a multi-qubit commitment scheme. This can be found in Section 6.2. In this scheme the size of the public key and the size of the commitments grow with the length of the quantum state committed to. Finally, in Section 6.3 we show how to construct a succinct multi-qubit commitment scheme, where the size of the public key as well as the size of the commitment grows only with the security parameter (and poly-logarithmically with the length of the quantum state committed to). We analyze these schemes in Section 7.

6.1 Construction for Single Qubit States

In this subsection, we describe our commitment scheme for a quantum state that consists of a single qubit, denoted by $\alpha_0 |0\rangle + \alpha_1 |1\rangle$. We use as a building block the commitment algorithm Commit_W from [Mah18] for the *multi-qubit* case. This algorithm makes use of a NTCF family

$$(\text{Gen}_{\text{NTCF}}, \text{Eval}_{\text{NTCF}}, \text{Invert}_{\text{NTCF}}, \text{Check}_{\text{NTCF}}, \text{Good}_{\text{NTCF}}).$$

The public key pk used by Commit_W to commit to an ℓ -qubit state is of the form $\text{pk} = (\text{pk}_1, \dots, \text{pk}_\ell)$ where each pk_j is a public key generated by $\text{Gen}_{\text{NTCF}}(1^\lambda)$.²⁹ The QPT algorithm

$$\text{Commit}_W \left((\text{pk}_1, \dots, \text{pk}_\ell), \sum_{\mathbf{s} \in \{0,1\}^\ell} \alpha_{\mathbf{s}} |\mathbf{s}\rangle_S \right)$$

outputs the following:

1. A measurement outcome $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_\ell)$, where each $\mathbf{y}_j \in R_{\text{pk}_j}$.
2. A state $|\varphi\rangle$ such that

$$|\varphi\rangle \equiv \sum_{\mathbf{s} \in \{0,1\}^\ell} \alpha_{\mathbf{s}} |\mathbf{s}\rangle_S |\mathbf{x}_{\mathbf{s}}\rangle_Z, \quad (18)$$

where $\mathbf{x}_{\mathbf{s}} = (\mathbf{x}_{s_1}, \dots, \mathbf{x}_{s_\ell})$ where each $\mathbf{x}_{s_j} \in D_{\text{pk}_j}$ and is such that

$$\mathbf{y}_j \in \text{Supp}(\text{Eval}(\text{pk}_j, s_j, \mathbf{x}_{s_j})).$$

²⁹We mention that [Mah18] used a dual mode NTCF family, where each pk_i is generated either in an injective mode in a two-to-one mode, depending on the opening basis which is assumed to fixed ahead of time.

Construction 6.1 (Commitment Scheme). *Our construction uses a noisy trapdoor claw-free (NTCF) function family $(\text{Gen}_{\text{NTCF}}, \text{Eval}_{\text{NTCF}}, \text{Invert}_{\text{NTCF}}, \text{Check}_{\text{NTCF}}, \text{Good}_{\text{NTCF}})$ and the algorithm Commit_W defined above. Our algorithms are defined as follows:*

- $\text{Gen}(1^\lambda)$:

1. For every $i \in \{0, 1, \dots, n+1\}$ sample $(\text{pk}_i, \text{sk}_i) \leftarrow \text{Gen}_{\text{NTCF}}(1^\lambda)$, where $n = n(\lambda)$ is such that the domain of each trapdoor claw-free function is a subset of $\{0, 1\}^n$.
2. Let $\text{pk} = (\text{pk}_0, \text{pk}_1, \dots, \text{pk}_{n+1})$ and $\text{sk} = (\text{sk}_0, \text{sk}_1, \dots, \text{sk}_{n+1})$.
3. Output (pk, sk) .

- $\text{Commit}(\text{pk}, \alpha_0 |0\rangle + \alpha_1 |1\rangle)$:

1. Parse $\text{pk} = (\text{pk}_0, \text{pk}_1, \dots, \text{pk}_{n+1})$
2. Compute $(\mathbf{y}_0, |\varphi_0\rangle) \leftarrow \text{Commit}_W(\text{pk}_0, \alpha_0 |0\rangle + \alpha_1 |1\rangle)$, where

$$|\varphi_0\rangle_{\mathcal{S}, \mathcal{Z}} \equiv \sum_{s \in \{0,1\}} \alpha_s |s\rangle_{\mathcal{S}} |\mathbf{x}_s\rangle_{\mathcal{Z}}$$

Here, $\mathbf{x}_s \in \{0, 1\}^n$ and $\mathbf{y}_0 \in \text{Supp}(\text{Eval}_{\text{NTCF}}(\text{pk}_0, s, \mathbf{x}_s))$ for every $s \in \{0, 1\}$. Note that register \mathcal{S} consists of 1 qubit and \mathcal{Z} consists of n qubits.

3. Apply the Hadamard unitary $H^{\otimes(n+1)}$ to $|\varphi_0\rangle$ to obtain

$$\begin{aligned} |\varphi_1\rangle_{\mathcal{S}, \mathcal{Z}} &= H^{\otimes(n+1)} |\varphi_0\rangle \\ &\equiv \frac{1}{\sqrt{2^{n+1}}} \sum_{\mathbf{d} \in \{0,1\}^{n+1}} \underbrace{(-1)^{\mathbf{d} \cdot (0, \mathbf{x}_0)} (\alpha_0 + (-1)^{\mathbf{d} \cdot (1, \mathbf{x}_0 \oplus \mathbf{x}_1)} \alpha_1)}_{\beta_{\mathbf{d}}} |\mathbf{d}\rangle_{\mathcal{S}, \mathcal{Z}} \end{aligned}$$

4. Apply the algorithm Commit_W with pk_1 to register \mathcal{S} of $|\varphi_1\rangle_{\mathcal{S}, \mathcal{Z}}$, and for every $i \in \{2, \dots, n+1\}$ apply Commit_W with pk_i to register \mathcal{Z}_i of $|\varphi_1\rangle_{\mathcal{S}, \mathcal{Z}}$. Denote the output by $(\mathbf{y}_1, \dots, \mathbf{y}_{n+1})$ and the resulting state by

$$|\varphi_2\rangle_{\mathcal{S}, \mathcal{Z}, \mathcal{Z}'} \equiv \frac{1}{\sqrt{2^{n+1}}} \sum_{\mathbf{d} \in \{0,1\}^{n+1}} \beta_{\mathbf{d}} |\mathbf{d}\rangle_{\mathcal{S}, \mathcal{Z}} \left| \mathbf{x}'_{1,d_1}, \mathbf{x}'_{2,d_2}, \dots, \mathbf{x}'_{n+1,d_{n+1}} \right\rangle_{\mathcal{Z}'}$$

where for every $i \in \{1, \dots, n+1\}$ and every $d_i \in \{0, 1\}$,

$$\mathbf{y}_i \in \text{Supp}(\text{Eval}_{\text{NTCF}}(\text{pk}_i, d_i, \mathbf{x}'_{i,d_i})).$$

Note that the \mathcal{Z}' register consists of $n \cdot (n+1)$ qubits, and we partition these qubits to registers $\mathcal{Z}'_1, \dots, \mathcal{Z}'_{n+1}$, each consisting of n qubits.

5. Rename the register \mathcal{S} to \mathcal{Z}_1 and split the register \mathcal{Z} into registers $\mathcal{Z}_2, \dots, \mathcal{Z}_{n+1}$ of 1 qubit each. Permute the registers to obtain a state $|\varphi_3\rangle$ such that

$$|\varphi_3\rangle \equiv \frac{1}{\sqrt{2^{n+1}}} \sum_{\mathbf{d} \in \{0,1\}^{n+1}} \beta_{\mathbf{d}} |d_1\rangle_{\mathcal{Z}_1} |\mathbf{x}'_{1,d_1}\rangle_{\mathcal{Z}'_1} \cdots |d_{n+1}\rangle_{\mathcal{Z}_{n+1}} |\mathbf{x}'_{n+1,d_{n+1}}\rangle_{\mathcal{Z}'_{n+1}}$$

6. Output $(\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_{n+1})$ and $|\varphi_3\rangle$.

- $\text{Open}(|\varphi\rangle, b)$:

If $b = 1$ (corresponding to an opening in the Hadamard basis) then output the measurement of $|\varphi\rangle$ in the standard basis, and if $b = 0$ (corresponding an opening in the standard basis) then output the measurement of $|\varphi\rangle$ in the Hadamard basis.

- $\text{Ver}(\text{sk}, \mathbf{y}, b, \mathbf{z})$:

1. Parse $\text{sk} = (\text{sk}_0, \text{sk}_1, \dots, \text{sk}_{n+1})$.

2. Parse $\mathbf{y} = (\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_{n+1})$.

3. If $b = 1$ then do the following:

(a) Parse $\mathbf{z} = (d_1, \mathbf{x}'_1, \dots, d_{n+1}, \mathbf{x}'_{n+1})$ and let $\mathbf{d} = (d_1, \dots, d_{n+1})$.

(b) Compute $((0, \mathbf{x}_0), (1, \mathbf{x}_1)) = \text{Invert}_{\text{NTCF}}(\text{sk}_0, \mathbf{y}_0)$.

(c) Verify that

– $\text{Check}_{\text{NTCF}}(\text{pk}_i, d_i, \mathbf{x}'_i, \mathbf{y}_i) = 1$ for every $i \in \{1, \dots, n+1\}$.

– $\mathbf{d} \in \text{Good}_{\mathbf{x}_0, \mathbf{x}_1}$.

If any of these checks does not hold output 0 and otherwise output 1.

4. If $b = 0$ then do the following:

(a) Parse $\mathbf{z} = (\mathbf{z}_1, \dots, \mathbf{z}_{n+1})$ where each $\mathbf{z}_i \in \{0, 1\}^{n+1}$.

(b) For every $i \in \{1, \dots, n+1\}$ compute $((0, \mathbf{x}'_{i,0}), (1, \mathbf{x}'_{i,1})) = \text{Invert}_{\text{NTCF}}(\text{sk}_i, \mathbf{y}_i)$.

(c) If there exists $i \in [n+1]$ such that $\mathbf{z}_i \notin \text{Good}_{\mathbf{x}'_{i,0}, \mathbf{x}'_{i,1}}$ then output 0.

(d) Else, for every $i \in [n+1]$ let $m_i = \mathbf{z}_i \cdot (1, \mathbf{x}'_{i,0} \oplus \mathbf{x}'_{i,1})$.

(e) If $\text{Check}_{\text{NTCF}}(\text{pk}_0, m_1, (m_2, \dots, m_{n+1}), \mathbf{y}_0) \neq 1$, output 0. Else, output 1

- $\text{Out}(\text{sk}, \mathbf{y}, b, \mathbf{z})$:

1. Parse $\text{sk} = (\text{sk}_0, \text{sk}_1, \dots, \text{sk}_{n+1})$.

2. Parse $\mathbf{y} = (\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_{n+1})$.

3. If $b = 1$:

(a) Compute $((0, \mathbf{x}_0), (1, \mathbf{x}_1)) = \text{Invert}_{\text{NTCF}}(\text{sk}_0, \mathbf{y}_0)$.

(b) Parse $\mathbf{z} = (d_1, \mathbf{x}'_1, \dots, d_{n+1}, \mathbf{x}'_{n+1})$ and let $\mathbf{d} = (d_1, \dots, d_{n+1})$.

(c) Output $m = \mathbf{d} \cdot (1, \mathbf{x}_0 \oplus \mathbf{x}_1)$.

4. If $b = 0$:

(a) Parse $\mathbf{z} = (\mathbf{z}_1, \dots, \mathbf{z}_{n+1})$.

(b) Compute $((0, \mathbf{x}'_{1,0}), (1, \mathbf{x}'_{1,1})) = \text{Invert}_{\text{NTCF}}(\text{sk}_1, \mathbf{y}_1)$.

(c) Output $m_1 = \mathbf{z}_1 \cdot (1, \mathbf{x}'_{1,0} \oplus \mathbf{x}'_{1,1})$.

6.2 Construction of Commitments for Multi-Qubit States

There are two ways one can extend our single-qubit commitment scheme to the multi-qubit setting. The first is to commit to an ℓ -qubit state qubit-by-qubit by generating ℓ key pairs and using the i 'th key pair to commit and open to the i 'th qubit. This construction results with key size and commitment size that grow linearly with ℓ , and is presented below. The second approach is to extend

our single qubit commitment scheme to a *succinct* multi-qubit commitment scheme. This is done in two steps. First, we construct a *semi-succinct* multi-qubit commitment scheme, which is the same as the non-succinct one, except that we generate a *single* key pair $(\mathbf{pk}, \mathbf{sk})$ and commit to each of the ℓ qubits using the same public key \mathbf{pk} . This results with a commitment string $(\mathbf{y}_1, \dots, \mathbf{y}_\ell)$. Then we show how to convert the semi-succinct commitment scheme into a succinct one. We elaborate on this approach in Section 6.3.

Construction 6.2 (Scheme for Multi-Qubit States). *Given any single-qubit commitment scheme $(\text{Gen}_1, \text{Commit}_1, \text{Open}_1, \text{Ver}_1, \text{Out}_1)$ we construct a multi-qubit commitment scheme consisting of algorithms*

$$(\text{Gen}, \text{Commit}, \text{Open}, \text{Ver}, \text{Out})$$

defined as follows, where we define $(\text{Open}, \text{Ver}, \text{Out})$ to operate one qubit at a time (see Remark 5.2):

- $\text{Gen}(1^\lambda, 1^\ell)$:
 1. For every $i \in [\ell]$ sample $(\mathbf{pk}_i, \mathbf{sk}_i) \leftarrow \text{Gen}_1(1^\lambda)$.
 2. Let $\mathbf{pk} = (\mathbf{pk}_1, \dots, \mathbf{pk}_\ell)$ and $\mathbf{sk} = (\mathbf{sk}_1, \dots, \mathbf{sk}_\ell)$.
 3. Output $(\mathbf{pk}, \mathbf{sk})$.
- $\text{Commit}(\mathbf{pk}, \sigma)$:
 1. Parse $\mathbf{pk} = (\mathbf{pk}_1, \dots, \mathbf{pk}_\ell)$.
 2. We assume that σ is an ℓ -qubit state, and we denote the ℓ registers of σ by $\mathcal{S}_1, \dots, \mathcal{S}_\ell$.
 3. Execute the following steps:
 - (a) Let $\rho_0 = \sigma$.
 - (b) For every $j \in \{1, \dots, \ell\}$, apply Commit_1 with key \mathbf{pk}_j to register \mathcal{S}_j of the state ρ_{j-1} , obtaining an outcome \mathbf{y}_j and a post-measurement state $(\rho_j)_{\mathcal{S}_1, \dots, \mathcal{S}_\ell, \mathcal{Z}_1, \dots, \mathcal{Z}_j}$.
 4. Output $(\mathbf{y}, (\rho_\ell)_{\mathcal{S}_1, \dots, \mathcal{S}_\ell, \mathcal{Z}_1, \dots, \mathcal{Z}_\ell})$, where $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_\ell)$.
- $\text{Open}(\rho_{\mathcal{S}_1, \dots, \mathcal{S}_\ell, \mathcal{Z}_1, \dots, \mathcal{Z}_\ell}, (j, b_j))$:
 1. Apply Open_1 with basis b_j to registers $\{\mathcal{S}_j, \mathcal{Z}_j\}$ of $\rho_{\mathcal{S}_1, \dots, \mathcal{S}_\ell, \mathcal{Z}_1, \dots, \mathcal{Z}_\ell}$, obtaining an outcome \mathbf{z}_j and post-measurement state ρ'_j .
 2. Output (\mathbf{z}_j, ρ'_j) .
- $\text{Ver}(\mathbf{sk}, \mathbf{y}, (j, b_j), \mathbf{z}_j)$:
 1. Parse $\mathbf{sk} = (\mathbf{sk}_1, \dots, \mathbf{sk}_\ell)$ and $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_\ell)$.
 2. Output $\text{Ver}_1(\mathbf{sk}_j, \mathbf{y}_j, b_j, \mathbf{z}_j)$.
- $\text{Out}(\mathbf{sk}, \mathbf{y}, (j, b_j), \mathbf{z}_j)$:
 1. Parse $\mathbf{sk} = (\mathbf{sk}_1, \dots, \mathbf{sk}_\ell)$, $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_\ell)$.
 2. Output $m_j \leftarrow \text{Out}_1(\mathbf{sk}_j, \mathbf{y}_j, b_j, \mathbf{z}_j)$.

We consider a *semi-succinct* version of the commitment scheme described above, which is used as a stepping stone for proving soundness of the fully succinct commitment scheme constructed in Section 6.3, below.

Definition 6.3. *A semi-succinct classical commitment scheme to quantum states is a commitment scheme obtained from a single qubit commitment scheme by applying the algorithms qubit-by-qubit, as in Construction 6.2, but with a single public key pk . Namely, it is similar to Construction 6.2 but where Gen generates a single key pair $(\text{pk}, \text{sk}) \leftarrow \text{Gen}_1(1^\lambda)$ (as opposed to ℓ such pairs), and sets $(\text{pk}_i, \text{sk}_i) = (\text{pk}, \text{sk})$ for every $i \in [\ell]$.*

Remark 6.4. *In Section 7.2, when we prove that our (non-succinct and semi-succinct) multi-qubit commitment scheme satisfies the binding property, we assume that $\text{C}^*. \text{Open}$ successfully opens all the ℓ qubits in the standard basis with high probability, and successfully opens all the ℓ qubits in the Hadamard basis with high probability $1 - \delta$. Namely, we assume that for every $b \in \{0, 1\}$, $\text{C}^*. \text{Open}(\rho, b^\ell)$ generates an accepting opening with probability $1 - \delta$. To ensure that this assumption holds, we repeat the commitment phase $O(\lceil 1/\delta \rceil)$ times.³⁰ For each of these commitments, we ask $\text{C}^*. \text{Open}$ with probability $1/3$ to open all the qubits in the standard basis, with probability $1/3$ to open all the qubits in the Hadamard basis, and with probability $1/3$ to open in desired basis $\mathbf{b} \in \{0, 1\}^\ell$.*

We emphasize that even if $\text{C}^. \text{Open}$ opens only a small subset of the qubits, we still require that $\text{C}^*. \text{Open}(\rho, b^\ell)$ generates a valid opening for every $b \in \{0, 1\}$. The reason is that our binding property states that there exists an extractor \mathcal{E} that uses any $\text{C}^*. \text{Open}$ to extract a state τ . We want the guarantee that even if $\text{C}^*. \text{Open}$ only opens a small subset, still the extractor can extract an ℓ -qubit state τ . This is important in some applications, such as compiling any quantum X/Z PCP into a succinct interactive argument (see Section 8.2). Importantly, we need to do this without increasing the communication complexity in the opening phase, and in particular it should not grow with ℓ . In what follows we show how this can be done succinctly.*

6.3 Construction of Succinct Multi-Qubit Commitments

Before we present our construction of a succinct multi-qubit commitment scheme, we define one of the main building blocks used in our construction.

State-Preserving Succinct Arguments of Knowledge Our scheme uses a state-preserving succinct argument of knowledge system, defined and constructed in [LMS22].

Definition 6.5. [LMS22] *A publicly verifiable argument system Π for an NP language \mathcal{L} (with witness relation \mathcal{R}) is an ϵ -state-preserving succinct argument-of-knowledge if it satisfies the following properties.*

- **Succinctness:** *when invoked on a security parameter λ , instance size n , and a relation \mathcal{R} decidable in time T , the communication complexity of the protocol is $\text{poly}(\lambda, \log T)$. The verifier computational complexity is $\text{poly}(\lambda, \log T) + \tilde{O}(n)$.*
- **State-Preserving Extraction.** *There exists an extractor \mathcal{E} , with oracle access to a cheating prover P^* and a corresponding quantum state ρ , with the following properties:*

³⁰We assume that C.Commit can generate $\sigma^{\otimes \lceil 1/\delta \rceil}$.

- **Efficiency:** \mathcal{E} on input $(\mathbf{x}, 1^\lambda, \epsilon)$ runs in time $\text{poly}(|\mathbf{x}|, \lambda, 1/\epsilon)$, and outputs a classical transcript \mathbb{T}_{Sim} , a classical string \mathbf{w} , and a residual state ρ_{Sim} .
- **State-preserving:** The following two games are ϵ -indistinguishable to any QPT distinguisher:
 - * **Game 0 (Real):** Generate a transcript \mathbb{T} by running $P^*(\rho, \mathbf{x})$ with the honest verifier V . Output \mathbb{T} along with the residual state ρ' .
 - * **Game 1 (Simulated):** Generate $((\mathbb{T}_{\text{Sim}}, \mathbf{w}), \rho_{\text{Sim}}) \leftarrow \mathcal{E}^{P^* \cdot \rho}(\mathbf{x}, 1^\lambda, \epsilon)$. Output $(\mathbb{T}_{\text{Sim}}, \rho_{\text{Sim}})$.
- **Extraction correctness:** for any P^* as above, the probability that \mathbb{T}_{Sim} is an accepting transcript but \mathbf{w} is not in $\mathcal{R}_{\mathbf{x}}$ is at most $\epsilon + \text{negl}(\lambda)$.

The following is an immediate corollary.

Corollary 6.6. *An ϵ -state-preserving succinct argument-of-knowledge protocol for an NP language \mathcal{L} (with witness relation \mathcal{R}) and extractor \mathcal{E} satisfies that for every cheating prover P^* and a corresponding quantum state ρ , and every \mathbf{x} , if $P^*(\rho, \mathbf{x})$ convinces the verifier V to accept with probability $1 - \delta$ then for*

$$((\mathbb{T}_{\text{Sim}}, \mathbf{w}), \rho_{\text{Sim}}) \leftarrow \mathcal{E}^{P^* \cdot \rho}(\mathbf{x}, 1^\lambda, \epsilon),$$

$$\Pr[(\mathbf{x}, \mathbf{w}) \in \mathcal{R}] \geq 1 - \delta - 2\epsilon - \text{negl}(\lambda).$$

Theorem 6.7 ([LMS22]). *Assuming the post-quantum $\text{poly}(\lambda, 1/\epsilon)$ hardness of learning with errors, there exists a (4-message, public coin) ϵ -state preserving succinct argument of knowledge for NP.*

6.3.1 Construction

We are now ready to present our construction of a succinct classical commitment for multi-qubit quantum states. Our construction uses the following ingredients:

- Collapsing hash family with local opening $(\text{Gen}_{\text{H}}, \text{Eval}_{\text{H}}, \text{Open}_{\text{H}}, \text{Ver}_{\text{H}})$, as defined in Definitions 3.4 and 3.6.
- A semi-succinct commitment scheme $(\text{Gen}_{\text{SS}}, \text{Commit}_{\text{SS}}, \text{Open}_{\text{SS}}, \text{Ver}_{\text{SS}}, \text{Out}_{\text{SS}})$, as defined in Definition 6.3, corresponding to an underlying single-qubit commitment scheme

$$(\text{Gen}_{\text{I}}, \text{Commit}_{\text{I}}, \text{Open}_{\text{I}}, \text{Ver}_{\text{I}}, \text{Out}_{\text{I}})$$

- An ϵ -state-preserving succinct argument of knowledge protocol (P, V) , as defined in Definition 6.5, for the NP languages \mathcal{L}^* and \mathcal{L}^{**} with a corresponding NP relations $\mathcal{R}_{\mathcal{L}^*}$ and $\mathcal{R}_{\mathcal{L}^{**}}$, respectively, defined as follows:

$$((\text{hk}, \text{rt}), \mathbf{y}) \in \mathcal{R}_{\mathcal{L}^*} \text{ if and only if } \text{Eval}_{\text{H}}(\text{hk}, \mathbf{y}) = \text{rt} \tag{19}$$

and

$$((\text{sk}_1, \text{hk}, \text{rt}, \text{rt}', b), (\mathbf{y}, \mathbf{z})) \in \mathcal{R}_{\mathcal{L}^{**}}$$

if and only if

$$\text{Eval}_{\text{H}}(\text{hk}, \mathbf{y}) = \text{rt} \wedge \text{Eval}_{\text{H}}(\text{hk}, \mathbf{z}) = \text{rt}' \wedge \text{Ver}_{\text{SS}}(\text{sk}_1, \mathbf{y}, b^\ell, \mathbf{z}) = 1.$$

Construction 6.8 (Succinct Commitment to Multi-Qubit Quantum States). *In what follows we use the ingredients above to construct a succinct commitment scheme to multi-qubit quantum states.*

- $\text{Gen}(1^\lambda)$:
 1. Sample $(\text{pk}_1, \text{sk}_1) \leftarrow \text{Gen}_1(1^\lambda)$.
 2. Sample $\text{hk} \leftarrow \text{Gen}_H(1^\lambda)$.
 3. Let $\text{pk} = (\text{pk}_1, \text{hk})$ and $\text{sk} = (\text{sk}_1, \text{hk})$.
 4. Output (pk, sk) .
- $\text{Commit}(\text{pk}, \sigma)$:
 1. Parse $\text{pk} = (\text{pk}_1, \text{hk})$.
 2. Compute $(\mathbf{y}, \rho) \leftarrow \text{Commit}_{\text{ss}}(\text{pk}_1, \sigma)$.
 3. Let $\text{rt} = \text{Eval}_H(\text{hk}, \mathbf{y})$.
 4. Output $(\text{rt}, \mathbf{y}, \rho)$.
- Ver.Commit runs the ϵ -state-preserving succinct argument of knowledge protocol for the NP language \mathcal{L}^* , where P and V take as input the instance (hk, rt) and P takes an additional input the witness \mathbf{y} .³¹ If V rejects then this commitment string rt is declared invalid, and the protocol aborts.
- Test is an interactive protocol between P_{Test} with input $(\text{pk}, \text{rt}, \mathbf{y}, \rho)$ and V_{Test} with input (sk, rt) that proceeds as follows:
 1. V_{Test} samples a uniformly random bit $b \leftarrow \{0, 1\}$, and sends b to P_{Test} .
 2. P_{Test} generates $\mathbf{z} \leftarrow \text{Open}_{\text{ss}}(\rho, b^\ell)$ and sends $\text{rt}' = \text{Eval}_H(\text{hk}, \mathbf{z})$ to V_{Test} .
 3. Run the ϵ -state-preserving succinct argument of knowledge protocol for the NP language \mathcal{L}^* , where P and V take as input the instance (hk, rt') and P takes the additional input \mathbf{z} . If V rejects then this commitment is declared invalid and the protocol aborts.
 4. V_{Test} sends sk_1 to the prover.
 5. Run the ϵ -preserving succinct argument of knowledge protocol (P, V) for the NP language \mathcal{L}^{**} , where P and V take as input the instance $(\text{sk}_1, \text{hk}, \text{rt}, \text{rt}', b)$ and P takes as additional input the witness (\mathbf{y}, \mathbf{z}) . If V rejects then this commitment is declared invalid.
 6. If V accepts in both steps 2 and 4 above then V_{Test} outputs 1, and otherwise it outputs 0.

Remark 6.9. *Note that the Test protocols runs the ϵ -state-preserving succinct argument of knowledge twice: once to prove knowledge of \mathbf{z} and once to prove knowledge of (\mathbf{y}, \mathbf{z}) . It may seem that it suffices to run this protocol once, since there is no need to prove knowledge of \mathbf{z} twice. However, in the first protocol the cheating prover does not know sk_1 and hence the \mathbf{z} that is extracted, using the extractor from Definition 6.5, can be efficiently computed without knowing sk_1 . Then we use the security guarantee of the underlying collision resistant hash family to argue that the vector \mathbf{z} extracted from the second protocol is identical to that extracted from the first protocol, and hence can be computed efficiently without knowing sk_1 . This fact is crucial for the soundness proof to go through.*

³¹Note that in this protocol both the prover and the verifier are classical.

- $\text{Open}((\boldsymbol{\rho}, \text{hk}, \mathbf{y}), (j, b_j))$:
 1. Parse $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_\ell)$.
 2. Compute $\mathbf{o}_j = \text{Open}_H(\text{hk}, \mathbf{y}, j)$.³²
 3. Compute $(\mathbf{z}_j, \boldsymbol{\rho}') \leftarrow \text{Open}_{\text{ss}}(\boldsymbol{\rho}, (j, b_j))$.
 4. Output $((\mathbf{y}_j, \mathbf{o}_j, \mathbf{z}_j), \boldsymbol{\rho}')$.
- $\text{Ver}(\text{sk}, \text{rt}, (j, b_j), (\mathbf{y}_j, \mathbf{o}_j, \mathbf{z}_j))$:
 1. Parse $\text{sk} = (\text{sk}_1, \text{hk})$.
 2. Let $v_0 = \text{Ver}_H(\text{hk}, \text{rt}, j, \mathbf{y}_j, \mathbf{o}_j)$.³³
 3. Let $v_1 = \text{Ver}_1(\text{sk}_1, \mathbf{y}_j, b_j, \mathbf{z}_j)$.
 4. Output $v_0 \wedge v_1$.
- $\text{Out}(\text{sk}, \text{rt}, (j, b_j), (\mathbf{y}_j, \mathbf{o}_j, \mathbf{z}_j))$:
 1. Parse $\text{sk} = (\text{sk}_1, \text{hk})$.
 2. Output $m_j = \text{Out}_1(\text{sk}_1, \mathbf{y}_j, b_j, \mathbf{z}_j)$.

7 Analysis of the Multi-Qubit Commitment Schemes from Section 6

7.1 Correctness

In this section we prove the correctness of Construction 6.2 and Construction 6.8.

Theorem 7.1. *The multi-qubit commitment scheme described in Construction 6.2 satisfies the correctness property given in Definition 5.4.*

Section 6.2 commits to each qubit of a multi-qubit state independently by using the single-qubit protocol given in construction Construction 6.1 as a black-box. Therefore, to prove Theorem 7.1 it suffices to prove the following theorem.

Theorem 7.2. *The single-qubit commitment scheme described in Construction 6.1 satisfies the correctness property given in Definition 5.4.*

We make use of the following lemma about Commit_W throughout the proof.

Lemma 7.3 (Correctness of Commit_W). *For any ℓ -qubit quantum state $|\varphi\rangle = \sum_{\mathbf{s} \in \{0,1\}^\ell} \alpha_{\mathbf{s}} |\mathbf{s}\rangle_{\mathcal{S}}$ and any basis $\mathbf{b} = (b_1, \dots, b_\ell) \in \{0\}^\ell \cup \{1\}^\ell$,*

$$\text{Real}_W(1^\lambda, |\varphi\rangle, \mathbf{b}) \stackrel{\text{negl}(\lambda)}{\equiv} \boldsymbol{\sigma}(\mathbf{b}) \quad (20)$$

where $\boldsymbol{\sigma}(\mathbf{b})$ is the distribution obtained by measuring each qubit j of $|\varphi\rangle$ in the basis specified by b_j (standard if $b_j = 0$, Hadamard if $b_j = 1$), and $\text{Real}_W(1^\lambda, |\varphi\rangle, \mathbf{b})$ is the distribution resulting from honestly opening the commitment. Specifically, $\text{Real}_W(1^\lambda, |\varphi\rangle, \mathbf{b})$ is defined by:

³² $\text{Open}_H(\text{hk}, \mathbf{y}, j)$ denotes an opening to the j 'th chunk of the preimage \mathbf{y} , consisting of \mathbf{y}_j which is the commitment to the j 'th qubit.

³³ $\text{Ver}_H(\text{hk}, \text{rt}, j, \mathbf{y}_j, \mathbf{o}_j)$ denotes the verification of the opening for \mathbf{y}_j , which is the j 'th chunk of the hashed preimage.

1. For every $i \in \{0, 1, \dots, \ell\}$, sample $(\mathbf{pk}_i, \mathbf{sk}_i) \leftarrow \text{Gen}_{\text{NTCF}}(1^\lambda)$.
2. Compute $(\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_\ell), |\varphi'\rangle) \leftarrow \text{Commit}_W((\mathbf{pk}_1, \dots, \mathbf{pk}_\ell), |\varphi\rangle)$, where $|\varphi'\rangle$ is of the same form as Equation (18).
3. If $\mathbf{b} = \{0\}^\ell$:
 - (a) Measure $|\varphi'\rangle$ in the standard basis to get $\mathbf{z} = (\mathbf{z}_1, \dots, \mathbf{z}_\ell)$. Parse each $\mathbf{z}_i = (s_i, \mathbf{x}_{i,s_i})$. If $\text{Check}_{\text{NTCF}}(\mathbf{pk}_i, s_i, \mathbf{x}_{i,s_i}, \mathbf{y}_i) \neq 1$ for some $i \in [\ell]$, output \perp . Otherwise, output $\mathbf{s} = (s_1, s_2, \dots, s_\ell)$.
4. If $\mathbf{b} = \{1\}^\ell$:
 - (a) Measure $|\varphi'\rangle$ in the Hadamard basis to get $\mathbf{d} = (\mathbf{d}_1, \dots, \mathbf{d}_\ell) \in \{0, 1\}^{\ell(n+1)}$. For each $j \in [\ell]$, compute $((0, \mathbf{x}_{j,0}), (1, \mathbf{x}_{j,1})) = \text{Invert}_{\text{NTCF}}(\mathbf{sk}_j, \mathbf{y}_j)$. If $\mathbf{d}_j \notin \text{Good}_{\mathbf{x}_{j,0}, \mathbf{x}_{j,1}}$ for some $j \in [\ell]$, output \perp . Otherwise, output $\mathbf{m} = (m_1, \dots, m_\ell)$, where each $m_j = \mathbf{d}_j \cdot (1, \mathbf{x}_{j,0} \oplus \mathbf{x}_{j,1})$.

Proof. This follows directly from the proof of Lemma 5.3 in [Mah18], where the $\mathbf{b} = \{0\}^\ell$ case corresponds to the *Test* round and the $\mathbf{b} = \{1\}^\ell$ case corresponds to the *Hadamard* round. \square

Recall that in Construction 6.1, the final state $|\phi_3\rangle$ is the result of applying Commit_W to the state

$$|\phi_2\rangle = \frac{1}{\sqrt{2^{n+1}}} \sum_{\mathbf{d} \in \{0,1\}^{n+1}} \beta_{\mathbf{d}} |\mathbf{d}\rangle.$$

in the commitment procedure (pre-measurement)

We now show that the outcome of opening $|\varphi_3\rangle$ in a basis $b \in \{0, 1\}$ is statistically indistinguishable from measuring the initial state, $|\psi\rangle$, in the basis b . We proceed with the proof for pure states, which extends to the case of mixed and entangled states by linearity, and show correctness for each basis separately. We treat the correctness of Commit_W as a black-box. Namely, we make use of Lemma 7.3 throughout the proof.

Lemma 7.4 (Opening in the Hadamard basis, $b = 1$). *For any pure single-qubit quantum state $|\psi\rangle = \alpha_0 |0\rangle + \alpha_1 |1\rangle$ and any NTCF family, the distribution over the outcomes of the following two experiments are statistically indistinguishable under Construction 6.1:*

- **Experiment 1.** Measure $|\psi\rangle$ in the Hadamard basis and report the outcome.
- **Experiment 2.** Execute $\text{Real}(1^\lambda, |\psi\rangle, b_1 = 1)$, as described in Definition 5.4.

Proof. By inspection, it can be seen that the distribution of outcomes obtained from $\text{Real}(1^\lambda, |\psi\rangle, b_1 = 1)$ here is the same as the outcome obtained from the following procedure:

1. Generate keys $(\mathbf{sk}_0, \mathbf{pk}_0)$.
2. Apply the weak commitment once to get $\mathbf{y}_0, |\phi_1\rangle \leftarrow \text{Commit}_W(\mathbf{pk}_0, |\psi\rangle)$.
3. Apply a Hadamard transform to the state to get $|\phi_2\rangle = H^{\otimes(n+1)} |\phi_1\rangle$.
4. Execute $\text{Real}_W(1^\lambda, |\phi_2\rangle, 0^{n+1})$ to obtain an outcome $\mathbf{d} = (d_0, \dots, d_{n+1})$.

5. Report an outcome $\mathbf{d} \cdot (1, \mathbf{x}_0 \oplus \mathbf{x}_1)$, where $\{(b, \mathbf{x}_b)\}_{b=0,1} = \text{Invert}_{\text{NTCF}}(\text{sk}_0, \mathbf{y}_0)$.

By Lemma 7.3 for standard basis openings, the distribution over \mathbf{d} is statistically close to the distribution obtained by measuring $|\phi_2\rangle$ in the standard basis. This, in turn, by construction is equal to the distribution obtained by measuring $|\phi_1\rangle$ in the Hadamard basis. Finally, by applying Lemma 7.3 again, this time for Hadamard basis openings, this implies that the distribution of $\mathbf{d} \cdot (1, \mathbf{x}_0 \oplus \mathbf{x}_1)$ is statistically close to the distribution obtained by measuring $|\psi\rangle$ in the Hadamard basis. \square

Lemma 7.5 (Opening in the standard basis, $b = 0$). *For any pure single-qubit quantum state $|\psi\rangle = \alpha_0 |0\rangle + \alpha_1 |1\rangle$ and any NTCF family, the distribution over the outcomes of the following two experiments are statistically indistinguishable under Construction 6.1:*

- **Experiment 1.** Measure $|\psi\rangle$ in the standard basis and report the outcome.
- **Experiment 2.** Execute $\text{Real}(1^\lambda, |\psi\rangle, b_1 = 0)$, as described in Definition 5.4.

Proof. By inspection, it can be seen that the distribution of outcomes obtained from $\text{Real}(1^\lambda, |\psi\rangle, b_1 = 1)$ here is the same as the outcome obtained from the following procedure:

1. Generate keys $(\text{sk}_0, \text{pk}_0)$.
2. Apply the weak commitment once to get $\mathbf{y}_0, |\phi_1\rangle \leftarrow \text{Commit}_W(\text{pk}_0, |\psi\rangle)$.
3. Apply a Hadamard transform to the state to get $|\phi_2\rangle = H^{\otimes(n+1)} |\phi_1\rangle$.
4. Execute $\text{Real}_W(1^\lambda, |\phi_2\rangle, 1^{n+1})$ to obtain an outcome $\mathbf{m} = (m_0, \dots, m_n)$.
5. If $\text{Check}_{\text{NTCF}}(\text{pk}_0, m_0, (m_1, \dots, m_n), \mathbf{y}_0) = 1$, output m_0 ; else, output \perp .

By Lemma 7.3, applied in the Hadamard basis case, the outcome \mathbf{m} has a distribution that is statistically close to the outcome of a Hadamard basis measurement of $|\phi_2\rangle$. By construction, this is equal to the distribution of the outcome of a standard basis measurement of $|\phi_1\rangle$. Finally, by Lemma 7.3, applied in the standard basis case, the distribution of a standard outcome of $|\phi_1\rangle$ will pass the check $\text{Check}_{\text{NTCF}}(\text{pk}_0, m_0, (m_1, \dots, m_n), \mathbf{y}_0)$ with probability negligibly close to 1, and the bit m_0 will be distributed close to the distribution obtained by measuring $|\psi\rangle$ in the standard basis. \square

Proof of Theorem 7.2. The theorem follows immediately from Lemma 7.4 and Lemma 7.5. \square

We now proceed with the proof of correctness for the succinct commitment scheme.

Theorem 7.6. *The succinct multi-qubit commitment scheme described in Construction 6.8 satisfies the correctness property given in Definition 5.5.*

Proof. By correctness of the state-preserving succinct argument of knowledge protocol [LMS22], Ver.Commit in Construction 6.8 accepts with probability $1 - \text{negl}(\lambda)$. The Commit algorithm in that construction uses $\text{Commit}_{\text{ss}}$ as a black box, which consists of applying the Commit_1 procedure to each qubit of σ under the same public key pk_1 . Therefore, for the $c = 0$ case in Definition 5.5, correctness holds by Theorem 7.2. The $c = 1$ case follows from the correctness of the state-preserving argument of knowledge [LMS22]. \square

7.2 Binding

In this section we prove the following two theorems.

Theorem 7.7. *The non-succinct commitment scheme described in Section 6.2 satisfies the binding property given in Definition 5.6 (assuming the existence of a NTCF family).*

Theorem 7.8. *The semi-succinct multi-qubit commitment scheme described in Section 6.2 (Definition 6.3) satisfies the binding property given in Definition 5.6 if the underlying single qubit commitment scheme is the one from Section 6.1 and the underlying NTCF family is the one from [BCM⁺18] and assuming it satisfies the distributional strong adaptive hardcore bit property (see Definition 4.1).³⁴*

To prove the above two theorems we need to prove that both the non-succinct and the semi-succinct commitment schemes described in Section 6.2 satisfy Equation (11) and Equation (12) of the binding property (Definition 5.6).

Remark 7.9. *We note that Equation (12) only relies on the fact that the underlying NTCF family is collapsing (as defined in Claim 3.12), whereas Equation (11) relies on the adaptive hardcore bit property for the non-succinct scheme and on the specific properties of the NTCF family from [BCM⁺18] (specifically, the distributional strong adaptive hardcore bit property) for the semi-succinct scheme.*

We start by proving that both the non-succinct and the semi-succinct commitment schemes satisfy Equation (12). We actually prove a stronger version of Equation (12), stated below.

Lemma 7.10. *[Stronger version of Equation (12)] For any QPT algorithm $C^*.Commit$ and quantum state σ , any purification $|\varphi\rangle$ of σ , any QPT algorithms $C_1^*.Open$ and $C_2^*.Open$, any $\mathbf{b} \in \{0, 1\}^\ell$, and any efficient unitaries V_1 and V_2 there exists a negligible function $\mu = \mu(\lambda)$ such that*

$$\mathbb{E}_{\substack{(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Gen}(1^\lambda) \\ (\mathbf{y}, |\psi\rangle) \leftarrow C^*.Commit(\mathbf{pk}, |\varphi\rangle)}}} \|U_1^\dagger U_{\text{Out}}^\dagger \text{CNOT}_{\text{copy, out}} U_{\text{Out}} U_1 V_{\text{ext}, 1} |\psi_{\text{ext}}\rangle - U_2^\dagger U_{\text{Out}}^\dagger \text{CNOT}_{\text{copy, out}} U_{\text{Out}} U_2 V_{\text{ext}, 2} |\psi_{\text{ext}}\rangle\|_2 \leq \eta + \epsilon + \mu$$

where

- $\epsilon = \mathbb{E}_{\substack{(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Gen}(1^\lambda) \\ (\mathbf{y}, |\psi\rangle) \leftarrow C^*.Commit(\mathbf{pk}, |\varphi\rangle)}}} \|V_1 |\psi\rangle - V_2 |\psi\rangle\|_2.$
- $\eta = \sum_{j=1}^\ell 2\sqrt{\delta_j}$ for

$$\delta_j \triangleq \mathbb{E}_{\substack{(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Gen}(1^\lambda) \\ (\mathbf{y}, |\psi\rangle) \leftarrow C^*.Commit(\mathbf{pk}, |\varphi\rangle) \\ (\mathbf{z}_i, \rho'_i) \leftarrow C_i^*.Open(\rho, \mathbf{b})}} \max_{i \in \{1, 2\}} \Pr[\text{Ver}(\mathbf{sk}, \mathbf{y}, (j, b_j), \mathbf{z}_{i,j}) = 0 \mid \text{Ver}(\mathbf{sk}, \mathbf{y}, (k, b_k), \mathbf{z}_{i,k}) = 1 \ \forall k \in [j-1]]$$

where $\mathbf{z}_i = (\mathbf{z}_{i,j})_{j=1}^\ell$.

³⁴We recall that [BCM⁺18] satisfies the distributional strong adaptive hardcore bit property under LWE (see Claim 4.4).

- $|\psi_{\text{ext}}\rangle = |0^\ell\rangle_{\text{copy}} \otimes |0^\ell\rangle_{\text{out}} \otimes |\mathbf{b}\rangle_{\text{basis}} \otimes |\psi\rangle$.
- For every $i \in \{1, 2\}$, $V_{\text{ext},i} = I_{\text{copy}} \otimes I_{\text{out}} \otimes I_{\text{basis}} \otimes V_i$.
- For every $i \in \{1, 2\}$, U_i is the unitary defined by applying $C_i^*.\text{Open}$ to the registers open and basis.
- U_{Out} is the unitary defined by first applying the unitary corresponding to $\text{Ver}(\text{sk}, \mathbf{y}, \cdot, \cdot)$ to registers open and basis, and controlled on Ver accepting, applying the unitary corresponding to $\text{Out}(\text{sk}, \mathbf{y}, \cdot, \cdot)$ to registers open and basis, and writing the output on the register out.
- $\text{CNOT}_{\text{copy,out}}$ applies a CNOT to registers copy and out (i.e., it copies register out to register copy).

Moreover, $C_1^*.\text{Open}$ and $C_2^*.\text{Open}$ can be QPT given $\text{sk}_1, \dots, \text{sk}_{n+1}$ when opening in the standard basis and QPT when opening in the Hadamard basis.³⁵ Alternatively, they can be QPT given sk_0 when opening in the Hadamard basis and QPT when opening in the standard basis.

Corollary 7.11. For any QPT algorithm $C^*.\text{Commit}$ and quantum state σ , any QPT algorithms $C_1^*.\text{Open}$ and $C_2^*.\text{Open}$, and any $\mathbf{b} \in \{0, 1\}^\ell$,

$$\text{Real}^{C_1^*.\text{Open}}(\lambda, \mathbf{b}, \sigma) \stackrel{2(\sqrt{\delta_0} + \sqrt{\delta_1})}{\approx} \text{Real}^{C_2^*.\text{Open}}(\lambda, \mathbf{b}, \sigma)$$

where denoting by $I_b = \{i \in [\ell] : \mathbf{b}_i = b\}$,

$$\delta_b = \mathbb{E}_{\substack{(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda) \\ (\mathbf{y}, \rho) \leftarrow C^*.\text{Commit}(\text{pk}, \sigma) \\ \mathbf{z}_i \leftarrow C_i^*.\text{Open}(\rho, \mathbf{b})}} \max_{i \in \{1, 2\}} \Pr \left[\text{Ver} \left(\text{sk}, \mathbf{y}, (I_b, b^{|I_b|}), \mathbf{z}_{i, I_b} \right) = 0 \right]$$

where $\mathbf{z}_{i, I_b} = (\mathbf{z}_{i, j})_{j \in I_b}$.

Moreover, $C_1^*.\text{Open}$ and $C_2^*.\text{Open}$ can be QPT given $\text{sk}_1, \dots, \text{sk}_{n+1}$ when opening in the standard basis and QPT when opening in the Hadamard basis.

Proof of Corollary 7.11 Fix any QPT algorithm $C^*.\text{Commit}$ and quantum state σ , any algorithms $C_1^*.\text{Open}$ and $C_2^*.\text{Open}$ as in the statement of Corollary 7.11, and any basis \mathbf{b} . For every $i \in \{1, 2\}$ we slightly change $C_i^*.\text{Open}$ to $C_i^{**}.\text{Open}$, as follows: $C_i^{**}.\text{Open}(\rho, (j, b))$ coherently computes $\mathbf{z} \leftarrow C_i^*.\text{Open}(\rho, \mathbf{b})$ and outputs \mathbf{z}_j if $\text{Ver}(\text{sk}, \mathbf{y}, (I_b, \mathbf{b}_{I_b}), \mathbf{z}) = 1$, and otherwise it outputs \perp .³⁶ Note that $C_i^{**}.\text{Open}$ remains a QPT algorithm when opening in the Hadamard basis since Ver does not use sk when verifying a Hadamard basis opening, whereas it uses $\text{sk}_1, \dots, \text{sk}_{n+1}$ when opening in the standard basis. Thus $C_1^{**}.\text{Open}$ and $C_2^{**}.\text{Open}$ satisfy the efficiency conditions of Lemma 7.10. In addition, note that for every $i \in \{1, 2\}$,

$$\text{Real}^{C_i^*.\text{Open}}(\lambda, \mathbf{b}, \sigma) \equiv \text{Real}^{C_i^{**}.\text{Open}}(\lambda, \mathbf{b}, \sigma).$$

³⁵This generalization is needed to obtain Corollary 7.11.

³⁶Note that $\mathbf{b}_{I_b} = b^{|I_b|}$

where $(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Gen}(1^\lambda)$ and $(\mathbf{y}, \rho) \leftarrow \mathbf{C}^*. \text{Commit}(\mathbf{pk}, \sigma)$. By Lemma 7.10 for any purification $|\varphi\rangle$ of σ there exists a negligible function μ such that

$$\begin{aligned} & \mathbb{E}_{\substack{(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Gen}(1^\lambda) \\ (\mathbf{y}, |\psi\rangle) \leftarrow \mathbf{C}^*. \text{Commit}(\mathbf{pk}, |\varphi\rangle)}} \|U_1^\dagger U_{\text{Out}}^\dagger \text{CNOT}_{\text{copy, out}} U_{\text{Out}} U_1 |\psi_{\text{ext}}\rangle - U_2^\dagger U_{\text{Out}}^\dagger \text{CNOT}_{\text{copy, out}} U_{\text{Out}} U_2 |\psi_{\text{ext}}\rangle\|_2 \\ & \leq \eta + \mu \end{aligned}$$

where U_i is the unitary defined by $\mathbf{C}_i^{**}. \text{Open}$, and U_{Out} and η are as defined in Lemma 7.10. It remains to observe that $\eta \leq 2\sqrt{\delta_0} + 2\sqrt{\delta_1}$, which follows from the definition of $\mathbf{C}_i^{**}. \text{Open}$, which asserts that $\delta_j = 0$ if there exists $k \in \{1, \dots, j-1\}$ for which $b_j = b_k$. \square

Proof of Lemma 7.10 We prove this lemma for the semi-succinct variant of the multi-qubit commitment scheme described in Section 6.2. The proof for the non-succinct variant is identical. The proof is by induction on ℓ .

Base case: $\ell = 1$. Fix any QPT algorithm $\mathbf{C}^*. \text{Commit}$, a quantum state σ , algorithms $\mathbf{C}_1^*. \text{Open}$ and $\mathbf{C}_2^*. \text{Open}$, basis $b \in \{0, 1\}$, and efficient unitaries V_1 and V_2 , as in the lemma statement. Also fix a purification $|\varphi\rangle$ of σ . Suppose for the sake of contradiction that there exists a non-negligible $\xi = \xi(\lambda)$ such that

$$\begin{aligned} & \mathbb{E}_{\substack{(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Gen}(1^\lambda) \\ (\mathbf{y}, |\psi\rangle) \leftarrow \mathbf{C}^*. \text{Commit}(\mathbf{pk}, |\varphi\rangle)}} \|U_1^\dagger U_{\text{Out}}^\dagger \text{CNOT}_{\text{copy, out}} U_{\text{Out}} U_1 V_{\text{ext}, 1} |\psi_{\text{ext}}\rangle - U_2^\dagger U_{\text{Out}}^\dagger \text{CNOT}_{\text{copy, out}} U_{\text{Out}} U_2 V_{\text{ext}, 2} |\psi_{\text{ext}}\rangle\|_2 \geq \\ & \eta + \epsilon + \xi \end{aligned}$$

We construct a QPT adversary \mathbf{A} that uses the QPT committer $\mathbf{C}^*. \text{Commit}$, its purified state $|\varphi\rangle$, and the unitaries $U_1, U_2, V_1, V_2, U_{\text{Out}}$ to break the collapsing property of the underlying NTCF family (Definition 3.10). We break the collapsing property as formulated in Remark 3.11. We distinguish between the case that $b = 0$ and the case that $b = 1$.

Case 1: $b = 0$. The adversary \mathbf{A} operates as follows:

1. **Adversary:** Upon receiving a public key \mathbf{pk}_0 from the challenger, where $(\mathbf{pk}_0, \mathbf{sk}_0) \leftarrow \text{Gen}_{\text{NTCF}}(1^\lambda)$:
 - (a) For every $i \in [n+1]$ generate $(\mathbf{pk}_i, \mathbf{sk}_i) \leftarrow \text{Gen}_{\text{NTCF}}(1^\lambda)$.
 - (b) Let $\mathbf{pk} = (\mathbf{pk}_0, \mathbf{pk}_1, \dots, \mathbf{pk}_{n+1})$.
 - (c) Compute $(\mathbf{y}, |\psi\rangle) \leftarrow \mathbf{C}^*. \text{Commit}(\mathbf{pk}, |\varphi\rangle)$.
 - (d) Parse $\mathbf{y} = (\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_{n+1})$
 - (e) Let $|\psi'\rangle = U(|+\rangle_{\text{coin}} \otimes |\psi_{\text{ext}}\rangle)$, where

$$U = |0\rangle\langle 0|_{\text{coin}} \otimes U_{\text{Out}} U_1 V_{\text{ext}, 1} + |1\rangle\langle 1|_{\text{coin}} \otimes U_{\text{Out}} U_2 V_{\text{ext}, 2}$$

Recall that U_{Out} first computes Ver which in Item 4d computes $\mathbf{m} \in \{0, 1\}^{n+1}$. U_{Out} stores in register out the output, which is the first bit of \mathbf{m} . We denote by preimage the registers that store the last n bits of \mathbf{m} .

(f) Send to the challenger the string \mathbf{y}_0 and the registers **out** and **preimage** of $|\psi'\rangle$.

Notice that since $b = 0$, U_{out} (as possibly U_1 and U_2) use only the secret keys $(\text{sk}_1, \dots, \text{sk}_{n+1})$, which A knows, and thus A can efficiently apply the unitary U to the state $|+\rangle_{\text{coin}} \otimes |\psi_{\text{ext}}\rangle$.

2. **Challenger:** Recall that the challenger applies in superposition the algorithm $\text{Check}_{\text{NTCF}}$ to the state it receives w.r.t. public key pk_0 and the image string \mathbf{y}_0 , and measures the bit indicating whether the output of $\text{Check}_{\text{NTCF}}$ is 1. If this is not the case it sends \perp . Otherwise, it chooses a random bit $u \leftarrow \{0, 1\}$ and measures this state if and only if $u = 1$. It then sends the resulting state to the adversary.

Note that by the two-to-one nature of the underlying NTCF family, measuring the entire state is equivalent to measuring only the first qubit of the state, i.e., register **out**. Thus, we can assume that the challenger measures only register **out** if and only if $u = 1$.

In addition, note that conditioned on the challenger not outputting \perp , the state is projected to $\Pi_{\text{Ver}} |\psi'\rangle$ (up to normalization), where $\Pi_{\text{Ver}} |\psi'\rangle$ is the state $|\psi'\rangle$ projected to the challenger accepting the state. Consider the state $\text{CNOT}_{\text{copy, out}}^u \Pi_{\text{Ver}} |\psi'\rangle$. Note that this state, with the **copy** register excluded, is indistinguishable from the state returned from the challenger conditioned on choosing the random bit u . Thus we think of the adversary as receiving this state.

3. **Adversary:** If the adversary receives \perp from the challenger, then it outputs a uniformly random u' . Note that this occurs with probability at most δ .

Otherwise, the adversary A receives the registers **out** and **preimage** from the challenger (either measured or not, depending on u). The joint state of the adversary and challenger at this point is $\text{CNOT}_{\text{copy, out}}^u \Pi_{\text{Ver}} |\psi'\rangle$, where all registers except **copy** are held by the adversary. The adversary does the following:

(a) Let

$$U' = |0\rangle\langle 0|_{\text{Coin}} \otimes U_1^\dagger U_{\text{out}}^\dagger + |1\rangle\langle 1|_{\text{Coin}} \otimes U_2^\dagger U_{\text{out}}^\dagger.$$

(b) Apply U' to the adversary's system, resulting in the joint state

$$\begin{aligned} & U' \text{CNOT}_{\text{copy, out}}^u \Pi_{\text{Ver}} |\psi'\rangle = \\ & U' \text{CNOT}_{\text{copy, out}}^u \Pi_{\text{Ver}} U |\psi_{\text{ext}}\rangle = \\ & (|0\rangle\langle 0|_{\text{Coin}} U_1^\dagger U_{\text{out}}^\dagger \text{CNOT}_{\text{copy, out}}^u \Pi_{\text{Ver}} U_{\text{out}} U_1 V_{\text{ext}, 1} + |1\rangle\langle 1|_{\text{Coin}} U_2^\dagger U_{\text{out}}^\dagger \text{CNOT}_{\text{copy, out}}^u \Pi_{\text{Ver}} U_{\text{out}} U_2 V_{\text{ext}, 2}) |\psi_{\text{ext}}\rangle \end{aligned}$$

(c) Output the measurement of the **Coin** register in the Hadamard basis, denoted by u' (i.e., $u' = 0$ if the measurement is $|+\rangle$ and is $u' = 1$ if the measurement is $|-\rangle$).

Consider the states

$$U_1^\dagger U_{\text{out}}^\dagger \text{CNOT}_{\text{copy, out}}^u \Pi_{\text{Ver}} U_{\text{out}} U_1 V_{\text{ext}, 1} |\psi_{\text{ext}}\rangle \quad \text{and} \quad U_2^\dagger U_{\text{out}}^\dagger \text{CNOT}_{\text{copy, out}}^u \Pi_{\text{Ver}} U_{\text{out}} U_2 V_{\text{ext}, 2} |\psi_{\text{ext}}\rangle$$

Note that for $u = 0$, these states are $(2\sqrt{\delta} + \epsilon)$ -close in $\|\cdot\|_2$ distance. This follows from the fact that by Lemma 3.1, together with the assumption that the probability that $|\psi_{\text{ext}}\rangle$ opens successfully is $\geq 1 - \delta$, it holds that for every $i \in \{1, 2\}$:

$$\mathbb{E}_{\substack{(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda) \\ (\mathbf{y}, |\psi\rangle) \leftarrow C^*. \text{Commit}(\text{pk}, |\varphi\rangle)}} \|U_i^\dagger U_{\text{out}}^\dagger \Pi_{\text{Ver}} U_{\text{out}} U_i V_{\text{ext}, i} |\psi_{\text{ext}}\rangle - U_i^\dagger U_{\text{out}}^\dagger U_{\text{out}} U_i V_{\text{ext}, i} |\psi_{\text{ext}}\rangle\|_2 \leq \sqrt{\delta},$$

and from our assumption that

$$\epsilon = \mathbb{E}_{\substack{(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Gen}(1^\lambda) \\ (\mathbf{y}, |\psi\rangle) \leftarrow \mathbf{C}^*. \text{Commit}(\mathbf{pk}, |\varphi\rangle)}} \|V_{\text{ext},1} |\psi_{\text{ext}}\rangle - V_{\text{ext},2} |\psi_{\text{ext}}\rangle\|_2.$$

This implies that there exists a negligible function μ such that

$$\begin{aligned} & \mathbb{E}_{\substack{(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Gen}(1^\lambda) \\ (\mathbf{y}, |\psi\rangle) \leftarrow \mathbf{C}^*. \text{Commit}(\mathbf{pk}, |\varphi\rangle)}} \|U_1^\dagger U_{\text{Out}}^\dagger \text{CNOT}_{\text{copy, out}} U_{\text{Out}} U_1 V_{\text{ext},1} |\psi_{\text{ext}}\rangle - U_2^\dagger U_{\text{Out}}^\dagger \text{CNOT}_{\text{copy, out}} U_{\text{Out}} U_2 V_{\text{ext},2} |\psi_{\text{ext}}\rangle\|_2 \leq \\ & 2\sqrt{\delta} + \epsilon + \mu \end{aligned}$$

On the other hand, by our contradiction assumption, for $u = 1$, these two states are $(2\sqrt{\delta} + \epsilon)$ -far. This, together with Claim 7.12 below, implies that \mathbf{A} indeed breaks the collapsing property of the underlying NTCF family.

Claim 7.12. *For any two states $|\psi_0\rangle$ and $|\psi_1\rangle$ such that $\| |\psi_0\rangle - |\psi_1\rangle \| = \epsilon$, and for $|\varphi\rangle = \frac{1}{\sqrt{2}} |0\rangle |\psi_0\rangle + \frac{1}{\sqrt{2}} |1\rangle |\psi_1\rangle$, it holds that*

$$\Pr[H[\varphi] \rightarrow 1] = \frac{\epsilon^2}{4}.$$

Proof. We calculate

$$\begin{aligned} \Pr[H[\varphi] \mapsto 1] &= \| \langle 1 | \otimes I \rangle H |\varphi\rangle \|^2 \\ &= \left\| \langle 1 | \otimes I \left(\frac{1}{\sqrt{2}} |+\rangle |\psi_0\rangle + \frac{1}{\sqrt{2}} |-\rangle |\psi_1\rangle \right) \right\|^2 \\ &= \left\| \frac{1}{2} |\psi_0\rangle - \frac{1}{2} |\psi_1\rangle \right\|^2 \\ &= \frac{1}{4} \epsilon^2. \quad \square \end{aligned}$$

Case 2: $b = 1$. We show how to use the adversary \mathbf{A} to break the extended collapsing game (see Claim 3.13). The adversary \mathbf{A} operates as follows:

1. Upon receiving public keys $(\mathbf{pk}_1, \dots, \mathbf{pk}_{n+1})$ from the challenger, where $(\mathbf{pk}_i, \mathbf{sk}_i) \leftarrow \text{Gen}_{\text{NTCF}}(1^\lambda)$ for every $i \in [n+1]$, do the following:
 - (a) Generate $(\mathbf{pk}_0, \mathbf{sk}_0) \leftarrow \text{Gen}_{\text{NTCF}}(1^\lambda)$.
 - (b) Let $\mathbf{pk} = (\mathbf{pk}_0, \mathbf{pk}_1, \dots, \mathbf{pk}_{n+1})$.
 - (c) Compute $(\mathbf{y}, |\psi\rangle) \leftarrow \mathbf{C}^*. \text{Commit}(\mathbf{pk}, |\varphi\rangle)$.
 - (d) Parse $\mathbf{y} = (\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_{n+1})$.
 - (e) Compute $((0, \mathbf{x}_0), (1, \mathbf{x}_1)) = \text{Invert}_{\text{NTCF}}(\mathbf{sk}_0, \mathbf{y}_0)$.
 - (f) Let $J = \{j \in \{2, \dots, n+1\} : x_{0,j-1} \oplus x_{1,j-1} = 1\} \cup \{1\}$.

(g) As in the $b = 0$ case, define

$$U = |0\rangle\langle 0|_{\text{coin}} \otimes U_{\text{Out}}U_1V_{\text{ext},1} + |1\rangle\langle 1|_{\text{coin}} \otimes U_{\text{Out}}U_2V_{\text{ext},2}$$

and prepare the state $|\psi'\rangle = U(|+\rangle_{\text{coin}} \otimes |\psi_{\text{ext}}\rangle)$.

Note that since $b = 1$ it holds that $|\psi'\rangle$ can be computed efficiently given sk_0

- (h) For every $j \in [J]$, denote by \mathcal{X}_j and \mathcal{Z}_j the registers in ρ' corresponding to d_j and x'_j , respectively.
- (i) Send J , $\{\mathbf{y}_j\}_{j \in J}$ and the registers $\{\mathcal{X}_j, \mathcal{Z}_j\}_{j \in J}$ of $|\psi'\rangle$.

2. Recall that the challenger applies in superposition the algorithm **Check** to the state it received w.r.t. the image strings $\{\mathbf{y}_j\}_{j \in J}$, where the j 'th check is w.r.t pk_j , and measures the bit indicating whether the output of **Check** is 1. If any of the outputs of **Check** are 0, the challenger immediately halts and sends \perp to the adversary. Otherwise, it chooses a random bit $u \leftarrow \{0, 1\}$ and applies Z^u to every \mathcal{X}_j register. It then sends the resulting state to the adversary.
3. If the adversary receives \perp , it returns a uniformly random u' . Otherwise, observe that once the adversary receives the state from the challenger, it is in possession of all the quantum registers. At this point, they are, up to normalization, in the state $Z_J^u \Pi_{\text{Ver}, J} |\psi'\rangle$, where $Z_J = \prod_{j \in J} Z_{\mathcal{X}_j}$ and $\Pi_{\text{Ver}, J} |\psi'\rangle$ is the state $|\psi'\rangle$ projected to an accepting state.

It then does the following:

(a) Let

$$U' = |0\rangle\langle 0|_{\text{Coin}} \otimes U_1^\dagger U_{\text{out}}^\dagger + |1\rangle\langle 1|_{\text{Coin}} \otimes U_2^\dagger U_{\text{out}}^\dagger.$$

(b) Apply U' to its registers, resulting in the state

$$U' Z_J^u \Pi_{\text{Ver}, J} |\psi'\rangle = \\ \left(|0\rangle\langle 0|_{\text{Coin}} \otimes U_1^\dagger U_{\text{out}}^\dagger Z_J^u \Pi_{\text{Ver}} U_{\text{out}} U_1 V_{\text{ext},1} + |1\rangle\langle 1|_{\text{Coin}} \otimes U_2^\dagger U_{\text{out}}^\dagger Z_J^u \Pi_{\text{Ver}} U_{\text{out}} U_2 V_{\text{ext},2} \right) (|+\rangle_{\text{Coin}} \otimes |\psi_{\text{ext}}\rangle)$$

(c) Output the measurement of the first register of this state in the Hadamard basis, denoted by u' .

Consider the states

$$U_1^\dagger U_{\text{out}}^\dagger Z_J^u \Pi_{\text{Ver}} U_{\text{out}} U_1 V_{\text{ext},1} |\psi_{\text{ext}}\rangle \quad \text{and} \quad U_2^\dagger U_{\text{out}}^\dagger Z_J^u \Pi_{\text{Ver}} U_{\text{out}} U_2 V_{\text{ext},2} |\psi_{\text{ext}}\rangle$$

Note that similarly to the $b = 0$ case, for $u = 0$ these states are $(2\sqrt{\delta} + \epsilon)$ -close in $\|\cdot\|_2$ distance. On the other hand, by our contradiction assumption, together with Lemma 3.3, for $u = 1$, these two states are $(2\sqrt{\delta} + \epsilon)$ -far in $\|\cdot\|_2$ distance. This together with Claim 7.12, implies that indeed **A** breaks the collapsing property of the underlying NTCF family.

Induction step: Suppose that the multi-qubit commitment scheme is sound for $\ell - 1$ and we prove that it is sound for ℓ . We need to prove that there exists a negligible function $\mu = \mu(\lambda)$ such that

$$\mathbb{E}_{\substack{(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda) \\ (\mathbf{y}, |\psi\rangle) \leftarrow \text{C}^*. \text{Commit}(\text{pk}, |\varphi\rangle)}} \left\| U_1^\dagger U_{\text{out}}^\dagger \text{CNOT}_{\text{copy}, \text{out}} U_{\text{out}} U_1 V_{\text{ext},1} |\psi_{\text{ext}}\rangle - U_2^\dagger U_{\text{out}}^\dagger \text{CNOT}_{\text{copy}, \text{out}} U_{\text{out}} U_2 V_{\text{ext},2} |\psi_{\text{ext}}\rangle \right\|_2 \leq \\ \eta + \epsilon + \mu$$

for $\eta = \sum_{j=1}^{\ell} 2\sqrt{\delta_j}$ and $\epsilon = \mathbb{E}_{\substack{(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda) \\ (\mathbf{y}, |\psi\rangle) \leftarrow \text{C}^*. \text{Commit}(\text{pk}, |\varphi\rangle)}} \|V_1 |\psi\rangle - V_2 |\psi\rangle\|_2$.

To this end, note that for every $i \in \{1, 2\}$

$$\begin{aligned} & U_i^\dagger U_{\text{Out}}^\dagger \text{CNOT}_{\text{copy}, \text{out}} U_{\text{Out}} U_i V_{\text{ext}, i} |\psi_{\text{ext}}\rangle = \\ & U_i^\dagger U_{\text{Out}}^\dagger \text{CNOT}_{\text{copy}_\ell, \text{out}_\ell} \text{CNOT}_{\text{copy}_{[1, \ell-1]}, \text{out}_{[1, \ell-1]}} U_{\text{Out}} U_i V_{\text{ext}, i} |\psi_{\text{ext}}\rangle = \\ & U_i^\dagger U_{\text{Out}_\ell}^\dagger \text{CNOT}_{\text{copy}_\ell, \text{out}_\ell} U_{\text{Out}_\ell} U_i V_{\text{ext}, i} V_{\text{ext}, i}^\dagger U_i^\dagger U_{\text{Out}_{[1, \ell-1]}}^\dagger \text{CNOT}_{\text{copy}_{[1, \ell-1]}, \text{out}_{[1, \ell-1]}} U_{\text{Out}_{[1, \ell-1]}} U_i V_{\text{ext}, i} |\psi_{\text{ext}}\rangle = \\ & U_i^\dagger U_{\text{Out}_\ell}^\dagger \text{CNOT}_{\text{copy}_\ell, \text{out}_\ell} U_{\text{Out}_\ell} U_i \underbrace{U_i^\dagger U_{\text{Out}_{[1, \ell-1]}}^\dagger \text{CNOT}_{\text{copy}_{[1, \ell-1]}, \text{out}_{[1, \ell-1]}} U_{\text{Out}_{[1, \ell-1]}} U_i V_{\text{ext}, i} |\psi_{\text{ext}}\rangle}_{V_i'} \end{aligned}$$

For every $i \in \{1, 2\}$, denote by

$$|\psi'_i\rangle = V_i' |\psi_{\text{ext}}\rangle$$

By the induction hypothesis, there exists a negligible function $\mu = \mu(\lambda)$ such that

$$\mathbb{E}_{\substack{(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda) \\ (\mathbf{y}, |\psi\rangle) \leftarrow \text{C}^*. \text{Commit}(\text{pk}, |\varphi\rangle)}} \| |\psi'_1\rangle - |\psi'_2\rangle \|_2 \leq \eta' + \mu$$

where $\eta' = \sum_{j=1}^{\ell-1} 2\sqrt{\delta_j} + \epsilon$. Denoting by $\epsilon' = \eta'$, our base case implies that there exists a negligible function $\nu = \nu(\lambda)$ such that

$$\mathbb{E}_{\substack{(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda) \\ (\mathbf{y}, |\psi\rangle) \leftarrow \text{C}^*. \text{Commit}(\text{pk}, |\varphi\rangle)}} \| U_1^\dagger U_{\text{Out}_\ell}^\dagger \text{CNOT}_{\text{copy}_\ell, \text{out}_\ell} U_{\text{out}_\ell} U_1 V_1' |\psi_{\text{ext}}\rangle - U_1^\dagger U_{\text{Out}_\ell}^\dagger \text{CNOT}_{\text{copy}_\ell, \text{out}_\ell} U_{\text{out}_\ell} U_1 V_2' |\psi_{\text{ext}}\rangle \|_2 \leq 2\sqrt{\delta_\ell} + \eta' + \nu$$

as desired. □

We next prove that both the non-succinct and the semi-succinct commitment schemes from Section 6.2 satisfy Equation (11).

Lemma 7.13. *The non-succinct commitment scheme described in Section 6.2 satisfies Equation (11) from Definition 5.6 assuming the underlying NTCF family has the adaptive hardcore bit property.*

Lemma 7.14. *The semi-succinct commitment scheme described in Section 6.2 (Definition 6.3) satisfies Equation (11) from Definition 5.6 assuming the underlying NTCF family is the one from [BCM⁺18] and assuming it has the distributional strong adaptive hardcore bit property (which is the case under LWE).*

Proof of Lemmas 7.13 and 7.14 We prove these two lemmas jointly since much of the proof is identical. In both cases we think of the public and secret keys as being

$$\text{pk} = (\text{pk}_1, \dots, \text{pk}_\ell) \quad \text{and} \quad \text{sk} = (\text{sk}_1, \dots, \text{sk}_\ell)$$

where in the non-succinct commitment each $(\mathbf{pk}_i, \mathbf{sk}_i) \leftarrow \text{Gen}_{\text{NTCF}}(1^\lambda)$ and in the semi-succinct commitment $(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Gen}_{\text{NTCF}}(1^\lambda)$, and for every $i \in [\ell]$

$$\mathbf{sk}_i = \mathbf{sk} \quad \text{and} \quad \mathbf{pk}_i = \mathbf{pk}$$

Fix any QPT cheating committer C^* .Commit with auxiliary quantum state σ that commits to an ℓ -qubit state. Denote by

$$(\mathbf{y}, \rho) \leftarrow C^*.\text{Commit}(\mathbf{pk}, \sigma),$$

where $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_\ell)$, each $\mathbf{y}_i = (\mathbf{y}_{i,0}, \mathbf{y}_1, \dots, \mathbf{y}_{i,n+1})$ and each $\mathbf{y}_{i,j}$ is in $\mathbb{R}_{\mathbf{pk}_j}$ which is the range of the NTCF function $\text{Eval}(\mathbf{pk}_j, \cdot)$. Fix any QPT algorithm $C^*.\text{Open}$. We start by defining the QPT extractor $\text{Ext}^{C^*.\text{Open}}(\mathbf{sk}, \mathbf{y}, \rho)$. We do so in two steps:

1. First, we define 2ℓ “operational observables” $\{P_{X_i}, P_{Z_i}\}_{i \in [\ell]}$ such that for every $i \in [\ell]$ and $b \in \{0, 1\}$,

$$(\mathbf{pk}, \mathbf{y}, \mathbf{m}_{\text{ideal}, i, b}) \equiv (\mathbf{pk}, \mathbf{y}, \mathbf{m}_{i, b})$$

where $(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Gen}(1^\lambda, 1^\ell)$,³⁷ $(\mathbf{y}, \rho) \leftarrow C^*.\text{Commit}(\mathbf{pk}, \sigma)$, $\mathbf{m}_{\text{ideal}, i, b}$ is obtained by measuring ρ in the P_{X_i} basis if $b = 1$ and measuring it in the P_{Z_i} basis if $b = 0$, and $\mathbf{m}_{i, b}$ is obtained by computing $\mathbf{z} \leftarrow C^*.\text{Open}(\rho, b^\ell)$ and setting $\mathbf{m}_{i, b} = \text{Out}(\mathbf{sk}, \mathbf{y}, (i, b), \mathbf{z}_i)$.

2. We then use these operational observables to extract a state τ . This is done following the approach of [Mah18, Vid20, BKL⁺22],

Defining the operational observables $\{P_{X_i}, P_{Z_i}\}_{i \in [\ell]}$. To define these operational observables formally, we add $L = \ell \cdot ((n+1)^2 + 1)$ ancilla registers to ρ , which we initialize to 0. We denote by

$$\rho_{\text{Ext}} = \rho \otimes |0^L\rangle\langle 0^L|,$$

where the first $\ell \cdot (n+1)^2$ ancilla registers are denoted by $\text{open} = (\text{open}_1, \dots, \text{open}_\ell)$, and these registers store the output $(\mathbf{z}_1, \dots, \mathbf{z}_\ell)$ generated by Open , where $\mathbf{z}_i \in \{0, 1\}^{(n+1)^2}$ is stored in open_i . The last ℓ ancilla registers are denoted by $\text{out} = (\text{out}_1, \dots, \text{out}_\ell)$, and these registers store the output (v_1, \dots, v_ℓ) generated by Out , where $v_i \in \{0, 1\}$ is stored in register out_i .

Definition 7.15. For any $(\mathbf{sk}, \mathbf{y})$ and any QPT algorithm $C^*.\text{Open}$ we define the operational observables $(P_{X_i}, P_{Z_i})_{i \in [\ell]}$ to be

$$P_{X_i} = U_1^\dagger \text{Out}_{i,1}^\dagger Z_{\text{out}_i} \text{Out}_{i,1} U_1$$

and

$$P_{Z_i} = U_0^\dagger \text{Out}_{i,0}^\dagger Z_{\text{out}_i} \text{Out}_{i,0} U_0$$

where for every $i \in [\ell]$ and every $b \in \{0, 1\}$,

- U_b is the unitary corresponding to $C^*.\text{Open}(\cdot, (b, \dots, b))$. The output is recorded in registers open .
- $\text{Out}_{i,b}$ computes $\text{Out}(\mathbf{sk}, \mathbf{y}, (i, b), \cdot)$ and records the output in the ancilla register out_i .
- Z_{out_i} is the Pauli Z operator applied only on the register out_i .

We next define the extractor Ext which uses the operational observables $\{P_{X_i}, P_{Z_i}\}_{i \in [\ell]}$, defined above. For the sake of simplicity, we define Ext to operate on pure states. The definition easily generalizes to mixed states by linearity.

³⁷In the semi-succinct setting $(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Gen}(1^\lambda)$.

$\text{Ext}^{\text{C}^*.\text{Open}}(\text{sk}, \mathbf{y}, |\varphi\rangle)$ operates as follows:

1. Consider the operational observables $\{P_{X_i}, P_{Z_i}\}_{i \in [\ell]}$ corresponding to (sk, \mathbf{y}) .

2. Prepare the state

$$\frac{1}{2^\ell} \sum_{\mathbf{r}, \mathbf{s} \in \{0,1\}^\ell} |\mathbf{r}, \mathbf{s}\rangle_{\text{Coin}} \otimes |0^\ell\rangle_{\mathcal{A}} \otimes |\varphi\rangle_{\mathcal{B}}.$$

3. Denote by

$$\mathbf{X}^{\mathbf{r}} = X_\ell^{r_\ell} \dots X_1^{r_1} \quad \text{and} \quad \mathbf{Z}^{\mathbf{s}} = Z_\ell^{s_\ell} \dots Z_1^{s_1}.$$

Similarly, denote by

$$P_{\mathbf{X}}^{\mathbf{r}} = P_{X_\ell}^{r_\ell} \dots P_{X_1}^{r_1} \quad \text{and} \quad P_{\mathbf{Z}}^{\mathbf{s}} = P_{Z_\ell}^{s_\ell} \dots P_{Z_1}^{s_1}.$$

4. Controlled on the values \mathbf{r}, \mathbf{s} of the Coin register, apply $\mathbf{Z}^{\mathbf{s}} \mathbf{X}^{\mathbf{r}}$ to the \mathcal{A} register and apply $P_{\mathbf{X}}^{\mathbf{r}} P_{\mathbf{Z}}^{\mathbf{s}}$ to the \mathcal{B} register to obtain the state

$$\frac{1}{2^\ell} \sum_{\mathbf{r}, \mathbf{s} \in \{0,1\}^\ell} |\mathbf{r}, \mathbf{s}\rangle_{\text{Coin}} \otimes \mathbf{Z}^{\mathbf{s}} \mathbf{X}^{\mathbf{r}} |0^\ell\rangle_{\mathcal{A}} \otimes P_{\mathbf{X}}^{\mathbf{r}} P_{\mathbf{Z}}^{\mathbf{s}} |\varphi\rangle_{\mathcal{B}}$$

5. Apply Hadamard gates $H^{\otimes 2\ell}$ to the Coin register in the to obtain the state

$$\frac{1}{4^\ell} \sum_{\mathbf{r}, \mathbf{s}, \mathbf{r}', \mathbf{s}' \in \{0,1\}^\ell} (-1)^{\mathbf{r} \cdot \mathbf{r}' + \mathbf{s} \cdot \mathbf{s}'} |\mathbf{r}', \mathbf{s}'\rangle_{\text{Coin}} \otimes \mathbf{Z}^{\mathbf{s}} \mathbf{X}^{\mathbf{r}} |0^\ell\rangle_{\mathcal{A}} \otimes P_{\mathbf{X}}^{\mathbf{r}} P_{\mathbf{Z}}^{\mathbf{s}} |\varphi\rangle_{\mathcal{B}}$$

where

$$\mathbf{r} \cdot \mathbf{r}' = \sum_{i=1}^{\ell} r_i \cdot r'_i \pmod{2} \quad \text{and} \quad \mathbf{s} \cdot \mathbf{s}' = \sum_{i=1}^{\ell} s_i \cdot s'_i \pmod{2}.$$

6. Apply $\mathbf{X}^{\mathbf{s}'} \mathbf{Z}^{\mathbf{r}'}$ to the \mathcal{A} register. Note that

$$\begin{aligned} \mathbf{X}^{\mathbf{s}'} \mathbf{Z}^{\mathbf{r}'} \mathbf{Z}^{\mathbf{s}} \mathbf{X}^{\mathbf{r}} |0^\ell\rangle &= \\ \mathbf{X}^{\mathbf{s}'} \mathbf{Z}^{\mathbf{s}} \mathbf{Z}^{\mathbf{r}'} \mathbf{X}^{\mathbf{r}} |0^\ell\rangle &= \\ (-1)^{\mathbf{r} \cdot \mathbf{r}'} \mathbf{X}^{\mathbf{s}'} \mathbf{Z}^{\mathbf{s}} \mathbf{X}^{\mathbf{r}} \mathbf{Z}^{\mathbf{r}'} |0^\ell\rangle &= \\ (-1)^{\mathbf{r} \cdot \mathbf{r}'} \mathbf{X}^{\mathbf{s}'} \mathbf{Z}^{\mathbf{s}} \mathbf{X}^{\mathbf{r}} |0^\ell\rangle &= \\ (-1)^{\mathbf{r} \cdot \mathbf{r}' + \mathbf{s} \cdot \mathbf{s}'} \mathbf{Z}^{\mathbf{s}} \mathbf{X}^{\mathbf{s}'} \mathbf{X}^{\mathbf{r}} |0^\ell\rangle &= \\ (-1)^{\mathbf{r} \cdot \mathbf{r}' + \mathbf{s} \cdot \mathbf{s}'} \mathbf{Z}^{\mathbf{s}} \mathbf{X}^{\mathbf{r}} \mathbf{X}^{\mathbf{s}'} |0^\ell\rangle &= \\ (-1)^{\mathbf{r} \cdot \mathbf{r}' + \mathbf{s} \cdot \mathbf{s}'} \mathbf{Z}^{\mathbf{s}} \mathbf{X}^{\mathbf{r}} |\mathbf{s}'\rangle. & \end{aligned}$$

Therefore the state obtained is

$$\frac{1}{4^\ell} \sum_{\mathbf{r}, \mathbf{s}, \mathbf{r}', \mathbf{s}' \in \{0,1\}^\ell} |\mathbf{r}', \mathbf{s}'\rangle_{\text{Coin}} \otimes \mathbf{Z}^{\mathbf{s}} \mathbf{X}^{\mathbf{r}} |\mathbf{s}'\rangle_{\mathcal{A}} \otimes P_{\mathbf{X}}^{\mathbf{r}} P_{\mathbf{Z}}^{\mathbf{s}} |\varphi\rangle_{\mathcal{B}}$$

which is equal to the state

$$\left(\frac{1}{\sqrt{2}^\ell} (|0\rangle + |1\rangle)^{\otimes \ell} \right) \otimes \frac{1}{(2\sqrt{2})^\ell} \sum_{\mathbf{r}, \mathbf{s}, \mathbf{s}' \in \{0,1\}^\ell} |\mathbf{s}'\rangle \otimes \mathbf{Z}^{\mathbf{s}} \mathbf{X}^{\mathbf{r}} |\mathbf{s}'\rangle_{\mathcal{A}} \otimes P_{\mathbf{X}}^{\mathbf{r}} P_{\mathbf{Z}}^{\mathbf{s}} |\varphi\rangle_{\mathcal{B}}.$$

7. Discard the first ℓ registers to obtain the state

$$\frac{1}{(2\sqrt{2})^\ell} \sum_{\mathbf{r}, \mathbf{s}, \mathbf{s}' \in \{0,1\}^\ell} |\mathbf{s}'\rangle \otimes \mathbf{Z}^{\mathbf{s}} \mathbf{X}^{\mathbf{r}} |\mathbf{s}'\rangle_{\mathcal{A}} \otimes P_{\mathbf{X}}^{\mathbf{r}} P_{\mathbf{Z}}^{\mathbf{s}} |\varphi\rangle_{\mathcal{B}}.$$

8. Output the state $\tau_{\mathcal{A}, \mathcal{B}}$ that is the reduced state of the above on registers \mathcal{A}, \mathcal{B} .

We next prove that

$$\text{Real}^{\text{C}^*. \text{Open}}(\lambda, \mathbf{b}, \boldsymbol{\sigma}) \stackrel{10\sqrt{\delta}}{\approx} \text{Ideal}^{\text{Ext}, \text{C}^*. \text{Open}}(\lambda, \mathbf{b}, \boldsymbol{\sigma}). \quad (21)$$

To this end, for a given $\mathbf{b} \in \{0,1\}^\ell$, denote by

$$I = \{i \in [\ell] : \mathbf{b}_i = 0\} \quad \text{and} \quad J = \{j \in [\ell] : \mathbf{b}_j = 1\},$$

so that I and J partition $[\ell]$.

Next we define a new opening algorithm $\text{C}^*. \text{Open}_{[\ell]}$. We first give a ‘‘buggy’’ definition of $\text{C}^*. \text{Open}_{[\ell]}$, and then show how to fix it in Remark 7.16. $\text{C}^*. \text{Open}_{[\ell]}$ on input $(\boldsymbol{\rho}, \mathbf{b})$ does the following:

1. Compute $\boldsymbol{\rho}_1 = U_0^\dagger \text{CNOT}_{\text{copy}_I, \text{open}_I} U_0[\boldsymbol{\rho}]$, where open_I is the register that contains the openings $\{\mathbf{z}_i\}_{i \in I}$, and $\text{CNOT}_{\text{copy}_I, \text{open}_I}$ copies the content of this register to a fresh register denoted by copy_I .
2. Measure the registers copy_I of $\boldsymbol{\rho}_1$ to obtain $\{\mathbf{z}_i\}_{i \in I}$ and denote the resulting state by $\boldsymbol{\rho}_2$.
3. Compute $\boldsymbol{\rho}_3 = U_1^\dagger \text{CNOT}_{\text{copy}_J, \text{open}_J} U_1[\boldsymbol{\rho}_2]$, where open_J is the register that contains the openings $\{\mathbf{z}_j\}_{j \in J}$, and $\text{CNOT}_{\text{copy}_J, \text{open}_J}$ copies the content of this register to a fresh register denoted by copy_J .
4. Measure the registers copy_J of $\boldsymbol{\rho}_3$ to obtain $\{\mathbf{z}_j\}_{j \in J}$ and denote the resulting state by $\boldsymbol{\rho}_4$.
5. Output $((\mathbf{z}_1, \dots, \mathbf{z}_\ell), \boldsymbol{\rho}_4)$.

Remark 7.16. *We remark that as written, $\text{C}^*. \text{Open}_{[\ell]}(\boldsymbol{\rho}, \mathbf{b})$ may be rejected with high probability. The reason is that, while the standard basis openings of $\text{C}^*. \text{Open}_{[\ell]}$ and $\text{C}^*. \text{Open}$ are identical, $\text{C}^*. \text{Open}_{[\ell]}$ can completely fail to open in the Hadamard basis, since after computing the standard basis openings its state becomes $U_0^\dagger \text{CNOT}_{\text{copy}_I, \text{open}_I} U_0[\boldsymbol{\rho}]$, with the copy_I registers measured. This is a disturbed state and it is no longer clear that computing the Hadamard basis opening on it will give an accepting opening.*

To ensure that $\text{C}^. \text{Open}_{[\ell]}(\boldsymbol{\rho}, \mathbf{b})$ is accepted with the same probability as $\text{C}^*. \text{Open}(\boldsymbol{\rho}, \mathbf{b})$, up to negligible factors, we need to ensure that the state after computing the standard basis openings remains undisturbed, or at least that this disturbance is undetected by the algorithms that compute*

the Hadamard basis opening and verify whether this opening is valid. To achieve this we slightly modify $\mathbf{C}^*.\text{Open}_{[\ell]}$ and allow it to compute the standard basis opening using $(\mathbf{sk}_1, \dots, \mathbf{sk}_{n+1})$. We note that such opening algorithms are allowed in Corollary 7.11 (which we will later use in our analysis).

Specifically, $\mathbf{C}^*.\text{Open}_{[\ell]}$, rather than placing the output of U_0 in the open_I registers, which when measured may disturb the state, we use $(\mathbf{sk}_1, \dots, \mathbf{sk}_{n+1})$ to apply the following post-opening unitary to each open_i register, to ensure that when measured the disturbance will not be noticed. Recall that open_i contains a vector $\mathbf{z} = (\mathbf{z}_1, \dots, \mathbf{z}_{n+1}) \in \{0, 1\}^{(n+1)^2}$ where each $\mathbf{z}_j \in \{0, 1\}^{n+1}$. The post-opening unitary does the following:

1. Coherently compute for every $j \in [n+1]$ the bit $m_j = \mathbf{z}_j \cdot (1, \mathbf{x}'_{j,0} \oplus \mathbf{x}'_{j,1})$, where $\mathbf{x}'_{j,0}$ and $\mathbf{x}'_{j,1}$ are the two preimages of $\mathbf{y}_{i,j}$ that are computed using \mathbf{sk}_j .
2. Let $\mathbf{m} = (m_1, \dots, m_{n+1}) \in \{0, 1\}^{n+1}$.
Note that if \mathbf{z} is a successful opening (i.e., it is accepted) then \mathbf{m} is a preimage of $\mathbf{y}_{i,0}$, and whether a preimage is measured or not is undetectable without knowing \mathbf{sk}_0 , due to the collapsing property of the underlying NTCF family.
3. On an ancilla register, compute a super-position over all $\mathbf{z}' = (\mathbf{z}'_1, \dots, \mathbf{z}'_{n+1}) \in \{0, 1\}^{(n+1)^2}$ such that for every $j \in [n+1]$ $m_j = \mathbf{z}'_j \cdot (1, \mathbf{x}'_{j,0} \oplus \mathbf{x}'_{j,1})$.
4. Swap register open_i with the ancilla register above, so that now $\mathbf{z}' = (\mathbf{z}'_1, \dots, \mathbf{z}'_{n+1})$ is in register open_i .

Now we can argue that the residual state after computing the standard basis opening seems undisturbed for anyone who does not know \mathbf{sk}_0 due to the collapsing property of the underlying NTCF family, and computing the Hadamard opening and verification of it does not use \mathbf{sk}_0 (and is done publicly given only \mathbf{pk}).

Note that since

$$\delta = \mathbb{E}_{\substack{(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Gen}(1^\lambda, 1^\ell) \\ (\mathbf{y}, \rho) \leftarrow \mathbf{C}^*.\text{Commit}(\mathbf{pk}, \sigma)}}} \max_{\mathbf{b}' \in \{\mathbf{b}, \mathbf{0}, \mathbf{1}\}} \Pr[\text{Ver}(\mathbf{sk}, \mathbf{y}, \mathbf{b}', \mathbf{C}^*.\text{Open}(\rho, \mathbf{b}')) = 0].$$

it holds that

$$\mathbb{E}_{\substack{(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Gen}(1^\lambda, 1^\ell) \\ (\mathbf{y}, \rho) \leftarrow \mathbf{C}^*.\text{Commit}(\mathbf{pk}, \sigma)}}} \max_{\mathbf{b}' \in \{\mathbf{b}, \mathbf{0}, \mathbf{1}\}} \Pr[\text{Ver}(\mathbf{sk}, \mathbf{y}, \mathbf{b}', \mathbf{C}^*.\text{Open}_{[\ell]}(\rho, \mathbf{b}')) = 0] \leq 2\delta. \quad (22)$$

This is the case since the probability that $\mathbf{C}^*.\text{Open}_{[\ell]}(\rho, \mathbf{b})$ is rejected is bounded by the sum of the probabilities that $\mathbf{C}^*.\text{Open}(\rho, \mathbf{0}^\ell)$ is rejected and $\mathbf{C}^*.\text{Open}(\rho, \mathbf{1}^\ell)$ is rejected.

By Corollary 7.11, we conclude that for every $\mathbf{b} \in \{0, 1\}^\ell$,

$$\text{Real}^{\mathbf{C}^*.\text{Open}_{[\ell]}}(\lambda, \mathbf{b}, \sigma) \stackrel{6\sqrt{\delta}}{\approx} \text{Real}^{\mathbf{C}^*.\text{Open}}(\lambda, \mathbf{b}, \sigma) \quad (23)$$

This implies that to prove Equation (21) it suffices to prove

$$\text{Real}^{\mathbf{C}^*.\text{Open}_{[\ell]}}(\lambda, \mathbf{b}, \sigma) \stackrel{4\sqrt{\delta}}{\approx} \text{Ideal}^{\text{Ext}, \mathbf{C}^*.\text{Open}}(\lambda, \mathbf{b}, \sigma) \quad (24)$$

To this end, we first compute the distribution of measurement outcomes on the extracted state. While in general the input to the extractor is a mixed state ρ , we will perform the calculations for a general pure state $|\varphi\rangle$ instead. The results we obtain will hold for any pure state $|\varphi\rangle$. Thus, they will extend by convexity to the post-commitment state ρ as well, since we can always write $\rho = \sum_k p_k |\varphi_k\rangle\langle\varphi_k|$ for some collection of pure states $\{|\varphi_k\rangle\}$.

As a first step in the computation, we remark that for every $i, j \in [\ell]$, it holds that P_{Z_i} and P_{Z_j} commute and P_{X_i} and P_{X_j} commute. This follows from the fact that we defined all the P_{Z_i} with respect to the same unitary U_0 and defined all the P_{X_i} with respect to the same unitary U_1 . Thus, for an input state $|\varphi\rangle$, the output of the extractor can be written as

$$\frac{1}{(2\sqrt{2})^\ell} \sum_{\mathbf{r}, \mathbf{s}, \mathbf{s}' \in \{0,1\}^\ell} |\mathbf{s}'\rangle_{\text{Coin}} \otimes \mathbf{Z}^{\mathbf{s}} \mathbf{X}^{\mathbf{r}} |\mathbf{s}'\rangle_{\mathcal{A}} \otimes P_{\mathbf{X}_I}^{\mathbf{r}_I} P_{\mathbf{X}_J}^{\mathbf{r}_J} P_{\mathbf{Z}_J}^{\mathbf{s}_J} P_{\mathbf{Z}_I}^{\mathbf{s}_I} |\varphi\rangle_{\mathcal{B}}.$$

Measuring the I registers of \mathcal{A} in the standard basis. Now, we imagine measuring the I registers of \mathcal{A} in the standard basis; we denote these registers by \mathcal{A}_I . When we measure them we obtain an outcome which we will denote \mathbf{a}_I . The unnormalized post-measurement state is obtained by applying the projector $I \otimes |\mathbf{a}_I\rangle\langle\mathbf{a}_I|_{\mathcal{A}_I}$ to the state, where the factor of identity acts on all registers other than \mathcal{A}_I . To calculate what happens, let us examine what happens when we apply the projector $|\mathbf{a}_I\rangle\langle\mathbf{a}_I|$ to $\mathbf{Z}_I^{\mathbf{s}_I} \mathbf{X}_I^{\mathbf{r}_I} |\mathbf{s}'_I\rangle$. Note that

$$\begin{aligned} \langle\mathbf{a}_I| \mathbf{Z}_I^{\mathbf{s}_I} \mathbf{X}_I^{\mathbf{r}_I} |\mathbf{s}'_I\rangle &= \\ \prod_{i \in I} \langle a_i | Z_i^{s_i} X_i^{r_i} |s'_i\rangle &= \\ \prod_{i \in I} \langle a_i | Z_i^{s_i} |s'_i \oplus r_i\rangle &= \\ \prod_{i \in I} \langle a_i | (-1)^{s_i \cdot a_i} |s'_i \oplus r_i\rangle &= \\ \prod_{i \in I} (-1)^{s_i \cdot a_i} \langle a_i | s'_i \oplus r_i\rangle, \end{aligned}$$

where for every $i \in I$, $\langle a_i | s'_i \oplus r_i\rangle$ is 1 if $a_i \oplus r_i = s'_i$, and 0 otherwise. This means that if we obtain an outcome \mathbf{a}_I , then we force the \mathbf{s}'_I register to be $\mathbf{a}_I \oplus \mathbf{r}_I$. This means that the sum over \mathbf{s}' collapses to a sum over \mathbf{s}'_J , since J is the complement of I . Thus, we obtain the unnormalized post-measurement state

$$\frac{1}{(2\sqrt{2})^\ell} \sum_{\mathbf{r}, \mathbf{s} \in \{0,1\}^\ell, \mathbf{s}'_J \in \{0,1\}^{|J|}} (-1)^{\mathbf{s}_I \cdot \mathbf{a}_I} |\mathbf{a}_I \oplus \mathbf{r}_I\rangle_{\text{Coin}_I} |\mathbf{s}'_J\rangle_{\text{Coin}_J} \otimes |\mathbf{a}_I\rangle_{\mathcal{A}_I} \otimes Z_J^{\mathbf{s}_J} X_J^{\mathbf{r}_J} |\mathbf{s}'_J\rangle_{\mathcal{A}_J} \otimes P_{\mathbf{X}_I}^{\mathbf{r}_I} P_{\mathbf{X}_J}^{\mathbf{r}_J} P_{\mathbf{Z}_J}^{\mathbf{s}_J} P_{\mathbf{Z}_I}^{\mathbf{s}_I} |\varphi\rangle_{\mathcal{B}}$$

Note that

$$\frac{1}{2^{|I|}} \sum_{\mathbf{s}_I \in \{0,1\}^{|I|}} (-1)^{\mathbf{s}_I \cdot \mathbf{a}_I} P_{Z_I}^{\mathbf{s}_I} = \frac{1}{2^{|I|}} \prod_{i \in I} \sum_{s_i \in \{0,1\}} (-1)^{s_i \cdot a_i} P_{Z_i}^{s_i} = \prod_{i \in I} \frac{I + (-1)^{a_i} P_{Z_i}}{2}$$

and thus the state we obtain after the projection is equal to

$$\frac{1}{2^{|J|} \cdot 2^{\ell/2}} \sum_{\mathbf{r} \in \{0,1\}^\ell, \mathbf{s}_J, \mathbf{s}'_J \in \{0,1\}^{|J|}} |\mathbf{a}_I \oplus \mathbf{r}_I\rangle \otimes |\mathbf{s}'_J\rangle \otimes |\mathbf{a}_I\rangle_{\mathcal{A}_I} \otimes Z_J^{\mathbf{s}_J} X_J^{\mathbf{r}_J} |\mathbf{s}'_J\rangle_{\mathcal{A}_J} \otimes P_{\mathbf{X}_I}^{\mathbf{r}_I} P_{\mathbf{X}_J}^{\mathbf{r}_J} P_{\mathbf{Z}_J}^{\mathbf{s}_J} \prod_{i \in I} \left(\frac{I + (-1)^{a_i} P_{Z_i}}{2} \right) |\varphi\rangle_{\mathcal{B}}.$$

Denoting by $\Pi_{P_{\mathbf{Z}_I, \mathbf{a}_I}} = \prod_{i \in I} \frac{I + (-1)^{a_i} P_{Z_i}}{2}$, the above projected state is equal to

$$|\Psi_{\mathbf{a}_I}\rangle = \frac{1}{2^{|\mathcal{J}|} \cdot 2^{\ell/2}} \sum_{\mathbf{r} \in \{0,1\}^{\ell}, \mathbf{s}_J, \mathbf{s}'_J \in \{0,1\}^{|\mathcal{J}|}} |\mathbf{a}_I \oplus \mathbf{r}_I\rangle_{\text{Coin}_I} \otimes |\mathbf{s}'_J\rangle_{\text{Coin}_J} \otimes |\mathbf{a}_I\rangle_{\mathcal{A}_I} \otimes Z_J^{\mathbf{s}_J} X_J^{\mathbf{r}_J} |\mathbf{s}'_J\rangle_{\mathcal{A}_J} \otimes P_{\mathbf{X}_I}^{\mathbf{r}_I} P_{\mathbf{X}_J}^{\mathbf{r}_J} P_{\mathbf{Z}_J}^{\mathbf{s}_J} \Pi_{P_{\mathbf{Z}_I, \mathbf{a}_I}} |\varphi\rangle_{\mathcal{B}}. \quad (25)$$

The square norm of this unnormalized state is the probability that the measurement returns outcome \mathbf{a}_I . We now calculate this:

$$\begin{aligned} \Pr[\mathbf{a}_I] &= \frac{1}{2^{2|\mathcal{J}|} \cdot 2^{\ell}} \sum_{\mathbf{r} \in \{0,1\}^{\ell}, \mathbf{s}'_J \in \{0,1\}^{|\mathcal{J}|}} \left\| \sum_{\mathbf{s}_J \in \{0,1\}^{|\mathcal{J}|}} Z_J^{\mathbf{s}_J} X_J^{\mathbf{r}_J} |\mathbf{s}'_J\rangle_{\mathcal{A}_J} \otimes P_{\mathbf{X}_I}^{\mathbf{r}_I} P_{\mathbf{X}_J}^{\mathbf{r}_J} P_{\mathbf{Z}_J}^{\mathbf{s}_J} \Pi_{P_{\mathbf{Z}_I, \mathbf{a}_I}} |\varphi\rangle_{\mathcal{B}} \right\|^2 \\ &= \frac{1}{2^{2|\mathcal{J}|} \cdot 2^{\ell}} \sum_{\mathbf{r} \in \{0,1\}^{\ell}, \mathbf{s}'_J \in \{0,1\}^{|\mathcal{J}|}} \left\| \sum_{\mathbf{s}_J \in \{0,1\}^{|\mathcal{J}|}} (-1)^{\mathbf{s}_J \cdot (\mathbf{s}'_J + \mathbf{r}_J)} |\mathbf{s}'_J + \mathbf{r}_J\rangle_{\mathcal{A}_J} \otimes P_{\mathbf{X}_I}^{\mathbf{r}_I} P_{\mathbf{X}_J}^{\mathbf{r}_J} P_{\mathbf{Z}_J}^{\mathbf{s}_J} \Pi_{P_{\mathbf{Z}_I, \mathbf{a}_I}} |\varphi\rangle_{\mathcal{B}} \right\|^2 \\ &= \frac{1}{2^{2|\mathcal{J}|} \cdot 2^{\ell}} \sum_{\mathbf{r} \in \{0,1\}^{\ell}, \mathbf{s}'_J \in \{0,1\}^{|\mathcal{J}|}} \sum_{\mathbf{s}_J, \mathbf{s}''_J \in \{0,1\}^{|\mathcal{J}|}} (-1)^{(\mathbf{s}_J + \mathbf{s}''_J) \cdot (\mathbf{s}'_J + \mathbf{r}_J)} \langle \varphi | \Pi_{P_{\mathbf{Z}_I, \mathbf{a}_I}} P_{\mathbf{Z}_J}^{\mathbf{s}''_J} P_{\mathbf{Z}_J}^{\mathbf{s}_J} \Pi_{P_{\mathbf{Z}_I, \mathbf{a}_I}} |\varphi\rangle_{\mathcal{B}} \\ &= \frac{1}{2^{2|\mathcal{J}|}} \sum_{\mathbf{s}'_J \in \{0,1\}^{|\mathcal{J}|}} \sum_{\mathbf{s}_J \in \{0,1\}^{|\mathcal{J}|}} \langle \varphi | \Pi_{P_{\mathbf{Z}_I, \mathbf{a}_I}} \Pi_{P_{\mathbf{Z}_I, \mathbf{a}_I}} |\varphi\rangle_{\mathcal{B}} \\ &= \langle \varphi | \Pi_{P_{\mathbf{Z}_I, \mathbf{a}_I}} \Pi_{P_{\mathbf{Z}_I, \mathbf{a}_I}} |\varphi\rangle_{\mathcal{B}}. \end{aligned} \quad (26)$$

Thus, we have shown that the outcome distribution from the extracted state is identical to the outcome distribution from measuring the original state $|\varphi\rangle$.

Measuring the J registers of \mathcal{A} in the Hadamard basis. Now, we imagine taking the standard basis post-measurement state $|\Psi_{\mathbf{a}_I}\rangle$, and then further measuring the J registers of \mathcal{A} in the Hadamard basis. We denote these registers by \mathcal{A}_J and the outcome by \mathbf{a}_J . To obtain the unnormalized post-measurement state after this measurement, we apply the projector $H^{\otimes |\mathcal{J}|} |\mathbf{a}_J\rangle \langle \mathbf{a}_J| H^{\otimes |\mathcal{J}|}$ to the J registers of \mathcal{A} . Note that

$$\begin{aligned} \langle \mathbf{a}_J | H^{\otimes |\mathcal{J}|} Z_J^{\mathbf{s}_J} X_J^{\mathbf{r}_J} |\mathbf{s}'_J\rangle &= \\ \prod_{j \in \mathcal{J}} \langle a_j | H Z_j^{\mathbf{s}_j} X_j^{\mathbf{r}_j} |s'_j\rangle &= \\ \prod_{j \in \mathcal{J}} \langle a_j | H Z_j^{\mathbf{s}_j} |s'_j \oplus r_j\rangle &= \\ \prod_{j \in \mathcal{J}} \langle a_j | H (-1)^{\mathbf{s}_j \cdot (\mathbf{s}'_j \oplus r_j)} |s'_j \oplus r_j\rangle &= \\ \prod_{j \in \mathcal{J}} \frac{(-1)^{\mathbf{s}_j \cdot (\mathbf{s}'_j \oplus r_j)}}{\sqrt{2}} \langle a_j | (|0\rangle + (-1)^{\mathbf{s}'_j \oplus r_j} |1\rangle) &= \\ \frac{1}{2^{|\mathcal{J}|/2}} \prod_{j \in \mathcal{J}} (-1)^{(\mathbf{s}_j \oplus a_j) \cdot (\mathbf{s}'_j \oplus r_j)} \triangleq \beta_J^{\mathbf{s}'_J} \end{aligned}$$

Thus we obtain the state

$$\frac{1}{2^{|J|} \cdot 2^{\ell/2}} \sum_{\mathbf{r} \in \{0,1\}^\ell, \mathbf{s}_J, \mathbf{s}'_J \in \{0,1\}^{|J|}} |\mathbf{a}_I \oplus \mathbf{r}_I\rangle \otimes \beta_J^{\mathbf{s}'_J} |\mathbf{s}'_J\rangle \otimes |\mathbf{a}_I\rangle_{\mathcal{A}_I} \otimes H^{\otimes |J|} |\mathbf{a}_J\rangle_{\mathcal{A}_J} \otimes P_{\mathbf{X}_I}^{\mathbf{r}_I} P_{\mathbf{X}_J}^{\mathbf{r}_J} P_{\mathbf{Z}_J}^{\mathbf{s}_J} \Pi_{P_{\mathbf{Z}_I, \mathbf{a}_I}} |\varphi\rangle_{\mathcal{B}}.$$

Next, we observe that

$$\sum_{\mathbf{s}'_J} \beta_J^{\mathbf{s}'_J} |\mathbf{s}'_J\rangle = \frac{1}{\sqrt{2^{|J|}}} \sum_{\mathbf{s}'_J} (-1)^{(\mathbf{s}_J \oplus \mathbf{a}_J) \cdot (\mathbf{s}'_J \oplus \mathbf{r}_J)} |\mathbf{s}'_J\rangle \quad (27)$$

$$= (-1)^{(\mathbf{s}_J \oplus \mathbf{a}_J) \cdot \mathbf{r}_J} H^{\otimes |J|} |\mathbf{s}_J \oplus \mathbf{a}_J\rangle \quad (28)$$

Thus, applying Equation (28) to simplify the sum over \mathbf{s}'_J , we can write this as

$$|\Psi_{\mathbf{a}_I, \mathbf{a}_J}\rangle = \frac{1}{2^{\ell/2} 2^{|J|}} \sum_{\mathbf{r} \in \{0,1\}^\ell, \mathbf{s}_J \in \{0,1\}^{|J|}} (-1)^{(\mathbf{s}_J \oplus \mathbf{a}_J) \cdot \mathbf{r}_J} |\mathbf{a}_I \oplus \mathbf{r}_I\rangle_{\text{Coin}_I} \otimes H^{\otimes |J|} |\mathbf{s}_J \oplus \mathbf{a}_J\rangle_{\text{Coin}_J} \\ \otimes |\mathbf{a}_I\rangle_{\mathcal{A}_I} \otimes H^{\otimes |J|} |\mathbf{a}_J\rangle_{\mathcal{A}_J} \otimes P_{\mathbf{X}_I}^{\mathbf{r}_I} P_{\mathbf{X}_J}^{\mathbf{r}_J} P_{\mathbf{Z}_J}^{\mathbf{s}_J} \Pi_{P_{\mathbf{Z}_I, \mathbf{a}_I}} |\varphi\rangle_{\mathcal{B}}.$$

Note that

$$\frac{1}{2^{|J|}} \sum_{\mathbf{r}_J \in \{0,1\}^{|J|}} (-1)^{(\mathbf{s}_J \oplus \mathbf{a}_J) \cdot \mathbf{r}_J} P_{\mathbf{X}_J}^{\mathbf{r}_J} = \prod_{j \in J} \frac{I + (-1)^{s_j \oplus a_j} P_{X_j}}{2}.$$

Let us define

$$\Pi_{P_{\mathbf{X}_J, \mathbf{s}_J \oplus \mathbf{a}_J}} = \prod_{j \in J} \frac{I + (-1)^{s_j \oplus a_j} P_{X_j}}{2}.$$

Then we can rewrite $|\Psi_{\mathbf{a}_I, \mathbf{a}_J}\rangle$ as

$$|\Psi_{\mathbf{a}_I, \mathbf{a}_J}\rangle = \frac{1}{2^{\ell/2}} \sum_{\mathbf{r}_I \in \{0,1\}^{\ell}, \mathbf{s}_J \in \{0,1\}^{|J|}} |\mathbf{a}_I \oplus \mathbf{r}_I\rangle_{\text{Coin}_I} \otimes H^{\otimes |J|} |\mathbf{s}_J \oplus \mathbf{a}_J\rangle_{\text{Coin}_J} \otimes |\mathbf{a}_I\rangle_{\mathcal{A}_I} \otimes H^{\otimes |J|} |\mathbf{a}_J\rangle_{\mathcal{A}_J} \\ \otimes P_{\mathbf{X}_I}^{\mathbf{r}_I} \Pi_{P_{\mathbf{X}_J, \mathbf{a}_J \oplus \mathbf{s}_J}} P_{\mathbf{Z}_J}^{\mathbf{s}_J} \Pi_{P_{\mathbf{Z}_I, \mathbf{a}_I}} |\varphi\rangle_{\mathcal{B}}.$$

The square norm of this unnormalized state is the probability that the measurement returns outcome \mathbf{a}_J . We now calculate this:

$$\Pr[\mathbf{a}_I, \mathbf{a}_J] = \frac{2^{|I|}}{2^\ell} \left\| \sum_{\mathbf{s}_J \in \{0,1\}^{|J|}} H^{\otimes |J|} |\mathbf{s}_J \oplus \mathbf{a}_J\rangle_{\text{Coin}_J} \otimes \Pi_{P_{\mathbf{X}_J, \mathbf{a}_J \oplus \mathbf{s}_J}} P_{\mathbf{Z}_J}^{\mathbf{s}_J} \Pi_{P_{\mathbf{Z}_I, \mathbf{a}_I}} |\varphi\rangle_{\mathcal{B}} \right\|^2 \\ = \frac{1}{2^{|J|}} \sum_{\mathbf{s}_J \in \{0,1\}^{|J|}} \|\Pi_{P_{\mathbf{X}_J, \mathbf{a}_J \oplus \mathbf{s}_J}} P_{\mathbf{Z}_J}^{\mathbf{s}_J} \Pi_{P_{\mathbf{Z}_I, \mathbf{a}_I}} |\varphi\rangle_{\mathcal{B}}\|^2 \\ = \mathbb{E}_{\mathbf{s}_J \in \{0,1\}^{|J|}} \|\Pi_{P_{\mathbf{X}_J, \mathbf{a}_J \oplus \mathbf{s}_J}} P_{\mathbf{Z}_J}^{\mathbf{s}_J} \Pi_{P_{\mathbf{Z}_I, \mathbf{a}_I}} |\varphi\rangle_{\mathcal{B}}\|^2. \quad (29)$$

This equation can be interpreted operationally as follows: the probability of obtaining an outcome $(\mathbf{a}_I, \mathbf{a}_J)$ by measuring the extracted state is *equal* to the probability of obtaining this outcome by the following procedure acting on $|\varphi\rangle$:

1. First, measure the observables P_{Z_i} for every $i \in I$ on $|\varphi\rangle$, obtaining an outcome \mathbf{a}_I .

2. Next, sample a string $\mathbf{s}_J \in \{0, 1\}^{|J|}$ uniformly at random and apply $P_{\mathbf{Z}_J}^{\mathbf{s}_J}$ to the state.
3. Next, measure the observables P_{X_j} for every $j \in J$ on the state, obtaining an outcome \mathbf{u}_J .
4. Set $\mathbf{a}_J = \mathbf{u}_J \oplus \mathbf{s}_J$ and return $(\mathbf{a}_I, \mathbf{a}_J)$.

The proof of Equation (24) We first define a new distribution, which we denote by $\widehat{\text{Real}}^{\text{C}^*. \text{Open}_{[\ell]}}(\lambda, \mathbf{b}, \sigma)$. This distribution is identical to $\text{Real}^{\text{C}^*. \text{Open}_{[\ell]}}(\lambda, \mathbf{b}, \sigma)$ except that it does not run the `Ver` algorithm (i.e., it does not run Item 2 of the definition of `Real` in Definition 5.6), and simply sets $\mathbf{m} = \text{Out}(\text{sk}, \mathbf{y}, \mathbf{b}, \mathbf{z})$. We note that by Lemma 3.2,

$$\widehat{\text{Real}}^{\text{C}^*. \text{Open}_{[\ell]}}(\lambda, \mathbf{b}, \sigma) \stackrel{\sqrt{2\delta}}{\approx} \text{Real}^{\text{C}^*. \text{Open}_{[\ell]}}(\lambda, \mathbf{b}, \sigma) \quad (30)$$

where recall 2δ is the probability that `Ver` rejects $\text{C}^*. \text{Open}_{[\ell]}(\lambda, \mathbf{b}, \sigma)$ (see Equation (22)). Therefore to prove Equation (24) it suffices to prove that

$$\widehat{\text{Real}}^{\text{C}^*. \text{Open}_{[\ell]}}(\lambda, \mathbf{b}, \sigma) \stackrel{2\sqrt{\delta}}{\approx} \text{Ideal}^{\text{Ext}, \text{C}^*. \text{Open}}(\lambda, \mathbf{b}, \sigma) \quad (31)$$

To this end, we first claim that

$$(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_{\text{Sim}, I}) \equiv (\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_{\widehat{\text{Real}}, I})$$

where $(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_{\text{Sim}, I})$ is distributed by generating

$$(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}) \leftarrow \text{Ideal}^{\text{Ext}, \text{C}^*. \text{Open}}(\lambda, \mathbf{b}, \sigma)$$

and outputting $(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I)$, and $(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_{\widehat{\text{Real}}, I})$ is distributed by generating

$$(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}) \leftarrow \widehat{\text{Real}}^{\text{C}^*. \text{Open}_{[\ell]}}(\lambda, \mathbf{b}, \sigma)$$

and outputting $(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I)$. To see why this is true, recall that Equation (26) implies that for a given $\text{pk}, \mathbf{y}, \mathbf{b}$, the outcome $\mathbf{m}_{\text{Sim}, I}$, which is equal to \mathbf{a}_I in the notation used in that equation, is distributed according to the outcome of measuring P_{Z_i} on the qubits $i \in I$ qubits of the post-commitment state. Moreover, P_{Z_i} is defined so that it exactly matches the action of $\widehat{\text{Real}}$ since both do not run `Ver`.

Remark 7.17. We note that the observable P_{Z_i} was defined with respect to the opening algorithm $\text{C}^*. \text{Open}$ and we are considering $\widehat{\text{Real}}$ with respect to the opening algorithm $\text{C}^*. \text{Open}_{[\ell]}$. The observable P_{Z_i} corresponding to $\text{C}^*. \text{Open}$, when viewed as a unitary, is different from observable corresponding to $\text{C}^*. \text{Open}_{[\ell]}$, denoted by P'_{Z_i} , when viewed as a unitary. In particular, recall that

$$P_{Z_i} = U_0^\dagger \text{Out}_{i,0}^\dagger Z_{\text{out}_i} \text{Out}_{i,0} U_0$$

whereas

$$P'_{Z_i} = U_0^\dagger U_{\text{post}}^\dagger \text{Out}_{i,0}^\dagger Z_{\text{out}_i} \text{Out}_{i,0} U_{\text{post}} U_0,$$

where U_{post} is the unitary that does some post-processing to the `openi` register to ensure that measuring it will not disturb the state in a detectable way. Despite the fact that P_{Z_i} and P'_{Z_i} are different unitaries, on the subspace where the ancilla registers are initialized to $|0\rangle$, they are identical operators. In particular, P'_{Z_i} preserves the subspace where the ancilla registers are initialized to $|0\rangle$.

To avoid cluttering of notation, from now on we denote $\mathbf{m}_{\widehat{\text{Real}},I}$ and $\mathbf{m}_{\text{Sim},I}$ by \mathbf{m}_I . Denote by

$$\rho'_I = \Pi_{P_{\mathbf{z}_I, \mathbf{m}_I}}[\rho]$$

where ρ is post-commitment state and \mathbf{m}_I is distributed as $\mathbf{m}_{\text{Sim},I}$. We note in the experiment $\widehat{\text{Real}}^{\text{C}^*. \text{Open}_{[\ell]}}(\lambda, \mathbf{b}, \sigma)$, the post state after measuring \mathbf{m}_I is ρ'_I . This follows from Remark 7.17. We prove that

$$(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_{\widehat{\text{Real}},J}) \stackrel{2\sqrt{\delta}}{\approx} (\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_{\text{Sim},J}) \quad (32)$$

where $\mathbf{m}_{\widehat{\text{Real}},J}$ is obtained as follows:

1. Compute $\mathbf{z}_J \leftarrow \text{C}^*. \text{Open}_{[\ell]}(\rho'_I, (J, \mathbf{b}_J))$.
2. For every $j \in J$ let $\mathbf{m}_{\widehat{\text{Real}},j} = \text{Out}(\text{sk}, \mathbf{y}, (j, 1), \mathbf{z}_j)$
3. Output $\mathbf{m}_{\widehat{\text{Real}},J} = \{\mathbf{m}_{\widehat{\text{Real}},j}\}_{j \in J}$.

To describe how $\mathbf{m}_{\text{Sim},J}$ is obtained, we take the procedure obtained immediately below Equation (29), and apply the definitions of the operational observable P_X , to obtain the following:

1. Sample at random $\mathbf{s}_J \leftarrow \{0, 1\}^{|J|}$.
2. Compute $\mathbf{z}_J \leftarrow \text{C}^*. \text{Open}_{[\ell]}(P_{Z_J}^{\mathbf{s}_J}[\rho'_I], (J, \mathbf{b}_J))$.
3. For every $j \in J$ let $u_j = \text{Out}(\text{sk}, \mathbf{y}, (j, 1), \mathbf{z}_j)$.
4. For every $j \in J$ let $\mathbf{m}_{\text{Sim},j} = u_j \oplus \mathbf{s}_j$.
5. Output $\mathbf{m}_{\text{Sim},J} = (\mathbf{m}_{\text{Sim},j})_{j \in J}$.

We prove Equation (32) separately for the non-succinct and the semi-succinct versions. For the non-succinct version we rely on the adaptive hardcore bit property and for the semi-succinct version we rely on the distributional strong adaptive hardcore bit property. In both cases, we assume without loss of generality that $\text{Open}_{[\ell]}$ opens in the Hadamard basis honestly, by measuring the relevant qubits in the standard basis. For every $j \in J$, we denote by \mathcal{O}_j the $n+1$ registers that are measured to obtain the opening of the j 'th committed qubit in the Hadamard basis.

Proof of Equation (32) in the non-succinct setting. Let $\Pi_{\text{Ver}}[\rho'_I]$ denote the state ρ'_I projected to

$$\text{Ver}(\text{sk}, \mathbf{y}, (J, 1^{|J|}), \text{Open}(\rho'_I, (J, 1^{|J|}))) = 1.$$

By Lemma 3.2,

$$\Pi_{\text{Ver}}[\rho'_I] \stackrel{\sqrt{\delta}}{\equiv} \rho'_I.$$

Therefore to prove Equation (32) it suffices to prove that

$$(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_{\widehat{\text{Real}},J}^*) \approx (\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_{\text{Sim},J}^*) \quad (33)$$

where $\mathbf{m}_{\widehat{\text{Real}},J}^*$ is distributed as $\mathbf{m}_{\widehat{\text{Real}},J}$ except that ρ'_I is replaced with $\Pi_{\text{Ver}}[\rho'_I]$. Similarly, $\mathbf{m}_{\text{Sim},J}^*$ is distributed as $\mathbf{m}_{\text{Sim},J}$ except that ρ'_I is replaced with $\Pi_{\text{Ver}}[\rho'_I]$.

To prove Equation (33) it suffices to prove that

$$(\mathbf{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \{\mathbf{m}_{j,0}\}_{j \in J}) \approx (\mathbf{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \{\mathbf{m}_{j,1} \oplus 1\}_{j \in J}) \quad (34)$$

where for every $j \in J$ and $u \in \{0, 1\}$,

$$(\mathbf{z}_{j,u}, \boldsymbol{\rho}'_{j,u}) = \mathbf{C}^*. \text{Open}(P_{Z_j}^u \Pi_{\text{Ver}}[\boldsymbol{\rho}'_I], (j, 1)) \quad \text{and} \quad \mathbf{m}_{j,u} = \text{Out}(\mathbf{sk}, \mathbf{y}, (j, 1), \mathbf{z}_{j,u}).$$

We prove that for every $j \in J$,

$$(\mathbf{pk}, \mathbf{sk}_{-(j,0)}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_{j,0}, \boldsymbol{\rho}'_{j,0}) \approx (\mathbf{pk}, \mathbf{sk}_{-(j,0)}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_{j,1} \oplus 1, \boldsymbol{\rho}'_{j,1}), \quad (35)$$

where $\mathbf{sk}_{-(j,0)}$ denotes all the secret keys except $\mathbf{sk}_{j,0}$. Namely,

$$\mathbf{sk}_{-(j,0)} \triangleq (\mathbf{sk}_{[\ell,1]}, \dots, \mathbf{sk}_{[\ell,n+1]}, \mathbf{sk}_{[\ell] \setminus \{j\}, 0}).$$

We next argue that Equation (35) implies Equation (34). To this end, we first notice that P_{Z_j} and $\mathbf{C}^*. \text{Open}(\cdot, (j, 1))$ only touch the registers corresponding to the j 'th committed qubit. This follows from our assumption that $\mathbf{C}^*. \text{Open}$ behaves honestly when opening in the Hadamard basis. This in turn implies that for every $u \in \{0, 1\}$ it holds that $\boldsymbol{\rho}'_{j,u}$ and $\Pi_{\text{Ver}}[\boldsymbol{\rho}'_I]$ are distributed identically on the registers that do not correspond to the j 'th committed qubit. Thus, Equation (35) implies that

$$\left(\mathbf{pk}, \mathbf{sk}_{-(j,0)}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_{j,0}, \Pi_{\text{Ver}}[\boldsymbol{\rho}'_I]_{\{\mathcal{O}_j\}_{j \in J \setminus \{j}\}} \right) \approx \left(\mathbf{pk}, \mathbf{sk}_{-(j,0)}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_{j,1} \oplus 1, \Pi_{\text{Ver}}[\boldsymbol{\rho}'_I]_{\{\mathcal{O}_j\}_{j \in J \setminus \{j}\}} \right). \quad (36)$$

We next note that $\mathbf{m}_{j,u}$ is a QPT function of $\Pi_{\text{Ver}}[\boldsymbol{\rho}'_I]_{\mathcal{O}_j}$ and \mathbf{sk}_j . This, together with a hybrid argument implies that indeed Equation (36) implies Equation (34), as desired.

Thus, we focus on proving Equation (35). Fix $j \in J$ and consider the mixed state

$$\boldsymbol{\rho}_{\text{mix},j} = \frac{1}{2} \Pi_{\text{Ver}}[\boldsymbol{\rho}'_I] + \frac{1}{2} P_{Z_j} \Pi_{\text{Ver}}[\boldsymbol{\rho}'_I]$$

Note that this state can be generated efficiently, with probability $1 - \delta$, from $\boldsymbol{\rho}$ given $(\mathbf{sk}_{[\ell,1]}, \dots, \mathbf{sk}_{[\ell,n+1]})$. In addition, note that $\boldsymbol{\rho}_{\text{mix},j}$ is identical to the state $\Pi_{\text{Ver}}[\boldsymbol{\rho}'_I]$ after measuring it in the P_{Z_j} basis. Recall that we assume that $\mathbf{C}^*. \text{Open}$ behaves honestly on the Hadamard basis. Thus, the $(n+1)^2$ qubits of this projected state $\boldsymbol{\rho}_{\text{mix},j}$ corresponding to the j 'th committed qubit are in superposition over $|d_1, \mathbf{x}'_1\rangle \dots |d_{n+1}, \mathbf{x}'_{n+1}\rangle$ such that for every $i \in [n+1]$, $\mathbf{y}_{j,i} = \text{Eval}(\mathbf{pk}_{j,i}, d_i, \mathbf{x}'_i)$.

Let

$$(\mathbf{z}^*, \boldsymbol{\rho}^*) \leftarrow \mathbf{C}^*. \text{Open}(\boldsymbol{\rho}_{\text{mix},j}, (j, 1)).$$

By the adaptive hardcore bit property (w.r.t. $\mathbf{pk}_{j,0}$), letting $m^* = \text{Out}(\mathbf{sk}, \mathbf{y}, (j, 1), \mathbf{z}^*)$,

$$(\mathbf{pk}, \mathbf{sk}_{-(j,0)}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, m^*, \boldsymbol{\rho}^*) \approx (\mathbf{pk}, \mathbf{sk}_{-(j,0)}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, U, \boldsymbol{\rho}^*) \quad (37)$$

where U is the uniform distribution over $\{0, 1\}$. Note that m^* is a random variable that with probability $\frac{1}{2}$ is distributed identically to $\mathbf{m}_{j,0}$ and with probability $\frac{1}{2}$ is distributed identically to $\mathbf{m}_{j,1}$. We next argue that Equation (37) implies that

$$\begin{aligned} (\mathbf{pk}, \mathbf{sk}_{-(j,0)}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_{j,0}, \boldsymbol{\rho}'_{j,0}) &\approx \\ (\mathbf{pk}, \mathbf{sk}_{-(j,0)}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_{j,1} \oplus 1, \boldsymbol{\rho}'_{j,1}), & \end{aligned}$$

as desired. To this end, suppose for contradiction that there exists a QPT algorithm A and a non-negligible $\epsilon > 0$ such that

$$\begin{aligned} & \Pr[A(\text{pk}, \text{sk}_{-(j,0)}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_{j,1} \oplus 1, \boldsymbol{\rho}'_{j,1}) = 1] - \\ & \Pr[A(\text{pk}, \text{sk}_{-(j,0)}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_{j,0}, \boldsymbol{\rho}'_{j,0}) = 1] \geq \epsilon. \end{aligned}$$

Denote by

$$p_u = \Pr[A(\text{pk}, \text{sk}_{-(j,0)}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_{j,u}, \boldsymbol{\rho}'_{j,u}) = 1]$$

and denote by

$$p'_1 = \Pr[A(\text{pk}, \text{sk}_{-(j,0)}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_{j,1} \oplus 1, \boldsymbol{\rho}'_{j,1}) = 1]$$

Note that

$$\Pr[A(\text{pk}, \text{sk}_{-(j,0)}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}^*, \boldsymbol{\rho}^*) = 1] = \frac{1}{2}p_0 + \frac{1}{2}p_1.$$

On the other hand

$$\Pr[A(\text{pk}, \text{sk}_{-(j,0)}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, U, \boldsymbol{\rho}'_{j,1}) = 1] = \frac{1}{2}p'_1 + \frac{1}{2}p_1,$$

which by the collapsing property implies that there exists a negligible function μ such that

$$\Pr[A(\text{pk}, \text{sk}_{-(j,0)}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, U, \boldsymbol{\rho}^*) = 1] = \frac{1}{2}p'_1 + \frac{1}{2}p_1 \pm \mu$$

This contradicts Equation (37) since

$$\left(\frac{1}{2}p'_1 + \frac{1}{2}p_1 \pm \mu \right) - \left(\frac{1}{2}p_0 + \frac{1}{2}p_1 \right) = \frac{1}{2}(p'_1 - p_0) - \mu \geq \frac{\epsilon}{2} \pm \mu,$$

which is non-negligible.

The proof of Equation (32) for the semi-succinct version. We prove Equation (32) holds if the underlying NTCF family satisfies the strong adaptive hardcore bit property (see Definition 3.8) and the distributional strong adaptive hardcore bit property (Definition 4.1). Let $\Pi_{\text{Ver}}[\boldsymbol{\rho}'_I]$ denote the state $\boldsymbol{\rho}'_I$ projected to

$$\text{Ver}(\text{sk}, \mathbf{y}, (J, 0^{|J|}), \text{C}^*. \text{Open}(\boldsymbol{\rho}'_I, (J, 0^{|J|}))) = 1. \text{ }^{38}$$

By Lemma 3.2,

$$\Pi_{\text{Ver}}[\boldsymbol{\rho}'_I] \stackrel{\sqrt{\delta}}{\equiv} \boldsymbol{\rho}'_I.$$

Therefore to prove Equation (32) it suffices to prove that

$$(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_{\text{Real},J}^*) \approx (\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_{\text{Sim},J}^*) \quad (38)$$

where $\mathbf{m}_{\text{Real},J}^*$ is distributed as $\mathbf{m}_{\widehat{\text{Real},J}}$ except that $\boldsymbol{\rho}'_I$ is replaced with $\Pi_{\text{Ver}}[\boldsymbol{\rho}'_I]$. Similarly, $\mathbf{m}_{\text{Sim},J}^*$ is distributed as $\mathbf{m}_{\text{Sim},J}$ except that $\boldsymbol{\rho}'_I$ is replaced with $\Pi_{\text{Ver}}[\boldsymbol{\rho}'_I]$.

³⁸Note that this is different from the non-succinct case where we used $\Pi_{\text{Ver}}[\boldsymbol{\rho}'_I]$ to denote the state $\boldsymbol{\rho}'_I$ projected to $\text{Ver}(\text{sk}, \mathbf{y}, (J, 1^{|J|}), \text{C}^*. \text{Open}(\boldsymbol{\rho}'_I, (J, 1^{|J|}))) = 1$.

We prove Equation (38) via a hybrid argument. Namely, denoting by $k = |J|$, we prove that for every $\alpha \in [k]$,

$$\left(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_J^{(\alpha-1)} \right) \approx \left(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_J^{(\alpha)} \right) \quad (39)$$

where $\mathbf{m}_J^{(\beta)}$ is distributed exactly as $\mathbf{m}_{\text{Sim},J}$ except that rather than choosing $\mathbf{s}_J \leftarrow \{0,1\}^{|J|}$ at random, we only choose the first β coordinates randomly and the rest we set to zero. Namely, we choose $\mathbf{s}_1, \dots, \mathbf{s}_\beta \leftarrow \{0,1\}$ and set $\mathbf{s}_{\beta+1} = \dots = \mathbf{s}_k = 0$. Equation (39) implies that

$$\left(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_J^{(0)} \right) \approx \left(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_J^{(k)} \right)$$

thus proving Equation (34) since

$$\left(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_J^{(0)} \right) \equiv \left(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \widehat{\mathbf{m}}_{\text{Real},J}^* \right) \quad \text{and} \quad \left(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_J^{(k)} \right) \equiv \left(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_{\text{Sim},J}^* \right).$$

Fix any $\alpha \in [k]$ and we prove Equation (39) for this α . Denote by $J = \{j_1, \dots, j_k\} \subseteq [\ell]$. For every $\mathbf{s}_{[\alpha-1]} = (s_1, \dots, s_{\alpha-1}) \in \{0,1\}^{\alpha-1}$ consider the states

$$\rho_{\mathbf{s}_{[\alpha-1]}} = \prod_{i=1}^{\alpha-1} P_{Z_{j_i}}^{s_i} \Pi_{\text{Ver}}[\rho'_I] \quad \text{and} \quad \rho_{\mathbf{s}_{[\alpha-1]}}^* = \frac{1}{2} \rho_{\mathbf{s}_{[\alpha-1]}} + \frac{1}{2} P_{Z_{j_\alpha}}[\rho_{\mathbf{s}_{[\alpha-1]}}]. \quad (40)$$

The main ingredient in the proof of Equation (39) is following claim.

Claim 7.18.

$$\begin{aligned} & \left(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, (\mathbf{x}_{j_i,0} \oplus \mathbf{x}_{j_i,1})_{i \in [k]}, \mathbf{d}_{J \setminus \{j_\alpha\}}^*, \mathbf{d}_{j_\alpha}^* \cdot (1, \mathbf{x}_{j_\alpha,0} \oplus \mathbf{x}_{j_\alpha,1}) \right) \approx \\ & \left(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, (\mathbf{x}_{j_i,0} \oplus \mathbf{x}_{j_i,1})_{i \in [k]}, \mathbf{d}_{J \setminus \{j_\alpha\}}^*, U \right) \end{aligned} \quad (41)$$

where \mathbf{d}_J^* is distributed by measuring the registers $\mathcal{O}_{j_1}, \dots, \mathcal{O}_{j_k}$ of $\rho_{\mathbf{s}_{[\alpha-1]}}^*$ in the standard basis.

In what follows, we use Claim 7.18 to prove Equation (39) and then we prove Claim 7.18. To this end, for every $u \in \{0,1\}$, denote by

$$\mathbf{d}_J^{(u)} = \left(\mathbf{d}_{j_1}^{(u)}, \dots, \mathbf{d}_{j_k}^{(u)} \right)$$

the values obtained by measuring registers $\mathcal{O}_{j_1}, \dots, \mathcal{O}_{j_k}$ of $P_{Z_{j_\alpha}}^{(u)}[\rho_{\mathbf{s}_{[\alpha-1]}}]$ in the standard basis. For every $i \in [k]$, denote by

$$m_{j_i}^{(u)} = \mathbf{d}_{j_i}^{(u)} \cdot (1, \mathbf{x}_{j_i,0} \oplus \mathbf{x}_{j_i,0}) \quad \text{and} \quad m_{j_i}^* = \mathbf{d}_{j_i}^* \cdot (1, \mathbf{x}_{j_i,0} \oplus \mathbf{x}_{j_i,0})$$

and let

$$\mathbf{m}^{(u)} = \left(m_{j_1}^{(u)}, \dots, m_{j_k}^{(u)} \right) \quad \text{and} \quad \mathbf{m}^* = \left(m_{j_1}^*, \dots, m_{j_k}^* \right).$$

Fix any QPT adversary A . For every $u \in \{0,1\}$ denote by

$$p_u = \Pr \left[A \left(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_J^{(u)} \right) = 1 \right],$$

and denote by

$$p'_1 = \Pr \left[\mathbf{A} \left(\mathbf{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \left\{ \mathbf{m}_{J \setminus \{\alpha\}}^{(1)}, m_\alpha^{(1)} \oplus 1 \right\} \right) = 1 \right].$$

To prove Equation (39) we need to prove that

$$|p'_1 - p_0| \leq \text{negl}(\lambda). \quad (42)$$

Note that

$$\Pr \left[\mathbf{A} \left(\mathbf{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_J^* \right) = 1 \right] = \frac{1}{2}p_0 + \frac{1}{2}p_1$$

which by Claim 7.18 implies that

$$\Pr \left[\mathbf{A} \left(\mathbf{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \left\{ \mathbf{m}_{J \setminus \{\alpha\}}^*, U \right\} \right) = 1 \right] = \frac{1}{2}p_0 + \frac{1}{2}p_1 \pm \text{negl}(\lambda)$$

In addition note that

$$\Pr \left[\mathbf{A} \left(\mathbf{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \left\{ \mathbf{m}_{J \setminus \{\alpha\}}^{(1)}, U \right\} \right) = 1 \right] = \frac{1}{2}p'_1 + \frac{1}{2}p_1.$$

It remains to note that

$$\left(\mathbf{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_{J \setminus \{\alpha\}}^* \right) \equiv \left(\mathbf{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_I, \mathbf{m}_{J \setminus \{\alpha\}}^{(1)} \right),$$

which together with the equation above, implies that Equation (42) indeed holds, thus proving Equation (39), as desired.

The rest of the proof is dedicated to proving Claim 7.18, which we prove assuming the underlying NTCF family has the strong adaptive hardcore bit property (see Definition 3.8) and the distributional strong adaptive hardcore bit property (Definition 4.1).

Proof of Claim 7.18. We start by defining QPT algorithms \mathbf{A} and \mathbf{C} for the distributional strong adaptive hardcore bit property.

Algorithm A. It takes as on input \mathbf{pk}_0 generated by $\text{Gen}_{\text{NTCF}}(1^\lambda)$, and does the following:

1. For every $i \in [n+1]$ generate $(\mathbf{pk}_i, \mathbf{sk}_i) \leftarrow \text{Gen}_{\text{NTCF}}(1^\lambda)$.
2. Set $\mathbf{pk} = (\mathbf{pk}_0, \mathbf{pk}_1, \dots, \mathbf{pk}_{n+1})$.
3. Compute $(\mathbf{y}, \boldsymbol{\rho}) \leftarrow \mathbf{C}^*. \text{Commit}(\mathbf{pk}, \boldsymbol{\sigma})$.
4. Use \mathbf{sk}_1 to generate $\boldsymbol{\rho}'_I = \Pi_{P_{Z_I}, \mathbf{m}_I}(\boldsymbol{\rho})$.
5. Use $\mathbf{sk}_1, \dots, \mathbf{sk}_{n+1}$ to compute $\Pi_{\text{Ver}}[\boldsymbol{\rho}'_I]$, which is the state $\boldsymbol{\rho}'_I$ projected to

$$\text{Ver} \left((\mathbf{sk}_1, \dots, \mathbf{sk}_{n+1}), \mathbf{y}, \left(J, 0^{|J|} \right), \mathbf{C}^*. \text{Open} \left(\boldsymbol{\rho}'_I, \left(J, 0^{|J|} \right) \right) \right) = 1.$$

This step fails with probability δ , in which case \mathbf{A} outputs \perp .

6. Use \mathbf{sk}_1 to compute $\boldsymbol{\rho}_{\mathbf{s}_{[\alpha-1]}} = \prod_{i=1}^{\alpha-1} P_{Z_{j_i}}^{s_i} \Pi_{\text{Ver}}[\boldsymbol{\rho}'_I]$.
7. Use \mathbf{sk}_1 to compute $\boldsymbol{\rho}_{\mathbf{s}_{[\alpha-1]}}^* = \frac{1}{2}\boldsymbol{\rho}_{\mathbf{s}_{[\alpha-1]}} + \frac{1}{2}P_{Z_\alpha}[\boldsymbol{\rho}_{\mathbf{s}_{[\alpha-1]}}]$.

8. Parse $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_\ell)$ and parse $\mathbf{y}_{j_\alpha} = (\mathbf{y}_{j_\alpha,0}, \mathbf{y}_{j_\alpha,1}, \dots, \mathbf{y}_{j_\alpha,n+1})$
9. Measure the $P_{Z_{j_\alpha}}$ observable of the state $\rho_{\mathbf{s}_{[\alpha-1]}}^*$ to obtain a preimage $(b, \mathbf{x}_{j_\alpha,b})$ of $\mathbf{y}_{j_\alpha,0}$. Denote the post-measurement state by ρ^{**} . We assume without loss of generality that this state includes $(\mathbf{pk}, (\mathbf{sk}_1, \dots, \mathbf{sk}_{n+1}), \mathbf{y}, \mathbf{m}_I)$.
10. Output $(\mathbf{y}_{j_\alpha,0}, b, \mathbf{x}_{j_\alpha,b}, \rho^{**})$.
We rename the register \mathcal{O}_{j_α} of ρ^{**} by \mathcal{O}_1 and rename all the other registers by \mathcal{O}_2 .

Algorithm C. It takes as input the secret vector $\mathbf{sk}_{0,\text{pre}}$, corresponding to \mathbf{sk}_0 , and the \mathcal{O}_2 registers of ρ^{**} , denoted by $\rho_{\mathcal{O}_2}^{**}$, and does the following:

1. For every $i \in [k] \setminus \{\alpha\}$ use $\mathbf{sk}_1, \dots, \mathbf{sk}_{n+1}$ to coherently compute a preimage of $\mathbf{y}_{i,0}$ as follows:
 - Coherently compute $\mathbf{C}^*. \text{Open}(\rho^{**}, (j_i, 0))$.
 - Coherently run $\text{Ver}((\mathbf{sk}_1, \dots, \mathbf{sk}_{n+1}), \mathbf{y}_{j_i}, (j_i, 0), \cdot)$ to coherently generate a preimage.

Note that since the state $\Pi_{\text{Ver}}[\rho'_I]$ is the state ρ'_I projected to the state where $\mathbf{C}^*. \text{Open}$ is always accepted when the opening is in the standard basis, we indeed get a coherent preimage of $\mathbf{y}_{j_i,0}$ with probability 1.

2. Use $\mathbf{sk}_{0,\text{pre}}$ to compute (and measure) $\mathbf{x}_{j_i,0} \oplus \mathbf{x}_{j_i,1}$.
Note that this is a deterministic quantity and hence does not disturb the state.
3. Measure all the registers in \mathcal{O}_2 corresponding to $\{\mathcal{O}_{j_i}\}_{i \in [k] \setminus \{\alpha\}}$ to obtain $\mathbf{d}_{J \setminus \{j_\alpha\}}^*$.
4. Output $\text{aux} = (\mathbf{pk}, \mathbf{y}, \mathbf{m}_I, (x_{j_i,0} \oplus x_{j_i,1})_{i \in [k] \setminus \{\alpha\}}, \mathbf{d}_{J \setminus \{j_\alpha\}}^*)$.
Denote the residual state by ρ_{aux}^{**} .

To finish the proof of Claim 7.18 it remains to prove that $\mathbf{d}_{j_\alpha}^*$ satisfies the desired min-entropy requirement. This is captured in the claim below.

Claim 7.19. *With overwhelming probability over $\text{aux} \leftarrow C(\mathbf{sk}_{0,\text{pre}}, \rho_{\mathcal{O}_2}^{**})$ it holds that for every $\mathbf{d}' \in \{0, 1\}^n$,*

$$\Pr[\text{Good}(\mathbf{x}_{j_\alpha,0}, \mathbf{x}_{j_\alpha,1}, \mathbf{d}_{j_\alpha}^*) = \mathbf{d}'] = \text{negl}(\lambda).$$

where $\mathbf{d}_{j_\alpha}^*$ is obtained by measuring the \mathcal{O}_1 registers of ρ_{aux}^{**} in the standard basis.

The proof of Claim 7.19 relies on the adaptive hardcore bit property of the underlying NTCF family. It also makes use of the following fact which follows immediately from the definition of P_{Z_j} .³⁹

Fact 7.20. *For every $\alpha \in [k]$ and every $\mathbf{s}_{[\alpha-1]} \in \{0, 1\}^{\alpha-1}$, the state $\rho_{\mathbf{s}_{[\alpha-1]}}$ and the state $\rho_{\mathbf{s}_{[\alpha-1]}}^*$ can be efficiently constructed from $(\mathbf{pk}, \mathbf{y}, \rho)$ and \mathbf{sk}_1 .*

³⁹Recall that P_{Z_j} does not check if $\text{Ver}(\mathbf{sk}, \mathbf{y}, (j, 0), \cdot) = 1$.

Proof of Claim 7.19. Suppose for the sake of contradiction that there is a non-negligible $\epsilon = \epsilon(\lambda)$ such that with probability ϵ over $(\mathbf{aux}, \boldsymbol{\rho}_{\mathbf{aux}}^{**}) \leftarrow C(\mathbf{sk}_{0,\text{pre}}, \boldsymbol{\rho}_{\mathcal{O}_2}^{**})$ it holds that there exists a vector $\mathbf{d}' = \mathbf{d}'_{\mathbf{aux}} \in \{0, 1\}^n$ such that

$$\Pr[\text{Good}(\mathbf{x}_{j_\alpha,0}, \mathbf{x}_{j_\alpha,1}, \mathbf{d}_{j_\alpha}^*) = \mathbf{d}'] \geq \epsilon \quad (43)$$

where $\mathbf{d}_{j_\alpha}^*$ is obtained by measuring the \mathcal{O}_1 registers of $\boldsymbol{\rho}_{\mathbf{aux}}^{**}$ in the standard basis. Denote the set of all \mathbf{aux} that satisfy Equation (43) by BAD. In the rest of this proof, for the sake of ease of notation, we denote by

$$\mathbf{d}^* \triangleq \mathbf{d}_{j_\alpha}^* \quad \text{and} \quad (\mathbf{x}_0, \mathbf{x}_1) \triangleq (\mathbf{x}_{j_\alpha,0}, \mathbf{x}_{j_\alpha,1}).$$

We next denote by

$$\mathbf{d}^* = (d_0^*, d_1^*, \dots, d_n^*)$$

and we argue that there exists a subset

$$S = \{\beta_1, \dots, \beta_m\} \subseteq [n] \quad (44)$$

of size $m = n^{0.1}$ for which there exists an (all powerful) algorithm B such that for every $\mathbf{aux} \in \text{BAD}$ and for every $i \in [m]$

$$\Pr\left[\mathbf{B}\left(\mathbf{aux}, \left(d_1^*, \dots, d_{j_{\beta_i}'}^*\right)\right) = d_{j_{\beta_i}'}^*\right] \geq 1 - \frac{1}{\sqrt{n}} \quad (45)$$

where the probability is over \mathbf{d}^* obtained by measuring the \mathcal{O}_1 registers of $\boldsymbol{\rho}_{\mathbf{aux}}^{**}$ in the standard basis.

The existence of such a set S follows from Equation (43), since otherwise for every \mathbf{d}' ,

$$\Pr[\text{Good}(\mathbf{x}_0, \mathbf{x}_1, \mathbf{d}^*) = \mathbf{d}'] \leq \left(1 - \frac{1}{\sqrt{n}}\right)^{n-m} = \text{negl}(\lambda),$$

contradicting Equation (43).

We construct QPT algorithm A that on input $(\mathbf{pk}_2, \dots, \mathbf{pk}_{n+1})$ generates with non-negligible probability images $\{\mathbf{y}_i^*\}_{i \in S}$ corresponding to $\{\mathbf{pk}_{i+1}\}_{i \in S}$, preimages $\{(d_i^*, \mathbf{x}_i^*)\}_{i \in S}$, and equations $\{\mathbf{z}_i, \mathbf{z}_i \cdot (1, \mathbf{x}_{i,0} \oplus \mathbf{x}_{i,1})\}_{i \in S}$, such that each $\mathbf{z}_i \in \text{Good}_{\text{Invert}(\mathbf{sk}_i, \mathbf{y}_i^*)}$ and $(\mathbf{x}_{i,0}, \mathbf{x}_{i,1}) = \text{Invert}(\mathbf{sk}_i, \mathbf{y}_i)$, thus breaking the adaptive hardcore bit property. Algorithm A $(\mathbf{pk}_2, \dots, \mathbf{pk}_{n+1})$ does the following:

1. Sample $(\mathbf{pk}_0, \mathbf{sk}_0), (\mathbf{pk}_1, \mathbf{sk}_1) \leftarrow \text{Gen}(1^\lambda)$.
2. Let $\mathbf{pk} = (\mathbf{pk}_0, \mathbf{pk}_1, \mathbf{pk}_2, \dots, \mathbf{pk}_{n+1})$.
3. Generate $(\mathbf{y}, \boldsymbol{\rho}) \leftarrow C^*. \text{Commit}(\mathbf{pk}, \boldsymbol{\sigma})$.
4. Parse $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_\ell)$ and parse $\mathbf{y}_{j_\alpha} = (\mathbf{y}_0^*, \mathbf{y}_1^*, \dots, \mathbf{y}_{n+1}^*)$.
5. Compute $((0, \mathbf{x}_0), (1, \mathbf{x}_1)) \leftarrow \text{Invert}(\mathbf{sk}_0, \mathbf{y}_0^*)$.
6. Use \mathbf{sk}_1 to generate $\boldsymbol{\rho}'_I = \Pi_{P_{Z_I}, \mathbf{m}_I}(\boldsymbol{\rho})$.
7. Use \mathbf{sk}_1 to compute $\boldsymbol{\rho}_{\mathbf{s}_{[\alpha-1]}} = \prod_{i=1}^{\alpha-1} P_{Z_{j_i}}^{s_i}[\boldsymbol{\rho}'_I]$.
8. Use \mathbf{sk}_1 to compute $\boldsymbol{\rho}_{\mathbf{s}_{[\alpha-1]}}^* = \frac{1}{2}\boldsymbol{\rho}_{\mathbf{s}_{[\alpha-1]}} + \frac{1}{2}P_{Z_\alpha}[\boldsymbol{\rho}_{\mathbf{s}_{[\alpha-1]}}]$.

9. Measure registers \mathcal{O}_{j_α} of $\rho_{S_{[\alpha-1]}}^*$ in the standard basis to obtain $\mathbf{d}^* = (d_0^*, d_1^*, \dots, d_n^*)$.
10. For every $i \in S$ measure the registers corresponding to \mathbf{x}_i^* . We note that with probability $1 - \delta$ it holds that for every $i \in S$,

$$\mathbf{y}_{i+1}^* = \text{Eval}(\text{pk}_{i+1}, d_i^*, \mathbf{x}_i^*).$$

Denote the resulting state by ρ' .

11. Compute $(\mathbf{z}, \rho'') \leftarrow \text{C}^*. \text{Open}(\text{pk}, (j_\alpha, 0), \rho')$.
12. Parse $\mathbf{z} = (\mathbf{z}_1, \dots, \mathbf{z}_{n+1}) \in \{0, 1\}^{(n+1)^2}$.
13. Compute $(0, \mathbf{x}'_{1,0}), (1, \mathbf{x}'_{1,1}) \leftarrow \text{Invert}(\text{sk}_1, \mathbf{y}_1^*)$.
14. Let $\gamma = \mathbf{z}_1 \cdot (1, \mathbf{x}'_{1,0} \oplus \mathbf{x}'_{1,1})$.
15. Output $\{\text{pk}_{i+1}, \mathbf{y}_{i+1}^*, d_i^*, \mathbf{x}_i^*, \mathbf{z}_{i+1}, \mathbf{x}_{\gamma,i}\}_{i \in S}$.

We next argue that with non-negligible probability it holds that for every $i \in S$:

$$\mathbf{y}_{i+1}^* = \text{Eval}(\text{pk}_{i+1}, (d_i^*, \mathbf{x}_i^*)) \wedge \mathbf{z}_{i+1} \cdot (1, \mathbf{x}'_{i,0} \oplus \mathbf{x}'_{i,1}) = \mathbf{x}_{\gamma,i} \wedge \mathbf{z}_{i+1} \in \text{Good}_{\mathbf{x}'_{i,0}, \mathbf{x}'_{i,1}},$$

where $((0, \mathbf{x}'_{i,0}), (1, \mathbf{x}'_{i,1})) = \text{Invert}(\text{sk}_{i+1}, \mathbf{y}_{i+1}^*)$, contradicting the adaptive hardcore bit property.

The fact that the underlying $\text{C}^*. \text{Open}$ algorithm is accepted by Ver with probability $\geq 1 - \delta$ implies that

$$\Pr[\forall i \in S : \mathbf{y}_{i+1}^* = \text{Eval}(\text{pk}_{i+1}, (d_i^*, \mathbf{x}_i^*))] \geq 1 - \delta.$$

Moreover, if we did not measure $\{(d_i^*, \mathbf{x}_i^*)\}_{i \in [n]}$ it would also imply that,

$$\Pr[\forall i \in S : \mathbf{z}_{i+1} \cdot (1, \mathbf{x}'_{i,0} \oplus \mathbf{x}'_{i,1}) = \mathbf{x}_{\gamma,i} \wedge \mathbf{z}_{i+1} \in \text{Good}_{\mathbf{x}'_{i,0}, \mathbf{x}'_{i,1}}] \geq 1 - \delta$$

By the collapsing property applied to $\{\text{pk}_i\}_{i \in \{2, \dots, n+1\} \setminus \{i+1 : i \in S\}}$, even if we measured $\{(d_i^*, \mathbf{x}_i^*)\}_{i \in [n] \setminus S}$,

$$\Pr[\forall i \in S : \mathbf{z}_{i+1} \cdot (1, \mathbf{x}'_{i,0} \oplus \mathbf{x}'_{i,1}) = \mathbf{x}_{\gamma,i} \wedge \mathbf{z}_i \in \text{Good}_{\mathbf{x}'_{i,0}, \mathbf{x}'_{i,1}}] \geq 1 - 2\delta - \text{negl}(\lambda)$$

We note however, that by Equation (45), for every $i \in S$ measuring (d_i^*, \mathbf{x}_i^*) disturbs the state by at most $n^{-1/4}$. Hence

$$\Pr[\forall i \in S : \mathbf{y}_{i+1}^* = \text{Eval}(\text{pk}_{i+1}, (d_i^*, \mathbf{x}_i^*)) \wedge \mathbf{z}_{i+1} \cdot (1, \mathbf{x}'_{i,0} \oplus \mathbf{x}'_{i,1}) = \mathbf{x}_{\gamma,i} \wedge \mathbf{z}_{i+1} \in \text{Good}_{\mathbf{x}'_{i,0}, \mathbf{x}'_{i,1}}] \geq 1 - 2\delta - \frac{n^{0.1}}{n^{1/4}} - \text{negl}(\lambda),$$

contradicting the adaptive hardcore bit property. □

□

Remark 7.21. We next argue that relying on some form of the adaptive hardcore bit is necessary. Specifically, if there exists a QPT algorithm A and a non-negligible function $\epsilon = \epsilon(\lambda)$ such that

$$\Pr[A(\mathbf{pk}_1, \dots, \mathbf{pk}_{n+1}) = (b_i, \mathbf{x}_i, \mathbf{d}_i, m_i)_{i=1}^{n+1} : \forall i \in [n+1] \quad \mathbf{d}_i \cdot (1, \mathbf{x}_{i,0} \oplus \mathbf{x}_{i,1}) = m_i] \geq \epsilon(\lambda)$$

where $(\mathbf{x}_{i,0}, \mathbf{x}_{i,1}) = \text{Invert}(\text{sk}_i, \text{Eval}(\mathbf{pk}_i, (b_i, \mathbf{x}_i)))$, then one can use this adversary A to attack the scheme, as follows:

1. Given $(\mathbf{pk}_0, \mathbf{pk}_1, \dots, \mathbf{pk}_{n+1})$, compute

$$(b_i, \mathbf{x}_i, \mathbf{d}_i, m_i)_{i=1}^{n+1} = A(\mathbf{pk}_1, \dots, \mathbf{pk}_{n+1}).$$

2. Let $\mathbf{m} = (m_1, \dots, m_{n+1})$ and set $\mathbf{y}_0 = \text{Eval}(\mathbf{pk}_0, \mathbf{m})$.

3. For every $i \in [n+1]$, set $\mathbf{y}_i = \text{Eval}(\mathbf{pk}_i, (b_i, \mathbf{x}_i))$.

4. Output $(\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_{n+1})$ as the commitment.

5. If asked to open in the Hadamard basis output $((b_1, \mathbf{x}_1), \dots, (b_{n+1}, \mathbf{x}_{n+1}))$.

6. If asked to open in the Standard basis output $(\mathbf{d}_1, \dots, \mathbf{d}_{n+1})$.

The adversary is accepted with probability ϵ and the openings are distinguishable from a qubit, since the standard basis opening is deterministic and the Hadamard opening is biased, both in a detectable way.

7.3 Binding for the Succinct Commitment Scheme

In this section we prove the following theorem.

Theorem 7.22. *The succinct multi-qubit commitment scheme described in Section 6.3 satisfies the binding condition defined in Definition 5.10.*

Proof. To prove soundness we need to prove that Equations (14) and (15) hold. To this end, fix any QPT algorithm $C^*.Commit$, a quantum state σ , a polynomial $\ell = \ell(\lambda)$, a QPT prover P^* for $Ver.Commit$. We start by defining a QPT algorithm $C_{ss}^*.Commit$ for the underlying semi-succinct commitment scheme. This algorithm is associated with a parameter ϵ_0 , and sometime to be explicit, we denote it by $C_{ss}^*.Commit_{\epsilon_0}$. It takes as input (\mathbf{pk}_1, σ) and commits to an ℓ -qubit state, as follows:

1. Sample $hk \leftarrow \text{Gen}_H(1^\lambda)$.

2. Set $\mathbf{pk} = (\mathbf{pk}_1, hk)$.

3. Compute $(rt, \rho) \leftarrow C^*.Commit(\mathbf{pk}, \sigma)$.

4. Use the state-preservation extractor \mathcal{E} (from Definition 6.5) for the NP language \mathcal{L}^* (defined in Equation (19)) to generate

$$(\mathbb{T}_{\text{Sim}}, \mathcal{Y}, \rho_{\text{post,Sim}}) \leftarrow \mathcal{E}^{P^*, \rho} \left((hk, rt), 1^\lambda, \epsilon_0 \right)$$

5. If $\text{Eval}(hk, \mathcal{Y}) \neq rt$ then output \perp .

6. Else, output $(\mathbf{y}, \rho_{\text{post,Sim}})$.

By Definition 6.5 and Corollary 6.6,

$$[\Pr[\text{Eval}_H(\text{hk}, \mathbf{y}) = \text{rt}] \geq 1 - \delta_0 - 2\epsilon_0 - \text{negl}(\lambda)]. \quad (46)$$

We use this QPT algorithm $C_{\text{ss}}^*. \text{Commit}_{\epsilon_0}$ to prove the soundness of the succinct scheme. We start with proving Equation (15). To this end, fix a subset $J \subseteq [\ell]$ and a basis $\mathbf{b}_J \in \{0, 1\}^{|J|}$, and two QPT algorithms $C_1^*. \text{Open}$ and $C_2^*. \text{Open}$. We need to prove that

$$\text{Real}^{C^*. \text{Commit}, P^*, C_1^*. \text{Open}}(\lambda, (J, \mathbf{b}_J), \sigma) \stackrel{\eta}{\approx} \text{Real}^{C^*. \text{Commit}, P^*, C_2^*. \text{Open}}(\lambda, (J, \mathbf{b}_J), \sigma) \quad (47)$$

where $\eta = O(\sqrt{\delta_0 + \delta})$, where δ_0 is defined in Equation (17) and δ is defined in Equation (16).

We next define two QPT opening algorithms $C_{\text{ss},1}^*. \text{Open}$ and $C_{\text{ss},2}^*. \text{Open}$ for the underlying semi-succinct commitment scheme, corresponding to $C_1^*. \text{Open}$ and $C_2^*. \text{Open}$, respectively. For every $i \in \{1, 2\}$, $C_{\text{ss},i}^*. \text{Open}(\rho_{\text{post,Sim}}, \mathbf{b}_J)$ does the following:⁴⁰

1. Run $(\mathbf{y}_{i,J}, \mathbf{o}_i, \mathbf{z}_{i,\text{Sim}}, \rho'_{i,\text{Sim}}) \leftarrow C_i^*. \text{Open}(\rho_{\text{post,Sim}}, (J, \mathbf{b}_J))$.
2. If $\text{Ver}_H(\text{hk}, \text{rt}, J, \mathbf{y}_{i,J}, \mathbf{o}_i) = 0$ then output \perp .
3. Else, output $(\mathbf{z}_{i,\text{Sim}}, \rho'_{i,\text{Sim}})$.

By Definition 6.5,

$$(\mathbb{T}_{\text{Sim}}, \rho_{\text{post,Sim}}, \text{sk}) \stackrel{\epsilon_0}{\approx} (\mathbb{T}, \rho_{\text{post}}, \text{sk}) \quad (48)$$

which implies that for every $i \in \{1, 2\}$,

$$\Pr[\text{Ver}(\text{sk}, \text{rt}, (J, \mathbf{b}_J), \mathbf{y}_{i,J}, \mathbf{o}_i, \mathbf{z}_{i,\text{Sim}}) = 0] \leq \delta + \epsilon_0. \quad (49)$$

For every $j \in J$ let

$$\mathbf{m}_{i,\text{Sim},j} = \text{Out}_1(\text{sk}_1, \mathbf{y}_j, b_j, \mathbf{z}_{i,\text{Sim},j}) \quad \text{and} \quad \mathbf{m}_{i,j} = \text{Out}_1(\text{sk}_1, \mathbf{y}_j, b_j, \mathbf{z}_{i,j}),$$

where $(\mathbf{z}_i, \rho'_i) \leftarrow C_i^*. \text{Open}(\rho_{\text{post}}, (J, \mathbf{b}_J))$. Let $\mathbf{m}_{i,\text{Sim},J} = (\mathbf{m}_{i,\text{Sim},j})_{j \in J}$ and $\mathbf{m}_{i,J} = (\mathbf{m}_{i,j})_{j \in J}$. Equation (48) implies that for every $i \in \{1, 2\}$,

$$(\text{pk}, \text{rt}, (J, \mathbf{b}_J), \mathbf{m}_{i,J}) \stackrel{\epsilon_0}{\approx} (\text{pk}, \text{rt}, (J, \mathbf{b}_J), \mathbf{m}_{i,\text{Sim},J}). \quad (50)$$

By the binding of the underlying semi-succinct commitment scheme,

$$\text{Real}^{C_{\text{ss}}^*. \text{Commit}, C_{\text{ss},1}^*. \text{Open}}(\lambda, \mathbf{b}_J, \sigma) \stackrel{\eta^*}{\approx} \text{Real}^{C_{\text{ss}}^*. \text{Commit}, C_{\text{ss},2}^*. \text{Open}}(\lambda, \mathbf{b}_J, \sigma). \quad (51)$$

where $\eta^* = O(\sqrt{\delta^*})$ is defined in Definition 5.6 and

$$\delta^* = \mathbb{E}_{\substack{(\text{pk}_1, \text{sk}_1) \leftarrow \text{Gen}_1(1^\lambda) \\ (\mathbf{y}, \rho) \leftarrow C_{\text{ss}}^*. \text{Commit}(\text{pk}_1, \sigma)}} \max_{\substack{i \in \{1, 2\}, \\ \mathbf{b}' \in \{\mathbf{b}_{|J|}, \mathbf{0}, \mathbf{1}\}}} \Pr[\text{Ver}_{\text{ss}}(\text{sk}, \mathbf{y}, \mathbf{b}', C_{\text{ss},i}^*. \text{Open}(\rho, \mathbf{b}')) = 0]. \quad (52)$$

We note that $\delta^* \leq \delta_0 + 3\epsilon_0 + \delta + \text{negl}(\lambda)$. This follows from Equations (46) and (49), together with the collision resistance property of the underlying hash family.

⁴⁰We assume without loss of generality that the state $\rho_{\text{post,Sim}}$ includes pk and \mathbf{y} .

We thus conclude that

$$\begin{aligned}
& \text{Real}^{\mathbf{C}^*. \text{Commit}, P^*, \mathbf{C}_1^*. \text{Open}}(\lambda, (J, \mathbf{b}_J), \boldsymbol{\sigma}) = \\
& (\text{pk}, \text{rt}, (J, \mathbf{b}_J), \mathbf{m}_{1,J}) \stackrel{\epsilon_0}{\approx} \\
& (\text{pk}, \text{rt}, (J, \mathbf{b}_J), \mathbf{m}_{1, \text{Sim}, J}) \stackrel{\eta^*}{\approx} \\
& (\text{pk}, \text{rt}, (J, \mathbf{b}_J), \mathbf{m}_{2, \text{Sim}, J}) \stackrel{\epsilon_0}{\approx} \\
& (\text{pk}, \text{rt}, (J, \mathbf{b}_J), \mathbf{m}_{2,J}) = \\
& \text{Real}^{\mathbf{C}^*. \text{Commit}, P^*, \mathbf{C}_2^*. \text{Open}}(\lambda, (J, \mathbf{b}_J), \boldsymbol{\sigma}),
\end{aligned}$$

where the second and fourth equations follow from Equation (50) and the third equation follows Equation (51). Setting $\epsilon_0 = \delta_0$ we conclude that Equation (15) holds.

It remains to prove Equation (14). To this end, we use the QPT extractor Ext_{ss} corresponding to the underlying semi-succinct commitment scheme, as well as the extractor \mathcal{E} corresponding to the underlying state-preserving argument-of-knowledge system, to construct the extractor Ext for the succinct commitment scheme. $\text{Ext}^{P^*, P_{\text{Test}}^*}(\text{sk}, \text{rt}, \boldsymbol{\rho}, 1^{\lceil 1/\epsilon \rceil})$ does the following:

1. Let $C \in \mathbb{N}$ be a constant that is larger than the constant from the definition of η in Equation (11) and in Equation (15). Namely, C is chosen so that in Equation (11) $\eta \leq C \cdot \sqrt{\delta}$, and in Equation (15) $\eta \leq C \cdot \sqrt{\delta_0 + \delta}$.
2. Set $\epsilon_0 = \left(\frac{\epsilon}{8C}\right)^2$.
3. Use the state-preservation extractor \mathcal{E} (from Definition 6.5) for the NP language \mathcal{L}^* (defined in Equation (19)) to generate

$$(\mathbb{T}_{\text{Sim}}, \mathbf{y}, \boldsymbol{\rho}_{\text{post}, \text{Sim}}) \leftarrow \mathcal{E}^{P^* \cdot \rho} \left((\text{hk}, \text{rt}), 1^\lambda, \epsilon_0 \right) \quad (53)$$

4. If $\text{Eval}_{\text{H}}(\text{hk}, \mathbf{y}) \neq \text{rt}$ then set $\mathbf{y} = \perp$.
5. Use P_{Test}^* to define $\mathbf{C}_{\text{ss}}^*. \text{Open}$, which is associated with a parameter ϵ_0 , and on input $(\boldsymbol{\rho}_{\text{post}, \text{Sim}}, (j, b))$,⁴¹ operates as follows:⁴²

- (a) Denote by U_b the unitary that does the following computation (coherently, using ancilla registers):
 - i. Compute the first message of $P_{\text{Test}}^*(\text{pk}, \text{rt}, \boldsymbol{\rho}_{\text{post}, \text{Sim}})$ upon receiving the bit $b \in \{0, 1\}$ from V_{Test} , to obtain rt' and a post state $\boldsymbol{\rho}'_{\text{Sim}}$.
 - ii. Use the state-preservation extractor \mathcal{E} (from Definition 6.5) for the NP language \mathcal{L}^* (defined in Equation (19)) to generate

$$(\mathbb{T}_{\text{Sim}}, \mathbf{z}, \boldsymbol{\rho}'_{\text{post}, \text{Sim}}) \leftarrow \mathcal{E}^{P_{\text{Test}}^* \cdot \rho'_{\text{Sim}}} \left((\text{hk}, \text{rt}'), 1^\lambda, \epsilon_0 \right) \quad (54)$$

- iii. Denote the ancilla registers where \mathbf{z} is stored by $(\text{open}_1, \dots, \text{open}_\ell)$.

⁴¹We assume without loss of generality that the state $\boldsymbol{\rho}_{\text{post}, \text{Sim}}$ includes $(\text{pk}, \mathbf{y}, \text{rt})$.

⁴² $\mathbf{C}_{\text{ss}}^*. \text{Open}$ is defined somewhat analogously to $\mathbf{C}^*. \text{Open}_{[\ell]}$ as defined in the proof of Lemmas 7.13 and 7.14.

- iv. If $b = 0$ (corresponding to a standard basis measurement), apply a post-processing unitary to each open_i register, to ensure that measuring this register would not disturb the state in a detectable way. This is done as in Remark 7.16. Specifically, Denoting $\text{sk}_1 = (\text{sk}_{1,0}, \text{sk}_{1,1}, \dots, \text{sk}_{1,n+1})$, the unitary U_0 uses $(\text{sk}_{1,1}, \dots, \text{sk}_{1,n+1})$ to apply the following post-processing unitary to each open_i register, to ensure that when measured the disturbance will not be noticed to a QPT algorithm which is not given $(\text{sk}_{1,1}, \dots, \text{sk}_{1,n+1})$. Recall that open_i contains a vector $\mathbf{z} = (\mathbf{z}_1, \dots, \mathbf{z}_{n+1}) \in \{0, 1\}^{(n+1)^2}$ where each $\mathbf{z}_j \in \{0, 1\}^{n+1}$. The post-processing unitary does the following:
- Coherently compute for every $j \in [n+1]$ the bit $m_j = \mathbf{z}_j \cdot (1, \mathbf{x}'_{j,0} \oplus \mathbf{x}'_{j,1})$, where $\mathbf{x}'_{j,0}$ and $\mathbf{x}'_{j,1}$ are the two preimages of $\mathbf{y}_{i,j}$ that are computed using sk_j .
 - Let $\mathbf{m} = (m_1, \dots, m_{n+1}) \in \{0, 1\}^{n+1}$. Note that if \mathbf{z} is a successful opening (i.e., it is accepted) then \mathbf{m} is a preimage of $\mathbf{y}_{i,0}$, and whether a preimage is measured or not is undetectable without knowing sk_0 , due to the collapsing property of the underlying NTCF family.
 - On an ancilla register, compute a super-position over all $\mathbf{z}' = (\mathbf{z}'_1, \dots, \mathbf{z}'_{n+1}) \in \{0, 1\}^{(n+1)^2}$ such that for every $j \in [n+1]$ $m_j = \mathbf{z}'_j \cdot (1, \mathbf{x}'_{j,0} \oplus \mathbf{x}'_{j,1})$.
 - Swap register open_i with the ancilla register above, so that now $\mathbf{z}' = (\mathbf{z}'_1, \dots, \mathbf{z}'_{n+1})$ is in register open_i .

- Compute $\rho' = U_b^\dagger \text{CNOT}_{\text{open}_j, \text{copy}_j} U_b[\rho_{\text{post, Sim}}]$
- Measure register copy_j in the standard basis to obtain \mathbf{z}_j .
- Output \mathbf{z}_j .

So far, we defined $\text{C}_{\text{ss}}^*. \text{Open}$ on a single coordinate (j, b) . We define $\text{C}_{\text{ss}}^*. \text{Open}$ on a set of coordinates (J, \mathbf{b}_J) to first apply $\text{C}_{\text{ss}}^*. \text{Open}$ on all the coordinates $j \in J$ such that $\mathbf{b}_j = 0$ (in order) and then apply it on all the coordinates $j \in J$ such that $\mathbf{b}_j = 1$ (in order).⁴³

6. Output $\tau_{A,B} \leftarrow \text{Ext}_{\text{ss}}^{\text{C}_{\text{ss}}^*. \text{Open}}(\text{sk}, \mathbf{y}, \rho_{\text{post, Sim}})$.

We need to argue that for every QPT algorithm $\text{C}^*. \text{Open}$,

$$\text{Real}^{\text{C}^*. \text{Commit}, P^*, \text{C}^*. \text{Open}}(\lambda, (J, \mathbf{b}_J), \sigma) \stackrel{\zeta}{\approx} \text{Ideal}^{\text{Ext}, \text{C}^*. \text{Commit}, P^*, P_{\text{Test}}^*}(\lambda, (J, \mathbf{b}_J), \sigma, \epsilon), \quad (55)$$

for

$$\zeta = O\left(\sqrt{\delta_0 + \delta'_0 + \delta}\right) + \epsilon.$$

To this end, we rely on the binding property of the underlying semi-succinct scheme (and in particular Equation (11)), which implies that

$$\text{Real}^{\text{C}_{\text{ss}}^*. \text{Commit}, \text{C}_{\text{ss}}^*. \text{Open}}(\lambda, (J, \mathbf{b}_J), \sigma) \stackrel{\eta^*}{\approx} \text{Ideal}^{\text{Ext}_{\text{ss}}, \text{C}_{\text{ss}}^*. \text{Commit}, \text{C}_{\text{ss}}^*. \text{Open}}(\lambda, (J, \mathbf{b}_J), \sigma) \quad (56)$$

⁴³This ordering is done for simplicity, as it allows us to rely on the analysis of $\text{C}^*. \text{Open}_{[\ell]}$ in the proof of Lemmas 7.13 and 7.14. In particular, we do not need to rely on the fact that measuring the Hadamard basis opening is not detectable when opening, and verifying the opening, in the standard basis.

where $\eta^* \leq C \cdot \sqrt{\delta^*}$ and

$$\delta^* = \mathbb{E}_{\substack{(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Gen}(1^\lambda) \\ (\mathbf{y}, \boldsymbol{\rho}) \leftarrow \mathbf{C}_{\text{SS}}^*. \text{Commit}(\mathbf{pk}, \boldsymbol{\sigma})}} \max_{\mathbf{b}' \in \{\mathbf{b}_J, \mathbf{0}^{|J|}, \mathbf{1}^{|J|}\}} \Pr[\text{Ver}_{\text{SS}}(\mathbf{sk}, \mathbf{y}, (J, \mathbf{b}'), \mathbf{C}_{\text{SS}}^*. \text{Open}(\boldsymbol{\rho}, \mathbf{b}')) = 0]. \quad (57)$$

By the definition of $\mathbf{C}_{\text{SS}}^*. \text{Open}$, and as explained in Remark 7.16 (and similarly to Equation (22)),

$$\delta^* \leq \epsilon_0^* + \epsilon_1^* + \text{negl}(\lambda) \quad (58)$$

where

$$\epsilon_b^* = \Pr[\text{Ver}_{\text{SS}}(\mathbf{sk}_1, \mathbf{y}, (J, b^{|J|}), \mathbf{z}_b) = 0],$$

and where \mathbf{y} is distributed as in Equation (53), and \mathbf{z}_b is distributed as in Equation (54) when computed coherently by U_b . We next argue that

$$\epsilon_0^* + \epsilon_1^* \leq \Pr[\mathbf{y} = \perp] + \Pr[\mathbf{z}_0 = \perp] + \Pr[\mathbf{z}_1 = \perp] + 2\delta'_0 + 8\epsilon_0 + \text{negl}(\lambda). \quad (59)$$

The reason Equation (59) holds is that after extracting \mathbf{y} the residual state is ϵ_0 -indistinguishable from the state obtained without extraction. After further extracting \mathbf{z}_b the residual state is $2\epsilon_0$ -indistinguishable from the state obtained without extraction. By Definition 6.5 and Corollary 6.6, this implies that the probability that in the third argument-of-knowledge, the extractor outputs a valid witness $(\mathbf{y}, \mathbf{z}_b)$, corresponding to the instance $(\mathbf{sk}_1, \mathbf{hk}, \mathbf{rt}, \mathbf{rt}', b)$, is at most $2\epsilon_0 + \delta'_{0,b} + 2\epsilon_0 = \delta'_{0,b} + 4\epsilon_0$, up to negligible factors, where $\delta'_{0,b}$ is the probability that P_{Test}^* is rejected given that the first message sent by V_{Test} is $b \in \{0, 1\}$. This, together with the collision resistance property of the underlying hash family, and with the fact that $\delta'_{0,0} + \delta'_{0,1} = 2\delta'_0$, implies that Equation (59) indeed holds.

Note that

$$\Pr[\mathbf{y} = \perp] \leq \delta_0 + 2\epsilon_0 + \text{negl}(\lambda) \quad (60)$$

This follows from the following calculation:

$$\begin{aligned} \Pr[\mathbf{y} = \perp] &= \\ \Pr[\mathbf{y} = \perp \wedge \mathbb{T}_{\text{Sim}} \text{ is rejecting}] &+ \Pr[\mathbf{y} = \perp \wedge \mathbb{T}_{\text{Sim}} \text{ is accepting}] \leq \\ \Pr[\mathbb{T}_{\text{Sim}} \text{ is rejecting}] &+ \Pr[\mathbf{y} = \perp \wedge \mathbb{T}_{\text{Sim}} \text{ is accepting}] \leq \\ \delta_0 + \epsilon_0 + \epsilon_0 &+ \text{negl}(\lambda) \end{aligned}$$

where the latter equation follows from the definition of δ_0 and from Definition 6.5. Similarly,

$$\Pr[\mathbf{z}_b = \perp] \leq \delta'_{0,b} + 3\epsilon_0 + \text{negl}(\lambda) \quad (61)$$

This follows from the following calculation:

$$\begin{aligned} \Pr[\mathbf{z}_b = \perp] &= \\ \Pr[\mathbf{z}_b = \perp \wedge \mathbb{T}_{\text{Sim}} \text{ is rejecting}] &+ \Pr[\mathbf{z}_b = \perp \wedge \mathbb{T}_{\text{Sim}} \text{ is accepting}] \leq \\ \Pr[\mathbb{T}_{\text{Sim}} \text{ is rejecting}] &+ \Pr[\mathbf{z}_b = \perp \wedge \mathbb{T}_{\text{Sim}} \text{ is accepting}] \leq \\ \delta'_{0,b} + 2\epsilon_0 + \epsilon_0 &+ \text{negl}(\lambda) \end{aligned}$$

This, together with Equations (58) and (59), implies that

$$\delta^* \leq (\delta_0 + 2\epsilon_0) + (2\delta'_0 + 6\epsilon_0) + 2\delta'_0 + 8\epsilon_0 = \delta_0 + 4\delta'_0 + 16\epsilon_0.$$

We conclude that

$$\eta^* \leq O(\sqrt{\delta_0 + \delta'_0}) + C \cdot \sqrt{16\epsilon_0} \leq O(\sqrt{\delta_0 + \delta'_0}) + \frac{\epsilon}{2}.$$

In order to use Equation (56), with η^* as above, we define

$$\text{Succ-Ideal}^{\text{Ext}_{\text{ss}}, \text{C}_{\text{ss}}^*. \text{Commit}_{\epsilon_0}, \text{C}_{\text{ss}}^*. \text{Open}}(\lambda, (J, \mathbf{b}_J), \boldsymbol{\sigma})$$

to be the distribution obtained by sampling

$$(\text{pk}_1, \mathbf{y}, (J, \mathbf{b}_J), \mathbf{m}) \leftarrow \text{Ideal}^{\text{Ext}_{\text{ss}}, \text{C}_{\text{ss}}^*. \text{Commit}_{\epsilon_0}, \text{C}_{\text{ss}}^*. \text{Open}}(\lambda, (J, \mathbf{b}_J), \boldsymbol{\sigma}),$$

sampling $\text{hk} \leftarrow \text{Gen}_{\text{H}}(1^\lambda)$, computing $\text{rt} = \text{Eval}_{\text{H}}(\text{hk}, \mathbf{y})$, and outputting

$$((\text{pk}_1, \text{hk}), \text{rt}, (J, \mathbf{b}_J), \mathbf{m}_J).$$

Similarly, we define

$$\text{Succ-Real}^{\text{C}_{\text{ss}}^*. \text{Commit}, \text{C}_{\text{ss}}^*. \text{Open}}(\lambda, (J, \mathbf{b}_J), \boldsymbol{\sigma})$$

to be the distribution obtained by sampling

$$(\text{pk}_1, \mathbf{y}, (J, \mathbf{b}_J), \mathbf{m}) \leftarrow \text{Real}^{\text{Ext}_{\text{ss}}, \text{C}_{\text{ss}}^*. \text{Commit}_{\epsilon_0}, \text{C}_{\text{ss}}^*. \text{Open}}(\lambda, (J, \mathbf{b}_J), \boldsymbol{\sigma}),$$

sampling $\text{hk} \leftarrow \text{Gen}_{\text{H}}(1^\lambda)$, computing $\text{rt} = \text{Eval}_{\text{H}}(\text{hk}, \mathbf{y})$, and outputting

$$((\text{pk}_1, \text{hk}), \text{rt}, (J, \mathbf{b}_J), \mathbf{m}_J).$$

Note that by the definition of the extractor Ext it holds that

$$\text{Ideal}^{\text{Ext}, \text{C}^*. \text{Commit}, P^*, P_{\text{Test}}^*}(\lambda, (J, \mathbf{b}_J), \boldsymbol{\sigma}, \epsilon) \equiv \text{Succ-Ideal}^{\text{Ext}_{\text{ss}}, \text{C}_{\text{ss}}^*. \text{Commit}_{\epsilon_0}, \text{C}_{\text{ss}}^*. \text{Open}}(\lambda, (J, \mathbf{b}_J), \boldsymbol{\sigma}).$$

This is the case since

$$\text{Ideal}^{\text{Ext}, \text{C}^*. \text{Commit}, P^*, P_{\text{Test}}^*}(\lambda, (J, \mathbf{b}_J), \boldsymbol{\sigma}, \epsilon)$$

extracts the state

$$\tau_{\mathcal{A}, \mathcal{B}} \leftarrow \text{Ext}_{\text{ss}}^{\text{C}_{\text{ss}}^*. \text{Open}}(\text{sk}, \mathbf{y}, \boldsymbol{\rho}_{\text{post}, \text{Sim}}),$$

where $(\mathbf{y}, \boldsymbol{\rho}_{\text{post}, \text{Sim}}) \leftarrow \text{C}_{\text{ss}}^*. \text{Commit}_{\epsilon_0}(\text{pk}_1, \boldsymbol{\sigma})$.

Therefore, to prove Equation (55) it suffices to prove that

$$\text{Real}^{\text{C}^*. \text{Commit}, P^*, \text{C}^*. \text{Open}}(\lambda, (J, \mathbf{b}_J), \boldsymbol{\sigma}) \stackrel{\zeta^*}{\approx} \text{Succ-Real}^{\text{C}_{\text{ss}}^*. \text{Commit}, \text{C}_{\text{ss}}^*. \text{Open}}(\lambda, (J, \mathbf{b}_J), \boldsymbol{\sigma}), \quad (62)$$

where $\zeta^* = O(\sqrt{\delta_0 + \delta'_0 + \delta}) + \frac{\epsilon}{2}$.

To this end, we use P_{Test}^* to define a QPT algorithm $\text{C}^{**}. \text{Open}$. We mention that $\text{C}^{**}. \text{Open}$ bears similarity to $\text{C}_{\text{ss}}^*. \text{Open}$ (defined in Item 5 of the definition of Ext), with the difference being that the latter was defined for the semi-succinct commitment, whereas $\text{C}^{**}. \text{Open}$ is defined for the succinct commitment. In particular, recall that for the succinct commitment, an opening to the j 'th qubit consists of a tuple $(\mathbf{y}_j, \mathbf{o}_j, \mathbf{z}_j)$. $\text{C}^{**}. \text{Open}$ uses P_{Test}^* to generate this opening as follows:

1. Use the state-preservation extractor \mathcal{E} (from Definition 6.5) for the NP language \mathcal{L}^* (defined in Equation (19)) to generate

$$(\mathbb{T}_{\text{Sim}}, \mathbf{y}, \rho_{\text{post,Sim}}) \leftarrow \mathcal{E}^{P_{\text{Test}}^* \cdot \rho} \left((\text{hk}, \text{rt}), 1^\lambda, \epsilon_0 \right).$$

Let $(\mathbf{y}_j, \mathbf{o}_j) = \text{Open}_{\text{H}}(\text{hk}, \mathbf{y}, j)$.

2. Let U_b be the unitary as defined in the definition of the extractor Ext above.
3. Compute $\rho' = U_b^\dagger \text{CNOT}_{\text{open}_j, \text{copy}_j} U_b[\rho_{\text{postSim}}]$.
4. Measure register copy_j in the standard basis to obtain \mathbf{z}_j .
5. Output $(\mathbf{y}_j, \mathbf{o}_j, \mathbf{z}_j)$.

Claim 7.23.

$$\text{Real}^{\text{C}^*. \text{Commit}, P^*, \text{C}^*. \text{Open}}(\lambda, (J, \mathbf{b}_J), \sigma) \stackrel{\zeta_1}{\approx} \text{Real}^{\text{C}^*. \text{Commit}, P^*, \text{C}^{**}. \text{Open}}(\lambda, (J, \mathbf{b}_J), \sigma)$$

where $\zeta_1 = O\left(\sqrt{\delta_0 + \delta'_0 + \delta}\right) + \frac{\epsilon}{2}$.

We note that Claim 7.23 completes the proof of Equation (62) since by the definition of $\text{C}^{**}. \text{Open}$ and $\text{C}_{\text{ss}}^*. \text{Open}$

$$\text{Real}^{\text{C}^*. \text{Commit}, P^*, \text{C}^{**}. \text{Open}}(\lambda, (J, \mathbf{b}_J), \sigma) \equiv \text{Succ-Real}^{\text{C}_{\text{ss}}^*. \text{Commit}_{\epsilon_0}, \text{C}_{\text{ss}}^*. \text{Open}}(\lambda, (J, \mathbf{b}_J), \sigma).$$

Proof of Claim 7.23. Equation (15) (which we proved above) implies that it suffices to prove the following:

$$\Pr[\text{Ver}(\text{sk}, \text{rt}, (J, b_J), (\mathbf{y}_J, \mathbf{o}_J, \mathbf{z}_J)) = 0] \leq O(\delta_0 + \delta'_0) + \left(\frac{\epsilon}{2C}\right)^2 + \text{negl}(\lambda) \quad (63)$$

where $(\mathbf{y}_J, \mathbf{o}_J, \mathbf{z}_J) = \text{C}^{**}. \text{Open}(\rho_{\text{post}}, (J, \mathbf{b}_J))$ and where $C \in \mathbb{N}$ is defined in Item 1 of the definition of Ext. We first note that by Definition 6.5 and Corollary 6.6,

$$\Pr[\text{Ver}_{\text{H}}(\text{hk}, \text{rt}, J, \mathbf{y}_J, \mathbf{o}_J) = 0] \leq \delta_0 + 2\epsilon_0 + \text{negl}(\lambda). \quad (64)$$

Moreover, the residual state, denoted by $\rho_{\text{post,Sim}}$ satisfies that

$$(\rho_{\text{post,Sim}}, \text{sk}) \stackrel{\epsilon_0}{\approx} (\rho_{\text{post}}, \text{sk}) \quad (65)$$

which implies that $P_{\text{Test}}^*(\text{pk}, \text{rt}, \rho_{\text{post,Sim}})$, upon receiving $b \in \{0, 1\}$ from V_{Test} is rejected with probability at most $\delta'_{0,b} + \epsilon_0 + \text{negl}(\lambda)$. By Definition 6.5 and Corollary 6.6, this implies that the tuple $(\mathbb{T}_{\text{Sim}}, \mathbf{z}_b, \rho'_{\text{post,Sim}})$ generated in Equation (54) satisfies that

$$\Pr[\text{Eval}_{\text{H}}(\text{hk}, \mathbf{z}_b) \neq \text{rt}'_b] \leq \delta'_{0,b} + 3\epsilon_0 + \text{negl}(\lambda).$$

By the union bound, we conclude that for every $b \in \{0, 1\}$,

$$\Pr[\text{Eval}_{\text{H}}(\text{hk}, \mathbf{y}) \neq \text{rt} \vee \text{Eval}_{\text{H}}(\text{hk}, \mathbf{z}_b) \neq \text{rt}'_b] \leq \delta_0 + \delta'_{0,b} + 5\epsilon_0 + \text{negl}(\lambda). \quad (66)$$

By Definition 6.5, for every $b \in \{0, 1\}$ it holds that the state $\rho'_{\text{post,Sim}}$, generated in Equation (54) as part of U_b , is ϵ_0 -indistinguishable from the state of $P_{\text{Test}}^*(\text{pk}, \text{rt}, \rho_{\text{post,Sim}}, b)$ after executing the first state-preserving argument-of-knowledge. This, together with Equation (65), implies that the state $\rho'_{\text{post,Sim}}$ is $2\epsilon_0$ -indistinguishable from the state of $P_{\text{Test}}^*(\text{pk}, \text{rt}, \rho_{\text{post}}, b)$ after executing the first state-preserving argument-of-knowledge. Since $P_{\text{Test}}^*(\text{pk}, \text{rt}, \rho_{\text{post}}, b)$ is accepted in both its state-preserving argument-of-knowledge protocols with probability at least $1 - \delta'_{0,b}$, it holds that it is accepted in the second state-preserving argument-of-knowledge protocol (w.r.t. the language \mathcal{L}^{**}) when it starts with the state $\rho'_{\text{post,Sim}}$ with probability at least $1 - \delta'_{0,b} - 2\epsilon_0$. This, together with Definition 6.5 and Corollary 6.6, implies that

$$\Pr[(\text{sk}_1, \text{hk}, \text{rt}, \text{rt}', b), (\mathbf{y}, \mathbf{z}) \in \mathcal{R}_{\mathcal{L}^{**}}] \geq 1 - \delta'_{0,b} - 4\epsilon_0,$$

which together with Equation (66) and the collision resistant property of the underlying hash family implies that

$$\Pr[\text{Ver}(\text{sk}, \text{rt}, (J, b_J), (\mathbf{y}_J, \mathbf{o}_J, \mathbf{z}_J)) = 0] \leq O(\delta_0 + \delta'_0) + 9\epsilon_0 + \text{negl}(\lambda)$$

Thus it remains to note that $9\epsilon_0 \leq (\frac{\epsilon}{2C})^2$, as desired. □

□

8 Applications

8.1 Succinct Interactive Arguments for QMA

In this section we construct a succinct interactive argument for **QMA**. To this end, we construct a *semi-succinct* interactive argument for **QMA**, where only the verifier's messages are short but the messages from the prover may be long. We then rely on a black-box transformation from [BKL⁺22] which shows a generic transformation for converting any semi-succinct interactive argument for **QMA** into a fully succinct one.

Ingredients Our semi-succinct interactive argument consists of the following three ingredients:

- A pseudorandom generator $\text{PRG} : \{0, 1\}^\lambda \rightarrow \{0, 1\}^\ell$, where $\ell = \ell(\lambda)$ is a polynomial specified in Lemma 8.1 below.
- A semi-succinct (qubit-by-qubit) commitment scheme ($\text{Gen}, \text{Commit}, \text{Open}, \text{Ver}, \text{Out}$), as defined in Section 5 and constructed in Section 6.
- The information-theoretic QMA verification protocol of Fitzsimons, Hajdušek, and Morimae [FHM18]. As in [BKL⁺22], we use an “instance-independent” version due to [ACGH20] and assume the soundness gap is $1 - \text{negl}(\lambda)$, where the latter can be achieved by standard **QMA** amplification.

Lemma 8.1 ([FHM18, ACGH20, BKL⁺22]). *For all languages $\mathcal{L} = (\mathcal{L}_{\text{yes}}, \mathcal{L}_{\text{no}}) \in \mathbf{QMA}$ there exists a polynomial $k(\lambda)$, a function $\ell(\lambda)$ that is polynomial in the time $T(\lambda)$ required to verify instances of size λ , a QPT algorithm P_{FHM} , and a PPT algorithm V_{FHM} such that the following holds.*

- $P_{\text{FHM}}(\mathbf{x}, |\psi\rangle) \rightarrow |\pi\rangle$: on input an instance $\mathbf{x} \in \{0, 1\}^\lambda$ and a quantum state $|\psi\rangle$, P_{FHM} outputs an $\ell(\lambda)$ -qubit state $|\pi\rangle$.
- **Completeness.** For all $\mathbf{x} \in \mathcal{L}_{\text{yes}}$ and $|\phi\rangle \in \mathcal{R}_{\mathcal{L}}(x)$ it holds that for a random $\mathbf{h} \leftarrow \{0, 1\}^{\ell(\lambda)}$

$$\Pr[V_{\text{FHM}}(\mathbf{x}, \mathbf{v}) = \text{acc} : |\pi\rangle \leftarrow P_{\text{FHM}}(\mathbf{x}, |\phi\rangle^{\otimes k(\lambda)})] \geq 1 - \text{negl}(\lambda)$$

where \mathbf{v} is the result of measuring $|\pi\rangle$ in basis \mathbf{h} .

- **Soundness.** For all $\mathbf{x} \in \mathcal{L}_{\text{no}}$ and all ℓ -qubit states $|\pi^*\rangle$ it holds that for a random $\mathbf{h} \leftarrow \{0, 1\}^{\ell(\lambda)}$,

$$\Pr[V_{\text{FHM}}(\mathbf{x}, \mathbf{v}^*) = \text{acc}] \leq \text{negl}(\lambda)$$

where \mathbf{v}^* is the result of measuring $|\pi^*\rangle$ in basis \mathbf{h} .

The semi-succinct interactive argument for QMA In the following protocol P and V are given an instance \mathbf{x} and P is given k copies of the **QMA** witness $|\psi\rangle$.

$V \rightarrow P$: Generate $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$, and send pk .

$P \rightarrow V$:

1. Compute $|\pi\rangle = P_{\text{FHM}}(\mathbf{x}, |\psi\rangle^{\otimes k})$.
2. Compute $(\mathbf{y}, \rho) \leftarrow \text{Commit}(\text{pk}, |\pi\rangle)$.
Denote by ℓ the number of qubits in $|\pi\rangle$, and denote by $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_\ell)$, where \mathbf{y}_i is a commitment to the i 'th qubit of σ .
3. Send \mathbf{y} .

$V \rightarrow P$: Send a random bit $b \in \{0, 1\}$.

If $b = 0$:⁴⁴

1. $V \rightarrow P$: Send a random bit $h \leftarrow \{0, 1\}$.
2. $P \rightarrow V$: Send $\mathbf{z} \leftarrow \text{Open}(\rho, h^\ell)$.
3. $V \rightarrow P$: Compute $v = \text{Ver}(\text{sk}, \mathbf{y}, h^\ell, \mathbf{z})$ and accept if $v = 1$ and otherwise, reject.

If $b = 1$:

1. $V \rightarrow P$: Send a random seed $\mathbf{s} \leftarrow \{0, 1\}^\lambda$.
2. $P \rightarrow V$: Compute $\mathbf{b} = \text{PRG}(\mathbf{s}) \in \{0, 1\}^\ell$ and send the openings $(\mathbf{z}_1, \dots, \mathbf{z}_\ell) \leftarrow \text{Open}(\rho, \mathbf{b})$.
3. V does the following:
 - (a) Compute $\mathbf{b} = \text{PRG}(\mathbf{s})$.
 - (b) For every $i \in [\ell]$ compute $u_i = \text{Ver}(\text{sk}, \mathbf{y}_i, b_i, \mathbf{z}_i)$ and $v_i = \text{Out}(\text{sk}, \mathbf{y}_i, b_i, \mathbf{z}_i)$.
 - (c) If there exists $i \in [\ell]$ such that $u_i = 0$ then reject.
 - (d) Else, accept if and only if V_{FHM} would accept $(\mathbf{x}, (b_1, \dots, b_\ell), (v_1, \dots, v_\ell))$.

Theorem 8.2. *The above scheme is a semi-succinct interactive argument for QMA.*

⁴⁴This should be thought of as a “test round.”

Proof of Theorem 8.2. The completeness property is straightforward and hence we focus on proving the binding property. Fix a QMA promise problem $\mathcal{L} = (\mathcal{L}_{\text{yes}}, \mathcal{L}_{\text{no}})$. Fix P^* , an input \mathbf{x}^* and an auxiliary state σ , such that $P^*(\mathbf{x}^*, \sigma)$ is accepted with probability $1 - \delta$, for $\delta \leq \frac{1}{\lambda^2}$. We argue that it must be the case that $\mathbf{x}^* \notin \mathcal{L}_{\text{no}}$. To this end, we use P^* to construct P_{FHM}^* that is accepted with high probability in the protocol $(P_{\text{FHM}}, V_{\text{FHM}})$ on input \mathbf{x}^* . The algorithm $P_{\text{FHM}}^*(\mathbf{x}^*, \sigma)$ proceeds as follows:

1. Generate $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$.
2. Generate $(\mathbf{y}, \rho) \leftarrow P^*(\text{pk}, \mathbf{x}^*, \sigma)$.
3. Use the extractor Ext from the binding property of the commitment scheme to extract a state $\tau \leftarrow \text{Ext}^{P^*}(\text{sk}, \mathbf{y}, \rho)$
4. Send τ .

We next argue that V_{FHM} accepts τ with high probability on a random basis. To this end, it suffices to argue that it accepts τ with high probability on a pseudorandom basis, since otherwise one can distinguish a pseudorandom string from a truly random one, thus breaking the underlying PRG. Denote by

$$\text{Good} = \{\mathbf{s} \in \{0, 1\}^\lambda : P^* \text{ is accepted w.p. } \geq 1 - \lambda\delta \text{ when } V \text{ sends } \mathbf{s}\}$$

Note that

$$p \triangleq \Pr[\mathbf{s} \in \text{Good}] \geq 1 - \frac{2}{\lambda} \tag{67}$$

which follows from the following Markov argument:

$$\begin{aligned} 1 - 2\delta &\leq \Pr[P^* \text{ is accepted} \mid b = 1] = \\ &\Pr[P^* \text{ is accepted} \mid b = 1 \wedge \mathbf{s} \in \text{Good}] \cdot \Pr[\mathbf{s} \in \text{Good} \mid b = 1] + \\ &\Pr[P^* \text{ is accepted} \mid b = 1 \wedge \mathbf{s} \notin \text{Good}] \cdot \Pr[\mathbf{s} \notin \text{Good} \mid b = 1] \leq \\ &p + (1 - \lambda\delta)(1 - p) = \\ &1 - \lambda\delta(1 - p) \end{aligned}$$

which implies that $-2\delta \leq -\lambda\delta(1 - p)$ and in turn that $\lambda(1 - p) \leq 2$, thus implying Equation (67). By the binding property of the underlying commitment scheme, for any basis $\mathbf{b} = \text{PRG}(\mathbf{s})$ such that $\mathbf{s} \in \text{Good}$, it holds that

$$(\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_{\text{Real}}) \stackrel{O(\sqrt{\lambda\delta})}{\approx} (\text{pk}, \mathbf{y}, \mathbf{b}, \mathbf{m}_{\text{Ideal}})$$

where $\mathbf{m}_{\text{Ideal}}$ is the result of measuring $\tau \leftarrow \text{Ext}^{P^*}(\text{sk}, \mathbf{y}, \rho)$ in basis \mathbf{b} , and \mathbf{m}_{Real} is the output corresponding to the opening of P^* . The fact that the measurements \mathbf{m}_{Real} are accepted by V_{FHM} with probability $\geq 1 - \lambda\delta$ (for any basis $\text{PRG}(\mathbf{s})$ such that $\mathbf{s} \in \text{Good}$) implies that $\mathbf{m}_{\text{Ideal}}$ is accepted by V_{FHM} with probability $\geq 1 - \lambda\delta - O(\sqrt{\lambda\delta})$ (for any such basis). This, together with Equation (67)

implies that τ is accepted by V_{FHM} on a pseudorandom basis with probability

$$\begin{aligned} & \Pr[V_{\text{FHM}} \text{ accepts } \tau \text{ on basis PRG}(\mathbf{s})] \geq \\ & \Pr[V_{\text{FHM}} \text{ accepts } \tau \text{ on basis PRG}(\mathbf{s}) \mid \mathbf{s} \in \text{Good}] \cdot \Pr[\mathbf{s} \in \text{Good}] \geq \\ & \Pr[V_{\text{FHM}} \text{ accepts } \tau \text{ on basis PRG}(\mathbf{s}) \mid \mathbf{s} \in \text{Good}] \cdot \left(1 - \frac{2}{\lambda}\right) \geq \\ & \left(1 - \lambda\delta - O(\sqrt{\lambda\delta})\right) \cdot \left(1 - \frac{2}{\lambda}\right) \end{aligned}$$

This, together with Lemma 8.1 and our assumption that $\delta \leq \frac{1}{\lambda^2}$, implies that $\mathbf{x}^* \notin \mathcal{L}_{\text{No}}$, as desired. \square

8.2 Succinct Interactive Arguments from X/Z Quantum PCPs

In this section we show how to convert any X/Z quantum PCP for a language \mathcal{L} into an succinct interactive argument (P, V) for \mathcal{L} . As in Section 8.1 we construct a semi-succinct interactive argument, and then use the black-box transformation from [BKL⁺22] to convert it into a fully succinct one.

Ingredients Our semi-succinct interactive argument consists of the following ingredients.

- A semi-succinct (qubit-by-qubit) commitment scheme $(\text{Gen}, \text{Commit}, \text{Open}, \text{Ver}, \text{Out})$, as defined in Section 5 and constructed in Section 6.
- An X/Z quantum PCP for the language \mathcal{L} , with verifier V_{QPCP} .

The semi-succinct interactive argument for \mathcal{L} In the following protocol (P, V) are given an instance \mathbf{x} and P is also given an X/Z quantum PCP $|\pi\rangle$.

$V \rightarrow P$: Generate $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$, and send pk .

$P \rightarrow V$: Compute $(\mathbf{y}, \rho) \leftarrow \text{Commit}(\text{pk}, |\pi\rangle)$ and send \mathbf{y}

Denote by ℓ the number of qubits in $|\pi\rangle$, and denote by $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_\ell)$, where \mathbf{y}_i is a commitment to the i 'th qubit of σ .

$V \rightarrow P$: Send a random bit $b \in \{0, 1\}$.

If $b = 0$:⁴⁵

1. $V \rightarrow P$: Send a random bit $h \leftarrow \{0, 1\}$.
2. $P \rightarrow V$: Send $\mathbf{z} \leftarrow \text{Open}(\rho, h^\ell)$.
3. $V \rightarrow P$: Compute $v = \text{Ver}(\text{sk}, \mathbf{y}, h^\ell, \mathbf{z})$ and accept if $v = 1$ and otherwise, reject.

If $b = 1$:

1. $V \rightarrow P$: Send a sample $(i_1, \dots, i_c, b_1, \dots, b_c) \leftarrow V_{\text{QPCP}}(\mathbf{x}, 1^\lambda)$.

⁴⁵This should be thought of as a “test round.”

2. $P \rightarrow V$: Send the openings $(\mathbf{z}_1, \dots, \mathbf{z}_c) \leftarrow \text{Open}(\boldsymbol{\rho}, (i_1, b_1), \dots, (i_c, b_c))$.
3. V does the following:
 - (a) For every $j \in [c]$ compute $v_j = \text{Ver}(\text{sk}, \mathbf{y}_{i_j}, b_j, \mathbf{z}_j)$ and $u_j = \text{Out}(\text{sk}, \mathbf{y}_{i_j}, b_j, \mathbf{z}_j)$.
 - (b) If there exists $j \in [c]$ such that $v_j = 0$ then reject.
 - (c) Else, accept if and only if V_{QPCP} would accept $(\mathbf{x}, (i_1, \dots, i_c), (b_1, \dots, b_c), (u_1, \dots, u_c))$.

Theorem 8.3. *The above scheme is a semi-succinct interactive argument for \mathcal{L} .*

Proof of Theorem 8.3. The completeness property is straightforward and hence we focus on proving the binding property. Fix a QPT cheating prover P^* , an input \mathbf{x}^* and an auxiliary state $\boldsymbol{\sigma}$, such that $P^*(\mathbf{x}^*, \boldsymbol{\sigma})$ is accepted with probability $1 - \delta$, for $\delta \leq \frac{1}{\lambda^2}$. We use P^* to extract an X/Z quantum PCP $\boldsymbol{\pi}$ for $x^* \in \mathcal{L}$ that is accepted with high probability, thus implying that indeed $x^* \in \mathcal{L}$ as desired. This is done as follows:

1. Generate $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$.
2. Generate $(\mathbf{y}, \boldsymbol{\rho}) \leftarrow P^*(\text{pk}, \mathbf{x}^*, \boldsymbol{\sigma})$.
3. Use the extractor Ext from the binding property of the commitment scheme to extract a state $\boldsymbol{\pi} \leftarrow \text{Ext}^{P^*}(\text{sk}, \mathbf{y}, \boldsymbol{\rho})$
4. Output $\boldsymbol{\pi}$.

The fact that P^* is accepted with probability $1 - \delta$ implies that for every $h \in \{0, 1\}$ it opens in an accepted way on h^ℓ with probability at least $1 - 4\delta$. Denote by Good the event that V_{QPCP} samples $(i_1, \dots, i_c, b_1, \dots, b_c)$ such that P^* is accepted with probability $\geq 1 - \lambda\delta$ when V sends $(i_1, \dots, i_c, b_1, \dots, b_c)$. Note that

$$p \triangleq \Pr[\text{Good}] \geq 1 - \frac{2}{\lambda} \quad (68)$$

which follows from the following Markov argument:

$$\begin{aligned} 1 - 2\delta &\leq \Pr[P^* \text{ is accepted} \mid b = 1] = \\ &\Pr[P^* \text{ is accepted} \mid b = 1 \wedge \text{Good}] \cdot \Pr[\text{Good} \mid b = 1] + \\ &\Pr[P^* \text{ is accepted} \mid b = 1 \wedge \neg\text{Good}] \cdot \Pr[\neg\text{Good} \mid b = 1] \leq \\ &p + (1 - \lambda\delta)(1 - p) = \\ &1 - \lambda\delta(1 - p) \end{aligned}$$

which implies that $-2\delta \leq -\lambda\delta(1 - p)$ and in turn that $\lambda(1 - p) \leq 2$, thus implying Equation (68). In what follows we say that $(i_1, \dots, i_c, b_1, \dots, b_c) \in \text{Good}$ if P^* is accepted when V sends $(i_1, \dots, i_c, b_1, \dots, b_c)$ with probability $\geq 1 - \lambda\delta$. By the binding property of the underlying commitment scheme, for any $(i_1, \dots, i_c, b_1, \dots, b_c) \in \text{Good}$, it holds that

$$(\text{pk}, \mathbf{y}, (i_1, \dots, i_c, b_1, \dots, b_c), \mathbf{m}_{\text{Real}}) \stackrel{O(\sqrt{\lambda\delta})}{\approx} (\text{pk}, \mathbf{y}, (i_1, \dots, i_c, b_1, \dots, b_c), \mathbf{m}_{\text{Ideal}})$$

where $\mathbf{m}_{\text{Ideal}}$ is the result of measuring $\boldsymbol{\pi} \leftarrow \text{Ext}^{P^*}(\text{sk}, \mathbf{y}, \boldsymbol{\rho})$ in locations (i_1, \dots, i_c) and basis (b_1, \dots, b_c) , and \mathbf{m}_{Real} is the output corresponding to the opening of P^* . The fact that the measurements \mathbf{m}_{Real} are accepted by V_{QPCP} with probability $\geq 1 - \lambda\delta$ (for any $(i_1, \dots, i_c, b_1, \dots, b_c) \in \text{Good}$)

implies that $\mathbf{m}_{\text{ideal}}$ is accepted by V_{QPCP} with probability $\geq 1 - \lambda\delta - O(\sqrt{\lambda\delta})$ (for any such basis). This, together with Equation (68) implies that

$$\begin{aligned} & \Pr[V_{\text{QPCP}} \text{ accepts } \boldsymbol{\pi}] \geq \\ & \Pr[V_{\text{QPCP}} \text{ accepts } \boldsymbol{\pi} \mid \text{Good}] \cdot \Pr[\text{Good}] \geq \\ & \Pr[V_{\text{QPCP}} \text{ accepts } \boldsymbol{\pi} \mid \text{Good}] \cdot \left(1 - \frac{2}{\lambda}\right) \geq \\ & \left(1 - \lambda\delta - O(\sqrt{\lambda\delta})\right) \cdot \left(1 - \frac{2}{\lambda}\right) \end{aligned}$$

This, together with our assumption that $\delta \leq \frac{1}{\lambda^2}$, implies that indeed $\boldsymbol{\pi}$ is an X/Z PCP that is accepted with high probability, and thus $\mathbf{x} \in \mathcal{L}$, as desired. □

9 Acknowledgements

Yael Kalai is supported by DARPA under Agreement No. HR00112020023. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Government or DARPA. Agi Villanyi was supported by a Doc Bedard fellowship from the Laboratory for Physical Sciences through the Center for Quantum Engineering. Sam Gunn is supported by a Google PhD Fellowship and the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator.

References

- [ACGH20] Gorjan Alagic, Andrew M. Childs, Alex B. Grilo, and Shih-Han Hung. Non-interactive classical verification of quantum computation. pages 153–180, 2020. [79](#)
- [BBCM93] Bob Blakley, G. R. Blakley, Agnes Hui Chan, and James L. Massey. Threshold schemes with disenrollment. pages 540–548, 1993. [18](#)
- [BCC88] Gilles Brassard, David Chaum, and Claude Crépeau. Minimum disclosure proofs of knowledge. *J. Comput. Syst. Sci.*, 37(2):156–189, 1988. [3](#)
- [BCM⁺18] Zvika Brakerski, Paul F. Christiano, Urmila Mahadev, Umesh V. Vazirani, and Thomas Vidick. A cryptographic test of quantumness and certifiable randomness from a single quantum device. In Mikkel Thorup, editor, *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 320–331. IEEE Computer Society, 2018. [7](#), [9](#), [10](#), [17](#), [18](#), [24](#), [25](#), [27](#), [29](#), [49](#), [55](#), [89](#), [91](#)
- [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, 2022. [3](#), [5](#), [7](#), [10](#), [11](#), [12](#), [13](#), [14](#), [19](#), [24](#), [31](#), [56](#), [79](#), [82](#), [87](#), [88](#)
- [BKNY23] James Bartusek, Fuyuki Kitagawa, Ryo Nishimaki, and Takashi Yamakawa. Obfuscation of pseudo-deterministic quantum circuits. Cryptology ePrint Archive, Paper 2023/252, 2023. <https://eprint.iacr.org/2023/252>. [11](#)
- [CCY20] Nai-Hui Chia, Kai-Min Chung, and Takashi Yamakawa. Classical verification of quantum computations with efficient verifier. pages 181–206, 2020. [10](#)
- [CDv88] David Chaum, Ivan Damgård, and Jeroen van de Graaf. Multiparty computations ensuring privacy of each party’s input and correctness of the result. pages 87–119, 1988. [3](#)
- [CMSZ21] Alessandro Chiesa, Fermi Ma, Nicholas Spooner, and Mark Zhandry. Post-quantum succinct arguments: Breaking the quantum rewinding barrier. Cryptology ePrint Archive, Paper 2021/334, 2021. <https://eprint.iacr.org/2021/334>. [5](#), [18](#), [23](#)
- [FHM18] Joseph F. Fitzsimons, Michal Hajdusek, and Tomoyuki Morimae. Post hoc verification of quantum computation. *Phys. Rev. Lett.*, 120:040501, Jan 2018. [20](#), [79](#)
- [FS87] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. pages 186–194, 1987. [3](#)
- [GJMZ22] Sam Gunn, Nathan Ju, Fermi Ma, and Mark Zhandry. Commitments to quantum states. Cryptology ePrint Archive, Paper 2022/1358, 2022. <https://eprint.iacr.org/2022/1358>. [3](#), [6](#), [11](#)
- [GMW86] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity and a methodology of cryptographic protocol design (extended abstract). pages 174–187, 1986. [3](#)

- [GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. pages 218–229, 1987. [3](#)
- [KLVY22] Yael Kalai, Alex Lombardi, Vinod Vaikuntanathan, and Lisa Yang. Quantum advantage from any non-local game, 2022. [11](#)
- [LMS22] Alex Lombardi, Fermi Ma, and Nicholas Spooner. Post-quantum zero knowledge, revisited or: How to do quantum rewinding undetectably. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 - November 3, 2022*, pages 851–859. IEEE, 2022. [5](#), [19](#), [43](#), [44](#), [48](#)
- [Mah18] Urmila Mahadev. Classical homomorphic encryption for quantum circuits. In Mikkel Thorup, editor, *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 332–338. IEEE Computer Society, 2018. [3](#), [7](#), [10](#), [11](#), [12](#), [13](#), [15](#), [20](#), [24](#), [31](#), [33](#), [39](#), [47](#), [56](#), [87](#)
- [Mer87] Ralph C Merkle. A digital signature based on a conventional encryption function. In *Conference on the theory and application of cryptographic techniques*, pages 369–378. Springer, 1987. [3](#)
- [Mer88] Ralph C. Merkle. A digital signature based on a conventional encryption function. pages 369–378, 1988. [23](#)
- [Mic94] Silvio Micali. CS proofs (extended abstracts). pages 436–453, 1994. [3](#)
- [MNZ24] Tony Metger, Anand Natarajan, and Tina Zhang. Succinct arguments for QMA from standard assumptions via compiled nonlocal games. 2024. To appear. [11](#)
- [MP11] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. *IACR Cryptol. ePrint Arch.*, page 501, 2011. [27](#)
- [Unr16] Dominique Unruh. Collapse-binding quantum commitments without random oracles. pages 166–195, 2016. [23](#), [26](#)
- [Vid20] Thomas Vidick. Interactions with quantum devices (course), 2020. <http://users.cms.caltech.edu/~vidick/teaching/fsmp/fsmp.pdf>. [9](#), [11](#), [13](#), [56](#)
- [VZ21] Thomas Vidick and Tina Zhang. Classical proofs of quantum knowledge. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology - EUROCRYPT 2021 - 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17-21, 2021, Proceedings, Part II*, volume 12697 of *Lecture Notes in Computer Science*, pages 630–660. Springer, 2021. [18](#)
- [Wil11] Mark Wilde. *From Classical to Quantum Shannon Theory*. 2011. arXiv:1106.1445v8. [22](#)

A Weak commitments to Quantum States (WCQ)

In this section, we recall the Measurement Protocol from Mahadev [Mah18], which was formalized by [BKL⁺22]. We refer to such a protocol as a *weak commitment to quantum states* (WCQ) protocol, and define it formally below.

Definition A.1 (Weak Commitment to Quantum States (WCQ)). *An ℓ -qubit WCQ protocol is specified by the five algorithms $(\text{Gen}_W, \text{Commit}_W, \text{Open}_W, \text{Test}_W, \text{Out}_W)$:*

1. Gen_W is a PPT algorithm that takes as input the security parameter λ (in unary) and a string $h \in \{0, 1\}^\ell$, and outputs a pair $(\text{pk}, \text{sk}) \leftarrow \text{Gen}_W(1^\lambda, h)$, where pk is referred to as the public key and sk is referred to as the secret key.
2. Commit_W is a QPT algorithm that takes as input a public key pk and a quantum state σ and outputs a pair $(\mathbf{y}, \rho) \leftarrow \text{Commit}_W(\text{pk}, \sigma)$, where \mathbf{y} is a classical string, referred to as the commitment string, and ρ is a quantum state.
3. Open_W is a QPT algorithm that takes as input a bit $c \in \{0, 1\}$ and a quantum state ρ and outputs a classical string $z \leftarrow \text{Open}_W(\rho, c)$, referred to as the opening string.
4. Test_W is a polynomial time algorithm that takes as input a public key pk and a pair (\mathbf{y}, z) , where \mathbf{y} is a commitment string and z is an opening string, and it outputs $\{\text{acc}, \text{rej}\} \leftarrow \text{Test}_W(\text{pk}, (\mathbf{y}, z))$.
5. Out_W is a polynomial time algorithm that takes as input a secret key sk and a pair (\mathbf{y}, z) , where \mathbf{y} is a commitment string and z is an opening string, and it outputs a classical string $m \in \{0, 1\}^\ell$.

The commitment protocol associated with the tuple $(\text{Gen}_W, \text{Commit}_W, \text{Open}_W, \text{Test}_W, \text{Out}_W)$ is a two party protocol between a QPT committer C which takes as input a quantum state σ , and a BPP verifier V which takes as input a classical string $h \in \{0, 1\}^\ell$. Both parties also take as input the unary security parameter λ . The protocol consists of two phases COMMIT and OPEN, proceeding as follows:

- COMMIT phase:
 1. $[C \leftarrow V]$: V samples $(\text{pk}, \text{sk}) \leftarrow \text{Gen}_W(1^\lambda, h)$ and sends the public key pk to C .
 2. $[C \rightarrow V]$: C computes $(\mathbf{y}, \rho) \leftarrow \text{Commit}_W(\text{pk}, \sigma)$ and sends the commitment string \mathbf{y} to the verifier.
- OPEN phase:
 1. $[C \leftarrow V]$: V samples a random challenge bit $c \leftarrow \{0, 1\}$ and sends c to C .
 2. $[C \rightarrow V]$: C sends $z \leftarrow \text{Open}_W(\rho, c)$ to V .
 3. If $c = 0$, V outputs $\{\text{acc}, \text{rej}\} \leftarrow \text{Test}_W(\text{pk}, \mathbf{y}, z)$. If $c = 1$, V outputs $m \leftarrow \text{Out}_W(\text{sk}, \mathbf{y}, z)$.

A WCQ protocol acts over registers $\mathcal{P}, \mathcal{Y}, \mathcal{Z}, \mathcal{W}$ where \mathcal{P} contains the public component of the output of Gen_W , \mathcal{Y} contains the output of Commit_W , \mathcal{Z} contains the output of Open_W , and \mathcal{W} are additional work registers. Additionally, the commitment protocol satisfies the following properties for *correctness* and *binding*.

Definition A.2 (WCQ correctness). Let $\text{RealW}(1^\lambda, \sigma, h)$ be the distribution resulting from running $(\text{pk}, \text{sk}) \leftarrow \text{Gen}_W(1^\lambda, h)$, $(\mathbf{y}, \rho) \leftarrow \text{Commit}_W(\text{pk}, \sigma)$, $z \leftarrow \text{Open}_W(\rho, 1)$, and outputting $m \leftarrow \text{Out}_W(\text{sk}, y, z)$. Let $\sigma(h)$ denote the distribution resulting from measuring each qubit i of a quantum state σ in the basis specified by h_i for $i \in [\ell]$. A WCQ protocol is correct if, for all ℓ -qubit quantum states σ and for every $h \in \{0, 1\}^\ell$, the following two properties are satisfied:

1. (Test Round Completeness):

$$\Pr \left[\begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{Gen}_W(1^\lambda, h); \\ \text{acc} \leftarrow \text{Test}_W(\text{pk}, y, z) : (\mathbf{y}, \rho) \leftarrow \text{Commit}_W(\text{pk}, \sigma); \\ z \leftarrow \text{Open}_W(\rho, 0) \end{array} \right] = 1 - \text{negl}(\lambda) \quad (69)$$

2. (Measurement Round Completeness):

$$\left\{ \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{Gen}_W(1^\lambda, h); \\ m \leftarrow \text{Out}_W(\text{sk}, y, z) : (\mathbf{y}, \rho) \leftarrow \text{Commit}_W(\text{pk}, \sigma); \\ z \leftarrow \text{Open}_W(\rho, 1) \end{array} \right\} \approx_c \sigma(h) \quad (70)$$

Definition A.3 (WCQ Binding). [BKL⁺22] A WCQ protocol is binding if there exists a PPT classical algorithm SimGen and a QPT oracle machine WExt such that, for any cheating QPT committer C^* with quantum state σ that satisfies that for every $h \in \{0, 1\}^\ell$:

$$\Pr \left[\begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{Gen}_W(1^\lambda, h); \\ \text{acc} \leftarrow \text{Test}_W(\text{pk}, y, z) : (\mathbf{y}, \rho) \leftarrow C^*. \text{Commit}_W(\text{pk}, \sigma); \\ z \leftarrow C^*. \text{Open}_W(\rho, 0) \end{array} \right] = 1 - \text{negl}(\lambda), \quad (71)$$

it holds that for every $h \in \{0, 1\}^\ell$,

$$\text{Sim}^{C^*}(1^\lambda, h) \approx_c \text{RealW}^{C^*}(1^\lambda, h)$$

where

- $\text{Sim}^{C^*}(1^\lambda, h)$ is the output distribution of the following procedure:
 1. Sample $(\text{pk}, \text{sk}) \leftarrow \text{SimGen}(1^\lambda)$.
 2. Execute the commitment round to obtain $(\mathbf{y}, \rho) \leftarrow C^*. \text{Commit}_W(\sigma)$.
 3. Execute $\tau \leftarrow \text{WExt}^{C^*}(\text{pk}, \text{sk}, y, \rho)$.
 4. Measure τ in the basis specified by h , where $h_i = 0$ corresponds to the standard basis and $h_i = 1$ corresponds to the Hadamard, and output these measurement values.
- $\text{RealW}^{C^*}(1^\lambda, h)$ is the output distribution of the following procedure:
 1. Sample $(\text{pk}, \text{sk}) \leftarrow \text{Gen}_W(1^\lambda, h)$.
 2. Emulate the commitment round to obtain $(\mathbf{y}, \rho) \leftarrow C^*. \text{Commit}_W(\sigma)$.
 3. Emulate the opening phase round corresponding to $c = 1$ to obtain $z \leftarrow C^*. \text{Open}_W(\rho, 1)$.
 4. Compute $m \leftarrow \text{Out}_W(\text{sk}, y, z)$ and output m .

B Proof of Lemma 4.5

To show our modified version of Lemma 4.6 of [BCM⁺18], it suffices to show a version of their Lemma 4.9 modified to handle two arbitrary binary strings $\mathbf{d}'_1, \mathbf{d}'_2$. Once this has been shown, the remaining argument proceeds unchanged. In this section will only present our modified version of Lemma 4.9 and its proof.

First, let us recall some basic properties of the discrete Fourier transform. Define the q th root of unity

$$\omega_q = e^{2\pi i/q}.$$

The “standard Fourier identity” is that

$$\sum_{x \in \mathbb{Z}_q} \omega_q^x = 0.$$

For a function $f : \mathbb{Z}_q^\ell \times \mathbb{Z}_2 \rightarrow \mathbb{C}$, the Fourier transform \hat{f} is defined by

$$\hat{f}(\mathbf{x}, y) = \sum_{\mathbf{v}, z} \omega_q^{\mathbf{v} \cdot \mathbf{x}} (-1)^{y \cdot z} f(\mathbf{v}, z).$$

With this normalization, we have the following version of Plancherel’s theorem:

$$\sqrt{2q^\ell} \|f\|_2 = \|\hat{f}\|_2.$$

Now, in the context of Lemma 4.9 of [BCM⁺18], we are given a random matrix $\mathbf{C} \in \mathbb{Z}_q^{\ell \times n}$, and arbitrary distinct nonzero binary vectors $\mathbf{d}'_1, \mathbf{d}'_2 \in \{0, 1\}^n$. Define

$$g(\mathbf{v}, z_1, z_2) = \Pr_{\mathbf{s} \in \{0, 1\}^n} [\mathbf{v} = \mathbf{C}\mathbf{s}, z_1 = \mathbf{d}'_1 \cdot \mathbf{s}, z_2 = \mathbf{d}'_2 \cdot \mathbf{s}].$$

Then, to prove the Lemma, it suffices to show that with high probability over the choice of the matrix \mathbf{C} , the probability distribution whose density is g is close to the uniform distribution over the space $\mathbb{Z}_q^\ell \times \mathbb{Z}_2 \times \mathbb{Z}_2$. Specifically, denoting by U the uniform distribution and denoting by D the Total Variation Distance, we wish to show that

$$D(g, U) \leq q^{\ell/2} \cdot 2^{-n/40}.$$

To do so, we relate the TVD distance to the L_2 norm of the difference $g - U$:

$$\begin{aligned} D(g, U) &= \frac{1}{2} \|g - U\|_1 \\ &\leq \frac{1}{2} \sqrt{2q^\ell} \|g - U\|_2 \\ &= \frac{1}{2} \|\hat{g} - \hat{U}\|_2, \end{aligned}$$

where the second line follows from Cauchy-Schwarz. Note that for *any* probability density g over $\mathbb{Z}_q^\ell \times \mathbb{Z}_2 \times \mathbb{Z}_2$, we have that $\hat{g}(0^\ell, 0, 0) = 1$. This is because

$$\hat{f}(0^\ell, 0, 0) = \sum_{\mathbf{v}, z_1, z_2} g(\mathbf{v}, z_1, z_2) = 1.$$

Moreover, for the uniform density U , we further have $\hat{U}(\mathbf{x}, y_1, y_2) = 0$ for all $(\mathbf{x}, y_1, y_2) \neq (0^\ell, 0, 0)$. This follows by the standard Fourier identity.

Thus, we get

$$\begin{aligned} \frac{1}{2} \|\hat{g} - \hat{U}\|_2 &= \frac{1}{2} \sqrt{\sum_{\mathbf{x}, y_1, y_2} |\hat{g}(\mathbf{x}, y_1, y_2) - \hat{U}(\mathbf{x}, y_1, y_2)|^2} \\ &= \frac{1}{2} \sqrt{\sum_{(\mathbf{x}, y_1, y_2) \neq (0^\ell, 0)} |\hat{g}(\mathbf{x}, y_1, y_2)|^2}. \end{aligned}$$

To bound this sum, we will now calculate \hat{g} , using the identities $(-1)^{yz} = (e^{\pi i})^{yz} = e^{(2\pi i/2) \cdot yz}$ to simplify the resulting sums.

$$\begin{aligned} \hat{g}(\mathbf{x}, y_1, y_2) &= \sum_{\mathbf{v}, z_1, z_2} \omega_q^{\mathbf{v} \cdot \mathbf{x}} (-1)^{y_1 z_1 + y_2 z_2} g(\mathbf{v}, z_1, z_2) \\ &= \sum_{\mathbf{v}, z_1, z_2} e^{2\pi i \cdot (\mathbf{v} \cdot \mathbf{x} / q + (y_1 z_1 + y_2 z_2) / 2)} g(\mathbf{v}, y_1, y_2) \\ &= \sum_{\mathbf{v}, z_1, z_2} e^{2\pi i \cdot (2\mathbf{v} \cdot \mathbf{x} + q(y_1 z_1 + y_2 z_2)) / 2q} g(\mathbf{v}, z_1, z_2) \\ &= \sum_{\mathbf{v}, z_1, z_2} \omega_{2q}^{2\mathbf{v} \cdot \mathbf{x} + q(y_1 z_1 + y_2 z_2)} g(\mathbf{v}, z_1, z_2) \\ &= \mathbb{E}_{\mathbf{s} \in \{0,1\}^n} \sum_{\mathbf{v}, z_1, z_2} \omega_{2q}^{2\mathbf{v} \cdot \mathbf{x} + q(y_1 z_1 + y_2 z_2)} \mathbf{1}[\mathbf{v} = \mathbf{C}\mathbf{s}, z_1 = \mathbf{d}'_1 \cdot \mathbf{s}, z_2 = \mathbf{d}'_2] \\ &= \mathbb{E}_{\mathbf{s} \in \{0,1\}^n} \omega_{2q}^{2\mathbf{x} \cdot (\mathbf{C}\mathbf{s}) + q((y_1 \mathbf{d}'_1 + y_2 \mathbf{d}'_2) \cdot \mathbf{s})}. \end{aligned}$$

Define $\mathbf{w} = 2\mathbf{C}^T \mathbf{x} + q(y_1 \mathbf{d}'_1 + y_2 \mathbf{d}'_2)$ so that $\mathbf{w}^T \mathbf{s}$ is equal to the exponent in the last line above. Then

$$\hat{g}(\mathbf{x}, y_1, y_2) = \mathbb{E}_{\mathbf{s} \in \{0,1\}^n} \omega_{2q}^{\mathbf{w}^T \mathbf{s}}.$$

Our goal is to show that \hat{g} , with the $(0^\ell, 0, 0)$ entry deleted, is small in 2-norm. We are going to do this by bounding the entries individually.

Case 1: $(\mathbf{x}, y) = (0^\ell, 1)$. In this case, we have

$$\begin{aligned} \hat{g}(0^\ell, y_1, y_2) &= \mathbb{E}_{\mathbf{s} \in \{0,1\}^n} \omega_{2q}^{q(y_1 (\mathbf{d}'_1)^T + y_2 (\mathbf{d}'_2)^T) \mathbf{s}} \\ &= \mathbb{E}_{\mathbf{s} \in \{0,1\}^n} (-1)^{(y_1 (\mathbf{d}'_1)^T + y_2 (\mathbf{d}'_2)^T) \mathbf{s}} \\ &= 0, \end{aligned}$$

where in the last line we used the fact that at least one of y_1, y_2 is nonzero, and that $\mathbf{d}'_1 \neq \mathbf{d}'_2$, to say that $y_1 \mathbf{d}'_1 + y_2 \mathbf{d}'_2$ is a nonzero binary vector.

Case 2: $\mathbf{x} \neq 0^\ell$. In this case, we will use the fact that \mathbf{C} is a random matrix. Specifically, in Lemma 4.8 of [BCM⁺18] it is shown that with probability $1 - q^\ell \cdot 2^{-n/8}$, \mathbf{C} is *moderate*. To define this, we start by defining a moderate *scalar*: for $x \in \mathbb{Z}_q$, let its *centered representative* be its unique representative in $(-q/2, q/2]$. Then we say x is *moderate* if its centered representative lies in the range $[-3q/8, -q/8) \cup (q/8, 3q/8]$. This is true for a uniformly random $x \in \mathbb{Z}_q$ with probability $1/2$. A *moderate vector* is one for which at least $1/4$ of the entries are moderate: for a uniformly random vector in \mathbb{Z}_q^n , the chance that it is moderate is exponentially close to 1, by a Chernoff bound. A *moderate matrix* is one for which every nonzero vector in the row span is moderate.

Now, observe:

$$\begin{aligned} \left| \mathbb{E}_{s \in \{0,1\}} \omega_q^{sx} \right| &= \left| \frac{1}{2} (\omega_q^0 + \omega_q^x) \right| \\ &= \left| \frac{1}{2} (\omega_q^{-x/2} + \omega_q^{x/2}) \right| \\ &= |\cos(\pi x/q)|. \end{aligned}$$

Thus, if x is moderate, then $|x/q| \in (1/8, 3/8]$, and $|\cos(\pi x/q)| \leq |\cos(\pi/8)|$. So for any moderate vector $\mathbf{r} \in \mathbb{Z}_q^n$, it holds that

$$\left| \mathbb{E}_{\mathbf{s} \in \{0,1\}^n} \omega_q^{\mathbf{r} \cdot \mathbf{s}} \right| \leq |\cos(\pi/8)|^{n/4}.$$

We will need a slightly refined version of this. Let r be a moderate scalar $\in \mathbb{Z}_q$, and $e_1, e_2 \in \{0, 1\}$ be arbitrary. Then

$$2r + q(e_1 + e_2) \in [-3q/4, -q/4) \cup [q/4, 3q/4) \cup [q+q/4, q+3q/4) \cup (q/4, 3q/4] \cup (q+q/4, q+3q/4] \cup (2q+q/4, 2q+3q/4].$$

Thus,

$$|(2r + q(e_1 + e_2))/q| \in [1/4, 3/4] \pmod{1}.$$

Therefore,

$$\left| \mathbb{E}_{s \in \{0,1\}} \omega_{2q}^{(2r+qe)s} \right| = \left| \cos\left(\frac{\pi}{2q}(2r + qe)\right) \right| \leq |\cos(\pi/8)|.$$

Thus, if \mathbf{r} is moderate and $\mathbf{e}_1, \mathbf{e}_2$ are arbitrary binary vectors, then by the same reasoning

$$\left| \mathbb{E}_{\mathbf{s} \in \{0,1\}^n} \omega_{2q}^{(2\mathbf{r} + q(\mathbf{e}_1 + \mathbf{e}_2)) \cdot \mathbf{s}} \right| \leq |\cos(\pi/8)|^{n/4}.$$

Now, to finish the argument, let's return to \hat{g} . For any $\mathbf{x} \neq 0^\ell$, since \mathbf{C} is moderate, we know $\mathbf{x}^T \mathbf{C}$ is a moderate vector. Thus, we have

$$\hat{g}(\mathbf{x}, y_1, y_2) = \mathbb{E}_{\mathbf{s} \in \{0,1\}^n} \omega_{2q}^{2(\mathbf{x}^T \mathbf{C})\mathbf{s} + q(y_1(\mathbf{d}'_1 \cdot \mathbf{s}) + y_2(\mathbf{d}'_2 \cdot \mathbf{s}))}$$

Now setting $\mathbf{r} = \mathbf{x}^T \mathbf{C}$ and $\mathbf{e}_{1,2} = y_{1,2} \mathbf{d}'_{1,2}$, we conclude that

$$|\hat{g}(\mathbf{x}, y_1, y_2)| \leq |\cos(\pi/8)|^{n/4}.$$

So in the end we get

$$\begin{aligned}
D(g, U) &\leq \frac{1}{2} \sqrt{\sum_{(\mathbf{x}, y_1, y_2) \neq (0^\ell, 0, 0)} |\hat{g}(\mathbf{x}, y_1, y_2)|^2} \\
&\leq \frac{1}{2} \sqrt{\sum_{\mathbf{x} \neq 0^\ell} \sum_{(y_1, y_2) \in \{0, 1\}^2} |\cos(\pi/8)|^{n/2}} \\
&= \frac{1}{2} \sqrt{4(q^\ell - 1) |\cos(\pi/8)|^{n/2}} \\
&\leq q^{\ell/2} \cdot 2^{-n/40}
\end{aligned}$$

This is the desired bound.