Publicly-Verifiable Deletion via Target-Collapsing Functions

James Bartusek[1], Dakshita Khurana[2], and Alexander Poremba[3]

[1] University of California, Berkeley
[2] University of Illinois Urbana-Champaign
[3] California Institute of Technology

Abstract. We build quantum cryptosystems that support publicly-verifiable deletion from standard cryptographic assumptions. We introduce target-collapsing as a weakening of collapsing for hash functions, analogous to how second preimage resistance weakens collision resistance; that is, target-collapsing requires indistinguishability between superpositions and mixtures of preimages of an honestly sampled image.

We show that target-collapsing hashes enable publicly-verifiable deletion (PVD), proving conjectures from [Poremba, ITCS'23] and demonstrating that the Dual-Regev encryption (and corresponding fully homomorphic encryption) schemes support PVD under the LWE assumption. We further build on this framework to obtain a variety of primitives supporting publicly-verifiable deletion from weak cryptographic assumptions, including:

- Commitments with PVD assuming the existence of injective one-way functions, or more generally, *almost-regular* one-way functions. Along the way, we demonstrate that (variants of) target-collapsing hashes can be built from almost-regular one-way functions.
- Public-key encryption with PVD assuming trapdoored variants of injective (or almost-regular) one-way functions. We also demonstrate that the encryption scheme of [Hhan, Morimae, and Yamakawa, Eurocrypt'23] based on pseudorandom group actions has PVD.
- X with PVD for $X \in$ {attribute-based encryption, quantum fully-homomorphic encryption, witness encryption, time-revocable encryption}, assuming X and trapdoored variants of injective (or almost-regular) one-way functions.

1 Introduction

Recent research has explored the exciting possibility of combining quantum information with computational hardness to enable classically infeasible cryptographic tasks. Beginning with proposals such as unforgeable money [28], this list has recently grown to include the possibility of provably deleting cryptographic information encoded into quantum states [27, 9, 17, 15, 16, 21, 6, 7, 2, 5].

In this work, we further investigate the task of provable deletion of information via destructive measurements. We focus on building primitives that satisfy *publicly-verifiable deletion* (PVD). This deletion property allows any participant in possession of a quantum encoding to publish a publicly-verifiable classical certificate proving that they deleted[4] the underlying plaintext. This is in contrast to the weaker *privately-verifiable deletion* property, where deletion can be verified only by parties that hold a secret verification key, and this key must remain hidden from the party holding the ciphertext. Public verification is more desirable due to its stronger security guarantee: secret verification keys do not need to be stored in hidden locations, and security continues to hold even when the verification key is leaked. Furthermore, clients can outsource verification of deletion by publishing the verification key itself.

Our approach to building publicly verifiable deletion departs from templates used in prior works on deletion. While most prior works, building on [27, 9], rely on the combination of a quantum information-theoretic tool such as Wiesner encodings/BB84 states [28, 8] and a cryptographic object such as an encryption scheme, our work enables publicly-verifiable deletion by directly using simple cryptographic properties of many-to-one hash functions.

The Template, in a Nutshell. When illustrating our approach to publicly-verifiable deletion, it will help to first consider enabling this for a simple cryptographic primitive: a commitment scheme. That is, we consider building a statistically binding non-interactive quantum bit commitment scheme where each commitment is accompanied by a classical, *public* verification key vk. A receiver holding the commitment may generate a classical proof that they deleted the committed bit b, and this proof can be publicly verified against vk. We would like to guarantee that as long as verification accepts, the receiver has information-theoretically removed b from their view and will be unable to recover it given unbounded resources, despite previously having the bit b determined by their view.

To allow verification to be a public operation, it is natural to imagine the certificate or proof of deletion to be a hard-to-find solution to a public puzzle. For instance, the public verification key could be an image y of a (one-way) function, and the certificate of deletion a valid pre-image $f^{-1}(y)$ of this key. Now, the commitment itself must encode the committed bit b in such a way

[4] In this work, we focus on *information-theoretic* deletion of computationally hidden secrets, where the guarantee is that after deletion, even an unbounded adversary cannot recover the plaintext that was previously determined by their view [6].

that the ability to generate $f^{-1}(y)$ given the commitment implies information-theoretic deletion of b. This can be enabled by encoding b in the *phase* of a state supported on multiple pre-images of y.

Namely, given an appropriate *two-to-one* function f, a commitment[5] to a bit b can be

$$\mathsf{Com}(b) = \left(y, |0, x_0\rangle_{\mathsf{A}} + (-1)^b |1, x_1\rangle_{\mathsf{A}}\right)$$

where $(0, x_0), (1, x_1)$ are the two pre-images of (a randomly sampled) image y.

Given an image y and a state on register A, a valid certificate of deletion of the underlying bit could be any pre-image of y, which for a well-formed commitment will be obtained by measuring the A register in the computational basis. It is easy to see that an immediate *honest* measurement of the A register implies information-theoretic erasure of the phase b. But a malicious adversary holding the commitment may decide to perform arbitrary operations on this state in an attempt to find a pre-image y without erasing b.

In this work, we analyze (minimal) requirements on the cryptographic hardness of f in the template above, so that the ability to computationally find any preimage of y given the commitment necessarily implies *information-theoretic* erasure of b. A useful starting point, inspired by recent conjectures in [21], is the *collapsing* property of hash functions. This property was first introduced in [26] as a quantum strengthening of collision-resistance.

Collapsing Functions. The notion of *collapsing* considers an experiment where a computationally bounded adversary prepares an arbitrary superposition of preimages of f on a register A, after which the challenger tosses a random coin c. If $c = 0$, the challenger measures register A, otherwise it measures a register containing the hash y of the value on register A, thus leaving A holding a superposition of preimages of y. The register A is returned to the adversary, and we say that f is collapsing if the adversary cannot guess c with better than negligible advantage. Constructions of collapsing hash functions are known based on LWE [25], low-noise LPN [30], and more generally on special types of collision-resistant hashes. They have played a key role in the design of post-quantum protocols, especially in settings where proofs of security of these protocols rely on *rewinding* an adversary.

It is easy to see that

$$\mathsf{Com}(b) = \left(y, |0, x_0\rangle + (-1)^b |1, x_1\rangle\right)$$

computationally hides the bit b as long as the function f used to build the commitment above is *collapsing*. Indeed, collapsing implies that the superposition $|0, x_0\rangle + (-1)^b |1, x_1\rangle$ is computationally indistinguishable from the result of measurement in the computational basis, and the latter perfectly erases the phase b. However, PVD requires something stronger: we must show that any adversary that generates a valid pre-image of y given the superposition

[5] Technically, it is only an appropriate purification of the scheme described here that will satisfy binding; we ignore this detail for the purposes of this overview.

$|0, x_0\rangle + (-1)^b |1, x_1\rangle$, must have *information-theoretically* deleted b from its view, despite b being information-theoretically present in the adversary's view before generating the certificate. We show via a careful proof that this is indeed the case for collapsing f. Proving this turns out to be non-trivial. Indeed, a similar construction in [21] based on the Ajtai hash function [3] relied on an unproven conjecture, which we prove in this work by developing new techniques.

In addition, we show how f in the template above can be replaced with functions that satisfy weaker properties than collapsing, yielding PVD from regular variants of one-way functions. We discuss these results below.

1.1 Our Results

We introduce new properties of (hash) functions, namely target-collapsing, generalized target-collision-resistance. We will show that hash functions satisfying these properties (1) can be based on (regular) variants of one-way functions and (2) imply publicly-verifiable deletion in many settings. Our results also use an intermediate notion, a variant of target-collapsing that satisfies certified everlasting security. Before discussing our results, we motivate and discuss these new definitions informally below.

Definitions

Target-Collapsing and Generalized Target-Collision-Resistant Functions. Towards better understanding the computational assumptions required for PVD, we observe that in the deletion experiment for the commitment above, the superposition $|x_0\rangle + (-1)^b |x_1\rangle$ is prepared by an *honest committer*. This indicates that the collapsing requirement, where security is required to hold even for an adversarial choice of superposition over preimages, may be overkill.

Inspired by this, we consider a natural weakening called *target-collapsing*, where *the challenger (as opposed to the adversary)* prepares a superposition of preimages of a random image y of f on register A. After this, the challenger tosses a random coin c. If $c = 0$, it does nothing to A, otherwise it measures A in the computational basis. The register A is returned to the adversary, and we say that a hash function is target-collapsing if a computationally bounded adversary cannot guess c with better than negligible advantage.

As highlighted above, this definition weakens collapsing to allow the challenger (instead of the adversary) to prepare the preimage register. The weakening turns out to be significant because we show that target-collapsing functions are realizable from relatively weak cryptographic assumptions – namely variants of one-way functions – which are unlikely to imply (standard) collapsing or collision-resistant hash functions due to known black-box separations [22].

To enable these instantiations from weaker assumptions, we first further generalize target-collapsing so that when $c = 1$, the challenger applies a *binary-outcome measurement* M to A (as opposed to performing a computational basis

measurement resulting in a singleton preimage). Thus, a template commitment with PVD from generalized target-collapsing hashes has the form:

$$\mathsf{Com}(b) = \left(y, \sum_{x:f(x)=y,M(x)=0} |x\rangle + (-1)^b \sum_{x:f(x)=y,M(x)=1} |x\rangle \right).$$

We show that this commitment satisfies PVD as long as f is target-collapsing w.r.t. the measurement M, and satisfies an additional property of "generalized" target-collision-resistance (TCR), that we discuss next.

Generalized target-collision-resistance is a quantum generalization of the (standard) cryptographic property of second pre-image resistance/target-collision-resistance. Very roughly, this considers an experiment where the challenger first prepares a superposition of preimages of a random image y of f on register A. After this, the challenger applies a measurement (e.g., a binary-outcome measurement) M on A to obtain outcome μ and sends A to the adversary. We require that no polynomially-bounded adversary given register A can output *any* preimage x' of y such that $M(x') \neq M(\mu)$ (except with negligible probability)[6].

Certified Everlasting Target-Collapsing. In order to show PVD, instead of directly relying on target-collapsing (which only considers computationally bounded adversaries), we introduce a stronger notion that we call *certified everlasting* target-collapsing. This considers the following experiment: as before, the challenger prepares a superposition of preimages of a random image y of f on register A. After this, the challenger tosses a random coin c. If $c = 0$, it does nothing to A, otherwise it applies measurement M to A. The register A is returned to the adversary, after which the adversary is required to return a pre-image of y as its "deletion certificate". While such a certificate can be obtained via an honest measurement of the register A, the *certified everlasting target-collapsing* property requires that the following *everlasting* security guarantee hold. As long as the adversary is computationally bounded at the time of generating a valid deletion certificate, verification of this certificate implies that the bit c is *information-theoretically* erased from the adversary's view, and cannot be recovered even given unbounded resources. That is, if the adversary indeed returns a valid pre-image, they will never be able to guess whether or not the challenger applied measurement M.

New Constructions and Theorems

Main Theorem. Now, we are ready to state the main theorem of our paper. In a nutshell, this says that any (hash) function f that satisfies both target-collapsing and (generalized) target-collision resistance also satisfies *certified everlasting* target-collapsing.

[6] We remark that this notion can also be seen as a generalization of "conversion hardness" defined in [14].

Theorem 1. *(Informal). If f satisfies target-collapsing and generalized target-collision-resistance with respect to measurement M, then f satisfies* certified everlasting target-collapsing *with respect to the measurement M.*

We also extend recent results from the collapsing literature [12, 30, 11] to show that for the case of binary-outcome (in fact, polynomial-outcome) measurements M, generalized TCR with respect to M actually implies target-collapsing with respect to M. Thus, we obtain the following corollary.

Corollary 1. *(Informal). If f satisfies generalized target-collision-resistance with respect to a* binary-outcome *measurement M, then f satisfies* certified everlasting target-collapsing *with respect to the measurement M.*

Resolving the Strong Gaussian Collapsing Conjecture [21]. We now apply the main theorem and its corollary to build various cryptographic primitives with PVD. First, we immediately **prove** the following "strong Gaussian-collapsing"[7] conjecture from [21], which essentially conjectures that the Ajtai hash function (based on the hardness of SIS) satisfies a certain form of key-leakage security after deletion. This follows from our main theorem because the Ajtai hash function is known to be collapsing [19, 21] and collision-resistant (which implies that it is target-collapsing and target-collision-resistant when preimages are sampled from the Gaussian distribution).

Conjecture 1 (Strong Gaussian-Collapsing Conjecture, [21]).
There exist parameters $n, m, q \in \mathbb{N}$ with $m \geq 2$ and $\sigma > 0$ such that, for every efficient quantum algorithm $\mathcal{A}$, it holds that

$$\Big|\Pr\big[\mathsf{StrongGaussCollapseExp}_{\mathcal{A},n,m,q,\sigma}(0)=1\big] - \Pr\big[\mathsf{StrongGaussCollapseExp}_{\mathcal{A},n,m,q,\sigma}(1)=1\big]\Big| \leq \mathrm{negl}(\lambda)$$

with respect to the experiment defined in Figure 1.

This conjecture, from [21] considers a slightly weaker notion of certified collapsing which resembles the notion of certified deletion first proposed by Broadbent and Islam [9]. Here, the adversary is not computationally unbounded once a valid deletion certificate is produced; instead, the challenger simply reveals some additional secret information (in the case of the strong Gaussian-collapsing experiment, the challenger reveals a short trapdoor vector for the Ajtai hash function[8]).

[7] Here, "Gaussian" refers to a quantum superposition of Gaussian-weighted vectors, where the distribution assigns probability proportional to $\rho_\sigma(\mathbf{x}) = \exp\big(-\pi\|\mathbf{x}\|^2/\sigma^2\big)$ for vectors $\mathbf{x} \in \mathbb{Z}^m$ and parameter $\sigma > 0$.

[8] In the strong Gaussian-collapsing experiment it is crucial that the trapdoor is only revealed after a valid certificate is presented; otherwise, the adversary can easily distinguish the collapsed from the non-collapsed world by applying the Fourier transform and using the trapdoor to distinguish LWE samples from uniformly random vectors [21].

$\underline{\mathsf{StrongGaussCollapseExp}_{\mathcal{A},n,m,q,\sigma}(b)}$:

1. The challenger samples $\bar{\mathbf{A}} \xleftarrow{\$} \mathbb{Z}_q^{n\times(m-1)}$ and prepares the Gaussian state

$$|\psi\rangle_{XY} = \sum_{\mathbf{x}\in\mathbb{Z}_q^m} \rho_\sigma(\mathbf{x})\,|\mathbf{x}\rangle_X \otimes |\mathbf{A}\cdot\mathbf{x} \pmod q\rangle_Y\,,$$

   where $\mathbf{A} = [\bar{\mathbf{A}} \,\|\, \bar{\mathbf{A}}\cdot\bar{\mathbf{x}} \pmod q] \in \mathbb{Z}_q^{n\times m}$ is a matrix with $\bar{\mathbf{x}} \xleftarrow{\$} \{0,1\}^{m-1}$.
2. The challenger measures Y in the computational basis, resulting in

$$|\psi_{\mathbf{y}}\rangle_{XY} = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^m:\\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \rho_\sigma(\mathbf{x})\,|\mathbf{x}\rangle_X \otimes |\mathbf{y}\rangle_Y\,.$$

3. If $b = 0$, the challenger does nothing. Else, if $b = 1$, the challenger measures system X in the computational basis. The challenger then sends system X to $\mathcal{A}$, together with the matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ and the string $\mathbf{y} \in \mathbb{Z}_q^n$.
4. $\mathcal{A}$ sends a classical witness $\mathbf{w} \in \mathbb{Z}_q^m$ to the challenger.
5. The challenger checks if $\mathbf{w}$ satisfies $\mathbf{A}\cdot\mathbf{w} = \mathbf{y} \pmod q$ and $\|\mathbf{w}\| \le \sigma\sqrt{m/2}$. If true, the challenger sends the trapdoor vector $\mathbf{t} = (\bar{\mathbf{x}}, -1) \in \mathbb{Z}^m$ to $\mathcal{A}$, where $\mathbf{A}\cdot\mathbf{t} = \mathbf{0} \pmod q$. Else, the challenger outputs a random bit $b' \leftarrow \{0,1\}$ and the game ends.
6. $\mathcal{A}$ returns a bit b', which is retured as the output of the experiment.

Fig. 1. The strong Gaussian-collapsing experiment [21].

Following results from [21], we obtain the following cryptosystems with PVD, for the first time from standard cryptographic assumptions.

Theorem 2. *(Informal) Assuming the hardness of* LWE *and* SIS *with appropriate parameters, there exists public-key encryption and (leveled) fully-homomorphic encryption with* PVD*.*

Next, we ask whether one necessarily needs to rely on concrete, highly structured assumptions such as LWE in order to achieve publicly-verifiable deletion, or whether weaker generic assumptions suffice. We present a more general approach to building primitives with PVD from weaker, generic assumptions.

Commitments with PVD *from Regular One-Way Functions.* We first formulate the notion of a *balanced binary-measurement* TCR hash, which is any function that is TCR with respect to some appropriately balanced binary-outcome measurement. By balanced, we mean that the set of preimages of a random image will have significant weight on preimages that correspond to both measurement outcomes (this will roughly be required to guarantee the binding property of

our commitment/correctness properties of our encryption schemes). By roughly following the template described above, we show that such hashes generically imply commitments with PVD. Next, we show that such "balanced" functions can be based on (almost-)regular one-way functions[9] By carefully instantiating this outline, we obtain the following results.

Theorem 3. *(Informal). Assuming the existence of almost-regular one-way functions, there exists a balanced binary-outcome TCR hash, and consequently there exist commitments with* PVD*.*

Public-Key Encryption with PVD *from Regular Trapdoor Functions.* Next, we take this framework to the public-key setting, showing that any balanced binary-outcome TCR hash with an additional "trapdoor" property generically implies a public-key encryption scheme with PVD. The additional property roughly requires the existence of a trapdoor for f that enables recovering the phase term from the quantum commitments discussed above: we call this *trapdoor phase-recoverability*. We show that balanced binary-outcome TCR, with trapdoor phase-recoverability, can be based on injective trapdoor one-way functions or pseudorandom group actions (the latter builds on [14]).

Theorem 4. *(Informal). Assuming the existence of injective trapdoor one-way functions or pseudorandom group actions, there exists a balanced binary-outcome TCR hash with trapdoor phase-recoverability, and consequently there exists public-key encryption with* PVD*.*

We also show that injectivity requirement on the trapdoor function can be further relaxed to a notion of "superposition-invertible" trapdoor regular one-way function for the results above. Informally, this is a regular one-way function, where a trapdoor allows one to obtain a uniform superposition over all preimages of a given image. This is an example of a *generic assumption* that is not known to, and perhaps is unlikely to, imply classical public-key encryption – but does imply PKE with quantum ciphertexts, and in fact even one that supports PVD. The only other assumption in this category is the concrete assumption that pseudorandom group actions exist [14].

Theorem 5. *(Informal). Assuming the existence of superposition-invertiable regular trapdoor functions, there exists a balanced binary-outcome TCR hash with trapdoor phase-recoverability and consequently, there exists public-key encryption with* PVD*.*

Advanced Encryption with PVD *from Weak Assumptions* Finally, we show that hybrid encryption gives rise to a generic compiler for encryption with PVD, obtaining the following results.

[9] This is a generalization of regular one-way functions where preimage sets for different images should be polynomially related in size.

Theorem 6. *(Informal). Assuming the existence of injective trapdoor one-way functions or pseudorandom group actions, and $X \in \{$attribute-based encryption, quantum fully-homomorphic encryption, witness encryption, timed-release encryption$\}$, there exists X with* PVD.

Prior to this work, while there existed encryption schemes with PVD from non-standard assumptions such as one-shot signatures [17], conjectured strong collapsing [21] or post-quantum indistinguishability obfuscation [7], no basic or advanced cryptosystems supporting PVD were known from standard assumptions. We provide a more detailed overview of prior work below.

1.2 Prior work

The first notion resembling *certified deletion* was introduced by Unruh [27] who proposed a (private-key) quantum timed-release encryption scheme that is *revocable*, i.e. it allows a user to *return* the ciphertext of a quantum timed-release encryption scheme, thereby losing all access to the data. Unruh's scheme uses conjugate coding [28, 8] and relies on the *monogamy of entanglement* in order to guarantee that revocation necessarily erases information about the plaintext. Broadbent and Islam [9] introduced the notion of *certified deletion* and constructed a private-key quantum encryption scheme with the aforementioned feature which is inspired by the quantum key distribution protocol [8, 24]. In contrast with Unruh's [27] notion of revocable quantum ciphertexts which are eventually returned and verified, Broadbent and Islam [9] consider certificates which are entirely classical. Moreover, the security definition requires that, once the certificate is successfully verified, the plaintext remains hidden even if the secret key is later revealed. Inspired by the notion of quantum copy-protection [1], Ananth and La Placa [4] defined a form of quantum software protection called *secure software leasing* whose anti-piracy notion requires that the encoded program is returned and verified.

Using a hybrid encryption scheme, Hiroka, Morimae, Nishimaki and Yamakawa [17] extended the scheme in [9] to both public-key and attribute-based encryption with privately-verifiable certified deletion via *receiver non-committing* encryption [18, 10]. Hiroka, Morimae, Nishimaki and Yamakawa [16] considered *certified everlasting zero-knowledge proofs* for QMA via the notion of *everlasting security* which was first formalized by Müller-Quade and Unruh [20]. Bartusek and Khurana [6] revisited the notion of certified deletion and presented a unified approach for how to generically convert any public-key, attribute-based, fully-homomorphic, timed-release or witness encryption scheme into an equivalent quantum encryption scheme with certified deletion. In particular, they considered a stronger notion called *certified everlasting security* which allows the adversary to be computationally unbounded once a valid deletion certificate is submitted. This is also the definition we consider in this work. In the same spirit, Hiroka, Morimae, Nishimaki and Yamakawa [15] gave a *certified everlasting* functional encryption scheme which allows the receiver of the ciphertext to obtain the outcome specific function applied the plaintext, but nothing else. In

other very recent work, Ananth, Poremba and Vaikuntanathan [5] used Gaussian superpositions to construct (key)-revocable cryptosystems, such as public-key encryption, fully homomorphic encryption and pseudorandom functions assuming the hardness of LWE, and Agarwal et al. [2] introduced a generic compiler for adding key-revocability to a variety of cryptosystems. In these systems, the cryptographic key consists of a quantum state which can later be *certifiably revoked* via a quantum channel – in contrast with the classical deletion certificates for ciphertexts considered in this work.

Cryptosystems with Publicly Verifiable Deletion. First, in addition to their results in the setting of private verification, [17] also gave a public-key encryption scheme with certified deletion which is *publicly verifiable* assuming the existence of one-shot signatures (which rely on strong black-box notions of obfucation) and extractable witness encryption. Using *Gaussian superpositions*, Poremba [21] proposed *Dual-Regev*-based public-key and fully homomorphic encryption schemes with certified deletion which are publicly verifiable and proven secure assuming the (then unproven) *strong Gaussian-collapsing conjecture* — a strengthening of the collapsing property of the Ajtai hash. Finally, a recent work [7] relies on post-quantum indistinguishability obfuscation (iO) to obtain both publicly verifiable deletion and publicly verifiable key revocation. This is a strong assumption for which we have candidates, but no constructions based on standard (post-quantum) assumptions at this time.

2 Technical Overview

In this overview, we begin by discussing the key ideas involved in proving our main theorem. We show how to prove publicly verifiable deletion for a toy protocol that relies on stronger assumptions than the ones that we actually rely on in our actual technical sections.

Next, we progressively relax these assumptions to instantiate broader frameworks, including the one from [21], obtaining public-key encryption and fully-homomorphic encryption with PVD from LWE/SIS.

Finally, we further generalize this to enable constructions from weak cryptographic assumptions – including commitments with PVD from variants of one-way functions and PKE with PVD from trapdoored variants of the same assumption. We also discuss a hybrid approach that enables a variety of advanced encryption schemes supporting PVD.

2.1 Proving Our Main Theorem

Consider the toy commitment

$$\mathsf{Com}(b) = \left(y, |0, x_0\rangle + (-1)^b |1, x_1\rangle\right)$$

where $(0, x_0), (1, x_1)$ are preimages of y under a structured two-to-one function f, where every image has a preimage that begins with a 0 and another that

begins with a 1. We note that this commitment can be efficiently prepared by first preparing a superposition over all preimages

$$\sum_{b\in\{0,1\},x\in\{0,1\}^\lambda} |b,x\rangle$$

on a register X, then writing the output of f applied on X to register Y, and finally measuring the contents of register Y to obtain image y. The register X contains $|0,x_0\rangle + |1,x_1\rangle$, which can be converted to $|0,x_0\rangle + (-1)^b |1,x_1\rangle$ via (standard) phase kickback.

To show that the commitment satisfies publicly-verifiable deletion, we consider an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ where $\mathcal{A}_1$ is (quantum) polynomial time and $\mathcal{A}_2$ is unbounded, participating in the following experiment.

- The challenger samples $b \leftarrow \{0,1\}$ and runs $\mathsf{Expmt}_0(b)$, described below.
  $\underline{\mathsf{Exmpt}_0(b):}$
  1. Prepare $\left(|0,x_0\rangle + (-1)^b |1,x_1\rangle, y\right)$ on registers A,B and send them to $\mathcal{A}_1$.
  2. $\mathcal{A}_1$ outputs a (classical) deletion certificate γ,[10] and left-over state ρ.
  3. If $f(\gamma) \neq y$, output a uniformly random bit $b' \leftarrow \{0,1\}$, otherwise output $b' = \mathcal{A}_2(\rho)$.
- The advantage of $\mathcal{A}$ is defined to be $\mathsf{Adv}_{\mathcal{A}}^{\mathsf{Expmt}_0} = \left|\Pr[b'=b] - \frac{1}{2}\right|$.

We discuss how to prove the following.

Claim. (Informal). For every $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ where $\mathcal{A}_1$ is (quantum) computationally bounded,

$$\mathsf{Adv}_{\mathcal{A}}^{\mathsf{Expmt}_0} = \mathrm{negl}(\lambda),$$

as long as f is target collapsing and target collision-resistant w.r.t. a computational basis measurement of the pre-image register.

Overview of the Proof of Claim 2.1. To prove this claim, we must show that b is information-theoretically *removed* from the leftover state of any $\mathcal{A}_1$ that generates a valid pre-image of y, despite the fact that the adversary's view contains b at the beginning of the experiment.

Proof techniques for this type of experiment were recently introduced in [6] in the context of *privately verifiable deletion* via BB84 states. Inspired by their method, our first step is to defer the dependence of the experiment on the bit b. In more detail, we will instead imagine sampling the distribution by guessing a uniformly random $c \leftarrow \{0,1\}$, and initializing the adversary with $(|x_0\rangle + (-1)^c |x_1\rangle, y)$. The challenger later obtains input b and aborts the experiment (outputs $\perp$) if $c \neq b$. Since c was a uniformly random guess, the trace distance between the $b=0$ and $b=1$ outputs of this modified experiment is at least half the trace distance between the outputs of the original experiment.

[10] If the $\mathcal{A}_1$ outputs a quantum state as their certificate, the state is measured in the computational basis to obtain a classical certificate γ.

Moreover, we can further delay the process of obtaining input b, and then abort or not until *after* the adversary outputs a certificate of deletion. That is, we can consider a *purification* where a register C contains a superposition $|0\rangle + |1\rangle$ of two choices for c, and is later measured to determine bit c. This experiment is discussed in detail below.

$\underline{\mathsf{Expmt}_1(b)}$: The experiment proceeds as follows.

1. Prepare the $|+\rangle$ state on an ancilla register C, and a superposition of preimages $|x_0\rangle + |x_1\rangle$ of a random y on register A.
2. Then, controlled on the contents of register C, do the following: if the control bit is 0, do nothing, and otherwise flip the phase on x_1 (via phase kickback), changing the contents of A to $|x_0\rangle - |x_1\rangle$. This means that the overall state is
$$\frac{1}{\sqrt{2}} \sum_{c\in\{0,1\}} |c\rangle_{\mathsf{C}} \otimes |0, x_0\rangle_{\mathsf{A}} + (-1)^c |1, x_1\rangle_{\mathsf{A}}$$
Send A to $\mathcal{A}_1$.
3. Obtain from $\mathcal{A}_1$ a purported certificate of deletion γ.
4. If $f(\gamma) \neq y$, abort, and otherwise measure register C to obtain output c, and abort if $c \neq b$. In the case of abort, output a uniformly random bit $b' \leftarrow \{0,1\}$.
5. If no aborts occurred, output $b' = \mathcal{A}_2(\rho)$.

We note that the event $c = b$ occurs with probability exactly $\frac{1}{2}$, and since measurements on separate subsystems commute, we have that

$$\mathsf{Adv}_{\mathcal{A}}^{\mathsf{Expmt}_1} \geq \frac{1}{2}\mathsf{Adv}_{\mathcal{A}}^{\mathsf{Expmt}_0}. \tag{1}$$

where $\mathsf{Adv}_{\mathcal{A}}^{\mathsf{Expmt}_1} = \left| \Pr[\mathsf{Expmt}_1(b) = b] - \frac{1}{2} \right|$ for $b \leftarrow \{0,1\}$.

Once the dependence of the experiment on b has been deferred, as above, we can consider another experiment (described below) where the challenger measures the contents of register A *before* sending it to $\mathcal{A}_1$. Intuitively, performing this measurement *removes* information about b from $\mathcal{A}_1$'s view in a manner that is computationally undetectable by $\mathcal{A}_1$ (due to the target-collapsing property of f).

$\underline{\mathsf{Expmt}_2(b)}$: The experiment proceeds as follows.

– Prepare the $|+\rangle$ state on an ancilla register C, and a superposition of preimages $|x_0\rangle + |x_1\rangle$ of a random y on register A. *Next, measure register* A *in the computational basis.*
Then, controlled on the contents of register C, do the following: if the control bit is 0, do nothing, and otherwise flip the phase on x_1. This means that the overall state is a uniform mixture of the states
$$\frac{1}{\sqrt{2}} \sum_{c\in\{0,1\}} |c\rangle_{\mathsf{C}} \otimes |0, x_0\rangle_{\mathsf{A}} \text{ and } \frac{1}{\sqrt{2}} \sum_{c\in\{0,1\}} (-1)^c |c\rangle_{\mathsf{C}} \otimes |1, x_1\rangle_{\mathsf{A}}$$

Finally, send A to $\mathcal{A}_1$.
- Obtain from $\mathcal{A}_1$ a purported certificate of deletion γ.
- If $f(\gamma) \neq y$, abort, otherwise measure register C to obtain output c, and abort if $c \neq b$. In the case of abort, output a uniformly random bit $b' \leftarrow \{0,1\}$.
- If no aborts occurred, output $b' = \mathcal{A}_2(\rho)$.

As described above, the target-collapsing property of f implies that $\mathcal{A}_1$ cannot (computationally) distinguish the register A obtained in $\mathsf{Expmt}_2(b)$ from the one obtained in $\mathsf{Expmt}_1(b)$. However, this is not immediately helpful: information about which experiment $\mathcal{A}_1$ participated in could potentially be encoded into $\mathcal{A}_1$'s left-over state ρ, so that it remains computationally hidden from $\mathcal{A}_1$ but can be extracted by (unbounded) $\mathcal{A}_2$. And it is after all the output of $\mathcal{A}_2$ that determines the advantage of $\mathcal{A}$. Because of $\mathcal{A}_2$ being unbounded and the experiments only being *computationally* indistinguishable, even if we could show that $\mathsf{Adv}_{\mathcal{A}}^{\mathsf{Expmt}_2} = \mathrm{negl}(\lambda)$, it is unclear how to use this to show our desired claim, i.e., $\mathsf{Adv}_{\mathcal{A}}^{\mathsf{Expmt}_0} = \mathrm{negl}(\lambda)$. It may appear that the proof is stuck.

To overcome this issue, we will aim to identify an *efficiently computable predicate* of the challenger's system, which will *imply* the following (inefficient) property: when $\mathcal{A}_1$ outputs a valid deletion certificate, even an unbounded $\mathcal{A}_2$ cannot determine whether it participated in $\mathsf{Expmt}_1(b)$ or $\mathsf{Expmt}_2(b)$, i.e., $\mathcal{A}_1$'s left-over state is information-theoretically independent of b.

Identifying an Efficiently Computable Predicate. Observe that in $\mathsf{Expmt}_2(b)$, the ancilla register C is *unentangled* with the rest of the experiment. In fact, the ancilla register is exactly $|+\rangle$ when we give the adversary $|0, x_0\rangle$ on register A, and $|-\rangle$ when we give the adversary $|1, x_1\rangle$ on register A. Moreover, in $\mathsf{Expmt}_2(b)$, the target-collision-resistance of f implies that the computationally-bounded $\mathcal{A}_1$ given x_0 cannot output x_1 as their deletion certificate (and vice-versa).

This, along with the fact that the certificate *must* be a pre-image of y means that the following guarantee holds (except with negligible probability) in $\mathsf{Expmt}_2(b)$:

When the adversary outputs a valid certificate γ, a projection of the pre-image register onto $|+\rangle$ succeeds if $\gamma = (0, x_0)$ and a projection of the pre-image register onto $|-\rangle$ succeeds if $\gamma = (1, x_1)$.

At this point, we can rely on the target-collapsing property of f to prove the following claim: the *efficient projection* described above also succeeds except with negligible probability in $\mathsf{Expmt}_1(b)$, when the adversary generates a valid deletion certificate. If this claim is not true, then since the experiments (including $\mathcal{A}_1$) run in quantum polynomial time until the point that the deletion certificate is generated, and the projection is efficient, one can build a reduction that contradicts target-collapsing of f. This reduction obtains a challenge (which is either a superposition when the challenger did not measure, or a mixture if the challenger did measure) on register A, prepares ancilla C as in $\mathsf{Expmt}_1(b)$, then follows steps 2, 3 identically to $\mathsf{Expmt}_1(b)$. Next, given a deletion certificate (β, x_β), the reduction projects C onto $|0\rangle + (-1)^\beta |1\rangle$, outputting 1 if the projection succeeds and 0 otherwise.

Introducing an Alternative Experiment. Having established that the projection above must succeed in $\mathsf{Expmt}_1(b)$ except with negligible probability, we can now consider an alternative experiment $\mathsf{Expmt}_{\mathsf{alt}}(b)$. This is identical to $\mathsf{Expmt}_1(b)$, except that the challenger *additionally* projects register C onto $|0\rangle + (-1)^{\beta}|1\rangle$ when the adversary generates a valid certificate (β, x_β). We established above that the projection is successful in $\mathsf{Expmt}_1(b)$ except with negligible probability, and this implies that

$$\mathsf{Adv}_{\mathcal{A}}^{\mathsf{Expmt}_{\mathsf{alt}}} \geq \mathsf{Adv}_{\mathcal{A}}^{\mathsf{Expmt}_1} - \mathrm{negl}(\lambda) \tag{2}$$

where as before, $\mathsf{Adv}_{\mathcal{A}}^{\mathsf{Expmt}_{\mathsf{alt}}} = \left|\Pr[\mathsf{Expmt}_{\mathsf{alt}}(b) = b] - \frac{1}{2}\right|$ for $b \leftarrow \{0,1\}$.

Crucially, in $\mathsf{Expmt}_{\mathsf{alt}}(b)$, the bit c is determined by a measurement on register C which is *unentangled* with the system and in either the $|+\rangle$ or $|-\rangle$ state (due to the projective measurement that we just applied). Thus, measuring C in the computational basis results in a uniformly random and independent c. By definition of the experiment (abort when $b \neq c$, continue otherwise) – this implies that the bit b is set in a way that is uniformly random and independent of the adversary's view, and thus

$$\mathsf{Adv}_{\mathcal{A}}^{\mathsf{Expmt}_{\mathsf{alt}}} = 0 \tag{3}$$

Now, equations (1, 2, 3) together yield the desired claim, that is, $\mathsf{Adv}_{\mathcal{A}}^{\mathsf{Expmt}_0} = \mathrm{negl}(\lambda)$.

This completes a simplified overview of our key ideas, assuming the existence of a perfectly 2-to-1 function f where every image y has preimages $((0,x_0),(1,x_1))$, and where f satisfies both target-collapsing and target-collision-resistance. Unfortunately, we do not know how to build functions satisfying these clean properties from simple generic assumptions. Instead, we will generalize the template above, where the first generalization will no longer require f be 2-to-1.

Generalizing the Template. First, note that we can replace $|0,x_0\rangle$ and $|1,x_1\rangle$ with superpositions over two disjoint sets of preimages of y separated via an efficient binary-outcome measurement, namely

$$\mathsf{Com}(b) = \sum_{x:f(x)=y,M(x)=0} |x\rangle + (-1)^b \sum_{x:f(x)=y,M(x)=1} |x\rangle$$

We can even consider measurements M that have arbitrarily many outcomes. Proof ideas described above also generalize almost immediately to show that for any M, Com satisfies PVD as long as f is target-collapsing and target-collision resistant w.r.t. M. In fact, we can generalize this even further (see our main results in the full version) to consider arbitrary (as opposed to uniform) distributions over pre-images, as well as to account for any auxiliary information that may be sampled together with the description of the hash function.

Certified Everlasting Target-Collapsing. As discussed in the results section, our actual technical proofs proceed in two parts. (1) Show that for any M, a function f that is target-collapsing and target-collision resistant w.r.t. M is also

certified everlasting target-collapsing w.r.t. M, and (2) show that f being *certified everlasting* target-collapsing implies that Com satisfies publicly verifiable deletion.

Recall that *certified* everlasting target collapsing requires that an adversary that outputs a valid deletion certificate information-theoretically loses the bit b determining whether they received a superposition or a mixture of preimages. Our proof of certified everlasting target-collapsing follows analogously to the proof sketched above. In short, we defer measurement of a bit b which decides whether the adversary is given a superposition or a mixture, and then rely on target-collapsing and target-collision-resistance to argue that an efficient projection on the challenger's state (almost) always succeeds when the adversary outputs a valid certificate. We finally show that success of this projection implies that the adversary's state is information-theoretically independent of b.

The certified everlasting target-collapsing property almost immediately implies certified deletion security of Com via a hybrid argument:

- In Hyb_0, the adversary obtains register A containing
$$\mathsf{Com}(0) = \sum_{x:f(x)=y,M(x)=0} |x\rangle + \sum_{x:f(x)=y,M(x)=1} |x\rangle$$
- In Hyb_1, the measurement M is applied to A before sending it to the adversary.
- In Hyb_2, the adversary obtains register A containing
$$\mathsf{Com}(1) = \sum_{x:f(x)=y,M(x)=0} |x\rangle - \sum_{x:f(x)=y,M(x)=1} |x\rangle$$

The certified everlasting hiding property of f guarantees that all hybrids are statistically close when the adversary outputs a valid deletion certificate. Moreover, these experiments abort and output a random bit when the adversary does not output a valid certificate, and it is easy to show (by computational indistinguishability) that the probability of generating a valid certificate remains negligibly close between experiments.

TCR Implies Target-Collapsing for Polynomial-Outcome Measurements We also show that when M has polynomially many possible outcomes, then TCR implies target-collapsing with respect to M. This follows from techniques that were recently developed in the literature on collapsing versus collision resistant hash functions [12, 30, 11]. In a nutshell, these works showed that any distinguisher that distinguishes mixtures from superpositions over preimages for an *adversarially chosen* image y, can be used to swap between pre-images, and therefore find a collision for y. We observe that their technique is agnostic to whether the image y is chosen randomly (in the targeted setting) or adversarially. Furthermore, it also extends to swapping superpositions over sets of pre-images to superpositions over other sets. These allow us to prove that TCR with respect to any polynomial-outcome measurement M implies target-collapsing with respect to M.

2.2 Publicly-Verifiable Deletion via Gaussian Superpositions

We revisit the *Dual-Regev* public-key and (leveled) fully homomorphic encryption schemes with publicly-verifiable deletion proposed by Poremba [21] and conjectured to be secure under the *strong Gaussian-collapsing property*. By applying our main theorem to the Ajtai hash function, we obtain a proof of the conjecture, which allows us to show the certified everlasting security of the aforementioned schemes assuming the hardness of the LWE assumption.

The constructions introduced in [21] exploit the the duality between LWE and SIS [23], and rely on the fact that one encode Dual-Regev ciphertexts via Gaussian superpositions. Below, we give a high-level sketch of the basic public-key construction.

- To generate a pair of keys $(\mathsf{sk}, \mathsf{pk})$, sample a random $\mathbf{A} \in \mathbb{Z}_q^{n\times(m+1)}$ together with a particular short trapdoor vector $\mathbf{t} \in \mathbb{Z}^{m+1}$ such that $\mathbf{A}\cdot\mathbf{t} = \mathbf{0} \pmod q$. Let $\mathsf{pk} = \mathbf{A}$ and $\mathsf{sk} = \mathbf{t}$.
- To encrypt a single bit $b \in \{0,1\}$ using $\mathsf{pk} = \mathbf{A}$, generate the following pair for a random $\mathbf{y} \in \mathbb{Z}_q^n$:

$$\mathsf{vk} \leftarrow (\mathbf{A}, \mathbf{y})$$
$$|\mathsf{CT}\rangle \leftarrow \sum_{\mathbf{s}\in\mathbb{Z}_q^n} \sum_{\mathbf{e}\in\mathbb{Z}_q^{m+1}} \rho_{q/\sigma}(\mathbf{e})\, \omega_q^{-\langle \mathbf{s},\mathbf{y}\rangle}\, |\mathbf{s}^\intercal\mathbf{A} + \mathbf{e}^\intercal + b\cdot(0,\ldots,0,\lfloor\frac{q}{2}\rfloor)\rangle\,,$$

  where vk is a public verification key and $|\mathsf{CT}\rangle$ is the ciphertext for $\sigma > 0$.
- To decrypt $|\mathsf{CT}\rangle$ using sk, measure in the computational basis to obtain $\mathbf{c} \in \mathbb{Z}_q^{m+1}$, and output 0, if $\mathbf{c}^\intercal \cdot \mathsf{sk} \in \mathbb{Z}_q$ is closer to 0 than to $\lfloor\frac{q}{2}\rfloor$, and output 1, otherwise. Here $\mathsf{sk} = \mathbf{t}$ is chosen such that $\mathbf{c}^\intercal \cdot \mathsf{sk}$ yields an approximation of $b\cdot\lfloor\frac{q}{2}\rfloor$ from which we can recover b.

To delete the ciphertext $|\mathsf{CT}\rangle$, perform a measurement in the Fourier basis. Poremba [21] showed that the Fourier transform of $|\mathsf{CT}\rangle$ results in the *dual* quantum state given by

$$|\widehat{\mathsf{CT}}\rangle = \sum_{\substack{\mathbf{x}\in\mathbb{Z}_q^{m+1}:\\ \mathbf{A}\mathbf{x}=\mathbf{y} \pmod q}} \rho_\sigma(\mathbf{x})\, \omega_q^{\langle \mathbf{x}, b\cdot(0,\ldots,0,\lfloor\frac{q}{2}\rfloor)\rangle}\, |\mathbf{x}\rangle\,.$$

In other words, a Fourier basis measurement of $|\mathsf{CT}\rangle$ will necessarily erase all information about the plaintext $b \in \{0,1\}$ and results in a *short* vector $\pi \in \mathbb{Z}_q^{m+1}$ such that $\mathbf{A}\cdot\pi = \mathbf{y} \pmod q$. To publicly verify a deletion certificate, simply check whether a certificate π is a solution to the (inhomogenous) SIS problem specified by $\mathsf{vk} = (\mathbf{A}, \mathbf{y})$. Due to the hardness of the SIS problem, it is computationally difficult to produce a valid deletion certificate from $(\mathbf{A}, \mathbf{y})$ alone.

Our approach to proving certified everlasting security of the Dual-Regev public-key and fully-homomorphic encryption schemes with publicly-verifiable deletion in [21] is as follows. First, we observe that the Ajtai hash function is both

target-collapsing and target-collision-resistant with respect to the discrete Gaussian distribution. Here, the former follows from LWE as a simple consequence of the *Gaussian-collapsing property* previously shown by Poremba [19, 21], whereas the latter follows immediately from the quantum hardness of SIS. Thus, our main theorem implies that the Ajtai hash function is certified-everlasting target-collapsing. Finally, as a simple corollary of our theorem, we obtain a proof of the *strong Gaussian-collapsing conjecture* in [21], which we stated in Conjecture 1. We also note that the aforementioned conjecture considers a weaker notion of certified collapsing which resembles the notion of certified deletion first proposed by Broadbent and Islam [9]. Here, the adversary is not computationally unbounded once a valid deletion certificate is produced; instead, the challenger simply reveals additional secret information (in the case of the strong Gaussian-collapsing experiment, this is a short trapdoor vector for the Ajtai hash function). Our notion of certified everlasting target-collapsing is significantly stronger; in particular, it implies the weaker collapsing scenario considered by Poremba [21]. This follows from the fact that the security reduction can simply brute-force search for a short trapdoor solution for the Ajtai hash once it enters the phase in which it is allowed to be computationally unbounded. We exploit this fact in the proof of Conjecture 1.

2.3 Weakening Assumptions for Publicly-Verifiable Deletion

Next, we look for instantiations of the above template from *generic* cryptographic assumptions, as opposed to structured specific assumptions such as LWE. Here, all of our instantiations only require us to consider functions that are target-collision-resistant and target-collapsing w.r.t. binary-outcome measurements (and as discussed above, TCR implies certified-everlasting target-collapsing in this setting). In addition, for the case of commitments, in order for the commitment to satisfy binding[11], we require that there is a measurement that can distinguish

$$\sum_{x:f(x)=y,M(x)=0}|x\rangle+\sum_{x:f(x)=y,M(x)=1}|x\rangle$$

from

$$\sum_{x:f(x)=y,M(x)=0}|x\rangle-\sum_{x:f(x)=y,M(x)=1}|x\rangle$$

with probability δ for any constant $0<\delta\leq 1$. For the case of public-key encryption, we similarly require that a trapdoor be able to recover the phase with

[11] We actually prove that a purification of the template commitment described above satisfies honest-binding [29]. Namely, the committer generates the state above but leaves registers containing the image y (and the key, if f is a keyed function) unmeasured, and holds on to these registers for the opening phase. It can later either open the commitment by sending these registers to a receiver, or request deletion, by measuring them and publishing y (and any keys for the function).

constant probability. We then resort to standard amplification techniques to boost correctness error from constant to (negligibly close to) 0. We note that this amplification would also work if the phase was recoverable with inverse-polynomial δ (as opposed to constant); however, we focus on constant δ because of simplicity, and because it suffices for our instantiations.

In the template above, we observe that a measurement can find the phase with inverse polynomial probability whenever the sets

$$\sum_{x:f(x)=y,M(x)=0} |x\rangle \text{ and } \sum_{x:f(x)=y,M(x)=1} |x\rangle$$

are somewhat "balanced", i.e. for a random image y, for sets $S_0 = \{x : f(x) = y, M(x) = 0\}$ and $S_1 = \{x : f(x) = y, M(x) = 1\}$, we have that $\frac{|S_0|}{|S_1|}$ is a fixed constant. We show in **??** and **??** that commitments and PKE with PVD can be obtained from appropriate variants of TCR functions following this template.

Now, our goal is to build such TCR functions from generic assumptions. A natural idea would be to start with any one-way function f and compose it with a random two-to-one hash h defined on its range[12]. Then, any output y of the composed function $(h \circ f)$ is associated with two elements $\{z_0, z_1\} = h^{-1}(y)$ in the range of f, and the binary-outcome measurement would measure one of z_0 or z_1. Recalling that we eventually want to prove target-collision-resistance, the hope would be that just given a superposition over the preimages of, say, z_0, the one-wayness of f would imply the difficulty of finding a preimage of z_1[13]. This could give the type of TCR property we need.

Technical Bottlenecks, and a Resolution. Unfortunately, there are two issues with the approach proposed above. First, f may be extremely unbalanced, so that the relative sizes of the sets of preimages of two random points y_1, y_2, i.e. $|\{x : f(x) = y_1\}|$ and $|\{x : f(x) = y_2\}|$ in its image may have very different sizes, that are not polynomially related with each other. There may even be many points in the co-domain/range that have *zero* preimages (for a general OWF, we cannot guarantee that its image is equal to its range). A second related issue is that the above sketched reduction to one-wayness may not work. Let's say we choose h to be a two-to-one function defined by a random shift Δ, i.e. $h(x) = h(x \oplus \Delta)$. Then we are essentially asking that it be hard to invert a random *range* element of f, as opposed to $f(x)$ for a random *domain* element x, which is the standard one-wayness assumption.

We don't know how to make this approach work from arbitrary one-way functions, which we leave as an open question. Instead, we appeal to a result

[12] The *co-domain* of a function $f : \{0,1\}^n \to \{0,1\}^m$ is $\{0,1\}^m$, and we will also refer to this as the *range* of the function in this paper. The *image* is the set of all actual output values of f, i.e. the set $\{y : \exists x \text{ such that } f(x) = y\}$. The co-domain/range may in general be a superset of the image of a function.

[13] More concretely, a purported reduction to one-wayness when given challenge image z_1, can sample a random image z_0 with its preimages, then find h s.t. $h(z_0) = h(z_1)$, thereby using a TCR adversary to find a preimage of the given challenge z_1.

of [13], who in the classical context of building statistically hiding commitments, show the following result. By appropriately combining an (almost)-*regular*[14] one-way function with universal hash functions, it *is* possible to obtain a function f with exactly the required properties: sufficiently balanced, and one-way over its *range*. The former property means that an overwhelming fraction of range elements have similar-sized preimage sets, while the latter property says that an element y sampled randomly from the range of the function cannot be inverted except with negligible probability. This resolves both the difficulties above.

Given such a balanced function f, we apply a random two-to-one hash h defined by a shift Δ to the range of this f. We prove that this implies the flavor of target-collision-restistant hash that we need to construct commitments with PVD.

Public-Key Encryption with PVD. Next, we note that the construction above *also* yields a public-key encryption scheme, as long as there is a trapdoor that allows recovery of the phase b given the state

$$y, \sum_{x:f(x)=y,M(x)=0} |x\rangle + (-1)^b \sum_{x:f(x)=y,M(x)=1} |x\rangle$$

We call this property "trapdoor phase-recoverability". We show that this property is achievable from generic assumptions, even those that are not known to imply classical PKE.

- Specifically, trapdoor phase-recoverability is implied by a trapdoored variant of (almost) regular one-way functions, for which a trapdoor to the function allows recovery of a uniform superposition over all preimages of any given image y. This then allows efficient projection onto $\sum_{x:f(x)=y,M(x)=0} |x\rangle + (-1)^b \sum_{x:f(x)=y,M(x)=1} |x\rangle$ for any efficient M. We also note that this property is satisfied by any (standard) trapdoored injective function. But it is also satisfied by functions such as the Ajtai function that are not necessarily injective. Indeed, it is unclear how to build classical public-key encryption, or even PKE with classical ciphertexts, given a general trapdoor phase-recoverable function. Nevertheless, we formalize the above ideas in the full version of this article in order to build PKE schemes with quantum ciphertexts, that also support PVD.
- Additionally, we show that a recent public-key encryption scheme of [14] from pseudorandom group actions also satisfies trapdoor phase-recoverability: in fact, the decryption algorithm in [14] relies on recovering the phase from a similar superposition, given a trapdoor.

[14] An almost regular one-way function generalizes regular one-way functions to require only that for any two images y_1, y_2 of the function, the sizes of preimage sets of y_1, y_2 are polynomially related. In particular, injective functions, and (standard) regular functions also satisfy almost-regularity.

Hybrid Encryption with PVD. Finally, we observe that we can use any encryption scheme Enc to encrypt the trapdoor td associated with the above construction, and security will still hold. That is, if Enc is semantically-secure, then our techniques extend to show that a ciphertext of the form

$$y, \sum_{x:f(x)=y,M(x)=0} |x\rangle + (-1)^b \sum_{x:f(x)=y,M(x)=1} |x\rangle \, , \mathsf{Enc}(\mathsf{td})$$

where td is the trapdoor for f, still supports publicly-verifiable deletion of the bit b. Thus, our approach can be seen as a way to *upgrade* cryptographic schemes Enc with special properties to satisfy PVD. In particular, we prove that instantiating Enc appropriately with attribute-based encryption, fully-homomorphic encryption, witness encryption, or timed-release encryption gives us the same scheme supporting PVD.

2.4 Discussion and Directions for Future Work

Our work demonstrates a strong relationship between weak security properties of (trapdoored) one-way functions and publicly-verifiable deletion. In particular, previous work [21] conjectured that collapsing functions, which are a quantum strengthening of collision-resistant hashes, lead to cryptosystems with publicly-verifiable deletion. Besides proving this conjecture, we also show that collapsing/collision-resistance, which are considered stronger assumptions than one-wayness, are actually not necessary for PVD.

Indeed, weakenings called target-collapsing and generalized-target-collision-resistance, that can be obtained from (regular) variants of one-way functions, suffice for publicly-verifiable deletion. Analogously to their classical counterparts, we believe that these primitives will be of independent interest. Indeed, a natural question that this work leaves open is whether variants of these primitives that suffice for publicly-verifiable deletion can be based on *one-way functions* without the regularity constraint. It is also interesting to further understand relationships and implications between target-collision-resistance and target-collapsing, including when these properties may or may not imply each other. It may also be useful to understand if these weaker properties can suffice in place of stronger properties such as collapsing and collision-resistance in other contexts, including the design of post-quantum protocols.

Finally, note that we rely on trapdoored variants of these primitives to build public-key encryption schemes. Here too, in addition to obtaining PKE with PVD from any injective trapdoor one-way function (TDF), it becomes possible to relax assumptions to only require (almost)-regularity and trapdoor phase-recoverability – properties that can plausibly be achieved from weaker concrete assumptions than injective TDFs. These are new examples of complexity assumptions that yield public-key encryption with quantum ciphertexts, but may be too weak to obtain PKE with classical ciphertexts. It is an interesting question to further investigate the weakest complexity assumptions that may imply public-key encryption, with or without PVD.

Acknowledgements

D.K. was supported in part by NSF CAREER CNS-2238718, NSF CNS-2247727 and DARPA SIEVE. This material is based upon work supported by the Defense Advanced Research Projects Agency through Award HR00112020024.

A.P. is partially supported by AFOSR YIP (award number FA9550-16-1-0495), the Institute for Quantum Information and Matter (an NSF Physics Frontiers Center; NSF Grant PHY-1733907) and by a grant from the Simons Foundation (828076, TV).

3 Main Theorem: Certified Everlasting Target-Collapsing

In this section, we prove our main theorem.

3.1 Definitions

First, we present our definitions of target-collapsing and (generalized) target-collision-resistance. We parameterize our definitions by a distribution $\mathcal{D}$ over preimages and a measurement function $\mathcal{M}$. Note that when $\mathcal{M}$ is the identity function, the notion of $(\mathcal{D},\mathcal{M})$-target-collapsing corresponds to a notion where the entire preimage register is measured in the computational basis. In this case we drop parameterization by $\mathcal{M}$ and just say $\mathcal{D}$-target-collapsing. Also, when $\mathcal{D}$ is the uniform distribution, we drop parameterization by $\mathcal{D}$ and just say $\mathcal{M}$-target-collapsing.

Definition 1 ($(\mathcal{D},\mathcal{M})$-Target-Collapsing Hash Function). *Let $\lambda\in\mathbb{N}$ be the security parameter. A hash function family given by $\mathcal{H}=\{H_\lambda:\{0,1\}^{m(\lambda)}\to\{0,1\}^{n(\lambda)}\}_{\lambda\in\mathbb{N}}$ is $(\mathcal{D},\mathcal{M})$-target-collapsing for some distribution $\mathcal{D}=\{D_\lambda\}_{\lambda\in\mathbb{N}}$ over $\{\{0,1\}^{m(\lambda)}\}_{\lambda\in\mathbb{N}}$ and family of functions $\mathcal{M}=\{\{M[h]:\{0,1\}^{m(\lambda)}\to\{0,1\}^{k(\lambda)}\}_{h\in H_\lambda}\}_{\lambda\in\mathbb{N}}$ if, for every QPT adversary $\mathcal{A}=\{\mathcal{A}_\lambda\}_{\lambda\in\mathbb{N}}$,*

$$\begin{aligned}&|\Pr\left[\mathsf{TargetCollapseExp}_{\mathcal{H},\mathcal{A},\mathcal{D},\mathcal{M},\lambda}(0)=1\right]-\\&\Pr\left[\mathsf{TargetCollapseExp}_{\mathcal{H},\mathcal{A},\mathcal{D},\mathcal{M},\lambda}(1)=1\right]|\leq\mathrm{negl}(\lambda).\end{aligned}$$

Here, the experiment $\mathsf{TargetCollapseExp}_{\mathcal{H},\mathcal{A},\mathcal{D},\mathcal{M},\lambda}(b)$ is defined as follows:

1. *The challenger prepares the state*

$$\sum_{x\in\{0,1\}^{m(\lambda)}}\sqrt{D_\lambda(x)}\,|x\rangle$$

on register X, and samples a random hash function $h\xleftarrow{\$}H_\lambda$. Then, it coherently computes h on X (into a fresh $n(\lambda)$-qubit register Y) and measures system Y in the computational basis, which results in an outcome $y\in\{0,1\}^{n(\lambda)}$.

2. *If $b = 0$, the challenger does nothing. Else, if $b = 1$, the challenger coherently computes $M[h]$ on X (into a fresh $k(\lambda)$-qubit register V) and measures system V in the computational basis. Finally, the challenger sends the outcome state in system X to $\mathcal{A}_\lambda$, together with the string $y \in \{0,1\}^{n(\lambda)}$ and a description of the hash function h.*
3. *$\mathcal{A}_\lambda$ returns a bit b', which we define as the output of the experiment.*

We also define an analogous notion of $(\mathcal{D}, \mathcal{M})$-target-collision-resistance, as follows. Similarly to above, we drop the parameterization by $\mathcal{M}$ in the case that it is the identity function, and we drop the parameterization by $\mathcal{D}$ in the case that it is the uniform distribution. Notice that target-collision-resistance (without parameterization) then coincides with the classical notion where a uniformly random input is sampled, and the adversary must find a collision with respect to this input (this is also sometimes called second-preimage resistance, or weak collision-resistance).

Definition 2 ($(\mathcal{D}, \mathcal{M})$-Target-Collision-Resistant Hash Function). *A hash function family $\mathcal{H} = \{H_\lambda : \{0,1\}^{m(\lambda)} \to \{0,1\}^{n(\lambda)}\}_{\lambda\in\mathbb{N}}$ is $(\mathcal{D}, \mathcal{M})$-target-collision-resistant for some distribution $\mathcal{D} = \{D_\lambda\}_{\lambda\in\mathbb{N}}$ over $\{\{0,1\}^{m(\lambda)}\}_{\lambda\in\mathbb{N}}$ and family of functions $\mathcal{M} = \{\{M[h] : \{0,1\}^{m(\lambda)} \to \{0,1\}^{k(\lambda)}\}_{h\in H_\lambda}\}_{\lambda\in\mathbb{N}}$ if, for every QPT adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda\in\mathbb{N}}$,*

$$|\Pr\left[\mathsf{TargetCollRes}_{\mathcal{H},\mathcal{A},\mathcal{D},\mathcal{M},\lambda} = 1\right]| \leq \mathrm{negl}(\lambda).$$

Here, the experiment $\mathsf{TargetCollRes}_{\mathcal{H},\mathcal{A},\mathcal{D},\mathcal{M},\lambda}$ is defined as follows:

1. *The challenger prepares the state*

$$\sum_{x\in\{0,1\}^{m(\lambda)}} \sqrt{D_\lambda(x)}\,|x\rangle$$

on register X, and samples a random hash function $h \xleftarrow{\$} H_\lambda$. Next, it coherently computes h on X (into a fresh $n(\lambda)$-qubit system Y) and measures system Y in the computational basis, which results in an outcome $y \in \{0,1\}^{n(\lambda)}$. Next, it coherently computes $M[h]$ on X (into a fresh $k(\lambda)$-qubit register V) and measures system V in the computational basis, which results in an outcome v. Finally, its sends the outcome state in system X to $\mathcal{A}_\lambda$, together with the string $y \in \{0,1\}^{n(\lambda)}$ and a description of the hash function h.
2. *$\mathcal{A}_\lambda$ responds with a string $x \in \{0,1\}^{m(\lambda)}$.*
3. *The experiment outputs 1 if $h(x) = y$ and $M[h](x) \neq v$.*

Finally, we define the notion of a *certified everlasting* target-collapsing hash.

Definition 3. *A hash function family $\mathcal{H} = \{H_\lambda : \{0,1\}^{m(\lambda)} \to \{0,1\}^{n(\lambda)}\}_{\lambda\in\mathbb{N}}$ is certified everlasting $(\mathcal{D}, \mathcal{M})$-target-collapsing for some distribution $\mathcal{D} = \{D_\lambda\}_{\lambda\in\mathbb{N}}$ over $\{\{0,1\}^{m(\lambda)}\}_{\lambda\in\mathbb{N}}$ and family of functions $\mathcal{M} = \{\{M[h] : \{0,1\}^{m(\lambda)} \to \{0,1\}^{k(\lambda)}\}_{h\in H_\lambda}\}_{\lambda\in\mathbb{N}}$ if for every two-part adversary $\mathcal{A} = \{\mathcal{A}_{0,\lambda}, \mathcal{A}_{1,\lambda}\}_{\lambda\in\mathbb{N}}$, where $\{\mathcal{A}_{0,\lambda}\}_{\lambda\in\mathbb{N}}$ is QPT and $\{\mathcal{A}_{1,\lambda}\}_{\lambda\in\mathbb{N}}$ is unbounded, it holds that*

$$|\Pr\left[\mathsf{EvTargetCollapseExp}_{\mathcal{H},\mathcal{A},\mathcal{D},\mathcal{M},\lambda}(0)=1\right]-\Pr\left[\mathsf{EvTargetCollapseExp}_{\mathcal{H},\mathcal{A},\mathcal{D},\mathcal{M},\lambda}(1)=1\right]|\leq \mathrm{negl}(\lambda).$$

Here, the experiment $\mathsf{EvTargetCollapseExp}_{\mathcal{H},\mathcal{A},\mathcal{D},\mathcal{M},\lambda}(b)$ *is defined as follows:*

1. *The challenger prepares the state*

$$\sum_{x\in\{0,1\}^{m(\lambda)}}\sqrt{D_\lambda(x)}\,|x\rangle$$

   on register X*, and samples a random hash function* $h\xleftarrow{\$}H_\lambda$*. Then, it coherently computes* h *on* X *(into a fresh* $n(\lambda)$*-qubit system* Y*) and measures system* Y *in the computational basis, which results in an outcome* $y\in\{0,1\}^{n(\lambda)}$*.*
2. *If* $b=0$*, the challenger does nothing. Else, if* $b=1$*, the challenger coherently computes* $M[h]$ *on* X *(into an auxiliary* $k(\lambda)$*-qubit system* V*) and measures system* V *in the computational basis. Finally, the challenger sends the outcome state in system* X *to* $\mathcal{A}_{0,\lambda}$*, together with the string* $y\in\{0,1\}^{n(\lambda)}$ *and a description of the hash function* h*.*
3. $\mathcal{A}_{0,\lambda}$ *sends a classical certificate* $\pi\in\{0,1\}^{m(\lambda)}$ *to the challenger and initializes* $\mathcal{A}_{1,\lambda}$ *with its residual state.*
4. *The challenger checks if* $h(\pi)=y$*. If true,* $\mathcal{A}_{1,\lambda}$ *is run until it outputs a bit* b'*. Otherwise,* $b'\leftarrow\{0,1\}$ *is sampled uniformly at random. The output of the experiment is* b'*.*

3.2 Main Theorem

Our main theorem is the following.

Theorem 7. *Let* $\mathcal{H}=\{H_\lambda\}_{\lambda\in\mathbb{N}}$ *be a hash function family that is both* $(\mathcal{D},\mathcal{M})$*-target-collapsing and* $(\mathcal{D},\mathcal{M})$*-target-collision-resistant, for some distribution* $\mathcal{D}$ *and efficiently computable family of functions* $\mathcal{M}$*. Then,* $\mathcal{H}$ *is certified everlasting* $(\mathcal{D},\mathcal{M})$*-target-collapsing.*

Proof. Throughout the proof, we will leave the security parameter implicit, defining $H\coloneqq H_\lambda, D\coloneqq D_\lambda, m\coloneqq m(\lambda), n\coloneqq n(\lambda)$, $k\coloneqq k(\lambda)$, $\mathcal{A}_0\coloneqq\mathcal{A}_{0,\lambda}$, and $\mathcal{A}_1\coloneqq\mathcal{A}_{1,\lambda}$. Next, we define

$$|\psi\rangle_X\coloneqq\sum_{x\in\{0,1\}^m}\sqrt{D(x)}\,|x\rangle\,.$$

For $h\in H, y\in\{0,1\}^m$, we define a unit vector

$$|\psi_{h,y}\rangle_X\propto\sum_{x\in\{0,1\}^m:h(x)=y}\sqrt{D(x)}\,|x\rangle\,.$$

Finally, for $h \in H, y \in \{0,1\}^m, v \in \{0,1\}^k$ we define a unit vector

$$|\psi_{h,y,v}\rangle_X \propto \sum_{x\in\{0,1\}^m : h(x)=y, M[h](x)=v} \sqrt{D(x)}\,|x\rangle\,.$$

We consider the following hybrids.

- $\mathsf{Exp}_0(b)$:
  1. The challenger prepares $|\psi\rangle_X$, samples a random hash function $h \xleftarrow{\$} H_\lambda$, coherently computes h on X into a fresh n-qubit register Y, and measures Y in the computational basis to obtain $y \in \{0,1\}^n$ and a left-over state $|\psi_{h,y}\rangle_X$.
  2. If $b = 0$, the challenger does nothing. Else, if $b = 1$, the challenger computes $M[h]$ on X into a fresh k-qubit register V, and measures V in the computational basis. Finally, the challenger sends the left-over state in system X to $\mathcal{A}_0$, together with the string $y \in \{0,1\}^n$ and a classical description of h.
  3. $\mathcal{A}_0$ sends a classical certificate $\pi \in \{0,1\}^m$ to the challenger and initializes $\mathcal{A}_1$ with its residual state.
  4. The challenger checks if $h(\pi) = y$. If true, $\mathcal{A}_1$ is run until it outputs a bit b'. Otherwise, $b' \leftarrow \{0,1\}$ is sampled uniformly at random. The output of the experiment is b'.
- $\mathsf{Exp}_1(b)$:
  1. The challenger prepares $|\psi\rangle_X$, samples a random hash function $h \xleftarrow{\$} H_\lambda$, coherently computes h on X into a fresh n-qubit register Y, and measures Y in the computational basis to obtain $y \in \{0,1\}^n$ and a left-over state $|\psi_{h,y}\rangle_X$.
  2. The challenger computes $M[h]$ on X into a fresh k-qubit register V to obtain a state

$$\propto \sum_{x\in\{0,1\}^m : h(x)=y} \sqrt{D(x)}\,|x\rangle_X\,|M[h](x)\rangle_V\,.$$

Then, the challenger samples a random string $z \xleftarrow{\$} \{0,1\}^k$, prepares a $|+\rangle$ state in system C, and applies a controlled-Z^z operation from C to V, which results in a state

$$\begin{aligned}
&\propto \sum_{c\in\{0,1\}} |c\rangle_C \otimes \sum_{x\in\{0,1\}^m : h(x)=y} \sqrt{D(x)}\,|x\rangle_X\,\mathsf{Z}^{c\cdot z}\,|M[h](x)\rangle_V \\
&= \sum_{c\in\{0,1\}} |c\rangle_C \otimes \sum_{x\in\{0,1\}^m : h(x)=y} \sqrt{D(x)}(-1)^{c\cdot\langle M[h](x),z\rangle}\,|x\rangle_X\,|M[h](x)\rangle_V\,.
\end{aligned}$$

Finally, the challenger uncomputes the V register by again computing $M[h]$ from X to V, and sends system X to $\mathcal{A}_0$, together with $y \in \{0,1\}^n$ and a classical description of h.

3. $\mathcal{A}_0$ sends a classical certificate $\pi \in \{0,1\}^m$ to the challenger and initializes $\mathcal{A}_1$ with its residual state.
4. The challenger checks if $h(\pi) = y$. Then, the challenger measures system C to obtain $c' \in \{0,1\}$ and checks that $c' = b$. If both checks are true, $\mathcal{A}_1$ is run until it outputs a bit b'. Otherwise, $b' \leftarrow \{0,1\}$ is sampled uniformly at random. The output of the experiment is b'.

– $\mathsf{Exp}_2(b)$:

1. The challenger prepares $|\psi\rangle_X$, samples a random hash function $h \xleftarrow{\$} H_\lambda$, coherently computes h on X into a fresh n-qubit register Y, and measures Y in the computational basis to obtain $y \in \{0,1\}^n$ and a left-over state $|\psi_{h,y}\rangle_X$.
2. The challenger computes $M[h]$ on X into a fresh k-qubit register V. Then, the challenger samples a random string $z \xleftarrow{\$} \{0,1\}^k$, prepares a $|+\rangle$ state in system C, applies a controlled-Z^z operation from C to V, and finally uncomputes the V register by again computing $M[h]$ from X to V. Note that this results in a state

$$\propto \sum_{c\in\{0,1\}} |c\rangle_C \otimes \sum_{x\in\{0,1\}^m : h(x)=y} (-1)^{c\cdot\langle M[h](x),z\rangle} |x\rangle_X .$$

   Finally, it sends system X to $\mathcal{A}_0$, together with $y \in \{0,1\}^n$ and a classical description of h.
3. $\mathcal{A}_0$ sends a classical certificate $\pi \in \{0,1\}^m$ and initializes $\mathcal{A}_1$ with its residual state.
4. The challenger checks if $h(\pi) = y$. Then, the challenger applies the following projective measurement to system C:

$$\left\{|\phi^z_\pi\rangle\langle\phi^z_\pi|, I - |\phi^z_\pi\rangle\langle\phi^z_\pi|\right\} \quad \text{where} \quad |\phi^z_\pi\rangle := \frac{1}{\sqrt{2}}\left(|0\rangle + (-1)^{\langle M[h](\pi),z\rangle}|1\rangle\right),$$

   and checks that the first outcome is observed. Finally, the challenger measures system C to obtain $c' \in \{0,1\}$ and checks that $c' = b$. If all three checks are true, $\mathcal{A}_1$ is run until it outputs a bit b'. Otherwise, $b' \leftarrow \{0,1\}$ is sampled uniformly at random. The output of the experiment is b'.

Finally, we also use the following hybrid which is convenient for the sake of the proof.

– $\mathsf{Exp}_3(b)$:

1. The challenger prepares $|\psi\rangle_X$, samples a random hash function $h \xleftarrow{\$} H_\lambda$, coherently computes h on X into a fresh n-qubit register Y, and measures Y in the computational basis to obtain $y \in \{0,1\}^n$ and a left-over state $|\psi_{h,y}\rangle_X$.
2. The challenger computes $M[h]$ on X into a fresh k-qubit register V. Then, the challenger measures V in the computational basis to obtain $v \in \{0,1\}^k$. Next, the challenger samples a random string $z \xleftarrow{\$} \{0,1\}^k$, prepares a $|+\rangle$ state in system C, applies a controlled-Z^z operation from

C to V, and finally uncomputes the V register by again computing $M[h]$ from X to V. Note that this results in the state

$$\frac{1}{\sqrt{2}}\left(|0\rangle_C + (-1)^{\langle v,z\rangle}|1\rangle_C\right) \otimes |\psi_{h,y,v}\rangle_X .$$

Finally, the challenger sends system X to $\mathcal{A}_0$, together with $y \in \{0,1\}^n$ and a classical description of h.

3. $\mathcal{A}_0$ sends a classical certificate $\pi \in \{0,1\}^m$ to the challenger and initializes $\mathcal{A}_1$ with its residual state.
4. The challenger checks if $h(\pi) = y$. Then, the challenger applies the following projective measurement to system C:

$$\left\{|\phi^z_\pi\rangle\langle\phi^z_\pi|, I - |\phi^z_\pi\rangle\langle\phi^z_\pi|\right\} \quad \text{where} \quad |\phi^z_\pi\rangle := \frac{1}{\sqrt{2}}\left(|0\rangle + (-1)^{\langle M[h](\pi),z\rangle}|1\rangle\right),$$

and checks that the first outcome is observed. Finally, the challenger measures system C to obtain $c' \in \{0,1\}$ and checks that $c' = b$. If all three checks are true, $\mathcal{A}_1$ is run until it outputs a bit b'. Otherwise, $b' \leftarrow \{0,1\}$ is sampled uniformly at random. The output of the experiment is b'.

Before we analyze the probability of distinguishing between the consecutive hybrids, we first show that the following statements hold for the final experiment $\mathsf{Exp}_3(b)$.

Claim. The probability that the challenger accepts the deletion certificate π in Step 4 of $\mathsf{Exp}_3(b)$ and $M[h](\pi) \neq v$ is negligible. That is,

$$\Pr_{h,y,v}\left[h(\pi) = y \ \wedge \ M[h](\pi) \neq v \ : \ \pi \leftarrow \mathcal{A}_0(h, y, |\psi_{h,y,v}\rangle)\right] \leq \mathrm{negl}(\lambda),$$

where the probability is over the challenger preparing $|\psi\rangle$, sampling h, and measuring y and v as described in $\mathsf{Exp}_3(b)$ to produce the left-over state $|\psi_{h,y,v}\rangle$.

Proof. This follows directly from the assumed $(\mathcal{D}, \mathcal{M})$-target-collision resistance of $\mathcal{H}$, since the above probability is exactly $\Pr\left[\mathsf{TargetCollRes}_{\mathcal{H},\mathcal{A},\mathcal{D},\mathcal{M},\lambda} = 1\right]$.

Claim. The probability that the challenger accepts the deletion certificate π in Step 4 of $\mathsf{Exp}_3(b)$ and the subsequent projective measurement on system C fails (returns the second outcome) is negligible.

Proof. This follows directly from Section 3.2, which implies that except with negligible probability, the register C is in the state

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

at the time the challenger applies the projective measurement.

For any experiment $\mathsf{Exp}_i(b)$, we define the advantage

$$\mathsf{Adv}(\mathsf{Exp}_i) \coloneqq |\Pr[\mathsf{Exp}_i(0) = 1] - \Pr[\mathsf{Exp}_i(1) = 1]|.$$

Claim.

$$\mathsf{Adv}(\mathsf{Exp}_2) = 0.$$

Proof. First note that in the case that the challenger rejects because either the deletion certificate is invalid or their projection fails, the experiment does not involve b, and thus the advantage of the adversary is 0. Second, in the case that the challenger's projection succeeds, the register C is either in the state

$$\frac{1}{\sqrt{2}}(|0\rangle + (-1)^{\langle \pi, z\rangle}|1\rangle) \qquad \text{or} \qquad \frac{1}{\sqrt{2}}(|0\rangle - (-1)^{\langle \pi, z\rangle}|1\rangle)$$

for some $z \in \{0,1\}^k$, and thereby completely unentangled from the rest of the system. Notice that the challenger's measurement of system C with outcome c' results in a uniformly random bit, which completely masks b. Therefore, the experiment is also independent of b in this case, and thus the adversary's overall advantage in Exp_2 is 0.

Next, we argue the following.

Claim.

$$|\mathsf{Adv}(\mathsf{Exp}_2) - \mathsf{Adv}(\mathsf{Exp}_1)| \leq \mathrm{negl}(\lambda).$$

Proof. Recall that Section 3.2 shows that the projective measurement performed by the challenger in Step 4 of Exp_3 succeeds with overwhelming probability. We now argue that the same is also true in Exp_2. Suppose for the sake of contradiction that there is a non-negligible difference between the success probabilities of the measurement. We now show that this implies the existence of an efficient distinguisher $\mathcal{A}'$ that breaks the $(\mathcal{D}, \mathcal{M})$-target-collapsing property of the hash family $\mathcal{H} = \{H_\lambda\}_{\lambda \in \mathbb{N}}$.

$\mathcal{A}'$ receives (y, h) and a state on register X from its challenger. Next, it computes $M[h]$ on X into a fresh k-qubit register V, samples a random string $z \xleftarrow{\$} \{0,1\}^k$, prepares a $|+\rangle$ state in system C, applies a controlled-Z^z operation from C to V, and then uncomputes register V by again applying $M[h]$ from X to V. Then, it runs $\mathcal{A}$ on (y, h, X), which outputs a certificate π.

Finally, $\mathcal{A}'$ applies the following projective measurement to system C:

$$\left\{|\phi^z_\pi\rangle\langle\phi^z_\pi|, I - |\phi^z_\pi\rangle\langle\phi^z_\pi|\right\} \qquad \text{where} \qquad |\phi^z_\pi\rangle \coloneqq \frac{1}{\sqrt{2}}\left(|0\rangle + (-1)^{\langle \pi, z\rangle}|1\rangle\right),$$

and outputs 1 if the measurement succeeds and 0 otherwise. If there is a non-negligible difference in success probabilities of this measurement between $\mathsf{Exp}_3(b)$ and $\mathsf{Exp}_2(b)$ (for any $b \in \{0,1\}$), then $\mathcal{A}'$ breaks $(\mathcal{D}, \mathcal{M})$-target-collapsing of $\mathcal{H}$.

Now, recall that $\mathsf{Exp}_2(b)$ is identical to $\mathsf{Exp}_1(b)$, except that the challenger applies an additional a measurement in Step 4. Because the measurement succeeds with overwhelming probability, it follows from Gentle Measurement that the advantage of the adversary must remain the same up to a negligible amount. This proves the claim.

Claim.

$$\mathsf{Adv}(\mathsf{Exp}_1) = \mathsf{Adv}(\mathsf{Exp}_0)/2.$$

Proof. First note that in $\mathsf{Exp}_1(b)$, we can imagine measuring register C to obtain c' and aborting if $c' \neq b$ before the challenger sends any information to the adversary. This follows because register C is disjoint from the adversary's registers. Next, by the random Pauli-Z twirl property, we have the following guarantees about the state on system X given to the adversary in $\mathsf{Exp}_1(b)$.

- In the case $c' = b = 0$, the reduced state on register X is $|\psi_{h,y}\rangle$.
- In the case that $c' = b = 1$, the reduced state on register X is a mixture over $|\psi_{h,y,v}\rangle$ where v is the result of measuring register V in the computational basis.

Thus, this experiment is identical to $\mathsf{Exp}_0(b)$, except that we decide to abort and output a uniformly random bit b' with probability $1/2$ at the beginning of the experiment.

Putting everything together, we have that $\mathsf{Adv}(\mathsf{Exp}_0) \leq \mathrm{negl}(\lambda)$, which completes the proof.

References

[1] Scott Aaronson. "Quantum copy-protection and quantum money". In: *2009 24th Annual IEEE Conference on Computational Complexity*. IEEE. 2009, pp. 229–242.

[2] Shweta Agarwal et al. "Public Key Encryption with Secure Key Leasing". In: *Eurocrypt 2023 (to appear)*. 2023.

[3] Miklós Ajtai. "Generating Hard Instances of Lattice Problems (Extended Abstract)". In: *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*. Ed. by Gary L. Miller. ACM, 1996, pp. 99–108. DOI: 10.1145/237814.237838. URL: https://doi.org/10.1145/237814.237838.

[4] Prabhanjan Ananth and Rolando L. La Placa. "Secure Software Leasing". In: *Advances in Cryptology – EUROCRYPT 2021*. Ed. by Anne Canteaut and François-Xavier Standaert. Cham: Springer International Publishing, 2021, pp. 501–530. ISBN: 978-3-030-77886-6.

[5] Prabhanjan Ananth, Alexander Poremba, and Vinod Vaikuntanathan. *Revocable Cryptography from Learning with Errors*. Cryptology ePrint Archive, Paper 2023/325. https://eprint.iacr.org/2023/325. 2023. URL: https://eprint.iacr.org/2023/325.

[6] James Bartusek and Dakshita Khurana. *Cryptography with Certified Deletion*. Cryptology ePrint Archive, Paper 2022/1178. https://eprint.iacr.org/2022/1178. 2022. URL: https://eprint.iacr.org/2022/1178.

[7] James Bartusek et al. *Obfuscation and Outsourced Computation with Certified Deletion*. Cryptology ePrint Archive, Paper 2023/265. 2023. URL: https://eprint.iacr.org/2023/265.

[8] C. H. Bennett and G. Brassard. "Quantum cryptography: Public key distribution and coin tossing". In: *Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing*. Bangalore, 1984, p. 175.

[9] Anne Broadbent and Rabib Islam. "Quantum Encryption with Certified Deletion". In: *Lecture Notes in Computer Science* (2020), pp. 92–122. ISSN: 1611-3349. DOI: 10.1007/978-3-030-64381-2_4. URL: http://dx.doi.org/10.1007/978-3-030-64381-2_4.

[10] Ran Canetti et al. "Adaptively Secure Multi-Party Computation". In: *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*. STOC '96. Philadelphia, Pennsylvania, USA: Association for Computing Machinery, 1996, pp. 639–648. ISBN: 0897917855. DOI: 10.1145/237814.238015. URL: https://doi.org/10.1145/237814.238015.

[11] Shujiao Cao and Rui Xue. "The Gap Is Sensitive to Size of Preimages: Collapsing Property Doesn't Go Beyond Quantum Collision-Resistance for Preimages Bounded Hash Functions". In: Springer-Verlag, 2022.

[12] Marcel Dall'Agnol and Nicholas Spooner. *On the necessity of collapsing*. Cryptology ePrint Archive, Paper 2022/786. https://eprint.iacr.org/2022/786. 2022. URL: https://eprint.iacr.org/2022/786.

[13] Iftach Haitner et al. "Reducing Complexity Assumptions for Statistically-Hiding Commitment". In: *Journal of Cryptology* 22.3 (2009), pp. 283–310. DOI: 10.1007/s00145-007-9012-8. URL: https://doi.org/10.1007/s00145-007-9012-8.

[14] Minki Hhan, Tomoyuki Morimae, and Takashi Yamakawa. "From the Hardness of Detecting Superpositions to Cryptography: Quantum Public Key Encryption and Commitments". In: *Eurocrypt 2023 (to appear)*. 2023.

[15] Taiga Hiroka et al. *Certified Everlasting Functional Encryption*. Cryptology ePrint Archive, Paper 2022/969. https://eprint.iacr.org/2022/969. 2022. URL: https://eprint.iacr.org/2022/969.

[16] Taiga Hiroka et al. "Certified Everlasting Zero-Knowledge Proof for QMA". In: *Advances in Cryptology - CRYPTO 2022 - 42nd Annual International Cryptology Conference, CRYPTO 2022, Santa Barbara, CA, USA, August 15-18, 2022, Proceedings, Part I*. Ed. by Yevgeniy Dodis and Thomas Shrimpton. Vol. 13507. Lecture Notes in Computer Science. Springer, 2022, pp. 239–268. DOI: 10.1007/978-3-031-15802-5_9. URL: https://doi.org/10.1007/978-3-031-15802-5%5C_9.

[17] Taiga Hiroka et al. "Quantum Encryption with Certified Deletion, Revisited: Public Key, Attribute-Based, and Classical Communication". In: *Advances in Cryptology - ASIACRYPT 2021 - 27th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 6-10, 2021, Proceedings, Part I*. Ed. by Mehdi Tibouchi and Huaxiong Wang. Vol. 13090. Lecture Notes in Computer

Science. Springer, 2021, pp. 606–636. DOI: 10.1007/978-3-030-92062-3_21. URL: https://doi.org/10.1007/978-3-030-92062-3%5C_21.

[18] Stanisław Jarecki and Anna Lysyanskaya. "Adaptively Secure Threshold Cryptography: Introducing Concurrency, Removing Erasures". In: *Proceedings of the 19th International Conference on Theory and Application of Cryptographic Techniques*. EUROCRYPT'00. Bruges, Belgium: Springer-Verlag, 2000, pp. 221–242. ISBN: 3540675175.

[19] Qipeng Liu and Mark Zhandry. "Revisiting Post-quantum Fiat-Shamir". In: *Advances in Cryptology – CRYPTO 2019*. Ed. by Alexandra Boldyreva and Daniele Micciancio. Cham: Springer International Publishing, 2019, pp. 326–355. ISBN: 978-3-030-26951-7.

[20] Jörn Müller-Quade and Dominique Unruh. "Long-Term Security and Universal Composability". In: *Theory of Cryptography*. Ed. by Salil P. Vadhan. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 41–60. ISBN: 978-3-540-70936-7.

[21] Alexander Poremba. "Quantum Proofs of Deletion for Learning with Errors". In: *14th Innovations in Theoretical Computer Science Conference, ITCS 2023, January 10-13, 2023, MIT, Cambridge, Massachusetts, USA*. Ed. by Yael Tauman Kalai. Vol. 251. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023, 90:1–90:14. DOI: 10.4230/LIPIcs.ITCS.2023.90. URL: https://doi.org/10.4230/LIPIcs.ITCS.2023.90.

[22] Daniel R. Simon. "Finding collisions on a one-way street: Can secure hash functions be based on general assumptions?" In: *Advances in Cryptology — EUROCRYPT'98*. Ed. by Kaisa Nyberg. Berlin, Heidelberg: Springer Berlin Heidelberg, 1998, pp. 334–345. ISBN: 978-3-540-69795-4.

[23] Damien Stehlé et al. "Efficient Public Key Encryption Based on Ideal Lattices". In: *Advances in Cryptology – ASIACRYPT 2009*. Ed. by Mitsuru Matsui. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 617–635. ISBN: 978-3-642-10366-7.

[24] Marco Tomamichel and Anthony Leverrier. "A largely self-contained and complete security proof for quantum key distribution". In: *Quantum* 1 (July 2017), p. 14. ISSN: 2521-327X. DOI: 10.22331/q-2017-07-14-14. URL: https://doi.org/10.22331/q-2017-07-14-14.

[25] Dominique Unruh. "Collapse-Binding Quantum Commitments Without Random Oracles". In: *Advances in Cryptology – ASIACRYPT 2016*. Ed. by Jung Hee Cheon and Tsuyoshi Takagi. Berlin, Heidelberg: Springer Berlin Heidelberg, 2016, pp. 166–195. ISBN: 978-3-662-53890-6.

[26] Dominique Unruh. "Computationally Binding Quantum Commitments". In: *Advances in Cryptology – EUROCRYPT 2016*. Ed. by Marc Fischlin and Jean-Sébastien Coron. Berlin, Heidelberg: Springer Berlin Heidelberg, 2016, pp. 497–527. ISBN: 978-3-662-49896-5.

[27] Dominique Unruh. "Revocable Quantum Timed-Release Encryption". In: *J. ACM* 62.6 (Dec. 2015). ISSN: 0004-5411. DOI: 10.1145/2817206. URL: https://doi.org/10.1145/2817206.

[28] Stephen Wiesner. "Conjugate Coding". In: *SIGACT News* 15.1 (Jan. 1983), pp. 78–88. ISSN: 0163-5700. DOI: 10.1145/1008908.1008920. URL: https://doi.org/10.1145/1008908.1008920.

[29] Jun Yan. "General Properties of Quantum Bit Commitments (Extended Abstract)". In: *Advances in Cryptology – ASIACRYPT 2022*. Ed. by Shweta Agrawal and Dongdai Lin. Cham: Springer Nature Switzerland, 2022, pp. 628–657. ISBN: 978-3-031-22972-5.

[30] Mark Zhandry. "New Constructions of Collapsing Hashes". In: *Advances in Cryptology - CRYPTO 2022 - 42nd Annual International Cryptology Conference, CRYPTO 2022, Santa Barbara, CA, USA, August 15-18, 2022, Proceedings, Part III*. Ed. by Yevgeniy Dodis and Thomas Shrimpton. Vol. 13509. Lecture Notes in Computer Science. Springer, 2022, pp. 596–624. DOI: 10.1007/978-3-031-15982-4_20. URL: https://doi.org/10.1007/978-3-031-15982-4%5C_20.