

Classical Proofs for the Quantum Collapsing Property of Classical Hash Functions

Serge Fehr

Centrum Wiskunde & Informatica – CWI (serge.fehr@cw.nl), and
Mathematical Institute, Leiden University, The Netherlands

Abstract. Hash functions are of fundamental importance in theoretical and in practical cryptography, and with the threat of quantum computers possibly emerging in the future, it is an urgent objective to understand the security of hash functions in the light of potential future quantum attacks. To this end, we reconsider the *collapsing property* of hash functions, as introduced by Unruh, which replaces the notion of *collision resistance* when considering quantum attacks. Our contribution is a formalism and a framework that offers significantly simpler proofs for the collapsing property of hash functions. With our framework, we can prove the collapsing property for hash domain extension constructions entirely by means of decomposing the iteration function into suitable elementary composition operations. In particular, given our framework, one can argue *purely classically* about the quantum-security of hash functions; this is in contrast to previous proofs which are in terms of sophisticated quantum-information-theoretic and quantum-algorithmic reasoning.

1 Introduction

Background. Given the threat of possible future quantum computing capabilities, it is an important and urgent objective to evaluate the security of classical cryptographic schemes against quantum attacks. There are different places where security can break down when using quantum computing techniques to attack a cryptographic scheme that was designed to withstand standard classical attacks. The most prominent place is the computational hardness *assumption*, which is typically well justified to hold for classical models of computation but may be false with respect to quantum computation. Another place is the security *proof*, which may use techniques that fail to work in the context of a quantum attacker, like proofs that rely on rewinding techniques. Finally, another place where things can go wrong is the security *definition*, which may not capture anymore what it is supposed to capture when allowing quantum attacks.

An example of the latter is the computational binding property of a commitment scheme. Our intuitive understanding of what a commitment should achieve is that once a commitment is “on the table” there should be no freedom left for the (computationally bounded) committer in choosing the value to which he can open the commitment. The formal definition of the binding property expresses this requirement by demanding that no (computationally bounded) dishonest

committer should be able to open a commitment in two distinct ways. While for classical committers this captures precisely what we want, it fails to do so for quantum committers. Indeed, a quantum committer can potentially open a commitment to *one* value that he *freely chooses* after he has put the commitment “on the table”, without contradicting the requirement of being unable to produce *two* distinct openings; this is because producing the opening information may involve a destructive quantum measurement that can only be applied once.

We stress that being able to open a given commitment to an arbitrary value that one can freely choose renders a commitment scheme useless in essentially all applications. So, when considering the security of commitment schemes against quantum attacks, it is essential that one uses a *stronger notion of security* than the standard computational binding property extended to quantum attackers.

A similar and related example is the collision resistance of hash functions. Also here, in the presence of a quantum attacker, the standard formal requirement that it should be computationally hard to produce two colliding inputs does not capture our intuitive understanding of a hash value as acting as a “fingerprint” that removes any freedom in the message to which it fits. As such, also here, when considering security against quantum attacks, the standard security notion, i.e. collision resistance, needs to be replaced by something stronger.

The Collapsing Property. Unruh [5] proposed the notion of *collapsing*; in the context of commitment schemes as a counterpart for the computational binding property when considering quantum attacks, and in the context of hash functions as a counterpart for collision resistance. In essence, for hash functions, the collapsing property requires that for any computationally bounded adversary that outputs a hash value together with a *quantum superposition* of corresponding preimages, he should not be able to tell if the superposition gets measured or not. The details of the notion, and why it indeed restores the right security properties when considering quantum attacks, are not so important for the discussion here. In terms of achievability, Unruh proved that the random oracle is collapsing as a hash function, and thus that simple hash-function-based commitment schemes are collapsing in the random oracle model. In the context of hash functions, he proved in a follow-up work [6] that the *Merkle-Damgård* construction for hash functions is collapsing (under some mild restriction on the padding) if the underlying compression function is. Given that the random oracle is collapsing, this in particular implies that the Merkle-Damgård construction is collapsing in the random oracle model, and thus gives heuristic evidence that certain practical hash functions like SHA-2 are collapsing. Recently, Czajkowski *et al.* [2] showed a similar result for the *Sponge* construction [1], which for instance underlies the hash function standard SHA-3: the Sponge construction is collapsing if both parts of the underlying round function, i.e., the so-called *inner* and *outer* parts, are collapsing, and if the inner part is “zero-preimage resistant”.¹

¹ This again implies security in the random oracle model, but a subtle issue here is that if the round function is efficiently invertible then the assumptions on the two parts are not satisfied. Hence, this is not so strong evidence yet that e.g. SHA-3 is collapsing.

Our Contribution. In this work, we introduce a new *formalism* and a new *framework* for arguing about the collapsing property of (hash) functions. The advantage of our new approach is that it allows for significantly simpler proofs compared to the previous work above.

At the heart of our new *formalism* is a pseudo-metric that abstracts away computational aspects, and which allows for an “algebraic” formulation of the collapsing property. This in turn allows for simple proofs of basic *composability results* for the collapsing property. Some of those have already been claimed and proven in the work mentioned above; however, our proofs are much simpler. For instance, proving that the collapsing property is preserved under nested composition takes 2 full pages in [5] (see Lemma 27 in the full version of [5]), with various quantum circuits depicted; our proof (see Lemma 5) is a few lines. The main reason for this difference lies in the “algebraic” nature of our formulation, compared to the “algorithmic” approach used in prior work. This means that instead of specifying quantum reduction algorithms and arguing that they “do the job”, our proofs are almost entirely by means of *term-manipulations*, where we manipulate the terms of interest by using a small set of basic rules that come along with our formalism. This not only results in very compact proofs, these proofs are also mathematically very clean in that in every term-manipulation step we can — and typically do — specify what basic rule was used.

These composability results for the collapsing property, together with a couple of basic features when “disallowing” certain inputs, form what we call our *framework*. With this framework, proving the collapsing property of hash domain extensions boils down to decomposing the iteration function under consideration into a few simple composition operations.

We demonstrate this new proof methodology on various examples. Applied to Merkle-Damgård, we obtain a proof of the collapsing property without any restriction on the padding as in [6], but with the additional assumption on the compression function to be “*iv*-preimage resistant” (which is satisfied in the random oracle model). We can also recover Unruh’s original result, which requires a restriction on the padding but avoids the “*iv*-preimage resistance”. By adding a counter and “salt” to the compression function but otherwise using the same kind of reasoning, we get a proof of the collapsing property of *HAIFA* [3], as proposed by Biham and Dunkelman. Applied to the *Sponge* construction, we recover the result from [2] up to an insignificant difference in the exact parameter.

The distinguishing feature of our proofs lies in their conceptual simplicity and low technical complexity. Our proofs are entirely in terms of decomposing the iteration function into elementary composition operations that are ensured to preserve the collapsing property. In particular, our proofs are purely classical. In contrast, the proofs provided in [6, 2] are in terms of lengthy hybrid arguments that consider sequences of “quantum games” and in terms of quantum information theoretic arguments and quantum reduction algorithms for reasoning that every game in the sequence behaves similarly to its predecessor.

As such, even though the collapsing property of *HAIFA* is new, we consider our main contribution more in terms of offering a simple understanding of *why*

certain hash function are collapsing, and in providing a tool to *easily* check if similar results also hold for other hash functions (as we demonstrate on HAIFA).

The Framework in Action. To give a better idea, we illustrate here on the various examples how our framework enables to argue for the collapsing property by means of decomposing the iteration function into suitable elementary decomposition operations, and thus in particular by means of purely classical reasoning. We challenge the reader to compare our proofs with those in [6, 2].

Merkle-Damgård. The Merkle-Damgård hash of a message x_1, \dots, x_i , consisting of i blocks, is given by $IH_i(x_1, \dots, x_i)$, where IH_i is iteratively defined as

$$IH_i(x_1, \dots, x_i) := f(IH_{i-1}(x_1, \dots, x_{i-1}), x_i)$$

with $IH_0() = iv$. The round function f is assumed to be collapsing. We observe that IH_i is the *nested composition* of f with the *concurrent composition* of IH_{i-1} with the identity $x_i \mapsto x_i$. Our framework ensures that these compositions preserve the collapsing property; thus, by recursive application, given that IH_0 is trivially collapsing, we get that IH_L is collapsing for every *fixed* L , and hence the Merkle-Damgård hash is collapsing when restricted to inputs of *fixed* size.

In order to deal with messages of variable size, we allow in the definition of $IH_i(x_1, \dots, x_i)$ the left-most message blocks to be “empty”, i.e., x_1 up to some x_j may be \perp , and we set $IH_i(\perp, \dots, \perp) := iv$ (for any i) and keep to recursive definition above if $x_i \neq \perp$. This extended version of IH_i is then the *disjoint union* of the trivial function $\{\perp^i\} \rightarrow \{iv\}$ and the restriction of IH_i to inputs different than \perp^i , if we “disallow” non- \perp^i inputs that are mapped to iv .² Thus, as long as we “disallow” such inputs (which is something our framework can capture), we still have that the recursive definition of IH_i decomposes into composition operations that are covered by our framework, and thus we can conclude that IH_L is collapsing for every fixed L , but now for inputs that may have \perp -prefixes, i.e., variable length. Finally, by the assumed “ iv -preimage resistant” of f , inputs ($\neq \perp^i$) that IH_i maps to iv are hard to find, and therefore “disallowing” those has no noticeable effect.

HAIFA. The HAIFA hash function is a variant of Merkle-Damgård that includes a counter in the iteration function, and it uses a “salt” (which we though treat as ordinary input). Formally,

$$IH_i(salt, x_1, \dots, x_i) := f(salt, IH_{i-1}(salt, x_1, \dots, x_{i-1}), x_i, i).$$

Here, we can reason exactly as above, except that now the iteration function is a *nested composition* of the function $f(\cdot, \cdot, \cdot, i)$, which is collapsing if f is, with the *parallel composition* of the projection function $(salt, x_1, \dots, x_i) \mapsto salt$ with the *concurrent composition* of IH_{i-1} with the identity function $x_i \mapsto x_i$. All

² The latter is because (our notion of) the disjoint union of two functions requires not only the two respective domains but also the two respective ranges to be disjoint.

these composition operations are covered by our framework, and so the collapsing property follows as for the original Merkle-Damgård construction, assuming again that f is “ iv -preimage resistance” in case of arbitrary length messages.

Sponge. The Sponge hash³ of a message x_1, \dots, x_i of i blocks is given by $S_i^0(x_1, \dots, x_i)$, where S_i^b is iteratively defined as

$$S_i^b(x_1, \dots, x_i) := f^b(S_{i-1}^0(x_1, \dots, x_{i-1}) \oplus x_i, S_{i-1}^1(x_1, \dots, x_{i-1}))$$

for $b \in \{0, 1\}$, with $S_0^0() = 0 = S_0^1()$, and it is assumed that both components of the round function $f = (f^0, f^1)$ are collapsing. Here, S_i^b is the *nested composition* of f^b with a function that is yet another *composition* of the functions S_{i-1}^0 and S_{i-1}^1 , and our framework immediately ensures that S_i^0 and S_i^1 stay collapsing as long as S_{i-1}^1 is. Thus, again, the iteration function decomposes into composition operations that are ensured to preserve the collapsing property, and so by recursive application we get that S_1^1, \dots, S_{L-1}^1 and eventually S_L^0 are collapsing. The only difference to above is that here, we have to set $S_i^b(\perp, \dots, \perp) := S_0^b() = 0$ to ensure that S_L^0 acts correctly on messages of smaller block size, i.e., that $S_j^b(x_1, \dots, x_j) = S_L^b(\perp, \dots, \perp, x_1, \dots, x_j)$. As a consequence, for the recursive reasoning, to have S_i^1 be the disjoint union of the trivial function $\{\perp^i\} \rightarrow \{0\}$ and the restriction of S_i^1 to non- \perp^i inputs, we need to “disallow” inputs ($\neq \perp^i$) which S_i^1 maps to 0; this has no noticeable effect though if f^1 is “zero-preimage resistant”.

2 Preliminaries

2.1 Basic Quantum Formalism

Knowledge of basic concepts of quantum information science is necessary in order to prove “correctness” of our framework (but not to apply the framework); we fix here some notation and conventions, which both are not fully standard.

Typically, the *state* of a quantum system with state space \mathcal{H} is given by a density matrix ρ , i.e., by a trace-1 positive-semidefinite matrix that acts on \mathcal{H} , and a *quantum operation* is expressed by a CPTP map \mathbb{T} which maps a state ρ to a new state $\mathbb{T}(\rho)$ over a possibly different state space. In this work, for technical reasons, we allow states to be *subnormalized*, and we consider the more general notion of completely-positive *trace-nonincreasing* (CPTN) maps, which are of the form $\mathbb{T} = \sum_i \mathbb{T}_i$ with $\mathbb{T}_i : \rho \mapsto T_i \rho T_i^\dagger$ and $\sum_i T_i^\dagger T_i \leq I$ (the identity on \mathcal{H}).⁴

For the purpose of this work, a *measurement* is a CPTN map $\mathbb{P} = \sum_i \mathbb{P}_i$ with $\mathbb{P}_i : \rho \mapsto P_i \rho P_i^\dagger$ as above, but with the restriction that the P_i ’s are mutually orthogonal Hermitian *projections* on \mathcal{H} . If \mathbb{P} is in fact a CPTP map, i.e., $\sum_i P_i =$

³ For simplicity, we consider *one* block of output only; multiple output blocks are argued by means of composition too.

⁴ This can be understood in that quantum operations may “abort”, and the trace $\text{tr}(\rho) \leq 1$ expresses the probability that the process that produces ρ does not abort.

I , then we speak of a *total* measurement, and otherwise of a *partial* measurement. The individual “components” \mathbb{P}_i of such a (partial or total) measurement are sometimes also referred to as *measurements with post-selection*.

We write $\Pr[\mathbb{P}(\rho) = i]$ for $\text{tr} \circ \mathbb{P}_i(\rho) = \text{tr}(P_i \rho P_i)$, i.e., the probability that “outcome i is observed”. An elementary property of any (*projective*, as considered here) measurement \mathbb{P} , is Winter’s “gentle-measurement lemma” [8], which captures that the measurement does not disturb the state much if the outcome is almost certain. Formally,⁵ for any state ρ and any $\beta \geq 0$:

$$\exists i : \Pr[\mathbb{P}(\rho) = i] \geq \text{tr}(\rho) - \beta \implies \delta(\mathbb{P}(\rho), \rho) \leq \sqrt{\beta} + \beta. \quad (1)$$

where δ is the *trace distance*, given by $\delta(\rho, \sigma) := \frac{1}{2} \|\rho - \sigma\|_{tr}$.

Different quantum systems are identified by means of “labels” X, Y etc., and we write ρ_X for the state of system X and \mathcal{H}_X for its state space, etc. For a CPTN map \mathbb{T} , we may write \mathbb{T}_X to emphasize that it acts on system X , and $\mathbb{T}_{X \rightarrow X'}$ to additionally emphasize that it maps into system X' . For simplicity, we tend to write $\rho_{\mathbb{T}(X)Y}$ rather than $(\mathbb{T}_X \otimes \mathbb{I}_Y)(\rho_{XY})$.

For any state space we consider a fixed orthonormal basis, referred to as the *computational* basis. For state spaces \mathcal{H}_X and \mathcal{H}_Y with respective computational bases $\{|x\rangle\}_{x \in \mathcal{X}}$ and $\{|y\rangle\}_{y \in \mathcal{Y}}$, we associate to any function $f : \mathcal{X} \rightarrow \mathcal{Y}$ the CPTP “evaluation” map $\mathbb{E}[f]_{X \rightarrow XY} : \rho \mapsto V[f] \rho V[f]^\dagger$ given by the isometry $V[f] : |x\rangle \mapsto |x\rangle|f(x)\rangle$. Here, we also write $\rho_{Xf(X)Z}$ instead of $\rho_{\mathbb{E}[f](X)Z}$.⁶ We note that $\mathbb{E}[f]$ admits a *left inverse*, i.e., a CPTP map $\mathbb{E}^{inv}[f]_{XY \rightarrow X}$ such that $\mathbb{E}^{inv}[f] \circ \mathbb{E}[f] = \mathbb{I}_X$.

The composition $\text{tr}_Y \circ \mathbb{E}[f]_{X \rightarrow XY}$ of a CPTP evaluation map with the partial trace tr_Y equals the measurement $\mathbb{M}[f] = \sum_y \mathbb{M}[f = y]$, where $\mathbb{M}[f = y]$ is the CPTN map given by the projection into the span of $\{|x\rangle \mid f(x) = y\}$. To simplify notation, we may also write ρ_{XfZ} instead of $\rho_{\mathbb{M}[f](X)Z}$, and, similarly, $\rho_{Xf=yZ}$ instead of $\rho_{\mathbb{M}[f=y](X)Z}$.

The usual “measurement in the computational basis”, given by the projections $|x\rangle\langle x|$, is simply denoted by \mathbb{M} . For lighter notation, we often use $\overline{(\cdot)}$ instead of \mathbb{M} and write $\rho_{\overline{XY}}$ instead of $\rho_{\mathbb{M}(X)Y}$. A quantum system X of a (possibly) joint state ρ_{XY} is called *classical* if $\rho_{\overline{XY}} = \rho_{XY}$.

When the state is clear from the context, then we may do the “arithmetic” on the labels. For instance, using this convention, we can then say that any state ρ_{XZ} satisfies

$$\overline{X}f(X)Z = \overline{X}f(\overline{X})Z = \overline{X}\overline{f(\overline{X})}Z = \overline{X}\overline{f(\overline{X})}Z, \quad (2)$$

to express that \mathbb{M}_X and $\mathbb{E}[f]_X$ commute, and that $f(\overline{X})$ is classical given that \overline{X} is. Similarly, we may then write $\Pr[\mathbb{M}[f](X) = y] = \Pr[\mathbb{M}(f(X)) = y] = \Pr[f(\overline{X}) = y]$, which may be interpreted differently but coincide.

⁵ This bound can e.g. be derived from [9]. [2] claims the bound $\sqrt{\beta}$, but their proof has a small flaw; fixing it gives $\sqrt{2\beta}$ instead (but only works for *total* measurements).

⁶ A subtle issue with this notation is that $\text{tr}_Y(\rho_{Xf(X)Z}) \neq \rho_{XZ}$, but rather $= \rho_{\mathbb{M}[f](X)Z}$ (see below).

2.2 Randomized Functions and States, and their Complexity

In Appendix A, we offer a formal discussion of *randomized functions*, *randomized CPTN maps*, and *randomized quantum states*. As one would expect, these are simply functions, CPTN maps and states that depend on some *global randomness* r , which is randomly chosen once and for all from some finite set \mathcal{R} .

Informally, when considering randomized functions, one can make the following distinction. In one case, r is given as input to the function f (or to the algorithm that computes f , if you prefer); one then typically speaks of *keyed* or *seeded* functions. In the other case, f makes queries to an *oracle* that computes every reply dependent on r , in which case one refers to f as an *oracle function*. A similar distinction can be made for randomized CPTN maps, and thus for randomized states, which are simply randomized CPTN maps that act on the trivial state space \mathbb{C} .

Formally, the way the two variants differ is by the way *complexity* is captured: for keyed functions one considers the *computational complexity* of computing the function whereas for oracle functions one considers the *query complexity*.

Our results apply to both variants in that we consider an *abstract* complexity measure \mathfrak{c} that assigns to every randomized function f a non-negative integer $\mathfrak{c}(f)$, also denoted \mathfrak{c}_f , and similarly for randomized CPTN maps, and which satisfies natural properties that one would expect from a complexity measure. The details of this are given in Appendix B. The computational complexity and the query complexity are then just specific instantiations.

2.3 The Distinguishing Advantage

The following parameterized indistinguishability measure, and our understanding of it as an abstract metric, is one of the central notions of our formalism.

Definition 1. For randomized states ρ_X and ρ_Y (with randomness r) over a common Hilbert space $\mathcal{H}_X = \mathcal{H}_Y$, and for any non-negative integer q , we set

$$\begin{aligned} \delta_q(\rho_X, \rho_Y) &:= \sup_{\mathbb{T}} \frac{1}{|\mathcal{R}|} \sum_r |\Pr[\mathbb{M}(\mathbb{T}(X))=0] - \Pr[\mathbb{M}(\mathbb{T}(Y))=0]| \\ &= \sup_{\mathbb{T}} \frac{1}{|\mathcal{R}|} \sum_r \delta(\mathbb{M} \circ \mathbb{T}(\rho_X), \mathbb{M} \circ \mathbb{T}(\rho_Y)), \end{aligned}$$

where the supremum is over all randomized CPTN maps \mathbb{T} (with randomness r) that map into the two-dimensional qubit state space and have complexity $\mathfrak{c}(\mathbb{T}) \leq q$ and, by convention, \mathbb{M} is the measurement in the computational basis.

Following the convention of doing the “arithmetic” on the labels, we typically write $\delta_q(X, Y)$ instead of $\delta_q(\rho_X, \rho_Y)$. Also, we write $\delta_q(X, Y|Z)$ as a short hand for $\delta_q(\rho_{XZ}, \rho_{YZ})$.

We emphasize that δ_q is a *pseudometric*: it is non-negative, symmetric, and satisfies triangle inequality, but it may potentially vanish for non-identical states.

Furthermore, δ_q is upper bounded by the trace distance δ , and it coincides with δ in case $q = \infty$, i.e., there is no restriction on $\mathfrak{c}(\mathbb{T})$. Finally, δ_q inherits several properties from the ordinary trace distance, which can easily be verified. For instance, it is *monotone* under randomized CPTN maps as

$$\delta_q(\mathbb{T}(X), \mathbb{T}(Y)) \leq \delta_{q+\mathfrak{c}(\mathbb{T})}(X, Y),$$

and for any randomized CPTN map $\mathbb{T} = \sum_i \mathbb{T}_i$, we have *subadditivity* as

$$\delta_q(\mathbb{T}(X), \mathbb{T}(Y)) \leq \sum_i \delta_q(\mathbb{T}_i(X), \mathbb{T}_i(Y)).$$

To simplify terminology, from now on we drop on the word “randomized” and take it as understood that functions, CPTN maps and states may be randomized, either in the form of keyed functions or as oracle functions, etc.

3 The Collapsing Property

We state here (a slight variation of) the definition of the collapsing property of functions, as proposed by Unruh [5], but using the formalism introduced above. In Section 3.2 we then discuss the straightforward extension to *partial* functions, which will turn out to be useful, and in Section 3.3 we show that the collapsing property behaves nicely under various composition operations. These composability results are all rather natural, and — with our formalism! — have simple short proofs. All together, this section then stands as “the framework” that we propose for arguing about the collapsing property of hash functions.

3.1 The Definition

The original formulation of the collapsing property for a function h is by means of two “games”, where an “adversary” produces a (normalized) state ρ_{XYE} of a certain form, namely Y must be classical and equal to $h(X)$, and then in one game X is measured in the computation basis whereas in the other game it is left untouched instead, and the definition requires that it should be hard for any “distinguisher” to distinguish between the two games.

As for the notion of collision resistance, the collapsing property is meaningful only for randomized functions h .⁷ In case of a *keyed* variant of such a function, one can aim for *conditional* results that state that h is collapsing (against computationally bounded adversaries) under some computational hardness assumption. In case of an *oracle* function and aiming for *unconditional* results, there is no exploitable effect in restricting the computational power of the parties, as long as the query complexity is limited. Our approach of using an abstract complexity notion allows us to cover *both* these settings simultaneously.

⁷ See the discussion in Appendix C for an exception to the rule.

Our formal definition of the collapsing property is given below. Compared to the original definition by Unruh (which comes in a couple of different flavors, which we discuss in Appendix C), we use a somewhat different terminology and formalism. For instance we do not explicitly speak of “games”, and instead of quantifying over the possible adversaries we quantify over the states that may possibly be prepared by an adversary, and the quantification over the distinguishers is absorbed into the pseudometric δ_q . These modifications to the mathematical language have obviously no effect on the notion. There are a few more differences compared to the definition proposed by Unruh, but they all have no more than a small quantitative effect, as we discuss below.

Definition 2. A function $h : \mathcal{X} \rightarrow \mathcal{Y}$ is called $\varepsilon(q)$ -collapsing if

$$\mathbf{cAdv}[h](q) := \sup_{\rho_{XYE}} \delta_q(X, \bar{X} | \bar{Y}E) \leq \varepsilon(q)$$

for all q , where the supremum is over all states⁸ $\rho_{XYE} = \rho_{Xh(X)E}$ with complexity $\mathbf{c}(\rho_{XYE}) \leq q$. The measure $\mathbf{cAdv}[h]$ is called the collapsing advantage of h .

Beyond the change in mathematical language, another difference is that in the original definition the system Y of the state ρ_{XYE} , as produced by the adversary, is required to be classical, whereas in Definition 2 we allow it to be non-classical but then “make it classical” by measuring it; this is obviously equivalent (given that measuring has zero complexity). A slightly more substantial difference is that we allow the state ρ_{XYE} to be *subnormalized*; i.e., we allow the adversary to abort. However, the collapsing advantage $\delta_q(X, \bar{X} | \bar{Y}E)$ of any subnormalized state ρ_{XYE} is the same as of the normalized state $\tilde{\rho}_{XYE} := \rho_{XYE} + (1 - \text{tr}(\rho_{XYE}))|x_\circ\rangle\langle x_\circ| \otimes |h(x_\circ)\rangle\langle h(x_\circ)| \otimes |0\rangle\langle 0|$ for an arbitrary choice of $x_\circ \in \mathcal{X}$ on which h is defined. Since $\mathbf{c}(\tilde{\rho}_{XYE}) \leq \mathbf{c}(\rho_{XYE}) + \mathbf{c}(h)$, this has only a small quantitative effect that is insignificant if $\mathbf{c}(h)$ is insignificant compared to q . In other words, we can easily transform an adversary that aborts into one that does not abort but outputs x_\circ and $y_\circ = h(x_\circ)$ instead.

Finally, in the original definition, the complexity of the adversary and the distinguisher *together* is bounded (by q), whereas we bound the individual complexities (both by q). This is merely for simplicity, and has only a factor-2 quantitative effect.

3.2 Partial versus Total Functions

In Definition 2, we implicitly considered the function $h : \mathcal{X} \rightarrow \mathcal{Y}$ to be a *total* function, i.e., a function that is defined on its entire domain \mathcal{X} . However, it will be useful to extend the definition to *partial* functions, which are defined only on a subset $\mathcal{X}_{\text{eff}} \subseteq \mathcal{X}$ of the domain.⁹ In the context of *randomized* functions,

⁸ We recall that the requirement $\rho_{XYE} = \rho_{Xh(X)E}$ is a shorthand for asking ρ_{XYE} to be equal to a state obtained by applying $\mathbb{E}[h]$ to system X .

⁹ This may be understood in that the computation of h “fails” on inputs not in \mathcal{X}_{eff} .

as considered here, we allow \mathcal{X}_{eff} to depend on the global randomness r ; this is what distinguishes such a partial function from a total function with a smaller domain, since the domain \mathcal{X} of a function is declared fixed and independent of r .

Definition 2 applies directly to such partial functions as well, given that the definition of the evaluation map $\mathbb{E}[h]$ is naturally extended to partial functions h by having the defining operator $V[h]$ map $|x\rangle$ to 0 for any $x \notin \mathcal{X}_{\text{eff}}$. The effect of this is that the requirement $\rho_{XYE} = \rho_{Xh(X)E}$ enforces X to contain no inputs from outside of \mathcal{X}_{eff} . Hence, considering partial functions in Definition 2 serves as a convenient way to “disallow” certain inputs.

Formally, consider a function $h : \mathcal{X} \rightarrow \mathcal{Y}$ (which may be partial but let us think of it as a total function for now), and let $\pi : \mathcal{X} \rightarrow \{0, 1\}$ be a predicate, which will always be understood to be a total function. Then, we define $h|_{\pi}$ to be the partial function $h|_{\pi} : \mathcal{X} \rightarrow \mathcal{Y}$ that is undefined for $x \in \mathcal{X}$ with $\pi(x) = 0$, and that coincides with h for the remaining $x \in \mathcal{X}$. The collapsing advantage of $h|_{\pi}$ then coincides with the collapsing advantage of h modified in that the quantification over ρ_{XYE} is restricted to states for which $\Pr[\pi(\bar{X})=0] = 0$.

Below, in Lemma 1 and 2, we show how $\text{cAdv}[h]$ and $\text{cAdv}[h|_{\pi}]$ relate to each other. Lemma 1 follows trivially from the above observation, i.e., that $\rho_{XYE} = \rho_{Xh|_{\pi}(X)E}$ implies $\rho_{XYE} = \rho_{Xh(X)E}$.

Lemma 1. *If h is $\varepsilon(q)$ -collapsing then so is $h|_{\pi}$, i.e., $\text{cAdv}[h|_{\pi}] \leq \text{cAdv}[h]$.*

Applied to h of the form $h|_{\tau}$, and noting that $(h|_{\tau})|_{\pi} = h|_{\pi \wedge \tau}$, we get the following, which captures that disallowing more inputs can only decrease the collapsing advantage.

Corollary 1. *For any predicates π and τ , it holds that $\text{cAdv}[h|_{\pi \wedge \tau}] \leq \text{cAdv}[h|_{\tau}]$. In particular, if π implies τ , i.e. $\pi(x) = 1 \Rightarrow \tau(x) = 1$, then $\text{cAdv}[h|_{\pi}] \leq \text{cAdv}[h|_{\tau}]$.*

For the other direction, disallowing some inputs has little effect if those are hard to find. For the formal statement, we need the following definition.

Definition 3. *A predicate $\pi : \mathcal{X} \rightarrow \{0, 1\}$ is called $\beta(q)$ -almost-certain if it holds that $\Pr[\pi(\bar{X})=0] \leq \beta(q)$ for any state ρ_X with complexity q .*

Lemma 2. *If π is $\beta(q)$ -almost-certain then*

$$\text{cAdv}[h](q) \leq \text{cAdv}[h|_{\pi}](q + \mathbf{c}_{\pi}) + \sqrt{\beta(q)} \cdot \min\{\sqrt{2}, 1 + \sqrt{\beta(q)}\}.$$

Proof. Let $\rho_{XYZ} = \rho_{Xh(X)E}$ be with complexity q . Consider the measurement $\mathbb{P} = \mathbb{P}_0 + \mathbb{P}_1$ given by $\mathbb{P}_0 := \mathbb{M} \circ \mathbb{M}[\pi = 0]$ and $\mathbb{P}_1 := \mathbb{M}[\pi = 1]$.¹⁰ By triangle inequality and since $\mathbb{M} = \mathbb{P} \circ \mathbb{M}$, we have

$$\begin{aligned} \delta_q(X, \bar{X} | \bar{Y}E) &\leq \delta_q(X, \mathbb{P}(X) | \bar{Y}E) + \delta_q(\mathbb{P}(X), \mathbb{P}(\bar{X}) | \bar{Y}E) \\ &\leq \delta(X, \mathbb{P}(X) | \bar{Y}E) + \delta_q(X^{\pi=1}, \bar{X}^{\pi=1} | \bar{Y}E) + \delta_q(\bar{X}^{\pi=0}, \bar{X}^{\pi=0} | \bar{Y}E) \\ &\leq \sqrt{\beta(q)} + \beta(q) + \text{cAdv}[h|_{\pi}](q + \mathbf{c}_{\pi}), \end{aligned}$$

¹⁰ I.e., \mathbb{P} first performs the measurement $\mathbb{M}[\pi]$, and then measures the resulting state in the computation basis if (and only if) the measurement outcome was 0.

where the second inequality is because $\delta_q \leq \delta$, and by subadditivity and choice of \mathbb{P} , and the last inequality is by the “gentle-measurement lemma” (1), plus footnote 5, given that $\Pr[\mathbb{P}(X)=1] = \Pr[\pi(\bar{X})=1] \geq \text{tr}(\rho_X) - \beta(q)$, plus the observation that $\rho_{X^{\pi=1}YE}$ has complexity $q + \mathbf{c}_\pi$. \square

We conclude with the following simple observation, which follows from the fact that under the given assumptions, $\Pr[\tau(\bar{X})=0] \leq \Pr[\pi \circ \lambda(\bar{X})=0] \leq \beta(q + \mathbf{c}(\lambda))$ for any state ρ_X with complexity q .

Lemma 3. *Consider predicates $\pi : \mathcal{X}' \rightarrow \{0, 1\}$ and $\tau : \mathcal{X} \rightarrow \{0, 1\}$ and a total function $\lambda : \mathcal{X}' \rightarrow \mathcal{X}$ such that $\pi \circ \lambda$ implies τ , i.e., $\pi(\lambda(x)) = 1 \Rightarrow \tau(x) = 1$. If π is $\beta(q)$ -almost-certain then τ is $\beta(q + \mathbf{c}(\lambda))$ -almost-certain.*

3.3 Composability Properties

We show composability of the collapsing property under different means of composing functions. In one or another form, some of these composability properties are also present in previous work (see e.g. Lemma 27 in the full version of [5] for the corresponding claim on nested composition); we cover them here for completeness and since our notion differs in minor ways, but also in order to demonstrate how succinctly these composability properties can be *phrased* and *proven* using our formalism.

We take it as understood that for partial functions g and h , the considered composition is defined whenever g and h are both defined on their respective inputs.

Lemma 4 (Concurrent composition). *For $g : \mathcal{X} \rightarrow \mathcal{Y}$ and $h : \mathcal{W} \rightarrow \mathcal{Z}$, the concurrent composition $g||h : \mathcal{X} \times \mathcal{W} \rightarrow \mathcal{Y} \times \mathcal{Z}$, $(x, w) \mapsto (g(x), h(w))$ satisfies*

$$\mathbf{cAdv}[g||h] \leq \mathbf{cAdv}[g] + \mathbf{cAdv}[h].$$

Proof. Let $\rho_{XWYZE} = \rho_{XWg(X)h(W)E}$ be with complexity q . Then, by triangle inequality,

$$\begin{aligned} \delta_q(XW, \bar{X}\bar{W}|\bar{Y}\bar{Z}E) &\leq \delta_q(XW, X\bar{W}|\bar{Y}\bar{Z}E) + \delta_q(X\bar{W}, \bar{X}\bar{W}|\bar{Y}\bar{Z}E) \\ &= \delta_q(W, \bar{W}|\bar{Z}X\bar{Y}E) + \delta_q(X, \bar{X}|\bar{Y}\bar{W}\bar{Z}E) \\ &\leq \mathbf{cAdv}[g](q) + \mathbf{cAdv}[h](q). \end{aligned} \quad \square$$

Lemma 5 (Nested composition). *For $g : \mathcal{X} \rightarrow \mathcal{Y}$ and $h : \mathcal{Y} \rightarrow \mathcal{Z}$, the nested (or sequential) composition $h \circ g : \mathcal{X} \rightarrow \mathcal{Z}$, $x \mapsto h(g(x))$ satisfies*

$$\mathbf{cAdv}[h \circ g](q) \leq \mathbf{cAdv}[g](q + \mathbf{c}_g) + \mathbf{cAdv}[h](q + \mathbf{c}_g).$$

Proof. Let $\rho_{XZE} = \rho_{X(h \circ g)(X)E}$ be with complexity q . Then, $\rho_{XYZE} = \rho_{Xg(X)ZE}$ has complexity at most $q + \mathbf{c}_g$. Recalling that ρ_{XZE} is recovered from ρ_{XYZE}

by applying $\mathbb{E}^{inu}[g]_{XY \rightarrow X}$, we get

$$\begin{aligned}
\delta_q(X, \bar{X}|\bar{Z}E) &\leq \delta_{q+c_g}(XY, \bar{X}Y|\bar{Z}E) && \text{(monotonicity)} \\
&\leq \delta_{q+c_g}(XY, X\bar{Y}|\bar{Z}E) + \delta_{q+c_g}(X\bar{Y}, \bar{X}Y|\bar{Z}E) && (\Delta \text{ inequality}) \\
&\leq \delta_{q+c_g}(Y, \bar{Y}|\bar{Z}XE) + \delta_{q+c_g}(X, \bar{X}|\bar{Y}\bar{Z}E) && (\bar{X}Y = \bar{X}\bar{Y} \text{ by (2)}) \\
&\leq c\text{Adv}[g](q + c_g) + c\text{Adv}[h](q + c_g). && \square
\end{aligned}$$

Lemma 6. For $g : \mathcal{X} \rightarrow \mathcal{Y}$ and $h : \mathcal{W} \times \mathcal{X} \rightarrow \mathcal{Z}$, where the latter function is such that $h(\cdot, x)$ is injective for any $x \in \mathcal{X}$, the composition $f : \mathcal{W} \times \mathcal{X} \rightarrow \mathcal{Y} \times \mathcal{Z}$, $(w, x) \mapsto (g(x), h(w, x))$ satisfies

$$c\text{Adv}[f] \leq c\text{Adv}[g].$$

We emphasize that the statement includes the special case where \mathcal{W} is empty, i.e., $h : \mathcal{X} \rightarrow \mathcal{Z}$, in which case the injectivity requirement becomes void, so that in particular the following holds.

Corollary 2 (Parallel composition). For $g : \mathcal{X} \rightarrow \mathcal{Y}$ and $h : \mathcal{X} \rightarrow \mathcal{Z}$, the parallel composition $(g, h) : \mathcal{X} \rightarrow \mathcal{Y} \times \mathcal{Z}$, $x \mapsto (g(x), h(x))$ satisfies

$$c\text{Adv}[(g, h)] \leq \min\{c\text{Adv}[g], c\text{Adv}[h]\}.$$

Proof (of Lemma 6). Let $\rho_{WXYZE} = \rho_{WXg(X)h(W,X)E}$ be with complexity q . Then, using that $\bar{W}\bar{X}\bar{Z} = W\bar{X}\bar{Z}$, which holds by (2) because w is a function of x and $z = h(w, x)$,

$$\delta_q(WX, \bar{W}\bar{X}|\bar{Y}\bar{Z}E) = \delta_q(X, \bar{X}|\bar{Y}\bar{Z}WE) \leq c\text{Adv}[g](q). \quad \square$$

Lemma 7 (Disjoint union). For $g : \mathcal{X} \rightarrow \mathcal{Y}$ and $h : \mathcal{W} \rightarrow \mathcal{Z}$ with disjoint domains and images, the disjoint union $g \sqcup h : \mathcal{X} \cup \mathcal{W} \rightarrow \mathcal{Y} \cup \mathcal{Z}$, which maps $x \in \mathcal{X}$ to $g(x)$ and $w \in \mathcal{W}$ to $h(w)$, satisfies

$$c\text{Adv}[g \sqcup h] \leq c\text{Adv}[g] + c\text{Adv}[h].$$

Proof. Let $\rho_{UVE} = \rho_{U(g \sqcup h)(U)E}$, and consider the ‘‘distinguishing function’’ $dis : \mathcal{X} \cup \mathcal{W} \rightarrow \{0, 1\}$ that maps $x \in \mathcal{X}$ to 1 and $w \in \mathcal{W}$ to 0. By our convention on function domains being recognizable, dis has zero complexity. Furthermore, $\rho_{U^{dis}VE} = \rho_{\mathbb{M}[dis](U)VE}$ is of the form

$$\rho_{U^{dis}VE} = \rho_{U^{dis=0}VE} + \rho_{U^{dis=1}VE} = \rho_{Xg(X)E} + \rho_{Wh(W)E}$$

and, by the disjointness of the images, $\rho_{U^{\bar{V}}E} = \rho_{U^{dis}\bar{V}E}$, and so it follows from subadditivity that

$$\delta_q(U, \bar{U}|\bar{V}E) = \delta_q(U^{dis}, \bar{U}^{dis}|\bar{V}E) \leq \delta_q(X, \bar{X}|\bar{Y}E) + \delta_q(W, \bar{W}|\bar{Z}E)$$

which is bounded by $c\text{Adv}[g] + c\text{Adv}[h]$. \square

4 Application I: Merkle-Damgård and HAIFA

We demonstrate the usefulness of our framework. Here, we do so by (re)proving the collapsing property Merkle-Damgård, and by showing that the proof trivially translates to the HAIFA variation [3]. In the subsequent section we analyze the Sponge construction [1]. Our proofs argue entirely by means of decomposing the iteration function under consideration into a few composition operations.

Here and in the remainder, for $b \in \{0, 1, \perp\}$ and positive integer $i \in \mathbb{N}$, we write $b^i \in \{0, 1, \perp\}^i$ for the i -fold concatenation (b, \dots, b) of b with itself.

4.1 The Construction

Let $f : \{0, 1\}^c \times \{0, 1\}^r \rightarrow \{0, 1\}^c$ be a (total) function, which will act as the round function in the Merkle-Damgård construction. For any positive integer i , we consider the function $IH_i : (\{0, 1\}^r)^i \rightarrow \{0, 1\}^c$ given recursively by

$$IH_i(x_1, \dots, x_i) := f(IH_{i-1}(x_1, \dots, x_{i-1}), x_i) \quad (3)$$

with $IH_0() := iv$, some fixed string in $\{0, 1\}^c$ called the *initialization vector*. The Merkle-Damgård hash function is then formally given by¹¹

$$MD : (\{0, 1\}^r)^* \rightarrow \{0, 1\}^c, (x_1, \dots, x_i) \mapsto IH_i(x_1, \dots, x_i).$$

For technical reasons, we extend the domain of IH_i above to

$$\mathcal{X}_i := \{(x_1, \dots, x_i) \in (\{\perp\} \cup \{0, 1\}^r)^i \mid x_j = \perp \Rightarrow x_1 = \dots = x_j = \perp\}$$

by setting $IH_i(\perp, \dots, \perp) := iv$ and keeping the recursive definition (3) for $x_i \neq \perp$. We can now apply IH_L to messages of size $i < L$ blocks by pre-padding it with \perp 's: $IH_i(x_1, \dots, x_i) = IH_{i+1}(\perp, x_1, \dots, x_i) = \dots = IH_L(\perp, \dots, \perp, x_1, \dots, x_i)$, and thus the restriction of MD to messages of block size $0 \leq i \leq L$ can be expressed as $MD^{\leq L}(x_1, \dots, x_i) = IH_L(\perp, \dots, \perp, x_1, \dots, x_i)$.

4.2 The Analysis

Using our framework, we will now prove the following security statement for Merkle-Damgård. The assumption on $c(f)$ is simply for normalization, and for f to be β -*iv-preimage-resistant* means, by definition, that the predicate $1_{f(y) \neq iv}$, which is 1 if y satisfies $f(y) \neq iv$ and 0 otherwise, is β -almost-certain.

Theorem 1. *If f has complexity $c(f) = 1$, is ε -collapsing and β -iv-preimage-resistant, then, for any integer $L \geq 0$, the function $MD^{\leq L}$ is γ -collapsing with*

$$\gamma(q) = L \cdot \varepsilon \left(q + \frac{1}{2}L(L+1) \right) + \sqrt{2\beta(q+L)}.$$

¹¹ Since the bit size of the input to MD must be an integer multiple of r , the Merkle-Damgård construction usually comes with a *padding* that maps a string of arbitrary size into a sequence of blocks of bit size r . We can safely ignore this since any injective padding preserves the collapsing property by Lemma 5.

For the purpose of the proof, we define for any i the predicate $\pi_i : \mathcal{X}_i \rightarrow \{0, 1\}$ as

$$\pi_i(x_1, \dots, x_i) = 1 \iff \forall j \in \{1, \dots, i\} : x_j = \perp \vee IH_j(x_1, \dots, x_j) \neq iv,$$

i.e., the bit is set unless the input is a non-trivial iv -preimage of some IH_j . In particular, if $\pi_i(x_1, \dots, x_i) = 0$ then it must be that $IH_j(x_1, \dots, x_j) = iv$ for some j with $x_j \neq \perp$, and thus $y := (IH_{j-1}(x_1, \dots, x_{j-1}), x_j)$ satisfies $f(y) = IH_j(x_1, \dots, x_j) = iv$ by (3). So, by Lemma 3, the following holds.

Lemma 8. *If f is β - iv -preimage-resistant then π_i is $\beta(q + \mathbf{c}_{IH_{i-1}})$ -almost-certain.*

Recall that $IH_i|_{\pi_i}$ is the partial function that is defined only for the inputs which satisfy π_i . The heart of the proof of Theorem 2 is the following recursive statement, which ensures that if $IH_{i-1}|_{\pi_{i-1}}$ is collapsing then so is $IH_i|_{\pi_i}$. By repeated application, we then get that $IH_L|_{\pi_L}$ is collapsing, and since π_L is almost-certain, IH_L is collapsing as well (by Lemma 2).

Proposition 1. *For any positive integer i :*

$$\mathbf{cAdv}[IH_i|_{\pi_i}](q) \leq \mathbf{cAdv}[IH_{i-1}|_{\pi_{i-1}}](q + \mathbf{c}_{IH_{i-1}}) + \varepsilon(q + \mathbf{c}_{IH_{i-1}}).$$

Proof. We let \dot{IH}_i and $\dot{\pi}_i$ be the respective restrictions of IH_i and π_i to the domain $\dot{\mathcal{X}}_i := \mathcal{X}_i \setminus \{\perp^i\}$. Then, we see that $IH_i|_{\pi_i}$ is the disjoint union of the trivial function $\{\perp^i\} \rightarrow \{iv\}$ and $\dot{IH}_i|_{\dot{\pi}_i}$; the crucial observation here is that the image of $\dot{IH}_i|_{\dot{\pi}_i}$ is disjoint with $\{iv\}$. Therefore, by Lemma 7,

$$\mathbf{cAdv}[IH_i|_{\pi_i}](q) \leq \mathbf{cAdv}[\dot{IH}_i|_{\dot{\pi}_i}](q) \leq \mathbf{cAdv}[\dot{IH}_i|_{\pi_{i-1}}](q),$$

where the latter inequality is by Lemma 1, given that $\dot{\pi}_i$ implies π_{i-1} .¹² Furthermore, since

$$\dot{IH}_i(x_1, \dots, x_i) = f(IH_{i-1}(x_1, \dots, x_{i-1}), x_i)$$

on its domain $\dot{\mathcal{X}}_i$, i.e., it is the nested composition of f with the concurrent composition of IH_{i-1} and the identity function $x_i \mapsto x_i$, Lemma 4 and 5 imply

$$\mathbf{cAdv}[\dot{IH}_i|_{\pi_{i-1}}](q) \leq \mathbf{cAdv}[IH_{i-1}|_{\pi_{i-1}}](q + \mathbf{c}_{IH_{i-1}}) + \mathbf{cAdv}[f](q + \mathbf{c}_{IH_{i-1}}),$$

which completes the proof. \square

Proof (of Theorem 1). $IH_0|_{\pi_0} = IH_0$ is trivially 0-collapsing. For convenience, we let n_i be the sum of integers $n_i := 1 + 2 + \dots + i = \frac{1}{2}i(i+1)$. Assuming by induction that $\mathbf{cAdv}[IH_i|_{\pi_i}](q) \leq i \cdot \varepsilon(q + n_{i-1})$, we get from Proposition 1 that

$$\mathbf{cAdv}[IH_{i+1}^1|_{\pi_{i+1}}](q) \leq \varepsilon(q + i) + i \cdot \varepsilon(q + n_{i-1} + i) \leq (i+1) \cdot \varepsilon(q + n_i),$$

using that $\mathbf{c}_{IH_i} = i \cdot \mathbf{c}_f = i$ and $n_{i-1} + i = n_i$. Hence, the induction assumption holds for all i , and

$$\begin{aligned} \mathbf{cAdv}[IH_L](q) &\leq \mathbf{cAdv}[IH_L|_{\pi_L}](q + L) + \sqrt{2\beta(q + L)} \quad (\text{Lemma 2 \& 8}) \\ &\leq L \cdot \varepsilon(q + L + n_{L-1}) + \sqrt{2\beta(q + L)}. \quad \square \end{aligned}$$

¹² Here, we understand π_{i-1} as $\pi_{i-1} : \dot{\mathcal{X}}_i \rightarrow \{0, 1\}$, $(x_1, \dots, x_i) \mapsto \pi_{i-1}(x_1, \dots, x_{i-1})$.

4.3 Instantiation with a Random Oracle

If f is a *random oracle*, which formally means that we consider the oracle \mathcal{O} that is a uniformly random function $\{0, 1\}^c \times \{0, 1\}^r \rightarrow \{0, 1\}^c$ and f is the trivial oracle function that outputs whatever \mathcal{O} outputs on the given input, then, as shown by Unruh in [5], f is $O(\sqrt{q^3/2^c})$ -collapsing.¹³ Furthermore, by the results on the hardness of quantum search from [4, Thm. 1], applied to the oracle function $F : \{0, 1\}^c \times \{0, 1\}^r \rightarrow \{0, 1\}$ given by $F(y) = 1$ if and only if $f(y) = iv$, we immediately get that f is $8(q+1)^{2/2^c}$ - iv -preimage-resistant. As such, we obtain that for messages of block-size at most L , the Merkle-Damgård hash function $MD^{\leq L}$ is ε -collapsing with

$$\varepsilon(q) = O\left(L\sqrt{(q+L^2)^3/2^c}\right).$$

As far as we understand, the results of [6] imply a collapsing advantage of $O(L\sqrt{(q+L)^3/2^c})$, which is slightly better because of the L^2 that we have in our bound, but this is insignificant in typical settings where $q \gg L$.

4.4 HAIFA

Along the very same lines as for the original Merkle-Damgård construction, we can easily show that also HAIFA, a variant proposed by Biham and Dunkelman [3], is collapsing, under the same assumptions. HAIFA works similarly to Merkle-Damgård except that

$$IH_i(\text{salt}, x_1, \dots, x_i) := f(\text{salt}, IH_{i-1}(\text{salt}, x_1, \dots, x_{i-1}), x_i, i)$$

i.e., the round function takes as additional inputs the round number i and some salt (that is the same for every round).¹⁴ Proposition 1 immediately extends to HAIFA; the only thing that changes in the proof is that f becomes $f_i = f(\cdot, \cdot, \cdot, i)$, which is collapsing if f is, and we also have to use Corollary 2 to argue that the parallel composition of $(\text{salt}, x_1, \dots, x_i) \mapsto \text{salt}$ with the concurrent composition of IH_{i-1} and $x_i \mapsto x_i$ stays collapsing. The collapsing property of HAIFA then follows easily by inductively applying this variation of Proposition 1 as in the proof of Theorem 1.

4.5 Merkle-Damgård Without iv -Preimage-Resistance

We can also recover Unruh's original result on MD , which does not require f to be iv -preimage-resistant but instead restricts the set of inputs to be *suffix-free*.

¹³ Even though our definition of the collapsing property differs slightly from the definition in [5], these differences disappear in such asymptotic statements, as discussed in Section 3. See also Appendix C.

¹⁴ For the purpose of collisions and the collapsing property, we can think of the salt simply as part of the input: we do not want collisions even for different choices of the salt.

For that, given a fixed integer $L > 0$ and arbitrary $0 \leq i \leq L$, consider the map IH_i^* given by

$$IH_i^* : (x_1, \dots, x_L) \mapsto (IH_i(x_1, \dots, x_i), x_{i+1}, \dots, x_L),$$

defined on the considered suffix-free inputs of size at most L blocks, left-padded with \perp 's, and we argue the following variant of Proposition 1: if $IH_i^*|_{x_i \neq \perp}$ is collapsing then $IH_{i+1}^*|_{x_{i+1} \neq \perp}$ is collapsing too (for $i < L$). This variant of Proposition 1 follows from the observation that the latter is obtained as the nested composition $IH_{i+1}^*|_{x_{i+1} \neq \perp} = (f \parallel id) \circ IH_i^*|_{x_{i+1} \neq \perp}$ of $IH_i^*|_{x_{i+1} \neq \perp}$ with the concurrent composition of f and the identity id acting on x_{i+2}, \dots, x_L . Furthermore,

$$IH_i^*|_{x_{i+1} \neq \perp}(x_1, \dots, x_L) = \begin{cases} IH_i^*|_{x_i \neq \perp}(x_1, \dots, x_L) = (IH_i(x_1, \dots, x_i), x_{i+1}, \dots, x_L) & \text{if } x_i \neq \perp \\ (iv, x_{i+1}, x_{i+2}, \dots, x_L) & \text{if } x_i = \perp \end{cases}$$

and therefore $IH_i^*|_{x_{i+1} \neq \perp}$ is the disjoint union of $IH_i^*|_{x_i \neq \perp}$ and the function $(\perp^i, x_{i+1}, \dots, x_L) \mapsto (iv, x_{i+1}, x_{i+2}, \dots, x_L)$. Here, we are using the suffix-freeness of the considered inputs x_1, \dots, x_L ; this ensures that not only the domains but also the images of the two functions are disjoint: if $(\perp^i, x_{i+1}, \dots, x_L)$ is ‘‘allowed’’ then (x_1, \dots, x_L) is not unless x_1 up to x_i are all \perp . The above variant of Proposition 1 then follows from the preservation of the collapsing property under the different compositions, and then, by inductively applying this variant of Proposition 1, we obtain that $IH_L^*|_{x_L \neq \perp}$ is collapsing, and thus $MD^{\leq L}$ is, given that the input is from a suffix-free set.

5 Application II: The Sponge

Here, we apply our framework to the Sponge construction [1]. As one can see, we follow the exact same blueprint as in Section 4.

5.1 The Construction

Let $f = (f^0, f^1) : \{0, 1\}^r \times \{0, 1\}^c \rightarrow \{0, 1\}^r \times \{0, 1\}^c$ be a (total) function, which will act as the round function in the Sponge construction. For any positive integer i , consider the function

$$S_i = (S_i^0, S_i^1) : (\{0, 1\}^r)^i \rightarrow \{0, 1\}^r \times \{0, 1\}^c$$

given recursively by

$$S_i(x_1, \dots, x_i) := f(S_{i-1}^0(x_1, \dots, x_{i-1}) \oplus x_i, S_{i-1}^1(x_1, \dots, x_{i-1})) \quad (4)$$

with $S_0() := 0$. The sponge function (with s rounds of ‘‘squeezing’’) is then formally given by¹⁵

$$\text{Sponge}[s] : (\{0, 1\}^r)^* \rightarrow (\{0, 1\}^r)^s \\ (x_1, \dots, x_i) \mapsto (S_i^0(x_1, \dots, x_i), S_{i+1}^0(x_1, \dots, x_i, 0^r), \dots, S_{i+s-1}^0(x_1, \dots, x_i, 0^r, \dots, 0^r)).$$

¹⁵ Like for Merkle-Damgård, we can safely ignore the padding here.

For technical reasons, we extend the domain of S_i above to

$$\mathcal{X}_i := \{(x_1, \dots, x_i) \in (\{\perp\} \cup \{0, 1\}^r)^i \mid x_j = \perp \Rightarrow x_1 = \dots = x_j = \perp\}$$

i.e., to strings that may have \perp -prefixes. We do so by setting

$$S_i(\perp, \dots, \perp) := 0^{r+c}$$

and keeping the recursive definition (4) for $x_i \neq \perp$. This extension allows us to apply S_L to messages $(x_1, \dots, x_i) \in (\{0, 1\}^r)^i$ of size $i < L$ blocks by pre-padding it with \perp 's: $S_i(x_1, \dots, x_i) = S_{i+1}(\perp, x_1, \dots, x_i) = \dots = S_L(\perp, \dots, \perp, x_1, \dots, x_i)$, and thus the restriction of $Sponge[s]$ to messages of block size $1 \leq i \leq L$ can be expressed as:

$$Sponge[s]^{\leq L}(x_1, \dots, x_i) = (S_L^0(\perp^{L-i}, x_1, \dots, x_i), S_{L+1}^0(\perp^{L-i}, x_1, \dots, x_i, 0^r), \dots) \quad (5)$$

where we note that we insist here on $i \geq 1$, i.e., the message is non-empty.

5.2 The Analysis

Here, we prove the following. Also here, the assumption on $\mathfrak{c}(f)$ is simply for normalization, and for f^1 to be β -zero-preimage-resistant means, by definition, that the predicate $1_{f^1(y) \neq 0^c}$ is β -almost-certain.

Theorem 2. *If f has complexity 1, and f^0 and f^1 are ε^0 - and ε^1 -collapsing, and f^1 is β -zero-preimage-resistant, then, for any integer $L \geq 0$, the Sponge function $Sponge[s]^{\leq L}$ is γ -collapsing with*

$$\gamma(q) \leq \varepsilon^0(q + 2L - 1) + (L - 1) \cdot \varepsilon^1(q + \frac{1}{2}L(L+1)) + \sqrt{2\beta(q + L)}.$$

For the purpose of the proof, we define for any i the predicate $\pi_i : \mathcal{X}_i \rightarrow \{0, 1\}$ as

$$\pi_i(x_1, \dots, x_i) = 1 \iff \forall j \in \{1, \dots, i\} : x_j = \perp \vee S_j^1(x_1, \dots, x_j) \neq 0^c,$$

i.e., the bit is set unless the input is a non-trivial zero-preimage of some S_j^1 . In particular, if $\pi_i(x_1, \dots, x_i) = 0$ then $S_j^1(x_1, \dots, x_j) = 0^c$ for some j with $x_j \neq \perp$, and thus $y := (S_{j-1}^0(x_1, \dots, x_{j-1}) \oplus x_j, S_{j-1}^1(x_1, \dots, x_{j-1}))$ satisfies $f^1(y) = S_j^1(x_1, \dots, x_j) = 0^c$ by (4). Thus, by Lemma 3, the following holds.

Lemma 9. *If f^1 is β -zero-preimage-resistant then π_i is $\beta(q + \mathfrak{c}_{S_{i-1}})$ -almost-certain, and the same holds for $\hat{\pi}_i$, defined as below.*

For any i , let \hat{S}_i^b and $\hat{\pi}_i$ be the respective restrictions of S_i^b and π_i to the domain $\hat{\mathcal{X}}_i := \mathcal{X}_i \setminus \{\perp^i\}$. The heart of the proof of Theorem 2 is the following recursive statement, which ensures that if $S_{i-1}^1|_{\pi_{i-1}}$ is collapsing then so are $\hat{S}_i^0|_{\hat{\pi}_i}$ and $S_i^1|_{\pi_i}$. By repeated application, we then get that $\hat{S}_L^0|_{\hat{\pi}_L}$ is collapsing, and since $\hat{\pi}_L$ is almost-certain, \hat{S}_L^0 is collapsing as well (by Lemma 2).

Proposition 2. For any positive integer i :

$$\text{cAdv}[\dot{S}_i^0|\dot{\pi}_i](q), \text{cAdv}[S_i^1|\pi_i](q) \leq \text{cAdv}[S_{i-1}^1|\pi_{i-1}](q + \mathbf{c}_{S_{i-1}}) + \varepsilon^b(q + \mathbf{c}_{S_{i-1}}).$$

Proof. We note that $S_i^1|\pi_i$ is the disjoint union of the trivial function $\{\perp^i\} \rightarrow \{0^c\}$ and $\dot{S}_i^1|\dot{\pi}_i$; the crucial observation here is that the image of \dot{S}_i^1 does not contain 0^c . Therefore, by Lemma 7,

$$\text{cAdv}[S_i^1|\pi_i](q) \leq \text{cAdv}[\dot{S}_i^1|\dot{\pi}_i](q) \leq \text{cAdv}[\dot{S}_i^1|\pi_{i-1}](q).$$

where the latter inequality is by Lemma 1, given that $\dot{\pi}_i$ implies π_{i-1} .¹⁶ Furthermore, since

$$\dot{S}_i^1(x_1, \dots, x_i) = f^1(S_{i-1}^0(x_1, \dots, x_{i-1}) \oplus x_i, S_{i-1}^1(x_1, \dots, x_{i-1}))$$

on its domain $\dot{\mathcal{X}}_i$, i.e., it is a nested composition of f^1 with a function that is obtained as a composition as considered in Lemma 6, Lemma 5 and 6 imply that

$$\text{cAdv}[\dot{S}_i^1|\pi_{i-1}](q) \leq \text{cAdv}[S_{i-1}^1|\pi_{i-1}](q + \mathbf{c}_{S_{i-1}}) + \text{cAdv}[f^b](q + \mathbf{c}_{S_{i-1}}),$$

which was to be proven. The reasoning for $\dot{S}_i^0|\dot{\pi}_i$ is exactly as for $\dot{S}_i^1|\dot{\pi}_i$ above. \square

Proof (of Theorem 2). $S_0^1|\pi_0 = S_0^1$ is trivially 0-collapsing. For convenience, we let n_i be the sum of integers $n_i := 1 + 2 + \dots + i = \frac{1}{2}i(i+1)$. Assuming by induction that $\text{cAdv}[S_i^1|\pi_i](q) \leq i \cdot \varepsilon^1(q + n_{i-1})$, we get from Proposition 2 that

$$\text{cAdv}[S_{i+1}^1|\pi_{i+1}](q) \leq \varepsilon^1(q + i) + i \cdot \varepsilon^1(q + n_{i-1} + i) \leq (i + 1) \cdot \varepsilon^1(q + n_i),$$

using that $\mathbf{c}_{S_i} = i \cdot \mathbf{c}_f = i$ and $n_{i-1} + i = n_i$. Hence, the induction assumption holds for all i , and

$$\begin{aligned} \text{cAdv}[\dot{S}_L^0](q) &\leq \text{cAdv}[\dot{S}_L^0|\dot{\pi}_L](q + L) + \sqrt{\beta(q + L)} \\ &\leq \varepsilon^0(q + 2L - 1) + \text{cAdv}[S_{L-1}^1|\pi_{L-1}](q + 2L - 1) + \sqrt{2\beta(q + L)} \\ &\leq \varepsilon^0(q + 2L - 1) + (L - 1) \cdot \varepsilon^1(q + n_L) + \sqrt{2\beta(q + L)}. \end{aligned}$$

where the first inequality is by Lemma 2 and 9, and the second by Proposition 2. The claim on $\text{Sponge}[s]^{\leq L}$ follows now from (5) and Corollary 2. \square

5.3 Instantiation with a Random Oracle

If $f = (f^0, f^1)$ is a *random oracle*, then it follows easily from the work of Unruh in [5] on the collapsing property of the random oracle that f^0 and f^1 are respectively $O(\sqrt{q^3/2^r})$ - and $O(\sqrt{q^3/2^c})$ -collapsing. Furthermore, as pointed out in [2], by the results on the hardness of quantum search from [4, Thm. 1] to the oracle function $F : \{0, 1\}^r \times \{0, 1\}^c \rightarrow \{0, 1\}$ given by $F(y) = 1$ if and

¹⁶ Here, we understand π_{i-1} as $\pi_{i-1} : \dot{\mathcal{X}}_i \rightarrow \{0, 1\}$, $(x_1, \dots, x_i) \mapsto \pi_{i-1}(x_1, \dots, x_{i-1})$.

only if $f^1(y) = 0^c$, we immediately get that the function f^1 is $8(q+1)^2/2^c$ -zero-preimage-resistant. Therefore, we get that for messages of block-size at most L , the sponge function $Sponge[s]^{\leq L}$, with the round function modeled by a random oracle, is ε -collapsing with

$$\varepsilon(q) = O\left(\sqrt{(q+L)^3/2^r} + L\sqrt{(q+L^2)^3/2^c}\right).$$

This matches with single-execution-variant (i.e. $t = 1$) of Theorem 33 of [2], except for the square in the L^2 term. When considering a t -fold parallel composition $Sponge[s]^{\leq L} \parallel \dots \parallel Sponge[s]^{\leq L}$, it follows immediately from Lemma 6 that the collapsing parameter grows linearly with t , i.e., as

$$O\left(t\sqrt{(q+L)^3/2^r} + tL\sqrt{(q+L^2)^3/2^c}\right),$$

which is comparable to Theorem 33 of [2] with a general t , which states a collapsing advantage of

$$O\left(t\sqrt{(q+tL)^3/2^r} + tL\sqrt{(q+tL)^3/2^c}\right).$$

6 Conclusion

We consider the quantum *collapsing* property of classical hash functions, which replaces the notion of *collision resistance* in the presence of quantum attacks, and we propose a formalism and a framework that enables to argue about the collapsing property of hash domain extension constructions simply by means of decomposing the iteration function under consideration into elementary composition operations. In particular, our framework allows us to argue by purely classical means that hash functions are secure against quantum attacks.

We demonstrate this proof methodology on several examples. For Merkle-Damgård and the Sponge construction, we recover what has already been proven in [6] and [2], up to insignificant differences, whereas our result for HAIFA is, strictly speaking, new. It is well possible that the respective proof provided in [6] extends to HAIFA as well; however, this is cumbersome to verify (we challenge the reader to do so). With our approach, on the other hand, it is *trivial* to see that our proof for Merkle-Damgård extends to this variation: the only thing that needs to be verified is that the modified iteration function still decomposes into composition operations that are covered by our framework.

We think it is fair to say that, compared to previous work which proves that some hash domain extension constructions *are* collapsing, our approach gives much more insight into *why* they are collapsing. Furthermore, our framework should be a helpful tool when designing new hash functions that are meant to withstand quantum attacks.

Last but not least, from a conceptual perspective, we find it particularly interesting to see that our simplified proofs are the result of departing from the common methodology of proving a conditional security statement by means of

an algorithmic reduction. Instead of assuming an attack against the construction and then building an attack against the underlying component, we argue *directly*—and in some sense “algebraically”—that if the underlying component is secure then so is the construction.

References

1. Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles van Assche. On the Indifferentiability of the Sponge Construction. In *EUROCRYPT 2008*, vol. 4965 of *LNCS*, pp. 181–197 (2008).
2. Jan Czajkowski, Leon Groot Bruinderink, Andreas Hülsing, Christian Schaffner, and Dominique Unruh. Post-quantum security of the sponge construction. To appear in *PQCrypto 2018*. Available at: <https://eprint.iacr.org/2017/771.pdf>.
3. Eli Biham, and Orr Dunkelman. A Framework for Iterative Hash Functions—HAIFA. *Second NIST Cryptographic Hash Workshop* (2006). Available at: <https://eprint.iacr.org/2007/278.pdf>.
4. Andreas Hülsing and Joost Rijneveld and Fang Song. Mitigating Multi-Target Attacks in Hash-based Signatures. In *PKC 2017*, vol. 9614 of *LNCS*, pp. 387–416 (2017). Available at: <https://eprint.iacr.org/2015/1256>.
5. Dominique Unruh. Computationally Binding Quantum Commitments. In *EUROCRYPT 2016*, vol. 9666 of *LNCS*, pp. 497–527 (2016).
6. Dominique Unruh. Collapse-Binding Quantum Commitments Without Random Oracles. In *ASIACRYPT 2016*, vol. 10032 of *LNCS*, pp. 166–195 (2016).
7. Hongjun Wu. The Hash Function JH. *NIST SHA-3 Competition Finalist* (2011). Available at: http://www3.ntu.edu.sg/home/wuhj/research/jh/jh_round3.pdf.
8. Andreas Winter. Coding Theorem and Strong Converse for Quantum Channels. *IEEE Trans. Inf. Theory* 45(7):2481-2485 (1999). Available at: <https://arxiv.org/abs/1409.2536>.
9. Mark Wilde. From Classical to Quantum Shannon Theory. Manuscript, 2nd edition (2016). Available at: <https://arxiv.org/abs/1106.1445>.

A Randomized Functions and CPTN Maps

In this work, we will consider two variants of the notion of a *randomized function*, and our techniques will apply to both. Formally, a randomized function is a function $f : \mathcal{R} \times \mathcal{X} \rightarrow \mathcal{Y}$ for a fixed choice of the (finite) set \mathcal{R} , and it is understood that $r \in \mathcal{R}$ is chosen uniformly at random once-and-for-all. Informally, we think of such a randomized function as a function $f : \mathcal{X} \rightarrow \mathcal{Y}$ that produces its output $f(x) = f(r, x)$ for any input x dependent on some “global randomness” r , which is the same for all inputs and all randomized functions considered at a time.

Informally, the two variants we consider in this work differ in the way the randomness r is accessed by the function. In one case, r is *explicitly* given as input to the function f (or to the algorithm that computes f , if you prefer); one then typically speaks of *keyed* (or *seeded*) functions. In the other case, r is not explicitly given to f but instead, f makes *oracle queries* to a designated

randomized function \mathcal{O} , called the *oracle*, which computes every reply dependent on r . This latter case is typically referred to as an *oracle function*.

We point out that from a mathematical perspective, there is no distinction yet between a keyed and an oracle function, in that both are merely functions that additionally act on some global randomness r . The way the two variants differ formally is by the way we capture *complexity*: for keyed functions we consider the *computational complexity* whereas for oracle functions we consider the *query complexity*.¹⁷ We address this in more detail in the subsequent section.

In line with the above, we can also consider the notion of a *randomized CPTN map* \mathbb{T} , which is a CPTN map whose action on a quantum state depends on the global randomness r , and we can distinguish between *keyed* CPTN maps that have *direct* access to r , and *oracle* CPTN maps that have *quantum oracle access* to a designated randomized function $\mathcal{O} : \mathcal{R} \times \mathcal{U} \rightarrow \mathcal{V}$. Here, “quantum oracle access” means that \mathbb{T} can query \mathcal{O} in superposition, i.e., it may ask to have the unitary $|u\rangle|v\rangle \mapsto |u\rangle|v + \mathcal{O}(r, u)\rangle$ applied to any state (of appropriate dimension). Again, the formal distinction between the two variants is in terms of the complexity measure.

We point out that by considering randomized CPTN maps \mathbb{T} (of either flavor) that act on the empty system with trivial one-dimensional state space, we may also speak of *randomized states* (of either flavor) as the states ρ produced as $\rho = \mathbb{T}(1)$ for such a randomized CPTN map. We take it as understood here that the description of such a randomized state includes the dependency on the randomness r .

B Complexity

We introduce here the abstract notion of *complexity* that we consider in this work and discuss below the two main instantiations that are relevant for us. In the context of randomized functions $f : \mathcal{X} \rightarrow \mathcal{Y}$, we consider a map that assigns to any such function a non-negative integer $\mathfrak{c}(f)$, called the *complexity* of f , which is meant to express how hard it is to “compute” f . We assume that our abstract notion satisfies natural properties, like that the identity function on any set has zero complexity, and that it behaves well under composition, so that

$$\mathfrak{c}(g \circ f) \leq \mathfrak{c}(f) + \mathfrak{c}(g) \quad \text{and} \quad \mathfrak{c}(f \parallel g) \leq \mathfrak{c}(f) + \mathfrak{c}(g)$$

for any f and g with appropriate domain/range, where $g \circ f : x \mapsto g(f(x))$ and $f \parallel g : (x, w) \mapsto (f(x), g(w))$. For simplicity, we additionally assume that certain “simple” functions have zero complexity. These are: constants, copying, deleting, swapping, checking equality, as well as bit-wise XOR. Also, to avoid certain technical complications, once \mathfrak{c} is fixed we only consider randomized

¹⁷ For the latter, one could actually consider both simultaneously.

functions $f : \mathcal{X} \rightarrow \mathcal{Y}$ for which \mathcal{X} can be recognized with zero complexity.¹⁸ For lighter notation, we may also write \mathfrak{c}_f instead of $\mathfrak{c}(f)$.

We also consider a notion of complexity for randomized CPTN maps, which, as above, assigns a non-negative integer $\mathfrak{c}(\mathbb{T})$ to any randomized CPTN map \mathbb{T} . Similarly to above, we assume that the identity \mathbb{I} has zero complexity, that

$$\mathfrak{c}(\mathbb{S} \circ \mathbb{T}) \leq \mathfrak{c}(\mathbb{T}) + \mathfrak{c}(\mathbb{S}) \quad \text{and} \quad \mathfrak{c}(\mathbb{T} \otimes \mathbb{S}) \leq \mathfrak{c}(\mathbb{T}) + \mathfrak{c}(\mathbb{S}),$$

and that certain “simple” maps have zero complexity, namely: the preparation of states in the computational basis, measurements (with or without post-selection) in the computational basis, partial traces, and swapping registers. On top, we assume the complexity notion for CPTN maps to be consistent with that of functions, in that we require that

$$\mathfrak{c}(\mathbb{E}[f]), \mathfrak{c}(\mathbb{E}^{inv}[f]), \mathfrak{c}(\mathbb{M}[f]), \mathfrak{c}(\mathbb{M}[f=y]) \leq \mathfrak{c}(f)$$

where the latter two actually follow from the first, given that partial traces and measurements in the computational basis are “for free”.

Given such a notion of complexity (for randomized CPTN maps), we can define the complexity of a randomized state ρ as $\mathfrak{c}(\rho) := \mathfrak{c}(\mathbb{T})$ where \mathbb{T} is the randomized CPTN map with minimal complexity that produces ρ as $\rho = \mathbb{T}(1)$. It obviously holds that $\mathfrak{c}(\mathbb{T}(\rho)) \leq \mathfrak{c}(\rho) + \mathfrak{c}(\mathbb{T})$ for any randomized CPTN map. A last requirement we pose onto our abstract complexity measure \mathfrak{c} is that $\mathfrak{c}(\rho + (1 - \text{tr}(\rho))\sigma) \leq \mathfrak{c}(\rho) + \mathfrak{c}(\sigma)$ for all randomized states.¹⁹

Example 1. An important class of examples for such a complexity measure arises by considering keyed functions and keyed CPTN maps and writing them as *circuits* that get the global randomness as additional input. The complexity can then be specified to be the minimal number of gates (of certain types) of any such circuit representation. For instance, writing any randomized function f as a binary circuit with AND and XOR gates, one may define $\mathfrak{c}^{\text{comp}}(f)$ to be the minimal necessary number of AND gates.²⁰ Similarly for randomized CPTN maps, where one could for instance count the number of gates (or the number of non-Clifford gates) with respect to a fixed universal set of gates. These kinds of notions of complexity are referred to as *computational complexity*.

Example 2. Another example that is relevant for us is the *query complexity* $\mathfrak{c}^{\text{query}}$ for oracle functions and oracle CPTN maps, which counts the number of (quantum) oracle queries that the function or CPTN map makes to the oracle \mathcal{O} .

We note that this abstract treatment of complexity allows us to *explicitly* cover computational complexity and query complexity in one go, using one language. Also, all results expressed using this language do *explicitly* not depend

¹⁸ Meaning that for any $\mathcal{X}' \supset \mathcal{X}$, the function $\mathcal{X}' \rightarrow \{0, 1\}$ that maps $x \in \mathcal{X}$ to 1 and $x \in \mathcal{X}' \setminus \mathcal{X}$ to 0 has zero complexity. This is trivially satisfied for *any* function f in case of $\mathfrak{c}^{\text{query}}$ (given in Example 2).

¹⁹ This is in line with our interpretation of $\text{tr}(\rho) < 1$ as capturing an “abort” of the preparation process.

²⁰ Remember that we want XOR’s to be “for free”.

on the technical details of any model of computation. Related to the latter, this approach allows us to reason about *functions* (and *CPTN* maps etc.), which are unambiguously defined objects that do not depend on any model of computation.

C On the Definitions of the Collapsing Property in [5, 6]

As already mentioned in Section 3, the definition of the collapsing property introduced by Unruh comes in a few different variations in [5, 6], which we want to briefly recall here, and we discuss how they compare to our definition. Some differences (like allowing non-normalized states), which have some minor quantitative impact, have already been discussed in Section 3; here, we focus on some technical differences that are orthogonal to those.

The definition originally proposed in [5, Def. 20] (respectively Def. 23 in the full version) is of asymptotic nature and for a deterministic function h (that depends on a security parameter κ): it requires that every (uniform) quantum-polynomial time adversary has a negligible collapsing advantage. Note that it makes sense to consider a *deterministic* hash function h in this case since a *uniform* model of computation is considered, i.e., the adversary that has a collision hard-wired into its code (for any κ) is not allowed. In the formal statement on the collapsing property of a random oracle [5, Thm. 31], which bounds the advantage by $O(q^3/\text{size of the range})$, an obvious variation (in terms of oracle algorithms, and with the randomness also over the oracle’s random choices) of this definition is then (implicitly) considered. It also remains implicit that the bound (and in particular the hidden constant) is independent of the running time of the adversary; indeed, the proof considers “ q -query adversaries”, which have bounded query complexity but (possibly) unbounded running time. As such, the bound $O(q^3/\text{size of the range})$ carries over to our non-asymptotic definition of the collapsing property (with the complexity measure being the query complexity), with the understanding that the bound includes a hidden constant and only applies for a large enough range of h (compared to q).

In terms of comparing Unruh’s original definition [5, Def. 20] with our Definition 2, though being similar in spirit (up small differences as discussed in Section 3), they are technically incomparable: an asymptotic definition as [5, Def. 20] makes no meaningful statement about a fixed instance, while, on the other hand, our non-asymptotic definition is meaningless for a deterministic hash function, for instance. If we consider an asymptotic variant of our Definition 2, which would ask $\text{cAdv}[h_\kappa](q(\kappa))$ to be negligible in the security parameter κ for any polynomially bounded function q , we get a *non-uniform* (and thus stronger) variant of [5, Def. 20]. It is easy to see that all our results directly carry over to such an asymptotic variant.

In [6, Def. 8], Unruh also considers a non-asymptotic “concrete security” variant of his original definition, which considers a *keyed* hash function (using our terminology) and defines the “collapsing advantage” of an arbitrary but fixed adversary in the obvious way, as the advantage of this adversary distinguishing

the two games.²¹ For a function h being $\varepsilon(q)$ -collapsing according to our definition thus immediately implies that the “collapsing advantage” according to [6, Def. 8] is bounded by $\varepsilon(q)$ (up to a small constant factor, as explained in Section 3) for any adversary that is bounded by q , and vice versa. In that sense, our Definition 2 and Unruh’s “concrete security” variant [6, Def. 8] are *equivalent* (again, up to negligible quantitative differences).

However, formalized as in [6], i.e., not as a security property of h but as a property of an arbitrary but fixed adversary that is attacking h , security statements are bound to be in reductionistic form. Indeed, all the concrete-security statements in [6] are like:

“Let A be a q -time (or query) adversary with collapsing advantage ε against function h , then there exists a $2q$ -time (or query) adversary A' with collapsing advantage ε^2 against h' ”

where we consider concrete example functions for the “security loss”. On the other hand, our definition allows us to express such a statement simply as:

“If the function h' is $\varepsilon'(q')$ -collapsing then h is $\sqrt{\varepsilon'(2q)}$ -collapsing”

or even more compactly as

$$\text{cAdv}[h](q) \leq \sqrt{\text{cAdv}[h](2q)}.$$

Again, these statements are *equivalent* (up to the minor quantitative differences discussed in Section 3), so it is merely a matter of taste what sort of language one prefers.

A more crucial difference is that we not only can state but also *prove* these kinds of security claims “in the forward direction”, i.e., not via the counter position of assuming an attacker A that breaks the target primitive and turning it into an attacker A' that breaks something else. Indeed, we have “algebraic” proofs that avoid reasoning about algorithms altogether. On the other hand, the proofs in [5, 6, 2] are typically constructive, in that the adversary A' is *explicitly* constructed from the adversary A . One can then for instance easily check that A' does not rely on any non-uniform auxiliary information, and thus that the reductions are *uniform*. Our proofs do not spell out the reductions; however, if desired, they could still be extracted by backtracking the proofs all the way down to the basic properties of the pseudometric δ_q upon which the proofs rely, and which all have simple — uniform — reduction proofs. This ensures that also our implicitly defined reductions are uniform.

²¹ As a matter of fact, [6, Def. 8] considers a t -fold parallel repetition of the game, where t is an additional parameter of the definition. We ignore this for the discussion here, recalling that we obtain immediately a similar variant of our definition by means of concurrent composition.