(Quantum) Indifferentiability and Pre-Computation

Joseph Carolan[1], Alexander Poremba[2], and Mark Zhandry[3]

[1] University of Maryland. jcarolan@umd.edu
[2] Massachusetts Institute of Technology. poremba@mit.edu
[3] NTT Research. mzhandry@gmail.com

Abstract. Indifferentiability is a popular cryptographic paradigm for analyzing the security of ideal objects—both in a classical as well as in a quantum world. It is typically stated in the form of a composable and simulation-based definition, and captures what it means for a construction (e.g., a cryptographic hash function) to be "as good as" an ideal object (e.g., a random oracle). Despite its strength, indifferentiability is not known to offer security against *pre-processing* attacks in which the adversary gains access to (classical or quantum) advice that is relevant to the particular construction. In this work, we show that indifferentiability is (generically) insufficient for capturing pre-computation. To accommodate this shortcoming, we propose a strengthening of indifferentiability which is not only composable but also takes arbitrary pre-computation into account. As an application, we show that the one-round sponge is indifferentiable (with pre-computation) from a random oracle. This yields the first (and tight) classical/quantum space-time trade-off for one-round sponge inversion.

1 Introduction

Hash functions are fundamental objects in cryptography, used in a multitude of applications such as password storage, integrity checks, and in digital signature schemes. Many real world cryptographic schemes can only be proven secure in the random oracle model [BR93] (ROM), in which a hash function is instead treated like an idealized perfectly random function. In this model, adversaries receieve only black-box access to the hash function, which enables one to prove query lower bounds justifying the security of a construction. Similarly, the quantum random oracle model [BDF+11] (QROM) models adversaries as having quantum query-access to an idealized random function. These tools have since become indispensable in analyzing real world cryptographic systems, both in the (post-)quantum and the classical setting.

Indifferentiability. In the real world, however, hash functions are built from lower-level building blocks, such as compression functions or publicly invertible permutations. The structure of these hashes can lead to attacks: length-extension attacks on Merkle-Damgård are a famous example; another example is circular-secure encryption when using Davies-Meyer [HK07]. These attacks work regardless of the lower-level building block, and simply exploit the way the building block is used in the higher-level protocol.

One possibility is to analyze a given hash construction when used in specific scenarios. A much better solution [CDMP05] is to ensure that the hash function is *indifferentiable* from a random oracle. Indifferentiability was first defined by Maurer, Renner, and Holenstein [MRH04], and is a composable, simulation-based definition. An indifferentiable hash function is "as good as" a random oracle, in that we can "lift" any single-stage security property of random oracles—which capture most of the standard properties—to conclude that the property also holds for an indifferentiable hash function, provided the underlying building block is modeled as an idealized object. Indifferentiability therefore ensures that no attacks were introduced in the conversion from the lower-level building block to the higher-level hash function. Therefore, rather than analyzing the hash function in every scenario of interest, we can simply prove that it is indifferentiable and immediately conclude its security.

A strengthening of this notion, called reset indifferentiability, was introduced by Ristenpart, Shacham, and Shrimpton [RSS11]. This notion requires a stateless simulator, but allows composable security in games with an arbitrary number of stages and adversaries. While numerous positive results are known for plain indifferentiability (such as domain extension of random oracles, and

equivalence between ideal ciphers and random oracles [CPS08, HKT11, DS16]), various barriers apply to reset indifferentiability [RSS11, LAMP12, DGHM13, BBM13]: in particular, domain extension is not possible. Despite this barrier, Zhandry [Zha21] has used reset indifferentiability to show, among other things, that ideal ciphers imply fixed size random oracles. In particular, it is shown that the single-round sponge is weakly reset-indifferentiable (quantum or classical) from a random oracle, when the rate does not exceed the capacity.

Adversaries with pre-computation. A common and desirable property of a cryptographic scheme is security against adversaries with some pre-computed advice. In the context of the random oracle model and other idealized models, this is usually considered in the auxiliary input model introduced by Unruh [Unr07]. In this model, adversaries are split into an inefficient offline and efficient online stage, where only a single (potentially small) advice message can be passed from offline to online. The online adversary then receives a challenge, or more generally must win some security game with the help of the advice. Many classical and quantum results are known in this model [CDG18, Yao90, DTT10, Hel80, FN00, CGK19, HXY19, CLQ20, CGLQ20].

The aforementioned works assume a random oracle. As mentioned above, however, hash functions are typically built from lower-level building blocks, and this structure may be exploited in attacks. A recent line of work therefore has investigated pre-computation attacks on structured hash functions, giving both attacks and lower-bounds in different settings [CDG18, ACDW20, GK22, FGK22, ADGL23, Aks24]. There are not, to our knowledge, any known space-time tradeoffs (either classical or quantum) for inverting the one-round sponge. A natural question is:

Why not just use the stong notion of indifferentiability to lift space-time trade-offs for random oracles to structured hash functions?

After all, the goal of indifferentiability is to avoid having to analyze a structured hash construction in every conceivable scenario, and instead simply lift existing random oracle results. The short answer, unfortunately, is that indifferentiability as currently defined simply does not work. Due to the pre-computation phase, space-time trade-offs are not single-stage games, and since the pre-computation is inefficient, it is also not a multi-stage game. Therefore, the lifting theorems for (reset) indifferentiability simply do not apply to space-time trade-offs. In fact, we show that this is inherent: there is a function which is (strongly, statistical) reset indifferentiable from a random oracle, but admits a pre-computation attack on function inversion with polynomial-size advice and polynomial computation. This leads us to ask the following question:

Does this mean that indifferentiability cannot help us understand space-time tradeoffs for structured hash functions?

2 Our contributions

We now give an overview of our main results.

2.1 (Quantum) indifferentiability with pre-computation

Motivated by the aforementioned counterexample, we introduce a notion of indifferentiability with pre-computation—a strengthening of plain indifferentiability, in Section 5. In this definition, a distinguisher is split into an offline and an online part, where the offline part is allowed unbounded access to some primitive. However, only a limited size message can be passed to the online adversary, which is bounded. We define indifferentiability by way of a simulator which is also split into two stages, an inefficient offline and an efficient online simulator, and again allow the simulators to pass some bounded size advice from offline to online. As our goal is to show (tight) space-time tradeoffs, rather than restricting to efficient (polynomial time) adversaries we instead phrase our definitions in a more fine-grained manner, parameterized by advice sizes, query count, and success probabilities. In strong indifferentiability, the simulator must simulate the adversaries' interface in both the offline and online phase, though in weak indifferentiability we instead allow the simulator to prepare both the advice for the online adversary, as well as for the online simulator.

In Section 6, we introduce a composition theorem under our definition of indifferentiability with pre-computation. Informally, this composition theorem says that if construction C is indifferentiable from construction R, then a security game with pre-computation (for instance a space-time tradeoff for function inversion) instantiated using C is as secure as one instantiated using R. This statement holds up to some loss incurred from indifferentiability, which depends on how much advice the simulator needs as well as how many queries.

Theorem 1 (Informal version of Theorem 4). *If construction C is indifferentiable with pre-computation from construction R, then any security game with a pre-computing adversary which is secure when instantiated with R remains secure when instantiated with C.*

2.2 (Quantum) space-time trade-offs for sponge inversion

In recent years, the National Institute of Standards and Technology (NIST) announced a new international hash function standard known as SHA-3. Unlike its predecessor SHA-2, which was rooted in the Merkle-Damgård construction [Mer88, Mer90, Dam87], the new hash function standard uses Keccak [BDPA11b]—a family of cryptographic functions based on the idea of *sponge hashing* [BDPA11a]. This particular approach allows for both variable input length and variable output length, which makes it particularly attractive towards the design of cryptographic hash functions. The internal state of a sponge function gets updated through successive applications of a so-called *block function* $\varphi : \{0,1\}^{r+c} \to \{0,1\}^{r+c}$ (which is typically modeled as a random permutation), where we call the parameters $r \in \mathbb{N}$ the *rate* and $c \in \mathbb{N}$ the *capacity* of the sponge.

One-round sponge. Suppose that $\varphi : \{0,1\}^{r+c} \to \{0,1\}^{r+c}$ is a permutation. In the special case when there is only a single application of the block function, the sponge function $\mathsf{Sp}^{\varphi} : \{0,1\}^{r} \to \{0,1\}^{r}$ takes a simple form which is illustrated in Figure 1; namely, on input $x \in \{0,1\}^{r}$, the output is given by $y = \mathsf{Sp}^{\varphi}(x)$, where y corresponds to the first r bits of $\varphi(x||0^{c})$. In other words, Sp^{φ} is defined as the restriction of φ onto the first r bits of its output.

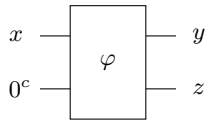

Fig. 1. The one-round sponge.

As an application of our results on indifferentiability, we show that the one-round sponge construction is both quantumly and classically indifferentiable with pre-computation from a random oracle $f : \{0,1\}^{r} \to \{0,1\}^{r}$ when the rate does not exceed the capacity (see Section 8). Our proof consists of two parts. First, in Section 7, we use *symmetrization techniques* and give a (strong, stateless) simulator with shared randomness which generates a permutation whose sponge-hash precisely matches a given random function. Next, we show how to remove the shared randomness at the cost of downgrading to a weak and stateful simulator with a single bit of advice.

Theorem 2 (Informal version of Theorem 7 and Corollary 3). *The single-round sponge is indifferentiable with pre-computation from a random oracle, both quantumly and classically, when the rate does not exceed the capacity. This holds with unbounded adversaries, either with stateless strong simulators with shared randomness, or stateful weak simulators that pass a single bit of advice.*

With shared randomness removed, we show in Section 8.2 how to derive tight quantum space-time tradeoffs for inverting the single-round sponge. To our knowledge, this is the first classical *or* quantum space-time tradeoff for sponge inversion. We summarize our results in Table 1.

Table 1. Summary of our space-time trade-offs in Section 8.2.

	Function inversion	Sponge inversion
Classical advice, classical queries	$ST = \tilde{\Omega}(\epsilon\, 2^r)$ Refs. [Yao90, DTT10]	$ST = \tilde{\Omega}(\epsilon\, 2^r)$ (**this work**)
Classical advice, quantum queries	$ST + T^2 = \tilde{\Omega}(\epsilon\, 2^r)$ Ref. [CGLQ20]	$ST + T^2 = \tilde{\Omega}(\epsilon\, 2^r)$ (**this work**)
Quantum advice, quantum queries	$ST + T^2 = \tilde{\Omega}(\epsilon^3\, 2^r)$ Ref. [CGLQ20]	$ST + T^2 = \tilde{\Omega}(\epsilon^3\, 2^r)$ (**this work**)

2.3 Related work

We now briefly discuss several related works on the topic of both (quantum) indifferentiability, pre-computation and the sponge construction.

Maurer, Renner, and Holenstein [MRH04] first proposed the notion of *indifferentiability* as a composable and simulation-based definition for what it means for a construction to be "as good as" as an ideal object. Ristenpart, Shacham, and Shrimpton [RSS11] observed that indifferentiability is insufficient for "multi-stage" games, and proposed the notion of *reset indifferentiability* instead which requires the simulator to be stateless. Bertoni, Daemen, Peeters and Van Assche [BDPVA08] proved the indifferentiability of the many-round sponge construction. Carstens, Ebrahimi, Tabia, and Unruh [CETU18] initiated the study of indifferentiability in the quantum setting, and analyzed the security of both Feistel networks and the sponge construction under conjectures. Zhandry [Zha18] introduced the compressed oracle technique, and used it to prove quantum indifferentiability of the Merkle-Damgård construction. Czajkowski, Majenz, Schaffner and Zur [CMSZ19] proved the quantum indifferentiability of the (many-round) sponge construction in the case when the block function is modeled as a random function or a random (non-invertible) permutation. Zhandry [Zha21] showed that the one-round sponge (in the special case when the message length is roughly half the block length) is quantumly reset-indifferentiable from a random oracle (even if the adversary has access to the inverse of the permutation). However, contrary to our work, none of the aforementioned works on indifferentiability take pre-computation into account.

Yao [Yao90] and De, Trevisan and Tulsiani [DTT10] gave (classical) space-time trade-offs for function inversion. Unruh [Unr07] introduced the auxiliary-input random oracle model. Nayebi, Aaronson, Belovs and Trevisan [NABT15] generalized space-time trade-offs for function inversion against quantum adversaries with classical advice. Later, Chung, Liao and Qian [CLQ20] generalized these bounds in the case of quantum advice. Hhan, Xagawa and Yamakawa [HXY19] gave space-time trade-offs for function (and permutation) inversion in the auxiliary-input quantum random oracle model. Chung, Guo, Liu and Qian [CGLQ20] gave the first tight quantum space-time trade-off for function inversion with both classical and quantum advice. Alagic, Bai, Poremba and Shi [ABPS24] showed quantum space-time trade-offs for two-sided permutation inversion—the task of inverting a random but invertible permutation, where the inverter also has access to a punctured inverse oracle. Freitag, Ghoshal and Komargodski [FGK22], and subsequently also Akshima, Duan, Guo and Liu [ADGL23, Aks24], gave space-time trade-offs for finding short collisions in the sponge construction. Carolan and Poremba [CP24] gave a (tight) quantum query lower bound for one-round sponge-inversion via symmetrization techniques. In concurrent work, Majenz, Malavolta and Walter [MMW24] also gave (non-tight) quantum query lower bounds for the task of sponge inversion (in a more general setting) via compressed oracle techniques. Ananth, Mutreja and Poremba [AMP24] recently gave space-time trade-offs for a simple query problem; namely that of finding elements in (one-round) sponge hash tables. Notably, none of the aforementioned works on function inversion result in (either classical or quantum) space-time trade-offs for the single-round sponge inversion task, as in our work.

Acknowledgements

The authors would like to thank Christian Majenz, Giulio Malavolta and Gorjan Alagic for useful discussions. JC is supported by the US Department of Energy grant no. DESC0020264. AP is supported by the National Science Foundation (NSF) under Grant No. CCF-1729369.

3 Preliminaries

Basic notation. For $N \in \mathbb{N}$, we use $[N] = \{1, 2, \ldots, N\}$ to denote the set of integers up to N. The symmetric group on $[N]$ is denoted by S_N. In slight abuse of notation, we oftentimes identify elements $x \in [N]$ with bit strings $x \in \{0,1\}^n$ via their binary representation whenever $N = 2^n$ and $n \in \mathbb{N}$. Similarly, we identify permutations $\pi \in S_N$ with permutations $\pi : \{0,1\}^n \to \{0,1\}^n$ over bit strings of length n.

Quantum computing. A finite-dimensional complex Hilbert space is denoted by $\mathcal{H}$, and we use subscripts to distinguish between different systems (or registers); for example, we let $\mathcal{H}_A$ be the Hilbert space corresponding to a system A. The tensor product of two Hilbert spaces $\mathcal{H}_A$ and $\mathcal{H}_B$ is another Hilbert space which we denote by $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$. We let $\mathcal{L}(\mathcal{H})$ denote the set of linear operators over $\mathcal{H}$. A quantum system over the 2-dimensional Hilbert space $\mathcal{H} = \mathbb{C}^2$ is called a *qubit*. For $n \in \mathbb{N}$, we refer to quantum registers over the Hilbert space $\mathcal{H} = \left(\mathbb{C}^2\right)^{\otimes n}$ as n-qubit states. We use the word *quantum state* to refer to both pure states (unit vectors $|\psi\rangle \in \mathcal{H}$) and density matrices $\rho \in \mathcal{D}(\mathcal{H})$, where we use the notation $\mathcal{D}(\mathcal{H})$ to refer to the space of positive semidefinite linear operators of unit trace acting on $\mathcal{H}$. A *unitary* $U : \mathcal{L}(\mathcal{H}_A) \to \mathcal{L}(\mathcal{H}_A)$ is a linear operator such that $U^\dagger U = UU^\dagger = I_A$, where the operator I_A denotes the identity operator on system $\mathcal{H}_A$. A quantum algorithm is a uniform family of quantum circuits $\{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$, where each circuit $\mathcal{A}_\lambda$ is described by a sequence of unitary gates and measurements; moreover, for each $\lambda \in \mathbb{N}$, there exists a deterministic Turing machine that, on input 1^λ, outputs a circuit description of $\mathcal{A}_\lambda$. We say that a quantum algorithm $\mathcal{A}$ has oracle access to a classical function $f : \{0,1\}^n \to \{0,1\}^m$, denoted by $\mathcal{A}^f$, if $\mathcal{A}$ is allowed to use a unitary gate O^f at unit cost in time. The unitary O^f acts as follows on the computational basis states of a Hilbert space $\mathcal{H}_X \otimes \mathcal{H}_Y$ of $n + m$ qubits:

$$O^f : \quad |x\rangle_X \otimes |y\rangle_Y \longrightarrow |x\rangle_X \otimes |y \oplus f(x)\rangle_Y ,$$

where the operation $\oplus$ denotes bit-wise addition modulo 2. Oracles with quantum query-access have been studied extensively, for example in the context of quantum complexity theory [BBBV97], as well as in cryptography [BDF+11, AHU18, AJOP20].

3.1 (Quantum) Indifferentiability

Our notation is based on [CMSZ19] which is close to the original definition of Maurer, Renner, and Holenstein [MRH04]. The basic idea is that an adversary who is interacting with some cryptographic system C has access to two *interfaces*:

- a *public* interface $\mathsf{C}_\lambda^{\text{pub}}$ (for example, a permutation φ) which is some public interface that takes as input a certain number of (qu)bits, and outputs a number of (qu)bits.
- a *private* interface $\mathsf{C}_\lambda^{\text{priv}}$ (for example, the sponge hash Sp^φ which uses the permutation φ internally) which is also some function that takes as input a certain number of (qu)bits and outputs a number of (qu)bits.

For the purposes of indifferentiability, we will consider constructions where the private interface is constructed from the public interface (i.e. $\mathsf{C}_\lambda^{\text{priv}}[\mathsf{C}_\lambda^{\text{pub}}]$ is an efficient algorithm with an oracle for $\mathsf{C}_\lambda^{\text{pub}}$). We can now define what it means for two interfaces to be indifferentiable.

Definition 1 (Indifferentiability).
Let $\lambda \in \mathbb{N}$ be the security parameter. A cryptographic system C is (T, ϵ)-indifferentiable from R, if there exists an efficient (classical or quantum) simulator $\mathcal{S}$ and a negligible function ϵ such that, for any efficient (classical or quantum) distinguisher $\mathcal{D}$ making at most T (classical or quantum) queries to C, it holds that

$$\left| \Pr\left[\mathcal{D}[\mathsf{C}_\lambda^{\text{priv}}[\mathsf{C}_\lambda^{\text{pub}}], \mathsf{C}_\lambda^{\text{pub}}] = 1\right] - \Pr\left[\mathcal{D}[\mathsf{R}_\lambda^{\text{priv}}, \mathcal{S}[\mathsf{R}_\lambda^{\text{pub}}]] = 1\right] \right| \leq \epsilon(\lambda).$$

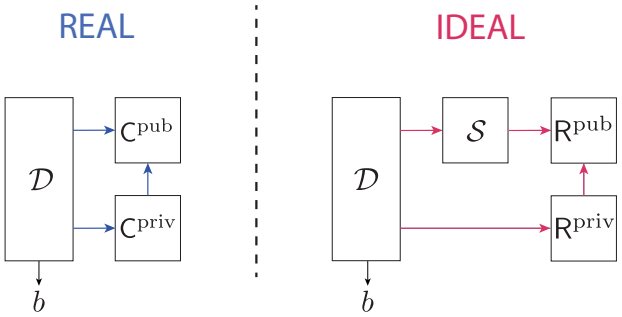

Fig. 2. Schematic representation of indifferentiability of construction C from idealized primitive R. The arrows denote "access to" the pointed system, and $\mathcal{D}$ is the distinguisher.

These notions are very general, encompassing many scenarios in which one would like to reason about building some idealized primitive R from a different idealized primitive C in a composably secure fashion. The composition theorem for plain indifferentiability states that a primitive instantiated using the private interface of C is *as secure as* a primitive instantiated using the private interface of R in many scenarios. More formally, any security game that is secure against a single adversary having oracle access to R^{pub} will also be secure against a single adversary having oracle access to C^{pub}, up to the indifferentiability loss [MRH04]. However, security games involving multiple rounds with distinct adversaries are not generically proven secure by indifferentiability [RSS11]. For this, one requires a stronger form called *reset* indifferentiability, in which the simulator $\mathcal{S}$ in Definition 1 is required to be stateless.

Often, one considers security of a cryptographic system that allows a pre-processing adversary. Security games in this setting consider an adversary split into two phases: an inefficient offline phase in which some advice is computed about the underlying primitive (in our case, the interface for C or R), and an efficient online adversary. The prepared advice is forwarded to the efficient online adversary, which then receives one or more challenge(s) that are independent of the advice. This is a more general setting than the single adversary setting of plain indifferentiability, but *a priori* it may seem like security of such a game would be implied by reset indifferentiability, as this is a security game with multiple stages and adversaries. However, we show that this is not the case: this is because the offline adversary is inefficient, which is not captured by reset indifferentiability.

In fact, even strong notions of indifferentiability do not imply non-trivial space-time trade-offs. At a high level, the counterexample is a random function with a trapdoor (as both the public and private interface). Such a function is (reset, statistical, strong, quantum or classical) indifferentiable from a random oracle for query bounded adversaries, but clearly admits pre-computation attacks for one-wayness and many other security games which would be secure with a random oracle.

4 Separating reset indifferentiability from pre-computaion

Let O_g be an oracle for a function $g : \{0,1\}^{2n} \to \{0,1\}^n$ drawn from the distribution which is uniform random, except on inputs of the form $x\|s$ for some random trapdoor $s \in \{0,1\}^n$. For such inputs, define $g(x\|s) = x$. Let oracle O_h be an oracle for a function $h : \{0,1\}^{2n} \to \{0,1\}^n$ drawn uniformly at random. Note that the distributions of g and h have large total variation distance, but are quantum query indistinguishable by the one-way to hiding lemma [Unr14]. We will consider two constructions, C and R, defined as follows.

1. Construction C is defined by $\mathsf{C}^{\text{pub}} := O_g$, and $\mathsf{C}^{\text{priv}}[\mathsf{C}^{\text{pub}}] := \mathsf{C}^{\text{pub}} = O_g$. In other words, both the public and private interface for C are the same oracle, and they are for a random function with a trapdoor.
2. Construction R is defined by $\mathsf{R}^{\text{pub}} := O_h$, and $\mathsf{R}^{\text{priv}}[\mathsf{R}^{\text{pub}}] := \mathsf{R}^{\text{pub}} = O_h$. In other words, both the public and private interface for R are the same oracle, and they are for a random function (with no trapdoor).

Now let $\mathcal{S}[\mathsf{R}^{\text{pub}}]$ be the trivial simulator which has access to h, and answers queries x as $h(x)$.

Theorem 3. *Construction* C *is* (T, ϵ) *strong quantum statistical reset indifferentiable from construction* R *for any* $\epsilon = O(T^2/2^n)$.

Proof. The interface for $\mathsf{R} = (\mathsf{R}^{\text{pub}}, \mathsf{R}^{\text{priv}}[\mathsf{R}^{\text{pub}}])$ is simply two oracles for the same random function. The simulated interface $(\mathcal{S}[\mathsf{C}^{\text{pub}}], \mathsf{C}^{\text{priv}}[\mathsf{C}^{\text{pub}}])$ is two oracles for a random function, but with a random trapdoor s such that inputs that end with s are easy to invert. We know that O_g and O_h are quantum query indistinguishable in T queries for any $\epsilon = O(T^2/2^n)$ [Unr14], and the interface above is exactly the interface exposed to the adversary in the indistinguishability game (except two copies of each oracle are exposed; this is straightforward to simulate with one copy). This shows quantum statistical indifferentiability. The simulator can be seen to not depend on the distinguisher and be stateless, hence strong and reset.

This proof straightforwardly implies the following corollary, for any choice of either bracketed item.

Corollary 1. *Construction* C *is* (T, ϵ) *strong* $\langle$ *classical* / *quantum* $\rangle$ *statistical reset indifferentiable* $\langle$ *with* / *without* $\rangle$ *shared randomness from construction* R *for any* $\epsilon = O(T^2/2^n)$.

Observe that this corollary is weaker than the one shown above, as the (classical) simulator can simply ignore it's randomness. The security bound can be strengthed in the case of a classical adversary. Observe also that, given advice $s \in \{0,1\}^n$ (which depends on C, e.g. found by an unbounded adversary querying C), it is straightforward to invert g with no queries. However, h remains hard to invert even with a much larger amount of advice, e.g. an adversary can succeed with constant probability having S qubits of advice and T quantum queries only when $ST + T^2 = \tilde{\Omega}(2^n)$ [CGLQ20].

5 Indifferentiability with Pre-computation

Given the prior counterexample, we see that all but the strongest notions of indifferentiability are insufficient to inherit security for games that allow for adversaries with pre-computed advice (e.g. space-time tradeoff lower bounds). We define in this section a notion which captures any security game allowing unbounded pre-computation.

Strong indifferentiability with pre-computation. To define the strong notion of indifferentiability with pre-computation, we will have both a pair of fixed simulators $\mathcal{S} = (\mathcal{S}_0, \mathcal{S}_1)$ and arbitrary distinguishers $\mathcal{D} = (\mathcal{D}_0, \mathcal{D}_1)$, where the first in each tuple is unbounded/offline and the second in each tuple is bounded/online. In the "real world", the offline distinguisher $\mathcal{D}_0$ receives unbounded access to some interface C. It then forwards S (qu)bits of advice to online distinguisher $\mathcal{D}_1$, which can make T queries to C^{priv} and C^{pub}, and then outputs a bit.

In the "ideal world", the offline simulator $\mathcal{S}_0$ receives unbounded access to the ideal interface R, which it then uses to implement an interface which offline distinguisher $\mathcal{D}_0$ has unbounded access to. Again $\mathcal{D}_0$ forwards S (qu)bits of advice to the online distinguisher $\mathcal{D}_1$, but we also allow $\mathcal{S}_0$ to forward S_{sim} (qu)bits of advice to the online simulator $\mathcal{S}_1$. The online distinguisher $\mathcal{D}_1$ makes T queries to R^{priv}, as well as an interface simulated by $\mathcal{S}_1$ which itself makes T_{sim} queries to R^{pub}. As before, the distinguisher outputs a bit.

Definition 2 (Strong Indifferentiability with Pre-Computation). *Let* $\lambda \in \mathbb{N}$ *be the security parameter. A cryptographic system* C *is strongly* $(S, T, S_{\text{sim}}, T_{\text{sim}}, \epsilon)$*-indifferentiable (with pre-computation) from a system* R*, if there exists a pair of simulators* $(\mathcal{S}_0, \mathcal{S}_1)$*, where*

- $\mathcal{S}_0$ *is an (classical/quantum) query and computation unbounded algorithm that outputs at most* S_{sim} *(qu)bits, and*
- $\mathcal{S}_1$ *is an efficient (classical/quantum) algorithm making* T_{sim} *(classical/quantum) queries,*

and a negligible function $\epsilon(\lambda)$ *such that, for any pair of algorithms* $(\mathcal{D}_0, \mathcal{D}_1)$*, where*

- $\mathcal{D}_0$ *is an unbounded (classical/quantum) algorithm which outputs* S*-many (qu)bits and*
- $\mathcal{D}_1$ *is an efficient (classical/quantum) making at most* T *queries to* C*,*

such that the following distinguishing property holds:

$$\Big|\Pr\Big[\mathcal{D}_1\Big[\mathsf{C}_\lambda^{\mathrm{priv}}[\mathsf{C}_\lambda^{\mathrm{pub}}],\mathsf{C}_\lambda^{\mathrm{pub}},\mathcal{D}_0[\mathsf{C}_\lambda]\Big]=1\Big]-\Pr\Big[\mathcal{D}_1\Big[\mathsf{R}_\lambda^{\mathrm{priv}},\mathcal{S}_1\big[\mathsf{R}_\lambda^{\mathrm{pub}},\mathcal{S}_0[\mathsf{R}_\lambda]_{\mathcal{S}}\big],\mathcal{D}_0[\mathcal{S}_0[\mathsf{R}_\lambda]_{\mathcal{D}}]\Big]=1\Big]\Big|\leq\epsilon(\lambda).$$

Here, we assume that $\mathcal{D}_1$ only has access to the interface D_0 via its output, i.e., it receives S many (qu)bits of advice.

The loss of the simulator, T_{sim} and S_{sim}, as well as the distinguishing advantage ϵ, will enter into the bounds inherited through this notion. Naturally, the smaller these quantities are, the tighter the bounds. We leave our definitions general so as to allow inheriting the tightest possible bounds.

Weak indifferentiability with pre-computation. The weak indifferentiability with pre-computation game is similar to strong, except the simulator $\mathcal{S}=(\mathcal{S}_0,\mathcal{S}_1)$ now depends on distinguisher $\mathcal{D}=(\mathcal{D}_0,\mathcal{D}_1)$. This allows us to consider offline simulators $\mathcal{S}_0$ which prepare both the advice for the online simulator $\mathcal{S}_1$, as well as for the online distinguisher $\mathcal{D}_1$. This can reduce to the notion of strong indifferentiability when the offline simulator $\mathcal{S}_0$ internally runs offline distinguisher $\mathcal{D}_0$. Our definition is depicted in Figure 3.

Definition 3 (Weak Indifferentiability with Pre-Computation). *Let $\lambda\in\mathbb{N}$ be the security parameter. A cryptographic system C is weakly $(S,T,S_{\mathsf{sim}},T_{\mathsf{sim}},\epsilon)$-indifferentiable (with pre-computation) from a system R, if, for any pair of algorithms $(\mathcal{D}_0,\mathcal{D}_1)$, where*

- *$\mathcal{D}_0$ is a query and computation unbounded (classical/quantum) algorithm which outputs S-many (qu)bits and*
- *$\mathcal{D}_1$ is an efficient distinguisher (with binary output) making at most T queries to C,*

there exists a pair of (classical/quantum) simulators $(\mathcal{S}_0,\mathcal{S}_1)$, where

- *$\mathcal{S}_0$ is a query and computation unbounded (classical/quantum) algorithm that outputs S_{sim} (qu)bits, and*
- *$\mathcal{S}_1$ is an efficient (classical/quantum) algorithm making T_{sim} (classical/quantum) queries,*

and a negligible function $\epsilon(\lambda)$ such that the following holds:

$$\Big|\Pr\Big[\mathcal{D}_1\Big[\mathsf{C}_\lambda^{\mathrm{priv}}[\mathsf{C}_\lambda^{\mathrm{pub}}],\mathsf{C}_\lambda^{\mathrm{pub}},\mathcal{D}_0[\mathsf{C}_\lambda]\Big]=1\Big]-\Pr\Big[\mathcal{D}_1\Big[\mathsf{R}_\lambda^{\mathrm{priv}},\mathcal{S}_1\big[\mathsf{R}_\lambda^{\mathrm{pub}},\mathcal{S}_0[\mathsf{R}_\lambda]_{\mathcal{S}}\big],\mathcal{S}_0[\mathsf{R}_\lambda]_{\mathcal{D}}\Big]=1\Big]\Big|\leq\epsilon(\lambda).$$

Here, we assume that $\mathcal{D}_1$ only has access to the interface D_0 via its output, i.e., it receives S many (qu)bits of advice.

Additional Variants. Moreover, we consider the following variants which apply to both weak and strong indifferentiability with pre-computation:

- **Computational/statistical/perfect:** *Computational* indifferentiability requires the distinguisher $\mathcal{D}_1$ to be a computationally efficient algorithm. *Statistical* indifferentiability requires $\mathcal{D}_1$ to make a bounded number of queries. Finally, *perfect* indifferentiability refers to the case when $\mathcal{D}_1$ is completely unbounded.
- **Computational/statistical simulation:** Indifferentiability with *computational simulation* requires $\mathcal{S}_1$ to be a computationally efficient simulator. Indifferentiability with *statistical simulation* requires $\mathcal{S}_1$ to be only query-efficient (but computation unbounded).
- **Shared randomness:** In *shared randomness* indifferentiability, the simulators $\mathcal{S}_0$ and $\mathcal{S}_1$ have access to the same arbitrary-sized set of random coins $\mathsf{SR}\in\{0,1\}^*$. Note that we do not reveal this randomness to the distinguisher.
- **Classical/quantum:** *Quantum* indifferentiability captures the security against quantum distinguishers (and also simulators) making classical or quantum queries to their oracles, whereas *classical* indifferentiability only considers classical algorithms that make classical queries. In the context of *quantum* indifferentiability, we also distinguish between *classical* and *quantum* pre-computation, i.e., whether $\mathcal{D}_1$ receives classical or quantum advice from $\mathcal{D}_0$.

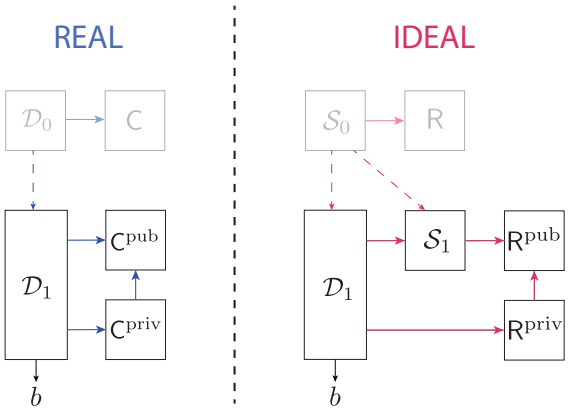

Fig. 3. Schematic representation of weak indifferentiability with pre-computation. Arrows denote access to the pointed to interface, and washed out colors denote inefficient pre-computation, i.e. entities with unbounded access. Dotted arrows denote forwarded advice.

5.1 Perfect reset indifferentiability suffices for pre-computation

In this section, we observe that if two constructions R, C are *perfectly* reset indifferentiable, even for query and computation unbounded adversaries, then any multi-stage security game with (some or all) adversaries unbounded that is secure in the R model will also be secure in the C model. This includes our security games with pre-computation as a special case. Further, it holds even in the case where an unbounded adversary can distinguish with only negligible advantage ϵ (whereas "perfect indifferentiability" usually requires exactly zero advantage)—we call this notion ϵ-perfect indifferentiability.

Lemma 1. *Suppose that construction* C *is* ϵ*-perfect reset indifferentiable from* R*, and any choice of the remaining variants. Suppose that the simulator makes (at most)* T_{sim} *queries to* R^{pub} *to implement* T *of the distinguisher's queries. Then* C *is* $(S, T, S_{\mathsf{sim}}, T_{\mathsf{sim}}, \epsilon)$ *indifferentiable with pre-computation from* R*, for the same choice of remaining variants, and for* $S_{\text{sim}} = 0$*,* T_{sim} *and* ϵ *as defined above, and any number of distinguisher queries* T *and distinguisher advice size* S*.*

Proof. We will prove the claim separately for strong and weak indifferentiability, though the proofs proceed analogously for any choice of the remaining variants.

- *(Case 1)* Strong indifferentiability with pre-computation. Let $\mathcal{S}[\mathsf{R}^{\text{pub}}]$ be the simulator that achieves ϵ-perfect reset indifferentiability from R (note that $\mathcal{S}$ may also be a function of some shared randomness). To construct a simulator in the strong indifferentiability with pre-computation setting $\mathcal{S}' = (\mathcal{S}'_0, \mathcal{S}'_1)$, we simply identify $\mathcal{S}'_0[\mathsf{R}] = (\mathcal{S}[\mathsf{R}^{\text{pub}}], \mathsf{R}^{\text{priv}})$ and $\mathcal{S}'_1[\mathsf{R}^{\text{pub}}] = \mathcal{S}[\mathsf{R}^{\text{pub}}]$. We know that $\mathcal{S}$ is stateless, so no advice needs to be passed from $\mathcal{S}'_0$ to $\mathcal{S}'_1$ while running $\mathcal{S}$. Any distinguisher $\mathcal{D}' = (\mathcal{D}'_0, \mathcal{D}'_1)$ in the strong indifferentiability with pre-computation game against $\mathcal{S}$ can now be identified with the distinguisher $\mathcal{D} = \mathcal{D}'$ in the strong ϵ-perfect reset indifferentiability game (i.e. $\mathcal{D}$ simply runs $\mathcal{D}'_0$ followed by $\mathcal{D}'_1$, passing advice internally as needed, and outputting the result). By the reset indifferentiability of C from R, this will distinguish with advantage at most ϵ and hence $\mathcal{D}'$ distinguishes with advantage at most ϵ.
- *(Case 2)* Weak indifferentiability with pre-computation. Let $\mathcal{D}' = (\mathcal{D}'_0, \mathcal{D}'_1)$ be a distinguisher for the indifferentiability with pre-computation game. Identify $\mathcal{D}'$ with a distinguisher $\mathcal{D}$ in the ϵ-perfect reset indifferentiability game of C from R, specifically the distinguisher $\mathcal{D}$ which runs $\mathcal{D}'$ internally in two stages, forwarding advice where necessary. Let $\mathcal{S}[\mathsf{R}^{\text{pub}}]$ be a stateless simulator such that $\mathcal{D}$ achieves advantage at most ϵ when run using $\mathcal{S}$.
  We construct $\mathcal{S}' = (\mathcal{S}'_0, \mathcal{S}'_1)$ in the indifferentiability with pre-computation game as $\mathcal{S}'_0[\mathsf{R}] = \mathcal{D}'_0[\mathcal{S}[\mathsf{R}^{\text{pub}}], \mathsf{R}^{\text{priv}}]_{\mathcal{D}}$ (meaning the offline simulator runs the offline distinguisher with simulator $\mathcal{S}$ and forwards the advice created to $\mathcal{D}'_1$, and nothing to online simulator $\mathcal{S}'_1$) and $\mathcal{S}'_1[\mathsf{R}^{\text{pub}}] = \mathcal{S}[\mathsf{R}^{\text{pub}}]$. By the ϵ-perfect reset indifferentiability achieved by $\mathcal{S}$ we know that $\mathcal{D}$ will distinguish with advantage at most ϵ, which implies that $\mathcal{D}'$ will distinguish with advantage at most ϵ.

In neither case did we place any constraints on the size of the advice from offline to online distinguisher, nor bound the number of queries of online distinguisher, so this holds for any S, T. The simulator overhead is clearly at most T_{sim}, as the simulator $\mathcal{S}$ receives T online queries in each case, and there is no advice passed between simulators so $S_{\mathrm{sim}} = 0$.

Our proof of indifferentiability with pre-computation and shared randomness of the one-round sponge from a random oracle—Section 8—will first show that the one-round sponge is perfect reset indifferentiability with shared randomness (a strengthening of [Zha21], Theorem 10), using a symmetrization argument. However, it is unclear how to remove the shared randomness from the reset indifferentiability game—in particular, techniques based on extracting randomness from an oracle such as used in [Zha21], do not work against an adversary that can learn the full oracle, as in our setting. Restricting to indifferentiability with pre-computation as defined in Section 5, we next show how to remove the shared randomness from this definition, as needed for the composition theorem.

5.2 Removing shared randomness from weak indifferentiability

A useful fact we will prove here is that weak (or strong) indifferentiability with pre-computation, statistical simulation, and shared randomness (and any choice of the remaining variants) implies weak indifferentiability with pre-computation *without* shared randomness and statistical simulation (for the same choice of remaining variants), up to a single bit of loss in the advice size.

Lemma 2. *Suppose that* C *is weakly* $(S, T, S_{\mathsf{sim}}, T_{\mathsf{sim}}, \epsilon)$ *indifferentiable with pre-computation from* R *with statistical simulation, shared randomness, as well as with a* $\langle$*computational / statistical / perfect*$\rangle$ $\langle$*classical / quantum*$\rangle$ *distinguisher. Then, the construction* C *is weakly* $(S, T, S_{\mathrm{sim}} + 1, T_{\mathrm{sim}}, \epsilon)$ *indifferentiable with precomputation from* R *with statistical simulation, no shared randomness, and* $\langle$*computational / statistical / perfect*$\rangle$ $\langle$*classical / quantum*$\rangle$ *distinguisher (for the same choice of variants as in the premise).*

The transformation is depicted in Figure 4.

Proof. Let $\mathcal{D} = (\mathcal{D}_0, \mathcal{D}_1)$ be a distinguisher for the indifferentiability with pre-computation security game between C and R, and let $\mathcal{S} = (\mathcal{S}_0, \mathcal{S}_1)$ be the simulator which witnesses the indifferentiability with pre-computation and shared randomness, incurring query loss T_{sim} and advice size loss S_{sim}. Let p be the probability that $\mathcal{D}$ outputs 1 when run in the "ideal" world, i.e. with simulator $\mathcal{S}$ and system R. From the premise, we know that p is at most ϵ away from the probability that $\mathcal{D}$ outputs 1 when run in the "real" world, with no simulator and system C. We will construct an $\mathcal{S}' = (\mathcal{S}'_0, \mathcal{S}'_1)$ which uses no shared randomness and incurs query loss T_{sim} and advice size loss $S_{\mathrm{sim}} + 1$, and causes $\mathcal{D}$ to output 1 with probability p as well. Note that this suffices because any distinguisher for the game with shared randomness is equally well a distinguisher for the game without shared randomness and vice versa; the distinguisher interfaces match syntactically.

Let us denote by $\mathcal{S}[\cdots, \mathsf{SR}]$ a simulator with shared randomness $\mathsf{SR} \in \{0,1\}^*$. We will consider simulators where this input is hard coded to value SR (recall that $\mathcal{S}$ is computationally unbounded, so this is valid). We then have two cases.

- *(Case 1)* For any $\mathsf{SR} \in \{0,1\}^*$, we have

$$\Pr\left[\mathcal{D}_1\left[\mathsf{R}_\lambda^{\mathrm{priv}}, \mathcal{S}_1\left[\mathsf{R}_\lambda^{\mathrm{pub}}, \mathcal{S}_0[\mathsf{R}_\lambda^{\mathrm{pub}}, \mathsf{SR}]_{\mathcal{S}}, \mathsf{SR}\right], \mathcal{S}_0[\mathsf{R}_\lambda^{\mathrm{pub}}, \mathsf{SR}]_{\mathcal{D}}\right] = 1\right] = p.$$

  In this case, we can simply choose any fixed value for the shared randomness SR and hard-code it into both offline and online simulators.
- *(Case 2)* There are two values of shared randomness, $\mathsf{SR}_0, \mathsf{SR}_1 \in \{0,1\}^*$, such that

$$\Pr\left[\mathcal{D}_1\left[\mathsf{R}_\lambda^{\mathrm{priv}}, \mathcal{S}_1\left[\mathsf{R}_\lambda^{\mathrm{pub}}, \mathcal{S}_0[\mathsf{R}_\lambda^{\mathrm{pub}}, \mathsf{SR}_0]_{\mathcal{S}}, \mathsf{SR}_0\right], \mathcal{S}_0[\mathsf{R}_\lambda^{\mathrm{pub}}, \mathsf{SR}_0]_{\mathcal{D}}\right] = 1\right] = p_0$$
$$\Pr\left[\mathcal{D}_1\left[\mathsf{R}_\lambda^{\mathrm{priv}}, \mathcal{S}_1\left[\mathsf{R}_\lambda^{\mathrm{pub}}, \mathcal{S}_0[\mathsf{R}_\lambda^{\mathrm{pub}}, \mathsf{SR}_1]_{\mathcal{S}}, \mathsf{SR}_1\right], \mathcal{S}_0[\mathsf{R}_\lambda^{\mathrm{pub}}, \mathsf{SR}_1]_{\mathcal{D}}\right] = 1\right] = p_1$$

and further $p_0 < p < p_1$. In this case, we hard-code both SR_0 and SR_1 into the offline and online simulators $\mathcal{S}_0', \mathcal{S}_1'$. Then, before running any other computation, $\mathcal{S}_0'$ samples a bit $s \in \{0,1\}$ such that

$$\Pr_{\mathcal{S}'}[s=1] = \frac{p-p_0}{p_1-p_0}.$$

After selecting bit s, the constructed offline simulator $\mathcal{S}_0'$ simply runs $\mathcal{S}[\cdots, \mathsf{SR}_s]$, i.e. the initial simulator with shared randomness hard-coded as dictated by s. In addition to whatever advice $\mathcal{S}_0$ would send, $\mathcal{S}_0'$ also sends the bit s to $\mathcal{S}_1'$. Then, $\mathcal{S}_1'$ runs $\mathcal{S}_1[\cdots, \mathsf{SR}_s]$, again hard-coding the shared randomness as dictated by s. This clearly incurs only an overhead of one (qu)bit of advice. We now have

$$\begin{aligned}&\Pr\left[\mathcal{D}_1\left[\mathsf{R}_\lambda^{\text{priv}}, \mathcal{S}_1'\left[\mathsf{R}_\lambda^{\text{pub}}, \mathcal{S}_0'[\mathsf{R}_\lambda^{\text{pub}}]_{\mathcal{S}}\right], \mathcal{S}_0'[\mathsf{R}_\lambda^{\text{pub}}]_{\mathcal{D}}\right] = 1\right] \\ &= p_0 \cdot \Pr_{\mathcal{S}'}[s=0] + p_1 \cdot \Pr_{\mathcal{S}'}[s=1] = p,\end{aligned}$$

which proves the claim.

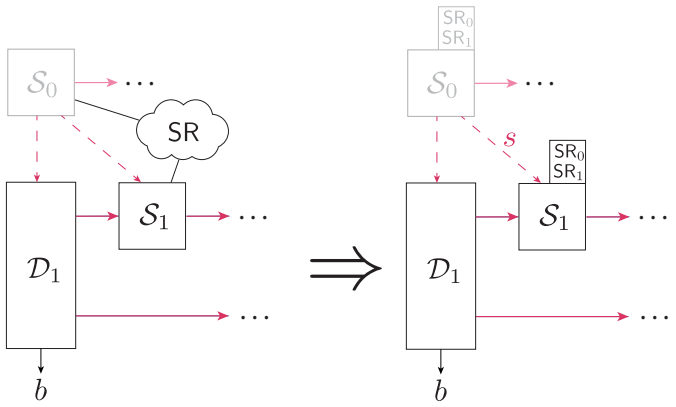

Fig. 4. The reduction for removing shared randomness for weak indifferentiability. Two values of shared randomness are hard-coded into the simulator, which then uses bit s to select between them.

6 Composition Framework for Indifferentiability with Pre-Computation

We show here that our proposed notion of indifferentiability with pre-computation implies composed security for a class of security games allowing pre-computing adversaries. This class of games includes the setting of most space-time tradeoffs as a special case.

6.1 Security games with pre-computation

For the remainder of this section, all objects will in fact be a family indexed by a security parameter λ; we drop this index for convenience. We will follow the indifferentiability model of Maurer et al. [MRH04]. In particular, let C be a construction with a private interface C^{priv} and a public interface C^{pub} as defined in Section 5. The components of a generic security game with pre-computation will be a cryptosystem $\mathcal{P}$, an enviroment $\mathcal{E}$, a pre-computation adversary $\mathcal{A}_0$, and an online adversary $\mathcal{A}_1$. For simplicity, we will quantify efficiency only in terms of number of queries to the interface C; all entities will be computationally unbounded. We further state our results for quantum queries and advice with weak indifferentiability, though one could define an analogous framework for classical queries and/or classical advice and/or quantum or classical computationally bounded entities and/or strong indifferentiability. The composition theorem proceeds similarly for all such cases. We do not consider composition theorems with shared randomness, as we will not use them.

In the following definitions, when we say that two objects "interact", we mean that they alternatively send data back and forth for some number of rounds. The number and order of rounds, as well as size and type (classical or quantum) of the data is a property of a specific security game. We will leave our definition general enough to capture any suitable security game. The only parameters we will need in a general game are the advice size S, and online query count T (which counts queries made by $\mathcal{A}_1$ and cryptosystem $\mathcal{P}$).

Definition 4. *An offline adversary $\mathcal{A}_0$ is a (computational and query) unbounded quantum algorithm, which interacts with C (both C^{priv} and C^{pub}) for an arbitrary number of rounds, and prepares an advice state α that is S qubits.*

Definition 5. *An online adversary $\mathcal{A}_1$ is an interactive quantum algorithm, which makes T_1 queries to the public interface C^{pub}, and interacts with the environment $\mathcal{E}$ and cryptosystem $\mathcal{P}$.*

Definition 6. *A cryptosystem $\mathcal{P}$ is an interactive quantum algorithm, which makes T_2 queries to the private interface C^{priv}, and interacts with the environment $\mathcal{E}$ and online adversary $\mathcal{A}_1$.*

Definition 7. *An environment $\mathcal{E}$ is an interactive quantum algorithm, which interacts with the cryptosystem $\mathcal{P}$ and online adversary $\mathcal{A}_1$. At the end of the experiment, $\mathcal{E}$ outputs a bit b.*

In Definition 5 and Definition 6, we require that $T = T_1 + T_2$. We call the tuple $(\mathcal{P}, \mathsf{C}, \mathcal{E})$ an instance of the $\mathcal{P}$ cryptosystem in model C. We can also consider a different interface R with a private interface R^{priv} that syntactically matches C^{priv} (e.g. if C^{priv} is an oracle for a function from $\{0,1\}^* \to \{0,1\}^n$, then R^{priv} is as well). We call the tuple $(\mathcal{P}, \mathsf{R}, \mathcal{E})$ an instance of the $\mathcal{P}$ cryptosystem in model R. We are now ready to state our composition theorem.

Theorem 4. *Suppose that construction C is $(S, T, S_{\mathsf{sim}}, T_{\mathsf{sim}}, \epsilon)$ indifferentiable with pre-computation from construction R. Let $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ be an attacker in the C model of $\mathcal{P}$ with advice size S and online query count T_1, in a game where $\mathcal{P}$ makes T_2 queries to C^{priv} such that $T_1 + T_2 = T$. Then there is an attacker $\mathcal{A}' = (\mathcal{A}'_0, \mathcal{A}'_1)$ in the R model of $\mathcal{P}$ with advice size $S + S_{\mathsf{sim}}$ and online query count T_{sim}. This attacker satisfies*

$$\Big| \Pr\left[\mathcal{E}\left[\mathcal{P}[\mathsf{C}^{\text{priv}}], \mathcal{A}_1[\mathsf{C}^{\text{pub}}, \mathcal{A}_0[\mathsf{C}]]\right] = 1\right] - \Pr\left[\mathcal{E}\left[\mathcal{P}[\mathsf{R}^{\text{priv}}], \mathcal{A}'_1[\mathsf{R}^{\text{pub}}, \mathcal{A}'_0[\mathsf{R}]]\right] = 1\right] \Big| \leq \epsilon$$

The construction of $\mathcal{A}$ is depicted in Figure 6, and the proof of correctness is depicted in Figure 7.

Proof. Let $\mathcal{S}[\mathsf{R}] = (\mathcal{S}_0[\mathsf{R}], \mathcal{S}_1[\mathsf{R}^{\text{priv}}])$ be a simulator which is $(S, T, S_{\mathsf{sim}}, T_{\mathsf{sim}}, \epsilon)$-indifferentiable in the indifferentiability with pre-computation game against interface C, for a certain distinguisher $\mathcal{D} = (\mathcal{D}_0, \mathcal{D}_1)$ which will be defined later. We can construct adversary $\mathcal{A}'$ in the R model of $\mathcal{P}$ from the adversary $\mathcal{A}$ in the C model of $\mathcal{P}$ and the simulator $\mathcal{S}$ in a black-box way. In particular, the pre-processing adversary $\mathcal{A}'_0$ is simply $\mathcal{S}_0[\mathsf{R}^{\text{priv}}]$, which prepares advice state $\alpha_{\mathcal{A}}$ of size S and $\alpha_{\mathcal{S}}$ of size S_{sim}. The online adversary is then the joint system of the online adversary and online simulator, $\mathcal{A}'_1 = \mathcal{A}_1[S_1[\mathsf{R}^{\text{priv}}, \alpha_{\mathcal{S}}], \alpha_{\mathcal{A}}]$. It is clear that this adversary $\mathcal{A}'$ uses $S + S_{\mathsf{sim}}$ qubits of advice and T_{sim} quantum queries (note that the original T_1 queries by the adversary are now queries to the simulator, and not to the construction; hence we only need count T_{sim}).

To show the success probability gap, let δ_1 be the probability that the environment $\mathcal{E}$ outputs 1 in the C model of $\mathcal{P}$ with adversary $\mathcal{A}$, and let δ_2 be the probability for the same environment $\mathcal{E}$ outputting 1 in the R model of $\mathcal{P}$ with adversary $\mathcal{A}'$. Note that the tuple $(\mathcal{E}, \mathcal{P}, \mathcal{A}_1)$ is a valid online distinguisher for the indifferentiability with pre-computation game, which we call $\mathcal{D}_1$. Similarly, $\mathcal{A}_0$ is a valid preprocessing distinguisher, which we call $\mathcal{D}_0$. Hence we have a distinguisher $\mathcal{D} = (\mathcal{D}_0, \mathcal{D}_1)$ for the indifferentiability with pre-computation game—this distinguisher $\mathcal{D}$ is the one which defines $\mathcal{S}$ earlier. Further, it will output 1 with probability δ_1 in the ideal world (from the definition of our security game), and with probability δ_2 in the real world (from the construction of $\mathcal{A}'$). Hence, we must have $|\delta_1 - \delta_2| \leq \epsilon$, demonstrating the claim.

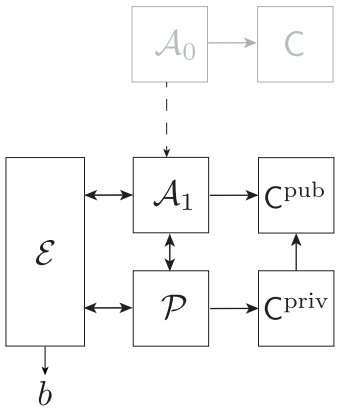

Fig. 5. The adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ in model C.

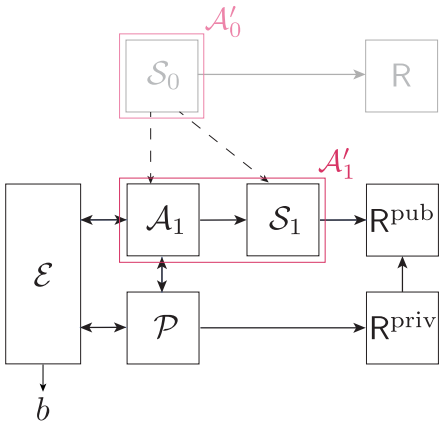

Fig. 6. The adversary $\mathcal{A}' = (\mathcal{A}'_0, \mathcal{A}'_1)$ in model R, constructed from $\mathcal{A}$ in model C and the simulator $\mathcal{S}$. The syntax of $\mathcal{A}$ is depicted in Figure 5.

7 Sponge symmetrization

In this section we introduce the relevant background building up to our symmetrization lemma. This lemma will be the main technical component in showing that the single round sponge is indifferentiable with pre-computation from a random oracle, when the rate is smaller than the capacity.

7.1 Group theory

We first define relevant notions in group theory, and recall known results about the symmetric group S_N on N elements. For a more complete overview of the subject, we refer the reader to the work of James [Jam84]. This presentation follows [CP24].

Let S_N denote the symmetric group consisting of permutations which act on the set $[N] := \{1, ..., N\}$. For a subset $A \subset [N]$, let S_A denote the maximal subgroup of S_N which fixes every element in the complement of A, i.e. $[N] \setminus A$. Now let $A_1, ..., A_l$ be a partition of $[N]$ such that the disjoint union satisfies $\bigsqcup_{i\in[\ell]} A_i = [N]$.

Definition 8 (Young subgroup). *A subgroup H of the symmetric group S_N is a Young subgroup if it can be expressed as $H = S_{A_1} \times ... \times S_{A_l}$, where $\times$ denotes the internal direct product and the collection of subsets $\{A_i\}_{i\in[\ell]}$ forms a partition of $[N]$.*

The concept of a double coset, which we review below, will also be relevant.

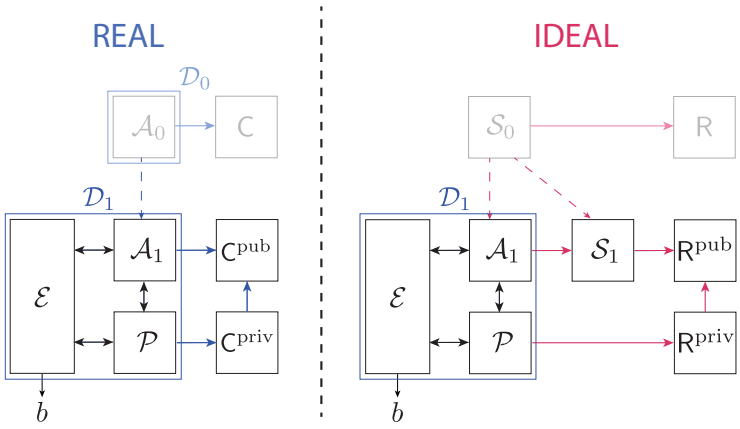

Fig. 7. The reduction from the constructed adversary in the R model to a distinguisher $\mathcal{D} = (\mathcal{D}_0, \mathcal{D}_1)$ for the indifferentiability game between C and R.

Definition 9 (Double cosets). *Let H, K be subgroups of a group G. The double cosets of G under (H, K), denoted $H \backslash G / K$, are the sets of elements which are invariant under left multiplication by H and right multiplication by K. In particular,*

$$H \backslash G / K = \big\{\{hxk : h \in H, k \in K\} : x \in G\big\}.$$

It is well-known that G is the disjoint union of its double cosets for any subgroups $H, K \leq G$. We focus on the double cosets of the symmetric group S_N, specifically those generated by Young subgroups. These subgroups admit the following characterization, adapted from Jones [Jon96] and James [Jam84].

Theorem 5 ([Jon96], Theorem 2.2). *Let H, K be Young subgroups of S_N, with corresponding partitions $A_1, ..., A_l$ (for H) and $B_1, ..., B_m$ (for K). Let $\pi \in S_N$ and $C = H\pi K$ be the corresponding double coset. Any other permutation $\pi' \in S_N$ is in C if and only if for all $i \leq l, j \leq m$ we have $|A_i \cap \pi' B_j| = |A_i \cap \pi B_j|$.*

Intuitively, the above characterization says that the double cosets defined by Young subgroups (H, K) correspond to sets of permutations which look the same if one only considers how they distribute the elements of each B_j among the different A_i's. Two permutations π, π' are in the same double coset if and only if for every A_i, B_j both π and π' send the same number of elements from B_j to A_i. This characterization combined with the following lemma is a key component of our reduction.

Theorem 6 ([Wil], Theorem 4.4). *For any subgroups H, K of a (finite) group G with $x, g \in G$ both in the same (H, K) double coset, there are exactly $|x^{-1}Hx \cap K|$ ways of choosing $h \in H$ and $k \in K$ such that $g = hxk$.*

As a corollary, it follows that selecting random symmetrizing elements from H and K suffice to give a random element of the double coset.

Lemma 3. *For any subgroups H, K of a (finite) group G with $x \in G$, let $h \sim H$ and $k \sim K$ be uniform random. Then $g = hxk$ is uniform random over the double coset HxK.*

Proof. Fix some $g \in HxK$ from the double coset of x. The number of ways to choose an $h \in H$ and a $k \in K$ such that $hxk = g$ is independent of g; hence for any g, the probability of obtaining $g = hxk$ is the same.

7.2 Symmetrization lemma

In this section, we construct a symmetrized permutation φ from a random function f such that the sponge hash of φ exactly matches the function f (at least when $r \leq c$), and φ is exponentially close to uniform random. Furthermore, a query to φ or φ^{-1} can be implemented with a single query to f. At a high level, we pick symmetrizing permutations such that the double cosets consist exactly of permutations with the same sponge hash. We use the notation $n = r + c$, and $\lambda = \min(r, c)$. We assume $r \leq c$ which implies $r = \lambda$.

Lemma 4 (Symmetrization of the one-round sponge). *Let $f : \{0,1\}^r \to \{0,1\}^r$ be a random function. Let $\mathcal{C} \subset \mathcal{S}_{2^n}$ be a subset of permutations on n-bit strings where $n = r + c$ and $r \leq c$, such that $\varphi \in \mathcal{C}$ if and only if $\mathsf{Sp}^\varphi = f$. Then there exists a (quantum or classical) algorithm that samples a random $\varphi \sim \mathcal{C}$, and can implement queries to φ and φ^{-1} each with a single query to f. Further, φ cannot be distinguished from a random permutation with advantage greater than $O(2^{-r/2})$.*

Proof. Define the transversal permutation π_f in terms of the random function $f : \{0,1\}^r \to \{0,1\}^r$. Then, for $x \in \{0,1\}^r$, $y \in \{0,1\}^r$, and $g \in \{0,1\}^{n-2r}$ (observing that g may be the empty string)

$$\pi_f(x\|g\|y) = y \oplus f(x)\|g\|x. \tag{1}$$

This construction is such that $\mathsf{Sp}^{\pi_f} = f$, e.g. the functions have the same truth table. Note also that both π_f and π_f^{-1} can be implemented using a single query to f. To construct the right symmetrizing subgroup, define the singleton sets B_z for all $z \in \{0,1\}^r$ as

$$B_z := \{z\|0^c\}, \tag{2}$$

and define $B_\perp$ as the complement of all B_z, e.g.

$$B_\perp := \{0,1\}^n \setminus \cup_{z\in\{0,1\}^r} B_z. \tag{3}$$

We now observe that the sets $\{B_z\}_{z\in\{0,1\}^r\cup\{\perp\}}$ define a partition of $\{0,1\}^n$. Thus, we can define the corresponding right symmetrizing Young subgroup $K \leq S_{2^{r+c}}$ as

$$K := \big\{\sigma \in S_{2^{r+c}} \,:\, \sigma(B_z) = B_z, \forall z \in \{0,1\}^r \cup \{\perp\}\big\}. \tag{4}$$

For the left symmetrizing subgroup, define the sets $A_x \subset \{0,1\}^n$ as

$$A_x := \{x\|y \,:\, y \in \{0,1\}^c\}, \tag{5}$$

which are easily seen to partition the set $\{0,1\}^n$. The $\{A_x\}_{x\in\{0,1\}^r}$ sets therefore define the left symmetrizing Young subgroup $H \leq S_{2^{r+c}}$ with

$$H := \big\{\omega \,:\, \omega(A_x) = A_x, \forall x \in \{0,1\}^r\big\}. \tag{6}$$

Suppose we now sample $\omega \sim H$ and $\sigma \sim K$ uniformly at random and symmetrize π_f to create a new permutation φ. In particular, we let

$$\omega \sim H, \qquad \sigma \sim K, \qquad \varphi := \omega \circ \pi_f \circ \sigma. \tag{7}$$

Remark 1. If we let $G := S_{2^n}$, then the double cosets $H\backslash G/K$ are exactly sets of permutations which have the same sponge hash. In particular, π, π' are in the same double coset if and only if $\mathsf{Sp}^\pi = \mathsf{Sp}^{\pi'}$

To see the above, recall the generic characterization that two permutations are in the same young double cosets if they distribute the elements of each B_i among the A_j in exactly the same way, Theorem 5. There are two directions to show.

($\to$) Suppose π, π' define the same sponge hash function. Then for any $z \in \{0,1\}^r$, we have $\pi(z\|0^c)[:r] = \pi'(z\|0^c)[:r]$, and hence if $\pi(B_z) \subset A_i$ then $\pi'(B_z) \subset A_i$. The B_z are singleton sets, so this gives the equation

$$|\pi(B_z) \cap A_x| = |\pi'(B_z) \cap A_x|, \text{ for all } x, z \in \{0,1\}^r. \tag{8}$$

The distribution of B_z determine the distribution of $B_\perp$ as

$$|\pi(B_\perp) \cap A_x| = |A_x| - \sum_{z \in \{0,1\}^r} |\pi(B_z) \cap A_x| \tag{9}$$

and similarly for π'. Therefore, the equality holds for all x, z including $z = \perp$.

($\leftarrow$) Similar to the above argument, but in reverse; the implications go both directions.

We have from Remark 1 above and the symmetrization lemma (Lemma 3) that φ is uniform random over all permutations having the same sponge hash function as π. From Lemma 5 in Appendix A, we have that the sponge hash of π can be distinguished from the distribution induced by the sponge hash of a truly random permutation with advantage $O(2^{-r/2})$, even given the full truth table of the sponge. It follows from our symmetrization argument that φ can be distinguished from a uniform random permutation with advantage $O(2^{-r/2})$, even given the whole truth table.

8 Indifferentiability (with Pre-Computation) of the One-Round Sponge

In this section we show that the one-round sponge is ϵ-perfect[4] strongly reset indifferentiable with shared randomness, for exponentially small ϵ, both quantumly and classically. This strengthens a result of [Zha21], which proved statistical instead of perfect indifferentiability, i.e. security only against query bounded adversaries. As a corollary, we then have that the one-round sponge is strongly indifferentiable with pre-computation and shared randomness from Lemma 1. We can then remove the shared randomness from this definition at the cost of switching to weak indifferentiability with computationally unbounded simulator through Lemma 2. With this result in place, we illustrate how our composition theorem implies a tight space-time tradeoff for the one-round sponge inversion. More broadly, our composition theorem allows the one-round sponge to inherit any space-time tradeoff of a random oracle.

8.1 Proof of Indifferentiability

Theorem 7. *Let $r, c \in \mathbb{N}$ with $r \leq c$ and $\lambda = \min(r, c)$ be parameters. Let $\varphi : \{0,1\}^{r+c} \to \{0,1\}^{r+c}$ be a random permutation. Then, the (one-round) sponge construction C with*

$$\mathsf{C}_\lambda^{\mathrm{priv}} = \mathsf{Sp}^\varphi \qquad \text{and} \qquad \mathsf{C}_\lambda^{\mathrm{pub}} = (\varphi, \varphi^{-1})$$

is strong, $\langle$ classical $|$ quantum $\rangle$, is ϵ-perfect reset indifferentiable (shared randomness) from a random oracle R, where for a random $f : \{0,1\}^r \to \{0,1\}^r$, the interfaces $\mathsf{R}^{\mathrm{priv}} = \mathsf{R}^{\mathrm{pub}}$ correspond to an oracle for f, and where $\epsilon = O(2^{-r/2})$.

Proof. We will prove the result for quantum distinguisher, though the proof proceeds similarly for classical. Let $\mathcal{D}$ be an unbounded quantum algorithm which makes queries to C. Consider the following sequence of hybrid experiments:

Game 1 : This hybrid corresponds to the *real world*. The adversary $\mathcal{D}$ receives access to the (single-round) sponge construction C with

$$\mathsf{C}_\lambda^{\mathrm{priv}} = \mathsf{Sp}^\varphi \quad \text{and} \quad \mathsf{C}_\lambda^{\mathrm{pub}} = (\varphi, \varphi^{-1})$$

where $\varphi : \{0,1\}^{r+c} \to \{0,1\}^{r+c}$ is a random permutation and $r \leq c$. Define the event

$$\textbf{Game 1} := \left[b = 1 \,:\, b \leftarrow \mathcal{D}[\mathsf{Sp}^\varphi, (\varphi, \varphi^{-1})]\right].$$

Game 2 : This corresponds to the following intermediate experiment. The adversary $\mathcal{D}$ receives access to the (single-round) sponge construction C with

$$\mathsf{C}_\lambda^{\mathrm{priv}} = \mathsf{Sp}^{\hat{\varphi}} \quad \text{and} \quad \mathsf{C}_\lambda^{\mathrm{pub}} = (\hat{\varphi}, \hat{\varphi}^{-1})$$

where $\hat{\varphi} : \{0,1\}^{r+c} \to \{0,1\}^{r+c}$ is a permutation which is generated as follows:

[4] This means that query unbounded adversaries have advantage at most ϵ

Algorithm 1: $\mathsf{Sym}_{r,c}(f)$

Input: Parameters $r, c \in \mathbb{N}$ and a truth table for a function $f : \{0,1\}^r \to \{0,1\}^r$.
Output: Truth table for a permutation $\hat{\varphi} : \{0,1\}^{r+c} \to \{0,1\}^{r+c}$.

Let $\pi : \{0,1\}^{r+c} \to \{0,1\}^{r+c}$ be the permutation with

$$\pi(x\|g\|y) := y \oplus f(x)\|g\|x.$$

For $z \in \{0,1\}^r$, let $B_z = \{z\|0^c\}$ and $B_\perp = \{0,1\}^n \setminus \cup_{z\in\{0,1\}^r} B_z$;
For $x \in \{0,1\}^r$, let $A_x = \{(x\|y) \,:\, y \in \{0,1\}^c\}$;
Sample a random permutation σ from the Young subgroup $K \leq S_{2^{r+c}}$ with

$$K = \{\sigma \in S_{2^{r+c}} \,:\, \sigma(B_z) = B_z, \forall z \in \{0,1\}^r \cup \{\perp\}\};$$

Sample a random permutation ω from the Young subgroup $H \leq S_{2^{r+c}}$ with

$$H = \{\omega \in S_{2^{r+c}} \,:\, \omega(A_x) = A_x, \forall x \in \{0,1\}^\lambda\};$$

Output a truth table for the permutation $\hat{\varphi} = \omega \circ \pi \circ \sigma$.

1. Sample a uniformly random function $f : \{0,1\}^r \to \{0,1\}^r$.
2. Run $\hat{\varphi} \leftarrow \mathsf{Sym}_{r,c}(f)$ using the symmetrization procedure in Algorithm 1.

We define the corresponding hybrid event by

$$\textbf{Game 2} := \left[b = 1 \,:\, b \leftarrow \mathcal{D}[\mathsf{Sp}^{\hat{\varphi}}, (\hat{\varphi}, \hat{\varphi}^{-1})]\right].$$

Game 3 : This hybrid corresponds to the *ideal world*. The adversary $\mathcal{D}$ receives access to interfaces which are simulated by the following stateless simulator $\mathcal{S}$which has access to a random oracle $f : \{0,1\}^r \to \{0,1\}^r$ and a common source of shared randomness $\mathsf{SR} \in \{0,1\}^*$. This is the simulator for the interface $(\hat{\varphi}, \hat{\varphi}^{-1})$: it answers queries using the permutation $\hat{\varphi} \leftarrow \mathsf{Sim}^f_{r,c}(\cdot\,; \mathsf{SR})$ which can be evaluated via oracle calls to the random oracle f and shared randomness $\mathsf{SR} \in \{0,1\}^*$, where $\mathsf{Sim}^f_{r,c}(\cdot\,; \mathsf{SR})$ is the procedure in Algorithm 2 which internally calls f.
We define the corresponding event for the ideal world by

$$\textbf{Game 3} := \left[b = 1 \,:\, b \leftarrow \mathcal{D}[f, \mathcal{S}[f, \mathsf{SR}]]\right].$$

First, we show the following:

Claim.

$$|\Pr[\textbf{Game 2}] - \Pr[\textbf{Game 1}]| \leq O(2^{-r/2}).$$

The claim follows from sponge symmetrization, Lemma 4. In particular, it was shown that distinguishing the truth table of a uniform random φ from the truth table of a symmetrized φ constructed from a random $f : \{0,1\}^r \to \{0,1\}^r$ can be done with advantage at most $O(2^{-r/2})$. The permutation φ uniquely determines Sp^φ in both **Game 1** and **Game 2**, so this suffices to prove the claim. Finally, we observe the following:

Claim.

$$\Pr[\textbf{Game 3}] = \Pr[\textbf{Game 2}].$$

Proof. Note that in the previous experiment, **Game 2**, we already introduced a perfectly random function, which is now featured as a random oracle. To prove the claim, it suffices to argue that a sufficiently long shared random strong $\mathsf{SR} \in \{0,1\}^*$ enables the of simulator $\mathcal{S}$ to run Algorithm 2 to generate the same symmetrizing permutations $\sigma \sim K$ and $\omega \sim H$ on each query (despite being stateless). Define $N = 2^{r+c}$ and recall that $H, K \leq S_N$ are both Young subgroups. Let $B = \{B_z\}_{z\in\{0,1\}^r \cup \{\perp\}}$ denote the invariant sets with respect to K, and let $A = \{A_x\}_{x\in\{0,1\}^r}$ denote the invariant sets for H. Then,

$$K \cong \prod_{A_i \in A} A_i \qquad \text{and} \qquad H \cong \prod_{B_j \in B} B_j.$$

Algorithm 2: $\mathsf{Sim}^{f}_{r,c}(x_{\mathsf{in}}\,;\mathsf{SR})$

Input: Parameters $r, c \in \mathbb{N}$, an oracle for a function $f : \{0,1\}^r \to \{0,1\}^r$, an input string $x_{\mathsf{in}} \in \{0,1\}^{r+c}$, and a random string $\mathsf{SR} \in \{0,1\}^*$.
Output: An output string $y_{\mathsf{out}} \in \{0,1\}^{r+c}$.

1 For $z \in \{0,1\}^r$, let $B_z = \{z \| 0^c\}$ and $B_\perp = \{0,1\}^n \setminus \cup_{z\in\{0,1\}^r} B_z$;
2 For $x \in \{0,1\}^r$, let $A_x = \{(x\|y) \,:\, y \in \{0,1\}^c\}$;
3 Use a subset of the random coins in $\mathsf{SR} \in \{0,1\}^*$ to assign a random permutation σ from the Young subgroup $K \leq S_{2^{r+c}}$ with

$$K = \{\sigma \in S_{2^{r+c}} \,:\, \sigma(B_z) = B_z, \forall z \in \{0,1\}^r \cup \{\perp\}\};$$

4 Use another subset of the random coins in $\mathsf{SR} \in \{0,1\}^*$ to assign a random permutation ω from the Young subgroup $H \leq S_{2^{r+c}}$ with

$$H = \{\omega \in S_{2^{r+c}} \,:\, \omega(A_x) = A_x, \forall x \in \{0,1\}^\lambda\};$$

5 Output $y_{\mathsf{out}} = \hat{\varphi}(x_{\mathsf{in}})$, where $\hat{\varphi}$ is the symmetrized permutation $\hat{\varphi} = \omega \circ \pi_f \circ \sigma$ and where $\pi_f : \{0,1\}^{r+c} \to \{0,1\}^{r+c}$ can be evaluated with an oracle call to f via

$$\pi_f(x\|g\|y) := y \oplus f(x)\|g\|x.$$

Hence, the claim is essentially just a consequence of the *coupon collector problem*[5].

Putting everything together, we get that

$$\left|\Pr\left[\mathcal{D}[\mathsf{Sp}^{\varphi}, (\varphi, \varphi^{-1})] = 1\right] - \Pr\left[\mathcal{D}\big[f, \mathcal{S}[f, \mathsf{SR}]\big] = 1\right]\right| \leq O(2^{-r/2}).$$

We can lift ϵ-perfect reset indifferentiability to the analogous notion of indifferentiability with pre-computation, as in the following corollary.

Corollary 2. *Let $r, c \in \mathbb{N}$ with $r \leq c$ and $\lambda = \min(r, c)$ be parameters. Let $\varphi : \{0,1\}^{r+c} \to \{0,1\}^{r+c}$ be a random permutation. Then, the (one-round) sponge construction C with*

$$\mathsf{C}^{\mathrm{priv}}_\lambda = \mathsf{Sp}^{\varphi} \qquad \text{and} \qquad \mathsf{C}^{\mathrm{pub}}_\lambda = (\varphi, \varphi^{-1})$$

is strong, $\langle$ classical $|$ quantum $\rangle$, and perfect $(S, T, S_{\mathsf{sim}}, T_{\mathsf{sim}}, \epsilon)$-indifferentiable (with pre-computation and shared randomness) from a random oracle $f : \{0,1\}^r \to \{0,1\}^r$ for any parameters S and T, where $S_{\mathsf{sim}} = 0$ and $T_{\mathsf{sim}} = T$, and where $\epsilon = O(2^{-r/2})$.

Proof. Follows from Theorem 7 and Lemma 1.

This further implies weak indifferentiability with shared randomness and a statistical simulator. Using Lemma 2, we can lift this to weak indifferentiability with pre-computation and without shared randomness at the cost of a single bit of loss in advice size, with the remaining variants set the same way.

Corollary 3. *Let $r, c \in \mathbb{N}$ and $r \leq c$, with $\lambda = r$ be parameters. Let $\varphi : \{0,1\}^{r+c} \to \{0,1\}^{r+c}$ be a random permutation. Then, the (one-round) sponge construction C with*

$$\mathsf{C}^{\mathrm{priv}}_\lambda = \mathsf{Sp}^{\varphi} \qquad \text{and} \qquad \mathsf{C}^{\mathrm{pub}}_\lambda = (\varphi, \varphi^{-1})$$

is weak, $\langle$ classical $|$ quantum $\rangle$ and perfect $(S, T, S_{\mathsf{sim}}, T_{\mathsf{sim}}, \epsilon)$-indifferentiable with pre-computation and a statistical simulator, but not *shared randomness, from a random oracle $f : \{0,1\}^r \to \{0,1\}^r$ for any parameters S and T, where $S_{\mathsf{sim}} = 1$, $T_{\mathsf{sim}} = T$, and $\epsilon = O(2^{-r/2})$.*

Proof. Follows from Corollary 2 and Lemma 2.

[5] In fact, to generate a random permutation $\pi \in S_N$ only $O(N \log N)$ random bits suffice on average. The probability of failure can be further suppressed with additional amounts of randomness. Since the shared random string $\mathsf{SR} \in \{0,1\}^*$ can in principle be unbounded, we did not analyze the length explicitly.

8.2 Space-Time Trade-Offs for Sponge Inversion

In this section we illustrate how a (tight) quantum space-time trade-off for single-round sponge inversion is a special case of our composition theorem. In other words, we consider $\mathsf{Sp}^{\varphi} : \{0,1\}^r \to \{0,1\}^r$, where $\varphi : \{0,1\}^{r+c} \to \{0,1\}^{r+c}$ is a random permutation, in the case of non-uniform quantum adversaries that make T quantum queries to φ, φ^{-1} and take S qubits of quantum advice (which may depend arbitrarily on φ). We will here assume that $r \leq c$. Chung et al. [CGLQ20] show that any quantum algorithm which finds a pre-image of a randomly generated image $f(x)$ with probability ϵ using S qubits of advice, where f is a random function $f : [N] \to [M]$, satisfies a space-time trade-off

$$\epsilon \leq \tilde{O}\left(\sqrt[3]{\frac{S \cdot T + T^2}{\min(N, M)}}\right).$$

If S is classical advice, then [CGLQ20] manage to get a slightly better bound, namely

$$\epsilon \leq \tilde{O}\left(\frac{S \cdot T + T^2}{\min(N, M)}\right).$$

From the characterization in Corollary 3, we can simply apply Theorem 4, our composition theorem, to obtain a time-space tradeoff for the one-round sponge by lifting a result which is known for random oracles. Note that this applies for both classical and quantum space-time tradeoffs, and both classical and quantum advice. We will illustrate this procedure for sponge inversion, though it applies more generally.

Sponge inversion. We prove the following space-time trade-off relations for one-round sponge inversion.

Theorem 8 (Space-time trade-off for sponge inversion). *Let $r, c \in \mathbb{N}$ be integers such that $r \leq c$. Any (classical or quantum) inverter $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ for the one-round sponge construction $\mathsf{Sp}^{\varphi} : \{0,1\}^r \to \{0,1\}^r$ which consists of a pair of algorithms, where*

- *$\mathcal{A}_0$ prepares S (qu)bits of advice α (depending arbitrarily on φ),*
- *$\mathcal{A}_1$ receives α and makes T (classical/quantum) queries to either φ or φ^{-1},*

and where $\mathcal{A}$ succeeds with non-trivial[6] probability $\epsilon = \omega(2^{-r/2})$ for a random permutation $\varphi : \{0,1\}^{r+c} \to \{0,1\}^{r+c}$, must obey the space-time trade-offs:

- (Classical advice and queries:)
$$\epsilon \leq O\left(\frac{ST}{2^r}\right)$$
- (Classical advice and quantum queries:)
$$\epsilon \leq \tilde{O}\left(\frac{ST + T^2}{2^r}\right).$$
- (Quantum advice and queries:)
$$\epsilon \leq \tilde{O}\left(\sqrt[3]{\frac{ST + T^2}{2^r}}\right).$$

Proof. Note that these bounds are all known to hold for inverting a random function $f : \{0,1\}^r \to \{0,1\}^r$. The first is due to Yao [Yao90] and De et al. [DTT10], and the second and third due to Chung et al [CGLQ20]. We will show how the function inversion game can be modelled as a security game with pre-computation, as defined in Section 6.

Let us consider quantum queries and quantum advice, though the proof proceeds similarly for all three cases.

[6] We remark that this requirement in our theorem statement can be relaxed at the cost of including an additional additive term of $O(2^{-r/2})$ in each of the space-time trade-offs.

Function inversion. Parties will receive oracle access to an idealized random function R with

$$\mathsf{R}^{\text{priv}} = f \quad \text{and} \quad \mathsf{R}^{\text{pub}} = f$$

where $f : \{0,1\}^r \to \{0,1\}^r$ is a random function. The adversary is denoted $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$. The offline adversary $\mathcal{A}_0$ receives unbounded access to R, and the online adversary receives access to R^{pub}. The cryptosystem $\mathcal{P}$ has access to the function through R^{priv}, and the environment $\mathcal{E}$ interacts with $\mathcal{P}$ (it will not need to interact with $\mathcal{A}$ in this game). The game proceeds as follows.

1. Offline adversary $\mathcal{A}_0$ receives unbounded access to R, which it uses to prepare an S qubit state $|\alpha_{\mathcal{A}}\rangle$, which is forwarded to $\mathcal{A}_1$.
2. Cryptosystem $\mathcal{P}$ samples a random $x \sim \{0,1\}^r$, and computes $y = f(x)$ using private interface R^{priv}. It forwards y to $\mathcal{A}_1$.
3. Online adversary $\mathcal{A}_1$ returns some $x' \in \{0,1\}^r$ to $\mathcal{P}$. The cryptosystem $\mathcal{P}$ then checks whether $f(x') = y$, puts the result in a bit $b \in \{0,1\}$.
4. Cryptosystem $\mathcal{P}$ forwards b to the environment $\mathcal{E}$, which then outputs b.

Sponge inversion. Parties will receive oracle access to the (single-round) sponge construction C with

$$\mathsf{C}^{\text{priv}} = \mathsf{Sp}^{\varphi} \quad \text{and} \quad \mathsf{C}^{\text{pub}} = (\varphi, \varphi^{-1})$$

where $\varphi : \{0,1\}^{r+c} \to \{0,1\}^{r+c}$ is a random permutation. The adversary is denoted $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$. The offline adversary $\mathcal{A}_0$ receives unbounded access to C, and the online adversary receives access to C^{pub}. The cryptosystem $\mathcal{P}$ has access to the sponge hash C^{priv}, and the environment $\mathcal{E}$ interacts with $\mathcal{P}$ (it will not need to interact with $\mathcal{A}$ in this game). The game proceeds as follows.

1. Offline adversary $\mathcal{A}_0$ receives unbounded access to C, which it uses to prepare an S qubit state $|\alpha_{\mathcal{A}}\rangle$, which is forwarded to online adversary $\mathcal{A}_1$.
2. Cryptosystem $\mathcal{P}$ samples a random $x \sim \{0,1\}^r$, and computes $y = \mathsf{Sp}^{\varphi}(x)$ using private interface C^{priv}. It forwards y to $\mathcal{A}_1$.
3. Online adversary $\mathcal{A}_1$ returns some $x' \in \{0,1\}^r$ to $\mathcal{P}$. The cryptosystem $\mathcal{P}$ then checks whether $\mathsf{Sp}^{\varphi}(x') = y$, puts the result in a bit $b \in \{0,1\}$.
4. Cryptosystem $\mathcal{P}$ forwards b to the environment $\mathcal{E}$, which then outputs b.

Observe that both games are instances of a security game with pre-computation as described in Section 6. Furthermore, the cryptosystem $\mathcal{P}$ and environment $\mathcal{E}$ are the same in each game, hence these two are the same game, with the first in the R model and the second in the C model. The cryptosystem $\mathcal{P}$ makes $T_2 = 2$ queries to the private interface. Now suppose that an S, T adversary wins the game **Sponge inversion** (i.e. inversion as defined by $\mathcal{P}, \mathcal{E}$ in model C) with probability ϵ. It follows from Theorem 4 that there is an $S+1, T+2$ adversary winning **Function inversion** (i.e. inversion as defined by $\mathcal{P}, \mathcal{E}$ in model R) with probability at least $\epsilon - O(2^{-r/2})$. This proves the claim for quantum advice and queries. A similar line of reasoning, using the appropriate notion of indifferentiability with pre-computation, proves the remaining two claims.

Note that significantly improving these bounds would imply new circuit lower bounds, by a result of Corrigan-Gibbs and Kogan [CGK19]. In particular, note that a time-space tradeoff lower bound for sponge inversion implies a similar time-space tradeoff lower bound for function inversion.

References

ABPS24. Gorjan Alagic, Chen Bai, Alexander Poremba, and Kaiyan Shi. On the two-sided permutation inversion problem. *IACR Communications in Cryptology*, 1(1), 2024. 4

ACDW20. Akshima, David Cash, Andrew Drucker, and Hoeteck Wee. Time-space tradeoffs and short collisions in merkle-damgård hash functions. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology – CRYPTO 2020*, pages 157–186, Cham, 2020. Springer International Publishing. 2

ADGL23. Akshima, Xiaoqi Duan, Siyao Guo, and Qipeng Liu. On time-space lower bounds for finding short collisions in sponge hash functions. In *Theory of Cryptography: 21st International Conference, TCC 2023, Taipei, Taiwan, November 29–December 2, 2023, Proceedings, Part III*, page 237–270, Berlin, Heidelberg, 2023. Springer-Verlag. 2, 4

AHU18. Andris Ambainis, Mike Hamburg, and Dominique Unruh. Quantum security proofs using semi-classical oracles. Cryptology ePrint Archive, Paper 2018/904, 2018. https://eprint.iacr.org/2018/904. 5

AJOP20. Gorjan Alagic, Stacey Jeffery, Maris Ozols, and Alexander Poremba. On quantum chosen-ciphertext attacks and learning with errors. *Cryptography*, 4(1), 2020. 5

Aks24. Akshima. Time-Space Tradeoffs for Finding Multi-Collisions in Merkle-Damgård Hash Functions. In Divesh Aggarwal, editor, *5th Conference on Information-Theoretic Cryptography (ITC 2024)*, volume 304 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 9:1–9:22, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. 2, 4

AMP24. Prabhanjan Ananth, Saachi Mutreja, and Alexander Poremba. Revocable encryption, programs, and more: The case of multi-copy security. Cryptology ePrint Archive, Paper 2024/1687, 2024. 4

BBBV97. Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and weaknesses of quantum computing. *SIAM Journal on Computing*, 26(5):1510–1523, October 1997. 5

BBM13. Paul Baecher, Christina Brzuska, and Arno Mittelbach. Reset indifferentiability and its consequences. In Kazue Sako and Palash Sarkar, editors, *Advances in Cryptology - ASIACRYPT 2013*, pages 154–173, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. 2

BDF+11. Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In Dong Hoon Lee and Xiaoyun Wang, editors, *Advances in Cryptology – ASIACRYPT 2011*, pages 41–69, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg. 1, 5

BDPA11a. G. Bertoni, J. Daemen, M. Peeters, and G. Van Assche. Cryptographic sponge functions. Submission to NIST (Round 3), 2011. 3

BDPA11b. G. Bertoni, J. Daemen, M. Peeters, and G. Van Assche. The keccak sha-3 submission. Submission to NIST (Round 3), 2011. 3

BDPVA08. Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. On the indifferentiability of the sponge construction. In Nigel Smart, editor, *Advances in Cryptology – EUROCRYPT 2008*, pages 181–197, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg. 4

BR93. Mihir Bellare and Phillip Rogaway. Random oracles are practical: a paradigm for designing efficient protocols. In *Proceedings of the 1st ACM Conference on Computer and Communications Security*, CCS '93, page 62–73, New York, NY, USA, 1993. Association for Computing Machinery. 1

CDG18. Sandro Coretti, Yevgeniy Dodis, and Siyao Guo. Non-uniform bounds in the random-permutation, ideal-cipher, and generic-group models. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018*, pages 693–721, Cham, 2018. Springer International Publishing. 2

CDMP05. Jean-Sébastien Coron, Yevgeniy Dodis, Cécile Malinaud, and Prashant Puniya. Merkle-damgård revisited: How to construct a hash function. In Victor Shoup, editor, *Advances in Cryptology – CRYPTO 2005*, pages 430–448, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg. 1

CETU18. Tore Vincent Carstens, Ehsan Ebrahimi, Gelo Noel Tabia, and Dominique Unruh. On quantum indifferentiability. Cryptology ePrint Archive, Paper 2018/257, 2018. 4

CGK19. Henry Corrigan-Gibbs and Dmitry Kogan. The function-inversion problem: Barriers and opportunities. In Dennis Hofheinz and Alon Rosen, editors, *Theory of Cryptography*, pages 393–421, Cham, 2019. Springer International Publishing. 2, 20

CGLQ20. Kai-Min Chung, Siyao Guo, Qipeng Liu, and Luowen Qian. Tight quantum time-space tradeoffs for function inversion, 2020. 2, 4, 7, 19

CLQ20. Kai-Min Chung, Tai-Ning Liao, and Luowen Qian. Lower Bounds for Function Inversion with Quantum Advice. In Yael Tauman Kalai, Adam D. Smith, and Daniel Wichs, editors, *1st Conference on Information-Theoretic Cryptography (ITC 2020)*, volume 163 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 8:1–8:15, Dagstuhl, Germany, 2020. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. 2, 4

CMSZ19. Jan Czajkowski, Christian Majenz, Christian Schaffner, and Sebastian Zur. Quantum lazy sampling and game-playing proofs for quantum indifferentiability. Cryptology ePrint Archive, Paper 2019/428, 2019. https://eprint.iacr.org/2019/428. 4, 5

CP24. Joseph Carolan and Alexander Poremba. Quantum one-wayness of the single-round sponge with invertible permutations. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology – CRYPTO 2024*, pages 218–252, Cham, 2024. Springer Nature Switzerland. 4, 13

CPS08. Jean-Sébastien Coron, Jacques Patarin, and Yannick Seurin. The random oracle model and the ideal cipher model are equivalent. In David Wagner, editor, *Advances in Cryptology – CRYPTO 2008*, pages 1–20, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg. 2

Dam87. Ivan Damgård. Collision free hash functions and public key signature schemes. In *Advances in Cryptology - EUROCRYPT '87, Workshop on the Theory and Application of of Cryptographic Techniques, Amsterdam, The Netherlands, April 13-15, 1987, Proceedings*, volume 304 of *Lecture Notes in Computer Science*, pages 203–216. Springer, 1987. 3

DGHM13. Grégory Demay, Peter Gaži, Martin Hirt, and Ueli Maurer. Resource-restricted indifferentiability. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology – EUROCRYPT 2013*, pages 664–683, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. 2

DS16. Yuanxi Dai and John Steinberger. Indifferentiability of 8-round feistel networks. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology – CRYPTO 2016*, pages 95–120, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg. 2

DTT10. Anindya De, Luca Trevisan, and Madhur Tulsiani. Time space tradeoffs for attacks against one-way functions and prgs. In Tal Rabin, editor, *Advances in Cryptology – CRYPTO 2010*, pages 649–665, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg. 2, 4, 19

FGK22. Cody Freitag, Ashrujit Ghoshal, and Ilan Komargodski. Time-space tradeoffs for sponge hashing: Attacks and limitations for short collisions. Cryptology ePrint Archive, Paper 2022/1009, 2022. https://eprint.iacr.org/2022/1009. 2, 4

FN00. Amos Fiat and Moni Naor. Rigorous time/space trade-offs for inverting functions. *SIAM Journal on Computing*, 29(3):790–803, 2000. 2

GG21. Shoni Gilboa and Shay Gueron. The advantage of truncated permutations. *Discrete Applied Mathematics*, 294:214–223, 2021. 24

GK22. Ashrujit Ghoshal and Ilan Komargodski. On time-space tradeoffs for bounded-length collisions in merkle-damgård hashing. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022*, pages 161–191, Cham, 2022. Springer Nature Switzerland. 2

Hel80. M. Hellman. A cryptanalytic time-memory trade-off. *IEEE Transactions on Information Theory*, 26(4):401–406, 1980. 2

HK07. Shai Halevi and Hugo Krawczyk. Security under key-dependent inputs. In *Proceedings of the 14th ACM Conference on Computer and Communications Security*, CCS '07, page 466–475, New York, NY, USA, 2007. Association for Computing Machinery. 1

HKT11. Thomas Holenstein, Robin Künzler, and Stefano Tessaro. The equivalence of the random oracle model and the ideal cipher model, revisited. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, STOC '11, page 89–98, New York, NY, USA, 2011. Association for Computing Machinery. 2

HXY19. Minki Hhan, Keita Xagawa, and Takashi Yamakawa. Quantum random oracle model with auxiliary input. In *Advances in Cryptology – ASIACRYPT 2019: 25th International Conference on the Theory and Application of Cryptology and Information Security, Kobe, Japan, December 8–12, 2019, Proceedings, Part I*, page 584–614, Berlin, Heidelberg, 2019. Springer-Verlag. 2, 4

Jam84. James. *The Representation Theory of the Symmetric Group*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1984. 13, 14

Jon96. Andrew R. Jones. A combinatorial approach to the double cosets of the symmetric group with respect to young subgroups. *European Journal of Combinatorics*, 17(7):647–655, 1996. 14

LAMP12. Atul Luykx, Elena Andreeva, Bart Mennink, and Bart Preneel. Impossibility results for indifferentiability with resets. Cryptology ePrint Archive, Paper 2012/644, 2012. 2

Mer88. Ralph C. Merkle. A digital signature based on a conventional encryption function. In Carl Pomerance, editor, *Advances in Cryptology — CRYPTO '87*, pages 369–378, Berlin, Heidelberg, 1988. Springer Berlin Heidelberg. 3

Mer90. Ralph C. Merkle. A certified digital signature. In Gilles Brassard, editor, *Advances in Cryptology — CRYPTO' 89 Proceedings*, pages 218–238, New York, NY, 1990. Springer New York. 3

MMW24. Christian Majenz, Giulio Malavolta, and Michael Walter. Permutation superposition oracles for quantum query lower bounds, 2024. 4

MRH04. Ueli Maurer, Renato Renner, and Clemens Holenstein. Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In Moni Naor, editor, *Theory of Cryptography*, pages 21–39, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg. 1, 4, 5, 6, 11

NABT15. Aran Nayebi, Scott Aaronson, Aleksandrs Belovs, and Luca Trevisan. Quantum lower bound for inverting a permutation with advice. *Quantum Info. Comput.*, 15(11–12):901–913, September 2015. 4

RSS11. Thomas Ristenpart, Hovav Shacham, and Thomas Shrimpton. Careful with composition: Limitations of the indifferentiability framework. In Kenneth G. Paterson, editor, *Advances in Cryptology – EUROCRYPT 2011*, pages 487–506, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg. 1, 2, 4, 6

Unr07. Dominique Unruh. Random oracles and auxiliary input. Cryptology ePrint Archive, Paper 2007/168, 2007. 2, 4

Unr14. Dominique Unruh. Revocable quantum timed-release encryption. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology – EUROCRYPT 2014*, pages 129–146, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg. 6, 7

Wil. Mark Wildon. A model for the double cosets of young subgroups. Royal Holloway, University of London. 14

Yao90. A. C.-C. Yao. Coherent functions and program checkers. In *Proceedings of the Twenty-Second Annual ACM Symposium on Theory of Computing*, STOC '90, page 84–94, New York, NY, USA, 1990. Association for Computing Machinery. 2, 4, 19

Zha18. Mark Zhandry. How to record quantum queries, and applications to quantum indifferentiability. Cryptology ePrint Archive, Paper 2018/276, 2018. https://eprint.iacr.org/2018/276. 4

Zha21. Mark Zhandry. Redeeming reset indifferentiability and applications to post-quantum security. 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*, page 518–548, Berlin, Heidelberg, 2021. Springer-Verlag. 2, 4, 10, 16

A Technical lemmas

In this section we will show that the full truth table of Sp^{φ} for a random permutation φ, and the full truth table of a random function from r bits to r bits, can be distinguished with only exponentially small probability. This holds even when given a sample that is the entire truth table of these two functions; this will be necessary to show that our symmetrization procedure is sound.

Lemma 5. *Let $\mathcal{D}_1$ be the uniform distribution on functions from $\{0,1\}^r \to \{0,1\}^r$. Let $\mathcal{D}_2$ be the distribution on functions from $\{0,1\}^r$ to $\{0,1\}^r$ induced by sampling $\varphi : \{0,1\}^n \to \{0,1\}^n$ uniformly at random and taking Sp^{φ}. Then the maximum distinguishing advantage of an algorithm $\mathcal{A}$ given a sample (i.e. a full truth table) from $\mathcal{D}_1$ or from $\mathcal{D}_2$ satisfies*

$$\left| \Pr_{f\sim D_1}[\mathcal{A}(f)=1] - \Pr_{f\sim D_2}[\mathcal{A}(f)=1] \right| \leq O\left(2^{-\lambda/2}\right).$$

Proof. We begin with an alternative characterization of $\mathcal{D}_1$. Consider drawing a random $f' : \{0,1\}^n \to \{0,1\}^r$, and then defining $f : \{0,1\}^r \to \{0,1\}^r$ as $f(x) := f'(x||0^c)$; f is the sample. The restriction of a random function is still a random function, so this characterization is equivalent, i.e. f is a uniform random function from r bits to r bits.

Now we will use Lemma 6 to upper bound the maximum distinguishing advantage given a sample of one of the two truth tables. We can build a (classical) adversary for the game defined in Lemma 6 which queries the 2^r inputs of the form $x||0^c$ to obtain a truth table of size $2^r \times r$. We can information-theoretically distinguish a truth table that came from a truly random function (f' above, note that in this case we have the truth table of f) from those that come from a truncated permutation (φ above, note that in this case we have the truth table of Sp^{φ}) with advantage ADV. We use $q = 2^r$ classical queries to the extended function/truncated permutation f/φ to construct the truncated function, and have $m = n - \lambda$, so from Lemma 6 we have

$$\mathsf{ADV} = O\left(\frac{2^r}{2^{\frac{2n-\lambda}{2}}}\right) = O\left(2^{-\lambda/2}\right).$$

We have used the following lemma from Gilboa and Gueron.

Lemma 6 ([GG21]). *A (classical) adversary which makes q queries to either*

1. *a random function $f : \{0,1\}^n \to \{0,1\}^{n-m}$, or*
2. *a random permutation $\varphi : \{0,1\}^n \to \{0,1\}^n$ where the last m bits are discarded (i.e. not learned)*

can distinguish with advantage $O\left(\frac{q}{2^{\frac{n+m}{2}}}\right)$.