



Universiteit
Leiden
The Netherlands

Knowledge extraction in the quantum random-oracle model

Don, J.W.

Citation

Don, J. W. (2024, January 23). *Knowledge extraction in the quantum random-oracle model*. Retrieved from <https://hdl.handle.net/1887/3714359>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3714359>

Note: To cite this publication please use the final published version (if applicable).

Chapter 5

Efficient NIZKs and Signatures from Commit-and-Open Protocols

Chapter contents

5.1	Introduction	155
5.2	Preliminaries	161
5.2.1	The Compressed Oracle — Seen as Quantum Lazy Sampling	162
5.2.2	The Quantum Transition Capacity and Its Relevance	164
5.2.3	An Improved Variant of Zhandry’s Lemma	167
5.3	Some Background on (Non-)Interactive Proofs	170
5.3.1	Non-interactive Proofs and Online Extractability	170
5.3.2	$\mathcal{G}$ -soundness of C&O Σ -Protocols	171
5.3.3	The Fiat-Shamir Transformation of (C&O) Σ -Protocols . .	173
5.4	Online Extractability of the FS-Transformation:	
	The Case of Ordinary C&O Protocols	174
5.4.1	Technical Preface	175
5.4.2	Online Extractability of the Fiat-Shamir Transformation . .	179
5.4.3	The Unruh-Transformation with a Compressing Hash Function	181
5.5	Online Extractability of the FS-Transformation: The Case of Merkle-tree-based C&O Protocols	184
5.5.1	Merkle-tree-based C&O Protocols	185
5.5.2	Online Extractability of the Fiat-Shamir Transformation . .	187
5.5.3	Discussion: Application to Picnic, and Limiting the Proof Size	193
5.5.4	The Merkle-Tree-Based Unruh Transformation	195
	Bibliography	196
	Samenvatting	209
	Acknowledgements	211
	Curriculum Vitae	213

Section 5.1

Introduction

Some interactive proofs come with amazing properties like *zero-knowledge*, which intuitively allows a prover to convince a verifier that she knows the witness to an NP-statement without giving away any information about this witness. Such zero-knowledge proofs of knowledge are some of the most fascinating objects in cryptography, and possibly in all of theoretical computer science. One might suspect that their “magic” is rooted in the fact that the prover and verifier run an *interactive* protocol with each other, and that this interaction causes the verifier to be convinced. Surprisingly, if the interactive proof is of suitable form, e.g. a Σ -protocol (a 3-round public-coin protocol, Section 2.2.1), the Fiat-Shamir transformation [FS87] (Section 2.2.3) provides a natural way to remove the interaction from such protocols while preserving (most of) the security properties, resulting in *non-interactive zero-knowledge* proofs (NIZKs). The idea is to compute the challenge c as a hash $c = H(a)$ of the first message, rather than letting the verifier choose c . If the original Σ -protocol has additional soundness properties, the resulting NIZK after the Fiat-Shamir transformation is ideally suited to be turned into a *digital-signature scheme*, simply by hashing the message m to be signed together with the first message a in order to obtain the challenge c . The (former) candidates Picnic [CDG+17] and Dilithium [DKL+18b] in the NIST post-quantum cryptography competition follow this design paradigm.

This intuitive preservation of security properties under the Fiat-Shamir transformation can be formalized in the random-oracle model (ROM), where the hash function H is treated as a uniformly random function, and the security reduction gets *enhanced access* to anybody who queries the random oracle, by seeing which values are queried, and by possibly returning (random-looking) outputs. While this situation is conveniently easy to handle in a non-quantum world, complications arise in the context of post-quantum security. When studying the security of these non-quantum protocols against attackers equipped with large-enough quantum computers, it is natural to assume that such attackers have access to the public description of the employed hash function, and can therefore compute it in superposition on their quantum computers. Therefore, the proper notion of post-quantum security for random oracles is the *quantum-accessible random-oracle model (QROM)* as introduced in [BDF+11]. Due to the difficulty of recording adversarial random-oracle queries in superposition

(explained in more detail in Section 1.1.3), establishing post-quantum security in the QROM has turned out to be quite a bit more difficult compared to the regular ROM.

Previous results in [DFMS19] (Chapter 3) (and concurrently in [LZ19b]) establish that for any interactive Σ -protocol Π that is a proof of knowledge, the non-interactive $\text{FS}[\Pi]$ is a proof of knowledge in the QROM. [DFM20] simplified the technical proof and extended these results to multi-round interactive proofs (Section 3.5). However, the most desirable property from such a proof of knowledge is *online extractability*. Indeed, online extractability avoids *rewinding*, which typically causes a significant loss in the security reduction (see later for a comparison) and has other disadvantages. Thus, online extractability allows for the tightest security reductions.

Chailloux was the first to aim for showing online extractability of the Fiat-Shamir transformation in the QROM when considering the relevant class of *commit-and-open* (C&O) Σ -protocols and modelling the hash function used for the commitments (and for computing the challenge) as a random oracle. Indeed, the Fiat-Shamir transformation of such C&O Σ -protocols are known to be online extractable in the classical ROM (see e.g. discussion in [Fis05]). In a first attempt [Cha19], Chailloux tried to lift the argument to the quantum setting by means of Zhandry’s compressed-oracle technique [Zha19a] (Section 2.4), which offers a powerful approach for re-establishing ROM results in the QROM, that has been successful in many instances. Unfortunately, this first attempt contained a subtle flaw, which turned out to be unfixable, and despite changing the technical approach, the latest version [Cha21] of this work still contains a gap in the proof, which is put as an assumption.¹

In Chapter 4 we established online extractability of *interactive* C&O Σ -protocols Π in the QROM; the result applies as soon as Π satisfies some liberal notion of *special soundness*, which is typically satisfied. As pointed out in Section 4.5.5, one can use previous results from [DFMS19; LZ19a; DFM20] to reduce the extractability of the resulting non-interactive protocol $\text{FS}[\Pi]$ to

¹ Informally, quoting from [Cha21], the considered Assumption 2 is that the random oracle can be replaced with a random function of a particular form “*without harming too much the studied scheme*”. More formally, the security loss caused by the considered replacement is assumed to remain bounded by a given function of the number of oracle queries. This assumption is rather ad-hoc and non-standard in that it is very much tailored to the scheme and its proof. Furthermore, even though Assumption 2 is an assumption that could potentially be proven in future work, it is hard to judge whether proving the assumption is actually any easier than proving the security of the considered scheme *directly*, avoiding Assumption 2 — as a matter of fact, in this work we show that the latter is feasible, while Assumption 2 remains open.

the extractability of the interactive protocol Π . However, the resulting extraction error still scales as $O(\varepsilon/q^2)$, which results in a prohibitive loss for digital-signature schemes (see Table 1), leaving open the main question originally posed by Chailloux:

How to establish tight security reductions of the Fiat-Shamir transformation for commit-and-open Σ -protocols in the QROM?

As the technical quantum details of Zhandry’s compressed-oracle technique are rather complicated and only accessible for experts, a recent article by Chung, Fehr, Huang and Liao [CFHL21] attempts to give a comprehensive exposition of Zhandry’s technique. In addition, they establish a framework that allows researchers without extensive quantum knowledge to still deploy the compressed-oracle technique (in certain cases), basically by reasoning about classical quantities only. In short, the punchline of [CFHL21] is that, if applicable, one can prove *quantum* query complexity lower bounds (think of collision finding, for instance) by means of the following recipe, which is an abstraction of the technique developed in a line of works started by Zhandry [Zha19a; LZ19a; CGLQ20; HM21]. First, one considers the corresponding *classical* query complexity problem, analyzing it by simulating the random oracle using lazy sampling and showing that the database, which keeps track of the oracle queries and the responses, is unlikely to satisfy a certain property (e.g. to contain a collision) after a bounded number of queries. Then, one lifts the analysis to the quantum setting by plugging certain key observations from the classical analysis into generic theorems provided by the [CFHL21] framework.

Contributions. In this chapter, we extend the framework from [CFHL21], and use it in a conceptually new way to establish strong and tight security statements for a large, popular class of non-interactive zero-knowledge proofs and digital signature schemes. In broad strokes, our contributions are threefold.

Online extractability for a class of NIZKs in the QROM. We prove online extractability of the Fiat-Shamir transformation in the QROM for (a large class of) C&O Σ -protocols. This solves the problem considered and attacked by Chailloux. In more detail, we prove that if the considered C&O Σ -protocol satisfies some very liberal notion of special soundness, then the resulting NIZK is a proof of knowledge with online extractability in the QROM, i.e., when the hash function used for the commitments and the Fiat-Shamir transformation is modeled as a quantum-accessible random oracle.

Our security reduction is tight: Whenever a prover outputs a valid proof, the online-extractor succeeds, except with a small probability accounting for collision and preimage attacks on the involved hash functions. For previous reductions, the guaranteed extraction success probability was at least by a factor of q^2 smaller than the success probability of the prover subjected to extraction (see Table 1). This is our main technical contribution, see Theorem 5.17. Our result also applies to a variant of the Fiat-Shamir transformation where a digital signature scheme (DSS) is constructed. It thereby, for the first time, enables a multiplicatively tight security reduction for, e.g., DSS based on the MPC-in-the-head paradigm [IKOS07a], like Picnic [CDG+17], Banquet [BSK+21] and Rainier [DKR+21], in the QROM.

A more efficient Unruh transformation. When a Σ -protocol does not have the mentioned C&O structure, a non-interactive proof of knowledge with online extractability in the QROM can be obtained using the Unruh transformation [Unr15b]. For technical reasons, the Unruh transformation requires the hash function to be *length preserving*, which may result in large commitments, and thus large NIZKs and digital signature schemes. We revisit this transformation and show, by a rather direct application of our main result above, that the online extractability of the Unruh transform still holds when using a *compressing* hash function. The crucial observation is that the Unruh transformation can be viewed as the composition of a pre-Unruh transformation, which makes use of hash-based commitments and results in a C&O protocol, and the Fiat-Shamir transformation. By applying our security reduction, we obtain the tight online extractability without requiring the hash function to be length preserving.

More efficient NIZKs via Merkle tree based commitments. In real-world constructions based on C&O protocols, like e.g., the Picnic digital signature scheme, commitments and their openings are responsible for a significant fraction of the signature/proof size. For certain parameters, this cost can be reduced by using a collective commitment mechanism based on Merkle trees. This was observed in passing, e.g. in [Fis05], and is exploited in the most recent versions of Picnic. We formalize Merkle-tree-based C&O protocols and extend our main result to NIZKs constructed from them (see Theorem 5.23). Applications of this result include a security reduction of Picnic 3, the newest version of the Picnic digital signature scheme, that is significantly tighter than existing ones: An adversary against the Picnic 3 signature scheme in the QROM with success probability ε can now be used to break the underlying hard problem

with probability ε , up to some additive error terms, while previous reductions yielded at most ε^5/q^{10} , where q is the number of random oracle queries. We outline this reduction in Section 5.5.3.

We compare our reductions in detail to existing techniques in Table 1.

	$2\text{-s} \Rightarrow \text{PoK}$	$\text{PoK} \xrightarrow{\text{FS}} \text{NIZK-PoK},$ $\text{PoK} \xrightarrow{\text{FS}} \text{UF-NMA DSS}$	$2\text{-s} \xrightarrow{\text{FS}} \text{NIZK-PoK},$ $2\text{-s} \xrightarrow{\text{FS}} \text{UF-NMA DSS}$
Unruh rewinding [Unr12] + generic FS [DFMS19]	$O(\varepsilon^3)$	$O(\varepsilon/q^2)$	$O(\varepsilon^3/q^6)$
Σ -protocol OE [DFMS22a] + generic FS [DFMS19]	$\varepsilon - g(q, r, n)$	$O(\varepsilon/q^2)$	$O(\varepsilon/q^2) - g(q, r, n)$
this work: NIZK OE	-	-	$\varepsilon - h(q, r, n)$

Table 1. Comparison of the losses of different reductions for the construction of a NIZK proof of knowledge (NIZK-PoK) from a special-sound (Merkle tree based) C&O protocol with constant challenge space size C using r -fold parallel repetition and the Fiat-Shamir transformation. “OE” stands for online extraction, 2-s for special soundness, UF-NMA for plain unforgeability and DSS for digital signature scheme. If the content of a cell in row “security property A $\Rightarrow$ security property B” is $f(\varepsilon)$, this means that an adversary breaking property B with probability ε yields an adversary breaking property A with probability $f(\varepsilon)$. Grey text indicates results that do not apply to Merkle-tree-based C&O protocols like the one used to construct the digital signature schemes Picnic 2 [CDG+20] and Picnic 3 [KZ20]. The additive error terms are $g(q, r, n) = C^{-r} + O(rq2^{-n/2}) + O(q^32^{-n})$ and $h(q, r, n) = O(q^32^{-n}) + O(q^2C^{-r})$, where n is the output length of the random oracles, and q is the number of adversarial (quantum) queries to the random oracle. Finally, we note that the constants hidden by the big-O in $h(q, r, n)$ are reasonable, see Theorems 5.17 and 5.23.

Technical Overview. Our starting point is the fact that the compressed-oracle technique (Section 2.4) can be appreciated as a variant of the classical lazy-sampling technique that is applicable in the QROM. Namely, to some extent and informally described here, the compressed-oracle technique gives access to a database that contains the hash values that the adversary $\mathcal{A}$, who has interacted with the random oracle (RO), may know. In particular, up to a small error, for any claimed-to-be hash value y output by $\mathcal{A}$, one can find its preimage x by inspecting the database (and one can safely conclude that $\mathcal{A}$ does not know a preimage of y if there is none in the database). Recalling that a C&O Σ -protocol Π (formally defined in Section 2.2.1) is an interactive proof where the first message consists of hash-based commitments, and exploiting that typically some sort of special soundness property ensures that knowing sufficiently many

preimages of these commitments/ hashes allows one to efficiently compute a witness, constructing an online extractor for the Fiat-Shamir transformation $\text{FS}[II]$ then appears straightforward: The extractor $\mathcal{E}$ simply runs the (possibly dishonest) prover P^* , answering random oracle queries using the compressed oracle. Once P^* has finished and outputs a proof, $\mathcal{E}$ measures the compressed-oracle database and classically reads off any preimages of the commitments in the proof. Finally, $\mathcal{E}$ runs the special soundness extractor that computes a witness from the obtained preimages. It is, however, not obvious that the database contains the preimages of the commitments that are *not* opened in the proof, or that these preimages are correctly formed. Intuitively this should be the case: the random oracle used for the Fiat-Shamir transformation replaces interaction in that it forces the prover to choose a full set of commitments *before* knowing which ones need to be opened. The crux lies in replacing this intuition by a rigorous proof.

The main insight leading to our proof is that the event that needs to be controlled, namely that *the prover succeeds yet the extractor fails*, can be translated into a property **SUC** (as in “adversarial **SUC**cess”) of the compressed-oracle database, which needs to be satisfied for the event to hold. It is somewhat of a peculiar property though. The database properties that have led to query complexity lower bounds in prior work, e.g. for (multi-)collision finding [LZ19a; HM21; CFHL21] and similar problems [Zha19a; CGLQ20; BLZ21], require the database to contain some particular input-output pairs (e.g. pairs that collide), while the database property **SUC** additionally *forbids* certain input-output pairs to be contained.

Indeed, the framework from [CFHL21] is almost expressive enough to treat our problem. So, after a mild extension, we can apply it to prove that it is hard for any query algorithm to cause the compressed-oracle database to have property **SUC**. Analyzing the relevant classical statistical properties of **SUC** is somewhat tedious but can be done (see the proof of Lemma 5.22). The resulting bound on the probability for the database to satisfy **SUC** then gives us a bound on the probability of the event that the prover succeeds in producing a valid proof while at the same time fooling the extractor.

Whenever it is advantageous for communication complexity, a Merkle tree can be used to collectively commit to all required messages in a C&O protocol. This collective commitment is one of the optimizations that improve the performance of, e.g. Picnic 2 [CDG+20] over Picnic [CDG+17]. As the above-described argument for the extractability of C&O protocols already analyses iterated hashing (the hash-based commitments are hashed to compute the chal-

lenge), it generalizes to Merkle-tree-based C&O protocols without too much effort. We present this generalization in Section 5.5, and obtain similar bounds (see Theorem 5.23).

Additional Related Work. Besides the already mentioned work above, we note that Chiesa, Manohar and Spooner [CMS19] consider and prove security of various SNARG constructions, while we consider the Fiat-Shamir transformation of C&O protocols with a form of special soundness. Similar to [CFHL21], they also provide some tools for deducing security of certain oracle games against quantum attacks by bounding a natural classical variant of the game.

Section 5.2

Preliminaries

Our main technical proofs rely on the recently introduced framework by Chung, Fehr, Huang, and Liao [CFHL21] for proving query complexity bounds in the QROM. This framework exploits Zhandry’s compressed-oracle technique but abstracts away all the quantum aspects, so that the reasoning becomes purely classical. We give here an introduction to a simplified, and slightly adjusted version that does not consider parallel queries. We start with recalling (a particular view on) the compressed oracle. Along the way, we also give an improved version of Zhandry’s central lemma for the compressed oracle.

Before getting into this, we fix the following standard notation. For any positive integer $\ell > 0$, we set $[\ell] := \{1, 2, \dots, \ell\}$, and we let $2^{[\ell]}$ denote the power set of $[\ell]$, i.e., the set of all subsets of $[\ell]$. We write $\{0, 1\}^{\leq \ell}$ for the set of bit strings of size at most ℓ , including the empty string denoted $\emptyset$; similarly for $\{0, 1\}^{< \ell}$. Concatenation of two bit strings $v \in \{0, 1\}^m$ and $w \in \{0, 1\}^n$ is denoted by $v \| w \in \{0, 1\}^{m+n}$. For any finite non-empty set $\mathcal{Z}$, $\mathbb{C}[\mathcal{Z}]$ denotes the Hilbert space $\mathbb{C}^{|\mathcal{Z}|}$ together with a basis $\{|z\rangle\}$ labeled by the elements $z \in \mathcal{Z}$.

Finally, we consider a hash function $H : \mathcal{X} \rightarrow \mathcal{Y}$, to be modeled as a random oracle. For concreteness and simplicity, we assume that all relevant variables are encoded as bit strings, and that we can therefore choose $H : \{0, 1\}^{\leq B} \rightarrow \{0, 1\}^n$ for sufficiently large B and n .²

² B and n may depend on the security parameter $\lambda \in \mathbb{N}$. We will then assume that B and n can be computed from λ in polynomial time (in λ).

5.2.1 The Compressed Oracle—Seen as Quantum Lazy Sampling

The compressed oracle technique was formally introduced in Section 2.4. In the current chapter we use slightly different notation in order to better connect with [CFHL21], whose framework we build upon. For purposes of exposition it will also be useful to approach the technique explicitly from the perspective of (quantum) lazy sampling. In this section we give a self-contained introduction of the technique in the form that we will need it later on.

With the goal to analyze oracle algorithms that interact with a random oracle $H : \mathcal{X} \rightarrow \mathcal{Y}$, consider the set $\mathfrak{D}$ of all functions $D : \mathcal{X} \rightarrow \mathcal{Y} \cup \{\perp\}$, where $\perp$ is a special symbol. Such a function is referred to as a *database*. Later, we will fix $\mathcal{X} = \{0, 1\}^{\leq B}$ and $\mathcal{Y} = \{0, 1\}^n$. For $D \in \mathfrak{D}$, $x \in \mathcal{X}$ and $y \in \mathcal{Y} \cup \{\perp\}$, $D[x \mapsto y]$ denotes the database that maps x to y and otherwise coincides with D , i.e., $D[x \mapsto y](x) = y$ and $D[x \mapsto y](\bar{x}) = D(\bar{x})$ for all $\bar{x} \in \mathcal{X} \setminus \{x\}$.

Following the exposition of [CFHL21], the compressed-oracle technique is a quantum analogue of the classical lazy-sampling technique, commonly used to analyze algorithms in the classical ROM. In the classical lazy-sampling technique, the (simulated) random oracle starts off with the empty database, i.e., with $D_0 = \perp$, which maps any $x \in \mathcal{X}$ to $\perp$. Then, recursively, upon a query x , the current database D_i is updated to $D_{i+1} := D_i$ if $D_i(x) \neq \perp$, and to $D_{i+1} := D_i[x \mapsto y]$ for a randomly chosen $y \in \mathcal{Y}$ otherwise. This construction ensures that $|\{x \mid D_i(x) \neq \perp\}| \leq i$; after i queries thus, using standard sparse-encoding techniques, the database D_i can be efficiently represented and updated.

In the compressed-oracle quantum analogue of this lazy-sampling technique, the (simulated) random oracle also starts off with the empty database, but now considered as a quantum state $|\perp\rangle$ in the $|\mathfrak{D}|$ -dimensional state space $\mathbb{C}[\mathfrak{D}]$, and after i queries the state of the compressed oracle is then supported by databases $|D_i\rangle$ for which $|\{x \mid D_i(x) = \perp\}| \leq i$.³ Here, the update is given by a unitary operator $\mathbf{cO}$ acting on $\mathbb{C}[\mathcal{X}] \otimes \mathbb{C}[\mathcal{Y}] \otimes \mathbb{C}[\mathfrak{D}]$, i.e., on the query register, the response register, and the state of the compressed oracle. With respect to the computational basis $\{|x\rangle\}$ of $\mathbb{C}[\mathcal{X}]$ and the Fourier basis $\{|\hat{y}\rangle\}$ of $\mathbb{C}[\mathcal{Y}]$, $\mathbf{cO}$ is a *control* unitary, i.e., of the form $\mathbf{cO} = \sum_{x, \hat{y}} |x\rangle\langle x| \otimes |\hat{y}\rangle\langle \hat{y}| \otimes \mathbf{cO}_{x, \hat{y}}$, where $\mathbf{cO}_{x, \hat{y}}$ is a unitary on $\mathbb{C}[\mathcal{Y} \cup \{\perp\}]$, which in the above expression is understood to act on the register that carries the value of the database at the point x . More

³ This means that the density operator that describes the state of the compressed oracle has its support contained in the span of these $|D_i\rangle$.

formally, $\text{cO}_{x,\hat{y}}$ acts on register R_x when identifying $\mathbb{C}[\mathfrak{D}]$ with $\bigotimes_{x \in \mathcal{X}} \mathbb{C}[\mathcal{Y} \cup \{\perp\}]$ by means of the isomorphism $|D\rangle \mapsto \bigotimes_{x \in \mathcal{X}} |D(x)\rangle_{R_x}$. We refer to Lemma 4.3 in the full version of [CFHL21] for the full specification of $\text{cO}_{x,\hat{y}}$; it is not really relevant here.

The compressed oracle is tightly related to the *purified* oracle, which initiates its internal state with a uniform superposition $\sum_h |H\rangle \in \mathbb{C}[\mathfrak{D}]$ of all functions $H : \mathcal{X} \rightarrow \mathcal{Y}$, and then answers queries “in superposition”. Indeed, at any point in time during the interaction with an oracle quantum algorithm $\mathcal{A}$, the joint state of $\mathcal{A}$ and the compressed oracle coincides with the joint state of $\mathcal{A}$ and the purified oracle after “compressing” the latter.⁴ Formally, identifying $\mathbb{C}[\mathfrak{D}]$ with $\bigotimes_{x \in \mathcal{X}} \mathbb{C}[\mathcal{Y} \cup \{\perp\}]$ again, the compression of the state of the purified oracle works by applying the unitary **Comp** to each register R_x , where

$$\text{Comp} : |y\rangle \mapsto (|y\rangle + \frac{1}{\sqrt{|\mathcal{Y}|}}(|\perp\rangle - |\hat{0}\rangle))$$

for any $y \in \mathcal{Y}$, and $\text{Comp} : |\perp\rangle \mapsto |\hat{0}\rangle$. Here, $|\hat{0}\rangle$ is the $\hat{0}$ -vector from the Fourier basis $\{|\hat{y}\rangle\}$ of $\mathbb{C}[\mathcal{Y}]$.

Similarly to the classical case, by exploiting a quantum version of the sparse-encoding technique, both the internal state of the compressed oracle and the evolution cO can be efficiently computed. Furthermore, for any classical function $f : \mathfrak{D} \rightarrow \mathcal{T}$ that can be efficiently computed when given the sparse representation of $D \in \mathfrak{D}$, the corresponding quantum measurement given by the projections $P_t = \sum_{D: f(D)=t} |D\rangle\langle D|$ can be efficiently performed when given the sparse representation of the internal state of the compressed oracle. In particular, in Lemma 5.1 below, the condition $\mathbf{y} = D(\mathbf{x})$ for given $\mathbf{x}$ and $\mathbf{y}$ can be efficiently checked by a measurement. See Section 2.4 for more details on this technique.

In the classical lazy-sampling technique, if at the end of the execution of an oracle algorithm $\mathcal{A}$, having made q queries to the (lazy-sampled) RO, the database D_q is such that, say, $D_q(x) \neq 0$ for any $x \in \mathcal{X}$, then $\mathcal{A}$ ’s output is unlikely to be a 0-preimage, i.e., an x that is hashed to 0 upon one more query. $\mathcal{A}$ ’s best chance is to output an x that he has not queried yet, and thus $D_q(x) = \perp$, and then he has a $1/|\mathcal{Y}|$ -chance that $D_{q+1}(x) := D_q[x \mapsto y](x) = 0$, given that y is randomly chosen. Something similar holds in the quantum setting, with some adjustments. The general statement is given by the following result by Zhandry.

⁴ The terminology is somewhat misleading here; the actual compression takes place when invoking the sparse encoding (see below).

Lemma 5.1 (Lemma 5 in [Zha19a]). *Let $R \subseteq \mathcal{X}^\ell \times \mathcal{Y}^\ell \times \mathcal{Z}$ be a relation, and let $\mathcal{A}$ be an oracle quantum algorithm that outputs $\mathbf{x} \in \mathcal{X}^\ell$, $\mathbf{y} \in \mathcal{Y}^\ell$ and $z \in \mathcal{Z}$. Furthermore, let*

$$p = p(\mathcal{A}) := \Pr[\mathbf{y} = H(\mathbf{x}) \wedge (\mathbf{x}, \mathbf{y}, z) \in R]$$

be the considered probability when $\mathcal{A}$ has interacted with the standard RO, initialized with a uniformly random function H , and let

$$p' = p'(\mathcal{A}) := \Pr[\mathbf{y} = D(\mathbf{x}) \wedge (\mathbf{x}, \mathbf{y}, z) \in R]$$

be the considered probability when $\mathcal{A}$ has interacted with the compressed oracle instead and D is obtained by measuring its internal state (in the basis $\{|D\rangle\}_{D \in \mathfrak{D}}$). Then

$$\sqrt{p} \leq \sqrt{p'} + \sqrt{\frac{\ell}{|\mathcal{Y}|}}.$$

Remark 5.2. This bound is particular useful in case $\mathcal{Z} = \emptyset$ (or R does not depend on its third input z), since then p' is bounded by $\Pr[\exists \tilde{\mathbf{x}} : (\tilde{\mathbf{x}}, D(\tilde{\mathbf{x}})) \in R]$ and the latter is determined solely by the evolution of the compressed oracle (when interacting with $\mathcal{A}$) and does not depend on the actual output of $\mathcal{A}$.

In Section 5.2.3, Corollary 5.8, we will give an alternative such relation between the success probability of an algorithm interacting with the actual RO, and probabilities obtained by inspecting the compressed oracle instead. Strictly speaking, the results of Lemma 5.1 and Corollary 5.8 are incomparable, but in typical applications the latter gives a significantly better bound.

5.2.2 The Quantum Transition Capacity and Its Relevance

The above discussion shows that, in order to bound the success probability p of an oracle algorithm $\mathcal{A}$, it is sufficient to bound the probability of the database D , obtained by measuring the internal state of the compressed oracle after the interaction with $\mathcal{A}$, satisfying a certain property (e.g., the property of there existing an x such that $D(x) = 0$).

To facilitate that latter, Chung et al. [CFHL21] introduced a framework that, in certain cases, allows to bound this alternative figure of merit by means of purely classical reasoning. We briefly recall here some of the core elements of this framework, which are relevant to us. Note that [CFHL21] considers the parallel-query model, where in each of the q (sequential) interactions with the

RO, an oracle algorithm $\mathcal{A}$ can make k queries simultaneously in parallel with each interaction. Here, we consider the (more) standard model of one query per interaction, i.e., setting $k = 1$. On the other hand, we state and prove a slight generalization of Theorem 5.16 in [CFHL21] (when restricted to $k = 1$).

A subset $P \subseteq \mathfrak{D}$ is called a *database property*. We say that $D \in \mathfrak{D}$ *satisfies* P if $D \in P$, and the complement of P is denoted $\neg P = \mathfrak{D} \setminus P$. For such a database property P , [CFHL21] defines $\llbracket \perp \xRightarrow{q} P \rrbracket$ as the square-root of the maximal probability of D satisfying P when D is obtained by measuring the internal state of the compressed oracle after the interaction with $\mathcal{A}$, maximized over all oracle quantum algorithms $\mathcal{A}$ with query complexity q , i.e., in short

$$\llbracket \perp \xRightarrow{q} P \rrbracket := \max_{\mathcal{A}} \sqrt{\Pr[D \in P]} . \quad (43)$$

In the context of Lemma 5.1 for the case $\mathcal{Z} = \emptyset$ (see Remark 5.2), we can define the database property $P^R := \{D \in \mathfrak{D} \mid \exists \mathbf{x} \in \mathcal{X}^\ell : (\mathbf{x}, D(\mathbf{x})) \in R\}$ induced by R , and thus bound

$$p'(\mathcal{A}) \leq \Pr[(\mathbf{x}, D(\mathbf{x})) \in R] \leq \Pr[D \in P^R] \leq \llbracket \perp \xRightarrow{q} P^R \rrbracket^2 \quad (44)$$

for any oracle quantum algorithm $\mathcal{A}$ with query complexity q .

Furthermore, Lemma 5.6 in [CFHL21] shows that for any target database property P and for any sequence $P_0, P_1, \dots, P_q$ with $\neg P_0 = \{\perp\}$ and $P_q = P$,

$$\llbracket \perp \xRightarrow{q} P \rrbracket \leq \sum_{s=1}^q \llbracket \neg P_{s-1} \rightarrow P_s \rrbracket , \quad (45)$$

where, for any database properties P and P' , the definition of the *quantum transition capacity* $\llbracket P \rightarrow P' \rrbracket$ is recalled in Definition 5.3.

The nice aspect of the framework of [CFHL21] is that it provides means to manipulate and bound quantum transition capacities using purely classical reasoning, i.e., without the need to understand and work with the definition. Indeed, for instance Theorem 5.4 below, which is a variant of Theorem 5.17 in (the full version of) [CFHL21], shows how to bound $\llbracket P \rightarrow P' \rrbracket$ by means of a certain classical probability; furthermore, to facilitate the application of such theorems, [CFHL21] showed that the quantum transition capacity satisfies several natural manipulation rules, like $\llbracket P \rightarrow P' \rrbracket = \llbracket P' \rightarrow P \rrbracket$ (i.e., it is symmetric), and

$$\begin{aligned} \llbracket P \cap Q \rightarrow P' \rrbracket &\leq \min\{\llbracket P \rightarrow P' \rrbracket, \llbracket Q \rightarrow P' \rrbracket\} \quad \text{and} \\ \min\{\llbracket P \rightarrow P' \rrbracket, \llbracket P \rightarrow Q' \rrbracket\} &\leq \llbracket P \rightarrow P' \cup Q' \rrbracket \leq \llbracket P \rightarrow P' \rrbracket + \llbracket P \rightarrow Q' \rrbracket , \end{aligned} \quad (46)$$

which allow to decompose complicated capacities into simpler ones. Therefore, by means of the above series of inequalities with p from Lemma 5.1 on the left hand side, it is possible (in certain cases) to bound the success probability of any oracle quantum algorithm $\mathcal{A}$ in the QROM by means of the following recipe: (1) Choose suitable transitions $P_{s-1} \rightarrow P_s$, (2) decompose the capacities $\llbracket \neg P_{s-1} \rightarrow P_s \rrbracket$ into simpler ones using manipulation rules as above, and (3) bound the simplified capacities by certain classical probabilities, exploiting results like Theorem 5.4. We will closely follow this recipe.

In order to state and later use Theorem 5.4, we need to introduce the following additional concepts. As explained above, there is no need to actually spell out the definition of the quantum transition capacity in order to use Theorem 5.4; for completeness, and since it is needed for the proof of Theorem 5.4, we do provide it below.

For any database $D \in \mathfrak{D}$ and any $x \in \mathcal{X}$,

$$D|_x := \{D[x \mapsto y] \mid y \in \mathcal{Y} \cup \{\perp\}\}$$

denotes the set of all databases that coincide with D outside of x . Furthermore, for a database property P ,

$$P|_{D|x} := \{y \in \mathcal{Y} \cup \{\perp\} \mid D[x \mapsto y] \in P\} \subseteq \mathcal{Y} \cup \{\perp\}$$

denotes the set of values y for which $D[x \mapsto y]$ satisfies P . Following the convention used in [CFHL21], we identify the subset $P|_{D|x} \subseteq \mathcal{Y} \cup \{\perp\}$ with the projector $P|_{D|x} = \sum_y |y\rangle\langle y|$ acting on $\mathbb{C}[\mathcal{Y} \cup \{\perp\}]$, where the sum is over all $y \in P|_{D|x}$.

Definition 5.3 (Definition 5.5 of [CFHL21], case $k = 1$). *Let P, P' be two database properties. Then, the quantum transition capacity (of order 1) is defined as*

$$\llbracket P \rightarrow P' \rrbracket := \max_{\mathbf{x}, \hat{\mathbf{y}}, D} \|P'|_{D|\mathbf{x}} cO_{\mathbf{x}, \hat{\mathbf{y}}} P|_{D|\mathbf{x}}\|$$

where the max is over all $\mathbf{x} \in \mathcal{X}^k$, $\hat{\mathbf{y}} \in \hat{\mathcal{Y}}^k$, and $D \in \mathfrak{D}$.

The following is a variation of Theorem 5.17 in (the full version of) [CFHL21], obtained by restricting k to 1. On the other hand, we exploit and include some symmetry that is not explicit in the original statement. The proof is a small adjustment to the original proof.

Theorem 5.4. *Let P and P' be database properties with trivial intersection, i.e., $P \cap P' = \emptyset$, and for every $D \in \mathfrak{D}$ and $x \in \mathcal{X}$ let*

$$\mathbb{L}^{x,D} := \begin{cases} P|_{D|x} & \text{if } \perp \in P'|_{D|x} \\ P'|_{D|x} & \text{if } \perp \in P|_{D|x}, \end{cases}$$

with $\mathbb{L}^{x,D}$ being either of the two if $\perp \notin P|_{D|x} \cup P'|_{D|x}$.⁵ Then

$$\llbracket P \rightarrow P' \rrbracket \leq \max_{x,D} \sqrt{10P[U \in \mathbb{L}^{x,D}]},$$

where U is uniform over $\mathcal{Y}$, and the maximization can be restricted to $D \in \mathfrak{D}$ and $x \in \mathcal{X}$ for which both $P|_{D|x}$ and $P'|_{D|x}$ are non-empty.

Remark 5.5. Both, $P|_{D|x}$ and $P'|_{D|x}$, and thus also $\mathbb{L}^{x,D}$, do not depend on the value of $D(x)$, only on the values of D outside of x .

Proof. For any $D \in \mathfrak{D}$ and $x \in \mathcal{X}$, we observe that $\mathbf{c}O_{x,-\hat{y}} = (\mathbf{c}O_{x,\hat{y}})^\dagger$ and hence

$$\|P'|_{D|x} \mathbf{c}O_{x,\hat{y}} P|_{D|x}\| = \|(P|_{D|x})^\dagger \mathbf{c}O_{x,-\hat{y}} (P'|_{D|x})^\dagger\| = \|P|_{D|x} \mathbf{c}O_{x,-\hat{y}} P'|_{D|x}\|,$$

and so it is sufficient to argue for the case when $\mathbb{L}^{x,D}$ is set to $P'|_{D|x}$. By the disjointness requirement, as subsets of $\mathcal{Y} \cup \{\perp\}$, the complement of $\mathbb{L}^{x,D} = P'|_{D|x}$ is a superset of $P|_{D|x}$. Thus, as projections acting on $\mathbb{C}[\mathcal{Y} \cup \{\perp\}]$, $P|_{D|x} \leq \mathbb{I} - \mathbb{L}^{x,D}$. Therefore, the above norm is upper bounded by $\|\mathbb{L}^{x,D} \mathbf{c}O_{x,\hat{y}} (\mathbb{I} - \mathbb{L}^{x,D})\|$. Given that $\perp \notin \mathbb{L}^{x,D}$, the square norm $\|\mathbb{L}^{x,D} \mathbf{c}O_{x,\hat{y}} (\mathbb{I} - \mathbb{L}^{x,D})\|^2$ can be upper bounded exactly as in the proof of Theorem 5.17 in [CFHL21] by $10P[U \in \mathbb{L}^{x,D}]$, giving the claimed bound. $\square$

5.2.3 An Improved Variant of Zhandry's Lemma

We show here an alternative to Zhandry's lemma (Lemma 5.1), which offers a better bound in typical applications. To start with, note that Lemma 5.1 considers an algorithm $\mathcal{A}$ that not only outputs $\mathbf{x} = (x_1, \dots, x_\ell)$ but also $\mathbf{y} = (y_1, \dots, y_\ell)$, where the latter is supposed to be the point-wise hash of $\mathbf{x}$; indeed, this is what is being checked in the definition of the probability p , along with $(\mathbf{x}, \mathbf{y}, z) \in R$. This requirement is somewhat unnatural, in that an algorithm $\mathcal{A}$ for, say, finding a collision, i.e., $x_1 \neq x_2$ with $H(x_1) = H(x_2)$, does *not*

⁵ By the disjointness requirement, $\perp$ cannot be contained in both.

necessarily output the (supposed to be equal) hashes $y_1 = H(x_1)$ and $y_2 = H(x_2)$. Of course, this is no problem since one can easily transform such an algorithm $\mathcal{A}$ that does not output the hashes into one that does, simply by making a few more (classical) queries to the random oracle at the end of the execution, and then one can apply Lemma 5.1 to this tweaked algorithm $\tilde{\mathcal{A}}$.

We show below that if we anyway consider this tweaked algorithm $\tilde{\mathcal{A}}$, which is *promised* to query the random oracle to obtain and then output the hashes of $\mathbf{x} = (x_1, \dots, x_\ell)$, then we can actually improve the bound and avoid the square-roots in Lemma 5.1. On top, the proof is much simpler than Zhandry's proof for his lemma.

At the core is the following lemma; Corollary 5.8 below then puts it in a form that is comparable to Lemma 5.1 and shows the improvement.

Lemma 5.6. *Let $\mathcal{A}$ be an oracle quantum algorithm that outputs $\mathbf{x} = (x_1, \dots, x_\ell) \in \mathcal{X}^\ell$ and $z \in \mathcal{Z}$. Let $\tilde{\mathcal{A}}$ be the oracle quantum algorithm that runs $\mathcal{A}$, makes ℓ classical queries on the outputs x_i to obtain $\mathbf{y} = H(\mathbf{x})$, and then outputs $(\mathbf{x}, \mathbf{y}, z)$. When $\tilde{\mathcal{A}}$ interacts with the compressed oracle instead, and at the end D is obtained by measuring the internal state of the compressed oracle, then, conditioned on $\tilde{\mathcal{A}}$'s output $(\mathbf{x}, \mathbf{y}, z)$,*

$$\Pr[\mathbf{y} = D(\mathbf{x}) | (\mathbf{x}, \mathbf{y}, z)] \geq 1 - \frac{2\ell}{|\mathcal{Y}|}.$$

Proof. Consider first $\tilde{\mathcal{A}}$ interacting with the *purified* (yet uncompressed) oracle. Conditioned on $\tilde{\mathcal{A}}$'s output $(\mathbf{x}, \mathbf{y}, z)$, the state of the oracle is then supported by $|H\rangle$ with $H(x_i) = y_i$ for all $i \in \{1, \dots, \ell\}$, i.e., the registers labeled by $x_1, \dots, x_\ell$ are in state $|y_1\rangle \cdots |y_\ell\rangle$. Given that the compressed oracle is obtained by applying **Comp** to all the registers, we thus have that

$$\begin{aligned} \Pr[y_i = y'_i | (\mathbf{x}, \mathbf{y}, z)] &= |\langle y_i | \text{Comp} | y_i \rangle|^2 = \left| \langle y_i | \left(|y_i\rangle + \frac{1}{\sqrt{|\mathcal{Y}|}} (|\perp\rangle - |\hat{0}\rangle) \right) \right|^2 \\ &= \left| 1 - \frac{1}{\sqrt{|\mathcal{Y}|}} \langle y_i | \hat{0} \rangle \right|^2 = \left| 1 - \frac{1}{|\mathcal{Y}|} \right|^2 \geq 1 - \frac{2}{|\mathcal{Y}|}. \end{aligned}$$

Applying union bound concludes the claim. $\square$

The following generalization of Lemma 5.6 follows immediately by enhancing $\mathcal{A}$ so that it computes and outputs all the values x that need to be queried in order to compute $\mathcal{F}^H(z)$, and then apply Lemma 5.6 above.

Corollary 5.7. *Let $\mathcal{A}$ be an oracle quantum algorithm that produces an arbitrary output $z \in \mathcal{Z}$, and let $\mathcal{F}$ be an arbitrary classical ℓ -query oracle algorithm. Let $\tilde{\mathcal{A}} := \mathcal{F} \circ \mathcal{A}$ be the oracle quantum algorithm that first runs $\mathcal{A}$ to obtain z , then $\mathcal{F}$ to obtain $y := \mathcal{F}^H(z)$, and finally outputs (y, z) . When $\tilde{\mathcal{A}}$ interacts with the compressed oracle instead, and at the end D is obtained by measuring the internal state of the compressed oracle, then, conditioned on $\tilde{\mathcal{A}}$'s output (y, z) ,*

$$\Pr[y = \mathcal{F}^D(z) | (y, z)] \geq 1 - \frac{2\ell}{|\mathcal{Y}|}.$$

The following corollary of Lemma 5.6 is put in a form that can be nicely compared with Lemma 5.1, understanding that typically Lemma 5.1 is applied to $\tilde{\mathcal{A}}$.

Corollary 5.8. *Let $R \subseteq \mathcal{X}^\ell \times \mathcal{Y}^\ell \times \mathcal{Z}$ be a relation. Let $\mathcal{A}$ be an oracle quantum algorithm that outputs $\mathbf{x} \in \mathcal{X}^\ell$ and $z \in \mathcal{Z}$, and let $\tilde{\mathcal{A}}$ be as in Lemma 5.6. Let*

$$p_\circ(\mathcal{A}) := \Pr[(\mathbf{x}, H(\mathbf{x}), z) \in R]$$

be the considered probability when $\mathcal{A}$ has interacted with the RO. Furthermore, let $p(\tilde{\mathcal{A}})$ and $p'(\tilde{\mathcal{A}})$ be defined as in Lemma 5.1 (but now for $\tilde{\mathcal{A}}$). Then

$$p_\circ(\mathcal{A}) = p(\tilde{\mathcal{A}}) \leq p'(\tilde{\mathcal{A}}) + \frac{2\ell}{|\mathcal{Y}|}.$$

For convenience, we recall that

$$p'(\tilde{\mathcal{A}}) = \Pr[\mathbf{y} = D(\mathbf{x}) \wedge (\mathbf{x}, \mathbf{y}, z) \in R] \leq \Pr[(\mathbf{x}, D(\mathbf{x}), z) \in R].$$

Proof. The equality holds by construction of $\tilde{\mathcal{A}}$. For the first inequality, we observe that

$$\begin{aligned} p'(\tilde{\mathcal{A}}) &= \Pr[\mathbf{y} = D(\mathbf{x}) | (\mathbf{x}, \mathbf{y}, z) \in R] \Pr[(\mathbf{x}, \mathbf{y}, z) \in R] \\ &\geq \left(1 - \frac{2\ell}{|\mathcal{Y}|}\right) \Pr[(\mathbf{x}, \mathbf{y}, z) \in R] \geq \left(1 - \frac{2\ell}{|\mathcal{Y}|}\right) p(\tilde{\mathcal{A}}) \geq p(\tilde{\mathcal{A}}) - \frac{2\ell}{|\mathcal{Y}|}, \end{aligned}$$

where the first inequality is by Lemma 5.6. The second and last inequality in the statement holds trivially by definition of p' . $\square$

Section 5.3

Some Background on (Non-)Interactive Proofs

Let $\{\mathcal{I}_\lambda\}_{\lambda \in \mathbb{N}}$ and $\{\mathcal{W}_\lambda\}_{\lambda \in \mathbb{N}}$ be two families of sets, with the members being labeled by the security parameter $\lambda \in \mathbb{N}$. Let $R_\lambda \subseteq \mathcal{I}_\lambda \times \mathcal{W}_\lambda$ be a relation that is polynomial-time computable in λ . $w \in \mathcal{W}_\lambda$ is called a *witness* for $inst \in \mathcal{I}_\lambda$ if $R_\lambda(inst, w)$, and $L_\lambda := \{inst \in \mathcal{I}_\lambda \mid \exists w \in \mathcal{W}_\lambda : R_\lambda(inst, w)\}$.

Below, we recall some concepts in the context of interactive and non-interactive proofs for such families $\{R_\lambda\}_{\lambda \in \mathbb{N}}$ of relations. We start by discussing the aspired security definition for non-interactive proofs.

5.3.1 Non-interactive Proofs and Online Extractability

An *non-interactive proof in the random-oracle model* for a family $\{R_\lambda\}_{\lambda \in \mathbb{N}}$ of relations consists of a pair $(\mathcal{P}, \mathcal{V})$ of oracle algorithms, referred to as *prover* and *verifier*, both making queries to the random oracle $H : \mathcal{X} \rightarrow \mathcal{Y}$. The prover $\mathcal{P}$ takes as input $\lambda \in \mathbb{N}$ and an instance $inst \in L_\lambda$ and outputs a *proof* $\pi \in \Pi_\lambda$, and $\mathcal{V}$ takes as input $\lambda \in \mathbb{N}$ and a pair $(inst, \pi) \in \mathcal{I}_\lambda \times \Pi_\lambda$ and outputs a Boolean value, 0 or 1, or **accept** or **reject**. The verifier $\mathcal{V}$ is required to run in time polynomial in λ , while, *per-se*, $\mathcal{P}$ may have unbounded running time.⁶

By default, we require correctness and soundness, i.e., that for any $\lambda \in \mathbb{N}$ and any $inst \in L_\lambda$

$$\Pr[\mathcal{V}^H(\lambda, inst, \pi) : \pi \leftarrow \mathcal{P}^H(\lambda, inst)] \geq 1 - \varepsilon_{\text{cor}}(\lambda),$$

while for any $\lambda \in \mathbb{N}$ and any oracle quantum algorithm $\mathcal{P}^*$ (a *dishonest prover*) with query complexity q

$$\Pr[inst \notin L_\lambda \wedge \mathcal{V}^H(\lambda, inst, \pi) : (inst, \pi) \leftarrow \mathcal{P}^{*H}(\lambda)] \leq \varepsilon_{\text{snd}}(\lambda, q, n)$$

for certain ε_{cor} and ε_{snd} , respectively referred to as *correctness error* and *soundness error*. The fact that the instance $inst$, for which $\mathcal{P}^*$ tries to forge a proof, is not given as input to $\mathcal{P}^*$ but is instead chosen by $\mathcal{P}^*$ is referred to as $\mathcal{P}^*$ being *adaptive*.

⁶ Alternatively, one may consider a witness w for $inst$ to be given as additional input to $\mathcal{P}$, and then ask $\mathcal{P}$ to be polynomial-time as well.

We now move towards defining *online extractability* (for adaptive $\mathcal{P}^*$). For that purpose, let $\mathcal{P}^*$ be a dishonest prover as above, except that it potentially outputs some additional auxiliary (possibly quantum) output Z next to $(inst, \pi)$. We then consider an interactive algorithm $\mathcal{E}$, called *online extractor*, which takes $\lambda \in \mathbb{N}$ as input and simulates the answers to the oracle queries in the execution of $\mathcal{V}^H \circ \mathcal{P}^{*H}(\lambda)$, which we define to run $(inst, \pi, Z) \leftarrow \mathcal{P}^{*H}(\lambda)$ followed by $v \leftarrow \mathcal{V}^H(\lambda, inst, \pi)$; furthermore, at the end, $\mathcal{E}$ outputs $w \in \mathcal{W}_\lambda$. We denote the execution of $\mathcal{V}^H \circ \mathcal{P}^{*H}(\lambda)$ with the calls to H simulated by $\mathcal{E}$, and considering $\mathcal{E}$'s final output w as well, as $(inst, \pi, Z; v; w) \leftarrow \mathcal{V}^\mathcal{E} \circ \mathcal{P}^{*\mathcal{E}}(\lambda)$.

Definition 5.9. *A non-interactive proof in the (quantum-accessible) random-oracle model (QROM) for $\{R_\lambda\}_{\lambda \in \mathbb{N}}$ is a proof of knowledge with online extractability (PoK-OE) against adaptive adversaries if there exists an online extractor $\mathcal{E}$, and functions ε_{sim} (the simulation error) and ε_{ex} (the extraction error), with the following properties. For any $\lambda \in \mathbb{N}$ and for any dishonest prover $\mathcal{P}^*$ with query complexity q ,*

$$\delta([(inst, \pi, Z, v)]_{\mathcal{V}^H \circ \mathcal{P}^{*H}(\lambda)}, [(inst, \pi, Z, v)]_{\mathcal{V}^\mathcal{E} \circ \mathcal{P}^{*\mathcal{E}}(\lambda)}) \leq \varepsilon_{\text{sim}}(\lambda, q, n)$$

and

$$\Pr[v = \text{accept} \wedge (inst, w) \notin R : (inst, \pi, Z; v; w) \leftarrow \mathcal{V}^\mathcal{E} \circ \mathcal{P}^{*\mathcal{E}}(\lambda)] \leq \varepsilon_{\text{ex}}(\lambda, q, n).$$

Furthermore, the runtime of $\mathcal{E}$ is polynomial in $\lambda + q + n$, and $\varepsilon_{\text{sim}}(\lambda, q, n)$ and $\varepsilon_{\text{ex}}(\lambda, q, n)$ are negligible in λ whenever q and n are polynomial in λ .

Remark 5.10. In the classical definition of a proof of knowledge, the extractor $\mathcal{E}$ interacts with $\mathcal{P}^*$ only, and the verifier $\mathcal{V}$ is not explicitly involved, but would typically be run by $\mathcal{E}$. Here, in the context of online extractability, it is necessary to explicitly go through the verification procedure, which also makes oracle queries, to determine whether a proof is valid, i.e., for the event $v = \text{accept}$ to be well defined.

5.3.2 $\mathcal{G}$ -soundness of C&O Σ -Protocols

C&O protocols are a subclass of Σ -protocols for which the first message a consists of (hash based) commitments $y_1, \dots, y_\ell$ for messages $m_1, \dots, m_\ell \in \mathcal{M}$, and possibly an additional string a_o . The challenge c is chosen uniformly at random from a subset $\mathcal{C} \subseteq 2^{[\ell]}$ of indices, which point to the messages the prover opens in its response $z = \mathbf{m}_c = (m_i)_{i \in c}$. In this chapter we consider the

hash based version where $y_i = H(m_i)$, and we model the hash function H as a random oracle. See Section 2.2.1 for a complete formal specification of C&O protocols.

We briefly recall the notions of $\mathfrak{S}$ -soundness and $\mathfrak{S}$ -soundness* as developed in Section 4.5.2, which offer a convenient general notion of special soundness, or more generally k -soundness for C&O protocols.

Here and below, given a C&O protocol Π with challenge space $\mathcal{C} \subseteq 2^{[\ell]}$, we let $\mathfrak{S} \subseteq 2^{\mathcal{C}}$ be an arbitrary non-empty, monotone increasing set of subsets $S \subseteq \mathcal{C}$, where the monotonicity means that $S \in \mathfrak{S} \wedge S \subseteq S' \Rightarrow S' \in \mathfrak{S}$. We then also set $\mathfrak{S}_{\min} := \{S \in \mathfrak{S} \mid S_{\circ} \subsetneq S \Rightarrow S_{\circ} \notin \mathfrak{S}\}$ to be the minimal sets in $\mathfrak{S}$.

For simplicity, the reader can consider $\mathfrak{S} = \mathfrak{T}_k := \{S \subseteq \mathcal{C} \mid |S| \geq k\}$ for some threshold k , and thus $\mathfrak{S}_{\min} = \{S \subseteq \mathcal{C} \mid |S| = k\}$. This then corresponds to the notion of k -soundness for C&O protocols, which in turn means that the witness can be computed from valid responses to k (or more) distinct challenges for a given first message $y_1, \dots, y_{\ell}$, assuming the messages $m_1, \dots, m_{\ell}$ to be uniquely determined by their commitments.

Definition 5.11. *A C&O protocol Π is $\mathfrak{S}$ -sound if there exists an efficient deterministic algorithm $\mathcal{E}_{\mathfrak{S}}(\text{inst}, m_1, \dots, m_{\ell}, a_{\circ}, S)$ that takes as input an instance $\text{inst} \in \mathcal{I}$, messages $m_1, \dots, m_{\ell} \in \mathcal{M} \cup \{\perp\}$, a string $a_{\circ}$, and a set $S \in \mathfrak{S}_{\min}$, and outputs a witness for inst if $V(\text{inst}, c, \mathbf{m}_c, a_{\circ})$ for all $c \in S$.⁷*

We note the clash in terminology with Definition 4.22. However, the current definition applies exclusively to C&O Σ -protocols in the (Q)ROM, whereas Definition 4.22 applies exclusively to Σ -protocols in the standard model; so there should be no confusion. The two definitions are of course related: a $\mathfrak{S}$ -sound C&O Σ -protocol becomes a $\mathfrak{S}$ -sound plain Σ -protocol when the commitments are instantiated with a perfectly binding commitment scheme (rather than with a hash function).

A slightly stronger condition than $\mathfrak{S}$ -soundness is the following variant, which differs in that the extractor needs to work as soon as there *exists* a set S as specified, without the extractor being given S as input. We refer to Section 4.5.2 for a more detailed discussion of this aspect. As explained there, whether S is given or not often makes no (big) difference.

For instance, when $\mathfrak{S}_{\min}$ consists of a polynomial number of sets S then the extractor can do a brute-force search to find S , and so $\mathfrak{S}$ -soundness* is

⁷ The restriction for S to be in $\mathfrak{S}_{\min}$, rather than in $\mathfrak{S}$, is to avoid an exponentially sized input while asking $\mathcal{E}_{\mathfrak{S}}$ to be efficient.

then implied by $\mathfrak{S}$ -soundness. Also, the r -fold parallel repetition of a $\mathfrak{S}$ -sound protocol, which by default is a $\mathfrak{S}^{\vee r}$ -sound protocol (see Section 4.5.2), is automatically $\mathfrak{S}^{\vee}$ -sound* if $\mathfrak{S}_{\min}$ is polynomial in size: the extractor can then do a brute-force search in every repeated instance.

Definition 5.12. *A C&O protocol Π is $\mathfrak{S}$ -sound* if there exists an efficient deterministic algorithm $\mathcal{E}_{\mathfrak{S}}^*(inst, m_1, \dots, m_\ell, a_o)$ that takes as input an instance $inst \in \mathcal{I}$ and strings $m_1, \dots, m_\ell \in \mathcal{M} \cup \{\perp\}$ and a_o , and it outputs a witness for $inst$ if there exists $S \in \mathfrak{S}$ such that $V(inst, c, \mathbf{m}_c, a_o)$ for all $c \in S$.*

As for plain Σ -protocols, we define

$$p_{triv}^{\mathfrak{S}} := \frac{1}{|\mathcal{C}|} \max_{\hat{S} \notin \mathfrak{S}} |\hat{S}|, \quad (47)$$

capturing the “trivial” attack of picking a set $\hat{S} = \{\hat{c}_1, \dots, \hat{c}_m\} \notin \mathfrak{S}$ of challenges $\hat{c}_i \in \mathcal{C}$ and then prepare $\hat{\mathbf{m}} = (\hat{m}_1, \dots, \hat{m}_\ell)$ and a_o in such a way that $V(inst, c, \hat{\mathbf{m}}_c, a_o)$ holds if $c \in \hat{S}$. After committing to $\hat{m}_1, \dots, \hat{m}_\ell$, the prover can successfully answer to challenges $c \in \hat{S}$.

5.3.3 The Fiat-Shamir Transformation of (C&O) Σ -Protocols

The Fiat-Shamir (FS) transformation [FS87] turns arbitrary Σ -protocols into non-interactive proofs in the random oracle model by setting the challenge $c \in \mathcal{C}$ to be the hash of the instance and the first message a . For this transformation to work smoothly, it is typically assumed that $|\mathcal{C}|$ is a power of 2 and its elements are represented as bit strings of size $\log |\mathcal{C}|$, so that one can indeed set c to be (the first $\log |\mathcal{C}|$ bits of) the hash $H(inst, a)$. The assumption on $|\mathcal{C}|$ is essentially without loss of generality (WLOG), since one can always reduce the size of $|\mathcal{C}|$ to the next lower power of 2, at the cost of losing at most 1 bit of security. However, for a C&O Σ -protocol, where a challenge space $\mathcal{C}$ is a (typically strict) subset of $2^{[\ell]}$, there is not necessarily a natural way to represent $c \in \mathcal{C}$ as a bitstring of size $\log |\mathcal{C}|$. Therefore, we will make it explicit that the challenge-set $c \in \mathcal{C} \subset 2^{[\ell]}$ is computed from the “raw randomness” $H(inst, y_1, \dots, y_\ell, a_o)$ in a deterministic way as $c = \gamma \circ H(inst, y_1, \dots, y_\ell, a_o)$ for an appropriate function $\gamma : \mathcal{Y} \rightarrow \mathcal{C}$, mapping a uniformly random hash in $\mathcal{Y}$ to a random challenge-set in $\mathcal{C}$. Obviously, for $H(inst, y_1, \dots, y_\ell, a_o)$ to be defined, in addition to $\mathcal{M} \subseteq \mathcal{X}$ we also need that $\mathcal{I} \times \mathcal{Y}^\ell \subseteq \mathcal{X}$, which again just means that B needs to be large enough. We write $\text{FS}[\Pi]$ for the Fiat-Shamir transformation of a (C&O) Σ -protocol Π .

Remark 5.13. Additionally, we need that n is sufficiently large, so that there is a sufficient amount of randomness in the hash value $H(\text{inst}, y_1, \dots, y_\ell)$ in order to be mapped to a random $c \in \mathcal{C}$. The canonical choice for γ is then the function that the *interactive* verifier applies to his local randomness to compute the random challenge $c \in \mathcal{C}$. To simplify the exposition, we assume that n is indeed sufficiently large. Otherwise, one can simply set $\mathcal{Y} := \{0, 1\}^{n'}$ instead, for sufficiently large n' , and then let y_i be $H(m_i)$ truncated to the original number n of bits again. This truncation has no effect on our results.

Remark 5.14. We assume WLOG that the two kinds of inputs to H , i.e., m_i and $(\text{inst}, y_1, \dots, y_\ell, a_o)$, are differently formatted, e.g., bit strings of different respective sizes or prefixes (this is referred to as *domain separation*). In other words, we assume that $\mathcal{M}$ and $\mathcal{I} \times \mathcal{Y}^\ell$ are disjoint.

Remark 5.15. When considering the adaptive security of a Fiat-Shamir transformation $\text{FS}[\Pi]$ of a C&O protocol Π for a relation R , the additional string a_o , which may be part of the first message a of the original protocol Π , may WLOG be considered to be part of the instance *inst* instead.

Indeed, any dishonest prover $\mathcal{P}^*$ against $\text{FS}[\Pi]$, which (by Definition 5.9) outputs an instance *inst* and a proof $\pi = (a_o, y_1, \dots, y_\ell)$, can alternatively be parsed as a dishonest prover that outputs an instance $\text{inst}' = (\text{inst}, a_o)$ and a proof $\pi' = (y_1, \dots, y_\ell)$. Thus, $\mathcal{P}^*$ can be parsed as a dishonest prover against $\text{FS}[\Pi']$, where the C&O protocol Π' works as Π , except that a_o is considered as part of the instance, rather than as part of the first message, and thus Π' is a C&O protocol for the relation $((\text{inst}, a_o), w) \in R' :\Leftrightarrow (\text{inst}, w) \in R$.⁸ Therefore, security (in the sense of Definition 5.9) for $\text{FS}[\Pi']$ implies that of $\text{FS}[\Pi]$.

Section 5.4

Online Extractability of the FS-Transformation: The Case of Ordinary C&O Protocols

We now consider the Fiat-Shamir transformation $\text{FS}[\Pi]$ of an ordinary C&O protocol Π . Our goal is to show that $\text{FS}[\Pi]$ admits online extraction. We note

⁸ We do not specify the local computation of the honest prover $\mathcal{P}'$ in $\Pi' = (\mathcal{P}', \mathcal{V}')$, i.e., how to act when a_o is part of the input, and in general it might not be efficient, but this is fine since we are interested in the security against dishonest provers.

5.4. Online Extractability of the FS-Transformation: The Case of Ordinary C&O Protocols

that by exploiting Remark 5.15, we may assume WLOG that the first message of Π consists of the commitments $y_1, \dots, y_\ell$ only, and no additional string a_o . In Section 5.5, we then consider the case of Merkle-tree-based C&O protocols.

Our analysis of the online extractability of $\text{FS}[\Pi]$ uses the framework of Chung et al. [CFHL21], discussed and outlined in Section 5.2. Thus, at the core of our analysis is a bound on a certain quantum transition capacity. This is treated in the upcoming subsection.

5.4.1 Technical Preface

We first introduce a couple of elementary database properties (related to CoLlisions and the SiZe of the database) that will be useful for us:

$$\text{CL} := \{D \mid \exists x \neq x' : D(x) = D(x') \neq \perp\} \quad \text{and} \quad \text{SZ}_{\leq s} := \{D \mid \#\{z \mid D(z) \neq \perp\} \leq s\}.$$

Next, for an instance $\text{inst} \in \mathcal{I}$, we want to specify the database property that captures a cheating prover that succeeds in producing an accepting proof while fooling the extractor. For the purpose of specifying this database property, we introduce the following notation. For a given database $D \in \mathfrak{D}$ and for a commitment $y \in \mathcal{Y}$, we define $D^{-1}(y)$ to be the smallest $x \in \mathcal{X}$ with $D(x) = y$, with the convention that $D^{-1}(y) := \perp$ if there is no such x , as well as $D^{-1}(\perp) := \perp$. We note that by removing collisions, we ensure that there is at most one such x ; thus, taking the smallest one in case of multiple choices is not important but only for well-definedness.

The database property of interest can now be defined as

$$\text{SUC} := \left\{ D \mid \begin{array}{l} \exists \mathbf{y} \in \mathcal{Y}^\ell \text{ and } \text{inst} \in \mathcal{I} \text{ so that } \mathbf{m} := D^{-1}(\mathbf{y}) \text{ satisfies} \\ V(\text{inst}, c, \mathbf{m}_c) \text{ for } c := \gamma \circ D(\text{inst}, \mathbf{y}) \text{ and } (\text{inst}, \mathcal{E}^*(\text{inst}, \mathbf{m})) \notin R \end{array} \right\}. \quad (48)$$

Informally, assuming no collisions (i.e., restricting to $D \notin \text{CL}$), the database property **SUC** captures whether a database D admits a *valid* proof $\pi = (\mathbf{y}, \mathbf{m}_c)$ for an instance inst for which the (canonical) extractor, which first computes $\mathbf{m}$ by inverting D and then runs $\mathcal{E}^*$, *fails* to produce a witness.

Our (first) goal is to show that $\llbracket \perp \xrightarrow{q} \text{SUC} \cup \text{CL} \rrbracket$ is small, capturing that it is unlikely that after q queries the compressed database contains collisions or admits a valid proof upon which the extractor fails. Indeed, we show the following, where $p_{\text{triv}}^\mathfrak{S}$ is the trivial cheating probability of Π as defined in (47).

Lemma 5.16. $\llbracket \perp \xrightarrow{q} \text{SUC} \cup \text{CL} \rrbracket \leq 2eq^{3/2}2^{-n/2} + q\sqrt{10 \max(q\ell \cdot 2^{-n}, p_{\text{triv}}^\mathfrak{S})}.$

The formal proof is given below; we first give some informal outline here. In a first step, by using (45) and union-bound-like properties of the transition capacity, and additionally exploiting a bound from [CFHL21] to control the transition capacity of CL , we reduce the problem to bounding the quantum transition capacity $\llbracket \text{SZ}_{\leq s} \setminus \text{SUC} \rightarrow \text{SUC} \rrbracket$ for $s < q$. Informally, this capacity is a measure of the “likelihood”—but then in a *quantum*-sense—that a database $D \in \mathfrak{D}$ that is bounded in size and not in SUC turns into a database D' that is in SUC , when D is updated to $D' = D[x \mapsto U]$ with U uniformly random in $\mathcal{Y}$.

We emphasize that in the considered quantum setting, the state of the compressed oracle at any point is a *superposition* of databases, and a query is made up of a *superposition* of inputs; nevertheless, due to Theorem 5.4, the above classical intuition is actually very close to what needs to be shown to rigorously bound the considered quantum transition capacity. Formally, as will become clear in the proof below, we need to show that for any database $D \in \text{SZ}_{\leq s} \setminus \text{SUC}$ and for any $x \in \mathcal{X}$ with $D(x) = \perp$, the probability that $D[x \mapsto U] \in \text{SUC}$ is small. Below, this probability is bounded in the *Case 2* and *Case 3* parts of the proof, where the two cases distinguish between x being a “commit query” or a “challenge query”.

Informally, for D with $D(x) = \perp$, if x is a “commit query” then assigning a value to $D(x)$ can only make a difference, i.e., turn $D \notin \text{SUC}$ into $D[x \mapsto u] \in \text{SUC}$, if u is a coordinate of some $\mathbf{y} \in \mathcal{Y}^\ell$ for which $D(\text{inst}, \mathbf{y}) \neq \perp$ for some inst . Indeed, otherwise, $D[x \mapsto u]$ does not contribute to a valid proof π that did not exist before. Thus, given the bound $s < q$ on the size of D , this happens with probability at most $q\ell/2^n$ for a random u . Similarly, if x is a “challenge query”, i.e. of the form $x = (\text{inst}, \mathbf{y})$, then assigning a value u to $D(x)$ can only make a difference if $V(\text{inst}, c, \mathbf{m}_c)$ is satisfied for $c = \gamma(u)$ and $\mathbf{m} = D^{-1}(\mathbf{y})$, while $\mathcal{E}^*(\text{inst}, \mathbf{m})$ is not a witness for inst . However, for a random u , this is bounded by $p_{\text{triv}}^\mathcal{E}$.

But then, on top of the above, due to the quantum nature of the quantum transition capacity,⁹ Theorem 5.4 requires to also show the “reverse”, i.e., that for any $D \in \text{SUC}$ and for any $x \in \mathcal{X}$ with $D(x) \neq \perp$, the probability that $D[x \mapsto U] \in \text{SZ}_{\leq s} \setminus \text{SUC}$ is small; this is analyzed in *Case 1* below.

Thus, by exploiting the framework of [CFHL21], the core of the reasoning is purely classical, very closely mimicking how one would have to reason the classical setting with a classical RO. Due to the rather complex definition of SUC , the formal argument in each case is still somewhat cumbersome.

⁹ At the core, this is related to the reversibility of quantum computing and the resulting ability to “uncompute” a query.

5.4. Online Extractability of the FS-Transformation:
The Case of Ordinary C&O Protocols

Proof. We first observe that, by (45) (which is Lemma 5.6 in [CFHL21]) and basic properties of the quantum transition capacity as in (46),

$$\begin{aligned}
 \llbracket \perp \xRightarrow{q} \text{SUC} \cup \text{CL} \rrbracket &\leq \sum_{s=0}^{q-1} \llbracket \text{SZ}_{\leq s} \setminus \text{SUC} \setminus \text{CL} \rightarrow \text{SUC} \cup \text{CL} \cup \neg \text{SZ}_{\leq s+1} \rrbracket \\
 &\leq \sum_{s=0}^{q-1} (\llbracket \text{SZ}_{\leq s} \setminus \text{SUC} \setminus \text{CL} \rightarrow \neg \text{SZ}_{\leq s+1} \rrbracket + \llbracket \text{SZ}_{\leq s} \setminus \text{SUC} \setminus \text{CL} \rightarrow \text{CL} \rrbracket \\
 &\quad + \llbracket \text{SZ}_{\leq s} \setminus \text{SUC} \setminus \text{CL} \rightarrow \text{SUC} \rrbracket) \\
 &\leq \sum_{s=0}^{q-1} (\llbracket \text{SZ}_{\leq s} \rightarrow \neg \text{SZ}_{\leq s+1} \rrbracket + \llbracket \text{SZ}_{\leq s} \setminus \text{CL} \rightarrow \text{CL} \rrbracket + \llbracket \text{SZ}_{\leq s} \setminus \text{SUC} \rightarrow \text{SUC} \rrbracket).
 \end{aligned} \tag{49}$$

The first term, $\llbracket \text{SZ}_{\leq s} \rightarrow \neg \text{SZ}_{\leq s+1} \rrbracket$, vanishes, while the second term was shown to be bounded as

$$\llbracket \text{SZ}_{\leq s} \setminus \text{CL} \rightarrow \text{CL} \rrbracket \leq 2e\sqrt{(s+1)/|\mathcal{Y}|} \leq 2e\sqrt{q/2^n} \tag{50}$$

in Example 5.28 in [CFHL21]. Thus, it remains to control the third term, which we will do by means of Theorem 5.4 with $\mathbf{P} := \text{SZ}_{\leq s} \setminus \text{SUC}$ and $\mathbf{P}' := \text{SUC}$.

To this end, we consider arbitrary but fixed $D \in \mathfrak{D}$ and input $x \in \mathcal{X}$. By Remark 5.5, we may assume that $D(x) = \perp$. Furthermore, for $\mathbf{P}|_{D|x}$ to be non-empty, it must be that $D \in \text{SZ}_{\leq s}$, i.e., D is bounded in size. We now distinguish between the following cases for the considered D and x .

Case 1: $D \in \text{SUC}$. In particular, $\perp \in \text{SUC}|_{D|x} = \mathbf{P}'|_{D|x}$. So, Theorem 5.4 instructs us to set $\cdot := \mathbf{P}|_{D|x}$, where we leave the dependency of $\cdot$ on D and x implicit to simplify notation. Given that $D \in \text{SUC}$, we can consider *inst* and $\mathbf{y}$ as promised by the definition of SUC in (48), i.e., such that $V(\text{inst}, c, \mathbf{m}_c)$ and $(\text{inst}, \mathcal{E}^*(\text{inst}, \mathbf{m})) \notin R$ for

$$c := \gamma \circ D(\text{inst}, \mathbf{y}) \quad \text{and} \quad m_i := D^{-1}(y_i),$$

where it is understood that $\mathbf{m} = (m_1, \dots, m_\ell)$. Recall that $D(x) = \perp$; thus, by definition of the m_i 's, it must be that $x \neq m_i$ for all i , and the fact that $V(\text{inst}, c, \mathbf{m}_c)$ is satisfied for c as defined implies that $x \neq (\text{inst}, \mathbf{y})$. Furthermore,

$$u \in \mathbf{L} \iff D[x \mapsto u] \in \mathbf{P} \implies D[x \mapsto u] \notin \text{SUC} \implies u \in \{y_1, \dots, y_\ell\},$$

where the last implication is easiest seen by contraposition: Assume that $u \notin \{y_1, \dots, y_\ell\}$. Then, also recalling that $x \neq m_i$, we have that $m_i = D^{-1}(y_i) =$

$D[x \mapsto u]^{-1}(y_i)$. But also $c = \gamma \circ D(\text{inst}, \mathbf{y}) = \gamma \circ D[x \mapsto u](\text{inst}, \mathbf{y})$. Together, this implies that the defining property of **SUC** is also satisfied for $D[x \mapsto u]$, i.e., $D[x \mapsto u] \in \text{SUC}$, as was to be shown. Thus, we can bound

$$P[U \in] \leq P[U \in \{y_1, \dots, y_\ell\}] \leq \frac{\ell}{|\mathcal{Y}|}. \quad (51)$$

Case 2: $D \notin \text{SUC}$, and x is a “commit query”, i.e., $x = m \in \mathcal{M}$. In particular, $\perp \notin P'|_{D|x}$ (by the assumption that $D(x) = \perp$) and so in light of Theorem 5.4 we may choose $\mathbf{L} := P'|_{D|x}$. We then have

$$u \in \mathbf{L} \iff D[x \mapsto u] \in P' = \text{SUC} \implies \exists \text{inst}, \mathbf{y}, i : D(\text{inst}, \mathbf{y}) \neq \perp \wedge u = y_i. \quad (52)$$

This final implication can be seen as follows. By definition of **SUC**, the assumption $D[x \mapsto u] \in \text{SUC}$ implies the existence of inst and $\mathbf{y} = (y_1, \dots, y_\ell)$ with $V(\text{inst}, c, \mathbf{m}_c)$ and $(\text{inst}, \mathcal{E}^*(\text{inst}, \mathbf{m})) \notin R$ for

$$c := \gamma \circ D[x \mapsto u](\text{inst}, \mathbf{y}) = \gamma \circ D(\text{inst}, \mathbf{y}) \quad \text{and} \quad m_i := D[x \mapsto u]^{-1}(y_i),$$

where the equality in the definition of c exploits that x is not a “challenge” query. With the goal to reach a contradiction, assume that $u \neq y_i$ for all i . This assumption implies that $D[x \mapsto u](x) = u \neq y_i$. But also $D(x) = \perp \neq y_i$, and hence for all $\xi \in \mathcal{X}$ and $i \in \{1, \dots, \ell\}$: $D(\xi) = y_i \Leftrightarrow D[x \mapsto u](\xi) = y_i$. Therefore, $m_i = D[x \mapsto u]^{-1}(y_i) = D^{-1}(y_i)$ for all i , and the above then implies that $D \in \text{SUC}$, a contradiction. Thus, there exists i for which $u = y_i$; furthermore, $D(\text{inst}, \mathbf{y}) \neq \perp$ given that $V(\text{inst}, u, \mathbf{m}_c)$ is satisfied for $c = \gamma \circ D(\text{inst}, \mathbf{y})$. This shows the claimed implication.

Thus, we can bound

$$P[U \in] \leq P[\exists \text{inst}, \mathbf{y}, i : D(\text{inst}, \mathbf{y}) \neq \perp \wedge u = y_i] \leq \frac{s\ell}{|\mathcal{Y}|} \leq \frac{q\ell}{|\mathcal{Y}|}. \quad (53)$$

Case 3: $D \notin \text{SUC}$, and x is a “challenge query”, i.e., $x = (\text{inst}, \mathbf{y}) \in \mathcal{I} \times \mathcal{Y}^\ell$. Set $\mathbf{m} = (m_1, \dots, m_\ell)$ for $m_i := D^{-1}(y_i)$. Again, we have that $\perp \notin \text{SUC}|_{D|x} = P'|_{D|x}$, and so by Theorem 5.4 we may set $\mathbf{L} := P'|_{D|x}$. Here, we can argue that

$$\begin{aligned} u \in \mathbf{L} &\iff D[x \mapsto u] \in P' = \text{SUC} \\ &\implies V(\text{inst}, u, \mathbf{m}_{\gamma(u)}) \text{ and } (\text{inst}, \mathcal{E}^*(\text{inst}, \mathbf{m})) \notin R, \end{aligned}$$

where the final implication can be seen as follows. By definition of **SUC**, the assumption $D[x \mapsto u] \in \text{SUC}$ implies the existence of inst' and $\mathbf{y}' = (y'_1, \dots, y'_\ell)$

5.4. Online Extractability of the FS-Transformation:
The Case of Ordinary C&O Protocols

with $V(inst', u, \mathbf{m}'_c)$ and $\mathcal{E}^*(inst', \mathbf{m}') \neq w$ for

$$c := \gamma \circ D[x \mapsto u](inst', \mathbf{y}') \quad \text{and} \quad m'_i := D[x \mapsto u]^{-1}(y'_i) = D^{-1}(y'_i),$$

where the very last equality exploits that x is not a “commit” query. With the goal to come to a contradiction, assume that $(inst', \mathbf{y}') \neq (inst, \mathbf{y}) = x$. Then, $c = \gamma \circ D[x \mapsto u](inst', \mathbf{y}') = \gamma \circ D(inst', \mathbf{y}')$, and the above then implies that $D \in \text{SUC}$, a contradiction. Thus, $(inst', \mathbf{y}') = (inst, \mathbf{y}) = x$. In particular, $\mathbf{m}' = \mathbf{m}$ and $c = \gamma \circ D[x \mapsto u](inst', \mathbf{y}') = \gamma \circ D[x \mapsto u](x) = \gamma(u)$. Hence, the claimed implication holds.

Thus, we can bound

$$\begin{aligned} P[U \in] &\leq P[V(inst, \gamma(U), \mathbf{m}_{\gamma(U)}) \wedge \mathcal{E}^*(inst, \mathbf{m}) \neq w] \\ &\leq P[V(inst, \gamma(U), \mathbf{m}_{\gamma(U)}) \wedge S := \{c \mid V(inst, c, \mathbf{m}_c)\} \notin \mathfrak{S}] \\ &\leq P[\gamma(U) \in S := \{c \mid V(inst, c, \mathbf{m}_c)\} \notin \mathfrak{S}] \\ &\leq \max_{S \notin \mathfrak{S}} P[\gamma(U) \in S] \\ &\leq p_{triv}^{\mathfrak{S}}. \end{aligned} \tag{54}$$

By Theorem 5.4, we now get

$$\begin{aligned} \llbracket \text{SZ}_{\leq s} \setminus \text{SUC} \setminus \text{CL} \rightarrow \text{SUC} \rrbracket &\leq \max_{x, D} \sqrt{10 P[U \in \mathcal{L}^{x, D}]} \\ &\leq \sqrt{10} \sqrt{\max \left(\frac{\ell}{|\mathcal{Y}|}, \frac{q\ell}{|\mathcal{Y}|}, p_{triv}^{\mathfrak{S}} \right)} \\ &\leq \sqrt{10} \sqrt{\max (q\ell \cdot 2^{-n}, p_{triv}^{\mathfrak{S}})}, \end{aligned}$$

where we have used Equations (51), (53) and (54) in the second inequality. Combining with Equations (50) and (49) yields the desired bound. $\square$

5.4.2 Online Extractability of the Fiat-Shamir Transformation

We are now ready to state and proof the claimed online-extractability result for the Fiat-Shamir transformation of (ordinary) C&O protocols.

Theorem 5.17. *Let Π be a $\mathfrak{S}$ -sound* ordinary C&O protocol with challenge space $\mathcal{C}_\lambda$ and $\ell = \ell(\lambda)$ commitments, and set $\kappa = \kappa(\lambda) := \max_{c \in \mathcal{C}_\lambda} |c|$. Then,*

$\text{FS}[II]$ is a proof of knowledge with online extractability in the QROM (as in Definition 5.9), with $\varepsilon_{\text{sim}}(\lambda, q, n) = 0$ and

$$\begin{aligned} \varepsilon_{\text{ex}}(\lambda, q, n) &\leq 2(\kappa + 1) \cdot 2^{-n} + \left(2eq^{3/2}2^{-n/2} + q\sqrt{10 \max(q\ell \cdot 2^{-n}, p_{\text{triv}}^{\mathfrak{S}})} \right)^2 \\ &\leq (22\ell + 60)q^3 2^{-n} + 20q^2 p_{\text{triv}}^{\mathfrak{S}}. \end{aligned}$$

The runtime of the extractor is dominated by running the compressed oracle, which has complexity $O(q^2) \cdot \text{poly}(n, B)$, and running $\mathcal{E}^*$.

We note that the above bound on ε_{ex} is asymptotically tight, except for the factor ℓ . Indeed, the binding property of the hash-based commitment can be invalidated by means of a collision finding attack, which succeeds with probability $\Omega(q^3/2^n)$. Furthermore the trivial soundness attack, which potentially applies to a $\mathfrak{S}$ -sound* C&O protocol II , can be complemented with a Grover search, yielding an attack against $\text{FS}[II]$ that succeeds with probability $\Omega(q^2 p_{\text{triv}}^{\mathfrak{S}})$. The non-tightness by a factor of ℓ is very mild in most cases. In particular, the number of commitments ℓ is polynomial in λ and thus in n . For the most common case of a parallel repetition of a protocol with a constant number of commitments, using a hash function with output length linear in λ (e.g. $n = 3\lambda$) results in $\ell = O(n) = O(\lambda)$.

Proof. We consider an arbitrary but fixed $\lambda \in \mathbb{N}$. For simplicity, we assume that $|c|$ is the same for all $c \in \mathcal{C}_\lambda$, and thus equal to $\kappa = \kappa(\lambda)$. If it is not, we could always make the prover output a couple of dummy outputs m_i to match the upper bound on $|c|$. Let $\mathcal{P}^*$ be a dishonest prover that, after making q queries to a random oracle H , outputs $(\text{inst}, \pi) = (\text{inst}, \mathbf{y}, \mathbf{m}_o)$ plus some (possibly quantum) auxiliary output Z . In the experiment $\mathcal{V}^{\mathcal{E}} \circ \mathcal{P}^{*\mathcal{E}}(\lambda)$, our extractor $\mathcal{E}$ works as follows while simulating all queries to H (by $\mathcal{P}^*$ and $\mathcal{V}$) with the compressed oracle:

1. Run $\mathcal{P}^*(\lambda)$ to obtain (inst, π, Z) where $\pi = (\mathbf{y}, \mathbf{m}_o)$ with $\mathbf{m}_o = (m_1, \dots, m_\kappa)$.
2. Run $\mathcal{V}(\lambda, \text{inst}, \pi)$ to obtain v . In detail: obtain $h_0 := H(\text{inst}, \mathbf{y})$ and $h_j := H(m_j)$ for $j \in \{1, \dots, \kappa\}$, and set $v := \text{accept}$ if and only if the pair consisting of $\mathbf{x} = ((\text{inst}, \mathbf{y}), m_1, \dots, m_\kappa)$ and $\mathbf{h} = (h_0, h_1, \dots, h_\kappa)$ satisfies the relation $\tilde{R}$, defined to hold if and only if

$$(h_1, \dots, h_\kappa) = \mathbf{y}_c \quad \wedge \quad V(\text{inst}, c, \mathbf{m}_o) \quad \text{where} \quad c := \gamma(h_0).$$

3. Measure the internal state of the compressed oracle to obtain D .

5.4. Online Extractability of the FS-Transformation:
The Case of Ordinary C&O Protocols

4. Run $\mathcal{E}^*(inst, \mathbf{m})$ on input $inst$ and $\mathbf{m} := D^{-1}(\mathbf{y})$ to obtain w .

Note that in the views of both $\mathcal{P}^*$ and $\mathcal{V}$, the interaction with H and the interaction with $\mathcal{E}$ differ only in that their oracle queries are answered by a compressed oracle instead of a real random-oracle in the latter case. This simulation is perfect and therefore $\varepsilon_{\text{sim}}(\lambda, q, n) = 0$.

Considering $\mathcal{P}^*$ as the algorithm $\mathcal{A}$ in Lemma 5.6, the additional classical oracle queries that $\mathcal{V}$ performs in $\mathcal{V} \circ \mathcal{P}^*$ then match up with the algorithm $\tilde{\mathcal{A}}$, with $h_0, \dots, h_\kappa$ here playing the role of $y_1, \dots, y_\ell$ in Lemma 5.6. Thus,

$$\Pr[\mathbf{h} \neq D(\mathbf{x})] \leq 2(\kappa(\lambda) + 1) \cdot 2^{-n}.$$

Therefore, we can bound the figure of merit ε_{ex} as

$$\begin{aligned} \varepsilon_{\text{ex}}(\lambda, q, n) &= \Pr[v = \text{accept} \wedge (inst, w) \notin R] = \Pr[(\mathbf{x}, \mathbf{h}) \in \tilde{R} \wedge (inst, w) \notin R] \\ &\leq \Pr[(\mathbf{x}, D(\mathbf{x})) \in \tilde{R} \wedge (inst, w) \notin R] + 2(\kappa(\lambda) + 1) \cdot 2^{-n} \\ &\leq \Pr[(\mathbf{x}, D(\mathbf{x})) \in \tilde{R} \wedge (inst, w) \notin R \mid D \notin \text{SUC} \cup \text{CL}] + \Pr[D \in \text{SUC} \cup \text{CL}] + 2(\kappa(\lambda) + 1) \cdot 2^{-n} \end{aligned}$$

Using the definition of $\tilde{R}$, understanding that $c := \gamma \circ D(inst, \mathbf{y})$, we can write the first term as

$$\begin{aligned} &\Pr[D(\mathbf{m}_o) = \mathbf{y}_c \wedge V(\lambda, inst, c, \mathbf{m}_o) \wedge (inst, w) \notin R \mid D \notin \text{SUC} \cup \text{CL}] \\ &\leq \Pr[V(\lambda, inst, c, \mathbf{m}_c) \text{ for } \mathbf{m} := D^{-1}(\mathbf{y}) \wedge (inst, w) \notin R \mid D \notin \text{SUC} \cup \text{CL}] \\ &\leq \Pr[D \in \text{SUC} \mid D \notin \text{SUC} \cup \text{CL}] \\ &= 0, \end{aligned}$$

where the first equality exploits that $D(m) = y$ iff $m = D^{-1}(y)$ for $D \notin \text{CL}$.

We may thus conclude that

$$\begin{aligned} \varepsilon_{\text{ex}}(\lambda, q, n) &\leq (2\kappa(\lambda) + 1) \cdot 2^{-n} + \Pr[D \in \text{SUC} \cup \text{CL}] \\ &\leq (2\kappa(\lambda) + 1) \cdot 2^{-n} + \llbracket \perp \xrightarrow{q} \text{SUC} \cup \text{CL} \rrbracket^2, \end{aligned}$$

where the last inequality is by definition (43) of $\llbracket \perp \xrightarrow{q} \cdot \rrbracket$. The claimed bound now follows from Lemma 5.16. $\square$

5.4.3 The Unruh-Transformation with a Compressing Hash Function

We conclude this section by showing an improvement to the *Unruh transformation* [Unr15b], which follows directly from our result above. At the core of

the Unruh transformation is a generic technique to transform any Σ -protocol into a C&O protocol. In [Unr15b], this transformation is presented in combination with parallel repetition and the Fiat-Shamir transformation as a means to construct (online-extractable) NIZK proofs of knowledge in the QROM. The entire transformation was later dubbed the Unruh transformation.

In fact, the Unruh transformation was the first NIZK proof of knowledge in the QROM; the QROM security of the Fiat-Shamir transformation was only established several years later [DFMS19; LZ19a]. Despite being significantly less efficient than the Fiat-Shamir transformation, the Unruh transformation is still useful in certain cases because it puts weaker requirements on the underlying Σ -protocol.

Here, to allow for a modular analysis, we consider the first step of the Unruh transformation, i.e., the transformation from a Σ -protocol into a C&O protocol, as an individual transformation, which we refer to as the *pre-Unruh transformation*, formally defined below. We stress that we allow the random oracle H to be *compressing*, i.e. $|\mathcal{Y}| < |\mathcal{X}|$, while the extraction technique of [Unr15b] required H to be a *length-preserving* RO. This obviously has a significant positive impact on the efficiency of the Unruh transformation.

Let $\Sigma = (\mathcal{P}_o, \mathcal{V}_o)$ be a Σ -protocol. We write $a_o \leftarrow \mathcal{P}_o$ to denote the first message in Π_o as produced by $\mathcal{P}_o$ (for a given instance *inst*). Furthermore, we write $z(a_o, c)$ for $\mathcal{P}_o$'s response then upon receiving challenge $c \in \mathcal{C}$.¹⁰

Definition 5.18 (Pre-Unruh transformation). *Let $\Sigma = (\mathcal{P}_o, \mathcal{V}_o)$ be a Σ -protocol as above. Then, the pre-Unruh-transformation $\text{pU}[\Sigma] = (\mathcal{P}, \mathcal{V})$ of Π_o is the C&O protocol with first message*

$$a := (a_o, (y_i)_{i \in \mathcal{C}})$$

where $a_o \leftarrow \mathcal{P}_o$ and for each $i \in \mathcal{C}$, $y_i := H(z_i)$ for $z_i := z(a_o, i)$), and with response $z := z_c$ upon challenge $c \in \mathcal{C}$. To verify, $\mathcal{V}$ runs $\mathcal{V}_o$ on (a_o, c, z) and checks if $H(z) = y_c$; if both are true, it accepts, otherwise it rejects.

Clearly, $\text{pU}[\Sigma]$ is only efficient if Σ has at most polynomially many possible challenges (which can always be obtained by restricting the challenge space). As mentioned, the resulting C&O protocol can then be repeated in parallel and made non-interactive using the Fiat-Shamir transformation. We will now provide a fairly straightforward corollary to conclude the security of the more

¹⁰ We note that $z(a_o, c)$ may be a *randomized* function of a_o and c . Furthermore, $z(a_o, c)$ is typically computed by $\mathcal{P}_o$ by means of the *randomness used to produce a_o* .

5.4. Online Extractability of the FS-Transformation: The Case of Ordinary C&O Protocols

efficient variant of the (full) Unruh transformation that allows for a compressing RO, given by the composition of the pre-Unruh transformation introduced above, parallel repetition and the Fiat-Shamir transformation. In the following, denote the r -fold parallel repetition of a (C&O) Σ -protocol Π by Π^r and use the notation $\text{Unr}_r[\Sigma] := \text{FS}[\text{pU}[\Sigma]^r]$ for the Unruh transformation with r -fold parallel repetition.

Remark 5.19. A proof in $\Pi = \text{Unr}_r[\Sigma]$ can be generated in time $T_{\mathcal{P}}^{\Pi} = rT_{\mathcal{P}}^{\Sigma} + (\ell_0 r + 1)T_H$, and verified in time $T_{\mathcal{V}}^{\Pi} = rT_{\mathcal{V}}^{\Sigma} + (1 + r)T_H$, where $T_{\mathcal{P}}^{\Sigma}, T_{\mathcal{V}}^{\Sigma}$ and T_H are the prover and verifier runtime of Σ , and the time required for computing one hash, respectively.

It is straightforward to verify that the pre-Unruh transformation does not harm most security properties of the Σ -protocol. In particular, it tightly preserves soundness and honest-verifier zero-knowledge (in the QROM). It also preserves $\mathfrak{S}$ -soundness in a certain sense.

Proposition 5.20. *Let Σ be an $\mathfrak{S}$ -sound Σ -protocol with challenge space size $\ell = \ell(\lambda)$ with extractor runtime T . Then $\Pi := \text{pU}[\Sigma]$ is $\mathfrak{S}$ -sound as a C&O protocol with extractor runtime $T' \leq T + O(\ell)$. Furthermore, suppose that membership in $\mathfrak{S}$ is checkable in time $T_{\mathfrak{S}}$. Then Π is $\mathfrak{S}$ -sound* with extractor runtime $T'' \leq T' + \ell^2 T_{\mathfrak{S}} + \ell T_{\mathcal{V}}$, where $T_{\mathcal{V}}$ is the runtime of Π 's verification predicate $\mathcal{V}$.*

Proof. Let $\mathcal{E}_{\Sigma}$ be the extractor for Σ guaranteed to exist by Definition 4.22. Note that for $\Pi = \text{pU}[\Sigma]$ regarded as a C&O protocol, for each challenge exactly one of the commitments has to be opened. For such protocols, we use c and $\{c\}$ interchangeably (where c is a challenge in Π). We define an extractor $\mathcal{E}_{\Pi}$ as follows. On input $(\text{inst}, m_1, \dots, m_{\ell}, a_{\circ}, S)$, run $w = \mathcal{E}_{\Sigma}(\text{inst}, a_{\circ}, S, \{m_c\}_{c \in S})$, then output w . The only runtime overhead of $\mathcal{E}_{\Pi}$ results from having to parse its input and preparing the input for $\mathcal{E}_{\Sigma}$.

We continue to define an $\mathfrak{S}$ -soundness* extractor $\mathcal{E}_{\Pi}^*$ for Π as follows. On input $(\text{inst}, m_1, \dots, m_{\ell}, a_{\circ})$, compute $b_c = \mathcal{V}(\text{inst}, a_{\circ}, c, m_c)$ for all $c \in \mathcal{C}$, and set $\hat{S} = \{c \in \mathcal{C} \mid b_c = 1\}$. Using at most $\ell(\ell + 1)/2$ membership tests for $\mathfrak{S}$, find $S \subseteq \hat{S}$ such that $S \in \mathfrak{S}_{\min}$. Finally, run $w = \mathcal{E}_{\Pi}(\text{inst}, m_1, \dots, m_{\ell}, a_{\circ}, S)$ and output w . The runtime statement is straightforward. $\square$

Using Proposition 5.20 above and Lemma 5.3 from [DFMS22a] to argue $\mathfrak{S}^{\vee r}$ -soundness* of the parallel repetition of $\text{pU}[\Pi]$, and using Theorem 5.17 to argue online extractability of its Fiat-Shamir transformation, we obtain the

online-extractability of the Unruh transformation with computationally binding commitments, i.e., when using a *compressing* hash function for the commitments.

Corollary 5.21. *Let Σ be an $\mathfrak{S}$ -sound Σ -protocol with challenge space size ℓ_0 . Then $\Pi := \text{Unr}_r[\Sigma] = \text{FS}[\text{pU}[\Sigma]^r]$ is a proof of knowledge with online extractability in the QROM (as in Definition 5.9) with $\varepsilon_{\text{sim}} = 0$ and*

$$\varepsilon_{\text{ex}}(\lambda, q, n) \leq (22r\ell_0 + 60)q^3 2^{-n} + 20q^2 \left(p_{\text{triv}}^{\mathfrak{S}}\right)^r. \quad (55)$$

The online extractor for Π runs in time $T_{\mathcal{E}}^{\Pi} \leq rT_{\mathcal{E}}^{\text{pU}[\Sigma]} + O(q^2) \cdot \text{poly}(n, B)$, where $T_{\mathcal{E}}^{\text{pU}[\Sigma]}$ is the runtime of $\text{pU}[\Sigma]$'s $\mathfrak{S}$ -soundness* extractor as given in Proposition 5.20.

Section 5.5

Online Extractability of the FS-Transformation: The Case of Merkle-tree-based C&O Protocols

For an ordinary C&O protocol with reasonable concrete security (e.g., 128 bits), the number of commitments ℓ might be considerable. In this case, the communication complexity of the protocol (and thus the size of the non-interactive proof system, or digital-signature scheme, obtained via the Fiat-Shamir transformation) can be reduced by using a *Merkle tree* to collectively commit to the ℓ strings m_i . Such a construction is mentioned in [Fis05], and it is used in the construction of the digital-signature schemes Picnic2 and Picnic3 [KKW18; CDG+20; KZ20; CDG+19]. The Merkle-tree-based C&O mechanism shrinks the commitment information from $\ell \cdot n$ to n , at the expense of increasing the cost of opening $|c|$ values m_i by an additive term of about $\lesssim |c| \cdot n \cdot \log \ell$.

The cost of opening can, in fact, be slightly reduced again, by streamlining the opening information. When opening several leaves of a Merkle tree, the authentication paths overlap, so opening requires a number of hash values less than h per leaf, where h is the height of the tree. This overlap was observed and exploited in the octopus authentication algorithm which constitutes one of the optimizations of the stateless hash-based signature scheme gravity-SPHINCS [AE18], as well as in Picnic2 and Picnic3 [CDG+20; KZ20]. In the following

section, we formalize tree-based collective commitment schemes with “octopus” opening.

5.5.1 Merkle-tree-based C&O Protocols

As was noted in Section 2.2.1, we can consider C&O protocols with a different choice of commitment scheme, compared to the default choice of committing by element-wise hashing. Here, we discuss a particular choice of an alternative commitment scheme, which gives rise to more efficient C&O protocols in certain cases when ℓ is large. Informally, we consider C&O protocols where $m_1, \dots, m_\ell$ is committed to by using a *Merkle tree*, and individual m_i ’s are opened by announcing the corresponding authentication paths.

To make this more formal, we introduce the following notation. For simplicity, we assume that ℓ is a power of 2, and thus $\ell = 2^h$ for $h \in \mathbb{N}$. We then consider the *full binary tree* $\text{Tree} = \{0, 1\}^{\leq h}$ of depth h , where the vertices are identified by bit strings. The root is denoted by $\emptyset$; the i -th leaf is denoted by $\text{lf}(i) \in \{0, 1\}^h$ and is given by the binary representation of $i \in [\ell]$. The *authentication path* for the i -th leaf is the subtree that consists of all the ancestors of $\text{lf}(i)$ and their siblings:

$$\text{Auth}(i) := \text{Anc}(\text{lf}(i)) \cup \{\text{sib}(v) \mid \emptyset \neq v \in \text{Anc}(\text{lf}(i))\},$$

where $\text{Anc}(v) := \{u \in \text{Tree} \mid \exists w : u \parallel w = v\}$ and $\text{sib}(u \parallel b) := u \parallel (1 - b)$ for any $b \in \{0, 1\}$. Finally, for any subset $c \subseteq \{1, \dots, \ell\}$, we let $\text{Auth}(c) := \bigcup_{i \in c} \text{Auth}(i)$ be the union of the authentication paths of the considered leaves, and we define the *octopus* $\text{Octo}(c)$ to be the restriction of $\text{Auth}(c)$ to its leaves, but excluding the leaves $\text{lf}(i)$ for $i \in c$, i.e.,

$$\text{Octo}(c) := \text{leaves}(\text{Auth}(c)) \setminus \{\text{lf}(i) \mid i \in c\}$$

where, for any subtree T of Tree , $\text{leaves}(T) := \{v \in T \mid (v \parallel 0), (v \parallel 1) \notin T\}$.

Extending on the above notation, for a given hash function $H : \mathcal{X} \rightarrow \mathcal{Y}$, where $\mathcal{X} = \{0, 1\}^{\leq B}$ and $\mathcal{Y} = \{0, 1\}^n$ for sufficiently large B , we define the *Merkle tree* of $\mathbf{m} = (m_1, \dots, m_\ell) \in \mathcal{X}^\ell$ to be the *labeled* binary tree that has its leaves $\text{lf}(1), \dots, \text{lf}(\ell)$ labeled by $H(m_1), \dots, H(m_\ell)$, respectively, and each internal vertex is labeled by the hash of the labels of its two children. Formally,

$$\text{MTree}_H(\mathbf{m}) := \{(v, l_v(\mathbf{m})) \mid v \in \text{Tree}\}$$

with the labeling $l_v(\mathbf{m})$ recursively defined as

$$l_v(\mathbf{m}) := H(l_{v \parallel 0}(\mathbf{m}) \parallel l_{v \parallel 1}(\mathbf{m})) \text{ for } v \in \{0, 1\}^{< h}$$

and

$$l_{\text{lf}(i)}(\mathbf{m}) := H(m_i) \text{ for } i \in \{1, \dots, \ell\},$$

where we leave the dependency of the labeling on H , i.e., $l_v = l_v^H$, implicit. We also write $\text{MRoot}_H(\mathbf{m})$ then for the root label $l_\emptyset(\mathbf{m})$. In the same spirit, we write $\text{MAuth}_H(c, \mathbf{m}) := \{(v, l_v(\mathbf{m})) \mid v \in \text{Auth}(c)\}$ for the labeled authentication path and $\text{MOcto}_H(c, \mathbf{m}) := \{(v, l_v(\mathbf{m})) \mid v \in \text{Octo}(c)\}$ for the labeled octopus, using the same labeling function as for the Merkle tree.




Fig. 5.1. The Merkle tree $\text{MTree}_H(\mathbf{m})$ for $\mathbf{m} = (m_1, \dots, m_8)$ with $\text{MRoot}_H(\mathbf{m}) = y$. The yellow vertices mark the octopus $\text{MOcto}_H(\{1\}, \mathbf{m})$, which is revealed (along with m_1) when opening the commitment y to m_1 .

A *Merkle-tree-based C&O* protocol is now defined to be a variation of a C&O protocol, where the first message of the protocol, i.e., the commitment of $\mathbf{m} = (m_1, \dots, m_\ell)$, is computed as $y = \text{MRoot}_H(\mathbf{m})$, and the response z for challenge-set c then consists of the messages $\mathbf{m}_c = (m_i)_{i \in c}$ together with $O = \text{MOcto}_H(c, \mathbf{m})$. The verifier $\mathcal{V}$ then accepts if and only if $\mathbf{m}_c$ and O “hash down to” y and the predicate $V(\lambda, \text{inst}, c, \mathbf{m}_c, a)$ is satisfied. More formally, the former means that $\mathcal{V}$ computes $\text{MAuth}_H(c, \mathbf{m})$ from $O \cup \{(\text{lf}(i), H(m_i)) \mid i \in c\}$ in the obvious way, and then checks whether $l_\emptyset(\mathbf{m}) = y$. This verification is denoted by $\text{OctoVerify}^H(c, y, \mathbf{m}_c, O)$, see Figure 5.2.

Looking ahead, we may also consider a variation where the verifier resamples the challenge c if the resulting octopus is bigger than a given bound. Formally, this means that the challenge space of the Merkle-tree-based C&O protocol is restricted to those challenges $c \in [\ell]$ for which $\text{Octo}(c)$ is not too large.

5.5. Online Extractability of the FS-Transformation: The Case of Merkle-tree-based C&O Protocols




Fig. 5.2. A Merkle-tree based C&O Σ -protocol, formally introduced in Section 5.5.1.

5.5.2 Online Extractability of the Fiat-Shamir Transformation

The analysis in Section 5.4 can be generalized to the case of FS-transformed Merkle-tree-based C&O protocols. To that end, we generalize the notation from that section as follows. Let Π be a Merkle-tree-based C&O protocol with number of messages to be committed equal to $\ell = 2^h$ where h is the height of the commitment Merkle tree.¹¹

For a given database $D \in \mathfrak{D}$, recall from Section 5.4 the definition of D^{-1} ; applied to a tuple $\mathbf{y} = (y_1, \dots, y_\ell) \in \mathcal{Y}^\ell$ of commitments, D^{-1} attempts to recover the corresponding committed messages $m_1, \dots, m_\ell$. Here, in a similar spirit but now considering the Merkle-tree commitment, MRoot_D^{-1} attempts to recover the committed messages from the root label of the Merkle tree.

In more detail, for a commitment $y \in \mathcal{Y} = \{0, 1\}^n$ we reverse engineer the Merkle tree in the obvious way (see Figure 5.3 for an example); namely, accepting a small clash in notation with the labeling function $l_v(\mathbf{m})$ defined for a tuple $\mathbf{m} \in \mathcal{M}^\ell$, we set the root label $l_\emptyset(y) := y$, and recursively define

$$(l_{v\|0}(y), l_{v\|1}(y)) := \text{split} \circ D^{-1}(l_v(y)) \in \mathcal{Y} \times \mathcal{Y}$$

for $\emptyset \neq v \in \{0, 1\}^{\leq h}$, where split maps any $2n$ -bit string, parsed as $y_1\|y_2$ with $y_1, y_2 \in \{0, 1\}^n$, to the pair (y_1, y_2) of n -bit strings, while it maps anything else to $(\perp, \perp)$. Then, accepting a small clash in notation again, we set

$$\text{MTree}_D(y) := \{l_v(y) \mid v \in \{0, 1\}^{\leq h}\},$$

and finally

$$\text{MRoot}_D^{-1}(y) := (D^{-1}(l_{\text{f}(1)}(y)), \dots, D^{-1}(l_{\text{f}(\ell)}(y))).$$

¹¹ As in the previous section we assume that ℓ is a power of 2 for ease of exposition.

Following the strategy we used in Section 5.4, we define the database property

$$\text{SUC} := \left\{ D \mid \begin{array}{l} \exists y \in \mathcal{Y} \text{ and } \text{inst} \in \mathcal{I} \text{ so that } \mathbf{m} := \text{MRoot}_D^{-1}(y) \text{ satisfies} \\ V(\text{inst}, c, \mathbf{m}_c) \text{ for } c := \gamma \circ D(\text{inst}, y) \text{ and } (\text{inst}, \mathcal{E}^*(\text{inst}, \mathbf{m})) \notin R \end{array} \right\},$$

and our first goal is to show that $\llbracket \perp \xRightarrow{q} \text{SUC} \cup \text{CL} \rrbracket$ is small.

Lemma 5.22. *Let Π be an $\mathfrak{S}$ -sound C&O protocol with $p_{\text{triv}}^{\mathfrak{S}}$ as defined in (47). Then*

$$\llbracket \perp \xRightarrow{q} \text{SUC} \cup \text{CL} \rrbracket \leq 2eq^{3/2}2^{-n/2} + q\sqrt{10 \max(q\ell \cdot 2^{-n+1}, p_{\text{triv}}^{\mathfrak{S}})}.$$

The proof works exactly as the proof of Lemma 5.16, accounting for some syntactic differences due to the Merkle tree commitment. In particular, where in Case 1 and 2 of the proof of Lemma 5.16 we have to exclude U from falling on one of the hash values $y_1, \dots, y_\ell$ in order to keep the $\mathbf{m}$ that was constructed from the database intact, we now have a similar restriction for U , but with respect to the whole tree $\text{MTree}_D(y)$.

Proof. As in the proof of Lemma 5.16, we can bound

$$\llbracket \perp \xRightarrow{q} \text{SUC} \cup \text{CL} \rrbracket \leq \sum_{s=0}^{q-1} (\llbracket \text{SZ}_{\leq s} \setminus \text{CL} \rightarrow \text{CL} \rrbracket + \llbracket \text{SZ}_{\leq s} \setminus \text{SUC} \rightarrow \text{SUC} \rrbracket) \quad (56)$$

and use that

$$\llbracket \text{SZ}_{\leq s} \setminus \text{CL} \rightarrow \text{CL} \rrbracket \leq 2e\sqrt{(s+1)/2^n} \leq 2e\sqrt{q/2^n}. \quad (57)$$

Thus, it remains to control the second term, which we will do again by means of Theorem 5.4 with $\mathbf{P} := \text{SZ}_{\leq s} \setminus \text{SUC}$ and $\mathbf{P}' := \text{SUC}$.

To this end, we consider arbitrary but fixed $D \in \mathfrak{D}$ and input $x \in \mathcal{X}$. By Remark 5.5, we may assume that $D(x) = \perp$. Furthermore, for $\mathbf{P}|_{D|x}$ to be non-empty, it must be that $D \in \text{SZ}_{\leq s}$, i.e., D is bounded in size. We now distinguish between the following cases for the considered D and x .

Case 1: $D \in \text{SUC}$. In particular, $\perp \in \text{SUC}|_{D|x} = \mathbf{P}'|_{D|x}$. So, Theorem 5.4 instructs us to set $\gamma := \mathbf{P}|_{D|x}$, where we leave the dependency of γ on D and x implicit. Given that $D \in \text{SUC}$, we can consider inst and y as promised by the definition of SUC above, i.e., such that $V(\text{inst}, c, \mathbf{m}_c)$ and $(\text{inst}, \mathcal{E}^*(\text{inst}, \mathbf{m})) \notin R$ for

$$c := \gamma \circ D(\text{inst}, y) \quad \text{and} \quad \mathbf{m} := \text{MRoot}_D^{-1}(y). \quad (58)$$

5.5. Online Extractability of the FS-Transformation: The Case of Merkle-tree-based C&O Protocols

Note that, since $D(x) = \perp$ and $V(\text{inst}, c, \mathbf{m}_c)$ holds, which in particular means that c must be defined, it must be that $x \neq (\text{inst}, y)$. Therefore

$$\gamma \circ D(\text{inst}, y) = \gamma \circ D[x \mapsto u](\text{inst}, y). \quad (59)$$

Our goal now is to show the final implication in

$$u \in \mathbf{L} \iff D[x \mapsto u] \in \mathbf{P} \implies D[x \mapsto u] \notin \mathbf{SUC} \implies u \in \mathbf{MTree}_D(y).$$

We will do this by showing that $u \notin \mathbf{MTree}_D(y)$ implies

$$\mathbf{MRoot}_D^{-1}(y) = \mathbf{MRoot}_{D[x \mapsto u]}^{-1}(y). \quad (60)$$

Indeed, the contraposition $u \notin \mathbf{MTree}_D(y) \Rightarrow D[x \mapsto u] \in \mathbf{SUC}$ of the claimed implication then follows from the fact that (59) and (60) together imply that c and $\mathbf{m}$ remain unchanged when replacing D by $D[x \mapsto u]$ in (58), and so $D[x \mapsto u] \in \mathbf{SUC}$ as well.

Towards showing (60), exploiting again that $D(x) = \perp$, it follows by definition of the reverse engineered labeling function $l_v(y)$ that $x \neq (l_{v||0}(y), l_{v||1}(y))$ for any v with $l_{v||0}(y) \neq \perp \neq l_{v||1}(y)$, i.e., x is not equal to any pair of siblings in $\mathbf{MTree}_D(y)$ with non- $\perp$ labeling (see Figure 5.3). Due to a similar reasoning, $x \neq m_i$ for any i . It now follows by definition of the reverse engineered Merkle tree and of $\mathbf{MRoot}^{-1}$ that if $u \notin \mathbf{MTree}_D(y)$ then $\mathbf{MTree}_D(y) = \mathbf{MTree}_{D[x \mapsto u]}(y)$ and $\mathbf{MRoot}_D^{-1}(y) = \mathbf{MRoot}_{D[x \mapsto u]}^{-1}(y)$, as claimed.

Thus, we can bound

$$P[U \in \cdot] \leq P[U \in \mathbf{MTree}_D(y)] \leq \frac{2 \cdot 2^h - 1}{|\mathcal{Y}|} = \frac{2\ell - 1}{|\mathcal{Y}|}. \quad (61)$$

Case 2: $D \notin \mathbf{SUC}$, and x is a “commit query”, i.e., $x = m \in \mathcal{M}$ or $x = (l_{v||0}, l_{v||1})$ for two labels $l_{v||0}, l_{v||1} \in \mathcal{Y}$. In particular, $\perp \notin \mathbf{P}'|_{D|x}$ (given that $D(x) = \perp$) and so in the light of Theorem 5.4 we may choose $\mathbf{L} := \mathbf{P}'|_{D|x}$. We then have

$$u \in \mathbf{L} \iff D[x \mapsto u] \in \mathbf{P}' = \mathbf{SUC} \implies \exists \text{ inst}, y : D(\text{inst}, y) \neq \perp \wedge u \in \mathbf{MTree}_D(y)$$

where final implication can be seen as follows. By definition of $\mathbf{SUC}$, the assumption $D[x \mapsto u] \in \mathbf{SUC}$ implies the existence of inst and y with $V(\text{inst}, c, \mathbf{m}_c)$ and $(\text{inst}, \mathcal{E}^*(\text{inst}, \mathbf{m})) \notin R$ for

$$c := \gamma \circ D[x \mapsto u](\text{inst}, y) = \gamma \circ D(\text{inst}, y) \quad \text{and} \quad \mathbf{m} := \mathbf{MRoot}_{D[x \mapsto u]}^{-1}(y),$$




Fig. 5.3. Example of a reverse engineered Merkle tree $\text{MTree}_D(y)$, with the $\perp$ -children of the $\perp$ -labels omitted. Since $D(x) = \perp$, $x \neq (l_u(y), l_w(y))$ for any two siblings (u, w) in $\text{MTree}_D(y)$, i.e., nodes with the same color. Assuming that $u \notin \text{MTree}_D(y)$ then implies that reprogramming D to $D[x \mapsto u]$ does not affect the reverse engineered Merkle tree.

where the equality in the definition of c exploits that x is not a “challenge” query. The fact that $V(\text{inst}, c, \mathbf{m}_c)$ is satisfied for this c thus implies that $D(\text{inst}, y) \neq \perp$. Next, with the goal to reach a contradiction, assume that $u \notin \text{MTree}_D(y)$. Then for all $\perp \neq h \in \text{MTree}_D(y)$ we have that $D^{-1}(h) = D[x \mapsto u]^{-1}(h)$ except if $D(x) = h$, but this cannot be since $D(x) = \perp$. It follows that $\text{MTree}_D(y) = \text{MTree}_{D[x \mapsto u]}(y)$ and $\text{MRoot}_D^{-1}(y) = \text{MRoot}_{D[x \mapsto u]}^{-1}(y)$. The above then implies that $D \in \text{SUC}$, a contradiction.

Thus, we can bound

$$P[U \in] \leq P[\exists \text{ inst}, y : D(\text{inst}, y) \neq \perp \wedge U \in \text{MTree}_D(y)] \leq \frac{s(2\ell - 1)}{|\mathcal{Y}|} \leq \frac{q(2\ell - 1)}{|\mathcal{Y}|}. \quad (62)$$

Case 3: $D \notin \text{SUC}$, and x is a “challenge query”, i.e., $x = (\text{inst}, y) \in \mathcal{I} \times \mathcal{Y}$. Set $\mathbf{m} := \text{MRoot}_D^{-1}(y)$. Again, we have that $\perp \notin \text{SUC}|_{D|x} = \mathbf{P}'_{D|x}$, and so by Theorem 5.4 we may set $c := \mathbf{P}'_{D|x}$. Here, we can argue that

$$\begin{aligned} u \in \mathbf{L} &\iff D[x \mapsto u] \in \mathbf{P}' = \text{SUC} \\ &\implies V(\text{inst}, \gamma(u), \mathbf{m}_{\gamma(u)}) \text{ and } (\text{inst}, \mathcal{E}^*(\text{inst}, \mathbf{m})) \notin R, \end{aligned}$$

where the final implication can be seen as follows. By definition of SUC , the assumption $D[x \mapsto u] \in \text{SUC}$ implies the existence of inst' and y' with $V(\text{inst}', c, \mathbf{m}'_c)$ and $(\text{inst}', \mathcal{E}^*(\text{inst}', \mathbf{m}')) \notin R$ for

$$c := \gamma \circ D[x \mapsto u](\text{inst}', y') \quad \text{and} \quad \mathbf{m}' := \text{MRoot}_{D[x \mapsto u]}^{-1}(y') = \text{MRoot}_D^{-1}(y'),$$

where the very last equality exploits that x is not a “commit” query. With the goal to come to a contradiction, assume that $(inst', y') \neq (inst, y) = x$. Then, $c = \gamma \circ D[x \mapsto u](inst', y') = \gamma \circ D(inst', y')$, and the above then implies that $D \in \text{SUC}$, a contradiction. Thus, $(inst', y') = (inst, y) = x$. In particular, $\mathbf{m}' = \mathbf{m}$ and $c = \gamma \circ D[x \mapsto u](inst', y') = \gamma \circ D[x \mapsto u](x) = \gamma(u)$. Hence, the claimed implication holds.

Thus, we can bound

$$\begin{aligned}
 P[U \in] &\leq P[V(inst, \gamma(U), \mathbf{m}_{\gamma(U)}) \wedge (inst, \mathcal{E}^*(inst, \mathbf{m})) \notin R] \\
 &\leq P[V(inst, \gamma(U), \mathbf{m}_{\gamma(U)}) \wedge S := \{c \mid V(inst, c, \mathbf{m}_c)\} \notin \mathfrak{S}] \\
 &\leq P[\gamma(U) \in S := \{c \mid V(inst, c, \mathbf{m}_c)\} \notin \mathfrak{S}] \\
 &\leq \max_{S \notin \mathfrak{S}} P[\gamma(U) \in S] \\
 &\leq p_{triv}^{\mathfrak{S}}.
 \end{aligned} \tag{63}$$

By Theorem 5.4, we now get

$$\begin{aligned}
 \llbracket \text{SZ}_{\leq s} \setminus \text{SUC} \setminus \text{CL} \rightarrow \text{SUC} \rrbracket &\leq \max_{x, D} \sqrt{10 P[U \in \mathbb{L}^{x, D}]} \\
 &\leq \sqrt{10} \sqrt{\max \left(\frac{2\ell - 1}{|\mathcal{Y}|}, \frac{q(2\ell - 1)}{|\mathcal{Y}|}, p_{triv}^{\mathfrak{S}} \right)} \\
 &\leq \sqrt{10} \sqrt{\max (q\ell \cdot 2^{-n+1}, p_{triv}^{\mathfrak{S}})},
 \end{aligned}$$

where we have used Equations (61), (62) and (63) in the second inequality. Combining with Equations (57) and (56) yields the desired bound. $\square$

Similarly to Theorem 5.17, we now obtain the following.

Theorem 5.23. *Let Π be an $\mathfrak{S}$ -sound* Merkle-tree-based C&O protocol with challenge space $\mathcal{C}_\lambda$. Then $\text{FS}[\Pi]$ is a proof of knowledge with online extractability in the QROM (as in Definition 5.9), with $\varepsilon_{\text{sim}}(\lambda, q, n) = 0$ and*

$$\begin{aligned}
 \varepsilon_{\text{ex}}(\lambda, q, n) &\leq 2(\kappa \log \ell + 1) \cdot 2^{-n} + \left(2eq^{3/2}2^{-n/2} + q\sqrt{10 \max (q\ell \cdot 2^{-n+1}, p_{triv}^{\mathfrak{S}})} \right)^2 \\
 &\leq (22\ell \log \ell + 60) q^3 2^{-n} + 20q^2 p_{triv}^{\mathfrak{S}}
 \end{aligned}$$

where $\kappa = \kappa(\lambda) := \max_{c \in \mathcal{C}_\lambda} |c|$ and ℓ is the number of leaves of the Merkle-tree-based commitment. The running time of the extractor is dominated by running the compressed oracle, which has complexity $O(q^2) \cdot \text{poly}(n, B)$, and by computing $\text{MRoot}_D^{-1}(y)$ and running $\mathcal{E}^*$.

Here again the proof follows exactly the outline of its counterpart from Section 5.4.2, with some minor alterations to cope with the formalism of a Merkle-tree based C&O Σ -protocol. The difference in the bound is simply due to the difference between Lemmas 5.16 and 5.22.

Proof. We consider an arbitrary but fixed $\lambda \in \mathbb{N}$. Let $\mathcal{P}^*$ be a dishonest prover that, after making q queries to a random oracle H , outputs an instance *inst* and a proof $\pi = (y, \mathbf{m}_o, O)$ plus some (possibly quantum) auxiliary output Z , where O is an authentication octopus as defined in Section 5.5.1. For simplicity, we assume that $|c|$ is the same for all $c \in \mathcal{C}_\lambda$, and thus equal to κ . If it is not, we could always make the prover output a couple of dummy outputs m_i to match the upper bound on $|c|$. In the experiment $\mathcal{V}^\mathcal{E} \circ \mathcal{P}^{*\mathcal{E}}(\lambda)$, our extractor $\mathcal{E}$ works as follows while simulating all queries to H (by $\mathcal{P}^*$ and $\mathcal{V}$) with the compressed oracle:

1. Run $\mathcal{P}^*(\lambda)$ to obtain $(inst, \pi, Z)$ with $\pi = (y, \mathbf{m}_o, O)$.
2. Compute $v \leftarrow \mathcal{V}^H(inst, \pi)$, given by the truth value of

$$OctoVerify^H(c, y, \mathbf{m}_o, O) \quad \wedge \quad V(inst, c, \mathbf{m}_o) \quad \text{with} \quad c := \gamma(H(inst, y)).$$

3. Measure the internal state of the compressed oracle to obtain D .
4. Run $\mathcal{E}^*$ on input $\text{MRoot}_D^{-1}(y)$ to obtain w .

Note that in the views of both $\mathcal{P}^*$ and $\mathcal{V}$, the interaction with H and the interaction with $\mathcal{E}$ differ only in that their oracle queries are answered by a compressed oracle instead of a real random oracle in the latter case. This simulation is perfect and therefore $\varepsilon_{\text{sim}}(\lambda, q, n) = 0$.

Considering $\mathcal{P}^*$ as the algorithm $\mathcal{A}$ in Corollary 5.7, the composition $\mathcal{V} \circ \mathcal{P}^*$ then matches up with the algorithm $\tilde{\mathcal{A}}$ for $\mathcal{F} = \mathcal{V}$. Thus, noting that $\kappa(\log \ell + 1)$ is an upper bound on the amount of queries that *OctoVerify* makes,

$$\Pr[v \neq \mathcal{V}^D(inst, \pi)] \leq 2(\kappa \log \ell + 1) \cdot 2^{-n}.$$

Therefore, we can bound the figure of merit ε_{ex} as

$$\begin{aligned} \varepsilon_{\text{ex}}(\lambda, q, n) &= \Pr[v = 1 \wedge (inst, w) \notin R] \\ &\leq \Pr[\mathcal{V}^D(inst, \pi) \wedge (inst, w) \notin R] + 2(\kappa \log \ell + 1) \cdot 2^{-n} \\ &\leq \Pr[\mathcal{V}^D(inst, \pi) \wedge (inst, w) \notin R \mid D \notin \text{SUC} \cup \text{CL}] \\ &\quad + \Pr[D \in \text{SUC} \cup \text{CL}] + 2(\kappa \log \ell + 1) \cdot 2^{-n}. \end{aligned}$$

Using the definition of $\mathcal{V}^D(inst, \pi)$, understanding that $c := \gamma \circ D(inst, y)$, we can write the first term as

$$\begin{aligned} & \Pr[OctoVerify^D(c, y, \mathbf{m}_o, O) \wedge V(inst, c, \mathbf{m}_o) \wedge (inst, w) \notin R \mid D \notin \text{SUC} \cup \text{CL}] \\ & \leq \Pr[V(inst, c, \mathbf{m}_c) \text{ for } \mathbf{m} := \text{MRoot}_D^{-1}(y) \wedge (inst, w) \notin R \mid D \notin \text{SUC} \cup \text{CL}] \\ & \leq \Pr[D \in \text{SUC} \mid D \notin \text{SUC} \cup \text{CL}] \\ & = 0, \end{aligned}$$

where the first equality exploits that $D(m) = h$ iff $m = D^{-1}(h)$ for $D \notin \text{CL}$.

We may thus conclude that

$$\begin{aligned} \varepsilon_{\text{ex}}(\lambda, q, n) & \leq 2(\kappa \log \ell + 1) \cdot 2^{-n} \cdot 2^{-n} + \Pr[D \in \text{SUC} \cup \text{CL}] \\ & \leq 2(\kappa \log \ell + 1) \cdot 2^{-n} + \llbracket \perp \xRightarrow{q} \text{SUC} \cup \text{CL} \rrbracket^2, \end{aligned}$$

where the last inequality is by definition of $\llbracket \perp \xRightarrow{q} \cdot \rrbracket$. The claimed bound now follows from Lemma 5.22. $\square$

5.5.3 Discussion: Application to Picnic, and Limiting the Proof Size

Application to Picnic. A prominent use case of C&O protocols is the construction of digital signature schemes via the Fiat-Shamir transformation. An important example is Picnic [CDG+17] currently under consideration as an alternate candidate in the NIST standardization process for post-quantum cryptographic schemes. On a high level, the design of Picnic can be described as follows. A C&O Σ -protocol is constructed using the MPC-in-the-head paradigm [IKOS07a]. Then, the Fiat-Shamir transformation is applied in the usual way to obtain a digital signature scheme. There are three evolutions of Picnic: Picnic-FS, Picnic 2 and Picnic 3.¹² Picnic-FS uses plain hash-based commitments, while Picnic 2 and Picnic 3 use a Merkle-tree-based collective commitment.

All three evolutions enjoy provable post-quantum security when the hash function used for the Fiat-Shamir transformation is modeled as a (quantum-accessible) RO. The best reduction applying to all of them proceeds as follows. First, Unruh’s rewinding lemma [Unr12] is used to construct a knowledge extractor for the underlying Σ -protocol based on an appropriate $\mathfrak{S}$ -soundness notion. Then, the *generic* QROM reduction for the Fiat-Shamir transformation

¹² The original evolution also came with a variant using the Unruh transformation, Picnic-Ur.

We restrict our attention to the variants using the Fiat-Shamir transformation.

from Theorem 3.7 is used to construct a knowledge extractor for the signature scheme in the QROM from the extractor for the Σ -protocol. Finally, the technique from [GHHM21] is used for simulating the chosen-message oracle to reduce breaking NMA (no-message attack) security to breaking CMA (chosen-message attack) security. This final step connects to the previous one because for the signature scheme the witness extracted from an NMA attacker is the secret key.

The first two steps in this chain of reductions, i.e. Unruh’s rewinding and Theorem 3.7, are, however, not tight: The former loses at least a fifth power in the Picnic case, and the latter a factor of q^2 , where q is the number of random oracle queries. This means that an NMA attacker with success probability ϵ can be used to break the underlying hard problem with probability $\Omega(\epsilon^5/q^{10})$ (or worse, depending on the Picnic variant).

For Picnic-FS (only), when in addition modeling the hash function used for the commitments as a RO, Unruh’s rewinding can be replaced with the tight online extraction technique from Chapter 4. The remaining loss due to the Fiat-Shamir reduction is of order ϵ/q^2 , up to some additive terms accounting for search and collision finding in the RO, a sizable improvement over the above but still not tight.

By analyzing the Fiat-Shamir transformation of a C&O protocol (with or without Merkle tree commitments) directly, our results provide a tight alternative to the above lossy reductions. Using Theorems 5.17 (for Picnic-FS) and 5.23 (for Picnic 2 and Picnic 3) we can avoid all multiplicative/power losses in the reduction for NMA security. An NMA attacker with success probability ϵ can, in other words, be used to break the underlying hard problem with probability ϵ , up to some unavoidable additive terms accounting for search and collision finding in the RO.

An observation about octopus opening sizes. Depending on the parameters of the C&O protocol, the octopus opening information, $\text{MOcto}(c, \mathbf{m})$ can be significantly smaller than the concatenation of the individual authentication paths. On the other hand, it is also *variable in size* (namely dependent on the choice of the challenge c), and the variance can be significant (see e.g. the computations for gravity SPHINCS in [AE18]). In the context of a digital signature scheme constructed via the Fiat-Shamir transformation of a Merkle-tree-based C&O protocol, like, e.g., Picnic 2 and Picnic 3, this leads to the undesirable property of a variable signature size, where signatures can be quite a bit larger in the worst case than on average. This might, e.g., lead to problems when look-

ing for a drop-in replacement for quantum-broken digital signature schemes for use in a larger protocol, where signatures need to be stored in a data field of fixed size.

One option to mitigate this situation is to cut off the tail of the octopus size distribution, i.e. to restrict the challenge space of the Merkle-tree-based C&O protocol to the set of challenges whose octopus is not larger than some bound. This can be done before applying the Fiat-Shamir transformation, e.g. using rejection sampling. In that way, one obtains a digital signature scheme with significantly reduced worst case signature size, at the expense of a tiny security loss.

5.5.4 The Merkle-Tree-Based Unruh Transformation

The Merkle tree based commitment mechanism can replace plain random-oracle based commitments in *any* ordinary C&O protocol, in particular in $\Pi := \text{pU}[\Sigma]$ for any Σ -protocol Σ . The result is a Merkle-tree-based C&O protocol and we obtain a corollary analogous to Corollary 5.21.

Corollary 5.24. *Let Σ be an $\mathfrak{S}$ -sound Σ -protocol with challenge space size ℓ_0 . Then $\text{FS}[\text{MPpU}_r[\Sigma]]$ is online-extractable with*

$$\varepsilon_{\text{ex}} \leq (22r\ell_0 \log(r\ell_0) + 60) q^3 2^{-n} + 20q^2 \left(p_{\text{triv}}^{\mathfrak{S}} \right)^r \quad (64)$$

where $\text{MPpU}_r[\Sigma]$ is the *Merkle-tree-based, Parallel-repeated, pre-Unruh transformation* of Σ , i.e., the Merkle-tree-based C&O protocol obtained by replacing the commitments of $\text{pU}[\Pi]^r$ with a Merkle-tree-based collective commitment.