

# The Fiat–Shamir Transformation of $(\Gamma_1, \dots, \Gamma_\mu)$ -Special-Sound Interactive Proofs

Thomas Attema<sup>1,3</sup> , Serge Fehr<sup>1,2</sup>, Michael Kloos<sup>4</sup> , and Nicolas Resch<sup>5</sup> 

<sup>1</sup> CWI, Cryptology Group, Amsterdam, The Netherlands

`serge.fehr@cwi.nl`

<sup>2</sup> Leiden University, Mathematical Institute, Leiden, The Netherlands

<sup>3</sup> TNO, Applied Cryptography & Quantum Applications, The Hague, The Netherlands

`thomas.attema@tno.nl`

<sup>4</sup> Aalto University, Espoo, Finland

`michael.kloos@aalto.fi`

<sup>5</sup> University of Amsterdam, Informatics Institute, Amsterdam, The Netherlands

`n.a.resch@uva.nl`

December 22, 2023

**Abstract.** The Fiat-Shamir transformation is a general principle to turn any public-coin interactive proof into non-interactive one (with security then typically analyzed in the random oracle model). While initially used for 3-round protocols, many recent constructions use it for *multi-round* protocols. However, in general the soundness error of the Fiat-Shamir transformed protocol degrades exponentially in the number of rounds. On the positive side, it was shown that for the special class of  $(k_1, \dots, k_\mu)$ -special-sound protocols the loss is actually only *linear* in the number of random oracle queries, and *independent* of the number of rounds, which is optimal.

A natural next question is whether this positive result extends to the Fiat-Shamir transformation of so-called  $(\Gamma_1, \dots, \Gamma_\mu)$ -special-sound protocols, a notion recently defined and analyzed in the interactive case, with the aim to capture the most general notion of special-soundness.

We show in this work that this is indeed the case. Concretely, we show that the Fiat–Shamir transformation of any  $(\Gamma_1, \dots, \Gamma_\mu)$ -special-sound interactive proof is knowledge sound under the same condition under which the original interactive proof is knowledge sound. Furthermore, also here the loss is linear in the number of random-oracle queries and independent of the number of rounds.

In light of the above, one might suspect that our argument follows as a straightforward combination of the above mentioned prior works. However, this is not the case. The approach used for  $(k_1, \dots, k_\mu)$ -special-sound protocols, which is based on an extractor that samples without replacement, does not (seem to) generalize; on the other hand, the other approach, which uses an extractor based on sampling with replacement, comes with an additional loss that would blow up in the recursive multi-round analysis.

## 1 Introduction

### 1.1 Interactive Proofs and Special Sound Protocols

Interactive proofs play an important role in modern cryptography (and well beyond). They allow a prover  $\mathcal{P}$  to convince a verifier  $\mathcal{V}$  of the truth of a statement, i.e., that a certain instance  $x$  is an element of a given language  $L$ , without revealing any additional information (*zero-knowledge*), or with a proof that is (much) smaller than the statement (*succinctness*).

A crucial property of an interactive proof is *soundness*, which means that no dishonest prover  $\mathcal{P}^*$  can convince the verifier  $\mathcal{V}$  of a false statement (except with small probability), i.e., make  $\mathcal{V}$  accept if  $x \notin L$ . However, in many situations a stronger property is needed: *knowledge soundness*, which informally means that in order to make  $\mathcal{V}$  accept, not only has  $x$  to be in  $L$ , the prover actually needs to know a witness  $w$  attesting that  $x \in L$  (considering  $L$  to be an NP language). One then also speaks of a *proof of knowledge*.

More formally, knowledge soundness requires the existence of a *knowledge extractor*, i.e., an efficient algorithm that outputs a witness when given rewindable oracle access to a (possibly dishonest) prover  $\mathcal{P}^*$  that convinces a verifier  $\mathcal{V}$  to accept with sufficiently high probability. In certain cases, ordinary soundness could be meaningless; this happens, e.g., when the considered language  $L$  is trivial, i.e., when every

instance admits a witness. In such cases, knowledge soundness is the only meaningful soundness notion. For example, consider a prover claiming to know a hash collision for a hash function  $H : \{0, 1\}^* \rightarrow \{0, 1\}^m$ : here, the statement that  $H$  has a hash collision is vacuously true, and the nontrivial requirement is that the prover *actually knows* inputs  $x \neq x'$  for which  $H(x) = H(x')$ .

Proving knowledge soundness is tricky in general; typically much harder than proving ordinary soundness. The problem is that different interactive proofs may require different extraction strategies. Thus, for a given protocol it is oftentimes unclear how to design a successful knowledge extractor, and/or how to analyze it. Furthermore, the formal definition sets a rather stringent requirement on the success probability of the extractor, which needs to hold for any dishonest prover. It is thus desirable to identify classes of interactive proofs for which there exist generic knowledge extraction results. An important example are *special-sound*  $\Sigma$ -protocols: classic results guarantee they are knowledge sound (with a soundness error determined by the size of the challenge set).

While interactive proofs studied in the past tended to be special-sound  $\Sigma$ -protocols, this is not the case anymore for many of the more recently proposed protocols, like Bulletproofs [BCC<sup>+</sup>16, BBB<sup>+</sup>18]. Motivated by this, the original result on special-sound  $\Sigma$ -protocols was extended in [ACK21, AF22, Att23] to  $k$ -special-sound  $\Sigma$ -protocols and to their multi-round variants,  $(k_1, \dots, k_\mu)$ -special-sound protocols. However, many modern interactive proofs [BBHR18, RR22, ACY23] – especially those employing Merkle-tree commitments – do not satisfy  $(k_1, \dots, k_\mu)$ -special-soundness for reasonable parameters (in particular, the implied knowledge soundness would be far from optimal). Motivated by this, in the recent work [AFR23], the generalization of special-soundness is pushed further, and in some sense to the extreme, by introducing the notion of  $(\Gamma_1, \dots, \Gamma_\mu)$ -special-sound interactive proofs. In spirit, while for  $k$ -special-sound  $\Sigma$ -protocols the special-soundness is specified by a *threshold*  $k$ , it is specified by a *general access structure*  $\Gamma$  in the case of  $\Gamma$ -special-sound  $\Sigma$ -protocols, and correspondingly for the multi-round variants. Besides introducing the definition, [AFR23] proves knowledge soundness and strong parallel repetition of (certain)  $(\Gamma_1, \dots, \Gamma_\mu)$ -special-sound interactive proofs. One cannot expect efficient extractability to hold for all choices of the access structures  $\mathbf{\Gamma} = (\Gamma_1, \dots, \Gamma_\mu)$ ; thus, in more detail, [AFR23] identifies two relevant parameters  $\kappa_{\mathbf{\Gamma}}$  and  $t_{\mathbf{\Gamma}}$ , determined solely by  $\mathbf{\Gamma}$ , and prove knowledge soundness with knowledge error  $\kappa_{\mathbf{\Gamma}}$  for any  $\mathbf{\Gamma}$ -special-sound interactive proof for which  $t_{\mathbf{\Gamma}}$  is polynomial.

## 1.2 The Fiat–Shamir Transformation

The *Fiat–Shamir transformation* is a powerful technique for turning (certain) *interactive* proofs, as discussed above, into *non-interactive* ones, or for designing signature schemes. Although originally suggested for  $\Sigma$ -protocols, it has become popular to apply the Fiat–Shamir transformation also to *multi-round* public-coin interactive proofs. Indeed, there has been a recent focus on interactive proofs with succinct communication, and those tend to have a non-constant number of rounds; the Fiat–Shamir transformation is then often used to avoid the increased round complexity (which could then actually form the bottle neck) by making the proof entirely non-interactive. Furthermore, for certain applications, it is crucial that a proof be non-interactive, making the use of the Fiat–Shamir transformation necessary.

However, the Fiat–Shamir transformation is not free. First of all, the security of the Fiat–Shamir transformed scheme is typically “only” proven in the random oracle model. Furthermore, there is often a non-trivial security loss involved. Indeed, until recently, the best reduction had a security loss on the knowledge error that is exponential in the number of rounds. More precisely, if the interactive proof  $\Pi$  has a knowledge error  $\kappa$  then the knowledge error of the Fiat–Shamir transformed non-interactive proof  $\text{FS}[\Pi]$  can be as bad as (roughly)  $Q^\mu \cdot \kappa$ , where  $Q$  is the number of oracle queries performed by the attacker, and  $\mu$  the number of challenge rounds of the interactive proof  $\Pi$ . Only in the recent work [AFK22], it was shown that the class of  $(k_1, \dots, k_\mu)$ -special-sound interactive proofs avoid this exponential security loss under the Fiat–Shamir transformation. As a matter of fact, the knowledge error of the Fiat–Shamir transformation of such a  $(k_1, \dots, k_\mu)$ -special-sound protocol was shown to be  $(Q + 1) \cdot \kappa$ , independent of the number of rounds, where  $\kappa$  denotes the knowledge error of the interactive proof.

Unfortunately, while this notion of special-soundness covers certain interactive proofs with succinct communication, e.g., Bulletproofs, there are many interesting examples that fall outside this class. A notable example is a standard amortization technique, where a prover proves knowledge of  $n$  homomorphism preimages by proving knowledge of a random linear combination of these preimages. This amortization technique is used in many (lattice-based) interactive proofs to reduce the communication complexity,

e.g., [BBC<sup>+</sup>18, ALS20, ENS20, LNP22]. Unfortunately, it renders the corresponding threshold special-soundness parameter  $k$  too large to provide reasonable security guarantees. However, the amortization technique is  $\Gamma$ -special-sound for a  $\Gamma$  with small parameters. Second, a common design principle is to first design an *interactive oracle proof* (IOP) with succinct communication, and to then instantiate the oracle with a hash-based Merkle-tree commitment. By construction, an interactive proof obtained via this recipe is not  $(k_1, \dots, k_\mu)$ -special-sound, at least not for reasonable parameters  $k_i$ , whereas it *can* be cast as a  $(\Gamma_1, \dots, \Gamma_\mu)$ -special-sound protocol with reasonable parameters.<sup>6</sup> As a final example, we note that the  $t$ -fold parallel repetition  $\Pi^t$  of a  $k$ -special-sound  $\Sigma$ -protocol  $\Pi$  is  $((k-1)^t + 1)$ -special-sound, i.e., the threshold special-soundness parameter grows exponentially in  $t$ , and thus this property does not imply knowledge soundness. However,  $\Pi^t$  is also  $\Gamma$ -special-sound with parameter  $t_\Gamma$  growing only linearly in  $t$ , i.e., the  $\Gamma$ -special-soundness property immediately implies knowledge soundness.

Given the current situation on the Fiat–Shamir front, the natural question that we address in this work is whether the techniques and the results from [AFK22] on the Fiat–Shamir transformation of  $(k_1, \dots, k_\mu)$ -special-sound protocols extend to this general notion of  $(\Gamma_1, \dots, \Gamma_\mu)$ -special-soundness introduced in [AFR23].

The short answer is: the results generalize, but not the techniques; the long answer follows below and in the rest of the paper.

### 1.3 Our Results

Indeed, in this paper, we show that the results of [AFK22], which show the security of the Fiat–Shamir transform of multi-round  $(k_1, \dots, k_\mu)$ -special-sound protocols with a security loss independent of the number of rounds, carry over to the generalization of special-sound (multi-round) protocols considered in [AFR23]. Thus, in detail, we show that the Fiat–Shamir transformation  $\text{FS}[\Pi]$  of any  $(\Gamma_1, \dots, \Gamma_\mu)$ -special-sound interactive proof  $\Pi$ , where  $\Gamma := (\Gamma_1, \dots, \Gamma_\mu)$  is such that the parameter  $t_\Gamma$  is polynomial, is knowledge sound with knowledge error

$$(Q + 1) \cdot \kappa_\Gamma, \tag{1}$$

where  $Q$  is the number of random oracle calls made by the attacker, and  $\kappa_\Gamma$  is the knowledge error of the interactive proof  $\Pi$ .<sup>7</sup> We stress the independence of the bound on  $\mu$ , the number of challenge rounds of the protocol.

One might hope that this result can be obtained by generalizing the reasoning of [AFK22] to this generalized notion of  $(\Gamma_1, \dots, \Gamma_\mu)$ -special-soundness. But this is not the case. The extractor considered in [AFK22] is based on sampling *without replacement*, and then properties of the *negative hypergeometric distribution* are used for the analysis; however, this use of the negative hypergeometric distribution is tailored to the threshold case and does not (seem to) generalize. Thus, a new extraction strategy is needed.

Fortunately, the extraction strategy from [AF22, AFR23]—which was so far considered in the *interactive* setting—comes to the rescue. These extractors work by sampling *with replacement*, together with an appropriately chosen, randomized stopping criteria, and using properties of the *geometric distribution* to analyze the success probability and the run time.

In more detail, as a first step, we show that for the 3-round case (i.e., for  $\Sigma$ -protocols), we can actually reduce the analysis of the considered extractor for the non-interactive Fiat–Shamir transformed protocol  $\text{FS}[\Pi]$  to that of (some variant of) the underlying interactive proof  $\Pi$ . Simply relying on the extraction properties shown in [AF22, AFR23] for the interactive proof  $\Pi$ , would then settle the 3-round case, showing that the Fiat–Shamir transformation of a  $\Gamma$ -special-sound  $\Sigma$ -protocol is a proof of knowledge with knowledge soundness  $\kappa_\Gamma$  if the parameter  $t_\Gamma$  is polynomial.

However, the multi-round case turns out to be more subtle. A naive approach would be to follow the same strategy as in the 3-round case, and reduce the non-interactive analysis to the analysis of a multi-round interactive proof. As we will discuss in Section 6.2, this approach is too restrictive and will result

<sup>6</sup> We note that modeling the hash function used in the Merkle tree commitment as a random oracle allows for so-called *straight-line extraction*. Hence, in this model rewinding can be avoided, which may result in a more efficient knowledge extractor. In fact, straight-line extraction is the standard approach for analyzing IOP-based proof systems. The notion  $(\Gamma_1, \dots, \Gamma_\mu)$ -special-soundness, together with its knowledge extractors, provides an alternative for analyzing their knowledge soundness.

<sup>7</sup> The definition of  $\kappa_\Gamma$  is provided in (14).

in an exponential security loss. For this reason, our approach is to follow a similar recursive strategy as in [AFK22], which uses the 3-round extractor and applies it recursively over the different rounds to extract the required tree of transcripts, using an “early-aborting” trick to achieve the round-independent expected runtime. Even though we can use a similar recursive approach with the same early-abort trick, we encountered multiple barriers in the analysis that we had to solve.

One problem of using the extraction strategy from [AF22, AFR23] for the 3-round case is that the success probability of the extractor suffers a factor  $t_\Gamma$ -loss, compared to the 3-round extractor used in [AFK22]. Since we anyway have to require this parameter to be polynomial for the extractor to be efficient, this is fine for the 3-round case; however, in the recursive analysis from [AFK22] of the extractor’s success probability, the loss  $t_\Gamma$  negatively impacts the knowledge error. More precisely, this approach would introduce an additional factor  $t_\Gamma = \prod_{i=1}^\mu t_{\Gamma_i}$  loss in the knowledge error. Since  $t_\Gamma$  is required to be polynomial, this would actually be sufficient for proving knowledge soundness. However, in this work we aim to derive a *tight* knowledge error that is essentially equal to the trivial cheating probability of a dishonest prover.

Our solution is to apply the above recursive approach to a different base extractor for 3-round  $\Gamma$ -special-sound interactive proofs [KL23]. This base extractor is an improved variant of the ones presented in [AF22, AFR23], which does not suffer the factor  $t_\Gamma$  loss. Since the base extractor of [KL23] has not been published yet (it was communicated via personal communication), we provide a detailed description (Section 3.4) and complete analysis (Appendix B). This approach ensures the desired success probability of our knowledge extractor for the Fiat–Shamir transformation of  $(\Gamma_1, \dots, \Gamma_\mu)$ -special sound interactive proof  $\Pi$ .

To additionally achieve the required expected polynomial running time, we must however refine the running time analysis of this base extractor. This refinement is inspired by [AFK22], but requires new analysis techniques. As before, the difference is caused by the fact that our extractor uses sampling *with* replacement, whereas [AFK22] use sampling *without* replacement. The refined analysis allows us to apply the same kind of early-abort strategy as used in [AFK22], which significantly reduces the expected running time without compromising the extractor’s success probability. Altogether, this proves our main result: the Fiat–Shamir transformation of  $(\Gamma_1, \dots, \Gamma_\mu)$ -special sound interactive proof  $\Pi$  is knowledge sound with knowledge error  $(Q + 1) \cdot \kappa_\Gamma$ , if the parameter  $t_\Gamma$  is polynomial.<sup>8</sup>

## 1.4 Related Work

We already discussed the related work [AF22, AFK22, AFR23], and their relevance and relation to this work, above.

Another closely related work is [Wik18], in which Wikström introduces and studies an abstraction of the problem of knowledge extraction when considering a generalized notion of special-soundness. His generalization resembles that of  $(\Gamma_1, \dots, \Gamma_\mu)$ -special-soundness studied in [AFR23], except that the access structure  $\Gamma_i$  is restricted to be the set of bases of a matroid  $\mathbb{M}_i$ . The main technical result [Wik18, Theorem 1] is the existence of a so-called *accepting basis extractor* for a *matroid tree*, with bounds on (1) the *extraction error* (which corresponds to the knowledge error), (2) the expected number of queries, and (3) the *tail bound*.

A central goal of [Wik18] is establishing good tail bounds (instead of only expected query bounds) with a precise analysis. The presented bounds are “convoluted expressions” [Wik18, p. 12] (and depend on tweakable parameters of the extractor). While an interpretation of the bounds is given, no simplifications are provided. This, together with the non-standard language and formalism used in [Wik18], makes a comparison of [Wik18] and [AFR23] that goes beyond the high-level similarities challenging.

In [Wik21], Wikström extends the language and results of [Wik18] to the Fiat–Shamir transformation. This amplifies the difficulties in comparing our work (which extends [AFR23]) with [Wik21] (which extends [Wik18]). In fact, we argue that the relative simplicity of our language, definitions, and theorem claims is a major differentiation. For completeness, we mention that [Wik21, Sect. 7] claims similar result as ours, i.e., security loss linear in the number of queries.

<sup>8</sup> Our main result (Theorem 4) is expressed in terms of  $T_\Gamma := \prod_{i=1}^\mu (t_{\Gamma_i} + 1)$  instead of  $t_\Gamma$ . This is merely for notational convenience, since  $t_\Gamma$  is polynomial if and only if  $T_\Gamma$  (for nontrivial  $t_{\Gamma_i} > 1$ ).

## 2 Preliminaries: Interactive and Non-Interactive Proofs

We recall here some standard definitions and concepts related to interactive and non-interactive proofs.

### 2.1 Interactive Proofs

We begin by introducing the necessary terminology for our discussion of interactive proofs.

Let  $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$  be a binary relation, which we view as a set of statement-witness pairs  $(x; w)$ . We let  $R(x) = \{w : (x; w) \in R\}$  denote the set of valid witnesses for a statement  $x$ . Throughout, all relations are NP-relations, i.e., verifying  $(x; w) \in R$  on input  $(x; w)$  can be done in time polynomial in  $|x|$  (thus, without loss of generality we also have  $|w|$  polynomial in  $|x|$ ). An *interactive proof* is an (interactive) protocol wherein a prover attempts to convince a verifier that a public statement  $x$  admits a (secret) witness  $w \in R(x)$ , or even that the prover *knows* such a witness.

If the verifier publishes all its random coins, the protocol is called *public-coin*. In such a case we may assume without loss of generality that all the messages from verifier to prover are uniformly random elements from some finite challenge set  $\mathcal{C}$ . The special case in which a public-coin interactive proof consists of 3 communication rounds, and in which the prover speaks in the first and third rounds, is termed a  $\Sigma$ -*protocol*.

An interactive proof is *complete* if, on public input  $x$  and private prover input  $w \in R(x)$ , the protocol execution will result in an accepting transcript (with high probability; in many protocols this probability is in fact 1). An interactive proof is *sound* if, on public input  $x$  that does not admit a witness  $w$  (i.e.,  $R(x) = \emptyset$ ), even for a potentially cheating prover  $\mathcal{P}^*$  the probability that the protocol transcript is rejecting is large. The stronger notion of *knowledge soundness* informally requires that, on public input  $x$ , if a potentially cheating prover  $\mathcal{P}^*$  manages to convince the verifier to accept with high enough probability, then it in fact must “know” a witness  $w \in R(x)$ . This is formalized in terms of a *knowledge extractor*, which is an expected polynomial-time algorithm that is able to extract a witness given blackbox access to such a prover  $\mathcal{P}^*$ . This is the main property of interactive proofs that we study in this work, so we provide a precise definition.

**Definition 1 (Knowledge Soundness).** *An interactive proof  $\Pi = (\mathcal{P}, \mathcal{V})$  is knowledge-sound with knowledge error  $\kappa : \mathbb{N} \rightarrow [0, 1]$  if there exists a positive polynomial  $p$  and an algorithm  $\mathcal{E}$ , called a knowledge extractor, with the following properties. Given input  $x$  and blackbox access to  $\mathcal{P}^*$ ,  $\mathcal{E}$  runs in expected polynomial time (counting queries to  $\mathcal{P}^*$  as one time-step) and outputs a witness  $w \in R(x)$  with the following probability:*

$$\Pr \left( \mathcal{E}^{\mathcal{P}^*}(x) \in R(x) \right) \geq \frac{\epsilon(\mathcal{P}^*, x) - \kappa(|x|)}{p(|x|)} .$$

In the above,  $\epsilon(\mathcal{P}^*, x) = \Pr((\mathcal{P}^*, \mathcal{V})(x) = \text{accept})$  is the success probability of  $\mathcal{P}^*$  on input  $x$ .

*Black-box access* means in particular that  $\mathcal{E}$  can run  $\mathcal{P}^*$  multiple times. In case of a randomized  $\mathcal{P}^*$ ,  $\mathcal{E}$  can rerun  $\mathcal{P}^*$  with the same randomness as in the previous run. This is referred to as *rewinding*.

*Remark 1 (Interactive Arguments).* In cryptography, one is often concerned with interactive *arguments*, i.e., interactive proofs where the soundness only holds against computationally bounded adversaries. In particular, computationally unbounded provers could convince the verifier with high probability. It thus may appear that our study of interactive proofs is not relevant for this setting. However, in practice most interactive arguments can in fact be cast as interactive proofs for the following “or” relation:

$$R' = \{(x; w) : (x; w) \in R \text{ or } w \text{ solves computational problem } X\} .$$

That is, knowledge soundness in this case guarantees that a successful prover either knows a witness for the given instance, or that it can solve some presumably hard computational problem (e.g., it outputs a discrete logarithm, a factor of an RSA modulus, a hash collision, etc.). Therefore, knowledge extractors for interactive proofs can typically be repurposed to prove knowledge soundness for interactive arguments as well.

## 2.2 Non-Interactive Random Oracle Proofs (NIROPs)

In certain applications, it is essential for proofs to be non-interactive. Certain interactive proofs can be made non-interactive via the Fiat–Shamir transformation (see below); the resulting non-interactive proof is then typically analyzed in the *random oracle model* (ROM). Proofs in the ROM are also referred to as non-interactive *random oracle proofs* (NIROPs). Below, we briefly recall the ROM, the formal definition of NIROPs and knowledge-soundness for NIROPs.

In the *random oracle model* (ROM), all algorithms have black-box access to an oracle  $\text{ro}: \{0, 1\}^* \rightarrow \mathcal{Y}$ , called the *random oracle*, which is instantiated as a uniformly random function. Typically  $\mathcal{Y} = \{0, 1\}^\eta$  for some  $\eta \in \mathbb{N}$  related to the security parameter, but we find it more convenient to not make any restrictions on the codomain of the random oracle. For technical reasons, we limit the domain of the random oracle to be some arbitrary but finite message set  $\mathcal{M}$ .

One can naturally extend the ROM to allow  $\mathcal{A}$  to access multiple independent random oracles  $\text{ro}_1, \dots, \text{ro}_\mu$ , possibly with different codomains. This can be implemented from a single random oracle using standard techniques for domain separation and for any sampling algorithms, which admit reverse sampling or explainable sampling [LW22]. The following definitions and discussion may be extended to this multiple random oracle model in a natural way.

**Definition 2 (Non-Interactive Random Oracle Proof (NIROP)).** *A non-interactive random oracle proof (NIROP) for a relation  $R$  is a pair  $(\mathcal{P}, \mathcal{V})$  of (probabilistic) random-oracle algorithms, a prover  $\mathcal{P}$  and a polynomial-time verifier  $\mathcal{V}$ , such that the following holds. Given  $(x; w) \in R$  and access to a random oracle  $\text{ro}$ , the prover  $\mathcal{P}^{\text{ro}}(x; w)$  outputs a proof  $\pi$ . Given  $x \in \{0, 1\}^*$ , a purported proof  $\pi$ , and access to a random oracle  $\text{ro}$ , the verifier  $\mathcal{V}^{\text{ro}}(x, \pi)$  outputs 0 to reject or 1 to accept the proof.*

The definition of completeness is the natural one: honestly generated proofs indeed convince the verifier (with high probability). The basic stipulation of soundness is that it is infeasible for a prover to convince a verifier that a false statement is true, except with some small probability. In the non-interactive setting, the soundness error will depend on the number of queries that the cheating prover is permitted to make to the random oracle. The concept of knowledge soundness for NIROP’s is the natural analogue of knowledge soundness for interactive proofs (Definition 1). We now formally define it.

**Definition 3 (Knowledge Soundness - NIROP).** *A non-interactive random oracle proof  $(\mathcal{P}, \mathcal{V})$  for a relation  $R$  is knowledge sound with knowledge error  $\kappa: \mathbb{N} \times \mathbb{N} \rightarrow [0, 1]$  if there exists a positive polynomial  $p$  and an algorithm  $\mathcal{E}$  – called a knowledge extractor – with the following properties. The extractor, given input  $x$  and oracle access to any (potentially dishonest)  $Q$ -query random oracle prover  $\mathcal{P}^*$ , runs in an expected number of steps that is polynomial in  $|x|$  and  $Q$  and outputs a witness  $w$  such that, for all  $x \in \{0, 1\}^*$ ,*

$$\Pr(w \in R(x) : w \leftarrow \mathcal{E}^{\mathcal{P}^*}(x)) \geq \frac{\epsilon(\mathcal{P}^*, x) - \kappa(|x|, Q)}{p(|x|)}$$

where  $\epsilon(\mathcal{P}^*, x) = \Pr(\mathcal{V}^{\text{ro}}(x, \mathcal{P}^{*, \text{ro}}) = 1)$ . Here,  $\mathcal{E}$  implements  $\text{ro}$  for  $\mathcal{P}^*$ : in particular, it may arbitrarily program  $\text{ro}$ . Moreover, the randomness is over the randomness of  $\mathcal{E}$ ,  $\mathcal{V}$ ,  $\mathcal{P}^*$  and  $\text{ro}$ .

## 2.3 Fiat–Shamir Transformation

The Fiat–Shamir Transformation is a general-purpose operation that converts a given public-coin interactive proof into a non-interactive random oracle proof (NIROP). The idea is to replace the challenges from the verifier (which, recall, are without loss of generality uniformly random bit strings) by hashes of (some part of) the transcript up until that point. For concreteness, given a  $\Sigma$ -protocol with first message  $a$  and challenge  $c$ , the Fiat–Shamir transformed NIROP either sets  $c = \text{ro}(a)$  or  $c = \text{ro}(x, a)$ , where  $x$  is the public input. The former definition is sufficient for *static* security where a malicious prover is given the input and then must attempt to convince the verifier, while the latter is required for *adaptive* security [AFK23, Definition 10] where the malicious prover may first choose the input  $x$  and subsequently attempt to forge a false proof [BPW12].

For multi-round protocols, there are multiple variants that one could consider. In this work, we will focus on the most conservative choice where all the prior prover messages are hashed along with the current round’s message, i.e., the  $i$ -th challenge is computed as

$$c_i = \text{ro}_i(a_1, \dots, a_{i-1}, a_i)$$

where  $a_1, \dots, a_{i-1}, a_i$  are the  $i$  messages sent from the prover to the verifier so far, and  $\text{ro}_i$  is a random oracle with suitable codomain  $\mathcal{C}_i$ . Again in this multi-round setting the Fiat–Shamir transform comes with a statically secure and an adaptively secure variant: in the latter case, the statement  $x$  is included in the input for each hash function evaluation. As our results apply equally well to each variant, we will not explicitly spell out this distinction (static versus adaptive) in the rest of this work. A formal definition of the Fiat–Shamir transformation is included in Appendix A.

## 2.4 Geometric Distribution

Finally, as in prior works [AF22, AFR23], our extractor analysis relies on certain facts about the geometric distribution, which we now quickly recall.

A random variable  $B$  with two possible outcomes, denoted 0 (failure) and 1 (success), is said to follow a Bernoulli distribution with parameter  $p$  if  $p = \Pr(B = 1)$ . Sampling from a Bernoulli distribution is also referred to as running a Bernoulli trial. The probability distribution of the number  $X$  of independent and identical Bernoulli trials needed to obtain a success is called the geometric distribution with parameter  $p = \Pr(B = 1)$ . In this case,  $\Pr(X = k) = (1 - p)^{k-1}p$  for all  $k \in \mathbb{N}$  and we write  $X \sim \text{Geo}(p)$ . For two independent geometric distributions we have the following lemma.

**Lemma 1** ([AFR23, Lemma 1]). *Let  $X \sim \text{Geo}(p)$  and  $Y \sim \text{Geo}(q)$  be independently distributed. Then,*

$$\Pr(X \leq Y) = \frac{p}{p + q - pq}.$$

The following simple argument shows that, if in a geometric experiment each trial is associated with a cost whose expected value has a constant upper bound, then the expected cost of the experiment is upper bounded by the expected number of trials times the upper bound of the cost for each trial.

**Lemma 2.** *Let  $X \sim \text{Geo}(p)$ , let  $Z_i$  denote the cost of the  $i$ -th Bernoulli trial in the geometric experiment, and let  $Z = Z_1 + Z_2 + \dots + Z_X$  denote the total cost of running the geometric experiment. Then, if  $\mathbb{E}[Z_i \mid X \geq i] \leq \theta$  for all  $i$ , it holds that*

$$\mathbb{E}[Z] \leq \frac{\theta}{p}.$$

*Proof.* Without loss of generality, we may assume that  $\mathbb{E}[Z_i \mid X < i] = 0$ . Then,

$$Z = \sum_{i=1}^X Z_i = \sum_{i=1}^{\infty} Z_i,$$

and

$$\begin{aligned} \mathbb{E}[Z_i] &= \Pr(X < i) \cdot \mathbb{E}[Z_i \mid X < i] + \Pr(X \geq i) \cdot \mathbb{E}[Z_i \mid X \geq i] \\ &\leq \Pr(X < i) \cdot 0 + \Pr(X \geq i) \cdot \theta = (1 - p)^{i-1} \theta. \end{aligned}$$

Hence,

$$\mathbb{E}[Z] = \sum_{i=1}^{\infty} \mathbb{E}[Z_i] \leq \sum_{i=1}^{\infty} (1 - p)^{i-1} \theta = \frac{\theta}{p},$$

which completes the proof.  $\square$

## 3 Preliminaries: $\Gamma$ -Special-Sound Protocols

The standard concept of *special-soundness* for  $\Sigma$ -protocols [Cra96, CD98] has recently seen many generalizations. Firstly, a  $\Sigma$ -protocol is  $k$ -*special-sound* for  $k \in \mathbb{N}$  if one can efficiently construct a witness given  $k$  accepting transcripts with the same first message but pairwise distinct second messages (setting  $k = 2$  recovers standard special-soundness). There are also generalizations to multi-round public-coin interactive proofs; in this case, an efficient procedure for constructing a witness given an appropriate *tree* of accepting transcripts. Very recently, a very general notion of  $\Gamma$ -special-soundness (and  $(\Gamma_1, \dots, \Gamma_\mu)$ -special-soundness in the multi-round case) was defined and shown to imply knowledge soundness [AFR23] under certain conditions. In our work, we will study the knowledge-soundness of non-interactive proofs in the random oracle model obtained by applying the Fiat–Shamir transform to such  $(\Gamma_1, \dots, \Gamma_\mu)$ -special-sound protocols. We provide the precise definition below.

### 3.1 $\Gamma$ -out-of- $\mathcal{C}$ Special-Soundness

Firstly, we must recall the definition of *monotone structures*.

**Definition 4 (Monotone Structure).** Let  $\mathcal{C}$  be a nonempty finite set and let  $\Gamma \subseteq 2^{\mathcal{C}}$  be a family of subsets of  $\mathcal{C}$ . Then,  $(\Gamma, \mathcal{C})$ , or just  $\Gamma$ , is said to be a monotone structure if it is closed under taking supersets. That is,  $S \in \Gamma$  and  $S \subseteq T \subseteq \mathcal{C}$  implies  $T \in \Gamma$ .

Note that  $\emptyset \subseteq 2^{\mathcal{C}}$  and  $2^{\mathcal{C}}$  are monotone structures according to our definition (which differs from some textbook definitions). We now provide the definition of  $\Gamma$ -out-of- $\mathcal{C}$  special-soundness.

**Definition 5 ( $\Gamma$ -out-of- $\mathcal{C}$  Special-Soundness).** Let  $(\Gamma, \mathcal{C})$  be a monotone structure. A 3-round public-coin interactive proof  $(\mathcal{P}, \mathcal{V})$  for a relation  $R$ , with challenge set  $\mathcal{C}$ , is  $\Gamma$ -out-of- $\mathcal{C}$  special sound if there exists an algorithm that, on input a statement  $x$  and a set of accepting transcripts  $(a, c_1, z_1), \dots, (a, c_k, z_k)$  with common first message  $a$  and such that  $\{c_1, \dots, c_k\} \in \Gamma$ , runs in polynomial time and outputs a witness  $w \in R(x)$ . We also say  $(\mathcal{P}, \mathcal{V})$  is  $\Gamma$ -special-sound.

Note that this definition recovers  $k$ -special soundness by taking  $\Gamma$  to be the family of subsets of  $\mathcal{C}$  of size at least  $k$ .

*Remark 2.* Technically, the monotone structure  $(\Gamma, \mathcal{C})$  of Definition 5 may depend on the input  $x$ . We should therefore refer to a family  $(\Gamma_x, \mathcal{C}_x)_{x \in \{0,1\}^*}$  of monotone structures. For ease of notation, we will not make the dependency explicit and simply write  $(\Gamma, \mathcal{C})$ .

### 3.2 Some Concepts Related to $\Gamma$ -Special-Sound Protocols

In this subsection we introduce some of the concepts and technical tools developed in [AFR23] which are used to analyze  $(\Gamma, \mathcal{C})$ -special sound protocols. We refer to this paper – particularly Section 4 – for additional context and motivation for these ideas.

Firstly, we require the concept of useful elements, which informally are elements that “bring us closer” to finding a set of challenges for which a (potentially cheating) prover  $\mathcal{P}^*$  succeeds. That is, if the set of challenges the extractor has currently found is  $S \subseteq \mathcal{C}$ ,  $\mathcal{U}_\Gamma(S)$  is the set of challenges that could be useful in its quest to find a set of accepting transcripts  $(a, c_1, z_1), \dots, (a, c_k, z_k)$  with  $\{c_1, \dots, c_k\} \in \Gamma$ .

**Definition 6 (Useful Elements).** For a monotone structure  $(\Gamma, \mathcal{C})$  we define the following function:

$$\mathcal{U}_\Gamma: 2^{\mathcal{C}} \rightarrow 2^{\mathcal{C}}, \quad S \mapsto \{c \in \mathcal{C} \setminus S : \exists A \in \Gamma \text{ s.t. } S \subset A \wedge A \setminus \{c\} \notin \Gamma\}.$$

It is easily seen that

$$\mathcal{U}_\Gamma(B) \subseteq \mathcal{U}_\Gamma(A) \quad \text{for all } A \subseteq B. \quad (2)$$

Further, [AFR23, Lemma 3] shows that

$$\mathcal{C} \setminus \mathcal{U}_\Gamma(S) \notin \Gamma \quad \text{for all } S \notin \Gamma. \quad (3)$$

The efficiency of the knowledge extractor depends on how long it could take to find enough useful challenges. This is formalized by the  $t$ -value.

**Definition 7 ( $t$ -value).** Let  $(\Gamma, \mathcal{C})$  be a monotone structure and  $S \subseteq \mathcal{C}$ . Then

$$t_\Gamma(S) := \max \left\{ t \in \mathbb{N}_0 : \begin{array}{l} \exists c_1, \dots, c_t \in \mathcal{C} \text{ s.t.} \\ c_i \in \mathcal{U}_\Gamma(S \cup \{c_1, \dots, c_{i-1}\}) \quad \forall i \end{array} \right\}.$$

Further,

$$t_\Gamma := t_\Gamma(\emptyset).$$

Observe that  $t_\Gamma(S) = 0$  if and only if  $S \in \Gamma$  or  $\Gamma = \emptyset$ . The basic fact that we require is that adding an element  $c \in \mathcal{U}_\Gamma(S)$  to  $S$  decreases the  $t$ -value.

**Lemma 3 ([AFR23, Lemma 4]).** Let  $(\Gamma, \mathcal{C})$  be a nonempty monotone structure and let  $S \subseteq \mathcal{C}$  such that  $S \notin \Gamma$ . Then, for all  $c \in \mathcal{U}_\Gamma(S)$ ,

$$t_\Gamma(S \cup \{c\}) < t_\Gamma(S).$$



### 3.3 Knowledge Soundness of $\Gamma$ -Special Sound $\Sigma$ -Protocols

For any  $\Gamma$ -special sound  $\Sigma$ -protocol  $\Pi$ , Theorem 1 of [AFR23] proves the existence of a knowledge extractor that makes an expected number of at most  $2t_\Gamma - 1$  queries to the considered prover and successfully extracts a witness with probability at least

$$\frac{1}{t_\Gamma} \cdot \frac{\epsilon - \kappa_\Gamma}{1 - \kappa_\Gamma}$$

where  $\epsilon$  is the success probability of the considered prover, and

$$\kappa_\Gamma := \max_{S \notin \Gamma} \frac{|S|}{|\mathcal{C}|}. \quad (4)$$

Looking ahead, we will reduce the existence of a knowledge extractor for the Fiat–Shamir transformation of a  $\Gamma$ -special sound  $\Sigma$ -protocol to the existence of a knowledge extractor for the original protocol  $\Pi$ , which is provided by Theorem 1 of [AFR23] (if  $t_\Gamma$  is polynomial). However, it turns out that when trying to extend our result to the multi-round case (which we do by applying the 3-round case recursively), the additional denominator  $t_\Gamma$  in the above success probability is problematic; indeed, it results in an unwanted (and unnecessary) blowup of the knowledge error. For this reason, we will consider an improved version of the above extractor, proposed in [KL23], which avoids the  $t_\Gamma$  in the denominator. We discuss this improved extractor in detail below.

### 3.4 An Improved Extractor for $\Gamma$ -Special Sound $\Sigma$ -Protocols

Here, we discuss the improved knowledge extractor from the unpublished manuscript [KL23], shared in personal communication, that avoids the factor  $1/t_\Gamma$  loss in success probability.

Similar to the recent works on the topic, we now continue the discussion in a more abstract language. Consider an algorithm  $\mathcal{A}: \mathcal{C} \rightarrow \{0, 1\}^*$ , as well as a verification predicate  $\mathcal{V}: \mathcal{C} \times \{0, 1\}^* \rightarrow \{0, 1\}$ . Naturally, we would instantiate  $\mathcal{A}$  with a dishonest prover  $\mathcal{P}^*$ , attacking a  $\Gamma$ -out-of- $\mathcal{C}$  special-sound  $\Sigma$ -protocol, with the understanding that the output  $(a, z)$  of  $\mathcal{A}(c)$  consists of the two messages  $a$  and  $z$  produced by  $\mathcal{P}^*$  on challenge  $c$ . By a standard averaging argument we may assume the dishonest prover  $\mathcal{P}^*$ , and thus the algorithm  $\mathcal{A}$ , to be deterministic. See for instance [AF22] for a formal proof of this claim. Note that, under this assumption, the first message  $a$  is fixed and independent of the challenge  $c$ . We call an output  $y \leftarrow \mathcal{A}(c)$  *accepting* or *correct* if  $\mathcal{V}(c, y) = 1$ . For  $C$  uniformly random over  $\mathcal{C}$ , the *success probability* of  $\mathcal{A}$  is denoted as

$$\epsilon^\mathcal{V}(\mathcal{A}) := \Pr(\mathcal{V}(C, \mathcal{A}(C)) = 1);$$

this then obviously coincides with the success probability of  $\mathcal{P}^*$ . The goal is to design an extractor, with black-box access to  $\mathcal{A}$ , that finds accepting  $y$ 's for challenges  $c_1, \dots, c_k$  that form a set in  $\Gamma$ .

The extractor proposed in [KL23] is given in Figure 1 below. In spirit, it is quite similar to the extractor considered in [AFR23]: it first runs  $\mathcal{A}$  on a random challenge  $c$  and aborts if  $\mathcal{A}$  fails to produce an accepting transcript; otherwise, if  $\mathcal{A}$  has produced an accepting transcript, the extractor runs a geometric experiment with the goal of finding enough additional accepting transcripts. These two phases of the extractor are respectively denoted  $\mathcal{E}_{\text{init}, \Gamma}^{\mathcal{A}}$  and  $\mathcal{E}_{\text{search}, \Gamma}^{\mathcal{A}}$ .

A crucial difference is that, in contrast to the extractor of [AFR23], the coin toss in the geometric experiments of in [KL23] depends on the challenges collected so far. More precisely, in [AFR23] the probability that the coin returns 1 equals the success probability of the first  $\mathcal{A}$ -invocation of the extractor, and thus the probability of entering the second phase of the extraction. This yields a simple running time analysis, but incurs the factor  $t_\Gamma$  loss in the success probability. In [KL23], the probability that the coin returns 1 is chosen more carefully, which leads to a (much) smaller probability for the coin returning 1, thus reducing the probability that the extractor aborts before succeeding, which in turn improves the success probability of the extractor (while still being able to control the expected running time sufficiently).

Further, the extractor of [AFR23] recursively invokes a (sub)extractor for a larger monotone structure. As a consequence, the invocation of a (sub)extractor might fail even after some accepting transcripts have been found, in which case these accepting transcripts are “forgotten”, and either the extractor aborts or

**Fig. 1.** Extractor  $\mathcal{E}_\Gamma^{\mathcal{A}}$  of [KL23].

**Parameters:** a nonempty monotone structure  $(\Gamma, \mathcal{C})$ .

**Oracle access to:** algorithm  $\mathcal{A}: \mathcal{C} \rightarrow \{0, 1\}^*$  and verification predicate  $\mathcal{V}: \mathcal{C} \times \{0, 1\}^* \rightarrow \{0, 1\}$ .

**Output:**  $\perp$  or  $(c_1, \dots, c_k, y_1, \dots, y_k)$  with  $\{c_1, \dots, c_k\} \in \Gamma$  and  $\mathcal{V}(c_j, y_j) = 1 \quad \forall j$ .

1. Sample  $c_1 \in \mathcal{C}$  uniformly at random and evaluate  $y_1 \leftarrow \mathcal{A}(c_1)$ .
2. If  $\mathcal{V}(c_1, y_1) = 0$ , abort and output  $\perp$ .
3. If  $\mathcal{V}(c_1, y_1) = 1$  and  $\{c_1\} \in \Gamma$ , output  $(c_1, y_1) \in \mathcal{C} \times \{0, 1\}^*$ .

$\mathcal{E}_{\text{init}, \Gamma}^{\mathcal{A}}$

4. Else set  $S_1 = \{c_1\} \subseteq \mathcal{C}$ ,  $i = 1$  and  $\text{COIN} = 0$ .

$\mathcal{E}_{\text{search}, \Gamma}^{\mathcal{A}}(c_1, y_1)$

5. Repeat:

- sample  $c_{i+1} \in \mathcal{U}_\Gamma(S_i)$  uniformly at random and evaluate  $y_{i+1} \leftarrow \mathcal{A}(c_{i+1})$ ;
  - if  $\mathcal{V}(c_{i+1}, y_{i+1}) = 1$ , set  $S_{i+1} := S_i \cup \{c_{i+1}\}$  and then set  $i = i + 1$ ;
  - else sample  $d \in \mathcal{U}_\Gamma(S_{i-1})$ , respectively in  $d \in \mathcal{C}$  in case  $i = 1$ , uniformly at random and set  $\text{COIN} = 1$  if  $d \notin \mathcal{U}_\Gamma(S_i)$  and  $\mathcal{V}(d, \mathcal{A}(d)) = 1$ ;

until  $S_i \in \Gamma$  or until  $\text{COIN} = 1$ .

6. If  $\text{COIN} = 1$ , abort and output  $\perp$ .

7. If  $S_i = \{c_1, \dots, c_i\} \in \Gamma$ , output  $(c_1, \dots, c_i, y_1, \dots, y_i)$  with  $\mathcal{V}(c_j, y_j) = 1$  for all  $1 \leq j \leq i$ .

it invokes the subextractor again with fresh randomness. By contrast, the extractor of [KL23] collects the challenges iteratively, and accepting transcripts are never forgotten.

Additionally, when analyzing the geometric experiment, [KL23] expresses the relevant bounds in terms of the challenges collected so far, while [AFR23] uses a worst case bound over all possibilities.

Altogether, this allows [KL23] to argue the following extractability result, which, compared to [AFR23, Theorem 1], avoids the  $t_\Gamma$  in the denominator of the success probability. For completeness, we provide a proof of Theorem 1 in Appendix B.

**Theorem 1 (Extraction Algorithm -  $\Sigma$ -protocols [KL23]).** *Let  $(\Gamma, \mathcal{C})$  be a nonempty monotone structure and let  $\mathcal{V}: \mathcal{C} \times \{0, 1\}^* \rightarrow \{0, 1\}$ . Then there exists an oracle algorithm  $\mathcal{E}_\Gamma$  with the following properties: The algorithm  $\mathcal{E}_\Gamma^{\mathcal{A}}$ , given oracle access to a (probabilistic) algorithm  $\mathcal{A}: \mathcal{C} \rightarrow \{0, 1\}^*$ , requires an expected number of at most  $1 + t_\Gamma \cdot (1 + \kappa_\Gamma)$  queries to  $\mathcal{A}$  and, with probability at least*

$$\frac{\epsilon^{\mathcal{V}(\mathcal{A})} - \kappa_\Gamma}{1 - \kappa_\Gamma},$$

*it outputs pairs  $(c_1, y_1), (c_2, y_2), \dots, (c_k, y_k) \in \mathcal{C} \times \{0, 1\}^*$  with  $\mathcal{V}(c_i, y_i) = 1$  for all  $i$  and  $\{c_1, \dots, c_k\} \in \Gamma$ .*

*Remark 3.* In contrast to the extractors of [AFR23] and [KL23], the knowledge extractor of Figure 1 samples the first challenge uniformly at random from  $\mathcal{C}$  rather than from  $\mathcal{U}_\Gamma(\emptyset)$ . This minor adaptation may cause the first challenge  $c_1$  found by the extractor to be useless, i.e.,  $c_1$  may be in  $\mathcal{C} \setminus \mathcal{U}_\Gamma(\emptyset)$ . As a consequence, the expected number of  $\mathcal{A}$ -queries is bounded by  $1 + t_\Gamma \cdot (1 + \kappa_\Gamma)$  instead of the slightly smaller bound  $1 + (t_\Gamma - 1) \cdot (1 + \kappa_\Gamma)$ . The reason for this seemingly suboptimal design choice is to simplify the analysis of the Fiat–Shamir transformation. More precisely, we will see that, due to this design choice, Theorem 1 can be deployed in a black-box manner when analyzing the Fiat–Shamir transformation of  $\Gamma$ -special-sound  $\Sigma$ -protocols. Note that only for contrived examples of interactive proofs does it hold that  $\mathcal{C} \neq \mathcal{U}_\Gamma(\emptyset)$ . Hence, the above discussion can be avoided by the reasonable assumption that  $\mathcal{C} = \mathcal{U}_\Gamma(\emptyset)$ , as is done in [KL23].

## 4 Refined Running Time Analysis of the Extractor from Theorem 1

In this section, we will refine the running time analysis of the knowledge extractor of Figure 1. Instead of simply counting the (expected) number of  $\mathcal{A}$ -invocations, we associate  $\mathcal{A}$  with a cost function  $\theta: \mathcal{C} \rightarrow \mathbb{R}_{\geq 0}$ , such that  $\theta(c)$  denotes the cost of evaluating  $\mathcal{A}(c)$ . We can thereby give a tighter bound on the running time in scenarios where some  $\mathcal{A}$ -invocations are more costly than others. This refinement turns out to be essential when considering Fiat–Shamir transformations of multi-round interactive proofs.

**Lemma 4 (Refined Running Time Analysis - IP Extractor).** *Let  $(\Gamma, \mathcal{C})$  be a nonempty monotone structure and  $\mathcal{A}: \mathcal{C} \rightarrow \{0, 1\}^*$  an algorithm accompanied by a verification predicate  $\mathcal{V}: \mathcal{C} \times \{0, 1\}^* \rightarrow \{0, 1\}$ . Further, let  $\theta: \mathcal{C} \rightarrow \mathbb{R}_{\geq 0}$ , such that  $\theta(c)$  denotes the cost of evaluating  $\mathcal{A}(c)$ .*

*Then the expected cost of the  $\mathcal{A}$ -invocations of the extractor  $\mathcal{E}_\Gamma^{\mathcal{A}}$  of Figure 1 is at most*

$$\mathbb{E}[\theta(C)] + t_\Gamma \cdot \frac{\mathbb{E}[\theta(C)]}{1 - \kappa_\Gamma}$$

where  $C$  is distributed uniformly at random over  $\mathcal{C}$ .

Before we prove this lemma, we introduce some convenient notation. We also note that even though it is not necessary, it may be didactically helpful to first verify the non-refined analysis in Appendix B.

For any  $0 < k \in \mathbb{Z}$ ,  $\mathbf{c} = (c_1, \dots, c_k) \in \mathcal{C}^k$  and  $1 \leq i \leq k$ , we write  $\mathbf{c}_i = (c_1, \dots, c_i)$ . We will abuse notation by occasionally interpreting the vector  $\mathbf{c}$  as the subset of challenges containing the coordinates of  $\mathbf{c}$ , i.e., the statement  $\mathbf{c} \in \Gamma \subseteq 2^{\mathcal{C}}$  is interpreted as  $\{c_1, \dots, c_k\} \in \Gamma$ . Further, to simplify notation, we write  $V$  for the event  $\mathcal{V}(C, \mathcal{A}(C)) = 1$  and  $U_i$  for the event  $C \in \mathcal{U}_\Gamma(\mathbf{c}_i)$ , where  $C$  is distributed uniformly at random over  $\mathcal{C}$  and  $\mathbf{c}_i$  is given by the context. Additionally,  $U_0$  denotes the event  $C \in \mathcal{C}$ , i.e.,  $\Pr(U_0) = 1$ .<sup>9</sup> Recall that  $\mathcal{U}_\Gamma(\mathbf{c}) \subseteq \mathcal{U}_\Gamma(\mathbf{c}_{k-1}) \subseteq \dots \subseteq \mathcal{U}_\Gamma(\mathbf{c}_1) \subseteq \mathcal{C}$  (Equation (2)), or, in terms of probability events,

$$U_k \implies U_{k-1} \implies \dots \implies U_0. \quad (5)$$

The following three quantities will play a crucial role in the analysis:

$$\begin{aligned} \delta(\mathbf{c}_i) &= \Pr(V \mid U_i), \\ \Delta(\mathbf{c}_i) &= \Pr(V \wedge \neg U_i \mid U_{i-1}), \\ \tilde{\delta}(\mathbf{c}_i) &= \frac{\delta(\mathbf{c}_{i-1}) - \Delta(\mathbf{c}_i)}{1 - \Delta(\mathbf{c}_i)}. \end{aligned} \quad (6)$$

The first two quantities represent certain parameters of the considered extractor from Figure 1. Namely,  $\delta(\mathbf{c}_i)$  is the probability that  $\mathcal{V}(c_{i+1}, y_{i+1}) = 1$  in the repeat loop in step 5, given that the extractor has already found some set (or vector) of accepting challenges  $\mathbf{c}_i = (c_1, \dots, c_i)$ . On the other hand,  $\Delta(\mathbf{c}_i)$  is the probability for the considered coin to become 1 (which then means that the extractor stops unsuccessfully). Hence,  $\delta(\mathbf{c}_i)$  and  $\Delta(\mathbf{c}_i)$  are the parameters of the two geometric experiments that are run in parallel in step 5, considering the challenges collected so far. Additionally,  $\tilde{\delta}(\mathbf{c}_i)$  is an auxiliary quantity that will be relevant in the analysis; for instance, the following shows that it lower bounds  $\delta(\mathbf{c}_i)$ :

$$\begin{aligned} \delta(\mathbf{c}_i) &= \Pr(V \mid U_i) = \frac{\Pr(V \wedge U_i \mid U_{i-1})}{\Pr(U_i \mid U_{i-1})} = \frac{\Pr(V \wedge U_i \mid U_{i-1})}{1 - \Pr(\neg U_i \mid U_{i-1})} \geq \frac{\Pr(V \wedge U_i \mid U_{i-1})}{1 - \Pr(V \wedge \neg U_i \mid U_{i-1})} \\ &= \frac{\Pr(V \wedge U_i \mid U_{i-1})}{1 - \Delta(\mathbf{c}_i)} = \frac{\Pr(V \mid U_{i-1}) - \Pr(V \wedge \neg U_i \mid U_{i-1})}{1 - \Delta(\mathbf{c}_i)} = \frac{\delta(\mathbf{c}_{i-1}) - \Delta(\mathbf{c}_i)}{1 - \Delta(\mathbf{c}_i)} = \tilde{\delta}(\mathbf{c}_i). \end{aligned} \quad (7)$$

Furthermore, exploiting this inequality  $\delta(\mathbf{c}_i) \geq \tilde{\delta}(\mathbf{c}_i)$ , and using that  $\Delta(\mathbf{c}_i) \leq 1$ , we have

$$\delta(\mathbf{c}_i) + \Delta(\mathbf{c}_i) - \Delta(\mathbf{c}_i)\delta(\mathbf{c}_i) \geq \tilde{\delta}(\mathbf{c}_i) + \Delta(\mathbf{c}_i) - \Delta(\mathbf{c}_i)\tilde{\delta}(\mathbf{c}_i) = \delta(\mathbf{c}_{i-1}), \quad (8)$$

where the equality is obtained by solving the definition of  $\tilde{\delta}(\mathbf{c}_i)$  for  $\delta(\mathbf{c}_{i-1})$ . Similarly,

$$\frac{\delta(\mathbf{c}_i) + \Delta(\mathbf{c}_i) - \Delta(\mathbf{c}_i)\delta(\mathbf{c}_i)}{\delta(\mathbf{c}_i)} \leq \frac{\tilde{\delta}(\mathbf{c}_i) + \Delta(\mathbf{c}_i) - \Delta(\mathbf{c}_i)\tilde{\delta}(\mathbf{c}_i)}{\tilde{\delta}(\mathbf{c}_i)} = \frac{\delta(\mathbf{c}_{i-1})}{\tilde{\delta}(\mathbf{c}_i)}. \quad (9)$$

Equation (8) is used in the running time analysis (below and in Appendix B), while Equation (9) is used for the analysis of the success probability (in Appendix B).

We are now ready for the proof of Lemma 4.

<sup>9</sup> This convention reflects that the extractor sampling the first challenge from  $\mathcal{C}$ , rather than from  $\mathcal{U}_\Gamma(\emptyset)$ .

*Proof of Lemma 4.* Let us write  $t = t_\Gamma$ . For  $1 \leq i \leq t$ , let the random variable  $C_i$  denote the  $i$ -th successful challenge found by the extractor, where we let  $C_i = \perp$  if the extractor finishes before finding  $i$  challenges, i.e.,  $C_i$  has support in  $\mathcal{C} \cup \{\perp\}$ . Note that if  $\mathbf{C}_i := (C_1, \dots, C_i) \in \Gamma$  then  $C_{i+1} = \dots = C_t = \perp$ . Vice versa,  $C_i = \perp$  implies that the extractor was either successful before the  $i$ -th iteration, or that it aborted and failed before the  $i$ -th iteration.

For  $1 \leq i < t$ , let  $\mathbf{c}_i \in \mathcal{C}^i$  with  $\mathbf{c}_i \notin \Gamma$  and  $\Pr[C_i = \mathbf{c}_i] > 0$ . Then, conditioning on  $[C_i = \mathbf{c}_i]$  means that we consider a case where the extractor has found  $i$  challenges, but it needs at least one more. In order to do so, it runs two geometric experiments in parallel with parameters  $\delta(\mathbf{c}_i)$  and  $\Delta(\mathbf{c}_i)$ , trying to find the  $(i+1)$ -th challenge (step 5). The probability that at least one of the experiments finishes in a single trial equals

$$p := 1 - (1 - \delta(\mathbf{c}_i))(1 - \Delta(\mathbf{c}_i)) = \delta(\mathbf{c}_i) + \Delta(\mathbf{c}_i) - \Delta(\mathbf{c}_i)\delta(\mathbf{c}_i) \geq \delta(\mathbf{c}_{i-1}),$$

where the inequality follows from Equation (8). Thus, this run of two geometric experiments is again a geometric experiment, but now with parameter  $p$ .

Next, we evaluate the cost of this (combined) geometric experiment. For this purpose we note that to determine the value of the coin, the extractor only needs to invoke  $\mathcal{A}$  if  $d \notin \mathcal{U}_\Gamma(\mathbf{c}_i)$ . Hence, recalling that the distribution of  $c_{i+1}$  is that of the random variable  $C$  conditioned on  $U_i$ , and that  $d$  is distributed as  $C$  conditioned on  $U_{i-1}$ , in each trial the expected cost of the  $\mathcal{A}$ -invocations is at most

$$\begin{aligned} & \mathbb{E}[\theta(C) \mid U_i] + \Pr(\neg U_i \mid U_{i-1}) \cdot \mathbb{E}[\theta(C) \mid U_{i-1} \wedge \neg U_i] \\ &= \frac{\Pr(U_i)}{\Pr(U_i)} \cdot \mathbb{E}[\theta(C) \mid U_i] + \frac{\Pr(U_{i-1} \wedge \neg U_i)}{\Pr(U_{i-1})} \cdot \mathbb{E}[\theta(C) \mid U_{i-1} \wedge \neg U_i] \\ &\leq \frac{1}{\Pr(U_i)} \left( \Pr(U_i) \cdot \mathbb{E}[\theta(C) \mid U_i] + \Pr(U_{i-1} \wedge \neg U_i) \cdot \mathbb{E}[\theta(C) \mid U_{i-1} \wedge \neg U_i] \right) \\ &\leq \frac{\mathbb{E}[\theta(C)]}{1 - \Pr(\neg U_i)} \leq \frac{\mathbb{E}[\theta(C)]}{1 - \kappa_\Gamma}, \end{aligned}$$

where, in the first inequality, we use that  $U_i \implies U_{i-1}$  and thus  $\Pr(U_i) \leq \Pr(U_{i-1})$ . Moreover, the second inequality follows since  $\theta(c) \geq 0$  for all  $c \in \mathcal{C}$ , and the last inequality follows from the definition of  $\kappa_\Gamma$ , exploiting that  $\mathbf{c}_i \notin \Gamma$  and thus  $\mathcal{C} \setminus \mathcal{U}_\Gamma(\mathbf{c}_i) \notin \Gamma$  (Equation (3)).

Hence, if we let  $Y_i$  denote the cost of the  $\mathcal{A}$ -queries that the extractor makes in its  $i$ -th iteration, i.e., when trying to find the  $i$ -th challenge, then

$$\mathbb{E}[Y_1] = \mathbb{E}[\theta(C)] \quad \text{and} \quad \mathbb{E}[Y_{i+1} \mid \mathbf{C}_i = \mathbf{c}_i] \leq \frac{\mathbb{E}[\theta(C)]}{\delta(\mathbf{c}_{i-1}) \cdot (1 - \kappa_\Gamma)},$$

where the upper-bound is the product of the expected number of trials and the expected cost per trial (by Lemma 2). On the other hand, if  $\mathbf{c}_i$  is such that  $\mathbf{c}_i \in \Gamma$  or  $c_i = \perp$  (in either case, the extractor is done) then  $\mathbb{E}[Y_{i+1} \mid \mathbf{C}_i = \mathbf{c}_i] = 0$ .

Putting these observations together shows that for all  $1 \leq k \leq t$

$$\mathbb{E}[Y_{k+1}] = \sum_{\mathbf{c}_k \in T_k} \Pr(\mathbf{C}_k = \mathbf{c}_k) \cdot \mathbb{E}[Y_{k+1} \mid \mathbf{C}_k = \mathbf{c}_k] \leq \frac{\mathbb{E}[\theta(C)]}{1 - \kappa_\Gamma} \cdot \sum_{\mathbf{c}_k \in T_k} \frac{\Pr(\mathbf{C}_k = \mathbf{c}_k)}{\delta(\mathbf{c}_{k-1})}.$$

where  $T_k := \{\mathbf{c}_k \in \mathcal{C}^k : \mathbf{c}_k \notin \Gamma \wedge \Pr(\mathbf{C}_k = \mathbf{c}_k) > 0\}$ .

We will now expand the probability  $\Pr(\mathbf{C}_k = \mathbf{c}_k)$ , aiming for an upper bound on the above sum. For any  $\mathbf{c}_k \in T_k$  and  $1 \leq i < k$ , by Lemma 1 it holds that

$$\Pr(C_{i+1} \neq \perp \mid \mathbf{C}_i = \mathbf{c}_i) = \frac{\delta(\mathbf{c}_i)}{\delta(\mathbf{c}_i) + \Delta(\mathbf{c}_i) - \Delta(\mathbf{c}_i)\delta(\mathbf{c}_i)} \leq \frac{\delta(\mathbf{c}_i)}{\delta(\mathbf{c}_{i-1})},$$

where the last inequality applies Equation (8). Noting that  $c_{i+1} \neq \perp$  by definition of  $T_k$ , it follows that

$$\begin{aligned} \Pr(C_{i+1} = c_{i+1} \mid \mathbf{C}_i = \mathbf{c}_i) &= \Pr(C_{i+1} = c_{i+1} \wedge C_{i+1} \neq \perp \mid \mathbf{C}_i = \mathbf{c}_i) \\ &= \tau(c_{i+1} \mid \mathbf{c}_i) \cdot \Pr(C_{i+1} \neq \perp \mid \mathbf{C}_i = \mathbf{c}_i) \end{aligned}$$

$$\leq \tau(c_{i+1} \mid \mathbf{c}_i) \cdot \frac{\delta(\mathbf{c}_i)}{\delta(\mathbf{c}_{i-1})}.$$

where,

$$\tau(c_{i+1} \mid \mathbf{c}_i) := \Pr(C_{i+1} = c_{i+1} \mid C_{i+1} \neq \perp \wedge \mathbf{C}_i = \mathbf{c}_i).$$

Hence, it follows that

$$\Pr(\mathbf{C}_k = \mathbf{c}_k) = \prod_{i=1}^k \Pr(C_i = c_i \mid \mathbf{C}_{i-1} = \mathbf{c}_{i-1}) = \delta(\mathbf{c}_{k-1}) \cdot \prod_{i=1}^k \tau(c_i \mid \mathbf{c}_{i-1}).$$

Hence, for all  $k \geq 1$ ,

$$\begin{aligned} \mathbb{E}[Y_{k+1}] &\leq \frac{\mathbb{E}[\theta(C)]}{1 - \kappa_\Gamma} \cdot \sum_{\mathbf{c}_k \in \mathcal{T}_k} \prod_{i=1}^k \tau(c_i \mid \mathbf{c}_{i-1}) \\ &\leq \frac{\mathbb{E}[\theta(C)]}{1 - \kappa_\Gamma} \cdot \sum_{c_1 \in \mathcal{C}} \tau(c_1 \mid \mathbf{c}_0) \sum_{c_2 \in \mathcal{C}} \tau(c_2 \mid \mathbf{c}_1) \cdots \sum_{c_k \in \mathcal{C}} \tau(c_k \mid \mathbf{c}_{k-1}) \\ &\leq \frac{\mathbb{E}[\theta(C)]}{1 - \kappa_\Gamma}. \end{aligned}$$

Hence, the expected running time of the extractor is at most

$$\mathbb{E}[\theta(C)] + \sum_{k=1}^t \mathbb{E}[Y_{k+1}] \leq \mathbb{E}[\theta(C)] + t \cdot \frac{\mathbb{E}[\theta(C)]}{1 - \kappa_\Gamma},$$

which completes the proof of the lemma.  $\square$

## 5 The Fiat–Shamir Transformation of $\Gamma$ -Special-Sound $\Sigma$ -Protocols

In this section, we analyze the knowledge soundness of the Fiat–Shamir transformation of  $\Gamma$ -out-of- $\mathcal{C}$  special-sound  $\Sigma$ -protocols, i.e., we first restrict ourselves to 3-round interactive proofs. In Section 6, we will generalize the analysis to multi-round  $(\Gamma_1, \dots, \Gamma_\mu)$ -out-of- $(\mathcal{C}_1, \dots, \mathcal{C}_\mu)$  special-sound interactive proofs. In both cases, we show that the security loss of the Fiat–Shamir transformation is linear in the query complexity  $Q$  of a dishonest prover attacking the non-interactive random oracle proof; in particular, the security loss is independent of the number of rounds. In comparison, in general the security loss of the Fiat–Shamir transformation may be exponential in the number of rounds.

In Section 5.1, we will set the stage and introduce the required notation and auxiliary lemmas. In Section 5.2, we will present our knowledge extractor and analyze its properties. Subsequently, to prepare for our analysis of Fiat–Shamir transformations of multi-round interactive proofs, we will provide a refined running time analysis in Section 5.3.

### 5.1 Preliminary Discussion

Following the notation of prior works, and in line with the notation used in Section 3 for *interactive* proofs, we present our core results in an abstract language. More precisely, we define an abstract algorithm  $\mathcal{A}$  by which we capture the behavior of a (dishonest)  $Q$ -query prover  $\mathcal{P}^*$  attacking the non-interactive proof. A prover attacking the interactive proof receives a challenge  $c$  and aims to provide an accepting response. In the non-interactive setting, the prover  $\mathcal{P}^*$  receives (access to) a random oracle  $\text{ro}: \mathcal{M} \rightarrow \mathcal{C}$ , where  $\mathcal{M}$  is some finite set containing all potential first messages  $a$  of the interactive proof.

On input a statement  $x$ , and after making at most  $Q$  queries to  $\text{ro}$ , the prover  $\mathcal{P}^*$  outputs a proof  $\pi = (a, z)$ , which is accepting if and only if  $(a, c, z)$  is an accepting transcript, where  $c = \text{ro}(a)$ .<sup>10</sup>

<sup>10</sup> In the adaptive setting, the prover does not receive an input and outputs the statement  $x$  together with a proof  $\pi = (a, z)$ . In this case, one should define the challenge as  $c = \text{ro}(x, a)$ .

The algorithm  $\mathcal{A}$ , which captures the behavior of a prover attacking the non-interactive proof, is hence of the form

$$\mathcal{A}: \mathcal{C}^{\mathcal{M}} \rightarrow \mathcal{M} \times \{0, 1\}^*, \quad \mathbf{ro} \mapsto (a, z),$$

where  $\mathcal{C}^{\mathcal{M}}$  denotes the set of all functions with domain  $\mathcal{M}$  and codomain  $\mathcal{C}$ . Where convenient, we also write  $\mathcal{A} = (A, Z)$  such that  $A(\mathbf{ro})$  and  $Z(\mathbf{ro})$  denote the first and second output of  $\mathcal{A}(\mathbf{ro})$ , respectively. As before,

$$\mathcal{V}: \mathcal{M} \times \mathcal{C} \times \{0, 1\}^* \rightarrow \{0, 1\}, \quad (a, c, z) \mapsto v$$

denotes the verification predicate for the underlying *interactive* proof. This verification predicate naturally extends to a verification predicate

$$\mathcal{V}^{\mathbf{ro}}: \mathcal{M} \times \{0, 1\}^* \rightarrow \{0, 1\}, \quad (a, z) \mapsto \mathcal{V}(a, \mathbf{ro}(a), z),$$

for the non-interactive random oracle proof. By a slight abuse of notation, we sometimes reorder the inputs, so as to write  $\mathcal{V}(c, \mathcal{A}(\mathbf{ro}))$  for  $\mathcal{V}(a, c, z)$  with  $(a, z) \leftarrow \mathcal{A}(\mathbf{ro})$ . The algorithm  $\mathcal{A}$  now has a naturally defined success probability

$$\epsilon^{\mathcal{V}}(\mathcal{A}) = \Pr(\mathcal{V}^{\mathbf{ro}}(\mathcal{A}(\mathbf{RO})) = 1),$$

where  $\mathbf{RO}$  is distributed uniformly at random over the set of random oracles  $\mathcal{C}^{\mathcal{M}}$ .

We note that even though the above notation suggests that the entire function table  $\mathbf{ro}$  is given as input to  $\mathcal{A}$ , it is understood that  $\mathcal{A}$  represents a  $Q$ -query algorithm, and so accesses at most  $Q$  positions of  $\mathbf{ro}$ . Similarly,  $\mathcal{V}^{\mathbf{ro}}$  only needs to make a single query to the random oracle to verify the proof. This difference is only relevant when considering efficiency, but not when considering the success probability.

*Remark 4 (Probabilistic Algorithms).* As discussed in Section 3.4, we may assume  $\mathcal{P}^*$ , and thus  $\mathcal{A}$ , to be deterministic. However, in our (recursive) analysis of multi-round protocols, it will be essential to allow for probabilistic algorithm  $\mathcal{A}$ . For this reason, we do not restrict to deterministic algorithms  $\mathcal{A}$ . Further, given a probabilistic algorithm  $\mathcal{A}$ , we will write  $\mathcal{A}[r]$  for the deterministic algorithm that evaluates  $\mathcal{A}$  with fixed random coins  $r$ .

The goal of the knowledge extractor is to find accepting transcripts  $(a, c_1, z_1), \dots, (a, c_k, z_k)$ , with common first message  $a$  and such that  $\{c_1, \dots, c_k\} \in \Gamma$ . As we show below, this can be achieved by combining techniques from [AFR23, KL23] and [AFK22, AFK23]. Deriving its success probability and expected running time is more involved.

To aid in our analysis we observe that, for all  $\alpha \in \mathcal{M}$  and  $\mathbf{ro} \in \mathcal{C}^{\mathcal{M}}$ , the algorithm  $\mathcal{A}$  defines the algorithm

$$\mathcal{A}_{\alpha}^{\mathbf{ro}}: \mathcal{C} \rightarrow \mathcal{M} \times \{0, 1\}^*, \quad c \mapsto (a, z) := \mathcal{A}(\mathbf{ro}[c \mapsto \alpha]), \quad (10)$$

which takes as input a challenge  $c \in \mathcal{C}$ , reprograms the random oracle  $\mathbf{ro}$  so that it answers queries to  $\alpha$  with  $c$ , and then runs  $\mathcal{A}$  with the reprogrammed random oracle

$$\mathbf{ro}[c \mapsto \alpha]: \mathcal{M} \rightarrow \mathcal{C}, \quad m \mapsto \begin{cases} \mathbf{ro}(m), & \text{if } m \neq \alpha, \\ c, & \text{if } m = \alpha. \end{cases}$$

The corresponding verification predicate is

$$\mathcal{V}_{\alpha}(a, c, z) = \begin{cases} 1 & \text{if } \mathcal{V}(a, c, z) = 1 \text{ and } a = \alpha; \\ 0 & \text{otherwise.} \end{cases} \quad (11)$$

which is as  $\mathcal{V}$ , but additionally insists on  $a$  being  $\alpha$ . Again, we allow  $\mathcal{V}_{\alpha}$  to reorder its input so that we can write  $\mathcal{V}_{\alpha}(c, \mathcal{A}(\mathbf{ro}))$  for  $\mathcal{V}_{\alpha}(a, c, z)$  with  $(a, z) \leftarrow \mathcal{A}(\mathbf{ro})$ . By definition  $\mathcal{V}_{\alpha}^{\mathbf{ro}}(a, z) = \mathcal{V}_{\alpha}(a, \mathbf{ro}(a), z)$ , and it is easily seen that  $\mathcal{V}_{\alpha}(a, \mathbf{ro}(a), z) = \mathcal{V}_{\alpha}(a, \mathbf{ro}(\alpha), z)$ . It thus follows that  $\mathcal{V}_{\alpha}^{\mathbf{ro}}(\mathcal{A}(\mathbf{ro})) = \mathcal{V}_{\alpha}(\mathbf{ro}(\alpha), \mathcal{A}(\mathbf{ro}))$ .

*Remark 5 (Early-abort).* For later purposes, we assume that the computation of  $(a, z) \leftarrow \mathcal{A}(\mathbf{ro})$  is split into two steps:  $\mathcal{A}(\mathbf{ro})$  first computes  $a = A(\mathbf{ro})$ , and then continues to compute  $z = Z(\mathbf{ro})$ . This allows for an early-abort strategy for computing  $\mathcal{A}_{\alpha}^{\mathbf{ro}}(c) = \mathcal{A}(\mathbf{ro}[c \mapsto \alpha])$ , in which the  $\mathcal{A}$ -invocation is aborted if  $a \neq \alpha$  is output. This assumption is not well motivated for the 3-round case, where we cannot expect  $\mathcal{A}$  to spend significantly more time in computing  $z$  once it has decided on  $a$ , but it will be crucial in the multi-round running time analysis, where the 3-round case is recursively applied to  $\mathcal{A}$  being a (sub)extractor that indeed decides on  $a$  early on.

The algorithm  $\mathcal{A}_\alpha^{\text{ro}}$  can be understood as an attacker against the underlying *interactive* proof, with the additional requirement that the attacker must use a particular first message  $\alpha$ . Indeed, this algorithm is precisely of the form required by the extraction algorithm  $\mathcal{E}_I$  (for interactive proofs) of Section 3.4, i.e.,  $\mathcal{E}_I^{\mathcal{A}_\alpha^{\text{ro}}}$  is well-defined. The following lemma now relates the success probability of  $\mathcal{A}_\alpha^{\text{ro}}$  to that of  $\mathcal{A}$ .

**Lemma 5.** *Let  $\mathcal{A}: \mathcal{C}^{\mathcal{M}} \rightarrow \mathcal{M} \times \{0,1\}^*$  be an algorithm together with a verification predicate  $\mathcal{V}: \mathcal{M} \times \mathcal{C} \times \{0,1\}^* \rightarrow \{0,1\}$ . Then<sup>11</sup>*

$$\frac{1}{|\mathcal{C}^{\mathcal{M}}|} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} \sum_{\alpha \in \mathcal{M}} \epsilon^{\mathcal{V}_\alpha(\mathcal{A}_\alpha^{\text{ro}})} = \epsilon^{\mathcal{V}(\mathcal{A})}.$$

Intuitively, this is pretty clear. Averaged over a random choice of the random oracle  $\text{ro}$ , the reprogramming at  $\alpha$  has no effect, and then the summing over all  $\alpha$  removes the requirement on  $a$  being  $\alpha$ . The formal proof is spelled out as follows.

*Proof.* Let  $\text{RO}$  be distributed uniformly at random over  $\mathcal{C}^{\mathcal{M}}$ , and let  $C$  be distributed uniformly at random over  $\mathcal{C}$ . Then

$$\begin{aligned} \frac{1}{|\mathcal{C}^{\mathcal{M}}|} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} \sum_{\alpha \in \mathcal{M}} \epsilon^{\mathcal{V}_\alpha(\mathcal{A}_\alpha^{\text{ro}})} &= \sum_{\alpha \in \mathcal{M}} \Pr(\mathcal{V}_\alpha(C, \mathcal{A}_\alpha^{\text{RO}}(C)) = 1) \\ &= \sum_{\alpha \in \mathcal{M}} \Pr(\mathcal{V}_\alpha(C, \mathcal{A}(\text{RO}[\alpha \mapsto C])) = 1) \\ &= \sum_{\alpha \in \mathcal{M}} \Pr(\mathcal{V}_\alpha(\text{RO}[\alpha \mapsto C](\alpha), \mathcal{A}(\text{RO}[\alpha \mapsto C])) = 1) \\ &= \sum_{\alpha \in \mathcal{M}} \Pr(\mathcal{V}_\alpha(\text{RO}(\alpha), \mathcal{A}(\text{RO})) = 1) \\ &\stackrel{*}{=} \sum_{\alpha \in \mathcal{M}} \Pr(\mathcal{V}_\alpha^{\text{RO}}(\mathcal{A}(\text{RO})) = 1) \\ &= \Pr(\mathcal{V}^{\text{RO}}(\mathcal{A}(\text{RO})) = 1) = \epsilon^{\mathcal{V}(\mathcal{A})}, \end{aligned}$$

where, in the equality marked by  $*$ , we use that  $\mathcal{V}_\alpha^{\text{ro}}(\mathcal{A}(\text{ro})) = \mathcal{V}_\alpha(\text{ro}(\alpha), \mathcal{A}(\text{ro}))$ , as explained earlier. This completes the proof of the lemma.  $\square$

Before we define our knowledge extractor, we introduce the following crucial quantity, denoted  $q(\mathcal{A})$ . For a fixed random oracle  $\text{ro}$  and a fixed randomness of  $\mathcal{A}$ , we count the number of prover messages  $\alpha$  such that, after reprogramming  $\text{ro}$  in  $\alpha$  to a random value  $C$ ,  $\mathcal{A}$ 's first output equals  $\alpha$  with positive probability. The quantity  $q(\mathcal{A})$  is then defined as the expectation of this number, averaged over the choice of the the random oracle and of  $\mathcal{A}$ 's randomness. It turns out that both the expected running time and the success probability of our knowledge extractor depend on  $q(\mathcal{A})$ .

**Definition 8.** *Let  $\mathcal{A} = (A, Z): \mathcal{C}^{\mathcal{M}} \rightarrow \mathcal{M} \times \{0,1\}^*$ ,  $\text{ro} \mapsto (A(\text{ro}), Z(\text{ro}))$  be a (probabilistic) algorithm. Then*

$$q(\mathcal{A}) := \mathbb{E}_r \left[ \frac{1}{|\mathcal{C}^{\mathcal{M}}|} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} |\{\alpha \in \mathcal{M} : \Pr(A[r](\text{ro}[\alpha \mapsto C]) = \alpha) > 0\}| \right],$$

where  $C$  is distributed uniformly random over  $\mathcal{C}$  and the expectation is over the random coins  $r$  of  $\mathcal{A}$ .

The following lemma shows that we can control the value  $q(\mathcal{A})$  via the query complexity of  $\mathcal{A}$ .

**Lemma 6.** *If  $\mathcal{A} = (A, Z): \mathcal{C}^{\mathcal{M}} \rightarrow \mathcal{M} \times \{0,1\}^*$  is a  $Q$ -query algorithm, then  $q(\mathcal{A}) \leq Q + 1$ .*

<sup>11</sup> We note that we use  $\epsilon$  for the success probability in the interactive setting with a uniformly random challenge, as in  $\epsilon^{\mathcal{V}_\alpha(\mathcal{A}_\alpha^{\text{ro}})}$ , and in the non-interactive setting, as in  $\epsilon^{\mathcal{V}(\mathcal{A})}$ .

*Proof.* Without loss of generality, we may assume  $\mathcal{A}$  to be deterministic. Let us fix  $\mathbf{ro} \in \mathcal{C}^{\mathcal{M}}$  and  $a = \mathcal{A}(\mathbf{ro})$ . Further, let  $S(\mathbf{ro}) \subseteq \mathcal{C}$  be the set of messages queried by  $\mathcal{A}(\mathbf{ro})$ , then  $|S(\mathbf{ro})| \leq Q$ .

If  $\mathcal{A}(\mathbf{ro})$  does not query message  $\alpha$ , i.e.,  $\alpha \notin S(\mathbf{ro})$ , then  $\mathcal{A}_\alpha^{\mathbf{ro}}(c) = \mathcal{A}(\mathbf{ro}[\alpha \mapsto c])$  is oblivious to the input  $c \in \mathcal{C}$ . Hence, since  $\mathcal{A}$  is deterministic,  $\mathcal{A}(\mathbf{ro}[\alpha \mapsto c]) = a$  for all  $\alpha \notin S(\mathbf{ro})$  and all  $c \in \mathcal{C}$ . It therefore follows that

$$p(\alpha) := \Pr(\mathcal{A}(\mathbf{ro}[\alpha \mapsto C]) = \alpha) = 0$$

for all  $\alpha \notin S(\mathbf{ro}) \cup \{a\}$ , which implies that

$$|\{\alpha : p(\alpha) > 0\}| = |\{\alpha \in S(\mathbf{ro}) : p(\alpha) > 0\}| + |\{\alpha \notin S(\mathbf{ro}) : p(\alpha) > 0\}| \leq |S(\mathbf{ro})| + 1 \leq Q + 1.$$

The lemma now follows trivially.  $\square$

## 5.2 The Knowledge Extractor

With the above observations at hand, we can define a knowledge extractor  $\mathcal{F}_\Gamma$  for Fiat–Shamir transformations of  $\Gamma$ -out-of- $\mathcal{C}$  special-sound interactive proofs.

In the first step, the extractor  $\mathcal{F}_\Gamma$  evaluates  $(a, z) \leftarrow \mathcal{A}(\mathbf{ro})$  for a random oracle  $\mathbf{ro} \in \mathcal{C}^{\mathcal{M}}$  sampled uniformly at random. If  $\mathcal{V}^{\mathbf{ro}}(a, z) = 0$ , the extractor aborts. Otherwise, the extractor proceeds by running  $\mathcal{E}_{\text{search}, \Gamma}(c, a, z)$  on  $\mathcal{A}_a^{\mathbf{ro}}$ , i.e., it proceeds as in the interactive setting, but now using the algorithm  $\mathcal{A}_a^{\mathbf{ro}}$  (and using the same random coins for  $\mathcal{A}$  as in the first step). Furthermore, in this second phase of the extraction algorithm, every  $\mathcal{A}$ -invocation is early-aborted if its first output is incorrect, i.e.,  $(a', z') \leftarrow \mathcal{A}(\mathbf{ro}[a \mapsto c']) = \mathcal{A}_a^{\mathbf{ro}}(c')$  is early-aborted if  $a' \neq a$  (see also Remark 5). The extractor is formally described in Figure 2. Its properties are summarized in Theorem 2.

We note that the early-abort property is only exploited in the refined running time analysis of Section 5.3 and, subsequently, in the analysis of multi-round interactive proofs. For this reason, this property will not play a role in Theorem 2 and its proof.

**Fig. 2.** Expected polynomial time extractor  $\mathcal{F}_\Gamma$  for the Fiat–Shamir transformation of  $\Gamma$ -out-of- $\mathcal{C}$  special-sound  $\Sigma$ -protocols.

**Parameters:** a nonempty monotone structure  $(\Gamma, \mathcal{C})$ .

**Oracle access to:** algorithm  $\mathcal{A} : \mathcal{C}^{\mathcal{M}} \rightarrow \mathcal{M} \times \{0, 1\}^*$  with verification predicate  $\mathcal{V} : \mathcal{M} \times \mathcal{C} \times \{0, 1\}^* \rightarrow \{0, 1\}$ , such that  $\mathcal{A}_\alpha^{\mathbf{ro}}$  and  $\mathcal{V}_\alpha$  are defined as above (Equations (10) and (11)) for all  $\alpha \in \mathcal{M}$  and  $\mathbf{ro} \in \mathcal{C}^{\mathcal{M}}$ .

**Output:**  $\perp$  or  $(a, c_1, z_1), \dots, (a, c_k, z_k)$  with  $\{c_1, \dots, c_k\} \in \Gamma$  and  $\mathcal{V}(a, c_j, y_j) = 1$  for all  $j$ .

1. Sample  $\mathbf{ro} \in \mathcal{C}^{\mathcal{M}}$  uniformly at random and evaluate  $(a, z) \leftarrow \mathcal{A}(\mathbf{ro})$ .  $\mathcal{F}_{\text{init}, \Gamma}^{\mathcal{A}}$ 
  - From here on, all subsequent  $\mathcal{A}$ -invocations use the same random coins, i.e., from here on the first  $\mathcal{A}$ -invocation is rewound when  $\mathcal{A}$  is invoked.
2. If  $\mathcal{V}^{\mathbf{ro}}(a, z) = 0$ , abort and output  $\perp$ .
3. If  $\mathcal{V}^{\mathbf{ro}}(a, z) = 1$  and  $\{\mathbf{ro}(a)\} \in \Gamma$ , output  $(a, \mathbf{ro}(a), z)$ .

4. Else, set  $S_1 = \{\mathbf{ro}(a)\} \subseteq \mathcal{C}$ ,  $i = 1$  and COIN = 0.

$$\mathcal{F}_{\text{search}, \Gamma}^{\mathcal{A}}(a, z, \mathbf{ro}) = \mathcal{E}_{\text{search}, \Gamma}^{\mathcal{A}_a^{\mathbf{ro}}}(a, z, \mathbf{ro})$$

5. Repeat:

- sample  $c_{i+1} \in \mathcal{U}_\Gamma(S_i)$  uniformly at random and evaluate

$$(a', z_{i+1}) \leftarrow \mathcal{A}_a^{\mathbf{ro}}(c_{i+1}) = \mathcal{A}(\mathbf{ro}[a \mapsto c_{i+1}]),$$

early-aborting the  $\mathcal{A}$ -invocation if  $a' \neq a$ ;

- if  $\mathcal{V}_a(a', c_{i+1}, z_{i+1}) = 1$ , set  $S_{i+1} := S_i \cup \{c_{i+1}\}$  and then set  $i = i + 1$ ;
- else sample  $d \in \mathcal{U}_\Gamma(S_{i-1})$ , respectively  $d \in \mathcal{C}$  if  $i = 1$ , uniformly at random and set COIN = 1 if  $d \notin \mathcal{U}_\Gamma(S_i)$  and  $\mathcal{V}_a(a'', d, z'') = 1$  for  $(a'', z'') \leftarrow \mathcal{A}_a^{\mathbf{ro}}(d)$ , early-aborting the  $\mathcal{A}$ -invocation if  $a'' \neq a$ ;

until  $S_i \in \Gamma$  or until COIN = 1.

6. If COIN = 1, abort and output  $\perp$ .

7. If  $S_i = \{\mathbf{ro}(a), c_2, \dots, c_i\} \in \Gamma$ , output  $(a, \mathbf{ro}(a), z_1), (a, c_2, z_2), \dots, (a, c_i, z_i)$  with  $\mathcal{V}(a, \mathbf{ro}(a), z_1) = 1$  and  $\mathcal{V}(a, c_j, z_j) = 1$  for all  $2 \leq j \leq i$ .



**Theorem 2 (Extractor - Fiat–Shamir Transformation of  $\Sigma$ -Protocols).** *Let  $(\Gamma, \mathcal{C})$  be a nonempty monotone structure. Then there exists an oracle algorithm  $\mathcal{F}_\Gamma$  with the following properties: The algorithm  $\mathcal{F}_\Gamma^{\mathcal{A}}$ , given oracle access to an algorithm  $\mathcal{A} = (A, Z): \mathcal{C}^{\mathcal{M}} \rightarrow \mathcal{M} \times \{0, 1\}^*$  defined as above with verification predicate  $\mathcal{V}: \mathcal{M} \times \mathcal{C} \times \{0, 1\}^* \rightarrow \{0, 1\}$ , requires an expected number of at most  $1 + q(\mathcal{A}) \cdot t_\Gamma \cdot (1 + \kappa_\Gamma)$  queries to  $\mathcal{A}$  and, with probability at least*

$$\frac{\epsilon(\mathcal{A}) - q(\mathcal{A}) \cdot \kappa_\Gamma}{1 - \kappa_\Gamma}$$

*it outputs tuples  $(a, c_1, z_1), \dots, (a, c_k, z_k) \in \mathcal{M} \times \mathcal{C} \times \{0, 1\}^*$ , for some  $k \in \mathbb{N}$ , with  $\mathcal{V}(a, c_j, z_j) = 1$  for all  $j$  and  $\{c_1, \dots, c_k\} \in \Gamma$ .*

*Proof.* Note that, by linearity of expected value and the quantities we aim to bound, the running time and success bounds can be written as expressions  $\mathbb{E}_r[\dots]$  over the random coins of  $r$  of  $\mathcal{A}$ . Thus, without loss of generality, we may assume that  $\mathcal{A}$  is deterministic.

For simplicity, we drop the subscript  $\Gamma$  and write  $\mathcal{E}$  and  $\mathcal{F}$  for  $\mathcal{E}_\Gamma$  and  $\mathcal{F}_\Gamma$ . Let the random variable  $\text{RO}$  denote the random oracle sampled by the extractor for its first  $\mathcal{A}$ -invocation, i.e.,  $\text{RO}$  is distributed uniformly at random over  $\mathcal{C}^{\mathcal{M}}$ . By a small abuse of notation, let  $(A, Z) = \mathcal{A}(\text{RO})$ , i.e., the random variables  $A$  and  $Z$  denote the output of  $\mathcal{F}^{\mathcal{A}}$ 's first invocation of  $\mathcal{A}$ . Finally, for all  $\alpha \in \mathcal{M}$  and  $\text{ro} \in \mathcal{C}^{\mathcal{M}}$ , let  $\Psi_\alpha(\text{ro})$  denote the event

$$\text{RO}(m) = \text{ro}(m) \quad \forall m \in \mathcal{M} \setminus \{\alpha\},$$

i.e., the condition  $\Psi_\alpha(\text{ro})$  fixes  $\text{RO}$  everywhere outside  $\alpha$ . We are now ready to analyze our extractor.

**Success Probability.** Conditioned on  $\Psi_\alpha(\text{ro})$ , the extractor  $\mathcal{F}^{\mathcal{A}}$  first samples a challenge  $\text{RO}(\alpha) = c \in \mathcal{C}$  and evaluates  $(a, z) = \mathcal{A}(\text{ro}[\alpha \mapsto c]) = \mathcal{A}_\alpha^{\text{ro}}(c)$ . As such, conditioned on  $\Psi_\alpha(\text{ro})$ , the first steps of the extractors  $\mathcal{F}^{\mathcal{A}}$  and  $\mathcal{E}^{\mathcal{A}_\alpha^{\text{ro}}}$  are almost identical. The only difference between the two is that  $\mathcal{E}^{\mathcal{A}_\alpha^{\text{ro}}}$  aborts if its first  $\mathcal{A}$ -invocation outputs a message  $a \neq \alpha$ , whereas  $\mathcal{F}^{\mathcal{A}}$  would then proceed with running  $\mathcal{E}_{\text{search}}^{\mathcal{A}_\alpha^{\text{ro}}}$ .<sup>12</sup> It thus follows that

$$\Pr(\mathcal{F}^{\mathcal{A}} = 1 \wedge A = \alpha \mid \Psi_\alpha(\text{ro})) = \Pr(\mathcal{E}^{\mathcal{A}_\alpha^{\text{ro}}} = 1),$$

where we write  $\mathcal{F}^{\mathcal{A}} = 1$  and  $\mathcal{E}^{\mathcal{A}_\alpha^{\text{ro}}} = 1$  for the events that the extractors succeed.

By basic probability theory, we can now derive the following bound

$$\begin{aligned} \Pr(\mathcal{F}^{\mathcal{A}} = 1) &= \sum_{\alpha \in \mathcal{M}} \Pr(\mathcal{F}^{\mathcal{A}} = 1 \wedge A = \alpha) \\ &= \frac{1}{|\mathcal{C}|} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} \sum_{\alpha \in \mathcal{M}} \Pr(\mathcal{F}^{\mathcal{A}} = 1 \wedge A = \alpha \wedge \Psi_\alpha(\text{ro})) \\ &= \frac{1}{|\mathcal{C}|} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} \sum_{\alpha \in \mathcal{M}} \Pr(\Psi_\alpha(\text{ro})) \cdot \Pr(\mathcal{F}^{\mathcal{A}} = 1 \wedge A = \alpha \mid \Psi_\alpha(\text{ro})) \\ &= \frac{1}{|\mathcal{C}|} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} \sum_{\alpha \in \mathcal{M}} \frac{|\mathcal{C}|}{|\mathcal{C}^{\mathcal{M}}|} \cdot \Pr(\mathcal{E}^{\mathcal{A}_\alpha^{\text{ro}}} = 1) \\ &= \frac{1}{|\mathcal{C}^{\mathcal{M}}|} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} \sum_{\alpha \in \mathcal{M}} \Pr(\mathcal{E}^{\mathcal{A}_\alpha^{\text{ro}}} = 1) \\ &\geq \frac{1}{|\mathcal{C}^{\mathcal{M}}|} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} \sum_{\alpha \in \mathcal{M}} \max\left(\frac{\epsilon^{\mathcal{V}_\alpha(\mathcal{A}_\alpha^{\text{ro}})} - \kappa_\Gamma}{1 - \kappa_\Gamma}, 0\right) \\ &\geq \frac{1}{|\mathcal{C}^{\mathcal{M}}|} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} \sum_{\alpha: \epsilon^{\mathcal{V}_\alpha(\mathcal{A}_\alpha^{\text{ro}})} > 0} \frac{\epsilon^{\mathcal{V}_\alpha(\mathcal{A}_\alpha^{\text{ro}})} - \kappa_\Gamma}{1 - \kappa_\Gamma} \\ &\geq \frac{\epsilon^{\mathcal{V}}(\mathcal{A}) - q(\mathcal{A}) \cdot \kappa_\Gamma}{1 - \kappa_\Gamma}, \end{aligned}$$

<sup>12</sup> Here we are exploiting that  $\mathcal{E}$  samples  $c_1$  from  $\mathcal{C}$ , rather than from  $\mathcal{U}_\Gamma(\emptyset)$ , which would be more natural in the context of  $\mathcal{E}$ . See also Remark 3.

where the first inequality follows from Theorem 1 and the final inequality follows from Lemma 5 together with the observation that

$$\epsilon^{\mathcal{V}_\alpha}(\mathcal{A}_\alpha^{\text{ro}}) \leq \Pr(A(\text{ro}[\alpha \mapsto C]) = \alpha),$$

and thus

$$|\{\alpha : \epsilon^{\mathcal{V}_\alpha}(\mathcal{A}_\alpha^{\text{ro}}) > 0\}| \leq |\{\alpha : \Pr(A(\text{ro}[\alpha \mapsto C]) = \alpha) > 0\}|.$$

This proves the claimed bound on the success probability.

**Expected Running Time.** As in the success probability analysis, we will use the fact that after the first  $\mathcal{A}$ -invocation the extractor  $\mathcal{F}^{\mathcal{A}}$  proceeds exactly as in the interactive case.

For  $\text{ro} \in \mathcal{C}^{\mathcal{M}}$  and  $\alpha \in \mathcal{M}$ , let  $V_\alpha^{\text{ro}}$  denote the event that the first  $\mathcal{A}$ -invocation of the extractor  $\mathcal{E}^{\mathcal{A}_\alpha^{\text{ro}}}$ , when applied to  $\mathcal{A}_\alpha^{\text{ro}}$ , is successful. Further, we write  $X_\alpha^{\text{ro}}$  for the number of  $\mathcal{A}$ -queries made by  $\mathcal{E}^{\mathcal{A}_\alpha^{\text{ro}}}$ , not counting the first  $\mathcal{A}$ -query. Then, exploiting  $\mathbb{E}[X_\alpha^{\text{ro}} \mid \neg V_\alpha^{\text{ro}}] = 0$  and using Theorem 1, it follows that

$$\Pr(V_\alpha^{\text{ro}}) \cdot \mathbb{E}[X_\alpha^{\text{ro}} \mid V_\alpha^{\text{ro}}] = \mathbb{E}[X_\alpha^{\text{ro}}] \leq t_\Gamma \cdot (1 + \kappa_\Gamma). \quad (12)$$

Let us now do something similar for the extractor  $\mathcal{F}^{\mathcal{A}}$ . More precisely, we let  $V$  denote the event that  $\mathcal{F}^{\mathcal{A}}$ 's first  $\mathcal{A}$ -invocation is successful, and we let  $Y$  denote the number  $\mathcal{A}$ -queries made by  $\mathcal{F}^{\mathcal{A}}$  except for the first one. Then,

$$\Pr(V \wedge A = \alpha \mid \Psi_\alpha(\text{ro})) = \Pr(V_\alpha^{\text{ro}}),$$

and  $\mathbb{E}[Y \mid \neg V] = 0$ . Further, by construction of the extractors,

$$\mathbb{E}[Y \mid V \wedge A = \alpha \wedge \Psi_\alpha(\text{ro})] = \mathbb{E}[X_\alpha^{\text{ro}} \mid V_\alpha^{\text{ro}}].$$

Hence,

$$\begin{aligned} \mathbb{E}[Y] &= \Pr(V) \cdot \mathbb{E}[Y \mid V] \\ &= \frac{1}{|\mathcal{C}|} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} \sum_{\alpha \in \mathcal{M}} \Pr(\Psi_\alpha(\text{ro})) \cdot \Pr(V \wedge A = \alpha \mid \Psi_\alpha(\text{ro})) \cdot \mathbb{E}[Y \mid V \wedge A = \alpha \wedge \Psi_\alpha(\text{ro})] \\ &= \frac{1}{|\mathcal{C}^{\mathcal{M}}|} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} \sum_{\alpha \in \mathcal{M}} \Pr(V_\alpha^{\text{ro}}) \cdot \mathbb{E}[X_\alpha^{\text{ro}} \mid V_\alpha^{\text{ro}}] \\ &\leq \frac{1}{|\mathcal{C}^{\mathcal{M}}|} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} \sum_{\alpha: \epsilon^{\mathcal{V}_\alpha}(\mathcal{A}_\alpha^{\text{ro}}) > 0} t_\Gamma \cdot (1 + \kappa_\Gamma) \\ &\leq q(\mathcal{A}) \cdot t_\Gamma \cdot (1 + \kappa_\Gamma), \end{aligned}$$

where we use that  $\Pr(V_\alpha^{\text{ro}}) = \epsilon^{\mathcal{V}_\alpha}(\mathcal{A}_\alpha^{\text{ro}}) \leq \Pr(A(\text{ro}[\alpha \mapsto C]) = \alpha)$ . Hence, the expected running time of  $\mathcal{F}^{\mathcal{A}}$  is at most  $1 + q(\mathcal{A}) \cdot t_\Gamma \cdot (1 + \kappa_\Gamma)$ , which completes the proof.  $\square$

The following theorem now immediately follows from Lemma 6 and Theorem 2.

**Theorem 3 (Fiat–Shamir Transformation of a  $\Sigma$ -Protocol).** *Let  $\Pi$  be a  $\Gamma$ -out-of- $\mathcal{C}$  special-sound  $\Sigma$ -protocol  $\Pi$ , with  $t_\Gamma$  polynomial in the size  $|x|$  of the statement  $x$  and such that sampling from  $\mathcal{U}_\Gamma(S)$  takes polynomial time (in  $|x|$ ) for all  $S \subseteq \mathcal{C}$  with  $|S| \leq t_\Gamma$ . Then, the Fiat–Shamir transformation  $\text{FS}[\Pi]$  of a  $\Pi$  is knowledge sound with knowledge error*

$$(Q + 1) \cdot \kappa_\Gamma,$$

where  $\kappa_\Gamma = \max_{S \subseteq \mathcal{C}} \frac{|S|}{|\mathcal{C}|}$  is the knowledge error of the (interactive)  $\Sigma$ -protocol  $\Pi$ .

More precisely,  $\text{FS}[\Pi]$  admits a knowledge extractor that, on input a statement  $x$  and given oracle access to a  $Q$ -query prover  $\mathcal{P}^*$ , makes an expected number of at most  $1 + (Q + 1) \cdot t_\Gamma \cdot (1 + \kappa_\Gamma)$  queries to  $\mathcal{P}^*$  and succeeds to extract a witness for  $x$  with probability at least

$$\frac{\epsilon(\mathcal{P}^*, x) - (Q + 1) \cdot \kappa_\Gamma}{1 - \kappa_\Gamma}.$$

*Remark 6 (Generic Applicability of our Fiat–Shamir Analysis).* The knowledge extractor  $\mathcal{F}_\Gamma$  only mildly depends on the actual construction of the underlying extractor  $\mathcal{E}_\Gamma$  for  $\Gamma$ -special-sound *interactive* proofs. More precisely,  $\mathcal{F}_\Gamma$  uses the fact that  $\mathcal{E}_\Gamma$  first tries a random challenge and only continues if this challenge is accepted. However, to  $\mathcal{F}_\Gamma$  it is irrelevant how  $\mathcal{E}_\Gamma$  continues after this first phase; the second phase of  $\mathcal{E}_\Gamma$  is deployed in a black-box manner. This shows that our design principle allows one to transform any knowledge extractor for  $\Gamma$ -special-sound interactive proofs, containing the aforementioned two phases, into a knowledge extractor for the Fiat–Shamir transformation of  $\Gamma$ -special-sound interactive proofs.

### 5.3 A Refined Running Time Analysis

For the multi-round version of our result, we need a refined running time analysis, which considers an  $\mathcal{A} = (A, Z)$  for which it holds that computing the output  $A(\text{ro})$  is significantly cheaper than computing the full output  $\mathcal{A}(\text{ro}) = (A(\text{ro}), Z(\text{ro}))$ . This allows us to improve the extractor  $\mathcal{F}_\Gamma$  by early aborting, i.e., by only computing the output  $Z$  if it is really necessary. In this section, we make this early abort assumption, which was introduced in Remark 5, quantitative.

Looking ahead, in the multi-round extractor construction we will apply  $\mathcal{F}_\Gamma$  to a particular subextractor, which we cast as an algorithm  $\mathcal{A} = (A, Z)$ , and for which computing  $A$  is indeed cheap while computing  $Z$  is expensive. In this case, the cost will be measured by the number of calls the subextractor makes to the dishonest prover. Below, we just consider a generic cost measure  $\theta$ .

**Lemma 7 (Refined Running Time Analysis - NIROP Extractor).** *Let  $(\Gamma, \mathcal{C})$  be a nonempty monotone structure and*

$$\mathcal{A}: \mathcal{C}^\mathcal{M} \rightarrow \mathcal{M} \times \{0, 1\}^*, \quad \text{ro} \mapsto (A(\text{ro}), Z(\text{ro}))$$

*an algorithm with a verification predicate  $\mathcal{V}: \mathcal{M} \times \mathcal{C} \times \{0, 1\}^* \rightarrow \{0, 1\}$ . Let  $\theta: \mathcal{C}^\mathcal{M} \rightarrow \mathbb{R}_{\geq 1}$ , such that  $\theta(\text{ro})$  denotes the expected cost of evaluating  $(A(\text{ro}), Z(\text{ro})) = \mathcal{A}(\text{ro})$ . Additionally, let us assume that solely evaluating  $A(\text{ro})$  has constant cost 1.*

*Let  $\mathcal{F}_\Gamma^\mathcal{A}$  denote the knowledge extractor of Figure 2 equipped with the early-abort strategy. Then the expected cost of the  $\mathcal{A}$ -invocations of  $\mathcal{F}_\Gamma^\mathcal{A}$  is at most*

$$\mathbb{E}[\theta(\text{RO})] + t_\Gamma \cdot \frac{\mathbb{E}[\theta(\text{RO})] - 1 + q(\mathcal{A})}{1 - \kappa_\Gamma}.$$

*Proof.* By linearity of the expected value (over the random coins of  $\mathcal{A}$ ) and the quantities we aim to bound, we may, without loss of generality, assume that  $\mathcal{A}$  is deterministic.

As in the proof of Theorem 2, we will reduce the problem to the analysis of the extractor  $\mathcal{E}$  for the interactive case, but now using the refined running time analysis of Lemma 4.

We first fix  $\text{ro} \in \mathcal{C}^\mathcal{M}$  and  $\alpha \in \mathcal{M}$  and consider the execution of  $\mathcal{E}_\Gamma^{\mathcal{A}_\alpha^{\text{ro}}}$ . Here, in line with the early abort strategy,  $\mathcal{A}_\alpha^{\text{ro}}(c) = \mathcal{A}(\text{ro}[\alpha \mapsto c])$  first evaluates  $\alpha' \leftarrow A_\alpha^{\text{ro}}(c)$  and aborts if  $\alpha \neq \alpha'$ , otherwise it continues to additionally evaluate  $Z_\alpha^{\text{ro}}(c)$ , where here and below  $(A_\alpha^{\text{ro}}(c), Z_\alpha^{\text{ro}}(c))$  denotes the output of  $\mathcal{A}_\alpha^{\text{ro}}(c)$ . In particular, it holds that  $A_\alpha^{\text{ro}}(c) = A(\text{ro}[\alpha \mapsto c])$ . Thus,  $\mathcal{A}_\alpha^{\text{ro}}$  has cost function

$$\theta_\alpha^{\text{ro}}: \mathcal{C} \rightarrow \mathbb{R}_{\geq 0}, \quad c \mapsto \begin{cases} \theta(\text{ro}[\alpha \mapsto c]), & \text{if } A_\alpha^{\text{ro}}(c) = \alpha, \\ 1, & \text{otherwise.} \end{cases}$$

As in the proof of Theorem 2, let  $V_\alpha^{\text{ro}}$  denote the event that the first  $\mathcal{A}_\alpha^{\text{ro}}$ -invocation of the extractor  $\mathcal{E}_\Gamma^{\mathcal{A}_\alpha^{\text{ro}}}$  is successful, i.e.,  $\Pr(V_\alpha^{\text{ro}}) = \epsilon^{\mathcal{V}_\alpha(\mathcal{A}_\alpha^{\text{ro}})}$ , and let us write  $X_\alpha^{\text{ro}}$  for the cost of the  $\mathcal{A}_\alpha^{\text{ro}}$ -queries made by  $\mathcal{E}_\Gamma^{\mathcal{A}_\alpha^{\text{ro}}}$ , not counting the cost of the first  $\mathcal{A}_\alpha^{\text{ro}}$ -query. Then, recalling (12), we have

$$\Pr(V_\alpha^{\text{ro}}) \cdot \mathbb{E}[X_\alpha^{\text{ro}} \mid V_\alpha^{\text{ro}}] = \mathbb{E}[X_\alpha^{\text{ro}}].$$

In the proof of Theorem 2, we used Theorem 1 to bound  $\mathbb{E}[X_\alpha^{\text{ro}}]$ . Here we use the refined running time bound of Lemma 4, from which it follows that

$$\mathbb{E}[X_\alpha^{\text{ro}}] \leq t_\Gamma \cdot \frac{\mathbb{E}[\theta_\alpha^{\text{ro}}(C)]}{1 - \kappa_\Gamma} = \frac{t_\Gamma}{1 - \kappa_\Gamma} \left( \Pr(A_\alpha^{\text{ro}}(C) = \alpha) \mathbb{E}[\theta_\alpha^{\text{ro}}(C) \mid A_\alpha^{\text{ro}}(C) = \alpha] + \Pr(A_\alpha^{\text{ro}}(C) \neq \alpha) \right) \quad (13)$$

where  $C$  is uniformly random in  $\mathcal{C}$ . Note that we ignore the cost of the first  $\mathcal{A}$  invocation, and thus, in Equation (13), omitted the first summand  $\mathbb{E}[\theta_\alpha^{\text{ro}}(C)]$  from the running time bound of Lemma 4.

Let us now proceed similarly for the extractor  $\mathcal{F}_r^A$ . More precisely, we let  $V$  denote the event that the first  $\mathcal{A}$ -query is successful, and we let  $Y$  denote the cost of the  $\mathcal{A}$ -queries made by  $\mathcal{F}_r^A$  not counting the cost of the first  $\mathcal{A}$ -query. Clearly,  $\mathbb{E}[Y \mid \neg V] = 0$ . Further, we let the random variable  $\text{RO}$  denote the initial choice of the random oracle by  $\mathcal{F}_r^A$  and, as before,  $\Psi_\alpha(\text{ro})$  denotes the event

$$\text{RO}(m) = \text{ro}(m) \quad \forall m \in \mathcal{M} \setminus \{\alpha\}.$$

The crucial observation is that, by construction of  $\mathcal{F}_r$ ,

$$\mathbb{E}[Y \mid V \wedge A(\text{RO}) = \alpha \wedge \Psi_\alpha(\text{ro})] = \mathbb{E}[X_\alpha^{\text{ro}} \mid V_\alpha^{\text{ro}}],$$

and

$$\Pr(V \wedge A(\text{RO}) = \alpha \mid \Psi_\alpha(\text{ro})) = \Pr(V_\alpha^{\text{ro}}) = \epsilon^{\mathcal{V}_\alpha(\mathcal{A}_\alpha^{\text{ro}})}.$$

Therefore,

$$\begin{aligned} \mathbb{E}[Y] &= \Pr(V) \cdot \mathbb{E}[Y \mid V] \\ &= \frac{1}{|\mathcal{C}|} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} \sum_{\alpha \in \mathcal{M}} \Pr(\Psi_\alpha(\text{ro})) \Pr(V \wedge A(\text{RO}) = \alpha \mid \Psi_\alpha(\text{ro})) \mathbb{E}[Y \mid V \wedge A(\text{RO}) = \alpha \wedge \Psi_\alpha(\text{ro})] \\ &= \frac{1}{|\mathcal{C}^{\mathcal{M}}|} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} \sum_{\alpha: \epsilon^{\mathcal{V}_\alpha(\mathcal{A}_\alpha^{\text{ro}})} > 0} \Pr(V_\alpha^{\text{ro}}) \mathbb{E}[X_\alpha^{\text{ro}} \mid V_\alpha^{\text{ro}}] \\ &= \frac{1}{|\mathcal{C}^{\mathcal{M}}|} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} \sum_{\alpha: \epsilon^{\mathcal{V}_\alpha(\mathcal{A}_\alpha^{\text{ro}})} > 0} \mathbb{E}[X_\alpha^{\text{ro}}]. \end{aligned}$$

Hence, by Equation (13),

$$\mathbb{E}[Y] \leq \frac{1}{|\mathcal{C}^{\mathcal{M}}|} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} \sum_{\alpha: \epsilon^{\mathcal{V}_\alpha(\mathcal{A}_\alpha^{\text{ro}})} > 0} \frac{t_r}{1 - \kappa_r} \left( \Pr(A_\alpha^{\text{ro}}(C) = \alpha) \mathbb{E}[\theta_\alpha^{\text{ro}}(C) \mid A_\alpha^{\text{ro}}(C) = \alpha] + 1 - \Pr(A_\alpha^{\text{ro}}(C) = \alpha) \right).$$

As before, observe that  $\epsilon^{\mathcal{V}_\alpha(\mathcal{A}_\alpha^{\text{ro}})} \leq \Pr(A(\text{ro}[\alpha \mapsto C]) = \alpha)$ , and thus

$$|\{\alpha : \epsilon^{\mathcal{V}_\alpha(\mathcal{A}_\alpha^{\text{ro}})} > 0\}| \leq |\{\alpha : \Pr(A(\text{ro}[\alpha \mapsto C]) = \alpha) > 0\}|,$$

which shows that

$$\frac{1}{|\mathcal{C}^{\mathcal{M}}|} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} \sum_{\alpha: \epsilon^{\mathcal{V}_\alpha(\mathcal{A}_\alpha^{\text{ro}})} > 0} 1 \leq q(\mathcal{A}).$$

It therefore follows that

$$\begin{aligned} \mathbb{E}[Y] &\leq \frac{t_r}{1 - \kappa_r} \left( q(\mathcal{A}) + \frac{1}{|\mathcal{C}^{\mathcal{M}}|} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} \sum_{\alpha: \epsilon^{\mathcal{V}_\alpha(\mathcal{A}_\alpha^{\text{ro}})} > 0} \right. \\ &\quad \left. \left( \Pr(A_\alpha^{\text{ro}}(C) = \alpha) \cdot \mathbb{E}[\theta_\alpha^{\text{ro}}(C) \mid A_\alpha^{\text{ro}}(C) = \alpha] - \Pr(A_\alpha^{\text{ro}}(C) = \alpha) \right) \right). \end{aligned}$$

Further, since  $\theta_\alpha^{\text{ro}}(c) \geq 1$  for all  $c \in \mathcal{C}$ , it holds that all the terms in the above summation are nonnegative. Thus, we can extend the summation from  $\{\alpha : \epsilon^{\mathcal{V}_\alpha(\mathcal{A}_\alpha^{\text{ro}})} > 0\} \subseteq \mathcal{M}$  to all of  $\mathcal{M}$ . It therefore follows that

$$\begin{aligned} &\frac{1}{|\mathcal{C}^{\mathcal{M}}|} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} \sum_{\alpha: \epsilon^{\mathcal{V}_\alpha(\mathcal{A}_\alpha^{\text{ro}})} > 0} \left( \Pr(A_\alpha^{\text{ro}}(C) = \alpha) \cdot \mathbb{E}[\theta_\alpha^{\text{ro}}(C) \mid A_\alpha^{\text{ro}}(C) = \alpha] - \Pr(A_\alpha^{\text{ro}}(C) = \alpha) \right) \\ &\leq \frac{1}{|\mathcal{C}^{\mathcal{M}}|} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} \sum_{\alpha \in \mathcal{M}} \left( \Pr(A_\alpha^{\text{ro}}(C) = \alpha) \cdot \mathbb{E}[\theta_\alpha^{\text{ro}}(C) \mid A_\alpha^{\text{ro}}(C) = \alpha] - \Pr(A_\alpha^{\text{ro}}(C) = \alpha) \right) \\ &= \sum_{\alpha \in \mathcal{M}} \frac{1}{|\mathcal{C}^{\mathcal{M}}|} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} \left( \Pr(A_\alpha^{\text{RO}}(C) = \alpha \mid \text{RO} = \text{ro}) \cdot \mathbb{E}[\theta_\alpha^{\text{RO}}(C) \mid A_\alpha^{\text{RO}}(C) = \alpha \wedge \text{RO} = \text{ro}] \right. \\ &\quad \left. - \Pr(A_\alpha^{\text{RO}}(C) = \alpha \mid \text{RO} = \text{ro}) \right) \end{aligned}$$

$$\begin{aligned}
&= \sum_{\alpha \in \mathcal{M}} \sum_{\text{ro} \in \mathcal{C}^{\mathcal{M}}} \left( \Pr(A_{\alpha}^{\text{RO}}(C) = \alpha \wedge \text{RO} = \text{ro}) \cdot \mathbb{E}[\theta_{\alpha}^{\text{RO}}(C) \mid A_{\alpha}^{\text{RO}}(C) = \alpha \wedge \text{RO} = \text{ro}] \right. \\
&\qquad\qquad\qquad \left. - \Pr(A_{\alpha}^{\text{RO}}(C) = \alpha \wedge \text{RO} = \text{ro}) \right) \\
&= \sum_{\alpha \in \mathcal{M}} \left( \Pr(A_{\alpha}^{\text{RO}}(C) = \alpha) \cdot \mathbb{E}[\theta_{\alpha}^{\text{RO}}(C) \mid A_{\alpha}^{\text{RO}}(C) = \alpha] - \Pr(A_{\alpha}^{\text{RO}}(C) = \alpha) \right) \\
&= \sum_{\alpha \in \mathcal{M}} \left( \Pr(A(\text{RO}[\alpha \mapsto C]) = \alpha) \cdot \mathbb{E}[\theta(\text{RO}[\alpha \mapsto C]) \mid A(\text{RO}[\alpha \mapsto C])] - \Pr(A(\text{RO}[\alpha \mapsto C]) = \alpha) \right) \\
&\stackrel{*}{=} \sum_{\alpha \in \mathcal{M}} \left( \Pr(A(\text{RO}) = \alpha) \cdot \mathbb{E}[\theta(\text{RO}) \mid A(\text{RO}) = \alpha] - \Pr(A(\text{RO}) = \alpha) \right) \\
&= \mathbb{E}[\theta(\text{RO})] - 1,
\end{aligned}$$

where equality  $\stackrel{*}{=}$  uses that that  $\text{RO}$  and  $\text{RO}[\alpha \mapsto C]$  are identically distributed for all  $\alpha \in \mathcal{M}$ . Thus,

$$\mathbb{E}[Y] \leq \frac{t_{\Gamma}}{1 - \kappa_{\Gamma}} (q(\mathcal{A}) + \mathbb{E}[\theta(\text{RO})] - 1).$$

Hence, the expected cost of the  $\mathcal{A}$ -queries made by  $\mathcal{F}_{\Gamma}^{\mathcal{A}}$  is at most

$$\mathbb{E}[\theta(\text{RO})] + t_{\Gamma} \cdot \frac{\mathbb{E}[\theta(\text{RO})] - 1 + q(\mathcal{A})}{1 - \kappa_{\Gamma}},$$

which completes the proof.  $\square$

## 6 The Fiat–Shamir Transformation of Multi-Round Interactive Proofs

In this section, we extend our result to the multi-round setting. First, following prior works, we formally introduce multi-round special-soundness, and then proceed to describe and analyze our extractor in the multi-round setting.

### 6.1 $(\Gamma_1, \dots, \Gamma_{\mu})$ -out-of- $(\mathcal{C}_1, \dots, \mathcal{C}_{\mu})$ Special-Soundness

Let us write  $\mathbf{\Gamma} := (\Gamma_1, \dots, \Gamma_{\mu})$  and  $\mathcal{C} := (\mathcal{C}_1, \dots, \mathcal{C}_{\mu})$ . Then, towards defining multi-round special-soundness as defined in [AFR23], we recall the following.

**Definition 9 (Tree of Transcripts).** *Given a monotone structure  $(\Gamma, \mathcal{C})$ , a  $\Gamma$ -tree of transcripts with trunk  $a$  is a set  $T$  of triples  $(a, c, z) \in \{a\} \times \mathcal{C} \times \{0, 1\}^*$  such that*

$$\{c \in \mathcal{C} \mid \exists z \in \{0, 1\}^* : (a, c, z) \in T\} \in \Gamma.$$

*We write  $\text{TREE}_{\Gamma}(a)$  and  $\text{TREE}_{\Gamma}$  for the set of  $\Gamma$ -trees with trunk  $a$  and with arbitrary trunk, respectively.*

*Given  $(\Gamma_1, \mathcal{C}_1), \dots, (\Gamma_{\mu}, \mathcal{C}_{\mu})$ , a  $\mathbf{\Gamma}$ -tree of transcripts with trunk  $a_1$  is recursively defined to be a set  $T \in \text{TREE}_{\mathbf{\Gamma}}$  of triples  $(a_1, c_1, T') \in \{a_1\} \times \mathcal{C}_1 \times \text{TREE}_{\Gamma_2, \dots, \Gamma_{\mu}}$  such that*

$$\{c_1 \in \mathcal{C}_1 \mid \exists T' \in \text{TREE}_{\Gamma_2, \dots, \Gamma_{\mu}} : (a_1, c_1, T') \in T\} \in \Gamma_1,$$

*where  $\text{TREE}_{\Gamma_2, \dots, \Gamma_{\mu}}(a_2)$  and  $\text{TREE}_{\Gamma_2, \dots, \Gamma_{\mu}}$  denote the set of all  $(\Gamma_2, \dots, \Gamma_{\mu})$ -trees with trunk  $a_2$  and with arbitrary trunk, respectively.*

*Remark 7.* In the above recursive definition of a  $\mathbf{\Gamma}$ -tree, we naturally identify the set  $T$  of triples  $(a_1, c_1, T')$ , where each  $T'$  is a set again (namely a set of triples  $(a_2, c_2, T'')$  etc.), with the union of  $\{(a_1, c_1)\} \times T'$  over the triples in  $T$ . This way, a  $\mathbf{\Gamma}$ -tree of transcripts with trunk  $a_1$  becomes a set of tuples  $(a_1, c_1, a_2, \dots, a_{\mu}, c_{\mu}, a_{\mu+1})$ .

**Definition 10 (Tree of Accepting Transcripts).** *Let*

$$\mathcal{V}: \mathcal{M} \times \mathcal{C}_1 \times \mathcal{M} \times \cdots \times \mathcal{M} \times \mathcal{C}_\mu \times \{0, 1\}^* \rightarrow \{0, 1\},$$

*be a verification predicate. A  $\Gamma$ -tree  $T$  of transcripts (with arbitrary trunk  $a_1$ ) is called a  $\Gamma$ -tree of  $\mathcal{V}$ -accepting transcripts if*

$$\mathcal{V}(a_1, c_1, \dots, a_\mu, c_\mu, a_{\mu+1}) = 1$$

*for every  $(a_1, c_1, \dots, a_\mu, c_\mu, a_{\mu+1}) \in T$ . As above, we write  $\text{TREE}_\Gamma^\mathcal{V}$  and  $\text{TREE}_\Gamma^\mathcal{V}(a_1)$  for the set of all  $(\Gamma_1, \dots, \Gamma_\mu)$ -trees of  $\mathcal{V}$ -accepting transcripts (with trunk  $a_1$ ).*

The definition of the generalized special-soundness notion for multi-round protocols can now be provided.

**Definition 11 ( $(\Gamma_1, \dots, \Gamma_\mu)$ -out-of- $(\mathcal{C}_1, \dots, \mathcal{C}_\mu)$  Special-Soundness).** *Let  $(\Gamma_i, \mathcal{C}_i)$  be monotone structures for  $1 \leq i \leq \mu$ , and let  $\Gamma := (\Gamma_1, \dots, \Gamma_\mu)$  and  $\mathcal{C} := (\mathcal{C}_1, \dots, \mathcal{C}_\mu)$ . A  $2\mu+1$ -round public-coin interactive proof  $(\mathcal{P}, \mathcal{V})$  for relation  $R$ , with challenge sets  $\mathcal{C}_1, \dots, \mathcal{C}_\mu$ , is  $\Gamma$ -out-of- $\mathcal{C}$  special-sound if there exists a polynomial time algorithm that, on input a statement  $x$  and a  $\Gamma$ -tree of accepting transcripts, outputs a witness  $w \in R(x)$ . We also say that  $(\mathcal{P}, \mathcal{V})$  is  $\Gamma$ -special-sound.*

Our goal is to prove that, under certain mild assumptions, the Fiat–Shamir transformation of a  $\Gamma$ -special-sound interactive proof is knowledge sound. For context, we note that if the product of the  $t$ -values  $t_\Gamma := \prod_{i=1}^\mu t_{\Gamma_i}$  is polynomially bounded, then [AFR23] establishes that  $\Gamma$ -out-of- $\mathcal{C}$  special-sound interactive proofs are knowledge sound with knowledge error

$$\kappa_\Gamma = 1 - \prod_{i=1}^\mu (1 - \kappa_{\Gamma_i}), \quad (14)$$

where the  $\kappa_{\Gamma_i}$  are defined as in (4). However, we don't need this multi-round result in our work here; we only need the (improved) 3-round case of Section 3.4.

## 6.2 Preliminary Discussion

Given how we reduce the non-interactive to the interactive extractor for the 3-round case, a natural approach would be to do the same here for the multi-round case. Unfortunately, the running time would blow up too much. Indeed, in the 3-round case, having found one accepting transcript  $(a, c, z)$ , we can hope for finding more transcripts *with the same  $a$*  by running  $\mathcal{A}_a^\circ$  with different challenges (indeed, we have shown that  $Q+1$  tries are sufficient in expectation); however, in the multi-round case, having found a transcript  $(a_1, c_1, a_2, \dots, a_\mu, c_\mu, a_{\mu+1})$  it is too much to hope for finding more transcripts with the same  $(a_1, a_2, \dots, a_\mu)$ . It is also not necessary, since in different branches of a tree of transcripts the  $a_i$ 's may well be different.

Another approach, also followed in prior works [ACK21, AF22, AFK22, AFR23, AFK23, KL23], is to handle the multi-round setting recursively. Informally, a  $(2\mu+1)$ -round extractor invokes a 3-round extractor, but replaces the algorithm  $\mathcal{A}$  by appropriate instantiations of a  $(2\mu-1)$ -round extractor.

However, the analysis of this recursive approach exposes certain subtle issues, especially when considering Fiat–Shamir transformations. More precisely, when analyzing Fiat–Shamir transformations, the expected running time of a naive recursion scales linearly in  $Q^\mu$ , where  $Q$  is the query complexity of the prover attacking the non-interactive proof, and thus exponentially in the number of rounds  $2\mu+1$ . For a non-constant number of rounds, this renders the expected running time of the extractor superpolynomial.<sup>13</sup>

The same problem was encountered in the analysis of Fiat–Shamir transformations of  $(k_1, \dots, k_\mu)$ -special-sound interactive proofs [AFK22, AFK23]. There, it is solved by making the following observation. An invocation of the  $(2\mu-1)$ -round (sub)extractor is successful if it outputs a (sub)tree of accepting transcripts, *and* the (unique) first message of the tree corresponds to the first messages output by earlier invocations of this  $(2\mu-1)$ -round extractor. The second requirement ensures that sufficiently many (sub)trees can be combined into a single larger tree.

<sup>13</sup> Assuming  $Q$  is polynomially-large, as is standard.

Hence, every invocation of a subextractor can fail for two reasons: (1) because it fails to output a subtree; or (2) because the first message output by the subextractor is incorrect. Since the first message is already determined in the first step of the (sub)extractor, a failure of type (2) can be identified before completing the execution of the subextractor. In other words, this observation allows for an *early-abort* strategy, where the subextractor aborts its execution after the first step if it outputs the wrong first message (cf. Remark 5). Hence, failures of type (2) are far less costly than failures of type (1); the latter can only be identified after the subextractor has finished. By applying this early-abort strategy and refining the running time analysis, it was shown that the recursively defined extractor of [AFK22, AFK23] can indeed be made to run in expected polynomial time. More precisely, the expected running time of their extractor scales linearly in  $Q$ .

We will use the same early-abort strategy to obtain an efficient knowledge extractor. To take this strategy into account, we exploit the refined running time analysis for the 3-round FS-extractor, shown in Section 5.3.

### 6.3 Setting Up the Stage

Let us now move to the analysis of multi-round interactive proofs. In this section, assume we are given  $\mu$  monotone structures  $\Gamma_1, \dots, \Gamma_\mu$ , each  $\Gamma_i \subseteq 2^{\mathcal{C}_i}$  for challenge sets  $\mathcal{C}_i$ . As before, we let  $\Gamma := (\Gamma_1, \dots, \Gamma_\mu)$  and  $\mathcal{C} := (\mathcal{C}_1, \dots, \mathcal{C}_\mu)$ . Further we denote by  $\vec{\text{ro}} = (\text{ro}_1, \dots, \text{ro}_\mu)$  the vector of  $\mu$  random oracles, where the understanding is that the  $i$ -th round challenge is determined via the  $i$ -th random oracle (cf. Definition 13). We define  $\vec{\mathcal{RO}} := \mathcal{C}_1^{\mathcal{M}} \times \mathcal{C}_2^{\mathcal{M}^2} \times \dots \times \mathcal{C}_\mu^{\mathcal{M}^\mu}$ , the set from which the  $\mu$ -tuple of random oracles  $\vec{\text{ro}}$  is sampled.<sup>14</sup>

To capture the behavior of a dishonest prover  $\mathcal{P}^*$  attacking the non-interactive Fiat–Shamir transformation of a  $(\Gamma_1, \dots, \Gamma_\mu)$ -out-of- $(\mathcal{C}_1, \dots, \mathcal{C}_\mu)$  special-sound interactive proof (on input  $x$ ), we consider an abstract algorithm of the form:

$$\mathcal{A}: \vec{\mathcal{RO}} \rightarrow \mathcal{M}^\mu \times \{0, 1\}^*, \quad \vec{\text{ro}} = (\text{ro}_1, \dots, \text{ro}_\mu) \mapsto \mathbf{a} = (a_1, \dots, a_\mu, a_{\mu+1}).$$

Hence, we model the attacker as an algorithm  $\mathcal{A}$  that, on input  $\mu$  random oracles  $\text{ro}_1 \in \mathcal{C}_1^{\mathcal{M}}$ ,  $\text{ro}_2 \in \mathcal{C}_2^{\mathcal{M}^2}$ ,  $\dots$ ,  $\text{ro}_\mu \in \mathcal{C}_\mu^{\mathcal{M}^\mu}$ , outputs a proof  $\pi = (a_1, \dots, a_\mu, a_{\mu+1})$ . As for the 3-round case, we treat the random oracles  $\text{ro}_1, \dots, \text{ro}_\mu$  for the different rounds as input with the understanding that  $\mathcal{A}$  has oracle access only, and when considering an extractor that runs  $\mathcal{A}$ , the random oracle queries are answered using standard lazy sampling. When referring to the number of queries made by  $\mathcal{A}$ , we mean the total number of queries to the random oracles  $\text{ro}_1, \dots, \text{ro}_\mu$ .

The algorithm  $\mathcal{A}$  is accompanied by a verification predicate

$$\mathcal{V}: \mathcal{M} \times \mathcal{C}_1 \times \mathcal{M} \times \dots \times \mathcal{M} \times \mathcal{C}_\mu \times \{0, 1\}^* \rightarrow \{0, 1\},$$

and the success of  $\mathcal{A}$  is measured by checking if

$$\mathcal{V}(a_1, c_1, a_2, c_2, \dots, a_\mu, c_\mu, a_{\mu+1}) = 1 \quad \text{for } c_i = \text{ro}_i(a_1, \dots, a_i).$$

We write the above as

$$\mathcal{V}^{\vec{\text{ro}}}(a_1, a_2, \dots, a_\mu, a_{\mu+1}) = 1,$$

and the success probability of  $\mathcal{A}$  (as an attacker against the Fiat–Shamir transformed NIROP) is then given by

$$\epsilon^{\mathcal{V}}(\mathcal{A}) := \Pr(\mathcal{V}^{\vec{\text{ro}}}(\mathcal{A}(\vec{\mathcal{RO}})) = 1),$$

where  $\vec{\mathcal{RO}}$  is distributed uniformly at random over  $\vec{\mathcal{RO}}$ . The natural instantiation of  $\mathcal{V}$  is the verification predicate of the underlying interactive proof.

Below, we show the existence of a ( $\mathcal{V}$ -dependent, randomized) extractor for  $\mathcal{A}$ , denoted  $\mathcal{F}_\mu^{\mathcal{A}}$ , that aims to extract a tree  $T \in \text{TREE}_T^{\mathcal{V}}(a_1)$  of  $\mathcal{V}$ -accepting transcripts  $(a_1, c_1, \dots, a_\mu, c_\mu, a_{\mu+1})$  for some trunk  $a_1$ . Thus, the success probability of the extractor is given by  $\Pr(\mathcal{W}_\mu(\mathcal{F}_\mu^{\mathcal{A}}) = 1)$ , where  $\mathcal{W}_\mu(T)$  checks

<sup>14</sup> Alternatively, the  $\mu$  different random oracles could be implemented by (or modeled as) a single random oracle. The latter approach was followed in the multi-round Fiat–Shamir analysis of [AFK22]. Here, the above notation turned out to be more convenient.

whether  $T \in \text{TREE}_{\Gamma}^{\mathcal{V}}$ . Note that, without loss of generality, we may assume that  $\mathcal{F}_{\mu}^{\mathcal{A}}$ 's output is indeed such a tree, or  $\perp$  otherwise.

In order for the inductive analysis to go through, we need that  $\mathcal{F}_{\mu}$  additionally satisfies the following technical property, which resembles the property on  $\mathcal{A}$  necessary for Lemma 7 to be applicable (and this is of course no coincidence).

**Definition 12 (Early-Choice Property).** *An extractor  $\mathcal{F}_{\mu}^{\mathcal{A}}$  satisfies the early-choice property, if*

1. *it runs  $(a_1, \dots, a_{\mu+1}) \leftarrow \mathcal{A}(\vec{r\delta})$  as first step for uniformly random  $\vec{r\delta}$ , and*
2. *if it does not abort (i.e.,  $\mathcal{F}_{\mu}^{\mathcal{A}}$  does not output  $\perp$ ), it outputs  $T \in \text{TREE}_{\Gamma}^{\mathcal{V}}(a_1)$  with trunk  $a_1$ .*

The above early-choice property implies that, already after its first  $\mathcal{A}$ -invocation, the extractor  $\mathcal{F}_{\mu}^{\mathcal{A}}$  knows the trunk  $a_1$  of the tree it is going to output if it does not abort.

## 6.4 Recursive Extractor Construction

We show the existence of the extractor  $\mathcal{F}_{\mu}$  by means of a recursive construction. For  $\mu = 1$ ,  $\mathcal{F}_1^{\mathcal{A}}$  is simply the extractor from the 3-round case (Figure 2). The early-abort property is satisfied by construction.

For  $\mu > 1$ , we consider

$$\mathcal{A}[\text{ro}_1] : (\text{ro}_2, \dots, \text{ro}_{\mu}) \mapsto ((a_1, a_2), a_3, \dots, a_{\mu}, a_{\mu+1}) = \mathcal{A}(\text{ro}_1, \text{ro}_2, \dots, \text{ro}_{\mu})$$

which acts as  $\mathcal{A}$  but has the random oracle  $\text{ro}_1$  fixed, and where the output  $(a_1, a_2, a_3, \dots, a_{\mu}, z)$  of  $\mathcal{A}$  is parsed as indicated, i.e., the first two messages are combined so that  $\mathcal{A}[\text{ro}_1]$  is viewed as an algorithm that outputs  $\mu$  messages (instead of  $\mu + 1$ ). The algorithm  $\mathcal{A}[\text{ro}_1]$  comes with the obvious verification predicate  $\mathcal{V}[\text{ro}_1]$  that has  $\text{ro}_1$  fixed as well, i.e.,

$$\mathcal{V}[\text{ro}_1]((a_1, a_2), c_2, a_3, \dots, c_{\mu}, a_{\mu+1}) = \mathcal{V}(a_1, \text{ro}_1(a_1), a_2, c_2, a_3, \dots, c_{\mu}, a_{\mu+1}).$$

*Remark 8.* Note that, by exploiting our particular choice of the multi-round Fiat–Shamir transformation where the second challenge is computed as  $c_2 = \text{ro}_2(a_1, a_2)$ , it holds that

$$\mathcal{V}[\text{ro}_1]^{\text{ro}_2, \dots, \text{ro}_{\mu}}((a_1, a_2), a_3, \dots, a_{\mu+1}) = \mathcal{V}^{\text{ro}_1, \dots, \text{ro}_{\mu}}(a_1, a_2, a_3, \dots, a_{\mu+1}).$$

The above inequality would not hold if the second challenge would be computed differently, e.g., as  $\text{ro}_2(a_2)$ .

By applying recursion to  $\mathcal{A}[\text{ro}_1]$  and  $\mathcal{V}[\text{ro}_1]$ , there exists a  $\mathcal{V}[\text{ro}_1]$ -dependent, and thus  $\text{ro}_1$ -dependent, extractor

$$\mathcal{F}_{\mu-1}^{\mathcal{A}[\text{ro}_1]} \rightarrow T'$$

that (unless it aborts) outputs a tree  $T' \in \text{TREE}_{\Gamma_2, \dots, \Gamma_{\mu}}^{\mathcal{V}[\text{ro}_1]}(a_1, a_2)$  for some trunk  $(a_1, a_2)$ ; the corresponding ( $\text{ro}_1$ -dependent) verification predicate  $\mathcal{W}_{\mu-1}[\text{ro}_1]$  thus checks if the output of  $\mathcal{F}_{\mu-1}$  is indeed such a tree.

We then consider the (probabilistic) algorithm

$$\mathcal{B}^{\mathcal{A}} : \mathcal{C}_1^{\mathcal{M}} \rightarrow (\mathcal{M} \times \text{TREE}_{\Gamma_2, \dots, \Gamma_{\mu}}) \cup \{\perp\}, \quad \text{ro}_1 \mapsto \mathcal{F}_{\mu-1}^{\mathcal{A}[\text{ro}_1]},$$

that, on input  $\text{ro}_1$ , runs  $\mathcal{F}_{\mu-1}^{\mathcal{A}[\text{ro}_1]}$  and where, unless it aborts, the output  $T' = \{(a_1, a_2)\} \times T'' \in \text{TREE}_{\Gamma_2, \dots, \Gamma_{\mu}}(a_1, a_2)$  is parsed as  $(a_1, z) = (a_1, \{a_2\} \times T'')$ , where then  $\{a_2\} \times T'' \in \text{TREE}_{\Gamma_2, \dots, \Gamma_{\mu}}(a_2)$ . We note that, here and below, we merely do some obvious reformatting, but spelling it out is somewhat cumbersome.

Now note that the algorithm  $\mathcal{B}^{\mathcal{A}}$  is naturally equipped with a verification predicate, which we denote by  $\mathcal{V}_1$ . This predicate accepts a transcript  $(a_1, c_1, z) = (a_1, c_1, \{a_2\} \times T'')$  if  $\{a_2\} \times T'' \in \text{TREE}_{\Gamma_2, \dots, \Gamma_{\mu}}(a_2)$  and all transcripts  $(a_1, c_1, a_2, c_2, \dots) \in \{(a_1, c_1, a_2)\} \times T''$  are  $\mathcal{V}$ -accepting.

We note that, by construction, the success probability of  $\mathcal{B}^{\mathcal{A}}$  on input  $\text{ro}_1$  equals the success probability of the extractor  $\mathcal{F}_{\mu-1}^{\mathcal{A}[\text{ro}_1]}$ , i.e., formally,  $\mathcal{V}_1^{\text{ro}_1}(a_1, \{a_2\} \times T'') = \mathcal{W}_{\mu-1}[\text{ro}_1](\{(a_1, a_2)\} \times T'')$ .



Note that the algorithm  $\mathcal{B}^{\mathcal{A}}$  and its verification predicate  $\mathcal{V}_1$  are of precisely the form required by our base case extractor  $\mathcal{F}_1$ . For this reason, the extractor  $\mathcal{F}_\mu$  is given by

$$\mathcal{F}_\mu^{\mathcal{A}} := \mathcal{F}_1^{\mathcal{B}^{\mathcal{A}}},$$

with the early-abort strategy as specified in Figure 2. The early-abort strategy can be exploited in Lemma 7), using that, by induction,  $\mathcal{F}_{\mu-1}^{\mathcal{A}[\text{ro}_1]}$  has the early-choice property (see Lemma 8), and therefore  $\mathcal{B}^{\mathcal{A}}$  satisfies the requirement of having unit cost (measured in the number of queries to  $\mathcal{A}$ ) to compute  $a_1$ .

By construction,  $\mathcal{F}_1^{\mathcal{B}^{\mathcal{A}}}$  aims to output a  $\Gamma_1$ -tree of  $\mathcal{V}_1$ -accepting transcripts with a trunk  $a_1$  (but different challenges  $c_1$ ), which then forms a  $\Gamma$ -tree of  $\mathcal{V}$ -accepting transcripts.

## 6.5 Success Probability

We begin by inductively arguing that our extractor succeeds with the desired probability. First, we require a lemma bounding the  $q$ -value  $q(\mathcal{B}^{\mathcal{A}})$ , as defined in Definition 8. Recall that, for the base case  $\mu = 1$ , i.e., when considering Fiat–Shamir transformations of  $\Sigma$ -protocols, Lemma 6 shows that  $q(\mathcal{A}) \leq Q + 1$ , where  $Q$  denotes the number of random oracle queries made by  $\mathcal{A}$ . More generally, we wish to show that for any  $\mu$  and for  $\mathcal{B}^{\mathcal{A}}$  defined as above it holds that  $q(\mathcal{B}^{\mathcal{A}}) \leq Q + 1$ , where again  $Q$  denotes the query complexity of  $\mathcal{A}$ . In fact, this bound will be a consequence of Lemma 6, as our multi-round extractor satisfies the *early-choice property*.

**Lemma 8.** *For any  $\mu \geq 1$ , the extractor  $\mathcal{F}_\mu^{\mathcal{A}}$  ( $= \mathcal{F}_1^{\mathcal{B}^{\mathcal{A}}}$ ) satisfies the early-choice property of Definition 12.*

*Proof.* For  $\mu = 1$ ,  $\mathcal{F}_1^{\mathcal{A}}$  is simply the extractor from the 3-round case (Figure 2). The early-choice property is therefore satisfied by construction.

Via induction, the claim for  $\mu > 1$  then follows from transitivity. Namely, the extractor  $\mathcal{F}_1^{\mathcal{B}^{\mathcal{A}}}$  runs  $\mathcal{B}^{\mathcal{A}}$  as a first step, with  $\text{ro}_1$  sampled uniformly at random, and eventually outputs the obtained  $a_1$ . Further, by the induction hypothesis  $\mathcal{B}^{\mathcal{A}}(\text{ro}_1) = \mathcal{F}_{\mu-1}^{\mathcal{A}[\text{ro}_1]}$  runs  $\mathcal{A}$  as a first step, with the given  $\text{ro}_1$  and uniformly random  $\text{ro}_2, \dots, \text{ro}_\mu$ , and eventually outputs the obtained  $(a_1, a_2)$ . It thus follows that  $\mathcal{F}_1^{\mathcal{B}^{\mathcal{A}}}$  runs  $\mathcal{A}(\text{ro}_1, \dots, \text{ro}_\mu)$ , with uniformly random  $\text{ro}_1, \dots, \text{ro}_\mu$ , as a first step and eventually outputs the obtained  $a_1$ .  $\square$

**Lemma 9.** *For any algorithm  $\mathcal{A}: \mathcal{C}_1^{\mathcal{M}} \times \mathcal{C}_2^{\mathcal{M}^2} \times \dots \times \mathcal{C}_\mu^{\mathcal{M}^\mu} \rightarrow \mathcal{M}^\mu \times \{0, 1\}^*$  making at most  $Q$  queries (in total) to its random oracles, and for  $\mathcal{B}$  as in the above recursive construction,*

$$q(\mathcal{B}^{\mathcal{A}}) \leq Q + 1.$$

*Proof.* By definition,  $\mathcal{B}^{\mathcal{A}}$  takes only one random oracle  $\text{ro}_1$  as input. Further, by the early-choice property of  $\mathcal{F}_{\mu-1}$  (Lemma 8),  $\mathcal{B}^{\mathcal{A}}(\text{ro}_1) = \mathcal{F}_{\mu-1}^{\mathcal{A}[\text{ro}_1]}$  either aborts or its first output equals the first output of  $\mathcal{A}(\text{ro}_1, \text{ro}_2, \dots, \text{ro}_\mu)$ , for  $\text{ro}_2, \dots, \text{ro}_\mu$  sampled uniformly at random. In particular,  $\text{ro}_2, \dots, \text{ro}_\mu$  are sampled independently from the input  $\text{ro}_1$ , and thus uniquely determined by the random coins  $r$  of  $\mathcal{B}^{\mathcal{A}}$ . Let us now fix these random coins  $r$  and write  $\mathcal{B}^{\mathcal{A}}[r](\text{ro}_1)$  for the evaluation of  $\mathcal{B}^{\mathcal{A}}$  with random coins  $r$  and input  $\text{ro}_1$ .

Then, by the above, the first output of  $\mathcal{B}^{\mathcal{A}}[r](\text{ro}_1)$  equals  $\mathcal{A}(\text{ro}_1, \text{ro}_2^r, \dots, \text{ro}_\mu^r)$  for some fixed random oracles  $\text{ro}_2^r, \dots, \text{ro}_\mu^r$ . For this reason,

$$q(\mathcal{B}^{\mathcal{A}}[r]) \leq q(\mathcal{A}(\cdot, \text{ro}_2^r, \dots, \text{ro}_\mu^r)),$$

where  $\mathcal{A}(\cdot, \text{ro}_2^r, \dots, \text{ro}_\mu^r): \mathcal{C}_1^{\mathcal{M}} \rightarrow \mathcal{M}^\mu \times \{0, 1\}^*$  makes at most  $Q$  queries to the random oracles  $\text{ro}_1, \text{ro}_2^r, \dots, \text{ro}_\mu^r$ ; in particular, it makes at most  $Q$  queries to  $\text{ro}_1$ .<sup>15</sup> Finally, by Lemma 6, the right-hand side of the above inequality can thus be bounded by  $Q + 1$ . The lemma then follows by taking the expectation over the random coins  $r$ .  $\square$

We are now ready to provide a lower-bound for success probability of the recursive extractor.

<sup>15</sup> It is thus easily seen that, in the lemma statement, a bound  $Q$  on the queries to each random oracle (instead of bound on the total number of queries) would have sufficed.

**Proposition 1.** Let  $\mathcal{A}: \mathcal{C}_1^{\mathcal{M}} \times \mathcal{C}_2^{\mathcal{M}^2} \times \dots \times \mathcal{C}_\mu^{\mathcal{M}^\mu} \rightarrow \mathcal{M}^\mu \times \{0, 1\}^*$  be a  $Q$ -query algorithm. Then, the success probability of the extractor  $\mathcal{F}_\mu^{\mathcal{A}}$  satisfies

$$\Pr(\mathcal{F}_\mu^{\mathcal{A}} \neq \perp) = \Pr(\mathcal{W}_\mu(\mathcal{F}_\mu^{\mathcal{A}}) = 1) \geq \frac{\epsilon^{\mathcal{V}}(\mathcal{A}) - (Q+1) \cdot \kappa_{\mathbf{\Gamma}}}{1 - \kappa_{\mathbf{\Gamma}}}.$$

*Proof.* The proof is by induction. The base case  $\mu = 1$  follows from Theorem 2 and Lemma 6. Thus, consider now  $\mu > 1$ . First, setting  $\mathbf{\Gamma}' = (\Gamma_2, \dots, \Gamma_\mu)$  and using that  $\mathcal{B}^{\mathcal{A}}(\text{ro}_1) = \mathcal{F}_{\mu-1}^{\mathcal{A}[\text{ro}_1]}$ , we note that by the induction hypothesis

$$\Pr(\mathcal{V}_1^{\text{ro}_1}(\mathcal{B}^{\mathcal{A}}(\text{ro}_1)) = 1) = \Pr(\mathcal{W}_{\mu-1}[\text{ro}_1](\mathcal{F}_{\mu-1}^{\mathcal{A}[\text{ro}_1]}) = 1) \geq \frac{\epsilon^{\mathcal{V}[\text{ro}_1]}(\mathcal{A}[\text{ro}_1]) - (Q+1) \cdot \kappa_{\mathbf{\Gamma}'}}{1 - \kappa_{\mathbf{\Gamma}'}}.$$

Hence,

$$\begin{aligned} \epsilon^{\mathcal{V}_1}(\mathcal{B}^{\mathcal{A}}) &= \Pr(\mathcal{V}_1^{\text{RO}_1}(\mathcal{B}^{\mathcal{A}}(\text{RO}_1)) = 1) = \Pr(\mathcal{W}_{\mu-1}[\text{RO}_1](\mathcal{F}_{\mu-1}^{\mathcal{A}[\text{RO}_1]}) = 1) \\ &= \mathbb{E}_{\text{ro}_1} \left[ \Pr(\mathcal{W}_{\mu-1}[\text{ro}_1](\mathcal{F}_{\mu-1}^{\mathcal{A}[\text{ro}_1]}) = 1) \right] \\ &\geq \frac{\mathbb{E}_{\text{ro}_1} [\epsilon^{\mathcal{V}[\text{ro}_1]}(\mathcal{A}[\text{ro}_1])] - (Q+1) \cdot \kappa_{\mathbf{\Gamma}'}}{1 - \kappa_{\mathbf{\Gamma}'}}. \end{aligned}$$

By using that

$$\begin{aligned} \epsilon^{\mathcal{V}[\text{ro}_1]}(\mathcal{A}[\text{ro}_1]) &= \Pr[\mathcal{V}[\text{ro}_1]^{\text{RO}_2, \dots, \text{RO}_\mu}(\mathcal{A}[\text{ro}_1](\text{RO}_2, \dots, \text{RO}_\mu)) = 1] \\ &= \Pr[\mathcal{V}^{\text{ro}_1, \text{RO}_2, \dots, \text{RO}_\mu}(\mathcal{A}(\text{ro}_1, \text{RO}_2, \dots, \text{RO}_\mu)) = 1], \end{aligned}$$

where the second equality exploits the particular definition of the Fiat–Shamir transformation (see Remark 8), we see that  $\mathbb{E}_{\text{ro}_1} [\epsilon^{\mathcal{V}[\text{ro}_1]}(\mathcal{A}[\text{ro}_1])] = \epsilon^{\mathcal{V}}(\mathcal{A})$ . Hence,

$$\epsilon^{\mathcal{V}_1}(\mathcal{B}^{\mathcal{A}}) \geq \frac{\epsilon^{\mathcal{V}}(\mathcal{A}) - (Q+1) \cdot \kappa_{\mathbf{\Gamma}'}}{1 - \kappa_{\mathbf{\Gamma}'}}.$$

Furthermore, by Theorem 2, exploiting the bound on the  $q$ -value from Lemma 9,

$$\Pr(\mathcal{F}_\mu^{\mathcal{A}} \neq \perp) = \Pr(\mathcal{F}_1^{\mathcal{B}^{\mathcal{A}}} \neq \perp) \geq \frac{\epsilon^{\mathcal{V}_1}(\mathcal{B}^{\mathcal{A}}) - (Q+1) \cdot \kappa_{\Gamma_1}}{1 - \kappa_{\Gamma_1}}.$$

Thus, writing  $\epsilon$  for  $\epsilon^{\mathcal{V}}(\mathcal{A})$  and exploiting that  $(1 - \kappa_{\Gamma_1})(1 - \kappa_{\mathbf{\Gamma}'}) = 1 - \kappa_{\mathbf{\Gamma}}$ ,

$$\begin{aligned} \Pr(\mathcal{F}_\mu^{\mathcal{A}} \neq \perp) &\geq \frac{\frac{\epsilon - (Q+1) \cdot \kappa_{\mathbf{\Gamma}'}}{1 - \kappa_{\mathbf{\Gamma}'}} - (Q+1) \cdot \kappa_{\Gamma_1}}{1 - \kappa_{\Gamma_1}} \\ &= \frac{\epsilon - (Q+1) \cdot \kappa_{\mathbf{\Gamma}'} - (Q+1) \cdot \kappa_{\Gamma_1} \cdot (1 - \kappa_{\mathbf{\Gamma}'})}{1 - \kappa_{\mathbf{\Gamma}}} \\ &= \frac{\epsilon - (Q+1)(\kappa_{\mathbf{\Gamma}'} + \kappa_{\Gamma_1} \cdot (1 - \kappa_{\mathbf{\Gamma}'}))}{1 - \kappa_{\mathbf{\Gamma}}} \\ &= \frac{\epsilon - (Q+1)(1 - (1 - \kappa_{\mathbf{\Gamma}'})(1 - \kappa_{\Gamma_1}))}{1 - \kappa_{\mathbf{\Gamma}}} \\ &= \frac{\epsilon - (Q+1) \cdot \kappa_{\mathbf{\Gamma}}}{1 - \kappa_{\mathbf{\Gamma}}}, \end{aligned}$$

which completes the proof. □

## 6.6 Running Time Analysis

Let us now analyze the expected running time of our knowledge extractor.

**Proposition 2.** *Let  $\Gamma_1, \dots, \Gamma_\mu$  be monotone structures and let  $\Gamma = (\Gamma_1, \dots, \Gamma_\mu)$ . Let  $\mathcal{A}: \overrightarrow{\mathcal{RO}} \rightarrow \mathcal{M}^\mu \times \{0, 1\}^*$  be a  $Q$ -query algorithm, and let the extractor  $\mathcal{F}_\mu$  be defined as above. Let the random variable  $X_\Gamma$  denote the number of  $\mathcal{A}$ -queries extractor  $\mathcal{F}_\mu^{\mathcal{A}}$  makes. Then*

$$\mathbb{E}[X_\Gamma] \leq \frac{T_\Gamma + (T_\Gamma - 1)Q}{1 - \kappa_\Gamma},$$

where  $T_\Gamma = \prod_{j=1}^\mu (t_{\Gamma_j} + 1)$ .

*Proof.* The proof is by induction over  $\mu$ . The base case  $\mu = 1$  follows from Theorem 2 together with the bound  $q(\mathcal{A}) \leq Q + 1$  from Lemma 6. These results namely state that the expected number of  $\mathcal{A}$ -queries made by  $\mathcal{F}_1^{\mathcal{A}} = \mathcal{F}_{\Gamma_1}^{\mathcal{A}}$  is at most

$$\begin{aligned} 1 + q(\mathcal{A}) \cdot t_{\Gamma_1} \cdot (1 + \kappa_{\Gamma_1}) &\leq 1 + \frac{q(\mathcal{A}) \cdot t_{\Gamma_1}}{1 - \kappa_{\Gamma_1}} \leq \frac{1 + q(\mathcal{A}) \cdot t_{\Gamma_1}}{1 - \kappa_{\Gamma_1}} \\ &\leq \frac{1 + (Q + 1) \cdot t_{\Gamma_1}}{1 - \kappa_{\Gamma_1}} = \frac{t_{\Gamma_1} + 1 + t_{\Gamma_1}Q}{1 - \kappa_{\Gamma_1}} = \frac{T_{\Gamma_1} + (T_{\Gamma_1} - 1)(Q + 1)}{1 - \kappa_{\Gamma_1}}, \end{aligned}$$

where we use that  $1/(1 - \kappa_1) \geq 1 + \kappa_1 \geq 1$ . We note that this bound could also have been obtained by using the refined running time analysis of Lemma 7, then using that  $\theta(\text{ro}) = 1$  is constant.

For the induction step, let  $\Gamma' = (\Gamma_2, \dots, \Gamma_\mu)$  and thus  $T_{\Gamma'} = \prod_{j=2}^\mu (t_{\Gamma_j} + 1)$ . By construction,  $\mathcal{F}_\mu^{\mathcal{A}} = \mathcal{F}_1^{\mathcal{B}^{\mathcal{A}}}$  where  $\mathcal{B}^{\mathcal{A}}(\text{ro}_1) = \mathcal{F}_{\mu-1}^{\mathcal{A}[\text{ro}_1]}$ , up to some reformatting of the output.

By the early-choice property of  $\mathcal{F}_{\mu-1}^{\mathcal{A}[\text{ro}_1]}$ , the algorithm  $\mathcal{B}^{\mathcal{A}}$  satisfies the requirement on  $\mathcal{A}$  in Lemma 7, with the cost  $\theta$  of  $\mathcal{B}^{\mathcal{A}}$  measuring the expected number of calls to  $\mathcal{A}$ . Indeed, finding  $a_1$  as in  $(a_1, z) \leftarrow \mathcal{B}^{\mathcal{A}}(\text{ro}_1)$  can be done with one query to  $\mathcal{A}$ , while finding a subtree requires a full run of  $\mathcal{F}_{\mu-1}^{\mathcal{A}[\text{ro}_1]}$ , and thus costs (in expectation)  $\theta(\text{ro}_1)$  calls to  $\mathcal{A}$ , where, by induction hypothesis,

$$\mathbb{E}[\theta(\text{RO}_1)] = \mathbb{E}[X_{\Gamma'}] \leq \frac{T_{\Gamma'} + (T_{\Gamma'} - 1)Q}{1 - \kappa_{\Gamma'}}.$$

Thus, by Lemma 7, and using the bound on  $q(\mathcal{B}^{\mathcal{A}})$  from Lemma 9, we find

$$\begin{aligned} \mathbb{E}[X_\Gamma] &\leq \mathbb{E}[\theta(\text{RO}_1)] + t_{\Gamma_1} \cdot \frac{\mathbb{E}[\theta(\text{RO}_1)] - 1 + q(\mathcal{B}^{\mathcal{A}})}{1 - \kappa_{\Gamma_1}} \\ &\leq \frac{(t_{\Gamma_1} + 1) \mathbb{E}[\theta(\text{RO}_1)] + t_{\Gamma_1}Q}{1 - \kappa_{\Gamma_1}} \\ &\leq \frac{(t_{\Gamma_1} + 1) \left( \frac{T_{\Gamma'} + (T_{\Gamma'} - 1)Q}{1 - \kappa_{\Gamma'}} \right) + t_{\Gamma_1}Q}{1 - \kappa_{\Gamma_1}} \\ &\leq \frac{(t_{\Gamma_1} + 1)(T_{\Gamma'} + (T_{\Gamma'} - 1)Q) + t_{\Gamma_1}Q}{1 - \kappa_\Gamma} \\ &= \frac{(t_{\Gamma_1} + 1)T_{\Gamma'} + ((t_{\Gamma_1} + 1)(T_{\Gamma'} - 1) + t_{\Gamma_1})Q}{1 - \kappa_\Gamma} \\ &= \frac{T_\Gamma + (T_\Gamma - 1)Q}{1 - \kappa_\Gamma}, \end{aligned}$$

which completes the proof.  $\square$

## 6.7 Knowledge Soundness of the Fiat–Shamir Transformation

The following theorem summarizes the properties of the recursive extraction algorithm defined in Section 6.4. In fact, from the success probability bound derived in Section 6.5 and the running time bound derived in Section 6.6, it immediately follows that the Fiat–Shamir transformation of a  $(\Gamma_1, \dots, \Gamma_\mu)$ -out-of- $(\mathcal{C}_1, \dots, \mathcal{C}_\mu)$  special-sound interactive proof is knowledge sound (under certain conditions). Moreover, the knowledge error only grows linearly in the query complexity  $Q$  of dishonest provers attacking the non-interactive proof, and is independent of the number of rounds. This is the main result of this paper.

**Theorem 4 (Knowledge Soundness - Fiat–Shamir Transformation of Multi-Round Interactive Proofs).** Let  $(\Gamma_i, \mathcal{C}_i)$  be nonempty monotone structures, and let  $\Gamma := (\Gamma_1, \dots, \Gamma_\mu)$  and  $\mathcal{C} := (\mathcal{C}_1, \dots, \mathcal{C}_\mu)$ . Let  $\Pi$  be a  $\Gamma$ -out-of- $\mathcal{C}$  special-sound interactive proof. Suppose that  $T_\Gamma = \prod_{i=1}^\mu (t_{\Gamma_i} + 1)$  is polynomial in the size  $|x|$  of the statement  $x$  and that, for all  $i$ , sampling from  $\mathcal{U}_{\Gamma_i}(S_i)$  takes polynomial time (in  $|x|$ ) for all  $S_i \subseteq \mathcal{C}_i$  with  $|S_i| \leq t_{\Gamma_i}$ . Then the Fiat–Shamir transformation  $\text{FS}[\Pi]$  of  $\Pi$  is knowledge sound with knowledge error

$$(Q + 1) \cdot \kappa_\Gamma,$$

where  $\kappa_\Gamma = 1 - \prod_{i=1}^\mu (1 - \kappa_{\Gamma_i})$  is the knowledge error<sup>16</sup> of the interactive proof  $\Pi$  and  $Q$  denotes the query complexity of a prover attacking  $\text{FS}[\Pi]$ .

More precisely,  $\text{FS}[\Pi]$  admits a knowledge extractor that, on input a statement  $x$  and given oracle access to a  $Q$ -query prover  $\mathcal{P}^*$  makes at most

$$\frac{T_\Gamma + (T_\Gamma - 1)Q}{1 - \kappa_\Gamma}$$

queries to  $\mathcal{P}^*$  in expectation, and succeeds to extract a witness for  $x$  with probability at least

$$\frac{\epsilon(\mathcal{P}^*, x) - (Q + 1) \cdot \kappa_\Gamma}{1 - \kappa_\Gamma}.$$

*Proof.* This follows immediately from Propositions 1 and 2. □

## Acknowledgments

The first author has been supported by TNO’s Early Research Program - Next Generation Cryptography, the Vraaggestuurd Programma Cyber Security & Resilience, part of the Dutch Top Sector High Tech Systems and Materials program, and by the Smart Networks and Services Joint Undertaking (SNS JU) under the European Union’s Horizon Europe Research and Innovation Programme in the scope of the CONFIDENTIAL6G project under Grant Agreement 101096435. The contents of this publication are the sole responsibility of the authors and do not in any way reflect the views of the EU. The third author has been supported by Helsinki Institute for Information Technology (HIIT).

## References

- ACK21. Thomas Attema, Ronald Cramer, and Lisa Kohl. A compressed  $\Sigma$ -protocol theory for lattices. In Tal Malkin and Chris Peikert, editors, *CRYPTO*, volume 12826 of *Lecture Notes in Computer Science*, pages 549–579. Springer, 2021.
- ACY23. Gal Arnon, Alessandro Chiesa, and Eylon Yogev. Iops with inverse polynomial soundness error. *IACR Cryptol. ePrint Arch.*, page 1062, 2023.
- AF22. Thomas Attema and Serge Fehr. Parallel repetition of  $(k_1, \dots, k_\mu)$ -special-sound multi-round interactive proofs. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO*, volume 13507 of *Lecture Notes in Computer Science*, pages 415–443. Springer, 2022.
- AFK22. Thomas Attema, Serge Fehr, and Michael Kloof. Fiat-shamir transformation of multi-round interactive proofs. In Eike Kiltz and Vinod Vaikuntanathan, editors, *Theory of Cryptography Conference (TCC)*, volume 13747 of *Lecture Notes in Computer Science*, pages 113–142. Springer, 2022.
- AFK23. Thomas Attema, Serge Fehr, and Michael Kloof. Fiat-shamir transformation of multi-round interactive proofs (extended version). *Journal of Cryptology*, 36(4):36, 2023.
- AFR23. Thomas Attema, Serge Fehr, and Nicolas Resch. Generalized special-sound interactive proofs and their knowledge soundness. In Guy N. Rothblum and Hoeteck Wee, editors, *Theory of Cryptography - 21st International Conference, TCC 2023, Taipei, Taiwan, November 29 - December 2, 2023, Proceedings, Part III*, volume 14371 of *Lecture Notes in Computer Science*, pages 424–454. Springer, 2023.
- ALS20. Thomas Attema, Vadim Lyubashevsky, and Gregor Seiler. Practical product proofs for lattice commitments. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO*, volume 12171 of *Lecture Notes in Computer Science*, pages 470–499. Springer, 2020.
- Att23. Thomas Attema. *Compressed  $\Sigma$ -Protocol Theory*. PhD thesis, Leiden University, 2023.

<sup>16</sup> Recall that  $\kappa_{\Gamma_i} = \max_{S_i \notin \Gamma_i} |S_i| / |\mathcal{C}|$ .

- BBB<sup>+</sup>18. Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Gregory Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *IEEE Symposium on Security and Privacy (S&P)*, pages 315–334. IEEE Computer Society, 2018.
- BBC<sup>+</sup>18. Carsten Baum, Jonathan Bootle, Andrea Cerulli, Rafaël del Pino, Jens Groth, and Vadim Lyubashevsky. Sub-linear lattice-based zero-knowledge arguments for arithmetic circuits. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO*, volume 10992 of *Lecture Notes in Computer Science*, pages 669–699. Springer, 2018.
- BBHR18. Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Fast reed-solomon interactive oracle proofs of proximity. In Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella, editors, *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic*, volume 107 of *LIPICs*, pages 14:1–14:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
- BCC<sup>+</sup>16. Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Jens Groth, and Christophe Petit. Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT*, volume 9666 of *Lecture Notes in Computer Science*, pages 327–357. Springer, 2016.
- BPW12. David Bernhard, Olivier Pereira, and Bogdan Warinschi. How not to prove yourself: Pitfalls of the fiat-shamir heuristic and applications to helios. In Xiaoyun Wang and Kazue Sako, editors, *Advances in Cryptology - ASIACRYPT 2012 - 18th International Conference on the Theory and Application of Cryptology and Information Security, Beijing, China, December 2-6, 2012. Proceedings*, volume 7658 of *Lecture Notes in Computer Science*, pages 626–643. Springer, 2012.
- CD98. Ronald Cramer and Ivan Damgård. Zero-knowledge proofs for finite field arithmetic; or: Can zero-knowledge be for free? In Hugo Krawczyk, editor, *CRYPTO*, volume 1462 of *Lecture Notes in Computer Science*, pages 424–441. Springer, 1998.
- Cra96. Ronald Cramer. *Modular Design of Secure yet Practical Cryptographic Protocols*. PhD thesis, CWI and University of Amsterdam, 1996.
- ENS20. Muhammed F. Esgin, Ngoc Khanh Nguyen, and Gregor Seiler. Practical exact proofs from lattices: New techniques to exploit fully-splitting rings. In Shiho Moriai and Huaxiong Wang, editors, *Advances in Cryptology - ASIACRYPT 2020 - 26th International Conference on the Theory and Application of Cryptology and Information Security, Daejeon, South Korea, December 7-11, 2020, Proceedings, Part II*, volume 12492 of *Lecture Notes in Computer Science*, pages 259–288. Springer, 2020.
- KL23. Michael Kloß and Russel Lai. Personal communication: Adaptive special-soundness, 2023.
- LNP22. Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Maxime Plançon. Lattice-based zero-knowledge proofs and applications: Shorter, simpler, and more general. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology - CRYPTO 2022 - 42nd Annual International Cryptology Conference, CRYPTO 2022, Santa Barbara, CA, USA, August 15-18, 2022, Proceedings, Part II*, volume 13508 of *Lecture Notes in Computer Science*, pages 71–101. Springer, 2022.
- LW22. George Lu and Brent Waters. How to sample a discrete gaussian (and more) from a random oracle. In Eike Kiltz and Vinod Vaikuntanathan, editors, *Theory of Cryptography - 20th International Conference, TCC 2022, Chicago, IL, USA, November 7-10, 2022, Proceedings, Part II*, volume 13748 of *Lecture Notes in Computer Science*, pages 653–682. Springer, 2022.
- RR22. Noga Ron-Zewi and Ron D. Rothblum. Proving as fast as computing: succinct arguments with constant prover overhead. In Stefano Leonardi and Anupam Gupta, editors, *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, pages 1353–1363. ACM, 2022.
- Wik18. Douglas Wikström. Special soundness revisited. *IACR Cryptology ePrint Archive*, 2018.
- Wik21. Douglas Wikström. Special soundness in the random oracle model. *IACR Cryptology ePrint Archive*, 2021.

# Appendix

## A Definition - Fiat–Shamir Transformation

We provide a formal definition of the Fiat–Shamir transform, as discussed in Section 2.3. For simplicity, we assume access to multiple independent random oracles  $\text{ro}_1, \dots, \text{ro}_\mu$ , where  $\text{ro}_i$  maps tuples of  $i$  prover messages to the appropriate challenge space.

**Definition 13 (Static Fiat–Shamir Transformation).** *Let  $\Pi = (\mathcal{P}, \mathcal{V})$  be a  $2\mu + 1$ -move public-coin interactive proof for a relation  $R$  with  $i$ -th challenge set  $\mathcal{C}_i$ , and such that all potential prover messages are contained in  $\mathcal{M}$ . Let  $\text{ro}_1, \dots, \text{ro}_\mu$  be independent random oracles such that  $\text{ro}_i: \mathcal{M}^i \rightarrow \mathcal{C}_i$  for all  $i = 1, \dots, \mu$ .*

*The static Fiat–Shamir transformation  $\text{FS}[\Pi] = (\mathcal{P}_{\text{fs}}, \mathcal{V}_{\text{fs}})$  of the protocol  $\Pi$  is the following NIROP. Given  $(x; w) \in R$ , the NIROP prover  $\mathcal{P}_{\text{fs}}^{\text{ro}_1, \dots, \text{ro}_\mu}$  simulates the interactive proof prover  $\mathcal{P}(x; w)$  such that, after it outputs the  $i$ -th message  $a_i$ , it is provided with the challenge*

$$c_i = \text{ro}_i(a_1, \dots, a_{i-1}, a_i) \tag{15}$$

*for all  $i$ . The prover  $\mathcal{P}_{\text{fs}}^{\text{ro}_1, \dots, \text{ro}_\mu}$  then outputs the proof  $\pi = (a_1, a_2, \dots, a_{\mu+1})$ . On input a statement  $x$  and a proof  $\pi = (a_1, a_2, \dots, a_{\mu+1})$ , the verifier  $\mathcal{V}_{\text{fs}}^{\text{ro}_1, \dots, \text{ro}_\mu}$  accepts if and only if  $\mathcal{V}$  accepts the transcript  $(a_1, c_1, a_2, \dots, a_\mu, c_\mu, a_{\mu+1})$  on input  $x$ , where each  $c_i$  is computed as in (15).*

*Remark 9 (Adaptive Fiat–Shamir Transformation).* In the adaptive setting, a dishonest prover may adaptively choose the statement  $x$ . Thus, in the adaptive Fiat–Shamir transformation, the challenges from (15) are instead computed as

$$c_i = \text{ro}_i(x, a_1, \dots, a_{i-1}, a_i) \tag{16}$$

for  $i = 1, \dots, \mu$ . Proving and verification proceed mutatis mutandis as before.

## B Analysis of the Extractor for Interactive Proofs

This appendix provides the proof of Theorem 1, which describes the properties of the improved extractor for  $\Gamma$ -out-of- $\mathcal{C}$  special-sound  $\Sigma$ -protocols of [KL23]. We have merely adapted the results from [KL23] to our notation. Before we prove Theorem 1, we present two auxiliary lemmas. Further, this section makes use of the notation introduced in Section 4.

The following lemma provides an useful bound for a specific composition of consecutive  $\Delta$ -values  $\Delta(\mathbf{c}_0), \dots, \Delta(\mathbf{c}_k)$ , as defined in Equation (6). This lemma crucially relies on the fact that  $U_i \implies U_{i-1}$  for all  $i$  (Equation (5)). Recall that  $U_i$  denotes the event  $C \in \mathcal{U}_\Gamma(\mathbf{c}_i)$ , where  $C$  is uniformly distributed in  $\mathcal{C}$ . To give some intuition, we note that for a trivial  $V$  that occurs with certainty, so that  $\Delta(\mathbf{c}_k) = \Pr(\neg U_k \mid U_{k-1})$ , the inequality becomes the (trivial) equality

$$\Pr(U_k) = \Pr(U_k \wedge \dots \wedge U_1) = \prod_{i=1}^k \Pr(U_i \mid U_1 \wedge \dots \wedge U_{i-1}) = \prod_{i=1}^k \Pr(U_i \mid U_{i-1}).$$

For a non-trivial  $V$ , the proof is less straightforward.

**Lemma 10.** *Let  $\mathbf{c} \in \mathcal{C}^k$  for some  $k$  and let  $\mathbf{c}_i = (c_1, \dots, c_i)$  for all  $1 \leq i \leq k$ . Then*

$$\prod_{i=1}^k (1 - \Delta(\mathbf{c}_i)) \leq 1 - \Pr(V \wedge \neg U_k).$$

*Proof.* The proof of the lemma will now proceed inductively over  $k$ . The lemma trivially holds for the base case  $k = 1$ . So let us assume the lemma is proven for  $k' = k - 1$ . Then, by the induction hypothesis

$$\begin{aligned}
\prod_{i=1}^k (1 - \Delta(\mathbf{c}_i)) &\leq (1 - \Pr(V \wedge \neg U_{k-1})) \cdot (1 - \Delta(\mathbf{c}_k)) \\
&= (1 - \Pr(V \wedge \neg U_{k-1})) \cdot (1 - \Pr(V \wedge \neg U_k \mid U_{k-1})) \\
&= \Pr(\neg V \vee U_{k-1}) \cdot \Pr(\neg V \vee U_k \mid U_{k-1}) \\
&= (\Pr(U_{k-1}) + \Pr(\neg V \wedge \neg U_{k-1})) \cdot \Pr(\neg V \vee U_k \mid U_{k-1}) \\
&\leq \Pr((\neg V \vee U_k) \wedge U_{k-1}) + \Pr(\neg V \wedge \neg U_{k-1}) \\
&= \Pr((\neg V \wedge U_{k-1}) \vee (U_k \wedge U_{k-1})) + \Pr(\neg V \wedge \neg U_{k-1}) \\
&= \Pr((\neg V \wedge U_{k-1}) \vee U_k \vee (\neg V \wedge \neg U_{k-1})) \\
&= \Pr(\neg V \vee U_k) \\
&= 1 - \Pr(V \wedge \neg U_k),
\end{aligned}$$

which completes the proof of the lemma.  $\square$

The following auxiliary lemma provides a lower bound in terms of  $\mathcal{A}$ 's success probability  $\epsilon(\mathcal{A})$  and the knowledge error  $\kappa_\Gamma$ . The left-hand side of this inequality will appear in the analysis of the success probability of the knowledge extractor.

**Lemma 11.** *Let  $\mathbf{c}_k \in \mathcal{C}^k$  for some  $k$  with  $\mathbf{c}_k \notin \Gamma$  and let  $\mathbf{c}_i = (c_1, \dots, c_i)$  for all  $1 \leq i \leq k$ . Then*

$$\delta(\mathbf{c}_0) \prod_{i=1}^k \frac{\tilde{\delta}(\mathbf{c}_i)}{\delta(\mathbf{c}_{i-1})} = \tilde{\delta}(\mathbf{c}_1) \prod_{i=2}^k \frac{\tilde{\delta}(\mathbf{c}_i)}{\delta(\mathbf{c}_{i-1})} \geq \frac{\epsilon^{\mathcal{V}}(\mathcal{A}) - \kappa_\Gamma}{1 - \kappa_\Gamma},$$

where  $\kappa_\Gamma := \max_{S \notin \Gamma} \frac{|S|}{|\mathcal{C}|}$ .

*Proof.* First note that, for all  $1 \leq i \leq k$ ,

$$\begin{aligned}
\delta(\mathbf{c}_{i-1}) - \Delta(\mathbf{c}_i) &= \Pr(V \mid U_{i-1}) - \Pr(V \wedge \neg U_i \mid U_{i-1}) = \Pr(V \wedge U_i \mid U_{i-1}) \\
&= \Pr(V \mid U_i \wedge U_{i-1}) \Pr(U_i \mid U_{i-1}) = \delta(\mathbf{c}_i) \cdot \Pr(U_i \mid U_{i-1}),
\end{aligned}$$

where we use that  $U_i \implies U_{i-1}$ .

Hence,

$$\tilde{\delta}(\mathbf{c}_i) = \frac{\delta(\mathbf{c}_{i-1}) - \Delta(\mathbf{c}_i)}{1 - \Delta(\mathbf{c}_i)} = \delta(\mathbf{c}_i) \frac{\Pr(U_i \mid U_{i-1})}{1 - \Delta(\mathbf{c}_i)},$$

and

$$\begin{aligned}
\tilde{\delta}(\mathbf{c}_1) \prod_{i=2}^k \frac{\tilde{\delta}(\mathbf{c}_i)}{\delta(\mathbf{c}_{i-1})} &= \delta(\mathbf{c}_k) \prod_{i=1}^k \frac{\Pr(U_i \mid U_{i-1})}{1 - \Delta(\mathbf{c}_i)} \\
&= \delta(\mathbf{c}_k) \Pr(U_k \mid U_0) \prod_{i=1}^k \frac{1}{1 - \Delta(\mathbf{c}_i)} \\
&= \delta(\mathbf{c}_k) \Pr(U_k) \prod_{i=1}^k \frac{1}{1 - \Delta(\mathbf{c}_i)} \\
&= \Pr(V \wedge U_k) \prod_{i=1}^k \frac{1}{1 - \Delta(\mathbf{c}_i)},
\end{aligned} \tag{17}$$

where the second equality follows again because  $U_j \implies U_i$  for all  $i \leq j$ . By Lemma 10, it thus follows that

$$\tilde{\delta}(\mathbf{c}_1) \prod_{i=2}^k \frac{\tilde{\delta}(\mathbf{c}_i)}{\delta(\mathbf{c}_{i-1})} \geq \frac{\Pr(V \wedge U_k)}{1 - \Pr(V \wedge \neg U_k)} = \frac{\Pr(V) - \Pr(V \wedge \neg U_k)}{1 - \Pr(V \wedge \neg U_k)}. \tag{18}$$

Now note that, since  $\mathcal{C} \setminus \mathcal{U}_\Gamma(\mathbf{c}) \notin \Gamma$  for all  $\mathbf{c} \notin \Gamma$  (see Equation (3)),

$$\Pr(V \wedge \neg U_k) \leq \Pr(\neg U_k) = \Pr(C \notin \mathcal{U}_\Gamma(\mathbf{c}_k)) \leq \max_{S \notin \Gamma} \frac{|S|}{|\mathcal{C}|} = \kappa_\Gamma.$$

Hence, by Equation (18) and the monotonicity (decreasing) of the function  $x \mapsto \frac{q-x}{1-x}$  for all  $0 \leq q \leq 1$ , it follows that

$$\tilde{\delta}(\mathbf{c}_1) \prod_{i=2}^k \frac{\tilde{\delta}(\mathbf{c}_i)}{\delta(\mathbf{c}_{i-1})} \geq \frac{\Pr(V) - \kappa_\Gamma}{1 - \kappa_\Gamma} \geq \frac{\epsilon^\mathcal{V}(\mathcal{A}) - \kappa_\Gamma}{1 - \kappa_\Gamma},$$

which completes the proof of the lemma.  $\square$

We are not ready to present the proof of Theorem 1.

*Proof of Theorem 1.* The extractor  $\mathcal{E}_\Gamma^{\mathcal{A}}$  is formally defined in Figure 1. It iteratively tries to find new challenges  $c_i$ , together with outputs  $y_i$ , such that  $\mathcal{V}(c_i, y_i) = 1$ , until it has collected a subset  $\{c_1, \dots, c_k\} \in \Gamma$ . Further, for  $i \geq 2$ , every  $c_i$  is sampled from  $\mathcal{U}_\Gamma(\{c_1, \dots, c_{i-1}\})$ , i.e., for all extractor outputs it holds that  $k \leq t_\Gamma + 1$ . Note that,  $k = t_\Gamma + 1$  can only occur if the first challenge  $c_1$  is in  $\mathcal{C} \setminus \mathcal{U}_\Gamma(\emptyset)$ , i.e., if  $c_1$  is useless. Further, only in contrived interactive proofs  $\mathcal{C} \neq \mathcal{U}_\Gamma(\emptyset)$ , hence typically it even holds that  $k \leq t_\Gamma$ . See also Remark 3.

As before, for any  $\mathbf{c} = (c_1, \dots, c_k) \in \mathcal{C}^k$  and  $1 \leq i \leq k$ , we write  $\mathbf{c}_i = (c_1, \dots, c_i)$ . Let us now analyze the success probability and the expected number of  $\mathcal{A}$ -queries of the extractor.

**Success Probability.** Let us write  $t = t_\Gamma$ . Let the random variable  $C_i$  denote the  $i$ -th successful challenge found by the extractor, where we let  $C_i = \perp$  if the extractor finishes before finding  $i$  challenges, i.e.,  $C_i$  has support in  $\mathcal{C} \cup \{\perp\}$ . Note that if  $(C_1, \dots, C_i) \in \Gamma$  then  $C_{i+1} = \dots = C_t = \perp$ . Vice versa,  $C_i = \perp$  implies that the extractor was either successful before the  $i$ -th iteration, or that it aborted and failed before the  $i$ -th iteration. Further, we write  $\mathbf{C}_i = (C_1, \dots, C_i)$  and  $\mathbf{C} = \mathbf{C}_{t+1}$ . Additionally, for  $k \in \mathbb{N}$ , we define

$$\mathcal{G}_k = \{\mathbf{c} \in (\mathcal{C} \cup \{\perp\})^{t+1} : \Pr(\mathbf{C} = \mathbf{c}) > 0 \wedge \mathbf{c}_{k-1} \notin \Gamma \wedge \mathbf{c}_k \in \Gamma\}$$

and then set  $\mathcal{G} = \mathcal{G}_1 \cup \dots \cup \mathcal{G}_{t+1}$ . Note that,  $\mathcal{G}_k$  contains those challenge sequences  $\mathbf{c}$  which occur during a *successful* run of  $\mathcal{E}$  and for which  $\mathbf{c}_k \in \Gamma$  for the first time. Moreover, for all  $\mathbf{c} \in \mathcal{G}_k$ , it holds that  $c_i \neq \perp$  for all  $i \leq k$  and  $c_i = \perp$  for all  $i > k$ . Further, observe that  $\mathcal{G}$  is a disjoint union, and that the extractor succeeds if and only if  $\mathbf{C} \in \mathcal{G}$ .

We make the following case distinction. If  $T := \{c \in \mathcal{C} : \Pr(\mathcal{V}(c, \mathcal{A}(c)) = 1) > 0\} \notin \Gamma$  then

$$0 = \Pr(V \mid C \notin T) \geq \Pr(V) - \Pr(C \in T) \geq \epsilon^\mathcal{V}(\mathcal{A}) - \kappa_\Gamma,$$

i.e.,  $\epsilon^\mathcal{V}(\mathcal{A}) \leq \kappa_\Gamma$ . So, in this case, the extractor trivially succeeds with the claimed probability. In the remainder, we thus consider the case where  $T \in \Gamma$ , which in particular implies that for all  $i$  and all  $\mathbf{c}_i \in \mathcal{C}^i$  with  $\Pr(\mathbf{C}_i = \mathbf{c}_i) > 0$  and  $\mathbf{c}_i \notin \Gamma$  it holds that  $\Pr(C_{i+1} \neq \perp \mid \mathbf{C}_i = \mathbf{c}_i) > 0$ , and so conditioning on the event " $C_{i+1} \neq \perp \wedge \mathbf{C}_i = \mathbf{c}_i$ " is well defined then.

Towards analyzing the success probability  $\Pr(\mathbf{C} \in \mathcal{G})$  of the extractor  $\mathcal{E}$  in Fig. 1, it will be convenient to consider the variant  $\mathcal{E}'$  of  $\mathcal{E}$  that is not stopped by the coin, but keeps on searching for a next challenge  $c_{i+1} \in \mathcal{U}_\Gamma(\mathbf{c}_i)$  for which  $\mathcal{A}$  succeeds. Formally, for  $i \in \{1, \dots, t\}$  and  $\mathbf{c}_i \in \mathcal{C}^i$  with  $\Pr(\mathbf{C}_i = \mathbf{c}_i) > 0$  and  $\mathbf{c}_i \notin \Gamma$ ,  $\mathcal{E}'$  finds  $c_{i+1} \in \mathcal{C}$  with probability

$$\tau(c_{i+1} \mid \mathbf{c}_i) := \Pr(C_{i+1} = c_{i+1} \mid C_{i+1} \neq \perp \wedge \mathbf{C}_i = \mathbf{c}_i).$$

In all other cases, we set  $\tau(c_{i+1} \mid \mathbf{c}_i) := 0$  for  $\mathbf{c}_i \in (\mathcal{C} \cup \{\perp\})^i$  and  $c_{i+1} \in \mathcal{C}$ , and  $\tau(\perp \mid \mathbf{c}_i) := 1$  instead, to indicate that  $c_{i+1}$  is set to  $\perp$  then by  $\mathcal{E}'$ . Thus, by construction,  $\mathcal{E}'$  keeps finding useful challenges  $c_1, c_2, \dots \in \mathcal{C}$  until it has collected  $\mathbf{c}_i \in \Gamma$  which happens at the latest for  $i = t+1$ . Thus, the distribution

$$D(\mathbf{c}) := \prod_{i=0}^t \tau(c_{i+1} \mid \mathbf{c}_i)$$



of the list  $\mathbf{c}$  of challenges produced by  $\mathcal{E}'$  has its support contained in  $\mathcal{G}$ .

Below, we show how we can relate the probability of finding a certain  $\mathbf{c}$  among the two extractors.

Let us now go back to analyzing the extractor. We observe that, for  $\mathbf{c} \in \mathcal{G}_k$  and  $i < k$ , i.e.,  $\mathbf{c}_i \notin \Gamma$ , conditioned on  $[\mathbf{C}_i = \mathbf{c}_i]$  the extractor tries to find an  $(i+1)$ -th  $c_{i+1} \in \mathcal{U}_\Gamma(\mathbf{c}_i)$  with  $\mathcal{V}(c_{i+1}, \mathcal{A}(c_{i+1})) = 1$ . To this end, it starts running two geometric experiments in parallel until either of them finishes. The first geometric experiment repeatedly runs  $y_{i+1} \leftarrow \mathcal{A}(c_{i+1})$  for  $c_{i+1} \in \mathcal{U}_\Gamma(\mathbf{c}_i)$  sampled uniformly at random, and thus has parameter  $\delta(\mathbf{c}_i)$ . The second geometric experiment repeatedly runs  $y \leftarrow \mathcal{A}(d)$  for  $d \leftarrow_R \mathcal{U}_\Gamma(\mathbf{c}_{i-1})$  until  $d \notin \mathcal{U}_\Gamma(\mathbf{c}_i)$  and  $\mathcal{V}(d, y) = 1$ , i.e., it has parameter  $\Delta(\mathbf{c}_i)$ . The extractor succeeds in finding the  $(i+1)$ -th challenge if the first geometric experiment finishes before the second. Hence, by Lemma 1 and Equation (9),

$$\Pr(C_{i+1} \neq \perp \mid \mathbf{C}_i = \mathbf{c}_i) = \frac{\delta(\mathbf{c}_i)}{\delta(\mathbf{c}_i) + \Delta(\mathbf{c}_i) - \Delta(\mathbf{c}_i)\delta(\mathbf{c}_i)} \geq \frac{\tilde{\delta}(\mathbf{c}_i)}{\delta(\mathbf{c}_{i-1})},$$

and (noting that  $c_{i+1} \neq \perp$  by definition of  $\mathcal{G}$ )

$$\begin{aligned} \Pr(C_{i+1} = c_{i+1} \mid \mathbf{C}_i = \mathbf{c}_i) &= \Pr(C_{i+1} = c_{i+1} \wedge C_{i+1} \neq \perp \mid \mathbf{C}_i = \mathbf{c}_i) \\ &= \tau(c_{i+1} \mid \mathbf{c}_i) \cdot \Pr(C_{i+1} \neq \perp \mid \mathbf{C}_i = \mathbf{c}_i) \\ &\geq \tau(c_{i+1} \mid \mathbf{c}_i) \cdot \frac{\tilde{\delta}(\mathbf{c}_i)}{\delta(\mathbf{c}_{i-1})}. \end{aligned}$$

It thus follows that for  $k$  such that  $\mathbf{c} \in \mathcal{G}_k$

$$\begin{aligned} \Pr(\mathbf{C} = \mathbf{c}) &= \Pr(\mathbf{C}_k = \mathbf{c}_k) = \prod_{i=1}^k \Pr(\mathbf{C}_i = \mathbf{c}_i \mid \mathbf{C}_{i-1} = \mathbf{c}_{i-1}) \\ &\geq \delta(\mathbf{c}_0) \tau(c_1 \mid \mathbf{c}_0) \prod_{i=1}^{k-1} \frac{\tilde{\delta}(\mathbf{c}_i)}{\delta(\mathbf{c}_{i-1})} \tau(c_{i+1} \mid \mathbf{c}_i) \\ &\geq \frac{\epsilon^{\mathcal{V}(\mathcal{A})} - \kappa_\Gamma}{1 - \kappa_\Gamma} \prod_{i=1}^k \tau(c_i \mid \mathbf{c}_{i-1}), \end{aligned} \tag{19}$$

where the final inequality follows from Lemma 11. Hence, for  $\mathbf{c} \in \mathcal{G}_k$ , exploiting  $\tau(c_i \mid \mathbf{c}_{i-1}) = \tau(\perp \mid \mathbf{c}_{i-1}) = 1$  for  $i > k$ , we find

$$\Pr(\mathbf{C} = \mathbf{c}) \geq \frac{\epsilon^{\mathcal{V}(\mathcal{A})} - \kappa_\Gamma}{1 - \kappa_\Gamma} \prod_{i=1}^{t+1} \tau(c_i \mid \mathbf{c}_{i-1}) = \frac{\epsilon^{\mathcal{V}(\mathcal{A})} - \kappa_\Gamma}{1 - \kappa_\Gamma} \cdot D(\mathbf{c}).$$

Thus, the extractor's success probability equals

$$\Pr(\mathbf{C} \in \mathcal{G}) = \sum_{\mathbf{c} \in \mathcal{G}} \Pr(\mathbf{C} = \mathbf{c}) \geq \frac{\epsilon^{\mathcal{V}(\mathcal{A})} - \kappa_\Gamma}{1 - \kappa_\Gamma} \cdot \sum_{\mathbf{c} \in \mathcal{G}} D(\mathbf{c}) = \frac{\epsilon^{\mathcal{V}(\mathcal{A})} - \kappa_\Gamma}{1 - \kappa_\Gamma}$$

where the final equality follows from the support of  $D$  being contained in  $\mathcal{G}$ . This proves the claimed success probability.

**Expected Number of  $\mathcal{A}$ -Queries.** Let us now continue with the expected running time analysis. For  $1 \leq i < t$ , let  $\mathbf{c}_i \in \mathcal{C}^i$  with  $\mathbf{c}_i \notin \Gamma$  and  $\Pr[\mathbf{C}_i = \mathbf{c}_i] > 0$ . As before, we use that in its  $(i+1)$ -th iteration, conditioned on  $[\mathbf{C}_i = \mathbf{c}_i]$ , the extractor runs two geometric experiments in parallel with parameters  $\delta(\mathbf{c}_i)$  and  $\Delta(\mathbf{c}_i)$ , trying to find the  $(i+1)$ -th challenge. The probability that one of the experiments finishes in a single trial equals

$$1 - (1 - \delta(\mathbf{c}_i))(1 - \Delta(\mathbf{c}_i)) = \delta(\mathbf{c}_i) + \Delta(\mathbf{c}_i) - \Delta(\mathbf{c}_i)\delta(\mathbf{c}_i) \geq \delta(\mathbf{c}_{i-1}),$$

where the inequality follows from Equation (8).

Now note that, to determine the value of the coin, the extractor only needs to invoke  $\mathcal{A}$  if  $d \notin \mathcal{U}_\Gamma(\mathbf{c}_i)$ . Hence, in expectation each trial invokes  $\mathcal{A}$  at most

$$1 + \Pr(C \notin \mathcal{U}_\Gamma(\mathbf{c}_i) \mid C \in \mathcal{U}_\Gamma(\mathbf{c}_{i-1})) \leq 1 + \Pr(C \notin \mathcal{U}_\Gamma(\mathbf{c}_i)) \leq 1 + \kappa_\Gamma$$

times, where the first inequality follows since  $\mathcal{U}_\Gamma(\mathbf{c}_i) \subseteq \mathcal{U}_\Gamma(\mathbf{c}_{i-1})$ , and the second since  $\mathcal{C} \setminus \mathcal{U}_\Gamma(\mathbf{c}_i) \notin \Gamma$  for all  $\mathbf{c}_i \notin \Gamma$  (Equation (3)). Hence, if we let  $Y_i$  denote the number of  $\mathcal{A}$ -queries that the extractor requires in its  $i$ -th iteration, i.e., when trying to find the  $i$ -th challenge  $C_i$ , then (by Lemma 2)

$$\mathbb{E}[Y_1] = 1 \quad \text{and} \quad \mathbb{E}[Y_{i+1} \mid \mathbf{C}_i = \mathbf{c}_i] \leq \frac{1 + \kappa_\Gamma}{\delta(\mathbf{c}_{i-1})}.$$

On the other hand, if  $\mathbf{c}_i$  is such that  $\mathbf{c}_i \in \Gamma$  or  $c_i = \perp$  (in either case, the extractor is done) then  $\mathbb{E}[Y_{i+1} \mid \mathbf{C}_i = \mathbf{c}_i] = 0$ . Putting these observations together shows that, for all  $1 \leq k \leq t$ ,

$$\mathbb{E}[Y_{k+1}] = \sum_{\mathbf{c}_k \in T_k} \Pr(\mathbf{C}_k = \mathbf{c}_k) \cdot \mathbb{E}[Y_{k+1} \mid \mathbf{C}_k = \mathbf{c}_k] \leq (1 + \kappa_\Gamma) \cdot \sum_{\mathbf{c}_k \in T_k} \frac{\Pr(\mathbf{C}_k = \mathbf{c}_k)}{\delta(\mathbf{c}_{k-1})},$$

where  $T_k := \{\mathbf{c}_k \in \mathcal{C}^k : \mathbf{c}_k \notin \Gamma \wedge \Pr(\mathbf{C}_k = \mathbf{c}_k) > 0\}$ .

As in the success probability analysis, we will now expand the probability  $\Pr(\mathbf{C}_k = \mathbf{c}_k)$ , but now aiming for an upper bound. Then, for  $\mathbf{c}_k \in T_k$  and  $i < k$ , by Equation (8)

$$\Pr(C_{i+1} \neq \perp \mid \mathbf{C}_i = \mathbf{c}_i) = \frac{\delta(\mathbf{c}_i)}{\delta(\mathbf{c}_i) + \Delta(\mathbf{c}_i) - \Delta(\mathbf{c}_i)\delta(\mathbf{c}_i)} \leq \frac{\delta(\mathbf{c}_i)}{\delta(\mathbf{c}_{i-1})},$$

and (noting that  $c_{i+1} \neq \perp$  by definition of  $\mathcal{G}$ )

$$\begin{aligned} \Pr(C_{i+1} = c_{i+1} \mid \mathbf{C}_i = \mathbf{c}_i) &= \Pr(C_{i+1} = c_{i+1} \wedge C_{i+1} \neq \perp \mid \mathbf{C}_i = \mathbf{c}_i) \\ &= \tau(c_{i+1} \mid \mathbf{c}_i) \cdot \Pr(C_{i+1} \neq \perp \mid \mathbf{C}_i = \mathbf{c}_i) \\ &\leq \tau(c_{i+1} \mid \mathbf{c}_i) \cdot \frac{\delta(\mathbf{c}_i)}{\delta(\mathbf{c}_{i-1})}. \end{aligned}$$

where, as before,

$$\tau(c_{i+1} \mid \mathbf{c}_i) := \Pr(C_{i+1} = c_{i+1} \mid C_{i+1} \neq \perp \wedge \mathbf{C}_i = \mathbf{c}_i).$$

Similar to Equation (19), but now using the upper bound, it therefore follows that

$$\begin{aligned} \Pr(\mathbf{C}_k = \mathbf{c}_k) &= \prod_{i=1}^k \Pr(C_i = c_i \mid \mathbf{C}_{i-1} = \mathbf{c}_{i-1}) \\ &\leq \delta(\mathbf{c}_0) \cdot \tau(c_1 \mid \mathbf{c}_0) \cdot \prod_{i=1}^{k-1} \frac{\delta(\mathbf{c}_i)}{\delta(\mathbf{c}_{i-1})} \cdot \tau(c_{i+1} \mid \mathbf{c}_i) \\ &= \delta(\mathbf{c}_{k-1}) \cdot \prod_{i=1}^k \tau(c_i \mid \mathbf{c}_{i-1}). \end{aligned}$$

Hence, for all  $k \geq 1$ ,

$$\begin{aligned} \mathbb{E}[Y_{k+1}] &\leq (1 + \kappa_\Gamma) \cdot \sum_{\mathbf{c}_k \in T_k} \prod_{i=1}^k \tau(c_i \mid \mathbf{c}_{i-1}) \\ &\leq (1 + \kappa_\Gamma) \cdot \sum_{c_1 \in \mathcal{C}} \tau(c_1 \mid \mathbf{c}_0) \sum_{c_2 \in \mathcal{C}} \tau(c_2 \mid \mathbf{c}_1) \cdots \sum_{c_k \in \mathcal{C}} \tau(c_k \mid \mathbf{c}_{k-1}) \\ &\leq (1 + \kappa_\Gamma), \end{aligned}$$

where the ill-defined terms  $\tau(c_i \mid \mathbf{c}_{i-1})$ , i.e., with  $\Pr(\mathbf{C}_{i-1} = \mathbf{c}_{i-1}) = 0$ , are simply defined to be 0.

Hence, the expected running time of the extractor is at most

$$1 + \sum_{k=1}^t \mathbb{E}[Y_{k+1}] \leq 1 + t \cdot (1 + \kappa_\Gamma),$$

which completes the proof of the lemma. □