



Generalized Special-Sound Interactive Proofs and Their Knowledge Soundness

Thomas Attema^{1,3} , Serge Fehr^{1,2}, and Nicolas Resch⁴ 

¹ CWI, Cryptology Group, Amsterdam, The Netherlands
serge.fehr@cw.nl

² Leiden University, Mathematical Institute, Leiden, The Netherlands

³ TNO, Cyber Security and Robustness, The Hague, The Netherlands
thomas.attema@tno.nl

⁴ University of Amsterdam, Informatics Institute, Amsterdam, The Netherlands
n.a.resch@uva.nl

Abstract. A classic result in the theory of interactive proofs shows that a *special-sound* Σ -protocol is automatically a *proof of knowledge*. This result is very useful to have, since the latter property is typically tricky to prove from scratch, while the former is often easy to argue — if it is satisfied. While classic Σ -protocols often are special-sound, this is unfortunately not the case for many recently proposed, highly efficient interactive proofs, at least not in this strict sense. Motivated by this, the original result was recently generalized to k -special-sound Σ -protocols (for arbitrary, polynomially bounded k), and to multi-round versions thereof. This generalization is sufficient to analyze (e.g.) Bulletproofs-like protocols, but is still insufficient for many other examples.

In this work, we push the relaxation of the special soundness property to the extreme, by allowing an *arbitrary* access structure Γ to specify for which subsets of challenges it is possible to compute a witness, when given correct answers to these challenges (for a fixed first message). Concretely, for any access structure Γ , we identify parameters t_Γ and κ_Γ , and we show that any Γ -special-sound Σ -protocol is a proof of knowledge with knowledge error κ_Γ if t_Γ is polynomially bounded. Similarly for multi-round protocols.

We apply our general result to a couple of simple but important example protocols, where we obtain a tight knowledge error as an immediate corollary. Beyond these simple examples, we analyze the FRI protocol. Here, showing the general special soundness notion is non-trivial, but can be done (for a certain range of parameters) by recycling some of the techniques used to argue ordinary soundness of the protocol (as an IOP). Again as a corollary, we then derive that the FRI protocol, as an interactive proof by using a Merkle-tree commitment, has a knowledge extractor with almost optimal knowledge error, with the caveat that the extractor requires (expected) quasi-polynomial time.

Finally, building up on the technique for the parallel repetition of k -special-sound Σ -protocols, we show the same strong parallel repetition result for Γ -special-sound Σ -protocol and its multi-round variant.

1 Introduction

Background. A key feature of an interactive proof is *soundness*, which requires that the verifier will not accept a false statement, i.e., an instance x that is not in the considered language, except with bounded probability. In many situations however, a stronger notion of soundness is needed: *knowledge soundness*. Informally, knowledge soundness requires the prover to *know* a witness w that certifies that x is a true statement, in order for the verifier to accept (except with bounded probability). More formally, this is captured by the existence of an efficient *extractor*, which has (rewindable) oracle access to any, possibly dishonest, prover, and which outputs a witness w for the considered statement x with a probability that is tightly related to the probability of the prover making the verifier accept.

Since their introduction, interactive proofs that satisfy knowledge soundness, typically referred to *proofs of knowledge* then, have found a myriad of applications. However, showing that an interactive proof satisfies knowledge soundness is typically non-trivial—often significantly more involved than showing ordinary soundness. By default, it involves *designing* the extractor, and *proving* that it “does the job.” We got spoiled in the past, where most of the considered interactive proofs were Σ -protocols, i.e., public-coin 3-round interactive proofs, and had the additional property of being *special-sound*. Indeed, this made life rather easy since special-soundness is a property that is usually quite easy to prove, and that implies ordinary and knowledge soundness via a general classical result. Thus, knowledge soundness was often obtained (almost) for free. However, this has changed in recent years, where the focus has shifted towards finding *highly efficient* interactive proofs (where efficiency is typically measured via the communication complexity, verification time, etc.); many of these highly efficient solutions are *not* special-sound, and thus require a knowledge-soundness proof from scratch.

Given this situation, it would be desirable to have stronger versions of the generic “*special-soundness* \Rightarrow *knowledge soundness*” result that applies to a *weaker* notion of special-soundness, which then hopefully is satisfied by these new cutting-edge interactive proofs. One step in this direction was recently made in [2,3], where the above implication was extended to k -special-sound interactive proofs, and, even more generally, to (k_1, \dots, k_μ) -special-sound *multi-round* public-coin interactive proofs, for arbitrary positive integer parameters, subject to being suitably bounded from above (e.g., $k \leq \text{poly}(|x|)$). Rather naturally, k -special-soundness means that from accepting responses to k pairwise distinct challenges for one fixed message, a witness can be efficiently computed (so that 2-special-soundness coincides with the classical special-soundness property); for the multi-round version, a suitable tree of transcripts is needed for computing a witness. This weaker notion of special-soundness is in particular sufficient to analyze Bulletproofs-like protocols, and so we directly obtain knowledge soundness for these protocols.

However, this weaker notion still falls short of capturing many of the recent highly-efficient interactive proofs. For instance, a commonly used amortization

technique, where the prover proves a *random linear combination* of n statements (instead of proving all the statements individually), requires correct responses for n *linearly independent* challenge vectors in order to compute a witness. Another example comes from the design principle to first construct a highly efficient probabilistically checkable proof (PCP) or interactive *oracle* proof (IOP), and then to compile it into a standard (public-coin) interactive proof in the natural way by means of a Merkle-tree commitment [11–13]. Also here, one does not obtain a special-sound protocol in the above generalized sense (or then only for a too large parameter); instead, one requires challenges that correspond to sets whose union covers all (or sufficiently many of) the leaves of the Merkle tree, in order to obtain a witness.

Our Technical Results. In this paper, we push the weakening of the special-soundness property to the extreme. For Σ -protocols, in the spirit of ordinary or k -special-soundness, the notion of special-soundness that we consider in this work requires that a witness can be efficiently computed from accepting responses to *sufficiently many* pairwise distinct challenges, but now “sufficiently many” is captured by an arbitrary monotone (access) structure Γ , i.e., an arbitrary monotone set of subsets of the challenge set. This gives rise to the notion of Γ -special-soundness, which coincides with k -special-soundness in the special case where Γ is the threshold access structure with threshold k . This naturally extends to multi-round public-coin interactive proofs, leading to the notion of $(\Gamma_1, \dots, \Gamma_\mu)$ -special-soundness. Similar notions were considered in [9, 10] in the setting of *commit-and-open* Σ -protocols, and in some more constrained form, where the monotone structures are replaced by matroids, in [14, 15].

We cannot expect for every Γ that a Γ -special-sound protocol is a proof of knowledge. Instead, we identify parameters t_Γ and κ_Γ , determined by the structure Γ , and for any Γ -special-sound Σ -protocol we prove existence of an extractor that has a knowledge error κ_Γ and an expected running time that scales with t_Γ . Thus, as long as $t_\Gamma \leq \text{poly}(|x|)$, Γ -special-soundness implies knowledge soundness. Similarly for $(\Gamma_1, \dots, \Gamma_\mu)$ -special-sound multi-round protocols.

The construction of our extractor for Γ -special-sound protocols (and its multi-round generalization) is inspired by the extractor construction from [3]. As a nice consequence, we can recycle the line of reasoning from [3] to prove strong parallel repetition and extend it to our general notion of special-soundness, showing that also here the knowledge error of a parallel repetition decreases exponentially with the number of repetitions. For this result, we refer to the full version [1].

Applications. Our general technique gives immediate, tight results for simple but important example protocols. For example, applied to the above mentioned amortization technique of proving a random linear combination, we directly obtain knowledge extraction with a knowledge error that matches the trivial cheating probability. Similarly, applied to the natural interactive proof for a Merkle commitment, where the prover is challenged to open a random subset (of a certain size), we obtain a knowledge error that matches the probability of one faulty node not being opened.

In order to demonstrate the usefulness of our result beyond the above simple examples, we analyze the (interactive) FRI protocol [5].¹ We prove that for a certain range of parameters, when instantiated with a Merkle tree commitment using a collision resistant hash function (or with any non-interactive, computationally binding vector commitment scheme with local openings), the protocol admits a knowledge extractor with knowledge error essentially matching the trivial cheating probability with the following caveat: the knowledge extractor runs only in (expected) quasi-polynomial time. (At least, this is true if the protocol is run for logarithmically many rounds, as is typically done. For a natural constant-round variant, which requires more total communication, we can obtain nearly optimal knowledge soundness, i.e., here the knowledge extractor runs in expected polynomial time.) In more detail, for any proximity parameter δ up to $\delta < \frac{1-\rho}{4}$, where ρ is the relative rate of the considered code, we establish the existence a knowledge extractor running in expected time $N^{O(\log N)}$ which, when given oracle access to a (potentially dishonest) prover \mathcal{P}^* , succeeds with probability at least $\epsilon(\mathcal{P}^*) - ((1-\delta)^t + O(N/|\mathbb{F}|))$, where N is the length of the code, t is the number of repetitions of a certain verification step, and $\epsilon(\mathcal{P}^*)$ is the probability \mathcal{P}^* convinces the verifier to accept.² For context, the trivial cheating probability for the protocol is $\max\{(1-\delta)^t, 1/|\mathbb{F}|\}$. In contrast to the above simple examples, arguing that the FRI protocol is $(\Gamma_1, \dots, \Gamma_\mu)$ -special-soundness is not trivial; however, technical results from [5] can be recycled in order to show this, and then the existence of the knowledge extractor follows immediately from our generic result. While proving the existence of a quasi-polynomial time extractor does not suffice for establishing the standard notion of knowledge soundness, we believe that it still offers a nontrivial guarantee with the potential for practical relevance.

A final example, which we would like to briefly discuss, is parallel repetition. This example shows that our generic technique does not always work. For simplicity, consider a k -special-sound Σ -protocol with $k > 2$ (but the discussion also applies to multi-round protocols, and to our generalized notion of special soundness). Then, its t -fold parallel repetition is *not* k -special-sound anymore (unless $k = 2$). One can argue that it is $((k-1)^t + 1)$ -special-sound — but this parameter is exponential in t , and thus one cannot directly conclude knowledge soundness. On the other hand, equipped with our generalized notion, one can observe that the parallel repetition is Γ -special-sound for Γ being the structure that accepts a list of challenge vectors, each vector of length t , if there is one position where the challenge vectors feature at least k different values. Unfortunately, also here,

¹ We point out that, when considering the FRI protocol for an actual hash function (rather than the random oracle), ordinary soundness is meaningless: the *existence* of an opening of a Merkle commitment with a certain (not too obscure) property holds trivially. Thus, it is crucial to argue knowledge soundness in this case.

² For the constant-round variant, we obtain genuine knowledge-soundness $(1-\delta)^t + O(N/|\mathbb{F}|)$, that is here the knowledge extractor's expected running time is $N^{O(1)}$. We also point out that a simple argument can be used to show knowledge soundness $(1-2^\mu\delta)^t + O(N/|\mathbb{F}|)$, where μ is the number of rounds; however, this result is only meaningful for fairly extreme parameters.

the crucial parameter t_T turns out to be exponential for this structure Γ , and so our generic result does not imply knowledge soundness. Fortunately, for this particular and important example, the parallel repetition result from [3] applies in case of k -special-sound protocols (and its multi-round generalization), and our extension (see the full version [1]) of the parallel repetition applies in case of arbitrary $(\Gamma_1, \dots, \Gamma_\mu)$ -special-sound protocols. Thus, after all, we can still argue (optimal) knowledge soundness in this case.

In conclusion, we expect that with our generic result for $(\Gamma_1, \dots, \Gamma_\mu)$ -special-sound protocols (which requires control over certain parameters to be applicable), and with our general parallel repetition result, our work offers powerful tools for proving knowledge soundness of many sophisticated proofs of knowledge.

2 Preliminaries

We write $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$ for the set of nonnegative integers. Further, for any $q \in \mathbb{Z}$, $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$ denotes the ring of integers modulo q .

2.1 Interactive Proofs

Let us now introduce some standard terminology and definitions with respect to interactive proofs. We follow standard conventions as presented in [4].

Let $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$ be a binary relation, containing statement-witness pairs $(x; w)$. We assume all relations to be NP-relations, i.e., verifying that $(x; w) \in R$ takes time polynomial in $|x|$. An interactive proof for a relation R aims to allow a prover \mathcal{P} to convince a verifier \mathcal{V} that a public statement x admits a (secret) witness w , i.e., $(x; w) \in R$, or even that the knows a witness w for x .

An interactive proof with three communication rounds, where we may assume the prover to send the first and final message, is called a Σ -protocol. Further, an interactive proof is said to be *public-coin* if the verifier publishes all its random coins. In this case, we may assume all the verifier's messages to be sampled uniformly at random from finite (challenge) sets.

An interactive proof is said to be *complete* if for any statement witness pair $(x; w)$ an honest execution results in an accepting transcript (with high probability). It is *sound* if a dishonest prover cannot convince an honest verifier on public inputs x that do not admit a witness w , i.e., on false statements x . More precisely, $(\mathcal{P}, \mathcal{V})$ is sound if \mathcal{V} rejects false statements x with high probability. The stronger notion of *knowledge soundness* requires that (potentially dishonest) provers that succeed in convincing the verifier with large enough probability must actually “know” a witness w . We will mainly be interested in analyzing the knowledge soundness of interactive proofs. For this reason, we formally define this property below.

Definition 1 (Knowledge Soundness). *An interactive proof $(\mathcal{P}, \mathcal{V})$ for relation R is knowledge sound with knowledge error $\kappa: \mathbb{N} \rightarrow [0, 1]$ if there exists a*

positive polynomial q and an algorithm \mathcal{E} , called a knowledge extractor, with the following properties. Given input x and black-box oracle access to a (potentially dishonest) prover \mathcal{P}^* , the extractor \mathcal{E} runs in an expected number of steps that is polynomial in $|x|$ (counting queries to \mathcal{P}^* as a single step) and outputs a witness $w \in R(x)$ with probability

$$\Pr((x; \mathcal{E}^{\mathcal{P}^*}(x)) \in R) \geq \frac{\epsilon(\mathcal{P}^*, x) - \kappa(|x|)}{q(|x|)},$$

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

Remark 1 (Interactive Arguments). In some cases, soundness and knowledge soundness only hold with respect to computationally bounded provers, i.e., unbounded provers can falsely convince a verifier. Computationally (knowledge) sound protocols are referred to as interactive *arguments*. Proving soundness of interactive arguments can be significantly more complicated than proving soundness of interactive proofs. However, in the context of knowledge soundness, an interactive argument for relation R can oftentimes be cast as an interactive proof for a modified relation

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

Hence, in this case the knowledge extractor will either output a witness w with respect to the original relation w , or it will output the solution to some computational problem, e.g., a discrete logarithm relation. In fact, our analysis of the FRI protocol in Sect. 7 exemplifies this general principle. For this reason, knowledge soundness of interactive arguments can typically be analyzed via knowledge extractors that are originally defined for interactive proofs. Therefore, we will focus on the analyzes of interactive proofs.

Proving knowledge soundness of Σ -protocols directly is a nontrivial task, as it requires the construction of an efficient knowledge extractor. It is typically much easier to prove a related *threshold* special-soundness property, which states that a witness can be extracted from a sufficiently large set of colliding and accepting transcripts.

Definition 2 (k -out-of- N Special-Soundness). Let $k, N \in \mathbb{N}$. A 3-round public-coin interactive proof $\Pi = (\mathcal{P}, \mathcal{V})$ for relation R , with challenge set of cardinality $N \geq k$, is k -out-of- N special-sound if there exists an algorithm that, on input a statement x and k accepting transcripts $(a, c_1, z_1), \dots, (a, c_k, z_k)$ with common first message a and pairwise distinct challenges c_1, \dots, c_k , runs in polynomial time and outputs a witness w such that $(x; w) \in R$. We also say Π is k -special-sound and, if $k = 2$, it is simply said to be special-sound.

It is known that k -out-of- N special-soundness implies knowledge soundness with knowledge error $(k - 1)/N$. Recently, the multi-round generalization (k_1, \dots, k_μ) -out-of- (N_1, \dots, N_μ) special-soundness has become relevant. It is now

known that also this generalization tightly implies knowledge soundness [2]. For a formal definition, we refer either to [2] or to Sect. 6 where we generalize this (multi-round) notion beyond the threshold setting.

2.2 Geometric Distribution

This work adapts the extractor of [3]. For this reason, we also need the following preliminaries on the geometric distribution from their work.

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 $p = \Pr(X = 1)$. 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. *Let $X \sim \text{Geo}(p)$ and $Y \sim \text{Geo}(q)$ be independently distributed. Then,*

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

3 A Generalized Notion of Special-Soundness for Σ -Protocols

In this section, we define a generalized notion of special-soundness. To this end, we first recall the definition of *monotone structures*.

Definition 3 (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, Γ or (Γ, \mathcal{C}) is said to be a monotone structure if it is closed under taking supersets, i.e., $S \in \Gamma$ and $S \subseteq T \subseteq \mathcal{C}$ implies $T \in \Gamma$.*

In some textbooks monotone structures Γ do not contain the empty set \emptyset by definition, which is equivalent to $\Gamma \neq 2^{\mathcal{C}}$, and they are required to be nonempty, which is equivalent to $\mathcal{C} \in \Gamma$. For convenience, we also consider $\Gamma = \emptyset$ and $\Gamma = 2^{\mathcal{C}}$ to be monotone structures. Then, for any $\mathcal{D} \subseteq \mathcal{C}$, the restriction

$$\Gamma|_{\mathcal{D}} = \{S \subseteq \mathcal{D} : S \in \Gamma\} \subseteq 2^{\mathcal{D}}$$

defines a monotone structure $(\Gamma|_{\mathcal{D}}, \mathcal{D})$.

Definition 4 (Minimal Set). *Let (Γ, \mathcal{C}) be a monotone structure. A set $S \in \Gamma$ is minimal if none of its proper subsets are in Γ , i.e., for all $T \subsetneq S$ it holds that $T \notin \Gamma$. Further, $M(\Gamma) \subseteq \Gamma$ denotes the set of minimal elements of Γ .*

Definition 5 (Distance to a Monotone Structure). For a nonempty monotone structure (Γ, \mathcal{C}) , we define the following distance function:

$$d_\Gamma: 2^{\mathcal{C}} \rightarrow \mathbb{N}_0, \quad S \mapsto \min_{T \in \Gamma} |T \setminus S|.$$

Equivalently,

$$d_\Gamma: 2^{\mathcal{C}} \rightarrow \mathbb{N}_0, \quad S \mapsto \min_{T \subseteq \mathcal{C}} \{|T| : S \cup T \in \Gamma\}.$$

If $\Gamma = \emptyset$, we define d_Γ to be identically equal to ∞ .

The value $d_\Gamma(S) \in \mathbb{N}_0$ equals the minimum number of elements that have to be added to the set S to obtain an element of Γ . In particular, $d_\Gamma(S) = 0$ if and only if $S \in \Gamma$. Hence, it shows how close S is to the monotone structure Γ .

The key observation is now that typical knowledge extractors for interactive proofs proceed by extracting some set of accepting transcripts from a dishonest prover attacking the interactive proof. Subsequently, the knowledge extractor computes a witness from this set of accepting transcripts. Clearly, the set of sets of accepting transcripts from which a witness can be computed is closed under taking supersets, i.e., it is a monotone structure. Therefore, the following special-soundness notion for 3-round Σ -protocols follows naturally.

Definition 6 (Γ -out-of- \mathcal{C} Special-Soundness). Let (Γ, \mathcal{C}) be a monotone structure. A 3-round public-coin interactive proof $(\mathcal{P}, \mathcal{V})$ for relation R , with challenge set \mathcal{C} , is Γ -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 Γ -special-sound.

The above definition is a generalization of k -out-of- N special-soundness, where the extractability is guaranteed when given k colliding accepting transcripts with common first message a and pairwise distinct challenges c_i that are elements of a challenge set with cardinality N . Hence, when Γ contains all sets of cardinality at least k , i.e., it is a *threshold* monotone structure, Γ -out-of- \mathcal{C} special-soundness reduces to k -out-of- N special-soundness, where $N = |\mathcal{C}|$.

Remark 2. Formally, the monotone structure (Γ, \mathcal{C}) of Definition 6 may depend on the size $|x|$ of the public input x , i.e., it should actually be replaced by an ensemble $(\Gamma_\lambda, \mathcal{C}_\lambda)$ of monotone structures indexed by the size $\lambda \in \mathbb{N}$ of the public input of $(\mathcal{P}, \mathcal{V})$. For simplicity, we will abuse notation by ignoring this dependency and simply writing (Γ, \mathcal{C}) .

4 Knowledge Extraction for Γ -out-of- \mathcal{C} Special-Sound Σ -Protocols

Our goal is to prove that, for certain monotone structures (Γ, \mathcal{C}) , Γ -out-of- \mathcal{C} special-soundness (tightly) implies knowledge soundness, and to determine the

corresponding knowledge error. In order to prove this, we construct a knowledge extractor that, by querying a prover \mathcal{P}^* attacking the interactive proof, obtains a set of accepting transcripts with common first message and for which the challenges form a set in Γ . Without loss of generality we may assume \mathcal{P}^* to be deterministic,³ i.e., \mathcal{P}^* always outputs the same first message a . Hence, \mathcal{P}^* can be viewed as a (deterministic) function

$$\mathcal{P}^* : \mathcal{C} \rightarrow \{0, 1\}^* \quad c \mapsto y = (a, c, z),$$

that on input a challenge $c \in \mathcal{C}$ outputs a protocol transcript $y = (a, c, z)$.

Let $A \subseteq \mathcal{C}$ be the set of challenges for which \mathcal{P}^* succeeds, i.e., $A = \{c \in \mathcal{C} : V(\mathcal{P}^*(c)) = 1\}$. Then the goal of the extractor is to find a set $B \in \Gamma|_A$. The difficulty is that the extractor is only given oracle access to \mathcal{P}^* and therefore does not know the set A . For this reason, extractors typically proceed recursively as follows: if at some point the extractor has found some $S \subseteq A$ with $S \notin \Gamma$, it will try new challenges $c \in \mathcal{C}$ until \mathcal{P}^* succeeds. The hope is then that $S \cup \{c\} \subseteq A$ is “closer” to $\Gamma|_A$ than S . More precisely, the extractor tries to find a $c \in A \subseteq \mathcal{C}$ such that $d_{\Gamma|_A}(S \cup \{c\}) < d_{\Gamma|_A}(S)$. Note that not all challenges c shorten the distance to $\Gamma|_A$, e.g., $d_{\Gamma|_A}(S \cup \{c\}) = d_{\Gamma|_A}(S)$ for all $c \in S$. Since the extractor does not know the set A , it cannot evaluate this distance function.

However, for any S , the challenge set \mathcal{C} can be partitioned into a partition of “useless” challenges and a partition of “potentially useful” challenges. The useless challenges are the $c \in \mathcal{C}$ such that $d_{\Gamma|_A}(S \cup \{c\}) = d_{\Gamma|_A}(S)$ for all $A \subseteq \mathcal{C}$ containing S , i.e., for all A useless challenges will not shorten the distance to $\Gamma|_A$. For instance, all $c \in S$ are useless challenges for any S and any Γ . However, in some settings the set of useless challenges is larger than S , and in general this observation is crucial for the extractor to be efficient. In fact, this is the case for all interactive proofs that warrant a generalization of the existing threshold special-soundness notion. All challenges $c \in \mathcal{C}$ that are not useless are potentially useful, i.e., for these challenges there exist an $A \subseteq \mathcal{C}$ containing S such that $d_{\Gamma|_A}(S \cup \{c\}) < d_{\Gamma|_A}(S)$. The set of useful challenges is denoted $U_\Gamma(S)$, where the function U_Γ is formally defined below.

Definition 7 (Useful Elements). *For a monotone structure (Γ, \mathcal{C}) , we define the following function:*

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

Note that $\Gamma = \emptyset$ implies $U_\Gamma(S) = \emptyset$ for all $S \subseteq \mathcal{C}$. Moreover, if Γ is nonempty, $U_\Gamma(S) = \emptyset$ if and only if $S \in \Gamma$.

The following lemma shows that for any $c \in U_\Gamma(S)$, there exists an $A \in \Gamma$ containing $S \cup \{c\}$ such that

$$d_{\Gamma|_A}(S \cup \{c\}) < d_{\Gamma|_A}(S),$$

i.e., the challenges $c \in U_\Gamma(S)$ are indeed potentially useful to the extractor. Even more so, it is essential that the extractor considers *all* challenges $c \in U_\Gamma(S)$. For

³ See [3] for a proof of this claim.

every $c \in U_\Gamma(S)$, it might namely be the case that the $A \in \Gamma$ that “certifies” c , i.e., the A such that $S \subset A$ and $A \setminus \{c\} \notin \Gamma$, corresponds to the challenges for which the prover \mathcal{P}^* succeeds. Since $A \setminus \{c\} \notin \Gamma$, the extractor can only succeed if it considers the challenge $c \in U_\Gamma(S)$ at some point.

The same lemma shows that challenges $c \notin U_\Gamma(S)$ will never decrease the distance, i.e., they are indeed useless to the extractor. More precisely, if $c \notin U_\Gamma(S)$, for every $A \in \Gamma$ containing $S \cup \{c\}$ it holds that

$$d_{\Gamma|_A}(S \cup \{c\}) = d_{\Gamma|_A}(S).$$

Lemma 2. *Let (Γ, \mathcal{C}) be a monotone structure and $S \subset \mathcal{C}$. Then $c \in U_\Gamma(S)$ if and only if there exists an $A \in \Gamma$ containing $S \cup \{c\}$ such that*

$$d_{\Gamma|_A}(S \cup \{c\}) < d_{\Gamma|_A}(S).$$

For the proof of Lemma 2, we refer the to the full version [1].

We also derive the following lemma, which shows that even if all useless challenges $c \in \mathcal{C} \setminus U_\Gamma(S)$ are added to the set $S \in 2^\mathcal{C} \setminus \Gamma$, the resulting subset is still not in Γ .

Lemma 3. *Let (Γ, \mathcal{C}) be a monotone structure and $S \in 2^\mathcal{C} \setminus \Gamma$. Then, $(\mathcal{C} \setminus U_\Gamma(S)) \cup S \notin \Gamma$.*

For the proof of Lemma 3, we refer the to the full version [1].

The knowledge extractor will be restricted to sampling challenges that are potentially useful. The value t_Γ defines the maximum number of accepting transcripts that the extractor has to find, before it succeeds and obtains the accepting transcripts for a set $S \in \Gamma$. The efficiency of our knowledge extractor will depend on t_Γ . A formal definition is given below. Further, in Sect. 5, we describe the monotone structure and corresponding k -values for three (classes of) interactive proofs and explain their relevance.

Definition 8 (t -value). *Let (Γ, \mathcal{C}) be a monotone structure and $S \subseteq \mathcal{C}$. Then*

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

Further,

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

It is easily seen that $t_\Gamma(S) = 0$ if and only if $S \in \Gamma$ or $\Gamma = \emptyset$. Further, the following lemma shows that adding an element $c \in U_\Gamma(S)$ to S decreases the corresponding k -value. This lemma plays a pivotal role in our recursive extraction algorithm.

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

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

For the proof of Lemma 4, we refer the to the full version [1].

As in [3], we describe our technical results in a more abstract language. This will later allow us to easily derive composition results and handle more complicated scenarios, such as multi-round interactive proofs and parallel compositions. To this end, let us consider a finite set \mathcal{C} , a probabilistic algorithm $\mathcal{A}: \mathcal{C} \rightarrow \{0, 1\}^*$ and a verification function $V: \mathcal{C} \times \{0, 1\}^* \rightarrow \{0, 1\}$. An output $y \leftarrow \mathcal{A}(c)$ of the algorithm \mathcal{A} on input $c \in \mathcal{C}$ is said to be *accepting* or *correct* if $V(c, y) = 1$. The success probability of \mathcal{A} is denoted as

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

where C is uniformly random in \mathcal{C} . The obvious instantiation of \mathcal{A} is given by a deterministic dishonest prover \mathcal{P}^* attacking an interactive proof Π on input x . Note that even though it is sufficient to consider deterministic provers \mathcal{P}^* , we allow the algorithm \mathcal{A} to be probabilistic. This generalization is essential when considering multi-round interactive proofs and parallel repetitions [3].

Now let $\Gamma \subseteq 2^{\mathcal{C}}$ be a nonempty monotone structure. Then, for any $S \subset \mathcal{C}$ with $U_\Gamma(S) \neq \emptyset$, we define

$$\epsilon_\Gamma(\mathcal{A}, S) := \Pr(V(C, \mathcal{A}(C)) = 1 \mid C \in U_\Gamma(S)).$$

Typically, $U_\Gamma(\emptyset) = \mathcal{C}$ and thus $\epsilon(\mathcal{A}) = \epsilon_\Gamma(\mathcal{A}, \emptyset)$, i.e., all challenges $c \in \mathcal{C}$ are potentially useful. However, this is not necessarily the case.

Given oracle access to \mathcal{A} , the goal of the extractor is to find *correct* outputs y_1, \dots, y_k for challenges $c_1, \dots, c_k \in \mathcal{C}$ such that $\{c_1, \dots, c_k\} \in \Gamma$, i.e., such that $V(c_i, y_i) = 1$ for all i . If \mathcal{A} corresponds to a dishonest prover attacking a Γ -out-of- \mathcal{C} special-sound interactive proof on some input x , a witness w for statement x can be efficiently computed from the outputs y_1, \dots, y_k .

Let us further define the following quality measure for the algorithm \mathcal{A} :

$$\delta_\Gamma(\mathcal{A}) := \min_{S \notin \Gamma} \Pr(V(C, \mathcal{A}(C)) = 1 \mid C \notin S). \tag{1}$$

The value $\delta_\Gamma(\mathcal{A})$ defines a “punctured” success probability of \mathcal{A} , i.e., it equals the success probability of \mathcal{A} when the challenge c is sampled uniformly at random from some set $\mathcal{C} \setminus S \supseteq U_\Gamma(S)$ such that S is not in the monotone structure. We will show that the value $\delta_\Gamma(\mathcal{A})$ measures how well we can extract from the algorithm \mathcal{A} . The value $\delta_\Gamma(\mathcal{A})$ is a generalization of the measure

$$\delta_k(\mathcal{A}) := \min_{S \subseteq \mathcal{C}: |S|=k-1} \Pr(V(C, \mathcal{A}(C)) = 1 \mid C \notin S),$$

defined in [3].⁴ However, when restricting to threshold monotone structures, there is a syntactic difference between the definitions of $\delta_k(\mathcal{A})$ and $\delta_\Gamma(\mathcal{A})$. To see this, let \mathcal{T}_k denote the monotone structure containing all subsets of \mathcal{C} with cardinality at least k . Then, in the definition of $\delta_k(\mathcal{A})$ the minimum is over all

⁴ In the original version of [3], the restriction was $|S| < k$. Here, when considering 3-round protocols, this makes no difference, but it does for the multi-round case.

sets of cardinality *exactly* $k - 1$, whereas the corresponding $\delta_{\mathcal{T}_k}(\mathcal{A})$ is a minimum over all sets of size *at most* $k - 1$. In the threshold case this makes no difference: it is easily seen that there always exists a (maximal) set of size $k - 1$ that minimizes $\delta_{\mathcal{T}_k}(\mathcal{A})$ and so indeed $\delta_{\mathcal{T}_k}(\mathcal{A}) = \delta_k(\mathcal{A})$. A similar result does not hold for arbitrary access structures, i.e., in general the minimum may not be attained by a maximal set $S \notin \Gamma$. This issue will reoccur in a more substantial way when addressing multi-round protocols.

For any set $T \in 2^{\mathcal{C}} \setminus \Gamma$, we also define

$$\delta_{\Gamma}(\mathcal{A}, T) := \min_{S: S \cup T \notin \Gamma} \Pr(V(C, \mathcal{A}(C)) = 1 \mid C \notin S).$$

Since $S \cup T \notin \Gamma$ implies $S \cup T' \notin \Gamma$ for all $T' \subseteq T$, it follows that

$$\delta_{\Gamma}(\mathcal{A}, T') \leq \delta_{\Gamma}(\mathcal{A}, T), \quad \forall T' \subseteq T. \tag{2}$$

Further, by Lemma 3, it follows that $(\mathcal{C} \setminus U_{\Gamma}(T)) \cup T \notin \Gamma$ for all $T \notin \Gamma$. Hence,

$$\begin{aligned} \delta_{\Gamma}(\mathcal{A}, T) &= \min_{S: S \cup T \notin \Gamma} \Pr(V(C, \mathcal{A}(C)) = 1 \mid C \notin S) \\ &\leq \Pr(V(C, \mathcal{A}(C)) = 1 \mid C \notin \mathcal{C} \setminus U_{\Gamma}(T)) \\ &= \Pr(V(C, \mathcal{A}(C)) = 1 \mid C \in U_{\Gamma}(T)) \\ &= \epsilon_{\Gamma}(\mathcal{A}, T). \end{aligned} \tag{3}$$

We are now ready to define and analyze our extraction algorithm for Γ -out-of- \mathcal{C} special-sound interactive Σ -protocols. The extractor is defined in Fig. 1 and its properties are summarized in the following lemma.

Lemma 5 (Extraction Algorithm - Σ -protocols). *Let (Γ, \mathcal{C}) be a nonempty monotone structure and let $V: \mathcal{C} \times \{0, 1\}^* \rightarrow \{0, 1\}$. Then there exists an oracle algorithm \mathcal{E}_{Γ} 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 $2t_{\Gamma} - 1$ queries to \mathcal{A} and, with probability at least $\delta_{\Gamma}(\mathcal{A})/t_{\Gamma}$, it outputs pairs $(c_1, y_1), (c_2, y_2), \dots, (c_k, y_k) \in \mathcal{C} \times \{0, 1\}^*$ with $V(c_i, y_i) = 1$ for all i and $\{c_1, \dots, c_k\} \in \Gamma$.*

Proof. The extractor $\mathcal{E}_{\Gamma}^{\mathcal{A}}(S)$ is formally defined in Fig. 1. It takes as input a subset $S \in 2^{\mathcal{C}} \setminus \Gamma$. The input S represents the set of accepting challenges that the extractor has already found, i.e., the goal of $\mathcal{E}_{\Gamma}^{\mathcal{A}}(S)$ is to find pairs (c_i, y_i) such that $V(c_i, y_i) = 1$ and $\{c_1, \dots, c_k\} \cup S \in \Gamma$. Further, we define

$$\mathcal{E}_{\Gamma}^{\mathcal{A}} := \mathcal{E}_{\Gamma}^{\mathcal{A}}(\emptyset).$$

First note that, since $\Gamma \neq \emptyset$ and thus $U_{\Gamma}(S) \neq \emptyset$ for all $S \notin \Gamma$, the extractor is well-defined. Let us now analyze the success probability and the expected number of \mathcal{A} -queries of the extractor.

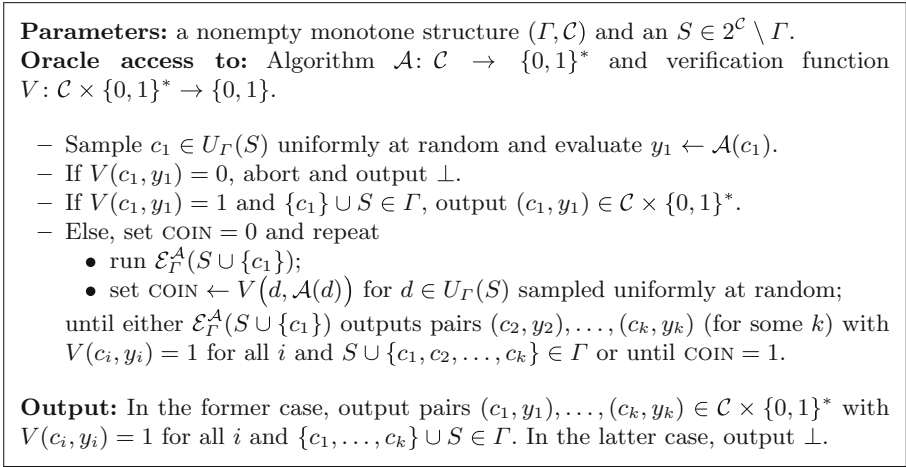


Fig. 1. Recursive Expected Polynomial Time Extractor $\mathcal{E}_{\Gamma}^{\mathcal{A}}(S)$.

Success Probability. By induction over $t_{\Gamma}(S)$, we will prove that $\mathcal{E}_{\Gamma}^{\mathcal{A}}(S)$ succeeds with probability at least

$$\frac{\delta_{\Gamma}(\mathcal{A}, S)}{t_{\Gamma}(S)}.$$

We first consider the base case. To this end, let $S \subseteq \mathcal{C}$ with $t_{\Gamma}(S) = 1$. Then, by Lemma 4, for all $c_1 \in U_{\Gamma}(S)$, it holds that $t_{\Gamma}(S \cup \{c_1\}) = 0$ and thus $S \cup \{c_1\} \in \Gamma$. Therefore, the extractor succeeds if and only if $V(c_1, \mathcal{A}(c_1)) = 1$ for the c_1 sampled from $U_{\Gamma}(S)$. Hence, the success probability of the extractor equals

$$\epsilon_{\Gamma}(\mathcal{A}, S) \geq \delta_{\Gamma}(\mathcal{A}, S),$$

where the inequality follows from Eq. 3. This proves the bound on the success probability for the base case $t_{\Gamma}(S) = 1$.

Let us now consider an arbitrary subset $S \subseteq \mathcal{C}$ with $t_{\Gamma}(S) > 1$ and assume that the claimed bound holds for all subsets $T \subseteq \mathcal{C}$ with $t_{\Gamma}(T) < t_{\Gamma}(S)$.

In the first step, the extractor succeeds with probability $\epsilon_{\Gamma}(\mathcal{A}, S)$ in finding a $c_1 \in U_{\Gamma}(S)$ and $y_1 \leftarrow \mathcal{A}(c_1)$ with $V(c_1, y_1) = 1$. If $\{c_1\} \cup S \in \Gamma$, the extractor has successfully completed its task. If not, the extractor starts running two geometric experiments until one of them finishes. In the first geometric experiment the extractor repeatedly runs $\mathcal{E}_{\Gamma}^{\mathcal{A}}(S \cup \{c_1\})$. By Lemma 4, it holds that $t_{\Gamma}(S \cup \{c_1\}) < t_{\Gamma}(S)$. Hence, by the induction hypothesis, $\mathcal{E}_{\Gamma}^{\mathcal{A}}(S \cup \{c_1\})$ succeeds with probability

$$p \geq \frac{\delta_{\Gamma}(\mathcal{A}, S \cup \{c_1\})}{t_{\Gamma}(S \cup \{c_1\})} \geq \frac{\delta_{\Gamma}(\mathcal{A}, S)}{t_{\Gamma}(S) - 1},$$

where the second inequality follows from Eq. 2 and Lemma 4. In the second geometric experiment, the extractor tosses a coin that returns heads with probability

$$q := \epsilon_\Gamma(\mathcal{A}, S).$$

The second step of the extractor succeeds if the second geometric experiment does not finish before the first, and so by Lemma 1 this probability is lower bounded as follows

$$\begin{aligned} \Pr(\text{Geo}(p) \leq \text{Geo}(q)) &\geq \frac{p}{p+q} \geq \frac{\frac{\delta_\Gamma(\mathcal{A}, S)}{t_\Gamma(S)-1}}{\frac{\delta_\Gamma(\mathcal{A}, S)}{t_\Gamma(S)-1} + \epsilon_\Gamma(\mathcal{A}, S)} \\ &\geq \frac{\frac{\delta_\Gamma(\mathcal{A}, S)}{t_\Gamma(S)-1}}{\frac{\epsilon_\Gamma(\mathcal{A}, S)}{t_\Gamma(S)-1} + \epsilon_\Gamma(\mathcal{A}, S)} = \frac{\delta_\Gamma(\mathcal{A}, S)}{t_\Gamma(S) \cdot \epsilon_\Gamma(\mathcal{A}, S)}, \end{aligned}$$

where the second inequality follows from the monotonicity of the function $x \mapsto \frac{x}{x+q}$ and the third inequality follows from the fact that $\delta_\Gamma(\mathcal{A}, S) \leq \epsilon_\Gamma(\mathcal{A}, S)$ (Eq. 3).

Since the first step of the extractor succeeds with probability $\epsilon_\Gamma(\mathcal{A}, S)$, it follows that $\mathcal{E}_\Gamma^{\mathcal{A}}(S)$ succeeds with probability at least $\delta_\Gamma(\mathcal{A}, S)/t_\Gamma(S)$ for all $S \in 2^{\mathcal{C}} \setminus \Gamma$, which proves the claimed bound. In particular, $\mathcal{E}_\Gamma^{\mathcal{A}}$ succeeds with probability at least $\delta_\Gamma(\mathcal{A})/t_\Gamma$.

Expected Number of \mathcal{A} -Queries. By induction over $t_\Gamma(S)$, we will prove that the expected number of \mathcal{A} -queries $Q_\Gamma(S)$ made by $\mathcal{E}_\Gamma^{\mathcal{A}}(S)$ is upper bounded as follows:

$$Q_\Gamma(S) \leq 2t_\Gamma(S) - 1.$$

We first consider the base case. To this end, let $S \subseteq \mathcal{C}$ with $t_\Gamma(S) = 1$. In this case, $\{c_1\} \cup S \in \Gamma$ for all $c_1 \in U_\Gamma(S)$. Hence, $\mathcal{E}_\Gamma^{\mathcal{A}}(S)$ either succeeds or fails after making exactly one \mathcal{A} -query, i.e., $Q_\Gamma(S) = 1 = 2t_\Gamma(S) - 1$, which proves the base case.

Let us now consider an arbitrary subset $S \subseteq \mathcal{C}$ with $t_\Gamma(S) > 1$ and assume that $Q_\Gamma(T) \leq 2t_\Gamma(T) - 1$ for all subsets $T \subseteq \mathcal{C}$ with $t_\Gamma(T) < t_\Gamma(S)$.

The extractor $\mathcal{E}_\Gamma^{\mathcal{A}}(S)$ first samples $c_1 \leftarrow_R U_\Gamma(S)$ uniformly at random and evaluates $y_1 \leftarrow \mathcal{A}(c_1)$. This requires exactly one \mathcal{A} -query. After this step the extractor aborts with probability $1 - \epsilon_\Gamma(\mathcal{A}, S)$. Otherwise, and if $\{c_1\} \cup S \notin \Gamma$, it continues running the two geometric experiments until either one of them finishes. The second geometric experiment finishes in an expected number of $1/\epsilon_\Gamma(\mathcal{A}, S)$ trials and requires exactly one \mathcal{A} -query per trial. Hence, the total expected number of trials for both experiments is at most $1/\epsilon_\Gamma(\mathcal{A}, S)$. Further, since $t_\Gamma(S \cup \{c_1\}) < t_\Gamma(S)$ (Lemma 4) and by the induction hypotheses, the expected number of \mathcal{A} -queries of the first geometric experiment is at most

$$Q_\Gamma(S \cup \{c_1\}) \leq 2t_\Gamma(S \cup \{c_1\}) - 1 \leq 2t_\Gamma(S) - 3,$$

per iteration, where the second inequality follows again from Lemma 4. Hence, every iteration of the repeat loop requires an expected number of at most $2t_\Gamma(S) - 2$ \mathcal{A} -queries.

From this it follows that

$$Q_\Gamma(S) \leq 1 + \epsilon_\Gamma(\mathcal{A}, S) \frac{2t_\Gamma(S) - 2}{\epsilon_\Gamma(\mathcal{A}, S)} = 2t_\Gamma(S) - 1,$$

for all $S \in 2^{\mathcal{C}} \setminus \Gamma$. In particular, $\mathcal{E}_\Gamma^{\mathcal{A}}$ requires an expected number of at most $2t_\Gamma - 1$ \mathcal{A} -queries, which completes the proof of the lemma. □

By basic probability theory, for any $S \notin \Gamma$,

$$\begin{aligned} \Pr(V(C, \mathcal{A}(C)) = 1 \mid C \notin S) &= \frac{\Pr(V(C, \mathcal{A}(C)) = 1 \wedge C \notin S)}{\Pr(C \notin S)} \\ &\geq \frac{\Pr(V(C, \mathcal{A}(C)) = 1) - \Pr(C \in S)}{\Pr(C \notin S)} \\ &= \frac{\epsilon(\mathcal{A}) - \Pr(C \in S)}{1 - \Pr(C \in S)} \\ &= \frac{\epsilon(\mathcal{A}) - |S|/|\mathcal{C}|}{1 - |S|/|\mathcal{C}|}. \end{aligned}$$

Hence, taking the minimum over all $S \notin \Gamma$ shows that

$$\delta_\Gamma(\mathcal{A}) \geq \frac{\epsilon(\mathcal{A}) - \kappa_\Gamma}{1 - \kappa_\Gamma}, \tag{4}$$

where $\kappa_\Gamma = \max_{S \notin \Gamma} |S|/|\mathcal{C}|$. In Γ -out-of- \mathcal{C} special-sound interactive proofs, a dishonest prover can potentially take any $S \notin \Gamma$ and choose the first message so that it will succeed if the verifier chooses a challenge $c \in S$. Hence, κ_Γ equals the trivial cheating strategy for Γ -out-of- \mathcal{C} special-sound interactive proofs.

Since the extractor succeeds with probability at least $\delta_\Gamma(\mathcal{A})/t_\Gamma$, the following theorem follows.

Theorem 1. *Let $(\mathcal{P}, \mathcal{V})$ be a Γ -out-of- \mathcal{C} special-sound Σ -protocol such that t_Γ is polynomial in the size $|x|$ of the public input statement x of $(\mathcal{P}, \mathcal{V})$ and sampling from $U_\Gamma(S)$ takes polynomial time (in $|x|$) for all S with $|S| < t_\Gamma$. Then $(\mathcal{P}, \mathcal{V})$ is knowledge sound with knowledge error $\kappa_\Gamma = \max_{S \notin \Gamma} |S|/|\mathcal{C}|$.*

5 Examples

In this section, we describe three very simple interactive proofs and their special-soundness properties. The first example shows that for the special case of k -out-of- N special-soundness notion, we recover the known results. The second and third example present techniques that have found numerous applications, but cannot be analyzed via their threshold special-soundness properties, i.e., these interactive proofs require an alternative analysis. Our knowledge extractor offers the means to easily handle these interactive proof as well. Finally, the fourth

example shows that our generic techniques do not always suffice. In Sect. 7, we will consider a more complicated protocol and demonstrate how our techniques enable a knowledge soundness analysis of the multi-round protocol FRI [5].

Example 1 (Threshold Access Structures). Let \mathcal{C} be a finite set with cardinality N , and let Γ be the monotone structure that contains all subsets of \mathcal{C} of cardinality at least $k \leq N$. Then a Γ -out-of- \mathcal{C} special-sound interactive proof is also k -out-of- N special-sound. Moreover, $U_\Gamma(A) = \mathcal{C} \setminus A$ for all $A \notin \Gamma$, $t_\Gamma = k$, and $\kappa_\Gamma = (k - 1)/N$. Hence, in the case of k -out-of- N special-soundness, we recover the results from [3].

Example 2 (Standard Amortization Technique). Let \mathbb{F} be a finite field and let Ψ be an \mathbb{F} -linear map. The following amortization technique, known from Σ -protocol theory, allows a prover to prove knowledge of n Ψ -preimages x_1, \dots, x_n of P_1, \dots, P_n for essentially the cost of one. The amortization technique is a 2-round protocol that proceeds as follows. First, the verifier samples a challenge vector $\mathbf{c} = (c_1, \dots, c_n) \in \mathbb{F}^n$ uniformly at random. Second, upon receiving the challenge vector \mathbf{c} , the prover responds with the element $z = \sum_{i=1}^n c_i x_i$. Finally, the verifier checks that $\Psi(z) = \sum_{i=1}^n c_i P_i$. Hence, instead of sending n preimages the prover only has to send one preimage.

The n preimages x_1, \dots, x_n of P_1, \dots, P_n can be extracted from accepting transcripts $(\mathbf{c}_1, z_1), \dots, (\mathbf{c}_k, z_k)$ if the challenge vectors $\mathbf{c}_1, \dots, \mathbf{c}_k$ span the vector space \mathbb{F}^n . Hence, the amortization protocol is Γ -out-of- \mathbb{F}^n special-sound, where Γ is the monotone structure that contains all subsets spanning \mathbb{F}^n . Further, $t_\Gamma = n$, $U_\Gamma(A) = \mathbb{F}^n \setminus \text{span}(A)$ for all $A \notin \Gamma$; and $\kappa_\Gamma = 1/|\mathbb{F}|$; thus, we obtain optimal knowledge soundness.

At the same time, the amortization protocol is $(|\mathbb{F}|^{n-1} + 1)$ -out-of- $|\mathbb{F}|^n$ special-sound, i.e., the threshold special-soundness parameter of this protocol is $|\mathbb{F}|^{n-1} + 1$, which is much larger than $t_\Gamma = n$. In fact, the parameter $|\mathbb{F}|^{n-1} + 1$ is typically not polynomially bounded, in which case knowledge soundness can not be derived from this threshold special-soundness property.

Example 3 (Merkle Tree Commitments). Let us now consider an interactive proof for proving knowledge of the opening of a Merkle tree commitment P , i.e., P is the root of a Merkle tree and the prover claims to know all n leaves. To verify this claim, the verifier selects a subset S of k (distinct) indices between 1 and n uniformly at random. The prover sends the corresponding leaves together with their validation paths, which are checked by the verifier.

An opening of the commitment P can be extracted from accepting transcripts $(S_1, z_1), \dots, (S_\ell, z_\ell)$ if the subsets S_i cover $\{1, \dots, n\}$. Hence, this interactive proof is Γ -out-of- \mathcal{C} , where

$$\mathcal{C} = \{S \subseteq \{1, \dots, n\} : |S| = k\} \quad \text{and} \quad \Gamma = \{\mathcal{D} \subseteq \mathcal{C} : \bigcup_{S \in \mathcal{D}} S = \{1, \dots, n\}\}.$$

Further, $t_\Gamma = n - k + 1$, $U_\Gamma(\mathcal{D}) = \{A \in \mathcal{C} : A \not\subseteq \bigcup_{S \in \mathcal{D}} S\}$ for all $\mathcal{D} \notin \Gamma$, and $\kappa_\Gamma = (n - k)/n$; thus, we obtain optimal knowledge soundness.

The threshold special-soundness parameter of this protocol is $\binom{n-1}{k} + 1$ which is typically much larger than $t_\Gamma = n - k + 1$. Hence, also in this case our generalization provides a much more efficient knowledge extractor.

This simple interactive proof is an essential component in many more complicated protocols based on probabilistically checkable proofs (PCPs), interactive oracle proofs (IOPs) or MPC-in-the-head.

Example 4 (Parallel Repetition). Finally, we consider an example where our generic technique does not work. To this end, let Π^t be the t -fold parallel composition of a k -out-of- N special-sound interactive proof Π with challenge set \mathcal{C} , i.e., Π^t has challenge set \mathcal{C}^t . Then, as discussed in the introduction, Π^t is $((k - 1)^t + 1)$ -out-of- N^t special-sound, i.e., its threshold special-soundness parameter $(k - 1)^t + 1$ grows exponentially in t (if $k > 2$).

The parallel repetition Π^t is also Γ -out-of- \mathcal{C}^t special-sound, where Γ contains all subsets of challenge vectors $\mathbf{c} \in \mathcal{C}^t$ such that there is one position $1 \leq i \leq t$ where the challenge vectors feature at least k different values. Then, $\kappa_\Gamma = (k - 1)^t / N^t$. However, $t_\Gamma = (k - 1)^t + 1$, i.e., t_Γ equals the threshold special-soundness parameter and grows exponentially in t . Hence, in this particular example, the correct access structure does not yield an efficient extractor. Fortunately, here we can apply the parallel repetition result of [3].

6 Knowledge Extraction for Multi-round Interactive Proofs

Let us now move to the analysis of multi-round interactive proofs $(\mathcal{P}, \mathcal{V})$. To this end, we first generalize the notion of Γ -out-of- \mathcal{C} special-soundness to multi-round interactive proofs. A $2\mu + 1$ -round interactive proof is said to be $(\Gamma_1, \dots, \Gamma_\mu)$ -out-of- $(\mathcal{C}_1, \dots, \mathcal{C}_\mu)$ if there exists an efficient algorithm that can extract a witness from appropriate trees of transcripts. Before we formally define trees of transcripts, we first define the related trees of challenges.

Definition 9 (Tree of Challenges). *Let $(\Gamma_i, \mathcal{C}_i)$ be monotone structures for $1 \leq i \leq \mu$. A set containing a single challenge vector $(c_1, \dots, c_\mu) \in \mathcal{C}_1 \times \dots \times \mathcal{C}_\mu$ is also referred to as a $(1, \dots, 1)$ -tree of challenges. Further, for $1 \leq t \leq \mu$, a $(1, \dots, 1, \Gamma_t, \dots, \Gamma_\mu)$ -tree T_t of challenges is the union of several $(1, \dots, 1, \Gamma_{t+1}, \dots, \Gamma_\mu)$ -trees, such that*

- The first $t - 1$ coordinates of all $\mathbf{c} \in T_t \subseteq \mathcal{C}_1 \times \dots \times \mathcal{C}_\mu$ are equal;
- The t -th coordinates of the tree elements form an element in Γ_t , i.e.,

$$\{c \in \mathcal{C}_t : \exists (c_1, \dots, c_{t-1}, c, c_{t+1}, \dots, c_\mu) \in T_t\} \in \Gamma_t.$$

Trivially, the verifier’s messages in a transcript of a $2\mu + 1$ -round interactive proof with challenge sets $\mathcal{C}_1, \dots, \mathcal{C}_\mu$ form a $(1, \dots, 1)$ -tree of challenges. Hence, by adding the prover’s messages we obtain a $(1, \dots, 1)$ -tree of transcripts, and thus, in the obvious way, we obtain the notion of a tree of transcripts. The

only additional requirement is that the prover's messages *collide*, i.e., they are uniquely determined by the challenges received before sending the message. In particular, the first message of every transcript is the same. Note that if the transcripts are generated by a deterministic prover, this property is guaranteed to hold.

Definition 10 (Tree of Transcripts). *Let $(\Gamma_i, \mathcal{C}_i)$ be monotone structures for $1 \leq i \leq \mu$. Let $(\mathcal{P}, \mathcal{V})$ be a $2\mu + 1$ -round public-coin interactive proof with challenge sets $\mathcal{C}_1, \dots, \mathcal{C}_\mu$. A $(\Gamma_1, \dots, \Gamma_\mu)$ -tree of transcripts is a set of protocol transcripts, such that*

- *The corresponding set of challenge vectors, obtained by ignoring the prover's messages, is a $(\Gamma_1, \dots, \Gamma_\mu)$ -tree of challenges;*
- *The prover's messages collide, i.e., if two transcripts $(a_0, c_1, a_1, \dots, c_\mu, a_\mu)$ and $(a'_0, c'_1, a'_1, \dots, c'_\mu, a'_\mu)$ are both in the tree, and $c_i = c'_i$ for all $i \leq j$, then also $a_i = a'_i$ for all $i \leq j$.*

Prior works (e.g., [2, 7, 8]) considered (k_1, \dots, k_μ) -trees, where $k_i \in \mathbb{N}$ for all i . These are special cases of the above defined trees. More precisely, if $\Gamma_i = \{S \subseteq \mathcal{C}_i : |S| \geq k_i\}$, a (k_1, \dots, k_μ) -tree is the same as a $(\Gamma_1, \dots, \Gamma_\mu)$ -tree.

We are now ready to define a generalized multi-round special-soundness notion.

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 $1 \leq i \leq \mu$. 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 $(\Gamma_1, \dots, \Gamma_\mu)$ -out-of- $(\mathcal{C}_1, \dots, \mathcal{C}_\mu)$ special-sound if there exists a polynomial time algorithm that, on input a statement x and a $(\Gamma_1, \dots, \Gamma_\mu)$ -tree of accepting transcripts, outputs a witness $w \in R(x)$. We also say that $(\mathcal{P}, \mathcal{V})$ is $(\Gamma_1, \dots, \Gamma_\mu)$ -special-sound.*

Remark 3. The monotone access structure $(\Gamma_{\text{TREE}}(\mathbf{\Gamma}), \mathcal{C}_1 \times \dots \times \mathcal{C}_\mu)$, where $\mathbf{\Gamma} = (\Gamma_1, \dots, \Gamma_\mu)$ and

$$\Gamma_{\text{TREE}}(\Gamma_1, \dots, \Gamma_\mu) := \{S \subseteq \mathcal{C}_1 \times \dots \times \mathcal{C}_\mu : S \text{ contains a } (\Gamma_1, \dots, \Gamma_\mu)\text{-tree}\},$$

allows one to cast a multi-round $(\Gamma_1, \dots, \Gamma_\mu)$ -out-of- $(\mathcal{C}_1, \dots, \mathcal{C}_\mu)$ special-sound interactive proof as a Γ -out-of- \mathcal{C} special-sound interactive proof. Therefore, in principle, one could immediately apply the results from Sect. 4. However, typically, this results in an inefficient knowledge extractor. More precisely, the value $t_{\Gamma_{\text{TREE}}(\mathbf{\Gamma})}$, and thus the expected running time of the extractor, grows linearly in the product of the sizes of the challenge sets $\mathcal{C}_1, \dots, \mathcal{C}_{\mu-1}$. For this reason, our multi-round knowledge extractor will proceed recursively over the different rounds.

Our goal is now to prove that, for appropriate monotone structures, $(\Gamma_1, \dots, \Gamma_\mu)$ -out-of- $(\mathcal{C}_1, \dots, \mathcal{C}_\mu)$ special-soundness (tightly) implies knowledge soundness. As before, again borrowing the notation from [3], we present our

results in a more abstract language. To this end, let $\mathcal{A}: \mathcal{C}_1 \times \dots \times \mathcal{C}_\mu \rightarrow \{0, 1\}^*$ be a probabilistic algorithm and

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

a verification function. The success probability of \mathcal{A} is denoted as

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

where C is distributed uniformly at random over $\mathcal{C}_1 \times \dots \times \mathcal{C}_\mu$. The obvious instantiation of \mathcal{A} is again a deterministic prover \mathcal{P}^* attacking a $(\Gamma_1, \dots, \Gamma_\mu)$ -out-of- $(\mathcal{C}_1, \dots, \mathcal{C}_\mu)$ special-sound interactive proof.

It turns out that defining the multi-round version of δ_Γ is somewhat subtle. In the case of a \mathbf{k} -special sound protocol, it is defined in [3] as

$$\delta_{\mathbf{k}}^V(\mathcal{A}) := \min_{S_1 \dots S_\mu} \Pr(V(C, \mathcal{A}(C)) = 1 \mid C_1 \notin S_1, C_2 \notin S_2(C_1), C_3 \notin S_2(C_1, C_2), \dots)$$

where the minimum is over all sets $S_1 \subset 2^{\mathcal{C}_1}$ with $|S_1| = k_1 - 1$, all functions $S_2: \mathcal{C}_1 \rightarrow 2^{\mathcal{C}_2}$ with $|S_2(c_1)| = k_2 - 1$ for all $c_1 \in \mathcal{C}_1$, etc.⁵ Thus, the natural extension to $(\Gamma_1, \dots, \Gamma_\mu)$ -special-sound protocols would be to use the very same expression but minimize over all (maximal) sets $S_1 \subset 2^{\mathcal{C}_1}$ with $S_1 \not\subseteq \Gamma_1$, all functions $S_2: \mathcal{C}_1 \rightarrow 2^{\mathcal{C}_2}$ with $S_2(c_1)$ (maximal and) not in Γ_2 for all $c_1 \in \mathcal{C}_1$, etc.

However, writing $\Gamma = (\Gamma_1, \Gamma_2, \dots, \Gamma_\mu)$, it turns out that defining δ_Γ^V in this way will not lead to the desired results. In essence, the problem lies in the fact that the condition $C_2 \notin S_2(C_1)$ may bias the distribution of C_1 , namely when $S_2(c_1)$ has different cardinality for different choices of c_1 . This issue is avoided in the threshold case by requiring the S_i 's to be maximal sets; here in the general case, this does not work, since different maximal sets may have different cardinality.

Because of this reason, we define δ_Γ^V by the following, harder to comprehend, expression:

$$\delta_\Gamma^V(\mathcal{A}) := \min_{S_1 \dots S_\mu} \sum_{\mathbf{c}} \Pr(V(C, \mathcal{A}(C)) = 1 \wedge C = \mathbf{c} \mid C_1 \notin S_1, C_2 \notin S_2(c_1), \dots, C_\mu \notin S_\mu(c_1, \dots, c_\mu)), \tag{5}$$

where, as in the above approach, the minimum is over all sets $S_1 \subset 2^{\mathcal{C}_1}$ with $S_1 \not\subseteq \Gamma_1$, all functions $S_2: \mathcal{C}_1 \rightarrow 2^{\mathcal{C}_2}$ with $S_2(c_1) \not\subseteq \Gamma_2$ for all $c_1 \in \mathcal{C}_1$, etc.

⁵ In the original version of [3], the restriction was $|S_i| < k_i$, i.e., the sets were not required to be maximal (this makes no difference for $\mu = 1$, but it does for the multi-round case, where the min is not necessarily attained by maximal sets). However, in an updated version, this was changed to the above (in essence because of a similar issue as discussed below).

Remark 4. Note that, in the special case of 3-round interactive proofs, i.e., if $\mu = 1$, it holds that

$$\begin{aligned} \delta_{\Gamma}^V(\mathcal{A}) &= \min_{S \notin \Gamma} \sum_{\mathbf{c}} \Pr(V(C, \mathcal{A}(C)) = 1 \wedge C = \mathbf{c} \mid C \notin S) \\ &= \min_{S \notin \Gamma} \Pr(V(C, \mathcal{A}(C)) = 1 \mid C \notin S). \end{aligned}$$

Hence, the multi-round version of δ defined in Eq. 5, is indeed a generalization of the 3-round version defined in Eq. 1.

Remark 5. Let us consider the multi-round threshold case, i.e., let $\mathcal{T}_{\mathbf{k}} = (\mathcal{T}_{k_1}, \dots, \mathcal{T}_{k_\mu})$ with \mathcal{T}_{k_i} the monotone structure containing all subsets of \mathcal{C}_i with cardinality at least k_i for all i . Then, although not immediately obvious, it turns out that $\delta_{\mathcal{T}_{\mathbf{k}}}(\mathcal{A}) = \delta_{\mathbf{k}}(\mathcal{A})$ for all \mathcal{A} .

By observing that for the non-vanishing terms in the sum, exploiting the independence of $V(\mathbf{c}, \mathcal{A}(\mathbf{c}))$ and C for a fixed \mathbf{c} ,

$$\begin{aligned} \Pr(V(C, \mathcal{A}(C)) = 1 \mid C = \mathbf{c}, C_1 \notin S_1, C_2 \notin S_2(c_1), \dots) \\ = \Pr(V(\mathbf{c}, \mathcal{A}(\mathbf{c})) = 1 \mid C = \mathbf{c}, C_1 \notin S_1, C_2 \notin S_2(c_1), \dots) \\ = \Pr(V(\mathbf{c}, \mathcal{A}(\mathbf{c})) = 1), \end{aligned}$$

we can re-write the definition as

$$\begin{aligned} \delta_{\Gamma}^V(\mathcal{A}) &= \min_{S_1 \dots S_\mu} \sum_{\mathbf{c}} \Pr(V(\mathbf{c}, \mathcal{A}(\mathbf{c})) = 1) \Pr(C = \mathbf{c} \mid C_1 \notin S_1, C_2 \notin S_2(c_1), \dots) \\ &= \min_{S_1 \dots S_\mu} \sum_{\mathbf{c}} \Pr(V(\mathbf{c}, \mathcal{A}(\mathbf{c})) = 1) \Pr(C_1 = c_1 \mid C_1 \notin S_1) \dots \\ &\quad \dots \Pr(C_\mu = c_\mu \mid C_\mu \notin S_\mu(c_1, \dots, c_{\mu-1})). \end{aligned}$$

This shows that the definition captures the success probability of \mathcal{A} when the challenges are samples as follows (for given sets/functions S_1, S_2, \dots , over which the minimum is then taken): c_1 is sampled uniformly at random subject to being outside of S_1 . Then, c_2 is sampled uniformly at random subject to being outside of $S_2(c_1)$. And so forth. We repeat, in general this is *not* the same as sampling c_1, \dots, c_μ uniformly at random subject to $c_1 \notin S_1, c_2 \notin S_2(c_1)$, etc., which biases the choice of c_1 towards those for which $S_2(c_1)$ is small(er), while with the above sampling there is no bias on c_1 (beyond the exclusion from S_1). Defining δ_{Γ}^V in this way is crucial to our work. Oftentimes, the verification function V is clear from context, in which case we simply write $\delta_{\Gamma}(\mathcal{A})$ instead of $\delta_{\Gamma}^V(\mathcal{A})$.

Any choice of sets/functions S_1, \dots, S_μ considered in the minimization in Eq. 5 defines a subset

$$X = \{(c_1, \dots, c_\mu) \in \mathcal{C}_1 \times \dots \times \mathcal{C}_\mu \mid c_1 \in S_1 \vee \dots \vee c_\mu \in S_\mu(c_1, \dots, c_{\mu-1})\}$$

that does not contain a Γ -tree. Hence, again the success probability is punctured by removing some set X from which we cannot extract and thus, for

which a dishonest prover may (potentially) be successful. Moreover, every subset of $\mathcal{C}_1 \times \dots \times \mathcal{C}_\mu$ that does not contain a Γ -tree is contained in a set X of this form. Hence, if a prover has positive success probability outside all such subsets X , i.e., if $\delta_\Gamma^V(\mathcal{A}) > 0$, then extraction of a Γ -tree of accepting transcripts is in principle possible. However, it is far less obvious that extraction can also be done efficiently. The following lemma shows that, for appropriate monotone structures $(\Gamma_i, \mathcal{C}_i)$, an efficient extraction algorithm indeed exists. This is a generalization of [3, Lemma 4]. Using the notation we introduced here, their proof almost immediately carries over to this more generic setting. For completeness, we present the proof below.

Lemma 6 (Multi-round Extraction Algorithm). *Let $\Gamma = (\Gamma_1, \dots, \Gamma_\mu)$ and $\mathcal{C} = \mathcal{C}_1 \times \dots \times \mathcal{C}_\mu$ be such that $(\Gamma_i, \mathcal{C}_i)$ are nonempty monotone structures for all i . Further, let $T := \prod_{i=1}^\mu t_{\Gamma_i}$ and $V: \mathcal{C} \times \{0, 1\}^* \rightarrow \{0, 1\}$. Then, there exists an algorithm $\mathcal{E}^{\mathcal{A}}$ so that, given oracle access to any (probabilistic) algorithm $\mathcal{A}: \mathcal{C} \rightarrow \{0, 1\}^*$, $\mathcal{E}^{\mathcal{A}}$ requires an expected number of at most $2^\mu \cdot T$ queries to \mathcal{A} and, with probability at least $\delta_\Gamma(\mathcal{A})/T$, outputs pairs $(\mathbf{c}_i, y_i) \in \mathcal{C} \times \{0, 1\}^*$ such that $\{\mathbf{c}_i\}_i$ is a Γ -tree with $V(\mathbf{c}_i, y_i) = 1$ for all i .*

For the proof of Lemma 6, we refer to the full version [1].

Let us now derive a lower bound on the value $\delta_\Gamma^V(\mathcal{A})$. To this end, for $\mathbf{c} = (c_1, \dots, c_\mu) \in \mathcal{C}_1 \times \dots \times \mathcal{C}_\mu$, we write $V(\mathbf{c})$ as a shorthand for $V(\mathbf{c}, \mathcal{A}(\mathbf{c}))$. Furthermore, for any fixed choices of S_1, S_2, \dots, S_μ , as in the definition of $\delta_\Gamma(\mathcal{A})$ (Eq. 5), we introduce the event

$$\Omega(\mathbf{c}) := [C_1 \notin S_1 \wedge C_2 \notin S_2(c_1) \wedge \dots \wedge C_\mu \notin S_\mu(c_1, \dots, c_{\mu-1})].$$

Then,

$$\begin{aligned} \sum_{\mathbf{c}} \Pr(V(\mathbf{c}) = 1 \wedge C = \mathbf{c} \mid \Omega(\mathbf{c})) &\geq \sum_{\mathbf{c}} \Pr(V(\mathbf{c}) = 1 \wedge \Omega(\mathbf{c}) \wedge C = \mathbf{c}) \\ &= \sum_{\mathbf{c}} \Pr(V(C) = 1 \wedge \Omega(C) \wedge C = \mathbf{c}) \\ &= \Pr(V(C) = 1 \wedge \Omega(C)) \\ &\geq \Pr(V(C) = 1) - \Pr(\neg\Omega(C)). \end{aligned}$$

Now note that

$$\begin{aligned} \Pr(\neg\Omega(C)) &= 1 - \Pr(\Omega(C)) \\ &= 1 - \Pr(C_1 \notin S_1) \Pr(C_2 \notin S_2(C_1) \mid C_1 \notin S_1) \dots \\ &\leq 1 - \left(1 - \max_{S_1 \notin \Gamma_1} \frac{|S_1|}{|C_1|}\right) \left(1 - \max_{S_2 \notin \Gamma_1} \frac{|S_2|}{|C_2|}\right) \dots \\ &= \kappa_\Gamma, \end{aligned} \tag{6}$$

where

$$\kappa_\Gamma := \max_{S \notin \Gamma_{\text{Tree}}(\Gamma)} \frac{|S|}{|C|} = 1 - \prod_{i=1}^\mu \left(1 - \max_{S_i \notin \Gamma_i} \frac{|S_i|}{|C_i|}\right) = 1 - \prod_{i=1}^\mu (1 - \kappa_{\Gamma_i}).$$

We thus obtain that

$$\delta_{\Gamma}(\mathcal{A}) \geq \epsilon(\mathcal{A}) - \kappa_{\Gamma}. \tag{7}$$

These observations complete the proof of the following theorem.

Theorem 2. *Let $(\mathcal{P}, \mathcal{V})$ be a $(\Gamma_1, \dots, \Gamma_{\mu})$ -out-of- $(\mathcal{C}_1, \dots, \mathcal{C}_{\mu})$ special-sound interactive proof such that $T_{\Gamma} = \prod_{i=1}^{\mu} t_{\Gamma_i}$ is polynomial in the size $|x|$ of the public input statement x of $(\mathcal{P}, \mathcal{V})$ and sampling from $U_{\Gamma_i}(S_i)$ takes polynomial time (in $|x|$) for all $1 \leq i \leq \mu$ and $S_i \subset \mathcal{C}_i$ with $|S_i| < t_{\Gamma_i}$. Then $(\mathcal{P}, \mathcal{V})$ is knowledge sound with knowledge error*

$$\kappa_{\Gamma} = 1 - \prod_{i=1}^{\mu} \left(1 - \max_{S_i \notin \Gamma_i} \frac{|S_i|}{|\mathcal{C}_i|} \right).$$

7 Analysis of the FRI-Protocol

In this section we show how to use our generalized notion of special-soundness to demonstrate the existence of a quasi-polynomial time knowledge extractor with essentially optimal success probability for the Fast Reed-Solomon Interactive Oracle Proof of Proximity due to Ben-Sasson et al. [5], assuming it has been compiled into an interactive proof the natural way (i.e., the oracles are replaced by compact commitments to the vectors with a local opening functionality). We first provide the necessary background on the protocol before providing our analysis. We remark that we use ideas that were implicit in prior works; our main aim in this section is to demonstrate the utility of our generalized special-soundness notion and the accompanying knowledge extractor.

7.1 Preliminaries on Reed-Solomon Codes

Let \mathbb{F} be a finite field of cardinality q and $S \subseteq \mathbb{F}$. Given a polynomial $f(X) \in \mathbb{F}[X]$ we let $f(S) = (f(s))_{s \in S}$ denote the vector of evaluations of f over the domain S (given in some arbitrary, but fixed, order). For an integer ℓ we write S^{ℓ} for the set of ℓ -powers of elements in S , i.e. $\{s^{\ell} : s \in S\}$.⁶

For any $0 \leq \rho \leq 1$, the Reed-Solomon code $\text{RS}[\mathbb{F}, S, \rho] \subseteq \mathbb{F}^{|S|}$ consists of all evaluations over the domain S of polynomials $F(X) \in \mathbb{F}[X]$ of degree less than $\rho|S|$. In notation,

$$\text{RS}[\mathbb{F}, S, \rho] := \{F(S) : F(X) \in \mathbb{F}[X] \wedge \deg(F) < \rho|S|\}.$$

In the sequel we will assume S is a multiplicative subgroup of \mathbb{F}^* of order a power of 2, with the understanding that our analysis should generalize readily to other “smooth” evaluation domains for FRI protocols. We further set $\rho = 2^{-r}$

⁶ We use this somewhat cumbersome notation as we will later need to denote j -fold Cartesian products of sets, and for this operation we will use the standard notation S^j .

for an integer $r < \log_2(|S|)$, which implies $\rho|S| \in \mathbb{N}$ and that the dimension of $\text{RS}[\mathbb{F}, S, \rho]$ is precisely $\rho|S|$.

Letting $N = |S|$, we therefore have $S = \langle \omega \rangle = \{1, \omega, \omega^2, \dots, \omega^{N-1}\}$, where ω is a primitive N -th root of unity. Note then that $S^{\cdot 2} = \langle \omega^2 \rangle = \{1, \omega^2, \omega^4, \dots, \omega^{N-2}\}$ is a multiplicative subgroup of \mathbb{F}^* of order $N/2$. More generally, for any $j = 1, 2, \dots, \log_2(N)$, $S^{\cdot 2^j} = \langle \omega^{2^j} \rangle$ is multiplicative subgroup of \mathbb{F}^* of order $N/2^j$.

Given two polynomials $f(X), g(X) \in \mathbb{F}[X]$ we let $d_S(f, g) := |\{s \in S : f(s) \neq g(s)\}|$ denote the number of points $s \in S$ on which f and g differ. Equivalently, it denotes the (unnormalized) Hamming distance between the vectors $f(S)$ and $g(S)$.

Given a polynomial $f \in \mathbb{F}[X]$, we let

$$\delta_S(f) := \frac{\min_F \{d_S(f, F) : F \in \mathbb{F}[X], \deg(F) < \rho|S|\}}{|S|}.$$

In other words, $\delta_S(f)$ denotes the relative Hamming distance of $f(S)$ to a closest codeword in $\text{RS}[\mathbb{F}, S, \rho]$.

7.2 FRI-Protocol

Let \mathcal{O}^f be an oracle implementing some function $f : S \rightarrow \mathbb{F}$, which of course uniquely corresponds to a polynomial of degree less than $N = |S|$. We are interested in the situation where a prover claims that $f(X)$ is in fact a polynomial of degree $< \rho N$, i.e., that $f(S) \in \text{RS}[\mathbb{F}, S, \rho]$. In order to verify this, the verifier may make queries to \mathcal{O}^f , but it is easy to see that in order to catch a lying prover the verifier must query each $s \in S$ (or at least $\Omega(|S|)$ such points in order to catch the prover with good probability).

Thus, for soundness, we will be satisfied with rejecting oracles implementing functions that are *far* from low degree, i.e., such that $\delta_S(f) \geq \delta$. However, here as well we cannot hope to catch cheating verifiers without making at least $\rho N + 1$ queries (as ρN evaluations are always consistent with some polynomial of degree $< \rho N$). It turns out to be possible to make significantly less (i.e., just logarithmically many) oracle queries if we allow the verifier to *interact with* the prover.

The resulting protocols are referred to as *interactive oracle proofs of proximity* (IOPPs). In order to demonstrate the utility of our general special soundness notion, we will show how to analyze the Fast Reed-Solomon Interactive Oracle Proof of Proximity (FRI-protocol) [5].

In order to implement the oracle \mathcal{O}^f cryptographically, one makes use of a compact commitment scheme, typically via a Merkle tree [6]. In the following we denote the commitment to the vector $F(S) = (F(s))_{s \in S}$ with public parameters pp by $P \leftarrow \text{COM}_{\text{pp}}(F(S))$ and the local opening information for $s \in S$ as γ_s . For example, in the case of a Merkle tree the public parameters pp would be a description of the hash function used, while γ_s would give hash values for the co-path of the leaf corresponding to s . We also assume access to a procedure LOC_{pp}

which takes as input a commitment P , a domain element s , a value $y_s \in \mathbb{F}$ and the opening information γ_s and outputs 1 if and only if γ_s indeed certifies that P opens to y_s on the element s .

We can therefore view the (cryptographically compiled version of the) FRI-protocol as an interactive proof for the pair of relations $(\mathfrak{R}_0, \mathfrak{R}_\delta \cup \mathfrak{R}_{\text{coll}})$, where for a parameter $\beta \in [0, 1)$ we define

$$\mathfrak{R}_\beta := \left\{ (P, \text{pp}; F, B, (\gamma_s)_{s \in B}) : \deg(F) < \rho N \wedge |B| \geq (1 - \beta)N \wedge \forall s \in B, \text{LOC}_{\text{pp}}(P, s, F(s), \gamma_s) = 1 \right\},$$

while

$$\mathfrak{R}_{\text{coll}} := \left\{ (\text{pp}; s, y, y', \gamma, \gamma') : y \neq y' \wedge \text{LOC}_{\text{pp}}(P, s, y, \gamma) = 1 \wedge \text{LOC}_{\text{pp}}(P, s, y', \gamma') = 1 \right\}.$$

This means that completeness holds with respect to relation \mathfrak{R}_0 and soundness holds with respect to $\mathfrak{R}_\delta \cup \mathfrak{R}_{\text{coll}}$, where the latter refers to the “or-relation” which accepts a witness for one or the other instance. On the one hand, this says that a prover that committed to a low-degree polynomial will indeed convince the verifier of this fact. On the other hand, if a prover has a good probability of convincing the verifier then we can either extract a commitment to many coordinates that agree with a low-degree polynomial, or we can extract two distinct local openings from the same commitment (invalidating the binding property of the commitment).⁷

Folding. An important ingredient in the FRI-protocol is a folding operation. For our specific choice of S , it is defined as follows: for $f(X) \in \mathbb{F}[X]$ and $c \in \mathbb{F}$, we define

$$\text{Fold}(f(X), c) = g(X) \in \mathbb{F}[X]$$

such that

$$g(X^2) = \frac{f(X) + f(-X)}{2} + c \frac{f(X) - f(-X)}{2X}.$$

Intuitively, this folding operation considers the even-power monomials of $f(X)$ and the odd-power monomials separately, obtains from these terms two polynomials of degree $\deg(f)/2$, and takes a random linear combination of these polynomials. Importantly, the polynomial $g(X)$ can then naturally be viewed as having degree roughly $\deg(f)/2$ (i.e., the degree is halved) and its domain is naturally viewed as $S^2 = \langle \omega^2 \rangle$, which has order $N/2$. That is, the folded polynomial has its degree and domain halved.

A one round version of the FRI-protocol thus proceeds as follows. First, the prover commits to $F(S)$, where it promises that $F(S) \in \text{RS}[\mathbb{F}, S, \rho]$. The verifier

⁷ Observe that this is a concrete instantiation of the idea alluded to in Remark 1: we can either extract a witness to the desired relation, or a solution to a computationally hard problem.

picks a random challenge $c \in \mathbb{F}$, sends it to the prover, and the prover responds with the folding G of F around c . The verifier first checks that $\deg(G) < \rho N/2$. If yes, the verifier then chooses t points $s_1, \dots, s_t \in S$ (each uniformly at random and thus possibly colliding), and asks for the evaluations of F on all points $\pm s_i$. It then checks that these evaluations are consistent with G , i.e., that $G(s_i^2) = \frac{f(s_i)+f(-s_i)}{2} + c \frac{f(s_i)-f(-s_i)}{2s_i}$ for all $1 \leq i \leq t$, and of course that these are indeed the values the prover committed to initially.

7.3 Analyzing the FRI-Protocol

In order to analyze the FRI-protocol, we must create an extractor that takes as input folding challenges and then openings for various points $s \in S$ that are consistent with the folded polynomials (which are assumed to be low-degree). From two distinct folding challenges $c, c' \in \mathbb{F}$, if $G(X)$ and $G'(X)$ are the foldings around c and c' respectively of the function the prover committed to, then we can create the following polynomial:

$$F(X) = X \frac{G(X^2) - G'(X^2)}{c - c'} + \frac{cG'(X^2) - c'G(X^2)}{c - c'}.$$

Note that if G and G' have degree less than $\rho N/2$, then indeed F would have degree less than ρN .

The extractor may also rewind the second phase of the protocol to obtain sets A and A' covering at least $(1 - \delta)$ fraction of S . We can then conclude that we have consistent openings on their intersection $A \cap A'$ (assuming that we do not violate the binding property of the commitment, i.e., that we do not extract a witness for the relation $\mathfrak{R}_{\text{coll}}$). The intersection $A \cap A'$ covers a $(1 - 2\delta)$ fraction of S , so we have found a low-degree polynomial agreeing with the commitment on a $(1 - 2\delta)$ fraction of the points of S .

At this point, we could iterate this argument. However, iterating this argument over μ folding rounds would cause us to only prove that the prover committed to a function that agrees with a low-degree polynomial on a $(1 - 2^\mu \delta)$ -fraction of the coordinates (assuming that we did not extract a collision in the commitment). This is quite unsatisfactory, as we would like to have μ logarithmic in N and $\delta \in (0, 1)$ a constant. Fortunately, by relying on ideas from prior works (specifically, [5]) we can show that we can indeed extract a low-degree polynomial agreeing with the commitment on a $(1 - \delta)$ fraction of coordinates (or, of course, a violation to the binding property of the commitment).

In order to analyze the soundness of the FRI-protocol more effectively, we will need the following *coset-distance* from f to $\text{RS}[\mathbb{F}, S, \rho]$:

$$\Delta_S(f) := \min_{F \in \mathbb{F}[X], \deg(F) < \rho N} \frac{|\{s \in S : f(s) \neq F(s) \vee f(-s) \neq F(-s)\}|}{N}.$$

This distance notion has been used in prior works [5]. Observe that $\Delta_S(f) \geq \delta_S(f)$. Intuitively, this measure is useful because it allows for a more careful accounting of how the Hamming metric behaves under the folding operation than

the above naïve analysis. For this reason, our extractor will succeed assuming a bound on $\Delta_S(f)$ rather than just $\delta_S(f)$.

The following lemma quantifies this intuition, by characterizing the set of challenges c that could cause the Hamming metric to decrease when a function f is folded around c . These ideas are implicit in [5, Lemma 4.4]; we restate them in a language that is convenient for us. The full version of this work [1] includes a proof of the following lemma.

Lemma 7. *Let $f(X) \in \mathbb{F}[X]$ be such that $\Delta_S(f) < (1 - \rho)/2$. The number of choices for $c \in \mathbb{F}$ such that $\delta_{S^2}(\text{Fold}(f, c)) < \Delta_S(f)$ is at most N .*

In particular, if there exist pairwise distinct $c_0, \dots, c_N \in \mathbb{F}$ such that $\delta_{S^2}(\text{Fold}(f, c)) \leq \delta$ for all $i \in \{0, 1, \dots, N\}$, then $\Delta_S(f) \leq \delta$.

We now precisely define the notion of special-soundness that we will prove the FRI-protocol with one folding iteration satisfies. Informally, for the folding round the previous lemma tells us we need $N + 1$ challenges to extract, while for the second round we need enough local openings of the commitment to reveal a $(1 - \delta)$ -fraction of the values that the prover committed to. We now make this formal.

Let

$$\mathcal{C} := S^t = \{(s_1, s_2, \dots, s_t) : s_i \in S \ \forall i\} .$$

For a challenge $c = (s_1, \dots, s_t) \in \mathcal{C}$ we denote by

$$B(c) = \{s_1, -s_1, s_2, -s_2, \dots, s_t, -s_t\}$$

the set⁸ of elements of S that appear in the challenge tuple c , along with their negations. That is, it is the set of points that will be queried by the verifier if it samples (s_1, s_2, \dots, s_t) in the final verification step. Let $(\Gamma_{N+1}, \mathbb{F})$ be the monotone structure that contains all subsets of \mathbb{F} of cardinality at least $N + 1$, and let (Γ, \mathcal{C}) be the monotone structure that contains all subsets of \mathcal{C} that cover at least a $(1 - \delta)$ -fraction of S , i.e.,

$$A \in \Gamma \subset 2^{\mathcal{C}} \iff \left| \bigcup_{c \in A} B(c) \right| \geq (1 - \delta)N .$$

Theorem 3 (FRI-protocol (one folding iteration)). *Let $\rho = 2^{-r}$ for some $r \in \{0, 1, \dots, m\}$ and let $\delta \in (0, 1)$ be such that $\delta < \frac{1-\rho}{4}$. The FRI-Protocol is perfectly complete with respect to relation \mathfrak{R}_0 and (Γ_{N+1}, Γ) -out-of- $(\mathbb{F}, \mathcal{C})$ special-sound with respect to relation $\mathfrak{R}_\delta \cup \mathfrak{R}_{\text{coll}}$.*

Proof. Completeness: This is immediate from prior work (e.g., [5]). To make our proof self-contained, we note that this follows immediately from the following facts concerning a polynomial $F(X) \in \mathbb{F}[X]$:

- if F has degree $< \rho N$ then $\text{Fold}(F, c)$ has degree $< \rho N/2$ for any $c \in \mathbb{F}$; and

⁸ That is, we explicitly remove repetitions, so $B(c)$ is *not* interpreted as a multi-set.

– for any $s \in S$ and $c \in \mathbb{F}$, $\text{Fold}(F, c)(s^2) = \frac{F(s)+F(-s)}{2} + c \frac{F(s)-F(-s)}{2s}$.

Soundness: We must extract a witness for either the relation \mathfrak{R}_δ or the relation $\mathfrak{R}_{\text{coll}}$ given a (Γ_{N+1}, Γ) -tree of accepting transcripts. Such a tree of transcripts consists of the following:

- folding challenges $c_0, \dots, c_N \in \mathbb{F}$,
- polynomials $G_0, \dots, G_N \in \mathbb{F}[X]$ of degree less than $\rho \frac{N}{2}$,
- subsets $A_0, \dots, A_N \subseteq \mathcal{C}$, each satisfying $\left| \bigcup_{c \in A_j} B(c) \right| \geq (1 - \delta)N$, and
- for each $0 \leq j \leq N$, for each $s \in \bigcup_{c \in A_j} B(c)$, opening information γ_{sj} for the element s . Let $y_{sj} \in \mathbb{F}$ be the element for which $\text{LOC}_{\text{pp}}(P, s, y_{sj}, \gamma_{sj}) = 1$.

Let $B_j := \bigcup_{c \in A_j} B(c)$ for $0 \leq j \leq N$, and observe that these sets are closed under negation (i.e., $s \in B_j \iff -s \in B_j$).

Suppose there exists $j \neq j'$ such that, for some $s \in B_j \cap B_{j'}$, $y_{sj} \neq y_{sj'}$. Then, we may output the following witness for the relation $\mathfrak{R}_{\text{coll}}$: $(s, y_{sj}, y_{sj'}, \gamma_{sj}, \gamma'_{sj})$.

We may now assume that the above does not occur. In other words, for each $s \in \bar{B} := B_0 \cup \dots \cup B_N$ the set $\{y_{sj} : s \in B_j\}$ is in fact a singleton set; denote its unique element by y_s . We also let $\gamma_s := \gamma_{sj}$ where j is the smallest element in $\{0, 1, \dots, N\}$ such that $s \in B_j$ (this is just an arbitrary tie-breaking rule).

For each $j \in \{0, 1, \dots, N\}$, the polynomial G_j and the elements y_s for $s \in B_j$ satisfy the following relation:

$$G_j(s^2) = \frac{y_s + y_{-s}}{2} + c_j \frac{y_s - y_{-s}}{2s} .$$

Let $f(X) \in \mathbb{F}[X]$ be a polynomial consistent with the y_s 's, i.e., for all $s \in \bar{B}$ we have $f(s) = y_s$. Furthermore, for reasons to be clear later, we let f be different to the polynomial F_0 defined below outside of \bar{B} , i.e., $f(s) \neq F_0(s)$ for all $s \notin \bar{B}$. Then, for each $j \in \{0, 1, \dots, N\}$ and all s^2 such that $\{\pm s\} \subseteq B_j$, we have

$$G_j(s^2) = \text{Fold}(f, c_j)(s^2) .$$

We conclude that $\text{Fold}(f, c_j)$ and G_j agree on at least $(1 - \delta) \frac{N}{2}$ elements of S^2 . As $\deg(G_j) < \rho \frac{N}{2}$ it follows that

$$\delta_{S^2}(\text{Fold}(f, c_j)) \leq \delta .$$

By Lemma 7, if we establish that $\Delta_S(f) < \frac{1-\rho}{2}$, it in fact then follows that $\Delta_S(f) \leq \delta$, which in turn implies $\delta_S(f) \leq \delta$. As $2\delta < \frac{1-\rho}{2}$ by assumption, it suffices for us to show $\Delta_S(f) \leq 2\delta$. We focus on proving this now.

Consider the polynomial

$$F_0(X) := X \frac{G_0(X^2) - G_1(X^2)}{c_0 - c_1} + \frac{c_0 G_1(X^2) - c_1 G_0(X^2)}{c_0 - c_1} .$$

Since the degrees of G_0 and G_1 are smaller than $\rho \frac{N}{2}$, it follows that $\deg(F_0) < \rho N$. Furthermore, we note that for all $s \in B_0 \cap B_1$ we have $f(s) = F_0(s)$. Indeed,

$$\begin{aligned} F_0(s) &= s \cdot \frac{G_0(s^2) - G_1(s^2)}{c_0 - c_1} + \frac{c_0 G_1(s^2) - c_1 G_0(s^2)}{c_0 - c_1} \\ &= \frac{s}{c_0 - c_1} \left[\frac{f(s) + f(-s)}{2} + c_0 \frac{f(s) - f(-s)}{2s} \right. \\ &\quad \left. - \left(\frac{f(s) + f(-s)}{2} + c_1 \frac{f(s) - f(-s)}{2s} \right) \right] \\ &\quad + \frac{1}{c_0 - c_1} \left[c_0 \cdot \left(\frac{f(s) - f(-s)}{2} + c_1 \frac{f(s) - f(-s)}{2s} \right) \right. \\ &\quad \left. - c_1 \cdot \left(\frac{f(s) + f(-s)}{2} + c_0 \frac{f(s) - f(-s)}{2s} \right) \right] \\ &= \frac{s}{c_0 - c_1} \cdot (c_0 - c_1) \frac{f(s) - f(-s)}{2s} + \frac{1}{c_0 - c_1} \cdot (c_0 - c_1) \frac{f(s) + f(-s)}{2} \\ &= \frac{f(s) - f(-s)}{2} + \frac{f(s) + f(-s)}{2} = f(s). \end{aligned}$$

From this, we can conclude that f and F_0 agree on at least $(1 - 2\delta)N/2$ pairs $\{\pm s\}$: here, we use the fact that as B_0 and B_1 are closed under negation, so is $B_0 \cap B_1$. Thus, the number of $s \in S$ for which $f(s) \neq F_0(s)$ or $f(-s) \neq F_0(-s)$ is at most $2\delta N$. Recalling $\deg(F_0) < \rho N$, we conclude $\Delta_S(f) \leq 2\delta$, as desired.

Thus, we have found that $\Delta_S(f) \leq \delta$, which in particular means $\delta_S(f) \leq \delta$, as desired. Let $F(X)$ denote the (necessarily unique) polynomial of degree $< \rho N$ such that $d_S(F(S), f(S)) \leq \delta N$. As $d_S(F_0(S), f(S)) \leq 2\delta N$ it also follows that $d_S(F_0(S), F(S)) \leq 3\delta N < 1 - \rho$. As $F_0(S), F(S) \in \text{RS}[\mathbb{F}, S, \rho]$ and this code has minimum distance $1 - \rho$, it must be that $F_0(S) = F(S)$, which further implies $F_0(X) = F(X)$ (as polynomials).

We can therefore extract a polynomial of degree $< \rho N$ that agrees with the function $f(X)$ on a $(1 - \delta)$ fraction of coordinates: namely, the polynomial $F_0(X)$. Furthermore, since f differs from F_0 outside of $\bar{B} = B_0 \cup \dots \cup B_N$ (by the choice of f), we can find a subset $B \subseteq \bar{B}$ of size at least $(1 - \delta)N$ for which $f(s) = F_0(s)$ for all $s \in B$. We may therefore output the following witness for $\mathfrak{R}_\delta: (F_0(X), B, (\gamma_s)_{s \in B})$. \square

We are now in position to apply the machinery developed in Sect. 6 to conclude the following bound on the knowledge error.

Corollary 1 (Knowledge Error of FRI-protocol (one folding iteration)). *Let $\rho = 2^{-r}$ for some $r \in \{0, 1, \dots, m\}$ and let $\delta \in (0, 1)$ be such that $\delta < \frac{1-\rho}{4}$. The FRI-Protocol is knowledge sound with respect to relation $\mathfrak{R}_\delta \cup \mathfrak{R}_{\text{coll}}$ with knowledge error*

$$\kappa := 1 - \left(1 - \frac{N}{|\mathbb{F}|} \right) \left(1 - \frac{(\lceil (1 - \delta)N \rceil - 1)^t}{N^t} \right) \leq \frac{N}{|\mathbb{F}|} + (1 - \delta)^t.$$

Proof. Theorem 3 shows that the FRI-Protocol is (Γ_{N+1}, Γ) -out-of- $(\mathbb{F}, \mathcal{C})$ special-sound. To apply Theorem 2, we must first establish that $t_\Gamma \cdot t_{\Gamma_{N+1}} \leq N^{O(1)}$. And this is indeed the case, as

$$t_\Gamma \leq \lceil (1 - \delta)N \rceil \text{ and } t_{\Gamma_{N+1}} \leq N + 1 .$$

We now establish the knowledge error. For this, it suffices to note that $\max_{S \notin \Gamma_{N+1}} \frac{|S|}{|\mathbb{F}|} = \frac{N}{|\mathbb{F}|}$ while

$$\max_{A \notin \Gamma} \frac{|A|}{|\mathcal{C}|} = \frac{(\lceil (1 - \delta)N \rceil - 1)^t}{N^t} \leq (1 - \delta)^t .$$

To see the first equality, first note that if $A \notin \Gamma$ then $\bigcup_{c \in A} B(c)$ has cardinality less than $(1 - \delta)N$, so A the number of $s \in S$ which can appear in a challenge $c \in A$ is at most $(1 - \delta)N$; as this is an integer, it is at most $\lceil (1 - \delta)N \rceil - 1$. That is, for some subset $T \subseteq S$ with $|T| \leq \lceil (1 - \delta)N \rceil - 1$, $A \subseteq T^t$, and $|T^t| \leq ((1 - \delta)N/2)^t$. The equality holds as we can certainly choose $A = T^t$ for some $T \subseteq S$ of size $\lceil (1 - \delta)N \rceil - 1$. For the denominator, as $\mathcal{C} = S^t$ it has cardinality $|S|^t = N^t$. \square

7.4 Additional Folding Iterations

The above analysis can naturally be extended to handle more folding iterations. Let $F_0 := F$ be the low degree polynomial the prover commits to in the first round. We have folding rounds $i = 1, \dots, \mu$, and in round i the verifier sends a challenge $c_{i-1} \in \mathbb{F}$ and the prover provides a commitment to $F_i(S^{2^i})$ where $F_i(X) = \text{Fold}(F_{i-1}, c_{i-1})(X)$. After these folding iterations, the verifier picks t points $s_1, \dots, s_t \in S$ independently and uniformly at random and then checks that for all $i = 1, \dots, \mu$ and $j = 1, \dots, t$, we have

$$F_i(s_j^{2^i}) = \frac{F_{i-1}(s_j^{2^{i-1}}) + F_{i-1}(-s_j^{2^{i-1}})}{2} + c_{i-1} \frac{F_{i-1}(s_j^{2^{i-1}}) - F_{i-1}(-s_j^{2^{i-1}})}{2s_j} .$$

The recursive structure of the extractor implies that after μ folding iterations we obtain a protocol with the following generalized special-soundness guarantee.

Theorem 4 (FRI-protocol (μ folding iterations)). *Let $\rho = 2^{-r}$ for some $r \in \{0, 1, \dots, m\}$ and let $\delta \in (0, 1)$ be such that $\delta < \frac{1-\rho}{4}$. Let $\mu \in \mathbb{N}$ be such that $\mu \leq \log_2 N$, and for $i = 1, 2, \dots, \mu$ let $N_i := N/2^{i-1}$. The FRI-protocol with μ folding iterations is perfectly complete with respect to relation \mathfrak{R}_0 and $(\Gamma_{N_1+1}, \Gamma_{N_2+1}, \dots, \Gamma_{N_\mu+1}, \Gamma)$ -out-of- $(\mathbb{F}, \mathbb{F}, \dots, \mathbb{F}, \mathcal{C})$ special-sound with respect to relation $\mathfrak{R}_\delta \cup \mathfrak{R}_{\text{coll}}$.*

This yields the following corollary regarding the knowledge error. However, we note that for $\mu = \Omega(\log N)$ the knowledge extractor only runs in expected *quasi-polynomial* time, preventing us from being able to claim the standard notion of knowledge soundness. Nonetheless we believe that the guarantee is meaningful. For the proof, we refer to the full version [1].

Corollary 2 (Knowledge Error of FRI-protocol (μ folding iterations)).

Let $N = 2^n$ for some $n \in \mathbb{N}$, let $\rho = 2^{-r}$ for some $r \in \{0, 1, \dots, m\}$ and let $\delta \in (0, 1)$ be such that $\delta < \frac{1-\rho}{4}$. Let $\mu \in \mathbb{N}$ be such that $\mu \leq \log_2 N$, and for $i = 1, 2, \dots, \mu$ let $N_i := N/2^{i-1}$. There exists a function $q(N, \mu) = N^{O(\mu)}$ such that the following holds.

There exists an extraction algorithm that, when given oracle access to a (potentially dishonest prover) \mathcal{P}^* and input x of size N for the FRI-protocol, runs in time $\leq q(N, \mu)$ and outputs a witness in the relation $\mathfrak{R}_\delta \cup \mathfrak{R}_{\text{coll}}$ with probability at least

$$\frac{\epsilon(\mathcal{P}^*, x) - \kappa(N, \mu)}{q(N, \mu)}$$

where

$$\begin{aligned} \kappa(N, \mu) &:= 1 - \left(\prod_{i=1}^{\mu} \left(1 - \frac{N_i}{|\mathbb{F}|} \right) \right) \cdot \left(1 - \frac{(\lceil (1-\delta)N \rceil - 1)^t}{N^t} \right) \\ &\leq \sum_{i=1}^{\mu} \frac{N_i}{|\mathbb{F}|} + (1-\delta)^t \leq \frac{2N}{|\mathbb{F}|} + (1-\delta)^t . \end{aligned}$$

Acknowledgments. The first author has been supported by TNO's Early Research Program - Next Generation Cryptography. We would additionally like to thank Michael Kloof for pointing out to us that the extractor for the FRI-protocol only runs in quasi-polynomial time.

References

1. Full version of this paper. IACR ePrint 2023/818 (2023)
2. Attema, T., Cramer, R., Kohl, L.: A compressed Σ -protocol theory for lattices. In: CRYPTO, vol. 12826, pp. 549–579 (2021)
3. Attema, T., Fehr, S.: Parallel repetition of (k_1, \dots, k_μ) -special-sound multi-round interactive proofs. In: CRYPTO, vol. 13507, pp. 415–443 (2022)
4. Attema, T., Fehr, S., Kloof, M.: Fiat-shamir transformation of multi-round interactive proofs. In: TCC, vol. 13747, pp. 113–142 (2022)
5. Ben-Sasson, E., Bentov, I., Horesh, Y., Riabzev, M.: Fast reed-solomon interactive oracle proofs of proximity. In: ICALP, vol. 107, pp. 14:1–14:17 (2018)
6. Ben-Sasson, E., Chiesa, A., Spooner, N.: Interactive oracle proofs. In: TCC, vol. 9986, pp. 31–60 (2016)
7. Bootle, J., Cerulli, A., Chaidos, P., Groth, J., Petit, C.: Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In: EUROCRYPT, vol. 9666, pp. 327–357 (2016)
8. Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., Maxwell, G.: Bulletproofs: short proofs for confidential transactions and more. In: IEEE (S&P), pp. 315–334 (2018)
9. Don, J., Fehr, S., Majenz, C., Schaffner, C.: Efficient NIZKs and signatures from commit-and-open protocols in the QROM. In: CRYPTO, vol. 13508, pp. 729–757 (2022)

10. Don, J., Fehr, S., Majenz, C., Schaffner, C.: Online-extractability in the quantum random-oracle model. In: Dunkelman, O., Dziembowski, S. (eds.) EUROCRYPT, vol. 13277, pp. 677–706 (2022)
11. Kilian, J.: A note on efficient zero-knowledge proofs and arguments (extended abstract). In: Kosaraju, S.R., Fellows, M., Wigderson, A., Ellis, J.A. (eds.) STOC, pp. 723–732 (1992)
12. Micali, S.: CS proofs (extended abstracts). In: FOCS, pp. 436–453 (1994)
13. Micali, S.: Computationally sound proofs. *SIAM J. Comput.* **30**(4), 1253–1298 (2000)
14. Wikström, D.: Special soundness revisited. *IACR ePrint 2018/1157* (2018)
15. Wikström, D.: Special soundness in the random oracle model. *IACR ePrint 2021/1265* (2021)