

The Cost of Rewinding for Succinct Arguments: Strict-Time Barriers and Expected-Time Security

Alessandro Chiesa

alessandro.chiesa@epfl.ch
EPFL

Ziyi Guan

ziyiguan@mit.edu
MIT

William Wang

ww@priv.pub
NYU

Yuetian Wu

yuetianw@stanford.edu
Stanford University

September 25, 2026

Abstract

Succinct arguments are central cryptographic objects, underlying various applications such as blockchains and image authentication. Almost all succinct arguments used in practice follow the commit-and-open paradigm introduced by Kilian (STOC 1992): commit to a probabilistic proof, then open the few locations the verifier queries. Decades of work have generalized and optimized this paradigm, yet its exact provable security in the standard model remains unclear. Existing analyses proceed by rewinding, and they leave gaps: strict-time analyses lose an inverse-polynomial factor, while expected-time analyses are sharper but apply only to special cases of this paradigm.

We resolve both gaps. On the negative side, we show that the strict-time loss is inherent: no strict polynomial-time reduction relying on a black-box extractor can achieve negligible knowledge error (unless the underlying language is easy). On the positive side, we give a new expected-time analysis, achieving negligible soundness and knowledge errors via an expected polynomial-time reduction against expected polynomial-time adversaries, at the full generality of the paradigm, capturing multi-round and functional variants such as polynomial and linear IOPs compiled with the corresponding commitments. This is the first standard-model analysis to reach this regime for succinct arguments deployed in practice.

Keywords: black-box separation; expected-time analysis; succinct arguments

Contents

1	Introduction	3
1.1	Our results	4
1.2	The cost of rewinding	5
1.3	Related work	6
2	Techniques	8
2.1	Strict-time barriers for Kilian’s protocol	8
2.2	Expected-time security of Kilian’s protocol	9
2.3	Expected-time security of the IBCS protocol	15
2.4	Extension to functional queries	19
3	Preliminaries	21
3.1	Interactive arguments	21
3.2	Query classes	22
3.3	Functional commitments	24
3.4	Functional interactive oracle proofs	25
3.5	Funky protocol	26
4	Barriers for knowledge soundness of Kilian’s protocol	28
4.1	Vector commitments	28
4.2	Probabilistically checkable proofs	28
4.3	Non-adaptive interactive arguments	29
4.4	Kilian’s protocol	30
4.5	Toolbox for lower-bound analysis	30
4.6	Barrier from zero-knowledge PCPs	31
4.7	Barrier from low-query PCPs	34
4.8	Barrier for expected-time rewinding from low-query PCPs	36
5	Expected-time soundness of the Funky protocol	38
5.1	Adversary construction	38
5.2	Security analysis	39
6	Expected-time knowledge soundness of the Funky protocol	42
6.1	Extractor construction	43
6.2	Security analysis	44
A	Coverage factor bounds	50
	Acknowledgments	53
	AI disclosure	53
	References	53

1 Introduction

A succinct argument for a relation R is a computationally sound interactive proof in which, for a given instance x , the argument prover convinces the argument verifier that there exists a witness w with $(x, w) \in R$, while communicating far fewer than $|w|$ bits. Succinct arguments are central cryptographic primitives with numerous applications. We study succinct arguments for NP in the standard model from falsifiable assumptions.

The commit-and-open paradigm. Kilian’s protocol [Kil92], the first succinct argument, is the prototype for almost all practical constructions: the argument prover commits to a probabilistically checkable proof (PCP) via a vector commitment scheme (e.g., a Merkle tree) and sends the commitment to the verifier; the verifier replies with PCP randomness; and the prover opens the queried locations of the PCP. The verifier accepts if the openings are valid and the PCP verifier accepts. The same *commit-and-open* template has since been applied to interactive oracle proofs (IOPs), the multi-round analogue of PCPs, and to functional IOPs such as polynomial IOPs [BCS16; LM19; CHMMVW20; BFS20; CFFQR21; LPS24; CDGS23; LPS25; CGKY26]. In the standard model, security of every such protocol is proved by *rewinding*: the reduction replays the adversary against fresh verifier challenges, collecting a few opened locations per run, until it can reconstruct a full proof string and invoke the security of the underlying probabilistic proof. This paper studies how many rewinds are needed.

The strict-time setting. In the standard *strict-time setting*, all algorithms (the malicious prover, the extractor, the reductions between them) run in worst-case bounded time. The state of the art for Kilian’s protocol is as follows [CDGSY24]. Let PCP be a probabilistically checkable proof with proof length ℓ , soundness error ϵ_{PCP} , and knowledge soundness error κ_{PCP} . Let VC be a vector commitment scheme with position binding error ϵ_{VC} . Then, for every $\epsilon > 0$, the interactive argument $\text{ARG} := \text{Kilian}[\text{PCP}, \text{VC}]$ satisfies the following error bounds.

- *Soundness.*

$$\epsilon_{\text{ARG}}(t_{\text{ARG}}) \leq \epsilon_{\text{PCP}} + \epsilon_{\text{VC}}(t_{\text{VC}}) + \epsilon, \text{ where } t_{\text{VC}} = O\left(\frac{\ell}{\epsilon} \cdot t_{\text{ARG}}\right).$$

- *Knowledge soundness.*

$$\kappa_{\text{ARG}}(t_{\text{ARG}}) \leq \kappa_{\text{PCP}} + \epsilon_{\text{VC}}(t_{\text{VC}}) + \epsilon, \text{ where } t_{\text{VC}} = O\left(\frac{\ell}{\epsilon} \cdot t_{\text{ARG}}\right) \text{ with extraction time } O\left(\frac{\ell}{\epsilon} \cdot t_{\text{ARG}}\right).$$

The additive error ϵ affects the time t_{VC} of the position-binding adversary (and extractor) multiplicatively as ℓ/ϵ ; driving this error to negligible (for all polynomial-size adversaries) forces a *super-polynomial* rewinding budget.

Can this loss be avoided? There is partial evidence that suggests not: instantiating Kilian’s protocol with a tailored PCP and a vector commitment based on Schnorr’s identification scheme, [CDGSY24] showed that a sufficiently strong improvement to the argument’s soundness bound would carry over to the standard-model security bound of Schnorr, which has resisted improvement for over two decades. This, however, is only a reduction to the *unproven* optimality of Schnorr’s analysis.¹ We therefore ask:

Is the loss of ℓ/ϵ provably necessary for Kilian’s protocol?

The expected-time setting. The connection to Schnorr’s identification scheme suggests a potential way to sidestep the discussion above. Schnorr’s canonical extractor rewinds only after a first accepting run, and then until it finds a second; its running time is therefore a random variable, and its guarantee is naturally stated in *expected* polynomial time. In the *expected-time setting*, all algorithms (adversary, extractor, and the reductions between them) are bounded in expected rather than worst-case running time, and the underlying cryptography is assumed secure against expected-time adversaries [JT20]. Analyzed this way, the Schnorr identification scheme loses essentially nothing relative to (the expected-time variant of the) discrete logarithm problem.

¹The recent lower bounds for Schnorr against restricted reductions [HM26] do not apply; see Section 1.3.

There has been recent progress in analyzing the expected-time security of Kilian’s protocol. [CDGSY24] gives an expected-time soundness analysis that replaces the ℓ/ϵ time overhead with $\ell \cdot \log(q/\epsilon)$, where q is the query complexity of PCP; but the analysis does not extend to knowledge soundness. [AKLY24] obtains an essentially tight expected-time knowledge soundness bound via *adaptive special soundness*, but only for PCPs (that moreover satisfy a certain knowledge property); the analysis does not extend to the multi-round IOPs and functional IOPs on which deployed systems are built. This motivates our second question:

Can expected-time analysis achieve tight security at the full generality of the commit-and-open paradigm?

1.1 Our results

We answer both questions positively. First, the strict-time loss from rewinding is provably necessary: every natural strict polynomial-time extractor for Kilian’s protocol, for *every* vector commitment scheme, incurs at least inverse polynomial knowledge soundness error. Second, we obtain tight expected-time security for the commit-and-open paradigm in its full generality (multiple rounds and functional queries).

1.1.1 Strict-time barriers for Kilian’s protocol

We consider knowledge extractors that are *fully black-box* in the following sense: (i) there is a single “universal” extractor for all provers, running in *strict polynomial-time with an a priori rewind budget p* (fixed independently of the prover); and (ii) *VC-oblivious*, i.e., the extractor accesses the prover only through an interface that returns, for any desired verifier challenge, the opened query-answer pairs and the bit denoting whether the opening proof is valid. For brevity we call such extractors *admissible*; every rewinding extractor we are aware of is admissible.

Theorem 1 (Informal). *Let PCP be any PCP for an NP relation R with $L(R) \notin \text{BPP}$, let VC be any vector commitment scheme, and set $\text{ARG} := \text{Kilian}[\text{PCP}, \text{VC}]$. Let κ_{ARG} be the non-adaptive knowledge soundness error of ARG against $\text{poly}(n)$ -time provers, with respect to any admissible extractor that rewinds the prover at most $p = \text{poly}(n)$ times.*

- *If PCP has query complexity $q = O(\log n)$, then $\kappa_{\text{ARG}} > \frac{1}{8p^2}$.*
- *If PCP is zero-knowledge, then $\kappa_{\text{ARG}} > \frac{1}{p}$.*

The first bound holds for any PCP with small query complexity; the second is stronger but requires zero knowledge. Either way, no fixed polynomial budget achieves negligible knowledge soundness error: the strict-time extractor in [CDGSY24] loses roughly ℓ/p under a budget of p , and the $1/p$ dependence is necessary.

Remark 1. *Some restriction on how the extractor accesses the VC scheme in Theorem 1 (e.g., VC-obliviousness) is necessary. Indeed, if the VC scheme is extractable (e.g., Merkle commitments in the random oracle model [CY24] or various schemes proved under knowledge assumptions) then an extractor can obtain the committed PCP string without any rewinding (and in strict polynomial time). A bound quantifying over all VC schemes must therefore exclude extractors that exploit a scheme’s internals. Our barrier applies to every analysis that treats the VC scheme as a black-box position-binding primitive, which is every known standard-model analysis.*

1.1.2 Expected-time security of the commit-and-open paradigm

We provide a tight expected-time security analysis of the commit-and-open paradigm in its full generality. For simplicity, we state the theorem below for the IBCS protocol, which compiles a public-coin IOP and a VC scheme into an interactive argument; we discuss the “functional” extension, known as the “Funky” protocol, later (and in more detail in Section 2.4). We denote the IOP’s round complexity as k , (total) proof length as ℓ , soundness error as ϵ_{IOP} , and knowledge soundness error as κ_{IOP} (with associated extraction time t_{E}); moreover $\epsilon_{\text{VC}}^*(t_{\text{VC}}^*)$ denotes the position binding error of the VC scheme against adversaries running in *expected-time* t_{VC}^* .

Theorem 2 (Informal). $\text{ARG} := \text{IBCS}[\text{IOP}, \text{VC}]$ has expected-time soundness error

$$\epsilon_{\text{ARG}}^*(t_{\text{ARG}}^*) \leq \epsilon_{\text{IOP}} + \epsilon_{\text{VC}}^*(t_{\text{VC}}^*) \quad , \text{ where } t_{\text{VC}}^* = O(\ell \cdot t_{\text{ARG}}^*) \quad ;$$

and expected-time knowledge soundness error

$$\kappa_{\text{ARG}}^*(t_{\text{ARG}}^*) \leq (1 + \alpha) \cdot \kappa_{\text{IOP}} + \epsilon_{\text{VC}}^*(t_{\text{VC}}^*) \quad , \text{ where } t_{\text{VC}}^* = O\left(\frac{k \cdot \ell}{\log(1 + \alpha)} \cdot t_{\text{ARG}}^*\right) \text{ and extractor running time } t_{\text{VC}}^* + t_{\text{E}} \quad .$$

On tightness of knowledge soundness. For a PCP ($k = 1$), the extractor of Theorem 2 rewinds $O(\ell / \log(1 + \alpha))$ times in expectation, compared to $O(\ell / \epsilon)$ rewinds of the strict-time extractor of [CDGSY24] at additive error ϵ . The dependence on ℓ is essentially unavoidable: a modification of Theorem 1 shows that, unless the randomized exponential time hypothesis (rETH) fails, every expected-time VC-oblivious extractor for Kilian’s protocol must rewind $\Omega(\ell / \text{polylog}(\ell))$ times in expectation when the PCP has quasilinear length and small query complexity. Hence the extractor in Theorem 2 is optimal up to a polylogarithmic factor; eliminating this factor remains open, as does determining whether the cost is forced by VC-obliviousness.

We stress that this tightness holds for knowledge soundness *only*. Every known soundness analysis of the commit-and-open paradigm, including ours, proceeds by extracting a committed proof string and then invoking soundness of the underlying probabilistic proof; the reduction therefore pays the same $\Theta(\ell)$ rewinding cost as a knowledge extractor. Our lower bounds show that this cost is necessary for extraction, but not that extraction is necessary for soundness. Whether soundness can be established without extracting a full proof string, at a rewinding cost independent of ℓ , is an open problem that we leave for future work.

On functional IOPs. Many succinct arguments of interest are based on polynomial IOPs or other forms of *functional IOPs*; these yield distinct efficiency tradeoffs (e.g., between prover time and argument size). In a polynomial IOP, for example, the verifier treats the committed message as a polynomial and queries consist of desired evaluation points. The *Funky protocol* [CGKY26] captures the commit-and-open paradigm at this level of generality: a functional IOP and a functional commitment scheme both for the same family $\mathbf{Q}$ of query functions are combined into a succinct interactive argument. The *strict-time* security of the Funky protocol inherits the same rewinding overhead as Kilian’s protocol [CGKY26]. We establish tight expected-time security of the Funky protocol, generalizing Theorem 2; the statement and the main ideas are in Section 2.4.

Remark 2 (Adaptive security). *Our expected-time analyses are proved with adaptive security in the common reference string model, where a malicious argument prover chooses the instance after seeing the public parameters (as in [CDGSY24]); plain-model variants, where the verifier starts by sending public parameters, follow directly. Our barriers in Theorem 1 are for non-adaptive knowledge soundness, making the statements only stronger.*

1.2 The cost of rewinding

We discuss how the improvement on the number of rewindings in the expected-time regime affects concrete security. We instantiate the vector commitment scheme as a Merkle commitment scheme constructed from an ideal hash function with output length λ , and compare several *soundness* bounds of Kilian’s protocol. Throughout, b denotes the target number of bits of security in the *cryptographic* contribution to the argument error; the information-theoretic term of the probabilistic proof (the PCP soundness error) is budgeted separately.

For an ideal hash function, a t -time adversary finds a collision with probability at most $t^2 / 2^\lambda$ (a birthday bound), whereas an adversary with *expected* running time t^* succeeds with probability at most $t^* / 2^{\lambda/2}$ [JT20]. This yields estimates for the position binding errors in the strict-time and expected-time settings:

$$\epsilon_{\text{VC}}(\lambda, t_{\text{VC}}) \leq \frac{t_{\text{VC}}^2}{2^\lambda} \text{ and } \epsilon_{\text{VC}}^*(\lambda, t_{\text{VC}}^*) \leq \frac{t_{\text{VC}}^*}{2^{\lambda/2}} \quad .$$

We use these estimates to determine the minimal hash output length λ that yields b bits of cryptographic security in three cases for Kilian’s protocol: the strict-time setting with [CDGSY24]’s bound, the expected-time setting with our bound in Theorem 2, and the straightline extraction setting in the random oracle model (ROM) [CY24].

- *Strict-time.* The position-binding adversary runs in time $t_{\text{VC}} = O(\frac{\ell}{\epsilon} \cdot t_{\text{ARG}})$, so the cryptographic contribution to the error is, for every tolerance $\epsilon \in (0, 1]$, $E(\epsilon) := \frac{(4 \cdot \ell \cdot t_{\text{ARG}})^2}{\epsilon^2 \cdot 2^\lambda} + \epsilon$, where the multiplicative factor 4 is the constant from the reduction in [CDGSY24]. Setting $\bar{\epsilon} := (32 \cdot (\ell \cdot t_{\text{ARG}})^2 / 2^\lambda)^{1/3}$ minimizes E , with value $E(\bar{\epsilon}) = \frac{3}{2}\bar{\epsilon}$. Requiring $E(\bar{\epsilon}) \leq 2^{-b}$ and solving for λ gives $\lambda \geq 3b + 2 \log_2(\ell \cdot t_{\text{ARG}}) + \log_2 108$.
- *Expected-time.* The position-binding adversary underlying Theorem 2 has expected running time $t_{\text{VC}}^* \leq (\ell + 1) \cdot t_{\text{ARG}}^*$, so the error is at most $(\ell + 1) \cdot t_{\text{ARG}}^* / 2^{\lambda/2}$. This is at most 2^{-b} when $\lambda \geq 2b + 2 \log_2((\ell + 1) \cdot t_{\text{ARG}}^*)$.
- *Straightline extraction.* [CY24] shows that $\epsilon_{\text{ARG}}(t_{\text{ARG}}) \leq \epsilon_{\text{PCP}} + t_{\text{ARG}}^2 / 2^\lambda$, so the cryptographic term is at most 2^{-b} when $\lambda \geq b + 2 \log_2 t_{\text{ARG}}$.

Each requirement is affine in the target b with slope 3, 2, or 1: raising the security by one bit costs 3 extra bits of hash output under the strict-time analysis, 2 under the expected-time analysis, and 1 in the ROM. To see the improvement more concretely, we consider a target argument error of 2^{-40} against adversaries of size 2^{60} . Suppose that the underlying PCP has soundness error 2^{-42} and proof length 2^{30} . The table below collects the three requirements for the output length λ together with their values at this benchmark.

Analysis	required λ	at the benchmark
Strict-time [CDGSY24]	$3b + 2 \log_2(\ell \cdot t_{\text{ARG}}) + \log_2 108$	308
Expected-time (Theorem 2)	$2b + 2 \log_2((\ell + 1) \cdot t_{\text{ARG}}^*)$	261
Straightline in the ROM [CY24]	$b + 2 \log_2 t_{\text{ARG}}$	161

1.3 Related work

Strict vs. expected polynomial time. [BL04] proves that nontrivial constant-round zero-knowledge arguments admit neither strict-polynomial-time black-box simulators nor strict-polynomial-time black-box knowledge extractors. Theorem 1 is inspired by that work, but does not follow from it because Kilian’s protocol is not zero knowledge against malicious verifiers. Other studies of expected-time cryptography [Klo21; JT20; Wik24] do not study the commit-and-open paradigm.

Special soundness and IOPs. Special soundness of the commit-and-open compiler would immediately yield expected-time knowledge soundness, but Kilian’s protocol is not k -special-sound for any polynomial k [CMSZ21]. *Generalized* special-soundness frameworks weaken this requirement and have been pushed toward IOPs: FRI [BBHR18] satisfies one [AFR23], with almost-optimal knowledge error but expected *quasi*-polynomial extraction time. None of them are known to yield an expected-polynomial-time knowledge extractor for the general commit-and-open paradigm.

Strict-time analyses of commit-and-open arguments. Kilian’s original work [Kil92] gives no formal security analysis; later analyses cover increasingly general probabilistic proofs and commitment schemes: non-adaptive reverse-samplable PCPs [BG08], linear PCPs and linear commitments [LM19], the IOP-based compiler [CDGS23], and functional IOPs with functional commitments [CGKY26]; the same regime has also been analyzed against quantum adversaries, for PCP-based Kilian [CMSZ21] and for multi-round IOP-based arguments [CDDGS25]. All of them enforce a worst-case running-time bound and hence pay an additive loss in security. [CDGSY24] gives partial evidence that this loss is inherent, by reduction to the unproven optimality of Schnorr’s standard-model analysis; the recent Schnorr lower bounds [HM26] cover only restricted, uniform-challenge reductions and do not transfer. Our lower bounds instead make the loss precise for a natural class of black-box extractors that includes all extractors currently known in this line of work.

Expected-time analysis of Kilian’s protocol. [CDGSY24] gives the first expected-time soundness analysis of Kilian’s protocol, replacing the strict-time ℓ/ϵ rewinding overhead by $\ell \cdot \log(q/\epsilon)$; it still pays the additive error ϵ and gives no knowledge soundness. [AKLY24] introduces adaptive special soundness (which is in turn inspired by almost special soundness [BF23]) and obtains the first lossless expected-time analysis of Kilian’s

protocol: it preserves the PCP knowledge error exactly, where our bound incurs a factor $1 + \alpha$, and it matches our expected-time lower bound (Section 1.1.2). The two bounds are nevertheless *incomparable*, because theirs requires the PCP to be strongly extractable: a *deterministic* extractor $\mathbf{E}$ such that, for every instance $\mathbb{x}$ and proof string Π ,

$$\Pr_{\rho}[\mathbf{V}^{\Pi}(\mathbb{x}; \rho) = 1] > \kappa_{\text{PCP}} \implies (\mathbb{x}, \mathbf{E}(\mathbb{x}, \Pi)) \in R .$$

We assume only the standard notion (Definition 4.5), in which the extractor may depend on the verifier randomness and only the joint failure probability is bounded. This is what enables our multi-round analysis: for IOPs the round- i proof $\Pi_i = \Pi_i(\rho_1, \dots, \rho_{i-1})$ depends on earlier verifier randomness, so there is no verifier-independent tuple $(\Pi_1, \dots, \Pi_k)$ on which deterministic per-string extraction can even be defined, and supplying the whole prover strategy would force exponential-time extraction. (Randomized extraction already arises at the PCP level under randomized code decoding [IJK09].) Generic composition of their extractor across rounds also blows up its running time superpolynomially. In sum, their exact knowledge term rests on a stronger, deterministic-extraction assumption, whereas our theorem uses the standard randomized notion, composes in expected polynomial time across rounds, and extends to functional IOPs.

Orthogonal models. Complementary lines of work escape the standard-model rewinding barrier by other means: straightline extraction from scheme-specific extractability in idealized models [CY24], and extractability assumptions for polynomial commitments in the standard model [CHMMVW20; BFS20]. Our lower bounds suggest this may be the only way out, since a black-box extractor must pay the rewinding cost (Remark 1). Finally, [CGKY26] treats strict-time Fiat–Shamir security of the Funky protocol, leaving the expected-time case open. Table 1 compares the standard-model rewinding analyses discussed above.

Analysis	Model	Security	Scope	Near-lossless?
[BG08]	strict-time	S, KS	PCP [†]	✗
[LM19]	strict-time	S, KS	PCP/LPCP	✗
[CDGS23]	strict-time	S, KS	IOP	✗
[CDGSY24]	strict-time	S, KS	PCP	✗
[CDGSY24]	expected-time	S	PCP	✗
[AKLY24]	expected-time	S, KS	PCP [‡]	✓
[CGKY26]	strict-time	S, KS*	FIOP	✗
This work	expected-time	S, KS	FIOP	✓

[†] Non-adaptive, reverse-samplable PCPs.

[‡] Strongly κ_{PCP} -extractable PCPs.

* State-restoration soundness and state-restoration knowledge soundness.

Table 1: Standard-model rewinding analyses of commit-and-open succinct arguments. S and KS denote soundness and knowledge soundness, and a bound is *near-lossless* if it has no additive truncation error.

2 Techniques

We overview the main ideas underlying our results. Section 2.1 explains our barriers for strict-time knowledge extraction in Kilian’s protocol. After that we discuss our expected-time analyses in increasing generality: first for Kilian’s protocol in Section 2.2 (as a warmup); then for the IBCS protocol in Section 2.3 (which compiles IOPs instead of PCPs); and finally for the Funky protocol in Section 2.4 (which supports functional queries).

2.1 Strict-time barriers for Kilian’s protocol

We explain Theorem 1’s lower bound for zero-knowledge PCPs. Additionally, we sketch the lower bound for low-query PCPs. For simplicity, we assume perfect completeness of the underlying probabilistic proofs and commitment schemes; the formal statements and proofs in Section 4 allow for imperfect completeness.

Definition 1. *An interactive argument $\text{ARG} = (\mathcal{P}, \mathcal{V})$ for a relation R has **non-adaptive knowledge soundness error** κ_{ARG} **with extraction time** $t_{\mathcal{E}}$ if there exists a $t_{\mathcal{E}}$ -time probabilistic algorithm $\mathcal{E}$ such that, for every instance x and time-bounded prover $\tilde{\mathcal{P}}$,*

$$\Pr\left[(x, w) \in R \mid w \leftarrow \mathcal{E}^{\tilde{\mathcal{P}}}(x)\right] \geq \Pr\left[b = 1 \mid b \leftarrow \langle \tilde{\mathcal{P}}, \mathcal{V}(x) \rangle\right] - \kappa_{\text{ARG}} .$$

Let R be the NP relation stated in Theorem 1. Suppose, toward a contradiction, that there is an admissible extractor (see Section 1.1.1) that makes at most p rewinds and has small knowledge soundness error. We show how to use this extractor to construct a BPP decider for the language $L(R)$, contradicting the assumption that $L(R) \notin \text{BPP}$.

The main idea is to run the extractor against a prover that behaves like the honest argument prover except that, with probability $1 - \theta$, it refuses to answer the verifier’s challenge. The extractor sees only a few honest PCP views during its execution, which can be simulated by the zero-knowledge simulator. In other words, we can reproduce the extractor’s execution without knowing the witness. For properly chosen θ , an extractor with knowledge error at most $1/p$ outputs a valid witness with inverse-polynomial probability. Repeating the simulation and checking its outputs gives a BPP decider for $L(R)$, which is a contradiction.

Fix $(x, w) \in R$ and $\theta \in [0, 1]$. Consider the following malicious argument prover $\tilde{\mathcal{P}}_{\theta, x, w}^*$.

- $\tilde{\mathcal{P}}_{\theta, x, w}^*$:
 1. Run the honest argument prover to commit to the honest PCP string: $(\text{cm}, \text{aux}) \leftarrow \mathcal{P}(x, w)$.
 2. Send cm and keep aux as private prover state.
- $\tilde{\mathcal{P}}_{\theta, x, w}^*(\text{aux}, \rho)$:
 1. Flip a fresh coin with heads probability θ .
 2. If heads: output the honest response $\mathcal{P}(\text{aux}, \rho)$.
 3. If tails: output $\perp$.

By completeness of the PCP and the vector commitment, $\tilde{\mathcal{P}}_{\theta, x, w}^*$ convinces the argument verifier with probability exactly θ . Hence, if an admissible extractor has knowledge soundness error κ_{ARG} , then

$$\Pr\left[(x, w') \in R \mid w' \leftarrow \mathcal{E}^{\tilde{\mathcal{P}}_{\theta, x, w}^*}(x)\right] \geq \theta - \kappa_{\text{ARG}} .$$

The extractor $\mathcal{E}^{\tilde{\mathcal{P}}_{\theta, x, w}^*}(x)$ sees at most $\theta \cdot p$ honest PCP views in expectation. When θ is chosen small enough, these views can be simulated by the zero-knowledge simulator without knowing the witness.

In more detail, consider the following malicious PCP verifier $\tilde{\mathbf{V}}$, which uses the extractor $\mathcal{E}$.

$\tilde{\mathbf{V}}^{\Pi}(x)$:

1. Simulate $\mathcal{E}(\mathbb{x})$; whenever it sends a challenge ρ to its prover oracle, answer as follows.
 - (a) Flip a fresh coin with heads probability θ .
 - (b) If heads: compute the honest PCP verifier's query set Q corresponding to ρ , query the PCP string for $\text{ans} := \Pi[Q]$, and return $(Q, \text{ans}, 1)$ to the extractor.
 - (c) If tails: return $\perp$ to the extractor.
2. Output the witness produced by $\mathcal{E}$.

Note that Step 1b in $\tilde{V}$ sends $(Q, \text{ans}, 1)$ to $\mathcal{E}$, while $\mathcal{E}$ expects to receive a tuple $(Q, \text{ans}, \text{pf})$ that includes the VC opening proof pf . A PCP verifier only has oracle access to the PCP string and it cannot generate either the VC commitment or its corresponding opening proof. However, this is not a problem because, being admissible, $\mathcal{E}$ is VC-oblivious, which means that it does not use the VC opening proof except through whether it is valid. In other words, it suffices for $\mathcal{E}$ to access the prover only through an oracle that, on each challenge, returns either $\perp$ or a triple (Q, ans, b) consisting of the query set, the claimed answers, and the commitment verification bit. The extractor's execution with this oracle is identically distributed to its original execution. In sum, we can replace VC opening proofs with commitment verification bits, as in the malicious PCP verifier.

With $\theta := 2/p$, a Chernoff bound shows that $\tilde{V}$ reads more than $m^* = \omega(\log n)$ honest PCP views with negligible probability. Suppose that PCP is zero-knowledge against malicious verifiers that read at most m^* views. The zero-knowledge simulator generates the view of $\tilde{V}$ from $\mathbb{x}$ alone, up to the zero-knowledge error. Since $\tilde{V}$ outputs a valid witness whenever $\mathcal{E}$ does, the simulator outputs a valid witness with probability at least $\theta - \kappa_{\text{ARG}} - \text{negl}(n)$. Hence, if $\kappa_{\text{ARG}} \leq 1/p$, this probability is at least $1/p - \text{negl}(n)$. Repeating the simulation and checking its outputs yields a BPP decider for $L(R)$, contradicting the assumption that $L(R) \notin \text{BPP}$.

Remark 3 (Imperfect completeness). *Our actual construction in Section 4.6 for the malicious PCP verifier $\tilde{V}$ is slightly different from the construction presented above, as it allows for imperfect completeness in both the PCP and the vector commitment scheme. In this setting, the commitment verification bit is no longer always 1; see Section 4.6 for details.*

The low-query case. When the PCP is not zero-knowledge, a similar argument still goes through, at the cost of a slightly weaker bound, provided the PCP makes few queries: whenever $|\Sigma|^q = \text{poly}(n)$, a single PCP check has only polynomially many possible answer sequences. Instead of simulating the honest PCP views seen by the extractor, we can then afford to guess them. We postpone the details to Section 4.7.

What about expected-time extractors? The analyses above exploit the worst-case nature of a strict rewind budget: a strict extractor must obey the same p -rewind cap on every execution, including executions in which useful continuations occur only with probability θ . On the other hand, an expected-time extractor may instead enter a long rewinding phase only after an acceptance event of probability θ , and then spend $\Theta(1/\theta)$ rewinds on that branch: the conditional cost $1/\theta$ is canceled by the probability θ of entering the branch. Expected-time analyses thus sidestep the lower bound, as we discuss next.

2.2 Expected-time security of Kilian's protocol

We discuss our expected-time soundness and knowledge soundness analyses for Kilian's protocol. The starting point is the expected-time soundness proof of [CDGSY24], which we recall in Section 2.2.1. In Section 2.2.2 we explain our soundness analysis, and in Section 2.2.3 we explain our knowledge soundness analysis.

2.2.1 Soundness analysis of [CDGSY24]

The analysis of [CDGSY24] formalizes the following intuition. First, the argument adversary can always be as successful as the best PCP adversary, since committing to a PCP string and honestly opening it is one available strategy. Second, if the argument adversary succeeds with probability higher than that, then it must be able to open the commitment inconsistently (with a similar advantage).

Optimal PCP string. Fix an instance $\mathfrak{x} \notin L(R)$ and a malicious argument prover $\tilde{\mathcal{P}}$ with expected running time t_{ARG}^* . For verifier randomness ρ , let $(\mathcal{Q}_\rho, \text{ans}_\rho, \text{pf}_\rho)$ denote the response given by $\tilde{\mathcal{P}}(\text{aux}, \rho)$, and let b_ρ denote the resulting decision bit of the Kilian verifier; we say that $(\mathcal{Q}_\rho, \text{ans}_\rho, \text{pf}_\rho)$ is *accepting* if $b_\rho = 1$. [CDGSY24] defines the “best” PCP string as follows. For every location $q \in [\ell]$ and symbol $\sigma \in \Sigma$, define

$$p(q, \sigma) := \Pr_{\rho \leftarrow \{0,1\}^r} [q \in \mathcal{Q}_\rho \wedge \text{ans}_\rho[q] = \sigma \wedge b_\rho = 1] .$$

The PCP string $\tilde{\Pi}^*$ is defined location-wise to be

$$\tilde{\Pi}^*[q] \in \arg \max_{\sigma \in \Sigma} p(q, \sigma) .$$

In other words, $\tilde{\Pi}^*[q]$ is the most likely value with which the malicious prover opens location q in an accepting execution. (If the prover never opens location q , then $\tilde{\Pi}^*[q]$ is defined to be the default symbol 0.) Note that $\tilde{\Pi}^*$ is not efficiently computable. Nevertheless, it is a well-defined PCP string.

By the law of total probability, the success probability of the argument adversary decomposes according to whether the PCP verifier, run with the fixed proof $\tilde{\Pi}^*$ on the same verifier challenge ρ , would also accept:

$$\Pr \left[\langle \tilde{\mathcal{P}}, \mathcal{V}(\mathfrak{x}) \rangle = 1 \right] \leq \Pr \left[\mathbf{V}^{\tilde{\Pi}^*}(\mathfrak{x}; \rho) = 1 \right] + \Pr \left[\begin{array}{l} \langle \tilde{\mathcal{P}}, \mathcal{V}(\mathfrak{x}) \rangle = 1 \\ \wedge \mathbf{V}^{\tilde{\Pi}^*}(\mathfrak{x}; \rho) = 0 \end{array} \right] . \quad (1)$$

The first term can be upper-bounded by ϵ_{PCP} according to PCP soundness.

VC adversary. It remains to bound the second term, where the argument verifier accepts but the PCP verifier rejects $\tilde{\Pi}^*$ on the same challenge. This event implies the existence of a query $q \in \mathcal{Q}_\rho$ such that $\text{ans}_\rho[q] \neq \tilde{\Pi}^*[q]$. [CDGSY24] exploits this inconsistency to build the following position-binding adversary A_{VC} , which guesses the location of the disagreement.

$A_{\text{VC}}(\mathfrak{x})$:

1. Run the argument prover to get the commitment: $(\text{cm}, \text{aux}) \leftarrow \tilde{\mathcal{P}}$.
2. Sample $\rho \leftarrow \{0, 1\}^r$ and compute $(\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\text{aux}, \rho)$. If $b_\rho = 0$, output $\perp$.
3. Sample $i \leftarrow [|\mathcal{Q}|]$ uniformly and let q be the i -th location of $\mathcal{Q}$. Rewind the prover on uniformly sampled verifier randomness until M (to be chosen later) accepting openings whose query sets contain q are collected.
4. If some collected opening opens q to a value different from $\text{ans}[q]$, output the two inconsistent openings; otherwise output $\perp$.

The second term of Equation 1 can be decomposed as

$$\begin{aligned} \Pr \left[\begin{array}{l} \langle \tilde{\mathcal{P}}, \mathcal{V}(\mathfrak{x}) \rangle = 1 \\ \wedge \mathbf{V}^{\tilde{\Pi}^*}(\mathfrak{x}; \rho) = 0 \end{array} \right] &\leq \mathfrak{q} \cdot \Pr \left[A_{\text{VC}} \text{ guesses } i \text{ such that } \text{ans}[\mathcal{Q}[i]] \neq \tilde{\Pi}^*[\mathcal{Q}[i]] \right] \\ &\leq \mathfrak{q} \cdot \Pr \left[\begin{array}{l} A_{\text{VC}} \text{ guesses } i \text{ such that } \text{ans}[\mathcal{Q}[i]] \neq \tilde{\Pi}^*[\mathcal{Q}[i]] \\ \wedge A_{\text{VC}} \text{ outputs a binding violation} \end{array} \right] \\ &\quad + \mathfrak{q} \cdot \Pr \left[\begin{array}{l} A_{\text{VC}} \text{ guesses } i \text{ such that } \text{ans}[\mathcal{Q}[i]] \neq \tilde{\Pi}^*[\mathcal{Q}[i]] \\ \wedge \text{all } M \text{ collected openings containing } \mathcal{Q}[i] \text{ open it to } \text{ans}[\mathcal{Q}[i]] \end{array} \right] . \end{aligned}$$

To see the first inequality, condition on the event on the left-hand side, so that a disagreeing location q exists; the probability of A_{VC} correctly guessing such a position is at least $1/\mathfrak{q}$. The second inequality is a decomposition by total probability as follows. If one of the M collected openings opens q to a value different from $\text{ans}[q]$, then A_{VC} outputs a binding violation; this gives the first term. Otherwise, they all open q to $\text{ans}[q]$; this gives the second term.

To bound the first term, [CDGSY24] proves that A_{VC} rewinds at most $O(\ell M)$ times in expectation. Thus, assuming argument executions are the dominant cost,

$$\Pr \left[\begin{array}{l} A_{\text{VC}} \text{ guesses } i \text{ such that } \text{ans}[\mathcal{Q}[i]] \neq \tilde{\Pi}^*[\mathcal{Q}[i]] \\ \wedge A_{\text{VC}} \text{ outputs a binding violation} \end{array} \right] \leq \epsilon_{\text{VC}}^*(O(\ell M \cdot t_{\text{ARG}}^*)) ,$$

where $\epsilon_{\text{VC}}^*(t_{\text{VC}}^*)$ is the position-binding error of VC against adversaries of expected running time t_{VC}^* .

To bound the second term, observe that

$$\Pr \left[\begin{array}{l} A_{\text{VC}} \text{ guesses } i \text{ such that } \text{ans}[\mathcal{Q}[i]] \neq \tilde{\Pi}^*[\mathcal{Q}[i]] \\ \wedge \text{all } M \text{ collected openings containing } \mathcal{Q}[i] \text{ open it to } \text{ans}[\mathcal{Q}[i]] \end{array} \right] \leq 2^{-M} .$$

To see this, fix a queried location q and suppose that $\text{ans}[q] \neq \tilde{\Pi}^*[q]$. The probability that an accepting opening opens q to $\text{ans}[q]$ is at most $1/2$, since the distinct symbol $\tilde{\Pi}^*[q]$ is at least as likely. Hence, the probability that all M collected openings open q to $\text{ans}[q]$ is at most 2^{-M} .

Final bound. Fix a parameter $\epsilon \in (0, 1)$. Taking $M := \lceil \log_2(q/\epsilon) \rceil$, [CDGSY24] concludes that

$$\Pr \left[\langle \tilde{\mathcal{P}}, \mathcal{V}(\mathbf{x}) \rangle = 1 \right] \leq \epsilon_{\text{PCP}} + q \cdot \epsilon_{\text{VC}}^*(O(\ell \log(q/\epsilon) \cdot t_{\text{ARG}}^*)) + \epsilon . \quad (2)$$

The additive overhead. The above expected-time analysis of [CDGSY24] is not lossless: unlike the adaptive-special-soundness analysis of [AKLY24], it pays the additive error ϵ as shown in Eq. 2. Intuitively, the error stems from the fact that A_{VC} never computes $\tilde{\Pi}^*[q]$, but rather approximates it from samples. Any finite number of samples is not guaranteed to contain $\tilde{\Pi}^*[q]$, and the reduction can only conclude that $\tilde{\Pi}^*[q]$ is missed with probability at most 2^{-M} ; this is exactly the additive error ϵ .

2.2.2 Our lossless soundness analysis

Roughly speaking, we obtain Theorem 2's lossless bound by defining a *different* “best” PCP string and showing how to *exactly* compute its symbol at position q (in expected time), thereby avoiding any approximation error.

Optimal PCP string. We define $\tilde{\Pi}$, which is not a fixed string but rather a *random variable*. Consider an experiment which rewinds $\tilde{\mathcal{P}}$ on uniformly sampled verifier randomness indefinitely, keeping the accepting openings, and let $(\mathcal{Q}_i, \text{ans}_i, \text{pf}_i)$ be the i -th accepting opening:

$$(\mathcal{Q}_1, \text{ans}_1, \text{pf}_1), (\mathcal{Q}_2, \text{ans}_2, \text{pf}_2), \dots$$

Let $N_q := \min\{i \geq 1 : q \in \mathcal{Q}_i\}$ be the index of the first opening that contains location q , and define the random PCP string $\tilde{\Pi}$ by

$$\tilde{\Pi}[q] := \begin{cases} \text{ans}_{N_q}[q] & \text{if } N_q < \infty \\ 0 & \text{otherwise.} \end{cases}$$

Although $\tilde{\Pi}$ is random, PCP soundness still applies to it: the sequence of openings is sampled before, and independently of, the verifier randomness. Thus, averaging yields $\Pr[\mathbf{V}^{\tilde{\Pi}}(\mathbf{x}; \rho) = 1] \leq \epsilon_{\text{PCP}}$, and the decomposition in Eq. 1 carries over with $\tilde{\Pi}$ in place of $\tilde{\Pi}^*$:

$$\begin{aligned} \Pr \left[\langle \tilde{\mathcal{P}}, \mathcal{V}(\mathbf{x}) \rangle = 1 \right] &\leq \Pr \left[\mathbf{V}^{\tilde{\Pi}}(\mathbf{x}; \rho) = 1 \right] + \Pr \left[\begin{array}{l} \langle \tilde{\mathcal{P}}, \mathcal{V}(\mathbf{x}) \rangle = 1 \\ \wedge \mathbf{V}^{\tilde{\Pi}}(\mathbf{x}; \rho) = 0 \end{array} \right] \\ &\leq \epsilon_{\text{PCP}} + \Pr \left[\begin{array}{l} \langle \tilde{\mathcal{P}}, \mathcal{V}(\mathbf{x}) \rangle = 1 \\ \wedge \mathbf{V}^{\tilde{\Pi}}(\mathbf{x}; \rho) = 0 \end{array} \right] . \end{aligned}$$

VC adversary. To bound the second term, we consider a new position-binding adversary A'_{VC} , which differs from Section 2.2.1's adversary in two ways. First, A'_{VC} does not have a cutoff: it continues rewinding until it finds an accepting opening whose query set contains q . Second, A'_{VC} in fact searches for all queries $q \in \mathcal{Q}$, instead of a random one. We present the construction below.

$A'_{\text{VC}}(\mathbb{z})$:

1. Run the argument prover to get the commitment: $(\text{cm}, \text{aux}) \leftarrow \tilde{\mathcal{P}}$.
2. Sample $\rho \leftarrow \{0, 1\}^r$ and compute $(\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\text{aux}, \rho)$. If the Kilian verifier rejects, output $\perp$.
3. Rewind the prover on uniformly sampled verifier randomness, keeping the accepting openings, until every query in $\mathcal{Q}$ appears in some accepting opening.
4. If some collected opening opens a location $q \in \mathcal{Q}$ to a value different from $\text{ans}[q]$, output the two inconsistent openings; otherwise output $\perp$.

For each $q \in \mathcal{Q}$, consider the *first* opening collected by A'_{VC} that contains q ; let σ_q denote the symbol with which this opening opens q . The key fact which we leverage is that each σ_q is distributed as $\tilde{\Pi}[q]$. Thus, the probability that some σ_q disagrees with $\text{ans}[q]$ (which implies A'_{VC} outputs a binding violation) is exactly the probability that $\tilde{\Pi}[\mathcal{Q}]$ disagrees with ans . We conclude that

$$\Pr \left[\begin{array}{l} \langle \tilde{\mathcal{P}}, \mathcal{V}(\mathbb{z}) \rangle = 1 \\ \wedge \mathbf{V}^{\tilde{\Pi}}(\mathbb{z}; \rho) = 0 \end{array} \right] \leq \Pr[A'_{\text{VC}} \text{ outputs a binding violation}] .$$

Bounding the running time. Conditioned on the argument verifier accepting, the initial opening $(\mathcal{Q}, \text{ans}, \text{pf})$ is distributed identically to each rewind accepting opening. Thus, the probability that the first n rewind openings fail to *cover* $\mathcal{Q}$ (i.e., contain every query in $\mathcal{Q}$) equals the probability that the $(n + 1)$ -th opening is *new* (i.e., reveals a location unseen so far). Summing over $n \in \mathbb{N}$, the expected number of accepting openings needed to cover $\mathcal{Q}$ equals the expected number of new openings. But there are at most ℓ new openings across the sequence (every new opening reveals at least one of the ℓ locations), so we deduce that A'_{VC} collects at most ℓ accepting openings in expectation. Each accepting opening in turn costs $1/\epsilon$ rewinds in expectation, where

$$\epsilon := \Pr_{\rho \leftarrow \{0,1\}^r} [b_\rho = 1] ,$$

denotes the prover's success probability. Thus, conditioned on Step 3 occurring, A'_{VC} makes at most ℓ/ϵ rewinds in expectation. But Step 3 only occurs with probability ϵ , so the unconditional expected number of rewinds is at most $\epsilon \cdot (\ell/\epsilon) = \ell$. The adversary A'_{VC} therefore runs in expected time $O(\ell \cdot t_{\text{ARG}}^*)$. We conclude that

$$\Pr[\langle \tilde{\mathcal{P}}, \mathcal{V}(\mathbb{z}) \rangle = 1] \leq \epsilon_{\text{PCP}} + \epsilon_{\text{VC}}^*(O(\ell \cdot t_{\text{ARG}}^*)) .$$

The above analysis does not suffice for knowledge soundness. One may expect our expected-time soundness analysis to carry over to expected-time knowledge soundness. However, difficulties arise when implementing the corresponding argument extractor. Since the soundness reduction only needs to determine whether one PCP verifier execution accepts, it suffices to obtain the local view $\tilde{\Pi}[\mathcal{Q}]$ by sampling until $\mathcal{Q}$ is covered. In particular, values outside $\mathcal{Q}$ are never used. In contrast, the PCP extractor $\mathbf{E}$ must receive a *complete* PCP string, and may depend on locations that the PCP verifier did not query in its execution. A local view of $\tilde{\Pi}$ is therefore insufficient.

Yet we also cannot simply reconstruct the entire $\tilde{\Pi}$. An expected-time reduction can afford a long rewinding phase, but only when the cost of the phase is canceled by the probability of entering it, as in the analysis of A'_{VC} (where the ℓ/ϵ rewinds are paid for by the acceptance probability ϵ of the initial execution). It is not possible to reconstruct the entire proof in a similar fashion: revealing a location that appears in accepting openings

with probability p takes $1/p$ accepting openings in expectation, and p can be arbitrarily small, with no event of matching probability to cancel this cost.

A tempting alternative is to follow the soundness reduction: sample ρ , sample openings until its query set is covered, and then fill all unopened locations with a default value. Unfortunately, the reconstructed proof (which is *not* equal to $\tilde{\Pi}$) depends on ρ , and PCP knowledge soundness does not apply to PCP strings that depend on the verifier randomness.

In sum, a new analysis is needed to establish the knowledge soundness of Kilian's protocol.

2.2.3 Our knowledge soundness analysis

Our analysis draws inspiration from both of the soundness analyses presented above. From Section 2.2.1 we take the finite cutoff: fix a complete PCP string after finitely many rewinds. From Section 2.2.2 we take the formulation of $\tilde{\Pi}$: define the PCP string according to the first answers seen across the sequence of rewind openings.

We first recall the definition of expected-time argument knowledge soundness. In contrast to the lower bounds, our upper bounds hold for *adaptive* knowledge soundness, where the malicious prover chooses the instance (see Remark 2).

Definition 2 (Adaptive expected-time knowledge soundness). *An interactive argument $\text{ARG} = (\mathcal{P}, \mathcal{V})$ for a relation R has **adaptive expected-time knowledge soundness error** κ_{ARG}^* **with extraction time** t_{ε}^* if there exists an algorithm $\mathcal{E}$ with expected running time at most t_{ε}^* such that, for every algorithm $\tilde{\mathcal{P}}$ with expected running time at most t_{ARG}^* ,*

$$\Pr \left[\begin{array}{l} (\mathbb{x}, \mathbb{w}) \notin R \\ \wedge b = 1 \end{array} \middle| \begin{array}{l} (\mathbb{x}, \text{aux}) \leftarrow \tilde{\mathcal{P}} \\ b \xleftarrow{\text{tr}} \langle \tilde{\mathcal{P}}(\text{aux}), \mathcal{V}(\mathbb{x}) \rangle \\ \mathbb{w} \leftarrow \mathcal{E}^{\tilde{\mathcal{P}}(\text{aux})}(\mathbb{x}, \text{tr}, b) \end{array} \right] \leq \kappa_{\text{ARG}}^* .$$

Above, tr and b are the transcript and decision of the argument execution. The extractor receives both, along with rewinding access to $\tilde{\mathcal{P}}(\text{aux})$, where aux is the state saved just after the prover chooses $\mathbb{x}$.

Since there are various ways to define the PCP knowledge soundness (as noted in Section 1.3), we recall the standard definition to avoid confusion.

Definition 3 (PCP knowledge soundness). *PCP = $(\mathbf{P}, \mathbf{V})$ for a relation R has **knowledge soundness error** κ_{PCP} **with extraction time** $t_{\mathbf{E}}$ if there exists a $t_{\mathbf{E}}$ -time probabilistic algorithm $\mathbf{E}$ such that, for every time-bounded prover $\tilde{\mathbf{P}}$,*

$$\Pr \left[\begin{array}{l} (\mathbb{x}, \mathbb{w}) \notin R \\ \wedge b = 1 \end{array} \middle| \begin{array}{l} (\mathbb{x}, \tilde{\Pi}) \leftarrow \tilde{\mathbf{P}} \\ \rho \leftarrow \{0, 1\}^r \\ b := \mathbf{V}^{\tilde{\Pi}}(\mathbb{x}; \rho) \\ \mathbb{w} \leftarrow \mathbf{E}(\mathbb{x}, \tilde{\Pi}, \rho) \end{array} \right] \leq \kappa_{\text{PCP}} .$$

Argument extractor. Our goal is to construct a knowledge extractor $\mathcal{E}$ for Kilian's protocol. The challenge is to efficiently sample a complete PCP string before the verifier randomness. To do this, our extractor follows the main ideas of the soundness analyses: sample a list T of M accepting openings, reconstruct a PCP string $\tilde{\Pi}_T$ by taking the first answer seen at each location, and run the PCP extractor on $\tilde{\Pi}_T$. Note that using a finite cutoff risks introducing an additive error, as seen in Section 2.2.1. To avoid this, our extractor carefully samples a final verifier randomness ρ^* , which is passed into $\mathbf{E}$ alongside $\tilde{\Pi}_T$, as follows. It continues rewinding until it finds an accepting opening $(Q^*, \text{ans}^*, \text{pf}^*)$ where Q^* is *covered* by the previously sampled M openings; ρ^* is then the corresponding verifier randomness. We present the construction below.

$\mathcal{E}^{\tilde{\mathcal{P}}(\text{aux})}(\mathbb{x}, \text{tr}, b)$:

1. If $b = 0$, output $\perp$.

2. Read the commitment cm from tr . All subsequent rewinds return to the prover's state immediately after it sent this commitment.
3. Rewind the prover on uniformly sampled verifier randomness until M (to be chosen later) accepting transcripts $T = ((\rho_1, \mathcal{Q}_1, \text{ans}_1, \text{pf}_1), \dots, (\rho_M, \mathcal{Q}_M, \text{ans}_M, \text{pf}_M))$ are collected.
4. Define the complete PCP string $\tilde{\Pi}_T$ by taking the first answer seen at every location in $\mathcal{Q}_1 \cup \dots \cup \mathcal{Q}_M$ and the default symbol elsewhere.
5. Rewind further until an accepting transcript $(\rho^*, \mathcal{Q}^*, \text{ans}^*, \text{pf}^*)$ with $\mathcal{Q}^* \subseteq \mathcal{Q}_1 \cup \dots \cup \mathcal{Q}_M$ is found.
6. Output $\mathbf{E}(\mathbb{x}, \tilde{\Pi}_T, \rho^*)$.

By the law of total probability, we can decompose the target probability as follows:

$$\Pr \left[\begin{array}{l} (\mathbb{x}, \mathbb{w}) \notin R \\ \wedge b = 1 \end{array} \right] \leq \Pr \left[\begin{array}{l} (\mathbb{x}, \mathbf{E}(\mathbb{x}, \tilde{\Pi}_T, \rho^*)) \notin R \\ \wedge \mathbf{V}^{\tilde{\Pi}_T}(\mathbb{x}; \rho^*) = 1 \end{array} \right] + \Pr[\text{ans}^* \neq \tilde{\Pi}_T[\mathcal{Q}^*]] . \quad (3)$$

Reducing to PCP knowledge soundness. At first glance, the first term of Eq. 3 can be bounded by the PCP knowledge soundness error κ_{PCP} . This is not the case, however, since the verifier randomness ρ^* is *not* uniformly sampled.

Instead, our analysis proceeds as follows. First, condition on T in Step 3. We say that T *covers* ρ if $\mathcal{Q}_\rho \subseteq \mathcal{Q}_1 \cup \dots \cup \mathcal{Q}_M$. Define

$$\delta := \Pr_{\rho}[T \text{ covers } \rho \mid b_{\rho} = 1] ,$$

so that a uniformly sampled ρ is accepting and covered with probability $\epsilon \cdot \delta$. Let B_T be the set of bad randomness for $\tilde{\Pi}_T$: the randomness ρ on which the PCP verifier $\mathbf{V}^{\tilde{\Pi}_T}(\mathbb{x}; \rho)$ accepts but the PCP extractor $\mathbf{E}(\mathbb{x}, \tilde{\Pi}_T, \rho)$ fails to output a valid witness; by PCP knowledge soundness, it holds that $\Pr_{\rho}[\rho \in B_T] \leq \kappa_{\text{PCP}}$. Thus,

$$\Pr[\rho^* \in B_T \mid T] = \frac{\Pr_{\rho}[\rho \in B_T \wedge T \text{ covers } \rho]}{\epsilon \cdot \delta} \leq \frac{\kappa_{\text{PCP}}}{\epsilon \cdot \delta} .$$

Note that $\mathcal{E}$ only starts rewinding if the initial interaction accepts, i.e., $b = 1$, which happens with probability ϵ . Averaging over the M accepting openings,

$$\Pr \left[\begin{array}{l} (\mathbb{x}, \mathbf{E}(\mathbb{x}, \tilde{\Pi}_T, \rho^*)) \notin R \\ \wedge \mathbf{V}^{\tilde{\Pi}_T}(\mathbb{x}; \rho^*) = 1 \end{array} \right] \leq \epsilon \cdot \mathbb{E} \left[\frac{\kappa_{\text{PCP}}}{\epsilon \cdot \delta} \right] = \kappa_{\text{PCP}} \cdot \mathbb{E} \left[\frac{1}{\delta} \right] .$$

It remains to bound $\mathbb{E}[1/\delta]$, which we call the *coverage factor*.

Bounding the coverage factor. For $W \subseteq [\ell]$, write p_W for the probability that a uniformly sampled accepting opening has its query set contained in W . In other words, $\delta = p_{\mathcal{Q}_1 \cup \dots \cup \mathcal{Q}_M}$. For every $z \in (0, 1]$, a union bound over the possible covered domains gives

$$\Pr[\delta \leq z] \leq \sum_{\substack{W \subseteq [\ell]: \\ p_W \leq z}} \Pr[\mathcal{Q}_1 \cup \dots \cup \mathcal{Q}_M = W] \leq \sum_{\substack{W \subseteq [\ell]: \\ p_W \leq z}} (p_W)^M \leq 2^{\ell} \cdot z^M ,$$

where the second inequality holds because $\mathcal{Q}_1 \cup \dots \cup \mathcal{Q}_M = W$ requires all M independent query sets to land inside W , an event of probability $(p_W)^M$. Then,

$$\mathbb{E}[1/\delta] = \int_0^{\infty} \Pr[1/\delta \geq u] du \leq 2^{\ell/M} + \int_{2^{\ell/M}}^{\infty} 2^{\ell} \cdot u^{-M} du = \frac{M}{M-1} \cdot 2^{\ell/M} ,$$

where we bound $\Pr[1/\delta \geq u]$ by 1 for $u \leq 2^{\ell/M}$ and by $2^{\ell} \cdot u^{-M}$ beyond. Fix a parameter $\alpha > 0$. Setting $M := \lceil (\ell + 2)/\log(1 + \alpha) \rceil$ makes the right-hand side at most $1 + \alpha$, so the first term of Eq. 3 is at most $(1 + \alpha) \cdot \kappa_{\text{PCP}}^2$.

²In Lemma A.2 we prove the bound $\mathbb{E}[1/\delta] \leq 2^{\ell/M}$, which enables setting $M := \lceil \ell/\log(1 + \alpha) \rceil$.

VC adversary. To bound the second term of Eq. 3, observe that, whenever the final answers ans^* disagree with $\tilde{\Pi}_T$ at a location of $\mathcal{Q}^*$, the extractor has found a violation of the VC position-binding property. To bound the extractor's running time, observe the following:

- Conditioned on $b = 1$, Step 3 requires M/ϵ rewinds in expectation.
- Conditioned on $b = 1$ and T , Step 5 requires $1/(\epsilon \cdot \delta)$ rewinds in expectation.

It follows that the extractor needs at most

$$\epsilon \cdot \left(\frac{M}{\epsilon} + \frac{1}{\epsilon} \cdot \mathbb{E}[1/\delta] \right) = M + \mathbb{E}[1/\delta] \leq M + 1 + \alpha = O\left(\frac{\ell}{\log(1 + \alpha)}\right)$$

rewinds in expectation, and the second term of Eq. 3 is bounded by $\epsilon_{\text{VC}}^*(O(\ell/\log(1 + \alpha)) \cdot t_{\text{ARG}}^*)$.

Final bound. Altogether, for every constant $\alpha > 0$,

$$\kappa_{\text{ARG}}^* = (1 + \alpha) \cdot \kappa_{\text{PCP}} + \epsilon_{\text{VC}}^* \left(O\left(\frac{\ell}{\log(1 + \alpha)}\right) \cdot t_{\text{ARG}}^* \right),$$

with the argument extractor running in expected time $O(\ell/\log(1 + \alpha)) \cdot t_{\text{ARG}}^* + t_{\text{E}}$.

Tightness of the expected-time knowledge soundness analysis. Our extractor rewinds the prover $O(\ell)$ times in expectation. Is this necessary? The lower bounds of Section 2.1 for strict-time knowledge extractors do not answer this question. We show that, assuming the randomized exponential time hypothesis (rETH), this rewinding cost is necessary up to a polylogarithmic factor for PCPs of quasilinear length.

Consider a PCP for 3SAT with $|\Sigma|^q = \text{poly}(m)$ and $\ell = m \log^{O(1)} m$, where m is the number of variables [Din07]. Assuming rETH, every expected polynomial-time VC-oblivious extractor for Kilian's protocol with knowledge error bounded away from 1 must rewind the prover $\Omega(m/\log m)$ times in expectation. For any fixed $\alpha > 0$, our extractor in Theorem 2 matches this lower bound up to a polylogarithmic factor.

The proof recycles the enumeration idea from the low-query case, with one twist: honest answers are no longer made rare. We run the extractor against the honest prover with a witness hardwired, which convinces the argument verifier with probability 1, so the extractor must output a valid witness with constant probability. Let p^* be its expected number of rewinds. Markov's inequality lets us truncate the execution after $O(p^*)$ rewinds while keeping the success probability constant. Our decider then explores the extractor's computation tree: at each of the $O(p^*)$ rewinds, it branches over all $|\Sigma|^q$ possible answer vectors, returning each together with the verification bit 1 (which is legitimate because the extractor is VC-oblivious). One branch in the computation tree reproduces the real execution, so there must be some branch that outputs a valid witness.

The search tree has size $(|\Sigma|^q)^{O(p^*)}$. Since $|\Sigma|^q = \text{poly}(m)$, this is at most $2^{O(p^* \log m)}$. If $p^* \log m = o(m)$, our decider therefore decides satisfiability of 3-CNF formulas in time $2^{o(m)}$, contradicting rETH. The formal statement and proof appear in Section 4.8.

2.3 Expected-time security of the IBCS protocol

We discuss our expected-time security analyses for the IBCS protocol, which extends Kilian's protocol to work with any public-coin IOP (a multiple-round PCP). In a k -round IOP, the prover sends a proof string Π_i and receives the verifier randomness ρ_i in each round $i \in [k]$. The IBCS protocol transforms this into a k -round interactive argument where the prover instead sends a commitment cm_i to Π_i ; after the last round, the prover opens the queried positions of every committed string.

2.3.1 Soundness

The analysis of [CDGSY24] (recalled in Section 2.2.1) directly extends to the IBCS protocol, still incurring the additive error ϵ and the factor q . Likewise, our new lossless analysis (overviewed in Section 2.2.2) directly extends to the IBCS protocol.

Indeed, given a malicious argument prover $\tilde{\mathcal{P}}$ we define a “best” IOP strategy as follows. For each round i and verifier randomness prefix $\rho_{<i}$, independently sample an infinite sequence of continuations of $\tilde{\mathcal{P}}$ on uniformly sampled randomness for the remaining rounds, conditioned on the IBCS verifier accepting. Define the round- i proof string $\tilde{\Pi}_i(\rho_{<i})$ by taking, at each location, the first answer seen in the round- i openings of this sequence, and the default symbol if the location never appears.

The position-binding adversary first samples an execution of the IBCS protocol with $\tilde{\mathcal{P}}$; let $(\rho_1, \dots, \rho_k)$ be the execution’s verifier randomness. If the verifier rejects, the adversary halts. Otherwise, for each round i , it rewinds from the prover’s round- i state to reconstruct the view of $\tilde{\Pi}_i(\rho_{<i})$ on the locations queried in this execution. With a similar analysis as in Section 2.2.2, we show that the expected number of rewinds for the i -th round is the round- i proof length ℓ_i . The expected running time of the adversary is therefore $O(\ell \cdot t_{\text{ARG}}^*)$, where $\ell = \sum_{i=1}^k \ell_i$ is the IOP’s total proof length. Thus, we obtain the expected-time soundness error bound

$$\epsilon_{\text{ARG}}^* \leq \epsilon_{\text{IOP}} + \epsilon_{\text{VC}}^*(O(\ell \cdot t_{\text{ARG}}^*)) .$$

2.3.2 Knowledge soundness

Generalizing our expected-time knowledge soundness to multiple rounds is less straightforward.

Prior art: trees of accepting transcripts. Previous rewinding extractors for multi-round protocols assemble a *tree of accepting transcripts*, i.e., a depth- k tree where:

- each i -th layer node is labeled with a prover message and has m_i children;
- each edge is labeled with verifier randomness;
- each path down the tree represents an accepting argument transcript.

Observe that such a tree has $\prod_i m_i$ leaves. In the IBCS protocol, the i -th round would seemingly require m_i to be on the order of ℓ_i . Hence, producing a tree of accepting transcripts would require on the order of $\prod_i \ell_i$ prover executions, which is prohibitively large. Indeed, the authors of [AKLY24], whose *adaptive special soundness* framework works with such trees, are only able to give a one-round analysis of Kilian’s protocol, leaving the multi-round IBCS analysis open. We remark that running time is not the only obstruction: as discussed in Section 1.3, their extractor also requires strong PCP knowledge soundness, which has no natural analogue for IOPs.

Our approach. Recall that Section 2.2.3’s extractor performs two distinct steps: it samples a list T of M accepting openings, and then samples one more accepting opening, covered by T , to pick ρ^* . At a high level, our extractor runs this two-step procedure once per round. At the start of round i , it has fixed the verifier randomness $(\rho_1^*, \dots, \rho_{i-1}^*)$ and obtained the lists $T_1, \dots, T_{i-1}$. It proceeds by sampling M_i accepting executions to form T_i ; crucially, each execution’s queries to the first $i - 1$ proof strings must be covered by the existing samples $T_1, \dots, T_{i-1}$, respectively. It then samples one more accepting execution, whose queries to the first i proof strings are covered by $T_1, \dots, T_i$, to pick ρ_i^* . In this way the extractor ensures that the final IOP transcript $(\tilde{\Pi}_{T_1}, \rho_1^*, \dots, \tilde{\Pi}_{T_k}, \rho_k^*)$ is accepting. We present the construction below.

$\mathcal{E}^{\tilde{\mathcal{P}}(\text{aux})}(\mathbb{X}, \text{tr}, b)$:

1. If $b = 0$, output $\perp$.
2. For $i = 1, \dots, k$:
 - (a) Initialize an empty list T_i .
 - (b) Repeat M_i (to be chosen later) times: rewind the prover after the prefix $(\rho_1^*, \dots, \rho_{i-1}^*)$ and sample uniformly from all verifier randomness $(\rho'_i, \dots, \rho'_k)$ that satisfies the following:
 - $\langle \tilde{\mathcal{P}}(\text{aux}), \mathcal{V}(\mathbb{X}; \rho_1^*, \dots, \rho_{i-1}^*, \rho'_i, \dots, \rho'_k) \rangle = 1$; and
 - for every $j < i$, the round- j opening is covered by T_j .
Append the round- i opening to T_i .
 - (c) Define the round- i proof $\tilde{\Pi}_{T_i}$ by taking the first answer seen in T_i at every covered location and the default symbol elsewhere.




Figure 1: Visualization of the rewinding strategy in $\mathcal{E}$. The M_i executions sampled in Step 2b determine $\tilde{\Pi}_{T_i}$ but are not used otherwise. After T_i is assembled, the execution sampled in Step 2d determines ρ_i^* , which extends the main transcript.

- (d) Sample $(\rho'_i, \dots, \rho'_k)$ as in Step 2b, with the additional requirement that the round- i opening is covered by T_i . Fix the round- i verifier randomness as $\rho_i^* := \rho'_i$.
3. Output $\mathbf{E}(\mathbb{X}, \tilde{\Pi}_{T_1}, \rho_1^*, \dots, \tilde{\Pi}_{T_k}, \rho_k^*)$.

Our extractor does not follow the “tree of accepting transcripts” approach outlined earlier. Rather, it aims to assemble a “path of accepting transcripts” (Figure 1): once Step 2d fixes ρ_i^* , every later sample extends the prefix $(\rho_1^*, \dots, \rho_i^*)$. This enables reconstructing a single (accepting) IOP transcript, while keeping the total cost small (roughly a sum over ℓ_i).

Why does a path suffice? Essentially it is because the standard definition of IOP knowledge soundness, presented below, only gives the IOP extractor a single transcript.

Definition 4 (IOP knowledge soundness). IOP = $(\mathbf{P}, \mathbf{V})$ for a relation R has **knowledge soundness error** κ_{IOP} with **extraction time** t_E if there exists a t_E -time algorithm $\mathbf{E}$ such that, for every prover $\tilde{\mathbf{P}}$,

$$\Pr \left[\begin{array}{l} (\mathbb{X}, \mathbb{W}) \notin R \\ \wedge b = 1 \end{array} \middle| \begin{array}{l} (\mathbb{X}, \tilde{\Pi}_1, \mathbf{aux}_1) \leftarrow \tilde{\mathbf{P}} \\ \text{for } i = 1, \dots, k-1: \\ \quad \rho_i \leftarrow \{0, 1\}^{r_i} \\ \quad (\tilde{\Pi}_{i+1}, \mathbf{aux}_{i+1}) \leftarrow \tilde{\mathbf{P}}(\mathbf{aux}_i, \rho_i) \\ \rho_k \leftarrow \{0, 1\}^{r_k} \\ b := \mathbf{V}^{\tilde{\Pi}_1, \dots, \tilde{\Pi}_k}(\mathbb{X}; \rho_1, \dots, \rho_k) \\ \mathbb{W} \leftarrow \mathbf{E}(\mathbb{X}, \tilde{\Pi}_1, \rho_1, \dots, \tilde{\Pi}_k, \rho_k) \end{array} \right] \leq \kappa_{\text{IOP}} .$$

We now bound the adversary’s success probability. As in Section 2.2.3, we have the following decomposition:

$$\begin{aligned} \Pr \left[\begin{array}{l} (\mathbb{X}, \mathbb{W}) \notin R \\ \wedge b = 1 \end{array} \right] &\leq \Pr \left[\begin{array}{l} (\mathbb{X}, \mathbb{W}) \notin R \\ \wedge \mathbf{V}^{\tilde{\Pi}_{T_1}, \dots, \tilde{\Pi}_{T_k}}(\mathbb{X}; \rho_1^*, \dots, \rho_k^*) = 1 \end{array} \right] \\ &\quad + \Pr \left[\begin{array}{l} \text{an accepting opening sampled by the extractor disagrees} \\ \text{with some } \tilde{\Pi}_{T_i} \text{ at a location covered by } T_i \end{array} \right] . \end{aligned} \tag{4}$$

The second term is bounded by the position-binding error of VC against adversaries with the expected running time of the extractor, as before. It remains to bound the first term.

Reducing to IOP knowledge soundness. As in Section 2.2.3, the first term of Eq. 4 is not necessarily bounded by κ_{IOP} , since the verifier randomness $\rho_1^*, \dots, \rho_k^*$ is not uniformly distributed.

Instead, our analysis proceeds as follows. Define

$$\begin{aligned} \kappa_k(\Pi_1, \rho_1, \dots, \Pi_k, \rho_k) &:= \mathbb{1} \left[\begin{array}{l} \mathbf{V}^{\Pi_1, \dots, \Pi_k}(\mathbb{X}; \rho_1, \dots, \rho_k) = 1 \\ \wedge (\mathbb{X}, \mathbf{E}(\mathbb{X}, \Pi_1, \rho_1, \dots, \Pi_k, \rho_k)) \notin R \end{array} \right] , \\ \kappa_{i-1}(\Pi_1, \rho_1, \dots, \Pi_{i-1}, \rho_{i-1}) &:= \max_{\Pi_i \in \Sigma^{\ell_i}} \mathbb{E}_{\rho_i \leftarrow \{0, 1\}^{r_i}} [\kappa_i(\Pi_1, \rho_1, \dots, \Pi_i, \rho_i)] . \end{aligned}$$




Figure 2: The distribution of ρ_i^* in a toy example: the accepting covered continuations of the prefix $\rho_{<i}^*$, grouped by their round- i verifier randomness; the marked value ρ_i^l has six of the ten continuations. Step 2d picks one continuation uniformly, so each value is selected with probability proportional to its number of continuations, that is, proportional to ϵ_{i+1} .

In other words, $\kappa_i(\Pi_1, \rho_1, \dots, \Pi_i, \rho_i)$ is the largest probability of accepting without extracting a witness that an IOP prover can achieve from a fixed partial transcript $(\Pi_1, \rho_1, \dots, \Pi_i, \rho_i)$. By IOP knowledge soundness, it holds that $\kappa_0 \leq \kappa_{\text{IOP}}$. Conditioned on $b = 1$, the event in the first term of Eq. 4 occurs with probability $\mathbb{E}[\kappa_k(\tilde{\Pi}_{T_1}, \rho_1^*, \dots, \tilde{\Pi}_{T_k}, \rho_k^*)]$. It remains to relate this expectation to κ_0 .

Fix the instance, first commitment, and corresponding prover state $(x, \text{cm}_1, \text{aux}_1)$ from the initial interaction and, for round i , condition on the verifier randomness $\rho_1^*, \dots, \rho_{i-1}^*$ fixed in the earlier rounds and on $T_1, \dots, T_i$. Let ϵ_i be the probability that uniformly sampled verifier randomness $(\rho_i', \dots, \rho_k')$ satisfies the two conditions of Step 2b in $\mathcal{E}$: (i) the corresponding execution of the argument accepts; and (ii) for every $j < i$, the round- j opening is covered by T_j . Let δ_i be the probability, conditioned on these two conditions, that the round- i opening is covered by T_i , as required in Step 2d of $\mathcal{E}$. As in the analysis for Kilian's protocol, the verifier randomness ρ_i^* is not uniformly sampled; in particular, for every $\rho_i \in \{0, 1\}^{r_i}$, the probability that $\rho_i^* = \rho_i$ is proportional to the probability that ρ_i leads to an accepting covered continuation (Figure 2):

$$\Pr[\rho_i^* = \rho_i] = \frac{\epsilon_{i+1}(\rho_i)}{2^{r_i} \cdot \epsilon_i \cdot \delta_i},$$

where $\epsilon_{i+1}(\rho_i)$ is ϵ_{i+1} with the round- i verifier randomness set to ρ_i : the probability that, given the prefix $(\rho_1^*, \dots, \rho_{i-1}^*, \rho_i)$, uniformly sampled verifier randomness for the later rounds leads to an accepting execution whose round- j opening is covered by T_j for every $j \leq i$. Consequently, suppressing the transcript prefix $(\tilde{\Pi}_{T_1}, \rho_1^*, \dots, \tilde{\Pi}_{T_{i-1}}, \rho_{i-1}^*)$ from the arguments of κ_{i-1} and κ_i ,

$$\mathbb{E}_{\rho_i^*} \left[\frac{\kappa_i(\tilde{\Pi}_{T_i}, \rho_i^*)}{\epsilon_{i+1}(\rho_i^*)} \right] \leq \frac{\mathbb{E}_{\rho_i \leftarrow \{0,1\}^{r_i}} [\kappa_i(\tilde{\Pi}_{T_i}, \rho_i)]}{\epsilon_i \cdot \delta_i} \leq \frac{\kappa_{i-1}}{\epsilon_i \cdot \delta_i}. \quad (5)$$

Let

$$C_i := \frac{M_i}{M_i - 1} \cdot 2^{\ell_i/M_i} \geq \mathbb{E}[1/\delta_i],$$

where the inequality follows from the calculation in Section 2.2.3.

We apply Eq. 5 from round k down to round 1. Observe that below, $\epsilon_{k+1} = 1$ because the execution sampled in Step 2d at round k accepts and its round- j opening is covered by T_j for every $j \leq k$. Therefore, conditioned on $b = 1$, using Eq. 5, we have

$$\begin{aligned} \mathbb{E}[\kappa_k(\tilde{\Pi}_{T_1}, \rho_1^*, \dots, \tilde{\Pi}_{T_k}, \rho_k^*)] &= \mathbb{E} \left[\frac{\kappa_k(\tilde{\Pi}_{T_1}, \rho_1^*, \dots, \tilde{\Pi}_{T_k}, \rho_k^*)}{\epsilon_{k+1}} \right] \\ &\leq C_k \cdot \mathbb{E} \left[\frac{\kappa_{k-1}(\tilde{\Pi}_{T_1}, \rho_1^*, \dots, \tilde{\Pi}_{T_{k-1}}, \rho_{k-1}^*)}{\epsilon_k} \right] \end{aligned}$$

$$\begin{aligned}
&\leq C_k \cdot C_{k-1} \cdot \mathbb{E} \left[\frac{\kappa_{k-2}(\tilde{\Pi}_{T_1}, \rho_1^*, \dots, \tilde{\Pi}_{T_{k-2}}, \rho_{k-2}^*)}{\epsilon_{k-1}} \right] \\
&\leq \dots \\
&\leq \left(\prod_{i=1}^k C_i \right) \cdot \frac{\kappa_0}{\epsilon_1} .
\end{aligned}$$

Finally, ϵ_1 is exactly the probability that the initial interaction accepts. Since the bound above is conditioned on this event, the first term of Eq. 4 equals

$$\epsilon_1 \cdot \mathbb{E} \left[\kappa_k(\tilde{\Pi}_{T_1}, \rho_1^*, \dots, \tilde{\Pi}_{T_k}, \rho_k^*) \right] \leq \epsilon_1 \cdot \left(\prod_{i=1}^k C_i \right) \cdot \frac{\kappa_0}{\epsilon_1} \leq \left(\prod_{i=1}^k C_i \right) \cdot \kappa_{\text{IOP}} .$$

Note that C_i decreases to 1 as M_i grows. We set $M_i := \lceil k \cdot (\ell_i + 2) / \log_2(1 + \alpha) \rceil$, which gives $C_i \leq (1 + \alpha)^{1/k}$ and hence $\prod_i C_i \leq 1 + \alpha$, independent of the round complexity.

Bounding the running time. The same calculation as in Eq. 5 applies to any nonnegative functions $F_0, \dots, F_k$ satisfying $\mathbb{E}_{\rho_i}[F_i] \leq F_{i-1}$ for uniformly sampled ρ_i . Under the same conditioning on the preceding history and T_i , we obtain

$$\mathbb{E}_{\rho_i^*} \left[\frac{F_i(\rho_i^*)}{\epsilon_{i+1}(\rho_i^*)} \right] \leq \frac{\mathbb{E}_{\rho_i \leftarrow \{0,1\}^{r_i}} [F_i(\rho_i)]}{\epsilon_i \cdot \delta_i} \leq \frac{F_{i-1}}{\epsilon_i \cdot \delta_i} . \quad (6)$$

We now take F_i to be the prover's conditional expected running time given the round- i prefix, so that $F_0 = t_{\text{ARG}}^*$. Conditioned on $b = 1$, each entry of T_i costs expected time F_{i-1}/ϵ_i by rejection sampling, and the sample of Step 2d costs $F_{i-1}/(\epsilon_i \cdot \delta_i)$. Applying Eq. 6 as before,

$$\mathbb{E} \left[\frac{F_{i-1}}{\epsilon_i} \right] \leq C_{i-1} \cdot \mathbb{E} \left[\frac{F_{i-2}}{\epsilon_{i-1}} \right] \leq \dots \leq \left(\prod_{j < i} C_j \right) \cdot \frac{t_{\text{ARG}}^*}{\epsilon_1} ,$$

and averaging over T_i bounds the cost of the sample of Step 2d by $C_i \cdot \mathbb{E}[F_{i-1}/\epsilon_i]$. Summing over the rounds, the expected running time conditioned on $b = 1$ is at most

$$\sum_{i=1}^k (M_i + C_i) \cdot \left(\prod_{j < i} C_j \right) \cdot \frac{t_{\text{ARG}}^*}{\epsilon_1} \leq (1 + \alpha) \cdot \left(\sum_{i=1}^k (M_i + C_i) \right) \cdot \frac{t_{\text{ARG}}^*}{\epsilon_1} .$$

Multiplying by the success probability ϵ_1 cancels the $1/\epsilon_1$, and $\sum_{i=1}^k (M_i + C_i) = O(k \cdot \ell / \log(1 + \alpha))$ by our choice of M_i , so unconditionally the extractor spends expected time $O(k \cdot \ell / \log(1 + \alpha) \cdot t_{\text{ARG}}^*)$ running the prover.

Final bound. Altogether, for every constant $\alpha > 0$,

$$\kappa_{\text{ARG}}^* \leq (1 + \alpha) \cdot \kappa_{\text{IOP}} + \epsilon_{\text{VC}}^* \left(O \left(\frac{k \cdot \ell}{\log(1 + \alpha)} \cdot t_{\text{ARG}}^* \right) \right) ,$$

with the argument extractor running in expected time $O(k \cdot \ell / \log(1 + \alpha) \cdot t_{\text{ARG}}^*) + t_{\text{E}}$.

2.4 Extension to functional queries

We discuss how our expected-time analyses extend to the Funky protocol [CGKY26], which combines a functional IOP (FIOP) with a functional commitment scheme. In an FIOP, the verifier reads the round- i committed string Π_i only through a class $\mathbf{Q}$ of query functions. Thus, an opening no longer assigns symbols to locations; rather, it induces *constraints* (α, β) such that the committed string Π satisfies $\alpha(\Pi) = \beta$.

In Section 2.2.2, the sampled openings defined the proof through the first answer seen at every covered location; here, to construct the best FIOP string Π_i^* from the openings sampled in round i , process their constraints in order, append each constraint not already determined by the previous accepting openings, and let Π_i^* be any string that satisfies the appended constraints. This is the functional analogue of taking the first answer seen at every covered location.

For point queries, a contradiction among collected constraints is two valid openings that disagree at a location, which implies a position-binding violation. For general query classes, a contradiction means that no single string satisfies the constraints of all the collected openings, even though every one of them was accepted with respect to the same commitment; this constitutes a violation of the *function binding* property of the functional commitment (Definition 3.18).

Recall from Section 2.2.2 that the expected number of rewinds for Kilian's protocol is the proof length ℓ , because only ℓ openings can ever reveal a new location. For a general query class, our construction of the FIOP strings gives the same bound with the same proof: each appended constraint strictly shrinks a nonempty solution space, so the number of appends is at most the length of the longest strictly decreasing chain of solution spaces, the diameter d_Q of the class (Definition 3.12). (For example, for linear queries the appended constraints form a row basis and the diameter is again the proof length.)

Let $d := \sum_{i=1}^k d_Q(\ell_i)$ be the sum of diameters from each round's query class. We obtain, for every constant $\alpha > 0$,

$$\begin{aligned}\epsilon_{\text{ARG}}^* &\leq \epsilon_{\text{FIOP}} + \epsilon_{\text{FC}}^*(O(d \cdot t_{\text{ARG}}^*)) , \\ \kappa_{\text{ARG}}^* &\leq (1 + \alpha) \cdot \kappa_{\text{FIOP}} + \epsilon_{\text{FC}}^*(O_\alpha(k \cdot d \cdot t_{\text{ARG}}^*)) ,\end{aligned}$$

with the argument extractor running in expected time $O_\alpha(k \cdot d \cdot t_{\text{ARG}}^*) + t_{\text{E}}$. For point queries the diameter equals the proof length and we recover the bounds in Theorem 2. The formal theorems and proofs are in Sections 5 and 6.

3 Preliminaries

We define the interactive argument security notions used throughout the paper and the functional commitments and oracle proofs used in the expected-time analyses, and then recall the Funky compiler that combines them. The vector commitment and PCP definitions, which are specializations of the functional commitments and oracle proof definitions, are stated in Section 4.

Definition 3.1. A **relation** R is a set of pairs $(\mathbb{x}, \mathbb{w})$ where $\mathbb{x}$ is an instance and $\mathbb{w}$ a witness. The corresponding **language** $L(R)$ is the set of instances $\mathbb{x}$ for which there exists a witness $\mathbb{w}$ such that $(\mathbb{x}, \mathbb{w}) \in R$.

3.1 Interactive arguments

An *interactive argument* (in the common reference string model) for a relation R is a tuple of polynomial-time algorithms $\text{ARG} = (\mathcal{G}, \mathcal{P}, \mathcal{V})$ that satisfies the following properties.

Definition 3.2 (Perfect completeness). $\text{ARG} = (\mathcal{G}, \mathcal{P}, \mathcal{V})$ for a relation R has **perfect completeness** if for every security parameter $\lambda \in \mathbb{N}$, instance size bound $n \in \mathbb{N}$, and instance-witness pair $(\mathbb{x}, \mathbb{w}) \in R$ with $|\mathbb{x}| \leq n$,

$$\Pr \left[\langle \mathcal{P}(\text{pp}, \mathbb{x}, \mathbb{w}), \mathcal{V}(\text{pp}, \mathbb{x}) \rangle = 1 \mid \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \right] = 1 .$$

Definition 3.3 (Adaptive soundness). $\text{ARG} = (\mathcal{G}, \mathcal{P}, \mathcal{V})$ for a relation R has **(adaptive strict-time) soundness error** ϵ_{ARG} if for every security parameter $\lambda \in \mathbb{N}$, instance size bound $n \in \mathbb{N}$, auxiliary input distribution $\mathcal{D}$, adversary time bound $t_{\text{ARG}} \in \mathbb{N}$, and t_{ARG} -time algorithm $\tilde{\mathcal{P}}$,

$$\Pr \left[\begin{array}{l} |\mathbb{x}| \leq n \\ \wedge \mathbb{x} \notin L(R) \\ \wedge b = 1 \end{array} \mid \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \eta \leftarrow \mathcal{D} \\ (\mathbb{x}, \text{aux}) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \eta) \\ b \leftarrow \langle \tilde{\mathcal{P}}(\text{aux}), \mathcal{V}(\text{pp}, \mathbb{x}) \rangle \end{array} \right] \leq \epsilon_{\text{ARG}}(\lambda, n, t_{\text{ARG}}) .$$

Definition 3.4 (Adaptive expected-time soundness). $\text{ARG} = (\mathcal{G}, \mathcal{P}, \mathcal{V})$ for a relation R has **(adaptive) expected-time soundness error** ϵ_{ARG}^* if for every security parameter $\lambda \in \mathbb{N}$, instance size bound $n \in \mathbb{N}$, auxiliary input distribution $\mathcal{D}$, adversary expected-time bound $t_{\text{ARG}}^* \in \mathbb{N}$, and algorithm $\tilde{\mathcal{P}}$ with expected running time at most t_{ARG}^* ,

$$\Pr \left[\begin{array}{l} |\mathbb{x}| \leq n \\ \wedge \mathbb{x} \notin L(R) \\ \wedge b = 1 \end{array} \mid \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \eta \leftarrow \mathcal{D} \\ (\mathbb{x}, \text{aux}) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \eta) \\ b \leftarrow \langle \tilde{\mathcal{P}}(\text{aux}), \mathcal{V}(\text{pp}, \mathbb{x}) \rangle \end{array} \right] \leq \epsilon_{\text{ARG}}^*(\lambda, n, t_{\text{ARG}}^*) .$$

Definition 3.5 (Adaptive knowledge soundness). $\text{ARG} = (\mathcal{G}, \mathcal{P}, \mathcal{V})$ for a relation R has **(adaptive) knowledge soundness error** κ_{ARG} **with extraction time** $t_{\mathcal{E}}$ if there exists a probabilistic algorithm $\mathcal{E}$ such that for every security parameter $\lambda \in \mathbb{N}$, instance size bound $n \in \mathbb{N}$, auxiliary input distribution $\mathcal{D}$, adversary time bound $t_{\text{ARG}} \in \mathbb{N}$, and t_{ARG} -time algorithm $\tilde{\mathcal{P}}$,

$$\Pr \left[\begin{array}{l} |\mathbb{x}| \leq n \\ \wedge (\mathbb{x}, \mathbb{w}) \notin R \\ \wedge b = 1 \end{array} \mid \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \eta \leftarrow \mathcal{D} \\ (\mathbb{x}, \text{aux}) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \eta) \\ b \xleftarrow{\text{tr}} \langle \tilde{\mathcal{P}}(\text{aux}), \mathcal{V}(\text{pp}, \mathbb{x}) \rangle \\ \mathbb{w} \leftarrow \mathcal{E}^{\tilde{\mathcal{P}}(\text{aux})}(\text{pp}, \mathbb{x}, \text{tr}, b) \end{array} \right] \leq \kappa_{\text{ARG}}(\lambda, n, t_{\text{ARG}}) ;$$

moreover, $\mathcal{E}$ runs in time $t_{\mathcal{E}}(\lambda, n, t_{\text{ARG}})$.

Definition 3.6 (Adaptive expected-time knowledge soundness). $\text{ARG} = (\mathcal{G}, \mathcal{P}, \mathcal{V})$ for a relation R has **(adaptive) expected-time knowledge soundness error** κ_{ARG}^* **with extraction time** $t_{\mathcal{E}}^*$ if there exists a probabilistic algorithm $\mathcal{E}$ such that for every security parameter $\lambda \in \mathbb{N}$, instance size bound $n \in \mathbb{N}$, auxiliary input distribution $\mathcal{D}$, adversary expected-time bound $t_{\text{ARG}}^* \in \mathbb{N}$, and algorithm $\tilde{\mathcal{P}}$ with expected running time at most t_{ARG}^* ,

$$\Pr \left[\begin{array}{l} |\mathbb{x}| \leq n \\ \wedge (\mathbb{x}, \mathbb{w}) \notin R \\ \wedge b = 1 \end{array} \middle| \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \eta \leftarrow \mathcal{D} \\ (\mathbb{x}, \text{aux}) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \eta) \\ b \stackrel{\text{tr}}{\leftarrow} \langle \tilde{\mathcal{P}}(\text{aux}), \mathcal{V}(\text{pp}, \mathbb{x}) \rangle \\ \mathbb{w} \leftarrow \mathcal{E}^{\tilde{\mathcal{P}}(\text{aux})}(\text{pp}, \mathbb{x}, \text{tr}, b) \end{array} \right] \leq \kappa_{\text{ARG}}^*(\lambda, n, t_{\text{ARG}}^*) ;$$

moreover, $\mathcal{E}$ has expected running time at most $t_{\mathcal{E}}^*(\lambda, n, t_{\text{ARG}}^*)$.

Above, $b \stackrel{\text{tr}}{\leftarrow} \langle \tilde{\mathcal{P}}(\text{aux}), \mathcal{V}(\text{pp}, \mathbb{x}) \rangle$ denotes the fact that tr is the transcript of the interaction (i.e., messages exchanged between $\tilde{\mathcal{P}}$ and $\mathcal{V}$). Moreover, $\mathcal{E}^{\tilde{\mathcal{P}}}$ means that $\mathcal{E}$ has black-box access to (each next-message function of) $\tilde{\mathcal{P}}$; in particular $\mathcal{E}$ can send verifier messages to $\tilde{\mathcal{P}}$ in order to obtain the next message of $\tilde{\mathcal{P}}$ (for a partial interaction where $\mathcal{V}$ sent those messages).

For security parameter $\lambda \in \mathbb{N}$ and instance size bound $n \in \mathbb{N}$, we write $t_{\mathcal{V}} := t_{\mathcal{V}}(\lambda, n)$ to denote the worst-case running time of $\mathcal{V}$ (even when interacting with a malicious prover $\tilde{\mathcal{P}}$).

Remark 3.7. The argument generator $\mathcal{G}$ receives two inputs: the security parameter λ and an instance size bound n . This means that the public parameter sampled by $\mathcal{G}$ may work only for instances of size at most n . However, one could consider the stronger notion where the sampled public parameter works for all instance sizes; in this case $\mathcal{G}$ receives only λ as input. Our analysis works for both cases.

Remark 3.8 (Deterministic adversaries). Following prior work, we assume that $\tilde{\mathcal{P}}$ is deterministic when proving upper bounds on (expected-time) soundness and knowledge soundness. This assumption is purely for notational simplicity (specifically, it allows us to avoid tracking the prover's intermediate states), and our approach works for probabilistic adversaries as well.

Remark 3.9 (Plain model variant). The above definitions consider interactive arguments in the *common reference string model*, where a generator samples a public parameter used by the argument prover and the argument verifier. One could also consider interactive arguments in the *plain model*, where there is no generator. This latter notion is implied, at the cost of an additional verifier message, as we now explain.

Suppose that $(\mathcal{G}, \mathcal{P}, \mathcal{V})$ is an interactive argument in the common reference string model. We describe an interactive argument $(\mathcal{P}', \mathcal{V}')$ in the plain model with an additional verifier message. The argument prover $\mathcal{P}'$ receives as input an instance $\mathbb{x}$ and witness $\mathbb{w}$, and the argument verifier $\mathcal{V}'$ receives as input the instance $\mathbb{x}$; both also receive as input the security parameter λ (in unary). They interact as follows:

- $\mathcal{V}'$ samples a public parameter $\text{pp} \leftarrow \mathcal{G}(1^\lambda, |\mathbb{x}|)$ and sends pp to $\mathcal{P}'$;
- $\mathcal{P}'$ and $\mathcal{V}'$ simulate an interaction of $\mathcal{P}(\text{pp}, \mathbb{x}, \mathbb{w})$ and $\mathcal{V}(\text{pp}, \mathbb{x})$.

It is straightforward to see that $(\mathcal{P}', \mathcal{V}')$ satisfies the standard definitions of completeness, soundness, and knowledge soundness for interactive arguments in the plain model.³ In fact, it would suffice for $(\mathcal{G}, \mathcal{P}, \mathcal{V})$ to satisfy the non-adaptive relaxations of soundness and knowledge soundness.

3.2 Query classes

Definition 3.10. A **query class** $\mathbf{Q}$ is a set of functions of the form $\alpha: \Sigma^\ell \rightarrow \mathbb{D}$, where Σ is the input alphabet, $\ell \in \mathbb{N}$ is the input length, and $\mathbb{D}$ is the answer domain. Here, Σ and $\mathbb{D}$ are finite sets of cardinality ≥ 2 .

³These standard definitions can be derived from Definitions 3.2, 3.3 and 3.5 by setting pp to be empty.

Definition 3.11. A **constraint** of a query class $\mathbf{Q}$ is a pair $(\alpha, \beta) \in \mathbf{Q} \times \mathbb{D}$. For a set of constraints $\mathcal{C}$, its solution space is

$$\text{Sol}(\mathcal{C}) := \left\{ \Pi \in \Sigma^\ell : \forall (\alpha, \beta) \in \mathcal{C}, \alpha(\Pi) = \beta \right\} ;$$

we say $\mathcal{C}$ is **satisfiable** if $\text{Sol}(\mathcal{C}) \neq \emptyset$. For satisfiable $\mathcal{C}$, a constraint (α, β) is:

- **entailed** by $\mathcal{C}$ if $\alpha(\Pi) = \beta$ for every $\Pi \in \text{Sol}(\mathcal{C})$.
- **refuted** by $\mathcal{C}$ if $\alpha(\Pi) \neq \beta$ for every $\Pi \in \text{Sol}(\mathcal{C})$.
- **decided** by $\mathcal{C}$ if it is entailed or refuted, and **undecided** otherwise.

Definition 3.12. The **diameter** of a query class $\mathbf{Q}$, denoted $d_{\mathbf{Q}}$, is the largest integer $d \geq 0$ for which there exist constraints $p_1, \dots, p_d$ satisfying

$$\emptyset \neq \text{Sol}(\{p_1, \dots, p_i\}) \subsetneq \text{Sol}(\{p_1, \dots, p_{i-1}\}) \quad \text{for every } i \in [d] ,$$

or ∞ if no largest d exists.

Definition 3.13. For a satisfiable constraint set $\mathcal{C}$ and a constraint (α, β) , define

$$\mathcal{C} \oplus (\alpha, \beta) := \begin{cases} \mathcal{C} \cup \{(\alpha, \beta)\} & \text{if } (\alpha, \beta) \text{ is undecided by } \mathcal{C}, \\ \mathcal{C} & \text{otherwise.} \end{cases}$$

The **refinement** of a constraint list $L = (c_1, \dots, c_n)$ is $\mathcal{C}_L := \emptyset \oplus c_1 \oplus \dots \oplus c_n$.

Definition 3.14. A **constraint solver** for a query class $\mathbf{Q}$ is a stateful algorithm $\text{Solver}_{\mathbf{Q}}$ that supports the following operations on an (initially empty) list of constraints L :

- Given a constraint, report whether it is entailed, refuted, or undecided by $\mathcal{C}_L$.
- Append an undecided constraint to L .
- Output an element $\Pi \in \text{Sol}(\mathcal{C}_L)$; for simplicity, we assume the choice is deterministic.

We write $t_{\mathbf{Q}}$ to denote the running time of one operation.

Query sets. Let $\alpha \subseteq \mathbf{Q}$ be an (ordered) query set. We write $\alpha(\Pi) \in \mathbb{D}^\alpha$ to denote the map sending $\alpha \in \alpha$ to $\alpha(\Pi)$. Given an answer set $\beta \in \mathbb{D}^\alpha$, we call (α, β) a *query-answer tuple*. By convention, we say:

- “ $\mathcal{C}$ entails (α, β) ” if $\mathcal{C}$ entails every constraint in (α, β) .
- “ $\mathcal{C}$ refutes (α, β) ” if $\mathcal{C}$ refutes at least one constraint in (α, β) .
- “ $\mathcal{C}$ decides (α, β) ” if $\mathcal{C}$ decides every constraint in (α, β) .

We often interpret a query-answer tuple as a list of constraints. Similarly, we interpret a list of query-answer tuples as a list of constraints, i.e., the list obtained by concatenating the constraints from each tuple.

Families of query classes. We often say “query class” to mean a *family* of query classes $\mathbf{Q} = (\mathbf{Q}_\ell)_{\ell \in \mathbb{N}}$ over a common alphabet Σ and answer domain $\mathbb{D}$, where each $\mathbf{Q}_\ell$ is a set of functions $\alpha: \Sigma^\ell \rightarrow \mathbb{D}$. Consequently, $d_{\mathbf{Q}}$ and $t_{\mathbf{Q}}$ are functions of $\ell \in \mathbb{N}$.

Definition 3.15 (Point queries). The class of point queries $\mathbf{Q}_{\text{Point}}$ is defined to be

$$\mathbf{Q}_{\text{Point}}(\ell) := \{ \alpha_i: \Sigma^\ell \rightarrow \Sigma \}_{i \in [\ell]} ,$$

where α_i maps Π to its i -th symbol $\Pi[i]$.

3.3 Functional commitments

An (interactive) *functional commitment scheme* for a query class $\mathbf{Q} = (\mathbf{Q}_\ell)_{\ell \in \mathbb{N}}$ over Σ is a tuple of algorithms

$$\text{FC} = (\text{FC.Gen}, \text{FC.Commit}, \mathcal{P}_{\text{FC}}, \mathcal{V}_{\text{FC}})$$

with the following syntax.

- $\text{FC.Gen}(1^\lambda, \ell_{\max}) \rightarrow \text{pp}$: On input a security parameter $\lambda \in \mathbb{N}$ and message length bound $\ell_{\max} \in \mathbb{N}$, FC.Gen samples public parameter pp .
- $\text{FC.Commit}(\text{pp}, \Pi) \rightarrow (\text{cm}, \text{aux})$: On input a public parameter pp and a message $\Pi \in \Sigma^\ell$, FC.Commit produces a commitment cm and the corresponding auxiliary state aux .
- $\langle \mathcal{P}_{\text{FC}}(\text{aux}, \alpha, \beta), \mathcal{V}_{\text{FC}}(\text{pp}, \text{cm}, \ell, \alpha, \beta) \rangle \rightarrow b$: The *opening phase* is an interactive protocol between the prover $\mathcal{P}_{\text{FC}}$ and verifier $\mathcal{V}_{\text{FC}}$. The prover receives as input an auxiliary state aux , a query set $\alpha \subseteq \mathbf{Q}_\ell$, and an answer set $\beta \in \mathbb{D}^\alpha$. The verifier receives as input the public parameter pp , a commitment cm , and a message length ℓ , along with α and β ; after interacting, it accepts or rejects, which is represented by the bit $b \in \{0, 1\}$.

Definition 3.16 (Perfect completeness). *FC has perfect completeness if for every security parameter $\lambda \in \mathbb{N}$, message length bound $\ell_{\max} \in \mathbb{N}$, query bound $s \in \mathbb{N}$, message $\Pi \in \Sigma^\ell$ where $\ell \leq \ell_{\max}$, and query set $\alpha \subseteq \mathbf{Q}_\ell$ with $|\alpha| \leq s$,*

$$\Pr \left[b = 1 \mid \begin{array}{l} \text{pp} \leftarrow \text{FC.Gen}(1^\lambda, \ell_{\max}) \\ (\text{cm}, \text{aux}) \leftarrow \text{FC.Commit}(\text{pp}, \Pi) \\ b \leftarrow \langle \mathcal{P}_{\text{FC}}(\text{aux}, \alpha, \alpha(\Pi)), \mathcal{V}_{\text{FC}}(\text{pp}, \text{cm}, \ell, \alpha, \alpha(\Pi)) \rangle \end{array} \right] = 1 .$$

Definition 3.17 (Function binding game). *FCGame is an interactive algorithm that works as follows.*

$\text{FCGame}(\text{pp})$:

1. For $i = 1, 2, \dots$, until the interaction terminates:

(a) Receive $(\text{cm}_i, \ell_i, \alpha_i, \beta_i)$.

(b) Run $\mathcal{V}_{\text{FC}}(\text{pp}, \text{cm}_i, \ell_i, \alpha_i, \beta_i)$ and let b_i denote its decision bit.

At the end of the interaction, for every commitment cm , $\ell \in \mathbb{N}$, and $s \in \mathbb{N}$, we define

$$\mathcal{C}_{\text{cm}, \ell, s} := \left\{ (\alpha, \beta_i(\alpha)) : \begin{array}{l} i \geq 1, \alpha \in \alpha_i, \\ \text{cm}_i = \text{cm}, \ell_i = \ell, \\ b_i = 1, |\alpha_i| \leq s \end{array} \right\} .$$

Definition 3.18 (Function binding). *FC has (strict-time) function binding error ϵ_{FC} if for every security parameter $\lambda \in \mathbb{N}$, message length bound $\ell_{\max} \in \mathbb{N}$, query bound $s \in \mathbb{N}$, auxiliary input distribution $\mathcal{D}_{\text{FC}}$, adversary time bound t_{FC} , and t_{FC} -time algorithm $\mathcal{A}_{\text{FC}}$,*

$$\Pr \left[\exists \text{cm}, \ell \leq \ell_{\max} : \text{Sol}(\mathcal{C}_{\text{cm}, \ell, s}) = \emptyset \mid \begin{array}{l} \text{pp} \leftarrow \text{FC.Gen}(1^\lambda, \ell_{\max}) \\ \text{ai}_{\text{FC}} \leftarrow \mathcal{D}_{\text{FC}} \\ \langle \mathcal{A}_{\text{FC}}(\text{pp}, \text{ai}_{\text{FC}}), \text{FCGame}(\text{pp}) \rangle \end{array} \right] \leq \epsilon_{\text{FC}}(\lambda, \ell_{\max}, s, t_{\text{FC}}) .$$

Definition 3.19 (Expected-time function binding). *FC has expected-time function binding error ϵ_{FC}^* if for every security parameter $\lambda \in \mathbb{N}$, message length bound $\ell_{\max} \in \mathbb{N}$, query bound $s \in \mathbb{N}$, auxiliary input distribution $\mathcal{D}_{\text{FC}}$, adversary expected-time bound t_{FC}^* , and algorithm $\mathcal{A}_{\text{FC}}$ with expected running time at most t_{FC}^* ,*

$$\Pr \left[\exists \text{cm}, \ell \leq \ell_{\max} : \text{Sol}(\mathcal{C}_{\text{cm}, \ell, s}) = \emptyset \mid \begin{array}{l} \text{pp} \leftarrow \text{FC.Gen}(1^\lambda, \ell_{\max}) \\ \text{ai}_{\text{FC}} \leftarrow \mathcal{D}_{\text{FC}} \\ \langle \mathcal{A}_{\text{FC}}(\text{pp}, \text{ai}_{\text{FC}}), \text{FCGame}(\text{pp}) \rangle \end{array} \right] \leq \epsilon_{\text{FC}}^*(\lambda, \ell_{\max}, s, t_{\text{FC}}^*) .$$

Remark 3.20 (Monotonicity of ϵ_{FC}). We assume hereafter that the function binding error ϵ_{FC} is monotone in each coordinate in the natural direction:

- $\epsilon_{\text{FC}}(\cdot, \ell, s, t_{\text{FC}})$ is non-increasing (larger security parameters decrease an adversary's success);
- $\epsilon_{\text{FC}}(\lambda, \cdot, s, t_{\text{FC}})$ is non-decreasing (larger messages increase an adversary's success);
- $\epsilon_{\text{FC}}(\lambda, \ell, \cdot, t_{\text{FC}})$ is non-decreasing (more constraints increase an adversary's success); and
- $\epsilon_{\text{FC}}(\lambda, \ell, s, \cdot)$ is non-decreasing (the success of an adversary increases with its computational power).

The last condition is trivially satisfied, while the first should also hold in any reasonable commitment scheme. The remaining two assumptions are natural; moreover, we may forgo them by replacing expressions of the form $\epsilon_{\text{FC}}(\lambda, \ell_{\max}, s_{\max}, t_{\text{FC}})$, where $\ell_{\max} = \max_i \{\ell_i\}$ and $s_{\max} = \max_j \{s_j\}$, with

$$\max_{i,j} \{\epsilon_{\text{FC}}(\lambda, \ell_i, s_j, t_{\text{FC}})\} .$$

Analogously, we assume the expected-time function binding error ϵ_{FC}^* is monotone.

3.4 Functional interactive oracle proofs

A *functional interactive oracle proof* (FIOP) is an interactive proof in which the verifier accesses each prover string through a prescribed query class. The definition below is adapted from the definition of IOPs with special queries in [BCG20].

Definition 3.21. A (public-coin) FIOP with query class $\mathbf{Q} = (\mathbf{Q}_\ell)_{\ell \in \mathbb{N}}$ over alphabet Σ and answer domain $\mathbb{D}$ for a relation R is a tuple $\text{FIOP} = (\mathbf{P}, \mathbf{V})$. The prover receives an instance-witness pair $(\mathbf{x}, \mathbf{w})$ and the verifier receives the instance $\mathbf{x}$. The protocol has k rounds. In round $i \in [k]$, the prover sends a proof string $\Pi_i \in \Sigma^{\ell_i}$, after which the verifier sends uniformly random coins $\rho_i \in \{0, 1\}^{r_i}$. After the interaction, the verifier accesses each Π_i only through queries in $\mathbf{Q}_{\ell_i}$: the answer to $\alpha \in \mathbf{Q}_{\ell_i}$ is $\alpha(\Pi_i) \in \mathbb{D}$. It then outputs a decision bit. We write $\mathbf{tr} = (\Pi_1, \rho_1, \dots, \Pi_k, \rho_k)$ for the interaction transcript.

Definition 3.22 (Perfect completeness). $\text{FIOP} = (\mathbf{P}, \mathbf{V})$ for a relation R has **perfect completeness** if, for every instance size bound $n \in \mathbb{N}$ and every instance-witness pair $(\mathbf{x}, \mathbf{w}) \in R$ with $|\mathbf{x}| \leq n$,

$$\Pr[\langle \mathbf{P}(\mathbf{x}, \mathbf{w}), \mathbf{V}(\mathbf{x}) \rangle = 1] = 1 .$$

An IOP uses the same definition for the class of point queries.

Definition 3.23 (Soundness). $\text{FIOP} = (\mathbf{P}, \mathbf{V})$ for a relation R has **soundness error** ϵ_{FIOP} if, for every instance size bound $n \in \mathbb{N}$, (unbounded) circuit $\tilde{\mathbf{P}}$, and auxiliary input distribution $\mathbf{D}$,

$$\Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge \mathbf{x} \notin L(R) \\ \wedge b = 1 \end{array} \middle| \begin{array}{l} \mathbf{ai} \leftarrow \mathbf{D} \\ (\mathbf{x}, \mathbf{aux}) \leftarrow \tilde{\mathbf{P}}(\mathbf{ai}) \\ b \leftarrow \langle \tilde{\mathbf{P}}(\mathbf{aux}), \mathbf{V}(\mathbf{x}) \rangle \end{array} \right] \leq \epsilon_{\text{FIOP}}(n) .$$

Definition 3.24 (Knowledge soundness). $\text{FIOP} = (\mathbf{P}, \mathbf{V})$ for a relation R has **(straightline) knowledge soundness error** κ_{FIOP} **with extraction time** $t_{\mathbf{E}}$ if there exists a deterministic algorithm $\mathbf{E}$ such that, for every instance size bound $n \in \mathbb{N}$, adversary $\tilde{\mathbf{P}}$, and auxiliary input distribution $\mathbf{D}$,

$$\Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge (\mathbf{x}, \mathbf{w}) \notin R \\ \wedge b = 1 \end{array} \middle| \begin{array}{l} \mathbf{ai} \leftarrow \mathbf{D} \\ (\mathbf{x}, \mathbf{aux}) \leftarrow \tilde{\mathbf{P}}(\mathbf{ai}) \\ b \xleftarrow{\mathbf{tr}} \langle \tilde{\mathbf{P}}(\mathbf{aux}), \mathbf{V}(\mathbf{x}) \rangle \\ \mathbf{w} \leftarrow \mathbf{E}(\mathbf{x}, \mathbf{tr}) \end{array} \right] \leq \kappa_{\text{FIOP}}(n) ;$$

moreover, $\mathbf{E}$ runs in time $t_{\mathbf{E}}(n)$.

Remark 3.25. An *interactive oracle proof* (IOP) is the special case of an FIOP where $\mathbf{Q} = \mathbf{Q}_{\text{Point}}$ is the class of point queries.

We consider the following efficiency measures for an FIOP:

- the *round complexity* k is the number of rounds (back-and-forth interactions) between the prover and verifier;
- the *alphabet* Σ is the set over which symbols of proof strings are defined;
- the *proof length* ℓ is the total number of alphabet symbols sent by the prover; moreover, ℓ_i is the length of the round- i proof string and $\ell_{\max} := \max_{i \in [k]} \ell_i$;
- the *query complexity* q is the total number of queries issued by the verifier; moreover, q_i is the number of queries issued to the round- i proof string and $q_{\max} := \max_{i \in [k]} q_i$;
- the *randomness complexity* r is the total number of random bits sent by the verifier; moreover, r_i is the number of random bits sent in round i .

Each efficiency measure may depend on the instance size $|\mathbb{x}|$.

3.5 Funky protocol

The Funky protocol combines an FIOP with a functional commitment scheme to construct an interactive argument. In more detail, consider the following ingredients:

- $\text{FIOP} = (\mathbf{P}, \mathbf{V})$, an FIOP for a relation R with alphabet Σ , answer domain $\mathbb{D}$, round complexity k , and query class $\mathbf{Q}$. We use explicit state-passing notation for $\mathbf{P}$, via the auxiliary state $\mathbf{aux}$.
- $\text{FC} = (\text{FC.Gen}, \text{FC.Commit}, \mathcal{P}_{\text{FC}}, \mathcal{V}_{\text{FC}})$, a functional commitment scheme for $\mathbf{Q}$.

The construction $\text{ARG} := \text{Funky}[\text{FIOP}, \text{FC}]$, specified below, is adapted from [CGKY26].

Construction 3.26 (Funky protocol). The argument generator $\mathcal{G}$ receives as input a security parameter $\lambda \in \mathbb{N}$ (in unary) and an instance size bound $n \in \mathbb{N}$, and works as follows.

- $\mathcal{G}(1^\lambda, n)$:
1. Sample $\text{pp} \leftarrow \text{FC.Gen}(1^\lambda, \ell_{\max}(n))$.
 2. Return pp .

The argument prover $\mathcal{P}$ receives as input the public parameter pp , an instance $\mathbb{x}$, and a witness $\mathbb{w}$, and the argument verifier $\mathcal{V}$ receives as input the public parameter pp and the instance $\mathbb{x}$. Then $\mathcal{P}$ and $\mathcal{V}$ interact as follows.

1. **Commitments and FIOP challenges.** For each round $i \in [k]$:

- (a) The prover:
- i. Computes the i -th FIOP string and auxiliary state

$$(\Pi_i, \mathbf{aux}_i) \leftarrow \begin{cases} \mathbf{P}(\mathbb{x}, \mathbb{w}) & \text{if } i = 1, \\ \mathbf{P}(\mathbf{aux}_{i-1}, \rho_{i-1}) & \text{if } i > 1. \end{cases}$$

- ii. Computes a functional commitment to the FIOP string:

$$(\text{cm}_i, \mathbf{aux}_i) \leftarrow \text{FC.Commit}(\text{pp}, \Pi_i) .$$

- iii. Sends cm_i to the verifier.

- (b) The verifier samples $\rho_i \leftarrow \{0, 1\}^{r_i}$ and sends ρ_i to the prover.

2. **FIOP answers.** The prover:

- (a) Runs the FIOP verifier

$$\mathbf{V}^{\Pi_1, \dots, \Pi_k}(\mathbb{X}; \rho_1, \dots, \rho_k)$$

to determine $\alpha_1, \dots, \alpha_k$, where $\alpha_i \subseteq \mathbf{Q}_{\ell_i}$ is the set of queries issued by $\mathbf{V}$ to Π_i .

- (b) Computes the evaluation $\beta_i(\alpha) := \alpha(\Pi_i)$ for every $i \in [k]$ and $\alpha \in \alpha_i$.
(c) Sends $(\tau_1, \dots, \tau_k)$ to the verifier, where τ_i denotes the query-answer tuple (α_i, β_i) .
3. **FC openings.** For each $i \in [k]$, sequentially, the prover and verifier engage in the FC opening phase for the i -th commitment:

$$\langle \mathcal{P}_{\text{FC}}(\text{aux}_i, \alpha_i, \beta_i), \mathcal{V}_{\text{FC}}(\text{pp}, \text{cm}_i, \ell_i, \alpha_i, \beta_i) \rangle \rightarrow b_i .$$

Let pf_i denote the resulting transcript.

4. **Decision.** The verifier accepts if and only if the following hold:

- (a) The FIOP verifier $\mathbf{V}(\mathbb{X}; \rho_1, \dots, \rho_k)$, receiving answers according to $\tau_1, \dots, \tau_k$, accepts;
(b) $b_i = 1$ for every $i \in [k]$.

The resulting interaction transcript is

$$\text{tr} = (\text{cm}_1, \rho_1, \dots, \text{cm}_k, \rho_k, (\tau_1, \dots, \tau_k), \text{pf}_1, \dots, \text{pf}_k) .$$

4 Barriers for knowledge soundness of Kilian's protocol

We first specify vector commitments, PCPs, and Kilian's protocol, together with the non-adaptive argument security notions used below. We then prove the two lower bounds on strict-time knowledge soundness error and the lower bound on the expected number of rewinds.

4.1 Vector commitments

A (static) *vector commitment scheme* [CF13] over alphabet Σ is a tuple $\text{VC} = (\text{Gen}, \text{Commit}, \text{Open}, \text{Check})$ with the following syntax.

- $\text{VC.Gen}(1^\lambda, \ell) \rightarrow \text{pp}$ samples public parameters for security parameter $\lambda \in \mathbb{N}$ and message size bound $\ell \in \mathbb{N}$.
- $\text{VC.Commit}(\text{pp}, m) \rightarrow (\text{cm}, \text{aux})$ commits to $m \in \Sigma^\ell$ and retains auxiliary state aux .
- $\text{VC.Open}(\text{pp}, \text{aux}, \mathcal{Q}) \rightarrow \text{pf}$ produces an opening proof for the restriction $m[\mathcal{Q}]$, where $\mathcal{Q} \subseteq [\ell]$.
- $\text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) \rightarrow \{0, 1\}$ checks whether pf validly opens cm to $\text{ans} \in \Sigma^{\mathcal{Q}}$ on $\mathcal{Q} \subseteq [\ell]$.

Definition 4.1 (Completeness). $\text{VC} = (\text{Gen}, \text{Commit}, \text{Open}, \text{Check})$ has **completeness error** γ_{VC} if for every security parameter $\lambda \in \mathbb{N}$, message length $\ell \in \mathbb{N}$, query bound $s \leq \ell$, message $m \in \Sigma^\ell$, and query set $\mathcal{Q} \subseteq [\ell]$ with $|\mathcal{Q}| \leq s$,

$$\Pr \left[\text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, m[\mathcal{Q}], \text{pf}) = 1 \mid \begin{array}{l} \text{pp} \leftarrow \text{VC.Gen}(1^\lambda, \ell) \\ (\text{cm}, \text{aux}) \leftarrow \text{VC.Commit}(\text{pp}, m) \\ \text{pf} \leftarrow \text{VC.Open}(\text{pp}, \text{aux}, \mathcal{Q}) \end{array} \right] \geq 1 - \gamma_{\text{VC}}(\lambda, \ell, s) .$$

It has **perfect completeness** if γ_{VC} is zero.

Definition 4.2 (Position binding). $\text{VC} = (\text{Gen}, \text{Commit}, \text{Open}, \text{Check})$ has **(strict-time) position binding error** ϵ_{VC} if for every security parameter $\lambda \in \mathbb{N}$, message length $\ell \in \mathbb{N}$, query set size $s \in \mathbb{N}$ with $s \leq \ell$, auxiliary input distribution $\mathcal{D}$, adversary time bound $t_{\text{VC}} \in \mathbb{N}$, and t_{VC} -time algorithm A_{VC} ,

$$\Pr \left[\begin{array}{l} |\mathcal{Q}|, |\mathcal{Q}'| \leq s \\ \wedge \exists i \in \mathcal{Q} \cap \mathcal{Q}' : \text{ans}[i] \neq \text{ans}'[i] \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}', \text{ans}', \text{pf}') = 1 \end{array} \mid \begin{array}{l} \text{pp} \leftarrow \text{VC.Gen}(1^\lambda, \ell) \\ \eta \leftarrow \mathcal{D} \\ (\text{cm}, \text{ans}, \text{ans}', \\ \mathcal{Q}, \mathcal{Q}', \text{pf}, \text{pf}') \leftarrow A_{\text{VC}}(\text{pp}, \eta) \end{array} \right] \leq \epsilon_{\text{VC}}(\lambda, \ell, s, t_{\text{VC}}) .$$

4.2 Probabilistically checkable proofs

A *probabilistically checkable proof* (PCP) is an information-theoretic proof system where a probabilistic verifier has oracle access to a proof string.

Definition 4.3 (Completeness). $\text{PCP} = (\mathbf{P}, \mathbf{V})$ for a relation R has **completeness error** γ_{PCP} if, for every instance size bound $n \in \mathbb{N}$ and every instance-witness pair $(\mathbf{x}, \mathbf{w}) \in R$ with $|\mathbf{x}| \leq n$,

$$\Pr \left[\mathbf{V}^\Pi(\mathbf{x}; \rho) = 1 \mid \begin{array}{l} \Pi \leftarrow \mathbf{P}(\mathbf{x}, \mathbf{w}) \\ \rho \leftarrow \{0, 1\}^r \end{array} \right] \geq 1 - \gamma_{\text{PCP}}(n) .$$

It has **perfect completeness** if γ_{PCP} is zero.

Definition 4.4 (Soundness). $\text{PCP} = (\mathbf{P}, \mathbf{V})$ for a relation R has **soundness error** ϵ_{PCP} if, for every instance size bound $n \in \mathbb{N}$, (unbounded) circuit $\tilde{\mathbf{P}}$, and auxiliary input distribution $\mathbf{D}$,

$$\Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge \mathbf{x} \notin L(R) \\ \wedge \mathbf{V}^{\tilde{\Pi}}(\mathbf{x}; \rho) = 1 \end{array} \mid \begin{array}{l} \mathbf{ai} \leftarrow \mathbf{D} \\ (\mathbf{x}, \tilde{\Pi}) \leftarrow \tilde{\mathbf{P}}(\mathbf{ai}) \\ \rho \leftarrow \{0, 1\}^r \end{array} \right] \leq \epsilon_{\text{PCP}}(n) .$$

Definition 4.5 (Knowledge soundness). $\text{PCP} = (\mathbf{P}, \mathbf{V})$ for a relation R has **knowledge soundness error** κ_{PCP} with **extraction time** $t_{\mathbf{E}}$ if there exists a probabilistic algorithm $\mathbf{E}$ such that, for every instance size bound $n \in \mathbb{N}$, adversary $\tilde{\mathbf{P}}$, and auxiliary input distribution $\mathbf{D}$,

$$\Pr \left[\begin{array}{l} |\mathbb{x}| \leq n \\ \wedge (\mathbb{x}, \mathbb{w}) \notin R \\ \wedge \mathbf{V}^{\tilde{\Pi}}(\mathbb{x}; \rho) = 1 \end{array} \middle| \begin{array}{l} \mathbf{ai} \leftarrow \mathbf{D} \\ (\mathbb{x}, \tilde{\Pi}) \leftarrow \tilde{\mathbf{P}}(\mathbf{ai}) \\ \rho \leftarrow \{0, 1\}^r \\ \mathbb{w} \leftarrow \mathbf{E}(\mathbb{x}, \tilde{\Pi}, \rho) \end{array} \right] \leq \kappa_{\text{PCP}}(n) ;$$

moreover, $\mathbf{E}$ runs in time $t_{\mathbf{E}}(n)$.

The PCP has *alphabet* Σ , *proof length* ℓ in alphabet symbols, *query complexity* $q \in [\ell]$, and *randomness complexity* r in bits. Each query is an index $q \in [\ell]$ answered by $\Pi[q] \in \Sigma$. These measures may depend on the instance $\mathbb{x}$, for example on $|\mathbb{x}|$.

Definition 4.6 (PCP verifier view). Given $\text{PCP} = (\mathbf{P}, \mathbf{V})$, a (possibly malicious) verifier $\tilde{\mathbf{V}}$, and an instance-witness pair $(\mathbb{x}, \mathbb{w})$, the view of $\tilde{\mathbf{V}}$, denoted by $\text{View}(\mathbf{P}, \tilde{\mathbf{V}}, \mathbb{x}, \mathbb{w})$, is the random variable $(\rho, \mathcal{Q}, \text{ans})$ where:

- ρ is the internal randomness of $\tilde{\mathbf{V}}$;
- $\mathcal{Q}$ and ans are the queries and answers when running $\tilde{\mathbf{V}}^{\Pi}$, where $\Pi \leftarrow \mathbf{P}(\mathbb{x}, \mathbb{w})$.

Definition 4.7 (Strong honest-verifier auxiliary-input zero-knowledge PCP). $\text{PCP} = (\mathbf{P}, \mathbf{V})$ is *strong-honest-verifier-auxiliary-input-zero-knowledge (SHVZK)* with error ζ_{PCP} if there exists a probabilistic algorithm $\mathbf{S}$ such that for every $(\mathbb{x}, \mathbb{w}) \in R$, $m^* \in \mathbb{N}$, size bound $t_{\tilde{\mathbf{V}}} \in \mathbb{N}$, and $t_{\tilde{\mathbf{V}}}$ -size malicious verifier $\tilde{\mathbf{V}}$ that accesses the honest PCP string only by requesting at most m^* honest PCP views from the oracle described below, the following two random variables are $\zeta_{\text{PCP}}(\mathbb{x}, m^*, t_{\tilde{\mathbf{V}}})$ -close in statistical distance:

$$\text{View}(\mathbf{P}, \tilde{\mathbf{V}}, \mathbb{x}, \mathbb{w}) \text{ and } \mathbf{S}((\text{Desc}(\tilde{\mathbf{V}}), 1^{t_{\tilde{\mathbf{V}}}}), \mathbb{x}) .$$

The oracle uses the single string $\Pi \leftarrow \mathbf{P}(\mathbb{x}, \mathbb{w})$ sampled in the real experiment. On input $\rho \in \{0, 1\}^r$, it runs $\mathbf{V}^{\Pi}(\mathbb{x}; \rho)$ and returns the honest PCP view $(\mathcal{Q}, \Pi[\mathcal{Q}])$, where $\mathcal{Q}$ consists of the locations queried in this execution. The malicious verifier may choose each ρ adaptively from previous replies, including repeating earlier choices. Moreover, $\mathbf{S}$ runs in time at most $\text{poly}(|\mathbb{x}|, t_{\tilde{\mathbf{V}}})$.

Remark 4.8. The SHVZK notion above is weaker than standard malicious-verifier auxiliary-input zero knowledge (MVZK) for PCPs, in which the malicious verifier may make a bounded number of arbitrary adaptive queries to individual locations of the PCP string. To see that MVZK implies SHVZK, note that each honest PCP view can be obtained by running $\mathbf{V}^{\Pi}(\mathbb{x}; \rho)$, which makes at most q individual queries. Hence a verifier that requests at most m^* honest PCP views can be implemented by a verifier that makes at most $m^* \cdot q$ individual queries. Applying the MVZK simulator to this implementation gives an SHVZK simulator.

4.3 Non-adaptive interactive arguments

For the lower bounds, we fix the argument instance before sampling public parameters and allow imperfect completeness. We use the completeness and non-adaptive knowledge soundness notions defined below. Provers in this section may be probabilistic; the determinism assumption in Remark 3.8 is only for the upper bounds.

Definition 4.9 (Completeness). $\text{ARG} = (\mathcal{G}, \mathcal{P}, \mathcal{V})$ for a relation R has **completeness error** γ_{ARG} if for every security parameter $\lambda \in \mathbb{N}$ and instance-witness pair $(\mathbb{x}, \mathbb{w}) \in R$,

$$\Pr \left[\langle \mathcal{P}(\text{pp}, \mathbb{x}, \mathbb{w}), \mathcal{V}(\text{pp}, \mathbb{x}) \rangle = 1 \mid \text{pp} \leftarrow \mathcal{G}(1^\lambda, \mathbb{x}) \right] \geq 1 - \gamma_{\text{ARG}}(\lambda, \mathbb{x}) .$$

It has **perfect completeness** if γ_{ARG} is zero.

Definition 4.10 (Non-adaptive knowledge soundness). $\text{ARG} = (\mathcal{G}, \mathcal{P}, \mathcal{V})$ for a relation R has **non-adaptive knowledge soundness error** κ_{ARG} **with extraction time** $t_{\mathcal{E}}$ if there exists a probabilistic algorithm $\mathcal{E}$ such that for every security parameter $\lambda \in \mathbb{N}$, instance $\mathbb{x}$, auxiliary input distribution $\mathcal{D}$, adversary time bound $t_{\text{ARG}} \in \mathbb{N}$, and t_{ARG} -time algorithm $\tilde{\mathcal{P}}$,

$$\Pr \left[(\mathbb{x}, \mathbb{w}) \in R \mid \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, \mathbb{x}) \\ \eta \leftarrow \mathcal{D} \\ \mathbb{w} \leftarrow \mathcal{E}^{\tilde{\mathcal{P}}(\text{pp}, \eta)}(\text{pp}, \mathbb{x}) \end{array} \right] \geq \Pr \left[b = 1 \mid \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, \mathbb{x}) \\ \eta \leftarrow \mathcal{D} \\ b \leftarrow \langle \tilde{\mathcal{P}}(\text{pp}, \eta), \mathcal{V}(\text{pp}, \mathbb{x}) \rangle \end{array} \right] - \kappa_{\text{ARG}}(\lambda, \mathbb{x}, t_{\text{ARG}}) ;$$

moreover, $\mathcal{E}$ runs in time $t_{\mathcal{E}}(\lambda, \mathbb{x}, t_{\text{ARG}})$.

4.4 Kilian's protocol

Kilian's protocol commits to a PCP string and opens the locations requested by the verifier.

Construction 4.11. Let $(\mathcal{G}, \mathcal{P}, \mathcal{V}) := \text{Kilian}[\text{PCP}, \text{VC}]$. For security parameter $\lambda \in \mathbb{N}$ and instance $\mathbb{x}$, the generator works as follows.

$\mathcal{G}(1^\lambda, \mathbb{x})$:

1. Sample $\text{pp}_{\text{VC}} \leftarrow \text{VC.Gen}(1^\lambda, \ell(|\mathbb{x}|))$ and output $\text{pp} := \text{pp}_{\text{VC}}$.

The prover receives $(\text{pp}, \mathbb{x}, \mathbb{w})$ and the verifier receives $(\text{pp}, \mathbb{x})$.

$\langle \mathcal{P}, \mathcal{V} \rangle$:

1. The prover computes $\Pi \leftarrow \mathbf{P}(\mathbb{x}, \mathbb{w})$ and $(\text{cm}, \text{aux}) \leftarrow \text{VC.Commit}(\text{pp}, \Pi)$, then sends cm .
2. The verifier samples and sends $\rho \leftarrow \{0, 1\}^r$.
3. The prover runs $\mathbf{V}^\Pi(\mathbb{x}; \rho)$ to obtain its query set $\mathcal{Q} \subseteq [\ell]$, computes $\text{pf} \leftarrow \text{VC.Open}(\text{pp}, \text{aux}, \mathcal{Q})$ and $\text{ans} := \Pi[\mathcal{Q}]$, and sends $(\mathcal{Q}, \text{ans}, \text{pf})$.
4. The verifier accepts if $\mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbb{x}; \rho) = 1$ and $\text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1$.

4.5 Toolbox for lower-bound analysis

Definition 4.12. Let $\text{ARG} = \text{Kilian}[\text{PCP}, \text{VC}]$ and let $\tilde{\mathcal{P}}$ be an argument prover. The algorithm $\tilde{\mathcal{P}}^\circ$ defined below answers each verifier challenge with the query set and answers of $\tilde{\mathcal{P}}$ together with the commitment verification bit output by VC.Check .

- $\tilde{\mathcal{P}}^\circ(\text{pp}, \eta)$:
 1. Compute $(\text{cm}, \text{aux}) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \eta)$ and retain private state $(\text{pp}, \text{cm}, \text{aux})$.
- $\tilde{\mathcal{P}}^\circ((\text{pp}, \text{cm}, \text{aux}), \rho)$:
 1. Compute $Y \leftarrow \tilde{\mathcal{P}}(\text{aux}, \rho)$ with newly sampled internal coins.
 2. If $Y = (\mathcal{Q}, \text{ans}, \text{pf})$, output $(\mathcal{Q}, \text{ans}, \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}))$; otherwise output $\perp$.

Definition 4.13 (VC-oblivious extractor). Let $\text{ARG} = \text{Kilian}[\text{PCP}, \text{VC}]$. A **VC-oblivious extractor** is an oracle algorithm $\mathcal{E}$ that receives an instance $\mathbb{x}$ and rewinding access to $\tilde{\mathcal{P}}^\circ$ for an argument prover $\tilde{\mathcal{P}}$. It induces an extractor for ARG in the sense of Definition 4.10: on input $(\text{pp}, \mathbb{x})$ and with access to $\tilde{\mathcal{P}}(\text{pp}, \eta)$, the induced extractor implements $\tilde{\mathcal{P}}^\circ(\text{pp}, \eta)$ using pp and outputs $\mathcal{E}^{\tilde{\mathcal{P}}^\circ(\text{pp}, \eta)}(\mathbb{x})$. We say that $\mathcal{E}$ achieves non-adaptive knowledge soundness error κ_{ARG} if its induced extractor does. In other words, VC-oblivious extractors induce a subclass of extractors for ARG that use no information about VC except for the commitment verification bits of $\tilde{\mathcal{P}}$'s responses.

Construction 4.14. Fix PCP for R , VC , $\text{ARG} = \text{Kilian}[\text{PCP}, \text{VC}]$, $(\mathbb{x}, \mathbb{w}) \in R$, and $\theta \in [0, 1]$. The prover $\tilde{\mathcal{P}}_{\theta, \mathbb{x}, \mathbb{w}}^*$ hardwires $(\mathbb{x}, \mathbb{w})$ and works as follows.

- $\tilde{\mathcal{P}}_{\theta, \mathbb{X}, \mathbb{W}}^*(\text{pp}, \perp)$: Compute $(\text{cm}, \text{aux}) \leftarrow \mathcal{P}(\text{pp}, \mathbb{X}, \mathbb{W})$; send cm and retain private state aux.
- $\tilde{\mathcal{P}}_{\theta, \mathbb{X}, \mathbb{W}}^*(\text{aux}, \rho)$: Flip a coin with heads probability θ . On heads output $\mathcal{P}(\text{aux}, \rho)$; otherwise output $\perp$.

The running time of $\tilde{\mathcal{P}}_{\theta, \mathbb{X}, \mathbb{W}}^*$ is at most

$$t_{\tilde{\mathcal{P}}^*} = O(t_{\mathbf{P}}(|\mathbb{X}|) + t_{\mathbf{V}}(|\mathbb{X}|) + t_{\text{VC.Commit}}(\lambda, \ell) + t_{\text{VC.Open}}(\lambda, \ell, \mathbf{q})) = \text{poly}(|\mathbb{X}|, \lambda) ,$$

where $t_{\mathbf{P}}$, $t_{\mathbf{V}}$, $t_{\text{VC.Commit}}$, and $t_{\text{VC.Open}}$ denote the running times of $\mathbf{P}$, $\mathbf{V}$, VC.Commit, and VC.Open, respectively.

Lemma 4.15. Fix PCP for R , VC, ARG = Kilian[PCP, VC], $(\mathbb{X}, \mathbb{W}) \in R$, $\theta \in [0, 1]$, and security parameter λ . Let γ_{PCP} and γ_{VC} denote the completeness errors of PCP and VC, respectively. Set

$$h := (1 - \gamma_{\text{PCP}}(|\mathbb{X}|))(1 - \gamma_{\text{VC}}(\lambda, \ell(|\mathbb{X}|), \mathbf{q}(|\mathbb{X}|))) .$$

Let $\mathcal{E}$ be a VC-oblivious extractor that achieves non-adaptive knowledge soundness error κ_{ARG} . Then

$$\Pr \left[(\mathbb{X}, \mathbb{W}') \in R \mid \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, \mathbb{X}) \\ \mathbb{W}' \leftarrow \mathcal{E}(\tilde{\mathcal{P}}_{\theta, \mathbb{X}, \mathbb{W}}^{\circ}(\text{pp}, \perp)(\mathbb{X})) \end{array} \right] \geq \theta h - \kappa_{\text{ARG}}(\lambda, \mathbb{X}, t_{\tilde{\mathcal{P}}^*}) . \quad (7)$$

Proof. Let a be the probability that the honest prover $\mathcal{P}(\text{pp}, \mathbb{X}, \mathbb{W})$ convinces $\mathcal{V}(\text{pp}, \mathbb{X})$, over $\text{pp} \leftarrow \mathcal{G}(1^\lambda, \mathbb{X})$ and the coins of both parties. We first show that $a \geq h$. Condition on the honest PCP string Π and verifier randomness ρ , which are independent of all VC coins and determine $\mathcal{Q}$. By Definition 4.1, the conditional probability that the VC check accepts is at least $1 - \gamma_{\text{VC}}(\lambda, \ell(|\mathbb{X}|), \mathbf{q}(|\mathbb{X}|))$. Multiplying this bound by the indicator that the PCP verifier accepts and averaging gives $a \geq h$, by Definition 4.3.

The prover $\tilde{\mathcal{P}}_{\theta, \mathbb{X}, \mathbb{W}}^*$ answers honestly with probability θ , independently of all coins of the honest execution, so it convinces $\mathcal{V}(\text{pp}, \mathbb{X})$ with probability exactly $\theta a \geq \theta h$. Applying Definition 4.10 with auxiliary input $\perp$ to the extractor induced by $\mathcal{E}$ gives Eq. 7. $\square$

4.6 Barrier from zero-knowledge PCPs

Section 2.1 explains the proof assuming perfect completeness of the PCP and the vector commitment. Here we give the full proof, which allows imperfect completeness of both.

Theorem 4.16. Let $\text{PCP} = (\mathbf{P}, \mathbf{V})$ be an SHVZK PCP for an NP relation R with completeness error γ_{PCP} and zero-knowledge error ζ_{PCP} . Let VC be a VC scheme with completeness error γ_{VC} . Set

$$h := (1 - \gamma_{\text{PCP}}(|\mathbb{X}|))(1 - \gamma_{\text{VC}}(\lambda, \ell(|\mathbb{X}|), \mathbf{q}(|\mathbb{X}|))) .$$

Let $\text{ARG} := \text{Kilian}[\text{PCP}, \text{VC}]$. Let $\theta \in [0, 1]$. Suppose there exists a VC-oblivious extractor $\mathcal{E}$ with strict running time $t_{\mathcal{E}}$ that rewinds the prover at most p times and achieves non-adaptive knowledge soundness error κ_{ARG} . Then for every $m^* \in \mathbb{N}$, there exists a probabilistic algorithm $\mathcal{A}$ that runs in time $\text{poly}(|\mathbb{X}|, m^*, t_{\mathcal{E}})$ and decides $L(R)$ with probability at least

$$\theta h - \zeta_{\text{PCP}}(\mathbb{X}, m^*, O(m^* \cdot t_{\mathcal{E}})) - (1 - (1 - \theta)^p)^{m^*+1} - \kappa_{\text{ARG}}(\lambda, \mathbb{X}, t_{\tilde{\mathcal{P}}^*}) .$$

Corollary 4.17. In the setting of Theorem 4.16, assume that PCP and VC have perfect completeness. Suppose that $t_{\mathcal{E}} = \text{poly}(|\mathbb{X}|)$ and $p = \text{poly}(|\mathbb{X}|) \geq 4$. Also suppose $\zeta_{\text{PCP}}(\mathbb{X}, m^*, \text{poly}(|\mathbb{X}|)) \leq \text{negl}(|\mathbb{X}|)$ for some $\omega(\log|\mathbb{X}|) \leq m^* \leq \text{poly}(|\mathbb{X}|)$. Then unless $L(R) \in \text{BPP}$, the bound

$$\kappa_{\text{ARG}}(\lambda, \mathbb{X}, t_{\tilde{\mathcal{P}}^*}) \leq \frac{1}{p}$$

cannot hold for every $\mathbb{X} \in L(R)$.

Proof. Perfect completeness gives $h = 1$. Set $\theta := 2/p \leq 1/2$. Since $\ln(1 - \theta) \geq -2\theta$ for $\theta \leq 1/2$, we have $(1 - \theta)^p \geq e^{-4}$ and hence

$$(1 - (1 - \theta)^p)^{m^*+1} \leq \exp(-e^{-4}(m^* + 1)) \leq \text{negl}(|x|) .$$

Suppose the displayed bound on κ_{ARG} holds for every $x \in L(R)$. Applying Theorem 4.16, together with the assumed negligible zero-knowledge error ζ_{PCP} , gives a strict polynomial-time algorithm $\mathcal{A}$ that always rejects instances outside $L(R)$ and, for every sufficiently large $x \in L(R)$, satisfies

$$\Pr[\mathcal{A}(x) = 1] \geq \frac{2}{p} - \frac{1}{p} - \text{negl}(|x|) \geq \frac{1}{2p} .$$

Run $\mathcal{A}$ independently $4p$ times and accept if at least one run accepts. This still takes strict polynomial time and the error is reduced to at most $(1 - 1/(2p))^{4p} \leq e^{-2} < 1/3$, which implies $L(R) \in \text{BPP}$. $\square$

Proof of Theorem 4.16. Fix $m^* \in \mathbb{N}$. Recall that the VC-oblivious extractor $\mathcal{E}$ is supposed to be given the rewinding access to $(\tilde{\mathcal{P}}_{\theta, x, w}^*)^\circ(\text{pp}, \eta)$: each time it queries with a challenge ρ , it obtains either $\perp$ or the response $(\mathcal{Q}, \text{ans}, b)$, where b is the commitment verification bit of the honest output $(\mathcal{Q}, \text{ans}, \text{pf})$ of $\tilde{\mathcal{P}}_{\theta, x, w}^*$ for challenge ρ . For a random tape r and a sequence $s = (s_1, s_2, \dots, s_{|s|})$ of responses, let $\mathcal{E}(x; r)[s]$ denote the run of $\mathcal{E}$ on input x with random tape r when the i -th query to the prover is answered by the i -th entry of s ; the run either finishes with output or stops when the query count exceeds s 's length.

We construct a malicious PCP verifier that tries to extract the witness by simulating $\mathcal{E}$:

$\tilde{\mathbf{V}}^\Pi(x)$:

1. Sample random tape r . Set $S_0 = \{()\}$ and $S_i = \emptyset$ for $i = 1, \dots, p$.
2. For each $i = 0, 1, \dots, p$ and each $s = (s_1, \dots, s_i) \in S_i$:
 - (a) Run $\mathcal{E}(x; r)[s]$. If it finishes and returns w with $(x, w) \in R$, return w .
 - (b) If it stops at query ρ (due to insufficient response entries in s):
 - i. Flip a coin with heads probability θ .
 - ii. On heads, induce $(\mathcal{Q}, \text{ans} := \Pi[\mathcal{Q}])$ from $\mathbf{V}^\Pi(x; \rho)$. Add $s \parallel (\mathcal{Q}, \text{ans}, 0)$ and $s \parallel (\mathcal{Q}, \text{ans}, 1)$ to S_{i+1} .
 - iii. On tails, add $s \parallel \perp$ to S_{i+1} .
3. Return $\perp$.

We claim the following about $\tilde{\mathbf{V}}^\Pi(x)$:

Claim 4.18. For any $(x, w) \in R$,

$$\Pr\left[(x, w') \in R \mid w' \leftarrow \tilde{\mathbf{V}}^\Pi(x)\right] \geq \theta h - \kappa_{\text{ARG}}(\lambda, x, t_{\tilde{\mathcal{P}}^*}) .$$

Moreover, let M be the random variable for the number of times $\tilde{\mathbf{V}}^\Pi(x)$ enters Step 2(b)ii, i.e., the number of honest PCP views it requests. Then

$$\Pr[M > m^*] \leq (1 - (1 - \theta)^p)^{m^*+1} .$$

Simulation using the zero-knowledge property of the PCP. We first assume the above claim and show how to use it to finish proving Theorem 4.16. Let $\tilde{\mathbf{V}}_{\leq m^*}^\Pi$ be $\tilde{\mathbf{V}}^\Pi$ truncated before it requests its $(m^* + 1)$ -th honest PCP view (outputting $\perp$ on truncation). From Claim 4.18 we know

$$\Pr\left[(x, w) \in R \mid w \leftarrow \tilde{\mathbf{V}}_{\leq m^*}^\Pi(x)\right] \geq \theta h - \kappa_{\text{ARG}}(\lambda, x, t_{\tilde{\mathcal{P}}^*}) - (1 - (1 - \theta)^p)^{m^*+1} .$$

Moreover, we can easily optimize the actual execution of $\tilde{\mathbf{V}}^\Pi(\mathbb{x})$ by saving the state of $\mathcal{E}(\mathbb{x}; r)[s]$ when it stops at Step 2b and resuming it when executing $\mathcal{E}(\mathbb{x}; r)[s \parallel (\mathcal{Q}, \text{ans}, 0/1)]$ or $\mathcal{E}(\mathbb{x}; r)[s \parallel \perp]$ afterwards. Then $\tilde{\mathbf{V}}_{\leq m^*}^\Pi(\mathbb{x})$ performs at most $m^* + 1$ full executions of $\mathcal{E}$ and thus its running time can be bounded by $t_{\tilde{\mathbf{V}}} := O(m^* \cdot t_{\mathcal{E}})$. Now we can simulate $\tilde{\mathbf{V}}_{\leq m^*}$ using the zero-knowledge property of the PCP. The output of $\tilde{\mathbf{V}}_{\leq m^*}$ can be expressed as a function f of its view. The decision algorithm $\mathcal{A}$ simply uses the PCP zero-knowledge simulator $\mathbf{S}$ to sample

$$V \leftarrow \mathbf{S}((\text{Desc}(\tilde{\mathbf{V}}_{\leq m^*}), 1^{t_{\tilde{\mathbf{V}}}}), \mathbb{x})$$

and outputs 1 exactly when $w' := f(V)$ satisfies $(\mathbb{x}, w') \in R$. By Definition 4.7, $\mathcal{A}$ runs in strict time $\text{poly}(|\mathbb{x}|, m^*, t_{\mathcal{E}})$. Moreover, for any $\mathbb{x} \in L(R)$,

$$\Pr[\mathcal{A}(\mathbb{x}) = 1] \geq \theta h - \kappa_{\text{ARG}}(\lambda, \mathbb{x}, t_{\tilde{\mathbf{P}}^*}) - (1 - (1 - \theta)^p)^{m^*+1} - \zeta_{\text{PCP}}(\mathbb{x}, m^*, O(m^* \cdot t_{\mathcal{E}})) .$$

For $\mathbb{x} \notin L(R)$, no valid witness can be found, so $\mathcal{A}(\mathbb{x})$ always outputs 0. This finishes the proof of Theorem 4.16.

Analysis of the malicious PCP verifier. Now we prove Claim 4.18. We first bound the success probability. The idea is that $\tilde{\mathbf{V}}$ can compute every part of a response from Π except the commitment verification bit, for which it tries both values. So if $\tilde{\mathcal{P}}_{\theta, \mathbb{x}, w}^*$ uses the same coins as $\tilde{\mathbf{V}}$, the responses in the real execution of $\mathcal{E}$ form one of the sequences that $\tilde{\mathbf{V}}$ runs. More formally, fix $(\mathbb{x}, w) \in R$ and the sampled random tape r . For each sequence s , let C_s denote an independent Bernoulli- θ coin. We couple the real extraction $\mathcal{E}(\mathbb{x}; r)^{(\tilde{\mathcal{P}}_{\theta, \mathbb{x}, w}^*)^\circ(\text{pp}, \perp)}$ and the simulation $\tilde{\mathbf{V}}^\Pi(\mathbb{x})$ using coins $\{C_s\}_s$. Consider the real extraction $\mathcal{E}(\mathbb{x}; r)^{(\tilde{\mathcal{P}}_{\theta, \mathbb{x}, w}^*)^\circ(\text{pp}, \perp)}$ with $\text{pp} \leftarrow \mathcal{G}(1^\lambda, \mathbb{x})$, in which the honest prover inside $\tilde{\mathcal{P}}_{\theta, \mathbb{x}, w}^*$ computes the same Π and, whenever the first i responses form the sequence s , $\tilde{\mathcal{P}}_{\theta, \mathbb{x}, w}^*$ uses C_s as its coin for the $(i + 1)$ -th query. Let s^* be the random variable for the full sequence of responses in this real extraction. In $\tilde{\mathbf{V}}^\Pi(\mathbb{x})$'s execution, when $\mathcal{E}(\mathbb{x}; r)[s]$ for some sequence s stops at a challenge and enters the coin flip in Step 2(b)i, we let $\tilde{\mathbf{V}}^\Pi(\mathbb{x})$ also use C_s as its coin. Starting from $() \in S_0$, induction shows that $\tilde{\mathbf{V}}$ either outputs a valid witness earlier or eventually runs $\mathcal{E}(\mathbb{x}; r)[s^*]$, which is the real extraction. Hence, by Lemma 4.15,

$$\begin{aligned} \Pr[(\mathbb{x}, \tilde{\mathbf{V}}^\Pi(\mathbb{x})) \in R] &\geq \Pr[(\mathbb{x}, \mathcal{E}(\mathbb{x}; r)[s^*]) \in R] \\ &= \Pr\left[(\mathbb{x}, w') \in R \mid \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, \mathbb{x}) \\ w' \leftarrow \mathcal{E}^{(\tilde{\mathcal{P}}_{\theta, \mathbb{x}, w}^*)^\circ(\text{pp}, \perp)}(\mathbb{x}) \end{array}\right] \\ &\geq \theta h - \kappa_{\text{ARG}}(\lambda, \mathbb{x}, t_{\tilde{\mathbf{P}}^*}) . \end{aligned}$$

Next we prove the tail bound on the number M of honest PCP views requested by $\tilde{\mathbf{V}}^\Pi(\mathbb{x})$. The case where $\theta = 1$ is trivial, and we assume $\theta < 1$ below. We also assume without loss of generality that $\mathcal{E}$ always makes exactly p queries in total. Let Z_t be the final size of S_t . A sequence s of length t yields one sequence $s \parallel \perp$ of length $t + 1$ on tails and two sequences $s \parallel (\mathcal{Q}, \text{ans}, 0/1)$ on heads, so

$$Z_0 = 1, \quad Z_{t+1} = Z_t + X_t, \quad X_t \mid Z_t \sim \text{Bin}(Z_t, \theta) . \quad (8)$$

Each head increases $Z_{t+1} - Z_t$ by one and requests one honest PCP view, so $M = Z_p - 1$. To calculate a tail bound for Z_p , we couple $\{Z_t\}_{t=0, \dots, p}$ with a process $\{Y_t\}_{t=0, \dots, p}$ whose distribution we can compute exactly. Let $\{G_{t,i}\}_{0 \leq t < p, i \geq 1}$ be independent random variables with $\Pr[G_{t,i} = n] = (1 - \theta)\theta^{n-1}$ for $n = 1, 2, \dots$, and let $B_{t,i} := \mathbb{1}[G_{t,i} \geq 2]$. Then $\Pr[B_{t,i} = 1] = \theta$ and $1 + B_{t,i} \leq G_{t,i}$. We generate both processes from these variables:

$$Z_0 = Y_0 = 1, \quad Z_{t+1} = \sum_{i=1}^{Z_t} (1 + B_{t,i}), \quad Y_{t+1} = \sum_{i=1}^{Y_t} G_{t,i} .$$

Given Z_t , the sum $\sum_{i=1}^{Z_t} B_{t,i}$ has distribution $\text{Bin}(Z_t, \theta)$, so these Z_t follow the definition in Eq. 8. If $Z_t \leq Y_t$, then $Z_{t+1} \leq \sum_{i=1}^{Z_t} G_{t,i} \leq Y_{t+1}$, so $Z_t \leq Y_t$ for every t by induction. To find the distribution of Y_t , we use the fact

that a sum of a geometric number of independent geometric random variables is again geometric. Precisely, let $K, G_1, G_2, \dots$ be independent random variables with $\Pr[K = n] = q(1-q)^{n-1}$ and $\Pr[G_i = n] = r(1-r)^{n-1}$ for $n \geq 1$, where $q, r \in (0, 1]$. Then

$$\Pr\left[\sum_{i=1}^K G_i = n\right] = qr(1-qr)^{n-1} \quad (9)$$

for $n \geq 1$.⁴ Now $Y_0 = 1$, and Y_t depends only on the variables $G_{t',i}$ with $t' < t$, so it is independent of $G_{t,1}, G_{t,2}, \dots$. Substituting $q = (1-\theta)^t$ and $r = 1-\theta$, induction on t gives

$$\Pr[Y_t = n] = (1-\theta)^t (1 - (1-\theta)^t)^{n-1}$$

for $n \geq 1$ and $t \geq 1$. Hence

$$\Pr[M > m^*] = \Pr[Z_p \geq m^* + 2] \leq \Pr[Y_p \geq m^* + 2] = (1 - (1-\theta)^p)^{m^*+1}.$$

□

4.7 Barrier from low-query PCPs

Theorem 4.19. *Let $\text{PCP} = (\mathbf{P}, \mathbf{V})$ be a PCP for an NP relation R with completeness error γ_{PCP} , alphabet Σ , and query complexity q such that $|\Sigma|^q = \text{poly}(|\mathbb{X}|)$. Let VC be a VC scheme with completeness error γ_{VC} . Set*

$$h := (1 - \gamma_{\text{PCP}}(|\mathbb{X}|)) (1 - \gamma_{\text{VC}}(\lambda, \ell(|\mathbb{X}|), q(|\mathbb{X}|))) .$$

Let $\text{ARG} := \text{Kilian}[\text{PCP}, \text{VC}]$. Let $\theta \in [0, 1]$. Suppose there exists a VC-oblivious extractor $\mathcal{E}$ with strict running time $t_{\mathcal{E}}$ that rewinds the prover at most p times and achieves non-adaptive knowledge soundness error κ_{ARG} . Then there exists a probabilistic algorithm $\mathcal{A}$ that runs in time $\text{poly}(t_{\mathcal{E}}, |\mathbb{X}|)$ and decides $L(R)$ with probability at least

$$\theta h - p^2 \theta^2 - \kappa_{\text{ARG}}(\lambda, \mathbb{X}, t_{\bar{P}^*}) .$$

Corollary 4.20. *In the setting of Theorem 4.19, assume that PCP and VC have perfect completeness. Also suppose $p = \text{poly}(|\mathbb{X}|)$ and $t_{\mathcal{E}} = \text{poly}(|\mathbb{X}|)$. Then unless $L(R) \in \text{BPP}$, the bound*

$$\kappa_{\text{ARG}}(\lambda, \mathbb{X}, t_{\bar{P}^*}) \leq \frac{1}{8p^2}$$

cannot hold for every $\mathbb{X} \in L(R)$.

Proof. Perfect completeness gives $h = 1$. Set $\theta := 1/(2p^2)$. Suppose the displayed bound on κ_{ARG} holds for every $\mathbb{X} \in L(R)$. Applying Theorem 4.19 gives a strict polynomial-time algorithm $\mathcal{A}$ that always rejects instances outside $L(R)$ and, for every $\mathbb{X} \in L(R)$, satisfies

$$\Pr[\mathcal{A}(\mathbb{X}) = 1] \geq \frac{1}{2p^2} - p^2 \cdot \frac{1}{4p^4} - \frac{1}{8p^2} = \frac{1}{8p^2} .$$

Run $\mathcal{A}$ independently $16p^2$ times and accept if at least one run accepts. This still takes strict polynomial time and the error is reduced to at most $(1 - 1/(8p^2))^{16p^2} \leq e^{-2} < 1/3$, which implies $L(R) \in \text{BPP}$. □

⁴To see this, consider two coins A and B with heads probabilities r and q , respectively. Repeat the following round until B lands heads: flip A until it lands heads, then flip B once. The numbers of flips of A in successive rounds and the number of rounds are independent and distributed as $G_1, G_2, \dots$ and K , so the total number of flips of A is distributed as $\sum_{i=1}^K G_i$. Equivalently, flip B together with every flip of A and ignore B whenever A lands tails. The process then stops at the first flip where both coins land heads, which happens independently at each flip with probability qr . This gives Eq. 9.

Proof of Theorem 4.19. In this proof we use the notation $\mathcal{E}(\mathbb{x}; r)[s]$ from the proof of Theorem 4.16. The idea is that, except with probability at most $p^2\theta^2$, the prover $\tilde{\mathcal{P}}_{\theta, \mathbb{x}, \mathbb{w}}^*$ answers at most one query honestly and responds $\perp$ otherwise. The only response other than $\perp$ is then (Q, ans, b) , where (Q, ans) are the queries and answers of $\mathbf{V}^\Pi(\mathbb{x}; \rho)$ for the queried challenge ρ and b is a commitment verification bit. Since there are at most $|\Sigma|^q = \text{poly}(|\Sigma|)$ possible answers, we can try all such responses without access to Π .

Formally, we construct the following algorithm. Without loss of generality, we assume the extractor $\mathcal{E}$ always makes exactly p queries.

$\mathcal{A}(\mathbb{x})$:

1. Sample random tape r .
2. Run $\mathcal{E}(\mathbb{x}; r)[\perp^p]$, and for each $i \in [p]$, let ρ_i be its i -th query. If it returns $\mathbb{w}$ with $(\mathbb{x}, \mathbb{w}) \in R$, return 1.
3. For each $i \in [p]$:
 - (a) Enumerate each (Q, ans) obtained by emulating $\mathbf{V}(\mathbb{x}; \rho_i)$ and assigning a symbol in Σ to each queried position in every possible way.
 - (b) For each such (Q, ans) and each $b \in \{0, 1\}$, run $\mathcal{E}(\mathbb{x}; r)[\perp^{i-1} \parallel (Q, \text{ans}, b) \parallel \perp^{p-i}]$. If it returns $\mathbb{w}$ with $(\mathbb{x}, \mathbb{w}) \in R$, return 1.
4. Return 0.

Since the query complexity of PCP is q , there are at most $2|\Sigma|^q$ tuples (Q, ans, b) for each i . So $\mathcal{A}$ runs $\mathcal{E}$ at most $1 + 2p|\Sigma|^q$ times and $\mathbf{V}$ at most $p|\Sigma|^q$ times. Without loss of generality, $p \leq t_\mathcal{E}$. Also, we have $|\Sigma|^q = \text{poly}(|\Sigma|)$. Thus, $\mathcal{A}$ runs in time $\text{poly}(t_\mathcal{E}, |\Sigma|)$.

Now we lower bound the success probability of $\mathcal{A}$. Fix $(\mathbb{x}, \mathbb{w}) \in R$, and let r be the random tape sampled by $\mathcal{A}$. Let $C_1, \dots, C_p$ be independent Bernoulli- θ coins, and consider the real extraction $\mathcal{E}(\mathbb{x}; r)^{(\tilde{\mathcal{P}}_{\theta, \mathbb{x}, \mathbb{w}}^*)^o(\text{pp}, \perp)}$ with $\text{pp} \leftarrow \mathcal{G}(1^\lambda, \mathbb{x})$, in which $\tilde{\mathcal{P}}_{\theta, \mathbb{x}, \mathbb{w}}^*$ uses C_i as its coin for the i -th query. Let $\mathbf{s}^*$ be the random variable for the full sequence of responses in this real extraction, so $\mathcal{E}(\mathbb{x}; r)[\mathbf{s}^*]$ is the real extraction. Let **OneQuery** be the event that at most one of $C_1, \dots, C_p$ equals 1. Under **OneQuery**, at most one entry of $\mathbf{s}^*$ is not $\perp$. If there is none, then $\mathbf{s}^* = \perp^p$. If the i -th entry is not $\perp$, then the first $i-1$ entries are $\perp$, so the i -th query of the real extraction is ρ_i , and $\mathbf{s}^* = \perp^{i-1} \parallel (Q, \text{ans}, b) \parallel \perp^{p-i}$, where (Q, ans) are the queries and answers of $\mathbf{V}^\Pi(\mathbb{x}; \rho_i)$ and thus tried by $\mathcal{A}$. In both cases, $\mathcal{A}$ runs $\mathcal{E}(\mathbb{x}; r)[\mathbf{s}^*]$. Hence, by Lemma 4.15 and a union bound over pairs of coins,

$$\begin{aligned} \Pr[\mathcal{A}(\mathbb{x}) = 1] &\geq \Pr[(\mathbb{x}, \mathcal{E}(\mathbb{x}; r)[\mathbf{s}^*]) \in R \wedge \text{OneQuery}] \\ &\geq \theta h - \kappa_{\text{ARG}}(\lambda, \mathbb{x}, t_{\tilde{\mathcal{P}}^*}) - \Pr[\exists i \neq j \in [p] : C_i = C_j = 1] \\ &\geq \theta h - \kappa_{\text{ARG}}(\lambda, \mathbb{x}, t_{\tilde{\mathcal{P}}^*}) - p^2\theta^2. \end{aligned}$$

For $\mathbb{x} \notin L(R)$, no valid witness can be found, so $\mathcal{A}(\mathbb{x})$ always outputs 0. $\square$

Remark 4.21 (Deterministic provers). The prover $\tilde{\mathcal{P}}_{\theta, \mathbb{x}, \mathbb{w}}^*$ of Construction 4.14 samples a new coin on every query, which Definition 4.12 permits. Our barriers also hold for deterministic provers, as in Remark 3.8. Let $\mathcal{D}$ output the honest prover's coins and the key of a p -wise independent hash function H with $\Pr[H(\rho) = 1] = \theta$ for every ρ , and let the prover answer ρ honestly exactly when $H(\rho) = 1$. An extractor that makes at most p queries cannot distinguish H from a random function with independent Bernoulli- θ values, so distinct queries are answered independently with probability θ and a repeated query returns the same response. Lemma 4.15 and the proof of Theorem 4.19 apply unchanged. In the proof of Theorem 4.16, the malicious PCP verifier reuses the stored coin and response whenever a query repeats along one response sequence, and the same tail bound holds.

Remark 4.22 (Comparison with prior extractors). Our barriers do not contradict the knowledge soundness analyses of Kilian's protocol in [CMSZ21; LM19]. The extractor in [CMSZ21] runs in expected polynomial time and receives as input $1^{1/\epsilon}$, where ϵ is a lower bound on the probability that the prover convinces the verifier, and it outputs a witness with probability $\Omega(\epsilon - \kappa_{\text{PCP}})$. In [LM19], for every polynomial-time prover there is a polynomial-time extractor that outputs a witness with non-negligible probability whenever the prover convinces

the verifier with non-negligible probability, and the polynomial bounding the running time of the extractor depends on the prover. Neither is a single strict polynomial-time extractor whose running time and number of rewinds are independent of the prover and of its success probability, which is what Theorems 4.16 and 4.19 rule out.

4.8 Barrier for expected-time rewinding from low-query PCPs

Let 3SAT be the relation whose instances x are 3-CNF formulas with m variables and $\Theta(m)$ clauses and whose witnesses are satisfying assignments. Such an instance has encoding length $|x| = \Theta(m \log m)$.

Definition 4.23 (Randomized Exponential Time Hypothesis). *The randomized exponential time hypothesis (rETH) states that there exists a constant $\epsilon_{\text{rETH}} > 0$ such that no randomized algorithm decides $L(3\text{SAT})$, with error at most $1/3$, in time at most*

$$2^{\epsilon_{\text{rETH}} \cdot m} \cdot \text{poly}(m) .$$

rETH is usually stated for 3-CNF formulas with m variables and any number of clauses. By the sparsification lemma [IPZ01], that statement implies the one above.

Theorem 4.24. *Let $\text{PCP} = (\mathbf{P}, \mathbf{V})$ be a PCP for 3SAT with completeness error γ_{PCP} , alphabet Σ , and query complexity q such that $|\Sigma|^q = \text{poly}(m)$. Let VC be a VC scheme with completeness error γ_{VC} . Set*

$$h := (1 - \gamma_{\text{PCP}}(|x|)) (1 - \gamma_{\text{VC}}(\lambda, \ell(|x|), q(|x|))) .$$

Let $\text{ARG} := \text{Kilian}[\text{PCP}, \text{VC}]$ and set $\theta := 1$. Suppose there exists a VC-oblivious extractor $\mathcal{E}$ that, against $\tilde{\mathcal{P}}_{\theta, x, w}^*$, has expected running time $t_{\mathcal{E}}^*$, rewinds the prover at most p^* times in expectation, and achieves non-adaptive knowledge soundness error κ_{ARG} satisfying

$$\kappa_{\text{ARG}}(\lambda, x, t_{\tilde{\mathcal{P}}^*}) \leq h - c$$

for some constant $c > 0$ and every $(x, w) \in 3\text{SAT}$. Then there exists a probabilistic algorithm $\mathcal{A}$ that decides $L(3\text{SAT})$ with probability at least $c/2$ and runs in strict time at most

$$\text{poly}(t_{\mathcal{E}}^*, m) \cdot 2^{O(p^* \cdot q \cdot \log_2 |\Sigma|)} .$$

Corollary 4.25. *In the setting of Theorem 4.24, suppose that PCP and VC have perfect completeness, that $t_{\mathcal{E}}^* = \text{poly}(m)$, and that $\kappa_{\text{ARG}}(\lambda(m), x, t_{\tilde{\mathcal{P}}^*}) \leq 1 - c$ for some constant $c > 0$ and every $(x, w) \in 3\text{SAT}$. Then under rETH, the bound*

$$p^* = o\left(\frac{m}{q \log_2 |\Sigma|}\right)$$

cannot hold.

Proof. Suppose the displayed bound on p^* holds. Applying the theorem gives an algorithm $\mathcal{A}$ that always rejects any $x \notin L(3\text{SAT})$ and, for every $x \in L(3\text{SAT})$, satisfies $\Pr[\mathcal{A}(x) = 1] \geq c/2$. Its strict running time is at most

$$\text{poly}(t_{\mathcal{E}}^*, m) \cdot 2^{O(p^* \cdot q \cdot \log_2 |\Sigma|)} = \text{poly}(m) \cdot 2^{o(m)} = 2^{o(m)} .$$

Run $\mathcal{A}$ independently $\lceil 4/c \rceil$ times and accept if at least one run accepts. This reduces the error to at most $(1 - c/2)^{\lceil 4/c \rceil} \leq e^{-2} < 1/3$. Since c is constant, the running time remains $2^{o(m)}$, contradicting rETH. $\square$

Remark 4.26. The extractor of Corollary 6.2 rewinds the prover $O(\ell / \log(1 + \alpha))$ times in expectation, where $\alpha > 0$ is the constant such that its knowledge soundness error is $(1 + \alpha) \cdot \kappa_{\text{PCP}}$ plus the binding error. For constant α , this is $O(\ell)$ rewinds, and comparing with Corollary 4.25, the remaining multiplicative gap is

$$O\left(\frac{\ell \cdot q \log_2 |\Sigma|}{m}\right) .$$

For a PCP with $q = O(1)$, $|\Sigma| = O(1)$, and $\ell = m \cdot \text{polylog}(m)$ [Din07], the gap is $\text{polylog}(\ell)$. Whether a PCP with $q = O(1)$, $|\Sigma| = O(1)$, and $\ell = \Theta(m)$ exists is an open problem; for such a PCP, the bound of Corollary 4.25 would be $\Omega(\ell)$ and the extractor of Corollary 6.2 would be optimal.

Proof of Theorem 4.24. In this proof we use the notation $\mathcal{E}(\mathbb{x}; r)[s]$ from the proofs of Theorem 4.16 and Theorem 4.19. Let $k := \lceil 4p^*/c \rceil$ and $\tau := \lceil 4t_\varepsilon^*/c \rceil$. We construct the following algorithm that simulates $\mathcal{E}$ by enumerating all possible responses while truncating the number of queries at k and the running time at τ .

$\mathcal{A}(\mathbb{x})$:

1. Sample random tape r . Set $S_0 = \{()\}$ and $S_i = \emptyset$ for $i = 1, \dots, k$.
2. For each $i = 0, 1, \dots, k$ and each $s \in S_i$:
 - (a) Run $\mathcal{E}(\mathbb{x}; r)[s]$ for at most τ steps. If it returns w with $(\mathbb{x}, w) \in 3\text{SAT}$, return 1.
 - (b) If it stops at query ρ and $i < k$: for each (Q, ans) obtained by emulating $\mathbf{V}(\mathbb{x}; \rho)$ and assigning a symbol in Σ to each queried position in every possible way, and each $b \in \{0, 1\}$, add $s \parallel (Q, \text{ans}, b)$ to S_{i+1} .
3. Return 0.

Each S_i contains at most $(2|\Sigma|^q)^i$ sequences, so $\mathcal{A}$ processes at most $2(2|\Sigma|^q)^k$ sequences, and for each it runs $\mathcal{E}$ for at most τ steps and $\mathbf{V}$ at most $|\Sigma|^q$ times. Since c is a constant, $\tau = O(t_\varepsilon^*)$ and $k \leq 4p^*/c + 1$. Since $|\Sigma| \geq 2$, we have $\log_2(2|\Sigma|^q) \leq 2q \log_2 |\Sigma|$. Together with $|\Sigma|^q = \text{poly}(m)$, $\mathcal{A}$ runs in time at most

$$\text{poly}(t_\varepsilon^*, m) \cdot (2|\Sigma|^q)^k \leq \text{poly}(t_\varepsilon^*, m) \cdot 2^{8 \cdot p^* \cdot q \cdot \log_2 |\Sigma| / c}.$$

Now we lower bound the success probability of $\mathcal{A}$. Fix $(\mathbb{x}, w) \in 3\text{SAT}$, and let r be the random tape sampled by $\mathcal{A}$. Consider the real extraction $\mathcal{E}(\mathbb{x}; r)^{(\tilde{\mathcal{P}}_{\theta, \mathbb{x}, w}^*)^\circ(\text{pp}, \perp)}$ with $\text{pp} \leftarrow \mathcal{G}(1^\lambda, \mathbb{x})$. Let s^* be the random variable for the full sequence of responses in this real extraction, so $\mathcal{E}(\mathbb{x}; r)[s^*]$ is the real extraction, and let T^* be its number of steps. Suppose that $|s^*| \leq k$ and $T^* \leq \tau$. If the prefix s of s^* of length $i < |s^*|$ lies in S_i , then $\mathcal{E}(\mathbb{x}; r)[s]$ stops within τ steps at the same challenge ρ as the real extraction. Since $\theta = 1$, the next response is (Q, ans, b) , where (Q, ans) are the queries and answers of $\mathbf{V}^\Pi(\mathbb{x}; \rho)$; since $i < k$, $\mathcal{A}$ adds the prefix of length $i + 1$ to S_{i+1} . Starting from $() \in S_0$, induction shows that $\mathcal{A}$ either finds a valid witness and returns 1 earlier or runs $\mathcal{E}(\mathbb{x}; r)[s^*]$, which is the real extraction and finishes within τ steps. Hence, by Lemma 4.15 with $\theta = 1$, the assumption on κ_{ARG} , and Markov's inequality,

$$\begin{aligned} \Pr[\mathcal{A}(\mathbb{x}) = 1] &\geq \Pr[(\mathbb{x}, \mathcal{E}(\mathbb{x}; r)[s^*]) \in 3\text{SAT} \wedge |s^*| \leq k \wedge T^* \leq \tau] \\ &\geq h - \kappa_{\text{ARG}}(\lambda, \mathbb{x}, t_{\tilde{\mathcal{P}}^*}) - \Pr[|s^*| > k] - \Pr[T^* > \tau] \\ &\geq c - \frac{c}{4} - \frac{c}{4} = \frac{c}{2}. \end{aligned}$$

For $\mathbb{x} \notin L(3\text{SAT})$, no valid witness can be found, so $\mathcal{A}(\mathbb{x})$ always outputs 0. $\square$

5 Expected-time soundness of the Funky protocol

Theorem 5.1. *Let FIOp be a public-coin FIOp with query class $\mathbf{Q}$ for a relation R , with round complexity k , proof length ℓ , query complexity q , and soundness error ϵ_{FIOp} ; suppose that $\mathbf{Q}$ has finite diameter $d_{\mathbf{Q}}$ and admits a constraint solver with operation time $t_{\mathbf{Q}}$. Let FC be a public-coin interactive functional commitment scheme for $\mathbf{Q}$ over Σ with expected-time function binding error ϵ_{FC}^* .*

Then $\text{ARG} := \text{Funky}[\text{FIOp}, \text{FC}]$ (Construction 3.26) is a public-coin interactive argument for R whose expected-time soundness error ϵ_{ARG}^* satisfies the following for every $t_{\text{ARG}}^* \geq t_{\mathbf{V}} + \max_{i \in [k]} q_i \cdot t_{\mathbf{Q}}(\ell_i)$:

$$\epsilon_{\text{ARG}}^*(\lambda, n, t_{\text{ARG}}^*) \leq \epsilon_{\text{FIOp}}(n) + \epsilon_{\text{FC}}^*(\lambda, \ell_{\max}, q_{\max}, t_{\text{FC}}^*) ,$$

where $t_{\text{FC}}^* = O(N \cdot t_{\text{ARG}}^*)$ for $N := 1 + \sum_{i=1}^k d_{\mathbf{Q}}(\ell_i)$. If the verifier of FIOp always issues a q_i^* -independent query set (Definition 5.2) to the i -th proof string, then the above holds for $N := 1 + \sum_{i=1}^k (d_{\mathbf{Q}}(\ell_i) - \max\{q_i^* - 1, 0\})$.

Theorem 5.1 establishes slightly tighter bounds for FIOps with *independent queries*, which we now define.

Definition 5.2 (Independent queries). *A query-answer tuple τ is m -independent if $|\mathcal{C}_{\tau}| \geq m$. A query set α is m -independent if (α, β) is m -independent for every $\beta \in \mathbb{D}^{\alpha}$.*

Throughout, let $\text{FIOp} = (\mathbf{P}, \mathbf{V})$ and $\text{ARG} = (\mathcal{G}, \mathcal{P}, \mathcal{V})$. Let $\text{Solver}_{\mathbf{Q}}$ be a constraint solver for $\mathbf{Q}$ with operation time $t_{\mathbf{Q}}$. We prove Theorem 5.1's second bound, which implies the first bound by setting all $q_i^* = 0$.

5.1 Adversary construction

We construct an adversary for FC which will be used to analyze the expected-time soundness of ARG.

Construction 5.3 (Sampler). For a round $i \in [k]$, the sampler $\mathcal{S}_i$ works as follows.

$\mathcal{S}_i^{\tilde{\mathcal{P}}(\text{aux})}(\text{pp}, \mathbb{X}, \rho_1, \dots, \rho_{i-1}, \text{cm}_i)$:

1. Loop:

- (a) For each $j \in [i, k]$, sample verifier randomness $\rho_j \leftarrow \{0, 1\}^{r_j}$.
- (b) Sample an execution of the prover $\tilde{\mathcal{P}}(\text{aux})$ as follows. In the FIOp phase, use randomness $\rho_1, \dots, \rho_k$; the resulting (partial) interaction transcript is of the form

$$(\text{cm}_1, \rho_1, \dots, \text{cm}_k, \rho_k, (\tau_1, \dots, \tau_k)) ,$$

where $\tau_j = (\alpha_j, \beta_j)$ is the query-answer tuple for the j -th commitment. In the opening phase, for each $j \in [k]$, interact with FCGame: send $(\text{cm}_j, \ell_j, \alpha_j, \beta_j)$ and then relay messages between $\tilde{\mathcal{P}}(\text{aux})$ and FCGame.

- (c) If the above execution is accepting, return τ_i .

Construction 5.4 (Reductor). The reductor $\mathcal{R}$ works as follows.

$\mathcal{R}^{\tilde{\mathcal{P}}(\text{aux})}(\text{pp}, \mathbb{X}, \text{tr})$:

- 1. Parse the transcript $\text{tr} = (\text{cm}_1, \rho_1, \dots, \text{cm}_k, \rho_k, (\tau_1, \dots, \tau_k), \text{pf}_1, \dots, \text{pf}_k)$.
- 2. For each $i \in [k]$:
 - (a) Initialize an empty list $T_i := \emptyset$.
 - (b) Loop while $\mathcal{C}_{T_i}$ does not decide τ_i :
 - i. Sample $\tau'_i \leftarrow \mathcal{S}_i^{\tilde{\mathcal{P}}(\text{aux})}(\text{pp}, \mathbb{X}, \rho_1, \dots, \rho_{i-1}, \text{cm}_i)$.
 - ii. Append $T_i := T_i \parallel \tau'_i$.

Construction 5.5 (FC adversary). For an argument prover $\tilde{\mathcal{P}}$, the adversary $\mathcal{A}_{\text{FC}}^{\tilde{\mathcal{P}}}$ interacts with $\text{FCGame}(\text{pp})$ as follows.

$\mathcal{A}_{\text{FC}}^{\tilde{\mathcal{P}}}(\text{pp}, \eta)$:

1. Compute $(\mathbb{x}, \text{aux}) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \eta)$. If $|\mathbb{x}| > n$, halt.
2. Sample an execution of the prover $\tilde{\mathcal{P}}(\text{aux})$ as follows. In the FIOP phase, sample verifier randomness $\rho_1, \dots, \rho_k$. In the opening phase, interact with FCGame as in Step 1b. Let b and tr be the resulting decision bit and interaction transcript, respectively.
3. If $b = 1$, run $\mathcal{R}^{\tilde{\mathcal{P}}(\text{aux})}(\text{pp}, \mathbb{x}, \text{tr})$.

5.2 Security analysis

Throughout, fix an argument prover $\tilde{\mathcal{P}}$ and auxiliary input distribution $\mathcal{D}$. We consider the experiment

$$E := \left\{ \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \eta \leftarrow \mathcal{D} \\ (\mathbb{x}, \text{aux}) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \eta) \\ b \stackrel{\text{tr}}{\leftarrow} \langle \tilde{\mathcal{P}}(\text{aux}), \mathcal{V}(\text{pp}, \mathbb{x}) \rangle \end{array} \right\},$$

where $\text{tr} = (\text{cm}_1, \rho_1, \dots, \text{cm}_k, \rho_k, (\tau_1, \dots, \tau_k), \text{pf}_1, \dots, \text{pf}_k)$. Let t_{ARG}^* denote the expected running time of $\tilde{\mathcal{P}}$ above; throughout, we assume the lower bound on t_{ARG}^* from Theorem 5.1.

We also consider the experiment

$$E_{\text{FC}} := \left\{ \begin{array}{l} \text{pp} \leftarrow \text{FC.Gen}(1^\lambda, \ell_{\max}) \\ \eta \leftarrow \mathcal{D} \\ \langle \mathcal{A}_{\text{FC}}^{\tilde{\mathcal{P}}}(\text{pp}, \eta), \text{FCGame}(\text{pp}) \rangle \end{array} \right\}.$$

When $\mathcal{A}_{\text{FC}}^{\tilde{\mathcal{P}}}$ does not run $\mathcal{R}$, the lists $T_1, \dots, T_k$ are taken to be empty.

Definition 5.6. For a public parameter pp and auxiliary input η , we define:

$$\begin{aligned} t_{\text{pp}, \eta}^*(\rho_{<i}) &:= \mathbb{E}_E[\text{running time of } \tilde{\mathcal{P}} \mid \text{pp}, \eta, \rho_{<i}] \\ \epsilon_{\text{pp}, \eta}(\rho_{<i}) &:= \Pr_E[b = 1 \mid \text{pp}, \eta, \rho_{<i}], \end{aligned}$$

where $\rho_{<i}$ denotes $(\rho_1, \dots, \rho_{i-1})$. When $i = 1$, we write $t_{\text{pp}, \eta}^*$ and $\epsilon_{\text{pp}, \eta}$ for brevity.

5.2.1 Running time

We analyze the expected running time of the adversary $\mathcal{A}_{\text{FC}}$ (Construction 5.5).

Lemma 5.7. Let $\mathbf{Q}$ be a query class with finite diameter $d_{\mathbf{Q}}$, and let D be a distribution over query-answer tuples. Let $X_0, X_1, X_2, \dots$ be an infinite sequence of samples independently drawn from D , and, for $n \geq 0$, let $\mathcal{C}_n := \mathcal{C}_{(X_1, \dots, X_n)}$ (with $\mathcal{C}_0 = \emptyset$). Then the random variable

$$N := \min\{n \geq 0 : \mathcal{C}_n \text{ decides } X_0\}$$

satisfies $\mathbb{E}[N] \leq d_{\mathbf{Q}}$. If every query-answer tuple in the support of D is m -independent, then $\mathbb{E}[N] \leq d_{\mathbf{Q}} - m + 1$.

Proof. For $n \geq 0$, we say that X_{n+1} is *new* if $\mathcal{C}_n$ does not decide X_{n+1} . Observe that the event $N > n$ occurs if and only if some constraint of X_0 is undecided by $\mathcal{C}_n$. Since the samples are i.i.d. and $\mathcal{C}_n$ is a function of $(X_1, \dots, X_n)$ only,

$$\Pr[N > n] = \Pr[\mathcal{C}_n \text{ does not decide } X_0] = \Pr[\mathcal{C}_n \text{ does not decide } X_{n+1}] = \Pr[X_{n+1} \text{ new}] .$$

It follows that

$$\mathbb{E}[N] = \sum_{n \geq 0} \Pr[N > n] = \sum_{n \geq 0} \Pr[X_{n+1} \text{ new}] = \mathbb{E}[\text{number of new samples}] ,$$

where the final equality is by linearity of expectation. We conclude by observing that there are at most d_Q new samples across the sequence (each new sample appends at least one constraint, and at most d_Q constraints can be appended in total). If every query-answer tuple in the support of D is m -independent, then there are at most $d_Q - m + 1$ new samples (since X_1 always contributes at least m constraints). $\square$

Lemma 5.8 (Expected running time). *The expected running time of $\mathcal{A}_{\text{FC}}^{\tilde{\mathcal{P}}}$ in E_{FC} is at most $O(N \cdot t_{\text{ARG}}^*)$ for $N := 1 + \sum_{i=1}^k (d_Q(\ell_i) - \max\{q_i^* - 1, 0\})$.*

Proof. For simplicity, we focus on the time spent executing $\tilde{\mathcal{P}}$; the remaining time can be bounded with some additional accounting. Fix a public parameter pp and auxiliary input η , which together determine the instance $\mathbb{X}$ and prover state aux . Without loss of generality, assume $|\mathbb{X}| \leq n$.

Fix a round $i \in [k]$, and let N_i denote the number of times the sampler $\mathcal{S}_i$ is invoked. Fix verifier randomness $\rho_{<i} = (\rho_1, \dots, \rho_{i-1})$, which determines the commitment cm_i . If $\epsilon_{\text{pp}, \eta}(\rho_{<i}) = 0$, then $N_i = 0$; hence assume it is positive. The expected running time of an invocation

$$\mathcal{S}_i^{\tilde{\mathcal{P}}(\text{aux})}(\text{pp}, \mathbb{X}, \rho_1, \dots, \rho_{i-1}, \text{cm}_i)$$

is $t_{\text{pp}, \eta}^*(\rho_{<i}) / \epsilon_{\text{pp}, \eta}(\rho_{<i})$. By Wald's identity,

$$\mathbb{E}[\text{time spent in } \mathcal{S}_i \mid \text{pp}, \eta, \rho_{<i}] = \mathbb{E}[N_i \mid \text{pp}, \eta, \rho_{<i}] \cdot \frac{t_{\text{pp}, \eta}^*(\rho_{<i})}{\epsilon_{\text{pp}, \eta}(\rho_{<i})} .$$

Observe that $N_i = 0$ unless $b = 1$. Thus, by total probability,

$$\mathbb{E}[N_i \mid \text{pp}, \eta, \rho_{<i}] \leq \mathbb{E}[N_i \mid \text{pp}, \eta, \rho_{<i}, b = 1] \cdot \epsilon_{\text{pp}, \eta}(\rho_{<i}) .$$

Conditioned on $b = 1$, the initial query-answer tuple τ_i follows the output distribution of $\mathcal{S}_i$. By Lemma 5.7 (where D is the output distribution of $\mathcal{S}_i$),

$$\mathbb{E}[N_i \mid \text{pp}, \eta, \rho_{<i}, b = 1] \leq d_Q(\ell_i) - \max\{q_i^* - 1, 0\} .$$

Putting everything together and averaging over $\rho_{<i}$, we get

$$\mathbb{E}[\text{time spent in } \mathcal{S}_i \mid \text{pp}, \eta] \leq (d_Q(\ell_i) - \max\{q_i^* - 1, 0\}) \cdot t_{\text{pp}, \eta}^* ;$$

summing over $i \in [k]$ and averaging over pp and η bounds the expected running time of the reducer $\mathcal{R}$. Accounting also for the initial execution of $\tilde{\mathcal{P}}$ gives the lemma. $\square$

5.2.2 Success probability

Definition 5.9. For a (possibly infinite) list of query-answer tuples T , let $\tilde{\Pi}_T$ denote the lexicographically first element of $\text{Sol}(\mathcal{C}_T)$. (When $T = (\tau_n)_{n \in \mathbb{N}}$ is infinite, $\mathcal{C}_T$ is defined to be the eventual value of $\mathcal{C}_{(\tau_1, \dots, \tau_n)}$, which exists since d_Q is finite.)

Definition 5.10. For a public parameter pp and auxiliary input η , we define $\mathbf{X}_{\text{pp}, \eta}(\rho_{<i})$ to be an infinite sequence of samples independently drawn from the output distribution of

$$\mathcal{S}_i^{\tilde{\mathcal{P}}(\text{aux})}(\text{pp}, \mathbb{X}, \rho_1, \dots, \rho_{i-1}, \text{cm}_i) ,$$

where $\rho_{<i} = (\rho_1, \dots, \rho_{i-1})$ for some round $i \in [k]$. (If the sampler never halts, we take $\mathbf{X}_{\text{pp}, \eta}(\rho_{<i})$ to be empty.) When pp and η are clear from context, we write $\mathbf{X}(\rho_{<i})$.

Proof of Theorem 5.1. Fix an argument prover $\tilde{\mathcal{P}}$ and auxiliary input distribution $\mathcal{D}$. Let E^+ be the extension of E that additionally samples

$$b' \leftarrow \mathbf{V}^{\tilde{\Pi}_{\mathbf{X}(\rho_{<1})}, \dots, \tilde{\Pi}_{\mathbf{X}(\rho_{<k})}}(\mathbb{X}; \rho_1, \dots, \rho_k) .$$

By total probability,

$$\Pr_E \left[\begin{array}{l} |\mathbb{X}| \leq n \\ \wedge \mathbb{X} \notin L(R) \\ \wedge b = 1 \end{array} \right] \leq \underbrace{\Pr_{E^+} \left[\begin{array}{l} |\mathbb{X}| \leq n \\ \wedge \mathbb{X} \notin L(R) \\ \wedge b' = 1 \end{array} \right]}_A + \underbrace{\Pr_{E^+} \left[\begin{array}{l} |\mathbb{X}| \leq n \\ \wedge b = 1 \\ \wedge b' = 0 \end{array} \right]}_B .$$

Bounding A. Define the FIOP prover $\tilde{\mathbf{P}}$ which, on input pp and η , computes $(\mathbb{X}, \text{aux}) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \eta)$ and outputs the instance $\mathbb{X}$. In the i -th round, having received $\rho_{<i}$, it samples and sends a proof string according to the distribution $\tilde{\Pi}_{\mathbf{X}(\rho_{<i})}$. Observe that A is equal to the success probability of $\tilde{\mathbf{P}}$ (with an auxiliary input distribution $\mathcal{D}$ that samples public parameter $\text{pp} \leftarrow \mathcal{G}(1^\lambda, n)$, auxiliary input $\eta \leftarrow \mathcal{D}$, and additional random coins). By soundness of FIOP (Definition 3.23),

$$A \leq \epsilon_{\text{FIOP}}(n) .$$

Bounding B. Define the random variables $T'_1, \dots, T'_k$ in E^+ as follows. If $|\mathbb{X}| \leq n$ and $b = 1$, then T'_i is the shortest prefix of $\mathbf{X}(\rho_{<i})$ such that $\mathcal{C}_{T'_i}$ decides τ_i ; such a prefix exists by Lemma 5.7. Otherwise, T'_i is an empty list. Observe that the following distributions are identical:

$$\{\text{tr}, T'_1, \dots, T'_k \mid E^+, |\mathbb{X}| \leq n, b = 1\} \sim \{\text{tr}, T_1, \dots, T_k \mid E_{\text{FC}}, |\mathbb{X}| \leq n, b = 1\} .$$

In E^+ , suppose that $b = 1$ and $b' = 0$. Then in some round $i \in [k]$, it must be that $\tilde{\Pi}_{\mathbf{X}(\rho_{<i})}$ disagrees with τ_i , and hence $\mathcal{C}_{T'_i}$ refutes τ_i . Thus,

$$B \leq \Pr_{E^+} \left[\begin{array}{l} |\mathbb{X}| \leq n \\ \wedge b = 1 \\ \wedge \exists i, \mathcal{C}_{T'_i} \text{ refutes } \tau_i \end{array} \right] = \Pr_{E_{\text{FC}}} \left[\begin{array}{l} |\mathbb{X}| \leq n \\ \wedge b = 1 \\ \wedge \exists i, \mathcal{C}_{T_i} \text{ refutes } \tau_i \end{array} \right] ,$$

In E_{FC} , observe that $\mathcal{C}_{T_i}$ refuting τ_i implies $\text{Sol}(\mathcal{C}_{\text{cm}_i, \ell_i, \mathbf{q}_i}) = \emptyset$ (see Definition 3.17). By expected-time function binding of FC (Definition 3.19) and Lemma 5.8,

$$B \leq \epsilon_{\text{FC}}^*(\lambda, \ell_{\max}, \mathbf{q}_{\max}, t_{\text{FC}}^*) ,$$

where t_{FC}^* is as stated in the theorem. □

6 Expected-time knowledge soundness of the Funky protocol

Theorem 6.1. *Let FIOP be a public-coin FIOP with query class $\mathbf{Q}$ for a relation R , with round complexity k , proof length ℓ , query complexity q , and knowledge soundness error κ_{FIOP} ; suppose that $\mathbf{Q}$ has coverage factor $C_{\mathbf{Q}}$ (Definition 6.11) and admits a constraint solver with operation time $t_{\mathbf{Q}}$. Let FC be a public-coin interactive functional commitment scheme for $\mathbf{Q}$ over Σ with expected-time function binding error ϵ_{FC}^* .*

Then $\text{ARG} := \text{Funky}[\text{FIOP}, \text{FC}]$ (Construction 3.26) is a public-coin interactive argument for R whose expected-time knowledge soundness error κ_{ARG}^* satisfies the following for every $t_{\text{ARG}}^* \geq t_{\nu} + \sum_{i=1}^k q_i \cdot t_{\mathbf{Q}}(\ell_i)$ and every $M_1, \dots, M_k \in \mathbb{N}$:

$$\kappa_{\text{ARG}}^*(\lambda, n, t_{\text{ARG}}^*) \leq \left(\prod_{i=1}^k C_{\mathbf{Q}}(\ell_i, M_i) \right) \cdot \kappa_{\text{FIOP}}(n) + \epsilon_{\text{FC}}^*(\lambda, \ell_{\max}, q_{\max}, t_{\text{FC}}^*) ,$$

where $t_{\text{FC}}^* = O(N \cdot t_{\text{ARG}}^*)$ for

$$N := 1 + \sum_{i=1}^k (M_i + C_{\mathbf{Q}}(\ell_i, M_i)) \prod_{j=1}^{i-1} C_{\mathbf{Q}}(\ell_j, M_j) .$$

Moreover, the expected running time of the argument extractor is $O(N \cdot t_{\text{ARG}}^* + t_{\mathbf{E}})$.

Theorem 6.1's bounds depend on the *coverage factor* of $\mathbf{Q}$, a quantity which we define in the security analysis. In Section A we bound coverage factors for various query classes; these yield the following corollaries.

Corollary 6.2. *In the setting of Theorem 6.1, suppose that $\mathbf{Q} = \mathbf{Q}_{\text{Point}}$ is the class of point queries. Then ARG is a public-coin interactive argument for R whose expected-time knowledge soundness error κ_{ARG}^* satisfies the following for every $t_{\text{ARG}}^* \geq t_{\nu} + \sum_{i=1}^k q_i \cdot t_{\mathbf{Q}_{\text{Point}}}(\ell_i)$ and constant $\alpha > 0$:*

$$\kappa_{\text{ARG}}^*(\lambda, n, t_{\text{ARG}}^*) \leq (1 + \alpha) \cdot \kappa_{\text{FIOP}}(n) + \epsilon_{\text{FC}}^*(\lambda, \ell_{\max}, q_{\max}, t_{\text{FC}}^*) ,$$

where $t_{\text{FC}}^* = O(N \cdot t_{\text{ARG}}^*)$ for

$$N := 1 + \frac{\alpha \rho}{\rho - 1} + \sum_{i=1}^k \max \left\{ 1, \left\lceil \frac{k \cdot \ell_i}{\log(1 + \alpha)} \right\rceil \right\} \cdot \rho^{i-1} , \quad \rho := (1 + \alpha)^{1/k} .$$

Moreover, the expected running time of the argument extractor is $O(N \cdot t_{\text{ARG}}^* + t_{\mathbf{E}})$.

Corollary 6.3. *In the setting of Theorem 6.1, suppose that $\mathbf{Q}$ has finite positive diameter $d_{\mathbf{Q}}$. Then ARG is a public-coin interactive argument for R whose expected-time knowledge soundness error κ_{ARG}^* satisfies the following for every $t_{\text{ARG}}^* \geq t_{\nu} + \sum_{i=1}^k q_i \cdot t_{\mathbf{Q}}(\ell_i)$ and constant $\alpha > 0$:*

$$\kappa_{\text{ARG}}^*(\lambda, n, t_{\text{ARG}}^*) \leq (1 + \alpha) \cdot \kappa_{\text{FIOP}}(n) + \epsilon_{\text{FC}}^*(\lambda, \ell_{\max}, q_{\max}, t_{\text{FC}}^*) ,$$

where $t_{\text{FC}}^* = O_{\alpha}(k \cdot d \cdot t_{\text{ARG}}^*)$ for $d := \sum_{i=1}^k d_{\mathbf{Q}}(\ell_i)$. Moreover, the expected running time of the argument extractor is $O_{\alpha}(k \cdot d \cdot t_{\text{ARG}}^* + t_{\mathbf{E}})$.

Remark 6.4. Making explicit the dependence on α , Corollary 6.3's bound takes $t_{\text{FC}}^* = O\left(\frac{1+\alpha}{\log(1+\alpha)} \cdot k \cdot d \cdot t_{\text{ARG}}^*\right)$, and the expected running time of the argument extractor is $O\left(\frac{1+\alpha}{\log(1+\alpha)} \cdot k \cdot d \cdot t_{\text{ARG}}^* + t_{\mathbf{E}}\right)$.

Throughout, let $\text{FIOP} = (\mathbf{P}, \mathbf{V})$ and $\text{ARG} = (\mathcal{G}, \mathcal{P}, \mathcal{V})$. Let $\mathbf{E}$ be an extractor for FIOP with knowledge soundness error κ_{FIOP} . Let $\text{Solver}_{\mathbf{Q}}$ be a constraint solver for $\mathbf{Q}$ with operation time $t_{\mathbf{Q}}$.

6.1 Extractor construction

We construct an extractor $\mathcal{E}$ and adversary $\mathcal{A}_{\text{FC}}$ which will be used to analyze the expected-time knowledge soundness of ARG.

Construction 6.5 (Sampler). For a round $i \in [k]$, the sampler $\mathcal{S}_i$ works as follows.

$\mathcal{S}_i^{\tilde{\mathcal{P}}(\text{aux})}(\text{pp}, \mathbb{X}, T_1, \rho_1, \dots, T_{i-1}, \rho_{i-1})$:

1. Loop:

- (a) For each $j \in [i, k]$, sample verifier randomness $\rho_j \leftarrow \{0, 1\}^{r_j}$.
- (b) Sample an execution of the prover $\tilde{\mathcal{P}}(\text{aux})$ as follows. In the FIOP phase, use randomness $\rho_1, \dots, \rho_k$; the resulting (partial) interaction transcript is of the form

$$(\text{cm}_1, \rho_1, \dots, \text{cm}_k, \rho_k, (\tau_1, \dots, \tau_k)) ,$$

where $\tau_j = (\alpha_j, \beta_j)$ is the query-answer tuple for the j -th commitment. In the opening phase, for each $j \in [k]$, interact with FCGame: send $(\text{cm}_j, \ell_j, \alpha_j, \beta_j)$ and then relay messages between $\tilde{\mathcal{P}}(\text{aux})$ and FCGame. Let tr be the resulting interaction transcript.

- (c) If the above execution is accepting and $\mathcal{C}_{T_j}$ decides τ_j for every $j < i$, return tr .

Construction 6.6 (Terminal sampler). For a round $i \in [k]$, the sampler $\mathcal{S}_{i,\text{fin}}$ works as follows.

$\mathcal{S}_{i,\text{fin}}^{\tilde{\mathcal{P}}(\text{aux})}(\text{pp}, \mathbb{X}, T_1, \rho_1, \dots, T_{i-1}, \rho_{i-1}, T_i)$:

1. Loop:

- (a) Sample $\text{tr} \leftarrow \mathcal{S}_i^{\tilde{\mathcal{P}}(\text{aux})}(\text{pp}, \mathbb{X}, T_1, \rho_1, \dots, T_{i-1}, \rho_{i-1})$.
- (b) Parse the i -th-round randomness ρ_i and query-answer tuple τ_i from tr .
- (c) If $\mathcal{C}_{T_i}$ decides τ_i , return tr .

Construction 6.7 (Reductor). Fix repetition parameters $M_1, \dots, M_k \in \mathbb{N}$. The reductor $\mathcal{R}$ works as follows.

$\mathcal{R}^{\tilde{\mathcal{P}}(\text{aux})}(\text{pp}, \mathbb{X})$:

1. For each $i \in [k]$:

- (a) Initialize an empty list $T_i := \emptyset$.
- (b) Repeat M_i times:
 - i. Sample $\text{tr} \leftarrow \mathcal{S}_i^{\tilde{\mathcal{P}}(\text{aux})}(\text{pp}, \mathbb{X}, T_1, \rho_1, \dots, T_{i-1}, \rho_{i-1})$ and parse τ_i from tr .
 - ii. Append $T_i := T_i \parallel \tau_i$.
- (c) Sample $\text{tr}_i \leftarrow \mathcal{S}_{i,\text{fin}}^{\tilde{\mathcal{P}}(\text{aux})}(\text{pp}, \mathbb{X}, T_1, \rho_1, \dots, T_{i-1}, \rho_{i-1}, T_i)$ and parse ρ_i from tr_i .

2. Return $(T_1, \rho_1, \dots, T_k, \rho_k)$.

Construction 6.8 (Argument extractor). The extractor $\mathcal{E}$ works as follows.

$\mathcal{E}^{\tilde{\mathcal{P}}(\text{aux})}(\text{pp}, \mathbb{X}, \text{tr}, b)$:

- 1. If $|\mathbb{X}| > n$ or $b = 0$, output $\perp$.
- 2. Sample $(T_1, \rho_1, \dots, T_k, \rho_k) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}(\text{aux})}(\text{pp}, \mathbb{X})$.
- 3. Return $\mathbf{E}(\mathbb{X}, \Pi_{T_1}, \rho_1, \dots, \Pi_{T_k}, \rho_k)$, where each $\Pi_{T_i} \in \text{Sol}(\mathcal{C}_{T_i})$ is output by Solver_Q.

Construction 6.9 (FC adversary). For an argument prover $\tilde{\mathcal{P}}$, the adversary $\mathcal{A}_{\text{FC}}^{\tilde{\mathcal{P}}}$ interacts with FCGame(pp) as follows.

$\mathcal{A}_{\text{FC}}^{\tilde{\mathcal{P}}}(\text{pp}, \eta)$:

- 1. Compute $(\mathbb{X}, \text{aux}) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \eta)$. If $|\mathbb{X}| > n$, halt.

2. Sample an execution of the prover $\tilde{\mathcal{P}}(\text{aux})$ as follows. In the FIOP phase, sample verifier randomness $\rho_1, \dots, \rho_k$. In the opening phase, interact with FCGame as in Step 1b. Let b and tr be the resulting decision bit and interaction transcript, respectively.
3. If $b = 1$, run $\mathcal{R}^{\tilde{\mathcal{P}}(\text{aux})}(\text{pp}, \mathbb{X})$.

6.2 Security analysis

Throughout, fix an argument prover $\tilde{\mathcal{P}}$ and auxiliary input distribution $\mathcal{D}$. We consider the experiment

$$E := \left\{ \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \eta \leftarrow \mathcal{D} \\ (\mathbb{X}, \text{aux}) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \eta) \\ b \stackrel{\text{tr}}{\leftarrow} \langle \tilde{\mathcal{P}}(\text{aux}), \mathcal{V}(\text{pp}, \mathbb{X}) \rangle \end{array} \right\} ,$$

where $\text{tr} = (\text{cm}_1, \rho_1, \dots, \text{cm}_k, \rho_k, (\tau_1, \dots, \tau_k), \text{pf}_1, \dots, \text{pf}_k)$. Let t_{ARG}^* denote the expected running time of $\tilde{\mathcal{P}}$ above; throughout, we assume the lower bound on t_{ARG}^* from Theorem 6.1.

We also consider the experiment

$$E_{\text{FC}} := \left\{ \begin{array}{l} \text{pp} \leftarrow \text{FC.Gen}(1^\lambda, \ell_{\max}) \\ \eta \leftarrow \mathcal{D} \\ \langle \mathcal{A}_{\text{FC}}^{\tilde{\mathcal{P}}}(\text{pp}, \eta), \text{FCGame}(\text{pp}) \rangle \end{array} \right\} .$$

When $\mathcal{A}_{\text{FC}}^{\tilde{\mathcal{P}}}$ does not run $\mathcal{R}$, the lists $T_1, \dots, T_k$ are taken to be empty.

Definition 6.10. For a public parameter pp and auxiliary input η , define:

$$\begin{aligned} t_{\text{pp}, \eta}^*(\rho_{<i}) &:= \mathbb{E}_E \left[\text{running time of } \tilde{\mathcal{P}} \mid \text{pp}, \eta, \rho_{<i} \right] \\ \epsilon_{\text{pp}, \eta}(\mathbf{T}_{<i}, \rho_{<i}) &:= \Pr_E \left[b = 1 \wedge \forall j < i, \mathcal{C}_{T_j} \text{ decides } \tau_j \mid \text{pp}, \eta, \rho_{<i} \right] , \\ \delta_{\text{pp}, \eta}(\mathbf{T}_{<i+1}, \rho_{<i}) &:= \Pr_E \left[\mathcal{C}_{T_i} \text{ decides } \tau_i \mid \text{pp}, \eta, \rho_{<i}, b = 1 \wedge \forall j < i, \mathcal{C}_{T_j} \text{ decides } \tau_j \right] , \end{aligned}$$

where $\mathbf{T}_{<i}$ and $\rho_{<i}$ denote $(T_1, \dots, T_{i-1})$ and $(\rho_1, \dots, \rho_{i-1})$, respectively. We additionally define $\epsilon_{\text{pp}, \eta}(\mathbf{T}_{<i}, \rho_{<i})$ for $i := k + 1$. When $i = 1$, we write $t_{\text{pp}, \eta}^*$ and $\epsilon_{\text{pp}, \eta}$ for brevity.

6.2.1 Coverage factors

To facilitate the security analysis, we introduce coverage factors and prove a technical lemma.

Definition 6.11 (Coverage factor). Let $\ell, M \in \mathbb{N}$. For a distribution D over query-answer tuples of $\mathbf{Q}_\ell$ and a satisfiable constraint set $\mathcal{C}$, define

$$c_D(\mathcal{C}) := \Pr_{X \leftarrow D} [\mathcal{C} \text{ decides } X] .$$

The coverage factor of $\mathbf{Q}$ is

$$C_{\mathbf{Q}}(\ell, M) := \sup_D \mathbb{E}_{T=(X_1, \dots, X_M) \leftarrow D^M} \left[\frac{1}{c_D(\mathcal{C}_T)} \right] ,$$

where the supremum ranges over all such distributions D . (Note that $c_D(\mathcal{C}_T) > 0$ with probability one, since $\mathcal{C}_T$ decides each X_i and the support of D is finite.)

Lemma 6.12. Let pp be a public parameter and η be auxiliary input where $|\mathbb{X}| \leq n$ and $\epsilon_{\text{pp},\eta} > 0$. Let $F_0, F_1, \dots, F_i$ be non-negative functions for $i \leq k$, where F_m takes as input $(\mathbf{T}_{<m+1}, \boldsymbol{\rho}_{<m+1})$, such that, for every $m \in [i]$, all $\mathbf{T}_{<m+1}$, and all $\boldsymbol{\rho}_{<m}$,

$$\mathbb{E}_{\rho_m \leftarrow \{0,1\}^{r_m}} [F_m(\mathbf{T}_{<m+1}, \boldsymbol{\rho}_{<m+1})] \leq F_{m-1}(\mathbf{T}_{<m}, \boldsymbol{\rho}_{<m}) .$$

It holds that

$$\mathbb{E}_{E_{\text{FC}}} \left[\frac{F_i(\mathbf{T}_{<i+1}, \boldsymbol{\rho}_{<i+1})}{\epsilon_{\text{pp},\eta}(\mathbf{T}_{<i+1}, \boldsymbol{\rho}_{<i+1})} \mid \text{pp}, \eta, b = 1 \right] \leq \left(\prod_{m=1}^i C_{\mathbf{Q}}(\ell_m, M_m) \right) \cdot \frac{F_0}{\epsilon_{\text{pp},\eta}} .$$

Proof. It suffices to show, for all $j \leq i$,

$$\begin{aligned} & \mathbb{E}_{E_{\text{FC}}} \left[\frac{F_j(\mathbf{T}_{<j+1}, \boldsymbol{\rho}_{<j+1})}{\epsilon_{\text{pp},\eta}(\mathbf{T}_{<j+1}, \boldsymbol{\rho}_{<j+1})} \mid \text{pp}, \eta, \mathbf{T}_{<j}, \boldsymbol{\rho}_{<j}, b = 1 \right] \\ & \leq C_{\mathbf{Q}}(\ell_j, M_j) \cdot \frac{F_{j-1}(\mathbf{T}_{<j}, \boldsymbol{\rho}_{<j})}{\epsilon_{\text{pp},\eta}(\mathbf{T}_{<j}, \boldsymbol{\rho}_{<j})} ; \end{aligned} \tag{10}$$

the lemma follows by successively averaging over each round $j = i, \dots, 1$, applying Eq. 10 in each step.

Let $j \in [i]$. Fix (pp, η) as in the statement, and fix reachable $(\mathbf{T}_{<j}, \boldsymbol{\rho}_{<j})$. Let H denote the experiment E conditioned on $(\text{pp}, \eta, \boldsymbol{\rho}_{<j})$, and let H_{FC} denote the experiment E_{FC} conditioned on $(\text{pp}, \eta, \mathbf{T}_{<j}, \boldsymbol{\rho}_{<j})$ and the event $\{b = 1\}$. For brevity, write

$$\begin{aligned} F_{j-1} &:= F_{j-1}(\mathbf{T}_{<j}, \boldsymbol{\rho}_{<j}), & F_j(T_j, \rho_j) &:= F_j(\mathbf{T}_{<j+1}, \boldsymbol{\rho}_{<j+1}), \\ \epsilon_j &:= \epsilon_{\text{pp},\eta}(\mathbf{T}_{<j}, \boldsymbol{\rho}_{<j}), & \epsilon_{j+1}(T_j, \rho_j) &:= \epsilon_{\text{pp},\eta}(\mathbf{T}_{<j+1}, \boldsymbol{\rho}_{<j+1}), \\ \delta_j(T_j) &:= \delta_{\text{pp},\eta}(\mathbf{T}_{<j+1}, \boldsymbol{\rho}_{<j}) . \end{aligned}$$

Fix T_j reachable in H_{FC} . We have

$$\mathbb{E}_{H_{\text{FC}}} \left[\frac{F_j(T_j, \rho_j)}{\epsilon_{j+1}(T_j, \rho_j)} \mid T_j \right] = \sum_{\substack{\rho_j \in \{0,1\}^{r_j} \\ \epsilon_{j+1}(T_j, \rho_j) > 0}} \Pr_{H_{\text{FC}}}[\rho_j \mid T_j] \cdot \frac{F_j(T_j, \rho_j)}{\epsilon_{j+1}(T_j, \rho_j)} .$$

Bayes' rule gives

$$\begin{aligned} \Pr_{H_{\text{FC}}}[\rho_j \mid T_j] &= \Pr_H[\rho_j \mid b = 1 \wedge \forall m \leq j, \mathcal{C}_{T_m} \text{ decides } \tau_m] \\ &= \frac{\Pr_H[b = 1 \wedge \forall m \leq j, \mathcal{C}_{T_m} \text{ decides } \tau_m \mid \rho_j] \cdot 2^{-r_j}}{\Pr_H[b = 1 \wedge \forall m \leq j, \mathcal{C}_{T_m} \text{ decides } \tau_m]} \\ &= \frac{\epsilon_{j+1}(T_j, \rho_j) \cdot 2^{-r_j}}{\epsilon_j \cdot \delta_j(T_j)} , \end{aligned}$$

where the final equality is by the definitions of ϵ and δ . It follows that

$$\begin{aligned} \mathbb{E}_{H_{\text{FC}}} \left[\frac{F_j(T_j, \rho_j)}{\epsilon_{j+1}(T_j, \rho_j)} \mid T_j \right] &\leq \sum_{\rho_j \in \{0,1\}^{r_j}} \frac{F_j(T_j, \rho_j) \cdot 2^{-r_j}}{\epsilon_j \cdot \delta_j(T_j)} \\ &= \frac{\mathbb{E}_{\rho_j \leftarrow \{0,1\}^{r_j}} [F_j(T_j, \rho_j)]}{\epsilon_j \cdot \delta_j(T_j)} \leq \frac{F_{j-1}}{\epsilon_j \cdot \delta_j(T_j)} . \end{aligned}$$

Let D be the distribution of query-answer tuples defined by

$$D(\tau) := \Pr_H[\tau_j = \tau \mid b = 1 \wedge \forall m < j, \mathcal{C}_{T_m} \text{ decides } \tau_m] .$$

By construction, $T_j \leftarrow D^{M_j}$ in H_{FC} and $\delta_j(T_j) = c_D(\mathcal{C}_{T_j})$. Thus, by Definition 6.11,

$$\mathbb{E}_{H_{\text{FC}}} \left[\frac{1}{\delta_j(T_j)} \right] \leq C_{\mathbf{Q}}(\ell_j, M_j) ,$$

and Eq. 10 follows by averaging over T_j in H_{FC} . $\square$

6.2.2 Running time

We analyze the expected running time of the adversary $\mathcal{A}_{\text{FC}}$ (Construction 6.9).

Lemma 6.13 (Expected running time). *The expected running time of $\mathcal{A}_{\text{FC}}^{\tilde{\mathcal{P}}}$ in E_{FC} is at most $O(N \cdot t_{\text{ARG}}^*)$ for N as defined in Theorem 6.1.*

Proof. For simplicity, we focus on the time spent executing $\tilde{\mathcal{P}}$; the remaining time can be bounded with some additional accounting. Fix a public parameter pp and auxiliary input η , which together determine the instance $\mathfrak{x}$ and prover state aux . Without loss of generality, assume $|\mathfrak{x}| \leq n$ and $\epsilon_{\text{pp}, \eta} > 0$.

Claim 6.14. *For a round $i \in [k]$, let $\mathcal{T}_i$ denote the time spent executing $\tilde{\mathcal{P}}$ during one invocation of S_i in the i -th iteration of Loop 1. It holds that*

$$\mathbb{E}[\mathcal{T}_i \mid \text{pp}, \eta, b = 1] \leq \left(\prod_{j=1}^{i-1} C_{\mathbf{Q}}(\ell_j, M_j) \right) \cdot \frac{t_{\text{pp}, \eta}^*}{\epsilon_{\text{pp}, \eta}} .$$

Proof. By construction,

$$\mathbb{E}[\mathcal{T}_i \mid \text{pp}, \eta, b = 1, \mathbf{T}_{<i}, \boldsymbol{\rho}_{<i}] \leq \frac{t_{\text{pp}, \eta}^*(\boldsymbol{\rho}_{<i})}{\epsilon_{\text{pp}, \eta}(\mathbf{T}_{<i}, \boldsymbol{\rho}_{<i})} .$$

Averaging over $(\mathbf{T}_{<i}, \boldsymbol{\rho}_{<i})$ and applying Lemma 6.12 with

$$F_j(\mathbf{T}_{<j+1}, \boldsymbol{\rho}_{<j+1}) := t_{\text{pp}, \eta}^*(\boldsymbol{\rho}_{<j+1}) \quad \text{for } j \in \{0, \dots, i-1\}$$

yields the claim. $\square$

Claim 6.15. *For a round $i \in [k]$, let $\mathcal{T}_{i, \text{fin}}$ denote the time spent executing $\tilde{\mathcal{P}}$ during the invocation of $S_{i, \text{fin}}$ in the i -th iteration of Loop 1. It holds that*

$$\mathbb{E}[\mathcal{T}_{i, \text{fin}} \mid \text{pp}, \eta, b = 1] \leq \left(\prod_{j=1}^i C_{\mathbf{Q}}(\ell_j, M_j) \right) \cdot \frac{t_{\text{pp}, \eta}^*}{\epsilon_{\text{pp}, \eta}} .$$

Proof. By construction,

$$\mathbb{E}[\mathcal{T}_{i, \text{fin}} \mid \text{pp}, \eta, b = 1, \mathbf{T}_{<i+1}, \boldsymbol{\rho}_{<i}] \leq \frac{t_{\text{pp}, \eta}^*(\boldsymbol{\rho}_{<i})}{\epsilon_{\text{pp}, \eta}(\mathbf{T}_{<i}, \boldsymbol{\rho}_{<i}) \cdot \delta_{\text{pp}, \eta}(\mathbf{T}_{<i+1}, \boldsymbol{\rho}_{<i})} .$$

Averaging over $\mathbf{T}_i$ and applying the definition of $C_{\mathbf{Q}}$ in Definition 6.11, we obtain

$$\mathbb{E}[\mathcal{T}_{i, \text{fin}} \mid \text{pp}, \eta, b = 1, \mathbf{T}_{<i}, \boldsymbol{\rho}_{<i}] \leq C_{\mathbf{Q}}(\ell_i, M_i) \cdot \frac{t_{\text{pp}, \eta}^*(\boldsymbol{\rho}_{<i})}{\epsilon_{\text{pp}, \eta}(\mathbf{T}_{<i}, \boldsymbol{\rho}_{<i})} .$$

Averaging over $(\mathbf{T}_{<i}, \boldsymbol{\rho}_{<i})$ and applying Lemma 6.12 (as in Claim 6.14) yields the claim. $\square$

By Claims 6.14 and 6.15,

$$\mathbb{E}[\text{running time of } \mathcal{R} \mid \text{pp}, \eta, b = 1] \leq \left(\sum_{i=1}^k (M_i + C_{\mathbf{Q}}(\ell_i, M_i)) \cdot \prod_{j=1}^{i-1} C_{\mathbf{Q}}(\ell_j, M_j) \right) \cdot \frac{t_{\text{pp}, \eta}^*}{\epsilon_{\text{pp}, \eta}} .$$

Since the reducer only executes if $b = 1$, we conclude that

$$\begin{aligned} \mathbb{E}[\text{running time of } \mathcal{R} \mid \text{pp}, \eta] &= \mathbb{E}[\text{running time of } \mathcal{R} \mid \text{pp}, \eta, b = 1] \cdot \Pr[b = 1 \mid \text{pp}, \eta] \\ &\leq \left(\sum_{i=1}^k (M_i + C_{\mathbf{Q}}(\ell_i, M_i)) \cdot \prod_{j=1}^{i-1} C_{\mathbf{Q}}(\ell_j, M_j) \right) \cdot t_{\text{pp}, \eta}^* ; \end{aligned}$$

averaging over pp and η bounds the expected running time of the reducer $\mathcal{R}$. Accounting also for the initial execution of $\tilde{\mathcal{P}}$ gives the lemma. $\square$

6.2.3 Success probability

We analyze the probability that the extractor $\mathcal{E}$ fails to output a witness.

When the reducer executes in E_{FC} , let $(\text{cm}_1, \rho_1, \dots, \text{cm}_k, \rho_k, (\tau_1, \dots, \tau_k))$ denote the transcript tr_k produced in the final loop iteration of the round- k terminal sampler. Define the following events:

- Ext: the reducer executes and $(\mathbb{X}, \mathbf{E}(\mathbb{X}, \tilde{\Pi}_{T_1}, \rho_1, \dots, \tilde{\Pi}_{T_k}, \rho_k)) \in R$.
- Coll: the reducer executes and, for some $i \in [k]$, $\mathcal{C}_{T_i}$ refutes τ_i .

Definition 6.16. For an instance $\mathbb{X}$, inductively define

$$\begin{aligned} \kappa_k(\mathbb{X}, \Pi_1, \rho_1, \dots, \Pi_k, \rho_k) &:= \mathbb{1} \left[\begin{aligned} &\mathbf{V}^{\Pi_1, \dots, \Pi_k}(\mathbb{X}; \rho_1, \dots, \rho_k) = 1 \\ &\wedge (\mathbb{X}, \mathbf{E}(\mathbb{X}, \Pi_1, \rho_1, \dots, \Pi_k, \rho_k)) \notin R \end{aligned} \right] , \\ \kappa_{i-1}(\mathbb{X}, \Pi_1, \rho_1, \dots, \Pi_{i-1}, \rho_{i-1}) &:= \max_{\Pi_i \in \Sigma^{\ell_i}} \mathbb{E}_{\rho_i \leftarrow \{0,1\}^{\ell_i}} [\kappa_i(\mathbb{X}, \Pi_1, \rho_1, \dots, \Pi_i, \rho_i)] \quad \text{for } i = k, \dots, 1 . \end{aligned}$$

Lemma 6.17 (Success probability). *It holds that*

$$\Pr \left[\begin{aligned} &|\mathbb{X}| \leq n \\ &\wedge \neg \text{Ext} \\ &\wedge \neg \text{Coll} \\ &\wedge b = 1 \end{aligned} \right] \leq \left(\prod_{i=1}^k C_{\mathbf{Q}}(\ell_i, M_i) \right) \cdot \kappa_{\text{FIOP}}(n) .$$

Proof. Fix a public parameter pp and auxiliary input η , which together determine the instance $\mathbb{X}$ and prover state aux . Without loss of generality, assume $|\mathbb{X}| \leq n$.

Conditioned on $b = 1$, the following hold:

- If Ext does not occur, then

$$(\mathbb{X}, \mathbf{E}(\mathbb{X}, \tilde{\Pi}_{T_1}, \rho_1, \dots, \tilde{\Pi}_{T_k}, \rho_k)) \notin R .$$

- Recall that each $\mathcal{C}_{T_i}$ decides τ_i . Thus, if Coll does not occur, then each $\mathcal{C}_{T_i}$ entails τ_i , and hence each $\tilde{\Pi}_{T_i}$ agrees with τ_i . Since tr_k is accepting,

$$\mathbf{V}^{\tilde{\Pi}_{T_1}, \dots, \tilde{\Pi}_{T_k}}(\mathbb{X}; \rho_1, \dots, \rho_k) = 1 .$$

Therefore,

$$\begin{aligned} \Pr[\neg \text{Ext} \wedge \neg \text{Coll} \mid \text{pp}, \eta, b = 1] &\leq \mathbb{E} \left[\kappa_k(\mathbb{X}, \tilde{\Pi}_{T_1}, \rho_1, \dots, \tilde{\Pi}_{T_k}, \rho_k) \mid \text{pp}, \eta, b = 1 \right] \\ &\leq \mathbb{E} \left[\frac{\kappa_k(\mathbb{X}, \tilde{\Pi}_{T_1}, \rho_1, \dots, \tilde{\Pi}_{T_k}, \rho_k)}{\epsilon_{\text{pp}, \eta}(\mathbf{T}_{<k+1}, \boldsymbol{\rho}_{<k+1})} \mid \text{pp}, \eta, b = 1 \right] , \end{aligned}$$

where the final quantity is well-defined since $\epsilon_{\text{pp}, \eta}(\mathbf{T}_{<k+1}, \boldsymbol{\rho}_{<k+1}) > 0$ for any $(\mathbf{T}_{<k+1}, \boldsymbol{\rho}_{<k+1})$ reachable in E_{FC} . Applying Lemma 6.12 with

$$F_j(\mathbf{T}_{<j+1}, \boldsymbol{\rho}_{<j+1}) := \kappa_j(\mathbb{X}, \tilde{\Pi}_{T_1}, \rho_1, \dots, \tilde{\Pi}_{T_j}, \rho_j) \quad \text{for } j = 0, \dots, k ,$$

we obtain

$$\mathbb{E} \left[\kappa_k(\mathbb{X}, \tilde{\Pi}_{T_1}, \rho_1, \dots, \tilde{\Pi}_{T_k}, \rho_k) \mid \text{pp}, \eta, b = 1 \right] \leq \left(\prod_{i=1}^k C_{\mathbf{Q}}(\ell_i, M_i) \right) \cdot \frac{\kappa_0(\mathbb{X})}{\epsilon_{\text{pp}, \eta}} .$$

By knowledge soundness of FIOP (Definition 3.24), it holds that $\kappa_0(\mathbb{X}) \leq \kappa_{\text{FIOP}}(n)$. We conclude that

$$\begin{aligned} \Pr[\neg \text{Ext} \wedge \neg \text{Coll} \wedge b = 1 \mid \text{pp}, \eta] &= \Pr[\neg \text{Ext} \wedge \neg \text{Coll} \mid \text{pp}, \eta, b = 1] \cdot \Pr[b = 1 \mid \text{pp}, \eta] \\ &\leq \left(\prod_{i=1}^k C_{\mathbf{Q}}(\ell_i, M_i) \right) \cdot \kappa_{\text{FIOP}}(n) , \end{aligned}$$

and the lemma follows by averaging over pp and η . $\square$

Proof of Theorem 6.1. Instantiate the reducer $\mathcal{R}$ with repetition parameters $M_1, \dots, M_k$ as in the statement. Let $E_{\mathcal{E}}$ be the extension of E that additionally samples

$$\mathbb{W} \leftarrow \mathcal{E}^{\tilde{\mathcal{P}}(\text{aux})}(\text{pp}, \mathbb{X}, \text{tr}, b) .$$

The reducer execution in this experiment is distributed as in E_{FC} . The argument extractor runs this reducer and then $\mathbf{E}$, so by Lemma 6.13, its expected running time is at most

$$O(N \cdot t_{\text{ARG}}^* + t_{\mathbf{E}}) .$$

By total probability,

$$\begin{aligned} \Pr_{E_{\mathcal{E}}} \left[\begin{array}{l} |\mathbb{X}| \leq n \\ \wedge (\mathbb{X}, \mathbb{W}) \notin R \\ \wedge b = 1 \end{array} \right] &= \Pr_{E_{\text{FC}}} \left[\begin{array}{l} |\mathbb{X}| \leq n \\ \wedge \neg \text{Ext} \\ \wedge b = 1 \end{array} \right] \\ &\leq \underbrace{\Pr_{E_{\text{FC}}} \left[\begin{array}{l} |\mathbb{X}| \leq n \\ \wedge \neg \text{Ext} \\ \wedge \neg \text{Coll} \\ \wedge b = 1 \end{array} \right]}_A + \underbrace{\Pr_{E_{\text{FC}}} \left[\begin{array}{l} |\mathbb{X}| \leq n \\ \wedge \text{Coll} \\ \wedge b = 1 \end{array} \right]}_B . \end{aligned}$$

Bounding A. By Lemma 6.17,

$$A \leq \left(\prod_{i=1}^k C_{\mathbf{Q}}(\ell_i, M_i) \right) \cdot \kappa_{\text{FIOP}}(n) .$$

Bounding B. In E_{FC} , observe that Coll implies $\text{Sol}(\mathcal{C}_{\text{cm}_i, \ell_i, \mathbf{q}_i}) = \emptyset$ for some $i \in [k]$ (see Definition 3.17). By expected-time function binding of FC (Definition 3.19) and Lemma 6.13,

$$B \leq \epsilon_{\text{FC}}^*(\lambda, \ell_{\max}, \mathbf{q}_{\max}, t_{\text{FC}}^*) ,$$

where t_{FC}^* is as stated in the theorem. $\square$

Proof of Corollary 6.2. We instantiate Theorem 6.1 with

$$M_i := \max \left\{ 1, \left\lceil \frac{k \cdot \ell_i}{\log(1 + \alpha)} \right\rceil \right\} \quad \text{for } i \in [k] .$$

By Lemma A.2, this yields

$$C_{\mathbf{Q}_{\text{Point}}}(\ell_i, M_i) \leq 2^{\ell_i/M_i} \leq (1 + \alpha)^{1/k} ,$$

Writing $\rho := (1 + \alpha)^{1/k}$, we have

$$\begin{aligned} N &= 1 + \sum_{i=1}^k (M_i + C_{\mathbf{Q}_{\text{Point}}}(\ell_i, M_i)) \prod_{j=1}^{i-1} C_{\mathbf{Q}_{\text{Point}}}(\ell_j, M_j) \\ &\leq 1 + \sum_{i=1}^k (M_i + \rho) \rho^{i-1} \\ &= 1 + \frac{\alpha \rho}{\rho - 1} + \sum_{i=1}^k M_i \rho^{i-1} . \end{aligned}$$

□

Proof of Corollary 6.3. By Lemma A.3, there exist repetition parameters $M_1, \dots, M_k$ satisfying

$$M_i = O_\alpha(k \cdot d_{\mathbf{Q}}(\ell_i)) \quad \text{and} \quad C_{\mathbf{Q}}(\ell_i, M_i) \leq (1 + \alpha)^{1/k} \quad \text{for } i \in [k] .$$

Instantiating Theorem 6.1 with these parameters yields

$$\begin{aligned} N &= 1 + \sum_{i=1}^k (M_i + C_{\mathbf{Q}}(\ell_i, M_i)) \prod_{j=1}^{i-1} C_{\mathbf{Q}}(\ell_j, M_j) \\ &= O_\alpha \left(k \cdot \sum_{i=1}^k d_{\mathbf{Q}}(\ell_i) \right) = O_\alpha(k \cdot d) . \end{aligned}$$

□

A Coverage factor bounds

Fact A.1 (Hölder's inequality for sums). For every finite set I , non-negative reals $(a_i)_{i \in I}$ and $(b_i)_{i \in I}$, and reals $p, q > 1$ satisfying $1/p + 1/q = 1$,

$$\sum_{i \in I} a_i b_i \leq \left(\sum_{i \in I} a_i^p \right)^{1/p} \left(\sum_{i \in I} b_i^q \right)^{1/q}.$$

Lemma A.2 (Point query bound). For every $\ell, M \in \mathbb{N}$ with $M \geq 1$,

$$C_{\mathbf{Q}_{\text{Point}}}(\ell, M) \leq 2^{\ell/M}.$$

Proof. Fix a distribution D over query-answer tuples of $\mathbf{Q}_{\text{Point}}$, and let Q be the induced distribution over their query sets, viewed as subsets of $[\ell]$. Draw $\tau_1, \dots, \tau_M$ independently from D , let $T := (\tau_1, \dots, \tau_M)$, and let U be the union of their query sets. For every $S \subseteq [\ell]$, write

$$c(S) := \Pr_{X \leftarrow Q}[X \subseteq S] \text{ and } p(S) := \Pr[U = S].$$

The event $U = S$ implies that all M sampled query sets are contained in S , and hence

$$p(S) \leq c(S)^M.$$

Observe that $c(U) = c_D(\mathcal{C}_T)$. By Fact A.1, the following holds for $M \geq 2$:

$$\begin{aligned} \mathbb{E} \left[\frac{1}{c_D(\mathcal{C}_T)} \right] &= \sum_{\substack{S \subseteq [\ell] \\ p(S) > 0}} \frac{p(S)}{c(S)} \\ &\leq \sum_{S \subseteq [\ell]} p(S)^{1-1/M} \\ &\leq \left(\sum_{S \subseteq [\ell]} 1^M \right)^{1/M} \left(\sum_{S \subseteq [\ell]} p(S) \right)^{1-1/M} \\ &= 2^{\ell/M}. \end{aligned}$$

For $M = 1$, the same bound follows from $p(S)/c(S) \leq 1$ and there being at most 2^ℓ possible sets S . $\square$

Lemma A.3 (Diameter-based bound). For every $\ell, M \in \mathbb{N}$ with $d := d_{\mathbf{Q}}(\ell)$ positive and $M \geq d$,

$$C_{\mathbf{Q}}(\ell, M) \leq 1 + \frac{2d}{M+1} + \sum_{j=0}^{d-1} \binom{M}{j} 2^{-(M-j-1)},$$

and, for all $M \geq 5d$,

$$C_{\mathbf{Q}}(\ell, M) \leq 1 + O\left(\frac{d}{M}\right).$$

Proof. Fix a distribution D over query-answer tuples of $\mathbf{Q}$. Draw $\tau_1, \dots, \tau_M, \tau_{M+1}$ independently from D and let $T := (\tau_1, \dots, \tau_M)$. Set $Z := 1/c_D(\mathcal{C}_T)$ and fix $\theta \in (0, 1)$; we have

$$\mathbb{E}[Z] = \mathbb{E}[Z \cdot \mathbb{1}[c_D(\mathcal{C}_T) \geq \theta]] + \mathbb{E}[Z \cdot \mathbb{1}[c_D(\mathcal{C}_T) < \theta]].$$

For $n \in \{0, \dots, M\}$, let $\mathcal{C}_n := \mathcal{C}_{(\tau_1, \dots, \tau_n)}$, with $\mathcal{C}_0 := \emptyset$. As in Lemma 5.7, say that τ_{n+1} is *new* if $\mathcal{C}_n$ does not decide τ_{n+1} .

Claim A.4.

$$\mathbb{E}\left[Z \cdot \mathbb{1}[c_D(\mathcal{C}_T) \geq \theta]\right] \leq 1 + \frac{d}{\theta(M+1)} .$$

Proof. On $c_D(\mathcal{C}_T) \geq \theta$,

$$Z = 1 + \frac{1 - c_D(\mathcal{C}_T)}{c_D(\mathcal{C}_T)} \leq 1 + \frac{1 - c_D(\mathcal{C}_T)}{\theta} .$$

It follows that

$$\begin{aligned} \mathbb{E}\left[Z \cdot \mathbb{1}[c_D(\mathcal{C}_T) \geq \theta]\right] &\leq \mathbb{E}\left[\left(1 + \frac{1 - c_D(\mathcal{C}_T)}{\theta}\right) \cdot \mathbb{1}[c_D(\mathcal{C}_T) \geq \theta]\right] \\ &\leq 1 + \frac{1}{\theta} \cdot \mathbb{E}[1 - c_D(\mathcal{C}_T)] . \end{aligned}$$

For $n \in \{0, \dots, M\}$, let $a_n := \mathbb{E}[1 - c_D(\mathcal{C}_n)]$. Since τ_{n+1} is independent of $\mathcal{C}_n$, a_n is the probability that τ_{n+1} is new. The sequence $(a_n)_{n=0}^M$ is non-increasing, and at most d samples can be new. Hence

$$\mathbb{E}[1 - c_D(\mathcal{C}_T)] = a_M \leq \frac{1}{M+1} \sum_{n=0}^M a_n \leq \frac{d}{M+1} .$$

□

Claim A.5.

$$\mathbb{E}\left[Z \cdot \mathbb{1}[c_D(\mathcal{C}_T) < \theta]\right] \leq \sum_{j=0}^{d-1} \binom{M}{j} \theta^{M-j-1} .$$

Proof. Let $R(T) \subseteq [M]$ be the set of indices j for which τ_j is new. Observe that $c_D(\mathcal{C}_T) < \theta$ can occur only when $|R(T)| \leq d-1$. Accordingly, fix $I \subseteq [M]$ with $|I| \leq d-1$, and let T_I denote the sublist of T indexed by I . Conditional on T_I , the remaining tuples are independent samples from D . The event $R(T) = I$ implies that $\mathcal{C}_T = \mathcal{C}_{T_I}$ and that $\mathcal{C}_{T_I}$ decides every tuple outside T_I . Therefore,

$$\Pr[R(T) = I \mid T_I] \leq c_D(\mathcal{C}_{T_I})^{M-|I|} .$$

On the event $R(T) = I$, we have $Z = 1/c_D(\mathcal{C}_{T_I})$. Thus, for every such I ,

$$\begin{aligned} \mathbb{E}\left[Z \cdot \mathbb{1}[c_D(\mathcal{C}_T) < \theta] \cdot \mathbb{1}[R(T) = I] \mid T_I\right] \\ \leq \mathbb{1}[c_D(\mathcal{C}_{T_I}) < \theta] \cdot c_D(\mathcal{C}_{T_I})^{M-|I|-1} \leq \theta^{M-|I|-1} . \end{aligned}$$

Averaging over T_I and summing over the possible values of $R(T)$ proves the claim. □

Applying the two claims, we get

$$\mathbb{E}[Z] \leq 1 + \frac{d}{\theta(M+1)} + \sum_{j=0}^{d-1} \binom{M}{j} \theta^{M-j-1} .$$

Setting $\theta = 1/2$ establishes the statement's first bound. The second bound follows from the fact that setting $M \geq 5d$ yields

$$\sum_{j=0}^{d-1} \binom{M}{j} 2^{-(M-j-1)} \leq 2^{-(M-d)} \sum_{j=0}^d \binom{M}{j}$$

$$\begin{aligned}
&\leq 2^{-(M-d)} \left(\frac{eM}{d} \right)^d \\
&\leq \left(\frac{(5e)^{1/5}}{2^{4/5}} \right)^M < 0.968^M .
\end{aligned}$$

□

Acknowledgments

Ziyi Guan thanks Christian Knabenhans, Nicholas Spooner and Eylon Yogev for helpful discussions in the early stages of this work.

AI disclosure

This paper has three main technical components: the lower bounds (Section 4), the soundness analysis of the Funky protocol (Section 5), and the knowledge soundness analysis of Funky (Section 6). The ideas and proofs in the lower bound section are entirely the work of the authors. For both analyses of Funky, the authors first established slightly weaker bounds without AI assistance. The bounds presented here were subsequently tightened through iterative discussions with AI models, drawing on the authors’ insights. The proofs of the coverage factor bounds in Section A were produced primarily by an AI model, with subsequent edits by the authors. AI tools also provided assistance in polishing the text throughout the paper.

References

- [AFR23] Thomas Attema, Serge Fehr, and Nicolas Resch. “Generalized Special-Sound Interactive Proofs and Their Knowledge Soundness”. In: *Proceedings of the 21st Theory of Cryptography Conference*. TCC ’23. 2023, pp. 424–454. DOI: 10.1007/978-3-031-48621-0_15.
- [AKLY24] Thomas Attema, Michael Klooß, Russell W. F. Lai, and Pavlo Yatsyna. *Adaptive Special Soundness: Improved Knowledge Extraction by Adaptive Useful Challenge Sampling*. Cryptology ePrint Archive, Paper 2024/2038. 2024. URL: <https://eprint.iacr.org/2024/2038>.
- [BBHR18] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. “Fast Reed–Solomon Interactive Oracle Proofs of Proximity”. In: *Proceedings of the 45th International Colloquium on Automata, Languages and Programming*. ICALP ’18. 2018, 14:1–14:17. DOI: 10.4230/LIPIcs.ICALP.2018.14.
- [BCG20] Jonathan Bootle, Alessandro Chiesa, and Jens Groth. “Linear-Time Arguments with Sublinear Verification from Tensor Codes”. In: *Proceedings of the 18th Theory of Cryptography Conference*. TCC ’20. 2020, pp. 19–46. DOI: 10.1007/978-3-030-64378-2_2.
- [BCS16] Eli Ben-Sasson, Alessandro Chiesa, and Nicholas Spooner. “Interactive Oracle Proofs”. In: *Proceedings of the 14th Theory of Cryptography Conference*. TCC ’16-B. 2016, pp. 31–60. DOI: 10.1007/978-3-662-53644-5_2.
- [BF23] Benedikt Bünz and Ben Fisch. “Multilinear Schwartz-Zippel Mod N and Lattice-Based Succinct Arguments”. In: *Proceedings of the 21st Theory of Cryptography Conference*. TCC ’23. 2023, 394–423. DOI: 10.1007/978-3-031-48621-0_14.
- [BFS20] Benedikt Bünz, Ben Fisch, and Alan Szepieniec. “Transparent SNARKs from DARK Compilers”. In: *Proceedings of the 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques*. EUROCRYPT ’20. 2020, pp. 677–706. DOI: 10.1007/978-3-030-45721-1_24.
- [BG08] Boaz Barak and Oded Goldreich. “Universal Arguments and their Applications”. In: *SIAM Journal on Computing* 38.5 (2008). Preliminary version appeared in CCC ’02., pp. 1661–1694. DOI: 10.1137/070709244.
- [BL04] Boaz Barak and Yehuda Lindell. “Strict Polynomial-Time in Simulation and Extraction”. In: *SIAM Journal on Computing* 33.4 (2004), pp. 783–818. DOI: 10.1137/S0097539703427975.
- [CDDGS25] Alessandro Chiesa, Marcel Dall’Agnol, Zijing Di, Ziyi Guan, and Nicholas Spooner. “Quantum Rewinding for IOP-Based Succinct Arguments”. In: *Proceedings of the 23rd Theory of Cryptography Conference*. TCC ’25. 2025, pp. 460–479. DOI: 10.1007/978-3-032-12296-4_16.

- [CDGS23] Alessandro Chiesa, Marcel Dall’Agnol, Ziyi Guan, and Nicholas Spooner. *On the Security of Succinct Interactive Arguments from Vector Commitments*. Cryptology ePrint Archive, Report 2023/1737. 2023. URL: <https://eprint.iacr.org/2023/1737>.
- [CDGSY24] Alessandro Chiesa, Marcel Dall’Agnol, Ziyi Guan, Nicholas Spooner, and Eylon Yogev. “Untangling the Security of Kilian’s Protocol: Upper and Lower Bounds”. In: *Proceedings of the 22nd Theory of Cryptography Conference*. TCC ’24. 2024, pp. 158–188. DOI: 10.1007/978-3-031-78011-0_6.
- [CF13] Dario Catalano and Dario Fiore. “Vector Commitments and Their Applications”. In: *Proceedings of the 16th International Conference on Practice and Theory in Public Key Cryptography*. PKC ’13. 2013, pp. 55–72. DOI: 10.1007/978-3-642-36362-7_5.
- [CFFQR21] Matteo Campanelli, Antonio Faonio, Dario Fiore, Anaïs Querol, and Hadrián Rodríguez. “Lunar: A Toolbox for More Efficient Universal and Updatable zkSNARKs and Commit-and-Prove Extensions”. In: *Proceedings of the 27th International Conference on the Theory and Application of Cryptology and Information Security*. ASIACRYPT ’21. 2021, pp. 3–33. DOI: 10.1007/978-3-030-92078-4_1.
- [CGKY26] Alessandro Chiesa, Ziyi Guan, Christian Knabenhans, and Zihan Yu. “On the Fiat–Shamir Security of Succinct Arguments from Functional Commitments”. In: *Proceedings of the 46th Annual International Cryptology Conference*. CRYPTO ’26. 2026, pp. 189–213. DOI: 10.1007/978-3-032-35424-2_7.
- [CHMMVW20] Alessandro Chiesa, Yuncong Hu, Mary Maller, Pratyush Mishra, Noah Vesely, and Nicholas Ward. “Marlin: Preprocessing zkSNARKs with Universal and Updatable SRS”. In: *Proceedings of the 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques*. EUROCRYPT ’20. 2020, pp. 738–768. DOI: 10.1007/978-3-030-45721-1_26.
- [CMSZ21] Alessandro Chiesa, Fermi Ma, Nicholas Spooner, and Mark Zhandry. “Post-Quantum Succinct Arguments: Breaking the Quantum Rewinding Barrier”. In: *Proceedings of the 62nd Annual IEEE Symposium on Foundations of Computer Science*. FOCS ’21. IEEE, 2021, pp. 49–58. DOI: 10.1109/FOCS52979.2021.00014.
- [CY24] Alessandro Chiesa and Eylon Yogev. *Building Cryptographic Proofs from Hash Functions*. 2024. URL: <https://github.com/hash-based-snarks-book/hash-based-snarks-book>.
- [Din07] Irit Dinur. “The PCP Theorem by Gap Amplification”. In: *Journal of the ACM* 54.3, 12 (2007). DOI: 10.1145/1236457.1236459.
- [HM26] Iftach Haitner and Nikolaos Makriyannis. “Tight Bounds on Uniform-Challenge Reductions from Σ Protocols”. In: *Proceedings of the 45th Annual International Conference on the Theory and Applications of Cryptographic Techniques*. EUROCRYPT ’26. 2026, pp. 125–153. DOI: 10.1007/978-3-032-25330-9_5.
- [IJK09] Russell Impagliazzo, Ragesh Jaiswal, and Valentine Kabanets. “Approximate List-Decoding of Direct Product Codes and Uniform Hardness Amplification”. In: *SIAM Journal on Computing* 39.2 (2009), pp. 564–605. DOI: 10.1137/070683994.
- [IPZ01] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. “Which Problems Have Strongly Exponential Complexity?” In: *Journal of Computer and System Sciences* 63.4 (2001), pp. 512–530. DOI: 10.1006/jcss.2001.1774.
- [JT20] Joseph Jaeger and Stefano Tessaro. “Expected-Time Cryptography: Generic Techniques and Applications to Concrete Soundness”. In: *Proceedings of the 18th Theory of Cryptography Conference*. TCC ’20. 2020, pp. 414–443. DOI: 10.1007/978-3-030-64381-2_15.
- [Kil92] Joe Kilian. “A note on efficient zero-knowledge proofs and arguments (extended abstract)”. In: *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*. STOC ’92. ACM, 1992, pp. 723–732. DOI: 10.1145/129712.129782.
- [Klo21] Michael Klooß. “On Expected Polynomial Runtime in Cryptography”. In: *Proceedings of the 19th Theory of Cryptography Conference*. TCC ’21. 2021, pp. 558–590. DOI: 10.1007/978-3-030-90459-3_19.
- [LM19] Russell W. F. Lai and Giulio Malavolta. “Subvector Commitments with Application to Succinct Arguments”. In: *Proceedings of the 39th Annual International Cryptology Conference*. CRYPTO ’19. 2019, pp. 530–560. DOI: 10.1007/978-3-030-26948-7_19.

- [LPS24] Helger Lipmaa, Roberto Parisella, and Janno Siim. “Constant-Size zk-SNARKs in ROM from Falsifiable Assumptions”. In: *Proceedings of the 43rd Annual International Conference on the Theory and Applications of Cryptographic Techniques*. EUROCRYPT ’24. 2024, pp. 34–64. DOI: 10.1007/978-3-031-58751-1_2.
- [LPS25] Helger Lipmaa, Roberto Parisella, and Janno Siim. “On Knowledge-Soundness of Plonk in ROM from Falsifiable Assumptions”. In: *Proceedings of the 45th Annual International Cryptology Conference*. CRYPTO ’25. 2025, pp. 362–395. DOI: 10.1007/978-3-032-01907-3_12.
- [Wik24] Douglas Wikström. “Special Soundness Revisited”. In: *IACR Communications in Cryptology* 1.3 (2024). DOI: 10.62056/aep2c3w9p.