

Concrete Security for Succinct Arguments from Vector Commitments

Alessandro Chiesa

alessandro.chiesa@epfl.ch

EPFL

Marcel Dall’Agnol

dallagnol@cs.princeton.edu

Princeton University

Ziyi Guan

ziyi.guan@epfl.ch

EPFL

Nicholas Spooner

nicholas.spooner@warwick.ac.uk

University of Warwick & NYU

November 9, 2023

Abstract

We study the concrete security of a fundamental family of succinct interactive arguments, stemming from the works of Kilian (1992) and Ben-Sasson, Chiesa, and Spooner (“BCS”, 2016). These constructions achieve succinctness by combining probabilistic proofs and vector commitments.

Our first result concerns the succinct interactive argument of Kilian, realized with any probabilistically-checkable proof (PCP) and any vector commitment. We establish the tightest known bounds on the security of this protocol. Prior analyses incur large overheads unsuitable for concrete security, or assume special (and restrictive) properties of the underlying PCP.

Our second result concerns an interactive variant of the BCS succinct non-interactive argument, which here we call IBCS, realized with any public-coin interactive oracle proof (IOP) and any vector commitment. We establish tight bounds on the security of this protocol. While this variant has been informally discussed in the literature, no prior security analysis, even asymptotic, existed before this work.

Finally, we study the capabilities and limitations of succinct arguments based on vector commitments. We show that a generalization of the IBCS protocol, which we call the *Finale protocol*, is secure when realized with any *public-query* IOP (a notion that we introduce) that satisfies a natural “random continuation sampling” (RCS) property. We also show a partial converse: if the Finale protocol satisfies the RCS property (which in particular implies its security), then so does the underlying public-query IOP.

Keywords: succinct interactive arguments; concrete security; vector commitment schemes

Contents

1	Introduction	3
1.1	Our results	4
1.2	Discussion	6
1.3	Related work	8
2	Techniques	10
2.1	Succinct arguments based on PCPs	10
2.2	Succinct arguments based on public-coin IOPs	16
2.3	Why doesn't the same analysis work for public-query IOPs?	19
2.4	Succinct interactive arguments with adaptive security	23
3	Preliminaries	25
3.1	Interactive arguments	25
3.2	Vector commitments	26
3.3	Probabilistically checkable proofs	27
3.4	Interactive oracle proofs	28
4	Interactive arguments based on PCPs	30
4.1	Construction	30
4.2	Security reduction	31
4.3	Adaptive soundness	35
4.4	Adaptive knowledge soundness	37
5	Interactive arguments based on public-coin IOPs	39
5.1	Construction	39
5.2	Security reduction	40
5.3	Adaptive soundness	44
5.4	Adaptive knowledge soundness	47
6	Interactive arguments based on public-query IOPs	50
6.1	Construction	50
6.2	Random continuation samplers	51
6.3	Equivalence of IOP-RCS and ARG-RCS	55
6.4	Security reduction	56
	Acknowledgments	62
	References	62

1 Introduction

A *succinct argument* for a relation R is a computationally-sound interactive proof where, for a given instance x , an argument prover seeks to convince the argument verifier that there exists a witness w such that $(x, w) \in R$, while communicating much less than $|w|$ bits. Succinct arguments are a central cryptographic object, with numerous theoretical and practical applications. Especially for the latter, it is crucial to understand **concrete security**: how the security of a given succinct argument quantitatively relates to the security of underlying building blocks. *In this work, we study the concrete security of a fundamental family of succinct arguments, namely, those obtained by combining probabilistic proofs and vector commitments (VCs).*

This family includes the first known succinct argument, due to Kilian [Kil92], which informally works as follows. The argument prover commits to a probabilistically checkable proof (PCP) string via a vector commitment (Kilian’s presentation uses a Merkle tree, a VC built from collision-resistant hash functions), and sends the resulting commitment to the argument verifier; the argument verifier sends PCP verifier randomness to the argument prover; and finally the argument prover reveals the values of the queried locations of the PCP string and accompanies these values with opening information. The argument verifier accepts if the opening information is valid and the PCP verifier accepts.

Despite this simple template, the security of Kilian’s protocol is not well understood. The original work of [Kil92] gave only an informal analysis. Barak and Goldreich [BG08] gave a detailed analysis, but with certain limitations. First, their analysis applies only to PCPs that satisfy certain (restrictive) properties; second, their goal was not concrete security, and their analysis is not tight; and third, they restrict their attention to Merkle trees rather than general VCs. This state of affairs motivates the first question that we ask in this work:

What is the concrete security of Kilian’s protocol, for any given PCP and VC?

While Kilian’s protocol relies on only basic (and practically efficient) cryptography, its reliance on PCPs means that it is of limited practical value, as known PCP constructions have poor concrete efficiency. To address this, Ben-Sasson, Chiesa, and Spooner [BCS16] introduced a new paradigm for constructing succinct arguments that builds on Kilian’s idea. Their insight was that, to construct a succinct argument, it is not necessary to start with a *non-interactive* object such as a PCP. Instead, they introduce an interactive generalization of a PCP, called an *interactive oracle proof* (IOP)¹ and show how to compile (via the “BCS protocol”) a public-coin IOP into a succinct *non-interactive* argument (SNARG). Since then, extensive study of IOPs has resulted in highly efficient IOP constructions, which have found widespread use in practice.

While this perspective was not taken in the original paper, the BCS protocol is generally understood as the composition of two transformations: first, convert the public-coin IOP into an *interactive* succinct argument using a vector commitment (analogously to Kilian’s protocol); then “collapse” the latter to a non-interactive argument (in the random oracle model) using the Fiat–Shamir transformation for multi-round protocols. In this work, we refer to the first component as “interactive BCS” (IBCS). While IBCS has appeared informally in the literature [BCG20; RR22; HR22], prior work contains no analysis of even *asymptotic* security, let alone concrete security. This motivates our second question:

What is the concrete security of IBCS, for any given public-coin IOP and VC?

Finally, we consider potential generalizations of the above result. We observe that one can formulate a generalization of IBCS that, syntactically, can be realized using *any* IOP (i.e., including private-coin IOPs). This brings us naturally to our final question: *for what class of IOPs is this VC-based approach secure?*

¹Independently introduced by [RRR16].

1.1 Our results

Throughout this section, we fix a vector commitment scheme VC and denote by ϵ_{VC} its position binding error, an upper bound on the probability that an adversary outputs two openings for the same commitment that disagree in at least one position. In general, ϵ_{VC} is a function of the security parameter λ , length ℓ of the committed vector, number s of opened entries of the vector, and bound t_{VC} on the adversary size.

Succinct arguments from PCPs. Our first result provides concrete security bounds for Kilian’s protocol [Kil92], a seminal result that combines a PCP system and a vector commitment scheme to obtain a succinct interactive argument. Let PCP be a PCP system for a relation R with proof length l , query complexity q , and soundness error ϵ_{PCP} (resp. knowledge soundness error κ_{PCP}); all can be functions of the given instance \mathbb{x} .

Theorem 1 (informal). *The soundness error ϵ_{ARG} of Kilian[PCP, VC] satisfies the following for every security parameter λ , instance $\mathbb{x} \notin L(R)$, adversary size bound t_{ARG} , and error tolerance $\epsilon > 0$:*

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

Similarly, the knowledge soundness error satisfies $\kappa_{\text{ARG}}(\lambda, \mathbb{x}, t_{\text{ARG}}) \leq \kappa_{\text{PCP}}(\mathbb{x}) + \epsilon_{\text{VC}}(\lambda, l(\mathbb{x}), q(\mathbb{x}), t_{\text{VC}}) + \epsilon$.

In particular, if ϵ_{VC} is negligible when $t_{\text{VC}} = \text{poly}(\lambda)$, then $\epsilon_{\text{ARG}}(\lambda, \mathbb{x}, \text{poly}(\lambda)) \leq \epsilon_{\text{PCP}}(\mathbb{x}) + \text{negl}(\lambda)$ (and similarly for κ_{ARG}).

The above bound is essentially tight. An adversary that commits to the PCP string $\tilde{\Pi}$ with maximal acceptance probability (and opens accordingly) succeeds with probability at least ϵ_{PCP} in Kilian[PCP, VC]. Moreover, an adversary that then tries to find a collision when \mathbf{V} rejects $\tilde{\Pi}$ achieves (under some mild conditions) a success probability of $\epsilon_{\text{PCP}} + (1 - \epsilon_{\text{PCP}}) \cdot \epsilon_{\text{VC}}$. The multiplicative loss $\frac{1}{\epsilon}$ in t_{VC} compared to t_{ARG} has an explanation: one needs to rewind the malicious argument prover enough times to reconstruct an almost full PCP string from small PCP fragments revealed in each (accepting) execution.

Succinct arguments from public-coin IOPs. Interaction often leads to dramatic gains in efficiency for proof systems. PCPs are no exception: *interactive oracle proofs* (IOPs) enable such gains by communicating oracle proof strings across multiple rounds. (In particular, a PCP system is an IOP with a single round.)

Let IOP be an IOP system for a relation R with round complexity k and soundness error ϵ_{IOP} (resp. knowledge soundness error κ_{IOP}). Owing to its multiple proof strings, the proof length and query complexity in an IOP may vary round by round, and we let l_{max} and q_{max} denote the maximum proof length and query complexity and let l and q denote the total proof length and query complexity.²

Our second main result provides a (concrete) security analysis for the *interactive BCS* protocol, inspired by [BCS16], which extends Kilian’s protocol from the case of PCPs to the case of *public-coin* IOPs. (An IOP is public-coin if each verifier message is a uniformly random string.)

Theorem 2 (informal). *If IOP is public-coin, the soundness error ϵ_{ARG} of IBCS[IOP, VC] satisfies the following for every security parameter λ , instance $\mathbb{x} \notin L(R)$, adversary size bound t_{ARG} , and error tolerance $\epsilon > 0$:*

$$\epsilon_{\text{ARG}}(\lambda, \mathbb{x}, t_{\text{ARG}}) \leq \epsilon_{\text{IOP}}(\mathbb{x}) + \epsilon_{\text{VC}}(\lambda, l_{\text{max}}(\mathbb{x}), q_{\text{max}}(\mathbb{x}), t_{\text{VC}}) + \epsilon, \text{ where } t_{\text{VC}} = O\left(\frac{k \cdot l}{\epsilon} \cdot t_{\text{ARG}}\right).$$

Similarly, the knowledge soundness error is $\kappa_{\text{ARG}}(\lambda, \mathbb{x}, t_{\text{ARG}}) \leq \kappa_{\text{IOP}}(\mathbb{x}) + \epsilon_{\text{VC}}(\lambda, l_{\text{max}}(\mathbb{x}), q_{\text{max}}(\mathbb{x}), t_{\text{VC}}) + \epsilon$.

In particular, if ϵ_{VC} is negligible when $t_{\text{VC}} = \text{poly}(\lambda)$, then $\epsilon_{\text{ARG}}(\lambda, \mathbb{x}, \text{poly}(\lambda)) \leq \epsilon_{\text{IOP}}(\mathbb{x}) + \text{negl}(\lambda)$ (and similarly for κ_{ARG}).

²If the IOP verifier, in round i , makes q_i queries to a proof string of length l_i , then $l_{\text{max}} := \max_i \{l_i\}$ and $q_{\text{max}} := \max_i \{q_i\}$, and $l := \sum_i l_i$ and $q := \sum_i q_i$.

The above parameters are analogous to those in Theorem 1 for PCPs, except for a dependence on the round complexity k of the IOP. Again, the error is essentially tight up to the loss in t_{VC} compared to t_{ARG} .

How general is the VC-based approach? Kilian’s protocol realizes query access to a PCP string via a vector commitment scheme: the values of positions queried by the PCP verifier are revealed alongside an opening proof attesting consistency with the commitment. In the IBCS protocol, this VC-based approach is applied round by round to a public-coin IOP. In general, however, an IOP may not be public-coin. We ask: *which IOPs can be compiled into a secure interactive argument using VCs?*

It is not hard to see that there exist private-coin IOPs which do not yield sound proof systems under the “natural” generalization of IBCS.³ For example, consider a 2-round IOP where the verifier, after receiving the first proof string from the prover, queries the first proof string at a few locations in order to determine its first message to the prover but *must keep the queried locations secret* from the prover in order to ensure soundness.⁴ In the VC-based approach, the argument verifier sends the desired queries to the argument prover in order to receive corresponding answers and opening proofs. A malicious argument prover would then be able to use this information to break IOP soundness when choosing which second proof string to commit to.

The case of public-query IOPs. The above discussion implies that a *necessary* condition for an IOP to be compatible with the VC-based approach is being **public-query** (a notion that we introduce): soundness (or knowledge soundness) of the IOP holds even if, whenever the verifier makes a query, the prover learns that this particular query has been made. (Of course, the IOP verifier’s own randomness remains secret.)

A public-coin IOP is, in particular, public-query. One may conjecture that the VC-based approach works not only for every public-coin IOP but also for every public-query IOP. Indeed, we describe a natural protocol that we call $\text{Finale}[\text{IOP}, \text{VC}]$ that, intuitively, ought to be secure for every public-query IOP.

However, challenges arise when we attempt to prove security. The “rewinding” proof of Theorem 2 crucially depends on the existence of an efficient algorithm for sampling IOP verifier randomness *conditioned on a partial transcript of interaction* (which includes prior verifier messages and queries). We call this algorithm a *random continuation sampler* (RCS). Public-coin IOPs have a (trivial) RCS, but there exist public-query IOPs that do not. (E.g., consider an IOP where the verifier derives its i -th message as $\text{PRG}(r_i)$ for secret randomness r_i .)

We show that if IOP possesses an efficient RCS then we can prove security of $\text{Finale}[\text{IOP}, \text{VC}]$ (for any VC). The proof proceeds in two steps: first, we show that if IOP has an RCS then $\text{Finale}[\text{IOP}, \text{VC}]$ has an RCS; then, we show that if $\text{Finale}[\text{IOP}, \text{VC}]$ has an RCS then it is (knowledge) sound. We also show a partial converse: if $\text{Finale}[\text{IOP}, \text{VC}]$ has an RCS, then so does IOP. This suggests that known proof approaches will not suffice to show the security of $\text{Finale}[\text{IOP}, \text{VC}]$ for general public-query IOPs.

Theorem 3 (informal). *Let IOP be a public-query IOP that admits a RCS with running time t_S . The soundness error ϵ_{ARG} of $\text{Finale}[\text{IOP}, \text{VC}]$ satisfies the following for every security parameter λ , instance $\mathbb{x} \notin L(R)$, adversary size bound t_{ARG} , and error tolerance $\epsilon > 0$:*

$$\epsilon_{ARG}(\lambda, \mathbb{x}, t_{ARG}) \leq \epsilon_{IOP}(\mathbb{x}) + \epsilon_{VC}(\lambda, l_{\max}, q_{\max}, t_{VC}) + \epsilon, \text{ where } t_{VC} = O\left(\frac{k \cdot l}{\epsilon} \cdot (t_{ARG} + t_S)\right).$$

Similarly, the knowledge soundness error is $\kappa_{ARG}(\lambda, \mathbb{x}, t_{ARG}) \leq \kappa_{IOP}(\mathbb{x}) + \epsilon_{VC}(\lambda, l_{\max}, q_{\max}, t_{VC}) + \epsilon$.

Moreover, $\text{Finale}[\text{IOP}, \text{VC}]$ admits an RCS if and only if IOP admits an RCS.

We actually prove a stronger version of Theorem 3: our analysis works for every public-query IOP that admits a RCS with an arbitrary sampling error α , in which case the (knowledge) soundness error bound of $\text{Finale}[\text{IOP}, \text{VC}]$ depends on both ϵ and α . (See Section 6 for the technical details of this case.)

³Of course, other approaches (e.g., using additional cryptography beyond VC schemes) could in principle work with every IOP.

⁴For example, one can modify a sound IOP so that the verifier always accepts if the prover guesses the query locations.

Remark 1 (adaptive choice of \mathbb{x}). The results of Theorems 1, 2 and 3 are stated, for simplicity, in the plain model (no trusted setups), where the argument verifier is responsible for sampling and sending to the argument prover public parameters pp for VC. However, we actually *prove* these results in the (adaptive) common reference string model, wherein public parameters pp for VC are sampled by a trusted party and *a malicious argument prover may adaptively choose the instance \mathbb{x} after learning pp* . Since in these stronger theorems there is no pre-set instance \mathbb{x} , the analogous statements (for corresponding security properties) in the common reference model replace \mathbb{x} with a size bound n (and hold for all instances such that $|\mathbb{x}| \leq n$). The plain model variants are straightforwardly implied (see Section 2.4 and Remark 3.6).

1.2 Discussion

Limitations of asymptotic security. The standard asymptotic notion of computational soundness for an interactive argument system is as follows:⁵ for every $a, b, c \in \mathbb{N}$ there exists a threshold $\lambda_0 \in \mathbb{N}$ such that for every security parameter $\lambda \in \mathbb{N}$ with $\lambda \geq \lambda_0$, malicious argument prover $\tilde{\mathcal{P}}$ of size at most λ^a , and instance $\mathbb{x} \notin L(R)$ of size at most λ^b ,

$$\Pr \left[\langle \tilde{\mathcal{P}}, \mathcal{V}(1^\lambda, \mathbb{x}) \rangle = 1 \right] \leq \lambda^{-c} .$$

For deployments in the real world, where we must set λ to achieve a desired security level, such security guarantees are not useful without a concrete estimate of λ_0 in terms of a, b, c (with explicit constants).

Setting parameters via concrete security bounds. Alternatively, consider the *concrete* bound of Theorem 1 for Kilian’s protocol: for every malicious argument prover $\tilde{\mathcal{P}}$ of size t_{ARG} and instance $\mathbb{x} \notin L(R)$,

$$\Pr \left[\langle \tilde{\mathcal{P}}, \mathcal{V}(1^\lambda, \mathbb{x}) \rangle = 1 \right] \leq \epsilon_{\text{PCP}}(\mathbb{x}) + \epsilon_{\text{VC}}(\lambda, l(\mathbb{x}), \mathbf{q}(\mathbb{x}), t_{\text{VC}}) + \epsilon,$$

where $t_{\text{VC}} \leq 4 \cdot \frac{1}{\epsilon} \cdot t_{\text{ARG}}$. (The constant 4 is what our analysis in the technical sections shows.)

Say that we are targeting a soundness error of 2^{-40} against adversaries of size $t_{\text{ARG}} = 2^{60}$. Suppose that the PCP underlying Kilian’s protocol, for the chosen instance size, achieves soundness error $\epsilon_{\text{PCP}} = 2^{-42}$ with a proof consisting of 2^{30} symbols. Further suppose that the vector commitment scheme has binding error $\epsilon_{\text{VC}}(\lambda, \ell, s, t_{\text{VC}}) \leq t_{\text{VC}}^2 / 2^\lambda$ (the bound achieved by an ideal Merkle tree). Setting $\epsilon = 2^{-42}$ results in a bound of

$$t_{\text{VC}} \leq 4 \cdot \frac{2^{30}}{2^{-42}} \cdot t_{\text{ARG}} < 2^{80} \cdot t_{\text{ARG}}$$

so

$$\epsilon_{\text{VC}} < \frac{(2^{80} \cdot t_{\text{ARG}})^2}{2^\lambda} = 2^{160-\lambda} \cdot t_{\text{ARG}}^2 = 2^{280-\lambda} .$$

Then Theorem 1 indicates that we can set $\lambda = 322$ to achieve the desired security level (soundness error of 2^{-40} against adversaries of size 2^{60}). Observe that this is noticeably larger than what would be required for equivalent security of the underlying collision-resistant hash function (there, $\lambda = 160$ would suffice) because of losses arising from (the rewinding in) the security analysis. Our results precisely quantify these losses.

Of course, known PCPs have poor concrete efficiency; in practice, one would use a public-coin IOP. The concrete security bound in (the detailed analysis underlying) Theorem 2 for the IBCS protocol enables reasoning similarly to the above for the case of public-coin IOPs.

⁵Once we open notions such as “bounded by a polynomial” and “bounded by a negligible function”.

Remark 1.1 (concrete security of components). We derive concrete security bounds for argument systems *as a function of the concrete security bounds of the underlying components*. In short, we take ϵ_{VC} , ϵ_{PCP} , κ_{PCP} (and ϵ_{IOP} , κ_{IOP}) as given. While statistical soundness bounds on PCPs and IOPs can be calculated (they are information-theoretic components), the position binding errors for VC must be derived from some (concrete) computational assumption. For example, if VC is a Merkle tree obtained from a collision-resistant hash function $h_\lambda: \{0, 1\}^{2\lambda} \rightarrow \{0, 1\}^\lambda$ computable in time t_{h_λ} whose collision probability against t_{CRH} -size adversaries is bounded by $\epsilon_{CRH}(\lambda, t_{CRH})$ then VC has binding error $\epsilon_{VC}(\lambda, \ell, s, t_{VC}) \leq \epsilon_{CRH}(\lambda, t_{CRH})$ where $t_{CRH} = t_{VC} + O(t_{h_\lambda} \cdot q \cdot \log l)$ for a small hidden constant that can be derived from the security reduction. (The reduction transforms a t_{VC} -size adversary A_{VC} against the Merkle tree into a t_{CRH} -size adversary A_{CRH} against the collision-resistant hash function. Briefly, A_{CRH} runs A_{VC} and then looks for a collision among the authentication paths output by A_{VC} , resulting in the additive increase of $O(t_{h_\lambda} \cdot q \cdot \log l)$ in size.)

On the price of rewinding. We compare the soundness of Kilian’s protocol when analyzed via: (i) a rewinding extractor based on a collision resistant hash function; or (ii) a straightline extractor based on an ideal hash function (a random oracle). This comparison makes explicit the price of rewinding: the cost of a more expensive security reduction that works while assuming “less” about the underlying cryptography.

- (i) *Rewinding extractor.* Suppose that the vector commitment scheme VC is obtained from a collision-resistant hash function (via a Merkle tree) with security $\epsilon_{CRH}(\lambda, t_{CRH})$. By Remark 1.1 this means that

$$\epsilon_{VC}(\lambda, \ell, s, t_{VC}) \leq \epsilon_{CRH}\left(\lambda, t_{CRH} = t_{VC} + O(t_{h_\lambda} \cdot q \cdot \log l)\right) .$$

Suppose that $\epsilon_{CRH}(\lambda, t_{CRH}) \leq t_{CRH}^2/2^\lambda$, which is what would be achieved by an ideal hash function. In this case, Theorem 1 gives the following upper bound on the soundness error for Kilian[PCP, VC]:

$$\epsilon_{ARG}(\lambda, \mathbb{X}, t_{ARG}) \leq \epsilon_{PCP}(\mathbb{X}) + O\left(\frac{1}{2^\lambda} \cdot \left(\frac{1}{\epsilon} \cdot t_{ARG} + t_{h_\lambda} \cdot q \cdot \log l\right)^2\right) + \epsilon .$$

In particular, this bound cannot be smaller than $\epsilon_{PCP}(\mathbb{X}) + \Theta(l \cdot t_{ARG} \cdot 2^{-\lambda/2})$.

- (ii) *Straightline extractor.* Suppose that we model the collision-resistant hash function as an ideal hash function, and analyze Kilian[PCP, VC] in the random oracle model. Prior work analyzing such protocols [BCS16] implies the following upper bound on soundness error:

$$\epsilon_{ARG}(\lambda, \mathbb{X}, t_{ARG}) \leq \epsilon_{PCP}(\mathbb{X}) + O\left(\frac{1}{2^\lambda} \cdot t_{ARG}^2\right) .$$

This smaller upper bound is achieved thanks to a straightline (i.e., non-rewinding) extractor for the vector commitment scheme (which is a Merkle tree in the random oracle model).

Improved understanding of [BCS16]. Theorem 2 is a step towards a fuller understanding of the (non-interactive) BCS protocol. While [BCS16] gives a detailed concrete security analysis of the BCS protocol as a whole, they crucially consider only the specific setting where the VC is implemented as a Merkle tree *and* the hash function used in the tree is modeled as a random oracle. It is believed that the BCS protocol is secure when realized with any vector commitment (in particular, with VCs in the standard model) but this has not been shown thus far. A natural way to prove this would be to combine (an appropriate strengthening of) Theorem 2 with an analysis of the Fiat–Shamir transformation applied to the resulting argument.⁶

⁶The appropriate strengthening refers to showing that IBCS is a succinct interactive argument that satisfies (a suitable computational notion of) state-restoration soundness (which makes the interactive argument compatible with the Fiat–Shamir transformation) if the underlying IOP also satisfies state-restoration soundness.

A reflection on succinct arguments. Succinct arguments are a rare example of an “advanced” cryptographic primitive that can be achieved from simple cryptography. Indeed, it is remarkable that, based solely on the existence of a collision resistant hash function (even given as a black box), one can achieve cryptographic proof systems with such exceptional efficiency. On the other hand, the security reduction of a succinct argument is tasked with a challenging goal: find a “long” witness when given a malicious argument prover that only outputs “short” messages in any given interaction. This naturally leads to *rewinding*, a fundamental method of analysis in cryptography. In light of this, Kilian’s protocol occupies a central place in cryptography: the simplest succinct argument, and its security analysis is a prominent example of extracting a long witness from sufficiently many short messages obtained via rewinding. This work contributes an overdue concrete security analysis of a central cryptographic paradigm.

The preprocessing case. We omit discussion of the *preprocessing setting* for succinct arguments, and believe that future work can build on our security analyses to cover this case. Informally, the preprocessing setting is an offline-online model that enables succinct verification even for “non-uniform” computations. The canonical way to construct succinct arguments in the preprocessing setting is to combine a *holographic* probabilistic proof and a vector commitment scheme [CHMMVW20; COS20] — this is a direct generalization of the VC-based approach that we study in this paper. Specifically, each protocol that we study in this paper (Kilian’s protocol, the IBCS protocol, the Finale protocol) has a straightforward extension to the preprocessing setting, where the vector commitment scheme is also used in the offline computation phase to commit to the holographic part of the probabilistic proof. We leave to future work extending our analyses to cover this case.

1.3 Related work

A vast literature studies many aspects of succinct arguments in different settings. Below we summarize only the most relevant prior work: security analyses for succinct arguments constructed from hash functions (more generally, vector commitment schemes).

Succinct arguments from collision-resistant functions. The first construction of a succinct argument is due to Kilian [Kil92], and is based on collision-resistant hash functions. The security reduction in [Kil92] is informal, and does not provide any asymptotic (nor concrete) security bounds.

Barak and Goldreich [BG08] provide a formal analysis of a variant of Kilian’s construction, towards their goal of constructing zero-knowledge arguments with a non-black-box simulator. Due to their setting, they restrict their result to the case where the PCP is *non-adaptive* and *reverse-samplable*. While the former restriction is mild (many known PCP constructions are non-adaptive, with few exceptions such as [KPT97]), the latter restriction is a non-standard strong property of the query algorithm, which has not been shown to hold for many PCP constructions of interest (e.g. the short PCPs in [BS06; BKKMS13]). Under these conditions, they establish that Kilian’s protocol achieves non-adaptive knowledge soundness, with a constant factor loss in soundness versus the PCP. In contrast, our analysis applies to *all* PCPs (including adaptive PCPs) and establishes tight adaptive knowledge soundness.⁷

Lai and Malavolta [LM19, Appendix C] prove that Kilian’s protocol is secure when realized with any PCP and vector commitment. In fact, they prove a more general result: a variant of Kilian’s protocol is secure when realized with any *linear PCP* and *linear map commitment*. This generality can lead to shorter proofs. However, they prove only *asymptotic* security; their concrete security estimates are heuristic.

⁷Formally, our result is incomparable with the one of Barak and Goldreich, as their knowledge extractor has the additional feature of providing query access to the extracted PCP. This is necessary in their setting, as the extracted PCP may have superpolynomial size. In contrast, our extractor outputs the extracted PCP in full; this suffices when the size of the PCP is bounded by a fixed polynomial, as is the case in most applications.

Chiesa, Ma, Spooner, and Zhandry [CMSZ21] prove *post-quantum* security of Kilian’s protocol. As part of their analysis, they give a proof of security for Kilian that also applies to the classical setting. Their analysis differs significantly from ours due to challenges unique to the quantum setting, and incurs a multiplicative soundness loss. In this work we consider soundness against classical adversaries only.

Several works [BCG20; RR22; HR22] obtain succinct arguments with very strong efficiency properties by applying the IBCS protocol to their highly-efficient IOP constructions (along with a Merkle tree based on linear-time CRHFs). They do not show soundness of the IBCS protocol itself (their focus is on novel IOPs).

Succinct arguments from ideal hash functions. A line of work studies security reductions for succinct arguments in the random oracle model (ROM) [Mic00; Val08; BCS16; CMS19; CY21a; CY21b; BGTZ23]. The focus of these works is on *non-interactive* succinct arguments, and benefits from the fact that the security reduction can observe the adversary’s queries to the random oracle. This enables, among other things, *straightline* (non-rewinding) extraction, which is highly efficient. Moreover, since these constructions are *unconditionally* secure in the ROM, it is possible to give very precise soundness bounds.

In contrast, our understanding of succinct arguments with standard-model security (the subject of this work) is much more limited. It is known that succinct *non-interactive* arguments for NP cannot be proven secure in the standard model via a black-box security reduction to a falsifiable assumption [GW11]. As such, we focus on succinct interactive arguments. Without access to the “queries” the adversary makes to the hash function, we must rely on a rewinding reduction, which leads to a quantitative loss in security.

Special-sound protocols. Interactive protocols with *special soundness* are an important and well-studied family of public-coin protocols. In the three-message public-coin (Σ -protocol) setting, k -special soundness means that a witness can be efficiently extracted from any k accepting protocol transcripts with distinct verifier challenges. A line of works extends this notion to multiple rounds [AC20; ACK22; AF22]. The concrete security of general special sound protocols is relatively well-understood.

However, as observed in [CMSZ21], for reasonable choices of PCP, Kilian’s protocol is *not* k -special sound for any polynomial k (for example, one can find a set of transcripts that includes only queries to a small fraction of the PCP).⁸ We are therefore not able to apply results about special soundness directly.

In concurrent work, Attema, Fehr, and Resch [AFR23] relax the notion of special soundness further: for an access structure Γ , a protocol is Γ -special sound if there is an efficient procedure to compute a witness from any set of transcripts whose challenges form an authorized set in Γ . They use this framework to show that the IBCS protocol applied to the FRI protocol (an IOP of proximity for Reed–Solomon codes [BBHR18]) is a proof of knowledge of a committed codeword (with a quasipolynomial-time extractor). In contrast, rather than analyzing any specific instantiation, our goal is to give a *general* security guarantee for the Kilian (resp. IBCS) protocol, which applies to any PCP (resp. IOP). It seems that Γ -special soundness may not be the right approach for such generality, as Γ will depend on the underlying PCP/IOP. Our guarantee depends, of course, on the *soundness* of the PCP/IOP, which must be analyzed separately.

⁸Towards a tighter security proof for Kilian in the post-quantum setting, Lombardi, Ma, and Spooner [LMS22] introduce the notion of *probabilistic special soundness* (PSS), a relaxation of special soundness, and show that Kilian’s protocol is PSS. We do not follow this approach, as we do not expect it to yield tight concrete security bounds in the classical setting.

2 Techniques

We overview the main ideas underlying our results. In Section 2.1 we review Kilian’s protocol and sketch our security analysis for it. In Section 2.2 we review the IBCS protocol and discuss our security analysis for it, comparing it to that for Kilian’s protocol. In Section 2.3 we overview the Finale protocol (an extension of the IBCS protocol that is formulated for *public-query* IOPs) and discuss capabilities and limitations of the VC-based approach for succinct interactive arguments. In Section 2.4 we discuss adaptive security.

Vector commitment schemes. We fix a vector commitment scheme VC throughout this technical overview, whose interface and properties are sketched below; see Section 3.2 for formal definitions. Here we omit the algorithm that samples public parameters (and suppress these parameters in the interfaces of VC).⁹

- VC.Commit: On input a message m , VC.Commit outputs a commitment cm and auxiliary state aux .
- VC.Open: On input the auxiliary state aux and a query set \mathcal{Q} , VC.Open outputs an opening proof pf .
- VC.Check: On input a commitment cm , query set \mathcal{Q} , answers ans , and opening proof pf , VC.Check determines if pf is valid for ans being the restriction to \mathcal{Q} of the message committed in cm .

The property of *perfect completeness* ensures that VC.Check always accepts if pf is output by VC.Open given the auxiliary information produced by VC.Commit. The security property of VC is *position binding*: VC has *position binding error* $\epsilon_{\text{VC}}(\lambda, \ell, s, t_{\text{VC}})$ if, when VC is instantiated with security parameter λ for messages of length ℓ , every adversary of size t_{VC} that outputs $(\text{cm}, \text{ans}, \text{ans}', \mathcal{Q}, \mathcal{Q}', \text{pf}, \text{pf}')$ with $|\mathcal{Q}| = |\mathcal{Q}'| = s$ satisfies the following predicate with probability at most $\epsilon_{\text{VC}}(\lambda, \ell, s, t_{\text{VC}})$ (over VC’s public parameters):

$$\begin{aligned} & \exists i \in \mathcal{Q} \cap \mathcal{Q}' : \text{ans}[i] \neq \text{ans}'[i] \\ & \wedge \text{VC.Check}(\text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \\ & \wedge \text{VC.Check}(\text{cm}, \mathcal{Q}', \text{ans}', \text{pf}') = 1 \end{aligned} .$$

In other words, position binding makes it hard to produce two incompatible openings to the same commitment.

Stateful algorithms. Throughout this section, the interactive algorithms that participate in protocols are stateful. When it is important to distinguish different computation phases of a stateful algorithm, we make explicit the state passed from one phase to the next.

2.1 Succinct arguments based on PCPs

We review Kilian’s protocol and sketch the main ideas behind Theorem 1; see Section 4 for details.

2.1.1 Kilian’s protocol

Kilian’s protocol [Kil92] obtains a succinct interactive argument by combining two ingredients: a probabilistically checkable proof (PCP) and a vector commitment scheme VC (fixed above). Let $\text{PCP} = (\mathbf{P}, \mathbf{V})$ be a PCP system for a relation R with alphabet Σ , proof length l , query complexity q , and verifier randomness complexity r . Kilian[PCP, VC] is an interactive argument $\text{ARG} = (\mathcal{P}, \mathcal{V})$ in which the argument prover \mathcal{P} receives an instance \mathfrak{x} and a witness \mathfrak{w} , and the argument verifier \mathcal{V} receives the instance \mathfrak{x} . Then \mathcal{P} and \mathcal{V} interact, exchanging 3 messages, as follows.

1. \mathcal{P} computes the PCP string $\Pi \leftarrow \mathbf{P}(\mathfrak{x}, \mathfrak{w})$, computes the commitment $(\text{cm}, \text{aux}) \leftarrow \text{VC.Commit}(\Pi)$, and sends cm to \mathcal{V} .

⁹For example, if VC is based on a Merkle tree, the public parameters are the (randomly sampled) collision-resistant function to be used for hashing the given message down to the Merkle root.

2. \mathcal{V} samples PCP verifier randomness $\rho \leftarrow \{0, 1\}^r$ and sends it to \mathcal{P} .
3. \mathcal{P} deduces the set \mathcal{Q} of queries that $\mathbf{V}(\mathbb{x}; \rho)$ makes to Π , sets the query answers $\text{ans} := \Pi[\mathcal{Q}]$, generates an opening proof $\text{pf} \leftarrow \text{VC.Open}(\text{aux}, \mathcal{Q})$, and sends the tuple $(\mathcal{Q}, \text{ans}, \text{pf})$ to \mathcal{V} .
4. \mathcal{V} performs the following checks.
 - (a) $\text{VC.Check}(\text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1$ (i.e., ans are valid answers for positions \mathcal{Q} relative to cm);
 - (b) $\mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbb{x}; \rho) = 1$ (i.e., the PCP verifier $\mathbf{V}(\mathbb{x}; \rho)$ accepts the answers ans on \mathcal{Q}).

Above, the notation $\mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbb{x}; \rho)$ refers to the decision bit of the PCP verifier \mathbf{V} , given instance \mathbb{x} and PCP randomness ρ , when each query $j \in \mathcal{Q}$ is answered with $\text{ans}[j] \in \Sigma$. (If \mathbf{V} queries outside the set \mathcal{Q} then $\mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbb{x}; \rho) = 0$.)

2.1.2 Security reduction

Intuitively, the soundness error of Kilian[PCP, VC] should be at most the (statistical) soundness error of PCP plus the position binding error of VC. The key lemma below formalizes this intuition.

Consider a malicious argument prover $\tilde{\mathcal{P}}$ whose first message is the commitment cm. Intuitively, by the position binding property of VC, $\tilde{\mathcal{P}}$ is “bound” to open locations of at most a single underlying PCP string $\tilde{\Pi}$. By *rewinding* $\tilde{\mathcal{P}}$ sufficiently many times to recover the underlying PCP string $\tilde{\Pi}$, we can relate the probability of $\tilde{\mathcal{P}}$ convincing the argument verifier \mathcal{V} to the probability of $\tilde{\Pi}$ convincing the PCP verifier \mathbf{V} .

Lemma 1 (informal). *There exists a probabilistic algorithm \mathcal{R} (the **reductor**) that, for every instance \mathbb{x} , error parameter $\epsilon > 0$, adversary size bound $t_{\text{ARG}} \in \mathbb{N}$, and t_{ARG} -size adversary $\tilde{\mathcal{P}}$, satisfies*

$$\Pr \left[\begin{array}{l} \mathbf{V}^{[\tilde{\mathcal{Q}}, \tilde{\Pi}]}(\mathbb{x}; \rho) \neq 1 \\ \wedge \mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbb{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \middle| \begin{array}{l} \text{cm} \leftarrow \tilde{\mathcal{P}} \\ (\tilde{\mathcal{Q}}, \tilde{\Pi}) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}}(\text{cm}, \epsilon) \\ \rho \leftarrow \{0, 1\}^r \\ (\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\rho) \end{array} \right] \leq \epsilon_{\text{VC}}(\lambda, l, \mathbf{q}, t_{\text{VC}}) + \epsilon ,$$

where $t_{\text{VC}} = O\left(\frac{1}{\epsilon} \cdot t_{\text{ARG}}\right)$.

The reductor \mathcal{R} handles the aforementioned rewinding process. It constructs a proof string $\tilde{\Pi} \in \Sigma^l$ whose convincing probability is approximately the same as that of the argument prover (up to the position binding error of VC and an arbitrary error term ϵ). Note that \mathcal{R} requires only black-box access to $\tilde{\mathcal{P}}$.

In the lemma above, the PCP verifier and the argument verifier are “coupled” in that they receive the same randomness ρ . The lemma states that it is unlikely, for a randomly-chosen ρ , that the argument verifier \mathcal{V} accepts the answers provided by $\tilde{\mathcal{P}}$ but the PCP verifier \mathbf{V} rejects $\tilde{\Pi}$ under the same randomness. Intuitively, this will allow us to approximately equate the probability that $\tilde{\mathcal{P}}$ convinces the argument verifier \mathcal{V} to the probability that $\tilde{\Pi}$ convinces the PCP verifier \mathbf{V} .

First we discuss how to use Lemma 1 to establish soundness and knowledge soundness of Kilian[PCP, VC] in Sections 2.1.3 and 2.1.4. Then in Section 2.1.5 we sketch the proof of Lemma 1.

2.1.3 Soundness analysis

We wish to upper bound the soundness error of Kilian[PCP, VC]. As claimed in Theorem 1, we argue that for every instance $\mathbb{x} \notin L(R)$, size bound $t_{\text{ARG}} \in \mathbb{N}$, and t_{ARG} -size adversary $\tilde{\mathcal{P}}$,

$$\Pr \left[\langle \tilde{\mathcal{P}}, \mathcal{V}(\mathbb{x}) \rangle = 1 \right] \leq \epsilon_{\text{PCP}}(\mathbb{x}) + \epsilon_{\text{VC}}(\lambda, l, \mathbf{q}, t_{\text{VC}}) + \epsilon .$$

By construction of the argument verifier \mathcal{V} , the above probability is equivalent to the following:

$$\Pr \left[\begin{array}{l} \mathbf{V}^{[Q, \text{ans}]}(\mathbb{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{cm}, Q, \text{ans}, \text{pf}) = 1 \end{array} \middle| \begin{array}{l} \text{cm} \leftarrow \tilde{\mathcal{P}} \\ \rho \leftarrow \{0, 1\}^r \\ (Q, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\rho) \end{array} \right] .$$

Below we fix the following experiment (which is the experiment above augmented with an invocation of \mathcal{R}):

$$\left[\begin{array}{l} \text{cm} \leftarrow \tilde{\mathcal{P}} \\ (\tilde{Q}, \tilde{\Pi}) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}}(\text{cm}, \epsilon) \\ \rho \leftarrow \{0, 1\}^r \\ (Q, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\rho) \end{array} \right]$$

Using the law of total probability,

$$\begin{aligned} & \Pr \left[\begin{array}{l} \mathbf{V}^{[Q, \text{ans}]}(\mathbb{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{cm}, Q, \text{ans}, \text{pf}) = 1 \end{array} \right] \\ &= \Pr \left[\begin{array}{l} \mathbf{V}^{[\tilde{Q}, \tilde{\Pi}]}(\mathbb{x}; \rho) = 1 \\ \wedge \mathbf{V}^{[Q, \text{ans}]}(\mathbb{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{cm}, Q, \text{ans}, \text{pf}) = 1 \end{array} \right] + \Pr \left[\begin{array}{l} \mathbf{V}^{[\tilde{Q}, \tilde{\Pi}]}(\mathbb{x}; \rho) \neq 1 \\ \wedge \mathbf{V}^{[Q, \text{ans}]}(\mathbb{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{cm}, Q, \text{ans}, \text{pf}) = 1 \end{array} \right] . \end{aligned}$$

The term on the right is bounded from above by $\epsilon_{\text{VC}}(\lambda, l, q, t_{\text{VC}}) + \epsilon$, due to Lemma 1.

The term on the left is bounded by $\epsilon_{\text{PCP}}(\mathbb{x})$ (the soundness error of PCP). Indeed, we can view the first message of $\tilde{\mathcal{P}}$ (cm in the experiment above) and the reductor \mathcal{R} as a malicious PCP prover $\tilde{\mathbf{P}}$ that outputs a PCP string $\tilde{\Pi}$:

$\tilde{\mathbf{P}}$:

1. Run $\text{cm} \leftarrow \tilde{\mathcal{P}}$.
2. Run $(\tilde{Q}, \tilde{\Pi}) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}}(\text{cm}, \epsilon)$.
3. Output $\tilde{\Pi}$.

Since $\mathbb{x} \notin L(R)$, by the definition of soundness error of PCP,

$$\Pr \left[\begin{array}{l} \mathbf{V}^{[\tilde{Q}, \tilde{\Pi}]}(\mathbb{x}; \rho) = 1 \\ \wedge \mathbf{V}^{[Q, \text{ans}]}(\mathbb{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{cm}, Q, \text{ans}, \text{pf}) = 1 \end{array} \right] \leq \Pr \left[\mathbf{V}^{\tilde{\Pi}}(\mathbb{x}; \rho) = 1 \middle| \begin{array}{l} \rho \leftarrow \{0, 1\}^r \\ \tilde{\Pi} \leftarrow \tilde{\mathbf{P}} \end{array} \right] \leq \epsilon_{\text{PCP}}(\mathbb{x}) .$$

2.1.4 Knowledge soundness analysis

We wish to upper bound the knowledge soundness error of Kilian[PCP, VC]. As claimed in Theorem 1, we argue that there exists a polynomial-time probabilistic extractor \mathcal{E} such that for every instance \mathbb{x} , size bound $t_{\text{ARG}} \in \mathbb{N}$, and t_{ARG} -size adversary $\tilde{\mathcal{P}}$,

$$\Pr \left[\begin{array}{l} b = 1 \\ \wedge (\mathbb{x}, \mathbb{w}) \notin R \end{array} \middle| \begin{array}{l} b \leftarrow \langle \tilde{\mathcal{P}}, \mathcal{V}(\mathbb{x}) \rangle \\ \mathbb{w} \leftarrow \mathcal{E}^{\tilde{\mathcal{P}}}(\mathbb{x}) \end{array} \right] \leq \kappa_{\text{PCP}}(\mathbb{x}) + \epsilon_{\text{VC}}(\lambda, l, q, t_{\text{VC}}) + \epsilon .$$

By construction of the argument verifier \mathcal{V} , the above probability is equivalent to the following:

$$\Pr \left[\begin{array}{l} \mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbf{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \\ \wedge (\mathbf{x}, \mathbf{w}) \notin R \end{array} \middle| \begin{array}{l} \text{cm} \leftarrow \tilde{\mathcal{P}} \\ \rho \leftarrow \{0, 1\}^r \\ (\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\rho) \\ \mathbf{w} \leftarrow \mathcal{E}^{\tilde{\mathcal{P}}}(\mathbf{x}) \end{array} \right] .$$

We construct \mathcal{E} using the PCP prover $\tilde{\mathbf{P}}$ described in Section 2.1.3 and the PCP extractor \mathbf{E} (which is given by the underlying PCP system):

$\mathcal{E}^{\tilde{\mathcal{P}}}(\mathbf{x})$:

1. Run $\tilde{\Pi} \leftarrow \tilde{\mathbf{P}}$.
2. Output $\mathbf{w} \leftarrow \mathbf{E}(\mathbf{x}, \tilde{\Pi})$.

Using the law of total probability,

$$\begin{aligned} & \Pr \left[\begin{array}{l} \mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbf{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \\ \wedge (\mathbf{x}, \mathbf{w}) \notin R \end{array} \middle| \begin{array}{l} \text{cm} \leftarrow \tilde{\mathcal{P}} \\ \rho \leftarrow \{0, 1\}^r \\ (\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\rho) \\ \mathbf{w} \leftarrow \mathcal{E}^{\tilde{\mathcal{P}}}(\mathbf{x}) \end{array} \right] \\ &= \Pr \left[\begin{array}{l} \mathbf{V}^{[\tilde{\mathcal{Q}}, \tilde{\Pi}]}(\mathbf{x}; \rho) = 1 \\ \wedge \mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbf{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \\ \wedge (\mathbf{x}, \mathbf{w}) \notin R \end{array} \right] + \Pr \left[\begin{array}{l} \mathbf{V}^{[\tilde{\mathcal{Q}}, \tilde{\Pi}]}(\mathbf{x}; \rho) \neq 1 \\ \wedge \mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbf{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \\ \wedge (\mathbf{x}, \mathbf{w}) \notin R \end{array} \right] , \end{aligned}$$

where the last two probabilities are with respect to

$$\left[\begin{array}{l} \text{cm} \leftarrow \tilde{\mathcal{P}} \\ \rho \leftarrow \{0, 1\}^r \\ (\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\rho) \\ \tilde{\Pi} \leftarrow \tilde{\mathbf{P}} \\ \mathbf{w} \leftarrow \mathbf{E}(\mathbf{x}, \tilde{\Pi}) \end{array} \right] .$$

The term on the right is bounded by $\epsilon_{\text{VC}}(\lambda, l, q, t_{\text{VC}}) + \epsilon$ due to Lemma 1.

The term on the left is bounded by $\kappa_{\text{PCP}}(\mathbf{x})$ (the knowledge soundness error of PCP) as shown below:

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

2.1.5 Proof sketch of Lemma 1

We are left to sketch the proof of Lemma 1. To do so, we present a reducer algorithm \mathcal{R} .

The goal of \mathcal{R} is to piece together a PCP string $\tilde{\Pi}$ obtained from the argument prover $\tilde{\mathcal{P}}$. Intuitively, $\tilde{\Pi}$ is “fixed” after $\tilde{\mathcal{P}}$ outputs a commitment cm , and \mathcal{R} attempts to obtain information about $\tilde{\Pi}$ by *rewinding* the second phase of $\tilde{\mathcal{P}}$, when given freshly sampled choices of PCP randomness ρ . Each such execution (if it outputs a valid opening) reveals a fragment of $\tilde{\Pi}$. By repeating this process sufficiently many times, \mathcal{R} obtains enough locations of the string $\tilde{\Pi}$. Below we denote by $N = N(\epsilon)$ the number of samples (set later).

$\mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}, \cdot)}(\text{cm}, \epsilon)$:

1. Initialize a proof string: $\tilde{\Pi} := (\sigma)^l$, where σ is an arbitrary element in Σ .
2. Initialize an empty set $\tilde{\mathcal{Q}}$ to track which locations of $\tilde{\Pi}$ are filled in.
3. Repeat the following N times:
 - (a) Sample PCP verifier randomness $\rho \leftarrow \{0, 1\}^r$.
 - (b) Ask $\tilde{\mathcal{P}}$ for answers to this randomness: $(\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\text{aux}, \rho)$.
 - (c) If $\text{VC.Check}(\text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1$, set $\tilde{\Pi}[\mathcal{Q}] := \text{ans}$ and update $\tilde{\mathcal{Q}} := \tilde{\mathcal{Q}} \cup \mathcal{Q}$.
4. Output $(\tilde{\mathcal{Q}}, \tilde{\Pi})$.

We make explicit the two computation phases of the (stateful) malicious argument prover $\tilde{\mathcal{P}}$:

$$(\text{cm}, \text{aux}) \leftarrow \tilde{\mathcal{P}} \quad \text{and} \quad (\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\text{aux}, \rho) ,$$

where aux is the auxiliary state passed across the two computation phases of $\tilde{\mathcal{P}}$. The reducer \mathcal{R} needs to rerun only the second phase of $\tilde{\mathcal{P}}$, so the oracle for \mathcal{R} is $\tilde{\mathcal{P}}(\text{aux}, \cdot)$.

As stated in Lemma 1, with the above notation we wish to bound the following probability:

$$\Pr \left[\begin{array}{l} \mathbf{V}[\tilde{\mathcal{Q}}, \tilde{\Pi}](\mathbf{x}; \rho) \neq 1 \\ \wedge \mathbf{V}[\mathcal{Q}, \text{ans}](\mathbf{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \left| \begin{array}{l} (\text{cm}, \text{aux}) \leftarrow \tilde{\mathcal{P}} \\ (\tilde{\mathcal{Q}}, \tilde{\Pi}) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}, \cdot)}(\text{cm}, \epsilon) \\ \rho \leftarrow \{0, 1\}^r \\ (\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\text{aux}, \rho) \end{array} \right. \right. .$$

Observe that, if $\text{VC.Check}(\text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1$ then $\mathbf{V}[\tilde{\mathcal{Q}}, \tilde{\Pi}](\mathbf{x}; \rho) \neq 1 \wedge \mathbf{V}[\mathcal{Q}, \text{ans}](\mathbf{x}; \rho) = 1$ implies either: (i) $\tilde{\Pi}$ and ans disagree at a position $q \in \mathcal{Q} \cap \tilde{\mathcal{Q}}$; or (ii) there is query q in \mathcal{Q} but not in $\tilde{\mathcal{Q}}$. We analyze the two events separately, which bounds the probability above by a union bound.

(i) Valid openings with disagreeing answers. We informally argue that

$$\Pr \left[\begin{array}{l} \exists q \in \mathcal{Q} \cap \tilde{\mathcal{Q}} : \text{ans}[q] \neq \tilde{\Pi}[q] \\ \wedge \text{VC.Check}(\text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \left| \begin{array}{l} (\text{cm}, \text{aux}) \leftarrow \tilde{\mathcal{P}} \\ (\tilde{\mathcal{Q}}, \tilde{\Pi}) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}, \cdot)}(\text{cm}, \epsilon) \\ \rho \leftarrow \{0, 1\}^r \\ (\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\text{aux}, \rho) \end{array} \right. \right] \leq \epsilon_{\text{VC}}(\lambda, l, q, t_{\text{VC}}) .$$

The reducer \mathcal{R} checks the validity of the opening for each position it fills into $\tilde{\Pi}$. Therefore, the event above implies that there are valid openings to two different values at the same query position; equivalently, the following adversary A_{VC} , which has size $t_{\text{VC}} = O(N \cdot t_{\text{ARG}})$, breaks VC's position binding:

A_{VC} :

1. Run $(\text{cm}, \text{aux}) \leftarrow \tilde{\mathcal{P}}$.
2. Run the reducer $\mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}, \cdot)}(\text{cm}, \epsilon)$ and collect the valid query-answer-opening tuples $(\mathcal{Q}, \text{ans}, \text{pf})$ it sampled in a set \mathcal{K} .
3. Sample fresh randomness $\rho \leftarrow \{0, 1\}^r$ and obtain another tuple $(\mathcal{Q}', \text{ans}', \text{pf}') \leftarrow \tilde{\mathcal{P}}(\text{aux}, \rho)$.
4. For every tuple $(\mathcal{Q}, \text{ans}, \text{pf}) \in \mathcal{K}$, if there exists a query $q \in \mathcal{Q} \cap \mathcal{Q}'$ such that $\text{ans}[q] \neq \text{ans}'[q]$, then output $(\text{cm}, \mathcal{Q}, \text{ans}, \text{pf}, \mathcal{Q}', \text{ans}', \text{pf}')$.¹⁰

¹⁰To conform with the algorithm's syntax, A_{VC} outputs $(\text{cm}, \mathcal{Q}', \text{ans}', \text{pf}', \mathcal{Q}, \text{ans}, \text{pf})$ if no disagreeing openings are found.

Note that A_{VC} only checks inconsistencies between $(\mathcal{Q}', \text{ans}', \text{pf}')$ and tuples in \mathcal{K} . In other words, A_{VC} does not attempt to detect inconsistencies among tuples in \mathcal{K} . Even so, A_{VC} captures the bound we wish to prove:

$$\begin{aligned} & \Pr \left[\begin{array}{l} \exists q \in \mathcal{Q} \cap \tilde{\mathcal{Q}} : \text{ans}[q] \neq \tilde{\Pi}[q] \\ \wedge \text{VC.Check}(\text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \left| \begin{array}{l} (\text{cm}, \text{aux}) \leftarrow \tilde{\mathcal{P}} \\ (\tilde{\mathcal{Q}}, \tilde{\Pi}) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}, \cdot)}(\text{cm}, \epsilon) \\ \rho \leftarrow \{0, 1\}^r \\ (\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\text{aux}, \rho) \end{array} \right. \right] \\ & \leq \Pr \left[\begin{array}{l} \exists q \in \mathcal{Q} \cap \mathcal{Q}' : \text{ans}[q] \neq \text{ans}'[q] \\ \wedge \text{VC.Check}(\text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \\ \wedge \text{VC.Check}(\text{cm}, \mathcal{Q}', \text{ans}', \text{pf}') = 1 \end{array} \left| (\text{cm}, \mathcal{Q}, \text{ans}, \text{pf}, \mathcal{Q}', \text{ans}', \text{pf}') \leftarrow A_{\text{VC}} \right. \right] \\ & \leq \epsilon_{\text{VC}}(\lambda, l, q, t_{\text{VC}}) . \end{aligned}$$

(ii) **Missing position in $\tilde{\Pi}$.** We show that

$$\Pr \left[\begin{array}{l} \mathcal{Q} \setminus \tilde{\mathcal{Q}} \neq \emptyset \\ \wedge \text{VC.Check}(\text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \left| \begin{array}{l} (\text{cm}, \text{aux}) \leftarrow \tilde{\mathcal{P}} \\ (\tilde{\mathcal{Q}}, \tilde{\Pi}) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}, \cdot)}(\text{cm}, \epsilon) \\ \rho \leftarrow \{0, 1\}^r \\ (\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\text{aux}, \rho) \end{array} \right. \right] \leq \frac{1}{N} .$$

To upper bound the probability of a query $q \in \mathcal{Q}$ not having been filled in by \mathcal{R} , we use the probability that a given position $q \in [l]$ is queried. The *weight* $\delta(q)$ of a query $q \in [l]$ is the probability that it is queried by the argument verifier with uniformly sampled randomness. We can write:

$$\begin{aligned} & \Pr \left[\mathcal{Q} \setminus \tilde{\mathcal{Q}} \neq \emptyset \right] \\ & = \Pr \left[\exists q \in [l] : q \in \mathcal{Q} \wedge q \notin \tilde{\mathcal{Q}} \right] \\ & \leq \sum_{q \in [l]} \delta(q) \cdot (1 - \delta(q))^N , \end{aligned}$$

where the inequality follows from the fact that \mathcal{Q} and all query sets used to generate $\tilde{\mathcal{Q}}$ correspond to independently sampled verifier randomness. Note that, for every $\delta \in [0, 1]$, $\delta \cdot (1 - \delta)^N \leq 1/N$.¹¹ Hence, the target probability is upper bounded by $\frac{1}{N}$.

In fact, the proof for this case is more delicate than sketched above. If a position $q \in [l]$ has weight $\delta(q)$, we cannot conclude that $q \notin \tilde{\mathcal{Q}}$ with probability at most $(1 - \delta(q))^N$, because $\tilde{\mathcal{P}}$ may often output invalid openings for q while \mathcal{R} only includes valid openings. To fix this issue, we use a refined notion: $\delta(q)$ is the probability that during the execution of the interactive argument, the verifier \mathcal{V} samples randomness that corresponds to a query set containing q and the prover \mathcal{P} outputs a valid VC opening for the query set.

Setting parameters. By an union bound, the desired probability can be upper bounded by

$$\epsilon_{\text{VC}}(\lambda, l, q, t_{\text{VC}}) + \frac{1}{N} .$$

Setting $N := \frac{1}{\epsilon}$, we get $t_{\text{VC}} = O(N \cdot t_{\text{ARG}}) = O\left(\frac{1}{\epsilon} \cdot t_{\text{ARG}}\right)$ and $\frac{1}{N} = \epsilon$, yielding the bound stated in Lemma 1.

¹¹A simple derivation of the inequality is the following: with $f(x) = x \cdot (1 - x)^N$, we have $\frac{d}{dx} f(\delta) = 0 \iff \delta = \frac{1}{N+1}$. As $f(0) = f(1) = 0$ and δ is the only critical point in $[0, 1]$, it achieves the maximum: $\max_{x \in [0, 1]} \{f(x)\} = f(\delta) \leq 1/N$.

Remark 2.1. Superficially one might hope for an improved analysis showing that one only needs $\frac{1}{q \cdot \epsilon}$ rewindings rather than $\frac{1}{\epsilon}$. Indeed, each rewinding that leads to an accepting transcript yields a freshly sampled fragment of the PCP containing q locations. However such a bound is unrealistic because, in general, a PCP may have many dummy queries. For example, consider a PCP where only $O(1)$ of the q queries are “real”, while all others are dummy queries to fixed locations of the PCP string. That said, there may be other metrics through which the factor $\frac{1}{\epsilon}$ can be slightly improved. We leave this intriguing question to future work.

2.2 Succinct arguments based on public-coin IOPs

We review the IBCS protocol and sketch the main ideas behind Theorem 2; see Section 5 for details.

2.2.1 IBCS protocol

The IBCS protocol (the interactive variant of [BCS16]) obtains a succinct interactive argument by combining two ingredients: a public-coin interactive oracle proof (IOP) and a vector commitment scheme VC (fixed above). Let $\text{IOP} = (\mathbf{P}, \mathbf{V})$ be a public-coin IOP for a relation R with alphabet Σ , round complexity k , proof lengths $(l_i)_{i \in [k]}$ and query complexities $(q_i)_{i \in [k]}$ (with maximal values l_{\max}, q_{\max} and total values l, q) and verifier randomness complexity r (r_i per round). $\text{IBCS}[\text{IOP}, \text{VC}]$ is an interactive argument $\text{ARG} = (\mathcal{P}, \mathcal{V})$ in which the argument prover \mathcal{P} receives an instance \mathfrak{x} and a witness \mathfrak{w} , and the argument verifier \mathcal{V} receives the instance \mathfrak{x} . Then \mathcal{P} and \mathcal{V} interact, across $k + 1$ rounds (exchanging $2k + 1$ messages), as follows.

1. For every round $i \in [k]$ of the IOP system:
 - (a) \mathcal{P} computes the IOP string $\Pi_i \leftarrow \mathbf{P}(\mathfrak{x}, \mathfrak{w}, \rho_{i-1})$, computes the commitment $(\text{cm}_i, \text{aux}_i) \leftarrow \text{VC.Commit}(\Pi_i)$, and sends cm_i to \mathcal{V} .
 - (b) \mathcal{V} samples the i -th IOP verifier randomness $\rho_i \leftarrow \{0, 1\}^{r_i}$ and sends it to \mathcal{P} .
2. \mathcal{P} runs the IOP verifier $\mathbf{V}^{(\Pi_i)_{i \in [k]}}(\mathfrak{x}, (\rho_i)_{i \in [k]})$ to deduce the query sets $(\mathcal{Q}_i)_{i \in [k]} \subseteq [l_1] \times \cdots \times [l_k]$ of \mathbf{V} (where \mathcal{Q}_i is the set of queries to the i -th IOP string), computes opening proofs $\text{pf}_i \leftarrow \text{VC.Open}(\text{aux}_i, \mathcal{Q}_i)$, and sets $\text{ans}_i := \Pi_i[\mathcal{Q}_i]$ for each $i \in [k]$, then sends $((\mathcal{Q}_i, \text{ans}_i, \text{pf}_i))_{i \in [k]}$ to \mathcal{V} .
3. \mathcal{V} performs the following checks.
 - (a) $\text{VC.Check}(\text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) = 1$ for every $i \in [k]$ (i.e., all answers have valid openings);
 - (b) $\mathbf{V}^{([\mathcal{Q}_i, \text{ans}_i])_{i \in [k]}}(\mathfrak{x}; (\rho_i)_{i \in [k]}) = 1$ (i.e., the IOP verifier $\mathbf{V}(\mathfrak{x}; (\rho_i)_{i \in [k]})$ accepts the answers).

Similarly to Section 2.1.1, $\mathbf{V}^{([\mathcal{Q}_i, \text{ans}_i])_{i \in [k]}}(\mathfrak{x}; (\rho_i)_{i \in [k]})$ refers to the decision bit of the IOP verifier \mathbf{V} , given instance \mathfrak{x} and IOP randomness $(\rho_i)_{i \in [k]}$, when each query $q \in \mathcal{Q}_i$ to the i -th oracle is answered with $\text{ans}_i[q]$. (If \mathbf{V} queries an oracle outside the corresponding query set then the decision bit is 0.)

2.2.2 Security reduction

The analysis in Section 2.1 extends to the case of public-coin IOPs with a more complicated reductor \mathcal{R} , which is responsible for producing, for each round, an appropriate IOP string for that round, given a partial transcript of interaction so far. Analogously to Lemma 1, the key step of the security reduction is this lemma.

Lemma 2 (informal). *There exists a probabilistic algorithm \mathcal{R} (the **reductor**) that, for every instance \mathfrak{x} ,*

error parameter $\epsilon > 0$, adversary size bound t_{ARG} , and t_{ARG} -size adversary $\tilde{\mathcal{P}}$, satisfies

$$\Pr \left[\begin{array}{l} \mathbf{V}([\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i]_{i \in [k]}(\mathbb{x}; (\rho_i)_{i \in [k]})) \neq 1 \\ \wedge \mathbf{V}([\mathcal{Q}_i, \text{ans}_i]_{i \in [k]}(\mathbb{x}; (\rho_i)_{i \in [k]})) = 1 \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) \right) = 1 \end{array} \mid \begin{array}{l} \text{For } i \in [k] : \\ (\text{cm}_i, \text{aux}_i) \leftarrow \tilde{\mathcal{P}}(\text{aux}_{i-1}, \rho_{i-1}) \\ (\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}_i, \cdot)}((\text{cm}_j)_{j \leq i}, (\rho_j)_{j < i}, \epsilon) \\ \rho_i \leftarrow \{0, 1\}^{r_i} \\ ((\mathcal{Q}_i, \text{ans}_i, \text{pf}_i))_{i \in [k]} \leftarrow \tilde{\mathcal{P}}(\text{aux}_k, \rho_k) \end{array} \right] \\ \leq \epsilon_{\text{VC}}(\lambda, l_{\text{max}}, q_{\text{max}}, t_{\text{VC}}) + \epsilon ,$$

where $t_{\text{VC}} = O\left(\frac{k!}{\epsilon} \cdot t_{\text{ARG}}\right)$.

Observe that Lemma 2 generalizes Lemma 1, which corresponds to the case $k = 1$.

Similar to the reducer in the PCP case, \mathcal{R} rewinds the IOP adversary to extract an IOP string for each round. In particular, for each round i , the reducer \mathcal{R} is given oracle access to $\tilde{\mathcal{P}}(\text{aux}_i, \cdot)$ (where aux_i is the auxiliary state output along with the i -th commitment cm_i) and has the following syntax. For round i , \mathcal{R} receives as input commitments $(\text{cm}_j)_{j \leq i}$, (partial) verifier randomness $(\rho_j)_{j < i} \in \{0, 1\}^{r_1} \times \dots \times \{0, 1\}^{r_{i-1}}$, and error tolerance $\epsilon > 0$; then \mathcal{R} outputs an IOP string $\tilde{\Pi}_i$ whose entries in $\tilde{\mathcal{Q}}_i$ are valid openings.

With Lemma 2, we can obtain the soundness and knowledge soundness error bounds of Theorem 2 via an analysis analogous to Sections 2.1.3 and 2.1.4. In this overview we omit this straightforward generalization, and only sketch the proof of Lemma 2.

2.2.3 Proof sketch of Lemma 2

The reducer \mathcal{R} is faced with two challenges. On the one hand, for each round i , the IOP string $\tilde{\Pi}_i$ to be extracted depends on prior randomness $(\rho_j)_{j < i}$ (since the corresponding commitment cm_i depends on these). On the other hand, for each round i , queries to $\tilde{\Pi}_i$ may depend on prior randomness $(\rho_j)_{j < i}$ and subsequent randomness $(\rho_j)_{i \leq j \leq k}$. (Indeed, all queries by the IOP verifier may occur after the interaction because we consider public-coin IOPs.) Therefore, \mathcal{R} must sample $(\rho_j)_{i \leq j \leq k}$ (i.e., IOP verifier randomness for all rounds $j \geq i$) to see a fragment of $\tilde{\Pi}_i$. We use $N = N(\epsilon)$ to denote the number of samples that we set later.

For every $i \in [k]$, the reducer algorithm \mathcal{R} works as follows.

$\mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}_i, \cdot)}((\text{cm}_j)_{j \leq i}, (\rho_j)_{j < i}, \epsilon)$:

1. Initialize an IOP string with an arbitrary symbol $\sigma \in \Sigma$: set $\tilde{\Pi}_i := (\sigma)^{l_i}$.
2. Initialize an empty set $\tilde{\mathcal{Q}}_i$ to store filled locations of $\tilde{\Pi}_i$.
3. Repeat the following N times:
 - (a) Sample IOP verifier randomness from round i to round k : $(\rho_i, \dots, \rho_k) \leftarrow \{0, 1\}^{r_i + \dots + r_k}$.
 - (b) Run the prover $\tilde{\mathcal{P}}$ using the sampled randomness (ρ_i, \dots, ρ_k) until the end of interaction, and obtain the query-answer-opening tuple $(\mathcal{Q}_i, \text{ans}_i, \text{pf}_i)$ for the i -th IOP string.
 - (c) If $\text{VC.Check}(\text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) = 1$, set $\tilde{\Pi}_i[\mathcal{Q}_i] := \text{ans}_i$ and update $\tilde{\mathcal{Q}}_i := \tilde{\mathcal{Q}}_i \cup \mathcal{Q}_i$.
4. Output $(\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i)$.

With the above construction, we can prove Lemma 2 with a similar case analysis as Section 2.1.5: if the IOP verifier \mathbf{V} , upon querying the proof strings output by \mathcal{R} , disagrees with the argument verifier \mathcal{V} , then either:

- (i) for some $i \in [k]$, $\tilde{\Pi}_i$ and ans_i disagree at a position $q \in \mathcal{Q}_i \cap \tilde{\mathcal{Q}}_i$; or
- (ii) for some $i \in [k]$, there is query q in \mathcal{Q}_i but not in $\tilde{\mathcal{Q}}_i$.

We bound the probability of the two events separately. We fix this experiment for the rest of the proof:

$$\left[\begin{array}{l} \text{For } i \in [k] : \\ (\text{cm}_i, \text{aux}_i) \leftarrow \tilde{\mathcal{P}}(\text{aux}_{i-1}, \rho_{i-1}) \\ (\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}_i, \cdot)}((\text{cm}_j)_{j \leq i}, (\rho_j)_{j < i}, \epsilon) \\ \rho_i \leftarrow \{0, 1\}^{r_i} \\ ((\mathcal{Q}_i, \text{ans}_i, \text{pf}_i))_{i \in [k]} \leftarrow \tilde{\mathcal{P}}(\text{aux}_k, \rho_k) \end{array} \right] .$$

(i) Valid openings with disagreeing answers. We informally argue that

$$\Pr \left[\begin{array}{l} \exists i \in [k], q \in \mathcal{Q}_i \cap \tilde{\mathcal{Q}}_i : \text{ans}_i[q] \neq \tilde{\Pi}_i[q] \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) = 1 \right) \end{array} \right] \leq \epsilon_{\text{VC}}(\lambda, l_{\max}, \mathbf{q}_{\max}, t_{\text{VC}}) .$$

We construct an adversary A_{VC} for VC with \mathcal{R} and $\tilde{\mathcal{P}}$ as in Section 2.1.5.

A_{VC} :

1. Sample $(\rho_i)_{i \in [k]} \leftarrow \{0, 1\}^{r_1 + \dots + r_k}$.
2. For every $i \leq k$:
 - (a) Run $(\text{cm}_i, \text{aux}_i) \leftarrow \tilde{\mathcal{P}}(\text{aux}_{i-1}, \rho_{i-1})$.
 - (b) Run $\mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}_i, \cdot)}((\text{cm}_j)_{j \leq i}, (\rho_j)_{j < i}, \epsilon)$ and collect the valid query-answer-opening tuples in a set \mathcal{K}_i .
3. Run $((\mathcal{Q}'_i, \text{ans}'_i, \text{pf}'_i))_{i \in [k]} \leftarrow \tilde{\mathcal{P}}(\text{aux}_k, \rho_k)$.
4. If there exists $i \in [k]$ and $(\mathcal{Q}_i, \text{ans}_i, \text{pf}_i) \in \mathcal{K}_i$ such that $q \in \mathcal{Q}_i \cap \mathcal{Q}'_i$ and $\text{ans}_i[q] \neq \text{ans}'_i[q]$, output $(\text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i, \mathcal{Q}'_i, \text{ans}'_i, \text{pf}'_i)$. (Otherwise, simply output $(\text{cm}_1, \mathcal{Q}'_1, \text{ans}'_1, \text{pf}'_1, \mathcal{Q}'_1, \text{ans}'_1, \text{pf}'_1)$.)

Once again, the VC adversary A_{VC} only needs to check for inconsistent answers between $((\mathcal{Q}_i, \text{ans}_i, \text{pf}_i))_{i \in [k]}$ and $(\mathcal{K}_i)_{i \in [k]}$, rather than also checking inconsistencies within each \mathcal{K}_i .

The VC adversary A_{VC} simulates round i of $\tilde{\mathcal{P}}$ for $i \cdot (N + 1)$ times (see construction of \mathcal{R}), so A_{VC} runs at most kN full executions of $\tilde{\mathcal{P}}$. Since one full execution of $\tilde{\mathcal{P}}$ has size t_{ARG} , the size of A_{VC} is at most $O(kN \cdot t_{\text{ARG}})$. Hence setting $t_{\text{VC}} = O(kN \cdot t_{\text{ARG}})$ bounds the probability above by $\epsilon_{\text{VC}}(\lambda, l_{\max}, \mathbf{q}_{\max}, t_{\text{VC}})$.

(ii) Missing query in $\tilde{\mathcal{Q}}_i$. By extending the approach in Section 2.1.5, we show that, for every $\delta \in [0, 1]$,

$$\Pr \left[\begin{array}{l} \left(\exists i \in [k] : \mathcal{Q}_i \setminus \tilde{\mathcal{Q}}_i \neq \emptyset \right) \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) = 1 \right) \end{array} \right] \leq \frac{1}{N} .$$

For each $i \in [k]$ and a given position $q \in [l_i]$, q has weight $\delta(q)$ with respect to a randomness prefix $\rho_1, \dots, \rho_{i-1}$ if it is queried by the argument verifier \mathcal{V} with probability $\delta(q)$ (over $\rho_{i+1}, \dots, \rho_k$).¹² Here we consider the *residual* probability after fixing a randomness prefix; this is necessary as this is the distribution that \mathcal{R} samples from.

Therefore:

$$\begin{aligned} & \Pr \left[\exists i \in [k] : \mathcal{Q}_i \setminus \tilde{\mathcal{Q}}_i \neq \emptyset \right] \\ &= \Pr \left[\exists i \in [k], q \in [l_i] : q \in \mathcal{Q}_i \wedge q \notin \tilde{\mathcal{Q}}_i \right] \end{aligned}$$

¹²Similarly to Section 2.1.5, here the definition of weight is simplified. In the full proof, we need to consider the fact that $\tilde{\mathcal{P}}$ might output invalid openings.

$$\begin{aligned}
&\leq \sum_{i=1}^k \Pr \left[q \in [l_i] : q \in \mathcal{Q}_i \wedge q \notin \tilde{\mathcal{Q}}_i \right] \\
&\leq \sum_{i=1}^k \sum_{q \in [l_i]} \delta(q) \cdot (1 - \delta(q))^N \\
&\leq \frac{1}{N} .
\end{aligned}$$

Setting parameters. Setting $N := \frac{1}{\epsilon}$, we have that $t_{VC} = O(kN \cdot t_{ARG}) = O\left(\frac{k \cdot 1}{\epsilon} \cdot t_{ARG}\right)$ and $\frac{1}{N} = \epsilon$ as desired.

2.3 Why doesn't the same analysis work for public-query IOPs?

We have discussed how the VC-based approach works for PCPs (in Section 2.1) and for public-coin IOPs (in Section 2.2). Here we discuss the capabilities and limitations of this approach towards succinct interactive arguments. We begin by giving an informal definition of a general (private-coin) IOP.¹³

Definition 1 (informal). An IOP = (\mathbf{P}, \mathbf{V}) with round complexity k and randomness complexity r works as follows.

1. \mathbf{V} samples its private random string $\rho \leftarrow \{0, 1\}^r$.
2. For every round $i \in [k]$:
 - (a) \mathbf{P} computes the i -th IOP string $\Pi_i \leftarrow \mathbf{P}(\mathbb{x}, \mathbb{w}, m_{i-1})$ and sends it to \mathbf{V} .
 - (b) \mathbf{V} computes the i -th query set $\mathbf{Q}_i = (\mathcal{Q}_{i,j})_{j \leq i} \leftarrow \mathbf{V}_q(\mathbb{x}, \rho, (\mathbf{ans}_j)_{j < i})$, where $\mathcal{Q}_{i,j}$ is the set of queries to the j -th IOP string in round i .
 - (c) \mathbf{V} obtains the answers \mathbf{ans}_i to \mathbf{Q}_i by querying $(\Pi_j)_{j \leq i}$.
 - (d) \mathbf{V} computes the i -th message $m_i \leftarrow \mathbf{V}_m(\mathbb{x}, \rho, (\mathbf{ans}_j)_{j \leq i})$ and sends it to \mathbf{P} .
3. \mathbf{V} checks that $\mathbf{V}_d(\mathbb{x}, \rho, (\mathbf{ans}_i)_{i \in [k]}) = 1$.

We use $\mathbf{V}_q, \mathbf{V}_m, \mathbf{V}_d$ to denote the query, message, and decision functions of the IOP verifier \mathbf{V} , respectively.

We cannot hope for the VC-based approach to work for *every* IOP. The argument verifier needs the argument prover's help to obtain the answers to its queries; hence, the argument prover learns each round's queries as they are made. Unless the underlying IOP remains secure even when queries are revealed to the IOP prover, the VC-based approach is insecure.

Therefore, at best we can hope for the approach to work for *public-query IOPs*. These are IOPs in which soundness (and knowledge soundness) holds even if the IOP prover knows the queries made by the IOP verifier, or, equivalently, where the latter sends its i -th round query set \mathbf{Q}_i along with the message m_i . This leads to the following (narrowed) question: *Does the VC-based approach work for every public-query IOP?*

In Section 2.3.1, we describe a succinct interactive argument that appears secure provided that the underlying IOP is public-query. In Section 2.3.2, we explain the challenges towards a proof of security for this protocol. In Section 2.3.3, we identify a property of public-query IOPs (which is trivially satisfied by public-coin IOPs) that suffices for security. In Section 2.3.4, we partially characterize the class of public-query IOPs that admit this property. Technical details are in Section 6.

¹³For notational simplicity, we assume that, within each round of interaction, the IOP verifier queries non-adaptively. All discussions in this paper directly extend to adaptive queries within each round.

2.3.1 The Finale protocol

We describe the *Finale* protocol, a succinct interactive argument obtained from any public-query IOP. It is a generalization of the IBCS protocol in Section 2.2.1 from public-coin IOPs to public-query IOPs. The main difference is that, while in the IBCS protocol the argument prover answers queries by the IOP verifier after the IOP interaction, the argument prover in the *Finale* protocol answers queries within each round.

Let $\text{IOP} = (\mathbf{P}, \mathbf{V})$ be a public-query IOP for a relation R with alphabet Σ , round complexity k , maximal proof length l_{\max} , maximal query complexity q_{\max} , and verifier randomness complexity r . $\text{Finale}[\text{IOP}, \text{VC}]$ is an interactive argument $\text{ARG} = (\mathcal{P}, \mathcal{V})$ in which the argument prover \mathcal{P} receives an instance \mathbb{x} and a witness \mathbb{w} , and the argument verifier \mathcal{V} receives the instance \mathbb{x} . Then \mathcal{P} and \mathcal{V} interact, across $2k$ rounds (exchanging $4k$ messages), as follows.

1. \mathcal{V} samples the verifier randomness $\rho \leftarrow \{0, 1\}^r$ for \mathbf{V} .
2. For every round $i \in [k]$ of the IOP system:
 - (a) \mathcal{P} computes the i -th IOP string $\Pi_i \leftarrow \mathbf{P}(\mathbb{x}, \mathbb{w}, m_{i-1})$, computes the corresponding commitment $(\text{cm}_i, \text{aux}_i) \leftarrow \text{VC.Commit}(\Pi_i)$, and sends cm_i to \mathcal{V} .
 - (b) \mathcal{V} sends the i -th query set $\mathbf{Q}_i = (\mathcal{Q}_{i,j})_{j \leq i} \leftarrow \mathbf{V}_q(\mathbb{x}, \rho, (\mathbf{ans}_j)_{j < i})$ to \mathcal{P} , where $\mathcal{Q}_{i,j}$ is the set of queries to Π_j in round i .
 - (c) For each $j \leq i$, \mathcal{P} computes an opening proof $\text{pf}_{i,j} \leftarrow \text{VC.Open}(\text{aux}_j, \mathcal{Q}_{i,j})$.
 - (d) \mathcal{P} sends $(\mathbf{ans}_i := (\Pi_i[\mathcal{Q}_{i,j}])_{j \leq i}, \mathbf{pf}_i := (\text{pf}_{i,j})_{j \leq i})$ to \mathcal{V} .
 - (e) Then \mathcal{V} sends the i -th message $m_i \leftarrow \mathbf{V}_m(\mathbb{x}, \rho, (\mathbf{ans}_j)_{j \leq i})$ to \mathcal{P} .
3. \mathcal{V} checks if the following conditions hold:
 - (a) $\mathbf{V}_d(\mathbb{x}, \rho, (\mathbf{ans}_i)_{i \in [k]}) = 1$.
 - (b) For every $i \in [k]$ and $j \leq i$, $\text{VC.Check}(\text{cm}_j, \mathcal{Q}_{i,j}, \text{ans}_{i,j}, \text{pf}_{i,j}) = 1$.

2.3.2 Is there a security reduction that works for every public-query IOP?

It would be natural to conjecture that the security analysis of the IBCS protocol (for public-coin IOPs) directly extends to a security analysis for the *Finale* protocol (for public-query IOPs). However, that is not the case.

The key object in the security analysis is the redactor: \mathcal{R} rewinds the malicious argument prover $\tilde{\mathcal{P}}$ multiple times, each time simulating a fresh partial interaction between $\tilde{\mathcal{P}}$ and the (honest) argument verifier \mathcal{V} . Let us attempt to construct such a redactor \mathcal{R} for the *Finale* protocol.

Consider the goal of the redactor \mathcal{R} for the first round. The redactor \mathcal{R} receives oracle access to $\tilde{\mathcal{P}}$ and receives as input the first commitment cm_1 (and error parameter ϵ); the (first) goal of \mathcal{R} is to output a (partially filled) IOP string $\tilde{\Pi}_1$ consistent with cm_1 . As before, \mathcal{R} can rewind $\tilde{\mathcal{P}}$ for $N(\epsilon)$ times, each time sampling fresh private randomness for \mathcal{V} and simulating a full interaction between $\tilde{\mathcal{P}}$ and \mathcal{V} .

The problem arises in the second round, where \mathcal{R} must simulate partial interactions of the interactive argument. Now \mathcal{R} receives as input the first two commitments $(\text{cm}_1, \text{cm}_2)$, the first round's verifier queries \mathbf{Q}_1 , the corresponding query answers \mathbf{ans}_1 and opening proofs \mathbf{pf}_1 , and the first round's verifier message m_1 (as well as the error parameter ϵ). The new goal of \mathcal{R} is to output a (partially filled) IOP string $\tilde{\Pi}_2$ consistent with cm_2 . To do so, \mathcal{R} must simulate multiple partial interactions between $\tilde{\mathcal{P}}$ and \mathcal{V} *starting from the second round*. This means that \mathcal{R} needs to sample “consistent” private randomness ρ for the argument verifier \mathcal{V} : ρ such that \mathcal{V} 's first query set is \mathbf{Q}_1 and, given answers \mathbf{ans}_1 to these queries, \mathcal{V} 's first message to the argument prover is m_1 . More generally, \mathcal{R} must have the capability of sampling consistent private randomness starting from any round (defined by a partial transcript of the interactive argument).

How may we sample from this conditional distribution? For the IBCS protocol, sampling a random continuation for a partial transcript is trivial because it is public coin: given first round randomness ρ_1 , sample

randomness ρ_2, \dots, ρ_k for the remaining rounds; and similarly if starting from a later round. In contrast, in the Finale protocol, the argument verifier is private-coin: it samples all of its private randomness ρ at the beginning of the interaction and then uses ρ to deterministically compute queries and messages.

An attempt to sample from the conditional distribution would be for \mathcal{R} to repeatedly sample choices of ρ until one is consistent with $((\text{cm}_1, \text{cm}_2), (\mathbf{Q}_1, \text{ans}_1, m_1))$. However, such a sampling strategy would be inefficient, and would imply (in the resulting reduction) an adversary against VC whose similarly large runtime makes ϵ_{VC} trivial (recall that in Section 2.1.5 the VC adversary A_{VC} uses \mathcal{R} as a subroutine).

In sum, the bottleneck of the security reduction approach in Section 2.2.2 is the ability to *efficiently and consistently sample argument verifier randomness consistent with a partial interaction transcript*, which in turn reduces to the ability to sample consistent IOP verifier randomness for the underlying public-query IOP (as we discuss soon in Section 2.3.3). We are not aware of alternative approaches, and the problem of proving the Finale protocol secure for every public-query IOP (or proving it insecure for some public-query IOP) **remains open**. In the meantime, to achieve a security reduction, we restrict our attention to public-query IOPs that possess efficient *random continuation samplers*, which we discuss next.

2.3.3 IOP random continuation sampler leads to security reduction

We sketch how, as claimed in Theorem 3, if the underlying public-query IOP has an efficient random continuation sampler then the Finale protocol is secure.

Transcripts for general IOPs. A complete interaction transcript for an IOP has the following form:

$$\text{tr} := ((\mathbf{Q}_i, \text{ans}_i, m_i))_{i \in [k]} .$$

For every $i \in [k]$, a partial interaction transcript tr_i for an IOP can have the following forms:

- $\text{tr}_i = (((\mathbf{Q}_j, \text{ans}_j, m_j))_{j < i}, (\mathbf{Q}_i, \text{ans}_i))$, if the IOP verifier is about to send its i -th message m_i ;
- $\text{tr}_i = (((\mathbf{Q}_j, \text{ans}_j, m_j))_{j \leq i})$, if the IOP verifier is about to output its $(i + 1)$ -th query set \mathbf{Q}_{i+1} ($i < k$).

An IOP interaction transcript includes only the values of queried locations, and not entire IOP strings.

What is an IOP random continuation sampler? An IOP random continuation sampler (IOP-RCS) receives as input a partial interaction transcript tr_i for an IOP and samples the next message for tr_i at random among all next messages consistent with tr_i . (In the technical sections, we also allow for a sampling error α .) Formally, an IOP-RCS \mathbf{S} works as follows:

- If $\text{tr}_i = (((\mathbf{Q}_j, \text{ans}_j, m_j))_{j < i}, (\mathbf{Q}_i, \text{ans}_i))$, then $\mathbf{S}(\text{tr}_i)$ samples the next message m_i ;
- If $\text{tr}_i = (((\mathbf{Q}_j, \text{ans}_j, m_j))_{j \leq i})$ and $i < k$, then $\mathbf{S}(\text{tr}_i)$ samples the next query set \mathbf{Q}_{i+1} .

How does IOP-RCS help with proving security? As discussed in Section 2.3.2, the key step in the security reduction is efficiently sampling “consistent random continuations” of a given partial interaction transcript of the argument system. In other words, we need an *ARG random continuation sampler* (ARG-RCS) for the Finale protocol. The notion of ARG-RCS was implicit in prior settings: the construction of \mathcal{R} for Kilian’s protocol (in Section 2.1.5) and for the IBCS protocol (in Section 2.2.3) implicitly relies on an ARG-RCS, which is trivial for arguments built from PCPs or public-coin IOPs (it suffices to sample uniform random strings for future rounds). In the Finale protocol, based on public-query IOPs, an efficient ARG-RCS is not trivial. Fortunately, we can show that if IOP admits an IOP-RCS then $\text{Finale}[\text{IOP}, \text{VC}]$ admits an ARG-RCS.

Step 1: IOP-RCS implies ARG-RCS. An ARG-RCS can be defined analogously to an IOP-RCS. A complete interaction transcript for $\text{Finale}[\text{IOP}, \text{VC}]$ can be written as:

$$\text{tr} := ((\text{cm}_i, \mathbf{Q}_i, \text{ans}_i, \mathbf{pf}_i, m_i))_{i \in [k]} .$$

Such an interaction transcript is syntactically identical to an IOP interaction transcript, but for the additional inclusion of commitments cm_i and openings \mathbf{pf}_i . Moreover, the argument verifier's behavior in the interaction is independent of the commitments and openings (only after the interaction the argument verifier uses them to check the validity of the answers to the query sets). Hence the transformation from an IOP-RCS \mathbf{S} to an ARG-RCS \mathcal{S} is trivial: given a partial argument transcript tr_i , discard the commitments and openings and pass the resulting partial IOP transcript to \mathbf{S} . More specifically, \mathcal{S} is obtained from \mathbf{S} as follows:

$\mathcal{S}(\text{tr}_i)$:

1. If $\text{tr}_i = (((\text{cm}_j, \mathbf{Q}_j, \text{ans}_j, \mathbf{pf}_j, m_j))_{j < i}, (\text{cm}_i, \mathbf{Q}_i, \text{ans}_i, \mathbf{pf}_i))$, output $\mathbf{S}((((\mathbf{Q}_j, \text{ans}_j, m_j))_{j < i}, (\mathbf{Q}_i, \text{ans}_i)))$;
2. If $\text{tr}_i = (((\text{cm}_j, \mathbf{Q}_j, \text{ans}_j, \mathbf{pf}_j, m_j))_{j \leq i})$ and $i < k$, output $\mathbf{S}((((\mathbf{Q}_j, \text{ans}_j, m_j))_{j \leq i}))$.

It is clear that \mathcal{S} samples valid randomness if \mathbf{S} does, and the running time of \mathcal{S} is $O(t_{\mathbf{S}})$.

Step 2: ARG-RCS enables a security reduction. We can use an ARG-RCS \mathcal{S} for the Finale protocol to construct an efficient reductor \mathcal{R} for the Finale protocol. Indeed, we can use \mathcal{S} to randomly sample the next query set and verifier message in the transcript.

$\mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}_i, \cdot)}((\text{cm}_j)_{j \leq i}, ((\mathbf{Q}_j, \text{ans}_j, \mathbf{pf}_j, m_j))_{j < i}, \epsilon)$:

1. Initialize an IOP string with an arbitrary symbol $\sigma \in \Sigma$: set $\tilde{\Pi}_i := (\sigma)^i$.
2. Initialize an empty set $\tilde{\mathcal{Q}}_i$ to store filled locations of $\tilde{\Pi}_i$.
3. Compute $N := N(\epsilon)$.
4. Repeat the following N times: For $i \leq j \leq k$:
 - (a) If $j \neq i$, compute the j -th commitment: $(\text{cm}_j, \text{aux}_j) \leftarrow \tilde{\mathcal{P}}(\text{aux}_{j-1}, \mathbf{Q}_{j-1}, m_{j-1})$.
 - (b) Sample the j -th IOP query set: $\mathbf{Q}_j \leftarrow \mathcal{S}((((\text{cm}_\ell, \mathbf{Q}_\ell, \text{ans}_\ell, \mathbf{pf}_\ell, m_\ell))_{\ell < j}))$.
 - (c) Get the j -th answer and VC openings from $\tilde{\mathcal{P}}$: $(\text{ans}_j, \mathbf{pf}_j, \text{aux}_j) \leftarrow \tilde{\mathcal{P}}(\text{aux}_j, \mathbf{Q}_j)$.
 - (d) Sample the j -th IOP verifier message: $m_j \leftarrow \mathcal{S}((((\text{cm}_\ell, \mathbf{Q}_\ell, \text{ans}_\ell, \mathbf{pf}_\ell, m_\ell))_{\ell < j}, (\text{cm}_j, \mathbf{Q}_j, \text{ans}_j, \mathbf{pf}_j)))$.
 - (e) If $\text{VC.Check}(\text{cm}_i, \mathbf{Q}_{j,i}, \text{ans}_{j,i}, \mathbf{pf}_{j,i}) = 1$, set $\tilde{\Pi}_i[\mathbf{Q}_{j,i}] := \text{ans}_{j,i}$ and update $\tilde{\mathcal{Q}}_i := \tilde{\mathcal{Q}}_i \cup \mathbf{Q}_{j,i}$.
5. Output $(\tilde{\mathcal{Q}}, \tilde{\Pi}_i)$.

If the ARG-RCS \mathcal{S} is efficient then the reductor \mathcal{R} is efficient. We can then prove security by adapting Lemma 2 to this setting; we sketch how the proof changes to account for public queries.

As before (Sections 2.1.5 and 2.2.3), the security reduction claim can be split into two parts.

- *Valid openings with disagreeing answers.* We construct an adversary A_{VC} for VC with \mathcal{R} and $\tilde{\mathcal{P}}$ as in Section 2.1.5. The construction is similar, except that the verifier's messages are now denoted by m rather than ρ and the running time of A_{VC} needs to take into account the running time of the ARG-RCS \mathcal{S} (in \mathcal{R}).
- *Missing query in $\tilde{\Pi}_i$.* For every i -round partial transcript tr_i , we can define the weight of a position $q \in [l_i]$ with respect to tr_i analogously as in Section 2.2.3, and the same analysis follows:

$$\Pr \left[\exists q \in [l_i] : (\exists j \in [i, k] : q \in \mathbf{Q}_{j,i}) \wedge q \notin \tilde{\mathcal{Q}}_i \right] \leq \frac{l_i}{N} .$$

Using a union bound over the rounds, we upper bound the target probability of this case as before by

$$\sum_{i=1}^k \frac{l_i}{N} = \frac{l}{N} .$$

We conclude that when the Finale protocol (based on a public-query IOP) has an efficient random continuation sampler, we can show (knowledge) soundness via a similar strategy to Sections 2.1.3 and 2.1.4.

2.3.4 ARG-RCS implies IOP-RCS

We show that IOP has an efficient IOP-RCS if Finale[IOP, VC] admits an efficient ARG-RCS. Combining this with the discussion from Section 2.3.3, we conclude that the existence of an efficient ARG-RCS for the Finale protocol is equivalent to the existence of an efficient IOP-RCS for the underlying public-query IOP.

Let \mathcal{S} be an ARG-RCS for Finale[IOP, VC] with running time $t_{\mathcal{S}}$. We wish to construct an IOP-RCS \mathbf{S} for IOP. Let tr_i be a partial interaction transcript for the IOP. To invoke \mathcal{S} , the IOP sampler \mathbf{S} needs to augment the IOP transcript to an argument transcript that includes commitments to each IOP string and openings to the commitments. For example, in order for \mathcal{S} to output the i -th query set, the argument transcript given as input should include cm_j for every $j \leq i$. However, \mathbf{S} has no information about the i -th IOP string.

We circumvent the issue above by observing that the argument verifier does not use the commitments nor the VC openings until the interaction ends. Specifically, the argument verifier in the Finale protocol uses them only for checking whether the answers provided by the argument prover are valid. Therefore, the IOP-RCS \mathbf{S} can use empty commitments and empty openings to construct an argument transcript:

$\mathbf{S}(\text{tr}_i)$:

1. If $\text{tr}_i = (((\mathbf{Q}_j, \text{ans}_j, m_j))_{j < i}, (\mathbf{Q}_i, \text{ans}_i))$, outputs $\mathcal{S}(((\perp, \mathbf{Q}_j, \text{ans}_j, \perp, m_j))_{j < i}, (\perp, \mathbf{Q}_i, \text{ans}_i, \perp))$;
2. If $\text{tr}_i = (((\mathbf{Q}_j, \text{ans}_j, m_j))_{j \leq i})$ and $i < k$, outputs $\mathcal{S}(((\perp, \mathbf{Q}_j, \text{ans}_j, \perp, m_j))_{j \leq i})$.

It is clear that \mathbf{S} samples valid randomness if \mathcal{S} does, and the running time of \mathbf{S} is $O(t_{\mathcal{S}})$.

2.4 Succinct interactive arguments with adaptive security

For simplicity, all results and discussions in Sections 1 and 2 are in the *plain model*, where there are no public parameters available to all parties (so the argument verifier is responsible to sample and send VC's public parameters to the argument prover). However, in the technical sections (Sections 4 to 6) we show stronger versions of Theorems 1 to 3 that hold with adaptive security in the common reference string (CRS) model.

An interactive argument in the CRS model includes an additional algorithm: a trusted *generator algorithm* that samples public parameters pp for the argument prover and argument verifier (which can be used any number of times across different interactions). After that, based on pp , a malicious argument prover can choose the instance on which to interact with the argument verifier. This setting necessitates appropriate definitions of *adaptive* soundness and knowledge soundness (see Section 3.1), which require error bounds to hold for any instance \mathbb{x} chosen by the malicious argument prover up to an instance size bound n .¹⁴ In particular, the (soundness and knowledge soundness) error bounds depend on n rather than \mathbb{x} .

Our security analyses to achieve adaptive security in the CRS model follow the same structure as the discussions in the sections above, with only syntactic modifications due to the different target definitions.

¹⁴We also rely on formulations of soundness and knowledge soundness for PCPs and IOPs in which the malicious prover chooses the instance (see Sections 3.3 and 3.4). This is only for convenience, because, due to the information-theoretic setting, these definitions are straightforwardly implied by (standard) definitions for fixed instances.

(E.g., modifying the experiments to replace a fixed instance x with an instance size bound n , and letting the malicious argument prover choose the instance.)

Overall, the (formal) statements provided in the technical sections (Sections 4 to 6) are stronger than the (informal) statements in Theorems 1 to 3 because we achieve adaptive security in the CRS model.¹⁵

¹⁵Adaptive security in the CRS model directly implies security in the plain model. Since no CRS is allowed, the argument verifier can begin the interaction by running itself the generator algorithm and sending the public parameters for the argument system to the argument prover. See Remark 3.6.

3 Preliminaries

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 $n \in \mathbb{N}$, public parameter $\text{pp} \in \mathcal{G}(1^\lambda, n)$, and instance-witness pair $(\mathbb{x}, \mathbb{w}) \in R$ with $|\mathbb{x}| \leq n$,

$$\Pr [\langle \mathcal{P}(\text{pp}, \mathbb{x}, \mathbb{w}), \mathcal{V}(\text{pp}, \mathbb{x}) \rangle = 1] = 1 .$$

Definition 3.3 (Adaptive soundness). $\text{ARG} = (\mathcal{G}, \mathcal{P}, \mathcal{V})$ for a relation R has **(adaptive) 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} , circuit size bound $t_{\text{ARG}} \in \mathbb{N}$, and t_{ARG} -size circuit $\tilde{\mathcal{P}}$,

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

Definition 3.4 (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} , circuit size bound $t_{\text{ARG}} \in \mathbb{N}$, and t_{ARG} -size circuit $\tilde{\mathcal{P}}$,

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

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

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., public parameters and 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).

Moreover, we can assume, without loss of generality, that $\tilde{\mathcal{P}}$ is deterministic relative to auxiliary input ai (as the internal coin flips of a probabilistic $\tilde{\mathcal{P}}$ can be incorporated into the auxiliary input distribution \mathcal{D}).

Remark 3.5. 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; see Remark 4.4.

Remark 3.6 (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 \mathfrak{x} and witness \mathfrak{w} , and the argument verifier \mathcal{V}' receives as input the instance \mathfrak{x} ; both also receive as input the security parameter λ (in unary). They interact as follows:

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

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.¹⁶ 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 Vector commitments

A (static) *vector commitment scheme* [CF13] over alphabet Σ is a tuple of algorithms

$$\text{VC} = (\text{Gen}, \text{Commit}, \text{Open}, \text{Check})$$

with the following syntax.

- $\text{VC.Gen}(1^\lambda, \ell) \rightarrow \text{pp}$: On input a security parameter $\lambda \in \mathbb{N}$ and message size bound $\ell \in \mathbb{N}$, VC.Gen samples public parameter pp .
- $\text{VC.Commit}(\text{pp}, m) \rightarrow (\text{cm}, \text{aux})$: On input a public parameter pp and a message $m \in \Sigma^\ell$, VC.Commit produces a commitment cm and the corresponding auxiliary state aux .
- $\text{VC.Open}(\text{pp}, \text{aux}, \mathcal{Q}) \rightarrow \text{pf}$: On input a public parameter pp , an auxiliary state aux , and a query set $\mathcal{Q} \subseteq [\ell]$, VC.Open outputs an opening proof string pf attesting that $m[\mathcal{Q}]$ is a restriction of m to \mathcal{Q} .
- $\text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) \rightarrow \{0, 1\}$: On input a public parameter pp , a commitment cm , a query set $\mathcal{Q} \subseteq [\ell]$, an answer string $\text{ans} \in \Sigma^{\mathcal{Q}}$, and an opening proof string pf , VC.Check determines if pf is a valid proof for $\text{ans} \in \Sigma^{\mathcal{Q}}$ being a restriction of the message committed in cm to \mathcal{Q} .

The vector commitment scheme VC must satisfy perfect completeness and position binding.

Definition 3.7 (Completeness). $\text{VC} = (\text{Gen}, \text{Commit}, \text{Open}, \text{Check})$ has **perfect completeness** if for every security parameter $\lambda \in \mathbb{N}$, message length $\ell \in \mathbb{N}$, message $m \in \Sigma^\ell$, and query set $\mathcal{Q} \subseteq [\ell]$,

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

Definition 3.8 (Position binding). $\text{VC} = (\text{Gen}, \text{Commit}, \text{Open}, \text{Check})$ has **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 size bound $t_{\text{VC}} \in \mathbb{N}$, and t_{VC} -size circuit A_{VC} ,

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

¹⁶These standard definitions can be derived from Definitions 3.2 to 3.4 by setting pp to be empty.

Remark 3.9 (Monotonicity of ϵ_{VC}). We assume hereafter that the position binding error ϵ_{VC} is monotone in each coordinate in the natural direction:

- $\epsilon_{\text{VC}}(\cdot, \ell, s, t_{\text{VC}})$ is non-increasing (larger security parameters decrease an adversary's success);
- $\epsilon_{\text{VC}}(\lambda, \cdot, s, t_{\text{VC}})$ is non-decreasing (opening some set in a string is easier than opening in a substring);
- $\epsilon_{\text{VC}}(\lambda, \ell, \cdot, t_{\text{VC}})$ is non-decreasing (finding a collision in a set is easier than finding one in a subset); and
- $\epsilon_{\text{VC}}(\lambda, \ell, s, \cdot)$ 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 are natural (and satisfied in the case of Merkle trees); in any case, otherwise one may replace, in our computations, expressions of the type $\epsilon_{\text{VC}}(\lambda, \ell_{\text{max}}, s_{\text{max}}, t_{\text{VC}})$, when $\ell_{\text{max}} = \max_i \{\ell_i\}$ and $s_{\text{max}} = \max_j \{s_j\}$, with

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

3.3 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 3.10 (Completeness). $\text{PCP} = (\mathbf{P}, \mathbf{V})$ for a relation R has **perfect completeness** if, for every instance-witness pair $(\mathbf{x}, \mathbf{w}) \in R$,

$$\Pr [\mathbf{V}^{\Pi}(\mathbf{x}) = 1 \mid \Pi \leftarrow \mathbf{P}(\mathbf{x}, \mathbf{w})] = 1 .$$

Definition 3.11 (Soundness). $\text{PCP} = (\mathbf{P}, \mathbf{V})$ for a relation R has **soundness error** ϵ_{PCP} if, for every (unbounded) circuit $\tilde{\mathbf{P}}$ and auxiliary input distribution \mathbf{D} ,

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

Definition 3.12 (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 circuit $\tilde{\mathbf{P}}$ and auxiliary input distribution \mathbf{D} ,

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

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

We consider several efficiency measures for a PCP:

- the *proof alphabet* Σ is the alphabet over which a PCP string is written;
- the *proof length* l is the number of alphabet symbols in the PCP string;
- the *query complexity* $q \in [l]$ is the number of queries that the PCP verifier makes to the PCP string (each query is an index in $[l]$ and is answered by the corresponding symbol in Σ in the PCP string);
- the *randomness complexity* r is the number of random bits used by the PCP verifier.

An efficiency measure may be a function of the instance \mathbf{x} (e.g., of its size $|\mathbf{x}|$).

3.4 Interactive oracle proofs

An *interactive oracle proof* (IOP) system [BCS16; RRR16] is a generalization of PCPs to multiple rounds. An IOP system for a relation R is formally defined as follows.

Definition 3.13 (Completeness). $\text{IOP} = (\mathbf{P}, \mathbf{V})$ for a relation R has **perfect completeness** if for every instance-witness pair $(\mathbf{x}, \mathbf{w}) \in R$,

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

Definition 3.14 (Soundness). $\text{IOP} = (\mathbf{P}, \mathbf{V})$ for a relation R has **soundness error** ϵ_{IOP} if for every circuit $\tilde{\mathbf{P}}$,

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

Definition 3.15 (Knowledge soundness). $\text{IOP} = (\mathbf{P}, \mathbf{V})$ for a relation R has **knowledge soundness error** κ_{IOP} **with extraction time** $t_{\mathbf{E}}$ if there exists a probabilistic algorithm \mathbf{E} such that, for every circuit $\tilde{\mathbf{P}}$,

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

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

We consider several efficiency measures for an IOP:

- the *proof alphabet* Σ is the alphabet over which each IOP string is written;
- the *proof length* l is the total number of alphabet symbols across all IOP strings sent by the IOP prover; moreover, l_i is the length of the proof sent by \mathbf{P} in round i and $l_{\max} := \max_i \{l_i\}$;
- the *query complexity* $q \in [l]$ is the total number of queries that the IOP verifier makes to any IOP string (each query specifies a round a location in the IOP string of that round and is answered by the corresponding symbol in the IOP string);
- the *randomness complexity* r is the number of random bits used by the IOP verifier;
- the *round complexity* k is the number of interaction rounds (back-and-forth interactions) between the IOP prover and IOP verifier.

Any efficiency measure may be a function of the instance \mathbf{x} (e.g., of the instance size $|\mathbf{x}|$).

Public-coin IOPs. We focus on IOPs that are *public-coin*, which informally means that the IOP verifier is a public-coin interactive algorithm.

In other words, in every round, the verifier sends a uniformly random message, independent of other messages they send, to the prover. We provide the formal definition below.

Definition 3.16. $\text{IOP} = (\mathbf{P}, \mathbf{V})$ for a relation R is **public-coin** if, for every $i \in [k]$, the i -th message of the IOP verifier \mathbf{V} is a freshly-sampled uniform random string ρ_i of a prescribed length r_i (which may depend on the instance). In particular, during an interaction, the IOP prover knows all the randomness sampled by the IOP verifier so far (as the messages received by the IOP prover are all that randomness so far).

Queries in a public-coin IOP can be postponed until after the interaction. (This is without loss of generality as, during the interaction, the IOP verifier simply sends a fresh random message in each round.) Hence the IOP can be viewed in two parts: an interaction phase (the prover sends oracles and the verifier sends random messages) and then a query phase (the verifier queries the oracles and outputs a decision). We denote by $b = \mathbf{V}^{(\Pi_i)_{i \in [k]}}(\mathbb{x}; (\rho_i)_{i \in [k]})$ the IOP verifier's decision in the query phase, where k is the IOP's round complexity; then the two-step experiment $(\mathbb{x}, \mathbf{aux}) \leftarrow \tilde{\mathbf{P}}$ followed by $b \leftarrow \langle \tilde{\mathbf{P}}(\mathbf{aux}), \mathbf{V}(\mathbb{x}) \rangle$ can be written as

$$\left[\begin{array}{l} \mathbf{ai} \leftarrow \mathbf{D} \\ (\mathbb{x}, \mathbf{aux}_0) \leftarrow \tilde{\mathbf{P}}(\mathbf{ai}) \\ (\tilde{\Pi}_1, \mathbf{aux}_1) \leftarrow \tilde{\mathbf{P}}(\mathbf{aux}_0) \\ \rho_1 \leftarrow \{0, 1\}^{r_1} \\ \text{For } i \in [k-1] \setminus \{1\} : \\ \quad (\tilde{\Pi}_i, \mathbf{aux}_i) \leftarrow \tilde{\mathbf{P}}(\mathbf{aux}_{i-1}, \rho_{i-1}) \\ \quad \rho_i \leftarrow \{0, 1\}^{r_i} \\ \tilde{\Pi}_k \leftarrow \tilde{\mathbf{P}}(\mathbf{aux}_{k-1}, \rho_{k-1}) \\ b := \mathbf{V}^{(\Pi_i)_{i \in [k]}}(\mathbb{x}; (\rho_i)_{i \in [k]}) \end{array} \right] = \left[\begin{array}{l} \mathbf{ai} \leftarrow \mathbf{D} \\ (\mathbb{x}, \mathbf{aux}_0) \leftarrow \tilde{\mathbf{P}}(\mathbf{ai}) \\ \rho_0 := \perp \\ \text{For } i \in [k] : \\ \quad (\tilde{\Pi}_i, \mathbf{aux}_i) \leftarrow \tilde{\mathbf{P}}(\mathbf{aux}_{i-1}, \rho_{i-1}) \\ \quad \rho_i \leftarrow \{0, 1\}^{r_i} \\ b := \mathbf{V}^{(\Pi_i)_{i \in [k]}}(\mathbb{x}; (\rho_i)_{i \in [k]}) \end{array} \right]. \quad (1)$$

Remark 3.17. We often make use of simplifications as in Experiment 1, where the right-hand side sets ρ_0 to the empty string \perp . (And does so, implicitly, for \mathbf{aux}_k as well.)

We shall additionally make use of notation for *partial* (or full) executions of an interactive algorithm: for example, an execution of $\tilde{\mathbf{P}}$ until its second message is denoted $(\mathbb{x}, \tilde{\Pi}_1, (\tilde{\Pi}_2, \mathbf{aux}_2)) \leftarrow \tilde{\mathbf{P}}(\mathbf{ai}, \rho_1)$ (we omit auxiliary outputs that are not used in the remainder of the experiment). Therefore, Experiment 1 is also equivalent to

$$\left[\begin{array}{l} \mathbf{ai} \leftarrow \mathbf{D} \\ (\rho_i)_{i \in [k]} \leftarrow \{0, 1\}^{r_1 + \dots + r_k} \\ (\mathbb{x}, \tilde{\Pi}_1, \dots, \tilde{\Pi}_k) \leftarrow \tilde{\mathbf{P}}(\mathbf{ai}, \rho_1, \dots, \rho_{k-1}) \\ b := \mathbf{V}^{(\Pi_i)_{i \in [k]}}(\mathbb{x}; (\rho_i)_{i \in [k]}) \end{array} \right].$$

Public-query IOPs. In Section 6 we discuss IOPs that are *public-query*, which informally means that security (soundness or knowledge soundness) holds even if the IOP prover can “see” the queries that the IOP verifier makes. In contrast, the general definition of an IOP implicitly assumes that an (honest or malicious) IOP prover has no information about the queries that the IOP verifier makes during the interaction.

Definition 3.18. IOP $= (\mathbf{P}, \mathbf{V})$ for a relation R is **public-query** if the soundness (and knowledge soundness) condition holds even if the malicious IOP prover $\tilde{\mathbf{P}}$ receives, during the interaction, the index of each location queried by the IOP verifier \mathbf{V} (the moment it happens). In particular, for each $i \in [k]$, the IOP string $\tilde{\Pi}_i$ sent by $\tilde{\mathbf{P}}$ in the i -th round may depend on every query made by the IOP verifier in prior rounds (to prior IOP strings), in addition to depending on messages sent by the IOP verifier so far.

Note that a public-coin IOP is a public-query IOP, but the converse need not hold.

4 Interactive arguments based on PCPs

Theorem 4.1. Consider these two ingredients:

- PCP = (\mathbf{P}, \mathbf{V}) , a PCP system for a relation R with alphabet Σ , proof length l , and query complexity q ; and
- VC = (Gen, Commit, Open, Check), a vector commitment scheme over alphabet Σ .

Then ARG = $(\mathcal{G}, \mathcal{P}, \mathcal{V}) := \text{Kilian}[\text{PCP}, \text{VC}]$ (Construction 4.3) is a three-message public-coin interactive argument system for R , whose soundness error ϵ_{ARG} and knowledge soundness error κ_{ARG} satisfy the following for every $\epsilon > 0$ and $t_{\text{ARG}} \geq t_{\mathbf{V}} + t_{\text{VC,Check}} + \log|\Sigma| + \log l$:

$$\begin{aligned} \epsilon_{\text{ARG}}(\lambda, n, t_{\text{ARG}}) &\leq \epsilon_{\text{PCP}}(n) + \epsilon_{\text{VC}}(\lambda, l, q, t_{\text{VC}}) + \epsilon \quad \text{and} \\ \kappa_{\text{ARG}}(\lambda, n, t_{\text{ARG}}) &\leq \kappa_{\text{PCP}}(n) + \epsilon_{\text{VC}}(\lambda, l, q, t_{\text{VC}}) + \epsilon . \end{aligned}$$

Above, ϵ_{PCP} and κ_{PCP} are the soundness and knowledge soundness errors of PCP, and $t_{\text{VC}} = O\left(\frac{1}{\epsilon} \cdot t_{\text{ARG}}\right)$. Moreover, the knowledge extractor runs in time $t_{\mathcal{E}} = O(t_{\mathbf{E}} + t_{\text{VC}})$.

Corollary 4.2. Let ARG be as in Theorem 4.1. Assume that for any $n \in \mathbb{N}$, $\epsilon_{\text{VC}}(\cdot, \cdot, \cdot, t_{\text{VC}}) = \text{negl}(n)$ if $t_{\text{VC}} = \text{poly}(n)$. Then, given that $t_{\text{ARG}} = \text{poly}(n)$, we have

$$\begin{aligned} \epsilon_{\text{ARG}}(\lambda, n, t_{\text{ARG}}) &\leq \epsilon_{\text{PCP}}(n) + \text{negl}(n) \quad \text{and} \\ \kappa_{\text{ARG}}(\lambda, n, t_{\text{ARG}}) &\leq \kappa_{\text{PCP}}(n) + \text{negl}(n) . \end{aligned}$$

Proof. Let $p(n)$ be an arbitrary polynomial. We set ϵ to be $\frac{1}{2p(n)} > 0$. Hence, $t_{\text{VC}} = O\left(\frac{1}{\epsilon} \cdot t_{\text{ARG}}\right) = \text{poly}(n)$, which implies that $\epsilon_{\text{VC}}(\lambda, l, q, t_{\text{VC}}) = \text{negl}(n)$. Therefore,

$$\epsilon_{\text{ARG}}(\lambda, n, t_{\text{ARG}}) \leq \epsilon_{\text{PCP}}(n) + \text{negl}(n) + \frac{1}{2p(n)} < \epsilon_{\text{PCP}}(n) + \text{negl}(n) + \frac{1}{p(n)} .$$

Since p is an arbitrary polynomial, we conclude that

$$\epsilon_{\text{ARG}}(\lambda, n, t_{\text{ARG}}) \leq \epsilon_{\text{PCP}}(n) + \text{negl}(n) .$$

An analogous argument holds for κ_{ARG} . □

4.1 Construction

The construction of $(\mathcal{G}, \mathcal{P}, \mathcal{V}) := \text{Kilian}[\text{PCP}, \text{VC}]$ is specified below.

Construction 4.3. The argument generator \mathcal{G} receives as input a security parameter $\lambda \in \mathbb{N}$ and an instance size bound $n \in \mathbb{N}$, and works as follows.

$\mathcal{G}(\lambda, n)$:

1. Sample public parameter for the VC scheme: $\text{pp}_{\text{VC}} \leftarrow \text{VC.Gen}(1^\lambda, l(n))$.
2. Set public parameter for the interactive argument: $\text{pp} := \text{pp}_{\text{VC}}$.
3. Output pp .

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

1. \mathcal{P} 's commitment.
 - (a) Compute a PCP string: $\Pi \leftarrow \mathbf{P}(\mathbf{x}, \mathbf{w})$.
 - (b) Compute a vector commitment to the PCP string: $(\text{cm}, \text{aux}) \leftarrow \text{VC.Commit}(\text{pp}, \Pi)$.
 - (c) Send cm to \mathcal{V} .
2. \mathcal{V} 's challenge.
 - (a) Sample PCP verifier randomness: $\rho \leftarrow \{0, 1\}^r$.
 - (b) Send ρ to \mathcal{P} .
3. \mathcal{P} 's response.
 - (a) Run the PCP verifier $\mathbf{V}^\Pi(\mathbf{x}; \rho)$ to deduce its query set $\mathcal{Q} \subseteq [l]$.
 - (b) Compute a VC opening proof: $\text{pf} \leftarrow \text{VC.Open}(\text{pp}, \text{aux}, \mathcal{Q})$.
 - (c) Set $\text{ans} := \Pi[\mathcal{Q}]$.
 - (d) Send $(\mathcal{Q}, \text{ans}, \text{pf})$ to \mathcal{V} .
4. \mathcal{V} 's decision: check that $\mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbf{x}; \rho) = 1$ and $\text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1$.

The interactive argument consists of three messages: a prover message; a verifier message; and a prover message. The interactive argument is public-coin since the verifier's (only) message is a uniform random string. The efficiency measures of interactive arguments are as follows:

- the generator outputs public parameter of size $|\text{pp}_{\text{VC}}|$ bits;
- the prover-to-verifier communication consists of $|\text{cm}| + q \cdot (\log l + \log |\Sigma|) + |\text{pf}|$ bits;
- the verifier-to-prover communication consists of r bits;
- the time complexity of the argument generator is $t_{\text{VC.Gen}}$.
- the time complexity of the argument prover is $t_{\mathcal{P}} + t_{\text{VC.Commit}} + t_{\mathcal{V}} + t_{\text{VC.Open}}$;
- the time complexity of the argument verifier is $t_{\mathcal{V}} + t_{\text{VC.Check}}$.

Remark 4.4. There are vector commitments for which VC.Gen needs only the security parameter λ as input (i.e., VC.Gen works for every message size); for example, vector commitments based on Merkle trees have this property, because the public parameter consist of (the description of) a hash function, which suffices for every message size. In this case, the argument generator \mathcal{G} in Construction 4.3 also requires only λ as input and works for every instance size. This leads the notion of an interactive argument discussed in Remark 3.5.

Remark 4.5. In the plain-model variant of Construction 4.3 (see Remark 3.6), the public parameters $\text{pp} := \text{pp}_{\text{VC}}$ are sampled and sent by the argument verifier (resulting in a four-message protocol). Hence the plain-model variant is public-coin if (and only if) VC.Gen is a public-coin algorithm (its output includes all of its randomness).

4.2 Security reduction

To analyze the soundness and knowledge soundness for the argument system of Construction 4.3, it is important to understand how the argument system is related to the PCP system. The core of the security analysis is the construction of a PCP prover $\tilde{\mathbf{P}}$ from an argument prover $\tilde{\mathcal{P}}$ (which may or may not be malicious). More precisely, given a convincing argument prover $\tilde{\mathcal{P}}$, we want to obtain a convincing PCP prover $\tilde{\mathbf{P}}$, which we achieve via the *reductor* algorithm \mathcal{R} in Construction 4.8.

Recall that if \mathcal{V} accepts if and only if both \mathbf{V} and VC.Check accept. Hence, Lemma 4.6 shows that PCP strings generated by the reductor \mathcal{R} are, up to small errors, as convincing to the PCP verifier \mathbf{V} as the argument prover \mathcal{P} is to the argument verifier \mathcal{V} ; in other words, \mathcal{R} transforms an argument prover $\tilde{\mathcal{P}}$ into a

PCP prover $\tilde{\mathbf{P}}$.¹⁷ We later use this lemma to prove (adaptive) soundness and knowledge soundness of ARG in Sections 4.3 and 4.4, respectively.

Lemma 4.6. *There exists a probabilistic algorithm \mathcal{R} which, for every $\epsilon > 0$, auxiliary input distribution \mathcal{D} , size bound $t_{\text{ARG}} \geq t_{\mathbf{V}} + t_{\text{VC.Check}} + 2q \cdot (\log|\Sigma| + \log l)$, and t_{ARG} -size circuit $\tilde{\mathcal{P}}$, satisfies*

$$\Pr \left[\begin{array}{l} \mathbf{V}[\tilde{\mathcal{Q}}, \tilde{\Pi}](\mathbf{x}; \rho) \neq 1 \\ \wedge \mathbf{V}[\mathcal{Q}, \text{ans}](\mathbf{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \mid \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbf{x}, (\text{cm}, \text{aux})) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ (\tilde{\mathcal{Q}}, \tilde{\Pi}) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}, \cdot)}(\text{pp}, \text{cm}, \epsilon) \\ \rho \leftarrow \{0, 1\}^r \\ (\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\text{aux}, \rho) \end{array} \right] \\ \leq \epsilon_{\text{VC}}(\lambda, l, q, t_{\text{VC}}) + \epsilon ,$$

where $t_{\text{VC}} = O\left(\frac{1}{\epsilon} \cdot t_{\text{ARG}}\right)$. Moreover, \mathcal{R} makes l/ϵ queries to $\tilde{\mathcal{P}}$ and runs in $O(t_{\text{VC}})$ time.

We stress that in the experiment above the reducer \mathcal{R} is *independent* of ρ , since it does not receive the verifier randomness as input. (Otherwise, the lemma would be satisfied trivially with $\tilde{\mathcal{Q}} := \mathcal{Q}$ and $\tilde{\Pi}[\tilde{\mathcal{Q}}] := \text{ans}$.)

We now construct \mathcal{R} , which will be convenient to separate into two parts: a sampling subroutine \mathcal{S} followed by a post-processing layer $\mathcal{R}_{\text{post}}$ that deterministically pieces together a PCP string $\tilde{\Pi}$ from the samples obtained by \mathcal{S} (and outputs the set \mathcal{Q} of “filled-in” coordinates along with $\tilde{\Pi}$).

Construction 4.7. We construct the *sampler* \mathcal{S} as follows.

$\mathcal{S}^{\tilde{\mathcal{P}}(\text{aux}, \cdot)}(\text{pp}, \text{cm}, \mathbf{N})$:

1. Initialize $\mathcal{K} := \emptyset$.
2. Repeat the following \mathbf{N} times:
 - (a) Sample PCP verifier randomness: $\rho' \leftarrow \{0, 1\}^r$.
 - (b) Obtain $(\mathcal{Q}', \text{ans}', \text{pf}') \leftarrow \tilde{\mathcal{P}}(\text{aux}, \rho')$.
 - (c) If $\text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}', \text{ans}', \text{pf}') = 1$, add $(\mathcal{Q}', \text{ans}', \text{pf}')$ to \mathcal{K} .
3. Output \mathcal{K} .

The algorithm \mathcal{S} makes \mathbf{N} queries to $\tilde{\mathcal{P}}$, and runs in time¹⁸

$$t_{\mathcal{S}} \leq \mathbf{N} \cdot (t_{\text{ARG}} + t_{\mathbf{V}} + t_{\text{VC.Check}}) \leq 3\mathbf{N} \cdot t_{\text{ARG}} .$$

Construction 4.8. The *reducer* \mathcal{R} is defined as follows. (Below, σ is an arbitrary symbol in the alphabet Σ .)

$\mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}, \cdot)}(\text{pp}, \text{cm}, \epsilon)$:¹⁹

1. Set $\mathbf{N} := \frac{1}{\epsilon}$ and run $\mathcal{K} \leftarrow \mathcal{S}^{\tilde{\mathcal{P}}(\text{aux}, \cdot)}(\text{pp}, \text{cm}, \mathbf{N})$.
2. Run $(\tilde{\mathcal{Q}}, \tilde{\Pi}) \leftarrow \mathcal{R}_{\text{post}}(\mathcal{K}, l)$.

¹⁷Moreover, \mathcal{R} preserves uniformity: if $\tilde{\mathcal{P}}$ is a uniform algorithm, then so is $\tilde{\mathbf{P}}$.

¹⁸Note that the $2q(\log|\Sigma| + \log l)$ overhead incurred by copying $(\mathcal{Q}', \text{ans}', \text{pf}')$ into \mathcal{K} is accounted for in the difference between t_{ARG} and the other terms.

¹⁹We also denote by $\mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}, \cdot)}(\text{pp}, \text{cm}, \epsilon; \rho)$, where $\rho = (\rho^{(\ell)})_{\ell \in [\mathbf{N}]}$ (and similarly for \mathcal{S}) the deterministic algorithm that uses $\rho^{(\ell)}$ as the randomness for \mathcal{S} 's ℓ -th sample. This allows the PCP and IOP provers of Constructions 4.10 and 5.8 to be deterministic.

3. Output $(\tilde{\mathcal{Q}}, \tilde{\Pi})$.

The reducer's second step is an execution of the following (deterministic) *post-processing* algorithm.

$\mathcal{R}_{\text{post}}(\mathcal{K}, l)$:

1. Initialize $\tilde{\Pi} := \sigma^l$ and $\tilde{\mathcal{Q}} := \emptyset$.
2. For every $(\mathcal{Q}', \text{ans}', \text{pf}') \in \mathcal{K}$:
 - (a) Set $\tilde{\mathcal{Q}} := \tilde{\mathcal{Q}} \cup \mathcal{Q}'$.
 - (b) For every $q \in \mathcal{Q}'$, set $\tilde{\Pi}[q] := \text{ans}'[q]$.
3. Output $(\tilde{\mathcal{Q}}, \tilde{\Pi})$.

Note that \mathcal{R} makes $N = l/\epsilon$ queries to $\tilde{\mathcal{P}}$ by construction, whose total $N \cdot t_{\text{ARG}}$ time dominates that of \mathcal{R} .

Proof. Throughout this proof, probabilistic expressions are with respect to the following experiment unless explicitly denoted otherwise:

$$\left[\begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbf{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ (\text{cm}, \text{aux}_1) \leftarrow \tilde{\mathcal{P}}(\text{aux}_0) \\ \rho \leftarrow \{0, 1\}^r \\ (\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\text{aux}_1, \rho) \end{array} \right] = \left[\begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbf{x}, (\text{cm}, \text{aux})) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ (\tilde{\mathcal{Q}}, \tilde{\Pi}) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}, \cdot)}(\text{pp}, \text{cm}, \epsilon) \\ \rho \leftarrow \{0, 1\}^r \\ (\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\text{aux}, \rho) \end{array} \right]. \quad (2)$$

Our goal is to upper bound the probability of the following expression:

$$\left[\begin{array}{l} \mathbf{V}[\tilde{\mathcal{Q}}, \tilde{\Pi}](\mathbf{x}; \rho) \neq 1 \\ \wedge \mathbf{V}[\mathcal{Q}, \text{ans}](\mathbf{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \right]. \quad (3)$$

Observe that Eq. 3 implies that either (i) $\tilde{\Pi}$ and ans disagree at a position $q \in \mathcal{Q} \cap \tilde{\mathcal{Q}}$; or (ii) there is a query q in $\mathcal{Q} \setminus \tilde{\mathcal{Q}}$. We analyze the two cases separately.

Valid openings with disagreeing answers. Our goal is to prove the following bound:

$$\Pr \left[\begin{array}{l} \exists q \in \mathcal{Q} \cap \tilde{\mathcal{Q}} : \text{ans}[q] \neq \tilde{\Pi}[q] \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \right] \leq \epsilon_{\text{VC}}(\lambda, l, \mathbf{q}, t_{\text{VC}}),$$

where $t_{\text{VC}} \leq 4N \cdot t_{\text{ARG}}$.

Consider the following adversary A_{VC} against the vector commitment scheme, which follows Experiment 2 (without executing $\mathcal{R}_{\text{post}}$) and attempts to find a collision using the output \mathcal{K} of the sampler \mathcal{S} .

$A_{\text{VC}}(\text{pp}, \text{ai})$:

1. Run $(\mathbf{x}, (\text{cm}, \text{aux})) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai})$.
2. Sample $\rho \leftarrow \{0, 1\}^r$.
3. Run $(\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\text{aux}, \rho)$.
4. Run $\mathcal{K} \leftarrow \mathcal{S}^{\tilde{\mathcal{P}}(\text{aux}, \cdot)}(\text{pp}, \text{cm}, N)$, with $N = \frac{l}{\epsilon}$ (as in Construction 4.8).
5. If there are $(\mathcal{Q}', \text{ans}', \text{pf}') \in \mathcal{K}$ and $q \in \mathcal{Q}' \cap \mathcal{Q}$ with $\text{ans}'[q] \neq \text{ans}[q]$, output $(\text{cm}, \text{ans}, \text{ans}', \mathcal{Q}, \mathcal{Q}', \text{pf}, \text{pf}')$.
6. Otherwise, output (the “dummy” tuple) $(\text{cm}, \text{ans}, \text{ans}, \mathcal{Q}, \mathcal{Q}, \text{pf}, \text{pf})$.

The time complexity of the sampler \mathcal{S} is at most $3N \cdot t_{\text{ARG}}$ and the collision-finding check (Step 5) runs in time $N \cdot 2q(\log|\Sigma| + \log l) \leq N \cdot t_{\text{ARG}}$,²⁰ so the time complexity of A_{VC} is $t_{\text{VC}} \leq 4N \cdot t_{\text{ARG}}$.

Therefore, according to Definition 3.8,

$$\begin{aligned} & \Pr \left[\begin{array}{l} \exists q \in \mathcal{Q} \cap \tilde{\mathcal{Q}} : \text{ans}[q] \neq \tilde{\Pi}[q] \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \right] \\ &= \Pr \left[\begin{array}{l} |\mathcal{Q}| = |\mathcal{Q}'| = q \\ \wedge \exists q \in \mathcal{Q} \cap \mathcal{Q}' : \text{ans}[q] \neq \text{ans}'[q] \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}', \text{ans}', \text{pf}') = 1 \end{array} \middle| \begin{array}{l} \text{pp} \leftarrow \text{VC.Gen}(1^\lambda, l) \\ \text{ai} \leftarrow \mathcal{D} \\ (\text{cm}, \text{ans}, \text{ans}', \\ \mathcal{Q}, \mathcal{Q}', \text{pf}, \text{pf}') \leftarrow A_{\text{VC}}(\text{pp}, \text{ai}) \end{array} \right] \\ &\leq \epsilon_{\text{VC}}(\lambda, l, q, t_{\text{VC}}) . \end{aligned}$$

Missing positions in $\tilde{\Pi}$. We now upper bound the probability that there is a missing position in $\tilde{\Pi}$; we claim that

$$\Pr \left[\begin{array}{l} \mathcal{Q} \setminus \tilde{\mathcal{Q}} \neq \emptyset \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \right] \leq \epsilon .$$

Fix a public parameter-auxiliary input pair (pp, ai) (recall that these are obtained in the first steps of Experiment 2), and define the *weight* of a coordinate $q \in [l]$ with respect to (pp, ai) as

$$\delta_{\text{pp}, \text{ai}}(q) := \Pr \left[\begin{array}{l} q \in \mathcal{Q} \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \middle| \begin{array}{l} (\mathbf{x}, \text{cm}, \text{aux}) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ \rho \leftarrow \{0, 1\}^r \\ (\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\text{aux}, \rho) \end{array} \right] .$$

Then, by a union bound over $q \in [l]$,

$$\begin{aligned} & \Pr \left[\begin{array}{l} \mathcal{Q} \setminus \tilde{\mathcal{Q}} \neq \emptyset \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \right] \\ &= \Pr \left[\begin{array}{l} \exists q \in [l] : q \in \mathcal{Q} \wedge q \notin \tilde{\mathcal{Q}} \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \right] \\ &\leq \max_{(\text{pp}, \text{ai})} \left\{ \Pr \left[\begin{array}{l} \exists q \in [l] : q \in \mathcal{Q} \wedge q \notin \tilde{\mathcal{Q}} \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \middle| \begin{array}{l} (\mathbf{x}, \text{cm}, \text{aux}) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ \rho \leftarrow \{0, 1\}^r \\ (\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\text{aux}, \rho) \end{array} \right] \right\} \\ &= \max_{(\text{pp}, \text{ai})} \left\{ \sum_{q \in [l]} \delta_{\text{pp}, \text{ai}}(q) \cdot (1 - \delta_{\text{pp}, \text{ai}}(q))^N \right\} \\ &\leq \frac{1}{N} , \end{aligned}$$

where the last inequality follows from $\delta \cdot (1 - \delta)^N \leq 1/N$ for any $\delta \in [0, 1]$.²¹ Finally, plugging in $N = \frac{1}{\epsilon}$ bounds Eq. 3 as desired, concluding the proof:

$$\Pr \left[\begin{array}{l} \mathbf{V}^{[\tilde{\mathcal{Q}}, \tilde{\Pi}]}(\mathbf{x}; \rho) \neq 1 \\ \wedge \mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbf{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \right]$$

²⁰Searching for an intersection between \mathcal{Q}' and \mathcal{Q} takes $2q \log l$ time, while each check for a symbol mismatch takes $2 \log|\Sigma|$. We omit the time required to produce the output (which can be accounted for in the runtime of \mathcal{S}).

²¹A simple derivation of the inequality is the following: with $f(x) = x \cdot (1 - x)^N$, we have $\frac{d}{dx} f(\delta) = 0 \iff \delta = \frac{1}{N+1}$. As $f(0) = f(1) = 0$ and δ is the only critical point in $[0, 1]$, it achieves the maximum $f(\delta) \leq 1/N$.

$$\begin{aligned} &\leq \Pr \left[\begin{array}{l} \exists q \in \mathcal{Q} \cap \tilde{\mathcal{Q}} : \text{ans}[q] \neq \tilde{\Pi}[q] \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \right] + \Pr \left[\begin{array}{l} \mathcal{Q} \setminus \tilde{\mathcal{Q}} \neq \emptyset \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \right] \\ &\leq \epsilon_{\text{VC}}(\lambda, l, \mathbf{q}, t_{\text{VC}}) + \epsilon . \quad \square \end{aligned}$$

4.3 Adaptive soundness

Lemma 4.9. *For every $\epsilon > 0$, security parameter $\lambda \in \mathbb{N}$, instance size bound $n \in \mathbb{N}$, auxiliary input distribution \mathcal{D} , circuit size bound $t_{\text{ARG}} \geq t_{\text{V}} + t_{\text{VC.Check}} + 2\mathbf{q} \cdot (\log|\Sigma| + \log l)$ and t_{ARG} -size circuit $\tilde{\mathcal{P}}$, the soundness error of the argument system in Construction 4.3 satisfies*

$$\epsilon_{\text{ARG}}(\lambda, n, t_{\text{ARG}}) \leq \epsilon_{\text{PCP}}(n) + \epsilon_{\text{VC}}(\lambda, l, \mathbf{q}, t_{\text{VC}}) + \epsilon ,$$

where $t_{\text{VC}} = O\left(\frac{1}{\epsilon} \cdot t_{\text{ARG}}\right)$.

Proof. Recall, from Definition 3.3 and Construction 4.3, that our goal is to upper bound

$$\begin{aligned} &\Pr \left[\begin{array}{l} |\mathbb{x}| \leq n \\ \wedge \mathbb{x} \notin L(R) \\ \wedge b = 1 \end{array} \left| \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbb{x}, \text{aux}) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ b \leftarrow \langle \tilde{\mathcal{P}}(\text{aux}), \mathcal{V}(\text{pp}, \mathbb{x}) \rangle \end{array} \right. \right] \\ &= \Pr \left[\begin{array}{l} |\mathbb{x}| \leq n \\ \wedge \mathbb{x} \notin L(R) \\ \wedge \mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbb{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \left| \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbb{x}, (\text{cm}, \text{aux})) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ \rho \leftarrow \{0, 1\}^r \\ (\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\text{aux}, \rho) \end{array} \right. \right] . \end{aligned}$$

(Since $\tilde{\mathcal{P}}$ sends cm as the first message in Construction 4.3, the former experiment is equivalent to the latter, where we omit the auxiliary state of the choice of instance and aux denotes that of the first message.)

As in Lemma 4.6, we consider the following experiment (a restatement of Experiment 2), which augments the above by executing \mathcal{R} , and thus leaves the probability unchanged.

$$\left[\begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbb{x}, (\text{cm}, \text{aux})) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ (\tilde{\mathcal{Q}}, \tilde{\Pi}) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}, \cdot)}(\text{pp}, \text{cm}, \epsilon) \\ \rho \leftarrow \{0, 1\}^r \\ (\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\text{aux}, \rho) \end{array} \right] .$$

By total probability,

$$\Pr \left[\begin{array}{l} |\mathbb{x}| \leq n \\ \wedge \mathbb{x} \notin L(R) \\ \wedge \mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbb{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \right]$$

$$= \Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge \mathbf{x} \notin L(R) \\ \wedge \mathbf{V}^{[\tilde{\mathcal{Q}}, \tilde{\Pi}]}(\mathbf{x}; \rho) = 1 \\ \wedge \mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbf{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \right] + \Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge \mathbf{x} \notin L(R) \\ \wedge \mathbf{V}^{[\tilde{\mathcal{Q}}, \tilde{\Pi}]}(\mathbf{x}; \rho) \neq 1 \\ \wedge \mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbf{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \right].$$

We first bound the probability of the leftmost term by the PCP system's soundness error (i.e., Definition 3.11 with respect to PCP).

Construction 4.10. We define the auxiliary input distribution \mathbf{D} of the PCP prover $\tilde{\mathbf{P}}$ as follows:

\mathbf{D} :

1. Sample $\text{pp} \leftarrow \mathcal{G}(1^\lambda, n)$ followed by $\text{ai} \leftarrow \mathcal{D}$ and $\boldsymbol{\rho} := (\rho^{(\ell)})_{\ell \in [\mathbb{N}]} \leftarrow (\{0, 1\}^r)^{\mathbb{N}}$.
2. Output $\text{ai} := (\text{pp}, \text{ai}, \boldsymbol{\rho})$.

The PCP prover is then given by the following next message functions.

$\tilde{\mathbf{P}}(\text{ai})$:

1. Parse ai as $(\text{pp}, \text{ai}, \boldsymbol{\rho})$.
2. Run $(\mathbf{x}, \text{aux}) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai})$.
3. Set $\text{aux} := (\text{pp}, \text{aux}_0, \boldsymbol{\rho})$.
4. Output (\mathbf{x}, aux) .

$\tilde{\mathbf{P}}(\text{aux})$:

1. Parse aux as $(\text{pp}, \text{aux}_0, \boldsymbol{\rho})$.
2. Run $(\text{cm}, \text{aux}_1) \leftarrow \tilde{\mathcal{P}}(\text{aux}_0)$.
3. Run $(\tilde{\mathcal{Q}}, \tilde{\Pi}) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}_1, \cdot)}(\text{pp}, \text{cm}, \epsilon; \boldsymbol{\rho})$.
4. Output $\tilde{\Pi}$.

Using Definition 3.11,²²

$$\begin{aligned} \Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge \mathbf{x} \notin L(R) \\ \wedge \mathbf{V}^{[\tilde{\mathcal{Q}}, \tilde{\Pi}]}(\mathbf{x}; \rho) = 1 \\ \wedge \mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbf{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \right] &\leq \Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge \mathbf{x} \notin L(R) \\ \wedge \mathbf{V}^{[\tilde{\mathcal{Q}}, \tilde{\Pi}]}(\mathbf{x}; \rho) = 1 \end{array} \right] \\ &\leq \Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge \mathbf{x} \notin L(R) \\ \wedge \mathbf{V}^{\tilde{\Pi}}(\mathbf{x}) = 1 \end{array} \middle| \begin{array}{l} \text{ai} \leftarrow \mathbf{D} \\ (\mathbf{x}, \text{aux}) \leftarrow \tilde{\mathbf{P}}(\text{ai}) \\ \tilde{\Pi} \leftarrow \tilde{\mathbf{P}}(\text{aux}) \end{array} \right] \\ &\leq \epsilon_{\text{PCP}}(n). \end{aligned}$$

Lastly, an application of Lemma 4.6 yields

$$\Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge (\mathbf{x}, \mathbf{w}) \notin R \\ \wedge \mathbf{V}^{[\tilde{\mathcal{Q}}, \tilde{\Pi}]}(\mathbf{x}; \rho) \neq 1 \\ \wedge \mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbf{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \right] \leq \Pr \left[\begin{array}{l} \mathbf{V}^{[\tilde{\mathcal{Q}}, \tilde{\Pi}]}(\mathbf{x}; \rho) \neq 1 \\ \wedge \mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbf{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \right]$$

²²Note that the prover in Definition 3.11 corresponds to the sequential execution of both steps in Construction 4.10.

$$\leq \epsilon_{\text{VC}}(\lambda, l, q, t_{\text{VC}}) + \epsilon ,$$

where $t_{\text{VC}} \leq \frac{4l}{\epsilon} \cdot t_{\text{ARG}}$, which concludes the proof. \square

4.4 Adaptive knowledge soundness

Lemma 4.11. *For every $\epsilon > 0$, security parameter $\lambda \in \mathbb{N}$, instance size bound $n \in \mathbb{N}$, auxiliary input distribution \mathcal{D} , circuit size bound $t_{\text{ARG}} \geq t_{\text{V}} + t_{\text{VC.Check}} + 2q \cdot (\log|\Sigma| + \log l)$ and t_{ARG} -size circuit $\tilde{\mathcal{P}}$, the knowledge soundness error of the argument system obtained by Construction 4.3 satisfies*

$$\kappa_{\text{ARG}}(\lambda, n, t_{\text{ARG}}) \leq \kappa_{\text{PCP}}(n) + \epsilon_{\text{VC}}(\lambda, l, q, t_{\text{VC}}) + \epsilon ,$$

where $t_{\text{VC}} = O\left(\frac{1}{\epsilon} \cdot t_{\text{ARG}}\right)$. If the PCP extractor's runtime is t_{E} , the argument system's extractor is $t_{\mathcal{E}} = t_{\text{E}} + O(t_{\text{VC}})$.

Construction 4.12. Let \mathbf{E} be the extractor for PCP. We use $\tilde{\mathbf{P}}$ (Construction 4.10) and \mathbf{E} to construct the knowledge extractor \mathcal{E} for ARG as follows.

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

1. Sample $\rho \leftarrow (\{0, 1\}^r)^{\mathbb{N}}$.
2. Set $\text{aux} := (\text{pp}, \text{aux}, \rho)$ and run $\tilde{\Pi} \leftarrow \tilde{\mathbf{P}}(\text{aux})$.²³
3. Run $\mathbb{w} \leftarrow \mathbf{E}(\mathbb{x}, \tilde{\Pi})$.
4. Output \mathbb{w} .

Note that \mathcal{E} executes \mathbf{E} once and $\tilde{\mathcal{P}}$ for l/ϵ times (as $\tilde{\mathbf{P}}$ runs the reductor \mathcal{R} , which in turn sets $\mathbb{N} = l/\epsilon$ and repeats \mathbb{N} executions of $\tilde{\mathcal{P}}$). Therefore, $t_{\mathcal{E}} = t_{\text{E}} + O(t_{\text{VC}})$.

Proof. From Definition 3.4 and Construction 4.3, our goal is to upper bound

$$\begin{aligned} & \Pr \left[\begin{array}{l} |\mathbb{x}| \leq n \\ \wedge (\mathbb{x}, \mathbb{w}) \notin R \\ \wedge b = 1 \end{array} \middle| \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbb{x}, \text{aux}) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ b \stackrel{\text{tr}}{\leftarrow} \langle \tilde{\mathcal{P}}(\text{aux}), \mathcal{V}(\text{pp}, \mathbb{x}) \rangle \\ \mathbb{w} \leftarrow \mathcal{E}^{\tilde{\mathcal{P}}(\text{aux})}(\text{pp}, \mathbb{x}, \text{tr}) \end{array} \right] \\ &= \Pr \left[\begin{array}{l} |\mathbb{x}| \leq n \\ \wedge (\mathbb{x}, \mathbb{w}) \notin R \\ \wedge \mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbb{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \middle| \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbb{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ (\text{cm}, \text{aux}_1) \leftarrow \tilde{\mathcal{P}}(\text{aux}_0) \\ \rho \leftarrow \{0, 1\}^r \\ (\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\text{aux}_1, \rho) \\ \text{tr} := (\text{cm}, \rho, \mathcal{Q}, \text{ans}, \text{pf}) \\ \mathbb{w} \leftarrow \mathcal{E}^{\tilde{\mathcal{P}}(\text{aux}_0)}(\text{pp}, \mathbb{x}, \text{tr}) \end{array} \right] . \end{aligned}$$

Now, note that by Constructions 4.10 and 4.12, the experiment above is equivalent to the following.

²³Note that, by Construction 4.10, $\tilde{\mathbf{P}}$ only executes $\tilde{\mathcal{P}}(\text{aux})$ (calls to $\tilde{\mathcal{P}}(\text{aux}', \rho)$ are continuations of executions of $\tilde{\mathcal{P}}(\text{aux})$). Therefore, an equivalent construction gives $\tilde{\mathbf{P}}$ oracle access to $\tilde{\mathcal{P}}(\text{aux})$ (rather than to $\tilde{\mathcal{P}}$), excluding aux from aux .

$$\left[\begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbf{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ (\text{cm}, \text{aux}_1) \leftarrow \tilde{\mathcal{P}}(\text{aux}_0) \\ \rho \leftarrow \{0, 1\}^r \\ (\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\text{aux}_1, \rho) \\ \boldsymbol{\rho} \leftarrow (\{0, 1\}^r)^N \\ \mathbf{aux} := (\text{pp}, \text{aux}_0, \rho) \\ \tilde{\Pi} \leftarrow \tilde{\mathbf{P}}(\mathbf{aux}) \\ \mathbf{w} \leftarrow \mathbf{E}(\mathbf{x}, \tilde{\Pi}) \end{array} \right] = \left[\begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbf{x}, (\text{cm}, \text{aux})) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ \rho \leftarrow \{0, 1\}^r \\ (\mathcal{Q}, \text{ans}, \text{pf}) \leftarrow \tilde{\mathcal{P}}(\text{aux}, \rho) \\ (\tilde{\mathcal{Q}}, \tilde{\Pi}) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}, \cdot)}(\text{pp}, \text{cm}, \epsilon) \\ \mathbf{w} \leftarrow \mathbf{E}(\mathbf{x}, \tilde{\Pi}) \end{array} \right] .$$

We thus consider the above experiment, which augments that of Lemma 4.6 by appending an execution of \mathbf{E} , for the rest of the proof. Note that, as in Lemma 4.9, the experiment consisting of $(\mathbf{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai})$ followed by $(\text{cm}, \text{aux}_1) \leftarrow \tilde{\mathcal{P}}(\text{aux}_0)$ can be replaced by $(\mathbf{x}, (\text{cm}, \text{aux})) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai})$: since \mathcal{R} only uses the second auxiliary state (which is obtained deterministically from the first), the former can be omitted. Note, moreover, that the explicit randomness for $\tilde{\mathbf{P}}$ on the left-hand side is replaced with sampling by \mathcal{R} on the right-hand side.

By total probability,

$$\begin{aligned} & \Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge (\mathbf{x}, \mathbf{w}) \notin R \\ \wedge \mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbf{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \right] \\ &= \Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge (\mathbf{x}, \mathbf{w}) \notin R \\ \wedge \mathbf{V}^{[\tilde{\mathcal{Q}}, \tilde{\Pi}]}(\mathbf{x}; \rho) = 1 \\ \wedge \mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbf{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \right] + \Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge (\mathbf{x}, \mathbf{w}) \notin R \\ \wedge \mathbf{V}^{[\tilde{\mathcal{Q}}, \tilde{\Pi}]}(\mathbf{x}; \rho) \neq 1 \\ \wedge \mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbf{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \right] . \end{aligned}$$

With the PCP prover $\tilde{\mathbf{P}}$ in Construction 4.10 and using Definition 3.12, we have

$$\Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge (\mathbf{x}, \mathbf{w}) \notin R \\ \wedge \mathbf{V}^{[\tilde{\mathcal{Q}}, \tilde{\Pi}]}(\mathbf{x}; \rho) = 1 \\ \wedge \mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbf{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \right] \leq \Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge (\mathbf{x}, \mathbf{w}) \notin R \\ \wedge \mathbf{V}^{\tilde{\Pi}}(\mathbf{x}) = 1 \end{array} \middle| \begin{array}{l} \mathbf{ai} \leftarrow \mathbf{D} \\ (\mathbf{x}, \tilde{\Pi}) \leftarrow \tilde{\mathbf{P}} \\ \mathbf{w} \leftarrow \mathbf{E}(\mathbf{x}, \tilde{\Pi}) \end{array} \right] \leq \kappa_{\text{PCP}}(n) .$$

Finally, Lemma 4.6 implies

$$\begin{aligned} & \Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge \mathbf{V}^{[\tilde{\mathcal{Q}}, \tilde{\Pi}]}(\mathbf{x}; \rho) \neq 1 \\ \wedge \mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbf{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \right] \leq \Pr \left[\begin{array}{l} \mathbf{V}^{[\tilde{\mathcal{Q}}, \tilde{\Pi}]}(\mathbf{x}; \rho) \neq 1 \\ \wedge \mathbf{V}^{[\mathcal{Q}, \text{ans}]}(\mathbf{x}; \rho) = 1 \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \end{array} \right] \\ & \leq \epsilon_{\text{VC}}(\lambda, l, q, t_{\text{VC}}) + \epsilon , \end{aligned}$$

where $t_{\text{VC}} \leq \frac{4l}{\epsilon} \cdot t_{\text{ARG}}$, which concludes the proof. \square

5 Interactive arguments based on public-coin IOPs

Theorem 5.1. *Consider these two ingredients:*

- IOP = (\mathbf{P}, \mathbf{V}) , a public-coin IOP system for a relation R with round complexity k , alphabet Σ , proof length l , and query complexity q ; and
- VC = (Gen, Commit, Open, Check), a vector commitment scheme over alphabet Σ .

Then $\text{ARG} = (\mathcal{G}, \mathcal{P}, \mathcal{V}) := \text{IBCS}[\text{IOP}, \text{VC}]$ (Construction 5.3) is a $(2k + 1)$ -message public-coin interactive argument system for R whose soundness error ϵ_{ARG} and knowledge soundness error κ_{ARG} satisfy the following for every $\epsilon > 0$ and $t_{\text{ARG}} \geq t_{\mathbf{V}} + t_{\text{VC.Check}} + \log|\Sigma| + \log l_{\max}$:

$$\begin{aligned} \epsilon_{\text{ARG}}(\lambda, n, t_{\text{ARG}}) &\leq \epsilon_{\text{IOP}}(n) + \epsilon_{\text{VC}}(\lambda, l_{\max}, q_{\max}, t_{\text{VC}}) + \epsilon \quad \text{and} \\ \kappa_{\text{ARG}}(\lambda, n, t_{\text{ARG}}) &\leq \kappa_{\text{IOP}}(n) + \epsilon_{\text{VC}}(\lambda, l_{\max}, q_{\max}, t_{\text{VC}}) + \epsilon, \end{aligned}$$

where $t_{\text{VC}} = O\left(\frac{k \cdot l}{\epsilon} \cdot t_{\text{ARG}}\right)$. Moreover, the knowledge extractor runs in time $t_{\mathcal{E}} = t_{\mathbf{E}} + O(t_{\text{VC}})$.

Corollary 5.2. *Let ARG be as in Theorem 5.1. Assume that for any $n \in \mathbb{N}$, $\epsilon_{\text{VC}}(\cdot, \cdot, \cdot, t_{\text{VC}}) = \text{negl}(n)$ if $t_{\text{VC}} = \text{poly}(n)$. Then, given that $t_{\text{ARG}} = \text{poly}(n)$, we have*

$$\begin{aligned} \epsilon_{\text{ARG}}(\lambda, n, t_{\text{ARG}}) &\leq \epsilon_{\text{IOP}}(n) + \text{negl}(n) \quad \text{and} \\ \kappa_{\text{ARG}}(\lambda, n, t_{\text{ARG}}) &\leq \kappa_{\text{IOP}}(n) + \text{negl}(n). \end{aligned}$$

5.1 Construction

We describe below the construction of the interactive argument, which we denote $(\mathcal{P}, \mathcal{V}) := \text{IBCS}[\text{IOP}, \text{VC}]$.

Construction 5.3. The argument generator \mathcal{G} receives as input a security parameter $\lambda \in \mathbb{N}$ and an instance size bound $n \in \mathbb{N}$, and works as follows.

$\mathcal{G}(\lambda, n)$:

1. Sample public parameters for the VC scheme: $\text{pp}_{\text{VC}} \leftarrow \text{VC.Gen}(1^\lambda, l_{\max}(n))$.²⁴
2. Set public parameters for the interactive argument: $\text{pp} := \text{pp}_{\text{VC}}$.
3. Output pp .

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

1. \mathcal{P} 's commitments.

For $i \in [k]$:

- (a) \mathcal{P} 's i -th commitment.

- i. Compute the i -th IOP string $\Pi_i \in \Sigma^{l_i}$ and auxiliary state:²⁵

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

²⁴Alternatively, VC.Gen could sample one set of public parameters $\text{pp}_{\text{VC},i}$ per proof length l_i and set $\text{pp} := (\text{pp}_{\text{VC},i})_{i \in [k]}$. For simplicity, we consider a single one and assume $\tilde{\mathcal{P}}$ pads proofs where $l_i < l_{\max}$ with a fixed symbol $\sigma \in \Sigma$ where appropriate.

²⁵Note the implicit setting of $\mathbf{aux}_k := \perp$, since \mathbf{P} only outputs the last proof string Π_k at the end of the interaction.

- ii. Compute a VC commitment to the IOP string: $(\text{cm}_i, \text{aux}_i) \leftarrow \text{VC.Commit}(\text{pp}, \Pi_i)$.
- iii. Send cm_i to \mathcal{V} .
- (b) \mathcal{V} 's i -th challenge.
 - i. Sample the i -th IOP verifier randomness $\rho_i \leftarrow \{0, 1\}^{r_i}$.
 - ii. Send ρ_i to \mathcal{P} .
- 2. \mathcal{P} 's response.
 - (a) Run the IOP verifier $\mathbf{V}^{\Pi_1, \dots, \Pi_k}(\mathbb{x}; \rho_1, \dots, \rho_k)$ to deduce $\mathcal{Q}_1, \dots, \mathcal{Q}_k$, where $\mathcal{Q}_i \subseteq [l_i]$ is the query set of \mathbf{V} to Π_i .
 - (b) For every $i \in [k]$, compute an opening proof $\text{pf}_i \leftarrow \text{VC.Open}(\text{pp}, \text{aux}_i, \mathcal{Q}_i)$ and set $\text{ans}_i := \Pi_i[\mathcal{Q}_i]$.
 - (c) Send $((\mathcal{Q}_i, \text{ans}_i, \text{pf}_i))_{i \in [k]}$ to \mathcal{V} .
- 3. \mathcal{V} 's decision.
 - Check that $\mathbf{V}^{((\mathcal{Q}_i, \text{ans}_i))_{i \in [k]}}(\mathbb{x}; \rho_1, \dots, \rho_k) = 1$ and $\text{VC.Check}(\text{pp}, \text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) = 1$ for all $i \in [k]$.²⁶

The protocol has $k + 1$ rounds: the first k simulate the IOP, and in the last $\tilde{\mathcal{P}}$ sends the query set assignments along with their opening proofs. Moreover, the protocol is public coin because the verifier's messages consist of random strings. We comment on the protocol's efficiency measures:

- the generator communication to prover and verifier consists of $|\text{pp}_{\text{VC}}|$ bits;
- the prover-to-verifier communication consists of $\sum_{i \in [k]} (|\text{cm}_i| + q \cdot (\log l_i + \log |\Sigma|) + |\text{pf}_i|)$ bits;
- the verifier-to-prover communication consists of $r = \sum_{i \in [k]} r_i$ bits;
- the time complexity of the argument generator is $t_{\text{VC.Gen}}$.
- the time complexity of the argument prover is $t_{\mathcal{P}} + k \cdot (t_{\text{VC.Commit}} + t_{\text{VC.Open}}) + t_{\mathbf{V}}$;
- the time complexity of the argument verifier is $t_{\mathcal{V}} + k \cdot t_{\text{VC.Check}}$.

5.2 Security reduction

As in Section 4.2, our analysis relies on the following security reduction lemma that relates the acceptance probability of $\text{ARG} := \text{IBCS}[\text{IOP}, \text{VC}]$ with that of IOP.

Lemma 5.4. *There exists a probabilistic algorithm \mathcal{R} which, for every size bound $t_{\text{ARG}} \geq t_{\mathbf{V}} + t_{\text{VC.Check}} + \log |\Sigma| + \log l_{\max}$, circuit $\tilde{\mathcal{P}}$ of size t_{ARG} and $\epsilon > 0$, satisfies*

$$\Pr \left[\begin{array}{l} \mathbf{V}^{((\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i))_{i \in [k]}}(\mathbb{x}; (\rho_i)_{i \in [k]}) \neq 1 \\ \wedge \mathbf{V}^{((\mathcal{Q}_i, \text{ans}_i))_{i \in [k]}}(\mathbb{x}; (\rho_i)_{i \in [k]}) = 1 \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\text{pp}, \text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) \right) = 1 \end{array} \right] \leq \epsilon_{\text{VC}}(\lambda, l_{\max}, q_{\max}, t_{\text{VC}}) + \epsilon,$$

$$\left[\begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbb{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ \rho_0 := \perp \\ \text{For } i \in [k]: \\ \quad (\text{cm}_i, \text{aux}_i) \leftarrow \tilde{\mathcal{P}}(\text{aux}_{i-1}, \rho_{i-1}) \\ \quad (\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}_i, \cdot)}(\text{pp}, (\text{cm}_j)_{j \in [i]}, (\rho_j)_{j < i}, \epsilon) \\ \quad \rho_i \leftarrow \{0, 1\}^{r_i} \\ \quad ((\mathcal{Q}_i, \text{ans}_i, \text{pf}_i))_{i \in [k]} \leftarrow \tilde{\mathcal{P}}(\text{aux}_k, \rho_k) \end{array} \right]$$

where t_{VC} and the total runtime of all executions of \mathcal{R} are $O\left(\frac{k!}{\epsilon} \cdot t_{\text{ARG}}\right)$.

²⁶ \mathcal{V} also implicitly checks for appropriate padding, i.e., that $\text{ans}_i[q] = \sigma$ for every $q \in \mathcal{Q}_i \setminus [l_i]$.

We first construct the reductor \mathcal{R} , which, similarly to Construction 4.8, will include a sampling subroutine \mathcal{S} and a post-processing subroutine $\mathcal{R}_{\text{post}}$.

Construction 5.5. Given a number of iterations $N \in \mathbb{N}$, we construct the sampler \mathcal{S} as follows.

$\mathcal{S}^{\tilde{\mathcal{P}}(\text{aux}_i, \cdot)}(\text{pp}, (\text{cm}_j)_{j \in [i]}, (\rho_j)_{j < i}, N)$:

1. Initialize $\mathcal{K} := \emptyset$.
2. Repeat the following N times:
 - (a) Sample IOP verifier randomness: $(\rho'_1, \dots, \rho'_k) \leftarrow \{0, 1\}^{r_1} \times \dots \times \{0, 1\}^{r_k}$.
 - (b) Obtain $((\mathcal{Q}'_i, \text{ans}'_i, \text{pf}'_i), ((\text{cm}'_j, \mathcal{Q}'_j, \text{ans}'_j, \text{pf}'_j))_{i < j \leq k}) \leftarrow \tilde{\mathcal{P}}(\text{aux}_i, \rho_1, \dots, \rho_{i-1}, \rho'_1, \dots, \rho'_k)$.
 - (c) If $\text{VC.Check}(\text{pp}, \text{cm}_i, \mathcal{Q}'_i, \text{ans}'_i, \text{pf}'_i) = 1$, add $(\mathcal{Q}'_i, \text{ans}'_i, \text{pf}'_i)$ to \mathcal{K} .
3. Output \mathcal{K} .

Construction 5.6. The reductor \mathcal{R} is defined below.

$\mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}_i, \cdot)}(\text{pp}, (\text{cm}_j)_{j \in [i]}, (\rho_j)_{j < i}, \epsilon)$:

1. Set $N := 1/\epsilon$.
2. Run $\mathcal{K} \leftarrow \mathcal{S}^{\tilde{\mathcal{P}}(\text{aux}_i, \cdot)}(\text{pp}, (\text{cm}_j)_{j \in [i]}, (\rho_j)_{j < i}, N)$.
3. Run $(\tilde{\mathcal{Q}}, \tilde{\Pi}) \leftarrow \mathcal{R}_{\text{post}}(\mathcal{K}, l_i)$.²⁷
4. Output $(\tilde{\mathcal{Q}}, \tilde{\Pi})$.

Note that in each of its N iterations, \mathcal{S} runs VC.Check (once) and makes at most k queries to $\tilde{\mathcal{P}}$. Since $t_{\text{ARG}} \geq t_{\text{VC.Check}}$ (and the executions of $\tilde{\mathcal{P}}$ dominate the remaining steps of \mathcal{R}), the total runtime of \mathcal{R} across all k rounds is $O\left(\frac{k \cdot l}{\epsilon} \cdot t_{\text{ARG}}\right)$.

Proof. Throughout this proof, probabilistic expressions are with respect to the following experiment:

$$\left[\begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathfrak{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ \rho_0 := \perp \\ \text{For } i \in [k] : \\ \quad (\text{cm}_i, \text{aux}_i) \leftarrow \tilde{\mathcal{P}}(\text{aux}_{i-1}, \rho_{i-1}) \\ \quad (\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}_i, \cdot)}(\text{pp}, (\text{cm}_j)_{j \in [i]}, (\rho_j)_{j < i}, \epsilon) \\ \quad \rho_i \leftarrow \{0, 1\}^{r_i} \\ \quad ((\mathcal{Q}_i, \text{ans}_i, \text{pf}_i))_{i \in [k]} \leftarrow \tilde{\mathcal{P}}(\text{aux}_k, \rho_k) \end{array} \right],$$

²⁷Recall that $\mathcal{R}_{\text{post}}$ pieces together a proof string $\tilde{\Pi}$ from a set of samples \mathcal{K} (see Construction 4.8).

which by Construction 5.6 is equivalent to

$$\left[\begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbb{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ \rho_0 := \perp \\ \text{For } i \in [k] : \\ \quad (\text{cm}_i, \text{aux}_i) \leftarrow \tilde{\mathcal{P}}(\text{aux}_{i-1}, \rho_{i-1}) \\ \quad \mathcal{K}_i \leftarrow \mathcal{S}^{\tilde{\mathcal{P}}(\text{aux}_i, \cdot)}(\text{pp}, (\text{cm}_j)_{j \in [i]}, (\rho_j)_{j < i}, \mathbf{N}) \\ \quad (\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i) \leftarrow \mathcal{R}_{\text{post}}(\mathcal{K}_i, l_i) \\ \quad \rho_i \leftarrow \{0, 1\}^{r_i} \\ \quad ((\mathcal{Q}_i, \text{ans}_i, \text{pf}_i))_{i \in [k]} \leftarrow \tilde{\mathcal{P}}(\text{aux}_k, \rho_k) \end{array} \right]. \quad (4)$$

The argument is analogous to Lemma 4.6. We first note that

$$\left[\begin{array}{l} \mathbf{V}^{([\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i])_{i \in [k]}}(\mathbb{x}; (\rho_i)_{i \in [k]}) \neq 1 \\ \wedge \mathbf{V}^{([\mathcal{Q}_i, \text{ans}_i])_{i \in [k]}}(\mathbb{x}; (\rho_i)_{i \in [k]}) = 1 \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\text{pp}, \text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) \right) = 1 \end{array} \right]$$

implies, likewise, the existence of valid openings with disagreeing answers or a missing query in one of the IOP strings. We analyze the two cases separately.

Valid openings with disagreeing answers. We aim to show that

$$\Pr \left[\begin{array}{l} \exists i \in [k], q \in \mathcal{Q}_i \cap \tilde{\mathcal{Q}}_i : \text{ans}_i[q] \neq \tilde{\Pi}_i[q] \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\text{pp}, \text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) \right) = 1 \end{array} \right] \leq \epsilon_{\text{VC}}(\lambda, l_{\max}, \mathbf{q}_{\max}, t_{\text{VC}})$$

where $t_{\text{VC}} = O\left(\frac{k!}{\epsilon} \cdot t_{\text{ARG}}\right)$.

We consider the following adversary A_{VC} against the vector commitment scheme, which (essentially) follows Experiment 4:

$A_{\text{VC}}(\text{pp}, \text{ai})$:

1. Run $(\mathbb{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai})$.
2. Set $\mathbf{N} := \frac{1}{\epsilon}$, $\rho_0 := \perp$ and sample $(\rho_i)_{i \in [k]} \leftarrow \{0, 1\}^{r_1 + \dots + r_k}$.
3. For $i \in [k]$:
 - (a) Run $(\text{cm}_i, \text{aux}_i) \leftarrow \tilde{\mathcal{P}}(\text{aux}_{i-1}, \rho_{i-1})$.
 - (b) Run $\mathcal{K}_i \leftarrow \mathcal{S}^{\tilde{\mathcal{P}}(\text{aux}_i, \cdot)}(\text{pp}, (\text{cm}_j)_{j \in [i]}, (\rho_j)_{j < i}, \mathbf{N})$.
4. Run $((\mathcal{Q}_i, \text{ans}_i, \text{pf}_i))_{i \in [k]} \leftarrow \tilde{\mathcal{P}}(\text{aux}_k, \rho_k)$.
5. If there exist $i \in [k]$ and $(\mathcal{Q}'_i, \text{ans}'_i, \text{pf}'_i) \in \mathcal{K}_i$ with $q \in \mathcal{Q}_i \cap \mathcal{Q}'_i$ and $\text{ans}_i[q] \neq \text{ans}'_i[q]$, output $(\text{cm}_i, \text{ans}_i, \text{ans}'_i, \mathcal{Q}_i, \mathcal{Q}'_i, \text{pf}_i, \text{pf}'_i)$.
6. Otherwise, output (the “dummy” tuple) $(\text{cm}_1, \text{ans}_1, \text{ans}_1, \mathcal{Q}_1, \mathcal{Q}_1, \text{pf}_1, \text{pf}_1)$.

The time complexity of k executions of the sampler \mathcal{S} is $O(k\mathbf{N} \cdot t_{\text{ARG}})$ and the collision-finding check (Step 5) runs in time $O(k\mathbf{N} \cdot (\log|\Sigma| + \log l_i)) = O(k\mathbf{N} \cdot t_{\text{ARG}})$, so the total time complexity of A_{VC} is $t_{\text{VC}} = O(k\mathbf{N} \cdot t_{\text{ARG}}) = O\left(\frac{k!}{\epsilon} \cdot t_{\text{ARG}}\right)$.

Therefore, according to Definition 3.8,²⁸

$$\begin{aligned}
& \Pr \left[\begin{array}{l} \exists i \in [k], q \in \mathcal{Q}_i \cap \tilde{\mathcal{Q}}_i : \text{ans}_i[q] \neq \tilde{\Pi}_i[q] \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\text{pp}, \text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) = 1 \right) \end{array} \right] \\
& \leq \Pr \left[\exists i \in [k] : \left(\begin{array}{l} \exists q \in \mathcal{Q}_i \cap \tilde{\mathcal{Q}}_i : \text{ans}_i[q] \neq \tilde{\Pi}_i[q] \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) = 1 \end{array} \right) \right] \\
& \leq \Pr \left[\begin{array}{l} |\mathcal{Q}| = |\mathcal{Q}'| \leq \mathbf{q}_{\max} \\ \wedge \exists q \in \mathcal{Q} \cap \mathcal{Q}' : \text{ans}[q] \neq \text{ans}'[q] \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}', \text{ans}', \text{pf}') = 1 \end{array} \left| \begin{array}{l} \text{pp} \leftarrow \text{VC.Gen}(1^\lambda, l_{\max}) \\ \left(\begin{array}{l} \text{cm}, \text{ans}, \text{ans}' \\ \mathcal{Q}, \mathcal{Q}', \text{pf}, \text{pf}' \end{array} \right) \leftarrow A_{\text{VC}}(\text{pp}) \end{array} \right. \right] \\
& \leq \epsilon_{\text{VC}}(\lambda, l_{\max}, \mathbf{q}_{\max}, t_{\text{VC}}) .
\end{aligned}$$

Missing positions in $\tilde{\Pi}_i$. We are left to show

$$\Pr \left[\begin{array}{l} \left(\exists i \in [k] : \mathcal{Q}_i \setminus \tilde{\mathcal{Q}}_i \neq \emptyset \right) \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\text{pp}, \text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) = 1 \right) \end{array} \right] \leq \epsilon .$$

To this end, fix a partial sequence of prover inputs $(\text{pp}, \text{ai}, (\rho_j)_{j < i})$ of Eq. 4 and define the *weight* of a coordinate $q \in [l_i]$ with respect to $(\text{pp}, \text{ai}, (\rho_j)_{j < i})$ as

$$\begin{aligned}
& \delta_{(\text{pp}, \text{ai}, (\rho_j)_{j < i})}(q) \\
& := \Pr \left[\begin{array}{l} q \in \mathcal{Q}_i \\ \wedge \left(\bigwedge_{j \in [k]} \text{VC.Check} \left(\begin{array}{l} \text{pp}, \text{cm}_j, \\ \mathcal{Q}_j, \text{ans}_j, \text{pf}_j \end{array} \right) = 1 \right) \end{array} \left| \begin{array}{l} (\rho_j)_{i \leq j \leq k} \leftarrow \{0, 1\}^{r_i + \dots + r_k} \\ \left(\begin{array}{l} \mathbf{x}, \text{cm}_1, \dots, \text{cm}_k, \\ ((\mathcal{Q}_j, \text{ans}_j, \text{pf}_j))_{j \in [k]} \end{array} \right) \leftarrow \tilde{\mathcal{P}} \left(\begin{array}{l} \text{pp}, \text{ai}, \\ (\rho_1, \dots, \rho_k) \end{array} \right) \end{array} \right. \right] .
\end{aligned}$$

(We perform a partial execution of $\tilde{\mathcal{P}}$ as in Lemma 4.6, omitting intermediate auxiliary states.) A union bound over all $q \in [l_i]$ yields

$$\begin{aligned}
& \Pr \left[\mathcal{Q}_i \setminus \tilde{\mathcal{Q}}_i \neq \emptyset \right] \\
& = \Pr \left[\exists q \in [l_i] : q \in \mathcal{Q}_i \wedge q \notin \tilde{\mathcal{Q}}_i \right] \\
& \leq \max_{(\text{pp}, \text{ai}, (\rho_j)_{j < i})} \left\{ \Pr \left[\begin{array}{l} \exists q \in [l_i] : q \in \mathcal{Q}_i \\ \wedge q \notin \tilde{\mathcal{Q}}_i \end{array} \left| \begin{array}{l} (\rho_j)_{i \leq j \leq k} \leftarrow \{0, 1\}^{r_i + \dots + r_k} \\ \left(\begin{array}{l} \mathbf{x}, \text{cm}_1, \dots, \text{cm}_k, \\ ((\mathcal{Q}_j, \text{ans}_j, \text{pf}_j))_{j \in [k]} \end{array} \right) \leftarrow \tilde{\mathcal{P}} \left(\begin{array}{l} \text{pp}, \text{ai}, \\ (\rho_1, \dots, \rho_k) \end{array} \right) \end{array} \right. \right] \right\} \\
& \leq \max_{(\text{pp}, \text{ai}, (\rho_j)_{j < i})} \left\{ \sum_{q \in [l_i]} \delta_{(\text{pp}, \text{ai}, (\rho_j)_{j < i})}(q) \cdot \left(1 - \delta_{(\text{pp}, \text{ai}, (\rho_j)_{j < i})}(q) \right)^N \right\} \\
& \leq \frac{l_i}{N} ,
\end{aligned}$$

²⁸Note that A_{VC} , if successful, outputs a tuple in one of the \mathcal{K}_i . Denoting by E_i the event that it is found in \mathcal{K}_i , the right-hand side of the second inequality may be replaced by $\sum_i \Pr[E_i] \cdot \epsilon_{\text{VC}}(\lambda, l_i, \mathbf{q}_i, t_{\text{VC}})$, and the last follows from monotonicity of ϵ_{VC} (see Remark 3.9).

and another union bound over $i \in [k]$ gives

$$\Pr \left[\begin{array}{l} (\exists i \in [k] : \mathcal{Q}_i \setminus \tilde{\mathcal{Q}}_i \neq \emptyset) \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\text{pp}, \text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) = 1 \right) \end{array} \right] \leq \sum_{i \in [k]} \frac{l_i}{N} = \frac{l}{N} .$$

Since $N = l/\epsilon$, we have

$$\begin{aligned} & \Pr \left[\begin{array}{l} \mathbf{V}^{([\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i]_{i \in [k]}(\mathbb{x}; (\rho_i)_{i \in [k]})) \neq 1} \\ \wedge \mathbf{V}^{([\mathcal{Q}_i, \text{ans}_i]_{i \in [k]}(\mathbb{x}; (\rho_i)_{i \in [k]})) = 1} \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\text{pp}, \text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) = 1 \right) \end{array} \right] \\ & \leq \Pr \left[\begin{array}{l} \exists i \in [k], q \in \mathcal{Q}_i \cap \tilde{\mathcal{Q}}_i : \text{ans}_i[q] \neq \tilde{\Pi}_i[q] \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\text{pp}, \text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) = 1 \right) \end{array} \right] \\ & \quad + \Pr \left[\begin{array}{l} \exists i \in [k], q \in \mathcal{Q}_i : q \notin \tilde{\mathcal{Q}}_i \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\text{pp}, \text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) = 1 \right) \end{array} \right] \\ & \leq \epsilon_{\text{VC}}(\lambda, l_{\max}, \mathbf{q}_{\max}, t_{\text{VC}}) + \epsilon , \end{aligned}$$

which concludes the proof. \square

5.3 Adaptive soundness

Lemma 5.7. *For every security parameter $\lambda \in \mathbb{N}$, circuit size bound $t_{\text{ARG}} \geq t_{\text{V}} + t_{\text{VC.Check}} + \log|\Sigma| + \log l_{\max}$, circuit $\tilde{\mathcal{P}}$ of size t_{ARG} , instance size bound $n \in \mathbb{N}$ and $\epsilon > 0$, the soundness error of the argument system in Construction 5.3 satisfies*

$$\epsilon_{\text{ARG}}(\lambda, n, t_{\text{ARG}}) \leq \epsilon_{\text{IOP}}(n) + \epsilon_{\text{VC}}(\lambda, l_{\max}, \mathbf{q}_{\max}, t_{\text{VC}}) + \epsilon ,$$

where $t_{\text{VC}} = O(\frac{k!}{\epsilon} \cdot t_{\text{ARG}})$.

Proof. From Definition 3.3 and Construction 5.3, our goal is to upper bound

$$\begin{aligned} & \Pr \left[\begin{array}{l} |\mathbb{x}| \leq n \\ \wedge \mathbb{x} \notin L(R) \\ \wedge b = 1 \end{array} \left| \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbb{x}, \text{aux}) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ b \leftarrow \langle \tilde{\mathcal{P}}(\text{aux}), \mathcal{V}(\text{pp}, \mathbb{x}) \rangle \end{array} \right. \right] \\ & = \Pr \left[\begin{array}{l} |\mathbb{x}| \leq n \\ \wedge \mathbb{x} \notin L(R) \\ \wedge \mathbf{V}^{([\mathcal{Q}_i, \text{ans}_i]_{i \in [k]}(\mathbb{x}; (\rho_i)_{i \in [k]})) = 1} \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\text{pp}, \text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) = 1 \right) \end{array} \left| \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbb{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ \rho_0 := \perp \\ \text{For } i \in [k] : \\ \quad (\text{cm}_i, \text{aux}_i) \leftarrow \tilde{\mathcal{P}}(\text{aux}_{i-1}, \rho_{i-1}) \\ \quad \rho_i \leftarrow \{0, 1\}^{r_i} \\ ((\mathcal{Q}_i, \text{ans}_i, \text{pf}_i))_{i \in [k]} \leftarrow \tilde{\mathcal{P}}(\text{aux}_k, \rho_k) \end{array} \right. \right] . \end{aligned}$$

As in Lemma 5.4, we consider the following experiment, which augments the above by executing \mathcal{R} (multiple times) and leaves the probability unchanged.

$$\left[\begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbf{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ \rho_0 := \perp \\ \text{For } i \in [k] : \\ \quad (\text{cm}_i, \text{aux}_i) \leftarrow \tilde{\mathcal{P}}(\text{aux}_{i-1}, \rho_{i-1}) \\ \quad (\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}_i, \cdot)}(\text{pp}, (\text{cm}_j)_{j \in [i]}, (\rho_j)_{j < i}, \epsilon) \\ \quad \rho_i \leftarrow \{0, 1\}^{r_i} \\ \quad ((\mathcal{Q}_i, \text{ans}_i, \text{pf}_i))_{i \in [k]} \leftarrow \tilde{\mathcal{P}}(\text{aux}_k, \rho_k) \end{array} \right].$$

By total probability,

$$\begin{aligned} & \Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge (\mathbf{x}, \mathbf{w}) \notin R \\ \wedge \mathbf{V}^{([\mathcal{Q}_i, \text{ans}_i])_{i \in [k]}(\mathbf{x}; \rho)} = 1 \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\text{pp}, \text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) \right) = 1 \end{array} \right] \\ &= \Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge (\mathbf{x}, \mathbf{w}) \notin R \\ \wedge \mathbf{V}^{([\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i])_{i \in [k]}(\mathbf{x}; (\rho_i)_{i \in [k]})} = 1 \\ \wedge \mathbf{V}^{([\mathcal{Q}_i, \text{ans}_i])_{i \in [k]}(\mathbf{x}; (\rho_i)_{i \in [k]})} = 1 \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\text{pp}, \text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) \right) = 1 \end{array} \right] \\ &+ \Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge (\mathbf{x}, \mathbf{w}) \notin R \\ \wedge \mathbf{V}^{([\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i])_{i \in [k]}(\mathbf{x}; (\rho_i)_{i \in [k]})} \neq 1 \\ \wedge \mathbf{V}^{([\mathcal{Q}_i, \text{ans}_i])_{i \in [k]}(\mathbf{x}; (\rho_i)_{i \in [k]})} = 1 \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\text{pp}, \text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) \right) = 1 \end{array} \right]. \end{aligned}$$

We first bound the probability on the left-hand side by using the soundness error definition of IOP in Definition 3.14.

Construction 5.8. We will construct an IOP prover $\tilde{\mathcal{P}}$, and first define its auxiliary input distribution \mathbf{D} as follows:

D:

1. Sample $\text{pp} \leftarrow \mathcal{G}(1^\lambda, n)$ followed by $\text{ai} \leftarrow \mathcal{D}$.
2. Sample $\boldsymbol{\rho} := (\rho_i)_{i \in [k]} \leftarrow (\{0, 1\}^{r_1 + \dots + r_k})^N \times (\{0, 1\}^{r_2 + \dots + r_k})^N \times \dots \times (\{0, 1\}^{r_k})^N$.
3. Output $\text{ai} := (\text{pp}, \text{ai}, \boldsymbol{\rho})$.

The IOP prover is then defined as follows.

• $\tilde{\mathcal{P}}(\text{ai})$:

1. Parse ai as $(\text{pp}, \text{ai}, \boldsymbol{\rho})$.

2. Run $(\mathbf{x}, \mathbf{aux}_0) \leftarrow \tilde{\mathcal{P}}(\mathbf{pp}, \mathbf{ai})$.
3. Set $\mathbf{aux} := (\mathbf{pp}, \mathbf{aux}_0, \boldsymbol{\rho})$.
4. Output $(\mathbf{x}, \mathbf{aux})$.

• $\tilde{\mathcal{P}}(\mathbf{aux}, \rho)$:

1. Parse \mathbf{aux} as $(\mathbf{pp}, \mathbf{aux}_{i-1}, (\mathbf{cm}_j)_{j < i}, \boldsymbol{\rho})$.²⁹
2. Set $\rho_{i-1} := \rho$ and run $(\mathbf{cm}_i, \mathbf{aux}_i) \leftarrow \tilde{\mathcal{P}}(\mathbf{aux}_{i-1}, \rho_{i-1})$
3. Run $(\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}(\mathbf{aux}_i, \cdot)}(\mathbf{pp}, (\mathbf{cm}_j)_{j \in [i]}, (\rho_j)_{j < i}, \epsilon; \boldsymbol{\rho}_i)$.
4. Set $\mathbf{aux}_i := (\mathbf{pp}, \mathbf{aux}_i, (\mathbf{cm}_j)_{j \in [i]}, (\rho_j)_{j < i}, \boldsymbol{\rho})$.³⁰
5. Output $(\tilde{\Pi}_i, \mathbf{aux}_i)$.

Using Definition 3.14 (with Experiment 1, the explicit public-coin IOP), we obtain

$$\begin{aligned}
& \Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge \mathbf{x} \notin L(R) \\ \wedge \mathbf{V}^{([\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i]_{i \in [k]})(\mathbf{x}; (\rho_i)_{i \in [k]})} = 1 \\ \wedge \mathbf{V}^{([\mathcal{Q}_i, \mathbf{ans}_i]_{i \in [k]})(\mathbf{x}; (\rho_i)_{i \in [k]})} = 1 \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\mathbf{pp}, \mathbf{cm}_i, \mathcal{Q}_i, \mathbf{ans}_i, \mathbf{pf}_i) \right) = 1 \end{array} \right] \\
& \leq \Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge \mathbf{x} \notin L(R) \\ \wedge \mathbf{V}^{([\tilde{\Pi}_i]_{i \in [k]})(\mathbf{x}; (\rho_i)_{i \in [k]})} = 1 \end{array} \left| \begin{array}{l} \mathbf{ai} \leftarrow \mathbf{D} \\ (\mathbf{x}, \mathbf{aux}_0) \leftarrow \tilde{\mathcal{P}}(\mathbf{ai}) \\ \rho_0 := \perp \\ \text{For } i \in [k] : \\ \quad (\tilde{\Pi}_i, \mathbf{aux}_i) \leftarrow \tilde{\mathcal{P}}(\mathbf{aux}_{i-1}, \rho_{i-1}) \\ \quad \rho_i \leftarrow \{0, 1\}^{r_i} \end{array} \right. \right] \\
& \leq \epsilon_{\text{IOP}}(n) .
\end{aligned}$$

Lastly, an application of Lemma 5.4 yields

$$\begin{aligned}
& \Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge \mathbf{x} \notin L(R) \\ \wedge \mathbf{V}^{([\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i]_{i \in [k]})(\mathbf{x}; (\rho_i)_{i \in [k]})} \neq 1 \\ \wedge \mathbf{V}^{([\mathcal{Q}_i, \mathbf{ans}_i]_{i \in [k]})(\mathbf{x}; (\rho_i)_{i \in [k]})} = 1 \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\mathbf{pp}, \mathbf{cm}_i, \mathcal{Q}_i, \mathbf{ans}_i, \mathbf{pf}_i) \right) = 1 \end{array} \right] \\
& \leq \Pr \left[\begin{array}{l} \mathbf{V}^{([\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i]_{i \in [k]})(\mathbf{x}; (\rho_i)_{i \in [k]})} = 1 \\ \wedge \mathbf{V}^{([\mathcal{Q}_i, \mathbf{ans}_i]_{i \in [k]})(\mathbf{x}; (\rho_i)_{i \in [k]})} = 1 \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\mathbf{pp}, \mathbf{cm}_i, \mathcal{Q}_i, \mathbf{ans}_i, \mathbf{pf}_i) \right) = 1 \end{array} \right] \\
& \leq \epsilon_{\text{VC}}(\lambda, l_{\max}, \mathbf{q}_{\max}, t_{\text{VC}}) + \epsilon ,
\end{aligned}$$

which concludes the proof. □

²⁹Note that the round $i < k$ is determined by the contents of \mathbf{aux} .

³⁰We may assume $\mathbf{aux} = \perp$ in the case $i = k$, and the following step outputs $(\tilde{\Pi}_k, \mathbf{aux}_k) = \tilde{\Pi}_k$.

5.4 Adaptive knowledge soundness

Lemma 5.9. For every security parameter $\lambda \in \mathbb{N}$, circuit size bound $t_{\text{ARG}} \geq t_{\text{V}} + t_{\text{VC.Check}} + \log|\Sigma| + \log l_{\text{max}}$, circuit $\tilde{\mathcal{P}}$ of size t_{ARG} , instance size bound $n \in \mathbb{N}$ and $\epsilon > 0$, the knowledge soundness error of the argument system obtained by Construction 5.3 satisfies

$$\kappa_{\text{ARG}}(\lambda, n, t_{\text{ARG}}) \leq \kappa_{\text{IOP}}(n) + \epsilon_{\text{IOP}}(n) + \epsilon_{\text{VC}}(\lambda, l_{\text{max}}, \mathfrak{q}_{\text{max}}, t_{\text{VC}}) + \epsilon,$$

where $t_{\text{VC}} = O\left(\frac{k-1}{\epsilon} \cdot t_{\text{ARG}}\right)$. If the IOP extractor's runtime is t_{E} , the argument system's extractor is $t_{\text{E}} = t_{\text{E}} + O(t_{\text{VC}})$.

Construction 5.10. Let \mathbf{E} be the extractor for IOP. We use $\tilde{\mathbf{P}}$ (Construction 5.8) and \mathbf{E} to construct the knowledge extractor \mathcal{E} for ARG:

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

1. Parse tr as $((\text{cm}_i, \rho_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i))_{i \in [k]}$.
2. Sample $\boldsymbol{\rho} := (\rho_i)_{i \in [k]} \leftarrow (\{0, 1\}^{r_1 + \dots + r_k})^{\text{N}} \times (\{0, 1\}^{r_2 + \dots + r_k})^{\text{N}} \times \dots \times (\{0, 1\}^{r_k})^{\text{N}}$.
3. Set $\text{aux} := (\text{pp}, \text{aux}, \boldsymbol{\rho})$.
4. Run $(\tilde{\Pi}_i)_{i \in [k]} \leftarrow \tilde{\mathbf{P}}(\text{aux})$ and set $\text{tr} := ((\tilde{\Pi}_i, \rho_i)_{i < k}, \tilde{\Pi}_k)$.
5. Run $\mathfrak{w} \leftarrow \mathbf{E}^{\tilde{\mathcal{P}}(\text{aux})}(\mathfrak{x}, \text{tr})$.
6. Output \mathfrak{w} .

Proof. From Definition 3.4 and Construction 5.3, our goal is to upper bound

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

$$= \Pr \left[\begin{array}{l} |\mathfrak{x}| \leq n \\ \wedge (\mathfrak{x}, \mathfrak{w}) \notin R \\ \wedge \mathbf{V}^{(\{\mathcal{Q}_i, \text{ans}_i\})_{i \in [k]}}(\mathfrak{x}; (\rho_i)_{i \in [k]}) = 1 \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\text{pp}, \text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) \right) = 1 \end{array} \left| \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathfrak{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ \rho_0 := \perp \\ \text{For } i \in [k] : \\ \quad (\text{cm}_i, \text{aux}_i) \leftarrow \tilde{\mathcal{P}}(\text{aux}_{i-1}, \rho_{i-1}) \\ \quad \rho_i \leftarrow \{0, 1\}^{r_i} \\ \quad ((\mathcal{Q}_i, \text{ans}_i, \text{pf}_i))_{i \in [k]} \leftarrow \tilde{\mathcal{P}}(\text{aux}_k, \rho_k) \\ \text{tr} := ((\text{cm}_i, \rho_i, (\mathcal{Q}_i, \text{ans}_i, \text{pf}_i))_{i \in [k]}) \\ \mathfrak{w} \leftarrow \mathcal{E}^{\tilde{\mathcal{P}}(\text{aux})}(\text{pp}, \mathfrak{x}, \text{tr}) \end{array} \right. \right]$$

We consider the following experiment, which is equivalent to the above and augments that of Lemma 5.4 by appending an execution of \mathbf{E} , for the rest of the proof.

$$\left[\begin{array}{l}
\mathbf{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\
\mathbf{ai} \leftarrow \mathcal{D} \\
\rho_0 := \perp \\
(\rho_i)_{i \in [k]} \leftarrow (\{0, 1\}^{\sum_{j=i}^k r_j})^N \times \dots \times (\{0, 1\}^{r_k})^N \\
\text{For } i \in [k] : \\
\quad (\mathbf{cm}_i, \mathbf{aux}_i) \leftarrow \tilde{\mathcal{P}}(\mathbf{aux}_{i-1}, \rho_{i-1}) \\
\quad (\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}(\mathbf{aux}_i, \cdot)}(\mathbf{pp}, (\mathbf{cm}_\ell)_{\ell \in [i]}, (\rho_\ell)_{\ell < i}, \epsilon; \rho_i) \\
\quad \rho_i \leftarrow \{0, 1\}^{r_i} \\
\quad ((\mathcal{Q}_i, \mathbf{ans}_i, \mathbf{pf}_i))_{i \in [k]} \leftarrow \tilde{\mathcal{P}}(\mathbf{aux}_k, \rho_k) \\
\mathbf{aux} := (\mathbf{pp}, \mathbf{aux}, \rho) \\
\mathbf{tr} := ((\tilde{\Pi}_i, \rho_i)_{i \in [k-1]}, \tilde{\Pi}_k) \\
\mathbf{w} \leftarrow \mathbf{E}^{\tilde{\mathcal{P}}(\mathbf{aux})}(\mathbf{x}, \mathbf{tr})
\end{array} \right]$$

By total probability,

$$\begin{aligned}
& \Pr \left[\begin{array}{l}
|\mathbf{x}| \leq n \\
\wedge (\mathbf{x}, \mathbf{w}) \notin R \\
\wedge \mathbf{V}^{([\mathcal{Q}_i, \mathbf{ans}_i])_{i \in [k]}}(\mathbf{x}; (\rho_i)_{i \in [k]}) = 1 \\
\wedge \mathbf{VC.Check}(\mathbf{pp}, \mathbf{cm}, \mathcal{Q}, \mathbf{ans}, \mathbf{pf}) = 1
\end{array} \right] \\
&= \Pr \left[\begin{array}{l}
|\mathbf{x}| \leq n \\
\wedge (\mathbf{x}, \mathbf{w}) \notin R \\
\wedge \mathbf{V}^{([\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i])_{i \in [k]}}(\mathbf{x}; (\rho_i)_{i \in [k]}) = 1 \\
\wedge \mathbf{V}^{([\mathcal{Q}_i, \mathbf{ans}_i])_{i \in [k]}}(\mathbf{x}; (\rho_i)_{i \in [k]}) = 1 \\
\wedge \left(\bigwedge_{i \in [k]} \mathbf{VC.Check}(\mathbf{pp}, \mathbf{cm}_i, \mathcal{Q}_i, \mathbf{ans}_i, \mathbf{pf}_i) \right) = 1
\end{array} \right] \\
&+ \Pr \left[\begin{array}{l}
|\mathbf{x}| \leq n \\
\wedge (\mathbf{x}, \mathbf{w}) \notin R \\
\wedge \mathbf{V}^{([\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i])_{i \in [k]}}(\mathbf{x}; (\rho_i)_{i \in [k]}) \neq 1 \\
\wedge \mathbf{V}^{([\mathcal{Q}_i, \mathbf{ans}_i])_{i \in [k]}}(\mathbf{x}; (\rho_i)_{i \in [k]}) = 1 \\
\wedge \left(\bigwedge_{i \in [k]} \mathbf{VC.Check}(\mathbf{pp}, \mathbf{cm}_i, \mathcal{Q}_i, \mathbf{ans}_i, \mathbf{pf}_i) \right) = 1
\end{array} \right]
\end{aligned}$$

By the construction of $\tilde{\mathcal{P}}$ (Construction 5.8) and Definition 3.15, we have

$$\begin{aligned}
& \Pr \left[\begin{array}{l}
|\mathbf{x}| \leq n \\
\wedge (\mathbf{x}, \mathbf{w}) \notin R \\
\wedge \mathbf{V}^{([\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i])_{i \in [k]}}(\mathbf{x}; (\rho_i)_{i \in [k]}) = 1 \\
\wedge \mathbf{V}^{([\mathcal{Q}_i, \mathbf{ans}_i])_{i \in [k]}}(\mathbf{x}; (\rho_i)_{i \in [k]}) = 1 \\
\wedge \left(\bigwedge_{i \in [k]} \mathbf{VC.Check}(\mathbf{pp}, \mathbf{cm}_i, \mathcal{Q}_i, \mathbf{ans}_i, \mathbf{pf}_i) \right) = 1
\end{array} \right] \\
&\leq \Pr \left[\begin{array}{l}
|\mathbf{x}| \leq n \\
\wedge (\mathbf{x}, \mathbf{w}) \notin R \\
\wedge b = 1
\end{array} \left| \begin{array}{l}
\mathbf{ai} \leftarrow \mathcal{D} \\
(\mathbf{x}, \mathbf{aux}) \leftarrow \tilde{\mathcal{P}}(\mathbf{ai}) \\
b \leftarrow \mathbf{tr} \langle \tilde{\mathcal{P}}(\mathbf{aux}), \mathbf{V}(\mathbf{x}) \rangle \\
\mathbf{w} \leftarrow \mathbf{E}^{\tilde{\mathcal{P}}(\mathbf{aux})}(\mathbf{x}, \mathbf{tr})
\end{array} \right. \right]
\end{aligned}$$

$$\leq \kappa_{\text{IOP}}(n) .$$

Finally, Lemma 5.4 implies

$$\begin{aligned} & \Pr \left[\begin{array}{l} |\mathbf{x}| \leq n \\ \wedge (\mathbf{x}, \mathbf{w}) \notin R \\ \wedge \mathbf{V}^{([\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i]_{i \in [k]}(\mathbf{x}; (\rho_i)_{i \in [k]}))} \neq 1 \\ \wedge \mathbf{V}^{([\mathcal{Q}_i, \text{ans}_i]_{i \in [k]}(\mathbf{x}; (\rho_i)_{i \in [k]}))} = 1 \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\text{pp}, \text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) \right) = 1 \end{array} \right] \\ & \leq \Pr \left[\begin{array}{l} \mathbf{V}^{([\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i]_{i \in [k]}(\mathbf{x}; (\rho_i)_{i \in [k]}))} \neq 1 \\ \wedge \mathbf{V}^{([\mathcal{Q}_i, \text{ans}_i]_{i \in [k]}(\mathbf{x}; (\rho_i)_{i \in [k]}))} = 1 \\ \wedge \left(\bigwedge_{i \in [k]} \text{VC.Check}(\text{pp}, \text{cm}_i, \mathcal{Q}_i, \text{ans}_i, \text{pf}_i) \right) = 1 \end{array} \right] \\ & \leq \epsilon_{\text{VC}}(\lambda, l_{\max}, \mathbf{q}_{\max}, t_{\text{VC}}) + \epsilon , \end{aligned}$$

which concludes the proof. □

6 Interactive arguments based on public-query IOPs

We argue that the security analysis for interactive arguments based on public-coin IOPs generalizes to interactive arguments based on public-query IOPs with an additional property: the existence of a *random continuation sampler (RCS)*. More precisely, we show that the existence of an RCS for $\text{Finale}[\text{IOP}, \text{VC}]$ is equivalent to its existence for the underlying IOP; and that having an RCS for $\text{Finale}[\text{IOP}, \text{VC}]$ is a sufficient condition for the (natural extension of the) public-coin security reduction to hold.

Theorem 6.1. *Consider these three ingredients:*

- $\text{IOP} = (\mathbf{P}, \mathbf{V})$, a public-query IOP system for a relation R with round complexity k , alphabet Σ , proof length l , and query complexity q ; moreover, IOP has an IOP random continuation sampler \mathbf{S} with error α and running time $t_{\mathbf{S}}$; and
- $\text{VC} = (\text{Gen}, \text{Commit}, \text{Open}, \text{Check})$, a vector commitment scheme over alphabet Σ .

The $\text{ARG} = (\mathcal{G}, \mathcal{P}, \mathcal{V}) := \text{Finale}[\text{IOP}, \text{VC}]$ (Construction 6.2) is a $4k$ -message interactive argument system for R which, for every $\epsilon > 0$ and $t_{\text{ARG}} \geq t_{\mathbf{V}} + t_{\text{VC.Check}} + \log|\Sigma| + \log l_{\max}$, has soundness error ϵ_{ARG} and knowledge soundness error κ_{ARG} satisfying the following:

$$\epsilon_{\text{ARG}}(\lambda, n, t_{\text{ARG}}) \leq \epsilon_{\text{IOP}}(n) + \epsilon_{\text{VC}}(\lambda, l_{\max}, q_{\max}, t_{\text{VC}}) + \epsilon + \frac{l \cdot \alpha}{\epsilon} \quad \text{and}$$

$$\kappa_{\text{ARG}}(\lambda, n, t_{\text{ARG}}) \leq \kappa_{\text{IOP}}(n) + \epsilon_{\text{VC}}(\lambda, l_{\max}, q_{\max}, t_{\text{VC}}) + \epsilon + \frac{l \cdot \alpha}{\epsilon},$$

where $t_{\text{VC}} = O\left(\frac{k \cdot l}{\epsilon} \cdot (t_{\text{ARG}} + t_{\mathbf{S}})\right)$. Moreover, the knowledge extractor runs in time $t_{\mathcal{E}} = t_{\mathbf{E}} + O(t_{\text{VC}})$.

Observe that if IOP admits an IOP random continuation sampler with error $\alpha = 0$, we obtain the same security bounds as in Theorem 5.1; conversely, if $\alpha > 0$, then setting $\epsilon = \sqrt{\alpha \cdot l}$ achieves the minimum additional error of 2ϵ . Note, moreover, that ARG is public-coin if IOP is public-coin.

6.1 Construction

Below, we describe the construction of the interactive argument we denote $(\mathcal{P}, \mathcal{V}) := \text{Finale}[\text{IOP}, \text{VC}]$.

Construction 6.2 (Arguments from public-query IOPs). The argument generator \mathcal{G} receives as input a security parameter $\lambda \in \mathbb{N}$ and an instance size bound $n \in \mathbb{N}$, and works as follows.

$\mathcal{G}(\lambda, n)$:

1. Sample public parameter for the VC scheme: $\text{pp}_{\text{VC}} \leftarrow \text{VC.Gen}(1^\lambda, l_{\max}(n))$.
2. Set public parameter for the interactive argument: $\text{pp} := \text{pp}_{\text{VC}}$.
3. Output pp .

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

1. \mathcal{V} 's randomness: Sample $\rho \leftarrow \{0, 1\}^r$.
2. \mathcal{P} 's commitments.
 - For $i \in [k]$:
 - (a) \mathcal{P} 's i -th commitment.

i. Compute the i -th IOP string and auxiliary state:

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

ii. Compute a VC commitment to the IOP string: $(\mathbf{cm}_i, \mathbf{aux}_i) \leftarrow \text{VC.Commit}(\text{pp}, \Pi_i)$.

iii. Send $(\mathbf{cm}_i, \mathbf{aux}_i)$ to \mathcal{V} .

(b) \mathcal{V} 's i -th query.

i. Compute the i -th IOP verifier query set $\mathbf{Q}_i := \mathbf{V}_q(\mathbf{x}, \rho, (\mathbf{ans}_j)_{j < i})$.

ii. Send \mathbf{Q}_i to \mathcal{P} .

(c) \mathcal{P} 's i -th response.

i. Parse \mathbf{Q}_i as $(\mathcal{Q}_{i,1}, \dots, \mathcal{Q}_{i,i})$, where $\mathcal{Q}_i \subseteq [i]$ is the query set of \mathbf{V} to Π_i .

ii. For every $j \in [i]$, compute $\text{pf}_{i,j} \leftarrow \text{VC.Open}(\text{pp}, \mathbf{aux}_j, \mathcal{Q}_{i,j})$ and set $\mathbf{ans}_{i,j} := \Pi_j[\mathcal{Q}_{i,j}]$.

iii. Set $\mathbf{ans}_i := (\mathbf{ans}_{i,j})_{j \in [i]}$ and $\mathbf{pf}_i := (\text{pf}_{i,j})_{j \in [i]}$.

iv. Send $(\mathbf{ans}_i, \mathbf{pf}_i)$ to \mathcal{V} .

(d) \mathcal{V} 's i -th message.

i. Compute the i -th IOP verifier message $m_i := \mathbf{V}_m(\mathbf{x}, \rho, (\mathbf{ans}_j)_{j \in [i]})$.

ii. Send m_i to \mathcal{P} .

3. \mathcal{V} 's decision: Check if the following hold:

(a) $\mathbf{V}_d(\mathbf{x}, \rho, (\mathbf{ans}_i)_{i \in [k]}) = 1$.

(b) For every $i \in [k]$, $\bigwedge_{j \leq i} \text{VC.Check}(\text{pp}, \mathbf{cm}_j, \mathcal{Q}_{i,j}, \mathbf{ans}_{i,j}, \text{pf}_{i,j})$.

Above, we use \mathbf{V}_q , \mathbf{V}_m and \mathbf{V}_d to denote the query, message and decision algorithms, respectively, for the IOP verifier \mathbf{V} .

The protocol has $2k$ rounds: Each round of the IOP needs two rounds to simulate because \mathbf{V} 's message in the i -th round depends on the answers to the queries in the i -th round. We comment on the protocol's efficiency measures:

- the generator communication to prover and verifier consists of $|\text{pp}_{\text{VC}}|$ bits;
- the prover-to-verifier communication consists of $\sum_{i \in [k]} (|\mathbf{cm}_i| + q \cdot (\log l_i + \log |\Sigma|)) + \sum_{j \leq i} |\text{pf}_{i,j}|$ bits;
- the verifier-to-prover communication consists of $\sum_{i \in [k]} |m_i| + q \cdot \log l_i$ bits;
- the time complexity of the argument generator is $t_{\text{VC.Gen}}$.
- the time complexity of the argument prover is $t_{\mathcal{P}} + k \cdot t_{\text{VC.Commit}} + \frac{k \cdot (k+1)}{2} \cdot t_{\text{VC.Open}}$;
- the time complexity of the argument verifier is $t_{\mathcal{V}} + \frac{k \cdot (k+1)}{2} \cdot t_{\text{VC.Check}}$.

6.2 Random continuation samplers

Definition 6.3 (IOP transcript and partial execution). *An IOP (interaction) transcript tr for a public-query IOP has the following form:*

$$\text{tr} := ((\mathbf{Q}_i, \mathbf{ans}_i, m_i))_{i \in [k]}.$$

For every $i \in [k]$, a **partial IOP transcript tr_i for a public-query IOP** has one of the following forms:

- $\text{tr}_i := (((\mathbf{Q}_j, \mathbf{ans}_j, m_j))_{j < i}, (\mathbf{Q}_i, \mathbf{ans}_i))$;
- $\text{tr}_i := (((\mathbf{Q}_j, \mathbf{ans}_j, m_j))_{j \leq i})$ for $i < k$.

For every IOP prover $\tilde{\mathbf{P}}$, auxiliary input \mathbf{ai} , verifier randomness $\rho \in \{0, 1\}^r$, and integer $i \in [k]$, the i -round partial execution of the IOP with respect to $(\tilde{\mathbf{P}}(\mathbf{ai}), \rho)$, denoted by $\langle \tilde{\mathbf{P}}(\mathbf{ai}), \mathbf{V}(\rho) \rangle_i$, is

$$\langle \tilde{\mathbf{P}}(\mathbf{ai}), \mathbf{V}(\rho) \rangle_i := (\mathbb{x}, \mathbf{tr}_i) ,$$

where \mathbb{x} and the partial IOP transcript $\mathbf{tr}_i = (((\mathbf{Q}_j, \mathbf{ans}_j, m_j))_{j < i}, (\mathbf{Q}_i, \mathbf{ans}_i))$ are obtained via the experiment

$$\left[\begin{array}{l} (\mathbb{x}, \mathbf{aux}_0) \leftarrow \tilde{\mathbf{P}}(\mathbf{ai}) \\ (\mathbf{Q}_0, m_0) := (\perp, \perp) \\ \text{For } 1 \leq j < i : \\ \quad (\tilde{\Pi}_j, \mathbf{aux}_j) \leftarrow \tilde{\mathbf{P}}(\mathbf{aux}_{j-1}, \mathbf{Q}_{j-1}, m_{j-1}) \\ \quad \mathbf{Q}_j \leftarrow \mathbf{V}(\mathbb{x}, \rho, (\mathbf{ans}_\ell)_{\ell < j}) \\ \quad \mathbf{ans}_j := \text{GetAnswers}((\tilde{\Pi}_\ell)_{\ell \leq j}, \mathbf{Q}_j) \\ \quad m_j \leftarrow \mathbf{V}(\mathbb{x}, \rho, (\mathbf{ans}_\ell)_{\ell \leq j}) \\ (\tilde{\Pi}_i, \mathbf{aux}_i) \leftarrow \tilde{\mathbf{P}}(\mathbf{aux}_{i-1}, \mathbf{Q}_{i-1}, m_{i-1}) \\ \mathbf{Q}_i \leftarrow \mathbf{V}(\mathbb{x}, \rho, (\mathbf{ans}_j)_{j \leq i}) \\ \mathbf{ans}_i := \text{GetAnswers}((\tilde{\Pi}_j)_{j \leq i}, \mathbf{Q}_i) \end{array} \right] ,$$

and $\text{GetAnswers}((\Pi_j)_{j \in [i]}, \mathbf{Q}_i)$ denotes $((\Pi_j[\mathbf{Q}_{i,j}]))_{j \in [i]}$ where $((\mathbf{Q}_{i,j}))_{j \in [i]} = \mathbf{Q}_i$.

Definition 6.4 (ARG transcript and partial execution). An ARG (interaction) transcript \mathbf{tr} for a public-query IOP based argument system has the following form:

$$\mathbf{tr} := ((\mathbf{cm}_i, \mathbf{Q}_i, \mathbf{ans}_i, \mathbf{pf}_i, m_i))_{i \in [k]} .$$

Similarly, for every $i \in [k]$, a partial ARG transcript \mathbf{tr}_i has one of the following forms:

- $\mathbf{tr}_i := (((\mathbf{cm}_j, \mathbf{Q}_j, \mathbf{ans}_j, \mathbf{pf}_j, m_j))_{j < i}, (\mathbf{cm}_i, \mathbf{Q}_i, \mathbf{ans}_i, \mathbf{pf}_i))$;
- $\mathbf{tr}_i := (((\mathbf{cm}_j, \mathbf{Q}_j, \mathbf{ans}_j, \mathbf{pf}_j, m_j))_{j \leq i})$ for $i < k$.

For all public parameters \mathbf{pp} , argument prover $\tilde{\mathcal{P}}$, auxiliary input \mathbf{ai} , verifier randomness $\rho \in \{0, 1\}^r$, and integer $i \in [k]$, the i -round partial execution of ARG with respect to $(\mathbf{pp}, \tilde{\mathcal{P}}(\mathbf{ai}), \rho)$, denoted by $\langle \tilde{\mathcal{P}}(\mathbf{ai}), \mathcal{V}(\rho) \rangle_i$, is

$$\langle \tilde{\mathcal{P}}(\mathbf{ai}), \mathcal{V}(\rho) \rangle_i := (\mathbb{x}, \mathbf{tr}_i) ,$$

where \mathbb{x} and the partial ARG transcript $\mathbf{tr}_i = (((\mathbf{cm}_j, \mathbf{Q}_j, \mathbf{ans}_j, \mathbf{pf}_j, m_j))_{j < i}, (\mathbf{cm}_i, \mathbf{Q}_i, \mathbf{ans}_i, \mathbf{pf}_i))$ are obtained via the experiment

$$\left[\begin{array}{l} (\mathbb{x}, \mathbf{aux}_0) \leftarrow \tilde{\mathcal{P}}(\mathbf{pp}, \mathbf{aux}) \\ (\mathbf{Q}_0, m_0) := (\perp, \perp) \\ \text{For } 1 \leq j < i : \\ \quad (\mathbf{cm}_j, \mathbf{aux}_j) \leftarrow \tilde{\mathcal{P}}(\mathbf{aux}_{j-1}, \mathbf{Q}_{j-1}, m_{j-1}) \\ \quad \mathbf{Q}_j \leftarrow \mathcal{V}(\mathbb{x}, \rho, (\mathbf{ans}_\ell)_{\ell < j}) \\ \quad (\mathbf{ans}_j, \mathbf{pf}_j, \mathbf{aux}_j) \leftarrow \tilde{\mathcal{P}}(\mathbf{aux}_j, \mathbf{Q}_j) \\ \quad m_j \leftarrow \mathcal{V}(\mathbb{x}, \rho, (\mathbf{ans}_\ell)_{\ell \in [j]}) \\ (\mathbf{cm}_i, \mathbf{aux}_i) \leftarrow \tilde{\mathcal{P}}(\mathbf{aux}_{i-1}, \mathbf{Q}_{i-1}, m_{i-1}) \\ \mathbf{Q}_i \leftarrow \mathcal{V}(\mathbb{x}, \rho, (\mathbf{ans}_j)_{j < i}) \\ (\mathbf{ans}_i, \mathbf{pf}_i, \mathbf{aux}_i) \leftarrow \tilde{\mathcal{P}}(\mathbf{aux}_i, \mathbf{Q}_i) \end{array} \right] .$$

Definition 6.5. For every IOP prover $\tilde{\mathbf{P}}$ and IOP prover auxiliary input \mathbf{ai} , a partial IOP execution $(\mathbf{x}, \mathbf{tr}_i)$ has non-zero measure with respect to $\tilde{\mathbf{P}}(\mathbf{ai})$ if

$$\Pr \left[(\mathbf{x}, \mathbf{tr}_i) = \langle \tilde{\mathbf{P}}(\mathbf{ai}), \mathbf{V}(\rho) \rangle_i \mid \rho \leftarrow \{0, 1\}^r \right] > 0 .$$

A partial execution for ARG with non-zero measure is defined similarly.

Definition 6.6. Consider $\text{IOP} = (\mathbf{P}, \mathbf{V})$ with randomness complexity r . For any partial IOP execution $(\mathbf{x}, \mathbf{tr}_i)$, we define $\mathbf{R}_{(\mathbf{x}, \mathbf{tr}_i)}$, the set of random strings consistent with $(\mathbf{x}, \mathbf{tr}_i)$, as follows:

$$\mathbf{R}_{(\mathbf{x}, \mathbf{tr}_i)} := \left\{ \rho \in \{0, 1\}^r : \langle \tilde{\mathbf{P}}(\mathbf{ai}), \mathbf{V}(\rho) \rangle_i = \mathbf{tr}_i \right\} .$$

For a partial argument execution \mathbf{tr}_i , we define $\mathbf{R}_{(\mathbf{x}, \mathbf{tr}_i)}$ analogously.

For simplicity, we often drop the instance \mathbf{x} (which is fixed by $\tilde{\mathbf{P}}(\mathbf{aux})$) and refer to a partial transcript \mathbf{tr}_i with non-zero measure, or a set $\mathbf{R}_{\mathbf{tr}_i}$ of random strings consistent with \mathbf{tr}_i .

Definition 6.7 (IOP random continuation sampler). Let IOP be a public-query IOP with round complexity k , and let \mathbf{S} be a probabilistic algorithm with the following syntax:

- On input $(\mathbf{x}, \mathbf{tr}_i)$ where $\mathbf{tr}_i = (((\mathbf{Q}_j, \mathbf{ans}_j, m_j))_{j < i})$ for $i \in [k]$, \mathbf{S} outputs a query set \mathbf{Q}_i ;
- On input $(\mathbf{x}, \mathbf{tr}_i)$ where $\mathbf{tr}_i = (((\mathbf{Q}_j, \mathbf{ans}_j, m_j))_{j < i}, (\mathbf{Q}_i, \mathbf{ans}_i))$, \mathbf{S} outputs a message m_i .

Given $\alpha \geq 0$ and $t_{\mathbf{S}} \in \mathbb{N}$, we say \mathbf{S} is an **IOP random continuation sampler (IOP-RCS) with error α and running time $t_{\mathbf{S}}$** if it satisfies the following guarantee: for every size bound t_{IOP} , t_{IOP} -size circuit $\tilde{\mathbf{P}}$, IOP prover auxiliary input \mathbf{ai} , $i \in [k]$, and partial IOP execution $(\mathbf{x}, \mathbf{tr}_i)$ that has non-zero measure with respect to $(\tilde{\mathbf{P}}, \mathbf{ai})$, we have

$$\Delta \left(\mathcal{D} \left[\tilde{\mathbf{P}}(\mathbf{ai}), \mathbf{S}(\mathbf{x}, \mathbf{tr}_i) \right], \mathcal{U} \left[\tilde{\mathbf{P}}(\mathbf{ai}), \mathbf{R}_{(\mathbf{x}, \mathbf{tr}_i)} \right] \right) = \alpha ,$$

where

- $\Delta(\cdot, \cdot)$ is the statistical (total variation) distance between two distributions;
- $\mathcal{D} \left[\tilde{\mathbf{P}}(\mathbf{ai}), \mathbf{S}(\mathbf{x}, \mathbf{tr}_i) \right]$ is the distribution of IOP transcripts produced by \mathbf{S} :

$$\mathcal{D} \left[\tilde{\mathbf{P}}(\mathbf{ai}), \mathbf{S}(\mathbf{x}, \mathbf{tr}_i) \right] := \left[\left((\mathbf{Q}_j, \mathbf{ans}_j, m_j) \right)_{j \in [k]} \left| \begin{array}{l} (\mathbf{x}, \mathbf{aux}_0) \leftarrow \tilde{\mathbf{P}}(\mathbf{ai}) \\ (\mathbf{Q}_0, m_0) := (\perp, \perp) \\ \text{For } 1 \leq j \leq i : \\ \quad (\tilde{\Pi}_j, \mathbf{aux}_j) \leftarrow \tilde{\mathbf{P}}(\mathbf{aux}_{j-1}, \mathbf{Q}_{j-1}, m_{j-1}) \\ \quad m_i \leftarrow \mathbf{S}(\mathbf{x}, \mathbf{tr}_i) \\ \text{For } i < j \leq k : \\ \quad (\tilde{\Pi}_j, \mathbf{aux}_j) \leftarrow \tilde{\mathbf{P}}(\mathbf{aux}_{j-1}, \mathbf{Q}_{j-1}, m_{j-1}) \\ \quad \mathbf{Q}_j \leftarrow \mathbf{S}(\mathbf{x}, (((\mathbf{Q}_\ell, \mathbf{ans}_\ell, m_\ell))_{\ell < j})) \\ \quad \mathbf{ans}_j := \text{GetAnswers}((\tilde{\Pi}_\ell)_{\ell \leq j}, \mathbf{Q}_j) \\ \quad m_j \leftarrow \mathbf{S}(\mathbf{x}, (((\mathbf{Q}_\ell, \mathbf{ans}_\ell, m_\ell))_{\ell < j}, (\mathbf{Q}_j, \mathbf{ans}_j))) \end{array} \right. \right] ;$$

- $\mathcal{U}[\tilde{\mathcal{P}}(\mathbf{ai}), \mathbf{R}_{(\mathbb{x}, \mathbf{tr}_i)}]$ is the uniform distribution over all complete IOP transcripts consistent with $(\mathbb{x}, \mathbf{tr}_i)$:³¹

$$\mathcal{U}[\tilde{\mathcal{P}}(\mathbf{ai}), \mathbf{R}_{(\mathbb{x}, \mathbf{tr}_i)}] := \left[\begin{array}{l} ((\mathbf{Q}_j, \mathbf{ans}_j, m_j))_{j \in [k]} \end{array} \left| \begin{array}{l} (\mathbb{x}, \mathbf{aux}_0) \leftarrow \tilde{\mathcal{P}}(\mathbf{ai}) \\ \rho \leftarrow \mathbf{R}_{(\mathbb{x}, \mathbf{tr}_i)} \\ (\mathbf{Q}_0, m_0) := (\perp, \perp) \\ \text{For } j \in [k]: \\ (\tilde{\Pi}_j, \mathbf{aux}_j) \leftarrow \tilde{\mathcal{P}}(\mathbf{aux}_{j-1}, \mathbf{Q}_{j-1}, m_{j-1}) \\ \mathbf{Q}_j \leftarrow \mathbf{V}_q(\mathbb{x}, \rho, (\mathbf{ans}_\ell)_{\ell < j}) \\ \mathbf{ans}_j := \text{GetAnswers}((\tilde{\Pi}_\ell)_{\ell \in [j]}, \mathbf{Q}_j) \\ m_j \leftarrow \mathbf{V}_m(\mathbb{x}, \rho, (\mathbf{ans}_\ell)_{\ell \in [j]}) \end{array} \right. \right].$$

Definition 6.8 (ARG random continuation sampler). Let ARG be an argument system with round complexity $2k$ that follows the same message structure as defined in Construction 6.2, and let \mathcal{S} be a probabilistic algorithm with the following syntax.

- On input $(\text{pp}, \mathbb{x}, \mathbf{tr}_i)$ where $\mathbf{tr}_i = (((\text{cm}_j, \mathbf{Q}_j, \mathbf{ans}_j, \mathbf{pf}_j, m_j))_{j < i}, \text{cm}_i)$ for $i \in [k]$, \mathcal{S} outputs \mathbf{Q}_{i+1} ;
- On input $(\text{pp}, \mathbb{x}, \mathbf{tr}_i)$ where $\mathbf{tr}_i := (((\text{cm}_j, \mathbf{Q}_j, \mathbf{ans}_j, \mathbf{pf}_j, m_j))_{j < i}, (\text{cm}_i, \mathbf{Q}_i, \mathbf{ans}_i, \mathbf{pf}_i))$, \mathcal{S} outputs m_i .

Given $\alpha \geq 0$ and $t_S \in \mathbb{N}$, \mathcal{S} is an **ARG random continuation sampler (ARG-RCS) with error α and running time t_S** for ARG if it satisfies the following guarantee: for every size bound t_{ARG} , t_{ARG} -size circuit $\tilde{\mathcal{P}}$, ARG prover auxiliary input \mathbf{ai} , public parameter pp , $i \in [k]$, and partial ARG execution $(\mathbb{x}, \mathbf{tr}_i)$ that has non-zero measure with respect to $\tilde{\mathcal{P}}(\mathbf{ai})$, we have

$$\Delta \left(\mathcal{D}[\tilde{\mathcal{P}}(\mathbf{ai}), \mathcal{S}(\text{pp}, \mathbb{x}, \mathbf{tr}_i)], \mathcal{U}[\tilde{\mathcal{P}}(\mathbf{ai}), \mathbf{R}_{(\mathbb{x}, \mathbf{tr}_i)}] \right) = \alpha,$$

where

- $\mathcal{D}[\tilde{\mathcal{P}}(\mathbf{ai}), \mathcal{S}(\text{pp}, \mathbb{x}, \mathbf{tr}_i)]$ is a distribution of ARG transcripts produced by \mathcal{S} :

$$\mathcal{D}[\tilde{\mathcal{P}}(\mathbf{ai}), \mathcal{S}(\text{pp}, \mathbb{x}, \mathbf{tr}_i)] := \left[\begin{array}{l} ((\text{cm}_j, \mathbf{Q}_j, \mathbf{ans}_j, \mathbf{pf}_j, m_j))_{j \in [k]} \end{array} \left| \begin{array}{l} (\mathbb{x}, \mathbf{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \mathbf{aux}) \\ (\mathbf{Q}_0, m_0) := (\perp, \perp) \\ \text{For } 1 \leq j \leq i: \\ (\text{cm}_j, \mathbf{aux}_j) \leftarrow \tilde{\mathcal{P}}(\mathbf{aux}_{j-1}, \mathbf{Q}_{j-1}, m_{j-1}) \\ m_i \leftarrow \mathcal{S}(\text{pp}, \mathbb{x}, \mathbf{tr}_i) \\ \text{For } i < j \leq k: \\ (\text{cm}_j, \mathbf{aux}_j) \leftarrow \tilde{\mathcal{P}}(\mathbf{aux}_{j-1}, \mathbf{Q}_{j-1}, m_{j-1}) \\ \mathbf{Q}_j \leftarrow \mathcal{S}(\text{pp}, \mathbb{x}, (((\text{cm}_\ell, \mathbf{Q}_\ell, \mathbf{ans}_\ell, \mathbf{pf}_\ell, m_\ell))_{\ell < j})) \\ (\mathbf{ans}_j, \mathbf{pf}_j, \mathbf{aux}_j) \leftarrow \tilde{\mathcal{P}}(\mathbf{aux}_j, \mathbf{Q}_j) \\ m_j \leftarrow \mathcal{S}(\text{pp}, (((\text{cm}_\ell, \mathbf{Q}_\ell, \mathbf{ans}_\ell, \mathbf{pf}_\ell, m_\ell))_{\ell < j}, (\mathbf{Q}_j, \mathbf{ans}_j, \mathbf{pf}_j))) \end{array} \right. ;$$

³¹Note that we run the malicious IOP prover and the IOP verifier from scratch in the experiment. Alternatively, one may break it into two parts as in $\mathcal{D}[\tilde{\mathcal{P}}(\mathbf{ai}), \mathbf{S}(\mathbb{x}, \mathbf{tr}_i)]$: rounds $1 \leq j \leq i$ and $i + 1 \leq j \leq k$. We opt for the more concise notation.

- $\mathcal{U}\left[\tilde{\mathcal{P}}(\text{ai}), \mathbf{R}_{(\mathbb{x}, \text{tr}_i)}\right]$ is the uniform distribution over all complete ARG transcripts consistent with tr_i :

$$\mathcal{U}\left[\tilde{\mathcal{P}}(\text{ai}), \mathbf{R}_{(\mathbb{x}, \text{tr}_i)}\right] := \left[\begin{array}{l} ((\text{cm}_j, \mathbf{Q}_j, \mathbf{ans}_j, \mathbf{pf}_j, m_j))_{j \in [k]} \\ \left. \begin{array}{l} (\mathbb{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{aux}) \\ \rho \leftarrow \mathbf{R}_{(\mathbb{x}, \text{tr}_i)} \\ (\mathbf{Q}_0, m_0) := (\perp, \perp) \\ \text{For } j \in [k] : \\ (\text{cm}_j, \text{aux}_j) \leftarrow \tilde{\mathcal{P}}(\text{aux}_{j-1}, \mathbf{Q}_{j-1}, m_{j-1}) \\ \mathbf{Q}_j \leftarrow \mathcal{V}(\mathbb{x}, \rho, (\mathbf{ans}_\ell)_{\ell < j}) \\ (\mathbf{ans}_j, \mathbf{pf}_j, \text{aux}_j) \leftarrow \tilde{\mathcal{P}}(\text{aux}_j, \mathbf{Q}_j) \\ m_j \leftarrow \mathcal{V}(\mathbb{x}, \rho, (\mathbf{ans}_\ell)_{\ell \in [j]}) \end{array} \right\} \end{array} \right].$$

6.3 Equivalence of IOP-RCS and ARG-RCS

We now show that an IOP-RCS exists if and only if a corresponding ARG-RCS exists.

Lemma 6.9 (IOP sampler to ARG sampler). *For every $\alpha \geq 0$ and $t_S \in \mathbb{N}$, if IOP has an IOP-RCS with error α and running time t_S , then $\text{Finale}[\text{IOP}, \text{VC}]$ admits an ARG-RCS \mathcal{S} with error α and running time $t_S = O(t_S)$.*

Proof. We first define the ARG random continuation sampler \mathcal{S} from its IOP analogue \mathbf{S} the natural way:

$\mathcal{S}(\text{pp}, \mathbb{x}, \text{tr}_i)$:

1. Set $\text{tr}_i := (((\mathbf{Q}_j, \mathbf{ans}_j, m_j))_{j < i})$ if $\text{tr}_i = (((\text{cm}_j, \mathbf{Q}_j, \mathbf{ans}_j, \mathbf{pf}_j, m_j))_{j < i}, \text{cm}_i)$.
2. Set $\text{tr}_i := (((\mathbf{Q}_j, \mathbf{ans}_j, m_j))_{j < i}, (\mathbf{Q}_i, \mathbf{ans}_i))$ if $\text{tr}_i = (((\text{cm}_j, \mathbf{Q}_j, \mathbf{ans}_j, \mathbf{pf}_j, m_j))_{j < i}, (\text{cm}_i, \mathbf{Q}_i, \mathbf{ans}_i, \mathbf{pf}_i))$.
3. Output $\mathbf{S}(\mathbb{x}, \text{tr}_i)$.

It is clear from construction that \mathcal{S} has running time $O(t_S)$. Moreover, since the argument verifier's messages in Construction 6.2 *only depend on the query sets and IOP prover answers* (\mathcal{V} replies with the messages of \mathbf{V} ; an invalid opening interferes with the decision, but not the communication), the error of \mathcal{S} is exactly α , the error of \mathbf{S} . \square

We now show that given an ARG random continuation sampler, we can construct an IOP random continuation sampler.

Lemma 6.10 (ARG sampler to IOP sampler). *For every $\alpha \geq 0$ and $t_S \in \mathbb{N}$, if $\text{Finale}[\text{IOP}, \text{VC}]$ has an ARG-RCS with error α and running time t_S , then IOP admits an IOP-RCS \mathbf{S} with error α and running time $t_S = O(t_S)$.*

Proof. Note that the apparent difficulty is caused by the fact that IOP transcripts do not contain vector commitments nor VC openings, but to run the ARG-RCS, one needs to supply the commitments and openings as inputs.

However, we observe (as we did above) that the argument verifier's messages in Construction 6.2 are independent of the commitments and VC openings. Therefore, these are only syntactic requirements that an arbitrary sequence of strings can fulfil; in particular, it suffices for the ARG-RCS to receive empty strings as commitments and VC openings:

$\mathbf{S}(\mathbb{x}, \text{tr}_i)$:

1. Set $\text{tr}_i := (((\perp, \mathbf{Q}_j, \mathbf{ans}_j, \perp, m_j))_{j < i}, \perp)$ if $\text{tr}_i = (((\mathbf{Q}_j, \mathbf{ans}_j, m_j))_{j < i})$.

2. Set $\text{tr}_i := (((\perp, \mathbf{Q}_j, \mathbf{ans}_j, \perp, m_j))_{j < i}, (\perp, \mathbf{Q}_i, \mathbf{ans}_i, \perp))$ if $\text{tr}_i = (((\mathbf{Q}_j, \mathbf{ans}_j, m_j))_{j < i}, (\mathbf{Q}_i, \mathbf{ans}_i))$.
3. Output $\mathcal{S}(\perp, \mathbf{x}, \text{tr}_i)$.

It is again clear from construction and definition of \mathcal{S} that \mathbf{S} has running time $O(t_s)$ and error α . \square

6.4 Security reduction

We show that given a random continuation sampler for IOP (or, equivalently, a random continuation sampler for ARG), a security reduction claim similar to Lemma 5.4 can be shown.

Lemma 6.11. *Assume there is an ARG-RCS \mathcal{S} with error α and running time t_s . There exist a probabilistic algorithms \mathcal{R} which, for every $\epsilon > 0$, size bound $t_{\text{ARG}} \geq t_{\mathbf{V}} + t_{\text{VC.Check}} + \log|\Sigma| + \log l_{\text{max}}$, and t_{ARG} -size circuit $\tilde{\mathcal{P}}$, satisfy*

$$\Pr \left[\begin{array}{l} \mathbf{V}^{([\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i])_{i \in [k]}(\mathbf{x}; \rho)} \neq 1 \\ \wedge \mathbf{V}^{([\mathcal{Q}_i, \mathbf{ans}_i])_{i \in [k]}(\mathbf{x}; \rho)} = 1 \\ \wedge \left(\bigwedge_{i \in [k]} \left(\bigwedge_{j \leq i} \text{VC.Check}(\text{pp}, \text{cm}_j, \mathcal{Q}_{i,j}, \mathbf{ans}_{i,j}, \text{pf}_{i,j}) \right) \right) \end{array} \right] \leq \epsilon_{\text{VC}}(\lambda, l_{\text{max}}, q_{\text{max}}, t_{\text{VC}}) + \epsilon + \frac{l \cdot \alpha}{\epsilon},$$

with respect to

$$\left[\begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbf{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ \rho \leftarrow \{0, 1\}^r \\ ((\text{cm}_i, \mathbf{Q}_i, \mathbf{ans}_i, \text{pf}_i, m_i))_{i \in [k]} \leftarrow \langle \tilde{\mathcal{P}}(\text{aux}_0), \mathcal{V}(\mathbf{x}; \rho) \rangle \\ \text{For } i \in [k]: \\ \quad (\text{cm}_i, \text{aux}_i) \leftarrow \tilde{\mathcal{P}}(\mathbf{x}, \text{aux}_{i-1}, \mathbf{Q}_{i-1}, m_{i-1}) \\ \quad (\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}_i, \cdot)}(\text{pp}, \mathbf{x}, (\text{cm}_j)_{j \leq i}, ((\mathbf{Q}_j, \mathbf{ans}_j, \text{pf}_j, m_j))_{j < i}, \epsilon) \end{array} \right] \quad (5)$$

where $t_{\text{VC}} = O\left(\frac{k \cdot l}{\epsilon} \cdot (t_{\text{ARG}} + t_s)\right)$.

The reductor is defined as follows:

- $\mathcal{R}_{\text{post}}(\mathcal{K}, l)$:
 1. Initialize $\tilde{\Pi} := (\sigma)^l$, where σ is an arbitrary symbol in Σ .
 2. For all $(\mathcal{Q}, \mathbf{ans}, \text{pf}) \in \mathcal{K}$ and $q \in \mathcal{Q}$: Set $\tilde{\Pi}[q] := \mathbf{ans}[q]$.
 3. Set $\tilde{\mathcal{Q}} := \bigcup_{(\mathcal{Q}, \mathbf{ans}) \in \mathcal{K}} \mathcal{Q}$.
 4. Output $(\tilde{\mathcal{Q}}, \tilde{\Pi})$.
- $\mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}_i, \cdot)}(\text{pp}, \mathbf{x}, (\text{cm}_j)_{j \leq i}, ((\mathbf{Q}_j, \mathbf{ans}_j, \text{pf}_j, m_j))_{j < i}, \epsilon)$:³²
 1. Initialize $\mathcal{K} = \emptyset$.
 2. Set $N := \frac{1}{\epsilon}$ and $\text{cm}'_i := \text{cm}_i$.

³²Note that, when $i = 1$ (i.e., when the second argument is \perp) the reductor simply samples ρ and applies the functions \mathbf{V}_q and \mathbf{V}_m as appropriate; there is no need for a sampler in this case. Equivalently, we assume \mathcal{S} performs this trivial sampling when $i = 1$ and is only nontrivial when $i > 1$.

3. Repeat the following N times:
 - (a) Run $\mathbf{Q}'_i \leftarrow \mathcal{S}(\text{pp}, \mathbb{x}, (((\text{cm}_\ell, \mathbf{Q}_\ell, \mathbf{ans}_\ell, \mathbf{pf}_\ell, m_\ell))_{\ell < i}, \text{cm}_i))$.
 - (b) Run $(\mathbf{ans}'_i, \mathbf{pf}'_i, \text{aux}'_i) \leftarrow \tilde{\mathcal{P}}(\text{aux}_i, \mathbf{Q}_i)$.
 - (c) Run $m'_i \leftarrow \mathcal{S}(\text{pp}, \mathbb{x}, (((\text{cm}_\ell, \mathbf{Q}_\ell, \mathbf{ans}_\ell, \mathbf{pf}_\ell, m_\ell))_{\ell < i}, (\mathbf{Q}'_i, \mathbf{ans}'_i, \mathbf{pf}'_i)))$.
 - (d) For $i < j \leq k$:
 - i. Run $(\text{cm}'_j, \text{aux}'_j) \leftarrow \tilde{\mathcal{P}}(\text{aux}'_{j-1}, \mathbf{Q}'_{j-1}, m'_{j-1})$.
 - ii. Run $\mathbf{Q}'_j \leftarrow \mathcal{S}(\text{pp}, \mathbb{x}, (((\text{cm}_\ell, \mathbf{Q}_\ell, \mathbf{ans}_\ell, \mathbf{pf}_\ell, m_\ell))_{\ell < i}, ((\text{cm}'_\ell, \mathbf{Q}'_\ell, \mathbf{ans}'_\ell, \mathbf{pf}'_\ell, m'_\ell))_{i \leq \ell < j}))$.
 - iii. Run $(\mathbf{ans}'_j, \mathbf{pf}'_j, \text{aux}'_j) \leftarrow \tilde{\mathcal{P}}(\text{aux}'_j, \mathbf{Q}'_j)$.
 - iv. Run $m'_j \leftarrow \mathcal{S}(\text{pp}, \mathbb{x}, (((\text{cm}_\ell, \mathbf{Q}_\ell, \mathbf{ans}_\ell, \mathbf{pf}_\ell, m_\ell))_{\ell < i}, ((\text{cm}'_\ell, \mathbf{Q}'_\ell, \mathbf{ans}'_\ell, \mathbf{pf}'_\ell, m'_\ell))_{i \leq \ell < j}, (\mathbf{Q}'_j, \mathbf{ans}'_j, \mathbf{pf}'_j)))$.
 - v. If $\text{VC.Check}(\text{cm}_i, \mathcal{Q}_{j,i}, \mathbf{ans}'_{j,i}, \mathbf{pf}'_{j,i}) = 1$, add $(\mathcal{Q}'_{j,i}, \mathbf{ans}'_{j,i}, \mathbf{pf}'_{j,i})$ to \mathcal{K} .³³
4. Output $(\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i) \leftarrow \mathcal{R}_{\text{post}}(\mathcal{K}, l_i)$.

Proof. Throughout the proof, probabilistic expressions are with respect to the experiment in Eq. 5.

Our goal is to upper bound the following expression:

$$\Pr \left[\begin{array}{l} \mathbf{V}([\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i]_{i \in [k]}(\mathbb{x}; \rho) \neq 1 \\ \wedge \mathbf{V}([\mathcal{Q}_{i, \mathbf{ans}_i}]_{i \in [k]}(\mathbb{x}; \rho) = 1 \\ \wedge \left(\bigwedge_{i \in [k]} \left(\bigwedge_{j \leq i} \text{VC.Check}(\text{pp}, \text{cm}_j, \mathcal{Q}_{i,j}, \mathbf{ans}_{i,j}, \mathbf{pf}_{i,j}) \right) \right) \end{array} \right].$$

Similar to the proof of Lemma 5.9, the above even implies either there are valid openings with disagreeing answers or there is a missing query in one of the IOP strings. We analyze them separately.

Valid openings with disagreeing answers. We show that

$$\Pr \left[\begin{array}{l} \exists i \in [k], j \in [i, k], q \in \mathcal{Q}_{j,i} \cap \tilde{\mathcal{Q}}_i \\ \wedge \tilde{\Pi}_i[q] \neq \mathbf{ans}_{j,i}[q] \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}_i, \mathcal{Q}_{j,i}, \mathbf{ans}_{j,i}, \mathbf{pf}_{j,i}) = 1 \end{array} \right] \leq \epsilon_{\text{VC}}(\lambda, l_{\max}, \mathbf{q}_{\max}, t_{\text{VC}})$$

where $t_{\text{VC}} = O\left(\frac{k \cdot l}{\epsilon} \cdot (t_{\text{S}} + t_{\text{ARG}})\right)$.

We define a VC adversary A_{VC} as before:

$A_{\text{VC}}(\text{pp}, \text{ai}):$

1. Run $(\mathbb{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai})$.
2. Sample $\rho \leftarrow \{0, 1\}^r$.
3. Run $((\text{cm}_i, \mathbf{Q}_i, \mathbf{ans}_i, \mathbf{pf}_i, m_i))_{i \in [k]} \leftarrow \langle \tilde{\mathcal{P}}(\text{aux}_0), \mathcal{V}(\text{pp}, \mathbb{x}, \rho) \rangle$.
4. For every $i \in [k]$:
 - (a) Run $(\text{cm}_i, \text{aux}_i) \leftarrow \tilde{\mathcal{P}}(\text{aux}_{i-1}, \mathbf{Q}_{i-1}, \mathbf{ans}_{i-1})$.
 - (b) Run $\mathcal{K}_i \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}_i, \cdot)}(\text{pp}, \mathbb{x}, (\text{cm}_j)_{j \leq i}, ((\mathbf{Q}_j, \mathbf{ans}_j, m_j))_{j < i}, \epsilon)$.
5. If there exists $i \in [k], j \in [i, k]$, and $(\mathcal{Q}, \mathbf{ans}, \mathbf{pf}) \in \mathcal{K}_i$ with $q \in \mathcal{Q}_{j,i} \cap \mathcal{Q}$ and $\mathbf{ans}_{j,i}[q] \neq \mathbf{ans}[q]$, output $(\text{cm}_i, \mathbf{ans}_{j,i}, \mathbf{ans}, \mathcal{Q}_{j,i}, \mathcal{Q}, \mathbf{pf}_{j,i}, \mathbf{pf})$.
6. Otherwise, output (the “dummy” tuple) $(\text{cm}_1, \mathbf{ans}_1, \mathbf{ans}_1, \mathcal{Q}_1, \mathcal{Q}_1, \mathbf{pf}_1, \mathbf{pf}_1)$.

Note that in A_{VC} , we output \mathcal{K} when invoking \mathcal{R} . This is a notational overload where we partially execute \mathcal{R} without calling $\mathcal{R}_{\text{post}}$. The time complexity of A_{VC} is $O\left(\frac{k \cdot l}{\epsilon} \cdot (t_{\text{S}} + t_{\text{ARG}})\right)$. Therefore, according to

³³Note the order of indices: in order to fill in the i -th IOP proof, the reducer uses the i -th query sets and answers of all rounds $j > i$ (i.e., all $\mathcal{Q}_{j,i}$ and $\mathbf{ans}_{j,i}$).

Definition 3.8,

$$\begin{aligned}
& \Pr \left[\begin{array}{l} \exists i \in [k], j \in [i, k], q \in \mathcal{Q}_{j,i} \cap \tilde{\mathcal{Q}}_i \\ \wedge \tilde{\Pi}_i[q] \neq \text{ans}_{j,i}[q] \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}_i, \mathcal{Q}_{j,i}, \text{ans}_{j,i}, \text{pf}_{j,i}) = 1 \end{array} \right] \\
& \leq \Pr \left[\begin{array}{l} |\mathcal{Q}| = |\mathcal{Q}'| \leq q \\ \wedge \exists q \in \mathcal{Q} \cap \mathcal{Q}' : \text{ans}[q] \neq \text{ans}'[q] \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}, \text{ans}, \text{pf}) = 1 \\ \wedge \text{VC.Check}(\text{pp}, \text{cm}, \mathcal{Q}', \text{ans}', \text{pf}') = 1 \end{array} \mid \begin{array}{l} \text{pp} \leftarrow \text{VC.Gen}(1^\lambda, l_{\max}) \\ \text{ai} \leftarrow \mathcal{D} \\ (\text{cm}, \text{ans}, \text{ans}', \\ \mathcal{Q}, \mathcal{Q}', \text{pf}, \text{pf}') \leftarrow A_{\text{VC}}(\text{pp}) \end{array} \right] \\
& \leq \epsilon_{\text{VC}}(\lambda, l_{\max}, q_{\max}, t_{\text{VC}}) .
\end{aligned}$$

Missing positions in $\tilde{\Pi}_i$. We show that

$$\Pr \left[\exists i \in [k], j \in [i, k], q \in \mathcal{Q}_{j,i} \setminus \tilde{\mathcal{Q}}_i \right] \leq \epsilon + 2l \cdot \frac{\alpha}{\epsilon} .$$

For this case, we rely on the property of ARG-RCS. In particular, according to Definition 6.8, the ARG verifier's next message in tr_i sampled by \mathcal{S} is statistically close to the "ideal" distribution of ARG verifier's next message, namely the ARG execution with respect to verifier randomness in $\mathbf{R}_{(\mathbb{x}, \text{tr}_i)}$.

We first define the following *ideal reductor* $\hat{\mathcal{R}}$ with respect to $\mathcal{U}[\tilde{\mathcal{P}}(\text{ai}), \mathbf{R}_{(\mathbb{x}, \text{tr}_i)}]$. In particular, $\hat{\mathcal{R}}$ works as if there were a perfect ARG-RCS as in the public-coin case:

$$\hat{\mathcal{R}}^{\tilde{\mathcal{P}}(\text{aux}_i, \cdot)}(\text{pp}, \mathbb{x}, (\text{cm}_j)_{j \leq i}, ((\mathbf{Q}_j, \text{ans}_j, \text{pf}_j, m_j))_{j < i}, \epsilon):$$

1. Initialize $\mathcal{K} := \emptyset$ and set $N := \frac{1}{\epsilon}$.
2. Repeat the following N times:
 - (a) Sample $\rho \leftarrow \mathbf{R}_{((\text{cm}_j, \mathbf{Q}_j, \text{ans}_j, \text{pf}_j, m_j))_{j < i}}$.
 - (b) Run $\mathbf{Q}_i \leftarrow \mathcal{V}(\mathbb{x}, \rho, (\text{ans}_j)_{j < i})$.
 - (c) Run $(\text{ans}_i, \text{pf}_i, \text{aux}_i) \leftarrow \mathcal{P}(\text{aux}_i, \mathbf{Q}_i)$.
 - (d) Run $m_i \leftarrow \mathcal{V}(\mathbb{x}, \rho, (\text{ans}_j)_{j \in [i]})$.
 - (e) For $i < j \leq k$:
 - i. Run $(\text{cm}, \text{aux}'_j) \leftarrow \tilde{\mathcal{P}}(\text{aux}_{j-1}, m_{j-1})$.
 - ii. Run $\mathbf{Q}_j \leftarrow \mathcal{V}(\mathbb{x}, \rho, (\text{ans}_\ell)_{\ell < j})$.
 - iii. Run $(\text{ans}_j, \text{pf}_j, \text{aux}_j) \leftarrow \tilde{\mathcal{P}}(\text{aux}'_j, \mathbf{Q}_j)$.
 - iv. Run $m_j \leftarrow \mathcal{V}(\mathbb{x}, \rho, (\text{ans}_\ell)_{\ell \in [j]})$.
 - (f) If $\text{VC.Check}(\text{cm}_i, \mathcal{Q}_{j,i}, \text{ans}_{j,i}, \text{pf}_{j,i}) = 1$, add $(\mathcal{Q}_{j,i}, \text{ans}_{j,i}, \text{pf}_{j,i})$ to \mathcal{K} .
3. Output $(\hat{\mathcal{Q}}_i, \hat{\Pi}_i) \leftarrow \mathcal{R}_{\text{post}}(\mathcal{K}, l_i)$.

We can adapt the proof of Lemma 5.4 to bound the following probability:

$$\Pr \left[\begin{array}{l} \exists i \in [k], j \in [i, k], q \in \mathcal{Q}_{j,i} \setminus \hat{\mathcal{Q}}_i \end{array} \mid \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbb{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ \rho \leftarrow \{0, 1\}^r \\ ((\text{cm}_j, \mathbf{Q}_j, \text{ans}_j, \text{pf}_j, m_j))_{j \in [k]} \leftarrow \langle \tilde{\mathcal{P}}(\text{aux}_0), \mathcal{V}(\text{pp}, \mathbb{x}, \rho) \rangle \\ \text{For } i \in [k] : \\ (\text{cm}_i, \text{aux}_i) \leftarrow \tilde{\mathcal{P}}(\text{aux}_{i-1}, \mathbf{Q}_{i-1}, m_{i-1}) \\ (\hat{\mathcal{Q}}_i, \hat{\Pi}_i) \leftarrow \hat{\mathcal{R}}^{\tilde{\mathcal{P}}(\text{aux}_i, \cdot)}(\text{pp}, \mathbb{x}, (\text{cm}_j)_{j \leq i}, ((\mathbf{Q}_j, \text{ans}_j, \text{pf}_j, m_j))_{j < i}, \epsilon) \end{array} \right] . \tag{6}$$

Fix prover auxiliary input ai , public parameters pp and a partial transcript $\text{tr}_i = ((\text{cm}_j, \mathbf{Q}_j, \mathbf{ans}_j, \mathbf{pf}_j, m_j))_{j < i}$ that has non-zero measure with respect to $\tilde{\mathcal{P}}(\text{ai})$, we define the *weight* of a coordinate $q \in [l_i]$ with respect to $(\text{pp}, \text{ai}, \text{tr}_i)$ as

$$\delta_{(\text{pp}, \text{ai}, \text{tr}_i)}(q) := \Pr \left[\begin{array}{l} \exists j \in [i, k] : q \in \mathcal{Q}_{j,i} \\ \wedge \bigwedge_{j \in [k]} \left(\bigwedge_{\ell \in [j]} \text{VC.Check}(\text{pp}, \text{cm}_\ell, \mathcal{Q}_{j,\ell}, \mathbf{ans}_{j,\ell}, \mathbf{pf}_{j,\ell}) \right) \end{array} \middle| \begin{array}{l} \rho \leftarrow \mathbf{R}_{(\mathbf{x}, \text{tr}_i)} \\ b \stackrel{\text{tr}}{\leftarrow} \langle \tilde{\mathcal{P}}(\text{pp}, \text{ai}), \mathcal{V}(\text{pp}, \mathbf{x}, \rho) \rangle \\ ((\text{cm}_j, \mathbf{ans}_j, \mathbf{pf}_j, m_j))_{j \in [k]} := \text{tr} \end{array} \right] .$$

A union bound over all $q \in [l_i]$ yields

$$\begin{aligned} & \Pr \left[\begin{array}{l} \exists j \in [i, k], q \in \mathcal{Q}_{j,i} : q \notin \hat{\mathcal{Q}}_i \end{array} \middle| \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbf{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ \rho \leftarrow \{0, 1\}^r \\ ((\text{cm}_j, \mathbf{Q}_j, \mathbf{ans}_j, \mathbf{pf}_j, m_j))_{j \in [k]} \leftarrow \langle \tilde{\mathcal{P}}(\text{aux}_0), \mathcal{V}(\text{pp}, \mathbf{x}, \rho) \rangle \\ \text{For } j \in [k] : \\ \quad (\text{cm}_j, \text{aux}_j) \leftarrow \tilde{\mathcal{P}}(\text{aux}_{j-1}, \mathbf{Q}_{j-1}, m_{j-1}) \\ \quad (\hat{\mathcal{Q}}_j, \hat{\Pi}_j) \leftarrow \hat{\mathcal{R}}^{\tilde{\mathcal{P}}(\text{aux}_j, \cdot)}(\text{pp}, \mathbf{x}, (\text{cm}_\ell)_{\ell \leq j}, ((\mathbf{Q}_\ell, \mathbf{ans}_\ell, \mathbf{pf}_\ell, m_\ell))_{\ell < j}, \epsilon) \end{array} \right] \\ & \leq \max_{(\text{pp}, \text{ai}, \text{tr}_i)} \left\{ \Pr \left[\begin{array}{l} \exists q \in [l_i], j \in [i, k] : q \in \mathcal{Q}_{j,i} \wedge q \notin \tilde{\mathcal{Q}}_i \end{array} \middle| \begin{array}{l} \rho \leftarrow \mathbf{R}_{(\mathbf{x}, \text{tr}_i)} \\ b \stackrel{\text{tr}}{\leftarrow} \langle \tilde{\mathcal{P}}(\text{pp}, \text{ai}), \mathcal{V}(\text{pp}, \rho) \rangle \\ ((\text{cm}_j, \mathbf{ans}_j, \mathbf{pf}_j, m_j))_{j \in [k]} := \text{tr} \end{array} \right] \right\} \\ & \leq \max_{(\text{pp}, \text{ai}, \text{tr}_i)} \left\{ \sum_{q \in [l_i]} \delta_{(\text{pp}, \text{ai}, \text{tr}_i)}(q) \cdot (1 - \delta_{(\text{pp}, \text{ai}, \text{tr}_i)}(q))^N \right\} \\ & \leq \frac{l_i}{N} , \end{aligned}$$

and another union bound over $i \in [k]$ gives

$$\begin{aligned} & \Pr \left[\begin{array}{l} \exists i \in [k], j \in [i, k], q \in \mathcal{Q}_{j,i} : q \notin \hat{\mathcal{Q}}_i \end{array} \middle| \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbf{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ \rho \leftarrow \{0, 1\}^r \\ ((\text{cm}_i, \mathbf{Q}_i, \mathbf{ans}_i, \mathbf{pf}_i, m_i))_{i \in [k]} \leftarrow \langle \tilde{\mathcal{P}}(\text{aux}_0), \mathcal{V}(\text{pp}, \mathbf{x}, \rho) \rangle \\ \text{For } i \in [k] : \\ \quad (\text{cm}_i, \text{aux}_i) \leftarrow \tilde{\mathcal{P}}(\mathbf{x}, \text{aux}_{i-1}, \mathbf{Q}_{i-1}, m_{i-1}) \\ \quad (\hat{\mathcal{Q}}_i, \hat{\Pi}_i) \leftarrow \hat{\mathcal{R}}^{\tilde{\mathcal{P}}(\text{aux}_i, \cdot)}(\text{pp}, \mathbf{x}, (\text{cm}_j)_{j \leq i}, ((\mathbf{Q}_j, \mathbf{ans}_j, \mathbf{pf}_j, m_j))_{j < i}, \epsilon) \end{array} \right] \\ & \leq \sum_{i \in [k]} \frac{l_i}{N} = \frac{l}{N} . \end{aligned}$$

It remains to show the connection between Eq. 6 and $\Pr \left[\exists i \in [k], j \in [i, k], q \in \mathcal{Q}_{j,i} : q \notin \tilde{\mathcal{Q}}_i \right]$. According to Definition 6.8,

$$\Delta \left(\mathcal{D} \left[\tilde{\mathcal{P}}(\text{ai}), \mathcal{S}(\text{pp}, \mathbf{x}, \text{tr}_i) \right], \mathcal{U} \left[\tilde{\mathcal{P}}(\text{ai}), \mathbf{R}_{(\mathbf{x}, \text{tr}_i)} \right] \right) = \alpha .$$

Therefore, for any $N \in \mathbb{N}$,

$$\Delta \left(\left(\mathcal{D} \left[\tilde{\mathcal{P}}(\text{ai}), \mathcal{S}(\text{pp}, \mathbf{x}, \text{tr}_i) \right] \right)^N, \left(\mathcal{U} \left[\tilde{\mathcal{P}}(\text{ai}), \mathbf{R}_{(\mathbf{x}, \text{tr}_i)} \right] \right)^N \right) = N \cdot \alpha .$$

Consider the following distribution:

$$\left[\begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbf{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ \rho \leftarrow \{0, 1\}^r \\ ((\mathbf{Q}_i, \hat{\mathbf{Q}}_i, \hat{\Pi}_i))_{i \in [k]} \\ ((\text{cm}_i, \mathbf{Q}_i, \text{ans}_i, \text{pf}_i, m_i))_{i \in [k]} \leftarrow \langle \tilde{\mathcal{P}}(\text{aux}_0), \mathcal{V}(\text{pp}, \mathbf{x}, \rho) \rangle \\ \text{For } i \in [k] : \\ (\text{cm}_i, \text{aux}_i) \leftarrow \tilde{\mathcal{P}}(\mathbf{x}, \text{aux}_{i-1}, \mathbf{Q}_{i-1}, m_{i-1}) \\ (\hat{\mathbf{Q}}_i, \hat{\Pi}_i) \leftarrow \hat{\mathcal{R}}^{\tilde{\mathcal{P}}(\text{aux}_i, \cdot)}(\text{pp}, \mathbf{x}, (\text{cm}_j)_{j \leq i}, ((\mathbf{Q}_j, \text{ans}_j, \text{pf}_j, m_j))_{j < i}, \epsilon) \end{array} \right]$$

We can rewrite it equivalently in terms of $\left(\mathcal{U} \left[\tilde{\mathcal{P}}(\text{ai}), \mathbf{R}_{(\mathbf{x}, \cdot)} \right] \right)^N$:

$$\begin{aligned} & \left[\begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbf{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ \rho \leftarrow \{0, 1\}^r \\ ((\mathbf{Q}_i, \hat{\mathbf{Q}}_i, \hat{\Pi}_i))_{i \in [k]} \\ ((\text{cm}_i, \mathbf{Q}_i, \text{ans}_i, \text{pf}_i, m_i))_{i \in [k]} \leftarrow \langle \tilde{\mathcal{P}}(\text{aux}_0), \mathcal{V}(\text{pp}, \mathbf{x}, \rho) \rangle \\ \text{For } i \in [k] : \\ \text{For } j \in [N] : \\ ((\text{cm}_{j,\ell}, \mathbf{Q}_{j,\ell}, \text{ans}_{j,\ell}, \text{pf}_{j,\ell}, m_{j,\ell}))_{\ell \in [k]} \leftarrow \mathcal{U} \left[\tilde{\mathcal{P}}(\text{ai}), \mathbf{R}_{((\text{cm}_\ell, \mathbf{Q}_\ell, \text{ans}_\ell, \text{pf}_\ell, m_\ell))_{\ell < i}, \text{cm}_i} \right] \\ \mathcal{K}_i := \left\{ \begin{array}{l} (\mathcal{Q}_{j,\ell,i}, \text{ans}_{j,\ell,i}, \text{pf}_{j,\ell,i}) : j \in [N], i \leq \ell \leq k, \\ \text{VC.Check}(\text{pp}, \text{cm}_{j,i}, \mathcal{Q}_{j,\ell,i}, \text{ans}_{j,\ell,i}, \text{pf}_{j,\ell,i}) = 1 \end{array} \right\} \\ (\hat{\mathbf{Q}}_i, \hat{\Pi}_i) \leftarrow \mathcal{R}_{\text{post}}(\mathcal{K}_i, \mathbf{l}_i) \end{array} \right] \\ = & \left[\begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbf{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ \rho \leftarrow \{0, 1\}^r \\ ((\text{cm}_i, \mathbf{Q}_i, \text{ans}_i, \text{pf}_i, m_i))_{i \in [k]} \leftarrow \langle \tilde{\mathcal{P}}(\text{aux}_0), \mathcal{V}(\text{pp}, \mathbf{x}, \rho) \rangle \\ \text{For } i \in [k] : \\ ((\text{cm}_{j,\ell}, \mathbf{Q}_{j,\ell}, \text{ans}_{j,\ell}, \text{pf}_{j,\ell}, m_{j,\ell}))_{j \in [N], \ell \in [k]} \\ \leftarrow \left(\mathcal{U} \left[\tilde{\mathcal{P}}(\text{ai}), \mathbf{R}_{((\text{cm}_j, \mathbf{Q}_j, \text{ans}_j, \text{pf}_j, m_j))_{j < i}, \text{cm}_i} \right] \right)^N \\ \mathcal{K}_i := \left\{ \begin{array}{l} (\mathcal{Q}_{j,\ell,i}, \text{ans}_{j,\ell,i}, \text{pf}_{j,\ell,i}) : j \in [N], i \leq \ell \leq k, \\ \text{VC.Check}(\text{pp}, \text{cm}_{j,i}, \mathcal{Q}_{j,\ell,i}, \text{ans}_{j,\ell,i}, \text{pf}_{j,\ell,i}) = 1 \end{array} \right\} \\ (\hat{\mathbf{Q}}_i, \hat{\Pi}_i) \leftarrow \mathcal{R}_{\text{post}}(\mathcal{K}_i, \mathbf{l}_i) \end{array} \right] \end{aligned}$$

Similarly, we can write

$$\left[\begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbb{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ \rho \leftarrow \{0, 1\}^r \\ ((\mathbf{Q}_i, \tilde{\mathcal{Q}}_i, \tilde{\Pi}_i))_{i \in [k]} \\ ((\text{cm}_i, \mathbf{Q}_i, \mathbf{ans}_i, \mathbf{pf}_i, m_i))_{i \in [k]} \leftarrow \langle \tilde{\mathcal{P}}(\text{aux}_0), \mathcal{V}(\text{pp}, \mathbb{x}, \rho) \rangle \\ \text{For } i \in [k] : \\ (\text{cm}_i, \text{aux}_i) \leftarrow \tilde{\mathcal{P}}(\mathbb{x}, \text{aux}_{i-1}, \mathbf{Q}_{i-1}, m_{i-1}) \\ (\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i) \leftarrow \mathcal{R}^{\tilde{\mathcal{P}}(\text{aux}_i, \cdot)}(\text{pp}, \mathbb{x}, (\text{cm}_j)_{j \leq i}, ((\mathbf{Q}_j, \mathbf{ans}_j, \mathbf{pf}_j, m_j))_{j < i}, \epsilon) \end{array} \right]$$

in terms of $\left(\mathcal{D} \left[\tilde{\mathcal{P}}(\text{ai}), \mathcal{S}(\text{pp}, \mathbb{x}, \cdot) \right] \right)^{\mathbf{N}}$:

$$\left[\begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda, n) \\ \text{ai} \leftarrow \mathcal{D} \\ (\mathbb{x}, \text{aux}_0) \leftarrow \tilde{\mathcal{P}}(\text{pp}, \text{ai}) \\ \rho \leftarrow \{0, 1\}^r \\ ((\text{cm}_i, \mathbf{Q}_i, \mathbf{ans}_i, \mathbf{pf}_i, m_i))_{i \in [k]} \leftarrow \langle \tilde{\mathcal{P}}(\text{aux}_0), \mathcal{V}(\text{pp}, \mathbb{x}, \rho) \rangle \\ ((\text{cm}_{i,j}, \mathbf{Q}_{i,j}, \mathbf{ans}_{i,j}, m_{i,j}))_{i \in [\mathbf{N}], j \in [k]} \leftarrow \left(\mathcal{U} \left[\tilde{\mathcal{P}}(\text{ai}), \mathbf{R}_{\text{cm}_1} \right] \right)^{\mathbf{N}} \\ \mathcal{K}_1 := \left\{ \begin{array}{l} (\mathcal{Q}_{i,j,1}, \mathbf{ans}_{i,j,1}, \mathbf{pf}_{i,j,1}) : i \in [\mathbf{N}], 1 \leq j \leq k, \\ \text{VC.Check}(\text{pp}, \text{cm}_{i,1}, \mathcal{Q}_{i,j,1}, \mathbf{ans}_{i,j,1}, \mathbf{pf}_{i,j,1}) = 1 \end{array} \right\} \\ (\tilde{\mathcal{Q}}_1, \tilde{\Pi}_1) \leftarrow \mathcal{R}_{\text{post}}(\mathcal{K}_1, l_1) \\ \text{For } 1 < i \leq k : \\ ((\text{cm}_{j,\ell}, \mathbf{Q}_{j,\ell}, \mathbf{ans}_{j,\ell}, \mathbf{pf}_{j,\ell}, m_{j,\ell}))_{j \in [\mathbf{N}], \ell \in [k]} \\ \leftarrow \left(\mathcal{D} \left[\tilde{\mathcal{P}}(\text{ai}), \mathcal{S}(\text{pp}, \mathbb{x}, ((\text{cm}_j, \mathbf{Q}_j, \mathbf{ans}_j, \mathbf{pf}_j, m_j))_{j < i}, \text{cm}_i) \right] \right)^{\mathbf{N}} \\ \mathcal{K}_i := \left\{ \begin{array}{l} (\mathcal{Q}_{j,\ell,i}, \mathbf{ans}_{j,\ell,i}, \mathbf{pf}_{j,\ell,i}) : j \in [\mathbf{N}], i \leq \ell \leq k, \\ \text{VC.Check}(\text{pp}, \text{cm}_{j,i}, \mathcal{Q}_{j,\ell,i}, \mathbf{ans}_{j,\ell,i}, \mathbf{pf}_{j,\ell,i}) = 1 \end{array} \right\} \\ (\tilde{\mathcal{Q}}_i, \tilde{\Pi}_i) \leftarrow \mathcal{R}_{\text{post}}(\mathcal{K}_i, l_i) \end{array} \right]$$

Therefore, we conclude that

$$\begin{aligned} \Pr \left[\exists i \in [k], j \in [i, k], q \in \mathcal{Q}_{j,i} \setminus \tilde{\mathcal{Q}}_i \right] &\leq \Pr \left[\exists i \in [k], j \in [i, k], q \in \mathcal{Q}_{j,i} \setminus \hat{\mathcal{Q}}_i \right] + \mathbf{N} \cdot \alpha \\ &\leq \frac{1}{\mathbf{N}} + \mathbf{N} \cdot \alpha, \end{aligned}$$

where $\Pr \left[\exists i \in [k], j \in [i, k], q \in \mathcal{Q}_{j,i} \setminus \hat{\mathcal{Q}}_i \right]$ is with respect to the same experiment as in Eq. 6. \square

Acknowledgments

Alessandro Chiesa and Ziyi Guan are partially supported by the Ethereum Foundation. We thank Fermi Ma and Julius Vering for valuable discussions and participating in early stages of this work. We thank Zijing Di for valuable feedback and comments on earlier drafts of this paper.

References

- [AC20] Thomas Attema and Ronald Cramer. “Compressed Σ -Protocol Theory and Practical Application to Plug & Play Secure Algorithmics”. In: *Proceedings of the 40th Annual International Cryptology Conference*. CRYPTO ’20. 2020, pp. 513–543.
- [ACK22] Thomas Attema, Ronald Cramer, and Lisa Kohl. “A Compressed Σ -Protocol Theory for Lattices”. In: *Proceedings of the 41st Annual International Cryptology Conference*. CRYPTO ’21. 2022, pp. 549–579.
- [AF22] Thomas Attema and Serge Fehr. “Parallel Repetition of (k_1, \dots, k_μ) -Special-Sound Multi-Round Interactive Proofs”. In: *Proceedings of the 42nd Annual International Cryptology Conference*. CRYPTO ’22. 2022, pp. 415–443.
- [AFR23] Thomas Attema, Serge Fehr, and Nicolas Resch. *A Generalized Special-Soundness Notion and its Knowledge Extractors*. IACR Cryptology ePrint Archive, Report 2023/818. 2023.
- [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.
- [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.
- [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.
- [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.
- [BGTZ23] Alexander R. Block, Albert Garreta, Pratyush Ranjan Tiwari, and Michal Zajac. “On Soundness Notions for Interactive Oracle Proofs”. In: (2023), p. 1256.
- [BKKMS13] Eli Ben-Sasson, Yohay Kaplan, Swastik Kopparty, Or Meir, and Henning Stichtenoth. “Constant Rate PCPs for Circuit-SAT with Sublinear Query Complexity”. In: *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science*. FOCS ’13. 2013, pp. 320–329.
- [BS06] Eli Ben-Sasson and Madhu Sudan. “Robust locally testable codes and products of codes”. In: *Random Structures and Algorithms* 28.4 (2006), pp. 387–402.
- [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.
- [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.
- [CMS19] Alessandro Chiesa, Peter Manohar, and Nicholas Spooner. “Succinct Arguments in the Quantum Random Oracle Model”. In: *Proceedings of the 17th Theory of Cryptography Conference*. TCC ’19. Available as Cryptology ePrint Archive, Report 2019/834. 2019, pp. 1–29.

- [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. 2021, pp. 49–58.
- [COS20] Alessandro Chiesa, Dev Ojha, and Nicholas Spooner. “Fractal: Post-Quantum and Transparent Recursive Proofs from Holography”. In: *Proceedings of the 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques*. EUROCRYPT ’20. 2020, pp. 769–793.
- [CY21a] Alessandro Chiesa and Eylon Yogev. In: *Proceedings of the 41st Annual International Cryptology Conference*. CRYPTO ’21. 2021, pp. 711–741.
- [CY21b] Alessandro Chiesa and Eylon Yogev. “Tight Security Bounds for Micali’s SNARGs”. In: *Proceedings of the 19th Theory of Cryptography Conference*. TCC ’21. 2021, pp. 401–434.
- [GW11] Craig Gentry and Daniel Wichs. “Separating Succinct Non-Interactive Arguments From All Falsifiable Assumptions”. In: *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*. STOC ’11. 2011, pp. 99–108.
- [HR22] Justin Holmgren and Ron D. Rothblum. “Faster Sounder Succinct Arguments and IOPs”. In: *Proceedings of the 42rd Annual International Cryptology Conference*. CRYPTO ’22. 2022, pp. 474–503.
- [KPT97] Joe Kilian, Erez Petrank, and Gábor Tardos. “Probabilistically checkable proofs with zero knowledge”. In: *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*. STOC ’97. 1997, pp. 496–505.
- [Kil92] Joe Kilian. “A note on efficient zero-knowledge proofs and arguments”. In: *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*. STOC ’92. 1992, pp. 723–732.
- [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.
- [LMS22] Alex Lombardi, Fermi Ma, and Nicholas Spooner. “Post-Quantum Zero Knowledge, Revisited or: How to Do Quantum Rewinding Undetectably”. In: *Proceedings of the 63rd Annual IEEE Symposium on Foundations of Computer Science*. FOCS ’22. 2022, pp. 851–859.
- [Mic00] Silvio Micali. “Computationally Sound Proofs”. In: *SIAM Journal on Computing* 30.4 (2000). Preliminary version appeared in FOCS ’94., pp. 1253–1298.
- [RR22] Noga Ron-Zewi and Ron D. Rothblum. “Proving as fast as computing: succinct arguments with constant prover overhead”. In: *Proceedings of the 54th Annual ACM Symposium on Theory of Computing*. STOC ’22. 2022, pp. 1353–1363.
- [RRR16] Omer Reingold, Ron Rothblum, and Guy Rothblum. “Constant-Round Interactive Proofs for Delegating Computation”. In: *Proceedings of the 48th ACM Symposium on the Theory of Computing*. STOC ’16. 2016, pp. 49–62.
- [Val08] Paul Valiant. “Incrementally Verifiable Computation or Proofs of Knowledge Imply Time/Space Efficiency”. In: *Proceedings of the 5th Theory of Cryptography Conference*. TCC ’08. 2008, pp. 1–18.