

Hardness of Approximation for Stochastic Problems via Interactive Oracle Proofs

Gal Arnon
gal.arnon@weizmann.ac.il
Weizmann Institute

Alessandro Chiesa
alessandro.chiesa@epfl.ch
EPFL

Eylon Yogev
eylon.yogev@biu.ac.il
Bar-Ilan University

February 14, 2022

Abstract

Hardness of approximation aims to establish lower bounds on the approximability of optimization problems in NP and beyond. We continue the study of hardness of approximation for problems beyond NP, specifically for *stochastic* constraint satisfaction problems (SCSPs). An SCSP with k alternations is a list of constraints over variables grouped into $2k$ blocks, where each constraint has constant arity. An assignment to the SCSP is defined by two players who alternate in setting values to a designated block of variables, with one player choosing their assignments uniformly at random and the other player trying to maximize the number of satisfied constraints.

In this paper, we establish hardness of approximation for SCSPs based on interactive proofs. For $k \leq O(\log n)$, we prove that it is $\text{AM}[k]$ -hard to approximate, to within a constant, the value of SCSPs with k alternations and constant arity. Before, this was known only for $k = O(1)$.

Furthermore, we introduce a natural class of k -round interactive proofs, denoted $\text{IR}[k]$ (for *interactive reducibility*), and show that several protocols (e.g., the sumcheck protocol) are in $\text{IR}[k]$. Using this notion, we extend our inapproximability to all values of k : we show that for every k , approximating an SCSP instance with $O(k)$ alternations and constant arity is $\text{IR}[k]$ -hard.

While hardness of approximation for CSPs is achieved by constructing suitable PCPs, our results for SCSPs are achieved by constructing suitable IOPs (interactive oracle proofs). We show that every language in $\text{AM}[k \leq O(\log n)]$ or in $\text{IR}[k]$ has an $O(k)$ -round IOP whose verifier has *constant* query complexity (*regardless* of the number of rounds k). In particular, we derive a “sumcheck protocol” whose verifier reads $O(1)$ bits from the entire interaction transcript.

Keywords: hardness of approximation; interactive oracle proofs; stochastic satisfaction problems

Contents

1	Introduction	3
1.1	Our results	4
2	Techniques	7
2.1	On the hardness of approximating SCSPs via IOPs	7
2.2	Transforming IPs into IOPs	8
2.3	Interactive reducibility	11
3	Preliminaries	16
3.1	Interactive oracle proofs	16
3.2	Round-by-round soundness	17
3.3	Extractors	18
4	Interactive reducibility	19
4.1	Interactive proofs from interactive reductions	19
4.2	Error reduction for interactive reducibility	21
4.3	Relations with interactive reducibility	23
5	IOPs from interactive reducibility	32
5.1	Round-query IOPs from bounded-output-length interactive reductions	33
5.2	Round-query IOPs from public-coin IPs	36
6	Hardness of approximation for stochastic problems	38
6.1	Stochastic constraint satisfaction problems	38
6.2	Hardness of approximation for SCSP	38
7	From round-query IOPs to binary IOPs	40
7.1	Local access to interaction randomness	40
7.2	Local access to prover messages	44
	Acknowledgments	46
	References	46

1 Introduction

Many combinatorial optimization problems are NP-hard, and so there is little hope to solve them in polynomial time. This has motivated the study of polynomial-time *approximation algorithms* to solve optimization problems, which has revealed a surprising landscape. While they are equivalent as *decision* problems (under polynomial-time reductions), NP-complete problems behave radically different from the point of view of approximability. Specifically, known approximation algorithms for NP-complete problems sometimes achieve approximations to within any constant, sometimes only to within a certain constant, and sometimes do not even achieve a constant approximation.

This differing behavior is justified via results in the area of *hardness of approximation*. For a given NP-complete problem, the goal is to prove that it is NP-hard to approximate the problem to better than a certain approximation ratio (e.g., better than $1/2$). Ideally, this ratio would match the best known approximation algorithm, thereby ruling out better approximation algorithms.

The key tool used to establish hardness of approximation for NP-complete problems are probabilistic proofs [FGLSS96]. For example, the PCP theorem [AS98; ALMSS98] says that every language in NP can be decided via a verifier that reads $O(1)$ bits from a polynomial-length proof, and in turn this implies, e.g., that the value of 3SAT cannot be approximated to within an arbitrary constant.

More generally, improvements in PCP constructions imply hardness results for corresponding *constraint satisfaction problems* (CSPs). Yet, there are numerous problems of interest that are *not* CSPs, and for which we wish to understand their behavior with respect to inapproximability.

Hardness of approximation beyond NP. Prior work has investigated hardness of approximation for natural problems in other complexity classes. In one direction, PCP-like theorems for fine-grained complexity have been used to establish hardness results for problems within the complexity class P (see [AB17; ARW17; AR18; CW19; CGLRR19]). In the other direction, several works study the inapproximability of two-player CSPs. A CSP can be viewed as a one-player game where the player wishes to maximize the number of satisfied constraints; this view naturally leads to two-player CSPs played in moves. Ko and Lin [KL94] proved the inapproximability of two-player CSPs with k moves, based on the hardness of the k -th level of the Polynomial Hierarchy. Haviv, Regev, and Ta-Shma [HRT07] proved that this inapproximability result holds even when each variable occurs $O(1)$ times.

SCSPs and their hardness. In this paper, we study the hardness of approximating a natural class of problems known as *stochastic constraint satisfaction problems* (SCSPs), also known as *games against nature* [Pap83]. Informally, they are two-player CSPs where one player is an adversary and the other player is a (public-coin) referee that plays random moves.

Definition 1.1 (informal). *An SCSP Φ with k alternations is a list of constraints C_1, \dots, C_m over variables that are grouped into $2k$ blocks and take on values over an alphabet Σ . The SCSP has arity q if each constraint depends on at most q variables. An assignment for Φ is defined by two players who alternate in setting values to a designated block of variables with one player choosing their assignments uniformly at random and the other player trying to maximize the number of satisfied clauses. The value of Φ is the expected fraction of clauses satisfied in this process.*

The hardness of approximating the value of SCSPs, to within a constant factor, has been studied in a line of works. Let $AM[k]$ be the class of languages that have a k -round public-coin IP with constant soundness error. Drucker [Dru11] extended the PCP theorem to the stochastic setting, showing that it is $AM[1]$ -complete to approximate the value of SCSPs with one alternation ($k = 1$) and arity $q = O(1)$. Subsequently, [ACY22] showed that, for every k , it is $AM[k]$ -complete to

approximate the value of SCSPs with k alternations and arity $q = O(k)$. In the regime of many alternations, Condon, Feigenbaum, Lund, and Shor [CFLS97] showed that it is PSPACE-complete to approximate the value of SCSPs with $k = \text{poly}$ alternations and arity $q = O(1)$. This leaves open the approximation hardness of SCSPs with k alternations and *constant* arity, for general values of k :

How hard is approximating the value of SCSPs with k alternations and arity $q = O(1)$?

It seems reasonable to hypothesize that it is $\text{AM}[k]$ -hard to approximate SCSPs with k alternations.

Note that Goldreich, Vadhan, and Wigderson [GVW02] showed that $\text{AM}[k] \neq \text{AM}[o(k)]$ (under reasonable hardness assumptions), meaning that increasing the round complexity k adds more power to the complexity class $\text{AM}[k]$. For sufficiently many rounds, we know that $\text{IP} = \text{PSPACE}$ [LFKN92; Sha92]. Thus, it is imperative to understand the approximation hardness of SCSPs with k alternations while respecting the different regimes for k .

SCSP hardness from IOPs. The above results are derived (implicitly or explicitly) by leveraging the connection between SCSPs and the PCP analog of interactive proofs, called *interactive oracle proofs* (IOPs) [BCS16; RRR16]. A k -round (public-coin) IOP is a k -round (public-coin) IP where the verifier has PCP-like access to each prover message: after the interaction, the verifier probabilistically reads a small number of locations from the interaction transcript and then accepts or rejects.

A k -round public-coin IOP with query complexity q can be viewed as an SCSP with k alternations and arity q (see Section 2.1). Therefore, constructions of IOPs for hard languages imply corresponding hardness of approximation results for the resulting SCSPs. This leads us to ask:

Does every language in $\text{AM}[k]$ have a k -round IOP with constant query complexity?

1.1 Our results

In this paper, we establish hardness of approximation for SCSPs from (general) interactive proofs. Moreover, we also prove that tighter hardness results can be achieved for specific languages that are *interactively reducible* (a notion that we introduce).

On the AM hardness of SCSPs. We prove that it is $\text{AM}[k]$ -hard to approximate the value of binary-alphabet SCSPs with k alternations and arity $\max\{O(1), O(k/\log |\mathfrak{x}|)\}$ (\mathfrak{x} is the instance).

Theorem 1 (informal). *Let $L \in \text{AM}[k]$ be a language. There exists a deterministic polynomial-time reduction that maps an instance \mathfrak{x} for L to an SCSP instance Φ with binary alphabet, k alternations, and arity $\max\{O(1), O(k/\log |\mathfrak{x}|)\}$ such that:*

- *if $\mathfrak{x} \in L$ then the value of Φ is 1;*
- *if $\mathfrak{x} \notin L$ then the value of Φ is at most $1/2$.*

Our improvement in arity is particularly meaningful for logarithmic round complexity: Theorem 1 establishes that, for $k(|\mathfrak{x}|) = O(\log |\mathfrak{x}|)$, approximating the value of an SCSP instance with k alternations and arity $O(1)$ is as hard as deciding all of $\text{AM}[k]$. Previously, this was known only for constant k [Dru11; ACY22].

Many results on the hardness of approximating the value of (standard) CSPs are achieved by constructing suitable PCPs. Similarly (as was noted in [ACY22]), by constructing a suitable k -round IOP for a language L , one can show that approximating the value of SCSPs with k alternations up to a constant factor is as hard as deciding L . We establish Theorem 1 using this framework by providing a transformation that maps a k -round IP into a k -round IOP with small query complexity.

Lemma 1 (informal). *Let L be a language with a k -round public-coin IP. Then L has a k -round public-coin binary IOP where, on input \mathbf{x} , the verifier reads $\max\{O(1), O(k/\log |\mathbf{x}|)\}$ bits of the interaction transcript. All other parameters are polynomially related.*

Lemma 1 is surprising in light of the work of Goldreich, Vadhan, and Wigderson [GVW02], which shows a separation between $\text{AM}[k]$ and $\text{AM}[o(k)]$ (under relatively weak hardness assumptions). Since the query complexity is smaller than the round complexity, Lemma 1 implies that the IOP verifier does not make queries to every round of the protocol. This can be viewed as saying that the power of $\text{AM}[k]$ is unchanged even when the verifier only accesses $O(k/\log |\mathbf{x}|)$ of the k rounds. In other words, reducing round complexity of public-coin protocols reduces their power, but it is nevertheless possible to not read every round of interaction while preserving the power.

Hardness of SCSPs from interactive reducibility. Theorem 1 establishes the hardness of SCSPs with $k = O(\log |\mathbf{x}|)$ alternations and arity $O(1)$ (over the binary alphabet), but does not work for $O(1)$ -arity SCSPs with $k = \omega(\log |\mathbf{x}|)$ alternations.

We extend this by showing that approximating the value of $O(1)$ -arity SCSPs with k alternations is as hard as solving $\#\text{SAT}_k$.¹ In fact, we show this for a more general class of languages that are *interactively reducible*, a notion that we introduce in this work. Informally, the notion requires that it is possible to reduce, via an interactive protocol, multiple transcripts of an IP for the relation into a single transcript. We require that the probability of the verifier accepting conditioned on the reduced transcript be (roughly) the minimum probability of the verifier accepting conditioned on any of the original transcripts. See Section 2.3 for further details and discussion.

Interactive reducibility is a natural property; we show that general interactive proofs can be seen as interactive reductions (albeit ones with bad parameters), and show interactive reductions for the sumcheck protocol [LFKN92] and for Shamir’s protocol [Sha92].

The notion of interactive reducibility allows us to get an optimal version of Lemma 1 for additional languages. Let $\text{IR}[k]$ be the class of languages that have a (1-round) interactive reduction with k predicates (these predicates roughly align with rounds of an IP).

Lemma 2 (informal). *Let L be a language in $\text{IR}[k]$. Then L has an $O(k)$ -round public-coin IOP, with polynomial proof length, where the verifier reads $O(1)$ bits of the interaction transcript.*

In particular, applying Lemma 2 to the sumcheck protocol yields the following (perhaps surprising) conclusion: *any k -round sumcheck protocol can be transformed to a $O(k)$ -round IOP where the verifier has $O(1)$ query complexity (over the binary alphabet).* Notice that using standard PCP techniques it is only known how to achieve a similar (non-interactive) result with *exponential* proof length.

Using this improved lemma, we immediately get that for every k , deciding whether a binary-alphabet SCSP instance with $O(k)$ alternations and arity $O(1)$ has value 1 or value $1/2$ is $\text{IR}[k]$ -hard:

Theorem 2 (informal). *Let $L \in \text{IR}[k]$ be a language. There exists a deterministic polynomial-time reduction that maps an instance \mathbf{x} for L to an SCSP instance Φ with binary alphabet, $O(k)$ alternations, and arity $O(1)$ such that:*

- if $\mathbf{x} \in L$ then the value of Φ is 1;
- if $\mathbf{x} \notin L$ then the value of Φ is at most $1/2$.

Using our interactive reduction for sumcheck, we know that $\#\text{SAT}_k \in \text{IR}[k]$. Thus, by Theorem 2, we establish that approximating the value of SCSPs with $O(k)$ alternations and constant arity is

¹ $\#\text{SAT}_k$ is the restriction of $\#\text{SAT}$ to instances of size n and $k(n)$ variables.

$\#\text{SAT}_k$ -hard. Similarly, using our interactive reduction for Shamir’s protocol, we recover the result of [CFLS97], showing that approximating the value of SCSPs with polynomially-many alternations and constant arity is PSPACE-hard.

Summary of results and open questions. We construct $O(1)$ -query IOPs for every language in $\text{AM}[k \leq O(\log |\mathbb{x}|)]$ or in $\text{IR}[k]$. Moreover, we construct $O(k/\log |\mathbb{x}|)$ -query IOPs for every language in $\text{AM}[k]$. These results establish approximation hardness for SCSPs as follows: (a) for $k = O(\log |\mathbb{x}|)$, it is $\text{AM}[k]$ -hard to approximate the value of SCSPs with k alternations and constant arity; and (b) for $k = \omega(\log |\mathbb{x}|)$, it is $\text{IR}[k]$ -hard to approximate the value of SCSPs with $O(k)$ alternations and constant arity, and $\text{AM}[k]$ -hard when the SCSPs have arity $O(k/\log |\mathbb{x}|)$. Our results are summarized in Figure 1 together with previously known results.

Our work leaves open the AM -hardness of approximating the value of SCSPs with $k = \omega(\log |\mathbb{x}|)$ alternations and constant arity. From the perspective of IOPs, we also leave open the basic question that we raised in the introduction: *Does every language in $\text{AM}[k]$ have a k -round (public-coin) IOP with constant query complexity over the binary alphabet?* Our results contribute notable progress towards resolving this question (see paragraph above), but answering the question for every regime of k remains a fascinating challenge in the theory of probabilistic proofs.

	hardness	alternations	arity
[CFLS97]	PSPACE	unbounded	$O(1)$
[this work]	$\text{IR}[k]$	$O(k)$	$O(1)$
[this work]	$\#\text{SAT}_k$	$O(k)$	$O(1)$
[ACY22]	$\text{AM}[k]$	k	$O(k)$
[this work]	$\text{AM}[k]$	k	$\max\{O(1), O(k/\log \mathbb{x})\}$
[Dru11]	$\text{AM}[1]$	1	$O(1)$
[ALMSS98; AS98; Din07]	NP	n/a	$O(1)$

Figure 1: Summary of results for approximating the value of binary-alphabet SCSPs to within a constant factor. $\text{AM}[k]$ denotes the class of languages with k -round public-coin interactive proofs. $\text{IR}[k]$ denotes the class of languages with (1-round) interactive reductions with k predicates. $\#\text{SAT}_k$ is the restriction of $\#\text{SAT}$ to instances of size n and $k(n)$ variables.

2 Techniques

We outline the main ideas behind our results. In Section 2.1 we explain a generic connection between SCSPs and IOPs: in order to establish the hardness of approximating SCSPs, it suffices to construct IOPs with certain properties. This will be our goal in the remaining sections. In Section 2.2 we show how to transform k -round IPs into k -round IOPs with query complexity $\max\{O(1), O(k/\log |\mathbf{x}|)\}$. In Section 2.3 we show that for relations that are *interactively reducible* we can construct $O(1)$ -query IOPs even when those relations are only known to have IPs with round complexity $\omega(\log |\mathbf{x}|)$.

2.1 On the hardness of approximating SCSPs via IOPs

We review the generic connection between CSPs and PCPs, and then describe the analogous connection between SCSPs and IOPs. In both cases, efficient constructions of PCPs/IOPs imply hardness of approximation results for corresponding CSPs/SCSPs.

CSP hardness from PCPs. A CSP Φ is a list of boolean functions C_1, \dots, C_m over variables from a bounded alphabet Σ . The CSP has arity q if each constraint depends on at most q variables. The goal is to determine the maximum fraction of constraints that can be satisfied by any assignment.

We can map a non-adaptive PCP verifier \mathbf{V} for a language L and an instance \mathbf{x} into a CSP. The variables represent locations of the PCP string. Each choice of PCP verifier randomness induces a corresponding constraint, whose variables are those that the PCP verifier would have read from the PCP string. The constraint is satisfied if and only if the PCP verifier accepts when it receives the assignment of the variables as its query answers. The CSP's arity equals the PCP's query complexity.

By completeness of the PCP system, if $\mathbf{x} \in L$ then there exists a PCP string that makes the PCP verifier accept with probability 1, which in turn means that there is an assignment that simultaneously satisfies every constraint in the CSP. By the soundness of the PCP system, if $\mathbf{x} \notin L$ then every PCP string makes the PCP verifier accept with at most probability $1/2$, which in turn means that no assignment can satisfy more than half of the constraints of the CSP.

Thus, distinguishing whether the CSP's value is 1 or at most $1/2$ is as hard as deciding L .

SCSP hardness from IOPs. The general connection between SCSPs (Definition 1.1) and IOPs is stated in the lemma below. Recall that a k -round IOP is a k -round IP where the verifier has PCP-like access to each prover message: the prover and verifier interact for k rounds, and after the interaction, the verifier probabilistically reads a small number of bits from each prover message and decides to accept or reject based on the examined locations. The randomness used in the final phase is called *decision randomness* (which we distinguish from the random messages that the verifier sends to the prover during the interaction and may be queried at only few locations).

Lemma 3. *Let L be a language with a non-adaptive k -round public-coin IOP with alphabet Σ , polynomial proof length, query complexity q , decision randomness r_{dc} , and soundness error β .*

Then there exists a deterministic polynomial-time reduction that maps an instance \mathbf{x} for L to an SCSP instance Φ with alphabet Σ , k alternations, arity q , $2^{r_{\text{dc}}}$ constraints and a polynomial number of variables such that:

- *if $\mathbf{x} \in L$ then the value of Φ is 1;*
- *if $\mathbf{x} \notin L$ then the value of Φ is at most β .*

Proof sketch. Let \mathbf{V}_{IOP} be the (non-adaptive) IOP verifier for L and let l be the (per-round) proof length of the IOP. Given an instance \mathbf{x} , we construct the SCSP instance Φ as follows. The SCSP has

$2k$ blocks of l variables that align with the interaction transcript between the IOP prover and IOP verifier (the i -th variable of the j -th block corresponds to the i -th symbol of the j -th message of the protocol). The SCSP has a constraint C_ρ for each $\rho \in \{0, 1\}^{r_{dc}}$, whose input variables correspond to the locations of the transcript that \mathbf{V}_{IOP} queries given instance \mathbf{x} and randomness ρ ; the constraint C_ρ is satisfied if and only if $\mathbf{V}_{IOP}(\mathbf{x}; \rho)$ accepts if it reads the symbols assigned to the variables of C_ρ .

The SCSP instance Φ has k alternations (corresponding to the rounds of the IOP) and $l = \text{poly}(|\mathbf{x}|)$ variables per alternation (corresponding to the message length) that are assigned values in the alphabet Σ (the alphabet of the IOP). Each of its $2^{r_{dc}}$ constraints has arity q since each constraint has as many inputs as queries made by the IOP verifier \mathbf{V}_{IOP} .

Finally, we analyze the value of the SCSP. By construction, there exists an IOP prover strategy that causes the IOP verifier \mathbf{V}_{IOP} to accept with probability δ if and only if there exists a strategy for the existential player in the SCSP such that the expected fraction of constraints that are satisfied is δ (i.e., the value of Φ is at least δ). Therefore, by perfect completeness of the IOP, if $\mathbf{x} \in L$ then the value of Φ is 1. Conversely, by soundness of the IOP, if $\mathbf{x} \notin L$ then the value of Φ is at most β . \square

2.2 Transforming IPs into IOPs

We outline the proof of Lemma 1 (transforming an IP into an IOP with small query complexity). In Section 2.2.1, we show how to transform a logarithmic-round IP into a $O(1)$ -query IOP. Then in Section 2.2.2 we extend this idea to transform a k -round IP into a $O(k/\log |\mathbf{x}|)$ -query IOP.

2.2.1 From $O(\log |\mathbf{x}|)$ -round IP to $O(1)$ -query IOP

We show how to transform a k -round public-coin IP where $k = O(\log |\mathbf{x}|)$ into an $O(k)$ -round public-coin IOP with the following efficiency: polynomial proof length over the binary alphabet; constant query complexity; and logarithmic decision randomness.

First, we sketch how to transform a k -round public-coin IP $(\mathbf{P}_{IP}, \mathbf{V}_{IP})$ into an $O(k)$ -round public-coin IOP $(\mathbf{P}_{IOP}, \mathbf{V}_{IOP})$ where the verifier reads $O(1)$ rounds (in their entirety) from the interaction transcript. Then we explain how to ensure that the verifier queries $O(1)$ bits in total.

A strawman protocol. We describe a natural strategy for transforming the IP into an IOP where the verifier reads $O(1)$ rounds, albeit with high soundness error. The IOP prover \mathbf{P}_{IOP} and IOP verifier \mathbf{V}_{IOP} respectively simulate the IP prover and verifier $(\mathbf{P}_{IP}, \mathbf{V}_{IP})$, inducing an interaction transcript $\text{tr} = (\rho_1, a_1, \dots, \rho_k, a_k)$. At this time, however, \mathbf{V}_{IOP} *does not read any messages from* tr . After this interaction, \mathbf{P}_{IOP} sends a transcript tr' , which is allegedly equal to tr , as a single message. Then \mathbf{V}_{IOP} reads tr' and checks that tr' is an accepting transcript for the IP verifier \mathbf{V}_{IP} . Moreover, \mathbf{V}_{IOP} tests consistency between tr (the real interaction) and tr' (the alleged copy of the interaction sent as a single message): \mathbf{V}_{IOP} samples a random $i \in [k]$, reads the i -th prover message and i -th verifier message in tr , and checks that these equal the corresponding messages in tr' .

In this IOP, \mathbf{V}_{IOP} reads $O(1)$ messages from the interaction transcript, but the soundness error of the IOP is large (even when discounting the soundness error of the underlying IP). Indeed, it may be that a cheating IOP prover sends a malicious transcript tr' that is accepting but differs from the real transcript tr in one round only. In this case, \mathbf{V}_{IOP} catches the inconsistency only with probability $1/k$, which means that the soundness error could be as large as $1 - 1/k$.

Note that reducing this soundness error via parallel repetition would increase the number of rounds queried by \mathbf{V}_{IOP} . Achieving constant soundness error would require $O(k)$ repetitions, resulting in an IOP verifier that reads $O(k)$ rounds, taking us back to where we started.

Our transformation. We present a transformation that improves on the above strawman protocol, achieving a constant soundness error for any IP that has a logarithmic number of rounds.

A malicious IOP prover in the strawman protocol has two strategies: the transcript tr' sent as the last message either agrees with the real transcript tr on more than half of the rounds, or it does not. If tr' agrees with tr in less than half of the rounds, then the IOP verifier catches this inconsistency with probability at least $1/2$. Intuitively, the transformation that we sketch below ensures that if tr' is consistent with tr on at least half of the rounds, then the consistent rounds must contain within them a full execution of the underlying IP. Then, since this consistent part was generated interactively in tr , by the soundness property of the IP, this contained transcript will be rejected with high probability. We now describe our transformation in more detail.

Suppose that our public-coin IP has $k(|\mathbf{x}|) = O(\log |\mathbf{x}|)$ rounds and that the verifier message in each round is r bits long. The IOP $(\mathbf{P}_{\text{IOP}}, \mathbf{V}_{\text{IOP}})$ on a given instance \mathbf{x} works as follows.

1. \mathbf{P}_{IOP} sets $S_0 := \{\emptyset\}$ (i.e., S_0 consists of the empty transcript).
2. For $i = 1, \dots, 2k$:
 - \mathbf{P}_{IOP} sends S_{i-1} .
 - \mathbf{V}_{IOP} sends a random $\rho_i \in \{0, 1\}^r$ (corresponding to a message of \mathbf{V}_{IP}).
 - \mathbf{P}_{IOP} sets $S_i := S_{i-1} \cup \{(\text{tr} \parallel \rho_i \parallel a_{\text{tr},i})\}_{\text{tr} \in S_{i-1}}$ where $a_{\text{tr},i} := \mathbf{P}_{\text{IP}}(\mathbf{x}, \text{tr} \parallel \rho_i)$ for each $\text{tr} \in S_{i-1}$.
3. \mathbf{P}_{IOP} sends S_{2k} and, for every $i \in [2k]$, also sends $T_i := S_i$.
4. In the decision phase, \mathbf{V}_{IOP} performs the checks below.
 - (a) *Subset consistency:* For every $i \in [2k]$, check that $T_{i-1} \subseteq T_i$.
 - (b) *Transcript consistency:* Choose a random $i \in [2k]$. Check that $S_{i-1} = T_{i-1}$ and $S_i = T_i$. Additionally, check that for every $\text{tr} \in S_{i-1}$ there is a message $a_{\text{tr},i}$ such that $(\text{tr} \parallel \rho_i \parallel a_{\text{tr},i}) \in S_i$, where ρ_i is the verifier message sent during the i -th round of interaction.
 - (c) *Membership:* Check that for every transcript $\text{tr} \in T_{2k}$ that is complete (i.e., contains messages for all k rounds of the IP) it holds that $\mathbf{V}_{\text{IP}}(\mathbf{x}, \text{tr}) = 1$.

Efficiency. We briefly discuss the main efficiency measures of the transformation.

- *Query complexity.* The IOP verifier reads $O(1)$ rounds from the transcript.
- *Communication complexity.* We argue that all messages in the protocol have length $\text{poly}(|\mathbf{x}|)$. For every i , $|S_i| \leq 2|S_{i-1}|$ since S_i contains all transcripts in S_{i-1} and continuations of each of these transcripts. Since $k = O(\log |\mathbf{x}|)$, $|S_i| = \text{poly}(|\mathbf{x}|)$ for every i . Each transcript within S_i has polynomial length, so the length of these messages is $\text{poly}(|\mathbf{x}|)$. Finally, the sets T_1, \dots, T_{2k} have the same sizes as S_1, \dots, S_{2k} (respectively) and so the prover's final message has length $\text{poly}(|\mathbf{x}|)$.
- *Decision randomness.* The IOP verifier uses $O(\log k) = O(\log |\mathbf{x}|)$ bits of decision randomness.

Analysis. In this overview, we discuss soundness only, as completeness follows straightforwardly from the construction. Let β be the soundness error of $(\mathbf{P}_{\text{IP}}, \mathbf{V}_{\text{IP}})$. We show that the IOP $(\mathbf{P}_{\text{IOP}}, \mathbf{V}_{\text{IOP}})$ has soundness error

$$\beta_{\text{IOP}} = \max \left\{ \frac{1}{2}, \binom{2k}{k} \cdot \beta \right\} .$$

The above expression can be made constant by first applying $\text{poly}(|\mathbf{x}|)$ repetitions to $(\mathbf{P}_{\text{IP}}, \mathbf{V}_{\text{IP}})$, until it has soundness error $\beta' \leq \left(2 \cdot \binom{2k}{k}\right)^{-1}$ (and only then applying the IP-to-IOP transformation).

Fix $\mathbf{x} \notin L$ and a cheating IOP prover $\tilde{\mathbf{P}}_{\text{IOP}}$. A fixed transcript of the IOP has the structure:

$$(S_0, \rho_1, S_1, \dots, \rho_{2k_{\text{IP}}}, S_{2k}, (T_0, \dots, T_{2k})) .$$

Given such a transcript, we say that an index i is *consistent* if: (a) $S_{i-1} = T_{i-1}$ and $S_i = T_i$; and (b) for every $\text{tr} \in S_{i-1}$ there is a message $a_{\text{tr},i}$ such that $(\text{tr} \parallel \rho_i \parallel a_{\text{tr},i}) \in S_i$.

Conditioned on the event that the transcript generated during the interaction has less than k consistent indices i , \mathbf{V}_{IOP} rejects with probability at least $1/2$ due to its check in Item 4b. We are thus left to analyze the probability that \mathbf{V}_{IOP} rejects conditioned on the event that the generated transcript has at least k consistent indices.

Fix indices $i_1 < \dots < i_k$ and suppose that all these indices are consistent with respect to the transcript. By the definition of consistency, for every $j \in [k]$, $S_{i_j-1} = T_{i_j-1}$, $S_{i_j} = T_{i_j}$ and $(\text{tr} \parallel \rho_{i_j} \parallel a_{\text{tr},i_j}) \in S_{i_j}$ for every $\text{tr} \in S_{i_j-1}$. This implies that there exist a_{i_1}, \dots, a_{i_k} such that $(\rho_{i_1} \parallel a_{i_1} \parallel \dots \parallel \rho_{i_k} \parallel a_{i_k}) \in T_{i_k}$. By the subset consistency check, \mathbf{V}_{IOP} accepts only if $T_{i_k} \subseteq T_{2k}$, in which case $(\rho_{i_1} \parallel a_{i_1} \parallel \dots \parallel \rho_{i_k} \parallel a_{i_k}) \in T_{2k}$. This transcript was generated interactively by the prover and verifier and hence, by the soundness of the IP, $\mathbf{V}_{\text{IP}}(\mathbf{x}, \rho_{i_1} \parallel a_{i_1} \parallel \dots \parallel \rho_{i_k} \parallel a_{i_k}) = 1$ with probability at most β . If the transcript is rejecting, then this is detected by \mathbf{V}_{IOP} in its membership check.

The previous analysis holds for fixed indices $i_1 < \dots < i_k$. By applying the union bound to all choices of indices, we have that, conditioned on the transcript generated having at least k consistent rounds, \mathbf{V}_{IOP} accepts with probability at most $\binom{2k}{k} \cdot \beta$.

Putting together both (non-intersecting) events of the number of rounds consistent with the generated transcript, we conclude that \mathbf{V}_{IOP} accepts with probability at most $\max\{\frac{1}{2}, \binom{2k}{k} \cdot \beta\}$.

Achieving query complexity $O(1)$ over the binary alphabet. The verifier in the IOP described above reads $O(1)$ rounds (in their entirety) from the interaction transcript, rather than $O(1)$ bits in total. We additionally achieve this latter goal by building on a result in [ACY22].

In more detail, [ACY22] transforms a k -round public-coin IP into an $O(k)$ -round public-coin IOP with polynomial proof length over the binary alphabet and where the IOP verifier reads $O(1)$ bits from each round. We extend this to transform a k -round IOP whose verifier reads q of the k rounds into a $O(k)$ -round IOP whose verifier reads $O(1)$ bits from each of $O(q)$ rounds. See Section 7 for more details.

2.2.2 IOPs from general IPs

The transformation described in the previous section works for $O(\log |\mathbf{x}|)$ -round IPs. However, it cannot be directly applied to IPs with more rounds because the proof length (and thus also the verifier running time) would be more than polynomial.

Nevertheless, we extend the transformation to work for any public-coin IP while achieving a moderate improvement on the number of read rounds. In the main loop of the IOP, rather than advancing each IP transcript in S_{i-1} by one round, advance it by $O(k/\log |\mathbf{x}|)$ rounds before inserting the resulting transcripts into the set S_i . During its decision phase, the IOP verifier chooses an index i and reads the entire $O(k/\log |\mathbf{x}|)$ -round interaction done during this iteration of the IOP. Completeness and soundness of this new transformation are similar to the one presented in the previous section, but now the verifier reads $O(k/\log |\mathbf{x}|)$ rounds. The rest of the efficiency parameters are similar to the IOP described in the previous section, except that proof length is polynomial regardless of k .

After applying the (adapted) transformation of [ACY22], this process yields a $O(k)$ -round IOP with query complexity $O(k/\log |\mathbb{x}|)$ (over the binary alphabet). This concludes the proof sketch of Lemma 1.

Can we do better? The construction described in this section doubles the number of transcripts stored in the set S_i relative to S_{i-1} . This causes a blow-up in parameters and is the reason why this approach fails in constructing $O(1)$ -query IOPs from IPs with super-logarithmic round complexity. Intuitively, if we could reduce this doubling then we may be able to modify the transformation to get $O(1)$ -query IOPs. While we do not achieve this for *general* IP, we show that, using this intuition, we can construct $O(1)$ -query IOPs for a rich class of relations that are *interactively reducible*. We discuss this notion, and corresponding new IOP constructions, in the following section.

2.3 Interactive reducibility

In Section 2.2.1 we described how to transform an IP into an IOP with $O(1)$ query complexity. The new protocol kept track of a set containing all partial transcripts of the IP generated so far in the protocol. In every round, every partial transcript in the set was advanced by one round, and the newly advanced transcripts were added to the set held previously. This meant that in every round, the set contains twice as many transcripts as in the previous round. As we require polynomial proof length and verifier running time, this technique was unable to allow reading of $O(1)$ rounds for IPs with greater than $O(\log |\mathbb{x}|)$ rounds.

Intuitively, suppose it were possible to take multiple transcripts and reduce them into a single transcript that in some sense preserves soundness of all of the transcripts combined. In that case, this issue could be bypassed, and the protocol would work for IPs with super-logarithmic round complexity. In more detail, suppose we want to reduce transcripts $\text{tr}_1, \dots, \text{tr}_t$ into a new transcript tr' . The *acceptance probability* of a transcript prefix tr_i is the maximum over all prover strategies of the probability that the verifier will end up accepting when continuing interaction with the prover from transcript tr_i . Roughly, we require that: (a) if the acceptance probability of every tr_i is 1, then so is the acceptance probability of tr' ; and (b) if there exists some transcript tr_i with small acceptance probability, then (with high probability) the acceptance probability of tr' is also small.

In this section, we introduce the concept of *interactive reducibility*, which formally captures this intuition. We exemplify this in Section 2.3.1 by describing how to reduce multiple transcripts of the sumcheck protocol into a single transcript. Then, in Section 2.3.2, we formally define interactive reducibility. In Section 2.3.3, we show how to adapt the protocol described in Section 2.2.1 to work with interactive reductions and bypass the blow-up in the original protocol. Finally, in Section 2.3.4, we discuss relations known to have interactive reducibility.

2.3.1 An interactive reduction for sumcheck

We exemplify the notion of interactive reducibility in the case of the sumcheck protocol [LFKN92]. Below we review this protocol, and then explain how to reduce multiple transcripts into one transcript via an interactive reduction.

The sumcheck protocol. The verifier has query access to a n -variate polynomial p of individual degree d over some field \mathbb{F} . The goal of the verifier is to test, for a given field element γ , whether

$$\sum_{\alpha_1, \dots, \alpha_n \in \{0,1\}} p(\alpha_1, \dots, \alpha_n) = \gamma .$$

The protocol begins with the prover sending a polynomial \tilde{p}_1 of degree d , claimed to equal $p_1(X) := \sum_{\alpha_2, \dots, \alpha_n \in \{0,1\}} p(X, \alpha_2, \dots, \alpha_n)$. The verifier checks that $\tilde{p}_1(0) + \tilde{p}_1(1) = \gamma$ (rejecting if not), samples a random field element r_1 , and sends it to the prover. Both parties define $\gamma_1 := \tilde{p}_1(r_1)$.

This one-round interaction leads to a new sumcheck claim

$$\sum_{\alpha_2, \dots, \alpha_n \in \{0,1\}} p(r_1, \alpha_2, \dots, \alpha_n) = \gamma_1 \text{ ,}$$

that has the following properties: (a) if the original claim is true then the new claim is also true; and (b) if the original claim is false then with high probability the new claim is also false.

Next, the prover sends \tilde{p}_2 claimed to equal $p_2(X) := \sum_{\alpha_3, \dots, \alpha_n \in \{0,1\}} p(r_1, X, \alpha_3, \dots, \alpha_n)$ and the protocol repeats as before. This process continues until the n variables are fixed to some field elements (r_1, \dots, r_n) , and the problem has been reduced to checking that $p(r_1, \dots, r_n) = \gamma_n$, which the verifier can check via one query to the polynomial p .

One can associate a round j of the protocol with a list of field elements (r_1, \dots, r_j) and a claimed sum γ_j , and think of that round as “reducing” a claim $z = ((r_1, \dots, r_j), \gamma_j)$ that $\sum_{\alpha_{j+1}, \dots, \alpha_n \in \{0,1\}} p(r_1, \dots, r_j, \alpha_{j+1}, \dots, \alpha_n) = \gamma_j$ into a new claim $z' = ((r_1, \dots, r_{j+1}), \gamma_{j+1})$ that $\sum_{\alpha_{j+2}, \dots, \alpha_n \in \{0,1\}} p(r_1, \dots, r_{j+1}, \alpha_{j+2}, \dots, \alpha_n) = \gamma_{j+1}$.

Reducing multiple sumcheck claims. We are given claims z_1, \dots, z_t where each z_i consists of $(r_{i,1}, \dots, r_{i,j})$ and a claimed sum $\gamma_{i,j}$. We seek to reduce these t claims into a single claim $z' = ((r'_1, \dots, r'_{j+1}), \gamma'_{j+1})$ such that: (a) If each z_i is a true statement, then z' is a true statement; and (b) If there is some z_i that is a false statement, then with high probability z' is a false statement. Notice that in order to merge multiple sumcheck transcripts it suffices to merge multiple sumcheck claims z_1, \dots, z_t . Thus we will focus on merging such claims.

Before describing the reduction, we define a polynomial $I_{z_1, \dots, z_t} : \mathbb{F} \rightarrow \mathbb{F}^j$ that represents a curve through the t points described by the instances z_1, \dots, z_t . That is, $I_{z_1, \dots, z_t}(i) = (r_{i,1}, \dots, r_{i,j})$ for every $i \in [t]$ (here we implicitly associate the set $[t]$ with an arbitrary set $S \subseteq \mathbb{F}$ of size t , known to all parties). By interpolation, the degree of I_{z_1, \dots, z_t} is less than t .

Given this definition, we describe the interactive reduction for the sumcheck protocol.

- Prover: Send the polynomial $g \in \mathbb{F}[X_1, X_2]$ defined as:

$$g(X_1, X_2) := \sum_{\alpha_{j+2}, \dots, \alpha_n \in \{0,1\}} p(I_{z_1, \dots, z_t}(X_1), X_2, \alpha_{j+2}, \dots, \alpha_n) \text{ .} \quad (1)$$

- Verifier: Receive a bivariate polynomial $\tilde{g} \in \mathbb{F}[X_1, X_2]$ of degree at most $j \cdot d \cdot (t - 1)$ in X_1 and degree at most d in X_2 .

1. *Consistency:* Check that for every $i \in [t]$ it holds that $\sum_{\alpha \in \{0,1\}} \tilde{g}(i, \alpha) = \gamma_{i,j}$. (Reject if not.)

2. *Generate new instance:*

- Sample uniformly random field elements $\rho, r^* \leftarrow \mathbb{F}$ and send them to the prover.
- Set $(r'_1, \dots, r'_j) := I_{z_1, \dots, z_t}(\rho)$ and $\gamma_{j+1} := \tilde{g}(\rho, r^*)$ and output the new instance

$$z' := \left((r'_1, \dots, r'_j, r^*), \gamma_{j+1} \right) \text{ .}$$

Analysis. It follows straightforwardly from the protocol that, if z_1, \dots, z_t are all true statements and the prover acts honestly, then z' is a true statement. We show that if any one of the statements z_1, \dots, z_t is false then with high probability so is z' .

Let g be as defined in Equation (1) with respect to z_1, \dots, z_t . Suppose that z_i is a false claim (i.e., $\sum_{\alpha_{j+1}, \dots, \alpha_n \in \{0,1\}} p(r_{i,1}, \dots, r_{i,j}, \alpha_{j+1}, \dots, \alpha_n) \neq \gamma_{i,j}$). Then, by definition,

$$\sum_{\alpha \in \{0,1\}} g(i, \alpha) = \sum_{\alpha_{j+1}, \dots, \alpha_n \in \{0,1\}} p(r_{i,1}, \dots, r_{i,j}, \alpha_{j+1}, \dots, \alpha_n) \neq \gamma_{i,j} .$$

During its consistency check, the verifier checks that $\sum_{\alpha \in \{0,1\}} \tilde{g}(i, \alpha) = \gamma_{i,j}$. Thus, in order for the verifier to not reject, a cheating prover must send $\tilde{g} \neq g$. By the Schwartz-Zippel lemma, since g and \tilde{g} are low-degree polynomials (provided that the degree d and the number of instances being reduced t are small with respect to $|\mathbb{F}|$), the probability that the uniformly chosen ρ and r^* are such that $\tilde{g}(\rho, r^*) = g(\rho, r^*)$ is small. Whenever $\tilde{g}(\rho, r^*) \neq g(\rho, r^*)$ we have that

$$\sum_{\alpha_{j+2}, \dots, \alpha_n \in \{0,1\}} p(r'_1, \dots, r'_j, r^*, \alpha_{j+2}, \dots, \alpha_n) = g(\rho, r^*) \neq \tilde{g}(\rho, r^*) = \gamma_{j+1} ,$$

and so the resulting statement $z' := ((r'_1, \dots, r'_j, r^*), \gamma_{j+1})$ is false.

2.3.2 Defining interactive reducibility

We define interactive reducibility, which captures the capability of merging multiple transcripts/instances into a single transcript/instance while preserving correctness and soundness.

Definition 1. *An ℓ -round public-coin protocol $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$ where \mathbf{V}_{IR} runs in polynomial time is an **interactive reduction** for a relation R with k predicates and soundness error ε if there exists a sequence of predicates f_0, f_1, \dots, f_k such that the following holds.*

- **Completeness:** *For every $(\mathbf{x}, \mathbf{w}) \in R$, $j \in [k]$, and z_1, \dots, z_t such that $f_{j-1}(\mathbf{x}, z_i) = 1$ for every $i \in [t]$, it holds that:*

$$\Pr [f_j(\mathbf{x}, z') = 1 \mid z' \leftarrow \langle \mathbf{P}_{\text{IR}}(\mathbf{x}, \mathbf{w}, z_1, \dots, z_t), \mathbf{V}_{\text{IR}}(\mathbf{x}, z_1, \dots, z_t) \rangle] = 1 .$$

- **Soundness:** *For every $\mathbf{x} \notin L(R)$, $j \in [k]$, and z_1, \dots, z_t , if there exists $i \in [t]$ where $f_{j-1}(\mathbf{x}, z_i) = 0$ then for every (computationally unbounded) $\tilde{\mathbf{P}}_{\text{IR}}$ it holds that:*

$$\Pr [f_j(\mathbf{x}, z') = 1 \mid z' \leftarrow \langle \tilde{\mathbf{P}}_{\text{IR}}, \mathbf{V}_{\text{IR}}(\mathbf{x}, z_1, \dots, z_t) \rangle] \leq \varepsilon(\mathbf{x}, t) .$$

- **Relation identity:** $f_0(\mathbf{x}, z) = 1$ if and only if $\mathbf{x} \in L(R)$.
- **Triviality:** $f_k(\mathbf{x}, z)$ can be computed in time $\text{poly}(|\mathbf{x}|, |z|)$.

We call \mathbf{x} the base instance and z_1, \dots, z_t round instances.

An interactive reduction $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$ has (polynomially) **bounded output length** if there exists $c \in \mathbb{N}$ such that, for every base instance \mathbf{x} , witness \mathbf{w} , and round instances z_1, \dots, z_t , the new instance z' output by $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$ on inputs $(\mathbf{x}, \mathbf{w}, z_1, \dots, z_t)$ has length at most $|\mathbf{x}|^c$.

2.3.3 IOPs from interactive reducibility

We show that any relation with an ℓ -round interactive reduction with k predicates (and bounded output length) has a $(\ell \cdot k)$ -round public-coin IOP with query complexity $O(k)$. This is a variation of the protocol described in Section 2.2, adapted to work with interactive reducibility. For simplicity, in this overview, we present the protocol only for the case $\ell = 1$ (such as the sumcheck protocol).

To aid with notation, in the description of the protocol, we replace the set S_i (which in Section 2.2.1 contained the set of all transcripts generated up until the i -th iteration) with an array A_i where $A_i[j]$ contains all of the round instances generated in the i -th iteration that are associated with the j -th predicate of the interactive reduction. In iteration i , the interactive reduction will be run k times in parallel, where for every $j \in [k]$ we run given the round instances stored in $A_{i-1}[j]$.

The protocol. Let $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$ be a one-round interactive reduction for R with k predicates. The IOP prover \mathbf{P}_{IOP} receives as input an instance \mathfrak{x} and witness \mathfrak{w} , and the IOP verifier \mathbf{V}_{IOP} receives as input the instance \mathfrak{x} . They interact as follows.

1. For every $i \in \{0, \dots, 2k\}$, \mathbf{P}_{IOP} defines the $(k+1)$ -entry array A_i as follows

$$A_i[j] := \begin{cases} \{\perp\} & \text{if } j = 0 \\ \emptyset & \text{if } j \in \{1, \dots, k\} \end{cases} .$$

The set $A_i[j]$ will store all instances corresponding to f_j collected by iteration i of the protocol.

2. For $i = 1, \dots, 2k$:
 - (a) \mathbf{P}_{IOP} sends A_{i-1} to \mathbf{V}_{IOP} .
 - (b) \mathbf{V}_{IOP} sends a random $\rho_i \leftarrow \{0, 1\}^r$ (this corresponds to a message of \mathbf{V}_{IR}).
 - (c) \mathbf{P}_{IOP} sends $a_{i,j} := \mathbf{P}_{\text{IR}}(\mathfrak{x}, \mathfrak{w}, A_{i-1}[j-1], \rho_i)$ for all $j \in [k]$, and sets $A_i[j] := A_{i-1}[j] \cup \{z_{i,j}\}$ where $z_{i,j} := \mathbf{V}_{\text{IR}}(\mathfrak{x}, A_{i-1}[j-1], \rho_i, a_{i,j})$ is the output of the interactive reduction verifier given base instance \mathfrak{x} , round instances $A_{i-1}[j-1]$, verifier randomness ρ_i , and prover reply $a_{i,j}$.
3. \mathbf{P}_{IOP} sends A_{2k} and, for every $i \in \{0, \dots, 2k\}$, sends $B_i := A_i$. This concludes the interaction.
4. In the decision phase, \mathbf{V}_{IOP} is given oracle access to a transcript with the following structure:

$$(A_0, \rho_1, (a_{1,1}, \dots, a_{1,k}), A_1, \dots, \rho_{2k}, (a_{2k,1}, \dots, a_{2k,k}), A_{2k}, (B_0, \dots, B_{2k})) .$$

\mathbf{V}_{IOP} performs the checks below.

- (a) *Subset consistency.* Read the arrays B_0, B_1, \dots, B_{2k} in their entirety. For every $i \in [2k]$ and $j \in \{0, \dots, k\}$ check that $B_{i-1}[j] \subseteq B_i[j]$.
- (b) *Transcript consistency.* Sample a random $i \in [2\ell]$. Read the arrays A_{i-1} and A_i sent by \mathbf{P}_{IOP} and the interaction ρ_i and $(a_{i,1}, \dots, a_{i,\ell})$.
 - i. Check that $A_{i-1} = B_{i-1}$ and $A_i = B_i$.
 - ii. For every $j \in [k]$, check that $A_i[j] = A_{i-1}[j] \cup \{z'_{i,j}\}$ where

$$z'_{i,j} := \mathbf{V}_{\text{IR}}(\mathfrak{x}, A_{i-1}[j-1], \rho_i, a_{i,j}) ,$$

is the output of the interactive reduction verifier given base instance \mathfrak{x} , round instances $A_{i-1}[j-1]$, verifier randomness ρ_i , and prover reply $a_{i,j}$. (Reject if \mathbf{V}_{IR} rejects.)

- (c) *Final predicate holds.* Check that $f_k(\mathfrak{x}, z) = 1$ for every $z \in B_{2k}[k]$.

Analysis. The protocol has perfect completeness and soundness error $\max\{\frac{1}{2}, \binom{2k}{k} \cdot k \cdot \epsilon\}$ where k is the number of predicates and ϵ is the soundness of the interactive reduction respectively. This can be shown in a similar manner to that described in Section 2.2.1. The main difference between the two protocols is in the analysis of the proof length. In the protocol of Section 2.2.1 the number of sent transcripts doubled in each round. In contrast, in the new protocol, one round instance is added for each predicate. In more detail, for every $i \in [2k]$ and $j \in [k]$, we have $|A_i[j]| = |A_{i-1}[j]| + 1$. Since $k = \text{poly}(|\mathbb{x}|)$, the total number of round instances generated and sent is polynomial in $|\mathbb{x}|$. If the interactive reduction has bounded output length, then each of these round instances has polynomially-bounded length. We can therefore conclude that the overall proof length is $\text{poly}(|\mathbb{x}|)$.

2.3.4 Relations with interactive reducibility

Several relations of interest have interactive reductions.

General IPs. We show that any relation with a k -round interactive proof has an m -round interactive reduction with k/m predicates (for any m that divides k). To see this, consider round instances z that are sets of j -message partial transcripts of the IP. The interactive reduction advances each of the transcripts in the set z by m rounds, as in the IP. The predicates are defined with respect to the “state function” of the IP, which roughly denotes whether the prover has an accepting strategy with respect to this transcript or whether no strategy will cause the verifier to accept with high probability (over the remaining interaction). See Section 3.2 for a formal definition of the state function of an IP through the concept of round-by-round soundness.

Notice that if this interactive reduction is used in the protocol of Section 2.3.3, this yields the protocol described in Section 2.2.1. This interactive reduction does not have bounded output length since the new round instance stores all of the previous transcripts and their continuations. Therefore the resulting IOP does not achieve $O(1)$ total query complexity.

Sumcheck protocol. Using the ideas described in Section 2.3.2, we show that any relation that can be reduced into k -variate sumcheck has a one-round interactive reduction with k predicates and *bounded output length*. As a result, any relation that can be reduced into k -variate sumcheck has a $O(k)$ -round public-coin IOP with query complexity $O(1)$.

Shamir’s protocol. Shamir’s protocol [Sha92] gives an IP for all of PSPACE, thereby showing the $\text{IP} = \text{PSPACE}$ theorem. Extending the ideas developed in Section 2.3.2, we show a one-round interactive reduction with bounded output length and polynomially-many predicates for Shamir’s protocol. This establishes that every language in PSPACE has a $\text{poly}(|\mathbb{x}|)$ -round public-coin IOP with query complexity $O(1)$.

Future directions. We leave the exploration of what other relations have interactive reductions to future work. Following the extensive use of polynomials in both the sumcheck protocol and Shamir’s protocol, it seems likely that these techniques can be adapted to also work for low-depth circuits through the delegation protocol in [GKR15].

3 Preliminaries

We consider proof systems for *binary* relations. A binary relation R is a set of tuples (\mathbf{x}, \mathbf{w}) where \mathbf{x} is the instance and \mathbf{w} the witness. The corresponding language $L(R)$ is the set of \mathbf{x} for which there exists \mathbf{w} such that $(\mathbf{x}, \mathbf{w}) \in R$. In some theorems we highlight important parameters with a yellow background .

3.1 Interactive oracle proofs

Interactive Oracle Proofs (IOPs) [BCS16; RRR16] are information-theoretic proof systems that combine aspects of Interactive Proofs [Bab85; GMR89] and Probabilistically Checkable Proofs [BFLS91; FGLSS91; AS98; ALMSS98], and also generalize the notion of Interactive PCPs [KR08]. Below we describe *public-coin* IOPs.

A k_{IOP} -round public-coin IOP works as follows. For each round $i \in [k_{\text{IOP}}]$, the verifier sends a uniformly random message ρ_i to the prover; then the prover sends a proof string Π_i to the verifier. After k_{IOP} rounds of interaction, the verifier makes some queries to the proof strings $\Pi_1, \dots, \Pi_{k_{\text{IOP}}}$ sent by the prover, and then decides if to accept or to reject.

In more detail, let $\text{IOP} = (\mathbf{P}_{\text{IOP}}, \mathbf{V}_{\text{IOP}})$ be a tuple where \mathbf{P}_{IOP} (the prover) is an interactive algorithm, and \mathbf{V}_{IOP} (the verifier) is an interactive oracle algorithm. We say that IOP is a *public-coin IOP* for a binary relation R with k_{IOP} rounds and soundness error β_{IOP} if the following holds.

- **Completeness.** For every $(\mathbf{x}, \mathbf{w}) \in R$,

$$\Pr_{\rho_1, \dots, \rho_{k_{\text{IOP}}}, \rho_{\text{dc}}} \left[\mathbf{V}_{\text{IOP}}^{\Pi_1, \dots, \Pi_{k_{\text{IOP}}}, \rho_1, \dots, \rho_{k_{\text{IOP}}}}(\mathbf{x}; \rho_{\text{dc}}) = 1 \mid \begin{array}{c} \Pi_1 \leftarrow \mathbf{P}_{\text{IOP}}(\mathbf{x}, \mathbf{w}, \rho_1) \\ \vdots \\ \Pi_{k_{\text{IOP}}} \leftarrow \mathbf{P}_{\text{IOP}}(\mathbf{x}, \mathbf{w}, \rho_1, \dots, \rho_{k_{\text{IOP}}}) \end{array} \right] = 1 .$$

- **Soundness.** For every $\mathbf{x} \notin L(R)$ and unbounded malicious prover $\tilde{\mathbf{P}}_{\text{IOP}}$,

$$\Pr_{\rho_1, \dots, \rho_{k_{\text{IOP}}}, \rho_{\text{dc}}} \left[\mathbf{V}_{\text{IOP}}^{\tilde{\Pi}_1, \dots, \tilde{\Pi}_{k_{\text{IOP}}}, \rho_1, \dots, \rho_{k_{\text{IOP}}}}(\mathbf{x}; \rho_{\text{dc}}) = 1 \mid \begin{array}{c} \tilde{\Pi}_1 \leftarrow \tilde{\mathbf{P}}_{\text{IOP}}(\rho_1) \\ \vdots \\ \tilde{\Pi}_{k_{\text{IOP}}} \leftarrow \tilde{\mathbf{P}}_{\text{IOP}}(\rho_1, \dots, \rho_{k_{\text{IOP}}}) \end{array} \right] \leq \beta_{\text{IOP}}(|\mathbf{x}|) .$$

Complexity measures. We consider several complexity measures beyond soundness error. All of these complexity measures are implicitly functions of the instance \mathbf{x} .

- *proof length* l_{IOP} : the total number of bits in $\Pi_1, \dots, \Pi_{k_{\text{IOP}}}$.
- *queries* q_{IOP} : the number of bits read by the verifier from $\rho_1, \Pi_1, \dots, \rho_{k_{\text{IOP}}}, \Pi_{k_{\text{IOP}}}$.
- *interaction randomness length* r_{int} : the total number of bits in $\rho_1, \dots, \rho_{k_{\text{IOP}}}$.
- *decision randomness length* r_{dc} : The number of bits in ρ_{dc} .
- *prover time* pt_{IOP} : \mathbf{P}_{IOP} runs in time pt_{IOP} .
- *verifier time* vt_{IOP} : \mathbf{V}_{IOP} runs in time vt_{IOP} .
- *decision complexity* dt_{IOP} : Following the choice of queries, \mathbf{V}_{IOP} runs in time dt_{IOP} to decide whether to accept or reject.

Round-query IOPs and IPs. A round-query IOP with round-query complexity q_{rnd} is an IOP in which the verifier reads every symbol of q_{rnd} rounds from the interaction. An interactive proof (IP) is

a round-query IOP with round-query complexity equal to the number of rounds in the protocol (i.e., the verifier reads every symbol in the interaction). Unless explicitly stated otherwise, we assume that IPs have no decision randomness.

3.2 Round-by-round soundness

Definition 3.1. Let $(\mathbf{P}_{\text{IOP}}, \mathbf{V}_{\text{IOP}})$ be an IOP for a relation R . A **state function** for $(\mathbf{P}_{\text{IOP}}, \mathbf{V}_{\text{IOP}})$ with decision error δ_{dc} is a (possibly inefficient) Boolean function that receives as input an instance \mathbb{x} and a transcript tr and outputs a bit for which the following holds.

- *Empty transcript:* if $\text{tr} = \emptyset$ is the empty transcript then $\text{state}(\mathbb{x}, \text{tr}) = 0$ if and only if $\mathbb{x} \notin L(R)$.
- *Prover moves:* If $\mathbb{x} \notin L(R)$, tr is a transcript where the prover is about to move and $\text{state}(\mathbb{x}, \text{tr}) = 0$, then for any potential prover message a , $\text{state}(\mathbb{x}, \text{tr}||a) = 0$.
- *Full transcript:* if tr is a full transcript and $\text{state}(\mathbb{x}, \text{tr}) = 0$ then $\Pr_{\rho_{\text{dc}}}[\mathbf{V}_{\text{IOP}}^{\text{tr}}(\mathbb{x}; \rho_{\text{dc}}) = 1] \leq \delta_{\text{dc}}$.

Definition 3.2. An IOP $(\mathbf{P}_{\text{IOP}}, \mathbf{V}_{\text{IOP}})$ with k_{IOP} rounds for a relation R has $(\beta_{\text{rbr}}, \delta_{\text{dc}})$ -**round-by-round soundness** if there exists a state function state with decision error δ_{dc} such that for all $\mathbb{x} \notin L(R)$, every $i \in [k_{\text{IOP}}]$, and every transcript tr of the first i rounds where the verifier is about to move and $\text{state}(\mathbb{x}, \text{tr}) = 0$ it holds that

$$\Pr_{\rho}[\text{state}(\mathbb{x}, \text{tr}||\rho) = 1] \leq \beta_{\text{rbr}} .$$

We call β_{rbr} the interaction error and δ_{dc} the decision error. If $\delta_{\text{dc}} = 0$, we omit δ_{dc} and simply say that IOP has round-by-round soundness β_{rbr} which aligns with the standard definition of round-by-round soundness for IPs.

Notice that if a k -round IOP has $(\beta_{\text{rbr}}, \delta_{\text{dc}})$ -round-by-round soundness, then it has soundness error at most $k \cdot \beta_{\text{rbr}} + \delta_{\text{dc}}$. Additionally, if a k -round IOP has constant soundness error, then it has $(c^{1/k}, O(1))$ -round-by-round soundness error for some constant $0 < c < 1$.

The following lemma is a generalization of [CCHLRR18], Corollary 5.7, and follows from their techniques:

Lemma 3.3. Let $m, t \in \mathbb{N}$ be parameters and R be a relation with an IOP $(\mathbf{P}_{\text{IOP}}, \mathbf{V}_{\text{IOP}})$ with $(\beta_{\text{rbr}}, \delta_{\text{dc}})$ -round-by-round soundness. Then R has an IOP $(\mathbf{P}'_{\text{IOP}}, \mathbf{V}'_{\text{IOP}})$ with $(\beta_{\text{rbr}}^m, \delta_{\text{dc}}^t)$ -round-by-round soundness and the following parameters:

IOP $(\mathbf{P}_{\text{IOP}}, \mathbf{V}_{\text{IOP}})$ for R			IOP $(\mathbf{P}'_{\text{IOP}}, \mathbf{V}'_{\text{IOP}})$ for R	
Rounds	k_{IOP}		Rounds	k_{IOP}
Alphabet size	λ_{IOP}		Alphabet size	λ_{IOP}^m
Proof length	l_{IOP}		Proof length	l_{IOP}
Round queries	q_{rnd}	→	Round queries	$t \cdot q_{\text{rnd}}$
Queries per round	q_{IOP}		Queries per round	$t \cdot q_{\text{IOP}}$
Total interaction randomness	r_{int}		Total interaction randomness	$m \cdot r_{\text{int}}$
Decision randomness	r_{dc}		Decision randomness	$t \cdot r_{\text{dc}}$
Round-by-round soundness	$(\beta_{\text{rbr}}, \delta_{\text{dc}})$		Round-by-round soundness	$(\beta_{\text{rbr}}^m, \delta_{\text{dc}}^t)$
Verifier running time	vt_{IOP}		Verifier running time	$m \cdot t \cdot \text{vt}_{\text{IOP}}$

3.3 Extractors

Definition 3.4. The **min-entropy** of a random variable X is

$$H_{\min}(X) = \min_{x \in \text{supp}(X)} -\log \Pr[X = x]$$

Definition 3.5. A function $\text{Ext}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (k, ε) -**extractor** if for every X with min-entropy at least k , $\text{SD}(\text{Ext}(X, U_d), U_m) \leq \varepsilon$ (where SD is the statistical distance). An extractor is **explicit** if it is computable in polynomial time.

We use the following explicit construction of extractors with tight parameters.

Theorem 3.6 ([GUV09]). For every constant $\alpha > 0$, and all positive integers n, k and all $\varepsilon > 0$, there is an explicit construction of a (k, ε) -extractor $\text{Ext}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $d = O(\log n + \log(1/\varepsilon))$, and $m \geq (1 - \alpha)k$.

Setting specific parameters, we will use this simpler version of the theorem.

Theorem 3.7. For all positive integers m , and $\ell \geq \log m$ there is an explicit construction of a $(2m, 2^{-\ell})$ -extractor $\text{Ext}: \{0, 1\}^{3m} \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $d = O(\ell)$.

Fact 3.8. For all $n \in \mathbb{N}$, all $x \in \{0, 1\}^n$ and $0 < \gamma < 1$ we have that

$$|\{x' \in \{0, 1\}^n : \Delta(x, x') \leq \gamma\}| \leq 2^{n \cdot H(\gamma)} .$$

(here H is the entropy function $H(p) = -p \log(p) - (1 - p) \log(1 - p)$).

4 Interactive reducibility

We define the notion of interactive reducibility. Then we show basic properties of interactive reducibility: (a) in Section 4.1 we show that if a relation has an interactive reduction then it also has an IP; and (b) in Section 4.2 we show how to reduce the soundness error of an interactive reduction. Finally, in Section 4.3, we show an interactive reduction for any public-coin IP and bounded-output-length interactive reductions for the sumcheck protocol and Shamir's protocol.

Definition 4.1. An **interactive reduction** for a relation R with k_{IR} predicates and soundness error ε_{IR} is a ℓ_{IR} -round public-coin protocol $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$ (where \mathbf{V}_{IR} runs in polynomial time) for which there exists a list of predicates $f_0, f_1, \dots, f_{k_{\text{IR}}}$ such that the following holds.

- **Completeness.** For every $(\mathbb{x}, \mathbb{w}) \in R$, $j \in [k_{\text{IR}}]$, and z_1, \dots, z_t such that $f_{j-1}(\mathbb{x}, z_i) = 1$ for every $i \in [t]$, it holds that

$$\Pr [f_j(\mathbb{x}, z') = 1 \mid z' \leftarrow \langle \mathbf{P}_{\text{IR}}(\mathbb{x}, \mathbb{w}, z_1, \dots, z_t), \mathbf{V}_{\text{IR}}(\mathbb{x}, z_1, \dots, z_t) \rangle] = 1 .$$

- **Soundness.** For every $\mathbb{x} \notin L(R)$, $j \in [k_{\text{IR}}]$, and z_1, \dots, z_t , if there exists $i \in [t]$ such that $f_{j-1}(\mathbb{x}, z_i) = 0$ then for every (computationally unbounded) $\tilde{\mathbf{P}}_{\text{IR}}$ it holds that

$$\Pr [f_j(\mathbb{x}, z') = 1 \mid z' \leftarrow \langle \tilde{\mathbf{P}}_{\text{IR}}, \mathbf{V}_{\text{IR}}(\mathbb{x}, z_1, \dots, z_t) \rangle] \leq \varepsilon_{\text{IR}}(\mathbb{x}, t) .$$

- **Relation identity.** $f_0(\mathbb{x}, z) = 1$ if and only if $\mathbb{x} \in L(R)$.
- **Triviality.** $f_{k_{\text{IR}}}(\mathbb{x}, z)$ can be computed in time $\text{poly}(|\mathbb{x}|, |z|)$.

We call \mathbb{x} the base instance and z_1, \dots, z_t round instances.

A relation R is *interactively reducible* with k_{IR} predicates and ℓ_{IR} rounds if R has an interactive reduction with k_{IR} predicates and k_{IR} -rounds. An interactive reduction has soundness error that is *well-behaved* if ε_{IR} is (weakly) monotonically increasing with t and grows at most polynomially with t . All interactive reductions described in this paper have well-behaved soundness error.

Definition 4.2. An interactive reduction has (polynomially) **bounded output length** if there exists $c \in \mathbb{N}$ such that, for every base instance \mathbb{x} , witness \mathbb{w} , and round instances z_1, \dots, z_t , the new instance output by $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$ on inputs $(\mathbb{x}, \mathbb{w}, z_1, \dots, z_t)$ has length at most $|\mathbb{x}|^c$.

4.1 Interactive proofs from interactive reductions

We construct an IP for any relation that is interactively reducible.

Theorem 4.3. If a relation R has an interactive reduction $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$ then R has an IP $(\mathbf{P}_{\text{IP}}, \mathbf{V}_{\text{IP}})$ with the parameters indicated below.

Interactive reduction for R			IP for R	
Number of predicates	k_{IR}	→	Rounds	$k_{\text{IR}} \cdot \ell_{\text{IR}}$
Rounds	ℓ_{IR}		Communication	$k_{\text{IR}} \cdot l_{\text{IR}}$
Communication	l_{IR}		Randomness	$k_{\text{IR}} \cdot r_{\text{IR}}$
Randomness	r_{IR}		Soundness error	$k_{\text{IR}} \cdot \varepsilon_{\text{IR}}(\mathbb{x}, 1)$
Soundness error	ε_{IR}		Verifier running time	$k_{\text{IR}} \cdot \mathbf{vt}_{\text{IR}} + \mathbf{ft}_{\text{IR}}$
Verifier running time	\mathbf{vt}_{IR}			
Final predicate time	\mathbf{ft}_{IR}			

Moreover, if $\ell_{\text{IR}} = 1$ then the IP for R has round-by-round soundness error $\varepsilon_{\text{IR}}(\mathbf{x}, 1)$.

Construction 4.4. The IP prover \mathbf{P}_{IP} receives as input an instance \mathbf{x} and a witness \mathbf{w} , and the IP verifier \mathbf{V}_{IP} receives as input the instance \mathbf{x} . Letting $z_0 := \mathbf{x}$, they interact as follows. For $j = 1, \dots, k_{\text{IR}}$, \mathbf{P}_{IP} and \mathbf{V}_{IP} run the interactive reduction $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$ to obtain a new instance

$$z_j \leftarrow \langle \mathbf{P}_{\text{IR}}(\mathbf{x}, \mathbf{w}, z_{j-1}), \mathbf{V}_{\text{IR}}(\mathbf{x}, z_{j-1}) \rangle .$$

Finally, after the interaction, \mathbf{V}_{IP} checks that $f_{k_{\text{IR}}}(\mathbf{x}, z_{k_{\text{IR}}}) = 1$.

Proof of Theorem 4.3. We prove completeness, then prove soundness (and round-by-round soundness), and finally analyze complexity measures.

Completeness. Fix $(\mathbf{x}, \mathbf{w}) \in R$. By completeness of $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$, for every $j \in [k_{\text{IR}}]$, if $f_{j-1}(\mathbf{x}, z_{j-1}) = 1$, then

$$\Pr [f_j(\mathbf{x}, z_j) = 1 \mid z_j \leftarrow \langle \mathbf{P}_{\text{IR}}(\mathbf{x}, \mathbf{w}, z_{j-1}), \mathbf{V}_{\text{IR}}(\mathbf{x}, z_{j-1}) \rangle] = 1 .$$

Since $f_0(\mathbf{x}, z_0) = f_0(\mathbf{x}, \mathbf{x}) = 1$, by induction it holds that $f_{k_{\text{IR}}}(\mathbf{x}, z_{k_{\text{IR}}}) = 1$ with probability 1. We conclude that $\Pr[\langle \mathbf{P}_{\text{IP}}(\mathbf{x}, \mathbf{w}), \mathbf{V}_{\text{IP}}(\mathbf{x}) \rangle = 1] = 1$.

Soundness. Fix $\mathbf{x} \notin L(R)$ and a malicious IP prover $\tilde{\mathbf{P}}_{\text{IP}}$ and let $\varepsilon_{\text{IR}} := \varepsilon_{\text{IR}}(\mathbf{x}, 1)$. Let $\tilde{\mathbf{P}}_{\text{IR}}^{(j)}$ be the machine that describes $\tilde{\mathbf{P}}_{\text{IP}}$ during the j -th execution of the interactive reduction. Each $\tilde{\mathbf{P}}_{\text{IR}}^{(j)}$ can be seen as a (potentially malicious) prover for the interactive reduction. We show that $f_{k_{\text{IR}}}(\mathbf{x}, z_{k_{\text{IR}}}) = 1$ (a necessary condition for the IP verifier \mathbf{V}_{IP} to accept) with probability at most $k_{\text{IR}} \cdot \varepsilon_{\text{IR}}$.

We show by induction that $f_j(\mathbf{x}, z_j) = 1$ with probability at most $j \cdot \varepsilon_{\text{IR}}$. For $j = 0$ this holds because $f_0(\mathbf{x}, z_0) = f_0(\mathbf{x}, \mathbf{x}) = 0$. For $j > 0$, suppose that $f_{j-1}(\mathbf{x}, z_{j-1}) = 1$ with probability at most $(j-1) \cdot \varepsilon_{\text{IR}}$. It may be that whenever $f_{j-1}(\mathbf{x}, z_{j-1}) = 1$ it holds that $f_j(\mathbf{x}, z_j) = 1$. However, whenever $f_{j-1}(\mathbf{x}, z_{j-1}) = 0$, we can invoke the soundness property of $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$ to deduce that

$$\Pr \left[f_j(\mathbf{x}, z_j) = 1 \mid z_j \leftarrow \langle \tilde{\mathbf{P}}_{\text{IR}}^{(j)}, \mathbf{V}_{\text{IR}}(\mathbf{x}, z_{j-1}) \rangle \right] \leq \varepsilon_{\text{IR}} .$$

By applying the union bound, the probability that $f_j(\mathbf{x}, z_j) = 1$ is at most $(j-1) \cdot \varepsilon_{\text{IR}} + \varepsilon_{\text{IR}} = j \cdot \varepsilon_{\text{IR}}$.

Round-by-round soundness. Suppose that $\ell_{\text{IR}} = 1$ (that is, $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$ has one round) and let $\varepsilon_{\text{IR}} := \varepsilon_{\text{IR}}(\mathbf{x}, 1)$. We focus on the MA case (the prover moves first in the interactive reduction); the AM case can be shown in a similar manner. To show round-by-round soundness we first define a state function **state**. Let \mathbf{x} be an instance and \mathbf{tr} be a transcript of the protocol up to round j (i.e., j full executions of the interactive reduction), and let (z_0, \dots, z_j) be the round instances output by the reduction verifier \mathbf{V}_{IR} in an execution of the protocol according to \mathbf{tr} (if the transcript ends with a prover message, then z_j is the last round instance specified according to the transcript). Then we set $\mathbf{state}(\mathbf{x}, \mathbf{tr}) := 1$ if and only if $f_j(\mathbf{x}, z_j) = 1$. We show that this is indeed a state function.

- *Empty transcript.* If $\mathbf{tr} = \emptyset$ is the empty transcript then by definition $\mathbf{state}(\mathbf{x}, \mathbf{tr}) = 0$ if and only if $f_0(\mathbf{x}, \emptyset) = 0$, which occurs if and only if $\mathbf{x} \notin L(R)$.
- *Prover moves.* Given $\mathbf{x} \notin L(R)$ and a transcript \mathbf{tr} . As in the definition of the state function, let z_0, \dots, z_j be the instances generated by \mathbf{V}_{IR} . If $\mathbf{state}(\mathbf{x}, \mathbf{tr}) = 0$ then by definition $f_j(\mathbf{x}, z_j) = 0$. Since no new instance is generated given only the prover's message, $\mathbf{state}(\mathbf{x}, \mathbf{tr}||a) = 0$ for every a .
- *Full transcript.* If \mathbf{tr} is a full transcript and $\mathbf{state}(\mathbf{x}, \mathbf{tr}) = 0$, this implies that $f_{k_{\text{IR}}}(\mathbf{x}, z_{k_{\text{IR}}}) = 0$. If this is the case then the IP verifier \mathbf{V}_{IP} rejects.

We argue that $(\mathbf{P}_{\text{IP}}, \mathbf{V}_{\text{IP}})$ has round-by-round soundness ε_{IR} with respect to this state function. Fix an instance \mathbb{x} , a round number $j \in [k_{\text{IR}}]$, and a transcript tr that contains messages up to round j (where the next message is a verifier message). Let (z_0, \dots, z_j) be the instances implied by tr . Suppose that $\text{state}(\mathbb{x}, \text{tr}) = 0$. Since the next message is a verifier message, and the interactive reduction is public-coin, we have that the verifier's next message ρ is a uniformly random string. Thus, we show that

$$\Pr_{\rho} [\text{state}(\mathbb{x}, \text{tr} \parallel \rho) = 1] \leq \varepsilon_{\text{IR}} .$$

Since $\text{state}(\text{tr}) = 0$, we have that $f_j(\mathbb{x}, z_j) = 0$. The claim, then, is implied by soundness of $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$, stating that for any malicious prover $\tilde{\mathbf{P}}_{\text{IR}}$:

$$\Pr [f_{j+1}(\mathbb{x}, z_{j+1}) = 1 \mid z_{j+1} \leftarrow \langle \tilde{\mathbf{P}}_{\text{IR}}, \mathbf{V}_{\text{IR}}(\mathbb{x}, z_j) \rangle] \leq \varepsilon_{\text{IR}} .$$

Complexity measures. We discuss complexity measures of $(\mathbf{P}_{\text{IP}}, \mathbf{V}_{\text{IP}})$.

- *Round complexity.* The round complexity of the IP is $k_{\text{IR}} \cdot \ell_{\text{IR}}$ because the IP consists of k_{IR} interactive reductions, each with ℓ_{IR} rounds, ran in sequence.
- *Communication complexity.* The communication complexity of the IP is $k_{\text{IR}} \cdot l_{\text{IR}}$, because each of the k_{IR} interactive reductions has communication complexity l_{IR} .
- *Randomness complexity.* The IP verifier uses $k_{\text{IR}} \cdot r_{\text{IR}}$ random bits because each of the k_{IR} invocations of the reduction verifier uses r_{IR} random bits.
- *Verifier running time.* The IP verifier runs in time $k_{\text{IR}} \cdot vt_{\text{IR}} + ft_{\text{IR}}$ because it runs k_{IR} invocations of the reduction verifier (whose time complexity is vt_{IR}) and also runs the predicate $f_{k_{\text{IR}}}$ (whose time complexity is ft_{IR}).

□

4.2 Error reduction for interactive reducibility

We show that the soundness error of an interactive reduction can be amplified. In order to improve soundness from ε_{IR} to $\varepsilon_{\text{IR}}^{\tau}$ we have each “new” round instance contain τ “original” round instances for the original protocol. Then, when reducing t separate new round instances (each one containing τ old round instances), we run the original interactive reduction on all of the old round instances contained within the new round instances.

Theorem 4.5. *Let R be a relation and $\tau \in \mathbb{N}$ a parameter. If R has an interactive reduction $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$ then R has an interactive reduction $(\mathbf{P}_{\text{IR}}^{(\tau)}, \mathbf{V}_{\text{IR}}^{(\tau)})$ with the parameters indicated below.*

Interactive reduction for R			Interactive reduction for R with reduced error	
Number of predicates	k_{IR}		Number of predicates	k_{IR}
Rounds	ℓ_{IR}		Rounds	$\ell_{\text{IR}}(\mathbb{x}, \tau \cdot t)$
Output length	s_{IR}		Output length	$\tau \cdot s_{\text{IR}}(\mathbb{x}, \tau \cdot t)$
Communication	l_{IR}	→	Communication	$\tau \cdot l_{\text{IR}}(\mathbb{x}, \tau \cdot t)$
Randomness	r_{IR}		Randomness	$\tau \cdot r_{\text{IR}}(\mathbb{x}, \tau \cdot t)$
Soundness error	ε_{IR}		Soundness error	$\varepsilon_{\text{IR}}^{\tau}(\mathbb{x}, \tau \cdot t)$
Verifier running time	vt_{IR}		Verifier running time	$\tau \cdot vt_{\text{IR}}(\mathbb{x}, \tau \cdot t)$
Final predicate time	ft_{IR}		Final predicate time	$\tau \cdot ft_{\text{IR}}(\mathbb{x}, \tau \cdot t)$

Construction 4.6. We define the new predicates and then describe the new interactive reduction.

The predicates. Let f_j be the j -th predicate in $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$. Define the new predicate $f_j^{(\tau)}$ for $(\mathbf{P}_{\text{IR}}^{(\tau)}, \mathbf{V}_{\text{IR}}^{(\tau)})$ so that $f_j^{(\tau)}(\mathbf{x}, \mathbf{z}) = 1$ if and only if $\mathbf{z} = (z_1, \dots, z_\tau)$ and $f_j(\mathbf{x}, z_m) = 1$ for all $m \in [\tau]$.

The reduction. The reduction prover $\mathbf{P}_{\text{IR}}^{(\tau)}$ and reduction verifier $\mathbf{V}_{\text{IR}}^{(\tau)}$ receive as input a base instance \mathbf{x} and round instances z_1, \dots, z_t where $z_i := (z_{i,1}, \dots, z_{i,\tau})$. They interact as follows.

- They run the interactive reduction $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$ with base instance \mathbf{x} and round instances $(z_{1,1}, \dots, z_{1,\tau}, \dots, z_{t,1}, \dots, z_{t,\tau})$ for τ times in parallel with fresh randomness in each execution.
- Output $\mathbf{z}' := (z'_1, \dots, z'_\tau)$, where z'_i is the result of the i -th independent execution of $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$.

Proof of Theorem 4.5. First we argue completeness, then soundness of the interactive reduction and triviality of the predicate f_{kIR} , and finally analyze the other complexity measures of the protocol.

Completeness. Fix $(\mathbf{x}, \mathbf{w}) \in R$ and z_1, \dots, z_t with $f_j(\mathbf{x}, z_1) = \dots = f_j(\mathbf{x}, z_t) = 1$. Completeness of the interactive reduction $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$ implies that with probability 1 every output z'_i satisfies $f_{j+1}(\mathbf{x}, z'_i) = 1$; in turn, by definition of the new predicates, this means that $\mathbf{z}' := (z'_1, \dots, z'_\tau)$ is such that $f_j^{(\tau)}(\mathbf{x}, \mathbf{z}') = 1$.

Soundness. Fix \mathbf{x} and z_1, \dots, z_t where $f_j^{(\tau)}(\mathbf{x}, z_i) = 0$ for some $i \in [t]$. This implies that there exists $z_{i,m} \in z_i$ such that $f_j^{(\tau)}(\mathbf{x}, z_{i,m}) = 0$. This is one of the instances in the list $(z_{1,1}, \dots, z_{1,\tau}, \dots, z_{t,1}, \dots, z_{t,\tau})$ given as input to the interactive reduction $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$. Thus, for each execution in the protocol we have, by the definition of reducible soundness that for every $\tilde{\mathbf{P}}_{\text{IR}}$:

$$\Pr [f_{j+1}(\mathbf{x}, \mathbf{z}) = 1 \mid \mathbf{z} \leftarrow \langle \tilde{\mathbf{P}}_{\text{IR}}, \mathbf{V}_{\text{IR}}(\mathbf{x}, z_{1,1}, \dots, z_{1,\tau}, \dots, z_{t,1}, \dots, z_{t,\tau}) \rangle] \leq \varepsilon_{\text{IR}}(\mathbf{x}, \tau \cdot t) .$$

Since each of the τ executions is independent, with probability at least $1 - \varepsilon_{\text{IR}}^\tau(\mathbf{x}, \tau \cdot t)$ there exists $i \in [\tau]$ with $f_{j+1}(\mathbf{x}, z'_i) = 0$. This implies that $f_{j+1}^{(\tau)}(\mathbf{x}, \mathbf{z}') = 0$ with probability at least $1 - \varepsilon_{\text{IR}}^\tau(\mathbf{x}, \tau \cdot t)$.

Triviality. Triviality of $f_{\text{kIR}}^{(\tau)}$ follows directly from the triviality of f_{kIR} .

Complexity measures. We discuss complexity measures of the new interactive reduction $(\mathbf{P}_{\text{IR}}^{(\tau)}, \mathbf{V}_{\text{IR}}^{(\tau)})$. Running the new interactive reduction with t round instances is essentially τ executions of the interactive reduction $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$ with $\tau \cdot t$ round instances each. Thus all complexity measures depend on the complexity of $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$ with $\tau \cdot t$ round instances.

- *Output length.* The output of the protocol has length $\tau \cdot s_{\text{IR}}(\mathbf{x}, \tau \cdot t)$.
- *Round complexity.* The protocol has $\ell_{\text{IR}}(\mathbf{x}, \tau \cdot t)$ rounds.
- *Communication complexity.* The communication complexity of the protocol is $\tau \cdot l_{\text{IR}}(\mathbf{x}, \tau \cdot t)$.
- *Randomness complexity.* The verifier of the new interactive reduction uses $\tau \cdot r_{\text{IR}}(\mathbf{x}, \tau \cdot t)$ random bits.
- *Verifier running time.* The verifier of the new interactive reduction runs in time $\tau \cdot vt_{\text{IR}}(\mathbf{x}, \tau \cdot t)$.
- *Final predicate time.* The final predicate of the new interactive reduction runs in time $\tau \cdot ft_{\text{IR}}(\mathbf{x}, \tau \cdot t)$.

□

4.3 Relations with interactive reducibility

We present interactive reductions for various relations. In Section 4.3.1 we show that every language with an IP has an interactive reduction, albeit not one with bounded output length. Then we show how to achieve bounded output length for specific classes of relations: 1. the sumcheck protocol in Section 4.3.2; and 2. Shamir's protocol in Section 4.3.3;

4.3.1 Interactive reductions for any IP

We show that any IP can be transformed into an IP with interactive reducibility where the base instance is the instance for the IP, and the round instances are sets of partial transcripts. The length of a merged round instance grows with the number of partial transcripts in all of the reduction instances combined, and so this reduction does not have bounded output-length.

Theorem 4.7. *If a relation R has a k_{IP} -round public-coin IP $(\mathbf{P}_{\text{IP}}, \mathbf{V}_{\text{IP}})$ then R has an interactive reduction $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$ with the parameters below for every $m > 0$ that divides k_{IP} .*

IP for R		→	Interactive reduction for R for round j given z_1, \dots, z_t	
Rounds	k_{IP}		Number of predicates	k_{IP}/m
Per-round communication	l_{IP}	Rounds	m	
Per-round randomness	r_{IP}	Output length	$t \cdot m \cdot l_{\text{IP}} + \sum_{i \in [t]} z_i $	
Round-by-round soundness error	β_{rbr}	Communication	$t \cdot m \cdot l_{\text{IP}}$	
Per-round verifier running time	\mathbf{vt}	Randomness	$m \cdot r_{\text{IP}}$	
Decision time	\mathbf{d}_{IP}	Soundness error	$m \cdot \beta_{\text{rbr}}$	
		Verifier running time	$t \cdot m \cdot \mathbf{vt}_{\text{IP}}$	
		Final predicate time	\mathbf{d}_{IP}	

Construction 4.8. Let state be the state function of $(\mathbf{P}_{\text{IP}}, \mathbf{V}_{\text{IP}})$. We define the predicates and then describe the interactive reduction for the relation described by this IP.

The predicates. For each $j \in [k_{\text{IP}}/m]$, we define $f_j(\mathbf{x}, z) := 1$ if and only if the following are true:

1. z is a set of $(j \cdot m)$ -round partial transcript ending with a prover message.
2. For every $\text{tr} \in z$, $\text{state}(\mathbf{x}, \text{tr}) = 1$.
3. If $\mathbf{x} \in L(R)$ and $j \neq k_{\text{IP}}/m$, then we require that every $\text{tr} \in z$ is consistent with the honest prover: Suppose $\text{tr} = (a_1, \rho_1, \dots, a_{j \cdot m})$ then $a_1 := \mathbf{P}_{\text{IP}}(\mathbf{x})$ and for every $1 < k \leq j \cdot m$, $a_k := \mathbf{P}_{\text{IP}}(\mathbf{x}, a_1, \rho_1, \dots, a_{k-1}, \rho_{k-1})$.

The reduction. The reduction prover \mathbf{P}_{IR} and reduction verifier \mathbf{V}_{IR} receive as input a base instance \mathbf{x} and round instances z_1, \dots, z_t where each z_i is a set of j -round partial transcripts ending in a prover message. The reduction prover \mathbf{P}_{IR} additionally receives as input a witness w . The protocol is as follows:

1. For $k = 1$ to m :
 - (a) \mathbf{V}_{IR} : Send a uniformly random message ρ_k corresponding to the IP verifier's $(j \cdot m + k)$ -th message.
 - (b) \mathbf{P}_{IR} : For every $\text{tr} \in \bigcup_{i \in [t]} z_i$, send the prover message $a_{\text{tr}, k} := \mathbf{P}_{\text{IP}}(\mathbf{x}, w, \text{tr}, \rho_1, \dots, \rho_k)$.
2. \mathbf{V}_{IR} : Output $z' := \{\text{tr} \mid \|\rho_1\| \|a_{\text{tr}, 1}\| \dots \|\rho_m\| \|a_{\text{tr}, m}\| \}_{\text{tr} \in \bigcup_{i \in [t]} z_i}$.

Proof of Theorem 4.7. First we argue completeness, then soundness and then show triviality of the predicate $f_{k_{\text{IP}}/m}$.

Completeness. Fix $(\mathbf{x}, \mathbf{w}) \in R$, $j \in [k_{\text{IP}}/m]$, and z_1, \dots, z_t such that $f_{j-1}(\mathbf{x}, z_1) = \dots = f_{j-1}(\mathbf{x}, z_t) = 1$. This implies that $\text{state}(\mathbf{x}, \text{tr}) = 1$ for every $\text{tr} \in \bigcup_{i \in [t]} z_i$ and that every transcript is consistent with the honest prover. We now show that for every such transcript tr and ρ_1, \dots, ρ_m , setting $a_{\text{tr},k} := \mathbf{P}_{\text{IP}}(\mathbf{x}, \mathbf{w}, \text{tr}, \rho_1, a_{\text{tr},1}, \dots, a_{\text{tr},k-1}, \rho_k)$, we have

$$\text{state}(\mathbf{x}, \text{tr} \parallel \rho_1 \parallel a_{\text{tr},1} \parallel \dots \parallel \rho_m \parallel a_{\text{tr},m}) = 1 .$$

This implies that $f_{j+1}(\mathbf{x}, z') = 1$, since $f_{j+1}(\mathbf{x}, z') = 1$ if and only if for every tr the state of the resulting transcript after running the protocol is 1.

Notice that any for any transcript tr that is consistent with the honest prover and any ρ : $\text{state}(\mathbf{x}, \text{tr} \parallel \rho \parallel a_{\text{tr}}) = 1$ where $a_{\text{tr}} := \mathbf{P}_{\text{IP}}(\mathbf{x}, \text{tr} \parallel \rho)$. Suppose towards contradiction of perfect completeness of the underlying IP that this was not the case. Then for some choice of ρ the honest prover sends a message that sets the state to 0. Since the protocol has non-zero round-by-round soundness error, once the state is set to zero, there is a non-zero probability that the final transcript has state 0, which will imply that the verifier rejects. Since tr was consistent with the honest prover, there is a non-zero probability that it is generated in an honest execution of the protocol. Therefore, there is non-zero probability that the honest prover will not convince the verifier in a correct execution, in contradiction to perfect completeness of the IP.

The above argument shows that for every transcript, its state will remain 1. By construction, every new transcript $(\text{tr} \parallel \rho_1 \parallel a_{\text{tr},1} \parallel \dots \parallel \rho_m \parallel a_{\text{tr},m})$ is consistent with the honest prover. We therefore conclude that $f_j(\mathbf{x}, z') = 1$.

Soundness. Fix $\mathbf{x} \notin L(R)$, $j \in [k_{\text{IP}}/m]$, and z_1, \dots, z_t and a (computationally unbounded) malicious reduction prover $\tilde{\mathbf{P}}_{\text{IR}}$. Before we begin our analysis, recall that round-by-round soundness of the IP, for any transcript tr with $\text{state}(\mathbf{x}, \text{tr}) = 0$:

$$\Pr_{\rho}[\text{state}(\mathbf{x}, \text{tr} \parallel \rho) = 1] \leq \beta_{\text{rbr}} .$$

Moreover, by the properties of the state function, since $\mathbf{x} \notin L(R)$, if ρ is such that $\text{state}(\mathbf{x}, \text{tr} \parallel \rho) = 0$ then for any a we have that $\text{state}(\mathbf{x}, \text{tr} \parallel \rho_j \parallel a) = 0$. Using induction, we can conclude that for any tr with $\text{state}(\mathbf{x}, \text{tr}) = 0$:

$$\Pr_{\rho_1, \dots, \rho_m} \left[\text{state}(\mathbf{x}, \text{tr} \parallel \rho_1 \parallel a_{\text{tr},1} \parallel \dots \parallel \rho_m \parallel a_{\text{tr},m}) = 1 \mid \begin{array}{l} a_{\text{tr},1} := \tilde{\mathbf{P}}_{\text{IR}}(\text{tr}, \rho_1) \\ \vdots \\ a_{\text{tr},m} := \tilde{\mathbf{P}}_{\text{IR}}(\text{tr}, \rho_1, a_{\text{tr},1}, \dots, a_{\text{tr},m-1}, \rho_m) \end{array} \right] \leq m \cdot \beta_{\text{rbr}} . \quad (2)$$

Suppose that there exists $i \in [t]$ such that $f_{j-1}(\mathbf{x}, z_i) = 0$. This implies that there exists a transcript $\text{tr} \in \bigcup_{i \in [t]} z_i$ with $\text{state}(\mathbf{x}, \text{tr}) = 0$ (notice that the requirement of consistency with the honest prover described in Item 3 is only for $\mathbf{x} \in L(R)$). By Equation (2) and since $(\text{tr} \parallel \rho_1 \parallel a_{\text{tr},1} \parallel \dots \parallel \rho_m \parallel a_{\text{tr},m})$ is added to the new instance z' , this implies that with probability at least $1 - m \cdot \beta_{\text{rbr}}$, $f_j(\mathbf{x}, z') = 0$ as required.

Triviality. $f_{k_{\text{IP}}/m}(\mathbf{x}, \mathbf{z}) = 1$ if and only if all of the transcripts in \mathbf{z} are complete transcripts and have state 1 (notice that since this is the last round, we do not have the requirement of consistency with the honest prover described in Item 3). This can be verified by running \mathbf{V}_{IP} on every $\text{tr} \in \mathbf{z}$, since $\mathbf{V}_{\text{IP}}(\mathbf{x}, \text{tr}) = \text{state}(\mathbf{x}, \text{tr})$.

Complexity measures. We discuss complexity measures of $(\mathbf{P}_{\text{IP}}, \mathbf{V}_{\text{IR}})$.

- *Output length.* The new instance size is $t \cdot m \cdot l_{\text{IP}} + \sum_{i \in [t]} |z_i|$.
- *Round complexity.* The protocol has m rounds.
- *Communication complexity.* The communication complexity of the protocol is $t \cdot m \cdot l_{\text{IP}}$ where l_{IP} is the per-round communication complexity of $(\mathbf{P}_{\text{IP}}, \mathbf{V}_{\text{IP}})$.
- *Randomness complexity.* \mathbf{V}_{IR} uses r_{IP} random bits where $m \cdot r_{\text{IP}}$ is the per-round randomness complexity of \mathbf{V}_{IP} .
- *Verifier running time.* \mathbf{V}_{IR} runs the IP verifier for $t \cdot m$ times in total time $t \cdot m \cdot \text{vt}_{\text{IP}}$.
- *Final predicate time.* The final predicate can be checked in time d_{IP} .

□

4.3.2 Sumcheck language

Definition 4.9. A **sumcheck instance** has the form $(\mathbb{F}, H, n, d, p, \gamma)$ where \mathbb{F} is a field, $H \subseteq \mathbb{F}$, $n \in \mathbb{N}$ is a number of variables, $d \in \mathbb{N}$ is a degree bound, $p: \mathbb{F}^n \rightarrow \mathbb{F}$ is an n -variate polynomial of individual degree at most d (given to the verifier as an oracle) and $\gamma \in \mathbb{F}$ is a claimed sum. We define $(\mathbb{F}, H, n, d, p, \gamma) \in L_{\Sigma}$ if and only if

$$\sum_{\alpha_1, \dots, \alpha_n \in H} p(\alpha_1, \dots, \alpha_n) = \gamma .$$

Theorem 4.10. The sumcheck language L_{Σ} has an interactive reduction that, for instances $\mathfrak{x} = (\mathbb{F}, H, n, d, p, \gamma)$, has the following parameters.

Interactive reduction for L_{Σ}	
Number of predicates	n
Messages	2
Output length	$ \mathfrak{x} $
Communication	$O(ntd^2 \log \mathbb{F})$
Randomness	$2 \log \mathbb{F} $
Soundness error	$O(ntd/ \mathbb{F})$
Verifier running time	$\text{poly}(\mathfrak{x} , t)$
Final predicate time	1 call to p

Construction 4.11. We define the predicates and then describe the interactive reduction.

The predicates. For a base instance $\mathfrak{x} = (\mathbb{F}, H, n, d, p, \gamma)$, for each $j \in [n]$, we define $f_j(\mathfrak{x}, \mathfrak{z}) = 1$ if and only if, parsing $\mathfrak{z} = (r_1, \dots, r_j, \gamma_j) \in \mathbb{F}^{j+1}$, we have

$$\sum_{\alpha_{j+1}, \dots, \alpha_n \in H} p(r_1, \dots, r_j, \alpha_{j+1}, \dots, \alpha_n) = \gamma_j .$$

The reduction. The reduction prover \mathbf{P}_{IR} and reduction verifier \mathbf{V}_{IR} receive as input a base instance \mathfrak{x} and round instances z_1, \dots, z_t . Parse each $z_i := (r_1^{(i)}, \dots, r_j^{(i)}, \gamma_j^{(i)})$. Let $I_{z_1, \dots, z_t}: \mathbb{F} \rightarrow \mathbb{F}^j$ be the polynomial of degree less than t such that $I_{z_1, \dots, z_t}(i) = (r_1^{(i)}, \dots, r_j^{(i)})$ for every $i \in [t]$,² for any $\gamma \in \mathbb{F}$, $I_{z_1, \dots, z_t}(\gamma)$ can be computed in $\text{poly}(t)$ operations. The interactive reduction is as follows:

²Here we implicitly associate the set $[t]$ with an arbitrary set $S \subseteq \mathbb{F}$ of size t .

- \mathbf{P}_{IR} : Send the polynomial $g \in \mathbb{F}[X_1, X_2]$ defined as:

$$g(X_1, X_2) := \sum_{\alpha_{j+2}, \dots, \alpha_n \in H} p(I_{z_1, \dots, z_t}(X_1), X_2, \alpha_{j+2}, \dots, \alpha_n) . \quad (3)$$

- \mathbf{V}_{IR} : Receive a bivariate polynomial $\tilde{g} \in \mathbb{F}[X_1, X_2]$ of degree at most $j \cdot d \cdot (t-1)$ in X_1 and degree at most d in X_2 .

1. *Consistency*: Check that for every $i \in [t]$ it holds that $\sum_{\alpha \in H} \tilde{g}(i, \alpha) = \gamma_j^{(i)}$. (Reject if not.)
2. *Generate new instance*:
 - (a) Sample uniformly random field elements $\rho, r^* \leftarrow \mathbb{F}$ and send them to \mathbf{P}_{IR} .
 - (b) Output the new instance $z' := (r'_1, \dots, r'_j, r^*, \gamma_{j+1})$ where $(r'_1, \dots, r'_j) := I_{z_1, \dots, z_t}(\rho)$ and $\gamma_{j+1} := \tilde{g}(\rho, r^*)$.

Proof of Theorem 4.10. We argue completeness, then soundness, and then triviality of the predicate f_n .

Completeness. Fix $\mathbf{x} = (\mathbb{F}, H, n, d, p, \gamma) \in L_\Sigma$ and z_1, \dots, z_t such that $f_j(\mathbf{x}, z_1) = \dots = f_j(\mathbf{x}, z_t) = 1$ where $z_i := (r_1^{(i)}, \dots, r_j^{(i)}, \gamma_j^{(i)})$. We argue that

$$\Pr [f_j(\mathbf{x}, z') = 1 \mid z' \leftarrow \langle \mathbf{P}_{\text{IR}}(\mathbf{x}, z_1, \dots, z_t), \mathbf{V}_{\text{IR}}(\mathbf{x}, z_1, \dots, z_t) \rangle] = 1 .$$

Let g be defined as in Equation (3). Notice that for every $i \in [t]$,

$$g(i, X_2) = \sum_{\alpha_{j+2}, \dots, \alpha_n \in H} p(r_1^{(i)}, \dots, r_j^{(i)}, X_2, \alpha_{j+2}, \dots, \alpha_n) ,$$

which, since $f_j(\mathbf{x}, z_i) = 1$, implies that $\sum_{\alpha \in H} g(i, \alpha) = \gamma_j^{(i)}$. Therefore the reduction verifier \mathbf{V}_{IR} does not reject in the consistency test (in Item 1). Moreover, by the definition of g , for every ρ and r^* ,

$$\begin{aligned} \sum_{\alpha_{j+2}, \dots, \alpha_n \in H} p(r'_1, \dots, r'_j, r^*, \alpha_{j+2}, \dots, \alpha_n) &= \sum_{\alpha_{j+2}, \dots, \alpha_n \in H} p(I_{z_1, \dots, z_t}(\rho), r^*, \alpha_{j+2}, \dots, \alpha_n) \\ &= g(\rho, r^*) . \end{aligned}$$

Therefore we have $f_i(\mathbf{x}, z') = 1$ with probability 1.

Soundness. Fix $\mathbf{x} = (\mathbb{F}, H, n, d, p, \gamma) \notin L_\Sigma$ and z_1, \dots, z_t where $z_i := (r_1^{(i)}, \dots, r_j^{(i)}, \gamma_j^{(i)})$. Suppose that $f_j(\mathbf{x}, z_i) = 0$ for some $i \in [t]$. Fix a malicious reduction prover $\tilde{\mathbf{P}}_{\text{IR}}$. We show that

$$\Pr [f_j(\mathbf{x}, z') = 1 \mid z' \leftarrow \langle \tilde{\mathbf{P}}_{\text{IR}}, \mathbf{V}_{\text{IR}}(\mathbf{x}, z_1, \dots, z_t) \rangle] \leq \frac{d \cdot (1 + n \cdot (t-1))}{|\mathbb{F}|} .$$

If \tilde{g} does not pass the consistency test (in Item 1) then the reduction verifier outputs \perp , and we have $f_{j+1}(\mathbf{x}, \perp) = 0$. Thus, we can assume that \tilde{g} passes the consistency test (that is, for every $i \in [t]$ it holds that $\sum_{\alpha \in H} \tilde{g}(i, \alpha) = \gamma_j^{(i)}$). Let g be defined as in Equation (3) with respect to \mathbf{x} and z_1, \dots, z_t .

We have that $g \neq \tilde{g}$ due to the fact that $f_j(\mathbf{x}, z_i) = 0$, which implies that

$$\sum_{\alpha \in H} g(i, \alpha) = \sum_{\alpha_{j+1}, \dots, \alpha_n \in H} p(r_1^{(i)}, \dots, r_j^{(i)}, \alpha_{j+1}, \dots, \alpha_n) \neq \gamma_j^{(i)} = \sum_{\alpha \in H} \tilde{g}(i, \alpha) ,$$

Thus, by applying the Schwartz–Zippel lemma, and recalling that g and \tilde{g} both have degree $j \cdot d \cdot (t-1)$ in their first variable and d in their second variable, we get:

$$\Pr_{\rho, r^*} [g(\rho, r^*) = \tilde{g}(\rho, r^*)] \leq \frac{d \cdot (1 + j \cdot (t-1))}{|\mathbb{F}|} \leq \frac{d \cdot (1 + n \cdot (t-1))}{|\mathbb{F}|} .$$

For any choice of ρ and r^* such that $g(\rho, r^*) \neq \tilde{g}(\rho, r^*)$, we have that

$$\sum_{\alpha_{j+2}, \dots, \alpha_n \in H} p(r'_1, \dots, r'_j, r^*, \alpha_{j+2}, \dots, \alpha_n) = g(\rho, r^*) \neq \tilde{g}(\rho, r^*) = \gamma_{j+1} ,$$

and so $f_{j+1}(\mathbf{x}, (r'_1, \dots, r'_j, r^*, \gamma_{j+1})) = 0$ with probability at least $1 - \frac{d \cdot (1 + n \cdot (t-1))}{|\mathbb{F}|}$.

Triviality. By definition, $f_n(\mathbf{x}, \mathbf{z}) = f_n((\mathbb{F}, H, n, d, p, \gamma), (r_1, \dots, r_n, \gamma_n)) = 1$ if and only if $p(r_1, \dots, r_n) = \gamma_n$, which can be efficiently verified by querying p once.

Complexity measures. We discuss complexity measures of the protocol $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$.

- *Output length.* The output is a list of $j + 2 \leq n + 2$ field elements. Thus, the maximal length output is at most $O(n \cdot \log |\mathbb{F}|)$, which is a fixed polynomial in the input length.
- *Rounds.* The interactive reduction has one round.
- *Communication complexity.* The prover sends g which is bivariate and has degree at most $n \cdot d \cdot (t-1)$ in its first variable and d in its second variable. This requires sending $O(ntd^2)$ field elements. The verifier sends two field elements. Thus the communication complexity of the protocol is $O(ntd^2 \log |\mathbb{F}|)$ bits.
- *Randomness complexity.* \mathbf{V}_{IR} uses $2 \log |\mathbb{F}|$ random bits.
- *Verifier running time.* \mathbf{V}_{IR} runs in time $\text{poly}(n, \log |\mathbb{F}|) = \text{poly}(|\mathbf{x}|)$.
- *Final predicate time.* Computing the final predicate requires one call to p .

□

4.3.3 Shamir's protocol

We show an interactive reduction for a class of languages that contains all of PSPACE. For notational convenience, we define an operator that acts either as a sum or product, depending on the index of the variable it is applied over. For every $j \in \mathbb{N}$ let:

$$\Psi_{\alpha_j \in H} := \begin{cases} \sum_{\alpha_j \in H} & j \text{ is odd} \\ \prod_{\alpha_j \in H} & j \text{ is even} \end{cases} .$$

Given polynomials p_0, \dots, p_n where each p_j is a j -variate polynomial over a field \mathbb{F} , we derive polynomials h_0, \dots, h_n where $h_j: \mathbb{F}^j \rightarrow \mathbb{F}$ is defined recursively as follows:

- $h_n(X_1, \dots, X_n) := p_n(X_1, \dots, X_n)$.
- $h_j(X_1, \dots, X_j) := p_j(X_1, \dots, X_j) \cdot \Psi_{\alpha_{j+1}} \left(h_{j+1}(X_1, \dots, X_j, \alpha_{j+1}) \right)$.

Observe that using this notation:

$$h_0 = p_0 \cdot \sum_{\alpha_1 \in H} \left[p_1(\alpha_1) \cdot \prod_{\alpha_2 \in H} \left[p_2(\alpha_1, \alpha_2) \cdot \sum_{\alpha_3 \in H} \left[\dots \left[p_{n-1}(\alpha_1, \dots, \alpha_{n-1}) \cdot \prod_{\alpha_n \in H} p_n(\alpha_1, \dots, \alpha_n) \right] \dots \right] \right] \right] .$$

Definition 4.12. A **Shamir instance** has the form $(\mathbb{F}, H, n, d, p_0, \dots, p_n, \gamma)$ where \mathbb{F} is a field, $H \subseteq \mathbb{F}$, $n \in \mathbb{N}$ is a number of variables, $d \in \mathbb{N}$ is a degree bound, each $p_j: \mathbb{F}^j \rightarrow \mathbb{F}$ is an j -variate polynomial where the derived polynomials h_0, \dots, h_n have individual degree at most d and $\gamma \in \mathbb{F}$ is a claimed value. We define $(\mathbb{F}, H, n, d, p_0, \dots, p_n, \gamma) \in L_{\text{Shmr}}$ if and only if $h_0 = \gamma$.

Note that a sumcheck instance (as defined in Section 4.3.2) is a Shamir instance where p_0, \dots, p_{n-1} are equal to the constant 1 polynomial, and $p_n(X_1, \dots, X_n) \triangleq p'_n(X_1, X_3, \dots, X_{n-1})$ (assuming n is even) for some polynomial p'_n of individual degree at most d .

Theorem 4.13. The Shamir language L_{Shmr} has an interactive reduction that, for instances $\mathfrak{x} = (\mathbb{F}, H, n, d, p_0, \dots, p_n, \gamma)$, has the following parameters.

Interactive reduction for L_{Shmr}	
Number of predicates	n
Messages	2
Output length	$ \mathfrak{x} $
Communication	$O(ntd^2 \log \mathbb{F})$
Randomness	$2 \log \mathbb{F} $
Soundness error	$O(ntd/ \mathbb{F})$
Verifier running time	$\text{poly}(\mathfrak{x} , t)$
Final predicate time	1 call to p_n

Shamir [Sha92] showed how deciding every language in PSPACE can be reduced to computing the value of an arithmetic expression of the form of Shamir instances. Specifically, Shamir’s reduction takes a TQBF formula and transforms it into a “simple” TQBF formula, roughly defined as one in which every occurrence of every variable is separated from its quantification point by at most one universal quantifier (and arbitrarily many other symbols). This ensures that when arithmetizing the formula, one ends up with an expression in which the degree of every variable is polynomial.

Using the above theorem, and Shamir’s reduction we get the following corollary showing that every language in PSPACE has an interactive reduction with bounded output length albeit with polynomially-many predicates.

Corollary 4.14. Every $L \in \text{PSPACE}$ has an interactive reduction that, for instances \mathfrak{x} , has the following parameters:

Interactive reduction for L	
Number of predicates	$\text{poly}(\mathfrak{x})$
Messages	2
Output length	$\text{poly}(\mathfrak{x})$
Communication	$\text{poly}(\mathfrak{x})$
Randomness	$\text{poly}(\mathfrak{x})$
Soundness error	$O(t/2^{ \mathfrak{x} })$
Verifier running time	$\text{poly}(\mathfrak{x} , t)$
Final predicate time	$\text{poly}(\mathfrak{x})$

Construction 4.15. We define the predicates and then describe the interactive reduction.

The predicates. For a base instance $\mathbf{x} = (\mathbb{F}, H, n, d, p_0, \dots, p_n, \gamma)$, for each $j \in [n]$, we define $f_j(\mathbf{x}, \mathbf{z}) = 1$ if and only if, parsing $\mathbf{z} = (r_1, \dots, r_j, \gamma_j) \in \mathbb{F}^j$, we have $h_j(r_1, \dots, r_j) = \gamma_j$.

The reduction. The reduction prover \mathbf{P}_{IR} and reduction verifier \mathbf{V}_{IR} receive as input a base instance \mathbf{x} and round instances $\mathbf{z}_1, \dots, \mathbf{z}_t$. Parse each $\mathbf{z}_i := (r_1^{(i)}, \dots, r_j^{(i)}, \gamma_j^{(i)})$. Let $I_{\mathbf{z}_1, \dots, \mathbf{z}_t}: \mathbb{F} \rightarrow \mathbb{F}^j$ be the polynomial of degree less than t such that $I_{\mathbf{z}_1, \dots, \mathbf{z}_t}(i) = (r_1^{(i)}, \dots, r_j^{(i)})$ for every $i \in [t]$ (we implicitly associate the set $[t]$ with an arbitrary set $S \subseteq \mathbb{F}$ of size t); for any $\gamma \in \mathbb{F}$, $I_{\mathbf{z}_1, \dots, \mathbf{z}_t}(\gamma)$ can be computed in $\text{poly}(t)$ operations. The interactive reduction is as follows:

- \mathbf{P}_{IR} : Send the polynomial $g \in \mathbb{F}[X_1, X_2]$ defined as:

$$g(X_1, X_2) := h_{j+1}(I_{\mathbf{z}_1, \dots, \mathbf{z}_t}(X_1), X_2) . \quad (4)$$

- \mathbf{V}_{IR} : Receive a bivariate polynomial $\tilde{g} \in \mathbb{F}[X_1, X_2]$ of degree at most $j \cdot d \cdot (t-1)$ in X_1 and degree at most d in X_2 .

1. *Consistency*: Check that for every $i \in [t]$ it holds that

$$p_j(I_{\mathbf{z}_1, \dots, \mathbf{z}_t}(i)) \cdot \prod_{\alpha_{j+1} \in H} \tilde{g}(i, \alpha_{j+1}) = \gamma_j^{(i)} .$$

(Reject if not.)

2. *Generate new instance*:

- (a) Sample uniformly random field elements $\rho, r^* \leftarrow \mathbb{F}$ and send them to \mathbf{P}_{IR} .
- (b) Output the new instance $\mathbf{z}' := (r'_1, \dots, r'_j, r^*, \gamma_{j+1})$ where $(r'_1, \dots, r'_j) := I_{\mathbf{z}_1, \dots, \mathbf{z}_t}(\rho)$ and $\gamma_{j+1} := \tilde{g}(\rho, r^*)$.

Proof of Theorem 4.10. We argue completeness, then soundness, and then triviality of the predicate f_n .

Completeness. Fix $\mathbf{x} = (\mathbb{F}, H, n, d, p_0, \dots, p_n, \gamma) \in L_{\text{shmr}}$ and $\mathbf{z}_1, \dots, \mathbf{z}_t$ such that $f_j(\mathbf{x}, \mathbf{z}_1) = \dots = f_j(\mathbf{x}, \mathbf{z}_t) = 1$ where $\mathbf{z}_i := (r_1^{(i)}, \dots, r_j^{(i)}, \gamma_j^{(i)})$. We argue that

$$\Pr [f_j(\mathbf{x}, \mathbf{z}') = 1 \mid \mathbf{z}' \leftarrow \langle \mathbf{P}_{\text{IR}}(\mathbf{x}, \mathbf{z}_1, \dots, \mathbf{z}_t), \mathbf{V}_{\text{IR}}(\mathbf{x}, \mathbf{z}_1, \dots, \mathbf{z}_t) \rangle] = 1 .$$

For every $i \in [t]$, since $f_j(\mathbf{x}, \mathbf{z}_i) = 1$, we have that

$$\begin{aligned} p_j(I_{\mathbf{z}_1, \dots, \mathbf{z}_t}(i)) \cdot \prod_{\alpha_{j+1} \in H} g(i, \alpha_{j+1}) &= p_j(r_1^{(i)}, \dots, r_j^{(i)}) \cdot \prod_{\alpha_{j+1} \in H} h_{j+1}(r_1^{(i)}, \dots, r_j^{(i)}, \alpha_{j+1}) \\ &= h_j(r_1^{(i)}, \dots, r_j^{(i)}) \\ &= \gamma_j^{(i)} . \end{aligned}$$

Therefore the reduction verifier \mathbf{V}_{IR} does not reject in the consistency test (in Item 1). Moreover, by the definition of g , for every ρ and r^* ,

$$\begin{aligned} h_{j+1}(r'_1, \dots, r'_j, r^*) &= h_{j+1}(I_{\mathbf{z}_1, \dots, \mathbf{z}_t}(\rho), r^*) \\ &= g(\rho, r^*) \end{aligned}$$

$$= \gamma_{j+1} .$$

Therefore we have $f_i(\mathbf{x}, \mathbf{z}') = 1$ with probability 1.

Soundness. Fix $\mathbf{x} = (\mathbb{F}, H, n, d, p_0, \dots, p_n, \gamma) \notin L_{\text{Shmr}}$ and $\mathbf{z}_1, \dots, \mathbf{z}_t$ where $\mathbf{z}_i := (r_1^{(i)}, \dots, r_j^{(i)}, \gamma_j^{(i)})$. Suppose that $f_j(\mathbf{x}, \mathbf{z}_i) = 0$ for some $i \in [t]$. Fix a malicious reduction prover $\tilde{\mathbf{P}}_{\text{IR}}$. We show that

$$\Pr [f_j(\mathbf{x}, \mathbf{z}') = 1 \mid \mathbf{z}' \leftarrow \langle \tilde{\mathbf{P}}_{\text{IR}}, \mathbf{V}_{\text{IR}}(\mathbf{x}, \mathbf{z}_1, \dots, \mathbf{z}_t) \rangle] \leq \frac{d \cdot (1 + n \cdot (t-1))}{|\mathbb{F}|} .$$

If \tilde{g} does not pass the consistency test (in Item 1) then the reduction verifier rejects. Thus, we can assume that \tilde{g} passes the consistency test. That is, for every $i \in [t]$ it holds that

$$p_j(I_{\mathbf{z}_1, \dots, \mathbf{z}_t}(i)) \cdot \Psi_{\alpha_{j+1} \in H} \tilde{g}(i, \alpha_{j+1}) = \gamma_j^{(i)} .$$

Let g be defined as in Equation (4) with respect to \mathbf{x} and $\mathbf{z}_1, \dots, \mathbf{z}_t$.

We have that $g \neq \tilde{g}$ due to the fact that $f_j(\mathbf{x}, \mathbf{z}_i) = 0$, which implies that

$$\begin{aligned} p_j(I_{\mathbf{z}_1, \dots, \mathbf{z}_t}(i)) \cdot \Psi_{\alpha_{j+1} \in H} g(i, \alpha_{j+1}) &= h_j(r_1^{(i)}, \dots, r_j^{(i)}) \\ &\neq \gamma_j^{(i)} \\ &= p_j(I_{\mathbf{z}_1, \dots, \mathbf{z}_t}(i)) \cdot \Psi_{\alpha_{j+1} \in H} \tilde{g}(i, \alpha_{j+1}) . \end{aligned}$$

Thus, by applying the Schwartz–Zippel lemma, and recalling that g and \tilde{g} both have degree $j \cdot d \cdot (t-1)$ in their first variable and d in their second variable, we get:

$$\Pr_{\rho, r^*} [g(\rho, r^*) = \tilde{g}(\rho, r^*)] \leq \frac{d \cdot (1 + j \cdot (t-1))}{|\mathbb{F}|} \leq \frac{d \cdot (1 + n \cdot (t-1))}{|\mathbb{F}|} .$$

For any choice of ρ and r^* such that $g(\rho, r^*) \neq \tilde{g}(\rho, r^*)$, we have that

$$h_{j+1}(r'_1, \dots, r'_j, r^*) = g(\rho, r^*) \neq \tilde{g}(\rho, r^*) = \gamma_{j+1} .$$

and so $f_{j+1}(\mathbf{x}, (r'_1, \dots, r'_j, r^*, \gamma_{j+1})) = 0$ with probability at least $1 - \frac{d \cdot (1 + n \cdot (t-1))}{|\mathbb{F}|}$.

Triviality. By definition, $f_n(\mathbf{x}, \mathbf{z}) = f_n((\mathbb{F}, H, n, d, p_0, \dots, p_n, \gamma), (r_1, \dots, r_n, \gamma_n)) = 1$ if and only if

$$h_j(r_1, \dots, r_n) \triangleq p_n(r_1, \dots, r_n) = \gamma ,$$

which can be efficiently verified by querying p_n once.

Complexity measures. We discuss complexity measures of the protocol $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$.

- *Output length.* The output is a list of $j + 1 \leq n + 1$ field elements. Thus, the output length is at most $O(n \cdot \log |\mathbb{F}|)$, which is a fixed polynomial in the input length.
- *Rounds.* The interactive reduction has one round.

- *Communication complexity.* The prover sends g which is bivariate and has degree at most $n \cdot d \cdot (t-1)$ in its first variable and d in its second variable. This requires sending $O(ntd^2)$ field elements. The verifier sends two field elements. Thus the communication complexity of the protocol is $O(ntd^2 \log |\mathbb{F}|)$ bits.
- *Randomness complexity.* \mathbf{V}_{IR} uses $2 \log |\mathbb{F}|$ random bits.
- *Verifier running time.* \mathbf{V}_{IR} runs in time $\text{poly}(n, \log |\mathbb{F}|) = \text{poly}(|\mathbf{x}|)$.
- *Final predicate time.* Computing the final predicate requires one call to p_n .

□

5 IOPs from interactive reducibility

We show how to use interactive reducibility to construct IOPs with small query complexity. In Section 5.1, we show that every language with a bounded-output-length interactive reduction (with good enough soundness) has a round-query IOP with round-query complexity $O(1)$. Recall that a round-query IOP with round-query complexity q_{rnd} is an IOP in which the verifier reads q_{rnd} rounds (prover and verifier messages) in their entirety. In Section 5.2, we show that, while the interactive reduction for any IP (in Section 4.3.1) does not have bounded output length, a similar claim can be made for any k -round IP; in this case, the resulting verifier queries $\max\{O(1), O(k/\log |\mathbb{x}|)\}$ rounds.

Combining Theorem 5.3 and Theorem 7.1 (for transforming round-query IOPs into binary IOPs), we obtain the following corollary showing that languages with interactive reducibility have binary IOPs with constant query complexity.

Corollary 5.1 (restatement of Lemma 2). *Let R be a relation with a bounded-output-length interactive reduction with well-behaved soundness error, k_{IR} predicates, and ℓ_{IR} rounds. Then R has a non-adaptive public-coin IOP with the parameters below.*

IOP for R	
Rounds	$O(\ell_{\text{IR}} \cdot k_{\text{IR}})$
Proof length	$\text{poly}(\mathbb{x})$
Alphabet size	2
Queries	$O(\ell_{\text{IR}})$
Interaction randomness	$\text{poly}(\mathbb{x})$
Decision randomness	$O(\log \mathbb{x})$
Soundness error	$O(1)$
Verifier running time	$\text{poly}(\mathbb{x})$

Proof. We first amplify the soundness error of the interactive reduction using Theorem 4.5 until we have $\binom{2 \cdot k_{\text{IR}}}{k_{\text{IR}}} \cdot k_{\text{IR}} \cdot \varepsilon_{\text{IR}}(\mathbb{x}, 2 \cdot k_{\text{IR}}) < \frac{1}{2}$. Since ε_{IR} is well-behaved and $\binom{2 \cdot k_{\text{IR}}}{k_{\text{IR}}} < 2^{2k_{\text{IR}}}$, this incurs only a polynomial overhead. We then apply Theorem 5.3 with parameter $\tau = 2$ to get an IOP with $O(\ell_{\text{IR}} \cdot k_{\text{IR}})$ rounds in which the verifier reads only $O(\ell_{\text{IR}})$ of the rounds with alphabet size $2^{\text{poly}(|\mathbb{x}|)}$. We then apply the transformation described Theorem 7.1 to transform this IOP into a binary one in which the verifier queries $O(\ell_{\text{IR}})$ rounds and makes $O(1)$ queries to each round (including its own random messages). \square

Similarly, we get the following corollary for any public-coin IP.

Corollary 5.2 (restatement of Lemma 1). *Let R be a relation with a non-adaptive k_{IP} -round public-coin IP. Then R has an IOP with the parameters below.*

IOP for R	
Rounds	k_{IP}
Proof length	$\text{poly}(\mathbb{x})$
Alphabet size	2
Queries	$\max\{O(1), O(k_{\text{IP}}/\log \mathbb{x})\}$
Interaction randomness	$\text{poly}(\mathbb{x})$
Decision randomness	$O(\log \mathbb{x})$
Soundness error	$O(1)$
Verifier running time	$\text{poly}(\mathbb{x})$

Proof. Given an IP for the relation R , we first apply the Babai–Moran transformation [BM88] to reduce the number of rounds by a constant fraction (enough to counteract the constant added to the number of rounds by the next two transformations). Next, we apply Theorem 5.7 when setting $m := \max\{O(1), O(k'_{\text{IP}}/\log |\mathbb{x}|)\}$ (where k'_{IP} is the number of rounds of the IP following the Babai–Moran transformation) and then apply Theorem 7.1 to the resulting IOP. \square

5.1 Round-query IOPs from bounded-output-length interactive reductions

We show that any relation with a bounded-output-length interactive reduction has a round-query IOP with polynomial proof length with small round-query complexity.

Theorem 5.3. *Let $\tau \in \mathbb{N}$ and let R be a relation with a bounded-output-length interactive reduction $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$ with well-behaved soundness error. Then R has a public-coin round-query IOP $(\mathbf{P}_{\text{IOP}}, \mathbf{V}_{\text{IOP}})$ with the parameters indicated below.*

Interactive reduction for R		Round-query IOP for R	
Number of predicates	k_{IR}	Rounds	$\tau \cdot k_{\text{IR}} \cdot \ell_{\text{IR}}$
Rounds	ℓ_{IR}	Proof length (per round)	$\tau \cdot k_{\text{IR}}^2 \cdot (s_{\text{IR}} + l_{\text{IR}})$
Communication	l_{IR}	Round queries	$O(\ell_{\text{IR}})$
Output length	s_{IR}	Interaction randomness	$O(\tau \cdot k_{\text{IR}} \cdot r_{\text{IR}})$
Randomness	r_{IR}	Decision randomness	$O(\log(\tau \cdot k_{\text{IR}}))$
Soundness error	ε_{IR}	Soundness error	$\max\{1/\tau, \binom{\tau \cdot k_{\text{IR}}}{k_{\text{IR}}} \cdot k_{\text{IR}} \cdot \varepsilon_{\text{IR}}(\mathbb{x}, \tau \cdot k_{\text{IR}})\}$
Verifier running time	vt_{IR}	Verifier running time	$\text{poly}(\tau, k_{\text{IR}}, s_{\text{IR}}, l_{\text{IR}}, vt_{\text{IR}}, ft_{\text{IR}})$
Final predicate time	ft_{IR}		

Construction 5.4. Let $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$ be an interactive reduction for R with k_{IR} predicates.

The IOP prover \mathbf{P}_{IOP} receives as input an instance \mathbb{x} and witness \mathbb{w} , and the IOP verifier \mathbf{V}_{IOP} receives as input the instance \mathbb{x} . They interact as follows.

- For every $i \in \{0, \dots, \tau \cdot k_{\text{IR}}\}$, \mathbf{P}_{IOP} defines the $(k_{\text{IR}} + 1)$ -entry array A_i as follows

$$A_i[j] := \begin{cases} \{\perp\} & \text{if } j = 0 \\ \emptyset & \text{if } j \in \{1, \dots, k_{\text{IR}}\} \end{cases} .$$

The set $A_i[j]$ will be used to store all of the instances corresponding to f_j collected by iteration i of the protocol.

- For $i = 1, \dots, \tau \cdot k_{\text{IR}}$:
 - \mathbf{P}_{IOP} sends A_{i-1} to \mathbf{V}_{IOP} .
 - Do in parallel for every $j \in [k_{\text{IR}}]$ such that $A_{i-1}[j-1] \neq \emptyset$:
 - Execute the *interaction phase* of:

$$z'_{i,j} \leftarrow \langle \mathbf{P}_{\text{IR}}(\mathbb{x}, \mathbb{w}, A_{i-1}[j-1]), \mathbf{V}_{\text{IR}}(\mathbb{x}, A_{i-1}[j-1]) \rangle .$$

Here the verifier only sends messages and does not verify the correctness of the execution. Hence $z'_{i,j}$ is computed only by the prover. Moreover, the verifier's random messages are shared among the parallel executions of $\langle \mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}} \rangle$ over each choice of $j \in [k_{\text{IR}}]$.

- \mathbf{P}_{IOP} sets $A_i[j] := A_{i-1}[j] \cup \{z'_{i,j}\}$.

3. \mathbf{P}_{IOP} sends $A_{\tau \cdot k_{\text{IR}}}$. Additionally, for every $i \in \{0, \dots, \tau \cdot k_{\text{IR}}\}$, it sends $B_i := A_i$. This concludes the interaction phase.
4. In the decision phase, \mathbf{V}_{IOP} is given oracle access to a transcript with the following structure:

$$\left(A_0, \{\text{tr}_{1,j}\}_{j \in [k_{\text{IR}}]}, A_1, \dots, \{\text{tr}_{\tau \cdot k_{\text{IR}}, j}\}_{j \in [k_{\text{IR}}]}, A_{\tau \cdot k_{\text{IR}}}, (B_0, \dots, B_{\tau \cdot k_{\text{IR}}}) \right), \quad (5)$$

where $\text{tr}_{i,j}$ corresponds to the interaction between the prover and verifier in the j -th parallel execution of $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$ on the i -th iteration. \mathbf{V}_{IOP} performs the checks below.

- (a) *Subset consistency.* Read the arrays $B_0, B_1, \dots, B_{\tau \cdot k_{\text{IR}}}$ in their entirety. For every $i \in [\tau \cdot k_{\text{IR}}]$ and $j \in \{0, \dots, k_{\text{IR}}\}$ check that $B_{i-1}[j] \subseteq B_i[j]$.
- (b) *Transcript consistency.* Sample a random $i \in [\tau \cdot k_{\text{IR}}]$. Read the arrays A_{i-1} and A_i sent by \mathbf{P}_{IOP} and the entire ℓ_{IR} -round interaction $\{\text{tr}_{i,j}\}_{j \in [k_{\text{IR}}]}$ between the prover and the verifier done during the i -th iteration of the the interaction phase.
 - i. Check that $A_{i-1} = B_{i-1}$ and $A_i = B_i$.
 - ii. For every $j \in [k_{\text{IR}}]$, check that $A_i[j] = A_{i-1}[j] \cup \{z'_{i,j}\}$ where

$$z'_{i,j} := \mathbf{V}_{\text{IR}}(\mathbf{x}, A_{i-1}[j-1]; \text{tr}_{i,j}) .$$

(If \mathbf{V}_{IR} rejects then \mathbf{V}_{IOP} immediately rejects.)

- (c) *Final predicate holds.* Check that $f_{k_{\text{IR}}}(\mathbf{x}, \mathbf{z}) = 1$ for every $\mathbf{z} \in B_{\tau \cdot k_{\text{IR}}}[k_{\text{IR}}]$.

Proof of Theorem 5.3. We prove completeness and soundness; then we analyze the complexity measures of the protocol.

Completeness. Fix $(\mathbf{x}, \mathbf{w}) \in R$. We consider each of \mathbf{V}_{IOP} 's checks, and show that \mathbf{P}_{IOP} 's messages satisfy each check with probability 1.

- (a) *Subset consistency.* \mathbf{P}_{IOP} sets, in each round i and for every j , $A_i[j] := A_{i-1}[j] \cup \{z'_{i,j}\}$ where $z'_{i,j}$ is the output of the interactive reduction. Since \mathbf{P}_{IOP} is honest, $B_{i-1} = A_{i-1}$ and $B_i = A_i$. Together this implies that $B_{i-1}[j] \subseteq B_i[j]$ and so \mathbf{V}_{IOP} 's subset consistency check passes.
- (b) *Transcript consistency.* The conditions established in the previous item also directly imply that \mathbf{V}_{IOP} 's transcript consistency check passes.
- (c) *Final predicate holds.* Since $B_i := A_i$, if A_i contains only round instances for which the predicate holds, then so does B_i , and \mathbf{V}_{IOP} accepts. We prove by induction on $i \in \{0, \dots, \tau \cdot k_{\text{IR}}\}$ that for every $j \in \{0, \dots, k_{\text{IR}}\}$ and $\mathbf{z} \in A_i[j]$ it holds that $f_j(\mathbf{x}, \mathbf{z}) = 1$.

First we consider the case when $j = 0$ (and any i). For every i , we have that $A_i[0] = \{\perp\}$. Since $\mathbf{x} \in L(R)$, we have that $f_j(\mathbf{x}, \perp) = 1$.

Next we consider the case when $j > 0$.

- $i = 0$: For every $j \in [k_{\text{IR}}]$ we set $A_0[j] = \emptyset$ and so the property holds (trivially) for these values of j .
- $i > 0$: Suppose that, with probability 1, for every $j \in [k_{\text{IR}}]$ and $\mathbf{z}_i \in A_i[j]$ it holds that $f_j(\mathbf{x}, \mathbf{z}_i) = 1$. We argue that, with probability 1, for every $j \in [k_{\text{IR}}]$ and $\mathbf{z}_{i+1} \in A_{i+1}[j]$, it holds that $f_j(\mathbf{x}, \mathbf{z}_{i+1}) = 1$. Fix $j > 0$. Since for every $\mathbf{z}_i \in A_i[j]$, $f_j(\mathbf{x}, \mathbf{z}_i) = 1$ and $A_i[j] := A_{i-1}[j] \cup \{z'_{i,j}\}$, we need only show that $f(\mathbf{x}, z'_{i,j}) = 1$. $z'_{i,j}$ is generated by applying

the interactive reduction with inputs \mathbb{x} and $\{z_1, \dots, z_{t_{i,j}}\} := A_{i-1}[j-1]$. Since for every $z' \in A_{i-1}[j-1]$, $f_{j-1}(\mathbb{x}, z') = 1$, by completeness of the interactive reduction we have that

$$\Pr \left[f_j(\mathbb{x}, z'_{i,j}) = 1 \mid z'_{i,j} \leftarrow \langle \mathbf{P}_{\text{IR}}(\mathbb{x}, \mathbb{w}, z_1, \dots, z_{t_{i,j}}), \mathbf{V}_{\text{IR}}(\mathbb{x}, z_1, \dots, z_{t_{i,j}}) \rangle \right] = 1 .$$

Soundness. Fix $\mathbb{x} \notin R(L)$ and an IOP prover $\tilde{\mathbf{P}}_{\text{IOP}}$. We show that \mathbf{V}_{IOP} accepts with probability at most $\max\{1/\tau, \binom{\tau \cdot k_{\text{IR}}}{k_{\text{IR}}} \cdot k_{\text{IR}} \cdot \varepsilon_{\text{IR}}(\mathbb{x}, \tau \cdot k_{\text{IR}})\}$. Consider a transcript tr of the interaction of the IOP with the structure described in Equation (5). For fixed tr we say that an index $i \in [\tau \cdot k_{\text{IR}}]$ is *consistent* if:

- $A_{i-1} = B_{i-1}$ and $A_i = B_i$ and,
- For every $j \in [k_{\text{IR}}]$: $A_i[j] = A_{i-1}[j] \cup \{z'_{i,j}\}$ where $z'_{i,j} := \mathbf{V}_{\text{IR}}(\mathbb{x}, A_{i-1}[j]; \text{tr}_{i,j})$ (if \mathbf{V}_{IR} rejects then i is not consistent).

We give two claims analyzing the probability that the verifier accepts relative to how many indices of the protocol are consistent. Claim 5.5 shows that in any execution of the protocol with less than k_{IR} consistent indices, \mathbf{V}_{IOP} accepts with probability at most $1/\tau$. By Claim 5.6, for any fixed set of k_{IR} indices, conditioned on the interaction outputting a transcript in which these indices are consistent, \mathbf{V}_{IOP} accepts with probability at most $k_{\text{IR}} \cdot \varepsilon_{\text{IR}}(\mathbb{x}, \tau \cdot k_{\text{IR}})$. Since there are $\binom{\tau \cdot k_{\text{IR}}}{k_{\text{IR}}}$ choices of k_{IR} indices, we can conclude that, conditioned on the transcript having at least k_{IR} consistent indices, \mathbf{V}_{IOP} accepts with probability at most $\binom{\tau \cdot k_{\text{IR}}}{k_{\text{IR}}} \cdot k_{\text{IR}} \cdot \varepsilon_{\text{IR}}(\mathbb{x}, \tau \cdot k_{\text{IR}})$. Putting the two cases together, we conclude that \mathbf{V}_{IOP} accepts with probability at most $\max\{1/\tau, \binom{\tau \cdot k_{\text{IR}}}{k_{\text{IR}}} \cdot k_{\text{IR}} \cdot \varepsilon_{\text{IR}}(\mathbb{x}, \tau \cdot k_{\text{IR}})\}$.

Claim 5.5. *Conditioned on the transcript tr generated by $\langle \tilde{\mathbf{P}}_{\text{IOP}}, \mathbf{V}_{\text{IOP}}(\mathbb{x}) \rangle$ being consistent with less than k_{IR} indices, \mathbf{V}_{IOP} accepts with probability at most $1/\tau$ (over its decision randomness).*

Proof. \mathbf{V}_{IOP} accepts only if its choice of $i \in [\tau \cdot k_{\text{IR}}]$ in Item 4b is consistent. Since there are at most k_{IR} consistent indices, the probability of the verifier sampling one of these indices is at most $1/\tau$. \square

Claim 5.6. *Fix indices $1 \leq i_1 < \dots < i_{k_{\text{IR}}} \leq \tau \cdot k_{\text{IR}}$. Then, conditioned on the transcript tr generated by $\langle \tilde{\mathbf{P}}_{\text{IOP}}, \mathbf{V}_{\text{IOP}}(\mathbb{x}) \rangle$ being consistent with respect to indices $i_1, \dots, i_{k_{\text{IR}}}$, \mathbf{V}_{IOP} accepts with probability at most $k_{\text{IR}} \cdot \varepsilon_{\text{IR}}(\mathbb{x}, \tau \cdot k_{\text{IR}})$ (over its interaction randomness).*

Proof. Let $i_0 := 0$. We show by induction that for every $j \in \{0, \dots, k_{\text{IR}}\}$ there exists $z \in A_{i_j}[j]$ such that $f_j(\mathbb{x}, z) = 0$ except with probability $j \cdot \varepsilon_{\text{IR}}$. This implies the statement, since if $f_{k_{\text{IR}}}(\mathbb{x}, z) = 0$, then \mathbf{V}_{IOP} rejects in Item 4c.

This is immediate for $j = 0$, since $\mathbb{x} \notin L(R)$ and so $f_0(\mathbb{x}, z) = 0$ for any z . Fix some $j > 0$ and suppose that there exists $z \in A_{i_{(j-1)}}[j-1]$ such that $f_{j-1}(\mathbb{x}, z) = 0$ (which happens with probability $1 - (j-1) \cdot \varepsilon_{\text{IR}}$). We first show that $z \in A_{i_j}[j-1]$. This can be seen by the following:

$$z \in A_{i_{j-1}}[j-1] = B_{i_{j-1}}[j-1] \subseteq B_{i_j-1}[j-1] = A_{i_j-1}[j-1] ,$$

where the equalities are due to the fact that the indices $i_{(j-1)}$ and i_j are consistent. The containment of $B_{i_{(j-1)}}[j-1]$ in $B_{i_j-1}[j-1]$ is due to the fact that $i_{(j-1)} \leq i_j - 1$ and that, in order for the verifier to accept the transcript, the check in Item 4a must pass.

This implies that for any $\tilde{\mathbf{P}}_{\text{IR}}$ (and in particular for whatever $\tilde{\mathbf{P}}_{\text{IOP}}$ does in this stage):

$$\Pr \left[f_j(\mathbb{x}, z'_{i_j,j}) = 1 \mid z'_{i_j,j} \leftarrow \langle \tilde{\mathbf{P}}_{\text{IR}}, \mathbf{V}_{\text{IR}}(\mathbb{x}, A_{i_j-1}[j-1]) \rangle \right] \leq \varepsilon_{\text{IR}}(\mathbb{x}, |A_{i_j-1}[j-1]|) \leq \varepsilon_{\text{IR}}(\mathbb{x}, \tau \cdot k_{\text{IR}}) ,$$

where the final inequality is true since $|A_{i_j-1}[j-1]| \leq \tau \cdot k_{\text{IR}}$, when additionally noting that ε_{IR} is well-behaved. Using again the fact that index i_j is consistent, we have that $A_{i_j}[j] := A_{i_j-1}[j] \cup \{z'_{i_j,j}\}$

where $z'_{i_j,j}$ is a round instance generated by the actual interaction between $\tilde{\mathbf{P}}_{\text{IOP}}$ and \mathbf{V}_{IOP} during the j -th parallel execution of the interactive reduction in iteration i_j . Therefore, we have that, conditioned there existing $z \in A_{i_{(j-1)}}[j-1]$ where $f_{j-1}(\mathbf{x}, z) = 0$, there exists $z'_{i_j,j} \in A_{i_j}[j]$ where $f_j(\mathbf{x}, z'_{i_j,j}) = 0$ except with probability $\varepsilon_{\text{IR}}(\mathbf{x}, \tau \cdot k_{\text{IR}})$. By the induction hypothesis, the condition happens with probability at least $1 - (j-1) \cdot \varepsilon_{\text{IR}}(\mathbf{x}, \tau \cdot k_{\text{IR}})$ and so we have that with probability at least $1 - j \cdot \varepsilon_{\text{IR}}(\mathbf{x}, \tau \cdot k_{\text{IR}})$, there exists $z'_{i_j,j} \in A_{i_j}[j]$ where $f_j(\mathbf{x}, z'_{i_j,j}) = 0$. \square

Complexity measures. We analyze the efficiency parameters of the IOP:

- *Proof length.* The largest message sent by \mathbf{P}_{IOP} is the final round, which contains all of the instances generated in the previous interaction. In each of the $\tau \cdot k_{\text{IR}}$ rounds, one instance is added to the list sent to the verifier for every $j \in [k_{\text{IR}}]$. The output of the interactive reduction has length at most s_{IR} . Thus this requires sending $O(\tau \cdot k_{\text{IR}}^2 \cdot s_{\text{IR}})$ bits. In every execution of the interactive reduction an additional l_{IR} bits are sent, for a total of $k_{\text{IR}} \cdot l_{\text{IR}}$ per round. Thus the final proof is $O(\tau \cdot k_{\text{IR}}^2 \cdot (s_{\text{IR}} + l_{\text{IR}}))$ bits long.
- *Round complexity.* The IOP contains $\tau \cdot k_{\text{IR}}$ executions of the ℓ_{IR} -round interactive reduction $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$. Therefore the IOP has $\tau \cdot k_{\text{IR}} \cdot \ell_{\text{IR}}$ rounds.
- *Round queries.* \mathbf{V}_{IOP} reads in its entirety the last message in the protocol, reads A_{i-1} and A_i , and reads the execution of the interactive reduction between blocks $i-1$ and i . Therefore it reads $O(\ell_{\text{IR}})$ rounds.
- *Randomness complexity.* \mathbf{V}_{IOP} generates r_{IR} bits for every separate $i \in [\tau \cdot k_{\text{IR}}]$ since we share verifier randomness between each of the parallel executions of $(\mathbf{P}_{\text{IR}}, \mathbf{V}_{\text{IR}})$. The interaction randomness is therefore $O(\tau \cdot k_{\text{IR}} \cdot r_{\text{IR}})$ bits. There are $O(\log(\tau \cdot k_{\text{IR}}))$ bits of decision randomness because \mathbf{V}_{IOP} samples a random $i \in [\tau \cdot k_{\text{IR}}]$.
- *Verifier running time.* \mathbf{V}_{IOP} checks all of the arrays $B_1, \dots, B_{\tau \cdot k_{\text{IR}}}$ for consistency, executes the interactive reduction k_{IR} times (each in time vt_{IR}), and verifies the final predicate of at most $\tau \cdot k_{\text{IR}}$ elements in $B_{\tau \cdot k_{\text{IR}}}$ (each in time ft_{IR}). Overall \mathbf{V}_{IOP} runs in time $\text{poly}(\tau, k_{\text{IR}}, s_{\text{IR}}, l_{\text{IR}}, \text{vt}_{\text{IR}}, \text{ft}_{\text{IR}})$.
- *Adaptivity.* \mathbf{V}_{IOP} is non-adaptive.

\square

5.2 Round-query IOPs from public-coin IPs

The interactive reduction for any public-coin IP described in Section 4.3.1 does not have bounded output length. Hence we cannot use it directly in Theorem 5.3. Nonetheless, the output length grows slowly enough to achieve weaker results. While the transformation is essentially identical to the protocol described in Construction 5.4, we describe it explicitly.

Theorem 5.7. *Let R be a relation with a public-coin IP $(\mathbf{P}_{\text{IP}}, \mathbf{V}_{\text{IP}})$ with k_{IP} rounds. Let $\tau, m \geq 1$ be parameters where m divides k_{IP} . Then R has a round-query IOP $(\mathbf{P}_{\text{IOP}}, \mathbf{V}_{\text{IOP}})$ with the parameters indicated below.*

IP for R	
Rounds	k_{IP}
Total communication	l_{IP}
Randomness per round	r_{IP}
Round-by-round soundness	β_{rbr}
Verifier running time	\mathbf{vt}_{IP}

Round-query IOP for R	
Rounds	$\tau \cdot k_{\text{IP}}$
Proof length (per round)	$O(l_{\text{IP}} \cdot 2^{\tau \cdot k_{\text{IP}}/m})$
Round queries	$O(m)$
Interaction randomness	$O(\tau \cdot r_{\text{IP}} \cdot k_{\text{IP}})$
Decision randomness	$O(\log(\tau \cdot k_{\text{IP}}/m))$
Soundness error	$\max\{1/\tau, (\tau \cdot k_{\text{IP}}/m)^{\tau \cdot k_{\text{IP}}/m} \cdot m \cdot k_{\text{IP}} \cdot \beta_{\text{rbr}}\}$
Verifier running time	$\text{poly}(2^{\tau \cdot k_{\text{IP}}/m}, l_{\text{IP}}, \mathbf{vt}_{\text{IR}})$

Construction 5.8. Let $(\mathbf{P}_{\text{IP}}, \mathbf{V}_{\text{IP}})$ be a public-coin IP with k_{IP} rounds and where the verifier message at each round is r_{IP} bits long. Let $\tau > 1$ be a parameter. On input \mathbf{x} and with witness \mathbf{w} the IOP $(\mathbf{P}_{\text{IOP}}, \mathbf{V}_{\text{IOP}})$ is as follows:

- Prover sets $S_0 := \{\emptyset\}$ (i.e., S_0 contains only the empty transcript).
- For $i = 1$ to $\tau \cdot k_{\text{IP}}/m$:
 - Prover sends S_{i-1} .
 - For $j \in [m]$:
 1. Verifier chooses and sends $\rho_{i,j} \leftarrow \{0, 1\}^{r_{\text{IP}}}$ uniformly at random.
 2. For every $\text{tr} \in S_{i-1}$, prover sends $a_{\text{tr},j} := \mathbf{P}_{\text{IP}}(\mathbf{x}, \mathbf{w}, \text{tr} \parallel \rho_{i,1} \parallel a_{\text{tr},1} \parallel \dots \parallel \rho_{i,j})$.
 - Prover sets $S_i := S_{i-1} \cup \{(\text{tr} \parallel \rho_{i,1} \parallel a_{\text{tr},1} \parallel \dots \parallel \rho_{i,m} \parallel a_{i,m})\}_{\text{tr} \in S_{i-1}}$.
- For every $i \in [\tau \cdot k_{\text{IP}}/m]$, the Prover sends $T_i := S_i$.
- Verifier accepts if and only if:
 1. *Subset consistency*: For every $i \in [\tau \cdot k_{\text{IP}}/m]$, $T_{i-1} \subseteq T_i$.
 2. *Transcript consistency*: Choose a random $i \in [\tau \cdot k_{\text{IP}}/m]$. Check that $S_{i-1} = T_{i-1}$ and $S_i = T_i$. Additionally, assert that for every $\text{tr} \in S_{i-1}$, there are some messages $a_{\text{tr},1}, \dots, a_{\text{tr},m}$ such that $(\text{tr} \parallel \rho_{i,1} \parallel a_{\text{tr},1} \parallel \dots \parallel \rho_{i,m} \parallel a_{i,m}) \in S_i$, where $\rho_{i,1}, \dots, \rho_{i,m}$ are the verifier messages sent during the i -th round of interaction.
 3. *Membership*: Every transcript $\text{tr} \in T_{\tau \cdot k_{\text{IP}}/m}$ that is complete (i.e., contains messages for all k_{IP} rounds of the IP) has $\mathbf{V}_{\text{IP}}(\mathbf{x}, \text{tr}) = 1$.

Proof of Theorem 5.7. The proofs of completeness and soundness are identical to those for Theorem 5.3, using the interactive reduction for IPs described in Theorem 4.7. We therefore only need to analyze the complexity measures of the resulting IOP.

- *Proof length.* The proof length of the protocol is $O(l_{\text{IP}} \cdot 2^{\tau \cdot k_{\text{IP}}/m})$ bits.
- *Rounds.* The protocol has $\tau \cdot k_{\text{IP}}$ rounds.
- *Round queries.* The verifier reads $O(m)$ rounds.
- *Randomness complexity.* The IOP verifier uses $O(\tau \cdot r_{\text{IP}} \cdot k_{\text{IP}})$ bits of interaction randomness and $O(\log(\tau \cdot k_{\text{IP}}/m))$ bits of decision randomness.
- *Verifier running time.* The verifier running time is $\text{poly}(2^{\tau \cdot k_{\text{IP}}/m}, l_{\text{IP}}, \mathbf{vt}_{\text{IR}})$.
- *Adaptivity.* \mathbf{V}_{IOP} is non-adaptive.

□

6 Hardness of approximation for stochastic problems

We use the transformations described in Section 5 to construct IOPs with small query complexity and show that they imply hardness of approximation results for certain stochastic problems. In Section 6.1 we define stochastic constraint satisfaction problems (SCSP) and adapt a theorem of [ACY22] to this setting. Then in Section 6.2, we establish hardness of approximation results for SCSPs.

6.1 Stochastic constraint satisfaction problems

In a constraint satisfaction problem (CSP) with l variables over an alphabet Σ , we are given a set of m constraints, each of which takes as input q variables (out of the l) and outputs a bit denoting whether the constraint is satisfied. Works on hardness of approximation of CSPs generally study the complexity of distinguishing whether there exists an assignment to the variables that satisfies all of the constraints or whether every assignment satisfies at most a constant fraction of the constraints.

We extend CSPs to the stochastic setting, where the variables are split into two types: variables chosen randomly; and variables chosen existentially.

Definition 6.1. An (Σ, k, l, m, q) -SCSP (stochastic CSP) instance Φ with alphabet Σ , k alternations, l variables per quantifier, m constraints and q constraint arity is a list of m constraints C_1, \dots, C_m . Each constraint consists of q tuples $(i_1, j_1, k_1), \dots, (i_q, j_q, k_q) \subseteq [k] \times \{a, \rho\} \times [l]$ and a function $f: \Sigma^q \rightarrow \{0, 1\}$.

An assignment $z: [k] \times \{a, \rho\} \times [l] \rightarrow \Sigma$ satisfies a constraint C if

$$f(z(i_1, j_1, k_1), \dots, z(i_q, j_q, k_q)) = 1 .$$

An instance Φ is in the language (Σ, k, l, m, q) -SCSP if for uniformly random $z_{1,\rho}: [l] \rightarrow \Sigma$ there is a choice of $z_{1,a}: [l] \rightarrow \Sigma$ such that for random $z_{2,\rho}: [l] \rightarrow \Sigma$, and so on, the probability that the assignment $z(i, j, k) := z_{i,j}(k)$ satisfies all of the constraints in Φ is greater than $1/2$.

Definition 6.2. The value of a (Σ, k, l, m, q) -SCSP instance Φ is the expected fraction of satisfied constraints if the existential variables are chosen to maximize the number of satisfied constraints.

Theorem 6.3 (implied by [ACY22]). Let R be a relation with a non-adaptive k -round public-coin IOP with alphabet Σ , per-round message length l (for both the prover and verifier messages), query complexity q , decision randomness r_{dc} , decision complexity d , and soundness error β .

Then there exists a deterministic polynomial-time reduction that maps an instance \mathfrak{x} for L to an instance Φ for $(\Sigma, k, l, 2^{r_{dc}}, q)$ -SCSP such that:

- if $\mathfrak{x} \in L$ then the value of Φ is 1;
- if $\mathfrak{x} \notin L$ then the value of Φ is at most β .

Moreover, each constraint in Φ has circuit complexity $O(d)$.

6.2 Hardness of approximation for SCSP

By combining Corollary 5.1 with Theorem 6.3, we show that approximating the value of SCSPs to within a constant factor is at least as hard as solving all languages with interactive reducibility.

Theorem 6.4 (restatement of Theorem 2). *Let R be a relation with a bounded-output-length interactive reduction with well-behaved soundness error, k_{IR} predicates and ℓ_{IR} rounds. Then there exists a deterministic polynomial-time reduction that maps an instance \mathfrak{x} for L to an SCSP instance Φ such that:*

- *if $\mathfrak{x} \in L$ then the value of Φ is 1;*
- *if $\mathfrak{x} \notin L$ then the value of Φ is at most $1/2$.*

Moreover, Φ has a binary alphabet, $\ell_{\text{IR}}(|\mathfrak{x}|) \cdot k_{\text{IR}}(|\mathfrak{x}|)$ alternations, $\text{poly}(|\mathfrak{x}|)$ variables per quantifier, $\text{poly}(|\mathfrak{x}|)$ constraints and constraint arity $O(1)$.

Similarly, by combining Corollary 5.2 with Theorem 6.3, we achieve a theorem for any IP.

Theorem 6.5 (restatement of Theorem 1). *Let $L \in \text{AM}[k]$ be a language. Then there exists a deterministic polynomial-time reduction that maps an instance \mathfrak{x} for L to an SCSP instance Φ such that:*

- *if $\mathfrak{x} \in L$ then the value of Φ is 1;*
- *if $\mathfrak{x} \notin L$ then the value of Φ is at most $1/2$.*

Moreover, Φ has a binary alphabet, $k(|\mathfrak{x}|)$ alternations, $\text{poly}(|\mathfrak{x}|)$ variables per quantifier, $\text{poly}(|\mathfrak{x}|)$ constraints and constraint arity $\max\{O(1), \frac{k(|\mathfrak{x}|)}{\log |\mathfrak{x}|}\}$.

7 From round-query IOPs to binary IOPs

We show how to transform a round-query IOP with round-query complexity q_{rnd} into a binary IOP with query complexity $O(q_{\text{rnd}})$. This adapts a result of [ACY22] that transforms an IP into a binary IOP to also apply to round-query IOPs.

Theorem 7.1. *Let IOP be a non-adaptive public-coin round-query IOP for a relation $R = \{(\mathbf{x}, \mathbf{w})\}$. Then there exists a non-adaptive public-coin IOP for R with the parameters below.*

Round-query IOP for R		→	IOP for R	
Rounds	k_{IOP}		Rounds	$O(k_{\text{IOP}})$
Proof length	l_{IOP}	Alphabet size	2	
Round queries	q_{rnd}	Proof length	$\text{poly}(\mathbb{X} , l_{\text{IOP}})$	
Interaction randomness	r_{int}	Queries	$O(q_{\text{rnd}})$	
Decision randomness	r_{dc}	Interaction randomness	$\text{poly}(\mathbb{X} , r_{\text{int}})$	
Soundness error	$O(1)$	Decision randomness	$O(\log \mathbb{X} + r_{\text{dc}})$	
Verifier running time	vt_{IOP}	Soundness error	$O(1)$	
		Verifier running time	$\text{poly}(\text{vt}_{\text{IOP}})$	

The transformation of [ACY22] is obtained via two transformations:

- *Local access to interaction randomness.* The IP is transformed into an IOP in which the verifier reads only $O(1)$ bits from each of its random messages but still reads the prover's messages in their entirety. The transformation described in [ACY22] does not immediately apply to our setting (i.e., if we start out with a round-query IOP rather than an IP). In Section 7.1 we show how to adapt their proof to our setting.
- *Local access to prover messages.* The IOP where the verifier has local access to its interaction randomness is transformed into an IOP in which the verifier queries only a few bits of the entire transcript. In our case, we begin with an IOP in which the verifier has local access to its interaction randomness and only reads a small number of rounds overall. We discuss why the transformation of [ACY22] is sufficient for our setting in Section 7.2.

Putting together the two (modified) transformations yields Theorem 7.1.

7.1 Local access to interaction randomness

This transformation maps an IP into an IOP where the verifier reads the prover's messages exactly in their entirety but reads only a few bits of each one of its own random messages. In more detail, we prove that a k_{IOP} -round round-query public-coin IOP $(\mathbf{P}_{\text{IOP}}, \mathbf{V}_{\text{IOP}})$ with round-query complexity q_{rnd} and constant soundness error can be transformed into a $O(k_{\text{IOP}})$ -round round-query public-coin IOP $(\mathbf{P}'_{\text{IOP}}, \mathbf{V}'_{\text{IOP}})$ with round-query complexity $O(q_{\text{rnd}})$ and constant soundness error, in which the verifier reads $O(q_{\text{rnd}})$ bits from its interaction randomness, but still reads $O(q_{\text{rnd}})$ messages sent by the prover in their entirety.

Construction 7.2. On input \mathbf{x} , with parameters $n_z, n_s \in \mathbb{N}$ and $\gamma \in (0, 1)$ the protocol $(\mathbf{P}'_{\text{IOP}}, \mathbf{V}'_{\text{IOP}})$ works as follows, given an extractor $\text{Ext}: \{0, 1\}^{n_z} \times \{0, 1\}^{n_s} \rightarrow \{0, 1\}^{r_{\text{IOP}}}$ with error ε_{Ext} . The parameters $n_z, n_s, \varepsilon_{\text{Ext}}$ and γ will be fixed during the analysis.

1. Since the IOP $(\mathbf{P}_{\text{IOP}}, \mathbf{V}_{\text{IOP}})$ has constant soundness error, it has $(c^{1/k}, O(1))$ -round-by-round soundness error for some constant $0 < c < 1$. Using Lemma 3.3, augment $(\mathbf{P}_{\text{IOP}}, \mathbf{V}_{\text{IOP}})$ such that it has $(1/(|\mathbb{x}| + k_{\text{IOP}}^2), \gamma/4)$ -round-by-round soundness error, and interaction randomness complexity $r'_{\text{IOP}} = \text{poly}(|\mathbb{x}| + k_{\text{IOP}})$ and decision randomness complexity $r'_{\text{dc}} = O(r_{\text{dc}})$, where r_{dc} is the decision randomness of $(\mathbf{P}_{\text{IOP}}, \mathbf{V}_{\text{IOP}})$.
2. For $j = 1, \dots, k_{\text{IOP}}$:
 - (a) \mathbf{V}'_{IOP} : Send to the prover a random string $z_j \leftarrow \{0, 1\}^{n_z}$.
 - (b) $\mathbf{P}'_{\text{IOP}}(\mathbb{x}, \mathbb{w}, z_1, s_1, \dots, z_j)$: Respond with $z'_j \in \{0, 1\}^{n_z}$ where (honestly) $z'_j := z_j$.
 - (c) \mathbf{V}'_{IOP} : Send to the prover a random seed $s_j \leftarrow \{0, 1\}^{n_s}$.
 - (d) $\mathbf{P}'_{\text{IP}}(\mathbb{x}, \mathbb{w}, z_1, s_1, \dots, z_j, s_j)$:
 - i. Compute $\rho_j := \text{Ext}(z'_j, s_j)$.
 - ii. Compute $a_j \leftarrow \mathbf{P}_{\text{IOP}}(\mathbb{x}, \mathbb{w}, \rho_1, \dots, \rho_j)$.
 - iii. Send (s_j, a_j) to the verifier.
3. $\mathbf{V}'_{\text{IOP}}{}^{\text{tr}}(\mathbb{x})$ where $\text{tr} = (z_1, z'_1, s_1, s'_1, a_1, \dots, z_{k_{\text{IOP}}}, z'_{k_{\text{IOP}}}, s_{k_{\text{IOP}}}, s'_{k_{\text{IOP}}}, a_{k_{\text{IOP}}})$. Do the following:
 - (a) Sample decision randomness $\rho_{\text{dc}} \leftarrow \{0, 1\}^{r_{\text{dc}}}$ for \mathbf{V}_{IOP} and sample random indices $m_z \leftarrow [n_z]$ and $m_s \leftarrow [n_s]$.
 - (b) Emulate $\mathbf{V}_{\text{IOP}}^{\rho_1, a_1, \dots, \rho_{k_{\text{IOP}}}, a_{k_{\text{IOP}}}}(\mathbb{x}; \rho_{\text{dc}})$ where queries are answered as follows:
 - Any query made by \mathbf{V}_{IOP} to a prover message a_i is answered by querying the corresponding prover message sent during interaction.
 - If \mathbf{V}_{IOP} attempts to read ρ_j then: Check that $z_j[m_z] = z'_j[m_z]$ and $s_j[m_s] = s'_j[m_s]$. If the check fails then reject. Otherwise, pass the string $\rho_j := \text{Ext}(z'_j, s'_j)$ to \mathbf{V}_{IOP} .
 - (c) Accept if and only if $\mathbf{V}_{\text{IOP}}^{\rho_1, a_1, \dots, \rho_{k_{\text{IOP}}}, a_{k_{\text{IOP}}}}(\mathbb{x}; \rho_{\text{dc}}) = 1$.

Completeness. Fix $(\mathbb{x}, \mathbb{w}) \in R$. Let tr be a transcript generated in a random execution of the protocol. Since the prover is honest, we have that $z'_j = z_j$ and $s'_j = s_j$, and so the verifier's checks in Item 3b pass with probability 1. Moreover, since the original IP has perfect completeness, and for every j , $a_j := \mathbf{P}_{\text{IOP}}(\mathbb{x}, \mathbb{w}, \rho_1, \dots, \rho_j)$, we have that (always) $\mathbf{V}_{\text{IOP}}^{\rho_1, a_1, \dots, \rho_{k_{\text{IOP}}}, a_{k_{\text{IOP}}}}(\mathbb{x}) = 1$. Therefore, $\mathbf{V}'_{\text{IOP}}{}^{\text{tr}}(\mathbb{x})$ accepts with probability 1.

Soundness. Fix $\mathbb{x} \notin L(R)$ and a malicious prover $\tilde{\mathbf{P}}_{\text{IOP}}$. Let \mathbf{E} be the event over the verifier's random coins, both interaction and decision, $(z_1, s_1, \dots, z_{k_{\text{IOP}}}, s_{k_{\text{IOP}}}, \rho_{\text{dc}})$, that there exists some j where the emulated the verifier \mathbf{V}_{IOP} reads ρ_j and at least one of the following is true: (i) $\Delta(z'_j, z_j) \geq \gamma$ or; (ii) $\Delta(s'_j, s_j) \geq \gamma$, where $z'_j := \tilde{\mathbf{P}}_{\text{IOP}}(z_1, s_1, \dots, z_{j-1}, s_{j-1}, z_j)$ and $(s'_j, a_j) := \tilde{\mathbf{P}}_{\text{IOP}}(z_1, s_1, \dots, z_{j-1}, s_{j-1}, z_j, s_j)$.

We first show that the probability that the verifier accepts and the event \mathbf{E} happens is small:

Claim 7.3. *We have that:*

$$\Pr \left[\langle \tilde{\mathbf{P}}_{\text{IOP}}, \mathbf{V}'_{\text{IOP}}(\mathbb{x}) \rangle = 1 \wedge (z_1, s_1, \dots, z_{k_{\text{IOP}}}, s_{k_{\text{IOP}}}, \rho_{\text{dc}}) \in \mathbf{E} \right] \leq 1 - \gamma .$$

Proof. For every choice of verifier randomness $(z_1, s_1, \dots, z_{k_{\text{IOP}}}, s_{k_{\text{IOP}}}, \rho_{\text{dc}}) \in \mathbf{E}$, there exists some j where the emulated the verifier \mathbf{V}_{IOP} reads ρ_j in which either $\Delta(z_j, z'_j) \geq \gamma$ or $\Delta(s_j, s'_j) \geq \gamma$. As a result, one of the checks made by \mathbf{V}'_{IOP} in Item 3b, causes the verifier to reject with probability at least γ . We conclude the claim by noting that

$$\begin{aligned} & \Pr \left[\langle \tilde{\mathbf{P}}_{\text{IOP}}, \mathbf{V}'_{\text{IOP}}(\mathbb{x}) \rangle = 1 \wedge (z_1, s_1, \dots, z_{k_{\text{IOP}}}, s_{k_{\text{IOP}}}, \rho_{\text{dc}}) \in \mathbf{E} \right] \\ & \leq \Pr \left[\langle \tilde{\mathbf{P}}_{\text{IOP}}, \mathbf{V}'_{\text{IOP}}(\mathbb{x}) \rangle = 1 \mid (z_1, s_1, \dots, z_{k_{\text{IOP}}}, s_{k_{\text{IOP}}}, \rho_{\text{dc}}) \in \mathbf{E} \right] \leq 1 - \gamma . \end{aligned}$$

□

We introduce a new malicious prover $\tilde{\mathbf{P}}_{\text{IOP}}^*$ for the IOP system that “corrects” any z'_j and s'_j messages sent by the prover that are γ -far from what they are claimed to be. $\tilde{\mathbf{P}}_{\text{IOP}}^*$ has the following next-message function:

- $\tilde{\mathbf{P}}_{\text{IOP}}^*(z_1, s_1, \dots, z_{j-1}, s_{j-1}, z_j)$: Compute $z''_j := \tilde{\mathbf{P}}_{\text{IOP}}(z_1, s_1, \dots, z_{j-1}, s_{j-1}, z_j)$. If $\Delta(z''_j, z_j) < \gamma$ then output z''_j and otherwise output z_j .
- $\tilde{\mathbf{P}}_{\text{IOP}}^*(z_1, s_1, \dots, z_{j-1}, s_{j-1}, z_j, s_j)$: Compute $(s''_j, a_j) := \tilde{\mathbf{P}}_{\text{IOP}}(z_1, s_1, \dots, z_{j-1}, s_{j-1}, z_j, s_j)$. If $\Delta(s''_j, s_j) < \gamma$ then output (s''_j, a_j) and otherwise output (s_j, a_j) .

Notice that $\tilde{\mathbf{P}}_{\text{IOP}}^*$ and $\tilde{\mathbf{P}}_{\text{IOP}}$ send exactly the same messages in rounds where $\tilde{\mathbf{P}}_{\text{IOP}}$ outputs z'_j and s'_j with $\Delta(z'_j, z_j) < \gamma$ and $\Delta(s'_j, s_j) < \gamma$. Therefore we can reinterpret the event \mathbf{E} relative to $\tilde{\mathbf{P}}_{\text{IOP}}^*$ as follows: $(z_1, s_1, \dots, z_{k_{\text{IOP}}}, s_{k_{\text{IOP}}}, \rho_{\text{dc}}) \notin \mathbf{E}$ if for every j where \mathbf{V}_{IOP} reads ρ_j and $\tilde{\mathbf{P}}_{\text{IOP}}^*$ has not “corrected” z'_j or s'_j (by sending z_j or s_j respectively). Whenever the verifier does not read a corrected round, it acts identically to an interaction with $\tilde{\mathbf{P}}_{\text{IP}}$ where it does not read a round that was far from the messages sent by the verifier. Therefore, using this interpretation we have:

$$\begin{aligned} & \Pr \left[\langle \tilde{\mathbf{P}}_{\text{IOP}}, \mathbf{V}'_{\text{IOP}}(\mathbb{x}) \rangle = 1 \wedge (z_1, s_1, \dots, z_{k_{\text{IOP}}}, s_{k_{\text{IOP}}}, \rho_{\text{dc}}) \notin \mathbf{E} \right] \\ & = \Pr \left[\langle \tilde{\mathbf{P}}_{\text{IOP}}^*, \mathbf{V}'_{\text{IOP}}(\mathbb{x}) \rangle = 1 \wedge (z_1, s_1, \dots, z_{k_{\text{IOP}}}, s_{k_{\text{IOP}}}, \rho_{\text{dc}}) \notin \mathbf{E} \right] \\ & \leq \Pr \left[\langle \tilde{\mathbf{P}}_{\text{IOP}}^*, \mathbf{V}'_{\text{IOP}}(\mathbb{x}) \rangle = 1 \right] \end{aligned}$$

We now show that the probability of \mathbf{V}'_{IOP} accepting when interacting with $\tilde{\mathbf{P}}_{\text{IOP}}^*$ is small. We begin by showing that $\tilde{\mathbf{P}}_{\text{IOP}}^*$ -s messages z_j have high min-entropy.

Claim 7.4. *For every j : $H_{\min}(Z'_j \mid \neg \mathbf{E}) \geq 0.5n_z$, where Z'_j is the random variable describing the output z'_j of $\tilde{\mathbf{P}}_{\text{IP}}^*$ in a random execution of $\langle \tilde{\mathbf{P}}_{\text{IOP}}^*, \mathbf{V}'_{\text{IOP}}(\mathbb{x}) \rangle$.*

Proof. Fix a round number j and some string z_j^* . We have that

$$\Pr \left[Z'_j = z_j^* \right] \leq \Pr_{z_j} \left[\Delta(z_j^*, z_j) < \gamma \right] \tag{6}$$

$$\begin{aligned} & = |\{ x' \in \{0, 1\}^{n_z} : \Delta(x, x') \leq \gamma \}| / 2^{n_z} \\ & \leq 2^{-n_z + n_z H(\gamma)} \end{aligned} \tag{7}$$

$$< 2^{-0.5n_z} . \tag{8}$$

Above, Equation (6) is due to the fact the output z'_j of $\tilde{\mathbf{P}}_{\text{IP}}^*$ always has Hamming distance at less than γ from z_j . Equation (7) true due to Fact 3.8 and Equation (8) is true for a small enough constant γ . Then, we get that

$$H_{\min}(Z'_j) \triangleq \min_{z_j^*} -\log \Pr[Z'_j = z_j^*] > 0.5n_z .$$

□

Claim 7.5. *Suppose that, after amplification in Item 1, the IOP $(\mathbf{P}_{\text{IOP}}, \mathbf{V}_{\text{IOP}})$ has $(\beta_{\text{rbr}}, \delta_{\text{dc}})$ -round-by-round soundness with state function \mathbf{state} . Then for every transcript \mathbf{tr} generated by an interaction between $\tilde{\mathbf{P}}_{\text{IOP}}^*$ and \mathbf{V}'_{IOP} where the verifier is about to make its j -th move such that $\mathbf{state}(\mathbf{x}, \mathbf{tr}) = 0$:*

$$\Pr[\mathbf{state}(\mathbf{x}, \mathbf{tr} || \rho_j) = 1] \leq (\beta_{\text{rbr}} + \varepsilon_{\text{Ext}}) \cdot 2^{n_s \cdot H(\gamma)} ,$$

where ρ_j is drawn as in the protocol description.

Proof. Fix some j and a transcript as in the claim statement. By Claim 7.4, we have that z'_j has min-entropy at least $0.5n_z$. Thus, by definition of the extractor,

$$| \Pr[\mathbf{state}(\mathbf{x}, \mathbf{tr} || \text{Ext}(z'_j, U_{n_s})) = 1] - \Pr[\mathbf{state}(\mathbf{x}, \mathbf{tr} || U_{r'_{\text{IP}}}) = 1] | \leq \varepsilon_{\text{Ext}} ,$$

where U_{n_s} and $U_{r'_{\text{IP}}}$ are the uniform distributions over bit strings of length n_s and r'_{IP} respectively. Furthermore, by $(\beta_{\text{rbr}}, \delta_{\text{dc}})$ -round-by-round soundness of $(\mathbf{P}_{\text{IOP}}, \mathbf{V}_{\text{IOP}})$, we have that

$$\Pr[\mathbf{state}(\mathbf{x}, \mathbf{tr} || U_{r'_{\text{IP}}}) = 1] < \beta_{\text{rbr}} .$$

Therefore, the fraction of seeds that cause the state function to change from 0 to 1 is at most $\varepsilon_{\text{Ext}} + \beta_{\text{rbr}}$. Recall that in the protocol, $\rho_j := \text{Ext}(z'_j, s'_j)$, i.e., the seed of the extractor is s'_j rather than a uniformly random seed. Recall that $\tilde{\mathbf{P}}_{\text{IOP}}^*$ only outputs messages s'_j with $\Delta(s'_j, s_j) < \gamma$. We say that a seed s_j is *bad* if there exists some s'_j with $\Delta(s'_j, s_j) < \gamma$ such that $\mathbf{state}(\mathbf{x}, \mathbf{tr} || \text{Ext}(z'_j, s'_j)) = 1$. Every point s'_j that inhibits changing of the state function has a ball of size $2^{n_s \cdot H(\gamma)}$ of random seeds that have distance at most γ from it (see Fact 3.8). The total probability of landing on a bad seed is at most the probability that a random seed s_j falls within one of these balls. Therefore the probability that s_j bad is at most $(\varepsilon_{\text{Ext}} + \beta_{\text{rbr}}) \cdot 2^{n_s \cdot H(\gamma)}$. □

We set the extractor error to be $\varepsilon_{\text{Ext}} = \beta_{\text{rbr}}$, on a source with min entropy $0.5n_z$. We set $n_z = 4r'_{\text{IOP}}$ such that the source has min entropy $2r'_{\text{IOP}}$, and thus we extract r'_{IOP} random bits with error ε_{Ext} . The seed length is $n_s = O(\log 1/\varepsilon_{\text{Ext}}) = O(\log(1/\beta_{\text{rbr}}))$. We therefore have that:

$$\Pr[\langle \tilde{\mathbf{P}}_{\text{IOP}}^*, \mathbf{V}'_{\text{IP}}(\mathbf{x}) \rangle = 1] \leq \delta_{\text{dc}} + \Pr[\exists j : \mathbf{state}(\mathbf{x}, \mathbf{tr} || \rho_j) = 1] \tag{9}$$

$$\leq \delta_{\text{dc}} + k_{\text{IOP}} \cdot (\beta_{\text{rbr}} + \varepsilon_{\text{Ext}}) \cdot 2^{n_s \cdot H(\gamma)} \tag{10}$$

$$\leq \delta_{\text{dc}} + k_{\text{IOP}} \cdot 2\beta_{\text{rbr}} \cdot 2^{O(\log(1/\beta_{\text{rbr}})) \cdot H(\gamma)} \tag{11}$$

$$\leq \delta_{\text{dc}} + k_{\text{IOP}} \cdot \sqrt{\beta_{\text{rbr}}} \tag{12}$$

$$\leq \gamma/4 + k_{\text{IOP}} / \sqrt{|\mathbf{x}|} + k_{\text{IOP}}^2 < \gamma/2 . \tag{13}$$

Equation (9) follows from the fact that the verifier \mathbf{V}'_{IP} accepts only if \mathbf{V}_{IP} accepts given \mathbf{x} , prover messages $a_1, \dots, a_{k_{\text{IOP}}}$ and verifier randomness $\rho_1, \dots, \rho_{k_{\text{IOP}}}$. By the round-by-round soundness of

the original IP, since $\text{state}(\mathbb{x}, \emptyset) = 0$ (which follows from the fact that $\mathbb{x} \notin L$), in order for the verifier to accept, it must be that the value of the state function changed from 0 to 1 in some round. Equation (10) is true by applying the union bound and Claim 7.5. We have Equation (11) by noting that we set $\varepsilon_{\text{Ext}} = \beta_{\text{rbr}}$ and $n_s = O(\log(1/\beta_{\text{rbr}}))$. Equation (12) holds for a small enough constant $\gamma > 0$, and, finally, Equation (13) holds by the definitions of δ_{dc} and β_{rbr} and for large enough values of $|\mathbb{x}|$.

Thus we have that

$$\Pr \left[\langle \tilde{\mathbf{P}}_{\text{IOP}}, \mathbf{V}'_{\text{IOP}}(\mathbb{x}) \rangle = 1 \wedge (z_1, s_1, \dots, z_{k_{\text{IOP}}}, s_{k_{\text{IOP}}}, \rho_{\text{dc}}) \notin \mathbf{E} \right] \leq \Pr \left[\langle \tilde{\mathbf{P}}_{\text{IOP}}^*, \mathbf{V}'_{\text{IOP}}(\mathbb{x}) \rangle = 1 \right] \leq \gamma/2 .$$

Putting this fact together with Claim 7.3, we have that

$$\begin{aligned} \Pr \left[\langle \tilde{\mathbf{P}}_{\text{IOP}}, \mathbf{V}'_{\text{IOP}}(\mathbb{x}) \rangle = 1 \right] &= \Pr \left[\langle \tilde{\mathbf{P}}_{\text{IOP}}, \mathbf{V}'_{\text{IOP}}(\mathbb{x}) \rangle = 1 \wedge (z_1, s_1, \dots, z_{k_{\text{IOP}}}, s_{k_{\text{IOP}}}, \rho_{\text{dc}}) \in \mathbf{E} \right] \\ &\quad + \Pr \left[\langle \tilde{\mathbf{P}}_{\text{IOP}}, \mathbf{V}'_{\text{IOP}}(\mathbb{x}) \rangle = 1 \wedge (z_1, s_1, \dots, z_{k_{\text{IOP}}}, s_{k_{\text{IOP}}}, \rho_{\text{dc}}) \notin \mathbf{E} \right] \\ &\leq 1 - \gamma + \gamma/2 = 1 - \gamma/2 . \end{aligned}$$

Thus the verifier accepts with constant probability.

Complexity measures. We analyze the efficiency parameters of the resulting round-query IOP:

- *Rounds.* The round-query IOP has $2k_{\text{IOP}}$ rounds.
- *Proof length.* We first amplify the protocol, giving polynomial overhead to all messages. In addition to the original prover messages, the prover also sends z'_j and s'_j . Therefore the proof length is $\text{poly}(|\mathbb{x}|, l_{\text{IOP}})$.
- *Round queries.* \mathbf{V}'_{IOP} queries $2 \cdot \mathbf{q}_{\text{rnd}}$ rounds.
- *Query complexity to randomness.* The verifier queries s_j and z_j in $O(1)$ locations.
- *Randomness complexity.* \mathbf{V}'_{IP} generates $n_z + n_s = \text{poly}(r_{\text{IP}}, |\mathbb{x}|)$ bits in every round.
- *Decision randomness.* \mathbf{V}'_{IP} chooses decision randomness $r'_{\text{dc}} = O(r_{\text{dc}})$ and then uses $\log n_z + \log n_s = O(\log |\mathbb{x}| + \log k_{\text{IOP}})$ bits of decision randomness. All-in-all $O(r_{\text{dc}} + \log |\mathbb{x}|)$.
- *Verifier running time.* \mathbf{V}'_{IP} runs the original IOP verifier for polynomially many repetitions, generates a few random strings and runs the extractor. Its running time is therefore polynomially related to the running time of \mathbf{V}_{IP} .
- *Adaptivity.* \mathbf{V}'_{IP} makes non-adaptive queries to its interaction randomness. Therefore, if \mathbf{V}_{IOP} was non-adaptive, then so is \mathbf{V}'_{IOP} .

7.2 Local access to prover messages

In this transformation, we take a round-query IOP $(\mathbf{P}_{\text{IOP}}, \mathbf{V}_{\text{IOP}})$ with round-query complexity \mathbf{q}_{rnd} where the verifier reads $O(1)$ bits from the interaction randomness of the rounds that it queries but reads the entire prover messages, and transform it into a binary IOP $(\mathbf{P}'_{\text{IOP}}, \mathbf{V}'_{\text{IOP}})$ with query complexity $O(\mathbf{q}_{\text{rnd}})$. We sketch this transformation, assuming familiarity with [ACY22].

The transformation involves simulating the round-query IOP, except that \mathbf{P}'_{IOP} encodes the messages of \mathbf{P}_{IOP} using an “index-decodable” PCP (ID-PCP). In the final round, \mathbf{P}'_{IOP} sends, for every choice of decision randomness of \mathbf{V}_{IOP} , a proof convincing that \mathbf{V}_{IOP} would have accepted, having chosen this decision randomness. The new IOP verifier \mathbf{V}'_{IOP} then chooses decision randomness and runs the ID-PCP verifier to verify the proof. Executing the ID-PCP verifier involves reading the relevant bits of the verifier interaction randomness, $O(1)$ bits of the final prover message, and $O(1)$ bits from the encoding of each prover message that \mathbf{V}_{IOP} would have queried using this decision randomness. Hence, \mathbf{V}'_{IOP} has (total) query complexity $O(q_{\text{rnd}})$.

Acknowledgments

Gal Arnon is supported in part by a grant from the Israel Science Foundation (no. 2686/20) and by the Simons Foundation Collaboration on the Theory of Algorithmic Fairness. Alessandro Chiesa is funded by the Ethereum Foundation. Eylon Yogev is supported by the BIU Center for Research in Applied Cryptography and Cyber Security in conjunction with the Israel National Cyber Bureau in the Prime Minister’s Office, and by the Alter Family Foundation.

References

- [AB17] Amir Abboud and Arturs Backurs. “Towards Hardness of Approximation for Polynomial Time Problems”. In: *Proceedings of the 8th Innovations in Theoretical Computer Science Conference*. ITCS ’17. 2017, 11:1–11:26.
- [ACY22] Gal Arnon, Alessandro Chiesa, and Eylon Yogev. *A PCP Theorem for Interactive Proofs*. Cryptology ePrint Archive, Report 2021/915. To appear at EUROCRYPT 2022. 2022.
- [ALMSS98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. “Proof verification and the hardness of approximation problems”. In: *Journal of the ACM* 45.3 (1998). Preliminary version in FOCS ’92., pp. 501–555.
- [AR18] Amir Abboud and Aviad Rubinfeld. “Fast and Deterministic Constant Factor Approximation Algorithms for LCS Imply New Circuit Lower Bounds”. In: *Proceedings of the 9th Innovations in Theoretical Computer Science Conference*. ITCS ’18. 2018, 35:1–35:14.
- [ARW17] Amir Abboud, Aviad Rubinfeld, and R. Ryan Williams. “Distributed PCP Theorems for Hardness of Approximation in P”. In: *Proceedings of the 59th Annual IEEE Symposium on Foundations of Computer Science*. FOCS ’17. 2017, pp. 25–36.
- [AS98] Sanjeev Arora and Shmuel Safra. “Probabilistic checking of proofs: a new characterization of NP”. In: *Journal of the ACM* 45.1 (1998). Preliminary version in FOCS ’92., pp. 70–122.
- [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.
- [BFLS91] László Babai, Lance Fortnow, Leonid A. Levin, and Mario Szegedy. “Checking computations in polylogarithmic time”. In: *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing*. STOC ’91. 1991, pp. 21–32.
- [BM88] László Babai and Shlomo Moran. “Arthur-Merlin Games: A Randomized Proof System, and a Hierarchy of Complexity Classes”. In: *Journal of Computer and System Sciences* 36.2 (1988), pp. 254–276.
- [Bab85] László Babai. “Trading group theory for randomness”. In: *Proceedings of the 17th Annual ACM Symposium on Theory of Computing*. STOC ’85. 1985, pp. 421–429.
- [CCHLRR18] Ran Canetti, Yilei Chen, Justin Holmgren, Alex Lombardi, Guy N. Rothblum, and Ron D. Rothblum. *Fiat-Shamir From Simpler Assumptions*. Cryptology ePrint Archive, Report 2018/1004. 2018.
- [CFLS97] Anne Condon, Joan Feigenbaum, Carsten Lund, and Peter W. Shor. “Random Debaters and the Hardness of Approximating Stochastic Functions”. In: *SIAM Journal on Computing* 26.2 (1997), pp. 369–400.
- [CGLRR19] Lijie Chen, Shafi Goldwasser, Kaifeng Lyu, Guy N. Rothblum, and Aviad Rubinfeld. “Fine-grained Complexity Meets $IP = PSPACE$ ”. In: *Proceedings of the 30th Annual Symposium on Discrete Algorithms*. SODA ’19. 2019, pp. 1–20.

- [CW19] Lijie Chen and Ryan Williams. “An Equivalence Class for Orthogonal Vectors”. In: *Proceedings of the 30th Annual Symposium on Discrete Algorithms*. SODA ’19. 2019, pp. 21–40.
- [Din07] Irit Dinur. “The PCP theorem by gap amplification”. In: *Journal of the ACM* 54.3 (2007), p. 12.
- [Dru11] Andrew Drucker. “A PCP Characterization of AM”. In: *Proceedings of the 38th International Colloquium on Automata, Languages and Programming*. ICALP ’11. 2011, pp. 581–592.
- [FGLSS91] Uriel Feige, Shafi Goldwasser, László Lovász, Shmuel Safra, and Mario Szegedy. “Approximating clique is almost NP-complete (preliminary version)”. In: *Proceedings of the 32nd Annual Symposium on Foundations of Computer Science*. SFCS ’91. 1991, pp. 2–12.
- [FGLSS96] Uriel Feige, Shafi Goldwasser, Laszlo Lovász, Shmuel Safra, and Mario Szegedy. “Interactive proofs and the hardness of approximating cliques”. In: *Journal of the ACM* 43.2 (1996), pp. 268–292.
- [GKR15] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. “Delegating Computation: Interactive Proofs for Muggles”. In: *Journal of the ACM* 62.4 (2015), 27:1–27:64.
- [GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. “The knowledge complexity of interactive proof systems”. In: *SIAM Journal on Computing* 18.1 (1989). Preliminary version appeared in STOC ’85., pp. 186–208.
- [GUV09] Venkatesan Guruswami, Christopher Umans, and Salil P. Vadhan. “Unbalanced expanders and randomness extractors from Parvaresh–Vardy codes”. In: *Journal of the ACM* 56.4 (2009), 20:1–20:34.
- [GVW02] Oded Goldreich, Salil Vadhan, and Avi Wigderson. “On interactive proofs with a laconic prover”. In: *Computational Complexity* 11.1/2 (2002), pp. 1–53.
- [HRT07] Ishay Haviv, Oded Regev, and Amnon Ta-Shma. “On the Hardness of Satisfiability with Bounded Occurrences in the Polynomial-Time Hierarchy”. In: *Theory of Computing* 3.1 (2007), pp. 45–60.
- [KL94] Ker-I Ko and Chih-Long Lin. “Non-approximability in the polynomial-time hierarchy”. In: *Technical Report 94-2, Dept. of Computer Science, SUNY at Stony Brook* (1994).
- [KR08] Yael Kalai and Ran Raz. “Interactive PCP”. In: *Proceedings of the 35th International Colloquium on Automata, Languages and Programming*. ICALP ’08. 2008, pp. 536–547.
- [LFKN92] Carsten Lund, Lance Fortnow, Howard J. Karloff, and Noam Nisan. “Algebraic Methods for Interactive Proof Systems”. In: *Journal of the ACM* 39.4 (1992), pp. 859–868.
- [Pap83] Christos H. Papadimitriou. “Games Against Nature (Extended Abstract)”. In: *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*. STOC ’83. 1983, pp. 446–450.
- [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.
- [Sha92] Adi Shamir. “IP = PSPACE”. In: *Journal of the ACM* 39.4 (1992), pp. 869–877.