

Efficiently Batching Unambiguous Interactive Proofs

Bonnie Berger, Rohan Goyal, Matthew M. Hong, and Yael Tauman Kalai

MIT

Abstract

We show that if a language $\mathcal{L}$ admits a public-coin unambiguous interactive proof (UIP) with round complexity ℓ , where a bits are communicated per round, then the *batch language* $\mathcal{L}^{\otimes k}$, i.e. the set of k -tuples of statements all belonging to $\mathcal{L}$, has an unambiguous interactive proof with round complexity $\ell \cdot \text{polylog}(k)$, per-round communication of $a \cdot \ell \cdot \text{polylog}(k) + \text{poly}(\ell)$ bits, assuming the verifier in the UIP has depth bounded by $\text{polylog}(k)$. Prior to this work, the best known batch UIP for $\mathcal{L}^{\otimes k}$ required communication complexity at least $(\text{poly}(a) \cdot k^\epsilon + k) \cdot \ell^{1/\epsilon}$ for any arbitrarily small constant $\epsilon > 0$ (Reingold-Rothblum-Rothblum, STOC 2016).

As a corollary of our result, we obtain a *doubly efficient proof system*, that is, a proof system whose proving overhead is polynomial in the time of the underlying computation, for any language computable in polynomial space and in time at most $n^{O(\sqrt{\frac{\log n}{\log \log n}})}$. This expands the state of the art of doubly efficient proof systems: prior to our work, such systems were known for languages computable in polynomial space and in time $n^{(\log n)^\delta}$ for a small $\delta > 0$ significantly smaller than $1/2$ (Reingold-Rothblum-Rothblum, STOC 2016).

Contents

1	Introduction	1
1.1	Our Results	2
1.2	Additional Related Work	3
2	Technical Overview	4
2.1	Creating Distance	6
2.2	Instance Reduction	8
2.3	Efficiency Gain over the RRR16 Protocol	11
2.4	Doubly Efficient Interactive Proofs	12
3	Preliminaries	12
3.1	Low Degree Extensions and Polynomial Valuation (PVAL)	13
3.2	Succinct Descriptions of Sets and Predicates	15
3.3	Unambiguous Interactive Proof (UIP)	15
3.4	Almost d -wise independent permutations	19
3.5	Unambiguous Interactive Proofs of Proximity (IPP)	21
4	Our Batch-UIP	21
4.1	Our Ingredients	22
4.2	Our Construction of Batch-UIP	27
4.3	Proof of Lemma 3: Generating Δ_c -Distance via Checksums.	31
4.4	Proof of Lemma 4: Instance Reduction for Δ_c -Distance.	38
4.5	Analysis of Other Sub-protocols	43
5	Doubly-Efficient Interactive Proofs	45
5.1	A New UIP for Verifying Space-Bounded Computations	45
6	IPP for PVAL with Column Distance	50
6.1	Overview of the RR IPP for PVAL.	51
6.2	Gap Amplification Lemmas for Column Distance	53
6.3	Efficient IPP for PVAL with Column Distance	58
6.4	Succinct Descriptions	68
6.5	Construction of the Verifier's Verification Circuit	69

1 Introduction

Verification is one of the most fundamental concepts in computer science and is the basis for the definition of the complexity class NP. In the mid-eighties, a flurry of works expanded the notion of proofs beyond the classical notion of a mathematical proof to interactive proofs [GMR89, BM88], multi-prover interactive proofs [BGKW88, BFL90], probabilistically checkable proofs (PCPs) [BFLS91, FGL⁺91, AS92, ALM⁺92], interactive PCPs [KR08], and interactive oracle proofs [BCS16, RRR16]. This line of work has proven instrumental in cryptography and complexity theory, leading to breakthroughs in hardness of approximation and to the emergence of fundamental concepts in cryptography, such as zero-knowledge proofs [GMR89], which underlie many cryptographic primitives today. It is also instrumental in constructing “succinct proofs” used in many blockchain applications.

Interactive proofs Interactive proof systems were shown to be extremely powerful in the celebrated $\text{IP} = \text{PSPACE}$ theorem [LFKN92, Sha92]. Specifically, it was proven that the correctness of any time- T space- S computation could be verified by a $\text{poly}(S, n)$ time verifier, via an interactive proof. The main drawback of this result is that the time required by the prover to convince the verifier of the correctness is $2^{O(S \cdot \log S)}$, rendering it impractical for real-world applications.¹

Doubly efficient interactive proofs. The work of [GKR15] initiated the study of *doubly efficient* interactive proofs, in which the prover’s runtime is required to be at most $p(T)$, where T is the time needed to carry out the underlying computation and p is a polynomial independent of T , while the verifier runs in time significantly less than T . They showed that the correctness of any computation that can be performed by a (log-space uniform) circuit of depth D and size T can be proven via an interactive proof whose communication complexity is $D \cdot \text{polylog}(T)$, the verifier runs in time $D \cdot \text{polylog}(T) + \tilde{O}(n)$, and the prover runs in time $\text{poly}(T)$.² In a breakthrough result, Reingold, Rothblum, and Rothblum [RRR16] constructed a doubly efficient (constant round) interactive proof for every language in time $T = n^{(\log n)^\delta}$ and polynomial space, for a sufficiently small constant $\delta > 0$. Since this work which was posted nearly 10 years ago, no improvements have been made to this fundamental problem. One of the building blocks in [RRR16] is a doubly efficient *batch unambiguous interactive proof*, which is the focus of our work.

Doubly efficient batch interactive proofs. Another important line of work is the one that focuses on doubly efficient *batch interactive proofs* [RRR16, RRR18, RR20], where the goal is to prove many statements at the price of one. Concretely, suppose we want to verify statements $x_1, \dots, x_k$, each of which belongs to a language that admits an interactive proof (i.e., a language in PSPACE). Observing that the total space needed to run k computations grows only by an (additive) $\text{polylog}(k)$ factor, it follows, from $\text{IP} = \text{PSPACE}$, that all k statements can also be proven via a single (batch) interactive proof with $\text{polylog}(k)$ overhead in communication complexity. This naturally raises the question:

¹We note that in those works (and at that time) the primary focus was on the power of the proof systems themselves, where the verifier was assumed to run in polynomial time, while the prover was thought of as being all-powerful. In fact, in [BM88] the all-powerful prover was even named after the famous wizard Merlin.

²In particular, this implies an improved $\text{IP} = \text{PSPACE}$ theorem where the communication complexity is $\text{poly}(S)$, the verifier runs in time $\text{poly}(S) + \tilde{O}(n)$, and the prover runs in time $2^{O(S)}$.

Which languages in IP can be batched *doubly-efficiently*, with $\text{polylog}(k)$ overhead in communication complexity, while preserving the efficiency of the prover and verifier?

The work of [BKP⁺24] proved that if a language $\mathcal{L} \in \text{NP}$ has a batch interactive proof with polylogarithmic overhead, where the prover runs in polynomial time given the witnesses (we refer to such proof as a *doubly efficient batch interactive proof*), then $\mathcal{L}$ also has a *statistical witness indistinguishable* (WI) proof system. This result can be viewed as an indication that not all of NP has doubly efficient batch interactive proofs, since it seems likely that not all NP languages admit statistical WI proofs.

Indeed, all known positive results in this regime focus either on UP , which is the class of NP languages for which each instance $x \in \mathcal{L}$ has a *unique* witness, or more generally, on the class UIP , consisting of all languages $\mathcal{L}$ that have an *unambiguous* interactive proofs.

An unambiguous interactive proof for a language $\mathcal{L}$ is an interactive proof system in which, for every $x \in \mathcal{L}$ and for every possible verifier's message there is only one answer that a prover can send that will later allow it convince the verifier to accept conditioned on this answer (with non-negligible probability). In other words, even though the prover may be all-powerful, there is at most one way for it to produce an accepting proof. We refer the reader to Section 3.3 for the formal definition.

For the case of UP , it was shown in [RR20] that there exists a doubly efficient batch interactive proof for any language $\mathcal{L} \in \text{UP}$, where proving $x_1, \dots, x_k \in \mathcal{L}$ requires communication complexity $m \cdot \text{polylog}(k)$, where m is the length of a single witness, and the prover runs in polynomial time given all the witnesses.

It was shown by [RRR16] that for every language $\mathcal{L}$ that has a public-coin UIP consisting of ℓ rounds, where in each round the prover and verifier send a bits, there exists a batch unambiguous interactive proof for proving that $x_1, \dots, x_k \in \mathcal{L}$, where the communication complexity is at least $(\text{poly}(a) \cdot k^\epsilon + k) \cdot \ell^{1/\epsilon}$, the number of rounds increase only by a constant factor (that grows exponentially with $1/\epsilon$), and the prover's and verifier's runtime increase by $\text{poly}(k, n, \ell)$, where $n = |x_i|$. It is critical that the underlying protocol being batched is public-coin³, i.e. that the verifier's messages are simply uniformly random coins.

1.1 Our Results

We construct a batch unambiguous interactive proof for any language $\mathcal{L}$ that has an ℓ -round public-coin unambiguous interactive proof, where the communication complexity of the batch interactive proof for proving $x_1, \dots, x_k \in \mathcal{L}$ grows by a factor of $\ell \cdot \text{polylog}(k)$, the round complexity grows by a factor of $\text{polylog}(k)$, and the runtime of the prover and verifier grow polynomially in ℓ, n, k where $n = |x_i|$.

Theorem 1 (Our Batch-UIP, Informal). *Let $\mathcal{L}$ be a language with a public-coin unambiguous interactive proof system that on input x of length n (and where the prover may have an additional input w) runs in $\ell = \ell(n)$ rounds, in each round $a = a(n)$ bits are sent, and the verifier's verification circuit is a log-space uniform boolean circuit of depth at most $\text{polylog}(k)$. Then there is a public-coin unambiguous interactive proof $(P_{\text{Batch-UIP}}, V_{\text{Batch-UIP}})$ for $\mathcal{L}^{\otimes k}$, that on input $(x_1, \dots, x_k)$, each of length n (and where the prover may have additional input $w_1, \dots, w_k$), runs in $\ell \cdot \text{polylog}(k)$ rounds,*

³Note that the Goldwasser-Sipser transformation from any interactive proof to a public-coin interactive proof does not preserve prover efficiency [GS86, Vad00].

where in each round $a \cdot \ell \cdot \text{polylog}(k) + \text{poly}(\ell)$ bits are sent. Moreover, the prover's computational complexity increases multiplicatively by $\text{poly}(k, \ell, n)$ and the verifier's computational complexity becomes $(\ell \cdot \text{Vtime} + kna^2 \cdot \text{poly}(\ell)) \cdot \text{polylog}(k, n, a, \ell)$, where Vtime is the verifier runtime of the base protocol.

Remark 1. The assumption that the verification circuit is log-space uniform is without loss of generality by the Cook-Levin reduction [Coo71, Lev73] since any efficient verifier, as a Turing machine, can be simulated by a log-space uniform boolean circuit of a comparable size. The additional assumption that the verifier's verdict circuit is of low depth is relatively mild, since we can generically flatten out this circuit by having the prover send the values of all the wires in the verdict circuit. This increases the communication complexity (in a single round).

We provide a high-level overview of our batch UIP protocol in Section 2. We note that it uses a component of the batch UP protocol of [RR20], which we refer to as the Instance Reduction protocol. We need to generalize this protocol to hold with respect a distance measure, which we call execution-wise distance. We elaborate on this in Section 2.2.

Applications to doubly efficient interactive proofs We show how one can use Theorem 1 to improve the state-of-the-art on doubly efficient interactive proofs.

Theorem 2 (Our Doubly-efficient Interactive Proof, Informal). *Every language computable in time $T(n) \leq n^{O(\sqrt{\frac{\log n}{\log \log n}})}$ and space $S(n) = \text{poly}(n)$ has a (public-coin) doubly efficient interactive proof, with verifier runtime and communication complexity as $\text{poly}(n) \cdot S(n)$ and prover runtime $T \cdot \text{poly}(n)$.*

Corollary 1 (Our Doubly Efficient IP, Informal). *Every language decidable in time $T(n) = n^{O(\sqrt{\frac{\log n}{\log \log n}})}$ and polynomial space has a (public-coin) doubly efficient interactive proof.*

For more details, see Theorem 8 and corollary 2. This improves upon the doubly efficient interactive proof from [RRR16], which is only for languages computed in polynomial space and time $T = n^{(\log n)^\delta}$ for a small constant $\delta > 0$ (which is significantly smaller than $1/2$).

1.2 Additional Related Work

The area of proof systems has become a popular area of research with numerous papers published every year. Most of these works are in the computational setting, where soundness is guaranteed to hold only against computationally bounded cheating provers. We do not elaborate on these works here, since they are less relevant for our result. We emphasize that our work is in the information theoretic setting, where soundness holds against any (even computationally unbounded) cheating prover. Most of the work in this regime is focused on guaranteeing privacy, such as zero-knowledge or witness indistinguishability. Again, we do not elaborate on these works since these are less relevant for our result.

A proof model which turns out to be useful for us, is that of *interactive proofs of proximity* (IPP) [RVW13], where a verifier has oracle access to a long input, and it interacts with a prover to check whether the input is close to one that satisfies a given property, while running in sublinear time. The non-interactive variant was explored in [GR15], and the significance of the round complexity of IPPs was studied in [GR17]. Although we do not focus on sublinear verification in our work, such proofs serve as a stepping stone to obtain our result.

2 Technical Overview

Our batch UIP protocol takes inspiration from the Batch-UP protocols of [RRR18, RR20].

The batch UP protocol from [RRR18, RR20]. The intuition underlying these protocols stems from the simple observation that if we are given k instances $x_1, \dots, x_k$ that are d -far from being in $\mathcal{L}^{\otimes k}$, in the sense that at least d of the instances are not in $\mathcal{L}$, then by choosing $k \cdot \text{polylog}(k)/d$ of the instances at random, we have the guarantee that at least one of the selected instances is not in $\mathcal{L}$ with probability $1 - \text{negl}(k)$. In other words, distance from $\mathcal{L}^{\otimes k}$ allows us to shrink the number of instances. Indeed, in the batch UP protocol of [RR20], they create distance and then shrink the number of instances, and they do this iteratively. To create distance, they use the GKR protocol.

The GKR protocol The GKR protocol [GKR15] is an unambiguous interactive protocol between a prover who wishes to prove to the verifier that a circuit C on input x outputs a bit b . The communication complexity and the number of rounds grow linearly with the depth of C , denoted by D , and poly-logarithmically with the size of C , denoted by S . It has the desired properties of being public-coin and the verdict function requires access to the input x at a single point in the low-degree extension of x .⁴ In particular, if the verifier is given oracle access to the low-degree extension of the input x , then its runtime is $D \cdot \text{polylog}(S)$, which may be poly-logarithmic in the input length.

The GKR protocol can be thought of as a reduction from verifying that $C(x) = b$ to verifying that the low-degree extension of x at point j is equal to a value v , where j is a uniformly random point determined by the random messages sent by the verifier in the GKR protocol, and v is determined by the answers sent by the prover in GKR protocol. This is also denoted by $x \in \text{PVAL}(j, v)$, a notation that originated in [RVW13]. The GKR protocol has the guarantee that if $C(x) \neq b$ then with high probability $x \notin \text{PVAL}(j, v)$, where we can make this probability as close as we want to 1, and obtain probability $1 - \mu$ at the price of the communication complexity growing polynomially with $\log(\mu^{-1})$. Importantly, the GKR protocol is run in the *holographic* mode, namely that the verifier never access the input x directly — it just outputs claims about the low-degree extension of x .

The batch UP protocol of [RR20] starts by running the GKR protocol on the batch verification circuit that has the instance $\mathbf{x} = (x_1, \dots, x_k)$ hardcoded (supposedly in $\mathcal{L}^{\otimes k}$), and takes as input a witness $\mathbf{w} = (w_1, \dots, w_k)$, and outputs 1 if and only if $\mathbf{w}$ is a valid witness for $\mathbf{x} \in \mathcal{L}^{\otimes k}$. Note that the depth of the batch verification circuit grows logarithmically with k , and hence the GKR protocol has communication complexity that grows only logarithmically in k (as desired). It has the guarantee that if $\mathbf{x} \notin \mathcal{L}^{\otimes k}$ then for every $\mathbf{w} = (w_1, \dots, w_k)$, with high probability $\mathbf{w} \notin \text{PVAL}(j, v)$. The batch UP protocol runs this GKR protocol $T = d \cdot m \cdot \text{polylog}(k)$ times, where $d = \text{polylog}(k)$ is the desired distance, in the sense that d witnesses will need to change in order to satisfy the resulting PVAL constraint, and m is the length of a single witness. Running the protocol T times in parallel reduces the soundness error to be exponentially small in T . They obtain the guarantee that if $\mathbf{x} \notin \mathcal{L}^{\otimes k}$ then for any $\mathbf{w} = (w_1, \dots, w_k)$, with high probability, which is exponentially (in

⁴The low-degree extension is a specific linear error correcting code (similar to the Reed-Muller encoding). We define it formally in Section 3.1 but the details can be ignored here.

T) close to 1, it holds that $\mathbf{w} \notin \cap_{i=1}^T \text{PVAL}(j_i, v_i)$. From now on, we use the condensed notation

$$\text{PVAL}(\mathbf{j}, \mathbf{v}) = \cap_{i=1}^T \text{PVAL}(j_i, v_i),$$

where $\mathbf{j} = (j_1, \dots, j_T)$ and $\mathbf{v} = (v_1, \dots, v_T)$. They then rely on the union bound to argue that even if a single x_i is not in $\mathcal{L}$, with high probability, every $\mathbf{w}$ that is d -close to the unambiguous witness,⁵ in the sense that it differs from the unambiguous one on at most d witnesses, satisfies that $\mathbf{w} \notin \text{PVAL}(\mathbf{j}, \mathbf{v})$, and hence obtain the desired distance guarantee!

We emphasize that if $\mathbf{x} \notin \mathcal{L}^{\otimes k}$ then the batch verification circuit computes the identically zero function, and in particular outputs zero for every witness vector that is close to the unambiguous one. This is precisely what allows us to obtain distance from running the GKR protocol. As we shall see shortly, this will no longer be the case in the UIP setting, which brings with it new technical challenges.

Note that after creating distance by running the GKR protocol T times, the new instances are no longer batch instances, rather are instances of the form $\mathbf{x} = (x_1, \dots, x_k)$ together with a joint PVAL constraint $\Phi = \Phi_{\mathbf{j}, \mathbf{v}}$ on the corresponding witnesses $\mathbf{w} = (w_1, \dots, w_k)$. The guarantee they obtain is that (with high probability) the unambiguous witness $\mathbf{w}$ corresponding to $\mathbf{x}$ is d -far from satisfying the constraint Φ , where distance is defined in terms of the minimum number of witnesses that need to be modified in order for the constraint Φ to be satisfied. In other words, they consider a new UP language $\mathcal{L}'$, where $(\mathbf{x}, \Phi) \in \mathcal{L}'$ if and only if the unambiguous witness $\mathbf{w}$ corresponding to $\mathbf{x} \in \mathcal{L}^{\otimes k}$ satisfies the PVAL constraint $\Phi(\mathbf{w}) = 1$. They obtain the guarantee that $(\mathbf{x}, \Phi)$ is d -far from $\mathcal{L}'$ in the sense that $(\mathbf{x}', \Phi) \notin \mathcal{L}'$, for every $\mathbf{x}'$ that differs from $\mathbf{x}$ on at most d instances.

They then show how to convert any such instance $(\mathbf{x}, \Phi)$ that is d -far from $\mathcal{L}'$ into a new instance $(x_{i_1}, \dots, x_{i_{k/d}}, \Phi')$ that is not in $\mathcal{L}'$, but without any distance guarantee.⁶ Recall that d was chosen so that $d = \text{polylog}(k)$, and hence the number of instances decreased by a factor of $\text{polylog}(k)$.

They continue as above, and use the GKR protocol to convert this instance into a new instance of the form $(x_{i_1}, \dots, x_{i_{k/d}}, \Phi'')$, where Φ'' is a new PVAL instance, with the guarantee that if $(x_{i_1}, \dots, x_{i_{k/d}}, \Phi') \notin \mathcal{L}'$ then $(x_{i_1}, \dots, x_{i_{k/d}}, \Phi'')$ is d -far from $\mathcal{L}'$. To obtain this reduction, they run the GKR protocol (repeated in parallel T times) with respect to the circuit that has $(x_{i_1}, \dots, x_{i_{k/d}}, \Phi')$ hardwired and on input $(w_{i_1}, \dots, w_{i_{k/d}})$ outputs 1 if and only if these are valid witnesses of $(x_{i_1}, \dots, x_{i_{k/d}})$ and they satisfy the constraint Φ' . They continue to iterate between creating distance and shrinking the number of instances until they are left with $\text{polylog}(k)$ number of instances, at which point the prover simply sends them over.

Our batch UIP protocol We refer to the verifier messages in the UIP protocol as *queries* and the prover messages as *answers*. Our batch UIP protocol follows a similar blueprint to the batch UP protocol of [RRR18, RR20], but for UIP as opposed to UP, which brings with it several challenges that we elaborate on below. The protocol starts by creating distance.

⁵If $x_i \notin \mathcal{L}$, we define the unambiguous witness to be $w_i = 0^m$.

⁶We note that the instance reduction protocol from [RR20] needs the stronger distance guarantee that at least $d \cdot m$ bits of the unambiguous witness of $\mathbf{x}$ need to change in order to satisfy the constraint Φ . This distance guarantee is indeed promised by the GKR protocol. We chose to present distance in terms of the number of witnesses that need to be changed since it is consistent with our approach. See Section 2.1.

2.1 Creating Distance

In the UIP setting, the instances $x_1, \dots, x_k$ may lack witnesses certifying membership in $\mathcal{L}$, since $\mathcal{L}$ need not belong to NP. Hence, we cannot apply the GKR protocol on the batch verification circuit that takes as input the witnesses (as such witnesses do not exist). Instead, in our UIP setting, distance is created by running a checksum over the k UIP protocols.

Checksum over the UIPs We begin by performing a checksum over the k UIP instances, following the batch UIP construction of [RRR16].⁷ Namely, we run the k UIP protocols in parallel, where the verifier uses the same queries in all the executions and the prover sends its answers via a checksum.

We assume that the underlying UIP protocol is public-coin and hence it can be run in the following way. Denoting by ℓ the number of rounds in the underlying UIP protocol, in each round $r \in [\ell]$ of the batch UIP protocol, the verifier sends a random message $q_r \in \{0, 1\}^a$ as in a single UIP protocol. The prover generates the k answers $\mathbf{a}_r = (a_{r,1}, \dots, a_{r,k}) \in (\{0, 1\}^a)^k$, but rather than sending $\mathbf{a}_r$ (which is too long), it sends a checksum of $\mathbf{a}_r$ (which should be thought of as a “digest” of $\mathbf{a}_r$). Formally, the prover sends the low-degree extension of $\mathbf{a}_r$ at $T = d \cdot a \cdot \text{polylog}(k)$ random points, where $d = \text{polylog}(k)$, and a is the length of each of the prover’s messages in a single execution. These points are sampled by the verifier at the beginning of the protocol.

Now we show that distance is created: if the checksummed transcript of answers only deviate from the unambiguous one by at most d executions, then at least one execution in the transcript will be rejecting with overwhelming probability. Suppose the cheating prover sends checksums that correspond to answers deviating in at most d executions, then for every $r \in [\ell]$, by the checksum’s error-correcting property, it is possible to extract from the r -th round checksums the entire set of r -th round answers. By unambiguity, at least one of the extracted answers will lead to a rejecting transcript with overwhelming probability when $(x_1, \dots, x_k) \notin \mathcal{L}^{\otimes k}$.

However, note that the verifier cannot verify the validity of the answers given by the prover, since the answers were not sent in the clear, but rather in the form of a checksum. To this end, we use the GKR protocol, this time not to create distance, but rather to allow the verifier to verify the validity of the answers. Specifically, as we describe below, the GKR protocol is used to convert the check-summed transcript with the distance guarantee mentioned above, into a PVAL claim with an analogous distance guarantee.

Remark 2. We note that this is where our approach differs substantially from [RRR16], which verified correctness by having the prover open some of the bits in *all* k instances (as opposed to using a succinct interactive proof, such as GKR, to verify correctness). As a result, denoting the number of iterations in their protocol as $1/\epsilon$ (for some $\epsilon > 0$), their batch UIP incurs an additive factor of k as well as a multiplicative factor of $(1/\epsilon)^{\text{poly}(1/\epsilon)}$ to the communication complexity. We elaborate on this difference in Section 2.3.

The GKR protocol In order to verify the check-summed UIPs, the prover and verifier run the GKR protocol, this time on the batch verification circuit that has hardwired into it the instances $\mathbf{x} = (x_1, \dots, x_k)$, the verifier’s queries $\mathbf{q} = (q_1, \dots, q_\ell)$, and a constraint Φ on the prover’s answers (as defined by the checksum). The circuit takes as input all the prover’s answers $\mathbf{a} = (\mathbf{a}_1, \dots, \mathbf{a}_\ell)$

⁷We note that in [RRR16] the checksum was not done directly on the UIP but rather the UIP was converted to an IOP and the checksum was done on the IOP. We elaborate on this in Section 2.3.

and outputs 1 if and only if for every $j \in [k]$ the transcript $(q_1, a_{1,j}, \dots, q_\ell, a_{\ell,j})$ is accepted with respect to x_j , and the answers $\mathbf{a}$ satisfy Φ .

The distance guarantee of the check-summed protocol implies that with high probability, every transcript of answers $\mathbf{a} = (\mathbf{a}_1, \dots, \mathbf{a}_\ell)$ that is d -close to the unambiguous one is rejected by the batch verification circuit, where here (and throughout our work) we define d -closeness to mean that it differs in at most d executions.

Remark 3. We emphasize that our definition of d -closeness is not the standard Hamming distance (which is bit-wise distance), but rather it is an execution-wise distance. Since our definition of distance deviates from the one used in [RR20] (which is Hamming distance) we need to prove that the instance reduction protocol from [RR20] works with our notion of execution-wise distance. We elaborate on this in Section 2.2.

After running the GKR protocol (in parallel) the parties obtain $\mathbf{j}$ and $\mathbf{v}$ such that for every $\mathbf{a}$ that is d -close to the unambiguous one, with high probability, $\mathbf{a} \notin \text{PVAL}(\mathbf{j}, \mathbf{v})$. We then rely on the union bound, following the approach in the UP setting, to argue that with high probability for every $\mathbf{a}$ that is d -close to the unambiguous one it holds that $\mathbf{a} \notin \text{PVAL}(\mathbf{j}, \mathbf{v})$. Therefore, we can apply an instance reduction protocol (similar to the one from [RR20], which we elaborate on in Section 2.2), which reduces the problem to verifying that $(x_{i_1}, \dots, x_{i_{k/d}}, \mathbf{q}, \Phi') \in \mathcal{L}'$ for some Φ' , where the language $\mathcal{L}'$ is defined so that $(x_{i_1}, \dots, x_{i_{k/d}}, \mathbf{q}, \Phi') \in \mathcal{L}'$ if and only if the unambiguous answers, corresponding to the instances $x_{i_1}, \dots, x_{i_{k/d}}$ and queries $\mathbf{q}$, satisfy Φ' .⁸

The guarantee we have is that if $(x_1, \dots, x_k, \mathbf{q}, \Phi)$ is d -far from $\mathcal{L}'$, in the sense that every answer $\mathbf{a}'$ that is d -close to the unambiguous one, does not satisfy Φ , then with high probability $(x_{i_1}, \dots, x_{i_{k/d}}, \mathbf{q}, \Phi') \notin \mathcal{L}'$. We note that this instance does not have a distance guarantee, and generating distance for instances in $\mathcal{L}'$ introduces new challenges that we elaborate on below.

Obtaining distance for $\mathcal{L}'$ It is tempting to try and continue as in the UP setting, and apply the GKR protocol to create distance. The problem is that at this point the cheating prover knows all the verifier's queries $\mathbf{q}$ in advance, and thus can easily generate accepting answers $\mathbf{a}$. Moreover, a cheating prover can generate answers $\mathbf{a}$ that are close to the unambiguous ones and which cause the GKR circuit to accept! Hence, the GKR protocol cannot be used to create distance.

We emphasize that the problem stems from the fact that the prover knows all the queries $\mathbf{q}$ ahead of time, and thus can easily generate accepting answers even for $x \notin \mathcal{L}$. We get around this problem by adding the following random continuation protocol from [RRR16].

Random continuation to the rescue In the random continuation protocol, the verifier chooses fresh queries $q'_1, \dots, q'_\ell \leftarrow \{0, 1\}^a$, and the prover and verifier run a checksum of $\ell \cdot k/d$ UIP protocols, where in the $(i \cdot k/d + j)$ 'th protocol, where $i \in \{0, 1, \dots, \ell - 1\}$ and $j \in \{1, \dots, k/d\}$, the verifier sends messages $(q_1, \dots, q_i, q'_{i+1}, \dots, q'_\ell)$ and the prover replies with answers corresponding to the j 'th instance. Thus, the prover and verifier run ℓ UIP protocols per instance, and thus a total of $\ell \cdot k/d$ UIP protocols, all run in a checksum.

Each of the prover's checksums consist of $T = d' \cdot a \cdot \text{polylog}(k)$ points in the low-degree extension of each vector of answers, where here we take $d' = \ell \cdot d$.

Remark 4. The reason for this choice of distance d' is that the instance reduction protocol, applied after the distance is generated, reduces the number of instances by a factor of d' . However, the

⁸We note that $\mathcal{L}'$ is not necessarily in NP (as was not $\mathcal{L}$).

random continuation increases the number of protocols by a factor of ℓ , and hence to ensure that the number of instances goes down by a factor of d , we need to take $d' = \ell \cdot d$.

One can argue (as above) that if the prover deviates from the (unambiguous) random continuation protocol on at most d' of the executions (where in each execution it sends a message of length a) then one can efficiently extract the messages of the prover from the check-summed transcript. Therefore, if there exists $i \in [\ell]$ such that the i 'th round check-summed answers sent by the prover, corresponding to $(q_1, \dots, q_i)$, were not the unambiguous ones, then by the unambiguity property the prover will fail to provide accepting answers for random continuation $(q'_{i+1}, \dots, q'_\ell)$, even if the corresponding instance is in $\mathcal{L}$. This implies that if there exists $i \in [\ell]$ such that the i 'th check-summed answers corresponding to $(q_1, \dots, q_i)$, are not the unambiguous ones, and the prover deviates from the unambiguous answers in at most d' executions, then we can efficiently extract from the check-summed transcript a rejecting transcript of answers sent by the prover. This gives us the desired distance guarantee, and we can now use the GKR protocol, as above, to reduce this to a PVAL instance $\text{PVAL}(\mathbf{j}, \mathbf{v})$ with the guarantee for every $\mathbf{w}$ that is d' -close to the unambiguous one, $\mathbf{w} \notin \text{PVAL}(\mathbf{j}, \mathbf{v})$. Hence, we can use the instance reduction protocol (similar to the one given in [RR20]) to reduce the number of instances from $\ell \cdot k/d$ to $\ell \cdot k/d \cdot 1/d' = k/(d^2)$.

Remark 5. We note that in our actual batch UIP protocol, for the sake of simplicity, we use distance $d' = \ell \cdot \text{polylog}(k)$ from the beginning of the protocol, and do not treat the initial reduction, from a batch instance $(x_1, \dots, x_k) \in \mathcal{L}^{\otimes k}$ to an instance in $\mathcal{L}'$, differently.

2.2 Instance Reduction

The Need for Execution-wise Distance Recall that the distance guarantee provided by our distance generation protocol, described in Section 2.1, is that if we start with a false statement on say k' instances, then after the distance generation protocol (which includes the checksum protocol followed by the GKR protocol), we obtain a PVAL constraint $\text{PVAL}(\mathbf{j}, \mathbf{v})$ that is not satisfied by any transcript of answers that differs on at most d' executions from the prescribed transcript of answers (with respect to fixed verifier's queries). For the next step, we would like to use the instance reduction protocol from [RR20] to further reduce this to a false claim on a subset of k'/d' instances (at the cost of losing our distance guarantee).

Unfortunately, the [RR20] protocol is guaranteed to produce a false claim on a subset of k'/d' instances, only if the original PVAL constraint $\text{PVAL}(\mathbf{j}, \mathbf{v})$ is not satisfied by any transcript that differs from the prescribed transcript in at most $(d' \cdot \ell \cdot a)$ -bits! Specifically, the issue is that these two protocols use a different distance measure, and these are not compatible: One guarantees execution-wise distance while the other requires Hamming distance.

We note that being d' -far in execution distance trivially implies being d' -far in Hamming distance. However, to reduce the number of instances by a d' -factor, the [RR20] protocol requires the Hamming distance to be $d' \cdot \ell \cdot a$, where $\ell \cdot a$ is the length of a single transcript. This requirement is much stronger than requiring the execution distance to be d' . In other words, if we only have the guarantee that the execution distance is d' , then applying [RR20] will result with the number of instances being $k' \cdot \frac{d'}{\ell \cdot a}$, which offers no decrease in the number of instances, unless we take d' to be larger than $\ell \cdot a$, in which case we invest $d' \cdot \ell \cdot a \cdot \text{polylog}(k)$ in communication to decrease the number of instances by a factor of d' . This will result in a suboptimal batch UIP result. While further optimizations are possible here, it seems that any purely Hamming distance based approach will result in a larger overhead than what we achieve.

Instead, we choose to make the two distance measures compatible by changing the instance reduction protocol to guarantee that if the PVAL constraint is not satisfied by any transcript that is d' -close in execution distance to the prescribed one, then it results with a false claim on a subset of k'/d' instances. We next elaborate on how this is done.

Overview of the original instance reduction protocol for PVAL from [RR20] The verifier in the instance reduction protocol of [RR20] takes as input $\mathbf{j}, \mathbf{v}$ that defines the PVAL constraint. In the YES case, the claim $\mathbf{w} \in \text{PVAL}(\mathbf{j}, \mathbf{v})$ is true, where $\mathbf{w}$ is the implicit vector of unique witnesses $\mathbf{w} = (w_1, \dots, w_k) \in (\{0, 1\}^m)^k$. In the NO case, $\mathbf{w}$ as a string in $\{0, 1\}^{k \cdot m}$ is d -far in Hamming distance from $\text{PVAL}(\mathbf{j}, \mathbf{v})$ with $d = m \cdot \text{polylog}(k)$.

They show how to subsample a $\frac{1}{d}$ fraction of the instances together with a new false PVAL claim on the corresponding witnesses. This is not straightforward since the initial constraint is a global constraint, and cannot be applied to each witness w_i separately. The authors in [RR20] propose to transform the global constraint into a bunch of localized constraints, each defined over $\frac{1}{d}$ fraction of the instances. In order to obtain these localized claims, the [RR20] protocol proceeds in at most $\log k$ iterations as follows:

1. The multi-linearity of $\text{PVAL}(\mathbf{j}, \mathbf{v})$ allows the prover to break the PVAL claim into two halves: PVAL_0 and PVAL_1 , where each claim is over $k/2$ witnesses. Specifically, the claim PVAL_b is over $\mathbf{w}_b$ where $\mathbf{w}_0$ and $\mathbf{w}_1$ are defined by:

$$\mathbf{w}_0 = (w_1, \dots, w_{k/2}) \quad \text{and} \quad \mathbf{w}_1 = (w_{k/2+1}, \dots, w_k).$$

This results in $(\mathbf{w}_0, \mathbf{w}_1)$ being d -far from the new claim $\text{PVAL}_0 \times \text{PVAL}_1$. For this overview, assume that these $\mathbf{w}_0$ and $\mathbf{w}_1$ are both $d/2$ far from the respective PVAL instances.

2. The next step would be to try to fold these two claims into a single claim of “size $k/2$ ” *while keeping the absolute distance to be d .*

A naive folding by taking a random linear combination of $\mathbf{w}_0$ and $\mathbf{w}_1$ along with the corresponding linear combination of PVAL_0 and PVAL_1 may reduce the absolute distance to $d/2$ since the errors in $\mathbf{w}_0$ and $\mathbf{w}_1$ may overlap. Instead [RR20] applies a more sophisticated folding, as follows.

3. The verifier uses a *pairwise-independent* family π to transform $\mathbf{w}_1$ and PVAL_1 into a new claim PVAL'_1 about $\pi(\mathbf{w}_1)$. The purpose is to scramble the $d/2$ errors that prevent $\mathbf{w}_1$ from satisfying PVAL_1 .
4. The verifier “folds” the two claims into one, by taking a random linear combination of $\mathbf{w}_0$ and $\pi(\mathbf{w}_1)$, and a random linear combination of PVAL_0 and PVAL'_1 to get a new claim that is also roughly d -far.

Each iteration allows us to reduce a d -far claim to another d -far claim but on a “folded” instance of half the size. The verifier repeats this until it is left with a d -far claim over roughly d “folded” witnesses. Then the verifier kindly asks the prover to send the batch of “folded” witnesses $\mathbf{w}^* = (w_1^*, \dots, w_d^*)$ in the clear, which requires $d \cdot m$ communication. Then it checks whether the batch of the “folded” witnesses $\mathbf{w}^*$ satisfies the final “folded” PVAL claim. If not, it rejects immediately. Otherwise, in the NO case, $\mathbf{w}^* = (w_1^*, \dots, w_d^*)$ is still d -far from any $\mathbf{w}^*$ that satisfies the final

PVAL claim, so the sent $\mathbf{w}^*$ must differ from $\mathbf{w}^*$ on at least d instances. Therefore, the verifier can indeed subsample constantly many of these instances $\mathbf{w}_i^*$ to detect an inconsistency with high probability, and the inconsistency depends only on the (k/d) witnesses folded in $\mathbf{w}_i^*$.

Working with execution-wise distance We generalize (and tweak) the protocol from [RR20] for Hamming distance to work for the execution-wise distance. The modification of the folding step is illustrated in Figure 1. To this end, we view the batch of transcripts $\mathbf{a}$ as a matrix where each row corresponds to a transcript of a single instance (Figure 1a). We split $\mathbf{a}$ into its top half submatrix $\mathbf{a}_0$ and bottom half submatrix $\mathbf{a}_1$ (Figure 1b). Execution-wise distance is preserved in the folding step by choosing a sufficiently random permutation π to scramble each column of $\mathbf{a}_1$ independently (Figure 1c). (Note that this is different from scrambling the entire $\mathbf{a}_1$, which moves errors across columns, as in the original protocol from [RR20].)




Figure 1: Our folding step works with execution-wise distance by permuting each column randomly.

We show that this protocol (as well as the [RR20] protocol) also has unambiguous soundness when the underlying field is large. Thus, it can also be used as a building block for a UIP protocol, leading to many of our applications.

2.3 Efficiency Gain over the RRR16 Protocol

The Batch-UIP protocol of [RRR16] has communication complexity at least

$$k \cdot (\ell/\epsilon)^{\text{poly}(1/\epsilon)} + \text{poly}(a) \cdot (\ell/\epsilon)^{\text{poly}(1/\epsilon)} \cdot k^\epsilon$$

for an arbitrarily small constant $\epsilon > 0$. On a closer inspection, their protocol allows the choice of a slightly sub-constant $\epsilon = \epsilon(k)$, at the cost of creating a super-polynomial dependency on ℓ (i.e. the terms involving $\ell^{1/\epsilon}$) in the communication complexity. In contrast, our Batch-UIP always achieves a communication complexity of $O(a \cdot \ell^2 \cdot \text{polylog}(k) + \text{poly}(\ell))$.

Both Batch-UIP protocols use the idea of running many protocols in a checksum and running the random continuation protocol (also in a checksum). These are used to enforce that the cheating prover will be caught if it deviates from the protocol only in a few instances. Nevertheless, our protocol is very different in how the verifier catches the prover if it indeed only deviated in a few instances. While the verifier in [RRR16] checks explicitly on its own that this is the case, by asking the prover to reveal some of these check-summed messages, the verifier in our protocol delegates these checks back to the prover (via the GKR protocol). This reduces the overall costs significantly.

Specifically, in the protocol of [RRR16], the prover first encodes the base UIP into an (unambiguous) *Interactive Oracle Proof* (IOP). In such a system, the prover and verifier interact, with the verifier sending random coins and receiving answers from the prover. The verifier does not read the prover's messages immediately, and only after the interaction is complete, it makes its verdict based on a set of random queries into very few locations of the stored transcript. Similar to our usage, the prover and verifier run the k instances of the IOP implicitly, with the prover only sending checksums of the underlying messages.

Importantly, after the interaction, in [RRR16], the verifier sends the set of query locations that it would like to check in the IOP to the prover, and the prover responds with *all* the messages on these locations. In addition to checking that all k instances of the IOP are accepting, the verifier also confirms that the checksums on those messages agree with what the prover committed to during the interaction. Note that this requires communicating $\Omega(k)$ bits.

The guarantee they obtain is that if originally there was at least one false statement, then if the verifier accepts the openings then it must be the case that the prover deviated on many instances of the IOP, so the verifier can subsample the instances to reduce the number of instances.

At this point, the random queries of the verifier in the original IOP have been revealed through the interaction, and to remedy this, the verifier in [RRR16] runs the IOP corresponding to a random continuation of the base IOP (constructed from the base UIP) in the next distance-creation (checksum) phase. The random continuation always incurs a multiplicative factor of ℓ blowup in the IOP, so if this process is iterated for $1/\epsilon$ times, the transcript length of the final IOP is $\text{poly}(a) \cdot \ell^{1/\epsilon}$.

In our protocol, the verifier avoids this accumulation of the factor ℓ in the transcript length by instead “flattening” the transcripts produced by the random continuation. Namely, instead of treating the implicit transcripts as k transcripts of the random continuation protocol (where in each round the prover sends a message of length $a \cdot \ell$), it treats them as $k \cdot \ell$ instances of the *base UIP*. This is desirable because if we generate a distance of $d \geq \ell \cdot \text{polylog}(k)$ by the checksum, we

can offset this factor of ℓ in the number of instances in the instance reduction step. A caveat is that the verifier should additionally check a joint constraint over each of the ℓ hybrid transcripts derived from the same instance by the random continuation, namely that they have shared prefixes. This check can be delegated to the prover via the GKR protocol. This optimization is not possible under the framework of IOP of [RRR16], since a priori the IOP of the random continuation protocol cannot be flattened into ℓ independent IOP's.

2.4 Doubly Efficient Interactive Proofs

Our Batch-UIP protocol can be used to obtain new (and more powerful) doubly-efficient interactive proofs for bounded-space computations, by applying a similar technique as that presented in [RRR16]. Intuitively, any Batch-UIP protocol can be used to obtain a doubly-efficient interactive proof for poly-space computations.

Formally, let $\mathcal{L} \in \text{DTISP}(T, S)$ be a language that is decidable by a Turing machine $\mathcal{M}$ in time $T(n)$ and space $S(n)$. Let $t = t(n) \leq T(n)$ be a parameter, and let $\mathcal{L}_t$ be the set of all tuples $(x, w_1, w_2) \in \{0, 1\}^n \times \{0, 1\}^{O(S(n))} \times \{0, 1\}^{O(S(n))}$ such that the Turing machine $\mathcal{M}$ transitions from configuration w_1 to configuration w_2 on input x in exactly t steps.

Crucially, for every $k = k(n) \in \mathbb{N}$, the language $\mathcal{L}_{k \cdot t}$ can be expressed as a batch of $\mathcal{L}_t$. Specifically, for input $(x, w_0, w_{k \cdot t})$:

- The prover runs $\mathcal{M}$ on input x for $k \cdot t$ steps, and obtains the intermediate configurations $w_1, w_2, \dots, w_{k \cdot t}$. (Note this step can be skipped if the prover already has the intermediate configurations.)
- The prover sends the k intermediate configurations $w_t, w_{2t}, \dots, w_{k \cdot t}$ to the verifier.
- Both parties run the Batch-UIP protocol on the k instances $\{(x, w_{(i-1) \cdot t}, w_{i \cdot t})\}_{i=1}^k$ of $\mathcal{L}_t$, and the verifier accepts if and only the Batch-UIP protocol accepts.

Therefore, an efficient batching protocol such as our protocol allows the verifier to verify $(k \cdot t)$ -step transitions by batch-verifying t -step transitions, at the cost of a small (polylogarithmic) increase in complexity. By repeating this process with a careful balancing of the parameters, the verifier gradually increases the transition length t that it can verify, until it reaches $\mathcal{L}_{T(n)}$.

We show that with $k = \text{poly}(n)$, this process can be repeated for $O(\sqrt{\frac{\log n}{\log \log n}})$ times before the communication complexity becomes superpolynomial, obtaining a doubly-efficient interactive proof for languages in $\text{DTISP}(T, S)$ where $T = n^{O(\sqrt{\frac{\log n}{\log \log n}})}$ and $S = \text{poly}(n)$, where the prover runtime is $T \cdot \text{poly}(n)$, and the verifier runtime is $\text{poly}(n)$.

3 Preliminaries

We adhere to the following convention when possible.

- Lower case letters a, b mean scalars, while bolded lower case letters $\mathbf{a}, \mathbf{b}$ mean vectors. Upper case bold letters like $\mathbf{A}, \mathbf{B}$ are matrices.
- For an integer n , we denote by $[n]$ the set $\{1, 2, \dots, n\}$. We let $\mathbb{F}$ denote a finite field. $\mathbb{GF}(2)$ denotes the finite field with 2 elements.

- For a vector $\mathbf{u} \in \mathbb{F}^n$, and a subset $\mathcal{S} \subset [n]$, $\mathbf{u}|_{\mathcal{S}} \in \mathbb{F}^{|\mathcal{S}|}$ denotes the subvector of $\mathbf{u}$ indexed by $\mathcal{S}$. For a sequence of vectors $\{\mathbf{u}_i\}_{i \in \mathcal{I}}$, where each $\mathbf{u}_i \in \mathbb{F}^n$, the notation $(\mathbf{u}_i)_{i \in \mathcal{I}} \in \mathbb{F}^{n \times |\mathcal{I}|}$ represents the matrix whose columns are the vectors $\mathbf{u}_i$ for every $i \in \mathcal{I}$.
- Given a matrix $\mathbf{A} = (\mathbf{a}_1, \dots, \mathbf{a}_n) \in \mathbb{F}^{m \times n}$, $\mathbf{a}_j \in \mathbb{F}^m$ denotes its j -th column. We also use $\mathbf{A}[i, :]$ to denote the i -th row of $\mathbf{A}$. For a subset $\mathcal{S} \subset [m]$, $\mathbf{A}[\mathcal{S}, :]$ denotes the submatrix of $\mathbf{A}$ consisting of rows indexed by $\mathcal{S}$. We also consider matrices $A \in \mathbb{F}^{k \times n}$ whose rows are indexed by an ordered set of k pairs, $\mathcal{T} = \{(u^1, v^1), \dots, (u^k, v^k)\}$, and we denote by $\mathbf{A}[(u^i, v^i), :]$ the row of $\mathbf{A}$ indexed by $(u^i, v^i) \in \mathcal{T}$, and for $\mathcal{S} \subset \mathcal{T}$, $\mathbf{A}[\mathcal{S}, :]$ denotes the submatrix of $\mathbf{A}$ consisting of rows indexed by $\mathcal{S}$.
- $\Delta(\mathbf{u}, \mathbf{v})$ is the (absolute) Hamming distance between the vectors $\mathbf{u}$ and $\mathbf{v}$. Denote by $\Delta(\mathbf{u})$ the vector $\mathbf{u}$'s *Hamming weight*, defined to be the number of non-zero elements in $\mathbf{u}$. For any $d \in \mathbb{N}$, two vectors $\mathbf{u}, \mathbf{v}$ are d -close (in Hamming distance) if $\Delta(\mathbf{u}, \mathbf{v}) \leq d$. On a linear space $\mathcal{U} \subset \mathbb{F}^a$, let $\Delta(\mathcal{U}) := \min_{\mathbf{u} \in \mathcal{U}, \mathbf{u} \neq \mathbf{0}} \Delta(\mathbf{u})$.
- Given a Boolean circuit V , $\text{size}(V)$ is the number of gates in the circuit, and $\text{depth}(V)$ is the depth of the circuit.

An important distance notion that we consider is the “column distance,” denoted by Δ_c , between matrices.

Definition 1 (Δ_c -distance). Let $\mathbb{F}$ be a finite field. For matrices $\mathbf{A} = (\mathbf{a}_1, \dots, \mathbf{a}_L) \in \mathbb{F}^{M \times L}$ and $\mathbf{B} = (\mathbf{b}_1, \dots, \mathbf{b}_L) \in \mathbb{F}^{M \times L}$, the Δ_c -distance between $\mathbf{A}$ and $\mathbf{B}$, denoted by $\Delta_c(\mathbf{A}, \mathbf{B})$, is the maximum Hamming distance between corresponding columns of $\mathbf{A}$ and $\mathbf{B}$, i.e.

$$\Delta_c(\mathbf{A}, \mathbf{B}) = \max_{i \in [L]} \Delta(\mathbf{a}_i, \mathbf{b}_i).$$

If $\Delta_c(\mathbf{A}, \mathbf{B}) \leq d$, we say $\mathbf{A}$ and $\mathbf{B}$ are Δ_c - d -close,

Furthermore, let $d \in \mathbb{N}$ be some distance parameter, then given $\mathbf{A} \in \mathbb{F}^{M \times L}$, its Δ_c - d -ball is defined to be

$$\mathcal{B}_{d, \mathbb{F}}(\mathbf{A}) := \{\mathbf{A}' \in \mathbb{F}^{M \times L} : \Delta_c(\mathbf{A}, \mathbf{A}') \leq d\}.$$

Remark 6. Observe that Δ_c is a (non-tight) lower bound on how many rows have to be modified to transform one matrix into another.

3.1 Low Degree Extensions and Polynomial Valuation (PVAL)

All finite fields $\mathbb{F}$ considered in this work are always *constructible* in the following sense.

Definition 2 (Constructible Field Ensemble). We say a field ensemble $\mathbb{F} = (\mathbb{F}_n)_{n \in \mathbb{N}}$ is *constructible* if every element in $\mathbb{F}_n$ has a $O(\log |\mathbb{F}_n|)$ -bit representation, and addition, multiplication, and inverses can be computed in $\text{polylog}(|\mathbb{F}_n|)$ time given the representations.

It is well known that for every $S = S(n)$, constructible field ensembles $\mathbb{F} = (\mathbb{F}_n)$ with $|\mathbb{F}_n| = \Theta(S)$ exist. Moreover, we make the (mild) assumption that addition, multiplication and inverses can be computed in $\tilde{O}(\log |\mathbb{F}_n|)$ time, where $\tilde{O}$ omits $\text{polylog log}(|\mathbb{F}_n|)$ factors. This implies that degree- d

polynomials in $\mathbb{F}[X]$ can be evaluated in $\tilde{O}(d \log |\mathbb{F}_n|)$ time (with $\tilde{O}$ omitting $\text{polylog}(d, \log |\mathbb{F}_n|)$ factors)⁹. These properties are satisfied in fields that support FFT. (See table 8.6 in [vzGG13])

We utilize the Schwartz-Zippel lemma, a basic fact about low-degree polynomials over finite fields.

Lemma 1 (Schwartz-Zippel Lemma, [Zip79, Sch80]). *Let $\mathbb{F}$ be a field, let $m \in \mathbb{N}$, and let $f \in \mathbb{F}[X_1, \dots, X_m]$ be a non-zero polynomial of total degree d . Then*

$$\Pr_{\mathbf{r} \leftarrow \mathbb{F}^m} [f(\mathbf{r}) = 0] \leq \frac{d}{|\mathbb{F}|}.$$

Given a subset $H \subset \mathbb{F}$ and an integer $m \in \mathbb{N}$, the *low-degree extension* (LDE) of a function $f : H^m \rightarrow \mathbb{F}$ is the unique $(|H| - 1)$ -individual-degree polynomial $\hat{f} : \mathbb{F}^m \rightarrow \mathbb{F}$ such that for all $\mathbf{s} \in H^m$, $\hat{f}(\mathbf{s}) = f(\mathbf{s})$. In this work we focus on the special case when $H = \{0, 1\}$, where $\hat{f}$ is multilinear over $\mathbb{F}^m$, and is called the multilinear extension of f . Abusing the notation, given a string $x \in \mathbb{F}^{2^m}$, $\hat{x} : \mathbb{F}^m \rightarrow \mathbb{F}$ denotes the multilinear extension of the function $f_x : \{0, 1\}^m \rightarrow \mathbb{F}$ whose function table is x . On a sequence $\mathbf{j} = (\mathbf{j}_1, \dots, \mathbf{j}_T) \in (\mathbb{F}^m)^T$ of length T , we use the shorthand $\hat{x}(\mathbf{j}) := (\hat{x}(\mathbf{j}_1), \dots, \hat{x}(\mathbf{j}_T)) \in \mathbb{F}^T$.

Consider the following *Polynomial Valuation* (PVAL) set, which is an affine subspace over $\mathbb{F}$, first defined in [RVW13].

Definition 3 (The PVAL set). Let $m, T \in \mathbb{N}$. The set $\text{PVAL}_{\mathbb{F}}(\mathbf{j}, \mathbf{v}) \subset \mathbb{F}^{2^m}$ is parameterized by the sequences $\mathbf{j} = (\mathbf{j}_1, \dots, \mathbf{j}_T) \in (\mathbb{F}^m)^T$ and $\mathbf{v} = (v_1, \dots, v_T) \in \mathbb{F}^T$. It consists of all strings $x \in \mathbb{F}^{2^m}$ whose corresponding multilinear extension $\hat{x} : \mathbb{F}^m \rightarrow \mathbb{F}$ satisfies $\hat{x}(\mathbf{j}) = \mathbf{v}$.

Note that $\text{PVAL}(\mathbf{j}, \mathbf{0})$ is a linear subspace of $\mathbb{F}^{2^m}$, so we can define $\Delta(\text{PVAL}(\mathbf{j}, \mathbf{0}))$ to be the minimum Hamming distance of a non-zero vector in $\text{PVAL}(\mathbf{j}, \mathbf{0})$, and the following proposition holds.

Proposition 1 (Random PVAL kernel has large distance). *Let $\sigma, d, m \in \mathbb{N}$. Suppose $\mathbb{F}$ is a field of characteristic 2, $|\mathbb{F}| \geq 2 \cdot m$, and $T \geq 2d(m + \log |\mathbb{F}|) + \sigma$. Then,*

$$\Pr_{\mathbf{j} \leftarrow (\mathbb{F}^m)^T} [\Delta(\text{PVAL}(\mathbf{j}, \mathbf{0})) \leq 2d] \leq 2^{-\sigma}.$$

Proof. Fix a non-zero $\mathbf{x} \in \mathbb{F}^{2^m}$ with Hamming weight at most $2d$. Since $\hat{\mathbf{x}}$ is a multilinear polynomial with total-degree m in $\mathbb{F}[X_1, \dots, X_m]$, by Lemma 1, the probability that $\hat{x}(\mathbf{j}_i) = 0$ is at most $m/|\mathbb{F}|$. As $\mathbf{j}$ contains T independent samples, all of these are zero with probability at most $(m/|\mathbb{F}|)^T$. Therefore, $\mathbf{x} \in \text{PVAL}(\mathbf{j}, \mathbf{0})$ with probability at most $(m/|\mathbb{F}|)^T$.

There are at most $\binom{2^m}{2d} |\mathbb{F}|^{2d}$ choices for $\mathbf{x} \in \mathbb{F}^{2^m}$ with Hamming weight at most $2d$, so by a union-bound and the assumption $T \geq 2d(m + \log |\mathbb{F}|) + \sigma$, the probability that some of them are in $\text{PVAL}(\mathbf{j}, \mathbf{0})$ is at most

$$\binom{2^m}{2d} |\mathbb{F}|^{2d} \left(\frac{m}{|\mathbb{F}|} \right)^T \leq 2^{2dm} \cdot |\mathbb{F}|^{2d} \cdot 2^{-T} \leq 2^{-\sigma}. \quad \square$$

⁹Both these assumptions can also be relaxed by further delegating operations over large fields. For simplicity, we do not deal with that here.

3.2 Succinct Descriptions of Sets and Predicates

Definition 4 (Uniform Arithmetic Circuits). Let $\mathbb{F}$ be a field. Let $C = \{C_n : \mathbb{F}^n \rightarrow \mathbb{F}^m\}_{n \in \mathbb{N}}$ be a family of arithmetic circuits, consisting of fan-in 2 ADD and MULT gates over a field $\mathbb{F}$. For any $f = f(n)$, we say that C is f -space uniform if there exists a fixed $O(f(n))$ -space Turing machine $\mathcal{M}$ that on input 1^n outputs the full description of the circuit C_n . When n is clear from the context, we omit the subscript n and write C instead of C_n .

An important special case is when $f(n) = \log(n)$, in which case we say that C is log-space uniform.

Remark 7. Any f -space uniform Boolean circuit C can be trivially extended to by a f -space uniform arithmetic circuit C' of the same size and depth over any field $\mathbb{F}$.

The following *succinct descriptions* of sets can be used to recover the entire set.

Definition 5 (Descriptions of Sets). A bit string $\langle \mathcal{S} \rangle \in \{0, 1\}^B$ is a description of a set $\mathcal{S} = \{s_1, \dots, s_k\} \subset \{0, 1\}^p$ if there exists a (multi-output) p -space uniform $G : [k] \times \{0, 1\}^B \rightarrow \{0, 1\}^p$ of fan-in 2, such that $G(i, \langle \mathcal{S} \rangle) = s_i$ for all $i \in [k]$. The description is succinct if $|\langle \mathcal{S} \rangle| = B < k \cdot p$.

Similarly, *succinct descriptions* of functions can be used to configure a uniform circuit family to implement the function.

Definition 6 (Descriptions of Functions). Let $\mathbb{F}$ be a field. We say that $\langle \Phi \rangle \in \{0, 1\}^B$ is a description of a function $\Phi : \mathbb{F}^n \rightarrow \mathbb{F}$ if there exists a log-space uniform circuit $C : \mathbb{F}^{n+B} \rightarrow \mathbb{F}$ of fan-in 2 such that $C(x, \langle \Phi \rangle) = \Phi(x)$ for all $x \in \mathbb{F}^n$. The description is succinct if $|\langle \Phi \rangle| = B < \text{size}(\Phi)$.

The Turing machines that generate the uniform G and C in Definitions 5 and 6 take in their “shape parameters” $(k, 1^B, 1^p)$ and $(1^n, 1^B)$ as input, respectively.

3.3 Unambiguous Interactive Proof (UIP)

An $(\ell, a, b, \text{Ptime}, \text{Vtime}, \Sigma)$ -protocol is a *public-coin, ℓ -round interactive protocol, with alphabet Σ , per-round prover message length a and per-round verifier message length b , and prover runtime Ptime and verifier runtime Vtime* . Specifically, such a protocol consists of a pair of interacting Turing machines (P, V) each having as input a string x of length $n = |x|$. The machines may also take in other parameters as additional input, such as parameters controlling the soundness error,¹⁰ but we omit them from the notation for simplicity. The machine P is deterministic and runs in time Ptime , and is called *the prover*, while V is probabilistic, runs in time Vtime , and is called *the verifier*. P and V are the *two parties* of the protocol. We omit the specification of Σ when $\Sigma = \{0, 1\}$.

In each round $j \in [\ell]$ of the protocol:

1. V sends a random message $q_j \leftarrow_R \Sigma^b$ to P , referred to as *queries*.
2. P responds with a message $a_j \in \Sigma^a$ determined by the prescribed next-message function, referred to as *answers*, which (abusing notation) is denoted by

$$a_j := P(x, j, (q_1, \dots, q_j)) \in \Sigma^a.$$

¹⁰In particular, the prover might have some extra information that helps make it more efficient, for example the UP witnesses in case the language was in UP.

We make the simplifying assumption that V never rejects in the middle of an execution. We abuse notation and denote the verdict circuit of V , that reads the entire transcript and decides if to accept or reject, also by V :

$$V(x, (q_1, \dots, q_\ell), (a_1, \dots, a_\ell)) \in \{0, 1\}.$$

We denote the sequence of verifier random coins by $\mathbf{q} := (q_1, \dots, q_\ell)$, and for $j \in [\ell]$ we denote by

$$\mathbf{q}_{\leq j} := (q_1, \dots, q_j) \quad \text{and} \quad \mathbf{q}_{> j} := (q_{j+1}, \dots, q_\ell).$$

Finally, let $P(x, \mathbf{q}) := (a_1, \dots, a_\ell)$, where $a_j = P(x, j, \mathbf{q}_{\leq j})$ are the prescribed messages.

The *total communication complexity* of the protocol is the number of bits exchanged between the prover and the verifier, i.e. $(a + b)\ell \cdot \log(|\Sigma|)$.

Let $\mathcal{L}$ be a language. We next define the notion of an unambiguous interactive proof of $\mathcal{L}$.

Definition 7 (ϵ -unambiguous $(\ell, a, b, \text{Ptime}, \text{Vtime}, \Sigma)$ -UIP). An $(\ell, a, b, \text{Ptime}, \text{Vtime}, \Sigma)$ -protocol (P, V) is an *unambiguous interactive proof* (UIP) with an unambiguity error ϵ for a language $\mathcal{L}$, if it satisfies the following completeness and unambiguity conditions.

- **Prescribed Completeness:** For any $x \in \{0, 1\}^n$ and any verifier coins $\mathbf{q} = (q_1, \dots, q_\ell) \in (\Sigma^b)^\ell$, the answers $\mathbf{a} = P(x, \mathbf{q}) \in (\Sigma^a)^\ell$ satisfy that $V(x, \mathbf{q}, \mathbf{a}) = 1$ iff $x \in \mathcal{L}$.

Remark 8. Completeness is usually defined only for $x \in \mathcal{L}$. *Prescribed completeness* also requires the prescribed prover to send rejected messages on $x \notin \mathcal{L}$. We assume that when $x \notin \mathcal{L}$, P sends the default string $\mathbf{0} \in \Sigma^a$ in every round, and V rejects this transcript.

- **ϵ -Unambiguity:** For any $x \in \{0, 1\}^n$, any $j^* \in [\ell]$, any prefix $\mathbf{q}_{\leq j^*} = (q_1, \dots, q_{j^*})$, and any (computationally unbounded) prover strategy P^* that deviates from P firstly in round $j^* \in [\ell]$, namely

$$P^*(x, j^*, \mathbf{q}_{\leq j^*}) \neq P(x, j^*, \mathbf{q}_{\leq j^*}) \quad \text{and} \quad P^*(x, j, \mathbf{q}_{\leq j}) = P(x, j, \mathbf{q}_{\leq j}) \quad \text{for all } j < j^*,$$

we have that P^* is accepted with probability at most $\epsilon = \epsilon(n)$ over the remaining random coins of V . Namely,

$$\Pr[V(x, (\mathbf{q}_{\leq j^*}, \mathbf{q}_{> j^*}), \mathbf{a}) = 1] \leq \epsilon,$$

where the probability is taken over $\mathbf{q}_{> j^*} = (q_{j^*+1}, \dots, q_\ell)$, and where $\mathbf{a} = P^*(x, (\mathbf{q}_{\leq j^*}, \mathbf{q}_{> j^*}))$.

Note that by Remark 8, unambiguity implies standard soundness.

3.3.1 The GKR protocol

We restate the main result from [GKR15] (in the language of PVAL)¹¹. We note that unambiguity is proven in [JKKZ21].

Theorem 3 (The GKR protocol). *Let $\mathbb{F}$ be a field, and let $\Phi : \mathbb{F}^n \rightarrow \mathbb{F}$ be an arithmetic circuit with addition and multiplication gates of fan-in 2 over $\mathbb{F}$, with description $\langle \Phi \rangle$. Let G be the log-space uniform circuit that outputs $\Phi(x)$ on input $x \in \mathbb{F}^n$ and $\langle \Phi \rangle \in \{0, 1\}^{|\langle \Phi \rangle|}$. Denote the depth and size of G by $D = D(n) \geq \log n$, and $S = S(n) \geq n$.*

¹¹In what follows we formulate the GKR protocol as a reduction (as opposed to a proof system). Specifically, at the end of the protocol the verifier either rejects or outputs a PVAL claim.

There exists an $(\ell, a, b, \text{Ptime}, \text{Vtime})$ -protocol (P_{GKR}, V_{GKR}) such that the following holds. The prover and verifier get the description $\langle \Phi \rangle$, and the prover additionally gets the input $x \in \mathbb{F}^n$. Let $m = \lceil \log n \rceil$. At the end, either V_{GKR} rejects, or both parties output a point $\mathbf{j} \in \mathbb{F}^m$ and a value $v \in \mathbb{F}$, such that:

- **Prescribed Completeness:** If both parties follow the protocol then $\hat{x}(\mathbf{j}) = v$ iff $\Phi(x) = 1$.
- **ϵ_{GKR} -Unambiguity:** For any (unbounded) cheating prover strategy P^* that deviates from P_{GKR} firstly in round $j^* \in [\ell]$,

$$\epsilon_{GKR}(D, \log S, |\mathbb{F}|) := \Pr[V_{GKR} \text{ accepts} \wedge \hat{x}(\mathbf{j}) = v] = O\left(\frac{D \log S}{|\mathbb{F}|}\right)$$

where the probability is taken over the verifier's remaining coins (after round j^*). In other words, with probability at least $1 - \epsilon_{GKR}(D, \log S, |\mathbb{F}|)$, either V_{GKR} rejects or $x \notin \text{PVAL}(\mathbf{j}, v)$.

- Regardless of the prover's strategy, $\mathbf{j}$ is uniformly distributed in $\mathbb{F}^m$, and only depends on the random coins of V_{GKR} .

The complexity of the protocol is as follows (with $\tilde{O}$ ignoring polylogarithmic factors in $D, \log S$):

- $\ell = O(D \cdot \log S)$.
- $a = O(\log |\mathbb{F}|)$.
- $b = O(\log |\mathbb{F}|)$.
- $\text{Ptime} = \tilde{O}(\text{poly}(S) \cdot \text{polylog}|\mathbb{F}|)$.
- $\text{Vtime} = \tilde{O}(D \log S \cdot \log |\mathbb{F}| + |\langle \Phi \rangle| \cdot \log |\mathbb{F}|)$ (and V does not access x).

Furthermore, the verifier's verdict (on whether to output the strings $\mathbf{j}, v$ or reject at the end) has the following complexities (with $\tilde{O}$ ignoring polylogarithmic factors in $D, \log S, \log |\mathbb{F}|$):

- $\text{size}(V) = \tilde{O}(D \log S \cdot \log |\mathbb{F}| + |\langle \Phi \rangle| \cdot \log |\mathbb{F}|)$.
- $\text{depth}(V) = \tilde{O}(1)$.

Remark 9 (The role of the circuit description $\langle \Phi \rangle$). After running the original GKR protocol from [GKR15] to the circuit $G : \mathbb{F}^{n+|\langle \Phi \rangle|} \rightarrow \mathbb{F}$ and input $(x, \langle \Phi \rangle)$ (which verifies the claim $\Phi(x) = 1$), the verifier obtains $\mathbf{j}' \in \mathbb{F}^{\lceil \log(n+|\langle \Phi \rangle|) \rceil}$, $v' \in \mathbb{F}$ and has to verify the claim $\widehat{x||\langle \Phi \rangle}(\mathbf{j}') = v'$ (where $\widehat{x||\langle \Phi \rangle}$ denotes the LDE of the pair $(x, \langle \Phi \rangle)$ as a string under some appropriate encoding). The protocol in Theorem 3 additionally reduces this to verifying $\hat{x}(\mathbf{j}) = v$ for some $\mathbf{j}, v$. This is without additional communication as follows: using $\langle \Phi \rangle$, the verifier first computes $v_0 = \widehat{\mathbf{0}^n || \langle \Phi \rangle}(\mathbf{j}')$, and subtracts it from v' , which yields the claim $\widehat{x||\langle \Phi \rangle}(\mathbf{j}') = v' - v_0$. This is equivalent to checking that $\hat{x}(\mathbf{j}) = \frac{v' - v_0}{\chi(\mathbf{j}')} \in \mathbb{F}$ where $\chi(\mathbf{j}')$ is a factor that depends only on $\mathbf{j}'$ and computable by the verifier on its own. The post-processing procedure amounts to the additional $O(|\langle \Phi \rangle| \cdot \log |\mathbb{F}|)$ verifier runtime.

The GKR protocol is Δ_c -distance-preserving Let $(P_{\text{GKR}}, V_{\text{GKR}})$ be the GKR protocol from Theorem 3. In [RVW13], the authors observed that when $(P_{\text{GKR}}, V_{\text{GKR}})$ is parallel-repeated T times for a large enough T , it is Δ_c -distance-preserving in the following sense: suppose the input x is d -far from satisfying Φ , and $\mathbf{j} = (j_1, \dots, j_T)$ and $\mathbf{v} = (v_1, \dots, v_T)$ are the outputs of T parallel repetitions of $(P_{\text{GKR}}, V_{\text{GKR}})$, then x is d -far from the set $\text{PVAL}(\mathbf{j}, \mathbf{v})$. Note that this increases the overall cost of the protocol by a factor of T .

This observation generalizes to Δ_c -distance, captured in the following lemma.

Lemma 2 (GKR is Δ_c -Distance-Preserving). *Let $M, L, d \in \mathbb{N}$ and assume $M = 2^m$ for some $m \in \mathbb{N}$. Let $\mathbb{F}$ be a field of characteristic 2. Let $\Phi : \mathbb{F}^{M \times L} \rightarrow \{0, 1\}$ be a log-space uniform arithmetic circuit over $\mathbb{F}$, with addition and multiplication gates of fan-in 2, of size S and depth D , and with description $\langle \Phi \rangle$.*

Let $\mathbf{a} \in \mathbb{F}^{M \times L}$. Denote by $\mathcal{S} \subset \mathbb{F}^{M \times L}$ the set of matrices accepted by Φ , and suppose that

$$|\mathcal{S} \cap \mathcal{B}_{d, \mathbb{F}}(\mathbf{a})| \leq 1,$$

i.e. at most one element in $\mathcal{B}_{d, \mathbb{F}}(\mathbf{a})$ satisfies the circuit. (See Definition 1 for the definition of $\mathcal{B}_{d, \mathbb{F}}(\mathbf{a})$.)

When applying $(P_{\text{GKR}}, V_{\text{GKR}})$ to the claim $\Phi(\mathbf{a}) = 1$ ¹² with $T \in \mathbb{N}$ parallel repetitions, we obtain $\mathbf{j} \in (\mathbb{F}^{m + \log L})^T$, $\mathbf{v} \in \mathbb{F}^T$, with the following guarantee (in addition to the ones stated in Theorem 3, with all complexities but the round complexity multiplied by T).

Unambiguous Distance Preservation: *For every prover P^* , the distance $\Delta_c(\mathbf{a}, \mathcal{S})$ is preserved by the interaction (P^*, V_{GKR}) in the following sense:*

** If $\Delta_c(\mathbf{a}, \mathcal{S}) \leq d$ (i.e. $\mathcal{S} \cap \mathcal{B}_{d, \mathbb{F}}(\mathbf{a}) \neq \emptyset$) and P^* answers according to the prescribed strategy P_{GKR} corresponding to the unique $\mathbf{a}^* \in \mathcal{S} \cap \mathcal{B}_{d, \mathbb{F}}(\mathbf{a})$, then the output $(\mathbf{j}, \mathbf{v})$ satisfies*

$$\begin{aligned} \Pr[\Delta_c(\mathbf{a}, \text{PVAL}(\mathbf{j}, \mathbf{v})) \leq d \text{ and } \mathbf{a}^* \text{ remains the unique element in } \mathcal{B}_{d, \mathbb{F}}(\mathbf{a}) \cap \text{PVAL}(\mathbf{j}, \mathbf{v})] \\ \geq 1 - (\epsilon_{\text{GKR}}(D, \log S, |\mathbb{F}|))^T \cdot \left(\binom{M}{d} |\mathbb{F}|^d \right)^L. \end{aligned}$$

** On the other hand, if $\Delta_c(\mathbf{a}, \mathcal{S}) > d$ (i.e. $\mathcal{S} \cap \mathcal{B}_{d, \mathbb{F}}(\mathbf{a}) = \emptyset$) or P^* does not answer according to P_{GKR} corresponding to the unique $\mathbf{a}^* \in \mathcal{B}_{d, \mathbb{F}}(\mathbf{a})$, then the output $(\mathbf{j}, \mathbf{v})$ satisfies*

$$\begin{aligned} \Pr[\Delta_c(\mathbf{a}, \text{PVAL}(\mathbf{j}, \mathbf{v})) > d] \\ \geq 1 - \left(\epsilon_{\text{GKR}}(D, \log S, |\mathbb{F}|) + (\epsilon_{\text{GKR}}(D, \log S, |\mathbb{F}|))^T \cdot \left(\binom{M}{d} |\mathbb{F}|^d \right)^L \right). \end{aligned}$$

Here $\epsilon_{\text{GKR}}(D, \log S, |\mathbb{F}|) = O(\frac{D \log S}{|\mathbb{F}|})$ is the soundness error of one GKR protocol.

Proof. Fix $\mathbf{a}' \in \mathcal{B}_{d, \mathbb{F}}(\mathbf{a})$, and consider the following two cases:

Case 1 $\Phi(\mathbf{a}') = 0$.

The claim $\Phi(\mathbf{a}') = 1$ is false. For any prover strategy P^* , by GKR's standard soundness amplification, $\Pr[\mathbf{a}' \in \text{PVAL}(\mathbf{j}, \mathbf{v})] \leq (\epsilon_{\text{GKR}}(D, \log S, |\mathbb{F}|))^T$.

¹²Recall that both P_{GKR} and V_{GKR} get the description $\langle \Phi \rangle$. Additionally, P_{GKR} gets the input $\mathbf{a}$.

Case 2 $\Phi(\mathbf{a}') = 1$.

If the prover P^* deviates from the prescribed prover P with respect to $\mathbf{a}^*$ in some round j^* , then it must be deviating in one of the T parallel executions of GKR. By the unambiguity of GKR¹³, $\Pr[\mathbf{a}' \in \text{PVAL}(\mathbf{j}, \mathbf{v})] \leq \epsilon_{\text{GKR}}(D, \log S, |\mathbb{F}|)$.

Crucially, note that there are at most $\left(\binom{M}{d} |\mathbb{F}|^d\right)^L$ points $\mathbf{a}'$ in $\mathcal{B}_{d,\mathbb{F}}(\mathbf{a})$. Let E denote the event that $\exists \mathbf{a}' \in \mathcal{B}_{d,\mathbb{F}}(\mathbf{a})$ s.t. $\Phi(\mathbf{a}') = 0$ and $\mathbf{a}' \in \text{PVAL}(\mathbf{j}, \mathbf{v})$. By Case 1 and a union bound,

$$\Pr[E] \leq \epsilon' := \epsilon_{\text{GKR}}(D, \log S, |\mathbb{F}|)^T \cdot \left(\binom{M}{d} |\mathbb{F}|^d \right)^L.$$

Using this bound on E , we have

- * When $\mathcal{S} \cap \mathcal{B}_{d,\mathbb{F}}(\mathbf{a}) = \{\mathbf{a}^*\}$ and P^* answers according to the prescribed strategy P_{GKR} corresponding to $\mathbf{a}^*$:

In this case, any other $\mathbf{a}' \in \mathcal{B}_{d,\mathbb{F}}(\mathbf{a})$ satisfies $\Phi(\mathbf{a}') = 0$. Additionally, by the prescribed completeness of the GKR protocol, $\mathbf{a}^* \in \text{PVAL}(\mathbf{j}, \mathbf{v})$. Therefore,

$$\begin{aligned} & \Pr \left[\begin{array}{l} \Delta_c(\mathbf{a}, \text{PVAL}(\mathbf{j}, \mathbf{v})) \leq d \text{ and} \\ \mathbf{a}^* \text{ remains the unique element} \\ \text{in } \mathcal{B}_{d,\mathbb{F}}(\mathbf{a}) \cap \text{PVAL}(\mathbf{j}, \mathbf{v}) \end{array} \right] \\ & \geq \Pr[\mathbf{a} \in \text{PVAL}(\mathbf{j}, \mathbf{v}) \text{ and } \overline{E}] \\ & = 1 - \epsilon'. \end{aligned}$$

- * When $\mathcal{S} \cap \mathcal{B}_{d,\mathbb{F}}(\mathbf{a}) = \{\mathbf{a}^*\}$ but P^* does not answer according to P_{GKR} corresponding to $\mathbf{a}^*$:
By Case 2, the probability that $\mathbf{a}^* \in \text{PVAL}(\mathbf{j}, \mathbf{v})$ is at most $\epsilon_{\text{GKR}}(D, \log S, |\mathbb{F}|)$. Therefore,

$$\begin{aligned} & \Pr[\mathcal{B}_{d,\mathbb{F}}(\mathbf{a}) \cap \text{PVAL}(\mathbf{j}, \mathbf{v}) = \emptyset] \\ & \geq 1 - (\Pr[\mathbf{a}^* \in \text{PVAL}(\mathbf{j}, \mathbf{v})] + \Pr[E]) \\ & \geq 1 - (\epsilon_{\text{GKR}}(D, \log S, |\mathbb{F}|) + \epsilon'). \end{aligned}$$

- * When $\mathcal{S} \cap \mathcal{B}_{d,\mathbb{F}}(\mathbf{a}) = \emptyset$, no matter what P^* does,

$$\begin{aligned} \Pr[\mathcal{B}_{d,\mathbb{F}}(\mathbf{a}) \cap \text{PVAL}(\mathbf{j}, \mathbf{v}) = \emptyset] & \geq 1 - \Pr[E] \\ & \geq 1 - \epsilon'. \end{aligned}$$

□

3.4 Almost d -wise independent permutations

We begin by defining *almost d -wise independent permutations*:

Definition 8 (Almost d -wise Independent Permutations). A distribution Π on S_N ¹⁴ is η -almost d -wise independent if for all $1 \leq x_1 < x_2 < \dots < x_d \leq N$, the distribution $(\pi(x_1), \dots, \pi(x_d))_{\pi \sim \Pi}$ is at most η far from the uniform distribution on distinct d -tuples over $[N]$ in TV distance.

¹³Note that P^* does not necessarily deviate in all T executions, so the error term is not necessarily amplified.

¹⁴The symmetric group on N elements i.e. the set of all permutations

We will be using a family of almost d -wise independent permutations on $\{0, 1\}^m$ in Section 6, which has a *pairwise independence* property and has succinct descriptions. The properties we will desire from our permutation family are:

- Almost d -wise independence.
- $\pi \sim \Pi$ have succinct descriptions, short seeds, and can be efficiently implemented.
- π^{-1} for $\pi \sim \Pi$ have succinct descriptions, short seeds, and can be efficiently implemented.

The question of sampling from almost d -wise independent permutations has been well studied [Gow96, HMMR04, BH08, KNR05, GHKO25, GHP25]. From this line of work, we get the following result:

Theorem 4 (Theorem 2 in [GHP25]). *For any m , and $d \leq 2^{m/50}$, a random reversible circuit with $\tilde{O}(dm \cdot \log(1/\eta))$ width-2 gates computes an η -almost d -wise independent permutation where $\tilde{O}$ hides $\text{polylog}(d, m)$ factors.*

Observe that since the sampled circuit is a random reversible circuit, it is always a permutation, the inverse is easy to compute, and the seed is just the set of gates which can also be described with $\tilde{O}(nk \cdot \log(1/\eta))$ bits and the permutation is also efficient to compute.

Definition 9 ($\Pi_{m,d,\eta}$). We let $\Pi_{m,d,\eta}$ be the distribution from Theorem 4.

Definition 10 (Uniform distributions). Let $\mathcal{U}_{m,d}$ be the uniform distribution of distinct d -tuples over $[2^m]$. Let $\mathcal{U}'_{m,d}$ be the uniform distribution of d -tuples over $[2^m]$.

Theorem 5. *Let $A_0, A_1 \subseteq \{0, 1\}^m$ such that $|A_0|, |A_1| \leq d$ such that $d \leq 2^{m/50}$. Then,*

$$\Pr_{\pi \sim \Pi_{m,d,\eta}} [|A_0 \cap \pi(A_1)| \geq \epsilon d] \leq \eta + \exp(-\epsilon d/6)$$

if $\epsilon d \geq 2$.

Proof. Observe that since the indicator random variable for $|A_0 \cap \pi(A_1)| \geq \epsilon d$ is Bernoulli, and the TV distance, we have that

$$\Pr_{\pi \sim \Pi_{m,d,\eta}} [|A_0 \cap \pi(A_1)| \geq \epsilon d] \leq \eta + \Pr_{\pi \sim \mathcal{U}_{m,d}} [|A_0 \cap \pi(A_1)| \geq 1 + \epsilon d].$$

For any $a \in A_1$, let X_a, X'_a be the indicator random variables for $\pi(a) \in A_0$ where $\pi \sim \mathcal{U}_{m,d}$ and $\mathcal{U}'_{m,d}$ respectively. Additionally, let $X = \sum X_a$ and $X' = \sum X'_a$. Observe that for any $A \subset A_0$,

$$\mathbb{E}[\Pi_{a \in A} X_a] = \Pr[\forall a \in A, X_a = 1] = \prod_{j=0}^{|A|-1} \frac{|A_0| - j}{2^m - j} \leq \left(\frac{|A_0|}{2^m} \right)^{|A|} \leq \mathbb{E}[\Pi_{a \in A} X'_a].$$

Thus, for any integral $c \geq 0$, we have that $\mathbb{E}[X^c] \geq \mathbb{E}[X'^c]$. Now, since we get Chernoff bound by applying Markov Inequality on the moment generating function, we must have that

$$\Pr[X \geq \mu(1 + \frac{\epsilon d}{\mu} - 1)] \leq e^{(1-\epsilon d) \cdot \frac{\frac{\mu}{\epsilon d} - 1}{1 + \frac{\epsilon d}{\mu}}} \leq e^{-\epsilon d/6}$$

where $\mu = \mathbb{E}[X] = \frac{|A_0| \cdot |A_1|}{2^m} \leq \frac{d^2}{2^m} \leq 1$. □

3.5 Unambiguous Interactive Proofs of Proximity (IPP)

Defined in [EKR04, RVW13], Interactive Proofs of Proximity (IPP) for the PVAL problem are a key ingredient in several prior works (for example, [RVW13, RRR18, RR20]). They are also used in our protocols. Here, we define a slight variant of the standard IPP with the added unambiguity property.

Definition 11 ([EKR04, RVW13] (Unambiguous) Interactive Proof of Proximity). Let $d \in \mathbb{N}$. A pair of interacting Turing machines (P, V) , where P is deterministic and takes as input a string $x \in \{0, 1\}^n$ and V has oracle access to x , is an *unambiguous interactive proof of proximity with distance d* (d -IPP) for a language $\mathcal{L}$, under distance dist , with perfect completeness and unambiguity error $\epsilon = \epsilon(n)$ if:

- **Prescribed Completeness:** For all $x \in \{0, 1\}^n$, the verifier accepts in the interaction $(P(x), V^x(1^n))$ iff $\text{dist}(x, \mathcal{L}) < d$.
- **ϵ -Unambiguity:** For all $x \in \{0, 1\}^n$, and all provers P^* that deviates from P in some round j^* ,

$$\Pr[V \text{ accepts the interaction with } P^*] \leq \epsilon(n),$$

where the probability is taken over the verifier's random coins after round j^* .

The query complexity $q(n)$ is the number of queries the verifier makes to x , the communication and round complexity of the protocol are as usual.

4 Our Batch-UIP

Let $k \in \mathbb{N}$. For a language $\mathcal{L}$, let its batched version $\mathcal{L}^{\otimes k}$ be

$$\mathcal{L}^{\otimes k} := \{(x_1, \dots, x_k) : \forall i \in [k], x_i \in \mathcal{L} \text{ and } |x_1| = \dots = |x_k|\}.$$

If $\mathcal{L}$ has a UIP (Definition 7), then $\mathcal{L}^{\otimes k}$ has a UIP obtained by running the underlying UIP for $\mathcal{L}$ for each of the k instances in parallel. This introduces a factor of k overhead. Our main result is a better batching protocol. For simplicity, assume the alphabet is $\Sigma = \{0, 1\}$.

Theorem 6 (Formal version of Theorem 1). *Let $\mathcal{L}$ be a language with a (public-coin) ϵ -unambiguous $(\ell, a, b, \text{Ptime}, \text{Vtime})$ -UIP (P, V) , where the verifier verdict V is a log-space uniform boolean circuit with size S and depth D . We also make the simplifying assumptions that the verifier never rejects in the middle of the UIP, and that $b \leq a$.*

Let $\mathbf{x} = (x_1, \dots, x_k) \in (\{0, 1\}^n)^k$ be a batch of statements in $\mathcal{L}^{\otimes k}$ to be verified, let $\sigma = \sigma(k) \in \mathbb{N}$ be an unambiguity parameter. Then there exists an ϵ' -unambiguous $(\ell', a', b', \text{Ptime}', \text{Vtime}', \{0, 1\})$ -UIP for the language $\mathcal{L}^{\otimes k}$, denoted by (P', V') (Protocol 1), with the following properties, where $\tilde{O}$ hides $\text{poly}(\sigma) \cdot \text{polylog}(k, n, a, \ell)$ factors:

- *The unambiguity error is $\epsilon' = \tilde{O}(\epsilon + 2^{-\sigma})$.*
- *The round complexity is $\ell' = \tilde{O}(\ell \cdot D \log S)$.*
- *The prover and verifier message length per round is $a' = b' = \tilde{O}(a\ell + \text{poly}(\ell))$.*

- The prover runtime is $\text{Ptime}' = \tilde{O}(k \cdot \ell \cdot \text{Ptime} + \text{poly}(k, a, \ell, S))$.
- The verifier runtime is $\text{Vtime}' = \tilde{O}(\ell \cdot \text{Vtime} + a\ell^2(kn + D \log S) + a^2 \cdot \text{poly}(\ell))$.

Furthermore, the verifier's verdict circuit is log-space uniform and has the following properties:

- $\text{size}(\mathbf{V}') = \tilde{O}(\ell \cdot S + a\ell^2(kn + D \log S) + a^2 \cdot \text{poly}(\ell))$.
- $\text{depth}(\mathbf{V}') = D + \tilde{O}(1)$.

The unambiguity parameter σ controls the unambiguity error ϵ' , which we show to be concretely upper-bounded by $2 \log k \cdot (\epsilon + 2^{-\sigma})$ in Section 4.2. Increasing σ imposes an overhead factor of $O(2^\sigma)$ in the prover's message length and the prover's and verifier's runtime.

If $D \log S = \text{polylog}(k, n)$, then Theorem 1 follows by choosing $\sigma = \log \frac{1}{\epsilon}$. With $\tilde{O}$ hiding $\text{polylog}(\sigma^{-1}, k, n, a, \ell)$ factors,

- $\epsilon' = O(\epsilon \log k)$.¹⁵
- $\ell' = \tilde{O}(\ell)$,
- $a' = b' = \tilde{O}(a\ell)$,
- $\text{Ptime}' = \tilde{O}(k \cdot \ell \cdot \text{Ptime} + \text{poly}(k, a, \ell, S))$,
- $\text{Vtime}' = \tilde{O}(\ell \cdot \text{Vtime} + kna\ell^2 + a^2 \cdot \text{poly}(\ell))$. Note that Theorem 1 uses the looser bound $kna\text{poly}(\ell)$ for the second and third term.

4.1 Our Ingredients

Our Construction Our Batch-UIP protocol consists of two main ingredients: a *Distance Generation Protocol* and an *Instance Reduction Protocol*. It iteratively applies these two protocols, where in each iteration, the *Distance Generation Protocol* ($\mathbf{P}_{\text{Dist}}, \mathbf{V}_{\text{Dist}}$) is applied, followed by the application of *Instance Reduction Protocol* ($\mathbf{P}_{\text{Reduce}}, \mathbf{V}_{\text{Reduce}}$).

The Distance Generation Protocol (Lemma 3) ensures that if the k -fold input instance was not in the language, then after running it, the prover and verifier obtain a related instance that is far from the language. While the Reduction Protocol (Lemma 4) has the guarantee that as long as the given k -fold input instance is far from the language, it produces a significantly smaller k' -fold instance not in the language (but without any distance guarantees). Therefore, as long as the overall shrinkage is significant, say larger than a factor of 2, the parties can repeat the two sub-protocols $\rho = O(\log k)$ times, and end up having roughly $k \cdot 2^{-\rho} = \text{polylog}(k)$ instances that can be explicitly checked.

In the initial stage, the verifier samples a random $\mathbf{q} \leftarrow \{0, 1\}^{b\ell}$, and sends the entire $\mathbf{q}$ to the prover. This $\mathbf{q}$, along with the statements $(x_1, \dots, x_k)$, implicitly define a set of k “prescribed answers”: For each i , denote by $\mathbf{a}^{x_i, \mathbf{q}} = \mathbf{P}(x_i, \mathbf{q}) \in \{0, 1\}^{a\ell}$ the answers the base UIP's prescribed prover $\mathbf{P}$ would send, on input statement x_i upon receiving $\mathbf{q}$ from the verifier.

We denote by $\Phi^{(0)}$ the boolean circuit that has $\mathbf{q}$ and $(x_1, \dots, x_k)$ hardwired. It takes the input $\mathbf{a}^{\mathbf{q}} := \{\mathbf{a}^{x_i, \mathbf{q}}\}_{i \in [k]}$ and outputs 1 iff $\mathbf{V}(x_i, \mathbf{q}, \mathbf{a}^{x_i, \mathbf{q}}) = 1$ for all $i \in [k]$. We treat $\mathbf{a}^{\mathbf{q}}$ as a matrix in $\{0, 1\}^{k \times (a\ell)}$ whose i 'th row is $\mathbf{a}^{x_i, \mathbf{q}}$, and call $\mathbf{a}^{\mathbf{q}}$ the *prescribed matrix* corresponding to $\mathbf{q}$.

¹⁵We require $\epsilon = O(1/\log k)$ for the unambiguity error to be non-trivially small.

Let $\mathcal{L}'_x$ be the *associated language* that consists of all instances $(\mathcal{S}, \Phi)$, where

$$\mathcal{S} = \{(i^1, \mathbf{q}^1), \dots, (i^g, \mathbf{q}^g)\}$$

is a set of g pairs in $[k] \times \{0, 1\}^{b\ell}$, such that the prescribed matrix $\mathbf{a}^{\mathcal{S}} \in \{0, 1\}^{g \times (a\ell)}$, consisting of rows $\{\mathbf{a}^{x_{ij}, \mathbf{q}^j}\}_{j \in [g]}$ where $\mathbf{a}^{x_{ij}, \mathbf{q}^j} = \mathbf{P}(x_{ij}, \mathbf{q}^j)$, satisfies the constraint Φ . Note that the language $\mathcal{L}'_x$ is not necessarily in NP, but as we shall see in Section 4.5.2, there is a UIP for explicitly checking it. We mention that the prescribed matrix $\mathbf{a}^{\mathcal{S}}$ replaces the role of the implicit vector of unique witnesses in the case of Batch-UP.

It follows from the prescribed completeness of the base UIP that for $\mathcal{S}^{(0)} := \{(i, \mathbf{q})\}_{i \in [k]}$,

$$(\mathcal{S}^{(0)}, \Phi^{(0)}) \in \mathcal{L}'_x \text{ iff } \mathbf{x} \in \mathcal{L}^{\otimes k}.$$

In each iteration, the following two sub-protocols are run:

1. The Distance Generation Protocol is applied to $(\mathcal{S}^{(0)}, \Phi^{(0)}) \in \mathcal{L}'_x$ with distance parameter $d = \ell \cdot \text{polylog}(k, \ell) \in \mathbb{N}$. At the end of this protocol, the verifier obtains a set $\mathcal{S}_{\text{Dist}}$ containing $M = |\mathcal{S}^{(0)}| \cdot \ell$ pairs, where initially $|\mathcal{S}^{(0)}| = |\{(i, \mathbf{q})\}_{i \in [k]}| = k$, as well as a predicate Φ_{Dist} that is evaluated over the prescribed matrix $\mathbf{a}^{\mathcal{S}_{\text{Dist}}}$. The guarantee of our Distance Generation Protocol is that when $(\mathcal{S}^{(0)}, \Phi^{(0)}) \notin \mathcal{L}'_x$, the prescribed matrix $\mathbf{a}^{\mathcal{S}_{\text{Dist}}}$ is Δ_c - d -far (Definition 1) from satisfying the predicate Φ_{Dist} .
2. Then, the Instance Reduction protocol with shrinkage parameter $d = \ell \cdot \text{polylog}(k, \ell) \in \mathbb{N}$ transforms the claim $(\mathcal{S}_{\text{Dist}}, \Phi_{\text{Dist}}) \in \mathcal{L}'_x$ to a new claim $(\mathcal{S}^{(1)}, \Phi^{(1)}) \in \mathcal{L}'_x$, where $\mathcal{S}^{(1)}$ is a subset of $\mathcal{S}_{\text{Dist}}$ with size $M' \leq M/d$. The Instance Reduction Protocol guarantees that $(\mathcal{S}^{(1)}, \Phi^{(1)}) \notin \mathcal{L}'_x$ if the prescribed matrix $\mathbf{a}^{\mathcal{S}_{\text{Dist}}}$ is Δ_c - d -far (Definition 1) from satisfying Φ_{Dist} . Importantly, the verifier never needs to explicitly access the prescribed matrix $\mathbf{a}^{\mathcal{S}_{\text{Dist}}}$.

We set the distance/shrinkage parameter to be $d > 2\ell$ to offset the intermediate factor of ℓ . In other words,

$$|\mathcal{S}_{\text{Dist}}| \leq \ell \cdot |\mathcal{S}^{(0)}| \quad \text{and} \quad |\mathcal{S}^{(1)}| \leq \frac{|\mathcal{S}_{\text{Dist}}|}{d} < \frac{|\mathcal{S}^{(0)}|}{2},$$

and thus each iteration shrinks the number of instances at least by a factor of 2, which implies that the parties need to iterate at most $\rho = \lfloor \log k \rfloor$ times, until they can check the final claim $(\mathcal{S}^{(\rho)}, \Phi^{(\rho)}) \in \mathcal{L}'_x$, defined over $|\mathcal{S}^{(\rho)}| \leq k/2^\rho$ pairs. The final check can be done by explicitly running the UIP for verifying $(\mathcal{S}^{(\rho)}, \Phi^{(\rho)}) \in \mathcal{L}'_x$ (as formally proven in Lemma 5 below).

The associated language $\mathcal{L}'_x$ We formally set up the notations to define the language $\mathcal{L}'_x$.

Definition 12 (Prescribed transcript relation $\mathcal{R}$). Let $\mathcal{R}$ be the binary relation containing pairs $((x, \mathbf{q}), \mathbf{a}) \in (\{0, 1\}^n \times \{0, 1\}^{b\ell}) \times \{0, 1\}^{a\ell}$, such that

- $\forall(x, \mathbf{q}, \mathbf{a}) = 1$, and
- $\mathbf{a}$ is the *prescribed prover transcript*, i.e. $\mathbf{a} = \mathbf{P}(x, \mathbf{q}) \in \{0, 1\}^{a\ell}$ is the prescribed prover messages $(\mathbf{a}_1^{x, \mathbf{q}}, \dots, \mathbf{a}_\ell^{x, \mathbf{q}})$ sent by $\mathbf{P}$ on input x when interacting with $\mathbf{V}$ that uses $\mathbf{q}$ as its random coins.¹⁶

¹⁶Recall that if $x \notin \mathcal{L}$, $\mathbf{a}^{x, \mathbf{q}} = \mathbf{P}(x, \mathbf{q})$ is the default all-zero string, and $\forall(x, \mathbf{q}, \mathbf{a}^{x, \mathbf{q}}) = 0$.

Notations (Prescribed prover messages $\mathbf{a}^{x,q}$, $\mathbf{a}_j^{x,q}$, and matrices $\mathbf{a}^S, \mathbf{a}_j^S$) For any $x \in \{0,1\}^n$ and any $\mathbf{q} \in (\{0,1\}^b)^\ell$, the shorthand $\mathbf{a}^{x,q} \in \{0,1\}^{a\ell}$ stands for the (unique) prescribed prover messages corresponding to $(x, \mathbf{q})$. Namely, $\mathbf{a}^{x,q} := \mathbf{P}(x, \mathbf{q})$. Similarly, for every $j \in [\ell]$, the shorthand $\mathbf{a}_j^{x,q}$ means the j -th round prover message in $\mathbf{a}^{x,q} := (\mathbf{a}_1^{x,q}, \dots, \mathbf{a}_\ell^{x,q})$. More generally, given a batch of statements $\mathbf{x} = (x_1, \dots, x_k) \in (\{0,1\}^n)^k$, and a subset of g pairs $\mathcal{S} \subset [k] \times (\{0,1\}^b)^\ell$, we denote by $\mathbf{a}^S := ((\mathbf{a}^{x_i,q})_{(i,q) \in \mathcal{S}})^\top$ the matrix whose rows are $\mathbf{a}^{x_i,q} \in \{0,1\}^{a\ell}$. For every $j \in [\ell]$, $\mathbf{a}_j^S := ((\mathbf{a}_j^{x_i,q})_{(i,q) \in \mathcal{S}})^\top$ is the matrix whose rows are $\mathbf{a}_j^{x_i,q} \in \{0,1\}^a$, the j -th round prescribed message in $\mathbf{a}^{x_i,q}$.

Note that $\mathcal{R}$ is not necessary a UP (or even NP) relation, but for every $x \in \mathcal{L}$ and every $\mathbf{q} \in \{0,1\}^{b\ell}$, there is a single $\mathbf{a}$ that satisfies $((x, \mathbf{q}), \mathbf{a}) \in \mathcal{R}$. As we shall see in Proposition 2, when $x \in \mathcal{L}$, there exists a UIP for certifying that $\mathbf{a}$ is indeed the (unique) $\mathbf{a}^{x,q}$ for which $((x, \mathbf{q}), \mathbf{a}) \in \mathcal{R}$. Therefore, in this sense, we can treat $\mathbf{a}^{x,q}$ as the unique witness for the relation $((x, \mathbf{q}), \mathbf{a}) \in \mathcal{R}$.

We now formally define the associated language $\mathcal{L}'_x$. For $(\mathcal{S}, \Phi) \in \mathcal{L}'_x$ to hold, not only should all remaining x_i 's be in $\mathcal{L}$, a joint constraint Φ should also be satisfied on the prescribed matrix $\mathbf{a}^S$ associated with $\mathcal{S}$.

Definition 13 (The associated language $\mathcal{L}'_x$). Let $\mathbf{x} \in (\{0,1\}^n)^k$, $\mathcal{S} \subset [k] \times \{0,1\}^{b\ell}$ be a set of $g := |\mathcal{S}|$ pairs, and Φ be a boolean circuit that takes $\mathbf{a}^S$ as input. Let $\mathcal{I} := \{i : (i, \mathbf{q}) \in \mathcal{S}\}$. Then

$$(\mathcal{S}, \Phi) \in \mathcal{L}'_x \quad \text{iff} \quad \Phi(\mathbf{a}^S) = 1 \quad \wedge \quad \mathbf{x}|_{\mathcal{I}} \in \mathcal{L}^{\otimes |\mathcal{I}|}.$$

4.1.1 The Distance Generation Protocol

We give the full specification of $(\mathbf{P}_{\text{Dist}}, \mathbf{V}_{\text{Dist}})$ and its analysis in Section 4.3, which has the guarantees stated in Lemma 3 below.

The protocol takes as input $\langle \mathcal{S} \rangle, \langle \Phi \rangle$, an unambiguity parameter $\sigma \in \mathbb{N}$, a distance parameter $d \in \mathbb{N}$ and a field parameter $\mathbb{F}$ with $|\mathbb{F}| \geq 2^\sigma \cdot \text{polylog}(k, n, a, \ell)$, and outputs $\langle \mathcal{S}_{\text{Dist}} \rangle, \langle \Phi_{\text{Dist}} \rangle$ such that $(\mathcal{S}_{\text{Dist}}, \Phi_{\text{Dist}}) \in \mathcal{L}'_x$ iff $(\mathcal{S}, \Phi) \in \mathcal{L}'_x$. It has a slightly increased instance size $|\mathcal{S}_{\text{Dist}}| = |\mathcal{S}| \cdot \ell$, but has the guarantee that if $\Phi(\mathbf{a}^S) = 0$ then $\mathbf{a}^{\mathcal{S}_{\text{Dist}}}$ is Δ_c - d -far from satisfying the circuit Φ_{Dist} ,¹⁷ where $\mathbf{a}^S$ and $\mathbf{a}^{\mathcal{S}_{\text{Dist}}}$ are the prescribed matrices corresponding to $\mathcal{S}$ and $\mathcal{S}_{\text{Dist}}$, respectively.

Lemma 3 (Δ_c -Distance Generation Protocol). *Let $\mathcal{L}$ be a language with an ϵ -unambiguous $(\ell, a, b, \text{Ptime}, \text{Vtime})$ -UIP $(\mathbf{P}, \mathbf{V})$. Let S and D be the size and depth of the decision circuit $\mathbf{V}$.*

There exists a protocol $(\mathbf{P}_{\text{Dist}}, \mathbf{V}_{\text{Dist}})$ (Protocol 2) that takes as input $\mathbf{x}, \langle \mathcal{S} \rangle, \langle \Phi \rangle$, where $\mathbf{x} \in (\{0,1\}^n)^k$, $\mathcal{S} \subset [k] \times \{0,1\}^{b\ell}$, and $\Phi : \{0,1\}^{|\mathcal{S}| \times a\ell} \rightarrow \{0,1\}$, an unambiguity parameter $\sigma \in \mathbb{N}$, a distance parameter $d \in \mathbb{N}$, a field $\mathbb{F}$ with $|\mathbb{F}| \geq 2^\sigma \cdot \text{polylog}(k, n, a, \ell)$, and outputs $\mathbf{x}, \langle \mathcal{S}_{\text{Dist}} \rangle, \langle \Phi_{\text{Dist}} \rangle$ such that $|\mathcal{S}_{\text{Dist}}| = |\mathcal{S}| \cdot \ell$, and the following holds.

- **Prescribed Completeness:**

In an honest execution, $(\mathcal{S}_{\text{Dist}}, \Phi_{\text{Dist}}) \in \mathcal{L}'_x$ iff $(\mathcal{S}, \Phi) \in \mathcal{L}'_x$.

- **$(\epsilon + 2^{-\sigma})$ -Unambiguous Generation of Δ_c -Distance:** *For every cheating prover $\mathbf{P}^*$ that deviates from $\mathbf{P}_{\text{Dist}}$ first in round j^* , conditioned on the first j^* rounds of the protocol,*

$$\Pr \left[\exists \mathbf{a}^* \in \mathcal{B}_{d, \mathbb{F}}(\mathbf{a}^{\mathcal{S}_{\text{Dist}}}) \text{ s.t. } \Phi_{\text{Dist}}(\mathbf{a}^*) = 1 \right] \leq \epsilon + 2^{-\sigma},$$

¹⁷Recall that $\mathbf{a}^{\mathcal{S}_{\text{Dist}}}$ is Δ_c - d -far from satisfying Φ_{Dist} , if for every $\mathbf{a}^* \in \mathcal{B}_{d, \mathbb{F}}(\mathbf{a}^{\mathcal{S}_{\text{Dist}}})$, it holds that $\Phi_{\text{Dist}}(\mathbf{a}^*) = 0$.

where $\langle \mathcal{S}_{\text{Dist}} \rangle, \langle \Phi_{\text{Dist}} \rangle$ denotes the output of the interaction between V_{Dist} and P^* , and where the probability is over the remaining coins of V_{Dist} .

Furthermore, for any (honest or cheating) prover P^* , $\mathbf{a}^{\mathcal{S}_{\text{Dist}}}$ is likely the only input in $\mathcal{B}_{d,\mathbb{F}}(\mathbf{a}^{\mathcal{S}_{\text{Dist}}})$ that satisfies Φ_{Dist} . Formally,

$$\Pr \left[\exists \mathbf{a}^* \in \mathcal{B}_{d,\mathbb{F}}(\mathbf{a}^{\mathcal{S}_{\text{Dist}}}) : \mathbf{a}^* \neq \mathbf{a}^{\mathcal{S}_{\text{Dist}}} \wedge \Phi_{\text{Dist}}(\mathbf{a}^*) = 1 \right] \leq \epsilon + 2^{-\sigma},$$

where the probability is over the random coins of V_{Dist} .

Let $\tilde{O}$ be hiding $\text{polylog}(|\mathbb{F}|, k, n, a, \ell)$ factors. Denote by $|\mathcal{S}| = g$, then

- $\ell_{\text{Dist}} = O(\ell)$.
- $a_{\text{Dist}} = \tilde{O}(da)$.
- $b_{\text{Dist}} = \tilde{O}(\max(b, d)) = \tilde{O}(da)$, with the simplifying assumption that $b \leq a$.
- $\text{Ptime}_{\text{Dist}} = \tilde{O}(g\ell \cdot \text{Ptime} + dga\ell^2 + |\langle \Phi \rangle| + |\langle \mathcal{S} \rangle| + k \cdot n)$.
- $\text{Vtime}_{\text{Dist}} = \tilde{O}(dal + |\langle \Phi \rangle| + |\langle \mathcal{S} \rangle| + k \cdot n)$.

The bit lengths of the descriptions are $|\langle \mathcal{S}_{\text{Dist}} \rangle| = \tilde{O}(|\langle \mathcal{S} \rangle| + a\ell)$, $|\langle \Phi_{\text{Dist}} \rangle| = \tilde{O}(|\langle \Phi \rangle| + |\langle \mathcal{S} \rangle| + dal + k \cdot n)$. Let $G(i, \langle \mathcal{S} \rangle)$ be the circuit that returns the i -th element in $\mathcal{S}$, and $C(x, \langle \Phi \rangle)$ be the circuit that returns $\Phi(x)$, then the new circuits $G_{\text{Dist}}(i, \langle \mathcal{S}_{\text{Dist}} \rangle)$ and $C_{\text{Dist}}(x, \langle \Phi_{\text{Dist}} \rangle)$, satisfy the following.

- $\text{size}(G_{\text{Dist}}) = \text{size}(G) + b\ell + \tilde{O}(1)$.
- $\text{depth}(G_{\text{Dist}}) = \text{depth}(G) + \tilde{O}(1)$.
- $\text{size}(C_{\text{Dist}}) = \tilde{O}(g\ell \cdot \text{size}(G) + \text{size}(C) + g \cdot S + dga\ell^2)$.
- $\text{depth}(C_{\text{Dist}}) \leq \max(\text{depth}(C), \text{depth}(G) + D) + \tilde{O}(1)$.

The verifier never rejects this protocol, i.e. V_{Dist} always outputs 1.

4.1.2 The Instance Reduction Protocol

The full specification of $(P_{\text{Reduce}}, V_{\text{Reduce}})$ and its analysis are in Section 4.4. It has the following properties.

The protocol takes as input $\langle \mathcal{S}_{\text{Dist}} \rangle, \langle \Phi_{\text{Dist}} \rangle$, as well the parameters $\sigma, d \in \mathbb{N}$ and the field $\mathbb{F}$ (i.e. the same parameters as $(P_{\text{Dist}}, V_{\text{Dist}})$), and outputs $\langle \mathcal{S}' \rangle, \langle \Phi' \rangle$ such that $|\mathcal{S}'| \approx \frac{|\mathcal{S}_{\text{Dist}}|}{d}$, with the guarantee that $(\mathcal{S}', \Phi') \in \mathcal{L}'_{\mathbf{x}}$ only if $\mathbf{a}^{\mathcal{S}_{\text{Dist}}}$ is Δ_c - d -close to satisfying Φ_{Dist} . Importantly, V_{Reduce} never accesses $\mathbf{a}^{\mathcal{S}_{\text{Dist}}}$ explicitly during this protocol.

Lemma 4 (Instance Reduction Protocol for Δ_c -Distance). *Let $\mathcal{L}$ be a language with an ϵ -unambiguous $(\ell, a, b, \text{Ptime}, \text{Vtime})$ -UIP (P, V) , where the verifier's decision V is a log-space uniform boolean circuit of size S and depth D .*

There exists a constant C , and a protocol $(P_{\text{Reduce}}, V_{\text{Reduce}})$ (Protocol 3) that takes as input $\mathbf{x}, \langle \mathcal{S}_{\text{Dist}} \rangle, \langle \Phi_{\text{Dist}} \rangle$, where $\mathbf{x} \in (\{0, 1\}^n)^k$, $\mathcal{S}_{\text{Dist}} \subset [k] \times \{0, 1\}^{b\ell}$, $\Phi_{\text{Dist}} : \{0, 1\}^{|\mathcal{S}_{\text{Dist}}| \times a\ell} \rightarrow \{0, 1\}$, an unambiguity parameter $\sigma \in \mathbb{N}$, a distance parameter $d \in \mathbb{N}$, a field $\mathbb{F}$ that satisfies $|\mathbb{F}| \geq 32 \cdot 2^\sigma (dal m)^C$, and outputs $\langle \mathcal{S}' \rangle, \langle \Phi' \rangle$, with $|\mathcal{S}'| = \tilde{O}(\frac{|\mathcal{S}_{\text{Dist}}|}{d})$ (where $\tilde{O}$ hides $\text{polylog}(|\mathbb{F}|, k, n, a, \ell)$ factors). Let $M := |\mathcal{S}_{\text{Dist}}| \leq k \cdot \ell$, and suppose $M = 2^m$ for $m \in \mathbb{N}$. If additionally, the technical condition $d \geq 16m\sigma$ holds, then the protocol has the following guarantees.

- **Prescribed Completeness:**

If V_{Reduce} interacts with P_{Reduce} , then $(S', \Phi') \in \mathcal{L}'_x$ iff $(S_{Dist}, \Phi_{Dist}) \in \mathcal{L}'_x$.

- $2^{-\sigma}$ -**Unambiguous Distance Preservation:** Suppose for every $\mathbf{a}^* \in \mathcal{B}_{d, \mathbb{F}}(\mathbf{a}^{S_{Dist}}) \setminus \{\mathbf{a}^{S_{Dist}}\}$, $\Phi_{Dist}(\mathbf{a}^*) = 0$, then for every cheating prover P^* that deviates from P_{Reduce} first in round j^* , let the output of the interaction between V_{Reduce} and P^* be $\langle S' \rangle, \langle \Phi' \rangle$, then

$$\Pr[V_{Reduce} \text{ accepts} \wedge \Phi'(\mathbf{a}^{S'}) = 1] \leq 2^{-\sigma},$$

where the probability is over V_{Reduce} 's remaining coins.

The protocol has the following complexities.

- $\ell_{Reduce} = \tilde{O}(\ell \cdot \text{depth}(C_{Dist}) \log \text{size}(C_{Dist}))$.
- $a_{Reduce} = b_{Reduce} = \tilde{O}(da + \text{poly}(d))$.
- $\text{Ptime}_{Reduce} = \text{poly}(\log |\mathbb{F}|, k, a, \ell, d, \text{size}(C_{Dist}), |\langle S_{Dist} \rangle|, |\langle \Phi_{Dist} \rangle|, n)$.
- $\text{Vtime}_{Reduce} = \tilde{O}(dal \cdot (\text{depth}(C_{Dist}) \log \text{size}(C_{Dist}) + |\langle \Phi_{Dist} \rangle|) + \text{poly}(d) + |\langle S_{Dist} \rangle| + k \cdot n)$.

The bit lengths of the descriptions are $|\langle S' \rangle| = \tilde{O}(|\langle S_{Dist} \rangle| + \text{poly}(d))$, $|\langle \Phi' \rangle| = \tilde{O}(dal + \text{poly}(d))$. Furthermore, suppose:

- $G_{Dist}(i, \langle S_{Dist} \rangle)$ is the circuit that returns the i -th element in S_{Dist} , and
- $C_{Dist}(x, \langle \Phi_{Dist} \rangle)$ is the circuit that returns $\Phi_{Dist}(x)$.

The new circuits $G'(i, \langle S' \rangle)$ and $C'(x, \langle \Phi' \rangle)$ satisfy the following.

- $\text{size}(G') = \text{size}(G_{Dist}) + \tilde{O}(\text{poly}(d))$.
- $\text{depth}(G') = \text{depth}(G_{Dist}) + \tilde{O}(1)$.
- $\text{size}(C') = \tilde{O}(\frac{M}{d} \cdot a\ell)$.
- $\text{depth}(C') = \tilde{O}(1)$.

And the verifier's verdict circuit (which outputs 0 iff it rejects amidst the protocol) satisfies:

- $\text{size}(V_{Reduce}) = \tilde{O}(dal \cdot (\text{depth}(C_{Dist}) \log \text{size}(C_{Dist}) + |\langle \Phi_{Dist} \rangle|) + \text{poly}(d) + |\langle S_{Dist} \rangle| + k \cdot n)$.
- $\text{depth}(V_{Reduce}) = \tilde{O}(1)$.

The number of instances is concretely $|\mathcal{S}'| \leq \left\lceil 8\sigma \frac{|\mathcal{S}_{Dist}|}{d} \right\rceil$.

4.1.3 The Explicit UIP for the associated language

We give the full specification of the UIP $(P_{\mathcal{L}'_x}, V_{\mathcal{L}'_x})$ that explicitly checks $(\mathcal{S}, \Phi) \in \mathcal{L}'_x$ in Section 4.5.2. This protocol has a cost proportional to $|\mathcal{S}|$. Therefore, it is only run at the end of the Batch-UIP protocol, when $|\mathcal{S}|$ is small.

Lemma 5 (The UIP for $\mathcal{L}'_x$). *Let $\mathcal{L}$ be a language with an ϵ -unambiguous $(\ell, a, b, \text{Ptime}, \text{Vtime})$ -UIP (P, V) , where the verifier verdict V is a log-space uniform boolean circuit of size S and depth D .*

There exists an ϵ -unambiguous $(\ell', a', b', \text{Ptime}', \text{Vtime}')$ -UIP $(P_{\mathcal{L}'_x}, V_{\mathcal{L}'_x})$ for verifying $(\mathcal{S}, \Phi) \in \mathcal{L}'_x$. Both parties have input x and $\langle \mathcal{S} \rangle, \langle \Phi \rangle$, where $|\mathcal{S}| = g$, and in the end $V_{\mathcal{L}'_x}$ either accepts or rejects.

Let $G(i, \langle \mathcal{S} \rangle)$ be the circuit that returns the i -th element in $\mathcal{S}$, and $C(x, \langle \Phi \rangle)$ be the circuit that returns $\Phi(x)$. The protocol has the following complexities.

- $\ell_{\mathcal{L}'_x} = \ell + 1$.
- $a_{\mathcal{L}'_x} = g\ell \cdot a$.
- $b_{\mathcal{L}'_x} = b$.
- $\text{Ptime}_{\mathcal{L}'_x} = O(g\ell \cdot \text{Ptime})$.
- $\text{Vtime}_{\mathcal{L}'_x} = O(g\ell \cdot \text{Vtime} + \text{size}(C) + g\text{size}(G))$.

Let $\tilde{O}$ hide $\text{polylog}(g, \ell)$ factors. The verifier's decision circuit $V_{\mathcal{L}'_x}$ has the following complexities:

- $\text{size}(V_{\mathcal{L}'_x}) = \tilde{O}(g\ell \cdot S + \text{size}(C) + g\text{size}(G))$.
- $\text{depth}(V_{\mathcal{L}'_x}) = \max(\text{depth}(G) + D, \text{depth}(C)) + \tilde{O}(1)$.

4.2 Our Construction of Batch-UIP

We show how to prove Theorem 6 assuming Lemmas 3 to 5.

Proof of Theorem 6. Denote the sub-protocols as $(P_{\text{Dist}}, V_{\text{Dist}})$, $(P_{\text{Reduce}}, V_{\text{Reduce}})$, and $(P_{\mathcal{L}'_x}, V_{\mathcal{L}'_x})$. The full batching protocol (P', V') is given in Protocol 1.

Prescribed Completeness. Note that $(\mathcal{S}^{(0)}, \Phi^{(0)}) \in \mathcal{L}'_x$ iff $\Phi^{(0)}(\mathbf{a}^{\mathcal{S}^{(0)}}) = 1$ and $\mathbf{x} \in \mathcal{L}^{\otimes k}$. The former condition is true iff $V(x_i, \mathbf{q}, \mathbf{a}^{x_i, \mathbf{q}}) = 1$ for every $(i, \mathbf{q}) \in \mathcal{S}^{(0)}$, which is true iff $\mathbf{x} \in \mathcal{L}^{\otimes k}$ because of the prescribed completeness of (P, V) . By invoking the prescribed completeness of $(P_{\text{Dist}}, V_{\text{Dist}})$ and $(P_{\text{Reduce}}, V_{\text{Reduce}})$ and $(P_{\mathcal{L}'_x}, V_{\mathcal{L}'_x})$, we get that V' accepts iff $\mathbf{x} \in \mathcal{L}^{\otimes k}$.

Unambiguity. Let P^* be a prover strategy that deviates from P' in some round, then it must be deviating in one of the ρ applications of $(P_{\text{Dist}}, V_{\text{Dist}})$, one of the ρ applications of $(P_{\text{Reduce}}, V_{\text{Reduce}})$, or in the final $(P_{\mathcal{L}'_x}, V_{\mathcal{L}'_x})$. The probability that V' accepts is upper bounded by the probability that unambiguity breaks in any of the subsequent rounds, which, by a union bound, is at most $\epsilon' = \rho \cdot (\epsilon + 2^{-\sigma} + 2^{-\sigma}) + \epsilon < (\rho + 1)\epsilon + 2 \cdot \rho 2^{-\rho} < 2 \log k \cdot (\epsilon + 2^{-\sigma})$.

Protocol 1 Our Batch-UIP (P', V').

Input Parameters: $\sigma \in \mathbb{N}$.

Input: $k \in \mathbb{N}$, and $\mathbf{x} = (x_1, \dots, x_k)$ are statements.

Other Parameters: $d = \lceil 16 \cdot \log(k\ell) \cdot \sigma \rceil \cdot \ell$ is the distance parameter, $\mathbb{F}$ is a characteristic-2 finite field of size $|\mathbb{F}| \geq 32 \cdot 2^\sigma \cdot (16 \log(k) \cdot \sigma \ell m)^{C_0}$, (where C_0 is some universal constant), and $\rho = \lfloor \log k \rfloor$ is the number of iterations.

Ingredients:

- The *Distance Generation Protocol*, $(P_{\text{Dist}}, V_{\text{Dist}})$ (Lemma 3).
It takes in a claim $\langle \mathcal{S} \rangle, \langle \Phi \rangle$ over $|\mathcal{S}|$ active pairs and output an intermediate claim $\langle \mathcal{S}_{\text{Dist}} \rangle, \langle \Phi_{\text{Dist}} \rangle$ over $|\mathcal{S}_{\text{Dist}}| = M = g\ell$ active pairs.
 - The *Instance Reduction Protocol*, $(P_{\text{Reduce}}, V_{\text{Reduce}})$ (Lemma 4).
It takes in the claim $\langle \mathcal{S}_{\text{Dist}} \rangle, \langle \Phi_{\text{Dist}} \rangle$ over M active pairs and output a reduced claim $\langle \mathcal{S}' \rangle, \langle \Phi' \rangle$ over $|\mathcal{S}'| \leq M/d$ active pairs.
 - The UIP (P', V') for explicitly checking the final claim $(\mathcal{S}^{(\rho)}, \Phi^{(\rho)}) \in \mathcal{L}'_{\mathbf{x}}$ (Lemma 5).
- 1: V' samples a random $\mathbf{q} \leftarrow \{0, 1\}^{b\ell}$, and send it entirely to P' .
 - 2: Both parties let $\mathcal{S}^{(0)} \leftarrow \{(i, \mathbf{q})\}_{i \in [k]}$, and $\Phi^{(0)}$ be the predicate that checks every row $\mathbf{a}^{x_i, \mathbf{q}}$ in $\mathcal{A}^{\mathcal{S}^{(0)}}$ are accepted, i.e. $V(x_i, \mathbf{q}, \mathbf{a}^{x_i, \mathbf{q}}) = 1$ for every $i \in [k]$.
 - 3: Define $\langle \mathcal{S}^{(0)} \rangle = \langle \mathbf{q} \rangle$, and $G^{(0)}(\langle \mathcal{S}^{(0)} \rangle)$ which enumerates $\mathcal{S}^{(0)}$ is the circuit that outputs $(i, \mathbf{q})$ for every $i \in [k]$.
 - 4: Define $\langle \Phi^{(0)} \rangle = \langle \mathbf{x}, \mathbf{q} \rangle$, and $C^{(0)}(x, \langle \Phi^{(0)} \rangle)$ be the circuit that checks $\Phi^{(0)}(x) = 1$.
 - 5: **for** $j = 1, \dots, \rho$ **do**
 - 6: Run $(P_{\text{Dist}}, V_{\text{Dist}})$ on $\langle \mathcal{S}^{(j-1)} \rangle, \langle \Phi^{(j-1)} \rangle$ with parameters σ, d , and $\mathbb{F}$, obtaining $\langle \mathcal{S}_{\text{Dist}}^{(j-1)} \rangle, \langle \Phi_{\text{Dist}}^{(j-1)} \rangle$.
 - 7: Run $(P_{\text{Reduce}}, V_{\text{Reduce}})$ on $\langle \mathcal{S}_{\text{Dist}}^{(j-1)} \rangle$ and $\langle \Phi_{\text{Dist}}^{(j-1)} \rangle$ with parameters σ, d , and $\mathbb{F}$, obtaining $\langle \mathcal{S}^{(j)} \rangle, \langle \Phi^{(j)} \rangle$.
 - 8: Run the UIP $(P_{\mathcal{L}'_{\mathbf{x}}}, V_{\mathcal{L}'_{\mathbf{x}}})$ explicitly on $\mathbf{x}$ and $\langle \mathcal{S}^{(\rho)} \rangle, \langle \Phi^{(\rho)} \rangle$ to check $(\mathcal{S}^{(\rho)}, \Phi^{(\rho)}) \in \mathcal{L}'_{\mathbf{x}}$.
 - 9: V' accepts iff $V_{\mathcal{L}'_{\mathbf{x}}}$ accepts.
-

Number of Intermediate Instances We show by induction that $|\mathcal{S}^{(j)}| \leq \lceil k \cdot 2^{-j} \rceil$, for $j \in [\rho]$. For the base case, $|\mathcal{S}^{(0)}| = k$ by definition. Assuming the inductive hypothesis, by the guarantees of $(P_{\text{Dist}}, V_{\text{Dist}})$, $|\mathcal{S}_{\text{Dist}}^{(j)}| = |\mathcal{S}^{(j)}| \cdot \ell \leq \lceil k \cdot 2^{-j} \rceil \cdot \ell$. We conclude by the guarantees of $(P_{\text{Reduce}}, V_{\text{Reduce}})$, $|\mathcal{S}^{(j+1)}| \leq \lceil 8\sigma |\mathcal{S}_{\text{Dist}}^{(j)}| / d \rceil \leq \frac{\lceil 8\sigma \cdot k \cdot 2^{-j} \cdot \ell \rceil}{\lceil 16 \cdot \log(k\ell) \cdot \sigma \rceil \cdot \ell} < \lceil k \cdot 2^{-j-1} \rceil$.

Description Lengths We first bound the bit lengths of all descriptions. With $\tilde{O}$ hiding $\text{polylog}(|\mathbb{F}|, k, n, a, \ell)$ factors, for all $1 \leq j \leq \rho < \log k = \tilde{O}(1)$, and

- $|\langle \mathcal{S}^{(j)} \rangle| = \tilde{O}(|\langle \mathcal{S}_{\text{Dist}}^{(j-1)} \rangle| + \text{poly}(d))$
 $= \tilde{O}(|\langle \mathcal{S}^{(j-1)} \rangle| + a\ell + \text{poly}(d))$
 $= \dots = \tilde{O}(j \cdot (a\ell + \text{poly}(d))) = \tilde{O}(a\ell + \text{poly}(d))$, because $j = \tilde{O}(1)$.
- $|\langle \Phi^{(j)} \rangle| = \tilde{O}(a\ell + \text{poly}(d))$.
- $|\langle \Phi_{\text{Dist}}^{(j)} \rangle| = \tilde{O}(|\langle \Phi^{(j)} \rangle| + |\langle \mathcal{S}^{(j)} \rangle| + da\ell + k \cdot n) = \tilde{O}(da\ell + \text{poly}(d) + k \cdot n)$.

Note that $|\langle \Phi^{(0)} \rangle| = \tilde{O}(k \cdot n + a\ell)$, because $\langle \Phi^{(0)} \rangle = (\mathbf{x}, \mathbf{q})$.

Intermediate Circuit Sizes and Depths For all $1 \leq j \leq \rho \leq \log k = \tilde{O}(1)$, let $G^{(j)}(s, \langle \mathcal{S}^{(j)} \rangle)$ be the circuit that returns the s -th element in $\mathcal{S}^{(j)}$, and let $C^{(j)}(x, \langle \Phi^{(j)} \rangle)$ be the circuit that returns $\Phi^{(j)}(x)$. For the base case, we already have $\text{size}(G^{(0)}) = \tilde{O}(k + a\ell)$ and $\text{depth}(G^{(0)}) = O(1)$, and $\text{size}(C^{(0)}) = \tilde{O}(k \cdot S)$ and $\text{depth}(C^{(0)}) = D + \tilde{O}(1)$. By the guarantees of Lemmas 3 and 4, we calculate the following:

- $\text{size}(G^{(j)}) = \text{size}(G_{\text{Dist}}^{(j-1)}) + \tilde{O}(\text{poly}(d)) = \text{size}(G^{(j-1)}) + \tilde{O}(a\ell + \text{poly}(d)) = \dots = \tilde{O}(k + a\ell + \text{poly}(d))$.
- $\text{depth}(G^{(j)}) = \text{depth}(G_{\text{Dist}}^{(j-1)}) + \tilde{O}(1) = \dots = \tilde{O}(1)$.
- $\text{size}(C^{(j)}) = \tilde{O}(|\mathcal{S}^{(j)}| \cdot a\ell)$.
- $\text{depth}(C^{(j)}) = \tilde{O}(1)$.
- $\text{size}(C_{\text{Dist}}^{(j)}) = \tilde{O}(|\mathcal{S}^{(j)}| \cdot \ell \cdot \text{size}(G^{(j)}) + \text{size}(C^{(j)}) + |\mathcal{S}^{(j)}| \cdot S + gda\ell^2)$
 $= \tilde{O}(|\mathcal{S}^{(j)}| \cdot \ell \cdot (k + a\ell + \text{poly}(d)) + |\mathcal{S}^{(j)}| \cdot (a\ell + S) + gda\ell^2)$
 $= \tilde{O}(k^2\ell + ka\ell^2 + ka\ell + k \cdot \ell \cdot \text{poly}(d) + kS + gda\ell^2)$.
- $\text{depth}(C_{\text{Dist}}^{(j)}) = \max(\text{depth}(C^{(j-1)}), \text{depth}(G^{(j-1)}) + D) + \tilde{O}(1) = \dots = D + \tilde{O}(1)$.

Therefore, $\text{depth}(C_{\text{Dist}}^{(j)}) \log(\text{size}(C_{\text{Dist}}^{(j)})) = (D + \tilde{O}(1))(\log S + \tilde{O}(1)) = D \log S + \tilde{O}(1)$.

The technical conditions stated in Lemma 4 are satisfied by the choice of parameters:

1. Since we choose $|\mathbb{F}|$ to be at least $32 \cdot 2^\sigma \cdot (16 \log(k) \cdot \sigma a\ell m)^{C_0}$, it satisfies the field size requirement of $|\mathbb{F}| \geq 32 \cdot 2^\sigma (da\ell m)^C$ as long as $C_0 \geq 2C$, where C is the constant in Lemma 4.

2. Fix $j \in [\rho]$, by the above analysis, $|\mathcal{S}_{\text{Dist}}^{(j-1)}| = |\mathcal{S}^{(j-1)}| \cdot \ell \leq \lceil k \cdot 2^{1-j} \cdot \ell \rceil$. To apply $(\mathbf{P}_{\text{Reduce}}, \mathbf{V}_{\text{Reduce}})$, the condition $d \geq 16m\sigma$ must hold, where $m := \log |\mathcal{S}_{\text{Dist}}^{(j-1)}| \leq \log(k\ell)$. Indeed, $d = \lceil 16 \cdot \log(k\ell) \cdot \sigma \rceil \cdot \ell > 16m\sigma$.

We rely on $d = \lceil 16 \cdot \log(k\ell) \cdot \sigma \rceil \cdot \ell = \tilde{O}(\ell)$ and $\rho \leq \log k = \tilde{O}(1)$ to simplify the complexities as follows.

The overall cost equals the cost of running the two sub-protocols **Dist** and **Reduce** for ρ times, plus the cost of running the final UIP $(\mathbf{P}_{\mathcal{L}'_x}, \mathbf{V}_{\mathcal{L}'_x})$ once. Note that the size of the final $|\mathcal{S}^{(\rho)}|$ is at most $k \cdot 2^{-\rho} = \tilde{O}(1)$. We use notations such as $\mathbf{Ptime}_{\text{Dist},j}$ to denote the time complexity of running **Dist** in the j -th round (and similarly for **Reduce** and $\mathcal{L}'_x$).

- $\ell' = \sum_{j=1}^{\rho} (\ell_{\text{Dist},j} + \ell_{\text{Reduce},j}) + \ell_{\mathcal{L}'_x} = \tilde{O}(\rho \cdot \ell D \log S) = \tilde{O}(\ell D \log S)$.
- $a' = b' = \max(a_{\text{Dist}}, a_{\text{Reduce}}, a_{\mathcal{L}'_x}) = \tilde{O}(\max(da, da + \text{poly}(d), |\mathcal{S}^{(\rho)}| \ell \cdot a)) = \tilde{O}(a\ell + \text{poly}(\ell))$.
- $\mathbf{Ptime}' = \tilde{O}\left(\sum_{j=1}^{\rho} ((k \cdot 2^{1-j}) \cdot \ell \cdot \mathbf{Ptime} + d(k \cdot 2^{1-j})a\ell^2 + |\langle \Phi^{(j)} \rangle| + |\langle \mathcal{S}_{\text{Dist}} \rangle| + k \cdot n))\right. \\ \left. + \tilde{O}\left(\sum_{j=1}^{\rho} \text{poly}(d, k \cdot 2^{1-j}, a, \ell, \text{size}(C_{\text{Dist}}), |\langle \Phi^{(j)} \rangle|, |\langle \mathcal{S}_{\text{Dist}} \rangle|), n\right) \right. \\ \left. + \tilde{O}(\ell \cdot \mathbf{Ptime})\right) \\ = \tilde{O}(k \cdot \ell \cdot \mathbf{Ptime} + \text{poly}(k, a, \ell, S)).$
- $\mathbf{Vtime}' = \sum_{j=1}^{\rho} \mathbf{Vtime}_{\text{Dist},j} + \sum_{j=1}^{\rho} \mathbf{Vtime}_{\text{Reduce},j} + \mathbf{Vtime}_{\mathcal{L}'_x} \\ = \tilde{O}(\rho \cdot dal + |\langle \mathcal{S}^{(j)} \rangle| + |\langle \Phi^{(j)} \rangle| + k \cdot n) \\ + \tilde{O}\left(\sum_{j=0}^{\rho-1} (dal(\text{depth}(C_{\text{Dist}}) \log \text{size}(C_{\text{Dist}}) + |\langle \Phi_{\text{Dist}}^{(j)} \rangle|) + \text{poly}(d) + |\langle \mathcal{S}_{\text{Dist}}^{(j)} \rangle| + k \cdot n)\right) \\ + \tilde{O}(|\mathcal{S}^{(\rho)}| \ell \cdot \mathbf{Vtime} + \text{size}(\Phi^{(\rho)}) + |\mathcal{S}^{(\rho)}| \text{size}(G^{(\rho)})) \\ = \tilde{O}(\rho dal + k \cdot n + \rho dal(D \log S + dal + \text{poly}(d) + k \cdot n) + \ell \mathbf{Vtime} + k + a\ell) \\ = \tilde{O}(\ell \cdot \mathbf{Vtime} + a\ell^2(kn + D \log S) + a^2 \cdot \text{poly}(\ell)).$

Finally, the verifier's verdict circuit $\mathbf{V}'$ rejects iff any of the sub-protocols rejects. (Note that $\mathbf{V}_{\text{Dist}}$ never rejects, so there are no terms corresponding to these sub-protocols.)

- $\text{size}(\mathbf{V}') \leq \sum_{j=1}^{\rho} \text{size}(\mathbf{V}_{\text{Reduce},j}) + \text{size}(\mathbf{V}_{\mathcal{L}'_x}) \\ = \tilde{O}((\ell \cdot S + a\ell^2(kn + D \log S)) + a^2 \cdot \text{poly}(\ell)).$

- $\text{depth}(V') = \max(\{\text{depth}(V_{\text{Reduce},j})\}_{j \in [\rho]}, \text{depth}(V_{\mathcal{L}'_x})) + O(\log \rho)$
 $= \max(\{\tilde{O}(1)\}_{j \in [\rho]}, \max(\text{depth}(G^{(\rho)}) + D, \text{depth}(C^{(\rho)}))) + O(\log \rho)$
 $= D + \tilde{O}(1).$

And moreover, V' is log-space uniform because V' is simply the AND's of the ρ V_{Reduce} and the final $V_{\mathcal{L}'_x}$, and all of these are themselves log-space uniform. $\square$

4.3 Proof of Lemma 3: Generating Δ_c -Distance via Checksums.

4.3.1 Additional Ingredients

A helper UIP for verifying $\mathbf{a} = \mathbf{a}^{x,q}$ Recall that $\mathbf{a}^{x,q}$ is the prescribed transcript of (P, V) on the statement x when verifier uses random coin $\mathbf{q}$ (Definition 12). There exists a UIP, which we call a *Transcript Checker* protocol, that allows a verifier to certify that some transcript $\mathbf{a} \in \{0, 1\}^{a\ell}$ indeed equals $\mathbf{a}^{x,q}$.

The *Random Continuation Protocol* is such a Transcript Checker. In this protocol, the verifier samples a *fresh* random coin sequence $\mathbf{q}' = (q'_1, \dots, q'_\ell) \in (\{0, 1\}^b)^\ell$, and challenges the prover to succeed in answering ℓ hybrid runs of (P, V) on the input x_i , where the j -th hybrid run uses the hybrid coin sequence $\mathbf{q}_j^{\text{Hyb}} := (\mathbf{q}_{\leq j}, \mathbf{q}'_{>j})$. The proof of Proposition 2 is straightforward and deferred to Section 4.5.1.

Proposition 2 (Transaction Checker; c.f. *Random Continuation* in Lemma 6.3 of [RRR16]). *Let $\mathcal{L}$ be a language with an ϵ -unambiguous $(\ell, a, b, \text{Ptime}, \text{Vtime})$ -UIP (P, V) . There exists an ϵ -unambiguous $(\ell, \ell \cdot a, b, \ell \cdot \text{Ptime}, \ell \cdot \text{Vtime})$ -UIP $(P^{\text{tc}}, V^{\text{tc}})$, where both parties take $((x, \mathbf{q}), \mathbf{a})$ as input, and the verifier decides whether to accept or reject. We call $(P^{\text{tc}}, V^{\text{tc}})$ the Transcript Checker, and it satisfies the following properties:*

- As a UIP for checking $((x, \mathbf{q}), \mathbf{a}) \in \mathcal{R}$, it satisfies prescribed completeness and ϵ -unambiguity.
- **Transcript Structure.** *The messages of P^{tc} can be exactly broken down into ℓ prescribed messages of the base UIP (P, V) . Namely, the r -th round message of P^{tc} on input $((x, \mathbf{q}), \mathbf{a})$ is $\mathbf{m}_r = \left((\mathbf{a}_r^{x, \mathbf{q}_j^{\text{Hyb}}})_{j \in [\ell]} \right)^\top$, where $\mathbf{a}_r^{x, \mathbf{q}_j^{\text{Hyb}}} = (P(x, (\mathbf{q}_j^{\text{Hyb}})_{<r}))$ is the r -th round message of P . Therefore, the entire transcript of P^{tc} equals $(\mathbf{m}_1, \dots, \mathbf{m}_\ell) = \left((\mathbf{a}^{x, \mathbf{q}_j^{\text{Hyb}}})_{j \in [\ell]} \right)^\top$.*
- **Check Structure.** *The verifier V^{tc} on a given transcript $\mathbf{a} = \left((\mathbf{a}^{x, \mathbf{q}_j^{\text{Hyb}}})_{j \in [\ell]} \right)^\top$ outputs 1 iff (1) for every $j \in [\ell]$, the corresponding $V(x, \mathbf{q}_j^{\text{Hyb}}, \mathbf{a}^{x, \mathbf{q}_j^{\text{Hyb}}}) = 1$, and (2) for each $j \in [\ell]$, the $(j \cdot a)$ -th prefix of $\mathbf{a}^{x, \mathbf{q}_j^{\text{Hyb}}}$ is the same as the $(j \cdot a)$ -th prefix of $\mathbf{a}^{x,q}$.*

Remark 10. When a batch of inputs $\mathbf{x} = (x_1, \dots, x_k) \in (\{0, 1\}^n)^k$ is given, for any sequence of g pairs $\mathcal{S} = ((i^1, \mathbf{q}^1), \dots, (i^g, \mathbf{q}^g)) \subset [k] \times (\{0, 1\}^b)^\ell$, and g transcripts $\mathbf{a}^1, \dots, \mathbf{a}^g \in \{0, 1\}^{\ell \times (a\ell)}$, we consider a batch run of g instances of $(P^{\text{tc}}, V^{\text{tc}})$ on $((x_{i^1}, \mathbf{q}^1), \mathbf{a}^1), \dots, ((x_{i^g}, \mathbf{q}^g), \mathbf{a}^g)$ where V^{tc} uses a shared random coin $\mathbf{q}'$ for all the g instances. Let

$$\mathcal{S}^{\text{Hyb}} := ((i^1, (\mathbf{q}^1)^{\text{Hyb}}), \dots, (i^1, (\mathbf{q}^1)_\ell^{\text{Hyb}}), \dots, (i^g, (\mathbf{q}^g)^{\text{Hyb}}), \dots, (i^g, (\mathbf{q}^g)_\ell^{\text{Hyb}}))$$

be the set of $g\ell$ hybrid pairs constructed from every pair in $\mathcal{S}$ using the new $\mathbf{q}'$. By the transcript structure of $(\mathbf{P}^{\text{tc}}, \mathbf{V}^{\text{tc}})$, the r -th round messages of the g instances of $(\mathbf{P}^{\text{tc}}, \mathbf{V}^{\text{tc}})$ constitute exactly the rows of $\mathbf{a}_r^{\text{S}^{\text{Hyb}}} = \left((\mathbf{a}_r^{x_i, \mathbf{q}})_{(i, \mathbf{q}) \in \mathcal{S}^{\text{Hyb}}} \right)^\top$, and the full transcripts of the g instances is $\mathbf{a}^{\text{S}^{\text{Hyb}}}$.

Checksum Let $M = 2^m$ for some $m \in \mathbb{N}$, and $\mathbb{F}$ be a constructible field ensemble of characteristic 2, such that $\log |\mathbb{F}| = \tilde{O}(\sigma)$.

Let $d \in \mathbb{N}$ be some *distance parameter* and $T_{\text{cksum}} \in \mathbb{N}$ be some parameter to be specified. We utilize the checksum of a systematic error correcting code with message space $\mathbb{F}^M$ of distance $2d$, which we denoted by $\text{cksum}_{\mathbf{U}}(\cdot)$, where the parameter $\mathbf{U} = (\mathbf{U}_1, \dots, \mathbf{U}_{T_{\text{cksum}}}) \in (\mathbb{F}^m)^{T_{\text{cksum}}}$ is a sequence of points in $\mathbb{F}^m$.

Proposition 3 (LDE-instantiated checksum (c.f. Proposition 6.6 of [RRR16])). *Let $M = 2^m$ and $m \in \mathbb{N}$, $\mathbb{F}$ be a field of characteristic 2. Let $\sigma \in \mathbb{N}$ be an unambiguity parameter, and $T_{\text{cksum}} := 2d(m + \log |\mathbb{F}|) + \sigma$. There exists a function $\text{cksum}_{\mathbf{U}} : \mathbb{F}^M \rightarrow \mathbb{F}^{T_{\text{cksum}}}$, such that with probability $1 - 2^{-\sigma}$ over uniform $\mathbf{U} = (\mathbf{U}_1, \dots, \mathbf{U}_{T_{\text{cksum}}})$, $\text{cksum}_{\mathbf{U}}(\cdot)$ is the checksum of a systematic linear error-correcting code of distance $2d$ on the message space $\mathbb{F}^M$. In other words, for any $\mathbf{m} \in \mathbb{F}^M$, for any $\mathbf{m}', \mathbf{m}'' \in \mathbb{F}^M$ that are both d -close to $\mathbf{m}$ and $\mathbf{m}' \neq \mathbf{m}''$, $\text{cksum}_{\mathbf{U}}(\mathbf{m}') \neq \text{cksum}_{\mathbf{U}}(\mathbf{m}'')$.*

This allows unique decoding in the following sense: if we know $\mathbf{m}'$ is d -close to some (fixed) $\mathbf{m}$, then we can uniquely determine $\mathbf{m}'$ given $\text{cksum}_{\mathbf{U}}(\mathbf{m}')$.

Proof. Recall that $\hat{\mathbf{m}} : \mathbb{F}^m \rightarrow \mathbb{F}^{T_{\text{cksum}}}$ is the multilinear extension of $\mathbf{m}$ (Section 3.1). We can let $\text{cksum}_{\mathbf{U}}(\mathbf{m}) := \hat{\mathbf{m}}(\mathbf{U}) = (\hat{\mathbf{m}}(\mathbf{U}_1), \dots, \hat{\mathbf{m}}(\mathbf{U}_{T_{\text{cksum}}})) \in \mathbb{F}^{T_{\text{cksum}}}$.

To show unique decoding, assume that on the contrary, $\text{cksum}_{\mathbf{U}}(\mathbf{m}') = \text{cksum}_{\mathbf{U}}(\mathbf{m}'')$. Then $\widehat{\mathbf{m}' - \mathbf{m}''}(\mathbf{U}) = 0$ (by linearity). Note $\mathbf{m}' - \mathbf{m}'' \in \Delta(\text{PVAL}(\mathbf{j}, \mathbf{0}))$ has Hamming weight at most $2d$, so $\Delta(\text{PVAL}(\mathbf{U}, \mathbf{0})) \leq 2d$, but this happens with probability at most $2^{-\sigma}$ by Proposition 1. $\square$

Abusing the notation, given a matrix $\mathbf{M} \in \mathbb{F}^{M \times a}$ with columns $(\mathbf{m}_1, \dots, \mathbf{m}_a)$, we use $\text{cksum}_{\mathbf{U}}(\mathbf{M})$ to denote $\mathbf{S} = (\text{cksum}_{\mathbf{U}}(\mathbf{m}_1), \dots, \text{cksum}_{\mathbf{U}}(\mathbf{m}_a)) \in \mathbb{F}^{T_{\text{cksum}} \times a}$.

4.3.2 Construction of the Dist Protocol

Road map of our Construction. Recall that the input claim is $(\mathcal{S}, \Phi) \in \mathcal{L}'_x$, and given the parameters $\sigma, d \in \mathbb{N}$ and the field $\mathbb{F}$ where $\log |\mathbb{F}| = \tilde{O}(\sigma)$, the goal is to define $(\mathcal{S}_{\text{Dist}}, \Phi_{\text{Dist}})$ such that Φ_{Dist} rejects everything in the $\mathcal{B}_{d, \mathbb{F}}(\mathbf{a}^{\text{S}^{\text{Dist}}})$ around $\mathbf{a}^{\text{S}^{\text{Dist}}}$ if $(\mathcal{S}, \Phi) \notin \mathcal{L}'_x$, except with probability at most $(\epsilon + 2^{-\sigma})$.

To this end, both parties interact for ℓ rounds, to *implicitly* run a batch of $(\mathbf{P}^{\text{tc}}, \mathbf{V}^{\text{tc}})$ in parallel. These protocols are run implicitly in the following sense.

- $\mathbf{V}_{\text{Dist}}$ samples $T_{\text{cksum}} = \tilde{O}(\sigma d)$ random points that defines the checksum function $\text{cksum}_{\mathbf{U}} : \mathbb{F}^{g\ell} \rightarrow \mathbb{F}^{T_{\text{cksum}}}$. The checksums enable unique decoding except with probability $2^{-\sigma}$ over the choice of $\mathbf{U}$ (Proposition 3).
- $\mathbf{V}_{\text{Dist}}$ samples a new random coin sequence $\mathbf{q}' \in (\{0, 1\}^b)^\ell$, and send the coin to the prover round by round,
- $\mathbf{P}_{\text{Dist}}$ only responds by the checksums of the underlying messages.

Specifically, in round $r \in [\ell]$, the prover computes checksums $\mathbf{S}_r \in \mathbb{F}^{T_{\text{cksum}}}$ on the columns of the matrix $\mathbf{a}_r^{\mathcal{S}^{\text{Hyb}}} \in \{0, 1\}^{(g\ell) \times a}$, which consists of the $(g\ell)$ base-UIP r -th round answers that the prover would have sent if the protocols $(\mathbf{P}^{\text{tc}}, \mathbf{V}^{\text{tc}})$ were run explicitly. Note that we would pick $T_{\text{cksum}} = \tilde{O}(\sigma d) \ll g\ell$, so indeed the checksums can be represented using very few bits ($\tilde{O}(\sigma d \log |\mathbb{F}|) = \tilde{O}(\sigma^2 d)$ bits).

Finally, they let the resulting claim be $(\mathcal{S}_{\text{Dist}}, \Phi_{\text{Dist}}) \in \mathcal{L}'_x$ where $\mathcal{S}_{\text{Dist}} = \mathcal{S}^{\text{Hyb}}$ and Φ_{Dist} is a circuit that takes in $\mathbf{a}$ (purported to be $\mathbf{a}^{\mathcal{S}^{\text{Hyb}}}$) as input. In addition to checking that rows (transcripts) in $\mathbf{a}$ are indeed accepted by the base UIP and satisfy the original constraint Φ , the circuit also recomputes the checksums of the columns of $\mathbf{a}$ to make sure they are equal to the checksums committed by the prover. Due to the unique decoding guarantee of cksum (Proposition 3), if there exists $\mathbf{a}^*$ in $\mathcal{B}_{d, \mathbb{F}}(\mathbf{a}^{\mathcal{S}_{\text{Dist}}})$ with $\Phi_{\text{Dist}}(\mathbf{a}^*) = 1$, the cheating prover would be implicitly committing to this $\mathbf{a}^*$. Therefore, one can actually extract (via unique decoding) an accepting but deviating prover strategy for one of the base UIP. The base UIP's unambiguity guarantees that any deviating strategy should only be accepted with probability at most ϵ . Therefore, with high probability, there is no $\mathbf{a}^* \in \mathcal{B}_{d, \mathbb{F}}(\mathbf{a}^{\mathcal{S}_{\text{Dist}}})$ that would make $\Phi_{\text{Dist}}(\mathbf{a}^*) = 1$, i.e. distance is generated.

The Protocol. The protocol is given in Protocol 2.

4.3.3 Analysis of the Dist Protocol

Recall that the input claim $(\mathcal{S}, \Phi) \in \mathcal{L}'_x$ is true iff $(\Phi(\mathbf{a}^{\mathcal{S}}) = 1) \wedge (\mathbf{x}|_{\mathcal{I}} \in \mathcal{L}^{\otimes |\mathcal{I}|})$ where $\mathcal{S}$ consists of $g \leq k$ pairs of the form $(i, \mathbf{q}) \in [k] \times \{0, 1\}^{b\ell}$, $\log |\mathbb{F}| = \tilde{O}(\sigma)$, and Φ is a log-space uniform boolean circuit of size $\tilde{O}(gal)$ and depth $\tilde{O}(1)$.

Prescribed Completeness. The honest prover $\mathbf{P}_{\text{Dist}}$ answers the random coin $\mathbf{q}'$ sampled by the verifier by running the Random Continuation UIP $(\mathbf{P}_{i, \mathbf{q}}, \mathbf{V}_{i, \mathbf{q}})$ in parallel on all $g = |\mathcal{S}|$ instances. It honestly sends the checksums $\mathbf{S}$ that are computed on the columns of $\mathbf{a}^{\mathcal{S}_{\text{Dist}}}$, where $\mathcal{S}^{\text{Hyb}}$ is the set of hybrid pairs between $\mathcal{S}$ and $\mathbf{q}'$. Recall that $\Phi_{\text{Dist}}(\mathbf{a}^{\mathcal{S}_{\text{Dist}}}) = 1$ iff

1. $\Phi(\mathbf{a}[\mathcal{S}, :]) = 1$ is true. Note that the submatrix $\mathbf{a}[\mathcal{S}, :]$ is indexed by the subset $\mathcal{S}$ on $\mathbf{a}$. (This is valid because $\mathcal{S} \subset \mathcal{S}^{\text{Hyb}}$.)
2. For all $(i, \mathbf{q}) \in \mathcal{S}^{\text{Hyb}}$, $\mathbf{V}(x_i, \mathbf{q}, \mathbf{a}'[(i, \mathbf{q}), :]) = 1$, and that for every $j \in [\ell]$, the $(j \cdot a)$ -th prefix of $\mathbf{a}[(i, \mathbf{q}_j^{\text{Hyb}}), :]$ and $\mathbf{a}[(i, \mathbf{q}), :]$ are identical.
3. $\text{cksum}_U(\mathbf{a}) = (\mathbf{S}_1, \dots, \mathbf{S}_\ell)$.
4. $\mathbf{a} \in \{0, 1\}^{M \times L}$ (i.e. all elements strictly lies in $\mathbb{GF}(2)$.)

The third condition is always true because the checksums are computed honestly on the columns of $\mathbf{a}^{\mathcal{S}_{\text{Dist}}}$. On the other hand, the first condition is true iff $\Phi(\mathbf{a}^{\mathcal{S}}) = 1$, while the second condition is true iff $\mathbf{x}|_{\mathcal{I}} \in \mathcal{L}^{\otimes \mathcal{I}}$, so, $\Phi_{\text{Dist}}(\mathbf{a}^{\mathcal{S}_{\text{Dist}}}) = 1$ iff $(\mathcal{S}, \Phi) \in \mathcal{L}'_x$.

Unambiguous Generation of Δ_c -Distance. The verifier $\mathbf{V}_{\text{Dist}}$ sends the fresh random coins $\mathbf{q}' = (q'_1, q'_2, \dots, q'_\ell)$ round by round. Let $\mathbf{P}^*$ be a cheating strategy that deviates from $\mathbf{P}_{\text{Dist}}$. For each $r \in [\ell]$, $\mathbf{P}^*$'s behavior in the r -th round is fully determined by $\mathbf{q}'_{\leq r} = (q'_1, \dots, q'_r)$, so let $\mathbf{S}_r^*$ be the checksum produced by $\mathbf{P}^*$ on $\mathbf{q}'_{\leq r}$. Let $\mathbf{S}^* = (\mathbf{S}_1^*, \dots, \mathbf{S}_\ell^*)$ be all the checksum sent.

Protocol 2 ($P_{\text{Dist}}, V_{\text{Dist}}$) Generating Δ_c -Distance via Checksums..

Input Parameters: $\sigma, d \in \mathbb{N}$ and a field $\mathbb{F}$.

Input: $\mathbf{x} = (x_1, \dots, x_k)$ are the UIP statements. $\langle \mathcal{S} \rangle, \langle \Phi \rangle$ are the descriptions for the claim $(\mathcal{S}, \Phi) \in \mathcal{L}'$.

Other Parameters: $g = |\mathcal{S}|$, $M = g\ell = 2^m$, and $T_{\text{cksum}} = 2d(m + \sigma) + \sigma \in \mathbb{N}$.

Ingredients:

- A Transcript Checker UIP , (P^{tc}, V^{tc}) (Proposition 2), for certifying $((x_i, \mathbf{q}), \mathbf{a}) \in \mathcal{R}$.
- A Checksum Function, $\text{cksum}_U(\cdot)$, computing the checksum of a systematic error-correcting code with distance $2d$, where $\mathbf{U} = (\mathbf{U}_1, \dots, \mathbf{U}_{T_{\text{cksum}}}) \in (\mathbb{F}^m)^{T_{\text{cksum}}}$ is a sequence of points in $\mathbb{F}^m$.

Output: $\mathbf{x}, \langle \mathcal{S}_{\text{Dist}} \rangle, \langle \Phi_{\text{Dist}} \rangle$.

- 1: V_{Dist} samples $\mathbf{q}' = (q'_1, \dots, q'_\ell) \in (\{0, 1\}^b)^\ell$.
 - 2: V_{Dist} samples $\mathbf{U} \leftarrow_R (\mathbb{F}^m)^{T_{\text{cksum}}}$ and sends it to P_{Dist} .
 - 3: **for** $r = 1, \dots, \ell$ **do**
 - 4: V_{Dist} sends q'_r to P_{Dist} . $\triangleright$ Random coins are shared.
 - 5: P_{Dist} receives q'_r and does:
 - 6: 1. Compute $\mathbf{m}_r^{i, \mathbf{q}} := P^{tc}((x_i, \mathbf{q}), (\mathbf{m}_1^{i, \mathbf{q}}, \dots, \mathbf{m}_{r-1}^{i, \mathbf{q}}))$ for all $(i, \mathbf{q}) \in \mathcal{S}$,
 - 7: 2. By Remark 10, $\{\mathbf{m}_r^{i, \mathbf{q}}\}_{(i, \mathbf{q}) \in \mathcal{S}}$ exactly forms the rows of the matrix $\mathbf{a}_r := \mathbf{a}_r^{\mathcal{S}^{\text{Hyb}}} \in \{0, 1\}^{(g\ell) \times a}$.
 - 8: 3. Compute $\mathbf{S}_r \leftarrow \text{cksum}_U(\mathbf{a}_r) \in \{0, 1\}^{T_{\text{cksum}} \times a}$, and send it to V_{Dist} .
 - 9: Define an arithmetic circuit $\Phi_{\text{Dist}}(\mathbf{a})$, where $\mathbf{a} \in \mathbb{F}^{M \times (a\ell)}$ is purported to be $\mathbf{a}^{\mathcal{S}^{\text{Hyb}}}$, and the circuit checks the following set of conditions. $\triangleright$ Note that $|\mathcal{S}^{\text{Hyb}}| = M = g\ell$.
 1. $\Phi(\mathbf{a}[\mathcal{S}, :]) = 1$ is true. Note that the submatrix $\mathbf{a}[\mathcal{S}, :]$ is indexed by the subset $\mathcal{S}$ on $\mathbf{a}$. (This is valid because $\mathcal{S} \subset \mathcal{S}^{\text{Hyb}}$.)
 2. For all $(i, \mathbf{q}) \in \mathcal{S}^{\text{Hyb}}$, $V(x_i, \mathbf{q}, \mathbf{a}'[(i, \mathbf{q}), :]) = 1$, and that for every $j \in [\ell]$, the $(j \cdot a)$ -th prefix of $\mathbf{a}[(i, \mathbf{q}_j^{\text{Hyb}}), :]$ and $\mathbf{a}[(i, \mathbf{q}), :]$ are identical.
 3. $\text{cksum}_U(\mathbf{a}) = (\mathbf{S}_1, \dots, \mathbf{S}_\ell)$.
 4. $\mathbf{a} \in \{0, 1\}^{M \times L}$ (i.e. all elements strictly lies in $\mathbb{GF}(2)$.)
 - 10: **return** $\langle \mathcal{S}^{\text{Hyb}} \rangle, \langle \Phi_{\text{Dist}} \rangle$.
-

Fix some $\mathbf{q}'_{\leq r^*}$ such that $r^* \in [\ell]$ is indeed the first round when P^* deviates, i.e. when it first produces some $\mathbf{S}_{r^*}^* \neq \text{cksum}(\mathbf{a}_{r^*})$. As r^* is minimal, $\mathbf{S}_{r^*}^* \neq \text{cksum}(\mathbf{a}_{r^*})$ but $\mathbf{S}_r^* = \text{cksum}(\mathbf{a}_r)$ for all $r < r^*$. (Recall that $\mathbf{a}_r$ is the r -th round chunks of $\mathbf{a}^{\text{SDist}}$). Note that this implies $\Phi_{\text{Dist}}(\mathbf{a}^{\text{SDist}}) = 0$ because $\mathbf{a}^{\text{SDist}}$ does not have the checksum $\mathbf{S}_{r^*}^*$ committed by the prover in round r^* .

Let E be the event that after $\mathbf{q}'_{\leq r^*}$ is sampled, there exists $\mathbf{a}^* \in \mathcal{B}_{d,\mathbb{F}}(\mathbf{a}^{\text{SDist}}) \setminus \{\mathbf{a}^{\text{SDist}}\}$, such that $\Phi_{\text{Dist}}(\mathbf{a}^*) = 1$. Given that we already showed $\Phi_{\text{Dist}}(\mathbf{a}^{\text{SDist}}) = 0$, to prove generation of Δ_c -distance (with probability at least $1 - (\epsilon + 2^{-\sigma})$, for every $\mathbf{a}^* \in \mathcal{B}_{d,\mathbb{F}}(\mathbf{a}^{\text{SDist}})$, $\Phi_{\text{Dist}}(\mathbf{a}^*) = 0$) we only need to show that $\Pr[E] \leq \epsilon + 2^{-\sigma}$. If E is impossible on this $\mathbf{q}'_{\leq r^*}$, then we are done, so we consider the case when $\Pr[E] > 0$.

Since $\Pr[E] > 0$, there exists some continuation $\mathbf{q}'_{> r^*}$ and $\mathbf{a}^* \in \mathcal{B}_{d,\mathbb{F}}(\mathbf{a}^{\text{SDist}}) \setminus \{\mathbf{a}^{\text{SDist}}\}$ such that $\Phi_{\text{Dist}}(\mathbf{a}^*) = 1$. In this case, for every $r \leq r^*$, if we let $\mathbf{a}_r^*$ be the corresponding chunk of $\mathbf{a}^*$ used in round r , we have $\text{cksum}_U(\mathbf{a}_r^*) = \mathbf{S}_r^*$ because of Condition 3 of Φ_{Dist} .

Furthermore, since P^* is first cheating on round r^* , we have $\mathbf{S}_{r^*}^* = \text{cksum}(\mathbf{a}_{r^*}^*) \neq \text{cksum}(\mathbf{a}_{r^*})$, and thus $\mathbf{a}_{r^*}^* \neq \mathbf{a}_{r^*}$. This means there exists a row-index $(i, \mathbf{q}^{\text{Hyb}}) \in \mathcal{S}^{\text{Hyb}}$ such that $\mathbf{a}_{r^*}^*[i, \mathbf{q}^{\text{Hyb}}] \neq \mathbf{a}_{r^*}[i, \mathbf{q}^{\text{Hyb}}]$. By Remark 10, $(i, \mathbf{q}^{\text{Hyb}})$ is some hybrid derived from some $(i, \mathbf{q}) \in \mathcal{S}$. Let $\mathcal{S}^{i, \text{Hyb}} = \{(i, \mathbf{q}_1^{\text{Hyb}}), \dots, (i, \mathbf{q}_\ell^{\text{Hyb}})\}$ be the set of hybrid pairs derived from $(i, \mathbf{q})$.

We construct a cheating prover P_*^{tc} , that breaks the unambiguous soundness of the UIP $(P^{\text{tc}}, V^{\text{tc}})$ as follows. For each round $r \in [\ell]$, it does:

1. Sample $\mathbf{U} \leftarrow_R \{0, 1\}^{T_{\text{cksum}}}$ uniformly at random and send it to P^* .
2. Let $\mathbf{q}'_{\leq r} = (q'_1, \dots, q'_r)$ be the random coins sent by V^{tc} in the previous rounds,
3. Apply P^* on $\mathbf{x}|_{\mathcal{I}}$ and $\mathbf{q}'_{\leq r}$ to obtain $\mathbf{S}_r^*$,
4. Try to infer $\mathbf{a}_r^*$, committed by P^* , by finding at most one $\mathbf{a}_r^*$ that is Δ_c -close to $\mathbf{a}_r$ with checksum $\mathbf{S}_r^*$.
5. If such an $\mathbf{a}_r^*$ exists and $\mathbf{a}_r^* \in \{0, 1\}^{M \times a}$, send the message $\mathbf{a}_r^*[\mathcal{S}^{i, \text{Hyb}}, :]$ to V^{tc} .

Claim 1. *Conditioned on unique decoding¹⁸ holding for the sampled $\mathbf{U}$, when V^{tc} samples the random coin prefix $\mathbf{q}'_{\leq r^*}$, P_*^{tc} first deviates on the r^* -th round and convinces V^{tc} to accept with probability at least $\Pr[E]$.*

Proof. First note that by the minimality of r^* and the uniqueness of $\mathbf{a}_r^*$ in each round, P_*^{tc} deviates from the prescribed P^{tc} the first time exactly in round r^* .

Whenever E occurs, there exists some completion $\mathbf{a}' \in \mathcal{B}_{d,\mathbb{F}}(\mathbf{a})$ such that $\Phi_{\text{Dist}}(\mathbf{a}') = 1$. Let $(\mathbf{a}'_1, \dots, \mathbf{a}'_L)$ be the columns of $\mathbf{a}'$. Again by unique decoding, P_*^{tc} are *exactly* sending those $\mathbf{a}'_r$ ($\mathbf{a}'_r = \mathbf{a}_r^*$) for all $r \in [\ell]$, because their checksums are agreeing with $\mathbf{S}^*$. Furthermore, $(\mathbf{a}'_1[\mathcal{S}^{i, \text{Hyb}}, :], \dots, \mathbf{a}'_\ell[\mathcal{S}^{i, \text{Hyb}}, :])$ satisfies Conditions 2 and 4 of Φ_{Dist} , so it is also accepted by V^{tc} (Proposition 2). $\square$

By the ϵ -unambiguity of $(P^{\text{tc}}, V^{\text{tc}})$, the probability of P_*^{tc} breaking the unambiguity of $(P^{\text{tc}}, V^{\text{tc}})$ is upper-bounded by ϵ . Therefore, $\Pr[E] - 2^{-\sigma} \leq \epsilon$. Note the loss of $2^{-\sigma}$ comes from the (rare) event that unique decoding does not hold for the initially sampled $\mathbf{U}$ (Proposition 3).

The *furthermore* part follows from that when E does not happen, $\mathbf{a}^{\text{SDist}}$ is the only candidate solution to $\Phi_{\text{Dist}}(\mathbf{a}) = 1$ in $\mathcal{B}_{d,\mathbb{F}}(\mathbf{a}^{\text{SDist}})$. This holds for the prescribed P_{Dist} as well.

¹⁸Recall that unique decoding means for any pair $\mathbf{a}', \mathbf{a}'' \in \mathcal{B}_{d,\mathbb{F}}(\mathbf{a}^{\text{SDist}})$, if $\text{cksum}_U(\mathbf{a}') = \text{cksum}_U(\mathbf{a}'')$ then $\mathbf{a}' = \mathbf{a}''$.

Complexities. With $\tilde{O}$ hiding $\text{poly}(\sigma) \cdot \text{polylog}(k, n, a, \ell)$ factors, Note that we rely on the assumption that $\log |\mathbb{F}| = \tilde{O}(\sigma)$ to simplify the expressions.

- $\ell_{\text{Dist}} = O(\ell)$,
- $a_{\text{Dist}} = T_{\text{cksum}} \cdot a \cdot \log |\mathbb{F}| = \tilde{O}(da)$.
- $b_{\text{Dist}} = \max(b, T_{\text{cksum}} \log |\mathbb{F}|) = \tilde{O}(\max(b, \sigma^2 d)) = \tilde{O}(da)$, assuming $b \leq a$.
- $\text{Ptime}_{\text{Dist}} = \tilde{O}(g\ell \text{Ptime} + T_{\text{cksum}}(g\ell)(a\ell) \log |\mathbb{F}| + |\langle \mathcal{S} \rangle| + |\langle \Phi \rangle| + k \cdot n) = \tilde{O}(g\ell \text{Ptime} + dga\ell^2 + |\langle \mathcal{S} \rangle| + |\langle \Phi \rangle| + k \cdot n)$.
- $\text{Vtime}_{\text{Dist}} = \tilde{O}(T_{\text{cksum}} \cdot a \cdot \log |\mathbb{F}| \cdot \ell + |\langle \mathcal{S} \rangle| + |\langle \Phi \rangle| + k \cdot n) = \tilde{O}(dal + |\langle \mathcal{S} \rangle| + |\langle \Phi \rangle| + k \cdot n)$.

The runtimes are less immediate, so we explain them below.

- For the prover runtime, note that it has to run $g = |\mathcal{S}|$ instances of the UIP ($\text{P}^{\text{tc}}, \text{V}^{\text{tc}}$) in parallel, which is in turn $g\ell$ instances of the base (P, V), so that takes $O(g\ell \text{Ptime})$. It also computes the checksums (which can be done in $\tilde{O}(T_{\text{cksum}}(g\ell)(a\ell) \log |\mathbb{F}|)$ time using FFT). The term $|\langle \mathcal{S} \rangle| + |\langle \Phi \rangle| + k \cdot n$ comes from reading the input.
- The verifier does not perform any check in this sub-protocol, but it has to store the $T_{\text{cksum}} \cdot a$ points it receives every round, which requires $\tilde{O}(T_{\text{cksum}} \cdot a \cdot \log |\mathbb{F}| \ell)$ time. In the end, the time to compute $\langle \mathcal{S}_{\text{Dist}} \rangle$ is $O(|\langle \mathcal{S}_{\text{Dist}} \rangle|) = O(|\langle \mathcal{S} \rangle| + b\ell)$, and the time to compute $\langle \Phi_{\text{Dist}} \rangle$ is $O(|\langle \Phi_{\text{Dist}} \rangle|) = \tilde{O}(|\langle \Phi \rangle| + |\langle \mathcal{S} \rangle| + dal + k \cdot n)$, since both of these consist of strings the verifier already has access to.

Circuit G_{Dist} for generating $\mathcal{S}_{\text{Dist}}$. Let $\mathbf{q}'$ be the fresh random coin generated in Protocol 2. Let $p := \lceil \log k \rceil + b \cdot \ell$ be the bit length of a pair $(i, \mathbf{q}) \in \{0, 1\}^{\lceil \log k \rceil} \times \{0, 1\}^{b \cdot \ell}$. We can represent $\mathcal{S}_{\text{Dist}}$ as a subset of $\{0, 1\}^p$.

Claim 2. *The set $\mathcal{S}_{\text{Dist}}$ has a description $\langle \mathcal{S}_{\text{Dist}} \rangle = (\langle \mathcal{S} \rangle, \mathbf{q}')$, and $|\langle \mathcal{S}_{\text{Dist}} \rangle| = |\langle \mathcal{S} \rangle| + b\ell + O(1)$.*

Proof. Assume $\mathcal{S} = \{\mathcal{S}_1, \dots, \mathcal{S}_g\}$, where $\mathcal{S}_s = (i^s, \mathbf{q}^s)$ for every $s \in [g]$, then by the definition of the hybrid coins in Protocol 2,

$$\mathcal{S}_{\text{Dist}} = \mathcal{S}^{\text{Hyb}} = \left\{ (i^1, (\mathbf{q}^1)_1^{\text{Hyb}}), \dots, (i^1, (\mathbf{q}^1)_\ell^{\text{Hyb}}), \dots, (i^g, (\mathbf{q}^g)_1^{\text{Hyb}}), \dots, (i^g, (\mathbf{q}^g)_\ell^{\text{Hyb}}) \right\},$$

where $(\mathbf{q}^s)_r^{\text{Hyb}} := (\mathbf{q}_{\leq r}^s, \mathbf{q}_{> r}^s)$ for every $s \in [g], r \in [\ell]$.

Let $M = |\mathcal{S}_{\text{Dist}}| = g\ell$, and $\langle \mathcal{S}_{\text{Dist}} \rangle = (\langle \mathcal{S} \rangle, \mathbf{q}')$. We define a Turing machine $\mathcal{M}$ that takes in $(M, 1^{|\langle \mathcal{S}_{\text{Dist}} \rangle|}, 1^p)$ as input, and outputs the circuit $G_{\text{Dist}} : [M] \times \{0, 1\}^{|\langle \mathcal{S}_{\text{Dist}} \rangle|} \rightarrow \{0, 1\}^p$. Since $\langle \mathcal{S} \rangle$ is a succinct description of $\mathcal{S}$, there exists a Turing machine $\mathcal{M}'$ that takes in $(g, 1^{|\langle \mathcal{S} \rangle|}, 1^p)$ as input, and outputs the circuit $G : [g] \times \{0, 1\}^{|\langle \mathcal{S} \rangle|} \rightarrow \{0, 1\}^p$, which, given $(s, \langle \mathcal{S} \rangle)$ as input, outputs $\mathcal{S}_s = (i^s, \mathbf{q}^s)$.

$\mathcal{M}$ first runs $\mathcal{M}'$ to get G , then moves onto defining G_{Dist} , which, given $(s', \langle \mathcal{S}_{\text{Dist}} \rangle) \in [M] \times \{0, 1\}^{|\langle \mathcal{S}_{\text{Dist}} \rangle|}$ as input,

1. First, computes (g, r) such that $s' = s \cdot g + r$ where $s \in [g], r \in [\ell]$, e.g. by long division.
2. Then, evaluates $G(s, \langle \mathcal{S} \rangle)$ to get $(i^s, \mathbf{q}^s)$.

3. Finally, computes the corresponding hybrid coin $(\mathbf{q}^s)_r^{\text{Hyb}}$ using $\mathbf{q}^s$ and $\mathbf{q}'$ (which is just copying the first r bits of $\mathbf{q}^s$ and appending the remaining bits of $\mathbf{q}'$), and outputs $(i^s, (\mathbf{q}^s)_r^{\text{Hyb}})$.

Indeed,

- $\text{size}(G_{\text{Dist}}) = \text{size}(G) + b\ell + \tilde{O}(1)$.
- $\text{depth}(G_{\text{Dist}}) = \text{depth}(G) + \tilde{O}(1)$.

The Turing machine $\mathcal{M}$ uses $O(p) = O(\log k + b\ell)$ space because it only needs to simulate $\mathcal{M}'$ (which uses $O(p)$ space) and maintain pointers into the circuit G_{Dist} (which uses $O(p) + O(\log k + \log b + \log \ell)$ space). $\square$

Circuit C_{Dist} for Φ_{Dist} . By construction, $\Phi_{\text{Dist}}(\mathbf{a})$ checks four conditions, whose circuit complexities are:

1. Checking Φ on $\mathbf{a}[\mathcal{S}, :]$, requires $\text{size}(\Phi)$ gates and $\text{depth}(\Phi)$ depth.
2. For all $(i, \mathbf{q}) \in \mathcal{S}^{\text{Hyb}}$, $\mathbf{V}(x_i, \mathbf{q}, \mathbf{a}'[(i, \mathbf{q}), :]) = 1$, and that for every $j \in [\ell]$, the $(j \cdot a)$ -th prefix of $\mathbf{a}[(i, \mathbf{q}_j^{\text{Hyb}}), :]$ and $\mathbf{a}[(i, \mathbf{q}), :]$ are identical.

This requires $O(M \cdot S + ML) = \tilde{O}(g\ell S + dgal^2)$ gates and $D + O(\log(ML)) = D + \tilde{O}(1)$ depth (by a binary tree of AND gates).

3. $\text{cksum}_U(\mathbf{a}) = (\mathbf{S}_1, \dots, \mathbf{S}_\ell)$.

This requires re-computing the checksums for every column of $\mathbf{a} \in \{0, 1\}^{M \times L}$. To compute one checksum, T_{cksum} LDE evaluation are needed. One LDE evaluation requires $\tilde{O}(M) = \tilde{O}(g\ell)$ gates, and $\tilde{O}(1)$ depth. Therefore, we need $\tilde{O}(T_{\text{cksum}}ML) = \tilde{O}(dgal^2)$ gates and $\tilde{O}(1)$ depth to compute the checksums. Finally, we implement the consistency checks by equality tests. This requires an additional $\tilde{O}(T_{\text{cksum}}L) = \tilde{O}(dal)$ gates and $\tilde{O}(1)$ depth.

4. $\mathbf{a} \in \{0, 1\}^{M \times L}$ (i.e. all elements strictly lies in $\mathbb{GF}(2)$.)

To check $\mathbf{a} \in \mathbb{GF}(2)^{M \times L}$, we just need to test if every element is in $\mathbb{GF}(2)$, and compute an AND over them. This requires $\tilde{O}(ML) = \tilde{O}(g\ell)$ gates and $\tilde{O}(1)$ depth, because the test can be implemented by evaluating the polynomial $X^2 + X$ over $\mathbb{F}$ and checking if the result is 0.

In order to perform the second check, C_{Dist} also needs to first enumerate the pairs $\mathcal{S}^{\text{Hyb}}$ explicitly using $\langle \mathcal{S}^{\text{Hyb}} \rangle$. This incurs an additional overhead of $\text{depth}(G_{\text{Dist}}) + \tilde{O}(1) = \text{depth}(G) + \tilde{O}(1)$ and $g\ell \cdot \text{size}(G_{\text{Dist}}) = \tilde{O}(g\ell \cdot \text{size}(G) + gbl^2)$ (see Claim 2).

Therefore, we have the following overall bounds.

- $\text{size}(C_{\text{Dist}}) = \tilde{O}(g\ell \cdot \text{size}(G) + \text{size}(C) + galS + dgal^2)$.
- $\text{depth}(C_{\text{Dist}}) = \max(\text{depth}(C), \text{depth}(G) + D + \tilde{O}(1), \tilde{O}(1)) + \tilde{O}(1)$
 $= \max(\text{depth}(C), \text{depth}(G) + D) + \tilde{O}(1)$.

Similarly to Claim 2, the circuit Φ_{Dist} also has a succinct description, given by

$$\langle \Phi_{\text{Dist}} \rangle = (\langle \mathcal{S}_{\text{Dist}} \rangle, \langle \Phi \rangle, \mathbf{U}, (\mathbf{S}_1, \dots, \mathbf{S}_\ell), \mathbf{x}),$$

and it is indeed log-space uniform, because its components (the four checks, as well as the circuit $C(\langle \mathcal{S}^{\text{Hyb}} \rangle)$ that generates $\mathcal{S}^{\text{Hyb}}$) are, and the additional binary trees of AND as well as the equality checks are simple to describe. The bit length $|\langle \Phi_{\text{Dist}} \rangle| = \tilde{O}(|\langle \Phi \rangle| + |\langle \mathcal{S} \rangle| + dal + k \cdot n)$.

4.4 Proof of Lemma 4: Instance Reduction for Δ_c -Distance.

4.4.1 Additional Ingredients

Road map of our Construction. Recall that the intermediate claim is $\Phi_{\text{Dist}}(\mathbf{a}^{S_{\text{Dist}}}) = 1$. We also have the *distance guarantee* from Lemma 3: If $\Phi_{\text{Dist}}(\mathbf{a}^{S_{\text{Dist}}}) = 0$ then every matrix $\mathbf{a} \in \mathcal{B}_{d,\mathbb{F}}(\mathbf{a}^{S_{\text{Dist}}})$ is rejected by Φ_{Dist} . Our goal is to output some reduced claim $(\mathcal{S}', \Phi')$ such that $|\mathcal{S}'| \leq |\mathcal{S}_{\text{Dist}}|/d$.

The RR protocol provides a similar instance reduction, but only with respect to the Hamming distance.

Theorem 7 (The original RR Instance Reduction Protocol, informally rephrased; [RR20]). *Suppose $M \in \mathbb{N}$. There exists a protocol (P_{RR}, V_{RR}) , where both the prover and verifier gets the description $\langle \Phi \rangle$ of a log-space uniform circuit $\Phi : \{0, 1\}^M \rightarrow \{0, 1\}$, and the prover gets the auxiliary input $\mathbf{w} \in \{0, 1\}^M$, such that if both parties gets the parameter $d = d(n) \in \mathbb{N}$, then the protocol outputs a succinct description $\langle \mathcal{Q} \rangle$ of a set $\mathcal{Q} \subset [M]$ and a new circuit $\Phi' : \{0, 1\}^{|\mathcal{Q}|} \rightarrow \{0, 1\}$, such that $|\mathcal{Q}| \leq \frac{|M|}{d}$. The protocol is sound in the following sense: if $\Phi(\mathbf{w}) = 0$ for every $\mathbf{w}$ that is (Hamming) d -close to $\mathbf{w}$, then $\Phi'(\mathbf{w}|_{\mathcal{Q}}) = 0$. It has $\tilde{O}(1)$ rounds and communication complexity $\tilde{O}(d)$. The prover runtime is $\text{poly}(M)$ and the verifier runtime is $\tilde{O}(d)$.*

As explained in Section 2.2, $\mathcal{B}_{d,\mathbb{F}}(\mathbf{a}^{S_{\text{Dist}}})$ is a strict subset of the Hamming ball. The soundness of Theorem 7 applies only when *every Hamming-close string* is rejected, which is not always true in our case because our notion of distance is different.

The RR Instance Reduction Template Theorem 7 is originally proven by the following two steps:

1. Use the GKR protocol (Theorem 3) to transform the input claim $\Phi(\mathbf{w}) = 1$ into the claim $\mathbf{w} \in \text{PVAL}(\mathbf{j}, \mathbf{v})$. The parameters are chosen to ensure that *distance is preserved*: if every $\mathbf{w}'$ that is d -close to $\mathbf{w}$ is rejected by Φ , then $\mathbf{w}$ is d -far from $\text{PVAL}(\mathbf{j}, \mathbf{v})$.
2. Apply a special-purpose protocol for reducing PVAL instances, on the input $\mathbf{w} \in \text{PVAL}(\mathbf{j}, \mathbf{v})$. This outputs a subset of coordinates $\mathcal{Q} \subset [M]$ that specifies the positions of the coordinates of $\mathbf{w}$ to read, and $|\mathcal{Q}| \leq |M|/d$, along with a predicate Φ' , such that $\Phi'(\mathbf{w}|_{\mathcal{Q}}) = 1$ only if $\mathbf{w}$ is Hamming- d -close to $\text{PVAL}(\mathbf{j}, \mathbf{v})$.

It turns out that the first step can be modified to work for Δ_c -distance quite straightforwardly (Lemma 2). However, for the second step, we have to make lower level modifications in order to make it work for our notion of distance. This discussion appears in Section 6.

In the rest of this section, we formally recall the guarantees of the two sub-protocols (Lemmas 2 and 6), corresponding to the two steps mentioned above, and use them to prove Lemma 4.

The GKR Protocol is Distance-Preserving Let $L := al$. To preserve the distance, we choose the parameter $T = \tilde{\Theta}(dL)$ (see Lemma 2). In the GKR protocol, both parties take as input the description $\langle \Phi_{\text{Dist}} \rangle$, and the prover additionally has access to $\mathbf{a}^{S_{\text{Dist}}}$. After running the GKR protocol with parameter T , the verifier obtains $\mathbf{j} \in (\mathbb{F}^m)^T, \mathbf{v} \in \mathbb{F}^T$ as output, with the guarantee that $\text{PVAL}(\mathbf{j}, \mathbf{v})$ is Δ_c - d -close to $\Phi_{\text{Dist}}(\mathbf{a}^{S_{\text{Dist}}})$ iff $\mathbf{a}^{S_{\text{Dist}}}$ is Δ_c - d -close to $\Phi_{\text{Dist}}(\mathbf{a}^{S_{\text{Dist}}})$.¹⁹

Recall that $D_{\text{Dist}} := \text{depth}(\Phi_{\text{Dist}})$, $S_{\text{Dist}} := \text{size}(\Phi_{\text{Dist}})$. The complexity of the protocol is as follows, where $\tilde{O}$ hides $\text{polylog}(\log |\mathbb{F}|, M, L)$ factors:

¹⁹Recall that a matrix $\mathbf{a}$ is Δ_c - d -close to $\mathbf{a}^{S_{\text{Dist}}}$ iff $\mathbf{a} \in \mathcal{B}_{d,\mathbb{F}}(\mathbf{a}^{S_{\text{Dist}}})$ (Definition 1).

- $\epsilon_{\text{GKR}} = \frac{O(D_{\text{Dist}} \log S_{\text{Dist}})}{|\mathbb{F}|} + \left(O\left(\frac{D_{\text{Dist}} \log S_{\text{Dist}}}{|\mathbb{F}|}\right) \right)^T \left(\binom{M}{d} |\mathbb{F}|^d \right)^L < 2^{-\sigma+2}$, so long as $O\left(\frac{D_{\text{Dist}} \log S_{\text{Dist}}}{|\mathbb{F}|}\right) < 2^{-\sigma-3}$.
- $\ell_{\text{GKR}} = O(D_{\text{Dist}} \log S_{\text{Dist}})$.
- $a_{\text{GKR}} = b_{\text{GKR}} = \tilde{O}(T \log |\mathbb{F}|) = \tilde{O}(d\ell)$.
- $\text{Ptime}_{\text{GKR}} = \tilde{O}(T \cdot \text{poly}(S_{\text{Dist}}) \cdot \text{polylog}|\mathbb{F}|) = \tilde{O}(d\ell) \cdot \text{poly}(S_{\text{Dist}})$.
- $\text{Vtime}_{\text{GKR}} = \tilde{O}(T \cdot D_{\text{Dist}} \log S_{\text{Dist}} \cdot \log |\mathbb{F}| + T \cdot |\langle \Phi_{\text{Dist}} \rangle| \log |\mathbb{F}|) = \tilde{O}(d\ell \cdot (D_{\text{Dist}} \log S_{\text{Dist}} + |\langle \Phi_{\text{Dist}} \rangle|))$.

And the verifier's verdict circuit satisfies:

- $\text{size}(\text{V}_{\text{GKR}}) = \tilde{O}(T \cdot D_{\text{Dist}} \log S_{\text{Dist}} + T \cdot |\langle \Phi_{\text{Dist}} \rangle| \log |\mathbb{F}|) = \tilde{O}(d\ell \cdot (D_{\text{Dist}} \log S_{\text{Dist}} + |\langle \Phi_{\text{Dist}} \rangle|))$.
- $\text{depth}(\text{V}_{\text{GKR}}) = \tilde{O}(1)$.

The Instance Reduction Protocol for PVAL with Δ_c -distance $(\text{P}_{\Delta_c \text{RR}}, \text{V}_{\Delta_c \text{RR}})$ is a special-purpose protocol for reducing PVAL instances with Δ_c distances. Both parties in the protocol take as input the parameters $\mathbf{j} \in (\mathbb{F}^m)^T, \mathbf{v} \in \mathbb{F}^T$ that specify $\text{PVAL}(\mathbf{j}, \mathbf{v})$, and the prover additionally has the input $\mathbf{a}^{S_{\text{Dist}}}$. The protocol outputs $\langle \mathcal{S}' \rangle, \langle \Phi' \rangle$ such that $|\mathcal{S}'| \leq |S_{\text{Dist}}|/d$, with the guarantee that $(\Phi', \mathcal{S}') \in \mathcal{L}'_x$ iff $\mathbf{a}^{S_{\text{Dist}}}$ is Δ_c - d -close to $\text{PVAL}(\mathbf{j}, \mathbf{v})$.

The following is a restatement of Theorem 9. Its proof appears in Section 6.

Lemma 6 (The Instance Reduction Protocol for PVAL with Δ_c -distance (Theorem 9 in Section 6)). *Suppose $\sigma, d, m, \log L \in \mathbb{N}$, $M = 2^m$, and $\mathbb{F}$ is a field. Let $\tilde{O}$ be omitting $\text{polylog}(|\mathbb{F}|, M, L)$ factors. Suppose $T \geq 8dLm = \tilde{O}(dL)$. Let $\mathbf{j} = (\mathbf{j}_1, \dots, \mathbf{j}_T) \in (\mathbb{F}^{m+\log L})^T$ and $\mathbf{v} = (v_1, \dots, v_T) \in \mathbb{F}^T$.*

There exists a constant C , and a public-coin unambiguous IPP $(\text{P}_{\Delta_c \text{RR}}, \text{V}_{\Delta_c \text{RR}})$ for $\text{PVAL}(\mathbf{j}, \mathbf{v})$, where both the prover and verifier gets the input $\mathbf{j}, \mathbf{v}$, while the prover additionally gets the input $\mathbf{a} \in \mathbb{F}^{M \times L}$, such that if both parties additionally gets parameters $\sigma, d \in \mathbb{N}$ and $\mathbb{F}$, that satisfy the following preconditions:

- $d \geq 16m\sigma$,
- $|\mathbb{F}| \geq 32 \cdot 2^\sigma (T \log M \log L)^C$,

then the protocol outputs a succinct description $\langle \mathcal{Q} \rangle$ of descriptions $\langle \mathcal{Q} \rangle$ of a set of rows $\mathcal{Q} \subsetneq [M]$, and $\langle \Phi' \rangle$ of a predicate Φ' ,

- **Prescribed Completeness:** *If $\text{V}_{\Delta_c \text{RR}}$ interacts with $\text{P}_{\Delta_c \text{RR}}$, then $\Phi'(\mathbf{a}[\mathcal{Q}, :]) = 1$ iff $\mathbf{a} \in \text{PVAL}(\mathbf{j}, \mathbf{v})$.*
- **$2^{-\sigma-2}$ -Unambiguity:** *Suppose $\Delta_c(\text{PVAL}(\mathbf{j}, \mathbf{0})) \geq 4d$, then for any (unbounded) cheating prover strategy P^* that deviates from $\text{P}_{\Delta_c \text{RR}}$ first in round $r^* \in [\ell_{\Delta_c \text{RR}}]$, with probability at least $1 - 2^{-\sigma-2}$ over the verifier's remaining coins, either $\text{V}_{\Delta_c \text{RR}}$ rejects, or $\Phi'(\mathbf{a}[\mathcal{Q}, :]) = 0$.*
- **Reduced Query:** *The subset of rows $\mathcal{Q} \subsetneq [M]$ has size $|\mathcal{Q}| \leq \lceil 8\sigma \cdot \frac{M}{d} \rceil$.*
- *Regardless of the prover's strategy, the distribution of $\mathcal{Q}$ only depends on the verifier's random coins.*

The complexity of the protocol is as follows.

- $\ell_{\Delta_c RR} = \tilde{O}(1)$.
- $a_{\Delta_c RR} = b_{\Delta_c RR} = \tilde{O}(L + \text{poly}(d))$.
- $\text{Ptime}_{\Delta_c RR} = \text{poly}(ML, T \cdot \log |\mathbb{F}|)$.
- $\text{Vtime}_{\Delta_c RR} = \tilde{O}(L + \text{poly}(d))$.

The bit-lengths are $|\langle \mathcal{Q} \rangle| = \tilde{O}(\text{poly}(d))$ and $|\langle \Phi' \rangle| = \tilde{O}(L + \text{poly}(d))$. Let $G_{\mathcal{Q}}(i, \langle \mathcal{Q} \rangle)$ be the circuit that returns the i -th element in $\mathcal{Q}$, and $C'(\mathbf{a}, \langle \Phi' \rangle)$ be the circuit that computes $\Phi'(\mathbf{a})$, then they satisfy the following.

- $\text{size}(G_{\mathcal{Q}}) = \tilde{O}(\text{poly}(d))$.
- $\text{depth}(G_{\mathcal{Q}}) = \tilde{O}(1)$.
- $\text{size}(C') = \tilde{O}(|\mathcal{Q}| \cdot L)$.
- $\text{depth}(C') = \tilde{O}(1)$.

The verifier's verdict circuit (which outputs 0 iff it rejects amidst the protocol) satisfies

- $\text{size}(V_{\Delta_c RR}) = \tilde{O}(L + \text{poly}(d))$.
- $\text{depth}(V_{\Delta_c RR}) = \tilde{O}(1)$.

4.4.2 Construction of the Reduce Protocol

We give the construction in Protocol 3.

4.4.3 Analysis of the Reduce Protocol

We first verify that the preconditions for $(P_{\Delta_c RR}, V_{\Delta_c RR})$ are satisfied. Indeed, $d > 16m\sigma$, and $|\mathbb{F}| \geq 32 \cdot 2^\sigma (T \log M \log L)^C$ are assumptions of Lemma 4. By Lemma 6, $|\mathcal{Q}| \leq \lceil 8\sigma \cdot \frac{M}{d} \rceil = O(\sigma \frac{M}{d})$.

Prescribed Completeness. By the prescribed completeness of $(P_{\text{GKR}}, V_{\text{GKR}})$, $\mathbf{a}^{\mathcal{S}_{\text{Dist}}} \in \text{PVAL}(\mathbf{j}, \mathbf{v})$ iff $\Phi_{\text{Dist}}(\mathbf{a}^{\mathcal{S}_{\text{Dist}}}) = 1$ (which is equivalent to $(\mathcal{S}_{\text{Dist}}, \Phi_{\text{Dist}}) \in \mathcal{L}'_x$ by the construction of Φ_{Dist} in Lemma 3), and then by the prescribed completeness of $(P_{\Delta_c RR}, V_{\Delta_c RR})$, $\Phi'(\mathbf{a}^{\mathcal{S}_{\text{Dist}}}[\mathcal{Q}, :]) = 1$ iff $\mathbf{a}^{\mathcal{S}_{\text{Dist}}} \in \text{PVAL}(\mathbf{j}, \mathbf{v})$. Finally, $\mathcal{S}' = \mathcal{S}_{\text{Dist}}[\mathcal{Q}, :]$, so $(\mathcal{S}', \Phi') \in \mathcal{L}'_x$ iff $(\mathcal{S}_{\text{Dist}}, \Phi_{\text{Dist}}) \in \mathcal{L}'_x$.

Unambiguity. With the same proof as in Proposition 1, with all but $2^{-\sigma-1}$ probability, the following claim holds, so the unambiguity of $(P_{\Delta_c RR}, V_{\Delta_c RR})$ applies.

Claim 3. With probability at least $1 - 2^{-\sigma-1}$, $\Delta(\text{PVAL}(\mathbf{j}, \mathbf{0})) \geq 4d$.

Proof. $(P_{\text{GKR}}, V_{\text{GKR}})$ always outputs T truly random points in $(\mathbb{K}^{m+\log L})^T$. The number of points with Δ_c distance at most $4d$ from the origin is at most $\binom{M}{4d} \cdot |\mathbb{F}|^{4dL}$. Fixing such a point $x \in \{0, 1\}^{M \times L}$, since every $\mathbf{j}$ is uniformly random, by Lemma 1, it is a root of the LDE of x with

Protocol 3 ($P_{\text{Reduce}}, V_{\text{Reduce}}$) Instance Reduction for Δ_c -Distance.

Input Parameters: $\sigma, d \in \mathbb{N}$, and a field $\mathbb{F}$ with $|\mathbb{F}| \geq 32 \cdot 2^\sigma (d\ell m)^C$ (where C is some constant).

Input: $\mathbf{x}, \langle \mathcal{S}_{\text{Dist}} \rangle, \langle \Phi_{\text{Dist}} \rangle$.

Other Parameters: $M, L \in \mathbb{N}$, $T = 8dLm$.

Ingredients:

- *The GKR protocol* ($P_{\text{GKR}}, V_{\text{GKR}}$) (Theorem 3) for checking $\Phi(\mathbf{a}) = 1$.

The prover and verifier gets the description $\langle \Phi \rangle$, while the prover additionally gets the input $\mathbf{a}$. Both parties additionally takes in the parameter $T \in \mathbb{N}$, and in the end outputs $\mathbf{j} \in (\mathbb{F}^{m+\log L})^T, v \in \mathbb{F}^T$.

- *The Instance Reduction Protocol for PVAL* ($P_{\Delta_c\text{RR}}, V_{\Delta_c\text{RR}}$) (Lemma 6) for checking that $\mathbf{a}$ is Δ_c - d -close to $\text{PVAL}(\mathbf{j}, v)$.

Both the prover and verifier gets the input $\mathbf{j}, v$, and the prover additionally gets the input $\mathbf{a}$. The protocol outputs descriptions $\langle \mathcal{Q} \rangle$ of a set of rows $\mathcal{Q} \subsetneq [M]$, and $\langle \Phi' \rangle$ of a predicate Φ' .

Output: $\langle \mathcal{S}' \rangle, \langle \Phi' \rangle$, new descriptions for the claim $(\mathcal{S}', \Phi') \in \mathcal{L}'$.

- 1: P_{Reduce} compute the prescribed $\mathbf{a}^{\mathcal{S}_{\text{Dist}}}$.
- 2: Both parties run ($P_{\text{GKR}}, V_{\text{GKR}}$) (Theorem 3) with $T = 8dLm$ on the claim

$$\Phi_{\text{Dist}}(\mathbf{a}^{\mathcal{S}_{\text{Dist}}}) = 1.$$

The verifier obtains $\mathbf{j}, v$.

- 3: Both parties run the Instance Reduction Protocol for PVAL, ($P_{\Delta_c\text{RR}}, V_{\Delta_c\text{RR}}$) (Lemma 6), on the claim

$$\mathbf{a}^{\mathcal{S}_{\text{Dist}}} \in \text{PVAL}(\mathbf{j}, v),$$

obtaining descriptions $\langle \mathcal{Q} \rangle$ of a set of rows $\mathcal{Q} \subsetneq [M]$, and $\langle \Phi' \rangle$ of a predicate Φ' .

- 4: Let $\langle \mathcal{S}' \rangle = (\langle \mathcal{S}_{\text{Dist}} \rangle, \langle \mathcal{Q} \rangle)$. $\triangleright$ Implicitly define $\mathcal{S}' := \mathcal{S}_{\text{Dist}}[\mathcal{Q}, :]$ and thus $|\mathcal{S}'| = |\mathcal{Q}| \leq |\mathcal{S}_{\text{Dist}}|/d$.
 - 5: **return** $\langle \mathcal{S}' \rangle, \langle \Phi' \rangle$.
-

probability at most $\frac{m}{|\mathbb{F}|}$. Therefore, the probability $x \in \text{PVAL}(\mathbf{j}, \mathbf{v})$ is at most $\left(\frac{m}{|\mathbb{F}|}\right)^T$. Taking a union bound over all such points, and using that $T \geq 8dLm$,

$$\Pr[\Delta_c(\text{PVAL}(\mathbf{j}, \mathbf{v})) \geq 4d] \leq \left(\binom{M}{4d} \cdot |\mathbb{F}|^{4d}\right)^L \cdot \left(\frac{m}{|\mathbb{F}|}\right)^T \leq 2^{-\sigma-1}. \quad \square$$

Let $\mathbf{P}^*$ be a cheating prover that deviates from $\mathbf{P}_{\text{Reduce}}$, then it must be deviating in either of the two protocols. By the unambiguity of $(\mathbf{P}_{\text{GKR}}, \mathbf{V}_{\text{GKR}})$ and $(\mathbf{P}_{\Delta_c\text{RR}}, \mathbf{V}_{\Delta_c\text{RR}})$ and the last claim, the probability that the prover produces $\langle \mathcal{S}' \rangle, \langle \Phi' \rangle$ such that $\Phi'(\mathcal{S}') = 1$ is at most $\epsilon_{\text{GKR}} + \epsilon_{\Delta_c\text{RR}} + 2^{-\sigma-1} \leq 2^{-\sigma}$.

Complexities. The GKR protocol has $\ell_{\text{GKR}} = O(D_{\text{Dist}} \log S_{\text{Dist}})$ rounds. We can reduce the message length per-round (a_{GKR}) by increasing the number of rounds to $\ell'_{\text{GKR}} = \tilde{O}(\ell \cdot D_{\text{Dist}} \log S_{\text{Dist}})$, which changes its complexity to $a'_{\text{GKR}} = b'_{\text{GKR}} = \tilde{O}(da)$. Similarly, for $(\mathbf{P}_{\Delta_c\text{RR}}, \mathbf{V}_{\Delta_c\text{RR}})$, we increase its number of rounds to $\ell'_{\Delta_c\text{RR}} = \tilde{O}(\ell)$, and its message lengths reduce to $a'_{\Delta_c\text{RR}} = b'_{\Delta_c\text{RR}} = \tilde{O}(da + \text{poly}(d))$. This is a simplifying assumption to make per-round messages have comparable lengths as in $(\mathbf{P}_{\text{Dist}}, \mathbf{V}_{\text{Dist}})$.

- $\ell_{\text{Reduce}} = \ell'_{\text{GKR}} + \ell'_{\Delta_c\text{RR}} = \tilde{O}(\ell \cdot D_{\text{Dist}} \log S_{\text{Dist}})$.
- $a_{\text{Reduce}} = b_{\text{Reduce}} = \max(a'_{\text{GKR}}, a'_{\Delta_c\text{RR}}) = \tilde{O}(da + \text{poly}(d))$.
- $\text{Ptime}_{\text{Reduce}} = \text{Ptime}_{\text{GKR}} + \text{Ptime}_{\Delta_c\text{RR}}$
 $= \tilde{O}(dL) \cdot \text{poly}(S_{\text{Dist}}) + \text{poly}(ML, T \cdot \log |\mathbb{F}|)$
 $= \text{poly}(\log |\mathbb{F}|, M, L, d, S_{\text{Dist}})$.
- $\text{Vtime}_{\text{Reduce}} = \text{Vtime}_{\text{GKR}} + \text{Vtime}_{\Delta_c\text{RR}} + \tilde{O}(k \cdot n + |\langle \mathcal{S}_{\text{Dist}} \rangle| + |\langle \Phi_{\text{Dist}} \rangle|)$ Note that the
 $\leq \tilde{O}(dal \cdot (D_{\text{Dist}} \log S_{\text{Dist}} + |\langle \Phi_{\text{Dist}} \rangle|) + \text{poly}(d) + |\langle \mathcal{S}_{\text{Dist}} \rangle| + k \cdot n)$
 third term in the first summation is due to the overhead of reading the input $(\mathcal{S}_{\text{Dist}}, \mathbf{x})$.

Circuit G' for generating $\mathcal{S}'$. The sub-protocol $(\mathbf{P}_{\Delta_c\text{RR}}, \mathbf{V}_{\Delta_c\text{RR}})$ outputs the description $\langle \mathcal{Q} \rangle$. Our new description $\langle \mathcal{S}' \rangle = (\langle \mathcal{S}_{\text{Dist}} \rangle, \langle \mathcal{Q} \rangle)$. We represent every pair $(i, \mathbf{q})$ as elements in $\{0, 1\}^p$, where $p = \lceil \log k \rceil + b \cdot \ell$. Let $g' = |\mathcal{S}'|$ be the number of pairs to output. The Turing machine $\mathcal{M}$ takes in $(g', 1^{|\langle \mathcal{S}' \rangle|}, 1^p)$, and outputs the circuit $G' : [g'] \times \{0, 1\}^{|\langle \mathcal{S}' \rangle|} \rightarrow \{0, 1\}^p$. $G'(\langle \mathcal{S}' \rangle)$ enumerates $\mathcal{S}'$ as follows: for every $s \in [g']$,

- Compute the s -th element in $\mathcal{Q} \in [M]$ (using $G_{\mathcal{Q}}(s, \langle \mathcal{Q} \rangle)$), and let it be $\mathcal{Q}_s \in [M]$.
- Output the $\mathcal{Q}_s$ -th pair in $\mathcal{S}_{\text{Dist}}$ (using $G_{\text{Dist}}(\mathcal{Q}_s, \langle \mathcal{S}_{\text{Dist}} \rangle)$).

The circuit satisfies:

- $\text{size}(G') = \text{size}(G_{\mathcal{Q}}) + \text{size}(G_{\text{Dist}}) = \text{size}(G_{\text{Dist}}) + \tilde{O}(\text{poly}(d))$.
- $\text{depth}(G') = \text{depth}(G_{\text{Dist}}) + \tilde{O}(1)$.

The bit length satisfies $|\langle \mathcal{S}' \rangle| = \tilde{O}(\text{poly}(d) + |\langle \mathcal{S}_{\text{Dist}} \rangle|)$. G' is p -space uniform because $G_{\mathcal{Q}}$ is p -space uniform and $\mathcal{M}$ only needs to maintain additional pointers into the circuit G' .

Circuit C' for checking Φ' . Note that the sub-protocol $(P_{\Delta_cRR}, V_{\Delta_cRR})$ already outputs the description $\langle \Phi' \rangle$. By Lemma 6, $|\langle \Phi' \rangle| = \tilde{O}(dal + \text{poly}(d))$, and the complexity of $C'(\mathbf{a}, \langle \Phi' \rangle) = \Phi'(\mathbf{a})$ is the following:

- $\text{size}(C') = \tilde{O}(\frac{M}{d} \cdot al)$,
- $\text{depth}(C') = \tilde{O}(1)$.

Verdict Circuit V_{Reduce} . The verdict circuit V_{Reduce} does not reject iff neither V_{GKR} nor V_{Δ_cRR} rejects. Therefore, the circuit that implements this check would satisfy:

- $\text{size}(V_{\text{Reduce}}) = \tilde{O}(\text{size}(V_{\text{GKR}}) + \text{size}(V_{\Delta_cRR}) + |\langle S_{\text{Dist}} \rangle| + |\langle \Phi_{\text{Dist}} \rangle| + k \cdot n)$
 $= \tilde{O}(dal(D_{\text{Dist}} \log S_{\text{Dist}} + |\langle \Phi_{\text{Dist}} \rangle|) + \text{poly}(d) + |\langle S_{\text{Dist}} \rangle| + k \cdot n)$.
- $\text{depth}(V_{\text{Reduce}}) = \max(\text{depth}(V_{\text{GKR}}), \text{depth}(V_{\Delta_cRR}))$
 $= \tilde{O}(1)$.

4.5 Analysis of Other Sub-protocols

We present of the deferred sub-protocols and analyze their properties in this section.

4.5.1 Proof of Proposition 2: the Random Continuation Protocol

The protocol $(P^{\text{tc}}, V^{\text{tc}})$ is formally given in Protocol 4.

Prescribed Completeness. This is immediate from the prescribed completeness of (P, V) .

Unambiguity. We consider the following two cases.

- If the input $\mathbf{a} = \mathbf{a}^{x,q}$, let $r^* \in [\ell]$ be the first round when P^* deviates from P^{tc} . Then P^* sends a non-prescribed message $\tilde{\mathbf{a}}_{r^*}^{x,q_j^{\text{Hyb}}} \neq \mathbf{a}_{r^*}^{x,q_j^{\text{Hyb}}}$ for some $j \in [\ell]$. With probability at most ϵ over the remaining interaction, $V(x_i, \mathbf{q}_j^{\text{Hyb}}, \tilde{\mathbf{a}}_{r^*}^{x,q_j^{\text{Hyb}}}) = 1$ by the unambiguity of (P, V) , so P^* would be rejected by V^{tc} with probability at least $1 - \epsilon$.
- If the input $\mathbf{a} \neq \mathbf{a}^{x,q}$, let $j^* \in [\ell]$ be the first round for which $\mathbf{a}_{j^*} \neq \mathbf{a}_{j^*}^{x,q}$. Since V^{tc} checks that $\mathbf{a}$ and $\mathbf{a}^{x,q_{j^*}^{\text{Hyb}}}$ share a prefix of length $j^* \cdot a$, V^{tc} would reject at once if it is not the case that $\mathbf{a}_{j^*} = \mathbf{a}_r^{x,q_{j^*}^{\text{Hyb}}}$ for every $1 \leq r \leq j^*$. Therefore, we assume $\mathbf{a}_{j^*} = \mathbf{a}_r^{x,q_{j^*}^{\text{Hyb}}}$ for every $1 \leq r \leq j^*$. This means that $\mathbf{a}^{x,q_{j^*}^{\text{Hyb}}}$ is the transcript of an interaction where P^* deviates firstly in round j^* of the base protocol. By the unambiguity of (P, V) , with probability at most ϵ over the remaining interaction starting from the j^* -th round, $\mathbf{a}^{x,q_{j^*}^{\text{Hyb}}}$ would be rejected, which means P^* would be rejected by V^{tc} with probability at least $1 - \epsilon$.

Complexities. The number of rounds is ℓ . The prover message is $a^{\text{tc}} = \ell \cdot a$, and the verifier message is $b^{\text{tc}} = b$. The prover time is $\text{Ptime}^{\text{tc}} = \ell \cdot \text{Ptime}$, and the verifier time is $\text{Vtime}^{\text{tc}} = \text{Vtime} + \ell \cdot \text{size}(V)$.

Transcript Check Structures. By inspection, indeed the j -th message of $\mathbf{P}^{\text{tc}}$ is

$$\mathbf{m}_r = \left((\mathbf{a}^{x, \mathbf{q}_j^{\text{Hyb}}})_{r \in [\ell]} \right)^\top,$$

and the verifier $\mathbf{V}^{\text{tc}}$ simply checks all $\mathbf{a}^{x, \mathbf{q}_j^{\text{Hyb}}}$ are accepted, and that $\mathbf{a}^{x, \mathbf{q}}$ and $\mathbf{a}^{x, \mathbf{q}_j^{\text{Hyb}}}$ share a prefix of length $j \cdot a$.

Protocol 4 ($\mathbf{P}^{\text{tc}}, \mathbf{V}^{\text{tc}}$) Random Continuation.

Input: $x \in \{0, 1\}^n$, $\mathbf{q} \in (\{0, 1\}^b)^\ell$, and a transcript $\mathbf{a} \in \{0, 1\}^{a \times \ell}$ purported to be $\mathbf{a}^{x, \mathbf{q}}$.

Output: the verifier decides to accept or reject the claim that $\mathbf{a}$ is indeed $\mathbf{a}^{x, \mathbf{q}}$, and that $x \in \mathcal{L}$.

- 1: $\mathbf{V}^{\text{tc}}$ samples $\mathbf{q}' = (q'_1, \dots, q'_\ell) \in (\{0, 1\}^b)^\ell$.
 - 2: **for** $r = 1, \dots, \ell$ **do**
 - 3: $\mathbf{V}^{\text{tc}}$ sends q'_r to $\mathbf{P}^{\text{tc}}$.
 - 4: **for** $j = 1, \dots, \ell$ **do**
 - 5: $\mathbf{P}^{\text{tc}}$ let $\mathbf{a}_r^{x, \mathbf{q}_j^{\text{Hyb}}} \leftarrow \mathbf{P}(x, (\mathbf{q}_j^{\text{Hyb}})_{\leq r})$, where $\mathbf{q}_j^{\text{Hyb}} = (q_1, \dots, q_j, q'_{j+1}, \dots, q'_\ell)$.
 - 6: $\mathbf{P}^{\text{tc}}$ sends $\left\{ \mathbf{a}_r^{x, \mathbf{q}_j^{\text{Hyb}}} \right\}_{j \in [r]}$ to $\mathbf{V}^{\text{tc}}$.
 - 7: $\mathbf{V}^{\text{tc}}$ let $\mathbf{a}^{x, \mathbf{q}_j^{\text{Hyb}}} \leftarrow (\mathbf{a}_1^{x, \mathbf{q}_j^{\text{Hyb}}}, \dots, \mathbf{a}_\ell^{x, \mathbf{q}_j^{\text{Hyb}}})$ for every $j \in [\ell]$.
 - 8: $\mathbf{V}^{\text{tc}}$ accepts iff $\bigwedge_{j \in [\ell]} [\mathbf{V}(x, \mathbf{q}_j^{\text{Hyb}}, \mathbf{a}^{x, \mathbf{q}_j^{\text{Hyb}}}) = 1]$ and that for every $j \in [\ell]$, $\mathbf{a}$ and $\mathbf{a}^{x, \mathbf{q}_j^{\text{Hyb}}}$ share a prefix of length $j \cdot a$.
-

4.5.2 Proof of Lemma 5: the Explicit UIP for the associated language

We construct $(\mathbf{P}_{\mathcal{L}'_x}, \mathbf{V}_{\mathcal{L}'_x})$, an explicit UIP for the associated language $\mathcal{L}'_x$, which takes as input $\mathbf{x}$ and $\langle \mathcal{S} \rangle, \langle \Phi \rangle$ and verifies the claim $(\mathcal{S}, \Phi) \in \mathcal{L}'_x$. It consists of the following two steps:

1. $\mathbf{P}_{\mathcal{L}'_x}$ sends the prescribed transcript matrix $\mathbf{a}^{\mathcal{S}} \in \{0, 1\}^{|\mathcal{S}| \times (a\ell)}$ to $\mathbf{V}_{\mathcal{L}'_x}$.
2. $\mathbf{V}_{\mathcal{L}'_x}$ enumerates $G(s, \langle \mathcal{S} \rangle)$ for every $s \in [|\mathcal{S}|]$.
3. Both parties execute $(\mathbf{P}^{\text{tc}}, \mathbf{V}^{\text{tc}})$ explicitly on the input $((i, \mathbf{q}), \mathbf{a}^{\mathcal{S}}[(i, \mathbf{q}), :])$, for every $(i, \mathbf{q}) \in \mathcal{S}$ in parallel.

By Remark 10, the messages sent by $\mathbf{P}_{\mathcal{L}'_x}$ constitute the transcript matrix $\mathbf{a} = \mathbf{a}^{\mathcal{S}^{\text{Hyb}}}$, where $\mathcal{S}^{\text{Hyb}}$ is the set of hybrid pairs in constructed between $\mathcal{S}$ and $\mathbf{q}' \in \{0, 1\}^{b\ell}$, the random coin sampled by $\mathbf{V}_{\mathcal{L}'_x}$.

4. $\mathbf{V}_{\mathcal{L}'_x}$ accepts iff every execution of $(\mathbf{P}^{\text{tc}}, \mathbf{V}^{\text{tc}})$ in $\mathbf{a}$ is accepted, $G(\mathbf{a}^{\mathcal{S}}, \langle \Phi \rangle) = 1$ and $\mathbf{a}[\mathcal{S}, :] = \mathbf{a}^{\mathcal{S}}$.

Let $\mathcal{I} := \{(i, \mathbf{q}) : (i, \mathbf{q}) \in \mathcal{S}\}$. Recall that $(\mathcal{S}, \Phi) \in \mathcal{L}'_x$ iff $(\Phi(\mathbf{a}^{\mathcal{S}}) = 1) \wedge (\mathbf{x}|_{\mathcal{I}} \in \mathcal{L}^{\otimes |\mathcal{S}|})$.

Prescribed Completeness. $\mathbf{P}_{\mathcal{L}'_x}$ sends the verifier the prescribed $\mathbf{a}^{\mathcal{S}}$ in the first step, and executes $\mathbf{P}_{\mathcal{L}'_x}$ honestly for every $(i, \mathbf{q}) \in \mathcal{S}$ in parallel, so $\mathbf{a} = \mathbf{a}^{\mathcal{S}^{\text{Hyb}}}$ (and thus $\mathbf{a}[\mathcal{S}, :] = \mathbf{a}^{\mathcal{S}}$). Since $\mathbf{a}^{\mathcal{S}}$ is the prescribed transcript matrix, all executions of $(\mathbf{P}^{\text{tc}}, \mathbf{V}^{\text{tc}})$ in $\mathbf{a}$ are accepted iff $\mathbf{x}|_{\mathcal{I}} \in \mathcal{L}^{\otimes |\mathcal{S}|}$. Therefore, the verifier's check passes exactly when $(\mathcal{S}, \Phi) \in \mathcal{L}'_x$.

Unambiguity. Suppose that P^* deviates from $P_{\mathcal{L}'_x}$ in the first step, i.e. it sends some $\mathbf{a}^* \neq \mathbf{a}^S$ to $V_{\mathcal{L}'_x}$. This means for some row-index $(i, \mathbf{q}) \in \mathcal{S}$, $\mathbf{a}^*[(i, \mathbf{q})] \neq \mathbf{a}^{i, \mathbf{q}}$. If P^* does not deviate in the rest of the protocol, then by prescribed completeness of (P^{tc}, V^{tc}) (Proposition 2), $V_{\mathcal{L}'_x}$ rejects on the input $((i, \mathbf{q}), \mathbf{a}^*[(i, \mathbf{q})])$, and thus $V_{\mathcal{L}'_x}$ rejects. Otherwise, P^* deviates in one of the ℓ rounds while executing (P^{tc}, V^{tc}) in parallel. By the ϵ -unambiguity of (P^{tc}, V^{tc}) (Proposition 2), that execution of (P^{tc}, V^{tc}) is accepted with probability at most ϵ , so $V_{\mathcal{L}'_x}$ accepts with probability at most ϵ .

Complexities. Let $g = |\mathcal{S}|$. The number of rounds is $\ell_{\mathcal{L}'_x} = \ell + 1$. The prover message (per round) is of length $a_{\mathcal{L}'_x} = g\ell \cdot a$. This can be seen by noting that the message in the first step is of length $g \cdot a\ell$ and the rest are of length $g \cdot a\ell$ per round as well (since g instances of the protocol in Proposition 2 are in parallel). The verifier message (per round) is of length $b_{\mathcal{L}'_x} = b$. The prover runtime is $Ptime_{\mathcal{L}'_x} = O(g\ell \cdot Ptime)$, because it simply runs g instances P^{tc} in parallel (each of which has runtime $\ell \cdot Ptime$), and the verifier runtime is $Vtime_{\mathcal{L}'_x} = O(g\ell \cdot Vtime + gsize(G) + size(C))$. This is because it first computes $G(s, \langle \mathcal{S} \rangle)$ for every $s \in [g]$, then runs g instances of V^{tc} in parallel (each of which has runtime $\ell \cdot Vtime$), and finally evaluates $C(\mathbf{a}^S, \langle \Phi \rangle)$ (which takes time $size(C)$).

The verifier's decision circuit $V_{\mathcal{L}'_x}$ has the following complexities:

- $size(V_{\mathcal{L}'_x}) = \tilde{O}(g \cdot \ell \cdot size(V) + gsize(G) + size(C))$.
- $depth(V_{\mathcal{L}'_x}) = \max(depth(G) + D, depth(C)) + \tilde{O}(1)$.

5 Doubly-Efficient Interactive Proofs

5.1 A New UIP for Verifying Space-Bounded Computations

We demonstrate the power of Batch-UIP by constructing a doubly-efficient interactive proof for delegating space-bounded computations. Denote by $DTISP(T, S)$ the class of languages decidable by a Turing machine in time $T(n)$ and space $S(n)$.

Theorem 8 (Doubly-efficient Interactive Proof for $DTISP(T, S)$). *For any $T(n) \leq n^{O\left(\left(\frac{\log n}{\log \log n}\right)^{1/2}\right)}$ and $S(n) = \text{poly}(n)$, any $\sigma = \sigma(n) = \text{polylog}(n)$, any $\delta = \delta(n) \in (0, 1)$, and any language $L \in DTISP(T, S)$, there exists a UIP with the following properties:*

- The soundness error is $2^{-\sigma} \cdot (1/\delta) \cdot (\log n)^{O(1/\delta)}$.
- The number of rounds is $(\log n)^{O(1/\delta)}$.
- The prover and verifier message length per-round is at most $T^\delta \cdot S(n) \cdot (\log n)^{O(1/\delta^2)}$.
- The prover runtime is $T(n) \cdot \text{poly}(n) \cdot (\log n)^{O(1/\delta^2)}$.
- The verifier runtime is $T^{O(\delta)} \cdot S(n)^2 \cdot \text{poly}(n) \cdot (\log n)^{O(1/\delta^2)}$.

Setting $\delta = \sqrt{\frac{\log \log n}{\log n}}$ (so that $(\log n)^{O(1/\delta^2)} = \text{poly}(n)$), and $\sigma = \Theta(\sqrt{\log n \cdot \log \log n})$, we have the following corollary.

Corollary 2 (Formal statement of Corollary 1). *Let language $L \in DTISP(T, S)$ where $T(n) = n^{O\left(\left(\frac{\log n}{\log \log n}\right)^{1/2}\right)}$ and $S(n) = \text{poly}(n)$, there exists a UIP with the following properties:*

- The soundness error is $o(1)$.
- The number of rounds is $o(n)$.
- The prover and verifier message length per-round is $S(n) \cdot \text{poly}(n)$.
- The prover runtime is $T(n) \cdot \text{poly}(n)$.
- The verifier runtime is $S(n)^2 \cdot \text{poly}(n)$.

5.1.1 Augmenting Interactive Proofs of Transition of Turing Machines

Our approach is similar to that in [RRR16]. Fix $\mathcal{M}$ to be some Turing machine with $T(n)$ time and $S(n)$ space. For every $t \leq T(n)$, let the language $\mathcal{L}_t$ consist of all tuples $(x, w_1, w_2) \in \{0, 1\}^n \times \{0, 1\}^{O(S(n))} \times \{0, 1\}^{O(S(n))}$ such that $\mathcal{M}$ transitions from configuration w_1 to configuration w_2 in exactly t steps when the input is x .

The crucial observation is captured in the following *Augmentation Lemma* (Lemma 7): let $k = k(n) \in \text{poly}(n)$ be some batch size parameter. The language $\mathcal{L}_{k \cdot t}$ can be expressed as a batch of $\mathcal{L}_t$. Therefore, an efficient batching protocol such as Theorem 6 allows the verifier to verify $(k \cdot t)$ -step transitions by batch-verifying t -step transitions, at the cost of a small increase in complexity. By repeated application of Lemma 7, the verifier gradually increase the transition length t that it can verify, until it reaches $\mathcal{L}_{T(n)}$.

Lemma 7 (Augmentation Lemma). *Let $\mathcal{M}$ be a Turing machine with $T(n)$ time and $S(n) \geq n$ space. Let $t = t(n)$ be some time bound, and let $k = k(n) \in \text{poly}(n)$ be some batch size parameter. Assume that $\mathcal{L}_t$ has an ϵ -sound $(\ell, a, b, \text{Ptime}, \text{Vtime})$ UIP (P, V) . For every $\sigma \in \mathbb{N}$, there exists a protocol ϵ_A -unambiguous $(\ell', a', b', \text{Ptime}', \text{Vtime}')$ UIP (P', V') for $\mathcal{L}_{k \cdot t}$. With $\tilde{O}$ hiding $\text{poly}(\sigma) \cdot \text{polylog}(k, S(n), a, \ell)$ factors, it has the following properties:*

- $\epsilon' = 2 \log k(\epsilon + 2^{-\sigma})$.
- $\ell' = \tilde{O}(\ell \cdot \text{depth}(V) \log \text{size}(V))$.
- $a' = b' = \tilde{O}(a\ell + \text{poly}(\ell)) + O(k \cdot S(n))$.
- $\text{Ptime}' = \tilde{O}(k \cdot \ell \cdot \text{Ptime} + \text{poly}(k, a, \ell, \text{size}(V)))$.
- $\text{Vtime}' = \tilde{O}(\ell \cdot \text{Vtime} + a\ell^2(kS(n) + \text{depth}(V) \log \text{size}(V)) + a^2 \cdot \text{poly}(\ell))$.

Furthermore,

- $\text{size}(V') = \tilde{O}(\ell \cdot \text{size}(V) + a\ell^2(kS(n) + \text{depth}(V) \log \text{size}(V)) + a^2 \cdot \text{poly}(\ell))$.
- $\text{depth}(V') = \text{depth}(V) + \tilde{O}(1)$.

Proof Sketch. The protocol (P', V') on input $(x, w_0, w_{k \cdot t})$ is constructed as follows.

1. P' runs $\mathcal{M}$ on input x and starting configuration w_1 for $k \cdot t$ steps, obtaining all intermediate configurations $w_1, w_2, \dots, w_{k \cdot t}$. Note that if P' already has access to $w_1, w_2, \dots, w_{k \cdot t}$, then it can skip this step.

2. P' sends $w_t, w_{2t}, \dots, w_{kt}$ to $V_{k \cdot t}$.
3. Both parties run the protocol $(P_{\text{Batch-UIP}}, V_{\text{Batch-UIP}})$ to batch (P, V) , with security parameter σ , that batch-verifies $\mathcal{L}_t^{\otimes k}$, on input $\{(x, w_t, w_{2t})\}_{t=0}^k$.
4. The verifier V' accepts iff the batched protocol accepts.

Prescribed Completeness. This is immediate.

ϵ_A -Unambiguity. A deviating prover can either deviate by sending some incorrect configuration w_t , which leads to at least one incorrect t -step transition and hence a false statement in $\mathcal{L}_t$, or by deviating in the batched protocol. In either case, the unambiguity of $(P_{\text{Batch-UIP}}, V_{\text{Batch-UIP}})$ ensures that the verifier accepts with probability at most ϵ_A .

Complexities. The complexities remain the same as in Theorem 6, with the input (n in Theorem 6) being replaced by $O(S)$, except that $a_A = b_A \leq \tilde{O}(al + \text{poly}(\ell)) + O(k \cdot S)$, where the $O(k \cdot S)$ term comes from the prover sending the configurations $w_t, w_{2t}, \dots, w_{kt}$. $\square$

5.1.2 Our Construction of the Doubly-efficient UIP for $\text{DTISP}(T, S)$

For the base case, we have the following (trivial) proof system for verifying one transition. On input (x, w_0, w_1) , the prover and verifier does not communicate, and the verifier simply simulate $\mathcal{M}$ on x with w_0 by one step and verify it actually transitions to w_1 , which requires $O(S(n))$ time.

Proposition 4 (A UIP for verifying one transition). *There exists a $(\ell_0, a_0, b_0, \text{Ptime}_0, \text{Vtime}_0)$ UIP (P_0, V_0) for $\mathcal{L}_1$ with the following properties:*

- The unambiguity error is $\epsilon_0 = 0$.
- The number of rounds is $\ell_0 = 0$.
- The prover message length is $a_0 = 0$.
- The verifier message length is $b_0 = 0$.
- The prover does not do anything, so its runtime is 0.
- The verifier runs in time $O(S(n))$.

Furthermore, the verifier's decision circuit V_0 has the following properties:

- The circuit size is $\text{size}(V_0) = O(S(n))$.
- The circuit depth is $\text{depth}(V_0) = O(1)$.

Using Lemma 7, we can show the following claim by induction.

Proposition 5. *Let $k = k(n) \in \text{poly}(n)$, $\gamma = \gamma(n) \in O\left(\sqrt{\frac{\log n}{\log \log n}}\right)$, $\sigma = \sigma(n) \in \text{polylog}(n)$ be parameters. There exist universal constants C_1, C_2, C_3 such that for each $i \leq \gamma$, there exists a $(\ell_i, a_i, b_i, \text{Ptime}_i, \text{Vtime}_i)$ UIP (P_i, V_i) for $\mathcal{L}_{k^i}$, and for large n , the following holds:*

1. $\max(k, S(n), a_i, \ell_i, \text{size}(\mathbf{V}_i)) = O(n^c)$, where

$$c = \limsup_{n \rightarrow \infty} \lceil 8C_1C_3 + 2\log_n(kS(n)) + 1 \rceil$$

is a constant independent of n and i (since $k, S(n) \in \text{poly}(n)$).

2. Regarding the complexities of $(\mathbf{P}_i, \mathbf{V}_i)$:

- The unambiguity error is $\epsilon_i \leq i \cdot (2\log k)^i 2^{-\sigma}$.
- The number of rounds is $\ell_i \leq \log^{3C_1 i} n$.
- The prover and verifier message length per round is $a_i = b_i \leq kS(n) \cdot \log^{4C_1C_3 i^2} n$.
- The prover runtime is $\text{Ptime}_i = k^i \cdot n^{C_2C_3} \log^{3C_1 i^2} n$.
- The verifier runs in time $\text{Vtime}_i \leq (kS(n))^2 \log^{8C_1C_3 i^2} n$.

Furthermore, the verifier's decision circuit $\mathbf{V}_i$ has the following properties:

- The circuit size is $\text{size}(\mathbf{V}_i) \leq (kS(n))^2 \log^{8C_1C_3 i^2} n$.
- The circuit depth is $\text{depth}(\mathbf{V}_i) \leq i \cdot \log^{C_1} n$.

To simplify our calculation, we use the following loose bounds on the constants involved.

Remark 11 (Determining the constants C_1, C_2, C_3 .) The hidden constants in Lemma 7 determine the constants C_1, C_2, C_3 , as follows.

- Consider the upper bound of all $\text{poly}(\sigma) \cdot \text{polylog}(k, S(n), a_i, \ell_i)$ factors hidden in $\tilde{O}$ of Lemma 7. Assume that indeed $k, S(n), a_i, \ell_i$ are each upper-bounded by n^c for large n , and that $\sigma = \text{polylog}(n)$, then for some constant c_1 and large n , the following holds,

$$\text{poly}(\sigma) \text{polylog}(k, S(n), a_i, \ell_i) \leq (c \log n)^{c_1} \leq (\lceil 8C_1C_3 + 2\log_n(kS(n)) + 1 \rceil \log n)^{c_1}.$$

For large n ,

$$\log n \geq \lceil 8C_1C_3 + 2\log_n(kS(n)) + 1 \rceil,$$

so the above is upper-bounded by $\log^{2c_1} n$, irrespectively of the value of c (or C_1). Therefore, if we take $C_1 := 2c_1 + 1$, then we can replace all occurrence of $\text{poly}(\sigma) \cdot \text{polylog}(k, S(n), a_i, \ell_i)$ with $\log^{C_1} n$.

- There exists a constant c_2 such that the term $\text{poly}(k, a_i, \ell_i, \text{size}(\mathbf{V}_i)) < (ka_i\ell_i\text{size}(\mathbf{V}_i))^{c_2}$ in the Ptime recurrence relation. If $\max(k, S(n), a_i, \ell_i, \text{size}(\mathbf{V}_i)) \leq n^c$ for large n , then $(ka_i\ell_i\text{size}(\mathbf{V}_i))^{c_2} \leq n^{4cc_2}$ for large n .

Therefore, we choose $C_2 := 4cc_2$.

- C_3 is simply defined to be the power of ℓ in the terms $\text{poly}(\ell)$ appearing in the recurrence relations.

Proof of Proposition 5. Let $C_1 := 2c_1 + 1$, $C_2 := 4cc_2$ be the constants as defined in Remark 11. (Recall that $c = \limsup_{n \rightarrow \infty} \lceil 8C_1C_3 + 2\log_n(kS(n)) + 1 \rceil$.)

We prove the following two claims simultaneously by induction on i .

1. $\max(k, S(n), a_i, b_i, \text{size}(\mathbf{V}_i)) \leq n^c$ for large n .
2. The bounds about each of $\epsilon_i, \ell_i, a_i, b_i, \text{Ptime}_i, \text{Vtime}_i, \text{size}(\mathbf{V}_i), \text{depth}(\mathbf{V}_i)$ hold.

The base case, $i = 0$, corresponds to a transition of one step and holds trivially as in Proposition 4.

Inductive Step. Assume that we have a UIP (P_i, V_i) for $\mathcal{L}_{k^i}$ with the stated complexity bounds. Note that an important consequence of $i \leq \gamma \leq \sqrt{\frac{\log n}{\log \log n}}$ is that $i^2 \log \log n < \log n$.

By our second inductive hypothesis, the following hold for large n .

1. $\text{depth}(V_i) \leq i \cdot \log^{C_1} n$.
2. From $\text{size}(V_i) \leq O((kS(n))^2 \log^{8C_1 C_3 i^2} n)$ and $\log(kS(n)) = \log_n(kS(n)) \log n$ we deduce

$$\begin{aligned} \log \text{size}(V_i) &\leq O(1) + 2 \log(kS(n)) + 8C_1 C_3 i^2 \log \log n \\ &= O(1) + (2 \log_n(kS(n)) + 8C_1 C_3 i^2) \log n \\ &< (2 \log_n(kS(n)) + 8C_1 C_3 i^2 + 1) \log n < \log^2 n. \end{aligned}$$

Therefore, for large n , $\text{depth}(V_i) \log \text{size}(V_i) < i \cdot \log^{C_1} n \log^2 n < \log^{C_1+3} n$.

Complexities Applying Lemma 7 with batch size k and security parameter σ , we obtain a UIP (P_{i+1}, V_{i+1}) for $\mathcal{L}_{k^{i+1}}$. As in Remark 11, the first inductive hypothesis implies that all $\tilde{O}(1)$ terms hidden can be upper bounded by $\log^{C_1} n$ when n is large. We work out the recurrence relations to obtain the following.

- $\text{depth}(V_{i+1}) = \text{depth}(V_i) + \tilde{O}(1) \leq \text{depth}(V_i) + \log^{C_1} n \leq (i+1) \cdot \log^{C_1} n$.
- $\ell_{i+1} = \ell_i \cdot (\text{depth}(V_i) \log \text{size}(V_i) \log^{C_1} n) \leq \log^{3C_1 i + (2C_1+3)} n < \log^{3C_1(i+1)} n$, (w.l.o.g. $C_1 > 3$).
- $a_{i+1} = b_{i+1} \leq a_i \cdot \ell_i \cdot \log^{C_1} n + \ell_i^{C_3} \log^{C_1} n + O(k \cdot S(n))$
 $= k \cdot S(n) \cdot (\log^{4C_1 C_3 i^2 + C_1 i + C_1} n + \log^{3C_1 C_3 i + C_1} n + O(1))$
 $< k \cdot S(n) \cdot \log^{4C_1 C_3 (i+1)^2} n$. (w.l.o.g. $C_3 \geq 1$).
- $\text{Vtime}_{i+1} \leq \text{Vtime}_i \cdot (\ell_i \cdot \log^{C_1} n) + (a_i \ell_i^2 (kS(n) + \text{depth}(V_i) \log \text{size}(V_i)) + a_i^2 \ell_i^{C_3}) \log^{C_1} n$
 $\leq (kS(n))^2 \log^{8C_1 C_3 i^2 + 3C_1 i + C_1} n$
 $+ kS(n)(kS(n) + \log^{C_1+3} n) \log^{4C_1 C_3 i^2 + 6C_1 i} n \cdot \log^{C_1} n$
 $+ (kS(n))^2 \cdot \log^{8C_1 C_3 i^2} n \cdot \log^{C_1 C_3 i + C_1} n$
 $< (kS(n))^2 \cdot \log^{8C_1 C_3 (i+1)^2} n$.
- Similarly, $\text{size}(V_{i+1}) < (kS(n))^2 \cdot \log^{8C_1 C_3 (i+1)^2} n$.
- $\text{Ptime}_{i+1} \leq \text{Ptime}_i \cdot (k \cdot \ell_i \cdot \log^{C_1} n) + \text{poly}(k, a_i, \ell_i, \text{size}(V_i))$
 $\leq \text{Ptime}_i \cdot k \log^{3C_1 i^2 + C_1} n \log^{C_1} n + n^{C_2 C_3}$
 $< k^{i+1} n^{C_2 C_3} \log^{3C_1 (i+1)^2 + C_1} n$.

Recall that $c = \lceil \limsup_{n \rightarrow \infty} 8C_1 C_3 + 2 \log_n(kS(n)) + 1 \rceil$. The condition in the induction follows from that $\max(k, S(n), a_{i+1}, b_{i+1}, \text{size}(V_{i+1})) = (kS(n))^2 \cdot \log^{8C_1 C_3 (i+1)^2} n \leq n^c$ for large n , since $(i+1)^2 \leq \gamma^2 \leq \frac{\log n}{\log \log n}$ and $(\log n)^{\frac{\log n}{\log \log n}} = n$.

Finally, the unambiguity error satisfies $\epsilon_{i+1} = 2 \log k(\epsilon_i + 2^{-\sigma}) \leq 2 \log k(i(2 \log k)^i 2^{-\sigma} + 2^{-\sigma}) < (i+1)(2 \log k)^{i+1} 2^{-\sigma}$. $\square$

Let $\delta(n) \in (0, 1)$ be a parameter. Theorem 8 now follows from Proposition 5 by setting $k = T^\delta$ and $\gamma = \frac{1}{\delta}$. In this case, $\mathcal{L}_{k\gamma} = \mathcal{L}$ and it has a UIP with the following complexities:

- $\epsilon_\gamma \leq \gamma(2 \log n)^\gamma 2^{-\sigma} = 2^{-\sigma} \cdot (1/\delta) \cdot (\log n)^{O(1/\delta)}$.
- $\ell_\gamma \leq \log^{3C_1\gamma} n = (\log n)^{O(1/\delta)}$.
- $a_\gamma = b_\gamma \leq k \cdot S(n) \log^{4C_1C_3\gamma^2} n = T^\delta \cdot S(n) \cdot (\log n)^{O(1/\delta^2)}$.
- $\text{Ptime}_\gamma \leq k^\gamma n^{C_2C_3} \log^{3C_1\gamma^2} n = T^\delta \cdot \text{poly}(n) \cdot (\log n)^{O(1/\delta^2)}$.
- $\text{Vtime}_\gamma = (kS(n))^2 \log^{8C_1C_3\gamma^2} n = T^{2\delta} \cdot S(n)^2 \cdot (\log n)^{O(1/\delta^2)}$.

6 IPP for PVAL with Column Distance

The work of [RR20] constructed an IPP (same as Definition 11 without the unambiguous soundness property) for the PVAL problem with respect to the Hamming distance. We use their ideas to construct a UIPP for a slightly different formulation of the PVAL problem which has a different distance measure. The same ideas also result in an instance reduction protocol for PVAL. In particular, we use the column distance Δ_c (Definition 1), as opposed to the Hamming distance.

We impose the additional constraint on our IPP to be a UIP, as it is used as a building block in our batch UIP protocol.²⁰ Thus, we also prove the unambiguous soundness of our protocol. We now present our main result for UIPPs.

Theorem 9 (UIPP for PVAL with column distance; Restatement of Lemma 6). *Suppose $\sigma, d, m, \log L \in \mathbb{N}$, $M = 2^m$, and $\mathbb{F}$ is a field. Let $\tilde{O}$ be omitting $\text{polylog}(|\mathbb{F}|, M, L)$ factors. Suppose $T \geq 8dLm = \tilde{O}(dL)$. Let $\mathbf{j} = (\mathbf{j}_1, \dots, \mathbf{j}_T) \in (\mathbb{F}^{m+\log L})^T$ and $\mathbf{v} = (v_1, \dots, v_T) \in \mathbb{F}^T$.*

There exists a constant C , and a public-coin unambiguous IPP $(P_{\Delta_c RR}, V_{\Delta_c RR})$ for $PVAL(\mathbf{j}, \mathbf{v})$, where both the prover and verifier gets the input $\mathbf{j}, \mathbf{v}$, while the prover additionally gets the input $\mathbf{a} \in \mathbb{F}^{M \times L}$, such that if both parties additionally gets parameters $\sigma, d \in \mathbb{N}$ and $\mathbb{F}$, that satisfy the following preconditions:

- $d \geq 16m\sigma$,
- $|\mathbb{F}| \geq 32 \cdot 2^\sigma (T \log M \log L)^C$,

then the protocol outputs a succinct description $\langle \mathcal{Q} \rangle$ of descriptions $\langle \mathcal{Q} \rangle$ of a set of rows $\mathcal{Q} \subseteq [M]$, and $\langle \Phi' \rangle$ of a predicate Φ' ,

- **Prescribed Completeness:** *If $V_{\Delta_c RR}$ interacts with $P_{\Delta_c RR}$, then $\Phi'(\mathbf{a}[\mathcal{Q}, :]) = 1$ iff $\mathbf{a} \in PVAL(\mathbf{j}, \mathbf{v})$.*
- **$2^{-\sigma-2}$ -Unambiguity:** *Suppose $\Delta_c(PVAL(\mathbf{j}, \mathbf{0})) \geq 4d$, then for any (unbounded) cheating prover strategy P^* that deviates from $P_{\Delta_c RR}$ first in round $r^* \in [\ell_{\Delta_c RR}]$, with probability at least $1 - 2^{-\sigma-2}$ over the verifier's remaining coins, either $V_{\Delta_c RR}$ rejects, or $\Phi'(\mathbf{a}[\mathcal{Q}, :]) = 0$.*

²⁰We note that [RR20] did not require their IPP to be unambiguous. Their protocol, despite being slightly unambiguous, does not have an exponentially small unambiguity error. Given that parallel repetition does not reduce unambiguity error, we have to modify the protocol.

- **Reduced Query:** The subset of rows $\mathcal{Q} \subsetneq [M]$ has size $|\mathcal{Q}| \leq \lceil 8\sigma \cdot \frac{M}{d} \rceil$.
- Regardless of the prover's strategy, the distribution of $\mathcal{Q}$ only depends on the verifier's random coins.

The complexity of the protocol is as follows.

- $\ell_{\Delta_c RR} = \tilde{O}(1)$.
- $a_{\Delta_c RR} = b_{\Delta_c RR} = \tilde{O}(L + \text{poly}(d))$.
- $\text{Ptime}_{\Delta_c RR} = \text{poly}(ML, T \cdot \log |\mathbb{F}|)$.
- $\text{Vtime}_{\Delta_c RR} = \tilde{O}(L + \text{poly}(d))$.

The bit-lengths are $|\langle \mathcal{Q} \rangle| = \tilde{O}(\text{poly}(d))$ and $|\langle \Phi' \rangle| = \tilde{O}(L + \text{poly}(d))$. Let $G_{\mathcal{Q}}(i, \langle \mathcal{Q} \rangle)$ be the circuit that returns the i -th element in $\mathcal{Q}$, and $C'(\mathbf{a}, \langle \Phi' \rangle)$ be the circuit that computes $\Phi'(\mathbf{a})$, then they satisfy the following.

- $\text{size}(G_{\mathcal{Q}}) = \tilde{O}(\text{poly}(d))$.
- $\text{depth}(G_{\mathcal{Q}}) = \tilde{O}(1)$.
- $\text{size}(C') = \tilde{O}(|\mathcal{Q}| \cdot L)$.
- $\text{depth}(C') = \tilde{O}(1)$.

The verifier's verdict circuit (which outputs 0 iff it rejects amidst the protocol) satisfies

- $\text{size}(V_{\Delta_c RR}) = \tilde{O}(L + \text{poly}(d))$.
- $\text{depth}(V_{\Delta_c RR}) = \tilde{O}(1)$.

6.1 Overview of the RR IPP for PVAL.

The setup for the RR IPP is as follows:

- The verifier has oracle access to a function $\mathbf{a} : \{0, 1\}^m \mapsto \mathbb{F}$ (note the notation $\mathbf{a}$ often refers to the truth table of the function which is a string in $\mathbb{F}^{2^m}$), where $\mathbb{F}$ is a field of characteristic 2. The verifier also has elements $\mathbf{j} \in (\mathbb{F}^m)^T$ and $\mathbf{v} \in \mathbb{F}^T$.
- $\hat{\mathbf{a}} : \mathbb{F}^m \mapsto \mathbb{F}$ is the multilinear extension of $\mathbf{a}$.
- The prover and verifier want to run a UIP protocol to verify whether $\hat{\mathbf{a}}(\mathbf{j}) = \mathbf{v}$ or if $\hat{\mathbf{a}}$ is δ far from $\text{PVAL}(\mathbf{j}, \mathbf{v})$ by making only $\tilde{O}(\frac{1}{\delta})$ oracle queries to the function $\mathbf{a}$.

Since the verifier only cares about the case when the distance is at least δ , it could hope to query $\mathbf{a}$ on just $O(\frac{1}{\delta})$ uniformly random points to get a false claim. Unfortunately, the verifier does not begin with a claim about any subset of $\frac{1}{\delta}$ coordinates, but only global constraints imposed by $\mathbf{j}$ and $\mathbf{v}$.

The original idea for getting around this difficulty was given in [RVW13]. Their idea was to break $\mathbf{a}$ into two parts: $\mathbf{a}_0, \mathbf{a}_1$, each corresponding to a function in $\{0, 1\}^{m-1} \rightarrow \mathbb{F}$ such that $\mathbf{a}_b(x) = \mathbf{a}(b, x)$ for $b \in \{0, 1\}$, and thus

$$\mathbf{a}(x_1, \vec{x}) = (1 - x_1)\mathbf{a}_0(\vec{x}) + x_1\mathbf{a}_1(\vec{x}).$$

This, together with the fact that $\hat{\mathbf{a}}$ is the multilinear extension of $\mathbf{a}$, implies that

$$\hat{\mathbf{a}}(x_1, \vec{x}) = (1 - x_1)\hat{\mathbf{a}}_0(\vec{x}) + x_1\hat{\mathbf{a}}_1(\vec{x})$$

Exploiting this observation, one can ask the prover to reduce the PVAL claims about $\hat{\mathbf{a}}$ to $\hat{\mathbf{a}}_0$ and $\hat{\mathbf{a}}_1$. Assuming $\mathbf{a}$ was far from the PVAL set, regardless of the prover's answers, we would expect a random linear combination $\mathbf{a}_0 + c\mathbf{a}_1$ to be far from the resulting PVAL claim. Let's try to understand how far from the claim we can expect $\mathbf{a}_0 + c\mathbf{a}_1$ to be.

We know that the prover needs to change $\mathbf{a}$ in $\delta \cdot 2^m$ points to satisfy the PVAL instance. Thus, the prover could answer such that it would need to change the answers in both $\mathbf{a}_0$ and $\mathbf{a}_1$ in $\delta \cdot 2^{m-1}$ positions each. Therefore, the prover may need to change $\mathbf{a}_0 + c\mathbf{a}_1$ in only $\delta \cdot 2^{m-1}$ points to satisfy the PVAL instance if the error positions are aligned. Thus, even if the error positions are aligned, we could hope that $\mathbf{a}_0 + c\mathbf{a}_1$ remains $(\delta 2^{m-1})/2^{m-1} = \delta$ -far from satisfying the resulting PVAL claim.

[RVW13] proved that with high probability, $\mathbf{a}_0 + c\mathbf{a}_1$ is at least $\frac{\delta}{2}$ -far from the resulting PVAL since they did not have any assumption that the base PVAL sets themselves have any distance. In contrast, [BKS18] got improved upon [RVW13] by additionally assuming that every two elements in the PVAL set are also far apart. This allowed them to associate a unique point in the set PVAL to which $\mathbf{a}_0 + c\mathbf{a}_1$ is close.

Now we have the guarantee that the reduced instance has δ -distance from the PVAL set. However, observe that if we want to query $\mathbf{a}_0 + c\mathbf{a}_1$ on one point then we need to query $\mathbf{a}$ on two points. Even though the “instance size” is halved, the “effective instance size”, measured by the number of needed queries to $\mathbf{a}$, remains the same. Specifically, if we have to query $\mathbf{a}_0 + c\mathbf{a}_1$ on q points, then we need to query $\mathbf{a}$ on $2q$ points. To make progress, we would in fact like to *double* the fractional distance, i.e. maintain the absolute number of error positions in the new claim. If this is achieved, the number of verifier queries to $\mathbf{a}_0 + c\mathbf{a}_1$ can be halved while retaining the desired soundness. Through this, the number of queries to $\mathbf{a}$ remains the same, and the PVAL constraint is indeed only over the folded instance, which is half the original size.

The above construction is indeed quite lossy since 1 error each in $\mathbf{a}_0$ and $\mathbf{a}_1$ can add up to just count as 1 error in $\mathbf{a}_0 + c\mathbf{a}_1$ when they align in the truth tables. Thus, [RR20] had the idea of instead reducing not to $\mathbf{a}_0 + c\mathbf{a}_1$ but to $\mathbf{a}_0 + c(\mathbf{a}_1 \circ \pi)$ where π is a sufficiently random permutation. This step distributes the errors, and $\mathbf{a}_0 + c(\mathbf{a}_1 \circ \pi)$ is indeed about 2δ -far from the resulting instance.

In the next subsection, we show that this approach is compatible with not just the Hamming distance but also column distance. Essentially, this is the same as running L instances of RR in parallel with the same randomness and with a PVAL claim over the L instances together. Additionally, since we want to use this protocol as a sub-protocol of a UIP, we prove that the resulting protocol is unambiguous. This requires $\mathbb{F}$ to be a large enough field, which is without loss of generality since we can always just replace $\mathbb{F}$ with a sufficiently large extension.

6.2 Gap Amplification Lemmas for Column Distance

In the section above, we gave an overview of how [RR20] used ideas from [RVW13] and [BKS18]. These results are actively studied in the SNARKs literature [BSCI⁺23, AHIV23, ACFY25, Zei24, ACFY24] and are known as “divergence” and “correlated-agreement” theorems. We prove the Δ_c analogues of [RVW13] and [BKS18] here now.

We begin by giving an analogous version of the [RVW13] batch amplification lemma. The following lemma only shows that if $\mathbf{a}_0$ and $\mathbf{a}_1$ are δ -far from vector spaces V then $\mathbf{a}_0 + c\mathbf{a}_1$ is at least $\delta/2$ -far from V with high probability.²¹ This result might seem suboptimal, but it is important to note that we do not have any assumption on the distance of the space V itself.

Lemma 8 ([RVW13] gap amplification for Δ_c). *Let $\mathbb{F}$ be a field of characteristic 2 and integers $m, \log L, T > 0$. Let $V \subset \mathbb{F}^{\{0,1\}^{m+\log L}}$ be a non-empty $\mathbb{F}$ -linear vector space, and $\mathbf{a}_0, \mathbf{a}_1 : \{0,1\}^{m+\log L} \rightarrow \mathbb{F}$ be functions. Suppose*

$$d_i = \Delta_c(\mathbf{a}_i, V).$$

Then,

$$\Pr_{c \leftarrow \mathbb{F}^*} [\Delta_c(\mathbf{a}_0 + c\mathbf{a}_1, V) < \max(d_0, d_1)/2] \leq \frac{1}{|\mathbb{F}| - 1}.$$

Proof. We emulate the proof given in [RVW13]. For the sake of contradiction assume that

$$\Pr[\Delta_c(\mathbf{a}_0 + c\mathbf{a}_1, V) < d_0/2] > \frac{1}{|\mathbb{F}|}.$$

Call a c *bad* if $\Delta_c(\mathbf{a}_0 + c\mathbf{a}_1, V) < d_0/2$. Now, if the theorem is false, then there exist 2 bad values of c , let these be c' and c'' .

We would like to use triangle inequality to say that if $g' = \mathbf{a}_0 + c'\mathbf{a}_1$, $g'' = \mathbf{a}_0 + c''\mathbf{a}_1$ are both less than $\frac{d_0}{2}$ far from V then their linear combinations are less than δ -far as well. In particular, for any $r \in \mathbb{F}$, we have that

$$\Delta_c(r \cdot g' + (1-r) \cdot g'', V) \leq \Delta_c(g, V) + \Delta_c(g', V) < d_0$$

Now, if $rc' + (1-r)c'' = 0$ then $r \cdot g' + (1-r) \cdot g'' = \mathbf{a}_0$ which would imply that $\Delta_c(\mathbf{a}_0, V) < d_0$ which is a contradiction. Thus, there is no $r \in \mathbb{F}$ such that $r(c' - c'') = -c''$. But, then we must have $c' = c''$ and there is at most one *bad* value of c as desired. $\square$

6.2.1 Correlated-Agreement under Column Distance

[BKS18] improves the [RVW13] analysis by observing that if we have an additional assumption about distance on the base vector space V then we could get a better result. Their ideas also provide an improved gap amplification lemma which served as the point upon which [RR20] built and improved. We prove the version of [BKS18] in our Δ_c setting. We later use this directly when we prove the RR gap amplification lemma over Δ_c distance.

[BKS18] observed that if we could assume that if $\mathbf{a}_0$ is close to v_0 and $\mathbf{a}_1$ is close to v_1 where $v_0, v_1 \in V$ where V is a vector space then we could assume that the closest point to $\mathbf{a}_0 + c\mathbf{a}_1$ is

²¹Note that absolute distances are used in place of the relative forms.

$v_0 + cv_1$ as long as V has large distance. As soon as we have such an assumption, we can immediately improve the [RVW13] analysis.

Similarly to what is done in [RR20], instead of stating the main [BKS18] result on gap amplification directly, we will first prove a “correlated-agreement” lemma. Since the plan is to fold the functions as $\mathbf{a}_0 + c(\mathbf{a}_1 \circ \pi)$, we need to understand how the errors in the sum look more carefully.

[BKS18] observed that if V is some vector space with large distance and $\mathbf{a}_0, \mathbf{a}_1$ are different points in the ambient space of V , then if many points along the line $\mathbf{a}_0 + c\mathbf{a}_1$ are close to V then not only are both $\mathbf{a}_0$ and $\mathbf{a}_1$ close to V but the number of coordinates that both of them simultaneously agree with elements of V on is also large.

Lemma 9 (Correlated Agreement for Δ_c). *Let $V \subseteq \mathbb{F}^{\{0,1\}^{m+\log L}}$ be a non-empty vector space over $\mathbb{F}$ and $\Delta_c(V) = \Upsilon$. For every $\varepsilon, \delta > 0$ and $\mathbf{a}_0, \mathbf{a}_1 \in \mathbb{F}^{\{0,1\}^{m+\log L}}$ such that*

- $\delta(1 - \varepsilon) \leq \Upsilon/3$
- $|\{c \mid \Delta_c(\mathbf{a}_0 + c\mathbf{a}_1, V) < \delta(1 - \varepsilon)\}| > 1/\varepsilon$

there exist $v_0, v_1 \in V$ such that for all $t \in \{0, 1\}^{\log L}$,

$$|\{i \mid (\mathbf{a}_0[i, t] = v_0[i, t]) \wedge (\mathbf{a}_1[i, t] = v_1[i, t])\}| \geq (1 - \delta)2^m$$

Proof. Let $S \subseteq \{c \mid \Delta_c(\mathbf{a}_0 + c\mathbf{a}_1, V) < \delta(1 - \varepsilon)\}$ such that $|S| = 1 + \lfloor 1/\varepsilon \rfloor$. Now, for each $c \in S$, let $v^{(c)} \in V$ be such that $\Delta_c(\mathbf{a}_0 + c\mathbf{a}_1, v^{(c)}) < \delta(1 - \varepsilon)$.

Observe that now if all these points $(c, v^{(c)})$ lied on a line then we can get the result. In particular, we can write $v^{(c)} = v_0 + cv_1$ for some fixed v_0, v_1 independent of c . And, we would get

$$\Delta_c(\mathbf{a}_0 + c\mathbf{a}_1, v_0 + cv_1) < \delta(1 - \varepsilon)$$

Rearranging, we get for every $c \in S$,

$$\Delta_c(\mathbf{a}_0 - v_0, c(v_1 - \mathbf{a}_1)) < \delta(1 - \varepsilon)$$

Now, observe that for any $i \in \{0, 1\}^m, t \in \{0, 1\}^{\log L}$ if $\mathbf{a}_0[i, t] \neq v_0[i, t]$ then there is at most one value of c for which $\mathbf{a}_0[i, t] - v_0[i, t] = c(v_1[i, t] - \mathbf{a}_1[i, t])$ and similarly if $\mathbf{a}_1[i, t] \neq v_1[i, t]$.

Thus, fix any $t \in \{0, 1\}^{\log L}$ and let

$$T_t = \{i \mid (\mathbf{a}_0[i, t] = v_0[i, t]) \wedge (\mathbf{a}_1[i, t] = v_1[i, t])\}$$

we get that there exists a $c \in S$ such that

$$\frac{|\{i \mid (\mathbf{a}_0[i, t] - v_0[i, t]) \neq c(v_1[i, t] - \mathbf{a}_1[i, t])\}|}{2^m} \geq \left(1 - \frac{|T_t|}{2^m}\right) \left(1 - \frac{1}{|S|}\right)$$

However, since, for any T , we have that:

$$\Delta_c(\mathbf{a}_0 - v_0, c(v_1 - \mathbf{a}_1)) \geq \frac{|\{i \mid (\mathbf{a}_0[i, t] - v_0[i, t]) \neq c(v_1[i, t] - \mathbf{a}_1[i, t])\}|}{2^m}$$

Thus,

$$\left(1 - \frac{|T|}{2^m}\right) \left(1 - \frac{1}{|S|}\right) \leq \delta(1 - \varepsilon) \implies 1 - \frac{|T|}{2^m} < \delta \implies 2^m(1 - \delta) < |T|$$

The final result is since $|S| > \frac{1}{\epsilon}$.

Thus, we just need to show that indeed the $(c, v^{(c)})$ lie on a line. This is feasible to show because we additionally assume that V has large distance. We will now use the same triangle inequality trick to show that if all $v^{(c)}$ are not on a line then there are two distinct points of V which are $< \Upsilon$ far from each other.

Claim. All the points $(c, v^{(c)})$ lie on a line.

Proof. Observe that the above claim follows trivially when $\epsilon > 1/2$ since then $|S| = 2$.

We are interested in the regime when $\epsilon \leq 1/2$. Consider any $c_1, c_2, c_3 \in S$. Then, we can take a linear combination of $(\mathbf{a}_0 + c_1 \mathbf{a}_1)$ and $(\mathbf{a}_0 + c_2 \mathbf{a}_1)$ to get $(\mathbf{a}_0 + c_3 \mathbf{a}_1)$. Observe that

$$\Delta_c \left(\frac{(c_2 - c_3)(\mathbf{a}_0 + c_1 \mathbf{a}_1) + (c_3 - c_1)(\mathbf{a}_0 + c_2 \mathbf{a}_1)}{c_2 - c_1}, \frac{(c_2 - c_3)v^{(c_1)} + (c_3 - c_1)v^{(c_2)}}{c_2 - c_1} \right) < 2\delta(1 - \epsilon)$$

Let $v' = \frac{(c_2 - c_3)v^{(c_1)} + (c_3 - c_1)v^{(c_2)}}{c_2 - c_1}$. Thus, by the above we have

$$\Delta_c(v', \mathbf{a}_0 + c_3 \mathbf{a}_1) < 2\delta(1 - \epsilon)$$

and $\Delta_c(v^{(c_3)}, \mathbf{a}_0 + c_3 \mathbf{a}_1) < \delta(1 - \epsilon)$. Finally, by the triangle inequality, we have that:

$$\Delta_c(v', v^{(c_3)}) < 3\delta(1 - \epsilon) \leq \Upsilon \implies v' = v^{(c_3)}.$$

Hence, $v' = v^{(c_3)}$ so $v^{(c_1)}, v^{(c_2)}$, and $v^{(c_3)}$ are indeed collinear as desired. $\square$

$\square$

6.2.2 An RR-Style Gap Amplification

Now, we are ready to prove the RR style gap amplification lemma for Δ_c .

[RR20]'s idea is to amplify distance using not just $\mathbf{a}_0 + c\mathbf{a}_1$ as in [RVW13] but instead to work with $\mathbf{a}_0 + c(\mathbf{a}_1 \circ \pi)$ where π is an efficient to compute permutation. If we assume that $\pi \sim \Pi$ is η -close to being d -wise independent then we can hope that the errors in $\mathbf{a}_0$ and $\mathbf{a}_1 \circ \pi$ do not align with very high probability. The issue here is that even though we have that $\Delta_c(\mathbf{a}_0 \| \mathbf{a}_1, \text{PVAL}(j, v)) \geq \delta$, but that does not translate into any immediate claim about distance of $\mathbf{a}_1 \circ \pi$ with any immediate new PVAL claim.

Thus, we need to add a GKR based protocol first to create claims about $\text{PVAL}(\mathbf{a}_1 \circ \tau)$ where $\tau = \pi \times I^{22}$. The permutation π is sampled as $\pi \sim \Pi$ where Π is an η -almost d -wise independent distribution as defined in Theorem 4.

Definition 14. $\Phi_{\pi, j, v}$ is the circuit that takes in an input $\mathbf{a}$ and computes whether

$$\left(\widehat{\mathbf{a} \circ (\pi \times I)} \right) (j) = v$$

Here, we assume that π has a succinct description.

Lemma 10 (Complexity of $\Phi_{\pi, j, v}$). *The circuit $\Phi_{\pi, j, v}$ has size $\tilde{O}(T \cdot (2^{m+\log L}) \cdot \log |\mathbb{F}|)$ and depth $\text{poly}(m + \log L, \log(T \log |\mathbb{F}|))$.*

²²It implements π in each column of the matrix as in Figure 1c.

Proof. The permutation is only 1 layer of the circuit and size linear in the input. Beyond that, it is verifying T distinct multilinear evaluations which can be evaluated by simply taking the sum of the indicator variables. Thus, the result follows. $\square$

Now, consider the following sub-protocol:

Protocol 5 (P_Π, V_Π): GKR to get claims about $\mathbf{a} \circ \tau$.

Input Parameters: $m, L, T \in \mathbb{N}$ and $\mathbb{F}$ is a field.

Input: $\pi \in \Pi$, a permutation, $\mathbf{j} \in (\mathbb{F}^{m+\log L})^T$ and $\mathbf{v} \in \mathbb{F}^T$.

Prover Auxiliary Input: $\mathbf{a} \in \{0, 1\}^{m \times L} \rightarrow \mathbb{F}$.

Ingredients:

- The GKR protocol ($P_{\text{GKR}}, V_{\text{GKR}}$) (Theorem 3) for checking $\Phi(\mathbf{a}) = 1$.

The prover and verifier gets the description $\langle \Phi \rangle$, while the prover additionally gets the input $\mathbf{a}$. Both parties additionally takes in the parameter $T \in \mathbb{N}$, and in the end outputs $\mathbf{j} \in (\mathbb{F}^{m+\log L})^T, \mathbf{v} \in \mathbb{F}^T$.

Output: $\mathbf{j}' \in (\mathbb{F}^{m+\log L})^T, \mathbf{v}' \in \mathbb{F}^T$.

- 1: Let $\tau = \pi \times I$. Let $\langle \Phi_{\pi, \mathbf{j}, \mathbf{v}} \rangle$ be the description of the circuit $\Phi_{\pi, \mathbf{j}, \mathbf{v}}$ that takes in a function $\mathbf{a}$ and checks whether $(\mathbf{a} \circ \tau)(\mathbf{j}) = \mathbf{v}$.
 - 2: Run the GKR protocol T times on $\langle \Phi_{\pi, \mathbf{j}, \mathbf{v}} \rangle$ and prover auxiliary input $\mathbf{a}$ and output the resulting $\mathbf{j}', \mathbf{v}'$.
-

We can directly use this sub-protocol to create claims about $\mathbf{a}_1 \circ \tau$ and thus use [RR20]'s idea of folding the function $\mathbf{a}$ as $\mathbf{a}_0 + c(\mathbf{a}_1 \circ \tau)$.

Lemma 11 (RR gap amplification for Δ_c). *Let $m, \log L, T, \sigma, d > 0$ be integers and $0 < 4d \leq 2^{m/50}$. Let $\mathbb{F}$ be a field of characteristic 2. For $\mathbf{j} \in (\mathbb{F}^{m+\log L})^T$ suppose that has minimum column distance $\geq 4d$. Let $\mathbf{v}_0, \mathbf{v}_1 \in \mathbb{F}^T$ be vectors such that both $PVAL(\mathbf{j}, \mathbf{v}_0)$ and $PVAL(\mathbf{j}, \mathbf{v}_1)$ are non-empty. Let $\mathbf{a}_0, \mathbf{a}_1 : \{0, 1\}^{m+\log L} \rightarrow \mathbb{F}$ be functions. Suppose*

$$d \leq \Delta_c \left(\begin{bmatrix} \mathbf{a}_0 \\ \mathbf{a}_1 \end{bmatrix}, \begin{bmatrix} PVAL(\mathbf{j}, \mathbf{v}_0) \\ PVAL(\mathbf{j}, \mathbf{v}_1) \end{bmatrix} \right)$$

Consider a uniformly random $\pi \in \Pi_{m, 4d, 2^{-\sigma}}$ where Π is the set of permutations on $\{0, 1\}^m$ defined in Section 3.4. Let $\tau = \pi \times I$ and $\mathbf{j}' \in (\mathbb{F}^{m+\log L})^T, \mathbf{v}' \in \mathbb{F}^T$ such that $\mathbf{j}', \mathbf{v}'$ is outputted by the Protocol 5 protocol on input $\mathbf{a}_1 \circ \tau, \pi^{-1}, \mathbf{j}, \mathbf{v}_1$ and the same additional parameters i.e. $m, \log L, T$.

Now, for $c \leftarrow \mathbb{F}$ define:

$$g = \mathbf{a}_0 + c(\mathbf{a}_1 \circ \tau).$$

Let $S_\pi \subseteq \mathbb{F}^{2t}$ be the set of pairs of vectors (u_0, u_1) such that the sets $PVAL((\mathbf{j}, \mathbf{j}'), (u_0, u_1))$ and $PVAL((\mathbf{j}, \mathbf{j}'), (u_1, \mathbf{v}'))$ are not empty. For $(u_0, u_1) \in S_\pi$, define

$$\delta_{\pi, \mathbf{j}', \mathbf{v}', (u_0, u_1), c} = \Delta_c(g, PVAL((\mathbf{j}, \mathbf{j}'), (u_0, u_1) + c(u_1, \mathbf{v}')))$$

Then for every $\gamma < (2/d, 1)$ and $\epsilon \in (0, 1)$, taking

$$\delta' = d(1 - \gamma)(1 - \epsilon),$$

we have that

$$\begin{aligned} \Pr_{\pi \leftarrow \Pi, \text{Protocol 5}} \left[\exists (\mathbf{u}_0, \mathbf{u}_1) \in S_\pi \text{ s.t. } \Pr_{c_0, c_1 \leftarrow \mathbb{F}} [\delta_{\pi, (u_0, u_1), c, \mathbf{j}'} < \delta'] > \frac{1}{\epsilon |\mathbb{F}|} + \frac{1}{|\mathbb{F}|} \right] \\ < 2^{-\sigma} + \exp(-\gamma d/6) + \text{err}_{\text{GKR}} \end{aligned}$$

where err_{GKR} refers to the failure probability of Protocol 5 as described in Lemma 2. Thus, we have that

$$\text{err}_{GKR} \leq \left(\epsilon_{GKR}(D, \log S, |\mathbb{F}|) + (\epsilon_{GKR}(D, \log S, |\mathbb{F}|))^T \cdot \left(\binom{M}{4d} |\mathbb{F}|^{4d} \right)^L \right)$$

where D and S are the depth and size of $\Phi_{\pi, j, v}$ respectively.

Proof. For the sake of contradiction, assume that the lemma is false.

Observe that for any $\pi \in \Pi$, with probability $\geq 1 - \text{err}_{GKR}$, by Lemma 2 and Theorem 3, there exists a $g \in \text{PVAL}(\mathbf{j}', v') \cap B_{4d, \mathbb{F}}(\mathbf{a}_1 \circ \tau)$ iff the prover was honest with implicit input g and $(g \circ \tau^{-1}) \in \text{PVAL}(\mathbf{j}, v_1)$.

Thus, we condition on this event happening and that accounts for the err_{GKR} term in the lemma by union bound.

Let $\bar{\delta} = \frac{\delta'}{1-\epsilon} < d < \frac{4d}{3}$. Now, with probability $> p = 2^{-\sigma} + \exp(-\gamma \cdot d/6)$ over the choice of π , there exists $(u_0, u_1) \in S_\pi$ and $c_0 \neq 0$, such that

$$\Pr_{c_1 \leftarrow \mathbb{F}} [\delta_{\pi, (u_0, u_1), c, \mathbf{j}'} < \bar{\delta}(1-\epsilon)] > \frac{1}{\epsilon |\mathbb{F}|}$$

Since $\bar{\delta} < 4d/3$, with probability $> p$, over the choice of $\pi \in \Pi$, let the choice of $(u_0, u_1) \in S_\pi$ and $c \neq 0$ be as above. Thus, let $V = \text{PVAL}((\mathbf{j}, \mathbf{j}'), 0)$ and observe that $\Delta_c(V) \geq 4d$. Let

$$h_0 \in \text{PVAL}((\mathbf{j}, \mathbf{j}'), (v_0, u_0)), \quad h_1 \in \text{PVAL}((\mathbf{j}, \mathbf{j}'), (u_1, v')),$$

then

$$\Pr_{c \leftarrow \mathbb{F}} [\Delta_c(V + h_0 + c(V + h_1), \mathbf{a}_0 + c(\mathbf{a}_1 \circ \tau)) < \bar{\delta}(1-\epsilon)] > \frac{1}{\epsilon |\mathbb{F}|},$$

and thus

$$\Pr_{c \leftarrow \mathbb{F}} [\Delta_c(V, (\mathbf{a}_0 - h_0) + c(\mathbf{a}_1 \circ \tau - h_1)) < \bar{\delta}(1-\epsilon)] > \frac{1}{\epsilon |\mathbb{F}|}.$$

Letting $u^* = \mathbf{a}_0 - h_0$ and $u = \mathbf{a}_1 \circ \tau - h_1$ in Lemma 9, we get that there exist $y_0, y_1 \in V$ such that for all $r \in \{0, 1\}^{\log L}$

$$|\{i \in \{0, 1\}^m \mid (\mathbf{a}_0(i, r) - h_0(i, r) = y_0(i, r)) \wedge (\mathbf{a}_1 \circ \tau(i, r) - h_1(i, r) = y_1(i, r))\}| \geq 2^m - \bar{\delta},$$

or equivalently for all $r \in \{0, 1\}^{\log L}$,

$$|\{i \in \{0, 1\}^m \mid (\mathbf{a}_0(i, r) = (y_0 + h_0)(i, r)) \wedge (\mathbf{a}_1 \circ \tau(i, r) = (y_1 + h_1)(i, r))\}| \geq 2^m - \bar{\delta}$$

Now, observe that since $(y_0 + h_0) \in \text{PVAL}(\mathbf{j}, v_0)$, $(y_1 + h_1) \in \text{PVAL}(\mathbf{j}', v')$, and $\Delta_c(\text{PVAL}((\mathbf{j}, \mathbf{j}'), 0)) > 3\bar{\delta}$, we have that there is a unique element in $z_0 \in \text{PVAL}(\mathbf{j}, v_0)$ such that $\Delta_c(\mathbf{a}_0, z_0) \leq \bar{\delta}$ and similarly, there is a unique element $z_1 \in \text{PVAL}(\mathbf{j}', v')$ such that $\Delta_c(\mathbf{a}_1 \circ \tau, z_1) \leq \bar{\delta}$. Now, observe that since $z_1 \in \text{PVAL}(\mathbf{j}', v')$ and $\Delta_c(\mathbf{a}_1 \circ \tau, z_1) \leq 4d$, we must have that the prover was honest with implicit input z_1 in mind and $z_1 \circ \tau^{-1} \in \text{PVAL}(\mathbf{j}, v_1)$. Observe that $\Delta_c(\text{PVAL}(\mathbf{j}, v_1)) \geq 4d$. Thus, $z'_1 = z_1 \circ \tau^{-1}$ is the unique element such that $z'_1 \in \text{PVAL}(\mathbf{j}, v_1)$ and $\Delta_c(\mathbf{a}_1, z'_1) < \bar{\delta}$.

Plugging, this back in, we get that with probability $\geq p$ over the choice of $\pi \in \Pi$, we have that for all $r \in \{0, 1\}^{\log L}$

$$|\{i \in \{0, 1\}^m \mid (\mathbf{a}_0(i, r) = z_0(i, r)) \wedge (\mathbf{a}_1 \circ \tau(i, r) = z'_1 \circ \tau(i, r))\}| > 2^m - \bar{\delta}$$

Since $\Delta_c(\mathbf{a}_0 \parallel \mathbf{a}_1, z_0 \parallel z'_1) \geq d$, for each $r \in \{0, 1\}^{\log L}$, let

$$\delta_{r,0} = \Delta_c(\mathbf{a}_0(\cdot, r), z_0(\cdot, r)) \quad \delta_{r,1} = \Delta_c(\mathbf{a}_1(\cdot, r), z'_1(\cdot, r)).$$

By our assumption, we know that $\max_{r \in \{0,1\}^m} (\delta_{r,0} + \delta_{r,1}) = \Delta_c(\mathbf{a}_0 \parallel \mathbf{a}_1, z_0 \parallel z'_1) \geq d$.

Now, we fix our focus to a fixed $r \in \{0, 1\}^{\log L}$ such that $\delta_{r,0} + \delta_{r,1} \geq d$. Now, for this r , with probability $\geq p$ over the choice of π , we have that

$$|\{i \in \{0, 1\}^m \mid (\mathbf{a}_0(i, r) = z_0) \wedge (\mathbf{a}_1 \circ \tau(i, r) = z'_1 \circ \tau(i, r))\}| \geq 2^m - \bar{\delta}$$

Let

$$E_0 = \{i \mid (\mathbf{a}_0(i, r) \neq z_0(i, r))\} \quad E_1 = \{i \mid (\mathbf{a}_1(i, r) \neq z'_1(i, r))\}.$$

We know $|E_0| + |E_1| \geq d$. Thus, we can take subsets $A_0 \subseteq E_0$ and $E_1 \subseteq A_1$ such that $|A_0| + |A_1| \geq d$ and $|A_0|, |A_1| \leq d$. Therefore, probability $\geq p$ over the choice of π , we have that $|A_0 \cap \pi A_1| \geq d - (1 - \gamma)d \geq \gamma d$.

Since $\pi \sim \Pi_{m,4d,2^{-\sigma}}$ and $p = 2^{-\sigma} + \exp(-\gamma d/6)$, this would contradict Theorem 5. Thus, we have a contradiction. $\square$

6.3 Efficient IPP for PVAL with Column Distance

Now, we give an analogous efficient Interactive proof of proximity (IPP) for PVAL to [RR20] in the column distance setting. The key differences here from their work is that we add a soundness parameter as we want the protocol to be unambiguous, and our distance notion is different. Besides these differences which mostly just affect our choices of parameters, the protocol is essentially the same as RR.

For the applications that we use the protocol, we do not require it to be an IPP. We mention the IPP here for completeness and historical reasons.

Before presenting the protocol, we go over the ideas in the protocol first:

- We begin with a function $\mathbf{a} : \{0, 1\}^{m+\log L} \rightarrow \mathbb{F}$ and a $\text{PVAL}(\mathbf{j}, \mathbf{v})$ set. The verifier would like to check whether $\mathbf{a} \in \text{PVAL}(\mathbf{j}, \mathbf{v})$ or if $\Delta_c(\mathbf{a}, \text{PVAL}(\mathbf{j}, \mathbf{v})) > d$. The verifier would also like the final check to be efficient and succinctly described so that it could be further delegated.
- **Breaking into Multilinear Functions:** The verifier asks the prover to break $\mathbf{a}$ into multilinear functions $\mathbf{a}_0, \mathbf{a}_1 : \{0, 1\}^{m+\log L-1} \mapsto \mathbb{F}$ such that $\hat{\mathbf{a}}(\chi, x) = (1 - \chi)\hat{\mathbf{a}}_0(x) + \chi\hat{\mathbf{a}}_1(x)$ for $\chi \in \mathbb{F}$. This is possible due to the multilinearity of $\hat{\mathbf{a}}$. The verifier then asks the prover to break $j_i = (\chi_i, j'_i)$ and send $\mathbf{a}_0(j'_i)$ and $\mathbf{a}_1(j'_i)$. Let the resulting PVAL instances be $\text{PVAL}(\mathbf{j}', \mathbf{v}_0)$ and $\text{PVAL}(\mathbf{j}', \mathbf{v}_1)$.

The verifier can now use the linearity to check if these claims are consistent.

- **Permuting Errors:** Now, the prover uses Protocol 5 to convert the claim $\mathbf{a}_1 \in \text{PVAL}(\mathbf{j}', \mathbf{v}_1)$ into a new claim that $\mathbf{a}_1 \circ \tau \in \text{PVAL}(\mathbf{j}', \mathbf{v}')$.
- **Filling in Values and emptiness checks:** The verifier would want to create now a new PVAL claim about $\mathbf{a}_0 + c(\mathbf{a}_1 \circ \tau)$ but does not know $\hat{\mathbf{a}}_0(\mathbf{j}')$ or $\widehat{\mathbf{a}_1 \circ \tau}(\mathbf{j}')$. Thus, the verifier simply asks the prover to fill in these values. This is all compatible with the Lemma 11 gap amplification that we have proved above.

- **Curve fitting:** Assuming all this goes well, the verifier has reduced to a claim about a function $\mathbf{a}' : \{0,1\}^{m+\log L-1} \mapsto \mathbb{F}$ but the resulting PVAL instance now has $2T$ evaluations instead of T . Thus, the verifier can ask the prover to construct a low degree curve through j' and j' and ask the prover for values of $\mathbf{a}_0$ and $\mathbf{a}_1 \circ \tau$ on all the points on the curve.
- **Reducing number of evaluation points:** Now, the verifier just picks T uniformly random points on the curve. Since j' was uniformly random, these T points also end up being uniformly random elements of $\mathbb{F}^{m+\log L-1}$ and the distance guarantee continues to hold with high probability.

The protocol (Protocol 10) is essentially just the above with some parameter balancing to get the appropriate soundness error. We need to do all the above over larger fields than [RR20] did since we also prove unambiguity of the procedure and thus cannot use repetition to amplify the unambiguous soundness error. In subsequent subsections - we describe each of these components of the full protocol. We then combine all of these parts to get Protocol 10 and the following theorem:

Theorem 10 (IPP for PVAL with column distance). *Let $m, \log L, T, \sigma, d \in \mathbb{N}$. Let $\mathbb{F}$ be a constructible field ensembles such that $\mathbb{F}$ has characteristic 2. If*

- $|\mathbb{F}| \geq \Omega(2^\sigma \cdot T^2(m + \log L)^2),$
- $T > 8d \cdot 2^{\log L} \cdot m,$
- $d \geq 16 \cdot m \cdot \sigma$

Then for any $\mathbf{j} \in (\mathbb{F}^{m+\log L})^T, v \in \mathbb{F}^T$, the matrix $PVAL(\mathbf{j}, v)$ has a public coin unambiguous IPP as long as $\Delta_c(PVAL(\mathbf{j}, 0)) > 4d$:

- *Soundness error:* $m \cdot 2^{4-\sigma}$
- *Prover to Verifier communication complexity:* $T \log |\mathbb{F}| \cdot \text{poly}(m, \log L) + L \cdot \log |\mathbb{F}| \cdot \text{poly}(d).$
- *Verifier to Prover communication complexity:* $T \log |\mathbb{F}| \cdot \text{poly}(m, \log L)$
- *Round Complexity:* $\text{poly}(m, \log L, \log(T \log |\mathbb{F}|))$
- *Prover runtime:* $\text{poly}(2^{m+\log L}, T \log |\mathbb{F}|)$
- *Verifier runtime:* $T \log |\mathbb{F}| \cdot \text{poly}(m, \log L, \log(T \log |\mathbb{F}|), \log |\mathbb{F}|) + \tilde{O}(L \cdot \log |\mathbb{F}| \cdot \text{poly}(d))$
- *Query complexity :* $|\mathcal{Q}| = \lceil \frac{4\sigma \cdot 2^m}{d} \rceil \cdot L.$

The final verifier verdict is given by a succinctly described predicate $\Phi : \mathbb{F}^{\lceil \frac{4\sigma \cdot 2^m}{d} \rceil \cdot 2^{\log L}} \mapsto \{0,1\}$ along with a succinctly described set $Q \subseteq [2^m]$ such that $|Q| = \lceil \frac{4\sigma \cdot 2^m}{d} \rceil$, and the verdict is $\Phi(f|Q \times \{0,1\}^{\log L})$ with $|\langle Q \rangle| = \tilde{O}(m^2 d \sigma + m \log |\mathbb{F}| + \sigma \cdot \text{poly}(d))$ and $|\langle \Phi \rangle| = |\langle Q \rangle| + \sigma \cdot \text{poly}(d) \cdot L \cdot \log |\mathbb{F}|.$

Protocol 6 Breaking $\mathbf{a}$ into two multilinear functions

Input Parameters: $m, L, T \in \mathbb{N}$ and $\mathbb{F}$ is a field.

Input: $\mathbf{j} \in (\mathbb{F}^{m+\log L})^T$ and $\mathbf{v} \in \mathbb{F}^T$.

Prover Auxiliary Input: $\mathbf{a} : \{0, 1\}^{m+\log L} \rightarrow \mathbb{F}$.

Output: $\{\mathbf{j}'_i, \zeta_i^{(0)}, \zeta_i^{(1)}\}_{i \in [T]}$.

- 1: P lets $\mathbf{j}_i = (\chi_i, \mathbf{j}'_i) \in \mathbb{F} \times \mathbb{F}^{m-1+\log L}$.
 - 2: P sends $\{\zeta_i^{(b)} = \hat{\mathbf{a}}(b, \mathbf{j}'_i)\}_{i \in [T], b \in \{0,1\}}$ to V.
 - 3: **if** $\bigwedge_{i \in [T]} [v_i = (1 - \chi_i)\zeta_i^{(0)} + \chi_i\zeta_i^{(1)}]$ **then**
 - 4: V outputs $\{\mathbf{j}'_i, \zeta_i^{(b)}\}_{b \in \{0,1\}, i \in [T]}$.
 - 5: **else**
 - 6: V halts and rejects.
-

6.3.1 Breaking into Multilinear Functions and Consistency Checks

Lemma 12. *Let $T, d, m, \log L \in \mathbb{N}$ and $\mathbb{F}$, a finite field of characteristic 2. $\mathbf{j} \in (\mathbb{F}^{m+\log L})^T$ and $\mathbf{v} \in \mathbb{F}^T$. If, $\Delta_c(\text{PVAL}(\mathbf{j}, \mathbf{0})) \geq 4d$, then Protocol 6 has the following properties:*

- **Unambiguous Distance Preservation:**
 - If $\Delta_c(\mathbf{a}, \text{PVAL}(\mathbf{j}, \mathbf{v})) \leq d$ and the prover answers according to the prover's prescribed strategy for the unique $\mathbf{a}^* \in \text{PVAL}(\mathbf{j}, \mathbf{v}) \cap \mathcal{B}_{d, \mathbb{F}}(\mathbf{a})$ then $\mathbf{a}^*$ remains the unique element in $\mathcal{B}_{d, \mathbb{F}}(\mathbf{a}) \cap (\text{PVAL}(\mathbf{j}', \zeta^{(0)}) \parallel \text{PVAL}(\mathbf{j}', \zeta^{(1)}))$.
 - If $\Delta_c(\mathbf{a}, \text{PVAL}(\mathbf{j}, \mathbf{v})) > d$ or the prover is not honest with respect to the unique $\mathbf{a}^* \in \mathcal{B}_{d, \mathbb{F}}(\mathbf{a}) \cap \text{PVAL}(\mathbf{j}, \mathbf{v})$ then, $\Delta_c(\mathbf{a}, \text{PVAL}(\mathbf{j}', \zeta^{(0)}) \parallel \text{PVAL}(\mathbf{j}', \zeta^{(1)})) > d$.
- **Round Complexity:** 1
- **Communication Complexity:** $T \cdot \log \mathbb{F} \cdot (m + \log L + 1)$.
- **Verifier Time Complexity:** $O(T \text{polylog}(\mathbb{F}))$
- **Prover Runtime:** $\text{poly}(2^{m+\log L}, T, \log \mathbb{F})$.
- **In all cases:** $\Delta_c(\text{PVAL}(\mathbf{j}', \mathbf{0})) \geq 4d$.

Proof. The round, communication and time complexity follow directly. The final property is since $\text{PVAL}(\mathbf{j}', \mathbf{0}) \parallel \mathbf{0} \subseteq \text{PVAL}(\mathbf{j}, \mathbf{0})$ and the first has distance $\Delta_c(\text{PVAL}(\mathbf{j}', \mathbf{0}))$.

Now, if the verifier doesn't reject and $\mathbf{a}_0 \in \text{PVAL}(\mathbf{j}', \zeta^{(0)})$ and $\mathbf{a}_1 \in \text{PVAL}(\mathbf{j}', \zeta^{(1)})$ then, we have that for all $i \in [T]$:

$$\widehat{\mathbf{a}_0 \parallel \mathbf{a}_1}(\mathbf{j}_i) = \widehat{\mathbf{a}_0 \parallel \mathbf{a}_1}(\chi_i, \mathbf{j}'_i) = (1 - \chi_i)\widehat{\mathbf{a}_0}(\mathbf{j}'_i) + \chi_i\widehat{\mathbf{a}_1}(\mathbf{j}'_i) = (1 - \chi_i)\zeta_i^{(0)} + \chi_i \cdot \zeta_i^{(1)} = \mathbf{v}_i$$

by multilinearity. We get that $\mathbf{a}_0 \parallel \mathbf{a}_1 \in \text{PVAL}(\mathbf{j}, \mathbf{v})$. This clearly implies the honest prover case and the $\Delta_c(\mathbf{a}, \text{PVAL}(\mathbf{j}, \mathbf{v})) > d$ cases.

If the prover is dishonest with respect to $\mathbf{a}^*$, then we know that $\mathbf{a}^* \notin \text{PVAL}(\mathbf{j}', \zeta^{(0)}) \parallel \text{PVAL}(\mathbf{j}', \zeta^{(1)})$. Thus, $\Delta_c(\mathbf{a}^*, \text{PVAL}(\mathbf{j}', \zeta^{(0)}) \parallel \text{PVAL}(\mathbf{j}', \zeta^{(1)})) \geq 4d \implies \Delta_c(\mathbf{a}, \text{PVAL}(\mathbf{j}', \zeta^{(0)}) \parallel \text{PVAL}(\mathbf{j}', \zeta^{(1)})) > d$. $\square$

6.3.2 Permuting Errors and Curve Fitting

Protocol 7 Permuting Errors and Curve Fitting

Input parameters: $m, L, T, \sigma \in \mathbb{N}$ and $\mathbb{F}$ is a field.

Input: $\mathbf{j} \in (\mathbb{F}^{m+\log L})^T$ and $\mathbf{v}_0, \mathbf{v}_1 \in \mathbb{F}^T$.

Prover Auxiliary Input: $\mathbf{a}_0, \mathbf{a}_1 : \{0, 1\}^{m+\log L} \rightarrow \mathbb{F}$.

Other parameters: $\Lambda = (\Lambda_1, \Lambda_2, \dots, \Lambda_{2T})$ are distinct fixed canonical points in $\mathbb{F}$.

Output: A permutation $\tau : \{0, 1\}^{m+\log L} \rightarrow \{0, 1\}^{m+\log L}$, a curve $\mathcal{C} : \mathbb{F} \rightarrow \mathbb{F}^{m+\log L}$, $\mathbf{j}' \in (\mathbb{F}^{m+\log L})^T$, $\mathbf{v}' \in \mathbb{F}^T$, and $g^{(0)}$ and $g^{(1)} : \mathbb{F} \rightarrow \mathbb{F}$ which are univariate polynomials of degree at most $2T(m + \log L)$.

- 1: $\mathbf{V}$ samples $\pi \leftarrow \Pi_{m, 4d, 2^{-\sigma}}$ and sends it to the prover. They both let $\tau = \pi \times I$.
- 2: Both parties run Protocol 5 on $\mathbf{a}_1 \circ \tau, \pi^{-1}, \mathbf{j}, \mathbf{v}_1$, obtaining $\mathbf{j}', \mathbf{v}'$.
- 3: Let $\mathcal{C} : \mathbb{F} \rightarrow \mathbb{F}^{m+\log L}$ be the canonical low degree curve passing through $\mathbf{j}$ and $\mathbf{j}'$ i.e. the unique degree $\leq 2T - 1$ curve such that $\mathcal{C}(\Lambda) = (\mathbf{j}, \mathbf{j}')$.
- 4: $\mathbf{P}$ computes the univariate polynomials of degree $\leq (2T - 1)(m + \log L)$, denoted by $g^{(0)}$ and $g^{(1)}$, that satisfy

$$g^{(0)}(x) = \widehat{\mathbf{a}_0}(\mathcal{C}(x)) \quad g^{(1)}(x) = \widehat{\mathbf{a}_1 \circ \tau}(\mathcal{C}(x)),$$

and sends them to $\mathbf{V}$.

- 5: $\mathbf{V}$ accepts and outputs $\pi, \mathcal{C}, \mathbf{j}', \mathbf{v}', g^{(0)}, g^{(1)}$ if

$$g^{(0)}(\Lambda_1, \Lambda_2, \dots, \Lambda_T) = \mathbf{v}_0 \quad \wedge \quad g^{(1)}(\Lambda_{T+1}, \dots, \Lambda_{2T}) = \mathbf{v}'.$$

Lemma 13. Let $m, \log L, T, \sigma, r, d > 0$ be integers with $r < d/4$, $4d \leq 2^{m/50}$. Let $\mathbb{F}$ be a constructible field ensemble of characteristic 2. We also have the inputs $\mathbf{j}, \mathbf{v}_0, \mathbf{v}_1$ where $\mathbf{j} \in (\mathbb{F}^{m+\log L})^T$ and $\mathbf{v}_0, \mathbf{v}_1 \in \mathbb{F}^T$, with prover having access to functions $\mathbf{a}_0, \mathbf{a}_1 : \{0, 1\}^{m+\log L} \rightarrow \mathbb{F}$. Additionally, if $\Delta_c(\text{PVAL}(\mathbf{j}, 0)) \geq 4d$ then we have that the protocol either breaks and rejects or outputs univariate polynomials $g_0, g_1 : \mathbb{F} \rightarrow \mathbb{F}$ of degree at most $2T(m + \log L)$, a permutation $\pi : \{0, 1\}^m \rightarrow \{0, 1\}^m$, and a uniformly random $\mathbf{j}' \in (\mathbb{F}^{m+\log L})^T$. Then Protocol 7 has the following properties:

- **Unambiguous Distance Preservation:** Let $\mathbf{a} = \mathbf{a}_0 \| \mathbf{a}_1$. Now, with probability, $\geq 1 - (2^{-\sigma} + \text{err}_{\text{GKR}} + \exp(-d/12r))$, the following properties hold:
 - If $\Delta_c(\mathbf{a}, \text{PVAL}(\mathbf{j}, \mathbf{v}_0) \| \text{PVAL}(\mathbf{j}, \mathbf{v}_1)) \leq d$ and the prover answers according to the prover's prescribed strategy for the unique $\mathbf{a}_0^* \| \mathbf{a}_1^* = \mathbf{a}^* \in (\text{PVAL}(\mathbf{j}, \mathbf{v}_0) \| \text{PVAL}(\mathbf{j}, \mathbf{v}_1)) \cap \mathcal{B}_{d, \mathbb{F}}(\mathbf{a})$ then:
 - * For any $c \in \mathbb{F}$: $\mathbf{a}_0^* + c\mathbf{a}_1^* \circ \tau$ is the unique element in $\mathcal{B}_{d, \mathbb{F}}(\mathbf{a}_0 + c\mathbf{a}_1 \circ \tau) \cap (\text{PVAL}(\mathcal{C}(\mathbb{F}), g_0 + cg_1(\mathbb{F})))$.
 - * $\Pr_c [\Delta_c(\mathbf{a}_0 + c\mathbf{a}_1 \circ \tau, \mathbf{a}_0^* + c\mathbf{a}_1^* \circ \tau) < \Delta_c(\mathbf{a}, \mathbf{a}^*)(1 - \frac{1}{r})] < \frac{2r+1}{\mathbb{F}}$.
 - If $\Delta_c(\mathbf{a}, \text{PVAL}(\mathbf{j}, \mathbf{v}_0) \| \text{PVAL}(\mathbf{j}, \mathbf{v}_1)) > d$ or the prover is not honest with respect to the unique $\mathbf{a}^* \in \Delta_c(\mathbf{a}, \text{PVAL}(\mathbf{j}, \mathbf{v}_0) \| \text{PVAL}(\mathbf{j}, \mathbf{v}_1))$ then either:
 - * $\text{PVAL}((\mathbf{j}, \mathbf{j}'), g_0(\Lambda))$ is empty; or
 - * $\text{PVAL}((\mathbf{j}, \mathbf{j}'), g_1(\Lambda))$ is empty; or

$$* \Pr_c [\Delta_c(\mathbf{a}_0 + c\mathbf{a}_1 \circ \tau, \text{PVAL}(\mathcal{C}(\mathbb{F}), g_0 + cg_1(\mathbb{F}))) < d(1 - \frac{1}{r})] < \frac{2r+1}{\mathbb{F}}.$$

- **Round Complexity:** $O(D \log S)$
- **Message length per round:** $\tilde{O}(T(m + \log L) \log F + dm\sigma)$.
- **Verifier Runtime:** $\tilde{O}(D \log S \cdot \text{polylog}(\mathbb{F}) + dm\sigma + T(m + \log L))$.
- **Prover Runtime:** $\text{poly}(S, \log \mathbb{F}, T, 2^{m+\log L})$.

Here D and S are the depth and size respectively of $\Phi_{\pi, j, v}$ as defined in Definition 14.

Proof. We again don't focus on round, communication, and time complexities as they follow directly. The complexities of π , g_0 , g_1 and Protocol 5 directly imply the results. Focusing on the unambiguous distance preservation property, the result is essentially a restatement of Lemma 11 with $\epsilon = \gamma = \frac{1}{2r}$. In particular:

- If $\Delta_c(\mathbf{a}, \text{PVAL}(\mathbf{j}, \mathbf{v}_0) \parallel \text{PVAL}(\mathbf{j}, \mathbf{v}_1)) > d$, then the result follows by Lemma 11 with $\epsilon = \gamma = \frac{1}{2r}$.
- If $\Delta_c(\mathbf{a}, \text{PVAL}(\mathbf{j}, \mathbf{v}_0) \parallel \text{PVAL}(\mathbf{j}, \mathbf{v}_1)) \leq d$ and the prover is honest then:
 - The first condition follows since $\Delta_c(\text{PVAL}(\mathcal{C}(\mathbb{F}), g_0 + cg_1(\mathbb{F}))) \geq \Delta_c(\text{PVAL}(\mathbf{j}, 0)) \geq 4d$, and we know that by the honesty definition that $\mathbf{a}_0^* + c\mathbf{a}_1^* \circ \tau \in \mathcal{B}_{d, \mathbb{F}}(\mathbf{a}_0 + c\mathbf{a}_1 \circ \tau)$. Thus, it must also be unique.
 - The second part follows from Lemma 11.
- If the prover is dishonest first in:
 - The GKR i.e. Protocol 5 then with probability $\geq 1 - \text{err}_{\text{GKR}}$, we have that $\Delta_c(\mathbf{a}_1 \circ \pi, \text{PVAL}(\mathbf{j}', \mathbf{v}')) \geq 4d$. Now, by Lemma 8, the result follows.
 - In returning g : If the checks pass, but the prover lied in the value of g , then we know that either $\Delta_c(\mathbf{a}_0^*, \text{PVAL}(\mathcal{C}(\mathbb{F}), g_0(\mathbb{F}))) \geq 4d$ or $\Delta_c(\mathbf{a}_1^* \circ \tau, \text{PVAL}(\mathcal{C}(\mathbb{F}), g_1(\mathbb{F}))) \geq 4d$. Thus, since $\mathbf{a}_0^*$ and $\mathbf{a}_1^* \circ \tau$ are within d distance of $\mathbf{a}_0$ and $\mathbf{a}_1 \circ \tau$ by the GKR correctness and our assumption respectively. Thus, we can again simply conclude by the triangle inequality and Lemma 8.

□

6.3.3 Efficient UIP for PVAL emptiness

Since we have used non-emptiness conditions on the PVAL instances defined in Lemma 11, we need to be able to check the same. The following is a standard interactive proof for checking whether a given PVAL instance is empty. This is based on the observation that checking emptiness of a PVAL instance is simply a rank computation, and thus we can use Theorem 3 on such a circuit. We will make use of this lemma freely now.

Lemma 14 (Lemma 5.10 in [RR20]). *Let $T, m, \sigma \in \mathbb{N}$ and $\mathbb{F}$, a finite field of characteristic 2. There is a public coin unambiguous interactive proof for the language*

$$\mathcal{L} = \left\{ (\mathbf{j}, \mathbf{v}) \in \left((\mathbb{F}^m)^T, \mathbb{F}^T \right) \mid \text{PVAL}(\mathbf{j}, \mathbf{v}) \neq \emptyset \right\},$$

with perfect completeness and the following parameters:

- *Unambiguous Soundness:* $2^{-\sigma}$
- *Communication complexity:* $\text{poly}(m, \log(T \log |\mathbb{F}|), \sigma)$.
- *Round Complexity:* $\text{poly}(m, \log(T \log |\mathbb{F}|), \sigma)$
- *Verifier running time:* $T \cdot \text{poly}(m, \log |\mathbb{F}|, \sigma)$
- *Prover running time:* $\text{poly}(2^m, T \log |\mathbb{F}|, 2^\sigma)$.

Proof. Let matrix $M \in \mathbb{F}^{2^m \times T}$ where

$$M_{a,b} = \hat{x}^a(j_b)$$

where $x^a = \prod_{a_r=1} x_r$. Observe that each entry of the matrix can be computed using an arithmetic circuit over $\mathbb{F}$ of depth $O(\log m)$ and size m . Now, since each multiplication and addition over $\mathbb{F}$ as a boolean circuit is of depth $\log |\mathbb{F}|$ and size $\log^2 |\mathbb{F}|$. Thus, each coordinate of M can be computed in size $m \log |\mathbb{F}|^2$ and depth $O(\log |\mathbb{F}| \log m)$.

The rank of a matrix can be computed in depth $O(\log^2(T \cdot 2^m))$ as a circuit over $\mathbb{F}$ and in size $\text{poly}(2^m T)$. Now, each operation over $|\mathbb{F}|$ multiplies another $\log^2 |\mathbb{F}|$ to size and another $\log |\mathbb{F}|$ to the depth. Now, the verifier can simply check using GKR whether $\text{rank}(M) = \text{rank}(M \| \mathbf{v})$. Observe that if these two ranks are equal then $\mathbf{v} \in \mathbb{F}^T$ is in the $\mathbb{F}$ linear span of the columns of M i.e. there is a polynomial whose extension applied to $\mathbf{j}$ has the image $\mathbf{v}$ and the affine space has dimension > 0 if the rank is not full. Similarly, we could have also just checked if $\dim \text{PVAL}(\mathbf{j}, \mathbf{v}) > 0$ instead of just checking if $|\text{PVAL}(\mathbf{j}, \mathbf{v})| > 0$. $\square$

6.3.4 Reducing Number of Evaluation Points:

Lemma 15. $m, \log L, T, \sigma, d > 0$ are integers. $\mathbb{F}$ is a constructible field of characteristic 2. We also have $\mathbf{j} \in (\mathbb{F}^{m+\log L})^T$ and uniformly random points $\mathbf{j}' \in (\mathbb{F}^{m+\log L})^T$. $\mathbf{\Lambda} = (\Lambda_1, \dots, \Lambda_{2T})$ are canonical distinct fixed points in $\mathbb{F}$. Let $\mathbf{a} : \{0, 1\}^{m+\log L} \rightarrow \mathbb{F}$ and let $g : \mathbb{F} \rightarrow \mathbb{F}$ be a univariate polynomial of degree at most $2T(m + \log L)$. Additionally, assume that $\Delta_c(\text{PVAL}(\mathbf{j}, 0))$ is non-empty.

Then, we have the following unambiguous distance preservation property:

- If $\Delta_c(\mathbf{a}, \text{PVAL}(\mathcal{C}(\mathbb{F}), g(\mathbb{F}))) \leq d$ then:

$$\Pr_{\alpha=(\alpha_1, \dots, \alpha_T) \in \mathbb{F}^T} [\mathcal{B}_{d, \mathbb{F}}(\mathbf{a}) \cap \text{PVAL}(\mathcal{C}(\alpha), g(\alpha)) = \mathcal{B}_{d, \mathbb{F}}(\mathbf{a}) \cap \text{PVAL}(\mathcal{C}(\mathbb{F}), g(\mathbb{F}))] \geq p$$

- If $\Delta_c(\mathbf{a}, \text{PVAL}(\mathcal{C}(\mathbb{F}), g(\mathbb{F}))) > d$ then:

$$\Pr_{\alpha=(\alpha_1, \dots, \alpha_T) \in \mathbb{F}^T} [\Delta_c(\text{PVAL}(\mathcal{C}(\alpha), g(\alpha)), \mathbf{a}) > d] \geq p,$$

where $p = 1 - \left(\binom{M}{d} |\mathbb{F}|^d \right)^L \left(\frac{2T(m+\log L)}{\mathbb{F}} \right)^T$.

Additionally,

$$\Pr_{\alpha=(\alpha_1, \dots, \alpha_T) \in \mathbb{F}^T} [\Delta_c(\text{PVAL}(\mathcal{C}(\alpha), 0)) \geq 4d] \geq 1 - \left(\left(\frac{M}{4d} \right) |\mathbb{F}|^{4d} \right)^L \left(\frac{m + \log L}{\mathbb{F}} \right)^T$$

where $\mathcal{C} : \mathbb{F} \rightarrow \mathbb{F}^{m+\log L}$ is the canonical $2T - 1$ degree curve passing through $\mathbf{j}, \mathbf{j}'$.

Proof. For the first part, observe that for if $\mathbf{a}' \in \mathcal{B}_{d,\mathbb{F}}(\mathbf{a}) \setminus \text{PVAL}(\mathcal{C}(\mathbb{F}), g(\mathbb{F}))$ then we have that $\widehat{\mathbf{a}'} \circ \mathcal{C} \neq g$. Thus, $\Pr_{\alpha \in \mathbb{F}^T}[\widehat{\mathbf{a}'} \circ \mathcal{C}(\alpha) = g(\alpha)] \leq \left(\frac{2T(m+\log L)}{|\mathbb{F}|}\right)^T$ by the Schwartz-Zippel Lemma (Lemma 1). Thus, we finish by a union bound over all elements of $\mathcal{B}_{d,\mathbb{F}}(\mathbf{a})$.

For the second part, observe that since $\mathbf{j}'$ is uniformly random, thus so is $\mathcal{C}(\alpha)$. Thus, for any element $\mathbf{a}'$ such that $0 < \Delta_c(0, \mathbf{a}') < 4d$, we have that $\Pr_{\alpha \in \mathbb{F}^T}[\widehat{\mathbf{a}'}(\mathcal{C}(\alpha)) = 0] \leq \left(\frac{m+\log L}{|\mathbb{F}|}\right)^T$ by Lemma 1. We can now finish by the union bound over all elements.

Remark. In the proof of the second part, we used the fact that $\mathbf{j}'$ was uniformly random. If we do not want to use this fact then we can observe that $\widehat{\mathbf{a}'} \circ \mathcal{C} \neq 0$ for any $\mathbf{a}'$ since $\Delta_c(\text{PVAL}(\mathbf{j}, 0)) \geq 8d$ and have a loss of a factor of T , but this will not matter much anyway. $\square$

6.3.5 Putting together the Reduction Protocol:

Protocol 8 The Reduction Protocol

Input Parameters: $m, L, T, \sigma \in \mathbb{N}$ and a field $\mathbb{F}$.

Input: $\mathbf{j} \in (\mathbb{F}^{m+\log L})^T, \mathbf{v} \in \mathbb{F}^T$.

Prover Auxiliary Input: $\mathbf{a} : \{0, 1\}^{m+\log L} \rightarrow \mathbb{F}$.

Output: A permutation $\tau \in \{0, 1\}^{m-1+\log L} \rightarrow \{0, 1\}^{m-1+\log L}$, a scalar $c \in \mathbb{F}$, and a smaller instance: $\mathbf{a}_{\text{new}} : \{0, 1\}^{m-1+\log L} \rightarrow \mathbb{F}, \mathbf{j}_{\text{new}} \in (\mathbb{F}^{m-1+\log L})^T$ and $\mathbf{v}_{\text{new}} \in \mathbb{F}^T$.

- 1: Run Protocol 6 on input $\mathbf{j}, \mathbf{v}$ and $\mathbf{a}$ to get output $\mathbf{j}'', \zeta^{(0)}, \zeta^{(1)}$.
 - 2: Break $\mathbf{a}$ into $\mathbf{a}_0 \parallel \mathbf{a}_1$.
 - 3: Run Protocol 7 on input $\mathbf{j}'', \mathbf{v}_0 = \zeta^{(0)}, \mathbf{v}_1 = \zeta^{(1)}$, obtaining $\tau, \mathcal{C}, \mathbf{j}', \mathbf{v}', g_0, g_1$.
 - 4: **PVAL Emptiness Check:** Both parties run the Protocol in Lemma 14 with error parameter σ to check the non-emptiness of $\text{PVAL}((\mathbf{j}, \mathbf{j}'), g^{(0)}(\Lambda))$ and $\text{PVAL}((\mathbf{j}, \mathbf{j}'), g^{(1)}(\Lambda))$.
 - 5: $\mathbf{V}$ rejects if the check fails.
 - 6: $\mathbf{V}$ samples $\alpha_1, \dots, \alpha_T \leftarrow \mathbb{F}$ and $c \leftarrow \mathbb{F}$ and sends them to $\mathbf{P}$.
 - 7: $\mathbf{V}$ outputs $\tau, c, \mathbf{a}_{\text{new}} = \mathbf{a}_0 + c(\mathbf{a}_1 \circ \tau), \mathbf{j}_{\text{new}} = (\mathcal{C}(\alpha_i))_{i=1}^T$ and $\mathbf{v}_{\text{new}} = ((g_0 + cg_1)(\alpha_i))_{i=1}^T$.
-

Lemma 16. Let $m, \log L, T, \sigma, d, r \in \mathbb{N}$ such that $r < d/4$ and $4d < 2^{m/50}$ and let $\mathbb{F}$ be a constructible field. Let $\mathbf{a} : \{0, 1\}^{m+\log L} \rightarrow \mathbb{F}$ and $\mathbf{j} \in (\mathbb{F}^{m+\log L})^T, \mathbf{v} \in \mathbb{F}^T$ such that $\Delta_c(\text{PVAL}(\mathbf{j}, 0)) \geq 4d$. Define the following quantity.

$$p := 1 - \left(\frac{2r+1}{|\mathbb{F}|} + 2 \left(\left(\frac{M}{4d} \right)^{|\mathbb{F}|^{4d}} \right)^L \left(\frac{m+\log L}{|\mathbb{F}|} \right)^T \right) - \left(\left(\left(\frac{M}{d} \right)^{|\mathbb{F}|^d} \right)^L \left(\frac{2T(m+\log L)}{|\mathbb{F}|} \right)^T + 2^{2-\sigma} + \text{err}_{\text{GKR}} + \exp(-d/12r) \right).$$

If we run Protocol 8 and get outputs $\mathbf{a}_{\text{new}}, \mathbf{j}_{\text{new}}, \mathbf{v}_{\text{new}}$ then we have the following properties:

- **Completeness:** If $\mathbf{a} \in \text{PVAL}(\mathbf{j}, \mathbf{v})$ and the prover was honest then $\mathbf{a}_{\text{new}} \in \text{PVAL}(\mathbf{j}_{\text{new}}, \mathbf{v}_{\text{new}})$.
- **Soundness:** If $\mathbf{a} \notin \text{PVAL}(\mathbf{j}, \mathbf{v})$ with respect to $\mathbf{a}$ then

$$\Pr[\Delta_c(\mathbf{a}_{\text{new}}, \text{PVAL}(\mathbf{j}, \mathbf{v})) \geq d(1 - \frac{1}{r})] \geq p.$$

- **Unambiguity:** If $\mathbf{a} \in \text{PVAL}(\mathbf{j}, \mathbf{v})$ but the prover deviates from the honest strategy then

$$\Pr[\Delta_c(\mathbf{a}_{\text{new}}, \text{PVAL}(\mathbf{j}, \mathbf{v})) \geq d(1 - \frac{1}{r})] \geq p.$$

- With probability at least $1 - \left(\binom{M}{4d} |\mathbb{F}|^{4d}\right)^L \left(\frac{m + \log L}{|\mathbb{F}|}\right)^T$, we have that $\Delta_c(\text{PVAL}(\mathbf{j}_{\text{new}}, 0)) \geq 4d$.

Additionally, the protocol has the following complexities:

- **Round Complexity:** $\tilde{O}(D \log S \cdot \text{poly}(m + \log L, \log(T \log |\mathbb{F}|), \sigma))$
- **Communication Complexity:** $(T + d) \cdot \text{poly}(m + \log L, \log(T \log |\mathbb{F}|), \sigma)$
- **Verifier Runtime:** $(T + d)D \log S \cdot \text{poly}(m + \log L, \log(T \log |\mathbb{F}|), \sigma)$.
- **Prover Runtime:** $\text{poly}(S, T, 2^{m + \log L}, \log |\mathbb{F}|, \sigma)$.
- **Verifier's verification Circuit Size:** $\tilde{O}(T) \text{poly}(\log |\mathbb{F}|, m + \log L, \sigma)$.
- **Verifier's verification circuit depth:** $\text{polylog}(T, m + \log L, \log |\mathbb{F}|, \text{sec})$.

If we further assume that

- $|\mathbb{F}| \geq \tilde{\Omega}(2^\sigma \cdot T^2(m + \log L + \log T)^2)$.
- $T > 8d \cdot 2^{\log L} \cdot m$
- $d \geq 16 \cdot r \cdot \sigma$
- $r = o(|\mathbb{F}| \cdot 2^{-\sigma})$,

then

$$\begin{aligned} & \left(\frac{2r + 1}{|\mathbb{F}|} + 3 \left(\binom{M}{4d} |\mathbb{F}|^{4d} \right)^L \left(\frac{m + \log L}{|\mathbb{F}|} \right)^T \right) \\ & + \left(\left(\binom{M}{d} |\mathbb{F}|^d \right)^L \left(\frac{2T(m + \log L)}{|\mathbb{F}|} \right)^T + 2^{2-\sigma} + \text{err}_{\text{GKR}} + \exp(-d/12r) \right) \leq 2^{3-\sigma}. \end{aligned}$$

Proof. This statement is a combination of Lemmas 12 to 15. The only thing left to show is the final verifier's verification circuit and the setting of parameters.

The verifier's verification checks are:

- Checking Linear Combinations: Size: $T \text{polylog}(|\mathbb{F}|)$ and Depth $O(\log(T) + \log \log |\mathbb{F}|)$.
- Protocol 5: Depth D and size S .
- Checking $O(T)$ polynomial evaluations:
Size $\tilde{O}(T \cdot \log |\mathbb{F}|(m + \log L))$ and depth $\text{poly}(\log(T \log |\mathbb{F}|(m + \log L)))$ since we can run multipoint polynomial evaluation.
- Emptiness checks: Size $T \cdot \text{poly}(m + \log L, \sigma, \log |\mathbb{F}|)$ and depth $\text{polylog}(T, m + \log L, \log |\mathbb{F}|, \sigma)$.

Thus, putting all these together, we get the desired result.

We verify the parameters set:

- $\frac{2r+1}{|\mathbb{F}|} \leq 2^{-\sigma}$
- $\left(\binom{M}{4d} \cdot |\mathbb{F}|^d \right) \left(\frac{(m + \log L)}{|\mathbb{F}|} \right)^T \leq \left(\frac{(2m + 2\log L)^m \cdot |\mathbb{F}|}{|\mathbb{F}|^m} \right)^{4dL} \leq \frac{1}{|\mathbb{F}|} \leq 2^{-\sigma}.$
- $\left(\binom{M}{d} \cdot |\mathbb{F}|^d \right) \left(\frac{T(m + \log L)}{|\mathbb{F}|} \right)^T \leq \left(\frac{T(m + 2\log L)^{8m} \cdot |\mathbb{F}|}{|\mathbb{F}|^{8m}} \right)^{dL} \leq \frac{1}{|\mathbb{F}|} \leq 2^{-\sigma}$
- $\text{err}_{\text{GKR}} \leq \left(\epsilon_{\text{GKR}}(D, \log S, |\mathbb{F}|) + (\epsilon_{\text{GKR}}(D, \log S, |\mathbb{F}|))^T \cdot \left(\binom{M}{4d} |\mathbb{F}|^{4d} \right)^L \right)$
 $\leq O\left(\frac{D \log S}{|\mathbb{F}|} \right) + \left(\frac{(O(D \log S)^m \cdot |\mathbb{F}|)}{|\mathbb{F}|^m} \right)^{4dL} \leq \frac{1}{|\mathbb{F}|} \leq 2^{-\sigma}$
- $\exp(-d/12r) \leq 2^{-\sigma}.$

Combining these we get the desired result. $\square$

6.3.6 The Base Protocol

Protocol 9 Base protocol in RR.

Input Parameters: m, m_{new} (where $m \geq m_{\text{new}}$), $L, \sigma, d \in \mathbb{N}$ and $\mathbb{F}$ is a field.

Input: $\mathbf{j} \in (\mathbb{F}^{m_{\text{new}} + \log L})^T, \mathbf{v} \in \mathbb{F}^T.$

Other Parameters: Let $M := 2^m$ and $M^{\text{new}} := 2^{m_{\text{new}}}.$

Prover Auxiliary Input: A function $\mathbf{a}_{\text{new}} \in \{0, 1\}^{m_{\text{new}}} \times \{0, 1\}^{\log L} \rightarrow \mathbb{F}.$

Verifier Auxiliary Input: A description $\langle H \rangle$ of a $\mathbb{F}$ -linear function $H : (\mathbb{F}^L)^M \rightarrow (\mathbb{F}^L)^{M_{\text{new}}},$ such that its every output coordinate is a linear combination of exactly $\frac{M}{M_{\text{new}}}$ elements of $\mathbb{F}^L.$ (We assume that the set of rows that fold into any given row of M_{new} are easy to compute given the description $\langle H \rangle.$)

Verifier Query Access: Query access to $\mathbf{a} : \{0, 1\}^m \times \{0, 1\}^{\log L} \rightarrow \mathbb{F}.$

Output: A succinct description $\langle Q \rangle$ of $Q \subset \{0, 1\}^m$ of size $\lceil \frac{\sigma \cdot M}{d} \rceil$ and a succinct predicate $\langle \Phi \rangle : \mathbb{F}^{|Q| \times L} \rightarrow \{0, 1\}.$

- 1: P explicitly sends all of $\mathbf{a}_{\text{new}}$ i.e. the entire string in $\mathbb{F}^{2^{m_{\text{new}}} \times L}$ to V.
 - 2: V checks if all the extensions are correct i.e. $\mathbf{a}_{\text{new}} \in \text{PVAL}(\mathbf{j}, \mathbf{v})$ (and rejects otherwise).
 - 3: V randomly selects Q' to be a uniform subset of $\lceil \frac{\sigma \cdot M}{d} \rceil$ of $[M_{\text{new}}].$
 - 4: V outputs Q as the set of rows that the rows of $H|_{Q'}$ depends on. Φ checks if $H(\mathbf{a}|_Q) = \mathbf{a}_{\text{new}}|_{Q'}.$
-

Claim 4. If $\Delta_c(\text{PVAL}(\mathbf{j}, \mathbf{v}), H(\mathbf{a})) \geq d$ then with probability $\geq 1 - 2^{-\sigma},$ we have that $\Phi(\mathbf{a}|_Q) = 0.$ Additionally, Protocol 9 has the following complexities:

- Round Complexity: $O(m_{\text{new}})$

- *Communication Complexity:* $O(2^{m_{\text{new}}} \cdot L \cdot \log(\mathbb{F}) + T \cdot (m_{\text{new}} + \log L) \cdot \log(\mathbb{F}))$.
- *Verifier runtime:* $O(M_{\text{new}} \cdot L \cdot \text{poly}(\log L + m, \log \mathbb{F}) \cdot T)$
- *Prover runtime:* $M_{\text{new}} L \cdot \log |\mathbb{F}| \text{poly}(T, m_{\text{new}} + \log L)$

Proof. When the distance is at least d , the verification passes exactly when none of the rows in the sampled Q' are inconsistent with H . This probability is at most $(1 - \frac{d}{M_{\text{new}}})^{|Q'|} \leq 2^{-\sigma}$. $\square$

6.3.7 The complete protocol:

Theorem 11 (IPP for PVAL with column distance). *Let $m, \log L, T, \sigma, d \in \mathbb{N}$. Let $\mathbb{F}$ be a constructible field ensembles such that $\mathbb{F}$ has characteristic 2. If*

- $|\mathbb{F}| \geq \Omega(2^\sigma \cdot T^2(m + \log L)^2)$,
- $T > 8d \cdot 2^{\log L} \cdot m$,
- $d \geq 16 \cdot m \cdot \sigma$

Then for any $\mathbf{j} \in (\mathbb{F}^{m+\log L})^T, v \in \mathbb{F}^T$, the matrix $\text{PVAL}(\mathbf{j}, v)$ has a public coin unambiguous IPP (Protocol 10) as long as $\Delta_c(\text{PVAL}(\mathbf{j}, 0)) > 4d$:

- *Soundness error:* $m \cdot 2^{4-\sigma}$
- *Prover to Verifier communication complexity:* $T \log |\mathbb{F}| \cdot \text{poly}(m, \log L) + L \cdot \log |\mathbb{F}| \cdot \text{poly}(d)$.
- *Verifier to Prover communication complexity:* $T \log |\mathbb{F}| \cdot \text{poly}(m, \log L)$
- *Round Complexity:* $\text{poly}(m, \log L, \log(T \log |\mathbb{F}|))$
- *Verifier runtime:* $T \log |\mathbb{F}| \cdot \text{poly}(m, \log L, \log(T \log |\mathbb{F}|), \log |\mathbb{F}|) + \tilde{O}(L \cdot \log |\mathbb{F}| \cdot \text{poly}(d))$
- *Prover runtime:* $\text{poly}(2^{m+\log L}, T \log |\mathbb{F}|)$
- *Query complexity :* $|\mathcal{Q}| = \frac{4\sigma \cdot 2^{m+\log L}}{d}$.

The final verifier verdict is given by a succinctly described predicate $\Phi : \mathbb{F}^{\frac{4\sigma}{d} \cdot 2^{m+\log L}} \mapsto \{0, 1\}$ along with a succinctly described set $Q \subseteq [2^m]$ such that $|Q| = \frac{4\sigma \cdot 2^m}{d}$, and the verdict is $\Phi(f|Q \times \{0, 1\}^{\log L})$ with $|\langle Q \rangle| = \tilde{O}(m^2 d \sigma + m \log |\mathbb{F}| + \sigma \cdot \text{poly}(d))$ and $|\langle \Phi \rangle| = |\langle Q \rangle| + \sigma \cdot \text{poly}(d) \cdot L \cdot \log |\mathbb{F}|$.

Proof. Observe that the total error is at most the sum of the errors in each round which we know is at most $2^{3-\sigma}$ for the first m rounds if $r = m$ in Lemma 16. This would result in that if $\Delta_c(\mathbf{a}, \text{PVAL}(\mathbf{j}, v)) \geq d$ initially, then after i rounds, we have $\Delta_c(\mathbf{a}_{\text{prot}}, \text{PVAL}(\mathbf{j}, v)) \geq d(1 - \frac{1}{m})^i$. Thus, with probability at least $m \cdot 2^{3-\sigma}$, we have that when Protocol 9 is called then indeed $\Delta_c(\mathbf{a}_{\text{prot}}, \text{PVAL}(\mathbf{j}, v)) > d/4$ as required. Thus, the total error is at most $m \cdot 2^{4-\sigma}$ as required.

The total round, time, communication complexities are just the sum of m times the complexities of Protocol 8 and the complexity of Protocol 9, as desired. The final query complexity is as described in Protocol 9.

We discuss the succinct descriptions in the next subsection. $\square$

Protocol 10 Efficient IPP for PVAL with column distance

Input Parameters: $m, L \in \mathbb{N}$, $\sigma, d \geq 16\sigma \cdot m$, $\mathbb{F}$ is a field with $|\mathbb{F}| \geq 32 \cdot 2^\sigma (T \log M \log L)^C$.

Input: $\mathbf{j} \in (\mathbb{F}^{m+\log L})^T, \mathbf{v} \in \mathbb{F}^T$.

Other Parameters: Let $r := m$.

Prover Auxiliary Input: $\mathbf{a} : \{0, 1\}^{m+\log L} \rightarrow \mathbb{F}$.

Verifier Query Access: Query access to $\mathbf{a} : \{0, 1\}^{m+\log L} \rightarrow \mathbb{F}$.

Output: Succinct descriptions of $Q \subset [M]$ and $\Phi : \mathbb{F}^{|Q| \times L} \rightarrow \{0, 1\}$.

- 1: $\mathbf{V}$ maintains a linear function $H : (\mathbb{F}^{\log L})^M \rightarrow (\mathbb{F}^{\log L})^{M_{\text{prot}}}$.
 - 2: Set the variable parameters: $m_{\text{prot}} := m$, $H_{\text{prot}} := Id$ and $\mathbf{a}_{\text{prot}} := \mathbf{a}$.
 - 3: **while** True **do**
 - 4: **if** $2^{m_{\text{prot}}/50} \leq 4d$: **then**
 - 5: Both parties run Protocol 9 on input parameters $m, m_{\text{prot}}, L, \sigma, d/4, \mathbb{F}$, and input $\mathbf{j}, \mathbf{v}$, prover auxiliary input $\mathbf{a}_{\text{prot}}$, verifier auxiliary input $\langle H_{\text{prot}} \rangle$ and verifier query access to $\mathbf{a}$, obtaining $\langle Q \rangle$ and $\langle \Phi \rangle$.
 - 6: $\mathbf{V}$ outputs $\langle Q \rangle$ and $\langle \Phi \rangle$.
 - 7: **else**
 - 8: Both parties run Protocol 8 on input parameters $m_{\text{prot}}, L, \sigma$ and T , input $\mathbf{j}, \mathbf{v}$ and prover auxiliary input $\mathbf{a}_{\text{prot}} = \mathbf{a}_{0,\text{prot}} \parallel \mathbf{a}_{1,\text{prot}}$, obtaining $\tau, c, \mathbf{j}_{\text{new}}, \mathbf{v}_{\text{new}}$ and $\mathbf{a}_{\text{new}} = \mathbf{a}_0 + c\mathbf{a}_1 \circ \tau$.
 - 9: Update $\mathbf{j} = \mathbf{j}_{\text{new}}, \mathbf{v} = \mathbf{v}_{\text{new}}, m_{\text{prot}} = m_{\text{prot}} - 1$, and $\mathbf{a}_{\text{prot}} = \mathbf{a}_{\text{new}}$.
 - 10: Update $H_{\text{prot}} = G \circ H_{\text{prot}}$ where $G(x_0 \parallel x_1) = x_0 + cx_1 \circ \tau$. $\langle H_{\text{prot}} \rangle = \langle H_{\text{prot}} \rangle \parallel (\tau, c)$.
-

6.4 Succinct Descriptions

We verify that the output parameters (Q, Φ) have succinct descriptions. The size of $\langle Q \rangle$ is given by the randomness of the choice in Protocol 9 which is $\text{poly}(d) \cdot \sigma \cdot \log |\mathbb{F}|$ and the (c, π) pairs from the m rounds. Thus, giving us a total complexity of $\tilde{O}(m^2 d \sigma + m \log |\mathbb{F}| + \sigma \cdot \text{poly}(d))$ since the seed length of a single π is $\tilde{O}(m \cdot d \cdot \sigma)$.

Additionally, $|\langle \Phi \rangle|$ is just given by $|\langle Q \rangle| + |Q'| \cdot L \cdot \log |\mathbb{F}|$ since $\langle Q \rangle$ contains the entire details of $\langle H \rangle$. Thus, we only need $\mathbf{a}_{\text{prot}}|Q'|$ which is $\log L \cdot |Q'| \cdot \log |\mathbb{F}| = \log L \log |\mathbb{F}| \sigma \cdot \text{poly}(d)$ bits.

Circuit G for generating Q Let $m' = m_{\text{prot}}$ when Protocol 9 is called.

The set Q has succinct description

$$\langle Q \rangle = (\langle Q \rangle', \{\pi : \pi \text{ is used in the folding steps of the protocol}\}),$$

where $Q' \subset \{0, 1\}^{m'}$ are the $\frac{4\sigma \cdot 2^{m'}}{d}$ random rows selected in the end of the iteration and $\langle R \rangle \in \{0, 1\}^{|Q'| \cdot m'}$ are the random coins used to generate Q' .

Given $(i, \langle Q \rangle)$, G needs to generate the i -th element of Q . The idea is to carry out the folding steps in reverse, which can be done as follows. First, it generates the full set Q' using $\langle R \rangle$, and let $Q_0 := Q'$, which requires $\tilde{O}(|Q'|)$ gates and $\tilde{O}(1)$ depth. If $i \leq |Q|/2$, the i -th element of Q would lie in the first half before folding, so G prepends a 0 to every element in Q_0 . Otherwise, G prepends a 1 to every element in Q_0 . Since the indices in the second half were permuted by π , G also applies π^{-1} to every element in Q_0 to reconstruct their locations in the original index space.

The circuit G repeats this process $(m - m')$ times, until it obtains $|Q'|$ indices that lies in $\{0, 1\}^m$, and it outputs the $(i \bmod |Q'|)$ -th index.

Circuit $C(a|Q \times \{0, 1\}^{\log L}, \langle \Phi \rangle)$ **for checking** $\Phi(a|Q \times \{0, 1\}^{\log L})$ Let $Q' \subset \{0, 1\}^{m'}$ be the $\left\lceil \frac{4\sigma \cdot 2^{m'}}{d} \right\rceil$ random rows selected in the end of the iteration, and $\langle Q' \rangle' \in \{0, 1\}^{|Q'| \cdot m'}$ be the random coin used to generate it. The predicate Φ can be described by the pair

$$(\langle Q' \rangle', \{(c, \pi) : (c, \pi) \text{ is used in the folding steps of the protocol}\}, \psi),$$

where ψ is the set of values claimed by the prover on Q' .

The circuit first enumerates the entire Q using $\langle Q' \rangle'$ and all the π 's (as in the last paragraph), and then it applies the $\mathbb{F}$ -linear function $\phi : \mathbb{F}^{|Q| \times 2^{\log L}} \rightarrow \mathbb{F}^{2^{\log L}}$ determined by all the (c, π) used in the folding steps of the protocol, on the $2^{\log L}$ columns of the truth table of f in parallel. This can be done by iteratively applying the corresponding π and figuring out which of c to multiply in every folding step for each position in Q , and requires $\tilde{O}(|Q| \cdot 2^{\log L} + \sigma + 2^{\log L} \log |\mathbb{F}|)$ gates and $\tilde{O}(1)$ depth.

6.5 Construction of the Verifier's Verification Circuit

We now wish to understand the size and depth of the verifier's verification circuit $V_{\Delta_c RR}$. $V_{\Delta_c RR}$ on reading the transcript of the interaction and the verifier decides whether to reject or to check Φ .

Claim 5. *If the transcript of the interaction is of length L , then $V_{\Delta_c RR}$, the verification circuit, has the following properties:*

- Size of $V_{\Delta_c RR}$ is $\tilde{O}(L)$ where we ignore $\text{poly}(m + \log L)$ factors.
- Depth of $V_{\Delta_c RR}$ is $\tilde{O}(\log L)$ where we ignore $\text{poly}(m + \log L)$ factors.

Proof. Observe that the verifier conducts the following checks:

- (1) $\tilde{O}(Tm)$ linearity checks to see if the verifier broke the instances into two claims correctly.
- (2) $\tilde{O}(m)$ PVAL emptiness checks which are run via a GKR.
- (3) $\tilde{O}(m)$ runs of Protocol 5 to apply the pairwise independent map and convert f_1 into $f_1 \circ \tau$.
- (4) $\tilde{O}(m)$ recalculations of input points j for recursing via evaluating low degree curves.
- (5) Checking if $\tilde{O}(T)$ low degree extensions are correct.

Now, step (1) is clearly linear in its input. Step (2) and (3) can be done in nearly linear input size and logarithmic depth due to Theorem 3 and Lemma 2. We may also need to check $\tilde{O}(m)$ low degree extensions which can be done in nearly linear of their input length as well. Thus, we just need to show that interpolation and evaluation of univariate polynomials can be done in nearly linear number of field operations. We know that this is possible due to our assumptions about our fields and multipoint interpolation and evaluation results (Corollary 10.8 and 10.12 in [vzGG13]).

Thus, all the checks can be done in nearly linear size of the input transcript. The depth of the overall circuit is also clearly $\tilde{O}(\log L)$. $\square$

Acknowledgement

The authors thank Ron Rothblum, Guy Rothblum, and the anonymous FOCS reviewers for their helpful comments and suggestions.

M.M.H. is partially supported by National Institute of Health (NIH) R01HG010959 (to B.B.). M.M.H., R.G. and Y.T.K. are supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR0011-25-C-0300 (to Y.T.K.). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Defense Advanced Research Projects Agency (DARPA).

References

- [ACFY24] Gal Arnon, Alessandro Chiesa, Giacomo Fenzi, and Eylon Yogev. STIR: Reed-solomon proximity testing with fewer queries. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part X*, volume 14929 of *LNCS*, pages 380–413. Springer, Cham, August 2024.
- [ACFY25] Gal Arnon, Alessandro Chiesa, Giacomo Fenzi, and Eylon Yogev. WHIR: Reed-solomon proximity testing with super-fast verification. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part IV*, volume 15604 of *LNCS*, pages 214–243. Springer, Cham, May 2025.
- [AHIV23] Scott Ames, Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkatasubramanian. Liger: lightweight sublinear arguments without a trusted setup. *DCC*, 91(11):3379–3424, 2023.
- [ALM⁺92] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and hardness of approximation problems. In *33rd FOCS*, pages 14–23. IEEE Computer Society Press, October 1992.
- [AS92] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs; A new characterization of NP. In *33rd FOCS*, pages 2–13. IEEE Computer Society Press, October 1992.
- [BCS16] Eli Ben-Sasson, Alessandro Chiesa, and Nicholas Spooner. Interactive oracle proofs. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 31–60. Springer, Berlin, Heidelberg, October / November 2016.
- [BFL90] László Babai, Lance Fortnow, and Carsten Lund. Non-deterministic exponential time has two-prover interactive protocols. In *31st FOCS*, pages 16–25. IEEE Computer Society Press, October 1990.
- [BFLS91] László Babai, Lance Fortnow, Leonid A. Levin, and Mario Szegedy. Checking computations in polylogarithmic time. In *23rd ACM STOC*, pages 21–31. ACM Press, May 1991.
- [BGKW88] Michael Ben-Or, Shafi Goldwasser, Joe Kilian, and Avi Wigderson. Multi-prover interactive proofs: How to remove intractability assumptions. In *20th ACM STOC*, pages 113–131. ACM Press, May 1988.

- [BH08] Alex Brodsky and Shlomo Hoory. Simple permutations mix even better. *Random Structures & Algorithms*, 32(3):274–289, 2008.
- [BKP⁺24] Nir Bitansky, Chethan Kamath, Omer Paneth, Ron D. Rothblum, and Prashant Nalini Vasudevan. Batch proofs are statistically hiding. In Bojan Mohar, Igor Shinkar, and Ryan O’Donnell, editors, *56th ACM STOC*, pages 435–443. ACM Press, June 2024.
- [BKS18] Eli Ben-Sasson, Swastik Kopparty, and Shubhangi Saraf. Worst-case to average case reductions for the distance to a code. In *Proceedings of the 33rd Computational Complexity Conference, CCC ’18*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018.
- [BM88] László Babai and Shlomo Moran. Arthur-Merlin games: A randomized proof system, and a hierarchy of complexity classes. *J. Comput. Syst. Sci.*, 36(2):254–276, 1988.
- [BSCI⁺23] Eli Ben-Sasson, Dan Carmon, Yuval Ishai, Swastik Kopparty, and Shubhangi Saraf. Proximity gaps for Reed-Solomon codes. *Journal of the ACM*, 70(5), 2023.
- [Coo71] Stephen A Cook. The complexity of theorem-proving procedures. In *3rd ACM STOC*, pages 151–158, 1971.
- [EKR04] Funda Ergün, Ravi Kumar, and Ronitt Rubinfeld. Fast approximate probabilistically checkable proofs. *Information and Computation*, 189(2):135–159, 2004.
- [FGL⁺91] Uriel Feige, Shafi Goldwasser, László Lovász, Shmuel Safra, and Mario Szegedy. Approximating clique is almost NP-complete (preliminary version). In *32nd FOCS*, pages 2–12. IEEE Computer Society Press, October 1991.
- [GHKO25] William Gay, William He, Nicholas Kocurek, and Ryan O’Donnell. Pseudorandomness properties of random reversible circuits. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part I*, volume 16000 of *LNCs*, pages 651–678. Springer, Cham, August 2025.
- [GHP25] Lucas Gretta, William He, and Angelos Pelecanos. More efficient approximate k -wise independent permutations from random reversible circuits via log-sobolev inequalities. In Yossi Azar and Debmalya Panigrahi, editors, *36th SODA*, pages 5582–5598. ACM-SIAM, January 2025.
- [GKR15] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: Interactive proofs for muggles. *Journal of the ACM*, 62(4):27:1–27:64, 2015.
- [GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989.
- [Gow96] W. T. Gowers. An almost m -wise independent random permutation of the cube. *Combinatorics, Probability and Computing*, 5(2):119–130, 1996.
- [GR15] Tom Gur and Ron D. Rothblum. Non-interactive proofs of proximity. In Tim Roughgarden, editor, *ITCS 2015*, pages 133–142. ACM, January 2015.

- [GR17] Tom Gur and Ron D. Rothblum. A hierarchy theorem for interactive proofs of proximity. In Christos H. Papadimitriou, editor, *ITCS 2017*, volume 4266, pages 39:1–39:43, 67, January 2017. LIPIcs.
- [GS86] Shafi Goldwasser and Michael Sipser. Private coins versus public coins in interactive proof systems. In *18th ACM STOC*, pages 59–68. ACM Press, May 1986.
- [HMMR04] Shlomo Hoory, Avner Magen, Steven Myers, and Charles Rackoff. Simple permutations mix well. In Josep Díaz, Juhani Karhumäki, Arto Lepistö, and Donald Sannella, editors, *ICALP 2004*, volume 3142 of *LNCS*, pages 770–781. Springer, Berlin, Heidelberg, July 2004.
- [JKKZ21] Ruta Jawale, Yael Tauman Kalai, Dakshita Khurana, and Rachel Yun Zhang. SNARGs for bounded depth computations and PPAD hardness from sub-exponential LWE. In Samir Khuller and Virginia Vassilevska Williams, editors, *53rd ACM STOC*, pages 708–721. ACM Press, June 2021.
- [KNR05] Eyal Kaplan, Moni Naor, and Omer Reingold. Derandomized constructions of k -wise (almost) independent permutations. In Chandra Chekuri, Klaus Jansen, José D. P. Rolim, and Luca Trevisan, editors, *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques*, pages 354–365. Springer Berlin Heidelberg, 2005.
- [KR08] Yael Tauman Kalai and Ran Raz. Interactive PCP. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfssdóttir, and Igor Walukiewicz, editors, *ICALP 2008, Part II*, volume 5126 of *LNCS*, pages 536–547. Springer, Berlin, Heidelberg, July 2008.
- [Lev73] Leonid A Levin. Universal sequential search problems. *Problems of information transmission*, 9(3):265–266, 1973.
- [LFKN92] Carsten Lund, Lance Fortnow, Howard J. Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *Journal of the ACM*, 39(4):859–868, 1992.
- [RR20] Guy N. Rothblum and Ron D. Rothblum. Batch verification and proofs of proximity with polylog overhead. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part II*, volume 12551 of *LNCS*, pages 108–138. Springer, Cham, November 2020.
- [RRR16] Omer Reingold, Guy N. Rothblum, and Ron D. Rothblum. Constant-round interactive proofs for delegating computation. In Daniel Wichs and Yishay Mansour, editors, *48th ACM STOC*, pages 49–62. ACM Press, June 2016.
- [RRR18] Omer Reingold, Guy N. Rothblum, and Ron D. Rothblum. Efficient batch verification for UP. In *Proceedings of the 33rd Computational Complexity Conference, CCC '18*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018.
- [RVW13] Guy N. Rothblum, Salil P. Vadhan, and Avi Wigderson. Interactive proofs of proximity: delegating computation in sublinear time. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 793–802. ACM Press, June 2013.

- [Sch80] Jacob T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM*, 27(4):701–717, 1980.
- [Sha92] Adi Shamir. $IP = PSPACE$. *Journal of the ACM*, 39(4):869–877, 1992.
- [Vad00] Salil P. Vadhan. On transformation of interactive proofs that preserve the prover’s complexity. In *32nd ACM STOC*, pages 200–207. ACM Press, May 2000.
- [vzGG13] Joachim von zur Gathen and Juergen Gerhard. *Modern Computer Algebra*. Cambridge University Press, 3 edition, 2013.
- [Zei24] Hadas Zeilberger. Khatam: Reducing the communication complexity of code-based SNARKs. Cryptology ePrint Archive, Report 2024/1843, 2024.
- [Zip79] Richard Zippel. Probabilistic algorithms for sparse polynomials. In *EUROSAM ’79: International Symposium on Symbolic and Algebraic Computation*, volume 72 of *Lecture Notes in Computer Science*, pages 216–226. Springer, 1979.