# eSTARK: Extending STARKs with Arguments 

Héctor Masip-Ardevol ${ }^{1}$, Marc Guzmán-Albiol ${ }^{1}$, Jordi Baylina-Melé ${ }^{2}$ and Jose Luis Muñoz-Tapia ${ }^{1}$<br>${ }^{1}$ Universitat Politècnica de Catalunya (UPC), Spain hector.masip@upc.edu<br>marc.guzman.albiol@upc.edu<br>jose.luis.munoz@upc.edu<br>${ }^{2}$ Polygon zkEVM, Switzerland<br>jordi@baylina.cat

March 31, 2023


#### Abstract

STARK is a widely used transparent proof system that uses low-degree tests for proving the correctness of a computer program. STARK consumes an intermediate representation known as AIR that is more appropriate for programs with a relatively short and structured description. However, an AIR is not able to succinctly express non-equality constraints, leading to the incorporation of unwanted polynomials. We present the eSTARK protocol, a new probabilistic proof that generalizes the STARK family through the introduction of a more generic intermediate representation called eAIR. We describe eSTARK in the polynomial IOP model, which combines the optimized version of the STARK protocol with the incorporation of three arguments into the protocol. We also explain various techniques that enhance the vanilla STARK complexity, including optimizations applied to polynomial computations, and analyze the tradeoffs between controlling the constraint degree either at the representation of the AIR or inside the eSTARK itself.


## Contents

1 Introduction ..... 3
1.1 Organization ..... 5
2 Preliminaries ..... 5
2.1 Notation ..... 5
2.2 Interactive Oracle Proofs and STARKs ..... 6
2.3 FRI ..... 7
2.4 Vanilla STARK ..... 12
2.5 Arguments ..... 17
3 Our Techniques ..... 20
3.1 Commiting to Multiple Polynomial at Once ..... 20
3.2 Preprocessed Polynomials and Public Values ..... 21
3.3 Adding Selected Vector Arguments ..... 21
3.4 On the Quotient Polynomial ..... 27
3.5 Controlling the Constraint Degree wih Intermediate Polynomials ..... 27
3.6 FRI Polynomial Computation ..... 30
4 Our eSTARK Protocol ..... 31
4.1 Extended Algebraic Intermediate Representation (eAIR) ..... 31
4.2 The Setup Phase ..... 32
4.3 Our IOP for eAIR ..... 32
4.4 From a STIK to a Non-Interactive STARK ..... 36
4.5 Full Protocol Description ..... 36
5 Conclusions ..... 39
A Vanilla STARK Description ..... 43

## 1 Introduction

Since the term was coined and the first example was given in 1985 by Goldwasser et al. [GMR85], probabilistic proofs leverages a set of protocols, such as Interactive Proofs (IPs) [GMR85] and Probabilistic Checkable Proofs [AS92], that enable one party, known as the prover, to provide a guarantee to another party, known as the verifier, that the former performed a computation requested by the latter correctly. In practical terms, probabilistic proofs enable a computationally weak device to monitor the execution of a powerful but untrusted powerful server, allowing it to outsource a program execution to the server.

Probabilistic proofs can be even more useful when they also possess a property known as zero-knowledge [GMR85], which intuitively means that the proof reveals nothing but its validity. Equipped with zero-knowledge, probabilistic proofs let the prover demonstrate knowledge of a secret input satisfying a certain statement without revealing any information about the input beyond the validity of the statement. For example, a probabilistic proof satisfying zero-knowledge can be used to prove possession of a secret key associated with a particular public key or the knowledge of the preimage of a particular hash digest; without revealing any information about the secret key or the preimage. The latter could be used to log in to a website without typing a password, by simply sending proof of possession of the valid password.

Slightly more formal, our scenario is one in which the verifier, denoted as $\mathcal{V}$, sends a program description $f$ and an input $x$ for that program to the prover, denoted as $\mathcal{P}$. Then, $\mathcal{P}$ computes and returns the execution of program $f$ on input $x$, i.e., $y=f(x)$, along with a proof $\pi$ that the output $y$ is correct and consistent with the description $f$ and the input $x$. The proof $\pi$ should satisfy the following properties:

- Completeness. If $\mathcal{P}$ is honest, which means that he knows a pair $(x, y)$ such that the claim $f(x)=y$ is true and follows the protocol properly, then $\mathcal{P}$ should be able to compute a proof that convinces the verifier $\mathcal{V}$ of the validity of the claim.
- Soundness. A malicious prover $\mathcal{P}^{*}$ should be able to produce a proof that convinces the verifier $\mathcal{V}$ of a statement of a false claim, i.e., that $f(x) \neq y$, with negligible probability.
Completeness means that an honest verifier will always accept a proof if it is generated by an honest prover. Soundness deals with the verifier not accepting false proofs from malicious provers. These two properties protect the verifier against malicious provers. Moreover, the verification of the proof $\pi$ should be much more efficient than rerunning the program $f$ on $x$.

To also make probabilistic proofs privacy-preserving, $\mathcal{P}$ can provide his private input $w$ to the computation, known as the witness. Thus, $f$ now is written as a function of two inputs $(x, w)$ such that $f(x, w)=y$. If at the end of the protocol, $\mathcal{V}$ is convinced that the statement $y=f(x, w)$ is true without learning anything about $w$, then the scheme satisfies the zero-knowledge property and is typically called a Zero-Knowledge Proof (ZKP).

More formally, a ZKP satisfies completeness, soundness and zero-knowledge defined as follows:

- Zero-Knowledge: After interacting with $\mathcal{P}$ about the claim $f(x, w)=y, \mathcal{V}$ should learn nothing about the witness $w$ and still be convinced of the validity of the claim.
Compared with completeness and soundness, the zero-knowledge property guarantees the privacy of the prover's secret information against malicious verifiers.
zk-SNARKs. One family of ZKPs that has become of important consideration in recent years both in theory and in practice are zero-knowledge Succinct Non-interactive ARguments of Knowledge (zk-SNARKs) [BCCT11]. At a high level, zk-SNARKs is a generic
term to refer to ZKPs whose cryptographic assumptions and computational complexity are reasonable for practical use. More specifically, the acronym zk-SNARKs refers to probabilistic proofs such that:
- zk: Satisfy the zero-knowledge property, even though it is optional.
- Succinct: Have size sublinear compared to the size of the statement or the witness, i.e., the size of the computation itself.
- Non-Interactive: Do not require rounds of interaction between a particular prover and verifier and are publicly verifiable.
- Argument: Assume that are generated by provers having at most polynomial computational resources. In particular, provers are not able to break standard cryptographic assumptions.
- Of Knowledge: Satisfy knowledge soundness rather than plain soundness. That is, the proof demonstrates that the prover owns a witness for the statement, not only its existence.

In the literature it can be found many examples of zk-SNARKs [Gro16], $\left[\mathrm{BBB}^{+} 18\right]$, [GWC19], $\left[\mathrm{CHM}^{+} 19\right]$, each of them based on distinct cryptographic primitives and coming with different tradeoffs.

An application that makes use of zk-SNARKs as its core technology is Zcash [Wil16], a public blockchain based on bitcoin that uses these proofs in its core protocol for verifying that private transactions, named shielded transactions, have been computed correctly while revealing anything else to an external observer. More applications range from using zk-SNARKs to fight disinformation [KHSS22], to using zk-SNARKs for proving that convolutional neural networks have made predictions correctly without revealing the model itself [LXZ21].
zk-STARKs. Although zk-SNARKs have numerous applications, they suffer from a severe limitation: the requirement of a trusted party that generates random public parameters, known as the trusted setup, for the system to work that eliminates any knowledge of the randomness used in the process (typically known as the toxic waste). In fact, if a malicious party obtains access to that toxic waste then this party can easily forge false proofs.

As a result, systems that remove the previous requirement, and are therefore of transparent nature, have been developed. The most prominent family of transparent ZKPs are zero-knowledge Scalable Transparent ARguments of Knowledge (zk-STARKs) [BBHR18b]. Apart from satisfying the same properties of a zk-SNARKs, zk-STARKs are required to satisfy two additional properties:

- Scalability. Proving time is at most quasilinear and verification time is at most logarithmic in the size of the statement.
- Transparency. They do not require any trusted setup, i.e., any randomness used by transparent frameworks are public coins.

Examples of zk-STARKs are also present in the literature [WTs ${ }^{+}$17], [COS19], [Set19].
Intermediate Representations. Given a program $f$, to compute a proof of the validity of the execution of $f$, we typically agree upon an equivalent model amenable to probabilistic proof and then express $f$ as an algebraic constraint satisfaction problem of that model. This latter step is often called arithmetization, the output of such being an intermediate representation. For instance, (unstructured) arithmetic circuits are often encoded as a
system of quadratic equations over a finite field known as Rank-1 Constraint System (R1CS). In contrast, "structured" arithmetic circuits are often encoded as a set of repeated polynomial computations known as Algebraic Intermediate Representation (AIR). Here, by structured we are referring to circuits that are designed with a layered architecture, where each layer contains a small circuit that is replicated in every layer. In fact, an AIR is the standard intermediate representation consumed by a STARK [Sta21].

Since we will be dealing with structured arithmetic circuits, let us introduce an AIR more formally. An AIR A over a field $\mathbb{F}$ is defined by a set of multivariate constraint polynomials $C_{i} \in \mathbb{F}\left[X_{1}, \ldots, X_{2 M}\right]$. The idea is that $f$, represented as a set of univariate polynomials $p_{1}, \ldots, p_{2 M} \in \mathbb{F}[X]$, will be said to satisfy $A$ if and only if the following polynomial constraints:

$$
C_{i}\left(p_{1}(X), \ldots, p_{2 M}(X)\right)=0
$$

are satisfied over a subset of $\mathbb{F}$. Therefore, we translate the computation of $f$ into the satisfaction of every polynomial $C_{i}$ of an (equivalent) AIR A, being the latter ready to be consumed by an appropriate STARK.

Note that the definition of an AIR only captures polynomial constraints defined through equality. So, for instance, we are not able to directly express that the evaluations of a polynomial over the given subset are lower than a predetermined upper bound. For this, we would need to express this inequality as a set of equality with the incorporation of new polynomials such that the new constraints satisfy the equality if and only if the inequality is satisfied.

In this paper, we enlarge the type of constraints that can be captured by an AIR with the addition of non-identity constraints satisfying some requirements. We call the resulting constraint system an extended AIR (eAIR). Details will be explained in Section 4.1. Moreover, we provide a particular probabilistic proof that is ready to consume an eAIR, and because it will be defined as a generalization of a STARK, we will refer to it as an extended STARK (eSTARK).

### 1.1 Organization

This paper is organized as follows. In Section 2 we recall basic definitions related to the IOP model, introduce the non-identity constraints from which we enlarge the AIR expressiveness and provide an in-depth overview of STARKs. Moreover, we explain the low-degree test FRI that will be later on used by our eSTARK. In Section 3 we not only exhibit the main differences between STARKs and our eSTARK, but we also explain how to introduce the new type of constraints into the STARK protocol. We put everything together in Section 4 and proof that the eSTARK protocol is sound. Moreover, in Section 4.4 we provide a full description of our eSTARK in the practical realm. We conclude in Section 5.

## 2 Preliminaries

### 2.1 Notation

We denote by $\mathbb{F}$ to a finite field of prime order $p$ and $\mathbb{F}^{*}$ to its respective multiplicative group and define $k$ to be the biggest non-negative integer such that $2^{k} \mid(p-1)$. We also write $\mathbb{K}$ to denote a finite field extension of $\mathbb{F}$, of size $p^{e}, e \geq 2$. Furthermore, we write $\mathbb{F}[X]$ (resp. $\mathbb{K}[X]$ ) for the ring of polynomials with coefficients over $\mathbb{F}$ (resp. $\mathbb{K}$ ) and write $\mathbb{F}_{<d}[X]$ (resp. $\mathbb{K}_{<d}[X]$ ) to denote the set of polynomials of degree lower than $d$.

Although all the protocols presented in this article work over any prime order, we fix our attention on fields that contain a multiplicative subgroup of size a large power of two. This restriction is crucial to achieving a practical protocol.

Given a cyclic subgroup $S$ of $\mathbb{F}^{*}$, we denote by $L_{i} \in \mathbb{F}_{<|S|}[X]$ the $i$-th Lagrange polynomial for $S$. That is, $L_{i}$ satisfies $L_{i}\left(g^{i}\right)=1$ and $L_{i}\left(g^{j}\right)=0$ for $j \neq i$, where $g$ is used here to denote the generator of $S$. Moreover, we denote by $Z_{S}(X)=X^{|S|}-1 \in \mathbb{F}[X]$ to the polynomial that vanishes only over $S$ and call it the vanishing polynomial over $S$.

We denote by $G$ to a cyclic subgroup of $\mathbb{F}^{*}$ with order $n$ satisfying $n \mid 2^{k}$ and $1<n<2^{k}$, and let $g \in \mathbb{F}$ denote the generator of $G$. Similarly, we denote by $H$ to a nontrivial coset of a cyclic subgroup of $\mathbb{F}^{*}$ with order $m$ satisfying $m \mid 2^{k}$ and $n<m$.

Given a set of polynomials $p_{1}, p_{2}, \ldots, p_{N} \in \mathbb{K}[X]$ we denote by $\operatorname{MTR}\left(p_{1}, \ldots, p_{N}\right)$ to the Merkle root obtained after computing a Merkle tree [Mer88] whose leaves are the evaluations of $p_{1}, \ldots, p_{N}$ over the domain $H$. Additionally, given a set of elements $x_{1}, x_{2}, \ldots, x_{N} \in \mathbb{K}$, we also use $\operatorname{MTP}\left(x_{1}, \ldots, x_{N}\right)$ to denote the Merkle tree path (i.e., the Merkle proof) computed from the leaf containing these elements. If $X=\left\{x_{1}, \ldots, x_{N}\right\}$, then we use $\operatorname{MTP}(X)$ as a shorthand for $\operatorname{MTP}\left(x_{1}, \ldots, x_{N}\right)$.

Finally, in the description of the protocols, we use $\mathcal{P}$ to denote the prover entity and $\mathcal{V}$ to denote the verifier entity.

### 2.2 Interactive Oracle Proofs and STARKs

In this section, we define a polynomial IOP [BCS16] and the standard security notions associated with this model. Moreover, we introduce a popular polynomial IOP family of protocols known as STIK [BBHR19] and explain how a STIK can be compiled into a STARK.

Definition 1 (Polynomial IOP). Given a function $F$, a public coin polynomial interactive oracle proof (IOP) for $F$ is an interactive protocol between two parties, the prover $\mathcal{P}$ and the verifier $\mathcal{V}$, that comprises $k$ rounds of interaction. $\mathcal{P}$ is given an input $w$ and both $\mathcal{P}$ and $\mathcal{V}$ are given a common input $x$. At the start of the protocol, $\mathcal{P}$ provides to $\mathcal{V}$ a value $y$ and claims to him the existence of a $w$ satisfying $y=F(x, w)$.

In the $i$-th round, $\mathcal{V}$ sends a uniformly and independently random message $\alpha_{i}$ to $\mathcal{P}$. Then $\mathcal{P}$ replies with a message of one of the two following forms: (1) a string $m_{i}$ that $\mathcal{V}$ reads in full, or (2) an oracle to a polynomial $f_{i}$ that $\mathcal{V}$ can query (via random access) after the $i$-th round of interaction. At the end of the protocol, $\mathcal{V}$ outputs either accept or reject, indicating whether $\mathcal{V}$ accepts $\mathcal{P}$ 's claim.

The security notions for IOPs are similar to the security notions of other preceding models (e.g., interactive proofs). In the following definition, probabilities are taken over the internal randomness of $\mathcal{V}$.

Definition 2 (Completeness and (Knowledge) Soundness). We say that a polynomial IOP has perfect completeness and soundness error at most $\varepsilon_{s}$ if the following two conditions hold.

- Perfect Completeness. For every $x$ and every prover $\mathcal{P}$ sending a value $y$ satisfying $y=F(x, w)$ at the start of the protocol, it holds that:

$$
\operatorname{Pr}[\mathcal{V}(x, \mathcal{P}(x, w))=\text { accept }]=1
$$

where $\mathcal{V}(x, \mathcal{P}(x, w))$ denotes the $\mathcal{V}$ 's output after interacting with the prover on input $x$.

- Soundness. For every $x$ and every prover $\mathcal{P}^{*}$ sending a value $y$ at the start of the protocol, if it holds that:

$$
\operatorname{Pr}\left[\mathcal{V}\left(x, \mathcal{P}^{*}(x, w)\right)=\text { accept }\right] \geq \varepsilon_{s}
$$

then $y$ satisfies $y=F(x, w)$.

If the next condition holds as well, we say that the polynomial IOP has knowledge soundness error at most $\varepsilon_{k s}$.

- Knowledge Soundness. There exists a probabilistic polynomial-time algorithm $\mathcal{E}$, known as the knowledge extractor, such that for every $x$ and every prover $\mathcal{P}^{*}$ sending a value $y$ at the start of the protocol if it holds that:

$$
\operatorname{Pr}\left[\mathcal{V}\left(x, \mathcal{P}^{*}(x, w)\right)=\text { accept }\right] \geq \varepsilon_{k s},
$$

then $\mathcal{E}\left(x, \mathcal{P}^{*}(x, w)\right)=w$ and $y$ satisfies $y=F(x, w)$.
In words, soundness guarantees that a malicious prover cannot succeed with probability greater than $\varepsilon_{s}$ on the "existence" claim of $w$; whereas knowledge soundness guarantees that a malicious prover cannot succeed with probability greater than $\varepsilon_{k s}$ claiming "possession" of $w$ satisfying $y=F(x, w)$. A polynomial IOP satisfying knowledge soundness is known as a polynomial IOP of knowledge.

Naturally, knowledge soundness implies soundness. Surprisingly, the converse is also true for polynomial IOPs (see, e.g., Lemma 2.3 in [CBBZ22]). This means that proving the soundness of a polynomial IOP is sufficient for achieving knowledge soundness.
Definition 3 (STIK). A Scalable Transparent polynomial IOP of Knowledge (STIK) is a polynomial IOP that moreover satisfies the following properties:

- Transparent. They do not require a trusted setup before the execution of the protocol. This setup constitutes a single point of failure and could be exploited by powerful parties to forge false proofs.
- Doubly Scalable. The verifier runs in time $\mathcal{O}(\log (n))$ and the prover runs in time $\mathcal{O}(n \log (n))$, where $n$ informally denotes the size of the computation $F$.

When STIKs get instantiated for practical deployment, they result in protocols in which the prover is assumed to be computationally bounded. Protocols of such kind are known as argument systems (in contrast to proof systems), and consequently, instantiation of STIKs results in protocols satisfying soundness only against adversaries running in polynomial time.

Definition 4 (STARK). A scalable transparent argument of knowledge (STARK) is a realization of a STIK through a family of collision-resistant hash functions. More specifically, polynomial oracles sent from the prover to the verifier in the underlying polynomial IOP are substituted by Merkle roots (computed from polynomial evaluations over a set); and whenever the verifier asks queries to a polynomial $f$ at $v$, the prover answers with $f(v)$ together with the Merkle path associated with it.

Finally, we briefly explain how to remove the interaction between a specific prover and verifier of protocols and make them publicly verifiable. The Fiat-Shamir heuristic [FS87] can be used to compile a STIK into a non-interactive argument of knowledge in the random oracle model by substituting verifier's messages for queries to the random oracle on input the previous prover's messages until that point. The random oracle is modeled in practice by a cryptographic hash function. Specific details on this compilation will be explained in Section 4.4. Therefore, a realization of a non-interactive STIK is called a non-interactive STARK (but we abuse notation and refer to both as STARKs).

### 2.3 FRI

Fast Reed-Solomon Interactive Oracle Proof of Proximity (FRI) [ $\mathrm{BCI}^{+}$20] is a protocol for proving that a function $f: H \rightarrow \mathbb{F}$ is close to a polynomial of low degree $d$. Here,
by low degree, we mean that $d \ll|H|$. The FRI protocol consists of two phases. In the first phase, known as the commit phase, the prover commits to (via Merkle trees) a series of functions generated from $f$ and random elements $v_{0}, v_{1}, \ldots$ from $\mathbb{K}$ provided by the verifier at each round. Then, in the second phase, known as the query phase, the prover provides a set of evaluations of the previously committed functions at a point randomly chosen by the verifier. Following, we provide more details about how each phase works.

## The Commit Phase

Let's denote by $p_{0}$ the function $f$ of interest and assume for the simplicity of the exposition that the prover is honest (i.e., $p_{0}$ is a polynomial of low degree). In the commit phase, the polynomial $p_{0}$ is split into two other polynomials $g_{0,1}, g_{0,2}: H^{2} \rightarrow \mathbb{K}$ of degree lower than $d / 2$. These two polynomials satisfy the following relation with $p_{0}$ :

$$
\begin{equation*}
p_{0}(X)=g_{0,1}\left(X^{2}\right)+X \cdot g_{0,2}\left(X^{2}\right) \tag{1}
\end{equation*}
$$

Then, the verifier sends to the prover a uniformly sampled $v_{0} \in \mathbb{K}$, and asks the prover to commit to the polynomial:

$$
p_{1}(X):=g_{0,1}(X)+v_{0} \cdot g_{0,2}(X)
$$

Note that $p_{1}$ is a polynomial of degree less than $d / 2$ and the commitment of $p_{1}$ is not over $H$ but over $H^{2}=\left\{x^{2}: x \in H\right\}$, which is of size $|H| / 2$.

The prover then continues by splitting $p_{1}$ into $g_{1,1}$ and $g_{1,2}$ of degree lower than $d / 4$, then constructing $p_{2}$ with a uniformly sampled $v_{1} \in \mathbb{K}$ sent by the verifier. Again, $p_{2}$ is of degree $d / 2^{2}$ and committed over $H^{2^{2}}=\left\{x^{2}: x \in H^{2}\right\}$, whose size is $|H| / 2^{2}$. The whole derivation of $p_{i+1}$ from $p_{i}$ is often known as split-and-fold due to the prover splitting the initial polynomial into two and then folding it into one using a random value.

The previous process gets repeated a total of $k=\log _{2}(d)$ times, the point at which $\operatorname{deg}\left(p_{k}\right)=0$ and the prover sends a constant $p_{k}$, representing a polynomial of degree lower than 1 , to the verifier.

| $\mathcal{P}\left(p_{0}, d, G, H, \mathbb{F}, \mathbb{K}\right)$ |  | $\mathcal{V}(d, G, H, \mathbb{F}, \mathbb{K})$ |
| :---: | :---: | :---: |
|  | $\operatorname{MTR}\left(p_{0}\right)$ |  |
|  | $v_{0}$ |  |
|  | $\operatorname{MTR}\left(p_{1}\right)$ |  |
|  | $v_{1}$ |  |
|  |  |  |
|  | $\operatorname{MTR}\left(p_{k-1}\right)$ |  |
|  | $v_{k-1}$ |  |
|  | $p_{k}$ |  |

Figure 1: Skeleton description of FRI's commit phase.

## The Query Phase

In the query phase, the verifier sends a uniformly sampled $r \in H$ to the prover and queries the evaluations $p_{0}(r), p_{0}(-r)$ and $p_{1}\left(r^{2}\right)$. From $p_{0}(r)$ and $p_{0}(-r)$ the verifier computes
$p_{1}\left(r^{2}\right)$ and checks that the computed value matches with the third value $p_{1}\left(r^{2}\right)$ provided by the prover. To obtain $p_{1}\left(r^{2}\right)$ from $p_{0}(r)$ and $p_{0}(-r)$, the verifier first solves the following system of linear equations for $g_{0,1}\left(r^{2}\right), g_{0,2}\left(r^{2}\right)$ :

$$
\begin{aligned}
p_{0}(r) & =g_{0,1}\left(r^{2}\right)+r \cdot g_{0,2}\left(r^{2}\right), \\
p_{0}(-r) & =g_{0,1}\left(r^{2}\right)-r \cdot g_{0,2}\left(r^{2}\right),
\end{aligned}
$$

and then computes:

$$
p_{1}\left(r^{2}\right)=g_{0,1}\left(r^{2}\right)+v_{0} \cdot g_{0,2}\left(r^{2}\right)
$$

The verifier continues by querying for $p_{1}\left(-r^{2}\right)$ and $p_{2}\left(r^{4}\right)$. From $p_{1}\left(r^{2}\right)$ and $p_{1}\left(-r^{2}\right)$ computes $p_{2}\left(r^{4}\right)$ as before and checks that the computed value is consistent with $p_{2}\left(r^{4}\right)$. Each step locally checks the consistency between each pair ( $p_{i}, p_{i+1}$ ). The verifier continues in this way until it reaches the value of the constant $p_{k}$. The verifier checks that the value sent by the prover is indeed equal to the value that the verifier computed from the queries up until $p_{k-1}$. To fully ensure correctness, the prover must accompany the evaluations that he sends with a claim of their existence (via Merkle tree paths).


Figure 2: Skeleton description of one iteration of FRI's query phase.
Upon the completion of this process, the verifier has a first confirmation that the polynomials committed in the commit phase $p_{0}, p_{1}, \ldots, p_{k}$ are consistent with each other.

Finally, to achieve the required bounds for the soundness of the protocol, the query phase is repeated multiple times. We give the specific soundness bound of a more generic version of FRI in Theorem 1. The full skeleton description of FRI can be found in Figure 3a.

## The Batched FRI Protocol

In this version of FRI, the prover wants to prove closeness to low-degree polynomials of a set of functions $f_{0}, f_{1}, \ldots, f_{N}: H \rightarrow \mathbb{F}$ at once. We could run the FRI protocol for every function $f_{i}$ in parallel, but there is a more efficient way proposed in Section 8.2 of $\left[\mathrm{BCI}^{+} 20\right]$. In the batched FRI protocol, the prover instead applies the FRI protocol directly to a random linear combination of the function $f_{i}$. More specifically, assuming the prover has committed to functions $f_{0}, f_{1}, \ldots, f_{N}$ and the verifier has sent a uniformly sampled value $\varepsilon \in \mathbb{K}$, the prover computes the function:

$$
\begin{equation*}
f(X):=f_{0}(X)+\sum_{i=1}^{N} \varepsilon^{i} f_{i}(X) \tag{2}
\end{equation*}
$$

and applies the FRI protocol to it.
Remark 1. In $\left[\mathrm{BCI}^{+} 20\right]$, they compute $f$ as $f_{0}(X)+\sum_{i=1}^{N} \varepsilon_{i} \cdot f_{i}(X)$ instead, i.e., they use a random value $\varepsilon_{i} \in \mathbb{K}$ per function $f_{i}$ instead of powers of a single one $\varepsilon$. Even if secure, the soundness bound of this alternative version is linearly increased by the number of functions $N$, so we might assume from now on that $N$ is sublinear in $|\mathbb{K}|$ to ensure the security of protocols.

## The Batched Consistency Check

As an extra check in the batched version, the verifier needs to ensure the correct relationship between functions $f_{0}, \ldots, f_{N}$ and the first FRI polynomial $p_{0}=f$. The verifier will use the evaluations of $p_{0}$ it received from the prover in each FRI query phase invocation. To allow for this check, the prover also sends the evaluations of functions $f_{0}, f_{1}, \ldots, f_{N}$ at both $r$ and $-r$ so that the verifier can check that:

$$
\begin{gathered}
p_{0}(r)=f_{0}(r)+\sum_{i=1}^{N} \varepsilon^{i} f_{i}(r), \\
p_{0}(-r)=f_{0}(-r)+\sum_{i=1}^{N} \varepsilon^{i} f_{i}(-r),
\end{gathered}
$$

i.e., a local consistency check between $f_{0}, \ldots, f_{N}$ and $p_{0}$. The prover accompanies the newly sent evaluations with their respective Merkle tree path. The full skeleton description of batched FRI can be found in Figure 3b.

Similarly to the non-batched FRI protocol, both the query phase and the batched consistency check gets repeated multiple times to ensure the protocol is sound. More precisely, the soundness error is shown in the following theorem, which is a (somewhat informally) adaptation of Theorems $7.2,8.3$ from $\left[\mathrm{BCI}^{+} 20\right]$.

Theorem 1 (Batched FRI Soundness). Let $f_{0}, f_{1}, \ldots, f_{N}$ be functions defined over $H$ and let $m \geq 3$ be an integer. Suppose a batched FRI prover that interacts with a batched FRI verifier causes it to accept with probability greater than:
$\varepsilon_{F R I}=\varepsilon_{C}+\varepsilon_{Q}^{s}=\left(N \cdot \frac{\left(m+\frac{1}{2}\right)^{7}}{3 \rho^{3 / 2}} \cdot \frac{|H|^{2}}{|\mathbb{K}|}+N \cdot \frac{(2 m+1) \cdot(|H|+1)}{\sqrt{\rho}} \cdot \frac{\sum_{i=0}^{k-1} a_{i}}{|\mathbb{K}|}\right)+\left(\sqrt{\rho}\left(1+\frac{1}{2 m}\right)\right)^{s}$,
where $a_{i}=\left|H^{2^{i}}\right| /\left|H^{2^{i+1}}\right|$ is the ratio ${ }^{1}$ between consecutive prover messages in the commit phase, $\rho=|G| /|H|$ is the rate of the code, $\varepsilon_{C}$ and $\varepsilon_{Q}$ are respectively the soundness error for the commit and the query phases and $s$ is the number of repetitions of the query phase.

Then, functions $f_{0}, f_{1}, \ldots, f_{N}$ are close to polynomials of degree lower than $n$.

In [BBHR18a] it is shown that for the FRI protocol to achieve security parameter $\lambda$ (i.e., $\varepsilon_{\text {FRI }} \leq 2^{-\lambda}$ ), at least $s \geq \lambda / \log _{2} \rho^{-1}$ many queries are needed, and Theorem 1 shows that if $|\mathbb{K}| \gg|H|^{2}$ then $s \approx 2 \lambda / \log _{2} \rho^{-1}$.

[^0]| $\mathcal{P}\left(p_{0}, \mathrm{pp}\right)$ |  | $\mathcal{V}(\mathrm{pp})$ |
| :---: | :---: | :---: |
|  | $\operatorname{MTR}\left(p_{0}\right)$ |  |
|  | $v_{0}$ |  |
|  | $\operatorname{MTR}\left(p_{1}\right)$ |  |
|  | $v_{1}$ |  |
|  |  |  |
|  | $v_{k-1}$ |  |
|  | $p_{k}$ |  |
|  | $h_{1}$ |  |
|  | $\left\{p_{0}\left(h_{1}\right), \ldots, p_{k-1}\left(h_{1}^{2 k-1}\right)\right\}$ |  |
|  | $\left\{\operatorname{MTP}\left(p_{0}\left(h_{1}\right)\right), \ldots, \operatorname{MTP}\left(p_{k-1}\left(h_{1}^{2^{k-1}}\right)\right)\right\}$ |  |
|  | $\left\{p_{0}\left(-h_{1}\right), \ldots, p_{k-1}\left(-h_{1}^{2^{k-1}}\right)\right\}$ |  |
|  | $\overrightarrow{\left\{\operatorname{MTP}\left(p_{0}\left(-h_{1}\right)\right), \ldots, \operatorname{MTP}\left(p_{k-1}\left(-h_{1}^{2^{k-1}}\right)\right)\right\}}$ |  |
|  |  |  |
|  | $h_{s}$ |  |
|  | $\left\{p_{0}\left(h_{s}\right), \ldots, p_{k-1}\left(h_{s}^{2^{k-1}}\right)\right\}$ |  |
|  | $\xrightarrow[\left\{\operatorname{MTP}\left(p_{0}\left(h_{s}\right)\right), \ldots, \operatorname{MTP}\left(p_{k-1}\left(h_{s}^{2^{k-1}}\right)\right)\right\}]{ }$ |  |
|  | $\left\{p_{0}\left(-h_{s}\right), \ldots, p_{k-1}\left(-h_{s}^{2^{k-1}}\right)\right\}$ |  |
|  | $\overrightarrow{\left\{\operatorname{MTP}\left(p_{0}\left(-h_{s}\right)\right), \ldots, \operatorname{MTP}\left(p_{k-1}\left(-h_{s}^{2^{k-1}}\right)\right)\right\}}$ |  |

(a)
$\underline{\mathcal{P}\left(\left\{f_{i}\right\}_{i}, \mathrm{pp}\right)}$

|  | $\mathcal{V}(\mathrm{pp})$ |
| :---: | :---: |
| $\operatorname{MTR}\left(p_{0}\right)$ |  |
| $v_{0}$ |  |
| $\operatorname{MTR}\left(p_{1}\right)$ |  |
| $v_{1}$ |  |
|  |  |
|  |  |
| $p_{k}$ |  |
| $h_{1}$ |  |
| $\left\{p_{0}\left(h_{1}\right), \ldots, p_{k-1}\left(h_{1}^{2 k-1}\right)\right\}$ |  |
| $\left\{\operatorname{MTP}\left(p_{0}\left(h_{1}\right)\right), \ldots, \operatorname{MTP}\left(p_{k-1}\left(h_{1}^{2^{k-1}}\right)\right)\right\}$ |  |
| $\left\{p_{0}\left(-h_{1}\right), \ldots, p_{k-1}\left(-h_{1}^{2^{k-1}}\right)\right\}$ |  |
| $\overrightarrow{\left\{\operatorname{MTP}\left(p_{0}\left(-h_{1}\right)\right), \ldots, \operatorname{MTP}\left(p_{k-1}\left(-h_{1}^{2^{k-1}}\right)\right)\right\}}$ |  |
| $\left\{f_{0}\left(h_{1}\right), f_{1}\left(h_{1}\right), \ldots, f_{N}\left(h_{1}\right)\right\}$ |  |
| $\left\{\operatorname{MTP}\left(f_{0}\left(h_{1}\right)\right), \ldots, \operatorname{MTP}\left(f_{N}\left(h_{1}\right)\right)\right\}$ |  |
| $\left\{f_{0}\left(-h_{1}\right), f_{1}\left(-h_{1}\right), \ldots, f_{N}\left(-h_{1}\right)\right\}$ |  |
| $\left\{\operatorname{MTP}\left(f_{0}\left(-h_{1}\right)\right), \ldots, \operatorname{MTP}\left(f_{N}\left(-h_{1}\right)\right)\right\}$ |  |
| : |  |
| $h_{s}$ |  |
| $\left\{p_{0}\left(h_{s}\right), \ldots, p_{k-1}\left(h_{s}^{2^{k-1}}\right)\right\}$ |  |
| $\begin{gathered} \left\{\operatorname{MTP}\left(p_{0}\left(h_{s}\right)\right), \ldots, \operatorname{MTP}\left(p_{k-1}\left(h_{s}^{2^{k-1}}\right)\right)\right\} \\ \left\{p_{0}\left(-h_{s}\right), \ldots, p_{k-1}\left(-h_{s}^{2^{k-1}}\right)\right\} \end{gathered}$ |  |
|  |  |  |
| $\left\{\overrightarrow{\left.\operatorname{MTP}\left(p_{0}\left(-h_{s}\right)\right), \ldots, \operatorname{MTP}\left(p_{k-1}\left(-h_{s}^{2^{k-1}}\right)\right)\right\}}\right.$ |  |
| $\left\{f_{0}\left(h_{s}\right), f_{1}\left(h_{s}\right), \ldots, f_{N}\left(h_{s}\right)\right\}$ |  |
| $\left\{\operatorname{MTP}\left(f_{0}\left(h_{s}\right)\right), \ldots, \operatorname{MTP}\left(f_{N}\left(h_{s}\right)\right)\right\}$ |  |
| $\left\{f_{0}\left(-h_{s}\right), f_{1}\left(-h_{s}\right), \ldots, f_{N}\left(-h_{s}\right)\right\}$ |  |
| $\left\{\operatorname{MTP}\left(f_{0}\left(-h_{s}\right)\right), \ldots, \operatorname{MTP}\left(f_{N}(-h)\right)\right\}$ |  |

(b)

Figure 3: Skeleton description of FRI and batched FRI, respectively for (a) and (b). Here, $\mathrm{pp}=(d, H, \mathbb{F}, \mathbb{K})$.

## FRI as a Polynomial Commitment Scheme

Although FRI has a different setting, it can be converted to a Polynomial Commitment Scheme (PCS) (for a definition, see [KZG10]) without much overhead. The scheme is based on the following claim: if $f \in \mathbb{F}[X]$ is a polynomial of degree lower than $d$, then $f(z)$ is the evaluation of $f$ at the point $z$ if and only if $f(X)-f(z)=(X-z) \cdot q(X)$, where $q \in \mathbb{F}[X]$ is some polynomial of degree lower than $d-1$.

FRI can be converted to a polynomial commitment scheme as follows.
Protocol 1 (FRI-based PCS). The protocol starts with a function $f: H \rightarrow \mathbb{F}$ in possession of the prover. The verifier knows an upper bound $d$ on the degree of $f$.

1. As with FRI, the prover's first message is a commitment to $f$.
2. The verifier uniformly samples a challenging point $z \in \mathbb{K} \backslash H$ at which he asks the prover to compute and send the evaluation of $f$.
3. The prover outputs $f(z)$ along with a FRI proof $\pi_{\text {FRI }}$ that:

$$
q(X):=\frac{f(X)-f(z)}{X-z}
$$

is close to a polynomial of degree lower than $d-1$.


Figure 4: Skeleton description of FRI-based PCS.

If FRI passes, the verifier is convinced with high probability that the prover committed, in the first step, to a polynomial $f$ of degree lower than $d$ and that $f(z)$ is the evaluation of $f$ at $z$.

In [VP19], the authors prove that this scheme satisfies the standard notions of security related to polynomial commitment schemes: correctness, polynomial binding and evaluation binding.

### 2.4 Vanilla STARK

In this section, we review the STARK generation procedure from [Sta21] as applied to a particular statement.

## Constraints and Trace

Fix vectors $a, b, c, d, e$ in $\mathbb{F}^{n}$. Denote the elements of $a, b, c, d, e$ by $a_{i}, b_{i}, c_{i}, d_{i}, e_{i}$, for $i \in[n]$, respectively. Let's say we want to generate a STARK for the following statement:

$$
\text { "I know some } a_{i}, b_{i}, c_{i}, d_{i}, e_{i} \in \mathbb{F} \text { such that: }
$$

$$
\begin{align*}
& a_{i} b_{i} c_{i}=a_{i}+b_{i}+c_{i}, \\
& d_{i}^{2}+2 a_{i+1}=e_{i},  \tag{3}\\
& \quad \text { for all } i \in[n] . "
\end{align*}
$$

Denote by $\operatorname{tr}_{1}, \operatorname{tr}_{2}, \operatorname{tr}_{3}, \operatorname{tr}_{4}, \operatorname{tr}_{5} \in \mathbb{F}_{<n}[X]$ the polynomials that interpolate the values $a_{i}, b_{i}, c_{i}, d_{i}, e_{i}$ over the domain $G$, respectively. That is, $\operatorname{tr}_{1}\left(g^{i}\right)=a_{i}, \operatorname{tr}_{2}\left(g^{i}\right)=b_{i}, \operatorname{tr}_{3}\left(g^{i}\right)=$ $c_{i}, \operatorname{tr}_{4}\left(g^{i}\right)=d_{i}, \operatorname{tr}_{5}\left(g^{i}\right)=e_{i}$ for $i \in[n]$. From now on, we will refer to $G$ as the trace evaluation domain and to $\operatorname{tr}_{1}, \operatorname{tr}_{2}, \operatorname{tr}_{3}, \operatorname{tr}_{4}, \operatorname{tr}_{5}$ as the trace column polynomials. Hence, the above constraint system (of size $2 n$ ) can be "compressed" down into two polynomial constraints through the trace column polynomials. In particular, if for all $x \in G$ the following constraints are true:

$$
\begin{array}{rlc}
\operatorname{tr}_{1}(x) \cdot \operatorname{tr}_{2}(x) \cdot \operatorname{tr}_{3}(x) & = & \operatorname{tr}_{1}(x)+\operatorname{tr}_{2}(x)+\operatorname{tr}_{3}(x), \\
\operatorname{tr}_{4}(x)^{2}+2 \cdot \operatorname{tr}_{1}(g x) & = & \operatorname{tr}_{5}(x), \tag{4}
\end{array}
$$

then the original constraint system (3) must hold. The prover sends commitments to $\operatorname{tr}_{1}, \operatorname{tr}_{2}, \operatorname{tr}_{3}, \operatorname{tr}_{4}, \operatorname{tr}_{5}$ to the verifier.

In general, we will be in the situation of generating a STARK for the knowledge of some polynomials $\operatorname{tr}_{1}, \ldots, \operatorname{tr}_{N}: G \rightarrow \mathbb{F}$ that satisfy a system of polynomial constraints $\mathcal{C}=\left\{C_{1}, \ldots, C_{\ell}\right\}$, where:
(a) $C_{i} \in \mathbb{F}\left[X_{1}, \ldots, X_{N}, X_{1}^{\prime}, \ldots, X_{N}^{\prime}\right]$ for all $i \in[\ell]$.
(b) For all $x \in G$ and all $i \in[\ell]$, we have:

$$
\begin{equation*}
C_{i}\left(\operatorname{tr}_{1}(x), \ldots, \operatorname{tr}_{N}(x), \operatorname{tr}_{1}(g x), \ldots, \operatorname{tr}_{N}(g x)\right)=0 \tag{5}
\end{equation*}
$$

when variables $X_{j}$ are replaced by polynomials $\operatorname{tr}_{j}(X)$ and variables $X_{j}^{\prime}$ are replaced by polynomials $\operatorname{tr}_{j}(g X)$ in each $C_{i}$.

## From Polynomial Constraints to Rational Functions

Given a constraint $C_{i}$, a rational function is associated with each one of them:

$$
\begin{equation*}
q_{i}(X):=\frac{C_{i}\left(\operatorname{tr}_{1}(X), \ldots, \operatorname{tr}_{N}(X), \operatorname{tr}_{1}(g X), \ldots, \operatorname{tr}_{N}(g X)\right)}{Z_{G}(X)} \tag{6}
\end{equation*}
$$

where, recalling that $\operatorname{deg}\left(\operatorname{tr}_{i}\right) \leq n-1$, each $q_{i}$ is a polynomial of degree at most $\operatorname{deg}\left(C_{i}\right)$. $(n-1)-n$ if and only if the polynomial $Z_{G}$ divides $C_{i}$ (i.e., $C_{i}$ is satisfed over $G$ ).

Following with our particular example, we have:

$$
\begin{aligned}
& q_{1}(X):=\frac{\operatorname{tr}_{1}(X) \cdot \operatorname{tr}_{2}(X) \cdot \operatorname{tr}_{3}(X)-\operatorname{tr}_{1}(X)-\operatorname{tr}_{2}(X)-\operatorname{tr}_{3}(X)}{Z_{G}(X)}, \\
& q_{2}(X):=\frac{\operatorname{tr}_{4}(X)^{2}+2 \cdot \operatorname{tr}_{1}(g X)-\operatorname{tr}_{5}(X)}{Z_{G}(X)},
\end{aligned}
$$

where $\operatorname{deg}\left(C_{1}\right)=3, \operatorname{deg}\left(C_{2}\right)=2$ and, $q_{1}(X), q_{2}(X)$ are of degree at most $2 n-3$ and $n-2$, respectively. In fact, the constraints in expression (4) get satisfied if and only if $\operatorname{deg}\left(q_{1}\right) \leq 2 n-3$ and $\operatorname{deg}\left(q_{2}\right) \leq n-2$.

## The Quotient Polynomial

In the next step, polynomials $q_{i}$ are combined into a single polynomial $Q$ known as the quotient polynomial. In the STARK proposed in [Sta21], to generate $Q$, the degree of each $q_{i}$ is adjusted to a sufficiently large power of two. More precisely, we define $D_{i}:=$ $\operatorname{deg}\left(C_{i}\right)(n-1)-|G|$ and call $D$ to the first power of two for which $D>D_{i}$ for all $i \in[\ell]$. Then, we compute the adjusted version of the rational functions:

$$
\begin{equation*}
\hat{q}_{i}(X):=\left(\mathfrak{a}_{i} X^{D-D_{i}-1}+\mathfrak{b}_{i}\right) \cdot q_{i}(X) \tag{7}
\end{equation*}
$$

where $\mathfrak{a}_{i}, \mathfrak{b}_{i} \in \mathbb{K}$ for all $i \in[\ell]$.
In our example, we have $D_{1}=2 n-3, D_{2}=n-2$, and therefore we take $D=2 n$ (recall $n$ is a power of 2 ) and then:

$$
\begin{aligned}
& \hat{q}_{1}(X):=\left(\mathfrak{a}_{1} X^{2}+\mathfrak{b}_{1}\right) \cdot q_{1}(X) \\
& \hat{q}_{2}(X):=\left(\mathfrak{a}_{2} X^{n+1}+\mathfrak{b}_{2}\right) \cdot q_{2}(X)
\end{aligned}
$$

hence, $\operatorname{deg}\left(q_{1}\right) \leq 2 n-3$ and $\operatorname{deg}\left(q_{2}\right) \leq n-2$ if and only if $\operatorname{deg}\left(\hat{q}_{1}\right), \operatorname{deg}\left(\hat{q}_{2}\right)<2 n$.
Given polynomials $\hat{q}_{i}$, we can now compute the quotient polynomial as follows:

$$
\begin{equation*}
Q(X):=\sum_{i=1}^{\ell} \hat{q}_{i}(X)=\sum_{i=1}^{\ell}\left(\mathfrak{a}_{i} X^{D-D_{i}-1} \cdot \mathfrak{b}_{i}\right) \frac{C_{i}\left(\operatorname{tr}_{1}(X), \ldots, \operatorname{tr}_{N}(X), \operatorname{tr}_{1}(g X), \ldots, \operatorname{tr}_{N}(g X)\right)}{Z_{G}(X)} \tag{8}
\end{equation*}
$$

that satisfies $\operatorname{deg}(Q)<D$.
Then, $Q$ is split into $S:=D / n$ polynomials $Q_{1}, \ldots, Q_{S} \in \mathbb{K}[X]$ of degree lower than $n$ satisfying:

$$
\begin{equation*}
Q(X)=\sum_{i=1}^{S} X^{i-1} Q_{i}\left(X^{S}\right) \tag{9}
\end{equation*}
$$

Notice that the polynomials $Q_{i}$ are bounded by the same degree as the trace column polynomials, so we refer to them as the trace quotient polynomials.

Continuing with our example, the quotient polynomial is $Q(X):=\hat{q}_{1}(X)+\hat{q}_{2}(X)$ satisfying $\operatorname{deg}(Q)<2 n$. In this case, we represent the quotient polynomial $Q(X)$ as two polynomials $Q_{1}, Q_{2} \in \mathbb{F}_{<n}[X]$ such that $Q(X)=Q_{1}\left(X^{2}\right)+X \cdot Q_{2}\left(X^{2}\right)$.

## Trace Low Degree Extension

Since the quotient polynomial $Q$ is defined through rational functions $q_{i}$ that we are going to evaluate, for Cstr. (6) to be well-defined, we need to ensure that the denominators of these rational functions are never zero. Therefore, from now on, the polynomials will not be evaluated over the trace evaluation domain, but rather over a larger and disjoint domain, which we refer to as the evaluation domain.

More specifically, we introduce the evaluation domain to be:

1. Larger than $G$, so that the trace column polynomial has enough redundancy to ensure the soundness of the FRI protocol.
2. Disjoint of $G$, so that Cstr. (6) is well-defined.

In order to achieve the previous two requirements, we will choose the evaluation domain to be the coset $H$, where remember that $|H|=2^{k} \cdot n$, with $k \geq 1$. We will refer to $2^{k}$ as the blowup factor. Therefore, we need to evaluate all the trace column polynomials over the evaluation domain. We refer to the resulting set of polynomial evaluations as the trace Low Degree Extension (LDE).

The trace LDE is computed in two steps:

1. We calculate the interpolation polynomial on the trace evaluation domain of each trace column polynomial using the Inverse Fast Fourier Transform (IFFT).
2. We evaluate the polynomials that result from the previous step on the evaluation domain using the Fast Fourier Transform (FFT).

## Trace Consistency Check

At this point, the verifier has everything he needs to perform a local consistency check between the trace column polynomials and the trace quotient polynomials, referred to as the trace consistency check. Hence, after the prover commits to the trace quotient polynomials, the verifier uniformly samples a random $z$ and requests the prover to send the necessary polynomial evaluations on either $z$ or $g z$.

More specifically, for a given $z \in \mathbb{K} \backslash(G \cup \bar{H})$ (here, $\bar{H}=\left\{x \in \mathbb{K} \mid x^{S} \in H\right\}$ ) uniformly sampled by a verifier, the prover sends back the trace column polynomials evaluations $\operatorname{tr}_{i}(z), \operatorname{tr}_{j}(g z)$ and trace quotient polynomials evaluations $Q_{k}\left(z^{S}\right)$ for $i, j \in N$ and $k \in[S]$. Notice that the necessary evaluations of the trace column polynomials strictly depend on the particular polynomial expressions in (5). Denote by Evals $(z)$ the set of polynomial evaluations over $z$ and by $\operatorname{Evals}(g z)$ the set of polynomial evaluations over $g z$. Naturally, there could exist evaluations of a single polynomial in both sets.

With these evaluations, the verifier can check that:

$$
\begin{equation*}
\sum_{i=1}^{S} z^{i-1} Q_{i}\left(z^{S}\right)=\sum_{i=1}^{\ell}\left(\mathfrak{a}_{i} X^{D-D_{i}-1} \cdot \mathfrak{b}_{i}\right) \frac{C_{i}\left(\operatorname{tr}_{1}(z), \ldots, \operatorname{tr}_{N}(z), \operatorname{tr}_{1}(g z), \ldots, \operatorname{tr}_{N}(g z)\right)}{Z_{G}(z)} \tag{10}
\end{equation*}
$$

In our particular example, the prover sends back $\operatorname{tr}_{1}(z), \operatorname{tr}_{1}(g z), \operatorname{tr}_{2}(z), \operatorname{tr}_{3}(z), \operatorname{tr}_{4}(z), \operatorname{tr}_{5}(z)$ (so that $\operatorname{Evals}(z)=\left\{\operatorname{tr}_{1}(z), \operatorname{tr}_{2}(z), \operatorname{tr}_{3}(z), \operatorname{tr}_{4}(z), \operatorname{tr}_{5}(z)\right\}$ and $\operatorname{Evals}(g z)=\left\{\operatorname{tr}_{1}(g z)\right\}$ ) and $Q_{1}\left(z^{2}\right), Q_{2}\left(z^{2}\right)$, and the verifier checks that:

$$
\begin{align*}
Q_{1}\left(z^{2}\right)+z \cdot Q_{2}\left(z^{2}\right)=\left(\mathfrak{a}_{1} z^{2}+\right. & \left.\mathfrak{b}_{1}\right) \cdot \frac{\operatorname{tr}_{1}(z) \cdot \operatorname{tr}_{2}(z) \cdot \operatorname{tr}_{3}(z)-\operatorname{tr}_{1}(z)-\operatorname{tr}_{2}(z)-\operatorname{tr}_{3}(z)}{Z_{G}(z)}+ \\
& +\left(\mathfrak{a}_{2} z^{n+1}+\mathfrak{b}_{2}\right) \cdot \frac{\operatorname{tr}_{4}(z)^{2}+2 \cdot \operatorname{tr}_{1}(g z)-\operatorname{tr}_{5}(z)}{Z_{G}(z)} \tag{11}
\end{align*}
$$

The problem is that the verifier cannot be sure that the received values are actually the evaluations of previously committed polynomials. In fact, it is easy to obtain values from $\mathbb{K}$ that satisfy (10) but are not from the expected polynomials. Therefore, after the prover sends the evaluations to the verifier, they engage in the part of the protocol to ensure that the values are the actual evaluations of the corresponding polynomials.

## The FRI Polynomial

To ensure the validity of the values sent from the prover to the verifier, the prover proceeds to create another set of constraints, then translate them to a problem of low-degree testing, and finally, combines them through the use of random field elements. To this end, the prover computes the F polynomial:

$$
\begin{align*}
\mathrm{F}(X):= & \sum_{i \in I_{1}} \varepsilon_{i}^{(1)} \cdot \frac{\operatorname{tr}_{i}(X)-\operatorname{tr}_{i}(z)}{X-z}+\sum_{i \in I_{2}} \varepsilon_{i}^{(2)} \cdot \frac{\operatorname{tr}_{i}(g X)-\operatorname{tr}_{i}(g z)}{X-g z} \\
& +\sum_{i=1}^{S} \varepsilon_{i}^{(3)} \cdot \frac{Q_{i}(X)-Q_{i}\left(z^{S}\right)}{X-z^{S}} \tag{12}
\end{align*}
$$

where $I_{1}=\left\{i \in[N]: \operatorname{tr}_{i}(z) \in \operatorname{Evals}(z)\right\}, I_{2}=\left\{i \in[N]: \operatorname{tr}_{i}(g z) \in \operatorname{Evals}(g z)\right\}$ and $\varepsilon_{i}^{(1)}, \varepsilon_{j}^{(2)}, \varepsilon_{k}^{(3)} \in \mathbb{K}$ for all $i \in I_{1}, j \in I_{2}, k \in[S]$.

What we obtain is that the F polynomial is of degree lower than $n-1$ if and only if: (1) both the trace columns polynomials $\operatorname{tr}_{i}$ and the trace quotient polynomials $Q_{i}$ are of degree lower than $n$ and (2) all the values sent by the prover in the previous step are evaluations of the corresponding polynomials.

Finishing with our example, the prover computes the F polynomial as:

$$
\begin{aligned}
\mathrm{F}(X):= & \varepsilon_{1}^{(1)} \frac{\operatorname{tr}_{1}(X)-\operatorname{tr}_{1}(z)}{X-z}+\varepsilon_{2}^{(1)} \cdot \frac{\operatorname{tr}_{2}(X)-\operatorname{tr}_{2}(z)}{X-z}+\varepsilon_{3}^{(1)} \cdot \frac{\operatorname{tr}_{3}(X)-\operatorname{tr}_{3}(z)}{X-z} \\
& +\varepsilon_{4}^{(1)} \cdot \frac{\operatorname{tr}_{4}(X)-\operatorname{tr}_{4}(z)}{X-z}+\varepsilon_{5}^{(1)} \cdot \frac{\operatorname{tr}_{5}(X)-\operatorname{tr}_{5}(z)}{X-z}+\varepsilon_{1}^{(2)} \cdot \frac{\operatorname{tr}_{1}(X)-\operatorname{tr}_{1}(g z)}{X-g z} \\
& +\varepsilon_{1}^{(3)} \cdot \frac{Q_{1}(X)-Q_{1}\left(z^{2}\right)}{X-z^{2}}+\varepsilon_{2}^{(3)} \cdot \frac{Q_{2}(X)-Q_{2}\left(z^{2}\right)}{X-z^{2}} .
\end{aligned}
$$

## Proof Verification

To verify the STARK, the verifier performs the following checks:
(a) Trace Consistency. Checks that the trace quotient polynomials $Q_{1}, \ldots, Q_{S}$ are consistent with the trace column polynomials $\operatorname{tr}_{1}, \ldots, \operatorname{tr}_{N}$ by means of the evaluations of these polynomials at either $z, g z$ or $z^{S}$. I.e., the verifier checks that Eq. (10) is satisfied.
(b) Batched FRI Verification. It runs the batched FRI verification procedure on the polynomial F.

If either $(a)$ or (b) fails at any point, the verifier aborts and rejects. Otherwise, the verifier accepts.

The full skeleton description of the vanilla STARK protocol can be found in one-shot in Appendix A, in which we include the FRI message exchange in Figure 7.

## Soundness

We finally recall the soundness error of the vanilla STARK protocol as in Theorem 4 of [Sta21]. Recall that $\rho=|G| /|H|$ is the rate of the code.

Theorem 2 (Soundness). Fix integer $m \geq 3$. Suppose a prover that interacts with $a$ verifier in the vanilla STARK protocol causes it to accept with probability two times greater than:

$$
\begin{equation*}
\varepsilon_{S T A R K}:=\ell\left(\frac{1}{|\mathbb{K}|}+\frac{(D+|G|+S) \cdot \ell}{|\mathbb{K}|-S|H|-|G|}\right)+\varepsilon_{F R I} \tag{13}
\end{equation*}
$$

then Eq. (5) is satisfied (i.e., the original statement is true), where $\ell=m / \rho$ and $\varepsilon_{\text {FRI }}$ is as defined in Theorem 1.

We give some intuitions for the computation of this error bound. The first term in Eq. (13) corresponds to the probability of sets $\left\{\mathfrak{a}_{1}, \mathfrak{b}_{1}, \ldots, \mathfrak{a}_{T}, \mathfrak{b}_{T}\right\}$ being "good" under a dishonest prover. This means that if all $\mathfrak{a}_{i}, \mathfrak{b}_{i}$ are uniformly and independently randomly sampled, then the probability the quotient polynomial $Q$ is a polynomial is at most $\ell /|\mathbb{K}|$. The second term corresponds to the probability of sampling a $z \in \mathbb{K} \backslash(G \cup \bar{H})$ for which the FRI protocol passes with a probability greater than $\varepsilon_{\text {FRI }}$.

### 2.5 Arguments

In this section, we introduce the "arguments" that we will use to extend the vanilla STARK. Here, by argument, we mean a relation between polynomials that cannot be directly expressed through an identity. We often refer to these arguments as non-identity constraint. The three arguments we will introduce are multiset equality, connection and inclusion.

The protocols that instantiate these arguments are all based on the same idea of the computation of a grand product polynomial over the two (or more) vectors involved in the argument. Specifically, a polynomial is cumulatively computed as the quotient of a function of the first vector and a function of the second vector. Then, a set of identities is proposed as insurance for a verifier of the protocol that not only the grand product was correctly computed by a prover, but also that the specific intention of the protocol is satisfied. To ensure the soundness of the protocols, random values uniformly sampled by the verifier are used in such computation.

Recall that $G=\langle g\rangle$ is a cyclic subgroup of $\mathbb{F}^{*}$ of order $n$.

## Multiset Equality

Given two vectors $f=\left(f_{1}, \ldots, f_{n}\right)$ and $t=\left(t_{1}, \ldots, t_{n}\right)$ in $\mathbb{F}^{n}$, a multiset equality argument, denoted $f \doteq t$, is used for checking that $f$ is equal to $t$ as multisets (or equivalently, that $f$ and $t$ are a permutation of each other). The protocol that instantiates the multiset equality arguments works by computing the following grand product polynomial $Z \in \mathbb{K}_{<n}[X]$ :

$$
Z\left(g^{i}\right)= \begin{cases}1, & \text { if } \quad i=1 \\ \prod_{j=1}^{i-1} \frac{\left(f_{j}+\gamma\right)}{\left(t_{j}+\gamma\right)}, & \text { if } \quad i=2, \ldots, n\end{cases}
$$

where $\gamma \in \mathbb{K}$ is the value sent from the verifier.
The definition of the previous polynomial is based on the following lemma.
Lemma 1 (Soundness of Multiset Equality). Fix two vectors $f=\left(f_{1}, \ldots, f_{n}\right)$ and $t=$ $\left(t_{1}, \ldots, t_{n}\right)$ in $\mathbb{F}^{n}$. If the following holds with probability larger than $\varepsilon_{M u l E q}(n):=n /|\mathbb{K}|$ over a random $\gamma \in \mathbb{K}$ :

$$
\prod_{i=1}^{n}\left(f_{i}+\gamma\right)=\prod_{i=1}^{n}\left(t_{i}+\gamma\right)
$$

then $f \doteq t$.
Proof. Assume that $f \neq t$. Then, there must be some $i^{*} \in[n]$ such that $t_{i^{*}} \neq f_{i}$ for all $i \in[n]$. Define degree $n$ polynomials $F(X):=\prod_{i=1}^{n}\left(f_{i}+X\right)$ and $T(X):=\prod_{i=1}^{n}\left(t_{i}+X\right)$. By the assumption, we have that $F \neq T$ and therefore by the Schwartz-Zippel lemma $F(\gamma) \neq T(\gamma)$ except with probability $n /|\mathbb{K}|$.

As a consequence of Lemma 1, the identities that must be checked by the verifier for $x \in G$ are the following:

$$
\begin{align*}
& L_{1}(x) \cdot(Z(x)-1)=0  \tag{14}\\
& Z(x \cdot g) \cdot(t(x)+\gamma)=Z(x) \cdot(f(x)+\gamma) \tag{15}
\end{align*}
$$

where $f, t \in \mathbb{F}_{<n}[X]$ are the polynomials resulting from the interpolation of $\left\{f_{i}\right\}_{i \in[n]}$ and $\left\{t_{i}\right\}_{i \in[n]}$ over $G$,respectively.

## Connection

The protocol for a connection argument and the definitions and results we provide next are adapted from [GWC19].

Given some vectors $f_{1}, \ldots, f_{k} \in \mathbb{F}^{n}$ and a partition $\mathcal{T}=\left\{T_{1}, \ldots, T_{s}\right\}$ of the set [kn], a connection argument, denoted $\left(f_{1}, \ldots, f_{k}\right) \propto\left\{T_{1}, \ldots, T_{s}\right\}$, is used to check that the partition $\mathcal{T}$ divides the field elements $\left\{f_{i, j}\right\}_{i \in[k], j \in[n]}$ into sets with the same value. More specifically, if we define the sequence $f_{(1)}, \ldots, f_{(k n)} \in \mathbb{F}$ by:

$$
f_{((i-1) n+j)}:=f_{i, j}
$$

for each $i \in[k], j \in[n]$, then we have $f_{\left(\ell_{1}\right)}=f_{\left(\ell_{2}\right)}$ if and only if $\ell_{1}, \ell_{2}$ belong to the same block of $\mathcal{T}$.

In order to express the partition $\mathcal{T}$ within a grand product polynomial, we define a permutation $\sigma:[k n] \rightarrow[k n]$ as follows: $\sigma$ is such that for each block $T_{i}$ of $\mathcal{T}, \sigma(\mathcal{T})$ contains a cycle going over all elements of $T_{i}$. Then, the protocol that instantiates the connection arguments works by computing the following grand product polynomial $Z \in \mathbb{K}_{<_{n}}[X]$ :

$$
Z\left(g^{i}\right)= \begin{cases}1, & \text { if } i=1 \\ \prod_{\ell=1}^{k} \prod_{j=1}^{i-1} \frac{\left.\left(f_{\ell, j}+\gamma \cdot(\ell-1) \cdot n+j\right)+\delta\right)}{\left(f_{\ell, j}+\gamma \cdot \sigma((\ell-1) \cdot n+j)+\delta\right)}, & \text { if } i=2, \ldots, n\end{cases}
$$

where $\gamma, \delta \in \mathbb{K}$ are the values sent from the verifier.
The definition of the previous polynomial is based on the following lemma, a proof of which can be found ${ }^{2}$ in Claim A.1. of [GWC19] and is similar to the one of Lemma 1.

Lemma 2 (Soundness of Connection). Fix $f_{1}, \ldots, f_{k} \in \mathbb{F}^{n}$ and a partition $\mathcal{T}=\left\{T_{1}, \ldots, T_{s}\right\}$ of $[k n]$. If the following holds with probability larger than $\varepsilon_{C o n}(n):=k n /|\mathbb{K}|$ over randoms $\gamma, \delta \in \mathbb{K}$ :

$$
\prod_{\ell=1}^{k} \prod_{j=1}^{n}\left(f_{\ell, j}+\gamma \cdot((\ell-1) \cdot n+j)+\delta\right)=\prod_{\ell=1}^{k} \prod_{j=1}^{n}\left(f_{\ell, j}+\gamma \cdot \sigma((\ell-1) \cdot n+j)+\delta\right)
$$

then, $\left(f_{1}, \ldots, f_{k}\right) \propto\left\{T_{1}, \ldots, T_{s}\right\}$.
As a consequence of Lemma 2, the identities that must be checked by the verifier for $x \in G$ are the following:

$$
\begin{gather*}
L_{1}(x) \cdot(Z(x)-1)=0 \\
Z(x \cdot g)=Z(x) \cdot \frac{\left(f_{1}(x)+\gamma \cdot S_{\mathrm{ID}_{1}}(x)+\delta\right)}{\left(f_{1}(x)+\gamma \cdot S_{\sigma_{1}}(x)+\delta\right)} \cdots \cdots \frac{\left(f_{k}(x)+\gamma \cdot S_{\mathrm{ID}_{k}}(x)+\delta\right)}{\left(f_{k}(x)+\gamma \cdot S_{\sigma_{k}}(x)+\delta\right)} \tag{16}
\end{gather*}
$$

where $S_{\mathrm{ID}_{i}}\left(g^{j}\right)=(i-1) \cdot n+j$ is the polynomial mapping $G$-elements to indexes in $[k n]$ and $S_{\sigma_{i}}\left(g^{j}\right)=\sigma((i-1) \cdot n+j)$ is the polynomial defined by $\sigma$. Since the permutation $\sigma$ perfectly relates with the partition $\mathcal{T}$ it refers to, from now on we denote a connection argument between polynomials $f_{1}, \ldots, f_{k} \in \mathbb{F}[X]$ and a partition $\mathcal{T}$ as $\left(f_{1}, \ldots, f_{k}\right) \propto\left(S_{\sigma_{1}}, \ldots, S_{\sigma_{k}}\right)$. As we will see in later sections, this overloading notation will become very natural.

For more details see [GWC19].

[^1]
## Inclusion

The protocol for an inclusion argument and the definitions and results we provide next is adapted from the well-known Plookup protocol [GW20], with the "alternating method" provided in $\left[\mathrm{PFM}^{+} 22\right]$.

Given two vectors $f=\left(f_{1}, \ldots, f_{n}\right)$ and $t=\left(t_{1}, \ldots, t_{n}\right)$ in $\mathbb{F}^{n}$, a inclusion argument, denoted $f \in t$, is used for checking that the set $A$ formed with the values $\left\{f_{i}\right\}_{i \in[n]}$ is contained in the set $B$ formed with the values $\{t\}_{i \in[n]}$. Notice that $|A|,|B| \leq n$.

In the protocol, the prover has to construct an auxiliary vector $s=\left(s_{1}, \ldots, s_{2 n}\right)$ containing every element of $f$ and $t$ where the order of appearance is the same as in $t$. The main idea behind the protocol is that if $f \in t$, then $f$ contributes to $s$ with repeated elements. To check this fact, a vector $\Delta s$ is defined as follows:

$$
\Delta s=\left(s_{1}+\gamma s_{2}, s_{2}+\gamma s_{3}, \cdots, s_{2 n}+\gamma s_{1}\right)
$$

Then, the protocol essentially checks that $\Delta s$ is consistent with the elements of $f, t$ and $s$. To do so, the vector $s$ is split into two vectors $h_{1}, h_{2} \in \mathbb{F}^{n}$. In the protocol described in [GW20], $h_{1}$ and $h_{2}$ contain the lower and upper halves of $s$, while in our protocol in [ $\mathrm{PFM}^{+} 22$ ], we use $h_{1}$ to store elements with odd indexes and $h_{2}$ for even indexes, that is:

$$
\begin{equation*}
h_{1}=\left(s_{1}, s_{3}, s_{5}, \ldots, s_{2 n-1}\right) \quad \text { and } \quad h_{2}=\left(s_{2}, s_{4}, s_{6}, \ldots, s_{2 n}\right) \tag{17}
\end{equation*}
$$

With this setting in mind, the grand product polynomial is defined as:
$Z\left(g^{i}\right)= \begin{cases}1, & \text { if } i=1 \\ (1+\gamma)^{i-1} \prod_{j=1}^{i-1} \frac{\left(\delta+f_{j}\right)\left(\delta(1+\gamma)+t_{j}+\gamma t_{j+1}\right)}{\left(\delta(1+\gamma)+s_{2 j-1}+\gamma s_{2 j}\right)\left(\delta(1+\gamma)+s_{2 j}+\gamma s_{2 j+1}\right)}, & \text { if } i=2, \ldots, n\end{cases}$
where $\gamma, \delta \in \mathbb{K}$ are the values sent from the verifier.
The definition of the previous polynomial is based on the following lemma, which is a slight modification of Claim 3.1. of [GW20].

Lemma 3 (Soundness of Inclusion). Fix three vectors $f=\left(f_{1}, \ldots, f_{n}\right), t=\left(t_{1}, \ldots, t_{n}\right)$ and $s=\left(s_{1}, \ldots, s_{2 n}\right)$ with elements in $\mathbb{F}$. If the following holds with probability larger than $\varepsilon_{\operatorname{lnc}}(n):=(4 n-2) /|\mathbb{K}|$ over randoms $\gamma, \delta \in \mathbb{K}$ :

$$
(1+\gamma)^{n} \prod_{i=1}^{n}\left(\delta+f_{i}\right) \prod_{i=1}^{n-1}\left(\delta(1+\gamma)+t_{i}+\gamma t_{i+1}\right)=\prod_{i=1}^{2 n-1}\left(\delta(1+\gamma)+s_{i}+\gamma s_{i+1}\right)
$$

then $f \in t$ and $s$ is the sorted by $t$ concatenation of $f$ and $t$.
As a consequence of Lemma 3, the identities that must be checked by the verifier for $x \in G$ are the following:

$$
\begin{gather*}
L_{1}(x)(Z(x)-1)=0 \\
Z(x \cdot g)=Z(x) \frac{(1+\gamma)(\delta+f(x))(\delta(1+\gamma)+t(x)+\gamma t(g x))}{\left(\delta(1+\gamma)+h_{1}(x)+\gamma h_{2}(x)\right)\left(\delta(1+\gamma)+h_{2}(x)+\gamma h_{1}(x \cdot g)\right)} \tag{18}
\end{gather*}
$$

where $f, t \in \mathbb{F}_{<n}[X]$ are the polynomials resulting from the interpolation of $\left\{f_{i}\right\}_{i \in[n]}$ and $\left\{t_{i}\right\}_{i \in[n]}$ over $G$, respectively; and $h_{1}, h_{2} \in \mathbb{F}_{<n}[X]$ are the polynomials resulting from the interpolation of the values defined in Eq. (17) over $G$.

For more details see [GW20] and $\left[\mathrm{PFM}^{+} 22\right]$.

## 3 Our Techniques

We now explain the main differences between the polynomial computations carried on during the rounds of the vanilla STARK (Section 2.4) and our protocol. We also explain the tradeoffs arising from controlling the constraint degree either at the representation of the AIR or inside the protocol itself.

### 3.1 Commiting to Multiple Polynomial at Once

In our protocol, the prover sends Merkle tree commitments to multiple polynomials in each round. The naive way of proceeding is by sending one Merkle tree root per polynomial. In this section, we explain how we achieve a sound alternative by computing a single Merkle tree of all polynomials in each round. Our strategy not only reduces the amount of Merkle roots that $\mathcal{P}$ has to send to $\mathcal{V}$ in each round but also reduces the Merkle paths that $\mathcal{P}$ needs to send to $\mathcal{V}$ when $\mathcal{P}$ gets asked for evaluations of multiple polynomials at the same point.

Notation. As explained in Section 2.4, commitments are generated by computing Merkle trees over polynomial evaluations over a nontrivial coset $H$ of a cyclic subgroup of $\mathbb{F}^{*}$ with order $m$. For this section, explicitly set $H=\left\{h_{1}, h_{2}, h_{3}, \ldots, h_{m}\right\}$.

Say that $f_{1}, \ldots, f_{N} \in \mathbb{K}_{<n}[X]$ is the set of polynomials that we want to construct the Merkle Tree on. More precisely, we want to compute the Merkle tree over the following $m \times N$ matrix of polynomial evaluations:

$$
\left(\begin{array}{cccc}
f_{1}\left(h_{1}\right), & f_{2}\left(h_{1}\right), & \ldots, & f_{N}\left(h_{1}\right) \\
f_{1}\left(h_{2}\right), & f_{2}\left(h_{2}\right), & \ldots, & f_{N}\left(h_{2}\right) \\
\vdots & \vdots & \ldots & \vdots \\
f_{1}\left(h_{m}\right), & f_{2}\left(h_{m}\right), & \ldots, & f_{N}\left(h_{m}\right)
\end{array}\right)
$$

We start the construction of the Merkle Tree by grouping the evaluations of $f_{1}, \ldots, f_{N}$ at a single point of $H$. That is, the $i$-th leaf of the Merkle tree is formed by (the hash of) the $i$-th row of the previous matrix. This gives a total of $m$ leaves, which is by assumption a power of two. To be more precise, the leaf elements of the Merkle Tree, indexed by the corresponding $H$-value, will consist on:

$$
\begin{array}{|ccc|}
\hline \text { leaf } h_{1} & \Longrightarrow & \mathcal{H}\left(f_{1}\left(h_{1}\right), f_{2}\left(h_{1}\right), \ldots, f_{N}\left(h_{1}\right)\right) \\
\text { leaf } h_{2} & \Longrightarrow & \mathcal{H}\left(f_{1}\left(h_{2}\right), f_{2}\left(h_{2}\right), \ldots, f_{N}\left(h_{2}\right)\right) \\
\vdots & & \vdots \\
\text { leaf } h_{m} & \Longrightarrow & \mathcal{H}\left(f_{1}\left(h_{m}\right), f_{2}\left(h_{m}\right), \ldots, f_{N}\left(h_{m}\right)\right) \\
\hline
\end{array}
$$

where $\mathcal{H}$ is any collision-resistant hash function. Once all the leaves are computed, the rest of the Merkle tree is computed, as usual, by recursively hashing the concatenation of its two child nodes until the Merkle root is achieved. The commitments to $f_{1}, \ldots, f_{N}$ is this single Merkle root.

Now, notice how if $\mathcal{V}$ requests a Merkle proof for the evaluation of all $f_{1}, \ldots, f_{N}$ at a single point $h_{i}, \mathcal{P}$ can prove the consistency of all the evaluations $f_{1}\left(h_{i}\right), \ldots, f_{N}\left(h_{i}\right)$ with the Merkle root by simply sending the Merkle path corresponding to the leaf containing such evaluations. Compared to the naive version, this version improves the proof size from $O(N \log m)$ elements down to $O(\log m)$ elements. This will be convenient for the batched FRI execution of our protocol, where we group evaluations of a set of polynomials at a single point for succinctly answering each batched consistency check.

### 3.2 Preprocessed Polynomials and Public Values

Among the set of polynomials that are part of the polynomial constraint system representing the problem's statement, we differentiate between two types: committed polynomials and preprocessed polynomials.

Committed polynomials are those polynomials for which the verifier only has oracle access and are therefore committed (via Merkle trees) by the prover before the verifier starts querying them. In other words, these polynomials can only be known, in principle, in their entire form by the prover of the protocol. Consequently, the verifier is limited to knowing a "small fraction" of these polynomials' evaluations. In practice, this fraction is randomly chosen by the verifier and is proportional to the number of oracle queries that the verifier makes to the particular polynomial. For the shake of the protocols to be scalable, the number of queries made to committed polynomials should be at most logarithmic in their degree. An example of committed polynomials is trace columns polynomials $\operatorname{tr}_{i}$.

On the other hand, preprocessed polynomials are known by the verifier in their entirety even before the execution of the protocol. More precisely, once a polynomial constraint system $\mathcal{C}$ is fixed, the verifier has complete access (either in coefficient form or in evaluation form) to the set of preprocessed polynomials. As with committed polynomials, the verifier ends up needing only a small subset of evaluations of such polynomials during the protocol. An example of preprocessed polynomials is Lagrange polynomials $L_{i}$. We explain how we treat preprocessed polynomials in the protocol in Section 4.2.
Example 1. As an example, the polynomial constraints such that for all $x \in G$ satisfies:

$$
\begin{equation*}
L_{1}(x)\left(\operatorname{tr}_{1}(x)-7\right)=0, L_{n}(x)\left(\operatorname{tr}_{1}(x)-3\right)=0 \tag{19}
\end{equation*}
$$

is composed of one committed polynomial, namely $\operatorname{tr}_{1}$, and two preprocessed polynomial, namely $L_{1}, L_{n}$.

Finally, we define public values as the set of committed polynomial evaluations that are attested by some constraint. Public values are known to both the prover and the verifier and a particular polynomial can have many public values associated with it. In the previous example, the evaluation of $\operatorname{tr}_{1}$ at $g$ and $g^{n}$ are public values since Eq. (19) are satisfied if and only if $\operatorname{tr}_{1}(g)=7$ and $\operatorname{tr}_{1}\left(g^{n}\right)=3$.

### 3.3 Adding Selected Vector Arguments

In this section, we describe how to augment the type of available constraints with the arguments presented in Section 2.5. Recall that we will add three new types of arguments:

- Inclusion $(\in)$. The set constructed from the evaluations of a polynomial $f$ over a multiplicative subgroup $G$ is contained in an equally defined set of another polynomial $t$.
- Multiset Equality $(\dot{=})$. The multiset considered from the evaluations of a polynomial $f$ over a multiplicative subgroup $G$ is equal to the multiset considered from the evaluations of another polynomial $t$.
- Connection $(\propto)$. The vectors constructed from the evaluations of a set of polynomials $f_{1}, \ldots, f_{N}$ over a multiplicative subgroup $G$ does not vary after applying a particular permutation $\sigma$ to them.

To include non-identity constraints in the protocol, we will represent them through their succinct set of identity constraints. We denote by $M$ the number of inclusion instantiations, $M^{\prime}$ the number of multiset equality instantiations and $M^{\prime \prime}$ the number of connection instantiations.

As detailed in Section 2.5, for the inclusion argument, we need to compute and commit the associated polynomials $h_{1, j}$ and $h_{2, j}$ before being able to compute the corresponding grand product polynomial for each inclusion constraint $j \in[M]$. This sums up to $2 M$ polynomials. After this, for each non-identity constraint, we compute the associate grand product polynomial $Z$ and commit to it. This definition of this polynomial is different depending on which argument we are executing as shown in Section 2.5. This sums up to $M$ inclusion polynomials, $M^{\prime}$ multiset equality polynomials and $M^{\prime \prime}$ connection polynomials. Overall, adding non-identity constraints adds up to $3 M+M^{\prime}+M^{\prime \prime}$ committed polynomials and $2\left(M+M^{\prime}+M^{\prime \prime}\right)$ polynomial constraints to the STARK.

Following, we explain how we generalize both inclusions and multiset equalities to not only involving multiple polynomials but also to a subset of the resulting vector. Therefore, somewhat artificially, enlarge the expressiveness of our arguments and let us handle more generic non-identity constraints.

## From Vector Arguments to Simple Arguments

Let's explain first how we reduce vector inclusions or multiset equalities to simple (i.e., involving only one polynomial on each side) inclusions or multiset equalities.

Definition 5 (Vector Arguments). Given polynomials $f_{i}, t_{i} \in \mathbb{K}_{<n}[X]$ for $i \in[N]$, a vector inclusion, denoted $\left(f_{1}, \ldots, f_{N}\right) \in\left(t_{1}, \ldots, t_{N}\right)$, is the argument in which for all $x \in G$ there exists some $y \in G$ such that:

$$
\begin{equation*}
\left(f_{1}(x), \ldots, f_{N}(x)\right)=\left(t_{1}(y), \ldots, t_{N}(y)\right) \tag{20}
\end{equation*}
$$

A vector multiset equality, denoted $\left(f_{1}, \ldots, f_{N}\right) \doteq\left(t_{1}, \ldots, t_{N}\right)$, is the argument in which for all $y \in G$ there exists exactly one $x \in G$ for which Eq. (20) holds. That is, (vector) multiset equalities define a bijective mapping.

To reduce the previous vector arguments to simple ones, we make use of a uniformly sampled element $\alpha \in \mathbb{K}$. Namely, instead of trying to generate an argument for the vector relation, we define the following polynomials:

$$
\begin{equation*}
F^{\prime}(X):=\sum_{i=1}^{N} \alpha^{i-1} f_{i}(X), \quad T^{\prime}(X):=\sum_{i=1}^{N} \alpha^{i-1} t_{i}(X) \tag{21}
\end{equation*}
$$

and proceed to prove the relation $F^{\prime} \in T^{\prime}$ or $F^{\prime} \doteq T^{\prime}$. Notice that both $F^{\prime}$ and $T^{\prime}$ are in general polynomials with coefficients over the field extension $\mathbb{K}$ even if every coefficient of $f_{i}, t_{i}$ is precisely over the base field $\mathbb{F}$.

The previous reduction leads to the following result.
Lemma 4. Given polynomials $f_{i}, t_{i} \in \mathbb{K}_{<n}[X]$ for $i \in[N]$ and $F^{\prime}, T^{\prime} \in \mathbb{K}_{<n}[X]$ as defined by Eq. (21), if $F^{\prime} \in T^{\prime}$ (resp. $F^{\prime} \doteq T^{\prime}$ ), then $\left(f_{1}, \ldots, f_{N}\right) \in\left(t_{1}, \ldots, t_{N}\right)$ (resp. $\left.\left(f_{1}, \ldots, f_{N}\right) \doteq\left(t_{1}, \ldots, t_{N}\right)\right)$ except with probability $n \cdot(N-1) /|\mathbb{K}|$ over the random choice of $\alpha$.

Proof. Assume $\left(f_{1}, \ldots, f_{N}\right) \notin\left(t_{1}, \ldots, t_{N}\right)$, then there exists some $x \in G$ such that for every $y \in G$ we have $\left(f_{1}(x), \ldots, f_{N}(x)\right) \notin\left(t_{1}(y), \ldots, t_{N}(y)\right)$. This means that $F^{\prime}(x) \neq$ $T^{\prime}(y)$ for each $y \in G$ except with probability $(N-1) /|\mathbb{K}|$ over the random choice of $\alpha$. Since $|G|=n$, the lemma follows.

We generalize to vector arguments the protocols for (simple) inclusion arguments and multiset equality arguments explained in Section 2.5 by incorporating the previous reduction strategy. Therefore, we give next the soundness bounds for these protocols.

Lemma 5. Given polynomials $f_{i}, t_{i} \in \mathbb{K}_{<n}[X]$ for $i \in[N]$, we obtain:

1. Inclusion Protocol. Let $F, T \in \mathbb{K}_{<n}[X]$ as defined by Eq. (21). The prover sends oracle functions $\left[f_{i}\right],\left[t_{i}\right]$ for $i \in[N]$ to the verifier in the first round, who responds with a uniformly sampled $\alpha \in \mathbb{K}$. If a prover that interacts with a verifier in the inclusion protocol of Section 2.5 on input $F$ and $T$ causes it to accept with probability greater than:

$$
n \frac{N-1}{|\mathbb{K}|}+\varepsilon_{\operatorname{lnc}}(n)
$$

then $\left(f_{1}, \ldots, f_{N}\right) \in\left(t_{1}, \ldots, t_{N}\right)$.
2. Multiset Equality Protocol. Let $F, T \in \mathbb{K}_{<n}[X]$ as defined by Eq. (21). The prover sends oracle functions $\left[f_{i}\right],\left[t_{i}\right]$ for $i \in[N]$ to the verifier in the first round, who responds with a uniformly sampled $\alpha \in \mathbb{K}$. If a prover that interacts with a verifier in the multiset equality protocol of Section 2.5 on input $F$ and $T$ causes it to accept with probability greater than:

$$
n \frac{N-1}{|\mathbb{K}|}+\varepsilon_{M u \mid E q}(n)
$$

then $\left(f_{1}, \ldots, f_{N}\right) \doteq\left(t_{1}, \ldots, t_{N}\right)$.

## From Selected Vector Arguments to Simple Arguments

Now, let's go one step further by the introduction of selectors. Informally speaking, a selected inclusion (multiset equality) is an inclusion (multiset equality) not between the specified two polynomials $f, t$, but between the polynomials generated by the multiplication of $f$ and $t$ with (generally speaking) independently generated selectors. We generalize to the vector setting.

Definition 6 (Selected Vector Arguments). We are given polynomials $f_{i}, t_{i} \in \mathbb{K}_{<n}[X]$ for $i \in[N]$. Furthermore, we are also given two polynomials $f^{\text {sel }}, t^{\text {sel }} \in \mathbb{F}_{<n}[X]$ whose range over the domain $G$ is $\{0,1\}$. That is, $f^{\text {sel }}$ and $t^{\text {sel }}$ are selectors. A selected vector inclusion, denoted $f^{\text {sel }} \cdot\left(f_{1}, \ldots, f_{N}\right) \in t^{\text {sel }} \cdot\left(t_{1}, \ldots, t_{N}\right)$, is the argument in which for all $x \in G$ there exists some $y \in G$ such that:

$$
\begin{equation*}
f^{\mathrm{sel}}(x) \cdot\left(f_{1}(x), \ldots, f_{N}(x)\right)=t^{\mathrm{sel}}(y) \cdot\left(t_{1}(y), \ldots, t_{N}(y)\right) \tag{22}
\end{equation*}
$$

where $f^{\text {sel }}(x) \cdot\left(f_{1}(x), \ldots, f_{N}(x)\right)$ denotes the component-wise scalar multiplication between the field element $f^{\text {sel }}(x)$ and the vector $\left(f_{1}(x), \ldots, f_{N}(x)\right)$.

A selected vector multiset equality, denoted $f^{\text {sel }} \cdot\left(f_{1}, \ldots, f_{N}\right) \doteq t^{\text {sel }} \cdot\left(t_{1}, \ldots, t_{N}\right)$, is the argument in which for all $y \in G$ there exists exactly one $x \in G$ for which Eq. (22) holds.
Remark 1. Note that if $f^{\text {sel }}=t^{\text {sel }}=1$, then Eq. (22) is reduced to (20); if $f^{\text {sel }}=t^{\text {sel }}=0$ then the argument is trivial; and if either $f^{\text {sel }}$ or $t^{\text {sel }}$ is equal to the constant 1 , then we remove the need for $f^{\text {sel }}$ or $t^{\text {sel }}$, respectively.

To reduce selected vector inclusion to simple ones, we proceed in two steps. First, we use the reduction shown in Eq. (21) to reduce the inner vector of polynomials to a single one. This process outputs polynomials $F^{\prime}, T^{\prime} \in \mathbb{K}_{<n}[X]$. Second, we make use of another uniformly sampled $\beta \in \mathbb{K}$ as follows. Namely, we define the following polynomials:

$$
\begin{align*}
& T(X):=t^{\mathrm{sel}}(X)\left[T^{\prime}(X)-\beta\right]+\beta \\
& F(X):=f^{\mathrm{sel}}(X)\left[F^{\prime}(X)-T(X)\right]+T(X) \tag{23}
\end{align*}
$$

and proceed to prove the relation $F \in T$.
Importantly, the presentation "re-ordering" in Eq. (23) is relevant: if $\beta$ had been introduced in the definition of $F$ instead, then there would be situations in which we would end up having $\beta$ as an inclusion value and therefore the inclusion argument not being satisfied even if the selectors are correct. See Example 2 to see why this is relevant.

Example 2. Choose $N=1, n=2^{3}$. We compute the following values:

| $x$ |
| :---: | :---: | :---: | :---: | :---: |
| $g$ |
| $g^{2}$ |
| $g^{3}$ |
| $g^{4}$ |
| $g^{5}$ |
| $g^{6}$ |
| $g^{7}$ |
| $g^{8}$ |$\quad$| $f_{1}(x)$ | $F^{\prime}(x)$ | $f^{\text {sel }}(x)$ | $F(x)$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 3 | 3 | 0 | 1 |
| 7 | 7 | 1 | 7 |
| 4 | 4 | 0 | 7 |
| 1 | 1 | 1 | 1 |
| 5 | 5 | 1 | 5 |
| 1 | 1 | 0 | 5 |
| 2 | 2 | 1 | 2 |
| 5 | 5 | 1 | 5 |
| $t_{1}(x)$ | $T^{\prime}(x)$ | $t^{\text {sel }}(x)$ | $T(x)$ |
| 1 | 1 | 1 | 1 |
| 7 | 1 | 0 | $\beta$ |
| 7 | 7 | 1 | 7 |
| 6 | 6 | 0 | $\beta$ |
| 5 | 5 | 1 | 5 |
| 5 | 5 | 1 | 5 |
| 5 | 5 | 0 | $\beta$ |
| 7 | 2 | 1 | 2 |

Notice how $F \in T$. However, if we would have instead defined $F, T$ as $F(X)=f^{\text {sel }}(X)\left[F^{\prime}(X)-\right.$ $\beta]+\beta$ and $T(X)=t^{\text {sel }}(X)\left[T^{\prime}(X)-F(X)\right]+F(X)$ then we would end up having $\beta$ as a inclusion value, which implies that $F \notin T$ even though $f_{1}, t_{1}$ and $f^{\text {sel }}, t^{\text {sel }}$ are correct.

To reduce selected vector multiset equalities to simple ones, we follow a similar process as with selected vector inclusions. We also first use the reduction in Eq. (21) to reduce the inner vector argument to a simple one, but then we define the following polynomials:

$$
\begin{align*}
& F(X):=f^{\mathrm{sel}}(X)\left[F^{\prime}(X)-\beta\right]+\beta \\
& T(X):=t^{\mathrm{sel}}(X)\left[T^{\prime}(X)-\beta\right]+\beta \tag{24}
\end{align*}
$$

and proceed to prove the relation $F \doteq T$. Here, we have been able to first define $F$ since we are dealing with multiset equalities instead of inclusions.

Similarly to Lemma 4, we obtain the following result. by observing that $\beta$ do not grow the total degree of polynomials $F, T$ (either from Eq. (23) or Eq. (24)) over variables $\alpha, \beta$.
Lemma 6. Given polynomials $f_{i}, t_{i} \in \mathbb{K}_{<n}[X]$ for $i \in[N]$, selectors $f^{\text {sel }}, t^{\text {sel }} \in \mathbb{K}_{<n}[X]$ and $F, T \in \mathbb{K}_{<n}[X]$ as defined by Eq. (23) (resp. Eq. (24)), if $F \in T$ (resp. $F \doteq T$ ), then $f^{\text {sel }} \cdot\left(f_{1}, \ldots, f_{N}\right) \in t^{\text {sel }} \cdot\left(t_{1}, \ldots, t_{N}\right)\left(\right.$ resp. $\left.f^{\text {sel }} \cdot\left(f_{1}, \ldots, f_{N}\right) \doteq t^{\text {sel }} \cdot\left(t_{1}, \ldots, t_{N}\right)\right)$ except with probability $n \cdot(N-1) /|\mathbb{K}|$ over the random and independent choice of $\alpha$ and $\beta$.

We generalize to selected vector arguments the protocols for (simple) inclusion arguments and multiset equality arguments explained in Section 2.5 by incorporating the reduction strategies explained in this section. Therefore, we give next the soundness bounds for these protocols.
Lemma 7. Given polynomials $f_{i}, t_{j} \in \mathbb{K}_{<n}[X]$ for $i \in[N]$ and selectors $f^{\text {sel }}, t^{\text {sel }} \in$ $\mathbb{K}_{<n}[X]$, we obtain:

1. Inclusion Protocol. Let $T \in \mathbb{K}_{<2 n-1}[X]$ and $F \in \mathbb{K}_{<3 n-1}[X]$ as defined by Eq. (23). The prover sends oracle functions $\left[f_{i}\right],\left[t_{i}\right],\left[f^{\text {sel }}\right],\left[t^{\text {sel }}\right]$ for $i \in[N]$ to the verifier in the first round, who responds with uniformly sampled $\alpha, \beta \in \mathbb{K}$. Moreover, enlarge the set of identities that must be checked by the verifier in the inclusion protocol of Section 2.5 with:

$$
\begin{aligned}
& f^{s e l}(x)\left(f^{s e l}(x)-1\right)=0 \\
& f^{\text {sel }}(x)\left(f^{\text {sel }}(x)-1\right)=0
\end{aligned}
$$

for all $x \in G$, i.e., the verifier checks that polynomials $f^{\text {sel }}, t^{\text {sel }}$ are valid selectors. If a prover that interacts with a verifier in the (enlarged) inclusion protocol of Section 2.5 on input $F$ and $T$ causes it to accept with probability greater than:

$$
n \frac{N-1}{|\mathbb{K}|}+\varepsilon_{\operatorname{lnc}}(3 n-1)
$$

then $f^{\text {sel }} \cdot\left(f_{1}, \ldots, f_{N}\right) \in t^{\text {sel }} \cdot\left(t_{1}, \ldots, t_{N}\right)$.
2. Multiset Equality Protocol. Let $F, T \in \mathbb{K}_{<2 n-1}[X]$ as defined by Eq. (24). The prover sends oracle functions $\left[f_{i}\right],\left[t_{i}\right],\left[f^{\text {sel }}\right],\left[t^{\text {sel }}\right]$ for $i \in[N]$ to the verifier in the first round, who responds with uniformly sampled $\alpha, \beta \in \mathbb{K}$. Moreover, enlarge the set of identities that must be checked by the verifier in the multiset equality protocol of Section 2.5 with:

$$
\begin{aligned}
& f^{s e l}(x)\left(f^{s e l}(x)-1\right)=0 \\
& f^{s e l}(x)\left(f^{s e l}(x)-1\right)=0
\end{aligned}
$$

for all $x \in G$. If a prover that interacts with a verifier in the (enlarged) multiset equality protocol of Section 2.5 on input $F$ and $T$ causes it to accept with probability greater than:

$$
n \frac{N-1}{|\mathbb{K}|}+\varepsilon_{M u \mid E q}(2 n-1)
$$

then $f^{\text {sel }} \cdot\left(f_{1}, \ldots, f_{N}\right) \doteq t^{\text {sel }} \cdot\left(t_{1}, \ldots, t_{N}\right)$.
Example 3. Say that for all $x \in G$ the prover wants to prove that he knows some polynomials $\operatorname{tr}_{1}, \operatorname{tr}_{2}, \operatorname{tr}_{3}, \operatorname{tr}_{4}, \operatorname{tr}_{5} \in \mathbb{F}_{<n}[X]$ such that:

$$
\begin{align*}
& \operatorname{tr}_{1} \in \operatorname{tr}_{3} \\
& \operatorname{tr}_{3} \doteq \operatorname{tr}_{4}  \tag{25}\\
&\left(\operatorname{tr}_{2}, \operatorname{tr}_{1}, \operatorname{tr}_{5}\right) \propto\left(S_{\sigma_{1}}, S_{\sigma_{2}}, S_{\sigma_{3}}\right)
\end{align*}
$$

where we have used the notation $\doteq$ to denote that $c$ and $d$ are a permutation of each other, without specifying a particular permutation.

Following the previous section and Section 3.5, the polynomial constraint system (25) gets transformed to the following one, so that for all $x \in G$ :

$$
\begin{gathered}
L_{1}(x)\left(Z_{1}(x)-1\right)=0, \\
Z_{1}(g x)=Z_{1}(x) \frac{(1+\beta)\left(\gamma+\operatorname{tr}_{1}(x)\right)\left(\gamma(1+\beta)+\operatorname{tr}_{3}(x)+\beta \operatorname{tr}_{3}(g x)\right)}{\left(\gamma(1+\beta)+h_{1,1}(x)+\beta h_{1,2}(x)\right)\left(\gamma(1+\beta)+h_{1,2}(x)+\beta h_{1,1}(g x)\right)}, \\
L_{1}(x)\left(Z_{2}(x)-1\right)=0, \\
Z_{2}(g x)=Z_{2}(x) \frac{\left(\gamma+\operatorname{tr}_{3}(x)\right)}{\left(\gamma+\operatorname{tr}_{4}(x)\right)}, \\
L_{1}(x)\left(Z_{3}(x)-1\right)=0, \\
\operatorname{im}_{1}(x)=\left(\operatorname{tr}_{1}(x)+\beta k_{1} x+\gamma\right)\left(\operatorname{tr}_{5}(x)+\beta k_{2} x+\gamma\right), \\
\operatorname{im}_{2}(x)=\left(\operatorname{tr}_{1}(x)+S_{\sigma_{2}}(x)+\gamma\right)\left(\operatorname{tr}_{5}(x)+S_{\sigma_{3}}(x)+\gamma\right), \\
Z_{3}(g x)=Z_{3}(x) \frac{\left(\operatorname{tr}_{2}(x)+\beta x+\gamma\right) \operatorname{im}_{1}(x)}{\left(\operatorname{tr}_{2}(x)+S_{\sigma_{1}}(x)+\gamma\right) \operatorname{im}_{2}(x)},
\end{gathered}
$$

where we notice that the only type of argument that sometimes need to be adjusted is the connection argument.

## Parallel Execution of the Arguments

We end this section by explaining the protocol corresponding to multiple executions of the previous protocols combined.

Protocol 2. The protocol starts with a set of polynomials $f_{i, j}, t_{i, j} \in \mathbb{F}_{<n}[X]$ for $i \in[N]$ and $j \in\left[M+M^{\prime}+M^{\prime \prime}\right]$ known to the prover. Here, for each $j \in[M],\left\{f_{i, j}, t_{i, j}\right\}_{i}$
correspond to the polynomials of each $M$ inclusion invocations; for each $j \in[M+1, M+$ $\left.M^{\prime}\right],\left\{f_{i, j}, t_{i, j}\right\}_{i}$ correspond to the polynomials of each $M^{\prime}$ multiset equality invocations and for each $j \in\left[M+M^{\prime}+1, M+M^{\prime}+M^{\prime \prime}\right],\left\{f_{i, j}\right\}_{i}$ correspond to the polynomials of each $M^{\prime \prime}$ connection invocations and $\left\{t_{i, j}\right\}_{i}$ correspond to the polynomials $\left\{S_{i, \sigma_{j}}\right\}_{i}$ derived from each permutation $\sigma_{j}$. For each $j \in\left[M+M^{\prime}\right]$, the prover possibly also knows selectors $f_{j}^{\text {sel }}, t_{j}^{\text {sel }}$.

1. Execution Trace Oracles: The prover sends oracle functions $\left[f_{i, j}\right],\left[t_{i, j}\right],\left[f_{j}^{\text {sel }}\right],\left[t_{j}^{\text {sel }}\right]$ for $i \in[N]$ and $j \in\left[M+M^{\prime}+M^{\prime \prime}\right]$ to the verifier, who responds with uniformly sampled values $\alpha, \beta \in \mathbb{K}$.
2. Inclusion Oracles: The prover computes the inclusion polynomials $h_{1, j}, h_{2, j}$ for each inclusion invocation $j \in[M]$. Then, he sends oracle functions of them to the verifier, who answers with uniformly sampled values $\gamma, \delta \in \mathbb{K}$.
3. Grand Product Oracles: The prover computes the grand product polynomials $Z_{j}$ for each argument $j \in\left[M+M^{\prime}+M^{\prime \prime}\right]$ and sends oracle functions of them to the verifier.
4. Verification: For each $j \in[M]$ and all $x \in G$, the verifier checks that constraints in Eq. (18) hold; for each $j \in\left[M+1, M+M^{\prime}\right]$, constraints in Eq. (14) hold; and for each $j \in\left[M+M^{\prime}+1, M+M^{\prime}+M^{\prime \prime}\right]$, constraints in Eq. (16) hold. Finally, the verifier also confirms that for each $j \in\left[M+M^{\prime}\right]$, the polynomials $f_{j}^{\text {sel }}, t_{j}^{\text {sel }}$ are valid selectors by checking the following constraints also hold:

$$
\begin{aligned}
& f_{j}^{\mathrm{sel}}(x)\left(f_{j}^{\mathrm{sel}}(x)-1\right)=0, \\
& f_{j}^{\mathrm{sel}}(x)\left(f_{j}^{\mathrm{sel}}(x)-1\right)=0 .
\end{aligned}
$$

| $\mathcal{P}\left(\left\{f_{i, j}, t_{i, j}\right\}_{i, j}, \mathbb{F}, \mathbb{K}\right)$ |  | $\mathcal{V}(\mathbb{F}, \mathbb{K})$ |
| :---: | :---: | :---: |
|  | $\left\{\left[f_{i, j}\right],\left[t_{i, j}\right],\left[f_{j}^{\text {sel }]}\right],\left[t_{j}^{\text {sel }}\right]\right\}_{i, j}$ |  |
|  | $\{\alpha, \beta\}$ |  |
|  | $\left\{\left[h_{1,1}\right],\left[h_{2,1}\right], \ldots,\left[h_{1, M}\right],\left[h_{2, M}\right]\right\}$ |  |
|  | $\{\gamma, \delta\}$ |  |
|  | $\left\{\left[Z_{1}\right], \ldots,\left[Z_{M+M^{\prime}+M^{\prime \prime}}\right]\right\}$ |  |

Figure 5: Skeleton description of Protocol 2.

Using Theorem 7 and the Parallel Repetition Theorem for polynomial IOPs [BCS16], [Gol98] we obtain the following result. Use $M_{1}, M_{2}, M_{3}$ to refer to the number of simple, vector and selected vector inclusions. We have $M=M_{1}+M_{2}+M_{3}$. For the multiset equality scenario, analogously define $M_{1}^{\prime}, M_{2}^{\prime}, M_{3}^{\prime}$, which also satisfy $M^{\prime}=M_{1}^{\prime}+M_{2}^{\prime}+M_{3}^{\prime}$.

Lemma 8 (Soundness bound for Protocol 2). Let $\varepsilon_{I n c}, \varepsilon_{M u I E q}, \varepsilon_{C o n}$ be the soundness for a single invocation of the protocols asserting the inclusion, multiset equality and connection arguments, respectively. Then if the prover interacts with the verifier in Protocol 2 and causes it to accept with probability greater than:

$$
\begin{aligned}
\varepsilon_{\text {Args }}:= & \left(M_{2}+M_{3}+M_{2}^{\prime}+M_{3}^{\prime}\right) \frac{n(N-1)}{|\mathbb{K}|}+ \\
& +\varepsilon_{\operatorname{lnc}}(n)^{M_{1}+M_{2}} \varepsilon_{\operatorname{lnc}}(2 n-1)^{M_{3}} \varepsilon_{\text {MulEq }}(n)^{M_{1}+M_{2}} \varepsilon_{\text {MulEq }}(3 n-1)^{M_{3}} \varepsilon_{\mathrm{Con}}(n)^{M^{\prime \prime}},
\end{aligned}
$$

then each of the $M$ inclusion arguments, $M^{\prime}$ multiset equality arguments and $M^{\prime \prime}$ connection arguments get satisifed.

Proof Sketch. First, notice that $\varepsilon_{\text {Args }}$ is the soundness bound in the "best-case" scenario, that is, assuming that the prover is lying in all the $M+M^{\prime}+M^{\prime \prime}$ arguments. An analogous expression for $\varepsilon_{\text {Args }}$ (and its corresponding proof) can be easily obtained.

For each $i \in\left[M_{2}+M_{3}+M_{2}^{\prime}+M_{3}^{\prime}\right]$, denote by $E_{i}$ the event in which the reduction from either the (vector or selected) inclusion argument or the (vector or selected) multiset equality argument is correct. Also, denote by $E$ the event in which every constraint corresponding to each of the arguments are satisfied. Then, using the union bound, the first term of $\varepsilon_{\text {Args }}$ corresponds to an upper bound to the probability $\operatorname{Pr}\left[\cup_{i} E_{i}\right]$. Now, apply the Parallel Repetition Theorem for polynomial IOPs and the second term of $\varepsilon_{\text {Args }}$ is obtained.

### 3.4 On the Quotient Polynomial

In the vanilla STARK protocol, the quotient polynomial $Q$ (Eq. 8) is computed by adjusting the degree of the rational functions:

$$
q_{i}(X):=\frac{C_{i}\left(\operatorname{tr}_{1}(X), \ldots, \operatorname{tr}_{N}(X), \operatorname{tr}_{1}(g X), \ldots, \operatorname{tr}_{N}(g X)\right)}{Z_{G}(X)}
$$

to a sufficiently large power of two $D$ with the help of two random values $\mathfrak{a}_{i}, \mathfrak{b}_{i}$. The sum of the resulting polynomials $\hat{q}_{i}:=\left(\mathfrak{a}_{i} X^{D-\operatorname{deg}\left(q_{i}\right)-1}+\mathfrak{b}_{i}\right) \cdot q_{i}(X)$ is precisely $Q$.

There are two major issues with the previous definition of the quotient polynomial: (1) it leads to an amount of uniformly sampled values $\mathfrak{a}_{i}, \mathfrak{b}_{i}$ proportional to the number of constraints; and (2) [Sta21] (or any other source, as far as we know) does not provide a proof of why the degree adjustment is necessary at all.

Therefore, we instead obtain a single random value $\mathfrak{a} \in \mathbb{K}$ and define the quotient polynomial as a random linear combination of the rational functions $q_{i}$ as follows:

$$
Q(X):=\sum_{i=1}^{\ell} \mathfrak{a}^{i-1} q_{i}(X)
$$

Note that we not only remove the degree adjustment of the $q_{i}$ 's but also use powers of a uniformly sampled value $\mathfrak{a}$ instead of sampling one value per constraint. A proof that this alternative way of computing the quotient polynomial is sound was carefully analyzed in Theorem 7 of [Hab22] (and based on Theorem 7.2 of [ $\left.\mathrm{BCI}^{+} 20\right]$ ). Importantly, the soundness bound of this alternative version is linearly increased by the number of constraints $\ell$, so we might assume from now on that $\ell$ is sublinear in $|\mathbb{K}|$ to ensure the security of protocols.

### 3.5 Controlling the Constraint Degree wih Intermediate Polynomials

In the vanilla STARK protocol, the initial set of constraints that one attest to compute the proof over is of unbounded degree. However, when one arrives at the point after computing the quotient polynomial $Q$, it should be split into polynomials of degree lower than $n$ to ensure the same redundancy is added as with the trace column polynomials $\operatorname{tr}_{i}$ for a sound application of the FRI protocol. In this section, we explain an alternative for this process and propose the split to happen "at the beginning" and not "at the end" of the proof computation.

Therefore, we will proceed with this approach assuming that the arguments in Section 2.5 are included among the initial set of constraints. The constraints imposed by the grand products polynomials $Z_{i}$ of multiset equalities and inclusions are of known degree:
degree 2 for the former and degree 3 for the latter. Based on this information, we will propose a splitting procedure that allows for polynomial constraints up to degree 3 but will split any exceeding it.

Say the initial set of polynomial constraints $\mathcal{C}=\left\{C_{1}, \ldots, C_{\ell}\right\}$ contain a constraint of total degree greater or equal to 4 . For instance, say that we have $\mathcal{C}=\left\{C_{1}, C_{2}\right\}$ with:

$$
\begin{align*}
& C_{1}\left(X_{1}, X_{2}, X_{3}, X_{1}^{\prime}, X_{2}^{\prime}, X_{3}^{\prime}\right)=X_{1} \cdot X_{2} \cdot X_{2}^{\prime} \cdot X_{3}^{\prime}-X_{3}^{3} \\
& C_{2}\left(X_{1}, X_{2}, X_{3}, X_{1}^{\prime}, X_{2}^{\prime}, X_{3}^{\prime}\right)=X_{2}-7 \cdot X_{1}^{\prime}+X_{3}^{\prime} \tag{26}
\end{align*}
$$

Now, instead of directly computing the (unbounded) quotient polynomial $Q$ and then doing the split, we will follow the following process:

1. Split the constraints of degree $t \geq 4$ into $\lceil t / 3\rceil$ constraints of degree lower or equal than 3 through the introduction of one formal variable and one constraint per split.
2. Compute the rational functions $q_{i}$. Notice the previous step restricts the degree of the $q_{i}$ 's to be lower than $2 n$.
3. Compute the quotient polynomial $Q \in \mathbb{F}_{<2 n}[X]$ and then split it into (at most) two polynomials $Q_{1}$ and $Q_{2}$ of degree lower than $n$ as follows:

$$
\begin{equation*}
Q(X)=Q_{1}(X)+X^{n} \cdot Q_{2}(X) \tag{27}
\end{equation*}
$$

where $Q_{1}$ is obtained by taking the first $n$ coefficients of $Q$ and $Q_{2}$ is obtained by taking the last $n$ coefficients (filling with zeros if necessary).
Remark 2. Here, we might have that $Q_{2}$ is identically equal to 0 . This is in contrast with the technique used for the split in Eq. (9), where the quotient polynomial $Q$ is distributed uniformly across each of the trace quotient polynomials $Q_{i}$.
This process will "control" the degree of $Q$ so that it will be always of a degree lower than $2 n$.

Following with the example in Eq. (26), we rename $C_{2}$ to $C_{3}$ and introduce the formal variable $Y_{1}$ and the constraint:

$$
\begin{equation*}
C_{2}\left(X_{1}, X_{2}, X_{3}, X_{1}^{\prime}, X_{2}^{\prime}, X_{3}^{\prime}, Y_{1}\right)=X_{1} \cdot X_{2}-Y_{1} \tag{28}
\end{equation*}
$$

Now, to compute the rational functions $q_{i}$, we have to compose $C_{2}$ not only with the trace column polynomials $\operatorname{tr}_{i}$ but also with additional polynomials corresponding with the introduced variables $Y_{i}$. We will denote these polynomials as $\mathrm{im}_{i}$ and refer to them as intermediate polynomials.

Hence, the set of constraints in (26) gets augmented to the following set:

$$
\begin{aligned}
& C_{1}\left(X_{1}, X_{2}, X_{3}, X_{1}^{\prime}, X_{2}^{\prime}, X_{3}^{\prime}, Y_{1}\right)=Y_{1} \cdot X_{2}^{\prime} \cdot X_{3}^{\prime}-X_{3}^{3} \\
& C_{2}\left(X_{1}, X_{2}, X_{3}, X_{1}^{\prime}, X_{2}^{\prime}, X_{3}^{\prime}, Y_{1}\right)=X_{1} \cdot X_{2}-Y_{1} \\
& C_{3}\left(X_{1}, X_{2}, X_{3}, X_{1}^{\prime}, X_{2}^{\prime}, X_{3}^{\prime}, Y_{1}\right)=X_{2}-7 \cdot X_{1}^{\prime}+X_{3}^{\prime}
\end{aligned}
$$

where we include the variable $Y_{1}$ in $C_{3}$ for notation simplicity. Note that now what we have is two constraints of degree lower than 3 , but we have added one extra variable and constraint to take into account.

Discussing more in-depth the tradeoff generated between the two approaches, we have for one side that $\operatorname{deg}(Q)=\max _{i}\left\{\operatorname{deg}\left(q_{i}\right)\right\}=\max _{i}\left\{\operatorname{deg}\left(C_{i}\right)(n-1)-|G|\right\}$. Denote by $i_{\text {max }}$ the index of the $q_{i}$ where the maximum is attained. Then, the number of polynomials $S$ in the split of $Q$ is equal to:

$$
\left\lceil\frac{\operatorname{deg}(Q)}{n}\right\rceil=\left\lceil\frac{\operatorname{deg}\left(C_{i_{\max }}\right)(n-1)-|G|}{n}\right\rceil=\operatorname{deg}\left(C_{i_{\max }}\right)+\left\lceil-\frac{|G|}{n}\right\rceil \text {, }
$$

which is equal to either $\operatorname{deg}\left(C_{i_{\max }}\right)-1$ or $\operatorname{deg}\left(C_{i_{\max }}\right)$.
We must compare this number with the number of additional constraints (or polynomials) added in our proposal. So, on the other side, we have that the overall number of constraints $\tilde{\ell}$ is:

$$
\sum_{i=1}^{\ell}\left\lceil\frac{\operatorname{deg}\left(C_{i}\right)}{3}\right\rceil
$$

with $\tilde{\ell} \geq \ell$.
We conclude that the appropriate approach should be chosen based on the minimum value between $\tilde{\ell}-\ell$ and $S$. Specifically, if the goal is to minimize the number of polynomials in the proof generation, then the vanilla STARK approach should be taken if $\min \{\tilde{\ell}-\ell, S\}=S$, and our approach should be taken if $\min \{\tilde{\ell}-\ell, S\}=\tilde{\ell}-\ell$.
Example 4. To give some concrete numbers, let us compare both approaches using the following set of constraints:

$$
\begin{aligned}
C_{1}\left(X_{1}, X_{2}, X_{3}, X_{4}, X_{1}^{\prime}\right) & =X_{1} \cdot X_{2}^{2} \cdot X_{3}^{4} \cdot X_{4}-X_{1}^{\prime} \\
C_{2}\left(X_{1}, X_{2}, X_{3}\right) & =X_{1} \cdot X_{2}^{3}+X_{3}^{3} \\
C_{3}\left(X_{2}, X_{3}, X_{4}, X_{2}^{\prime}\right) & =X_{2}^{3} \cdot X_{3} \cdot X_{4}+X_{2}^{\prime}
\end{aligned}
$$

In the vanilla STARK approach, we obtain $S=8$. On the other side, using the early splitting technique explained before, by substituting $X_{1} \cdot X_{2}^{2}$ by $Y_{1}$ and $X_{2} \cdot X_{3} \cdot X_{4}$ by $Y_{2}$ we transform the previous set of constraints into an equivalent one having all constraints of degree less or equal than 3 . This reduction only introduces 2 additional constraints:

$$
\begin{aligned}
C_{1}\left(X_{1}^{\prime}, Y_{1}, Y_{2}\right) & =Y_{1}^{2} \cdot Y_{2}-X_{1}^{\prime}, \\
C_{2}\left(X_{2}, X_{3}, Y_{1}\right) & =Y_{1} \cdot X_{2}+X_{3}^{3}, \\
C_{3}\left(X_{2}, X_{2}^{\prime}, Y_{2}\right) & =Y_{2} \cdot X_{2}^{2}+X_{2}^{\prime}, \\
C_{4}\left(X_{1}, X_{2}, Y_{1}\right) & =Y_{1}-X_{1} \cdot X_{2}^{2} \\
C_{5}\left(X_{2}, X_{3}, X_{4}, Y_{2}\right) & =Y_{2}-X_{2} \cdot X_{3} \cdot X_{4}
\end{aligned}
$$

Henceforth, the early splitting technique is convenient in this case, introducing 3 new polynomials instead of the 7 that proposes the vanilla STARK approach.

However, early splittings are not unique. That is, we can reduce the degree of the constraints differently, giving more polynomials and worsening our previous splitting in terms of numbers of polynomials. For example, the following set of constraints (achieved by substituting $X_{1} \cdot X_{2}^{2}$ by $Y_{1}, X_{3}^{3}$ by $Y_{2}, X_{3} \cdot X_{4}$ by $Y_{3}$ and $X_{2}^{3}$ by $Y_{4}$ ) is equivalent to the former ones, but in this case we added 4 extra polynomial constraints:

$$
\begin{aligned}
C_{1}\left(X_{1}^{\prime}, Y_{1}, Y_{2}, Y_{3}\right) & =Y_{1} \cdot Y_{2} \cdot Y_{3}-X_{1}^{\prime} \\
C_{2}\left(X_{2}, Y_{1}, Y_{2}\right) & =Y_{1} \cdot X_{2}+Y_{2} \\
C_{3}\left(X_{2}^{\prime}, Y_{3}, Y_{4}\right) & =Y_{3} \cdot Y_{4}+X_{2}^{\prime} \\
C_{4}\left(X_{1}, X_{2}, Y_{1}\right) & =Y_{1}-X_{1} \cdot X_{2}^{2} \\
C_{5}\left(X_{3}, Y_{2}\right) & =Y_{2}-X_{3}^{3} \\
C_{6}\left(X_{3}, X_{4}, Y_{3}\right) & =Y_{3}-X_{3} \cdot X_{4} \\
C_{7}\left(X_{2}, Y_{4}\right) & =Y_{4}-X_{2}^{3}
\end{aligned}
$$

On the other side, a system of constraints composed by the following kind of constraints is not easily early-reducible:

$$
C_{i}\left(X_{i}, X_{i+1}, X_{i+2}\right)=X_{i}^{3} \cdot X_{i+1}+X_{i+1}^{3} \cdot X_{i}+X_{i+2}
$$

More specifically, each $C_{i}$ added into our constraints system will increase by 2 the number of polynomial constraints if the early splitting technique is used. Informally, these constraints do not have repetitions in the monomials composing them, not allowing to generate optimal substitutions as done before. Therefore, even having only one of such constraints, Vanilla STARK approach is preferable.

That being said, a careful and sophisticated analysis should be used in order to choose the optimal solution between both approaches. However, as a rule of thumb, our approach is preferable whenever only a few constraints exceed degree 3 or/and there exists several repetitions among the monomials of the exceeding constraints.

### 3.6 FRI Polynomial Computation

Recall from Section 2.4 the F polynomial was computed as follows:

$$
\begin{aligned}
\mathrm{F}(X):= & \sum_{i \in I_{1}} \varepsilon_{i}^{(1)} \cdot \frac{\operatorname{tr}_{i}(X)-\operatorname{tr}_{i}(z)}{X-z}+\sum_{i \in I_{2}} \varepsilon_{i}^{(2)} \cdot \frac{\operatorname{tr}_{i}(g X)-\operatorname{tr}_{i}(g z)}{X-g z} \\
& +\sum_{i=1}^{S} \varepsilon_{i}^{(3)} \cdot \frac{Q_{i}(X)-Q_{i}\left(z^{S}\right)}{X-z^{S}},
\end{aligned}
$$

where $I_{1}=\left\{i \in[N]: \operatorname{tr}_{i}(z) \in \operatorname{Evals}(z)\right\}, I_{2}=\left\{i \in[N]: \operatorname{tr}_{i}(g z) \in \operatorname{Evals}(g z)\right\}$ and $\varepsilon_{i}^{(1)}, \varepsilon_{j}^{(2)}, \varepsilon_{k}^{(3)} \in \mathbb{K}$ for all $i \in I_{1}, j \in I_{2}, k \in[S]$. This way of computing the F polynomial has again (see Section 3.4) the issue that the number of random values sent from the verifier is proportional to the number of polynomials involved in the previous sum.

We will therefore compute the F polynomial by requesting two random values $\varepsilon_{1}, \varepsilon_{2} \in$ $\mathbb{K}$ instead, using $\varepsilon_{1}$ to compute the part regarding evaluations at $z$ and $g z$ separately and finally mixing it all with $\varepsilon_{1}$.

Following with the previous example, we define polynomials $\mathrm{F}_{1}, \mathrm{~F}_{2} \in \mathbb{K}_{<n}[X]$ :

$$
\begin{aligned}
& \mathrm{F}_{1}(X):=\sum_{i \in I_{1}} \varepsilon_{2}^{i-1} \cdot \frac{\operatorname{tr}_{i}(X)-\operatorname{tr}_{i}(z)}{X-z}+\sum_{i=1}^{S} \varepsilon_{2}^{\left|I_{1}\right|+i-1} \cdot \frac{Q_{i}(X)-Q_{i}(z)}{X-z} \\
& \mathrm{~F}_{2}(X):=\sum_{i \in I_{2}} \varepsilon_{2}^{i-1} \cdot \frac{\operatorname{tr}_{i}(g X)-\operatorname{tr}_{i}(g z)}{X-g z}
\end{aligned}
$$

and then we set $\mathrm{F}(X):=\mathrm{F}_{1}(X)+\varepsilon_{1} \cdot \mathrm{~F}_{2}(X)$. Note that since $\varepsilon_{1}, \varepsilon_{2}$ are uniformly sampled elements, then so is $\varepsilon_{1} \cdot \varepsilon_{2}^{i}$ for all $i \geq 0$.

A commonly used alternative version of the $F$ polynomial computation in practice involves requesting a single random value $\varepsilon \in \mathbb{K}$ and directly computing

$$
\begin{aligned}
\widetilde{\mathrm{F}}(X):= & \sum_{i \in I_{1}} \varepsilon^{i-1} \cdot \frac{\operatorname{tr}_{i}(X)-\operatorname{tr}_{i}(z)}{X-z}+\sum_{i \in I_{2}} \varepsilon^{\left|I_{1}\right|+i-1} \cdot \frac{\operatorname{tr}_{i}(g X)-\operatorname{tr}_{i}(g z)}{X-g z} \\
& +\sum_{i=1}^{S} \varepsilon^{\left|I_{1}\right|+\left|I_{2}\right|+i-1} \cdot \frac{Q_{i}(X)-Q_{i}\left(z^{S}\right)}{X-z^{S}} .
\end{aligned}
$$

This version has the disadvantage of not being computable in parallel like the previous version, so we prefer the first option (even if it means increasing the proof size by one field element). Specifically, when the powers of $\varepsilon_{2}$ are being computed, it is possible to compute the polynomials $F_{1}$ and $F_{2}$ both sequentially and in parallel, while $\widetilde{F}$ can only be computed sequentially after the computation of the powers of $\varepsilon$.

## 4 Our eSTARK Protocol

### 4.1 Extended Algebraic Intermediate Representation (eAIR)

In this section, we introduce the notion of eAIRs and eAIR satisfiability as a natural extension of the well-studied AIRs [BBHR19]. Informally speaking, an eAIR is an AIR whose expressiveness is extended with more types of allowed constraints. In the following recall that $G=\langle g\rangle$ is a multiplicative subgroup of $\mathbb{F}$ of order $n$.

Definition 7 (AIR and AIR Satisfiability). Given polynomials $p_{1}, \ldots, p_{M} \in \mathbb{F}[X]$, an algebraic intermediate representation (AIR) A is a set of algebraic constraints $\left\{C_{1}, \ldots, C_{K}\right\}$ such that each $C_{i}$ is a polynomial over $\mathbb{F}\left[X_{1}, \ldots, X_{M}, X_{1}^{\prime}, \ldots, X_{M}^{\prime}\right]$. For each $C_{i}$, the first half variables $X_{1}, \ldots, X_{M}$ will be replaced by polynomials $p_{1}(X), \ldots, p_{M}(X)$, whereas the second half variables $X_{1}^{\prime}, \ldots, X_{M}^{\prime}$ will be replaced by polynomials $p_{1}(g X), \ldots, p_{M}(g X)$.

Moreover, we say that polynomials $p_{1}, \ldots, p_{M}$ satisfy a given $A I R \mathrm{~A}=\left\{C_{1}, \ldots, C_{K}\right\}$ if and only if for each $i \in[K]$ we have that:

$$
C_{i}\left(p_{1}(x), \ldots, p_{M}(x), p_{1}(g x), \ldots, p_{M}(g x)\right)=0, \forall x \in G .
$$

Remark 2. There are two main simplifications between our definition for AIR and the definition for AIR in [Sta21]: (1) we define constraints only over the "non-shifted" and "shifted-by-one" version of the corresponding polynomials, i.e., $p_{i}(X)$ and $p_{i}(g X)$, respectively; and (2) we enforce constraints to vanish over the whole $G$ and not over a subset of it. The following definitions can, however, support a more generic version.

Now, we extend the definition of an AIR by allowing the arguments defined in Section 3.3 as new types of available constraints.

Definition 8 (Extended AIR). Given polynomials $p_{1}, \ldots, p_{M} \in \mathbb{F}[X]$, an extended algebraic intermediate representation (eAIR) eA is a set of constraints eA $=\left\{C_{1}, \ldots, C_{K}\right\}$ such that each $C_{i}$ can be one of the following form:
(a) A polynomial over $\mathbb{F}\left[X_{1}, \ldots, X_{M}, X_{1}^{\prime}, \ldots, X_{M}^{\prime}\right]$ as in Def. 7 .
(b) A positive integer $R_{i}$, a set of $2 R_{i}$ polynomials $C_{i, j}$ over $\mathbb{F}\left[X_{1}, \ldots, X_{M}\right]$ and two selectors $f_{i}^{\text {sel }}, t_{i}^{\text {sel }}$ over $\mathbb{F}[X]$ (recall that $f_{i}^{\text {sel }}(x), t_{i}^{\text {sel }}(x) \in\{0,1\}$ for all $\left.x \in G\right)$.
(c) An integer $S_{i} \in[M]$, a subset of $S_{i}$ polynomials $p_{(1)}, \ldots, p_{\left(S_{i}\right)} \in \mathbb{F}[X]$ from the set $\left\{p_{1}, \ldots, p_{M}\right\}$ and $S_{i}$ more polynomials $S_{\sigma_{i}, 1}, \ldots, S_{\sigma_{i}, S_{i}}$ representing a permutation $\sigma_{i}$.

Finally, we refer to the set constraints of the form described in $(a)$ as the set of identity constraints of eA and to the set of constraints of the form described by either (b) or (c) as the set of non-identity constraints of eA.

Definition 8 aims to capture the arguments in Section 3.3 in a slightly more generic way. Here, the polynomials subject to these arguments are generated as a polynomial combination between $p_{1}, \ldots, p_{M}$.

In what follows, we use $\overline{\mathrm{P}}$ as a shorthand for $\left(p_{1}, \ldots, p_{M}\right)$ and denote by $C \circ \overline{\mathrm{P}}$ to the univariate polynomial over $\mathbb{F}[X]$ resulting from the substitution of each of the variables $X_{i}, X_{i}^{\prime}$ of the constraint $C \in \mathbb{F}\left[X_{1}, \ldots, X_{M}, X_{1}^{\prime}, \ldots, X_{M}^{\prime}\right]$ by $p_{i}(X), p_{i}(g X)$, respectively. That is $C \circ \overline{\mathrm{P}}$ is the polynomal $C\left(p_{1}(X), \ldots, p_{M}(X), p_{1}(g X), \ldots, p_{M}(g X)\right)$.

Definition 9 (Extended AIR Satisfiability). We say that polynomials $p_{1}, \ldots, p_{M} \in \mathbb{F}[X]$ satisfy a given eAIR eA $=\left\{C_{1}, \ldots, C_{K}\right\}$ if and only if for each $i \in[K]$ one and only one
of the following is true for all $x \in G$ :

$$
\begin{aligned}
&\left(C_{i} \circ \overline{\mathrm{P}}\right)(x)=0 \\
& f_{i}^{\mathrm{sel}}(x) \cdot\left(\left(C_{i, 1} \circ \overline{\mathrm{P}}\right)(x), \ldots,\left(C_{i, R_{i}} \circ \overline{\mathrm{P}}\right)(x)\right) \in t_{i}^{\mathrm{sel}}(x) \cdot\left(\left(C_{i, R_{i}+1} \circ \overline{\mathrm{P}}\right)(x), \ldots,\left(C_{i, 2 R_{i}} \circ \overline{\mathrm{P}}\right)(x)\right), \\
& f_{i}^{\mathrm{sel}}(x) \cdot\left(\left(C_{i, 1} \circ \overline{\mathrm{P}}\right)(x), \ldots,\left(C_{i, R_{i}} \circ \overline{\mathrm{P}}\right)(x)\right) \doteq t_{i}^{\mathrm{sel}}(x) \cdot\left(\left(C_{i, R_{i}+1} \circ \overline{\mathrm{P}}\right)(x), \ldots,\left(C_{i, 2 R_{i}} \circ \overline{\mathrm{P}}\right)(x)\right), \\
&\left(p_{(1)}(x), \ldots, p_{\left(S_{i}\right)}(x)\right) \propto\left(S_{\sigma_{i}, 1}(x), \ldots, S_{\sigma_{i}, S_{i}}(x)\right) .
\end{aligned}
$$

### 4.2 The Setup Phase

During the protocol of Section 4.3, both the prover and the verifier will need to have access to a set of preprocessed polynomials pre $i_{i} \in \mathbb{F}[X]$. In particular, the prover will need to have full access to them, either in coefficient or in the evaluation form, to be able to correctly generate the proof. On the other hand, the verifier will only need to have access to a subset of the evaluations of these polynomials over the domain $H$.

To this end, in our protocol we assume the existence of a phase, known as the setup phase, that is before the protocol message exchange but after the particular statement to be proven (or equivalently, the set of constraints that describe the statement) is fixed. In the setup phase, the preprocessed polynomials are computed and the prover and the verifier receive different information regarding them. Particularly, the setup phase, with input from a set of polynomial constraints, consists of the following steps:

1. The trace LDE of each preprocessed polynomial is computed.
2. The Merkle tree of the set of preprocessed polynomials is computed.
3. Finally, the complete tree is sent to the prover and its corresponding root is sent to the verifier. This way, when the verifier needs to compute the evaluation of any preprocessed polynomial over $h \in H$, he can request it from the prover, who will respond with the evaluation along with its corresponding Merkle tree path. The verifier then verifies the accuracy of the information received by using the root of the tree.

Remark 3. Since the computational effort of the setup phase is greater than $\mathcal{O}(\log (n))$, we cannot include this phase as part of the verifier description if we want our protocol to satisfy verifier scalability.

Note that the setup phase does not include a measure for the verifier to be sure that the computation of the Merkle tree of preprocessed polynomials is correct. However, as the setup phase input is the set of polynomial constraints representing the problem's statement (something that the verifier also knows), the verifier can run at any time during the setup phase to check the validity of the computations.

Moreover, both the prover and the verifier will need to have access to the evaluations over $H$ of the vanishing polynomial $Z_{G}(X):=X^{n}-1$ and the first Lagrange polynomial $L_{1}(X):=\frac{g\left(X^{n}-1\right)}{n(X-g)}$. However, $Z_{G}$ and $L_{1}$ will appear later on in the protocol, and although in principle we do not consider them preprocessed polynomials, they are publicly known and therefore included in the Merkle tree computation of the setup phase.

### 4.3 Our IOP for eAIR

Before the start of the protocol, we assume the prover and verifier have fixed a specific eAIR instance eA $=\left\{\widetilde{C}_{1}, \ldots, \widetilde{C}_{T^{\prime}}\right\}$ and that the constraints of eA are ordered as follows: first, the identity constraints, then the inclusion arguments, followed by the permutation
arguments, and finally, the connection arguments. Additionally, we assume that the setup phase has been successfully executed.

Throughout the description of the protocol, we use the following useful notation:

- Let $N, R$ be two non-negative integers. We set $N$ to be the number of trace column polynomials and $R$ to be the number of preprocessed polynomials.
- Among the set of polynomial constraints eA, we denote by $M, M^{\prime}, M^{\prime \prime} \in \mathbb{Z}_{\geq 0}$ the number of inclusion arguments, permutation arguments and connection arguments, respectively.
- The prover parameters pp is composed of the finite field $\mathbb{F}$, the domains $G$ and $H$, the field extension $\mathbb{K}$, the eAIR instance eA, all the public values, the set of committed polynomials and the Merkle tree of preprocessed polynomials.
- The verifier parameters vp is composed of the finite field $\mathbb{F}$, the domains $G$ and $H$, the field extension $\mathbb{K}$, the eAIR instance eA, all the public values and the Merkle tree's root of preprocessed polynomials.

Our IOP for eAIR, which can be seen as an extension of the DEEP-ALI protocol [BGKS19], is as follows.

Protocol 3 (IOP for eAIR). The protocol starts with a set of trace column polynomials $\operatorname{tr}_{1}, \ldots, \operatorname{tr}_{N} \in \mathbb{F}_{<n}[X]$ and preprocessed polynomials pre $1_{1}, \ldots, \operatorname{pre}_{R} \in \mathbb{F}_{<n}[X]$. The following protocol is used by a prover to prove to a verifier that polynomials $\operatorname{tr}_{1}, \ldots, \operatorname{tr}_{N}, \operatorname{pre}_{1}, \ldots, \operatorname{pre}_{R}$ satisfy eA:

1. Trace Column Oracles: The prover sets oracle functions $\left[\operatorname{tr}_{1}\right], \ldots,\left[\operatorname{tr}_{N}\right]$ to $\operatorname{tr}_{1}, \ldots, \operatorname{tr}_{N} \in$ $\mathbb{F}_{<n}[X]$ for the verifier, who responds with uniformly sampled values $\alpha, \beta \in \mathbb{K}$. During this step, the prover also computes the intermediate polynomials resulting from the subset of identity constraints. Let $K \in \mathbb{Z}_{\geq 0}$ be the number of these polynomials, denoted as $\operatorname{im}_{i} \in \mathbb{F}_{<n}[X]$, where $i \in[K]$. It is important to recall that new identity constraints must also be considered when introducing intermediate polynomials, as demonstrated in Eq. (28). As additional intermediate polynomials may be introduced in Round 3, the prover will set oracles for $\mathrm{im}_{1}, \ldots, \mathrm{im}_{K}$ in that round.
2. Inclusion Oracles: As explained by Section 3.3, the prover, if needed:

- Uses $\alpha$ to reduce both vector inclusions and vector permutations into simple (possibly selected) inclusions and permutations.
- Uses $\beta$ to reduce both selected inclusions and selected permutations into simple (non-selected) inclusions and permutations.

After the previous two reductions, the prover computes the inclusion polynomials $h_{i, 1}, h_{i, 2} \in \mathbb{K}_{<n}[X]$ for each inclusion argument, with $i \in[M]$. Then, he sets oracle functions $\left[h_{1,1}\right],\left[h_{1,2}\right], \ldots,\left[h_{M, 1}\right],\left[h_{M, 2}\right]$ for the verifier, who answers with uniformly sampled values $\gamma, \delta \in \mathbb{K}$.
3. Grand Product and Intermediate Oracles: The prover uses $\gamma, \delta$ to compute the grand product polynomials $Z_{i} \in \mathbb{K}_{<n}[X]$ for each argument, with $i \in\left[M+M^{\prime}+M^{\prime \prime}\right]$. Importantly, some identity constraints induced by the constraints asserting the validity of the connection argument's grand product polynomials might be of a degree greater or equal to 4 . Therefore, following Section 3.5, the prover split these constraints into multiple constraints of degree at most 3 by the introduction of intermediate polynomials. Let $K^{\prime} \in \mathbb{Z}_{\geq 0}$ be the number of introduced intermediate polynomials, denoted as $\operatorname{im}_{K+i} \in \mathbb{K}[\bar{X}]$, where $i \in\left[K^{\prime}\right]$.

The prover sets oracle functions $\left[Z_{1}\right], \ldots,\left[Z_{M+M^{\prime}+M^{\prime \prime}}\right]$ and $\left[\mathrm{im}_{1}\right], \ldots,\left[\mathrm{im}_{K+K^{\prime}}\right]$ for the verifier. The verifier answers with a uniformly sampled value $\mathfrak{a} \in \mathbb{K}$.
At this point, the original eAIR eA $=\left\{\widetilde{C}_{1}, \ldots, \widetilde{C}_{T^{\prime}}\right\}$ has been reduced to an AIR A $=\left\{C_{1}, \ldots, C_{T}\right\}$, with $T \geq T^{\prime}$, so we continue by executing the DEEP-ALI protocol over A with the modifications mentioned in Sections 3.4 and 3.5.

Remark 3. Rounds 2 and 3 are skipped by both the prover and the verifier if the eAIR instance eA is an AIR. In such case, Round 4 follows from Round 1.
4. Trace Quotient Oracles: The prover computes the polynomial $Q(X) \in \mathbb{K}[X]$ :

$$
\begin{equation*}
Q(X):=\sum_{i=1}^{T} \mathfrak{a}^{i-1} \frac{\left(C_{i} \circ \overline{\mathrm{P}}\right)(X)}{Z_{G}(X)} \tag{29}
\end{equation*}
$$

where we have now used $\overline{\mathrm{P}}$ to denote the sequence of polynomials containing $\operatorname{tr}_{1}, \ldots, \operatorname{tr}_{N}$, $\operatorname{pre}_{1}, \ldots, \operatorname{pre}_{R}, h_{1,1}, h_{1,2}, \ldots, h_{M, 1}, h_{M, 2}, Z_{1}, \ldots, Z_{M+M^{\prime}+M^{\prime \prime}}$ and finally $\operatorname{im}_{1}, \ldots, \operatorname{im}_{K+K^{\prime}}$. Therefore, $\left(C_{i} \circ \overline{\mathrm{P}}\right)(X)$ is the (univariate) polynomial resulting from the composition of the (multivariate) polynomial $C_{i}$ and the non-shifted and shifted version of the (univariate) polynomials in $\overline{\mathrm{P}}$. Then, the prover splits $Q$ into two trace quotient polynomials $Q_{1}$ and $Q_{2}$ of degree lower than $n$, and sets their oracles [ $Q_{1}$ ] and [ $Q_{2}$ ] for the verifier. Polynomials $Q_{1}$ and $Q_{2}$ satisfy the following:

$$
\begin{equation*}
Q_{1}(X)+X^{n} \cdot Q_{2}(X)=\sum_{i=1}^{T} \mathfrak{a}^{i-1} \frac{\left(C_{i} \circ \overline{\mathrm{P}}\right)(X)}{Z_{G}(X)} \tag{30}
\end{equation*}
$$

5. DEEP Query Answers: The verifier samples a uniformly sampled DEEP query $z \in \mathbb{K} \backslash(G \cup H)$. Note that the verifier prohibits $z \in G$ to enable evaluation of the right-hand side of Eq. (30) and prohibits $z \in H$ to enable evaluation of the polynomial $F$, defined in Round 6, during the FRI protocol. Then, the verifier queries either the oracles set by the prover in previous rounds or the oracles to the preprocessed polynomials to obtain the evaluation sets Evals(z) and Evals $(g z)$. Two observations should be made: first, it is possible for a polynomial to have evaluations in both $\operatorname{Evals}(z)$ and $\operatorname{Evals}(g z)$; and second, evaluations of preprocessed polynomials are also included within $\operatorname{Evals}(z)$ and Evals $(g z)$. The verifier then sends to the prover uniformly sampled values $\varepsilon_{1}, \varepsilon_{2} \in \mathbb{K}$.
6. FRI Protocol: Among the set of polynomals in $\overline{\mathrm{P}}$, respectively denote by $f_{i}$ and $h_{i}$ those whose evaluation respectively belong to Evals $(z)$ and $\operatorname{Evals}(g z)$. The prover computes the polynomials $\mathrm{F}_{1}, \mathrm{~F}_{2} \in \mathbb{K}[X]$ :

$$
\begin{aligned}
& \mathrm{F}_{1}(X):=\sum_{i=1}^{|\operatorname{Evals}(z)|} \varepsilon_{2}^{i-1} \frac{f_{i}(X)-f_{i}(z)}{X-z} \\
& \mathrm{~F}_{2}(X):=\sum_{i=1}^{|\operatorname{Evals}(g z)|} \varepsilon_{2}^{i-1} \frac{h_{i}(X)-h_{i}(g z)}{X-g z},
\end{aligned}
$$

after which he computes the polynomial $\mathrm{F}(X):=\mathrm{F}_{1}(X)+\varepsilon_{1} \cdot \mathrm{~F}_{2}(X)$. Finally, the prover and the verifier run the FRI protocol to prove the low degree of $F$, which starts by setting oracle access to $F$ for the verifier.
7. Verification: Similar to the vanilla STARK verifier, the verifier proceeds as follows:
(a) ALI Consistency. Checks that the trace quotient polynomials $Q_{1}$ and $Q_{2}$ are consistent with the trace column polynomials $\operatorname{tr}_{1}, \ldots, t r_{N}$, the preprocessed polynomials pre ${ }_{1}, \ldots$, pre $_{R}$, the inclusion-related polynomials $h_{1,1}, h_{1,2}$, $\ldots, h_{M, 1}, h_{M, 2}$, the grand product polynomials $Z_{1}, \ldots, Z_{M+M^{\prime}+M^{\prime \prime}}$ and the intermediate polynomials $\operatorname{im}_{1}, \ldots, \operatorname{im}_{K+K^{\prime}}$. The verifier achieves so by means of Eq. (30) and the evaluations in $\operatorname{Evals}(z)$ and $\operatorname{Evals}(g z)$.
(b) Batched FRI Verification. It runs the batched FRI verification procedure on the polynomial $F$.

If either $(a)$ or $(b)$ fails at any point, the verifier aborts and rejects. Otherwise, the verifier accepts.

It is very straightforward to give an upper bound for the soundness of Protocol 3.
Theorem 3 (STIK for eAIR). Protocol 3 constitutes a STIK for extended AIR satisfiability, i.e., this protocol can be used to prove possession of a set of polynomials $p_{1}, \ldots, p_{N+R} \in \mathbb{F}_{<|G|}[X]$ satisfying a given extended AIR instance eA $=\left\{\widetilde{C}_{1}, \ldots, \widetilde{C}_{T^{\prime}}\right\}$. In particular, denote by $\varepsilon_{\text {eSTARK }}$ to the soundness error of the protocol. Then:

$$
\varepsilon_{e S T A R K}=\varepsilon_{A r g s}+\ell\left(\frac{T-1}{|\mathbb{K}|}+\frac{(D+|G|+2) \cdot \ell}{|\mathbb{K}|-|H|-|G|}\right)+\varepsilon_{F R I}
$$

where $\ell=m / \rho, m \geq 3$ is an integer, $D$ is the maximal degree of the polynomials involved that are summed in the computation of the trace quotient polynomial $Q$ (29), $\varepsilon_{\text {Args }}$ corresponds to the soundness error for Protocol 2 and $\varepsilon_{F R I}$ corresponds to the soundness error of the batched FRI protocol over the FRI polynomial F.

Before starting with the proof, let us discuss the main differences between the common part of this soundness bound and the bound of Theorem 2:

1. In the computation of the quotient polynomial $Q$ we are using powers of a random value $\mathfrak{a}^{i-1}$ instead of a pair $\left(\mathfrak{a}_{i}, \mathfrak{b}_{i}\right)$ for each constraint $C_{i}$. This leads to a polynomial of degree $T-1$ in the variable $\mathfrak{a}$ and therefore the Schwartz-Zippel lemma increases the bound from $1 /|\mathbb{K}|$ to $(T-1) /|\mathbb{K}|$.
2. Since we are splitting the quotient polynomial $Q$ into polynomials $Q_{1}, Q_{2}$ that follow a linear relation with $Q$ (i.e., that $Q_{1}(X)+X^{n} Q_{2}(X)=Q(X)$ ) instead of exponential (i.e., that $\left.Q_{1}\left(X^{2}\right)+X Q_{2}\left(X^{2}\right)=Q(X)\right)$ we only need to restrict the sampling space for $z$ at $\mathbb{K} \backslash(G \cup H)$ instead of $\mathbb{K} \backslash(G \cup \bar{H})$.

Proof Sketch. We start the proof by restricting it to the case of eA being an AIR, that is, we assume that no arguments are among the set of constraints. This means that neither Round 2 nor Round 3 of the previous protocol are needed and are therefore skipped. The resulting protocol turns out to be the (modified) vanilla STARK resulting from the combination of the DEEP-ALI protocol and the batched version of the FRI protocol. A proof that this combination results in a sound protocol is typically divided into two parts: (1) first, the soundness analysis of the DEEP-ALI part can be found in Theorem 6.2 of [BGKS19]; (2) second, the proof that the batched FRI as to applied in Section 4.4 is sound can be found in Theorem 8.3 of $\left[\mathrm{BCI}^{+} 20\right]$. The proof that the combination of these two sound protocols leads to a secure protocol can be found in Corollary 2 of [Sta21]. Since in this case $\varepsilon_{\text {Args }}=0$, we conclude that $\varepsilon_{\text {STARK }}=\varepsilon_{\text {eStARK }}$.

What remains to prove is the expression for the soundness of the protocol when there are arguments present in eA. However, this is achieved by noticing that the first term of $\varepsilon_{\text {STARK }}$ is the soundness error for the incorporated arguments and the rest of the terms describe the soundness error of the DEEP-ALI part plus the batched FRI protocol. In other words, we are accounting for a prover being "lucky" either in the former or in the latter part of the protocol.

### 4.4 From a STIK to a Non-Interactive STARK

As described in Definition 4, transforming the STIK of Section 4.3 to a STARK is very straightforward. First, oracles sent from the prover to the verifier in each round are substituted by a single Merkle tree root as explained in Section 3.1. So, for instance, oracle access to $\left[\operatorname{tr}_{1}\right], \ldots,\left[\operatorname{tr}_{N}\right]$ set by the prover in the first round is substituted by the Merkle tree root of the Merkle tree containing the evaluations of $\operatorname{tr}_{1}, \ldots, \operatorname{tr}_{N}$ over the domain $H$. Second, instead of letting the verifier ask a query to a set of polynomial oracles $\left[f_{1}\right], \ldots,\left[f_{N}\right]$ at the same point $v$, the verifier asks for these evaluations to the prover and the prover answers with $f_{1}(v), \ldots, f_{N}(v)$ together with the Merkle path associated with them. We do not specify the specific hash function used in the computation of each Merkle tree, but we only state that this hash function does not have any relation to the hash function used to render the protocol non-interactive, as long as its output space is in the appropriate field.

To render the following protocol in a non-interactive manner, we use the Fiat-Shamir heuristic [FS87]. To this end, let $\mathcal{H}:\{0,1\}^{*} \rightarrow \mathbb{K}$ denote the hash function, modeled as a random oracle in the soundness analysis, that instantiates the heuristic. We emphasize that $\mathcal{H}$ takes any number of inputs and returns elements of $\mathbb{K}$.

We denote by seed to the concatenation of the initial eAIR instance eA $=\left\{\widetilde{C}_{1}, \ldots, \widetilde{C}_{T^{\prime}}\right\}$, all the public values and the Merkle tree's root of preprocessed polynomials. This value will act as the seed to the hash function $\mathcal{H}$ to simulate the first verifier's message.

Moreover, to simulate subsequent verifier's messages we use a technique called hash chaining:

Definition 10 (Hash Chain). Let $\mathcal{H}:\{0,1\}^{*} \rightarrow \mathbb{K}$ be a hash function and let $a_{1}, \ldots, a_{r}$ be the messages sent from the prover to the verifier in an interactive protocol, where each $a_{i}$ represent a set of $\mathbb{K}$ elements. A hash chain is the set of field elements $\beta_{1}, \beta_{2}, \ldots, \beta_{r}$ obtained as follows:

$$
\begin{aligned}
\beta_{1} & =\mathcal{H}\left(\text { seed }, a_{1}\right), \\
\beta_{2} & =\mathcal{H}\left(\beta_{1}, a_{2}\right), \\
\beta_{3} & =\mathcal{H}\left(\beta_{2}, a_{3}\right), \\
\vdots & \\
\beta_{r} & =\mathcal{H}\left(\beta_{r-1}, a_{r}\right) .
\end{aligned}
$$

A proof that the resulting non-interactive protocol is sound after applying the FiatShamir using hash chaining can be found, for example, in Theorem 4 of [AFK21].

Remark 4. In Definition 10 we have avoided writing the generalization of having multiple verifier's messages sent at the same round. In such case, we compute them as $\beta_{i, j}=$ $\mathcal{H}\left(\beta_{i-1}, a_{i}, j\right)$ for $i \in[r]$ and $j=1,2, \ldots$ As a consequence, in round $i+1$ we must include every generated $\beta_{i, j}$ to compute the verifier' messages, i.e., $\beta_{i+1, j}=\mathcal{H}\left(\left\{\beta_{i, j}\right\}_{j}, a_{i+1}, j\right)$ with $k=1,2, \ldots$

### 4.5 Full Protocol Description

We split the protocol's description between the prover algorithm and verifier algorithm and compose each round in the prover algorithm of the computation of the verifier's challenges (via Fiat-Shamir) and the actual messages computed by the prover. We reuse the notation and assumptions from Section 4.3.

## PROVER ALGORITHM

## Round 1: Trace Column Polynomials

Given the trace column polynomials $\operatorname{tr}_{1}, \operatorname{tr}_{2}, \ldots, \operatorname{tr}_{N} \in \mathbb{F}_{<n}[X]$, the prover commits to them, as explained in Section 3.1, computing their associated Merkle root.

During this step, the prover also computes the intermediate polynomials $\mathrm{im}_{1}, \ldots, \mathrm{im}_{K}$ resulting from the subset of identity constraints.

The first prover output is $\operatorname{MTR}\left(\operatorname{tr}_{1}, \ldots, \operatorname{tr}_{N}\right)$.

## Round 2: Icnlusion Polynomials

The prover computes the reducing challenges $\alpha, \beta \in \mathbb{K}$ :

$$
\alpha=\mathcal{H}\left(\text { seed }, \operatorname{MTR}\left(\left\{\operatorname{tr}_{i}\right\}_{i}\right), 1\right), \quad \beta=\mathcal{H}\left(\text { seed }, \operatorname{MTR}\left(\left\{\operatorname{tr}_{i}\right\}_{i}\right), 2\right)
$$

Using $\alpha$ and $\beta$, the prover computes the inclusion polynomials $h_{i, 1}, h_{i, 2} \in \mathbb{K}_{<n}[X]$ for each inclusion argument, with $i \in[M]$, and commits to them.

The second prover output is $\operatorname{MTR}\left(h_{1,1}, h_{1,2}, \ldots, h_{M, 1}, h_{M, 2}\right)$.

## Round 3: Grand Product and Intermediate Polynomials

The prover computes the permutation challenges $\gamma, \delta \in \mathbb{K}$ :

$$
\gamma=\mathcal{H}\left(\alpha, \beta, \operatorname{MTR}\left(\left\{h_{i, j}\right\}_{i, j}\right), 1\right), \quad \delta=\mathcal{H}\left(\alpha, \beta, \operatorname{MTR}\left(\left\{h_{i, j}\right\}_{i, j}\right), 2\right) .
$$

The prover uses $\gamma, \delta$ to compute the grand product polynomials $Z_{i} \in \mathbb{K}_{<n}[X]$ for each argument, with $i \in\left[M+M^{\prime}+M^{\prime \prime}\right]$. The prover also computes the remaining intermediate polynomials $\operatorname{im}_{K+1}, \ldots, \operatorname{im}_{K+K^{\prime}}$.

In this round, the prover commits to both the grand product polynomials and all the intermediate polynomials. To save one element in the proof, the prover uses the same Merkle tree for both the grand product and the intermediate polynomials.

The third prover output is $\operatorname{MTR}\left(Z_{1}, \ldots, Z_{M+M^{\prime}+M^{\prime \prime}}, \operatorname{im}_{1}, \ldots, \operatorname{im}_{K+K^{\prime}}\right)$.

## Round 4: Trace Quotient Polynomials

The prover computes the combination challenge $\mathfrak{a} \in \mathbb{K}$ :

$$
\left.\mathfrak{a}=\mathcal{H}\left(\gamma, \delta, \operatorname{MTR}\left(\left\{Z_{i}\right\}_{i}\right)\right), \operatorname{MTR}\left(\left\{\operatorname{im}_{i}\right\}_{i}\right)\right) .
$$

The prover uses $\mathfrak{a}$ to compute the quotient polynomial $Q \in \mathbb{K}_{<2 n}[X]$ :

$$
Q(X):=\sum_{i=1}^{T} \mathfrak{a}^{i-1} \frac{\left(C_{i} \circ \overline{\mathrm{P}}\right)(X)}{Z_{G}(X)},
$$

and splits it into two polynomials $Q_{1}$ and $Q_{2}$ of degree lower than $n$ as in Eq. (27). Then, the prover commits to these two polynomials.

Recall that $Q_{1}, Q_{2}$ satisfy the following relation with $Q$ :

$$
\begin{equation*}
Q_{1}(X)+X^{n} \cdot Q_{2}(X)=\sum_{i=1}^{T} \mathfrak{a}^{i-1} \frac{\left(C_{i} \circ \overline{\mathrm{P}}\right)(X)}{Z_{G}(X)} \tag{31}
\end{equation*}
$$

The fourth prover output is $\operatorname{MTR}\left(Q_{1}, Q_{2}\right)$.

## Round 5: DEEP Query Answers

The prover computes the evaluation challenge $z \in \mathbb{K}$ :

$$
\begin{equation*}
z=\mathcal{H}\left(\mathfrak{a}, \operatorname{MTR}\left(Q_{1}, Q_{2}\right)\right) \tag{32}
\end{equation*}
$$

If $z$ falls either in $G$ or $H$, then we introduce a counter as an extra input to the hash function $\mathcal{H}$ in Eq. (32) and keep incrementing it until $z \in \mathbb{K} \backslash(G \cup H)$. The probability that $z$ is not of the expected form is proportional to $(|G|+|H|) /|\mathbb{K}|$, which is sufficiently small for all practical purposes.

The prover computes the evaluation sets Evals $(z)$ and $\operatorname{Evals}(g z)$.
The fifth prover output is (Evals(z), Evals $(g z))$.

## Round 6: FRI Protocol

The prover computes the combination challenges $\varepsilon_{1}, \varepsilon_{2} \in \mathbb{K}$ :

$$
\varepsilon_{1}=\mathcal{H}(z, \operatorname{Evals}(z), \operatorname{Evals}(g z), 1), \quad \varepsilon_{2}=\mathcal{H}(z, \operatorname{Evals}(z), \operatorname{Evals}(g z), 2)
$$

Among the set of polynomials in $\overline{\mathrm{P}}$, respectively denote by $f_{i}$ and $h_{i}$ those whose evaluation respectively belong to $\operatorname{Evals}(z)$ and $\operatorname{Evals}(g z)$. The prover computes the polynomials $\mathrm{F}_{1}, \mathrm{~F}_{2} \in \mathbb{K}[X]:$

$$
\begin{aligned}
& \mathrm{F}_{1}(X):=\sum_{i=1}^{|\operatorname{Evals}(z)|} \varepsilon_{2}^{i-1} \frac{f_{i}(X)-f_{i}(z)}{X-z} \\
& \mathrm{~F}_{2}(X):=\sum_{i=1}^{|\operatorname{Evals}(g z)|} \varepsilon_{2}^{i-1} \frac{h_{i}(X)-h_{i}(g z)}{X-g z}
\end{aligned}
$$

after which he computes the polynomial $\mathrm{F}(X):=\mathrm{F}_{1}(X)+\varepsilon_{1} \cdot \mathrm{~F}_{2}(X)$. Finally, the prover executes the (non-interactive version of the) FRI protocol to prove the low degree of the F polynomials, after which he obtains a FRI proof $\pi_{\text {FRI }}$.

The sixth prover output is $\pi_{\text {FRI }}$.

| $\mathcal{P}(\mathrm{pp})$ |  |
| :---: | :---: |
|  | $\operatorname{MTR}\left(\operatorname{tr}_{1}, \ldots, \operatorname{tr}_{N}\right)$ |
|  | $\{\alpha, \beta\}$ |
|  | $\operatorname{MTR}\left(h_{1,1}, h_{1,2}, \ldots, h_{M, 1}, h_{M, 2}\right)$ |
|  | $\{\gamma, \delta\}$ |
|  | $\underline{\operatorname{MTR}\left(Z_{1}, \ldots, Z_{M+M^{\prime}+M^{\prime \prime}}, \operatorname{im}_{1}, \ldots, \operatorname{im}_{K+K^{\prime}}\right)}$ |
|  | $\mathfrak{a}$ |
|  | $\operatorname{MTR}\left(Q_{1}, Q_{2}\right)$ |
|  | $z$ |
|  | $\{\operatorname{Evals}(z), \operatorname{Evals}(g z)\}$ |
|  | $\varepsilon_{1}, \varepsilon_{2}$ |
|  | $\pi_{\text {FRI }}$ |

Figure 6: Skeleton description of the eSTARK protocol.

The prover returns:

$$
\pi_{\mathrm{eSTARK}}=\left(\begin{array}{c}
\operatorname{MTR}\left(\operatorname{tr}_{1}, \ldots, \operatorname{tr}_{N}\right), \operatorname{MTR}\left(h_{1,1}, h_{1,2}, \ldots, h_{M, 1}, h_{M, 2}\right), \\
\operatorname{MTR}\left(Z_{1}, \ldots, Z_{M+M^{\prime}+M^{\prime \prime}}, \operatorname{im}_{1}, \ldots, \operatorname{im}_{K+K^{\prime}}\right), \\
\operatorname{MTR}\left(Q_{1}, Q_{2}\right), \operatorname{Evals}(z), \operatorname{Evals}(g z), \pi_{\mathrm{FRI}}
\end{array}\right)
$$

## VERIFIER ALGORITHM

To verify the STARK, the verifier performs the following steps:

1. Checks that $\operatorname{MTR}\left(\operatorname{tr}_{1}, \ldots, \operatorname{tr}_{N}\right), \operatorname{MTR}\left(h_{1,1}, h_{1,2}, \ldots, h_{M, 1}, h_{M, 2}\right), \operatorname{MTR}\left(Z_{1}, \ldots, Z_{M+M^{\prime}+M^{\prime \prime}}\right.$, $\left.\operatorname{im}_{1}, \ldots, \operatorname{im}_{K+K^{\prime}}\right), \operatorname{MTR}\left(Q_{1}, Q_{2}\right)$ are all in $\mathbb{K}$.
2. Checks that every element in both $\operatorname{Evals}(z)$ and $\operatorname{Evals}(g z)$ is from $\mathbb{K}$.
3. Computes the challenges $\alpha, \beta, \gamma, \delta, \mathfrak{a}, z, \varepsilon_{1}, \varepsilon_{2} \in \mathbb{K}$ from the elements in seed and $\pi_{\text {eSTARK }}$ (except for $\pi_{\text {FRI }}$, which is used to compute the challenges within FRI).
4. ALI Consistency. Checks that the trace quotient polynomials $Q_{1}$ and $Q_{2}$ are consistent with the trace column polynomials $\operatorname{tr}_{1}, \ldots, t r_{N}$, the preprocessed polynomials pre ${ }_{1}, \ldots$, pre $_{R}$, the inclusion-related polynomials $h_{1,1}, h_{1,2}, \ldots, h_{M, 1}, h_{M, 2}$, the grand product polynomials $Z_{1}, \ldots, Z_{M+M^{\prime}+M^{\prime \prime}}$ and the intermediate polynomials $\operatorname{im}_{1}, \ldots, \operatorname{im}_{K+K^{\prime}}$. The verifier achieves so by means of Eq. (31) and the evaluations contained in Evals $(z)$ and Evals $(g z)$.
5. Batched FRI Verification. Using $\pi_{\text {FRI }}$, it runs the batched FRI verification procedure on the polynomial F.

If either of the previous steps fails at any point, the verifier aborts and rejects. Otherwise, the verifier accepts.

## 5 Conclusions

In this paper, we have presented a probabilistic proof that generalizes the STARK family by introducing a more generic intermediate representation that we have called eAIR. We first explained multiple techniques that enhance the vanilla STARK complexity in both proof size and verification time. In particular, we demonstrated many optimizations applied to some polynomial computations in vanilla STARK. Additionally, we showed the tradeoffs arising from controlling the constraint degree either at the representation of the AIR or inside the eSTARK itself, offering a rule to decide whether to choose the first option or the second one. We anticipate these techniques to be useful for other types of SNARKs.

Secondly, we described our protocol in the polynomial IOP model as a combination of the optimized version of the vanilla STARK and the addition of rounds concerning the incorporation of three arguments into the protocol. Following the description, we proved that the protocol is sound in the polynomial IOP model.

Lastly, we provided a full description of the protocol by replacing the oracle access to polynomials via Merkle trees and turning it non-interactive through the Fiat-Shamir heuristic. We expect this protocol to be further expanded with the addition of more types of arguments that could fit a wider range of applications.

## References

[AFK21] Thomas Attema, Serge Fehr, and Michael Klooß. Fiat-shamir transformation of multi-round interactive proofs. Cryptology ePrint Archive, Report 2021/1377, 2021. https://eprint.iacr.org/2021/1377.
[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.
$\left[\mathrm{BBB}^{+} 18\right]$ Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In 2018 IEEE Symposium on Security and Privacy, pages 315-334. IEEE Computer Society Press, May 2018.
[BBHR18a] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Fast reedsolomon interactive oracle proofs of proximity. In Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella, editors, ICALP 2018, volume 107 of LIPIcs, pages 14:1-14:17. Schloss Dagstuhl, July 2018.
[BBHR18b] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Scalable, transparent, and post-quantum secure computational integrity. Cryptology ePrint Archive, Report 2018/046, 2018. https://eprint.iacr.org/2018/ 046.
[BBHR19] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Scalable zero knowledge with no trusted setup. In Alexandra Boldyreva and Daniele Micciancio, editors, CRYPTO 2019, Part III, volume 11694 of $L N C S$, pages 701-732. Springer, Heidelberg, August 2019.
[BCCT11] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. Cryptology ePrint Archive, Report 2011/443, 2011. https://eprint.iacr.org/2011/443.
$\left[\mathrm{BCI}^{+} 20\right]$ Eli Ben-Sasson, Dan Carmon, Yuval Ishai, Swastik Kopparty, and Shubhangi Saraf. Proximity gaps for reed-solomon codes. Cryptology ePrint Archive, Report 2020/654, 2020. https://eprint.iacr.org/2020/654.
[BCS16] Eli Ben-Sasson, Alessandro Chiesa, and Nicholas Spooner. Interactive oracle proofs. Cryptology ePrint Archive, Report 2016/116, 2016. https://eprint. iacr.org/2016/116.
[BGKS19] Eli Ben-Sasson, Lior Goldberg, Swastik Kopparty, and Shubhangi Saraf. DEEP-FRI: Sampling outside the box improves soundness. Cryptology ePrint Archive, Report 2019/336, 2019. https://eprint.iacr.org/2019/336.
[CBBZ22] Binyi Chen, Benedikt Bünz, Dan Boneh, and Zhenfei Zhang. HyperPlonk: Plonk with linear-time prover and high-degree custom gates. Cryptology ePrint Archive, Report 2022/1355, 2022. https://eprint.iacr.org/2022/ 1355.
$\left[\mathrm{CHM}^{+} 19\right]$ Alessandro Chiesa, Yuncong Hu, Mary Maller, Pratyush Mishra, Noah Vesely, and Nicholas Ward. Marlin: Preprocessing zkSNARKs with universal and updatable SRS. Cryptology ePrint Archive, Report 2019/1047, 2019. https: //eprint.iacr.org/2019/1047.
[COS19] Alessandro Chiesa, Dev Ojha, and Nicholas Spooner. Fractal: Post-quantum and transparent recursive proofs from holography. Cryptology ePrint Archive, Report 2019/1076, 2019. https://eprint.iacr.org/2019/1076.
[FS87] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, CRYPTO'86, volume 263 of LNCS, pages 186-194. Springer, Heidelberg, August 1987.
[GMR85] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof-systems (extended abstract). In 17 th ACM STOC, pages 291-304. ACM Press, May 1985.
[Gol98] Oded Goldreich. Modern Cryptography, Probabilistic Proofs and Pseudorandomness, volume 17 of Algorithms and Combinatorics. Springer, 1998.
[Gro16] Jens Groth. On the size of pairing-based non-interactive arguments. Cryptology ePrint Archive, Report 2016/260, 2016. https://eprint.iacr.org/ 2016/260.
[GW20] Ariel Gabizon and Zachary J. Williamson. plookup: A simplified polynomial protocol for lookup tables. Cryptology ePrint Archive, Report 2020/315, 2020. https://eprint.iacr.org/2020/315.
[GWC19] Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. PLONK: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge. Cryptology ePrint Archive, Report 2019/953, 2019. https: //eprint.iacr.org/2019/953.
[Hab22] Ulrich Haböck. A summary on the FRI low degree test. Cryptology ePrint Archive, Report 2022/1216, 2022. https://eprint.iacr.org/2022/1216.
[KHSS22] Daniel Kang, Tatsunori Hashimoto, Ion Stoica, and Yi Sun. Zk-img: Attested images via zero-knowledge proofs to fight disinformation, 2022.
[KZG10] Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In Masayuki Abe, editor, ASIACRYPT 2010, volume 6477 of $L N C S$, pages 177-194. Springer, Heidelberg, December 2010.
[LXZ21] Tianyi Liu, Xiang Xie, and Yupeng Zhang. zkCNN: Zero knowledge proofs for convolutional neural network predictions and accuracy. Cryptology ePrint Archive, Report 2021/673, 2021. https://eprint.iacr.org/2021/673.
[Mer88] Ralph C. Merkle. A digital signature based on a conventional encryption function. In Carl Pomerance, editor, CRYPTO'87, volume 293 of LNCS, pages 369-378. Springer, Heidelberg, August 1988.
$\left[\mathrm{PFM}^{+} 22\right]$ Luke Pearson, Joshua Fitzgerald, Héctor Masip, Marta Bellés-Muñoz, and Jose Luis Muñoz-Tapia. PlonKup: Reconciling PlonK with plookup. Cryptology ePrint Archive, Report 2022/086, 2022. https://eprint.iacr.org/ 2022/086.
[Set19] Srinath Setty. Spartan: Efficient and general-purpose zkSNARKs without trusted setup. Cryptology ePrint Archive, Report 2019/550, 2019. https: //eprint.iacr.org/2019/550.
[Sta21] StarkWare. ethSTARK documentation. Cryptology ePrint Archive, Report 2021/582, 2021. https://eprint.iacr.org/2021/582.
[VP19] Alexander Vlasov and Konstantin Panarin. Transparent polynomial commitment scheme with polylogarithmic communication complexity. Cryptology ePrint Archive, Report 2019/1020, 2019. https://eprint.iacr.org/2019/ 1020.
[Wil16] Nathan Wilcox. Zcash protocol specification, version 2.0-alpha-1. 2016.
$\left[\mathrm{WTs}^{+}{ }^{17]}\right.$ Riad S. Wahby, Ioanna Tzialla, abhi shelat, Justin Thaler, and Michael Walfish. Doubly-efficient zkSNARKs without trusted setup. Cryptology ePrint Archive, Report 2017/1132, 2017. https://eprint.iacr.org/2017/1132.

## A Vanilla STARK Description

In this section, we will outline the STARK protocol step by step. The following definitions will be employed in the protocol:

- We call prover parameters, and denote them by pp, to the finite field $\mathbb{F}$, the domains $G$ and $H$, the field extension $\mathbb{K}$, the set of constraints, and the set of polynomial evaluations that the prover has access to.
- We call verifier parameters, and denote them by vp, to the finite field $\mathbb{F}$, the domains $G$ and $H$, the field extension $\mathbb{K}$, the set of constraints, and the set of polynomial evaluations that the verifier has access to.

Protocol 4. The protocol starts with a set of constraints $\mathcal{C}=\left\{C_{1}, \ldots, C_{T}\right\}$ and a set of polynomials $\operatorname{tr}_{1}, \ldots, \operatorname{tr}_{N} \in \mathbb{F}[X]$ (presumably) satisfying them. The prover and verifier proceed as follows:

1. The prover sends oracle functions $\left[\operatorname{tr}_{1}\right], \ldots,\left[\operatorname{tr}_{N}\right]$ for $\operatorname{tr}_{1}, \ldots, \operatorname{tr}_{N}$ to the verifier:

| $\mathcal{P}(\mathrm{pp})$ | $\left[\operatorname{tr}_{1}\right], \ldots,\left[\operatorname{tr}_{N}\right]$ | $\mathcal{V}(\mathrm{vp})$ |
| :---: | :---: | :---: |
|  |  |  |

2. The verifier samples uniformly random values $\mathfrak{a}_{1}, \mathfrak{b}_{1}, \ldots, \mathfrak{a}_{T}, \mathfrak{b}_{T} \in \mathbb{K}$ and sends them to the prover:

| $\mathcal{P}(\mathrm{pp})$ |  | $\mathcal{V}(\mathrm{vp})$ |
| :---: | :---: | :---: |
|  | $\left[\operatorname{tr}_{1}\right], \ldots,\left[\operatorname{tr}_{N}\right]$ |  |
|  | $\left\{\mathfrak{a}_{1}, \mathfrak{b}_{1}, \ldots, \mathfrak{a}_{T}, \mathfrak{b}_{T}\right\}$ |  |
|  |  |  |

3. The prover computes the quotient polynomial $Q$, after which he computes the trace quotient polynomials $Q_{1}, \ldots, Q_{S}$ and sends their oracle representatives $\left[Q_{1}\right], \ldots,\left[Q_{S}\right]$ to the verifier.

| $\mathcal{P}(\mathrm{pp})$ |  | $\mathcal{V}(\mathrm{vp})$ |
| :---: | :---: | :---: |
|  | $\left[\operatorname{tr}_{1}\right], \ldots,\left[\operatorname{tr}_{N}\right]$ |  |
|  | $\left\{\mathfrak{a}_{1}, \mathfrak{b}_{1}, \ldots, \mathfrak{a}_{T}, \mathfrak{b}_{T}\right\}$ |  |
|  | $\left[Q_{1}\right], \ldots,\left[Q_{S}\right]$ |  |

4. The verifier samples a uniformly random value $z \in \mathbb{K} \backslash(G \cup \bar{H})$ and sends it to the prover:

| $\mathcal{P}$ (pp) |  | $\mathcal{V}(\mathrm{vp})$ |
| :---: | :---: | :---: |
|  | $\left[\operatorname{tr}_{1}\right], \ldots,\left[\operatorname{tr}_{N}\right]$ |  |
|  | $\left\{\mathfrak{a}_{1}, \mathfrak{b}_{1}, \ldots, \mathfrak{a}_{T}, \mathfrak{b}_{T}\right\}$ |  |
|  | $\left[Q_{1}\right], \ldots,\left[Q_{S}\right]$ |  |
|  | $z$ |  |

5. The prover computes the sets Evals $(z)$ and $\operatorname{Evals}(g z)$ of trace column polynomials evaluations. Then, he sends Evals $(z)$, Evals $(g z)$ and $Q_{i}\left(z^{S}\right)$ for all $i \in[S]$ to the verifier:

| $\mathcal{P}$ (pp) |  | $\mathcal{V}(\mathrm{vp})$ |
| :---: | :---: | :---: |
|  | $\left[\operatorname{tr}_{1}\right], \ldots,\left[\operatorname{tr}_{N}\right]$ |  |
|  | $\left\{\mathfrak{a}_{1}, \mathfrak{b}_{1}, \ldots, \mathfrak{a}_{T}, \mathfrak{b}_{T}\right\}$ |  |
|  | $\left[Q_{1}\right], \ldots,\left[Q_{S}\right]$ |  |
|  | $z$ |  |
|  | $\xrightarrow{\left\{\operatorname{Evals}(z), \operatorname{Evals}(g z), Q_{1}\left(z^{S}\right), \ldots, Q_{S}\left(z^{S}\right)\right\}}$ |  |

6. Recalling that $I_{1}=\left\{i \in[N]: \operatorname{tr}_{i}(z) \in \operatorname{Evals}(z)\right\}$ and $I_{2}=\left\{i \in[N]: \operatorname{tr}_{i}(g z) \in\right.$ Evals $(g z)\}$, the verifier sends uniformly random values $\varepsilon_{i}^{(1)}, \varepsilon_{j}^{(2)}, \varepsilon_{k}^{(3)} \in \mathbb{K}$, where $i \in I_{1}, j \in I_{2}, k \in[S]$, to the prover:

| $\mathcal{P}(\mathrm{pp})$ |  | $\mathcal{V}(\mathrm{vp})$ |
| :---: | :---: | :---: |
|  | $\left[\operatorname{tr}_{1}\right], \ldots,\left[\operatorname{tr}_{N}\right]$ |  |
|  | $\left\{\mathfrak{a}_{1}, \mathfrak{b}_{1}, \ldots, \mathfrak{a}_{T}, \mathfrak{b}_{T}\right\}$ |  |
|  | $\left[Q_{1}\right], \ldots,\left[Q_{S}\right]$ |  |
|  | $z$ |  |
|  | $\left\{\operatorname{Evals}(z), \operatorname{Evals}(g z), Q_{1}\left(z^{S}\right), \ldots, Q_{S}\left(z^{S}\right)\right\}$ |  |
|  | $\left\{\left\{\varepsilon_{i}^{(1)}\right\}_{i \in I_{1}},\left\{\varepsilon_{j}^{(2)}\right\}_{j \in I_{2}},\left\{\varepsilon_{k}^{(3)}\right\}_{k \in[S]}\right\}$ |  |

7. The prover computes the F polynomial. Then, the prover initiates the FRI message exchange along with the verifier to prove the low degree of the $F$ polynomial to the verifier. The prover sends the resulting FRI proof $\pi_{\text {FRI }}$ to the verifier:

| $\mathcal{P}(\mathrm{pp})$ |  | $\mathcal{V}$ (vp) |
| :---: | :---: | :---: |
|  | $\left[\operatorname{tr}_{1}\right], \ldots,\left[\operatorname{tr}_{N}\right]$ |  |
|  | $\left\{\mathfrak{a}_{1}, \mathfrak{b}_{1}, \ldots, \mathfrak{a}_{T}, \mathfrak{b}_{T}\right\}$ |  |
|  | $\left[Q_{1}\right], \ldots,\left[Q_{S}\right]$ |  |
|  | $z$ |  |
|  | $\left\{\operatorname{Evals}(z), \operatorname{Evals}(g z), Q_{1}\left(z^{S}\right), \ldots, Q_{S}\left(z^{S}\right)\right\}$ |  |
|  | $\left\{\left\{\varepsilon_{i}^{(1)}\right\}_{i \in I_{1}},\left\{\varepsilon_{j}^{(2)}\right\}_{j \in I_{2}},\left\{\varepsilon_{k}^{(3)}\right\}_{k \in[S]}\right\}$ |  |
|  | $\pi_{\text {FRI }}$ |  |

8. The verifier runs the three checks commented in Section 2.4: trace consistency, FRI consistency and batch consistency. If any of the checks fail, the verifier rejects them. Otherwise, the verifier accepts.


Figure 7: Full STARK transcript of Section 2.4


[^0]:    ${ }^{1}$ In our case, $a_{i}=2$ for all $i$. We decide to not explicitly write it to capture a version of FRI in which some layers of the commit phase can be skipped.

[^1]:    ${ }^{2}$ The claim in [GWC19] is for a slightly more general protocol.

