A note on adding zero-knowledge to STARK

Ulrich Haböck* and Al Kindi
ulrich@starkware.co, al-kindi@polygon.technology

February 20, 2025

Abstract

We discuss zero-knowledge in the context of univariate argument systems which use the FRI proximity test for Reed-Solomon codes [BBHR18] as polynomial commitment scheme. We confine ourselves to small-field STARK, i.e. arguments with an arithmetization over a small finite field (the basefield), and we dwell on two techniques widely used in practice: Randomization by polynomials over the basefield, and decomposing the overall quotient into polynomials of smaller degree. In particular the latter is a source for mistakes, both in literature as well as in software implementations. The current, updated version further includes a separate discussion on perfect zero-knowledge in permutation arguments.

1 Introduction

Adding zero-knowledge to a STARK[1](for Scalable and Transparent ARgument of Knowledge) is a subject that is somewhat neglected in the field of applied cryptography. Most papers focus solely on the soundness of the underlying interactive oracle proof (IOP), and leave its modification for zero-knowledge to the reader, referring to [BSCR+19] as one of the few examples which treats the issue in full depth. One reason for this might be that STARK are built from polynomial IOPs, and it does not seem a subtle task to randomize polynomials properly, so that their queried values reveal nothing about the witness. However, appearances may be deceptive. We found gaps in the treatment of zero-knowledge in the codebases of Plonky2 [Plo], Risc-Zero [Ris], Triton [Tri], not to forget the learning resource [Hab22] on FRI. (In the meantime, the gaps have been patched, and the summary on FRI has been updated.) These gaps were in the protocol design, exposing misinterpretations in the measures necessary for

*Formerly at Polygon Labs, where most of this work was written.

[1]We keep with the "bad practice" prevalent among practitioners, and consider a STARK any univariate Reed-Solomon encoded argument system that uses the FRI low-degree test.

achieving perfect zero-knowledge.[2]They where either due to an oversimplified analysis of the underlying polynomial IOP, or an underestimated impact of FRI queries on the composite protocol, or both.

This note is devoted to the subtleties of turning STARK into zero-knowledge. We review the [BSCR$^+$19] construction and clarify the importance of its masking step, which was fundamentally misinterpreted in a previous version of [Hab22], and we discuss two techniques widely used in practice: First (1), the randomization of witness polynomials over the base field (while the protocol challenges are from an extension field), and second (2), the decomposition of the overall quotient into polynomials of smaller degree. The first issue is not too tricky; however, we did not find a formal treatment in literature. On the contrary, the quotient decomposition is a source of pitfalls, with an earlier version of Plonk [GWC19] as an prominent example. We pay particular attention to the FFT-type decomposition of the quotient,

$$q(x) = q_0(x^d) + x^2 \cdot q_1(x^d) + \ldots + x^{d-1} \cdot q_{d-1}(x^d),$$

which is prevalent in the context of STARK. Unlike the monomial decomposition, which splits the quotient into polynomials of consecutive powers of x, or the more efficient decomposition by value (dubbed Lagrange decomposition in [HLP24]), the FFT-type decomposition is not amenable for randomization of its component polynomials, and its simulator analysis is particularly delicate: Our transcript simulator relies on the explicit construction of the splitting field for

$$v_D(x^d) = x^{n \cdot d} - a^n,$$

where $v_D(x) = x^n - a^n$ is the vanishing polynomial of the evaluation domain of the Reed-Solomon code, and its runtime analysis is closely related to that of polynomial factorization over an cyclotomic extension. It is an open question whether one can avoid such an explicit construction, without imposing additional constraints on the field and the evaluation domain (which we find artificial for the mere purpose of the proof).

This revision includes the discussion of another issue, which came up during closing the gap in the plonky2 code base, namely (3) perfect zero-knowledge in permutation arguments. In the fully private case (Plonk's wiring argument, or a permutation argument in a memory consistency proof) the incompleteness of the grand product argument leaks information about the witness, and obtaining perfect zero-knowledge is more subtle than proposed in the Halo2 book [BLH$^+$, Section 4].

The document is arranged as follows. In Section 2 we review the [BSCR$^+$19] construction in the context of FRI as a "small field polynomial commitment scheme", where the randomized witness polynomials are over a small prime field, but the queries are from an extension field. In Section 3, we describe how

[2]However, we stress the fact that none of the gaps caused exploitable vulnerabilities.

to add zero-knowledge to a STARK of a simple algebraic intermediate representation (AIR) with transitional constraints, which uses the above mentioned FFT decomposition of the quotient. This restrictive example is for demonstration purposes; yet it is general enough to carve out the main difficulties. The remaining sections are of supplementary character: In Section 4 we quickly sketch the randomization of other types of quotient decomposition, and Section 5 concludes with few remarks on the computational overhead introduced by zero-knowledge. In particular, we argue that in many applications a greedy choice of randomization parameters causes negligible extra costs for the prover. Appendix A briefly covers perfect zero-knowledge in permutation arguments.

Disclaimer. We assume that the reader is familiar with the formal notions of zero-knowledge, hence we skip their explicit definitions. However, even without explicit definitions we do not lack rigor.

Acknowledgements. The first author likes to thank Ariel Gabizon for pointing out a misinterpretation of the role of the masking polynomial in [Hab22]. This motivated us to take a closer look into the technical subtleties of zero-knowledge.

2 BSCR randomization and extension fields

Formally, the [BSCR$^+$19] construction turns any Reed-Solomon encoded interactive oracle proof, which makes use of the FRI low-degree test, into a (straight-line simulable and perfect) zero-knowledge proof (under the honest-verifier assumption, which is sufficient for final non-interactive proof). In simple words, it describes the randomization steps needed to turn FRI into a perfectly hiding polynomial commitment scheme. In the context of the composite protocol, the queried polynomials are randomized outside the trace domain, but one needs to carefully take into account the oracle queries beyond the main protocol, i.e. the queries within FRI. For simplicity we review the [BSCR$^+$19] construction in the context of the following toy protocol, Protocol 1, a minimal environment with all the constraint enforcement logic of a STARK removed. Yet, it reflects the situation of a small-field STARK, in which the verifier queries are sampled from an extension field.

Here and in the sequel, $\mathbb{F}_p$ is a prime field, and $H \subset \mathbb{F}_p$ denotes a multiplicative subgroup of smooth order $|H| = 2^n$ (the trace domain), $n \geq 0$. The set $D \subset \mathbb{F}_p$, a disjoint union of cosets of H, is the evaluation domain for the committed words from the Reed-Solomon code words. We assume that

$$D \cap H = \varnothing.$$

(This is default in many FRI implementations, even when zero-knowledge is not targeted.)

Protocol 1 (Toy IOP using FRI)**.** *Let F be an extension field of $\mathbb{F}_p$, of degree $e = [F : \mathbb{F}_p]$, and take $h \leq |H|$ satisfying*

$$e \cdot n_F + n_D \leq h, \tag{1}$$

where $n_F, n_D \geq 1$ are the number of DEEP queries and FRI queries, respectively, as used below. The oracles are the values of the respective polynomials over D.

1. *Given polynomials $w_1(X), \ldots, w_M(X) \in \mathbb{F}_p[X]^{<|H|}$, the prover samples $r_i(X) \leftarrow\!\!\$\ \mathbb{F}_p[X]^{<h}$, $i = 1, \ldots, M$, and sends the oracles for*

$$\hat{w}_i(X) = w_i(X) + v_H(X) \cdot r_i(X),$$

   $i = 1, \ldots, M$, to the verifier.

2. *The verifier samples $n_F \geq 1$ queries $z_1, \ldots z_{n_F} \leftarrow\!\!\$\ F \setminus (D \cup H)$ and sends them to prover, which responds with the values*

$$\vec{v}_{i,F} = \hat{w}_i(z)|_{z \in Q_F} \in F^{Q_F}$$

   for every $i = 1, \ldots, M$, where $Q_F = \{z_1, \ldots, z_{n_F}\}$.

Then, both prover run Protocol 2 with n_D query rounds on the DEEP quotients of $\hat{w}_1(X), \ldots, \hat{w}_M(X)$ with respect to the claims $\vec{v}_{i,F}$.

We adapt FRI to handle polynomials of non-twoadic degree $|H| + h$ as efficient as possible, assuming that

$$h \ll |H|,$$

which is the case in most applications. That is, we do not choose a different evaluation domain as in the non-zk setting, and prove the non-twoadic degree bound by a simple decomposition into polynomials of twoadic degree. Without going into details, we note that this utilization of FRI does not have any negative impact to soundness.

Protocol 2 (FRI batch evaluation proof with zk [BSCR$^+$19])**.** *Given oracles $\hat{w}_1(X), \ldots, \hat{w}_M(X) \in \mathbb{F}_p[X]^{<|H|+h}$ and evaluation claims $\vec{v}_{i,F} = \hat{w}_i(z)|_{z \in Q_F}$ over a set $Q_F = \{z_1, \ldots, z_{|Q_F|}\} \subset F \setminus (D \cup H)$ of size $|Q_F| \leq n_F$.*

1. *The prover samples a mask polynomial $R(X) \leftarrow\!\!\$\ F[X]^{<|H|+h-1}$ and sends its oracle to the verifier, which responds with a batching randomness $\lambda \leftarrow\!\!\$\ F$.*

2. *The prover provides* $h_0(X), h_1(X) \in F[X]^{<|H|}$ *subject to*

$$\begin{aligned} h(X) &= h_0(X) + X^{h-1} \cdot h_1(X) \\ &= R(X) + \sum_{i=1}^{M} \lambda^{|Q_F|\cdot(i-1)} \cdot \sum_{j=1}^{|Q_F|} \lambda^j \cdot \frac{\hat{w}_i(X) - v_i(z_j)}{X - z_j}, \end{aligned} \tag{2}$$

and both prover and verifier run FRI on $h_0(X), h_1(X)$ *for* $\mathsf{RS}[F, D, |H|]$ *with* n_D *query rounds.*

Remark 1. The decomposition into $h_0(X)$ and $h_1(X)$ is not unique, due to the overlap of monomial powers in $F[X]^{<|H|}$ and $X^{h-1} \cdot F[X]^{<|H|}$. In practice, the prover chooses $h_0(X)$ corresponding to $1, X, \ldots, X^{|H|-1}$ in $h(X)$, and $h_1(X)$ corresponding to the remaining powers $X^{|H|}, \ldots, X^{|H|+h-1}$, but any other way is admissible.

Let us discuss why the mask polynomial $R(X)$ is crucial for zero-knowledge. Each round of FRI not only folds the witness polynomials $w_i(X)$, but also their randomizer polynomials $r_i(X) \cdot v_H(X)$, halving security against evaluations with each step. To see this, let us assume for simplicity that $h = |H|$, and we therefore run FRI on the twoadic degree bound $2 \cdot |H|$. (This is a common choice in many implementations. The argument however generalizes to arbitrary choices of h, and the decomposition as in Protocol 2.) Since $v_H(X)$ is an even function,

$$\begin{aligned} v_H(X) \cdot r_i(X) &= v_{H^2}(X^2) \cdot \big(r_{i,0}(X^2) + X \cdot r_{i,1}(X^2)\big) \\ &= v_{H^2}(X^2) \cdot r_{i,0}(X^2) + X \cdot v_{H^2}(X^2) \cdot r_{i,1}(X^2), \end{aligned}$$

with $r_i(X) = r_{i,0}(X^2) + X \cdot r_{i,1}(X^2)$ being the decomposition of $r_i(X)$. Thus during a FRI folding step, the space $v_H(X) \cdot F[X]^{<|H|}$ is folded into $v_{H^2}(X) \cdot F[X]^{<|H|/2}$ of the half dimension. Applying the same argument to the other folding steps, we see that the foldings of the randomizer polynomials are within the chain of subspaces

$$v_H(X) \cdot F[X]^{<|H|} \longrightarrow v_{H^2}(X) \cdot F[X]^{<|H|/2} \longrightarrow \ldots \longrightarrow v_{H^{2^r}}(X) \cdot F[X]^{<|H|/2^r},$$

whereas the foldings of the witness polynomials $w_i(X)$ are within

$$F[X]^{<|H|} \longrightarrow F[X]^{<|H|/2} \longrightarrow \ldots \longrightarrow F[X]^{<|H|/2^r}.$$

By the size of the folded domain, $|D^{2^r}| = |D|/2^r > 2 \cdot |H|/2^r$, we are not able to reveal the folded oracle in the last step, without telling the folding of the witness polynomials itself. [3]

[3] For this argument, the group structure of the randomization domain is essential. The same entropy-loss does not occur with on-domain randomization, which uses intervals of consecutive points for the randomness. This suggests that in this case, the openings of the FRI foldings could be secured even without mask polynomial $R(X)$. However, we did not succeed in proving this approach zero-knowledge.

We split the proof of zero-knowledge for Protocol 1 into two lemmas, Lemma 1 and Lemma 2 below. The first lemma one confirms the intuition behind the bound in Equality (1), namely that each extension field query amounts to e base field queries, where e is the degree of the extension.

Lemma 1 (Extension field evaluation)**.** *Fix the query sets $Q_F \subset F \setminus H$ and $Q_D \subseteq D$ of size $|Q_F| \leq n_F$ and $|Q_D| \leq n_D$. Then the joint distribution of $\vec{v}_i = \hat{w}_i(z)|_{z \in Q_F \cup Q_D}$, where $i = 1, \ldots, M$, is independent from the (non-randomized) witness polynomials $w_1(X), \ldots, w_M(X)$.*

Proof. We assume that the reader is familiar with basic Galois theory of finite fields. To study the distribution of the evaluations over the query set $Q = Q_F \cup Q_D$, we consider its closure under the action of the Galois group $\mathsf{Gal}(F/\mathbb{F}_p)$ of the extension F, that is the set

$$\bar{Q} = \bigcup_{\phi \in \mathsf{Gal}(F/\mathbb{F}_p)} \phi^k(Q_F) \cup Q_D.$$

The set $\bar{Q}$ is invariant under $\mathsf{Gal}(F/\mathbb{F}_p)$, and so is its vanishing polynomial $v(X) = \prod_{z \in \bar{Q}}(X - z)$, showing that $v(X)$ is actually from $\mathbb{F}_p[X]$. Since $|\mathsf{Gal}(F/\mathbb{F}_p)| = e$, we have

$$\deg v(X) = |\bar{Q}| \leq e \cdot n_F + n_D.$$

The evaluation mapping $E : \mathbb{F}_p[X] \longrightarrow F^{\bar{Q}}$, which sends a polynomial $p(X)$ over $\mathbb{F}_p$ to $p(\gamma)|_{\gamma \in \bar{Q}}$, is linear and has the kernel

$$\ker(E) = (v(X)) = v(X) \cdot \mathbb{F}_p[X].$$

The range of E is isomorphic to $\mathbb{F}_p[X]/(v(X))$, which is a $|\bar{Q}|$-dimensional $\mathbb{F}_p$-linear subspace of $F^{\bar{Q}}$.

Let us now investigate the image of the randomizer space $v_H(X) \cdot \mathbb{F}_p[X]^{<h}$ under E. Since $\bar{Q}$ is disjoint to H, the vanishing polynomials $v_H(X)$ and $v(X)$ are coprime, and

$$\ker(E) \cap v_H(X) \cdot \mathbb{F}_p[X]^{<h} = v_H(X) \cdot v(X) \cdot \mathbb{F}_p[X]^{<h-|\bar{Q}|}.$$

(In the edge case $h = |\bar{Q}|$, the set $\mathbb{F}_p[X]^{<h-|\bar{Q}|}$ is the empty set.) Again, the dimension of $E\left(v_H(X) \cdot \mathbb{F}_p[X]^{<h}\right)$ is equal to $|\bar{Q}|$, and we conclude the equality of the spaces

$$E(v_H(X) \cdot \mathbb{F}_p[X]^{<h}) = E(\mathbb{F}_p[X]^{<|H|+h}) = E(\mathbb{F}_p[X]).$$

From this and the linearity of E, it follows that by drawing $r_i(X)$ independently and uniformly from $\mathbb{F}_p[X]^{<h}$ the values of $\hat{w}_i(X) = w_i(X) + v_H(X) \cdot r_i(X)$ over $\bar{Q}$, $i = 1, \ldots, M$, are uniformly distributed over $E(\mathbb{F}_p[X])^M$, regardless of the concrete choice of witness polynomials $w_i(X)$. Restricting $F^{\bar{Q}}$ to F^Q yields the claim of the lemma. □

Remark 2. The proof of the lemma shows that the distribution of the queried values is uniform over the range of the evaluation map $E : \mathbb{F}_p[X] \longrightarrow F^Q$ over the set $Q = Q_F \cup Q_D$, and it can be efficiently simulated for example by sampling uniformly from $\mathbb{F}_p[X]^{<|H|+h}$ and applying E. In the terminology of [BSCR$^+$19] the lemma states that Protocol 1 is perfect honest-verifier zero-knowledge *against query bound* n_D, meaning that it is zero-knowledge even under further (at most) n_D queries beyond the protocol execution.

The second lemma demonstrates that the mask polynomial $R(X)$ in Protocol 2 acts as a perfect isolator in the information theoretic sense. It statistically separates the transcript following the commitment of $R(X)$ from everything that has happened before.

Lemma 2 (Decoupling Lemma)**.** *Fix Q_F and Q_D as in Lemma 1, and fix $\lambda \in F$ in Protocol 2. Given the values $\hat{w}_i(z)|_{z\in Q_F\cup Q_D}$, $i = 1, \ldots, M$, the joint distribution of $R(z)|_{z\in Q_D}$ and the batch polynomial $h(X)$ is independent of $\hat{w}_1(X), \ldots, \hat{w}_M(X)$.*

Remark 3. With $h(X)$ being independent of $\hat{w}_1(X), \ldots, \hat{w}_M(X)$, so is the distribution of the component polynomials $h_0(X)$ and $h_1(X)$, as well as the entire further transcript of FRI is also independent.

Proof. Given the values $\vec{v}_i = \hat{w}_i(z)|_{z\in Q_F\cup Q_D}$, $i = 1, \ldots, M$, and $\vec{r} = R(z)|_{z\in Q_D}$, we claim that the batch polynomial $h(X)$ is uniformly distributed over the affine subspace

$$L_{\vec{v}_1,\ldots,\vec{v}_M,\vec{r}} = \left\{ h(X) \in F[X]^{<|H|+h-1} \,:\, h(X) \text{ satisfies (2) at all } z \in Q_D \right\}.$$

Take any subset $Q' \subset D$ disjoint to $Q = Q_D$ and so that $|Q \cup Q'| = |H| + h - 1$. Since we draw $R(X)$ uniformly from $F[X]^{<|H|+h-1}$, the distribution of $\vec{r}' = R(z)|_{z\in Q'}$ conditional to $\vec{r} = R(z)|_{z\in Q}$ is uniform over $F^{Q'}$, and so are the values of $h(X)$ over Q', independent of the polynomials $\hat{w}_1(X), \ldots, \hat{w}_M(X)$. Since the evaluation map $E : L_{\vec{v}_1,\ldots,\vec{v}_M,\vec{r}} \longrightarrow F^{Q'}$, $h(X) \mapsto h(z)|_{z\in Q'}$ is bijective, we obtain uniform distribution of $h(X)$ over $L_{\vec{v}_1,\ldots,\vec{v}_M,\vec{r}}$.

The claim of the lemma now follows from that the distribution of $\vec{r} = R(z)|_{z\in Q_D}$ is independent of the concrete choice of $\hat{w}_1(X), \ldots, \hat{w}_M(X)$. □

With the two lemmas in place, zero-knowledge of the toy protocol is an immediate consequence.

Theorem 4. *The IOP from Protocol 1 is perfect honest-verifier zero-knowledge.*

Proof. Although the statement of the theorem is essentially covered by the preceding discussion, let us explicitly describe the simulator. It first samples the

query points $z_1, \ldots, z_{n_F} \leftarrow\$ F \setminus (D \cup H)$, $x_1, \ldots, x_{n_D} \leftarrow\$ D$ uniformly from the respective sets, and draws $\hat{w}_i(X) \leftarrow\$ \mathbb{F}_p[X]^{<|H|+h}$ uniformly at random, so that their values over $Q = Q_F \cup Q_D$ (comprised of the previously sampled points) are distributed as in an honest prover-verifier interaction (cf. Lemma 1 and Remark 2). With their oracles in place, the simulator runs Protocol 2, except that it uses the pre-sampled $x_1, \ldots, x_{n_D}$ from above for the query phase. By Lemma 2 together with Remark 3, the distribution of the transcript is identical to that of an honest prover-verifier interaction. □

3 Zero-knowledge for a simple AIR

Let us extend our discussion to a more meaningful example, a STARK for a simple, yet representative algebraic intermediate representation.

Consider a trace composed of $M > 0$ witness columns w_i, $i = 1, \ldots, M$, where $w_i : H \longrightarrow \mathbb{F}_p$ and H is, again, the trace domain i.e., a smooth multiplicative subgroup with generator g. We can view this trace as sequence of rows $(w_1(x), \cdots, w_M(x))$ for $x \in H$. An *algebraic intermediate representation* (AIR) [BSBHR18, BSGKS20] is a collection of algebraic constraints of the form

$$P_i\left(s_i(x), w_1(x), \cdots, w_M(x), w_1(g \cdot x), \cdots, w_M(g \cdot x)\right) = 0,$$

for all $x \in H$. Here, $s_i(x)$ is the selector polynomial of the enforcement domain H_i of P_i, i.e. a coset of a subgroup of H, and

$$P_1, \cdots, P_C \in \mathbb{F}_p[X, X_1, \cdots X_M, Y_1, \cdots Y_M],$$

where the degree in the selector variable is $\deg_X P_i \leq 1$. The degree of the AIR is the maximum total degree of its constraints,

$$d_{\text{AIR}} = \max_i \deg P_i.$$

Note that we use a simplified notation of an AIR, working with constraints between neighbouring rows only. For notational convenience we will rather work with the *reduced degree* $d := d_{\text{AIR}} - 1$ (which in many cases matches the "actual" degree of the constraint system, with selectors not taken into account).

In terms of the low-degree extensions $w_1(X), \ldots, w_M(X) \in \mathbb{F}_p[X]^{<|H|}$ of the trace columns, satisfiability of the AIR constraints over H is then equivalent to that, with noticable probability, any random linear combination of the constraints is divisible by the vanishing polynomial $Z_H(X) = X^{|H|} - 1$ of H. This yields the *overall identity*

$$\sum_{i=1}^{C} \lambda^i \cdot P_i\left(s_i(X), w_1(X), \cdots, w_M(X), w_1(g \cdot X), \cdots, w_M(g \cdot X)\right) \\ = q(X) \cdot Z_H(X),$$

where λ is drawn from the extension field F, and $q(X) \in F[X]$ is a low-degree polynomial of degree

$$\deg q(X) = d_{\text{AIR}} \cdot (|H| - 1) - |H| < d \cdot |H|,$$

except the pointless use case $d_{\text{AIR}} = 0$.

In our interactive oracle proof, Protocol 3, the prover decomposes[4]the overall quotient into polynomials $q_1(X), \ldots, q_d(X)$ of smaller degree, using the FFT-type decomposition

$$q(X) = q_1(X^d) + X \cdot q_2(X^d) + \ldots + X^{d-1} \cdot q_d(X^d),$$

where d is the reduced degree of the AIR, as above. The verifier gets oracle access to $q_1(X), \ldots, q_d(X)$, and the overall identity is then tested at one (or more) random points z_i, $i = 1, \ldots, n_F$, from the extension field F. The evaluation claims for the polynomials are then proven as in our toy protocol from Section 2, by showing the single-point quotients at z_i and $g \cdot z_i$ are low-degree, using FRI over a sufficiently large evaluation domain D.

For zero-knowledge, the prover randomizes the witness polynomials outside the trace domain H,

$$\hat{w}_i(X) := w_i(X) + v_H(X) \cdot r_i(X) \in \mathbb{F}_p[X]^{<|H|+h},$$

with $r_i(X) \leftarrow\!\$\ \mathbb{F}_p[X]^{<h}$, $i = 1, \ldots, M$, where the degree of freedom h is chosen so that

$$2 \cdot d \cdot (e \cdot n_F + n_D) + n_D \leq h \leq |H|, \tag{3}$$

where $e = [F : \mathbb{F}_p]$ is the degree of the extension F, $n_F \geq 1$ is the number of out-of-domain queries (without counting their translates by g), and $n_D \geq 1$ is the number of FRI query rounds. We discuss the rationale behind that bound below. As before, the evaluation domain $D \subset \mathbb{F}_p$ for the Reed-Solomon code words is a disjoint union of cosets of H, and $D \cap H = \varnothing$.

Protocol 3 (IOP for AIR using DEEP-ALI)**.** *Let $\mathbb{F}_p$, H, D, and F as above and let $\hat{w}_1(X), \cdots \hat{w}_M(X) \in \mathbb{F}_p[X]^{<|H|+h}$ be the randomized witness polynomials, satisfying the AIR constraints specified by s_i and P_i, $i = 1, \ldots, C$ over H. The verifier is given oracle access to $[\hat{w}_i]$, i.e. the values $\hat{w}_i(X)$ over D, for each $i = 1, \ldots, M$.*

1. The verifier challenges the prover with a random value $\lambda \leftarrow\!\$\ F$, for which

[4]In general, one computes the values of $q(X)$ over a sufficiently large FFT domain, and derives the component polynomials from the monomial representation of $q(X)$. If $\mathbb{F}_q$ contains all $(d \cdot |D|)$-th roots of unity, then the decomposition can be computed via the values of $q_1(X), \ldots, q_d(X)$ over D^d as in a FRI reduction step, with a subsequent extrapolation back to D.

the prover computes $q(X) \in F[X]^{<d\cdot|H|+(d+1)\cdot h}$ *such that*

$$\sum_{i=1}^{C} \lambda^{i-1} \cdot P_i\left(s_i(X), \hat{w}_1(X), \cdots, \hat{w}_M(X), \hat{w}_1(g\cdot X), \cdots, \hat{w}_M(g\cdot X)\right) = Z_H(X)\cdot q(X). \quad (4)$$

It splits it into the unique polynomials $q_j(X) \in F[X]^{<\hat{k}}$, $j = 1, \ldots, d$, *with* $\hat{k} = |H| + \left\lceil \frac{d+1}{d} \cdot h \right\rceil$, *subject to*

$$q(X) = \sum_{j=1}^{d} X^{(j-1)} \cdot q_j(X^d), \quad (5)$$

and provides the verifier oracle access to their values over D.

2. *The verifier sends the prover random DEEP queries* $z_j \leftarrow\!\!\$\ F \setminus \left(\bar{D} \cup H\right)$, $j = 1, \ldots, n_F$, *where* $\bar{D} := \{y \in F : y^d \in D\}$, *on which the prover responds with evaluation claims*

$$(v_{i,j,1}, v_{i,j,2}) = \left(\hat{w}_i(z_j), \hat{w}_i(g\cdot z_j)\right),$$

for $i = 1, \ldots, M$ *and* $v_i = q_i(z_j^d)$, $i = 1, \ldots, d$, *for each* j.

Both prover and verifier then run batch FRI on the DEEP quotients corresponding to the evaluation claims, i.e. Protocol 2, with h *replaced by* $\lceil (d+1)/d \cdot h \rceil$, *and using* n_D *query rounds. The verifier accepts if Protocol 2 passes and if the evaluation claims satisfy the overall identity* (4) *at each* $X = z_j$, $j = 1, \ldots, n_F$.

Let us explain the intuition behind the degree bound in

$$2 \cdot d \cdot (e \cdot n_F + n_D) + n_D \leq h.$$

First, each query on a polynomial $\hat{w}_i(z)$ reveals either e, or a single field element, depending on whether it is from F, or from D (cf. Lemma 1). Since the value of the quotient $q(X)$ at a point z is uniquely determined from the values $\hat{w}_i(z)$ and $\hat{w}_i(g\cdot z)$ via the overall constraint (4), it would be sufficient to randomize the witness polynomials against

$$2 \cdot (e \cdot n_F + n_D)$$

queries, assuming that the prover would work with the non-split $q(X)$. (The factor 2 is due to tranlates by g.) To take into account the additional information revealed by the component polynomials $q_1(X), \ldots, q_d(X)$, recall that

$$(q_1(z^d), \ldots, q_d(z^d)) = \mathsf{FFT}(q(X)|_{z\cdot U}), \quad (6)$$

where U is the subgroup of the d-th roots of unity. Hence each query of the component polynomials is covered by the values of $q(X)$ over $|U| = d$ times

as many points, which explains the factor d in our bound. (We refer to these points as the implicit queries.) The additional n_D is due to fact, that in general the FRI queries on $\hat{w}_i(X)$ do not overlap with the implicit queries.

Our proof of zero-knowledge is based on the following adaption of Lemma 1 to the current, more complex environment. Again, since the out-of-domain samples are from the extension field F, its proof uses basic Galois theory. However, to make use of Formula (6), the setup of a suitable environment for studying the distribution of the queried values is significantly more complicated. It relies on the splitting field of $v_D(X^d)$, a polynomial which is not guaranteed to fully factorize over $\mathbb{F}_p$ or F. (Here, and in the sequel, $v_D(X)$ is the vanishing polynomial of D.)

Lemma 3 (Quotient components)**.** *Fix $\lambda \in F$ and query sets Q_F and Q_D of size $|Q_F| \leq n_F$ and $|Q_D| \leq n_D$ such that inequality* (3) *holds. Then the joint distribution of*

$$(\hat{w}_1(z), \hat{w}_1(g \cdot z), \ldots, \hat{w}_M(z), \hat{w}_M(g \cdot z), q_1(z^d), \ldots, q_d(z^d))|_{z \in Q_F},$$
$$(\hat{w}_1(z), \ldots, \hat{w}_M(z), q_1(z), \ldots, q_d(z))|_{z \in Q_D},$$

is independent of the (non-randomized) witness polynomials $(w_1(X), \ldots, w_M(X))$.

Proof. The proof is similar to that of Lemma 1, yet with a more careful choice of the evaluation environment, as F might not support all the roots we need for the distribution analysis of the quotient component values. (Note that the domain D might contain elements which do not have d-th roots in F. This may happen even when $\mathbb{F}_q$ contains all $(d \cdot |H|)$-th roots of unity.) For this reason, we take K an extension of F over which $v_D(X^d)$ entirely splits into linear factors (higher multiplicities are possible). Such an extension comes with all the roots necessary for our analysis: It contains (1) the subgroup U of all d-th roots of unity, and moreover (2) all d-th roots of elements in D.

Consider the polynomial

$$p(X) = \prod_{w \in Q_1} (X^d - w^d) \cdot \prod_{x \in Q_2} (X^d - x) \cdot \prod_{x \in Q_3} (X - x),$$

where $Q_1 = \bigcup_{\phi \in \mathsf{Gal}(F/\mathbb{F}_p)} \phi\left(Q_F \cup g \cdot Q_F\right)$, $Q_2 = Q_D \cup g^d \cdot Q_D$, and $Q_3 = Q_D$. The polynomial belongs to $\mathbb{F}_p[X]$, since it is a polynomial from $F[X]$ which is invariant under $\mathsf{Gal}(F/\mathbb{F}_p)$, and it splits over K. We denote by $\bar{Q}$ its set of roots in K. Let us explain its components:

– The roots of the first product cover the cosets $z \cdot U$ for the Galois conjugates of every out-of-domain query, and their translates by g. This is the set of points over which the values of $\hat{w}_1, \ldots, \hat{w}_M$ determine the values of the q_j over the Galois closure of each DEEP query.

– the second product is for all d-th roots of a FRI query, and their g-translates. This is the set of points over which the values of $\hat{w}_1, \ldots, \hat{w}_M$ determine the values of the q_j at each FRI query.

– The third product is simply for the FRI queries themselves.

We note that there might be overlaps between the roots of these products, and for this reason we consider $v(X)$ the vanishing polynomial of $\bar{Q}$. By the invariance of $\bar{Q}$ under $\mathsf{Gal}(F/\mathbb{F}_p)$, the vanishing polynomial is again from $\mathbb{F}_p[X]$, and its degree is

$$\deg v(X) = |\bar{Q}| \leq 2 \cdot d \cdot (e \cdot n_F + n_D) + n_D \leq h.$$

By the assumption on the out-of-domain queries and D, none of the roots from $\bar{Q}$ are contained in H.

The rest of the proof is as in Lemma 1. The kernel of the evaluation map $E : \mathbb{F}_p[X] \longrightarrow K^{\bar{Q}}$ is the ideal generated by $v(X)$ of degree $|\bar{Q}|$, and hence its image is a $\mathbb{F}_p$-linear subspace of $K^{\bar{Q}}$ with $\dim E(\mathbb{F}_p[X]) = |\bar{Q}|$. Likewise, since $v_H(X)$ and $v(X)$ are coprime, the kernel within the randomizer space $v_H(X) \cdot \mathbb{F}_p[X]^{<h}$ is

$$v_H(X) \cdot \mathbb{F}_p[X]^{<h} \cap \ker E = v_H(X) \cdot v(X) \cdot \mathbb{F}_p[X]^{<h-|\bar{Q}|},$$

including the edge case $h - |\bar{Q}| = 0$, in which the intersection is empty. This shows that the image of $v_H(X) \cdot \mathbb{F}_p[X]^{<h}$ under E is of dimension $|\bar{Q}|$, yielding the equality of the spaces

$$E\left(v_H(X) \cdot \mathbb{F}_p[X]^{<h}\right) = E\left(\mathbb{F}_p[X]^{<|H|+h}\right) = E\left(\mathbb{F}_p[X]\right).$$

From the latter equality, and the linearity of E, we conclude that the distribution of M-fold evaluation map E^M, which evaluates each $\hat{w}_1(X), \ldots, \hat{w}_M(X)$ over $\bar{Q}$, is uniform over $E(\mathbb{F}_p[X])^M$ and independent of the non-randomized witness polynomials $w_1(X), \ldots, w_M(X)$.

Since the values of $q_1(X), \ldots, q_d(X)$ at the requested queries are uniquely determined from those of $\hat{w}_1(X), \ldots, \hat{w}_M(X)$ over $\bar{Q}$ via the overall constraint of the AIR, together with Formula (6), their distribution is also independent from the witness polynomials. Finally, restricting our view to the query points of the protocol yields the claim of the lemma. □

Remark 5. The usage of splitting fields in the proof of Lemma 3 raises another interesting aspect, namely the running time for simulating the distribution of the queried values. This can be efficiently done by means of the evaluation map E over $\bar{Q}$, but for that one has to explicitly construct the set of roots $\bar{Q}$. The splitting field K of $v_D(X^d)$ can be generated in probabilistic polynomial time:

1. Construct a cyclotomic extension K' of F, which contains all $(d \cdot |D|)$-th roots of unity. For this, determine the extension degree $n = [K' : F]$ as the smallest n, $1 \leq n \leq d$, so that $d \cdot |D|$ divides $|F|^n - 1$, and sample a monic random polynomial of degree n, and check it on irreducibility using Rabin's test. (By the Moreau necklace counting function, the probability of a random polynomial being irreducible is $> 1 - 1/|F|^{1/2}$.)

2. Build an extension K of K' which contains a d-root of any $a \in D$. For this factorize $X^d - a$ into irreducibles, choose the one of lowest degree for defining K'.

To obtain all roots from $\bar{Q} = Q_1 \cup Q_2 \cup Q_D$, determine the subgroup U of all $(d \cdot |D|)$-th roots of unity in K. (This is done in probabilistic polynomial time.) Compute the complete preimage of D under the d-th power map by multiplying the defining element of K/K' with all elements of U. The requested evaluation set Q_1 for each out-of-domain query is computed by means of U. With the preimage of D, the requested evaluation set Q_2 is easily derived, and we are finished. All of the above steps run in probabilistic polynomial time, with respect to the size of the AIR and the security parameter.

Theorem 6. *The IOP from Protocol 3 is perfect honest-verifier zero-knowledge.*

Proof. The simulator samples the verifier challenges $\lambda \leftarrow\!\$\ F$, $z_F \leftarrow\!\$\ F \backslash \big(\bar{D} \cup H\big)$, where $\bar{D} := \{y \in F : y^d \in D\}$, and $x_1, \ldots, x_{n_D} \leftarrow\!\$\ D$.

Then, it constructs the splitting field of $v_D(X^d)$ and the evaluation set $\bar{Q}$ as in Remark 5. It samples polynomials $\hat{w}_i(X) \leftarrow\!\$\ \mathbb{F}_p[X]^{<|H|+h}$ at random, and computes their values over $\bar{Q}$. From these values it determines the values of the decomposition polynomials at the protocol queries, using the overall constraint (4) and Formula (6), and interpolates them by arbitrary polynomials $q_1(X), \ldots, q_d(X)$ from $F[X]^{<\hat{k}}$, with $\hat{k}$ as above. The resulting distribution is that of an honest prover-verifier interaction as stated in Lemma 3.

The remaining transcript of Protocol 2 on the DEEP evaluation claims for $\hat{w}_1(X), \ldots, \hat{w}_M(X)$ and $q_1(X), \ldots q_d(X)$ is produced as in the proof of Theorem 4: The simulator runs the commit phase of the protocol, but uses the $x_1, \ldots, x_{n_D}$ sampled beforehand in the query phase of FRI. Again, by Lemma 2 and Remark 3, the resulting transcript has the same distribution as that of an honest prover-verifier interaction. □

4 Other types of quotient decompositions

In this section, we briefly sketch how to tackle other commonly used decompositions for the overall quotient of a STARK, namely, the canonical decomposition based on monomials, and the decomposition by value, based on the Lagrange basis of the evaluation domain.

4.1 Canonical decomposition

Let $q(X) \in F[X]^{<d\cdot|H|+(d+1)\cdot h}$ be the quotient polynomial as it appears in Equation 4 in Protocol 3 with h, the degree of freedom of the witness randomizer, to be specified later. By the *canonical decomposition*, we mean

$$q(X) = \sum_{i=1}^{d} X^{\hat{k}\cdot(i-1)} \cdot q_i(X), \tag{7}$$

where each $q_i(X) \in F[X]^{<\hat{k}}$ with $\hat{k} = |H| + \lceil (d+1)/d \cdot h \rceil$. Contrary to the FFT decomposition, the decomposition in (7) can be randomized, using a technique from [GWC19]: In order to maintain the identity

$$q(X) = \sum_{i=1}^{d} X^{\hat{k}\cdot(i-1)} \cdot \hat{q}_i(X), \tag{8}$$

on draws $t_i(X) \leftarrow\!\!\$\ F[X]^{<h_p}$, $i = 1, \ldots, d-1$, independently and according to the uniform distribution, and sets

$$\begin{aligned}
\hat{q}_1(X) &= q_1(X) + X^{\hat{k}} \cdot t_1(X), \\
\hat{q}_2(X) &= q_2(X) + X^{\hat{k}} \cdot t_2(X) - t_1(X), \\
&\vdots \\
\hat{q}_{d-1}(X) &= q_{d-1}(X) + X^{\hat{k}} \cdot t_{d-1}(X) - t_{d-2}(X),
\end{aligned}$$

and eventually

$$\hat{q}_d(X) = q_d(X) - t_{d-1}(X).$$

The degree of freedom h_p is chosen such that

$$n_F + n_D \leq h_p. \tag{9}$$

With the above modification, Protocol 3 goes through unchanged except for Equation (3) which becomes

$$2 \cdot (e \cdot n_F + n_D) \leq h \leq |H|, \tag{10}$$

and the common degree bound for the batch opening proof, which is adapted accordingly.

Lemma 4. *Fix $\lambda \in F$ and query sets Q_F and Q_D of size $|Q_F| \leq n_F$ and $|Q_D| \leq n_D$ such that Equations* (9) *and* (10) *hold. Then the joint distribution of*

$$(\hat{w}_1(z), \hat{w}_1(g \cdot z), \ldots, \hat{w}_M(z), \hat{w}_M(g \cdot z), \hat{q}_1(z), \ldots, \hat{q}_d(z))|_{z \in Q_F},$$
$$(\hat{w}_1(z), \ldots, \hat{w}_M(z), \hat{q}_1(z), \ldots, \hat{q}_d(z))|_{z \in Q_D},$$

is independent of the witness polynomials $(w_1(X), \ldots, w_M(X))$.

Proof sketch. Using the same approach as in the proof of Lemma 1 we have that the evaluations of the randomized witness polynomials at the queried points is independent of the witness polynomials. Now, since the randomizer polynomials $t_i(X)$, $i = 1, \ldots, d-1$, are independently and uniformly drawn from $F[X]^{<h_p}$, the queried values of $\hat{q}_i(X)$, $i = 1, \ldots, d-1$, are uniformly distributed, independent of the witness polynomials. Since the values of the last component polynomial $\hat{q}_d(X)$ are fully determined from the values of the randomized witness polynomials in addition to those of $\hat{q}_i(X)$, $i = 1, \ldots, d-1$, we get the claim. □

Theorem 7. *The IOP from Protocol 3, with quotient decomposition 8 instead of 5, is perfect honest-verifier zero-knowledge.*

Proof. The simulator samples the verifier challenges $\lambda \leftarrow\!\$\ F$, $z_1, \ldots, z_{n_F} \leftarrow\!\$\ F \backslash (D \cup H)$, and $x_1, \ldots, x_{n_D} \leftarrow\!\$\ D$. It then samples $\hat{w}_i(X) \leftarrow\!\$\ \mathbb{F}_p[X]^{<|H|+\hat{h}}$ and $\hat{q}_i(X) \leftarrow\!\$\ F[X]^{<\hat{k}+h_p}$, $i = 1, \ldots, d-1$, uniformly at random, and computes their values over $Q = Q_F \cup Q_D$. Given these values, it takes any $\hat{q}_d(X) \in F[X]^{<|H|+h_p}$ satisfying the overall constraint (4) and the decomposition identity (8) at the points from Q. Inspecting the proof of Lemma 4, we see that the distribution of the transcript is identical to that of an honest prover-verifier interaction.

The remaining transcript for the batch opening proof Protocol 2 is simulated as previously. □

4.2 Lagrange decomposition

The *Lagrange decomposition* of a polynomial $q(X) \in F[X]^{<d \cdot |H|}$ over $\bar{H} = \bigcup_{i=1}^{d} H_i$ a union of disjoint cosets of H, is the unique decomposition

$$q(X) = \sum_{i=1}^{d} L_{H_i}(X) \cdot q_i(X), \tag{11}$$

where

$$L_{H_i}(X) = c_i \cdot \prod_{j \neq i} v_{H_j}(X) \tag{12}$$

is the selector polynomial of the coset H_i (normalized so that $L_{H_i}(x) = 1$ over H_i), and every $q_i(X) \in F[X]^{<|H|}$. Each $q_i(X)$ is directly obtained from the values of $q(X)$ over H_i via an FFT of witness domain size $|H|$. Besides that the Lagrange decomposition is more efficient than the other two discussed in this document, we stress the fact that its main advantage is its flexibility: It does not demand the evaluation domain $\bar{H}$ to be (the coset of) a group again, and hence supports non-two-adic blow-up factors, which is useful in certain applications.

In the zero-knowledge setting, the quotient $q(X) \in F[X]^{<d\cdot|H|+(d+1)\cdot h}$ can be still decomposed as in (11), allowing the last component polynomial

$$q_d(X) \in F[X]^{<|H|+(d+1)\cdot h}.$$

Demanding $q_i(X) \in F[X]^{<|H|}$ for $i = 1, \ldots, d-1$, the decomposition is still unique, and can be randomized as follows: For $i = 1, \ldots, d-1$, the prover takes

$$\hat{q}_i(X) = q_i(X) + v_{H_i}(X) \cdot t_i(X) \in F[X]^{<|H|+h_p}, \tag{13}$$

with $t_i(X) \leftarrow\!\!\$\ F[X]^{<h_p}$ and h_p as specified below, and

$$\hat{q}_d(X) = q_d(X) - v_{H_d}(X) \cdot \sum_{k=1}^{d-1} c_d^{-1} \cdot c_i \cdot t_i(X), \tag{14}$$

with the normalizing coefficients c_i and c_d from (12). This choice still satisfies

$$q(X) = \sum_{i=1}^{d} L_{H_i}(X) \cdot \hat{q}_i(X). \tag{15}$$

The degrees of freedom h and h_d are as for the canonical decomposition, with

$$n_F + n_D \leq h_p, \tag{16}$$

and

$$2 \cdot (e \cdot n_F + n_D) \leq h \leq |H|. \tag{17}$$

The statement of Lemma 4 again holds, with its proof carried over verbatim.

Lemma 5. *Fix $\lambda \in F$ and query sets Q_F and Q_D of size $|Q_F| \leq n_F$ and $|Q_D| \leq n_D$ such that Equations* (16) *and* (17) *hold. Then the joint distribution of*

$$(\hat{w}_1(z), \hat{w}_1(g \cdot z), \ldots, \hat{w}_M(z), \hat{w}_M(g \cdot z), \hat{q}_1(z), \ldots, \hat{q}_d(z))|_{z \in Q_F},$$
$$(\hat{w}_1(z), \ldots, \hat{w}_M(z), \hat{q}_1(z), \ldots, \hat{q}_d(z))|_{z \in Q_D},$$

is independent of the witness polynomials $(w_1(X), \ldots, w_M(X))$.

The distribution in the lemma can be efficiently simulated, yielding zero-knowledge in the honest-verifier setting.

Theorem 8. *The IOP from Protocol 3, with quotient decomposition 15 instead of 5, is perfect honest-verifier zero-knowledge.*

Proof. Almost verbatim to that of Theorem 7, with only a slight change of degree bounds: The simulator samples $\hat{w}_i(X) \leftarrow\!\!\$\ \mathbb{F}_p[X]^{<|H|+h}$ and $\hat{q}_i(X) \leftarrow\!\!\$\ F[X]^{<|H|+h_p}$, $i = 1, \ldots, d-1$, and takes any $\hat{q}_d(X) \in F[X]$ of degree less than $|H| + \max\{(d+1) \cdot h, h_p\}$ satisfying the overall constraint (4) and the decomposition identity (8) at the points of the query set Q. □

5 Practical considerations

Let us quickly review the computational overhead when adding zero-knowledge with the most greedy parameters, where the randomization degree is taken as small as possible. We again restrict to the FFT decomposition from Section 3; the decompositions from Section 4 are treated similarly.

Let h taken as the smallest possible value in (3). We assume that $h \ll |H|$, which is met under moderate FRI parameters and not too short traces.

The size of the evaluation domain D can be kept, at the cost of only a single additional commitment in the batching step of FRI. (Recall that using $\mathsf{RS}[F, D, \hat{k}]$ with $\hat{k} = |H| + \lceil (d+1)/d \cdot h \rceil$ in the batching step of FRI does not have significant impact on the overall sampling parameter n_D.) Therefore, in the wide trace regime with many column polynomials to be commitment, the hashing costs remain essentially the same.

For evaluating the FFT costs, we assume that the proof system, before introducing zero-knowledge, is optimized for constraint evaluation, by choosing

$$d \cdot |H| = |D|.$$

Using oversized FFTs to handle the slight increase in the degree of the polynomials, the overall cost per witness column increases from $(B+1) \cdot \mathsf{FFT}(|H|)$ to

$$\mathsf{FFT}(|H|) + \frac{B}{2} \cdot \mathsf{FFT}(2 \cdot |H|) + \mathsf{Eval}_{|H|+h}((d+1) \cdot h) + O(h),$$

where $B = |D|/|H| \geq 2$ is the blowup factor, $\mathsf{Eval}_{|H|+h}((d+1) \cdot h)$ denotes the cost for evaluating a polynomial from $\mathbb{F}_p[X]^{<|H|+h}$ over a set of size $(d+1) \cdot h$, and the $O(h)$ term covers the cost of adding $r_i(X) \cdot (X^{|H|} - 1)$ to the Fourier transform $w_i(X)$ of the witness column. (There are more efficient ways to compute the values of the randomized polynomial. We shall not dwell on such, since the impact to our estimates is negligible.) Under the simplifying assumption that the target set of $\mathsf{Eval}_{|H|+h}((d+1) \cdot h)$ is a coset of a subgroup of size $(d+1) \cdot h$,

the evaluation cost amounts to $|H| + h$ multiplications and additions for the reduction modulo the coset vanishing polynomial, plus an FFT of size $h \cdot (d+1)$. We approximate the former cost by two layers of an $\mathsf{FFT}(|H|)$, yielding

$$\begin{aligned}\mathsf{FFT}(|H|) + \frac{B}{2} \cdot \mathsf{FFT}(2 \cdot |H|) + \frac{2}{\log |H|} \cdot \mathsf{FFT}(|H|) + O(h \cdot \log h) \\ = \left(B + 1 + \frac{B+2}{\log |H|}\right) \cdot \mathsf{FFT}(|H|) + O(h \cdot \log h),\end{aligned}$$

where in the latter we have used the rough rule of thumb, that $\mathsf{FFT}(2 \cdot |H|) = \left(1 + \frac{1}{\log |H|}\right) \cdot 2 \cdot \mathsf{FFT}(|H|)$. Hence the overall increase of the arithmetic cost per witness column is expected to be roughly

$$C_{\text{zk}}/C_{\text{non-zk}} \approx \left(1 + \frac{4}{3 \cdot \log |H|}\right),$$

neglecting the $O(h \cdot \log h)$ term. In configurations where $d \cdot |H| < |D|$, $\mathsf{Eval}(|H| + h, h)$ can be dropped and we obtain the estimate

$$C_{\text{zk}}/C_{\text{non-zk}} \approx \left(1 + \frac{1}{\log |H|}\right),$$

neglecting an $O(h)$ term. In wide trace AIRs, these two ratios are expected to be an upper bound for the overhead.

References

[BBHR18] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Fast Reed-Solomon interactive oracle proofs of proximity. In *ICALP 2018*, 2018. Full paper: `https://eccc.weizmann.ac.il/report/2017/134/`.

[BLH+] Sean Bowe, Ying Tong Lai, Daira Hopwood, Jack Grigg, and Steven Smith. The halo2 book. `https://zcash.github.io/halo2`.

[BSBHR18] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Scalable, transparent, and post-quantum secure computational integrity. In *IACR ePrint Archive 2018/046*, 2018. `https://eprint.iacr.org/2018/046`.

[BSCR+19] Eli Ben-Sasson, Alessandro Chiesa, Michael Riabzev, Nicholas Spooner, Madars Virza, and Nicholas P. Ward. Aurora: Transparent succinct arguments for R1CS. In Y. Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019*, volume 11476 of *LNCS*. Springer, 2019.

[BSGKS20] Eli Ben-Sasson, Lior Goldberg, Swastik Kopparty, and Shubhangi Saraf. DEEP-FRI: Sampling outside the box improves soundness. In *ITCS 2020*, 2020. Full paper: `https://eprint.iacr.org/2019/336`.

[GWC19] Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. PLONK: Permutations over Lagrange-bases for oecumenical non-interactive arguments of knowledge. In *IACR ePrint Archive 2019/953*, 2019. `https://eprint.iacr.org/2019/953`.

[Hab22] Ulrich Haböck. A summary on the FRI low-degree test. In *IACR ePrint Archive 2022/1216*, 2022. `https://eprint.iacr.org/2022/1216`.

[HLP24] Ulrich Haböck, David Levit, and Shahar Papini. Circle STARKs. In *IACR preprint archive*, 2024. `https://eprint.iacr.org/2024/278`.

[Plo] Plonky2 codebase. `https://github.com/0xPolygonZero/plonky2`.

[Pol] Polygon Labs / Polygon Zero. Plonky2: Fast recursive arguments with PLONK and FRI. `https://github.com/mir-protocol/plonky2/blob/main/plonky2/plonky2.pdf`.

[Ris] RISC Zero: a zero-knowledge verifiable general computing platform based on zk-starks and the risc-v microarchitecture. `https://github.com/risc0/risc0`.

[Tri] Triton VM. `https://github.com/TritonVM/triton-vm`.

A Permutation arguments

Permutation arguments are a special case in the treatment of zero-knowledge. Unlike other components of a STARK, the used protocols are not complete, and a single successful run reveals information on the witness data. In most cases, the amount of revealed data is not an obstacle for statistical zero-knowledge. (We found only one exception[5].) However, perfect zero-knowledge requires careful modifications, which override incompleteness while not harming soundness.

[5] In plonky2 [Pol] the prover draws several challenges from a small base field, instead of a single one from a cryptographically large extension. While larger sample sets are a reasonable measure for boosting soundness, this is not the case for statistical zero-knowledge. In fact, the verifier learns even more with each additional sample, resulting in a modest *increase* of the distance to uniform distribution.

For simplicity, we restrict ourselves to the case of two columns $(w_1(x))_{x\in H}$ and $(w_2(x))_{x\in H}$, which we wish to prove a permutation of one another. (The general case of several columns, or a permutation invariance argument as in Plonk, is treated similarly.) The target is to prove the polynomial identity

$$\prod_{x\in H}(Y - w_1(x)) = \prod_{x\in H}(Y - w_2(x)),$$

which is reduced to the equality at a random challenge $\alpha \leftarrow\!\!\$\ F$, paraphrased as

$$\prod_{x\in H}\frac{\alpha - w_1(x)}{\alpha - w_2(x)} = 1. \tag{18}$$

This product is shown by a single "grand product argument" in the style of [GWC19]. The prover commits the running product column $(Z(x))_{x\in H}$, which is initialized with $Z(1) = 1$, and recursively determined at the other points $g, g^2, \ldots, g^{|H|-1}$ of the domain, via

$$Z(g\cdot x) = Z(x)\cdot\frac{\alpha - w_1(x)}{\alpha - w_2(x)},$$

where, since the overall product is equal to one, the incremental equation holds also at the last point of the domain. (As before, g is the generator of the witness domain H.) The constraints imposed on $Z(X)$ are the domain identities

$$(Z(x) - 1)\cdot L_1(x) = 0, \tag{19}$$

$$Z(g\cdot x)\cdot(\alpha - w_2(x)) = Z(x)\cdot(\alpha - w_1(x)), \tag{20}$$

for all $x \in H$, which together overwhelmingly prove the product (18).

However, the prover is not able to succeed whenever the random point α falls into the range

$$R = \{w_1(x) : x \in H\} = \{w_2(x) : x \in H\},$$

as it runs into a quotient $(\alpha - w_1(x))/(\alpha - w_2(x))$ which is undefined (division by zero), or the final running product must be zero, conflicting the transitional constraint 20 at the last point of the domain. Thus with each successful run of the protocol, the verifier learns that α is not any of the witness values – a fact that does not change when moving to randomized polynomials.

One may attempt to modify how the prover crafts a polynomial $Z(X)$ compliant with the constraints (19) and (20), but this still would not meet the target, since the constraint system itself is not satisfiable in all cases. Lets call $x \in H$ a *critical point* if $\alpha \in \{w_1(x), w_2(x)\}$, i.e. either $\alpha - w_1(x) = 0$ or $\alpha - w_2(x) = 0$. We point out two types of critical points:

1. If $\alpha - w_1(x) = 0$, but $\alpha - w_2(x_0) \neq 0$, then the transitional constraint (20) enforces that $Z(g\cdot x) = 0$. The same constraint enforces the forward

propagation of zero, at least up to the next critical point. We therefore call x a *forward propagation point*.

2. If $\alpha - w_2(x) = 0$, but $\alpha - w_1(x) \neq 0$, then the incremental constraint (20) enforces that $Z(x) = 0$, and the same constraint backward propagates the zero, at least down to the next critical point. Thus we call x a *backward propagation point*.

(The third type, when both $\alpha - w_1(x) = 0$ and $\alpha - w_2(x) = 0$, does not enforce anything on $Z(x)$ and $Z(g \cdot x)$.) It now depends on the position of forward and backward propagation points (if present), whether the constraint system is satisfiable.

- If the first critical point on H is a backward propagation point, including the case $x = 1$, then this contradicts the initial boundary condition imposed by (19). Likewise, if the last critical point of the domain is a forward propagation point.
- If both first and last critical points are of the respective other types, then the prover may choose Z throughout zero between these points. On the remaining interval around $x = 1$ the values of Z are uniquely determined from (19) together with (20), and the resulting $Z(X)$ satisfies the constraints over the entire domain H.

To mitigate the propagation of zeros, both forward and backward, we mute the transitional constraint (20) if one of the two, either $\alpha - w_1(x)$ or $\alpha - w_2(x)$, are zero.[6] This leads to the following weakened, but complete system of constraints:

$$\begin{aligned} L_1(x) \cdot (Z(x) - 1) &= 0, \\ (\alpha - w_1(x)) \cdot (\alpha - w_2(x)) \cdot \big(Z(g \cdot x) \cdot (\alpha - w_2(x)) - Z(x) \cdot (\alpha - w_1(x))\big) &= 0, \end{aligned}$$

for all $x \in H$. We note that soundness is essentially not affected. Whenever α is not in the range R, the modified constraints imply the same domain product as then incomplete system. In this case, the "regular case", the prover computes the polynomial $Z(X)$ as before. In the presence of critical points, the prover may craft a satisfying polynomial $Z(X)$ as indicated above in the second case.

In certain applications (e.g. zero-knowledge virtual machines) permutation arguments are combined with lookup arguments into an overall argument based on fractional decompositions. For these combined arguments perfect zero-knowledge is often too costly, and we suggest to keep with statistical zero-knowledge.

[6] An earlier version of the Halo2 book [BLH+] misses to mitigate backward propagation of zeroes. As a consequence their proposed solution is still not complete, thwarting perfect zero-knowledge.