A note on adding zero-knowledge to STARKs

Ulrich Haböck and Al Kindi
{al.kindi, uhaboeck}@polygon.technology

April 2024

Abstract

We discuss zero-knowledge in the context of FRI-based STARKs using techniques desirable in practice: Randomization by polynomials over the basefield, and decomposing the overall quotient into polynomials of smaller degree.

Contents

1 Introduction

Adding zero-knowledge to a STARK[1]is a subject that seems somewhat neglected in the field of applied cryptography. Most papers focus solely on the soundness

[1] We keep with the "bad practice" prevalent amongst practitioners, and consider a STARK any Reed-Solomon (or more generally, Goppa) encoded argument system that uses the FRI low-degree test.

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. The reason for this might be the common design principle as a polynomial IOP, and it is not too difficult to randomize the polynomials so that their queried values perfectly hide the witness. Maybe because of this apparent plainness one tends to oversimplify the issue. Examples for gaps in the treatment[2]of zero-knowledge are Plonky2 [plo], Risc-Zero [Ris], Triton [tri], not to forget the summary on FRI low-degree test [Hab22].

In this note we discuss the subtleties of turning FRI-based STARKs into zero-knowledge. We explain the [BSCR+19] construction and clarify the importance of its masking step, and we discuss two techniques desirable in practice and which we did not find formally treated in literature: Randomizing the witness polynomials over the base field (while the verifier challenges are from an extension field), and decomposing the overall quotient into polynomials of smaller degree. We focus on the FFT-like 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 seems prevalent in practice, but we also sketch how to treat other types of decompositions.

The note is structured as follows. In Section 2 we describe the [BSCR+19] construction in the context of a simple toy protocol which uses FRI as a "polynomial commitment scheme". This toy protocol is already designed in respect to small field STARKs, and randomizes the witnesses with basefield polynomials. Then, in Section 3, we describe how to add zero-knowledge to a STARK of an algebraic intermediate representation (AIR) with transitional constraints, for the FFT decomposition as described above. (The other two types are quickly sketched in Section 4.) Our restrictive choice of IOP is for demonstration purposes; its generalization to more versatile protocols which use permutation or lookup arguments is straight-forward. In Section 5, we finally conclude with few remarks on the computational costs for the most greedy choice of randomization parameters.

We assume that the reader is familiar with the formal notions of zero-knowledge, and thus we skip their explicit definitions. However, even without explicit definitions we do not lack rigor. As we throughout assume an honest verifier (which is good enough for our purposes) the desired transcript simulator can be easily inferred from the proven properties of the involved probability distributions.

[2] In the meantime, the gaps in these repositories have been patched.

2 The BSCR construction

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, by introducing only a few randomization steps. As ususal, witness polynomials are randomized outside the trace domain, but unlike in IOPs with homomorphic properties one needs to take into account further oracle queries beyond the main protocol, in this case FRI proof of proximity.

For simplicity we discuss the [BSCR+19] construction in the context of the following toy protocol, summarized in Protocol 1 below: Let $\mathbb{F}_q$ be a finite field, $H \subset \mathbb{F}_q$ a multiplicative subgroup of smooth order $|H| = 2^n$, and F a finite extension field of $\mathbb{F}_q$, of degree

$$e = [F : \mathbb{F}_q].$$

Given witness polynomials[3]$w_1(X), \dots, w_M(X) \in \mathbb{F}_q[X]^{<|H|}$ the prover randomizes them outside the witness domain H by polynomial multiples of the vanishing polynomial,

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

for each $i = 1, \dots, M$, where $r_i(X) \hookleftarrow\!\$\ \mathbb{F}_q[X]^{<h}$ is of some appropriate degree of freedom h, where typically $h \ll |H|$, to be determined below. Then the verifier samples $n_{DEEP} \geq 1$ random queries from the extension field F, and asks for the values of the randomized polynomials over the query set. The prover claims them and both prover and the verifier engage in a FRI low-degree test with $n_{FRI} \geq 1$ query rounds on the evaluation claim quotients. To take into account the typically small increase of degree, we use a variant of the low-degree test to handle the strict degree bound $|H| + h$ without jumping to the next smooth bound, Protocol 2.

In order to ensure zero-knowledge the following measures are undertaken: First of all, the evaluation domain D for the Reed-Solomon code $\mathsf{RS}[\mathbb{F}_q, D, |H| + h]$, a coset of a smooth multiplicative subgroup of sufficiently large order, is chosen disjoint from H,

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

(This is default in many FRI implementations, even when zero-knowledge is not targeted.) Second, the degree of freedom in the randomization takes into account that a single extension field evaluation reveals the information of e base field elements,

$$e \cdot n_{\text{DEEP}} + n_{\text{FRI}} \leq h. \tag{1}$$

[3] Here, and in the sequel, $\mathbb{F}_q[X]^{<n}$ is short for the set of polynomials over $\mathbb{F}_q$ of degree less than n.

(This intuitive rule-of-thumb is proven sufficient by Lemma 1.) Third, a random mask polynomial

$$R(X) \leftarrow\$ F[X]^{<|H|+h-1}$$

is inserted into batching step of FRI, in order to prevent potential data leakage during the FRI folding cascade. (The decoupling property of this step is summarized by Lemma 2.)

We restrict to the case that $h \leq |H|$, which covers our use cases.

Protocol 1 (Toy IOP using FRI). *Let H and D be cosets of multiplicative subgroups of $\mathbb{F}_q$ as above, F be a degree-e extension of $\mathbb{F}_q$, and take $h \leq |H|$ satisfying* (1)*. The oracles are the values of the respective polynomials over D.*

1. *Given polynomials $w_1(X), \ldots, w_M(X) \in \mathbb{F}_q[X]^{<|H|}$, the prover samples $r_i(X) \leftarrow\$ \mathbb{F}_q[X]^{<h}$, $i = 1, \ldots, M$, and sends the oracles for*
$$\hat{w}_i(X) = w_i(X) + v_H(X) \cdot r_i(X), \tag{2}$$
$i = 1, \ldots, M$, to the verifier.
2. *The verifier samples $n_{DEEP} \geq 1$ queries $z_1, \ldots z_{n_{DEEP}} \leftarrow\$ F \setminus (D \cup H)$ and sends them to prover, which responds with the values*
$$\vec{v}_{i,DEEP} = \hat{w}_i(z)|_{z \in Q_{DEEP}} \in F^{Q_{DEEP}}$$
for every $i = 1, \ldots, M$, where $Q_{DEEP} = \{z_1, \ldots, z_{n_{DEEP}}\}$.

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

In our variant of FRI, we reduce proximity to $\mathsf{RS}[F, D, |H| + h]$ to that of a decomposition being close to $\mathsf{RS}[F, D, |H|]$ which has again a two-adic rate.

Protocol 2 (FRI batch evaluation proof with zk [BSCR+19]). *Under the same assumptions of Protocol 1, given oracles $\hat{w}_1(X), \ldots, \hat{w}_M(X) \in \mathbb{F}_q[X]^{<|H|+h}$ and evaluation claims $\vec{v}_{i,DEEP} = w_i(z)|_{z \in Q_{DEEP}}$ over a set $Q_{DEEP} = \{\gamma_1, \ldots, \gamma_{|Q_{DEEP}|}\} \subset F$ of size $|Q_{DEEP}| \leq n_{DEEP}$.*

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 both $h_0(X) \in F[X]^{<h-1}$, $h_1(X) \in F[X]^{<|H|}$ subject to*
$$\begin{aligned} h(X) &= h_0(X) + X^{h-1} \cdot h_1(X) \\ &= R(X) + \sum_{i=1}^{M} \lambda^{|Q_{DEEP}| \cdot (i-1)} \cdot \sum_{j=1}^{|Q_{DEEP}|} \lambda^j \cdot \frac{\hat{w}_i(X) - v_i(\gamma_j)}{X - \gamma_j}, \end{aligned} \tag{3}$$

and both prover and verifier run FRI on $h_0(X), h_1(X)$ for $\mathsf{RS}[F, D, |H|]$ with n_{FRI} query rounds. (This variant of FRI treats two polynomials in the first folding step.)

Let us discuss why the mask polynomial $R(X)$ is crucial for zero-knowledge. Otherwise it may happen that during the folding cascade of FRI, the entropy of the randomizer polynomials $v_H(X) \cdot r_i(X)$ is drastically reduced, so that the witness data in the folded oracles is not sufficiently secured. For simplicity we assume that $h = |H|$. The argument however generalizes to arbitrary choices of h. Since $v_H(X)$ is an even function, the decomposition of the randomizer polynomial is

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

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

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

halving in each step, 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 open the folded oracle in the last step, without revealing the folding of the witness polynomials itself. (This argument is independent of r the number of folding steps.)

Remark 1. In Plonk-like proof systems with non-succinct selector polynomials, one typically uses a randomization strategy which does not increase the degree beyond $|H|$. The witness domain is restricted to a (non-group) subset H' of H, and the remaining space $H \setminus H'$, which is not touched by any of the constraints, is used to place random values. In this context, the entropy loss by folding also holds whenever $H \setminus H'$ is contained in a coset of a non-trivial subgroup U of H. To see this, decompose $\hat{w}(X) \in \mathbb{F}_q[X]^{<|H|}$ into $\hat{w}(X) = w_{H\setminus U}(X) + v_{H\setminus U}(X) \cdot \hat{w}_U(X)$, where $w_{H\setminus U}(X)$ depends on the witness polynomial only, and $v_{H\setminus U}(X) \cdot \hat{w}_U(X)$ contains all the randomness. The vanishing polynomial of $H \setminus U$ is still even, and the space $v_{H\setminus U}(X) \cdot F[X]^{<|U|}$ halves in each of the folding steps, unless it is a singleton, in which case it remains a singleton space: If $U = \{a\}$, then $v_{H\setminus\{a\}}(X) \cdot F$ is folded into the span $v_{H^2\setminus\{a^2\}}(X) \cdot F$. However, if the randomization domain is not contained in a coset of a non-trivial subgroup, it might be that the mask polynomial is not needed for zero-knowledge.

We prove the zero-knowledge property of Protocol 1 by the following two lemmas, Lemma 1 and Lemma 2.

Lemma 1. *Fix the query sets Q_{DEEP} and Q_{FRI} of size $|Q_{DEEP}| \leq n_{DEEP}$ and $|Q_{FRI}| \leq n_{FRI}$. Then the joint distribution of $\vec{v}_i = \hat{w}_i(z)|_{z \in Q_{DEEP} \cup Q_{FRI}}$, where $i = 1, \ldots, M$, is independent from the witness polynomials $w_1(X), \ldots, w_M(X)$.*

Proof. The only subtlety here is that the randomizer polynomials $v_H(X) \cdot r_i(X)$ are over the base field $\mathbb{F}_q$ but the DEEP queries are from the extension field F.

Let e be the degree of the extension, and take the Galois group closure of the query set, $\bar{Q} = \bigcup_{\phi \in \mathsf{Gal}(F/\mathbb{F}_q)} \phi^k(Q_{\text{DEEP}}) \cup Q_{\text{FRI}}$. Since $\bar{Q}$ is invariant under $\mathsf{Gal}(F/\mathbb{F}_q)$, so is its vanishing polynomial $v(X) = \prod_{z \in \bar{Q}}(X - z)$, showing that $v(X)$ is actually from $\mathbb{F}_q[X]$. Since $|\mathsf{Gal}(F/\mathbb{F}_q)| = e$, we have

$$\deg v(X) = |\bar{Q}| \leq e \cdot n_{\text{DEEP}} + n_{\text{FRI}}.$$

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

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

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

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

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

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

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

From this and the linearity of E, it follows that by drawing $r_i(X)$ independently and uniformly from $\mathbb{F}_q[X]^{<h}$ 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 $E(\mathbb{F}_q[X])^M$, independent 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 $E : \mathbb{F}_q[X] \longrightarrow F^Q$ over the set $Q = Q_{\text{DEEP}} \cup Q_{\text{FRI}}$, and it can be efficiently simulated for example by sampling uniformly from $\mathbb{F}_q[X]^{<|H|+h}$ and applying E. In the terminology of [BSCR$^+$19] the lemma shows that Protocol 1 is perfect honest-verifier zero-knowledge *against query bound* n_{FRI}, meaning that it is zero-knowledge even under further (at most) n_{FRI} queries beyond the protocol execution.

Lemma 2 (Decoupling Lemma)**.** *Fix Q_{DEEP} and Q_{FRI} as in Lemma 3, and $\lambda \in F$. Given the values $\hat{w}_i(z)|_{z\in Q_{DEEP}\cup Q_{FRI}}$, $i = 1, \ldots, M$, the joint distribution of $R(z)|_{z\in Q_{FRI}}$ 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 the concrete form of the polynomials $\hat{w}_1(X), \ldots, \hat{w}_M(X)$ the distribution of the component polynomials $h_0(X)$ and $h_1(X)$, and the entire further transcript of FRI is also independent.

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

$$L_{\vec{v}_1,\ldots,\vec{v}_M,\vec{r}} = \left\{ h(X) \in F[X]^{<|H|+h-1} \,:\, h(X) \text{ satisfies (3) at all } z \in Q_{\mathrm{FRI}} \right\}.$$

Take any subset $Q' \subset D$ disjoint to $Q = Q_{\mathrm{FRI}}$ 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 $\vec{r} = R(z)|_{z\in Q_{\mathrm{FRI}}}$ is independent of the concrete choice of $\hat{w}_1(X), \ldots, \hat{w}_M(X)$. □

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 $z_1, \ldots, z_{n_{\mathrm{DEEP}}} \hookleftarrow\$ F \setminus (D \cup H)$, $x_1, \ldots, x_{n_{\mathrm{FRI}}} \hookleftarrow\$ D$ uniformly from the respective sets, and draws $\hat{w}_i(X) \hookleftarrow\$ \mathbb{F}_q[X]^{<|H|+h}$ uniformly at random, so that their values over $Q = Q_{\mathrm{DEEP}} \cup Q_{\mathrm{FRI}}$ (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 $x_1, \ldots, x_{n_{\mathrm{FRI}}}$ 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 Within a STARK, using FFT decomposition

Consider a trace composed of $M > 0$ witness columns w_i, $i = 1, \ldots, M$, where $w_i : H \longrightarrow \mathbb{F}_q$ 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) [BSGKS20, BSBHR18] 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

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

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,

$$\mathsf{d} = \max_i \deg P_i.$$

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* $d := \mathsf{d} - 1$.

In terms of the low-degree extensions $w_1(X), \ldots, w_M(X) \in \mathbb{F}_q[X]^{<|H|}$ of the trace columns (we overload notation here), 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*

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

for some low-degree polynomial $q(X) \in F[X]$, where λ is drawn from the extension field F.

In our interactive oracle proof, Protocol 3, the prover decomposes 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 a one (or more) random points z_i, $i = 1, \ldots, n_{\text{DEEP}}$, from the extension field F. The evaluation claims for the polynomials are then proven 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.

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

$$\hat{w}_i(X) := w_i(X) + v_H(X) \cdot r_i(X) \in \mathbb{F}_q[X]^{<|H|+h}, \tag{4}$$

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

$$2 \cdot d \cdot (e \cdot n_{\text{DEEP}} + n_{\text{FRI}}) + n_{\text{FRI}} \leq h \leq |H|. \tag{5}$$

Here, e is the degree of the extension F, $n_{\text{DEEP}} \geq 1$ the number of DEEP queries (excluding their translates by g), and $n_{\text{FRI}} \geq 1$ is the number of FRI query rounds. The evaluation domain $D \subset \mathbb{F}_q$ for the low-degree test is a coset of a smooth multiplicative subgroup, large enough so that the rate $\hat{\rho} = \hat{k}/|D|$ of the Reed-Solomon code $\mathsf{RS}[F, D, \hat{k}]$ with

$$\hat{k} = |H| + \left\lceil \frac{d+1}{d} \cdot h \right\rceil \tag{6}$$

is as small as desired. Again, the evaluation domain is disjoint from the trace domain, $D \cap H = \varnothing$.

Protocol 3 (IOP for AIR using DEEP-ALI)**.** *Let $\mathbb{F}_q$, H, D, and F as above and let $\hat{w}_1(X), \cdots \hat{w}_M(X) \in \mathbb{F}_q[X]^{<|H|+h}$ 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 the prover computes $q(X) \in F[X]^{<d\cdot|H|+(d+1)\cdot h}$ such that*

$$\sum_{i=1}^{C} \lambda^{i-1} \cdot P_i\left(s_i(X), \hat{w}_1(X), \cdots, \hat{w}_M(X), \hat{w}_1(g \cdot X), \cdots, \hat{w}_M(g \cdot X)\right)$$
$$= Z_H(X) \cdot q(X). \tag{7}$$

   It splits it into the unique polynomials $q_j(X) \in F[X]^{<\hat{k}}$, $j = 1, \ldots, d$, with $\hat{k}$ as above, and subject to

$$q(X) = \sum_{j=1}^{d} X^{(j-1)} \cdot q_j(X^d). \tag{8}$$

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

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

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

   $i = 1, \ldots, M$ 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, Protocol 2, with h replaced by $\lceil (d+1)/d \cdot h \rceil$, and using n_{FRI} query rounds. The verifier accepts if Protocol 2 passes and if the evaluation claims satisfy the overall identity (7) *at each $X = z_j$, $j = 1, \ldots, n_{DEEP}$.*

Let us explain the intuition behind the degree bound in Equation (5). 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 $\mathbb{F}_q$, 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 (7), it would be sufficient to randomize the witness polynomials against

$$2 \cdot (e \cdot n_{\mathrm{DEEP}} + n_{\mathrm{FRI}})$$

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

$$(q_1(z^d), \ldots, q_d(z^d)) = \mathsf{FFT}(q(X)|_{z \cdot U}), \tag{9}$$

where U is the subgroup of the d-th roots of unity. Hence each query of the component polynomials amounts to $|U| = d$ times as many queries of $q(X)$, which explains the factor d in (5). (The additional n_{FRI} is due to the fact that the FRI queries on $\hat{w}_i(X)$ and $q_1(X), \ldots, q_d(X)$ do not overlap.) This line of argument goes through whenever the extension field F contains all d-th roots of unity. In general one has to be more careful for maintaining this bound, as we will see in the proof of the following lemma.

Lemma 3. *Fix $\lambda \in F$ and query sets Q_{DEEP} and Q_{FRI} of size $|Q_{DEEP}| \leq n_{DEEP}$ and $|Q_{FRI}| \leq n_{FRI}$ such that Equation* (5) *holds. Then the joint distribution of*

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

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

Proof. The proof is similar to that of Lemma 1, but with a more careful choice of evaluation set, due to the fact that U the set of d-th roots of unity might not be contained in F. Let $v_D(X) = X^{|D|} - a$, with $a \in \mathbb{F}_q^*$, be the vanishing polynomial of the domain D.

Let K be an extension of F in which $v_D(X^d) = X^{d \cdot |D|} - a$ splits. (Such an extension contains all $(d \cdot |D|)$-th roots of unity, and the domain D has a preimage under the d-th power map which is mapped onto D in a d-to-1 manner.) 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 $Q_1 = \bigcup_{\phi \in \mathsf{Gal}(F/\mathbb{F}_q)} \phi\left(Q_{\mathrm{DEEP}} \cup g \cdot Q_{\mathrm{DEEP}}\right)$, $Q_2 = Q_{\mathrm{FRI}} \cup g^d \cdot Q_{\mathrm{FRI}}$, and $Q_3 = Q_{\mathrm{FRI}}$. By construction the polynomial belongs to $\mathbb{F}_q[X]$, since it is a polynomial from $F[X]$ which is invariant under $\mathsf{Gal}(F/\mathbb{F}_q)$, and it splits over K, having $\bar{Q} \subseteq K$ as set of roots. The radical of $p(X)$, i.e. the vanishing polynomial $v(X)$ of $\bar{Q}$, is again a polynomial from $\mathbb{F}_q[X]$, and its degree is

$$\deg v(X) = |\bar{Q}| \leq 2 \cdot d \cdot (e \cdot n_{\mathrm{DEEP}} + n_{\mathrm{FRI}}) + n_{\mathrm{FRI}} \leq h.$$

By the assumption on the DEEP 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 $E : \mathbb{F}_q[X] \longrightarrow K^{\bar{Q}}$ is the ideal generated by $v(X)$ of degree $|\bar{Q}|$, and hence its image is a $\mathbb{F}_q$-linear subspace of $K^{\bar{Q}}$ with $\dim E(\mathbb{F}_q[X]) = |\bar{Q}|$. Likewise, since $v_H(X)$ and $v(X)$ are coprime, the kernel within the randomizer space $v_H(X) \cdot \mathbb{F}_q[X]^{<h}$ is

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

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

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

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 $E(\mathbb{F}_q[X])^M$ and independent of the 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 and (9), their distribution is also independent from the witness polynomials. Finally, restricting $\bar{Q}$ to the query points of the protocol yields the claim of the lemma. □

Remark 5. Again, the proof of the Lemma shows that the distribution of the queried values is efficiently simulated by means of the evaluation map over $\bar{Q}$ in the splitting field K of $v_D(X^d)$ over F. Since F contains all $|D|$-th roots of unity, the extension degree $[K : F] \leq d$, and K can be constructed with overwhelming success in probabilistic polynomial time, with respect to the size of the AIR and the security parameter[4], and the same holds for the set $\bar{Q}$. In other words, the protocol "on top of batch FRI" as a polynomial IOP is (perfect) zero-knowledge against query bound n_{FRI}.

[4] For example, determine the extension degree $n = [K : F]$ as the smallest n, $1 \leq n \leq d$, so that $d \cdot |D|$ divides $|F|^n - 1$, 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 $> 1 - 1/|F|^{1/2}$.

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

Proof. The simulator samples the verifier challenges $\lambda \leftarrow\$ F$, $z_{\text{DEEP}} \leftarrow\$ F \setminus \big(\bar{D} \cup H\big)$, where $\bar{D} := \{y \in F : y^d \in D\}$, and $x_1, \ldots, x_{n_{\text{FRI}}} \leftarrow\$ D$.

Then, it constructs the splitting field of $X^{d \cdot |D|} - 1$ over F and the evaluation set $\bar{Q}$ as in the proof of the Lemma 3. (See Remark 5 for the computational complexity of this step.), samples polynomials $\hat{w}_i(X) \leftarrow\$ \mathbb{F}_q[X]^{<|H|+h}$ 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 (7) and (9), and interpolates them by arbitrary polynomials $q_1(X), \ldots, q_d(X)$ from $F[X]^{<|H|+h}$. 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 $x_1, \ldots, x_{n_{\text{FRI}}}$ 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. □

We finally remark that our restrictive variant of AIR is chosen for simplicity only. The generalization of Lemma 3 and Theorem 6 to higher offset constraints, and Plonk-like protocols with additional permutation and lookup arguments, is straight-forward.

4 Other types of decompositions

We sketch how to tackle the cases of the other two commonly used in practice decompositions, namely, the "canonical" decomposition based on monomials, and the Lagrange decomposition.

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 7 in Protocol 3 with h, the degree of freedom of the witness randomizer, to be specified later. By the canonical decomposition, we mean

$$q(X) = \sum_{i=1}^{d} X^{\hat{k} \cdot (i-1)} \cdot q_i(X), \tag{10}$$

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 (10) can be randomized, using a technique from [GWC19]: In order to maintain the identity

$$q(X) = \sum_{i=1}^{d} X^{\hat{k}\cdot(i-1)} \cdot \hat{q}_i(X), \tag{11}$$

on draws $t_i(X) \leftarrow\!\!\$\; F[X]^{<h_q}$, $i = 1, \ldots, d-1$, 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_q is chosen such that

$$n_{\mathrm{DEEP}} + n_{\mathrm{FRI}} \leq h_q. \tag{12}$$

With the above modification, Protocol 3 goes through unchanged except for Equation (5) which becomes

$$2 \cdot (e \cdot n_{\mathrm{DEEP}} + n_{\mathrm{FRI}}) \leq h \leq |H|, \tag{13}$$

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

Lemma 4. *Fix $\lambda \in F$ and query sets Q_{DEEP} and Q_{FRI} of size $|Q_{DEEP}| \leq n_{DEEP}$ and $|Q_{FRI}| \leq n_{FRI}$ such that Equations* (12) *and* (13) *hold. Then the joint distribution of*

$$\begin{aligned}
(\hat{w}_1(z), \hat{w}_1(g \cdot z), \ldots, \hat{w}_M(z), \hat{w}_M(g \cdot z), \hat{q}_1(z), \ldots, \hat{q}_d(z))|_{z \in Q_{DEEP}}, \\
(\hat{w}_1(z), \ldots, \hat{w}_M(z), \hat{q}_1(z), \ldots, \hat{q}_d(z))|_{z \in Q_{FRI}},
\end{aligned}$$

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 $F[X]^{<h_q}$, 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 11 instead of 8, is perfect honest-verifier zero-knowledge.*

Proof. The simulator samples the verifier challenges $\lambda \leftarrow\!\!\$\ F$, $z_1, \ldots, z_{n_{\mathrm{DEEP}}} \leftarrow\!\!\$\ F \setminus (D \cup H)$, and $x_1, \ldots, x_{n_{\mathrm{FRI}}} \leftarrow\!\!\$\ D$. It then samples $\hat{w}_i(X) \leftarrow\!\!\$\ \mathbb{F}_q[X]^{<|H|+h}$ and $\hat{q}_i(X) \leftarrow\!\!\$\ F[X]^{<\hat{k}+h_q}$, $i = 1, \ldots, d-1$, uniformly at random, and computes their values over $Q = Q_{\mathrm{DEEP}} \cup Q_{\mathrm{FRI}}$. Given these values, it takes any $\hat{q}_d(X) \in F[X]^{<|H|+h_q}$ satisfying the overall constraint (7) and the decomposition identity (11) 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

$$q(X) = \sum_{i=1}^{d} L_{H_i}(X) \cdot q_i(X), \tag{14}$$

where

$$L_{H_i}(X) = c_i \cdot \prod_{j \neq i} v_{H_j}(X) \tag{15}$$

is the selector polynomial of the coset H_i (normalized so that $L_{H_i}(x) = 1$ over H_i), and $q_i(X) \in F[X]^{<|H|}$. This type of decomposition is particularly efficient: Each $q_i(X)$ is directly obtained from the values of $q(X)$ over H_i via an FFT of witness domain size $|H|$, without any precomputation on the values as for the FFT decomposition, or an FFT of larger size, as often used by the canonical decomposition from the previous section.

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 (14), 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

$$\hat{q}_i(X) = q_i(X) + v_{H_i}(X) \cdot t_i(X) \in F[X]^{<|H|+h_q}, \tag{16}$$

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

$$\hat{q}_d(X) = q_d(X) - v_{H_d}(X) \cdot \sum_{k=1}^{d-1} c_d^{-1} \cdot c_i \cdot t_i(X), \tag{17}$$

with the normalizing coefficients c_i and c_d from (15). This choice still satisfies

$$q(X) = \sum_{i=1}^{d} L_{H_i}(X) \cdot \hat{q}_i(X). \tag{18}$$

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

$$n_{\text{DEEP}} + n_{\text{FRI}} \leq h_q, \tag{19}$$

and

$$2 \cdot (e \cdot n_{\text{DEEP}} + n_{\text{FRI}}) \leq h \leq |H|. \tag{20}$$

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

Lemma 5. *Fix $\lambda \in F$ and query sets Q_{DEEP} and Q_{FRI} of size $|Q_{DEEP}| \leq n_{DEEP}$ and $|Q_{FRI}| \leq n_{FRI}$ such that Equations* (19) *and* (20) *hold. Then the joint distribution of*

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

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 18 instead of 8, 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 $\hat{w}_i(X) \leftarrow\!\!\$\, \mathbb{F}_q[X]^{<|H|+h}$ and $\hat{q}_i(X) \leftarrow\!\!\$\, F[X]^{<|H|+h_q}$, $i = 1, \ldots, d-1$, and takes any $\hat{q}_d(X) \in F[X]$ of degree less than $|H| + \max\{(d+1) \cdot h, h_q\}$ satisfying the overall constraint (7) and the decomposition identity (11) 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 degrees are 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 (5). 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_{FRI}.) Therefore, in the wide trace regime, the hashing costs remain essentially the same.

If the non-zk parameters are optimized for constraint evaluation, which means that $d \cdot |H| = |D|$, the overall FFT costs 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 $\mathbb{F}_q[X]^{<|H|+h}$ over a set of size $(d+1) \cdot h$, and $O(h)$ covers the cost of adding $r_i(X) \cdot (X^{|H|} - 1)$ to the Fourier transform $w_i(X)$ of the witness column. Under the simplifying assumption that the target set of $\mathsf{Eval}_{|H|+h}((d+1) \cdot h)$ is a coset, 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 we have used 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.

Acknowledgements

The first author likes to thank Ariel Gabizon for pointing out a misinterpretation of the role of the masking polynomial in [Hab22].

References

[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 noninteractive 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`.

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

[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`.