

Advanced zk-STARKs

Alan Szepieniec

艾伦·余丕涅茨

alan@neptune.cash



neptune

<https://neptune.cash/>



Triton VM

<https://triton-vm.org/>

<https://asz.ink/presentations/2025-09-18-Advanced-zkSTARKs.pdf>

Table of Contents

- Retrospective

- Optimizations

 - Batching

 - Quotient Segmentation

 - Grinding

- Enhancements

 - Zero-Knowledge

 - Randomized AIR with Preprocessing

- VM Architecture

 - Overview

 - Communication Arguments

 - Memory

- Other Topics

Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

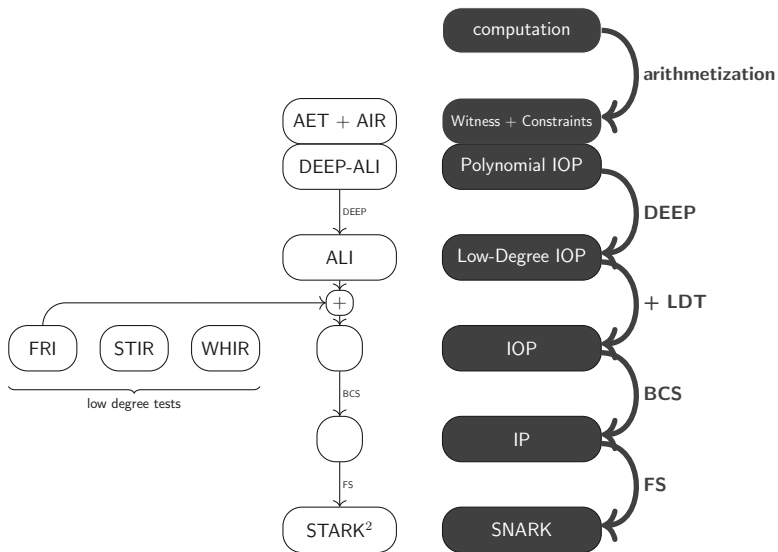
Overview

Communication Arguments

Memory

Other Topics

STARK Compilation Pipeline




STARK Diagram



algebraic execution trace


STARK Diagram



low-degree extension

low-degree extended trace

STARK Diagram




composition with AIR constraints

division by zero-finders

quotients

STARK Diagram




build Merkle tree

STARK Diagram




interact with verifier

STARK Diagram




sample out-of-domain point

STARK Diagram




produce out-of-domain rows


STARK Diagram




STARK Diagram



STARK Diagram




STARK Diagram




random linear combination

STARK Diagram




build Merkle tree

STARK Diagram




interact with verifier


STARK Diagram



STARK Diagram



STARK Diagram



obtain FRI indices
open indicated rows

STARK Diagram




Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory

Other Topics

Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory

Other Topics

Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory

Other Topics

Batching

– linear

$$\sum_{i=0}^{n-1} r_i \mathbf{c}_i$$

$$\epsilon = \epsilon_{\text{GAP}}(\delta_0)$$

Batching


– linear	$\sum_{i=0}^{n-1} r_i \mathbf{c}_i$	$\epsilon = \epsilon_{\text{GAP}}(\delta_0)$
– univariate	$\sum_{i=0}^{n-1} r^i \mathbf{c}_i$	$\epsilon = n \cdot \epsilon_{\text{GAP}}(\delta_0)$

Batching


– linear	$\sum_{i=0}^{n-1} r_i \mathbf{c}_i$	$\epsilon = \epsilon_{\text{GAP}}(\delta_0)$
– univariate	$\sum_{i=0}^{n-1} r^i \mathbf{c}_i$	$\epsilon = n \cdot \epsilon_{\text{GAP}}(\delta_0)$
– multilinear	$\sum_{i=0}^{n-1} \mathbf{r}^{\perp i \cdot \perp} \mathbf{c}_i$	$\epsilon = \log n \cdot \epsilon_{\text{GAP}}(\delta_0)$
	$= \sum \left(\prod_{j=0}^{\log n - 1} r_j^{b_j} (1 - r_j)^{(1-b_j)} \right) \mathbf{c}_i$	

Batching


- linear $\sum_{i=0}^{n-1} r_i \mathbf{c}_i$ $\epsilon = \epsilon_{\text{GAP}}(\delta_0)$
- univariate $\sum_{i=0}^{n-1} r^i \mathbf{c}_i$ $\epsilon = n \cdot \epsilon_{\text{GAP}}(\delta_0)$
- multilinear $\sum_{i=0}^{n-1} \mathbf{r}^{\perp i \perp} \mathbf{c}_i$ $\epsilon = \log n \cdot \epsilon_{\text{GAP}}(\delta_0)$



Univariate Batching



Batch Before DEEP



Batch Constraints




Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory

Other Topics

Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory


Other Topics

Quotient Segmentation

What if $\deg(\mathcal{C}) > 2$?


Quotient Segmentation

What if $\deg(\mathcal{C}) > 2$?




Quotient Segmentation

What if $\deg(\mathcal{C}) > 2$?




Quotient Segmentation

What if $\deg(\mathcal{C}) > 2$?



Quotient Segmentation

What if $\deg(\mathcal{C}) > 2$?



Quotient Segmentation

What if $\deg(\mathcal{C}) > 2$?




Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory

Other Topics

Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory


Other Topics

Grinding

NO GRINDING

Prover


Verifier



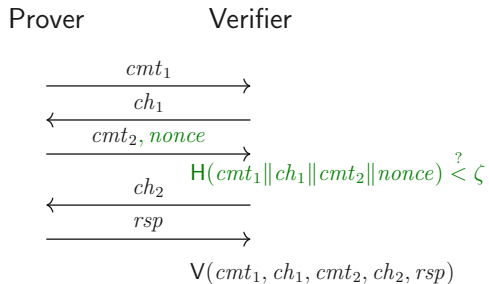
$V(cmt_1, ch_1, cmt_2, ch_2, rsp)$

Grinding

NO GRINDING




WITH GRINDING



Grinding

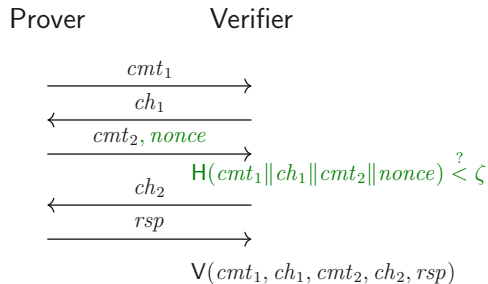
NO GRINDING



SOUNDNESS
ERROR

ϵ


WITH GRINDING



$\epsilon \cdot \zeta$

Grinding

NO GRINDING



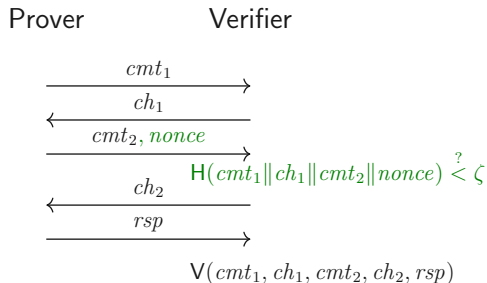
SOUNDNESS
ERROR

ϵ

PROVER
WORK

W

WITH GRINDING




$\epsilon \cdot \zeta$

$W + \zeta^{-1}$

Grinding

NO GRINDING




SOUNDNESS
ERROR

ϵ

PROVER
WORK

$W \approx 2^{30}$

WITH GRINDING



$\epsilon \cdot \zeta$

+ 10 bits

$W + \zeta^{-1}$

+ 0.1%

Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory

Other Topics

Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory

Other Topics

Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory

Other Topics

Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory

Other Topics

Zero Knowledge

zero-knowledge \Leftrightarrow transcript is independent of witness

Zero Knowledge

zero-knowledge \Leftrightarrow transcript is independent of witness

\rightarrow *mask with randomness*

Zero Knowledge

zero-knowledge \Leftrightarrow transcript is independent of witness

\rightarrow *mask with randomness*


1. salted Merkle leafs (optional)
2. batch randomizer polynomial
3. trace randomizer values

Zero Knowledge

zero-knowledge \Leftrightarrow transcript is independent of witness

\rightarrow *mask with randomness*

1. salted Merkle leafs (optional)
2. batch randomizer polynomial
3. trace randomizer values




Zero Knowledge

zero-knowledge \Leftrightarrow transcript is independent of witness

\rightarrow *mask with randomness*

1. salted Merkle leafs (optional)
2. batch randomizer polynomial
3. trace randomizer values




Zero Knowledge


zero-knowledge \Leftrightarrow transcript is independent of witness

\rightarrow *mask with randomness*

1. salted Merkle leafs (optional)
2. batch randomizer polynomial
3. trace randomizer values




Salted Merkle Leafs



$$\text{parent} = H(\text{left child} || \text{right child})$$

-  leafs
-  root
-  salts

Salted Merkle Leafs




$$\text{parent} = H(\text{left child} || \text{right child})$$



Motivation: make internal Merkle nodes ○ independent of un-opened data

Salted Merkle Leafs



$$\text{parent} = H(\text{left child} || \text{right child})$$



Motivation: make internal Merkle nodes  independent of un-opened data

→ in ROM: data = H(leaf) already independent

→ in standard model: concat-then-hash not enough ×

→ use perfectly-hiding + computationally-binding commitment scheme instead ✓

Batch Randomizer Polynomial

include uniformly random polynomial into batch

Batch Randomizer Polynomial

include uniformly random polynomial into batch

→ How: add *unconstrained* and *random trace* column

Batch Randomizer Polynomial

include uniformly random polynomial into batch

→ How: add *unconstrained* and *random trace* column


→ Why: make low-degree tested codeword independent of trace

Batch Randomizer Polynomial

include uniformly random polynomial into batch

→ How: add *unconstrained* and *random* trace column

→ Why: make low-degree tested codeword independent of trace



Trace Randomizer Values: Interleaving

interleave trace with random rows

→ Why: make observed rows independent of trace

Trace Randomizer Values: Interleaving

interleave trace with random rows

→ Why: make observed rows independent of trace


○ trace values

□ randomizer values

× evaluation domain

= coset of subgroup of order $2N/\rho$

code rate $\rho = \frac{\#\text{○} + \#\text{□}}{\#\text{×}}$ (!!!)



Trace Randomizer Values: Interleaving

interleave trace with random rows

→ Why: make observed rows independent of trace

$$w \cdot \#\square \geq \underbrace{t \cdot w}_{\text{FRI}} + \left(\underbrace{2 \cdot w}_{\text{DEEP}} + \underbrace{(t+1) \cdot k}_{\text{quotient segments}} \right) \cdot \underbrace{e}_{\text{extension degree}}$$


○ trace values

□ randomizer values

× evaluation domain

= coset of subgroup of order $2N/\rho$

code rate $\rho = \frac{\#\circ + \#\square}{\#\times}$ (!!!)



Trace Randomizer Values: Stingy

pad then concatenate random rows

→ Why: make observed rows independent of trace

$$w \cdot \#\square \geq \underbrace{t \cdot w}_{\text{FRI}} + \left(\underbrace{2 \cdot w}_{\text{DEEP}} + \underbrace{(t+1) \cdot k}_{\substack{\text{quotient} \\ \text{segments}}} \right) \cdot \underbrace{e}_{\substack{\text{extension} \\ \text{degree}}}$$

Trace Randomizer Values: Stingy

pad then concatenate random rows

→ Why: make observed rows independent of trace

$$w \cdot \#\square \geq \underbrace{t \cdot w}_{\text{FRI}} + \underbrace{(2 \cdot w)}_{\text{DEEP}} + \underbrace{(t+1) \cdot k}_{\text{quotient segments}} \cdot \underbrace{e}_{\text{extension degree}}$$


○ trace values

▣ randomizer values

✱ evaluation domain

= coset of subgroup of order N/ρ

$$\text{code rate } \rho = \frac{\#\circ + \#\square}{\#\times}$$



Trace Randomizer Values: Stingy

pad then concatenate random rows

→ Why: make observed rows independent of trace

$$w \cdot \#\square \geq \underbrace{t \cdot w}_{\text{FRI}} + \underbrace{(2 \cdot w)}_{\text{DEEP}} + \underbrace{(t+1) \cdot k}_{\text{quotient segments}} \cdot \underbrace{e}_{\text{extension degree}}$$


○ trace values

□ randomizer values

✱ evaluation domain

= coset of subgroup of order N/ρ

$$\text{code rate } \rho = \frac{\#\circ + \#\square}{\#\times}$$



complications:

- multiply quotients by $\deg \square - 1 Z(X)$
- quotient degree increases
 - more segments, or
 - worse rate \Rightarrow more indices

Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory

Other Topics

Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory

Other Topics

Two-Stage DEEP-ALI


0-stage:

1-stage:

2-stage:

Two-Stage DEEP-ALI


0-stage:



1-stage:


2-stage:

Two-Stage DEEP-ALI

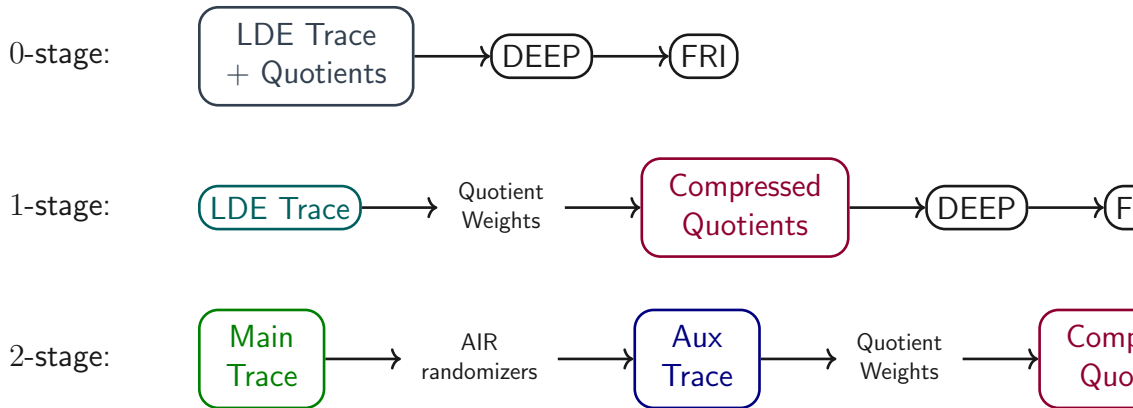


2-stage:

Two-Stage DEEP-ALI




Two-Stage DEEP-ALI



randomized AIR \gg deterministic AIR

Two-Stage DEEP-ALI (Diagram)



Preprocessing

pre-commit to separate “trace” table

- look-up tables ✓
- circuits ✓
- extra Merkle tree ✗
- need to know trace length beforehand ✗




Table Of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory

Other Topics

Table Of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory

Other Topics

Table Of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory

Other Topics

Table Of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory


Other Topics

Processor (Example)

clk	<i>clock / cycle counter</i>
ip	<i>instruction pointer</i>
ci	<i>current instruction</i>
arg ₀	<i>instruction argument 0</i>
arg ₁	<i>instruction argument 1</i>
arg ₂	<i>instruction argument 2</i>
ramp	<i>RAM pointer</i>
ramv	<i>RAM value</i>
reg ₀	<i>register 0</i>
reg ₁	<i>register 1</i>
reg ₂	<i>register 2</i>
reg ₃	<i>register 3</i>

Processor (Example)

clk	<i>clock / cycle counter</i>
ip	<i>instruction pointer</i>
ci	<i>current instruction</i>
arg ₀	<i>instruction argument 0</i>
arg ₁	<i>instruction argument 1</i>
arg ₂	<i>instruction argument 2</i>
ramp	<i>RAM pointer</i>
ramv	<i>RAM value</i>
reg ₀	<i>register 0</i>
reg ₁	<i>register 1</i>
reg ₂	<i>register 2</i>
reg ₃	<i>register 3</i>



Processor AIR Constraints (Example)

jmp a

jump to instruction reg_a

$$\text{dest} = \sum_{i=0}^3 \text{reg}_i \prod_{j \neq i} \frac{\text{arg}_0 - j}{i - j} \quad \text{value of } \text{reg}_a$$

$$\text{jump} = \text{dest} - \text{ip}^* \quad \text{update value of ip (jump case)}$$

$$\text{nojump} = \text{ip} + 1 - \text{ip}^* \quad \text{update value of ip (no jump)}$$

$$\text{selector} = \prod_{\text{instr} \in \mathcal{I} \setminus \{\text{jmp}\}} \frac{\text{instr} - \text{ci}}{\text{instr} - \text{jmp}} \quad 1 \text{ iff } \text{ci} = \text{jmp}$$


$$\text{selector} \cdot \text{jump} + (1 - \text{selector}) \cdot \text{nojump}$$

Communication Lines




Processor


Communication Lines



Communication Lines



Communication Lines



Communication Lines




Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory

Other Topics

Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory

Other Topics


Permutation Argument

$b = \sigma(a)$ for some permutation σ





Permutation Argument

$b = \sigma(a)$ for some permutation σ




Permutation Argument

$b = \sigma(a)$ for some permutation σ




Permutation Argument

$b = \sigma(a)$ for some permutation σ



Permutation Argument

$b = \sigma(a)$ for some permutation σ



Soundness

$$\Pr_x[p_a(x) = p_b(x) \mid p_a(X) \neq p_b(X)] \leq \frac{N}{\mathbb{F}}$$

Evaluation Argument



$$a[m_a] = b[m_b] \text{ for } m_a, m_b \subseteq \{0, \dots, N-1\}$$

a	m_a

b	m_b



Evaluation Argument

$$a[m_a] = b[m_b] \text{ for } m_a, m_b \subseteq \{0, \dots, N-1\}$$




Evaluation Argument

$$a[m_a] = b[m_b] \text{ for } m_a, m_b \subseteq \{0, \dots, N-1\}$$




Evaluation Argument

$$a[m_a] = b[m_b] \text{ for } m_a, m_b \subseteq \{0, \dots, N-1\}$$



Evaluation Argument

$$a[m_a] = b[m_b] \text{ for } m_a, m_b \subseteq \{0, \dots, N-1\}$$



Soundness

$$\Pr_x[f_a(x) = f_b(x) \mid f_a(X) \neq f_b(X)] \leq \frac{N}{\mathbb{F}}$$

Lookup Argument

$a \equiv b$ as sets

a	m_a

b	m_b

Lookup Argument

$a \equiv b$ as sets

a	m_a

b	m_b

$$\log \frac{d}{dX} [f(X)] = \frac{f'(X)}{f(X)}$$

Lookup Argument

a	m_a

$a \equiv b$ as sets

b	m_b

$$\log \frac{d}{dX} [f(X)] = \frac{f'(X)}{f(X)}$$

$$\begin{aligned} \log \frac{d}{dX} [\prod_i (X - a_i)^{m_i}] \\ = \sum_i \frac{m_i}{X - a_i} \end{aligned}$$

Lookup Argument

a	m_a	$\frac{m_a}{x-a}$


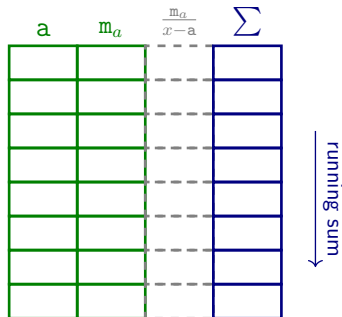
$a \equiv b$ as sets

b	m_b	$\frac{m_b}{x-b}$

$$\log \frac{d}{dX} [f(X)] = \frac{f'(X)}{f(X)}$$

$$\begin{aligned} \log \frac{d}{dX} [\prod_i (X - a_i)^{m_i}] \\ = \sum_i \frac{m_i}{X - a_i} \end{aligned}$$


Lookup Argument



$$\log \frac{d}{dX} [f(X)] = \frac{f'(X)}{f(X)}$$

$$\begin{aligned} \log \frac{d}{dX} [\prod_i (X - a_i)^{m_i}] \\ = \sum_i \frac{m_i}{X - a_i} \end{aligned}$$

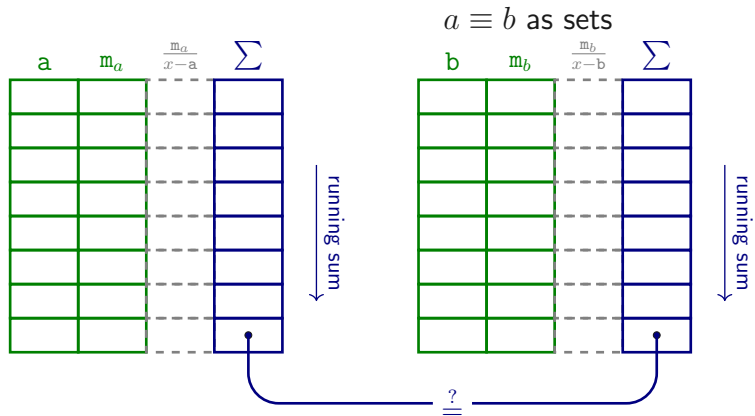
Lookup Argument



$$\log \frac{d}{dX} [f(X)] = \frac{f'(X)}{f(X)}$$

$$\begin{aligned} \log \frac{d}{dX} [\prod_i (X - a_i)^{m_i}] \\ = \sum_i \frac{m_i}{X - a_i} \end{aligned}$$

Lookup Argument



$$\log \frac{d}{dX} [f(X)] = \frac{f'(X)}{f(X)}$$

$$\begin{aligned} \log \frac{d}{dX} [\prod_i (X - a_i)^{m_i}] \\ = \sum_i \frac{m_i}{X - a_i} \end{aligned}$$

Soundness

$$\begin{aligned} & \Pr_x [S_a = S_b \mid a \not\equiv b] \\ &= \Pr_x [S_a \cdot \prod_i (x - a_i)^{m_{a,i}} \cdot \prod_i (x - b_i)^{m_{b,i}} = S_b \cdot \prod_i (x - a_i)^{m_{a,i}} \cdot \prod_i (x - b_i)^{m_{b,i}} \mid a \not\equiv b] \\ &\leq \frac{2N}{|\mathbb{F}|} \end{aligned}$$

Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory

Other Topics

Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture


Overview

Communication Arguments


Memory

Other Topics

Communication Lines Again



Memory — Problem Statement



Memory cells must have the same value as the previous time they were touched.



- ✓ random access
- ✓ read-write

Memory — Construction

Processor		
clk	ramp	ramv

RAM		
clk	ramp	ramv

Memory — Construction



Memory — Construction





1. same data, different order
2. within regions of constant ramp, correctly sorted by clk
3. correctly sorted by ramp

Memory — Construction



- memory integrity \Leftarrow {
1. same data, different order
 2. within regions of constant ramp, correctly sorted by clk
 3. correctly sorted by ramp

Memory — Construction




memory
integrity





1. same data, different order
2. within regions of constant ramp, correctly sorted by clk
3. ~~correctly sorted by ramp~~ regions of constant ramp are *contiguous*

Memory — Permutation




-
1. same data, different order
 2. within regions of constant `ramp`, correctly sorted by `clk`
 3. ~~correctly sorted by `ramp`~~ regions of constant `ramp` are *contiguous*

Memory — Permutation




-
1. same data, different order
 2. within regions of constant ramp, correctly sorted by clk
 3. ~~correctly sorted by ramp~~ regions of constant ramp are *contiguous*

Memory — Permutation



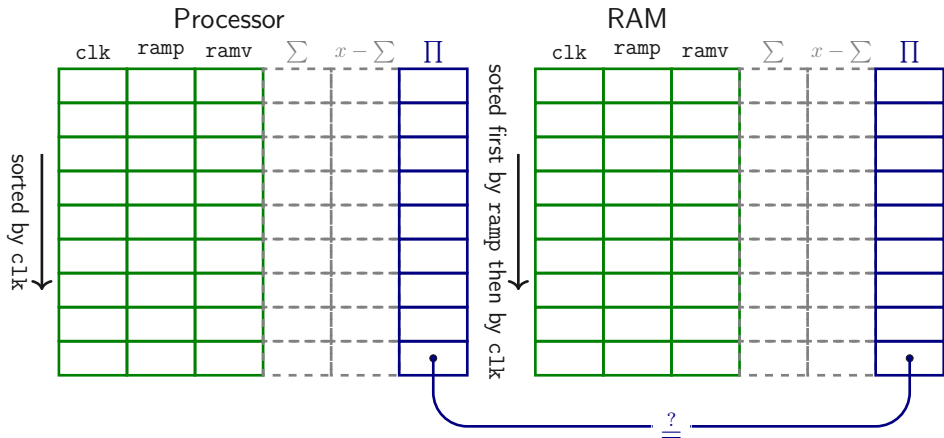
-
1. same data, different order
 2. within regions of constant $ramp$, correctly sorted by clk
 3. ~~correctly sorted by $ramp$~~ regions of constant $ramp$ are *contiguous*

Memory — Permutation





-
1. same data, different order
 2. within regions of constant $ramp$, correctly sorted by clk
 3. ~~correctly sorted by $ramp$~~ regions of constant $ramp$ are *contiguous*

Memory — Permutation




-
1. same data, different order
 2. within regions of constant $ramp$, correctly sorted by clk
 3. ~~correctly sorted by $ramp$~~ regions of constant $ramp$ are *contiguous*

Memory — Lookup



1. same data, different order
- 2. within regions of constant ramp, correctly sorted by clk
3. ~~correctly sorted by ramp~~ regions of constant ramp are *contiguous*

Memory — Lookup




1. same data, different order



2. within regions of constant $ramp$, correctly sorted by clk

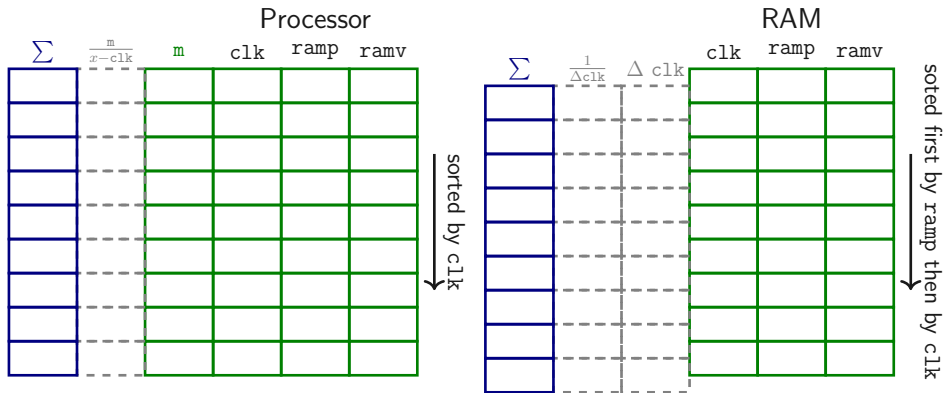
3. ~~correctly sorted by $ramp$~~ regions of constant $ramp$ are *contiguous*

Memory — Lookup



1. same data, different order
- 2. within regions of constant ramp, correctly sorted by clk
3. ~~correctly sorted by ramp~~ regions of constant ramp are *contiguous*

Memory — Lookup



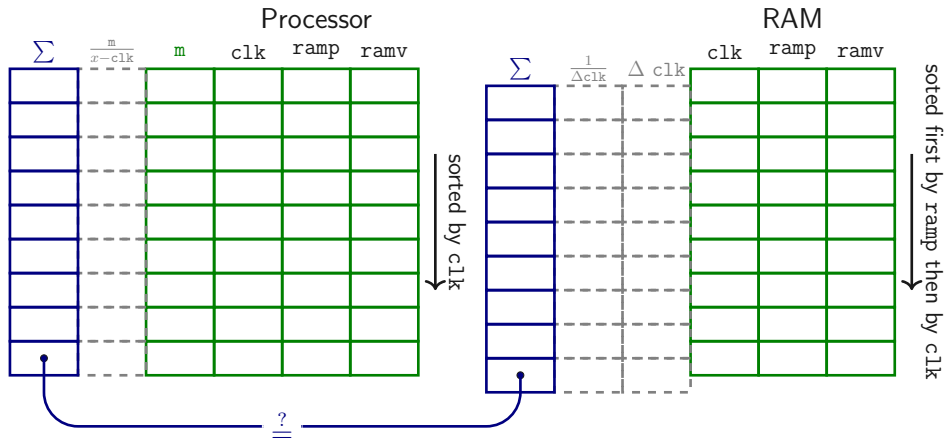
1. same data, different order



2. within regions of constant ramp, correctly sorted by clk

3. ~~correctly sorted by ramp~~ regions of constant ramp are *contiguous*

Memory — Lookup




1. same data, different order



2. within regions of constant ramp, correctly sorted by clk

3. ~~correctly sorted by ramp~~ regions of constant ramp are *contiguous*


Memory — Contiguity



1. same data, different order
2. within regions of constant ramp, correctly sorted by clk

→ 3. ~~correctly sorted by ramp~~ regions of constant ramp are *contiguous*

Memory — Contiguity




1. same data, different order

2. within regions of constant ramp, correctly sorted by clk

→ 3. ~~correctly sorted by ramp~~ regions of constant ramp are *contiguous*

Memory — Contiguity





1. same data, different order
2. within regions of constant ramp, correctly sorted by clk

→ 3. ~~correctly sorted by ramp~~ regions of constant ramp are *contiguous*

Memory — Contiguity

\prod accumulates one factor $X - \text{ramp}$
whenever $\text{ramp} \neq \text{ramp}^*$





1. same data, different order
2. within regions of constant ramp, correctly sorted by clk

→ 3. ~~correctly sorted by ramp~~ regions of constant ramp are *contiguous*

Memory — Contiguity

Π accumulates one factor $X - \text{ramp}$
whenever $\text{ramp} \neq \text{ramp}^*$





1. same data, different order
2. within regions of constant ramp, correctly sorted by clk

→ 3. ~~correctly sorted by ramp~~ regions of constant ramp are *contiguous*

Memory — Contiguity

\prod accumulates one factor $X - \text{ramp}$
whenever $\text{ramp} \neq \text{ramp}^*$



$$a(x) \cdot f(x) + b(x) \cdot f'(x) \stackrel{?}{=} 1$$

not contiguous

\Rightarrow repeated factors

$\Rightarrow \text{gcd} \neq 1$

1. same data, different order

2. within regions of constant ramp, correctly sorted by clk

3. ~~correctly sorted by ramp~~ regions of constant ramp are *contiguous*

Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory

Other Topics

Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory

Other Topics

Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview


Communication Arguments

Memory


Other Topics

Univariate versus Multilinear

Univariate




Multilinear




Univariate versus Multilinear

Univariate



DEEP-ALI


Multilinear



GKR


Univariate versus Multilinear

Univariate



DEEP-ALI

Multilinear



GKR

WHIR / GKR-logup




STIR




- asymptotically and concretely better than FRI
- simpler proof of soundness
- paves way for aggregation

ECFFT



Structured Fields



Arbitrary Fields

CircleSTARK

Table of Contents

Retrospective

Optimizations

Batching

Quotient Segmentation

Grinding

Enhancements

Zero-Knowledge

Randomized AIR with Preprocessing

VM Architecture

Overview

Communication Arguments

Memory

Other Topics

Table of Contents

- Retrospective

- Optimizations

 - Batching

 - Quotient Segmentation

 - Grinding

- Enhancements

 - Zero-Knowledge

 - Randomized AIR with Preprocessing

- VM Architecture

 - Overview

 - Communication Arguments

 - Memory

- Other Topics

Advanced zk-STARKs

Alan Szepieniec

艾伦·余丕涅茨

`alan@neptune.cash`



neptune

<https://neptune.cash/>



Triton VM

<https://triton-vm.org/>

<https://asz.ink/presentations/2025-09-18-Advanced-zkSTARKs.pdf>