

# Cirrus: Performant and Accountable Distributed SNARK\*

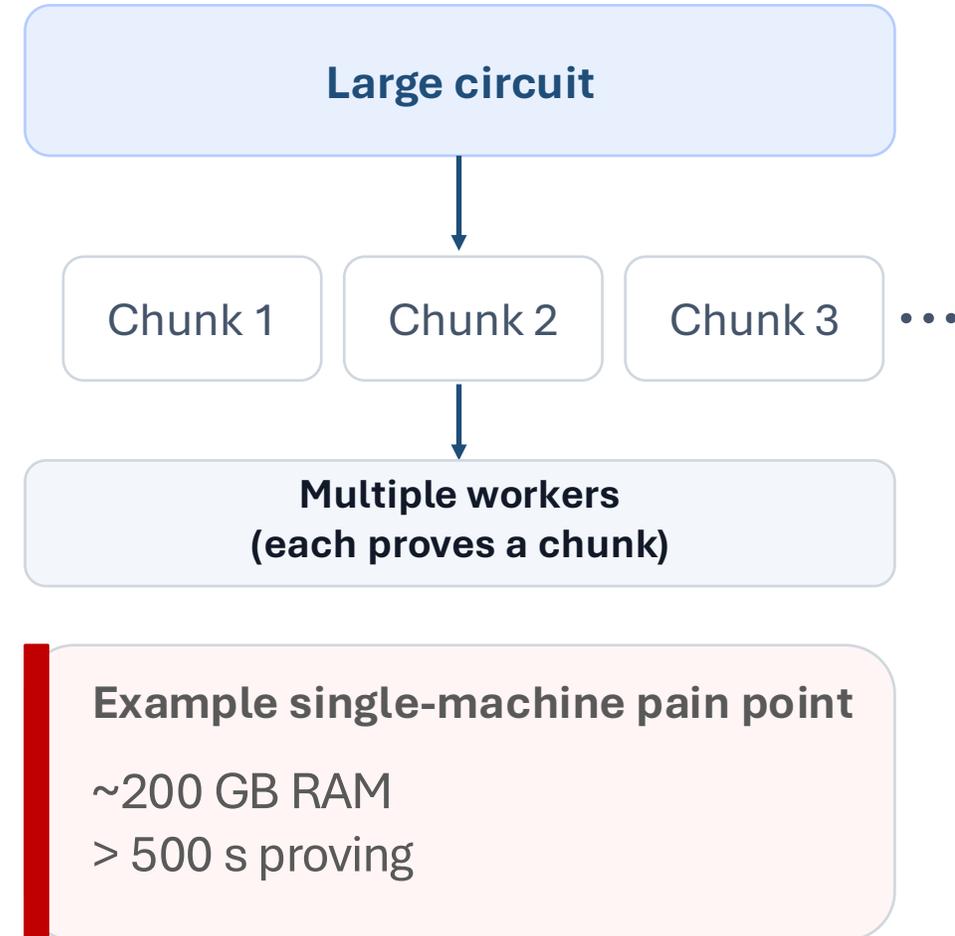
Wenhao Wang<sup>1,4</sup>, Fangyan Shi<sup>2</sup>, Dani Vilardell<sup>3,4</sup>, Fan Zhang<sup>1,4</sup>

<sup>1</sup>Yale University, <sup>2</sup>Tsinghua University, <sup>3</sup>Cornell University, <sup>4</sup>IC3



# SNARKs call for distribution

- Large real-world SNARK applications produce massive circuits
  - Proving becomes a bottleneck
- Single-machine proving hits **time** & **memory** limits
  - High memory footprint
  - Long proving time
- Idea: split the proving workload across many workers (horizontal scaling)



# What makes distributed SNARKs hard?

Not just “distributed”, but distributed well

## (1) Low overhead

- Linear scalability: Worker compute  $\propto$  circuit chunk size
- Low coordination cost
- Low communication per worker (sublinear in circuit chunk size)

## (2) Accountability

- Detect faults when proof fails
- Locate faulty workers
- Enables incentives (pay / slash) in prover markets
- A property that most prior works omit

## (3) Universal setup

- Avoid per-circuit trusted setup ceremonies
- Support multiple applications & varying worker counts at once
- Critical for real deployments and outsourcing

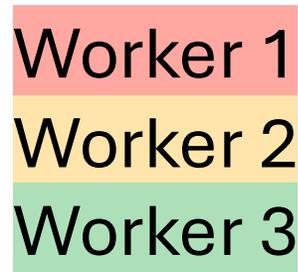
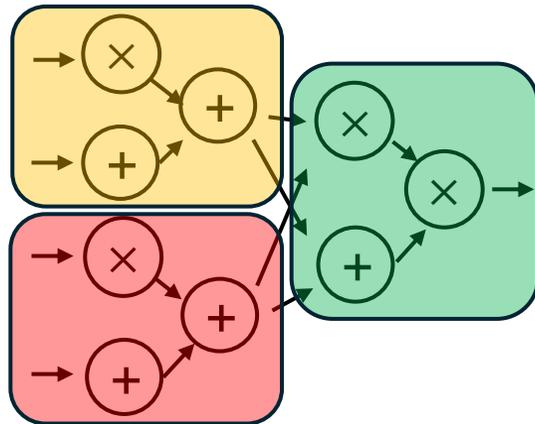
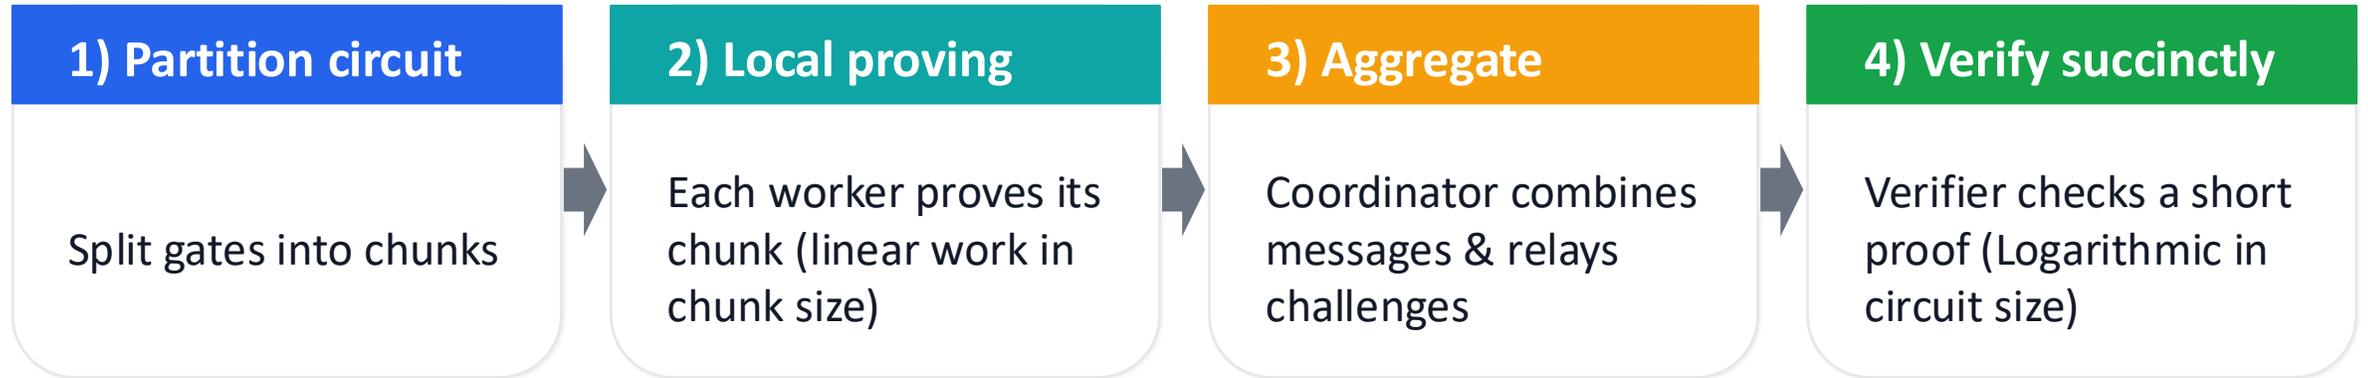
# Prior work: no one got all 3 properties

- Existing systems sacrifice at least one of scalability, accountability, or universal setup

Scheme	Overhead	Accountable?	Setup
DIZK	High Comm.	● NO	Circuit-specific
deVirgo	High Comm.	● NO	Transparent
Pianist	Low	● NO*	Universal
HyperPianist	Low	● NO	Universal
Hekaton	Low	● YES	Circuit-specific
<b>Cirrus</b>	<b>Low</b>	<b>● YES</b>	<b>Universal</b>

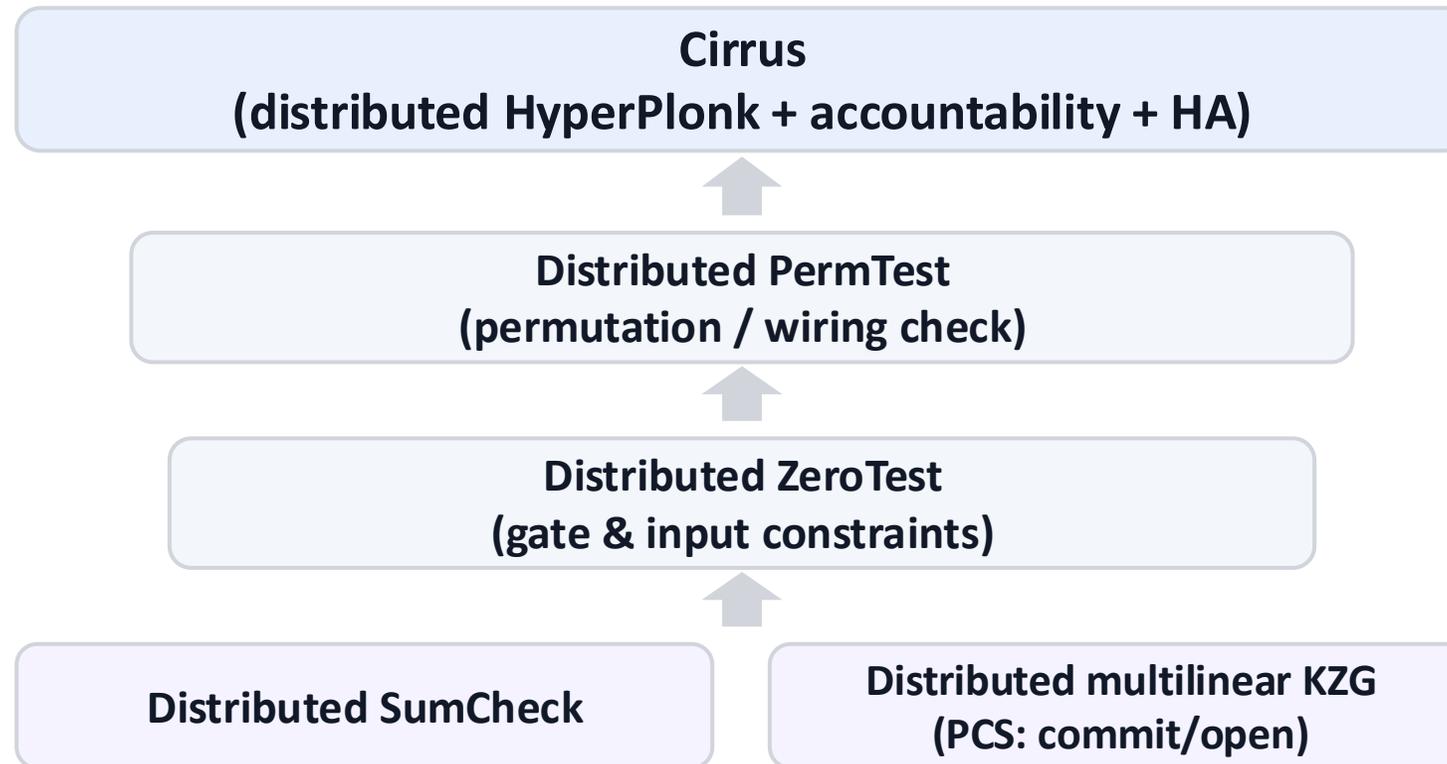
\* Pianist offers limited accountability for data-parallel circuits only.

# Cirrus: the high-level workflow



**Goal: linear work per worker + small communication, while retaining universal setup**

# Components of Cirrus



Built on HyperPlonk's universal setup; designed for scalable proving across many workers

# Distributed Multilinear KZG

- **Multilinear KZG (PCS):** commit to a multilinear polynomial and open it at one point succinctly

$$f(\alpha, \beta) = z$$

## Commit (distributed)

- Worker  $\mathcal{P}_b$  commits to its local chunk polynomial  $f^{(b)}(\mathbf{x})$
- Coordinator multiplies commitments to get a single commitment  $\text{com}_b$

$$\text{com} = \prod \text{com}_b$$

- Reuses HyperPlonk universal CRS

## Open at $(\alpha, \beta)$ (distributed)

- Worker computes local  $z_b = f_b(\alpha)$  and opening proof  $\pi_b$
- Coordinator aggregates  $(z_b, \pi_b)$  into  $(z, \pi)$

$$z = \sum z_b \quad \pi = \left( \underbrace{\prod \pi_b}_{\alpha \text{ coordinates}}, \underbrace{\pi'}_{\beta \text{ coordinates}} \right)$$

- Verifier checks a constant-size pairing equation

Each worker sends only  $O(\log T)$  group elements per opening; coordinator just aggregates.

# Distributed SumCheck

- **SumCheck:** Prove a sum over the Boolean hypercube without sending all evaluations

$$\sum_{z \in \{0,1\}^n} f(z) = v$$

## Standard SumCheck (one variable per round)

$$r_i(X) = \sum_{w \in \{0,1\}^{n-i}} f(\alpha_1, \dots, \alpha_{i-1}, X, w)$$

- Verifier checks:  $v_{i-1} = r_i(0) + r_i(1)$
- Verifier samples random challenge  $\alpha_i$
- Next target becomes  $v_i = r_i(\alpha_i)$
- At the end, open  $f$  at one random point

## Cirrus: distribute it across workers

$$r_i(X) = \sum_{b \in \{0,1\}^\xi} r_i^{(b)}(X)$$

- Worker  $b$  computes its local  $r_i^b(X)$  from its chunk
- Coordinator adds them up to form global  $r_i(X)$
- This repeats for  $\log T$  rounds ( $T = \text{chunk size}$ )
- At the end, workers open polynomials at the same random point

# Reducing polynomial degree in SumCheck

- Problem with distributed polynomial opening: only supports degree-1 polynomials

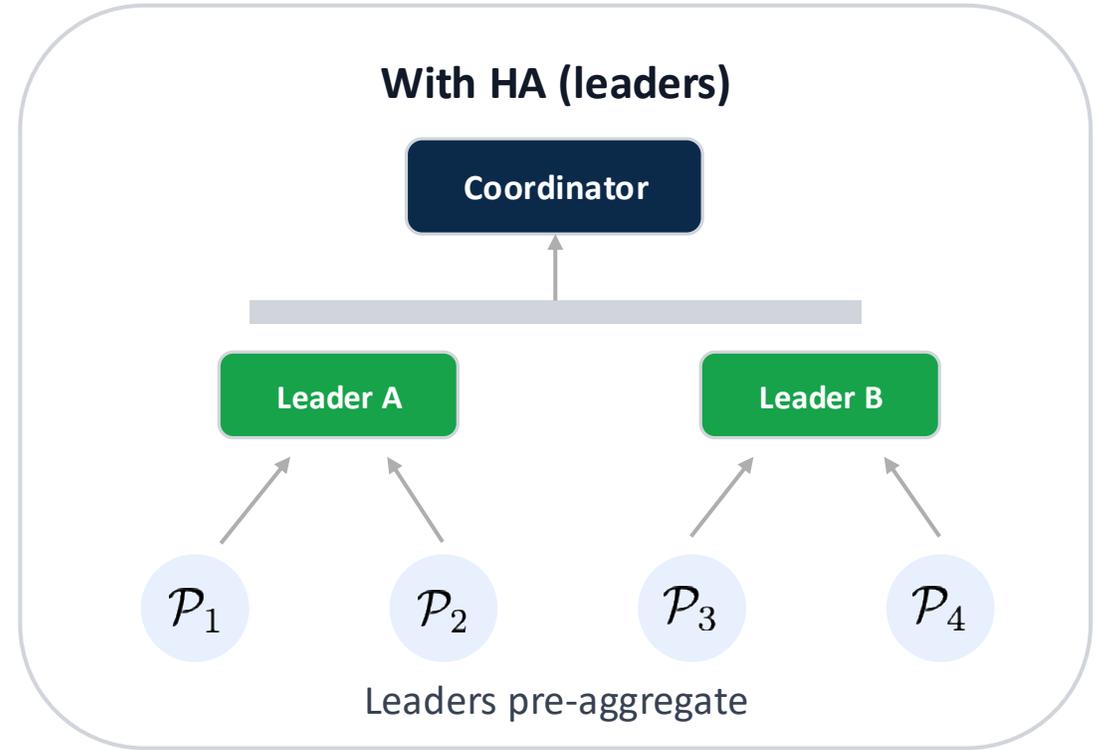
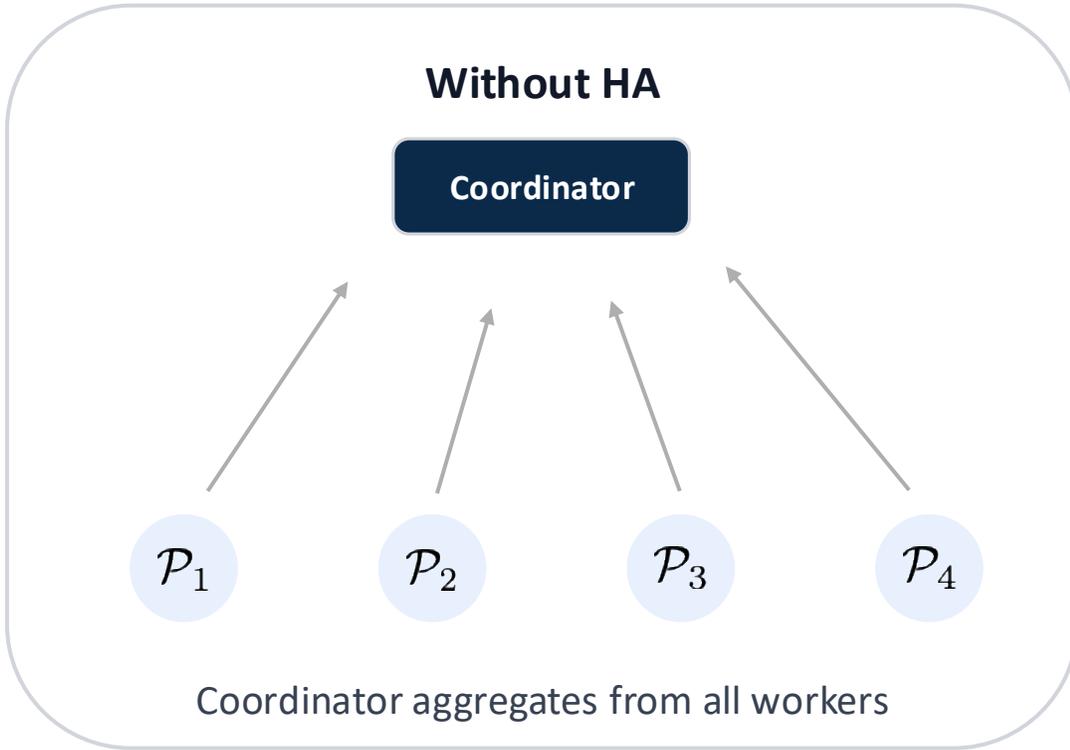
Our compatibility trick: reduce the higher-degree polynomial into a small number of multilinear polynomials, so multilinear KZG openings work cleanly

Given a high-degree polynomial  $h$ , we can split it into multiple linear polynomials, and then commit to each linear polynomial instead

$$\begin{aligned} \forall \mathbf{y} \in B_\xi, \quad \tilde{f}(\mathbf{y}) &= h(\tilde{g}_1(\mathbf{y}), \dots, \tilde{g}_c(\mathbf{y})) \\ &\equiv \sum_{\mathbf{b} \in B_\xi} h(g_1^{(\mathbf{b})}(\mathbf{y}), \dots, g_c^{(\mathbf{b})}(\mathbf{y})) \cdot \chi_{\mathbf{b}}(\mathbf{y}) = f(\alpha, \mathbf{y}) \end{aligned}$$

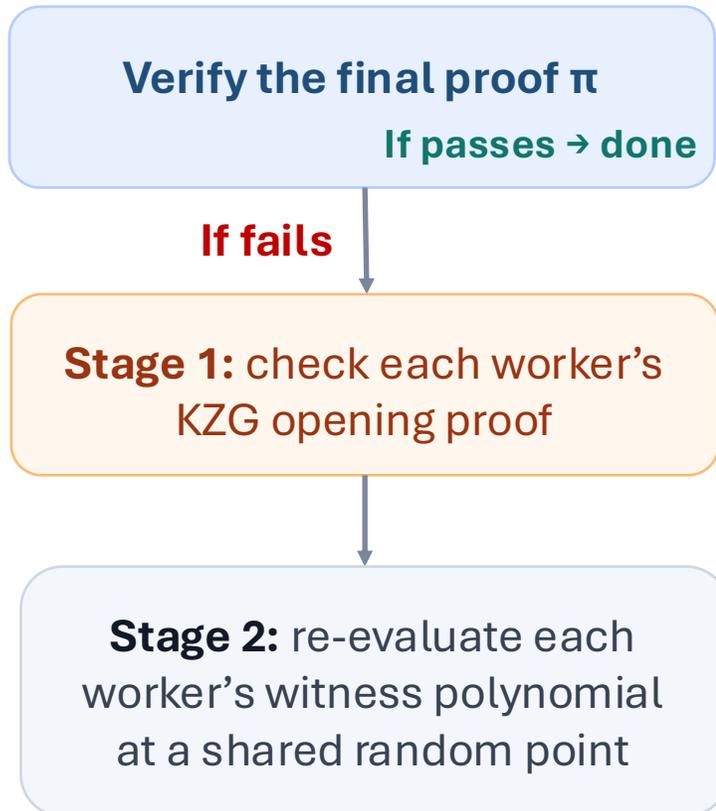
The equality showing that polynomial splitting does not break equality when evaluated on the Boolean hypercube

# Hierarchical Aggregation (HA)



Saves coordinator fan-in and aggregation work: from  $O(M \cdot \log T) \rightarrow O(M)$   
(Coordinator runtime independent of sub-circuit size)

# Accountability with low overhead



## Optimistic idea:

- If the final proof verifies  $\rightarrow$  All good
- Only localize faults when verification fails
- Goal: Make the accountability protocol fast
  - Naïve approach: redo all workers' computation

## Our solution: 2-Stage localization

- Observation: Verifying a few intermediate results suffices; no need to redo all work.
- Result: With just 2 checks, the coordinator detects faulty workers

# End-to-end scalability

HyperPlonk baseline: proves 4M gates in more than 100s (runs out of memory beyond that)

**33M-gate circuits proved in  $\approx 39$  seconds**

using 32 machines  $\times$  8 cores (AWS t3.2xlarge)

Per-worker communication

$\approx$  **40 KB**

Proof size

$\approx$  **17 KB**

Accountability

**< 4 s**

## Why distribution helps

- Cirrus keeps per-worker memory bounded by the subcircuit size  $T$
- Low coordination overhead  $\rightarrow$  near-linear scaling with more workers for large circuits

# Comparison with Hekaton

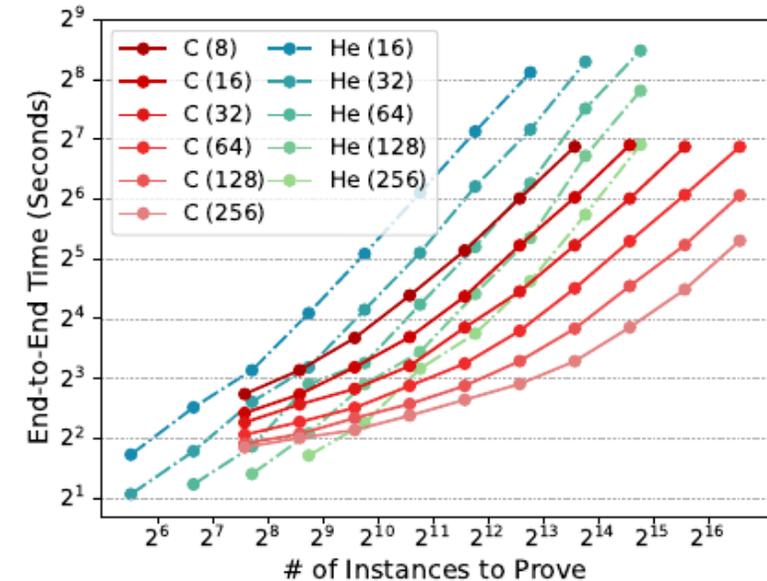
## Cirrus

PLONK • Universal setup

## Hekaton

R1CS • Circuit-specific setup

- **Apples-to-oranges:** different arithmetizations
- **But on PLONK-friendly tasks** Cirrus is faster
- **Pedersen hash circuits** > 7× faster proof generation



(a) Comparison between Cirrus and Hekaton to prove Pedersen

Fig. 3(a) from Cirrus paper: Pedersen hash benchmark  
(C = Cirrus, He = Hekaton; cores in parentheses)

# Takeaways

- Why this matters
  - Unlocks larger ZK workloads: zkML, zkVMs
  - Enables outsourced proving
- What Cirrus achieves
  - ✓ Linear-time workers and lightweight coordinator
  - ✓ Fast accountability (localize faulty workers)
  - ✓ Universal trusted setup

Personal Website: [wenhao-wang0.github.io](https://wenhao-wang0.github.io)



Link to full paper



Link to personal website

Thank you!