

New Modes on the Block: Security and Efficiency of Novel AO Compression Modes

ZKSC 2026, Vienna

Stefano Trevisani

Joint work with: E. Andreeva, R. Bhattacharyya, A. Roy



Hash functions and ZK-SNARKS

Cryptographic Hash Functions play a central role in SNARKs:

- In the Fiat-Shamir transform, enabling the 'N' in SNARK
- In the **Merkle Tree** commitment scheme:
 - Used in the ZK-STARK/FRI PCS.
 - In recursive SNARKs and IVC protocols.

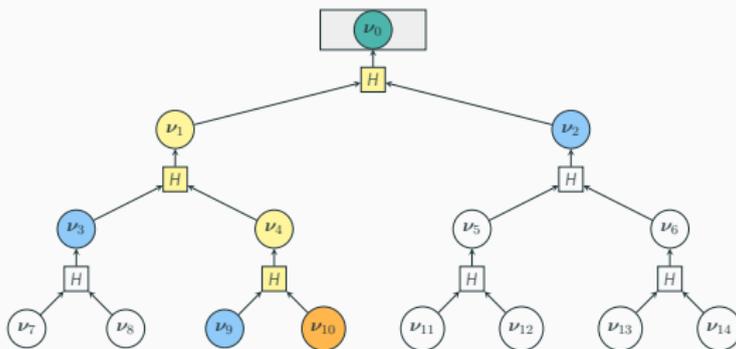
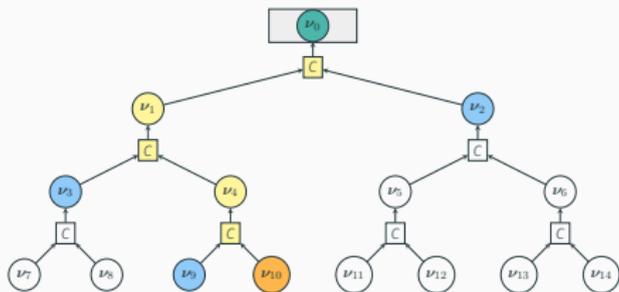


Figure 1: A binary Merkle Tree. Highlighted is an authentication path.

Merkle Trees Commitments

Merkle Trees used to instantiate a **vector commitment scheme**:

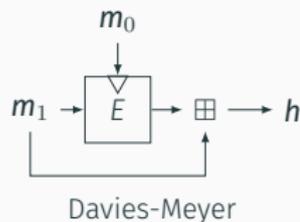
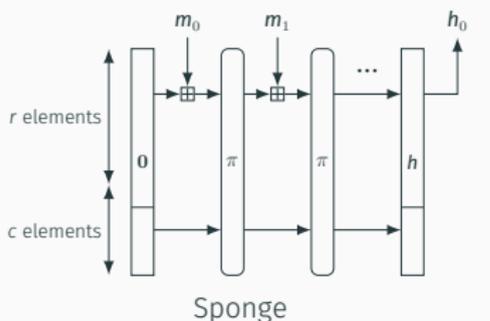
- The internal hash function is a $t : 1$ compression function.
 - Usually $t = 2$.
- $\text{Com}(\mathbf{m})$:
 - Split \mathbf{m} into leaf nodes $m_1 = \nu_{2^h-1}, \dots, m_{2^h} = \nu_{2^{h+1}-2}$.
 - Compute and output commitment $\mathbf{c} = \nu_0$.
- $\text{Open}(i)$: output m_i 's co-path to the root.
 - Importantly: opening ρ is logarithmic in $|\mathbf{m}|$.
- $\text{Ver}(\mathbf{c}, i, m_i, \rho)$: compute \mathbf{c}' from (i, m_i, ρ) , accept iff $\mathbf{c}' = \mathbf{c}$.



Compression Modes in the Wild

Which compression modes are used in practice?

- Blockcipher-based modes
 - PGV modes (e.g. Davies-Meyer, SHA-2).
- Permutation-based modes
 - Sponge (e.g. SHA-3).
 - Provably secure in the ideal-permutation model.
 - **Jive** (proposed with **Anemoi**).
 - **Trunc** (proposed with GRIFFIN, POSEIDON2, ...).
 - No provable security analysis.



Why Sponge?

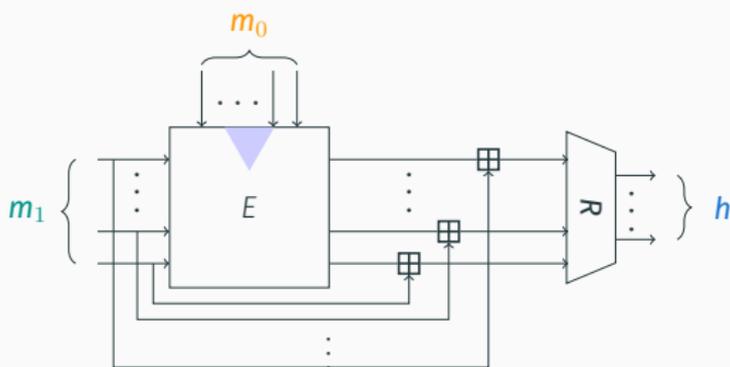
Sponge mode:

- RO-indifferentiable in the ideal permutation model.
- Single-round sponge (sponge-1):
 - CICO-1, CICO- k , CICO- (k_1, k_2) ...
- Limitations of Sponge:
 - Suboptimal security bounds (especially collision/preimage resistance).
 - Most efficient?
- Many AO permutation (POSEIDON- π , *Rescue*) are blockcipher-based.
 - Construct modes directly from the block cipher.
 - Better permutation-based modes.

The PGV-LC modes

At CSF'24, we introduced the PGV-LC **modes**:

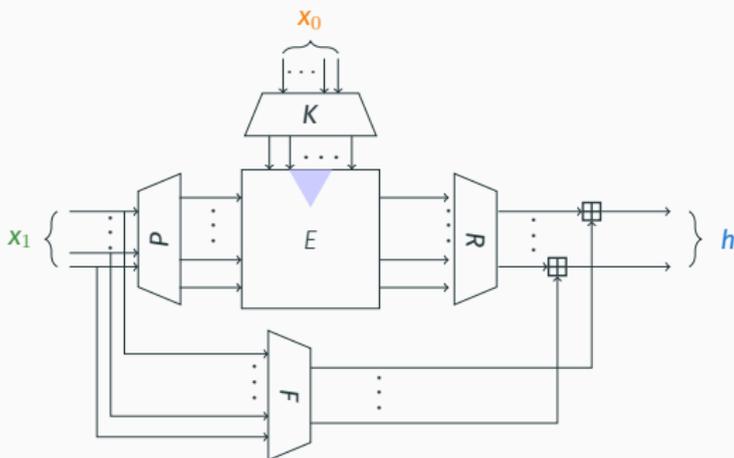
- Uses a blockcipher $E: \mathbb{F}_p^\kappa \times \mathbb{F}_p^n \rightarrow \mathbb{F}_p^n$.
- Matrix $R \in \mathbb{F}_p^{\ell \times n}$ parametrizes output size.
 - ◇ Compresses its input $\Rightarrow \ell \leq n$.
 - ◇ Algebraic generalization of e.g. truncation and chopping.



The PGV-ELC mode

And an immediate extension thereof, called PGV-ELC:

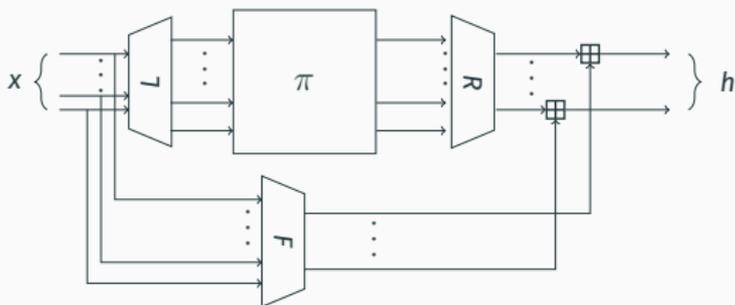
- Matrices $\mathbf{K} \in \mathbb{F}_p^{\kappa \times \kappa'}$ and $\mathbf{P} \in \mathbb{F}_p^{n \times n'}$ parametrize input size.
 - ◊ Expand their inputs $\Rightarrow \kappa' \leq \kappa$ and $n' \leq n$.
 - ◊ Algebraic generalization of e.g. zero-padding.
- Matrix $\mathbf{F} \in \mathbb{F}_p^{\ell \times n'}$ adapts input to output size.



The PAX modes

We are also proposing the PAX family of modes [ongoing work]:

- Based on a permutation $\pi: \mathbb{F}_p^n \rightarrow \mathbb{F}_p^n$.
- $\mathcal{C}_{L,F,R,\pi}$ maps $x \in \mathbb{F}_p^m$ to $h \in \mathbb{F}_p^\ell$, with $\ell \leq m$.
- Expansion matrix $L \in \mathbb{F}_p^{n \times m}$.
- Compression matrices $F \in \mathbb{F}_p^{\ell \times m}$ and $R \in \mathbb{F}_p^{\ell \times n}$.



PAX as a Generalization of Existing Modes

The PAX family generalizes most of the used constructions:

- **Trunc** $C(\mathbf{x}) = \lfloor \pi(\mathbf{x}) + \mathbf{x} \rfloor_\ell$:
 - Fix $n = m$, L, F, R are all (pseudo-)identities.
- **Jive** $C(\mathbf{x}) = \sum_i \mathbf{x}_i + \sum_i \mathbf{y}_i$, with $\mathbf{y} = \pi(\mathbf{x})$:
 - Fix $n = m$, L is the identity, F and R are circulant matrices on the repeated i th canonical vector of \mathbb{F}_p^ℓ .
- **First Sponge round**: $C(\mathbf{x}) = \lfloor \pi(\mathbf{x} \parallel \mathbf{0}) \rfloor_\ell$
 - L, R pseudo-identities, $R = \mathbf{0}$.
 - Note: R is (clearly) not right-invertible.

Black-box security of cryptographic constructions:

- Well-established model.
- Some primitives $\mathcal{P}_1, \dots, \mathcal{P}_n$ used in a *mode*.
 - Often one primitive \mathcal{P} .
 - \mathcal{P} is *ideal* (randomly sampled).
 - An adversary \mathcal{A} can query the primitive(s).
- Security of a mode with respect to some property is expressed via an advantage function $\mathbf{Adv}(\mathcal{A}, q)$.

Compression Modes Security

For a compression mode \mathcal{C} :

- Often \mathcal{P} is a blockcipher or a permutation.
 - Say over alphabet \mathbb{F}_p .
- Adversary makes q queries to \mathcal{P} .
- Some standard security properties:

- Collision resistance:

$$\mathbf{Adv}_{\mathcal{C}}^{\text{COL}}(\mathcal{A}, q) = \Pr[\mathcal{C}^{\mathcal{P}}(x) = \mathcal{C}^{\mathcal{P}}(x') \wedge x \neq x' \mid (x, x') \leftarrow \mathcal{A}^{\mathcal{P}}()]$$

- Preimage resistance:

$$\mathbf{Adv}_{\mathcal{C}}^{\text{PRE}}(\mathcal{A}, q) = \Pr[\mathcal{C}^{\mathcal{P}}(x) = \mathcal{C}^{\mathcal{P}}(x') \mid x \xleftarrow{\$} M; x' \leftarrow \mathcal{A}^{\mathcal{P}}(\mathcal{C}^{\mathcal{P}}(x))]$$

- Second-Preimage resistance (implied by collision resistance):

$$\mathbf{Adv}_{\mathcal{C}}^{\text{PRE}^2}(\mathcal{A}, q) = \Pr[\mathcal{C}^{\mathcal{P}}(x) = \mathcal{C}^{\mathcal{P}}(x') \wedge x \neq x' \mid x \xleftarrow{\$} M; x' \leftarrow \mathcal{A}^{\mathcal{P}}(x)]$$

Random-Oracle Indifferentiability

Random-oracle indifferentiability is the ideal notion:

- Let \mathcal{H} be a random oracle with compatible domain and range.
- Let \mathcal{S} be a simulator algorithm
 - Makes $q_{\mathcal{S}}$ queries to \mathcal{H} , with $q_{\mathcal{S}}$ small.
- Let \mathcal{D} be a differentiator:
 - Makes $q_{\mathcal{P}}$ queries to its left interface.
 - Makes $q_{\mathcal{H}}$ queries to its right interface.
- Then, letting $q = q_{\mathcal{P}} + q_{\mathcal{H}}$:

$$\mathbf{Adv}_{\mathcal{C}, \mathcal{S}}^{\text{INDIF}}(\mathcal{D}, q) = \left| \Pr \left[\mathcal{D}^{\mathcal{P}, \mathcal{C}^{\mathcal{P}}}() = \top \right] - \Pr \left[\mathcal{D}^{\mathcal{S}^{\mathcal{H}}, \mathcal{H}}() = \top \right] \right|$$

Merkle Tree Security

Which properties do we need for Merkle Tree commitments?

- Tree \mathcal{T} using $C = \mathcal{C}^{\mathcal{P}}$ as the internal compression function.
 - \mathcal{T}^C is also a compression function.
- Local opening security:

$$\mathbf{Adv}_{\mathcal{T}}^{\text{OPEN}}(\mathcal{A}, q) = \Pr\left[\text{Ver}\left(\mathcal{T}^C(m), i, m'_i, \rho'\right) = \top \wedge m'_i \neq m_i \mid (m, i, m'_i, \rho') \leftarrow \mathcal{A}^{\mathcal{P}}()\right]$$

- It can be shown that:

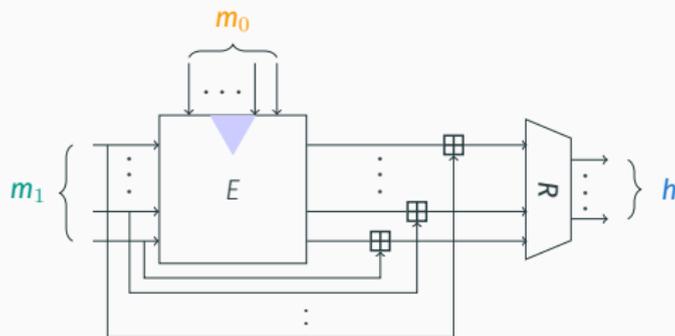
$$\mathbf{Adv}_{\mathcal{T}}^{\text{OPEN}}(q) \leq \mathbf{Adv}_{\mathcal{C}}^{\text{COL}}(q)^1$$

- Indifferentiability can however still be valuable.

¹Or $2q$ in a slightly different notion.

Collision resistance of PGV-LC:

1. Consider R right-invertible (full row rank).
2. \mathbb{F}_p^n is partitioned into p^ℓ equivalence classes.
3. \mathcal{A} can exploit partition unbalances from oracle replies.
4. Still, $\mathbf{Adv}_C^{\text{COL}}(q) \leq \frac{q^2+q}{p^\ell-q}$ (\approx birthday attack).
5. Similar reasoning for preimages: $\mathbf{Adv}_C^{\text{PRE}}(q) \leq \frac{q}{p^\ell-q}$.



Collision resistance of PGV-ELC:

1. Consider K and P left-invertible, F right-invertible.
 - Also induce partitions of respective spaces.
2. 'Meaningless' queries, can be exploited only indirectly.
3. Again, $\mathbf{Adv}_C^{\text{COL}}(q) \leq \frac{q^2+q}{p^\ell-q}$.
4. And for preimage resistance: $\mathbf{Adv}_C^{\text{PRE}}(q) \leq \frac{q}{p^\ell-q}$.

Security of PAX

Indifferentiability of PAX based on the following simulator:

```
function  $\mathcal{S}^{\mathcal{H}}(b \in \{-1, 1\}, z \in \mathbb{F}_p^n) \rightarrow \mathbb{F}_p^n$   
  static  $T \leftarrow \emptyset$   
  if  $b = 1$  then ▷ Forward query  
    if  $z \in \text{Dom}(T)$  then return  $T(z)$   
     $v \leftarrow L^+z$   
    if  $z = Lv$  then ▷ Queried point has a preimage  
       $h \leftarrow \mathcal{H}(v)$   
       $y \xleftarrow{\$} [h - Fv]_R \setminus \text{Ran}(T)$  ▷  $[x]_R$  equiv. class of  $x$   
    else ▷ Queried point is preimage-free  
       $y \xleftarrow{\$} \mathbb{F}_p^n \setminus \text{Ran}(T)$   
       $T \leftarrow T \cup \{(z, y)\}$   
      return  $y$   
  else ▷ Backward query  
    if  $z \in \text{Ran}(T)$  then return  $T^{-1}(z)$   
     $x \xleftarrow{\$} \mathbb{F}_p^n \setminus \text{Dom}(T)$   
     $T \leftarrow T \cup \{(x, z)\}$   
    return  $x$ 
```

Modes security comparison

Adversarial advantages for a compression function $\mathbb{F}_p^m \rightarrow \mathbb{F}_p^\ell$:

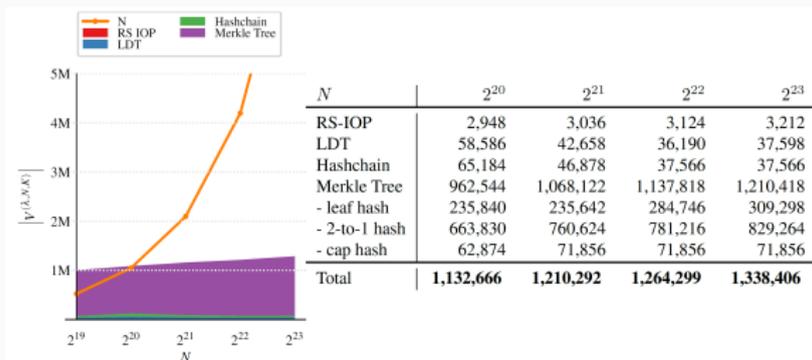
Mode	Primitive	COL	PRE	DIF
PGV-ELC	Block($m/2 + c, m/2 + c$)	q^2/p^ℓ	q/p^ℓ	q/p^c
PAX	Perm($m + c$)	q^2/p^ℓ	q/p^ℓ	q/p^c
Sponge-1	Perm($m + c$)	$\min\{q^2/p^\ell, q/p^c\}$	$q/p^{\min\{\ell, c\}}$	q/p^c
Sponge	Perm($r + c$)	$q^2/p^{\min\{\ell, c\}}$	$q/p^{\min\{\ell, c\}}$	q^2/p^c

- PAX and PGV-LC have optimal COL and PRE resistance.
- Single-iteration sponge (Sponge-1) has better bounds than multi-iteration.
- COL and PRE resistance of Sponge are sub-optimal.
- PGV-ELC indifferentiability does not depend on the key 'capacity'.

Impact of Hashing

How much is hashing relevant in a SNARK computation?

- In FRACTAL's verifier circuit [6], about 90% of constraints are from hash functions.
- In the RISC Zero chess example [13], about 37% of total time is spent on hashing.
- *Circuit* and *plain* performance are both important.
 - ZKProofMarket: by 2030 2^{27} proofs per day.
 - Even a 1ms saving per proof: $\approx 2^{24}$ seconds per day.



Caveat: Heuristic Security

Instantiating a compression functions:

- Target a specific security level (e.g. 128-bits for collision resistance).
- Choose a concrete underlying primitive:
 - parametrization based on cryptanalytical bounds.
- For PGV-ELC modes, this is not completely clear.
 - Need for cryptanalysis on AO blockcipher-based compression.
 - Experiments still relevant to establish efficiency margin.
- For PAX modes, results about previous permutation-based modes apply.

Arithmetization-Oriented Hash Functions

How to address the bottleneck?

- Plain performance: traditional solutions (e.g. SHA) work well.
- Circuit performance:
 - Arithmetic circuit over a field \mathbb{F}_{p^d} .
 - In SNARKs that rely on pairing-friendly EC: $p \approx 2^{256}$, $d = 1$.
 - For STARKs that rely on FRI IOPP: $p \approx 2^{32}$, $p \approx 2^{64}$, $d = 1$.
 - Multiplicative complexity/depth is an important starting point.
 - But cost is really tied to the *arithmetization technique*.
- *Arithmetization-Oriented* hash functions.

Arithmetization Techniques

Which arithmetization technique?

- R1CS arithmetization:

$$Ax \odot Bx = Cx$$

- (basic) \mathcal{PLonK} arithmetization:

$$\{q_x X + q_y Y + q_z Z + q_{xy} XY + q_c = 0\}$$

- AIR arithmetization:

$$T(\mathbf{x}_{\text{in}}, \mathbf{x}_{\text{out}}) \quad \{B(\mathbf{x})\}$$

- Lookup arguments: efficient range-checks and more.
- Traditional solutions become arithmetization-oriented?

PGV-LC Proof Generation

Time to generate a MT opening proof on Groth16:

- Scalar field of the BLS12-381 elliptic curve: $\log_2(p) \approx 255$.

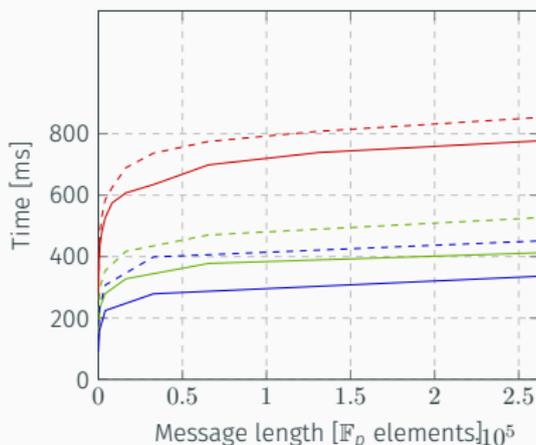


Figure 3: Dashed: POSEIDON, Solid: HADES with PGV-LC

PGV-LC Root Computation

Time to compute the root of a Merkle Tree:

- Scalar field of the BLS12-381 elliptic curve: $\log_2(p) \approx 255$.

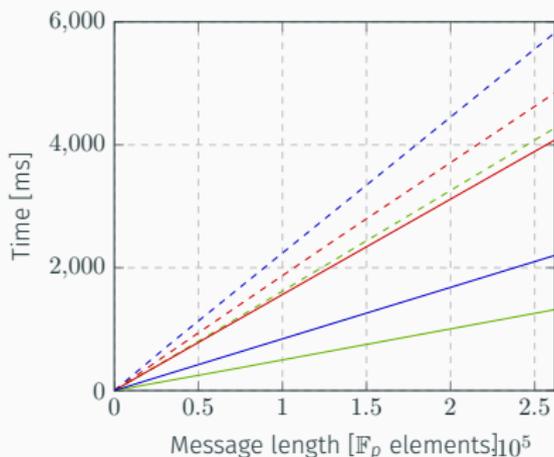


Figure 4: Dashed: POSEIDON, Solid: HADES with PGV-LC

Plain performance

Comp. ratio	BN254		Goldilocks	
	PAX	Sponge	PAX	Sponge
<i>POSEIDON-π</i>				
2:1	9.4 μ s	17.5 μ s	3.5 μ s	7.3 μ s
4:1	28.2 μ s	42.4 μ s	12.2 μ s	18.6 μ s
8:1	103.7 μ s	130.4 μ s	141.5 μ s	265.4 μ s
<i>Rescue</i>				
2:1	334.6 μ s	353.2 μ s	16.8 μ s	26.2 μ s
4:1	372.0 μ s	382.9 μ s	36.2 μ s	46.9 μ s
8:1	553.5 μ s	626.1 μ s	117.7 μ s	153.5 μ s

Groth16 benchmarks

Preimage circuit in Groth16:

Comp. ratio	# R1CS constraints		Proof Generation time	
	PAX	Sponge	PAX	Sponge
<i>POSEIDON-π</i>				
2 : 1	221	246	44.5 ms	47.3 ms
4 : 1	268	293	51.3 ms	54.9 ms
8 : 1	368	393	66.3 ms	71.5 ms
<i>Rescue</i>				
2 : 1	240	252	41.2 ms	42.3 ms
4 : 1	264	270	44.9 ms	45.1 ms
8 : 1	384	432	63.8 ms	67.7 ms

Preimage circuit in *Plonky2*:

Comp. ratio	# gates		Proof Generation time	
	PAX	Sponge	PAX	Sponge
<i>POSEIDON-π</i>				
2:1	122	259	11.3 ms	16.5 ms
4:1	439	668	26.0 ms	27.1 ms
8:1	2065	2864	90.8 ms	92.9 ms
<i>Rescue</i>				
2:1	91	175	10.9 ms	17.2 ms
4:1	284	418	16.8 ms	27.1 ms
8:1	976	1213	47.9 ms	49.0 ms

Fin

Thank you for your attention!
Any questions?



Martin Albrecht, Lorenzo Grassi, Christian Rechberger, Arnab Roy, and Tyge Tiessen.

Mimc: Efficient encryption and cryptographic hashing with minimal multiplicative complexity.

In Jung Hee Cheon and Tsuyoshi Takagi, editors, *Advances in Cryptology – ASIACRYPT 2016*, pages 191–219, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.



Elena Andreeva, Rishiraj Bhattacharyya, Arnab Roy, and Stefano Trevisani.

On Efficient and Secure Compression Functions for Arithmetization-Oriented Hashing.

In 2024 IEEE 37th Computer Security Foundations Symposium (CSF), pages 1–16, Los Alamitos, CA, USA, Jul 2024. IEEE Computer Society.



Amit Singh Bhati, Erik Pohle, Aysajan Abidin, Elena Andreeva, and Bart Preneel.

Let's go eevee! a friendly and suitable family of aead modes for iot-to-cloud secure computation.

In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security, CCS '23*, pages 2546–2560, New York, NY, USA, 2023. Association for Computing Machinery.



John Black, Phillip Rogaway, and Thomas Shrimpton.
Black-box analysis of the block-cipher-based hash-function constructions from pgv.

In Moti Yung, editor, *Advances in Cryptology — CRYPTO 2002*, pages 320–335, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg.

-  Joppe W. Bos and Peter L. Montgomery.
Montgomery arithmetic from a software perspective.
Cryptology ePrint Archive, Paper 2017/1057, 2017.
-  Alessandro Chiesa, Dev Ojha, and Nicholas Spooner.
Fractal: Post-quantum and transparent recursive proofs from holography.
In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology – EUROCRYPT 2020*, pages 769–793, Cham, 2020. Springer International Publishing.
-  Shafi Goldwasser, Silvio Micali, and Charles Rackoff.
The knowledge complexity of interactive proof systems.
SIAM Journal on Computing, 18(1):186–208, 1989.



Lorenzo Grassi, Dmitry Khovratovich, and Markus Schofnegger.

Poseidon2: A faster version of the poseidon hash function.

Cryptology ePrint Archive, Paper 2023/323, 2023.



Lorenzo Grassi, Reinhard Lüftenegger, Christian Rechberger, Dragos Rotaru, and Markus Schofnegger.

On a generalization of substitution-permutation networks: The hades design strategy.

Cryptology ePrint Archive, Paper 2019/1107, 2019.



Jens Groth.

On the size of pairing-based non-interactive arguments.

In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology – EUROCRYPT 2016*, pages 305–326, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.



Dmitry Khovratovich, Mario Marhuenda Beltrán, and Bart Mennink.

Generic security of the safe api and its applications.

In Jian Guo and Ron Steinfeld, editors, *Advances in Cryptology – ASIACRYPT 2023*, pages 301–327, Singapore, 2023. Springer Nature Singapore.



Ralph Charles Merkle.

Secrecy, Authentication, and Public Key Systems.

PhD thesis, Stanford University, Stanford, CA, USA, 1979.
AAI8001972.



Tomer Solberg.

Risc zero prover protocol & analysis.

[https://github.com/ingonyama-zk/papers/
blob/main/risc0_protocol_analysis.pdf](https://github.com/ingonyama-zk/papers/blob/main/risc0_protocol_analysis.pdf), 04 2023.

Accessed: 2026-02-08.