

A Cambrian Explosion of Cryptographic Proofs

Eli Ben-Sasson

◆ February 2023



Background

Background

Checking Computations in Polylogarithmic Time

1992

László Babai¹
Dept. Comp. Sci.
Univ. of Chicago⁶
Budapest, Hungary

Lance Fortnow²
Dept. Comp. Sci.
Univ. of Chicago⁶
Boston University⁴

Leonid A. Levin³
Dept. Comp. Sci.
Boston University⁴

Mario Szegedy⁵
Dept. Comp. Sci.
Univ. of Chicago

Abstract. Motivated by Manuel Blum's concept of *instance checking*, we consider new, very fast and generic mechanisms of checking computations. Our results exploit recent advances in interactive proof protocols [LFKN92], [Sha92], and especially the *MIP = NEXP* protocol from [BFL91].

We show that every nondeterministic computational task $S(x, y)$, defined as a polynomial time relation between the *instance* x , representing the input and output combined, and the *witness* y can be modified to a task S' such that: (i) the same instances remain accepted; (ii) each instance/witness pair becomes checkable in *polylogarithmic* Monte Carlo time; and (iii) a witness satisfying S' can be computed in polynomial time from a witness satisfying S .

Here the instance and the description of S have to be provided in error-correcting code (since the checker will not notice slight changes). A modification of the *MIP* proof was required to achieve polynomial time in (iii); the earlier technique yields $N^{O(\log \log N)}$ time only.

This result becomes significant if software and hardware *reliability* are regarded as a considerable cost factor. The polylogarithmic checker is the only part of the system that needs to be trusted; it can be *hard wired*. (We use just *one Checker* for all problems!) The checker is tiny and so presumably can be optimized and checked off-line at a modest cost.

In this setup, a single reliable PC can monitor the operation of a herd of supercomputers working with possibly extremely powerful but unreliable software and untested hardware.

In another interpretation, we show that in polynomial time, every formal mathematical proof can be transformed into a *transparent proof*, i.e. a proof verifiable in polylogarithmic Monte Carlo time, assuming the "theorem-candidate" is given in error-correcting code. In fact, for any $\varepsilon > 0$, we can transform any proof P in time $\|P\|^{1+\varepsilon}$ into a transparent proof, verifiable in Monte Carlo time $(\log \|P\|)^{O(1/\varepsilon)}$.

As a by-product, we obtain a binary error correcting code with very efficient error-correction. The code transforms messages of length N into codewords of length $\leq N^{1+\varepsilon}$; and for strings within 10% of a valid codeword, it allows to recover any bit of the unique codeword within that distance in polylogarithmic $((\log N)^{O(1/\varepsilon)})$ time.

Integrity* via Math

* Integrity means doing the right thing, even when no one is watching [C.S. Lewis]

Background

Integrity* via Math

“...a single reliable PC can monitor the operation of a herd of supercomputers with powerful but unreliable software and untested hardware ...”



Checking Computations in Polylogarithmic Time

László Babai¹
Dept. Comp. Sci.
Univ. of Chicago⁶
Budapest, Hungary

Lance Fortnow²
Dept. Comp. Sci.
Univ. of Chicago⁶
Dept. Comp. Sci.
Univ. of Chicago⁵

Leonid A. Levin³
Dept. Comp. Sci.
Boston University⁴

1992

Abstract. Motivated by Manuel Blum's concept of *instance checking*, we consider new, very fast and generic mechanisms of checking computations. Our results exploit recent advances in interactive proof protocols [LFKN92], [Sha92], and especially the *MIP = NEXP* protocol from [BFL91].

We show that every nondeterministic computational task $S(x, y)$, defined as a polynomial time relation between the *instance* x , representing the input and output combined, and the *witness* y can be modified to a task S' such that: (i) the same instances remain accepted; (ii) each instance/witness pair becomes checkable in *polylogarithmic* Monte Carlo time; and (iii) a witness satisfying S' can be computed in polynomial time from a witness satisfying S .

Here the instance and the description of S have to be provided in error-correcting code (since the checker will not notice slight changes). A modification of the *MIP* proof was required to achieve polynomial time in (iii); the earlier technique yields $N^{O(\log\log N)}$ time only.

This result becomes significant if software and hardware *reliability* are regarded as a considerable cost factor. The polylogarithmic checker is the only part of the system that needs to be trusted; it can be *hard wired*. (We use just *one Checker* for all problems!) The checker is tiny and so presumably can be optimized and checked off-line at a modest cost.

In this setup, a single reliable PC can monitor the operation of a herd of supercomputers working with possibly extremely powerful but unreliable software and untested hardware.

In another interpretation, we show that in polynomial time, every formal mathematical proof can be transformed into a *transparent proof*, i.e. a proof verifiable in polylogarithmic Monte Carlo time, assuming the “theorem-candidate” is given in an error-correcting code. In fact, for any $\varepsilon > 0$, we can transform any proof P in time $\|P\|^{1+\varepsilon}$ into a transparent proof, verifiable in Monte Carlo time $(\log\|P\|)^{O(1/\varepsilon)}$.

As a by-product, we obtain a binary error correcting code with very efficient error-correction. The code transforms messages of length N into codewords of length $\leq N^{1+\varepsilon}$; and for strings within 10% of a valid codeword, it allows to recover any bit of the unique codeword within that distance in polylogarithmic $((\log N)^{O(1/\varepsilon)})$ time.

* Integrity means doing the right thing, even when no one is watching [C.S. Lewis]

Background

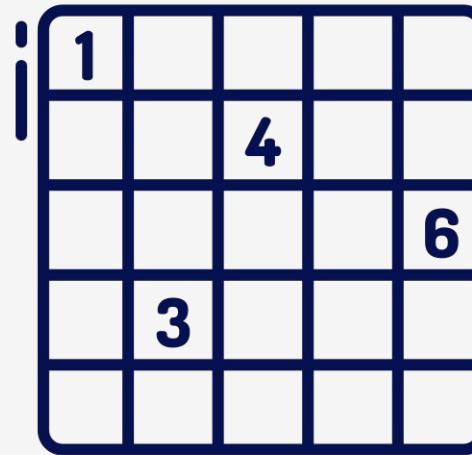
Claim: Starting @ state hash **x**, after **1,000,000** txs processed by program **P**, reached state hash **y**

“...a single reliable PC can monitor the operation of a herd of supercomputers with powerful but unreliable software and untested hardware ...”

Background

Claim: Starting @ state hash x , after **1,000,000** txs processed by program P , reached state hash y

A sudoku-like set of constraints is implied by the statement proved, by x , y , P , and $\#tx$ (=1,000,000)



PCP

“...a single reliable PC can monitor the operation of a herd of supercomputers with powerful but unreliable software and untested hardware ...”

Background

Claim: Starting @ state hash x , after **1,000,000** txs processed by program P , reached state hash y

A sudoku-like set of constraints is implied by the statement proved, by x , y , P , and $\#tx$ (=1,000,000)

1	1	3	2	6
4	1	4	7	3
8	5	8	3	6
6	3	5	9	7
5	2	8	7	6

PCP

“...a single reliable PC can monitor the operation of a herd of supercomputers with powerful but unreliable software and untested hardware ...”

Prover submits solution

Background

Claim: Starting @ state hash x , after **1,000,000** txs processed by program P , reached state hash y

“...a single reliable PC can monitor the operation of a herd of supercomputers with powerful but unreliable software and untested hardware ...”

A sudoku-like set of constraints is implied by the statement proved, by x , y , P , and $\#tx$ (=1,000,000)

1	1	3	2	6
4	1	4	7	3
8	5	8	3	6
6	3	5	9	7
5	2	8	7	6

PCP

Prover submits solution

Verifier samples and checks a single constraint

Background

Claim: Starting @ state hash x , after **1,000,000** txs processed by program P , reached state hash y

Magic (aka Math)

- Sampling constraints takes exponentially small time!
- Good proofs satisfy ALL constraints!
- A “proof” of a false claim satisfies < 1% of constraints!

“...a single reliable PC can monitor the operation of a herd of supercomputers with powerful but unreliable software and untested hardware ...”

A sudoku-like set of constraints is implied by the statement proved, by x , y , P , and $\#tx$ (=1,000,000)

1	1	3	2	6
4	1	4	7	3
8	5	8	3	6
6	3	5	9	7
5	2	8	7	6

PCP

Prover submits solution

Verifier samples and checks a single constraint

Background

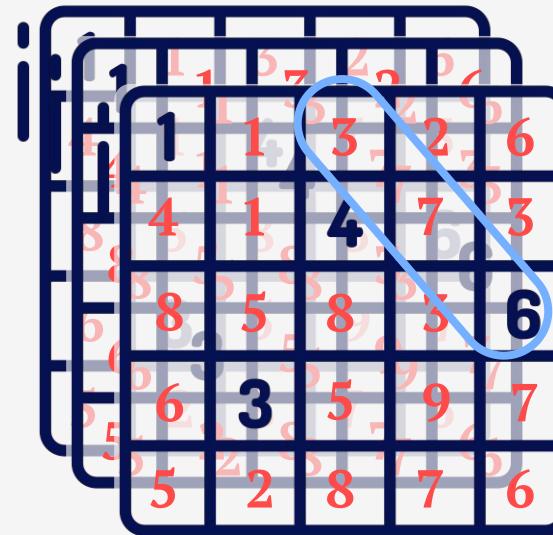
Claim: Starting @ state hash x , after **1,000,000** txs processed by program P , reached state hash y

Magic (aka Math)

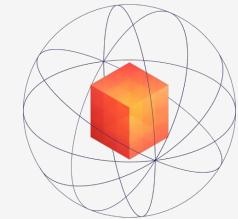
- Sampling constraints takes exponentially small time!
- Good proofs satisfy ALL constraints!
- A “proof” of a false claim satisfies < 1% of constraints!

“...a single reliable PC can monitor the operation of a herd of supercomputers with powerful but unreliable software and untested hardware ...”

A sudoku-like set of constraints is implied by the statement proved, by x , y , P , and $\#tx$ (=1,000,000)



STARK

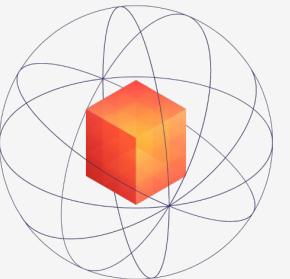


Prover submits solution

Verifier posts another (random) sudoku puzzle

Verifier samples and checks a single constraint

Background



STARK

Integrity* via Math
(impractical)

“...a single reliable PC can monitor the operation of a herd of supercomputers with powerful but unreliable software and untested hardware ...”

1992

Checking Computations in Polylogarithmic Time

László Babai¹
Dept. Comp. Sci.
Univ. of Chicago⁶ and
Univ., Budapest

Lance Fortnow²
Dept. Comp. Sci.
Univ. of Chicago⁶

Mario Szegedy⁵
Dept. Comp. Sci.
Univ. of Chicago

Leonid A. Levin³
Dept. Comp. Sci.
Boston University⁴

PCP

Abstract. Motivated by Manuel Blum's concept of *instance checking*, we consider new, very fast and generic mechanisms of checking computations. Our results exploit recent advances in interactive proof protocols [LFKN92], [Sha92], and especially the *MIP = NEXP* protocol from [BFKL91].

We show that every nondeterministic computational task $S(x, y)$, defined as a polynomial time relation between the *instance* x , representing the input and output combined, and the *witness* y can be modified to a task S' such that: (i) the same instances remain accepted; (ii) each instance/witness pair becomes checkable in *polylogarithmic* Monte Carlo time; and (iii) a witness satisfying S' can be computed in polynomial time from a witness satisfying S .

Here the instance and the description of S have to be provided in error-correcting code (since the checker will not notice slight changes). A modification of the *MIP* proof was required to achieve polynomial time in (iii); the earlier technique yields $N^{O(\log \log N)}$ time only.

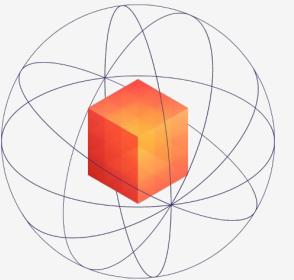
This result becomes significant if software and hardware *reliability* are regarded as a considerable cost factor. The polylogarithmic checker is the only part of the system that needs to be trusted; it can be *hard wired*. (We use just *one Checker* for all problems!) The checker is tiny and so presumably can be optimized and checked off-line at a modest cost.

In this setup, a single reliable PC can monitor the operation of a herd of supercomputers working with possibly extremely powerful but unreliable software and untested hardware.

In another interpretation, we show that in polynomial time, every formal mathematical proof can be transformed into a *transparent proof*, i.e. a proof verifiable in polylogarithmic Monte Carlo time, assuming the “theorem-candidate” is given in error-correcting code. In fact, for any $\varepsilon > 0$, we can transform any proof P in time $\|P\|^{1+\varepsilon}$ into a transparent proof, verifiable in Monte Carlo time $(\log \|P\|)^{O(1/\varepsilon)}$.

As a by-product, we obtain a binary error correcting code with very efficient error-correction. The code transforms messages of length N into codewords of length $\leq N^{1+\varepsilon}$; and for strings within 10% of a valid codeword, it allows to recover any bit of the unique codeword within that distance in polylogarithmic $((\log N)^{O(1/\varepsilon)})$ time.

Background



STARK

Integrity* via Math
(impractical)

“...a single reliable PC can monitor the operation of a herd of supercomputers with powerful but unreliable software and untested hardware ...”

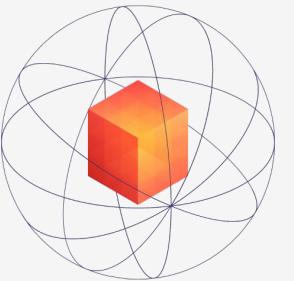
* Integrity means doing the right thing, even when no one is watching [C.S. Lewis]

The timeline consists of five academic papers arranged vertically, each with a large blue circular badge containing a year (1992, 2004, 2015, 2018) in the top right corner.

- 1992**
Title: Checking Computations in Polylogarithmic Time
Authors: László Babai¹, Lance Fortnow², Leonid A. Levin³
Institutions: Univ. of Chicago⁶ and Univ. of Chicago⁶, Univ. of Chicago⁶, Boston University⁴
Journal: SIAM J. COMPUT. Vol. 38, No. 2, pp. 551-607
Copyright: © 2008 Society for Industrial and Applied Mathematics
- 2004**
Title: SHORT PCPs WITH POLYLOG QUERY COMPLEXITY
Authors: ELI BEN-SASSON[†] AND MADHU SUDAN[‡]
- 2015**
Title: Interactive Oracle Proofs*
Authors: Eli Ben-Sasson¹, Alessandro Chiesa² and Nicholas Spooner³
- 2018**
Title: Fast Reed-Solomon Interactive Oracle Proofs of Proximity
Authors: Eli Ben-Sasson* Iddo Bentov[†] Yinon Horesh* Michael Riabzev*
Date: January 12, 2018
- 2018**
Title: Scalable, transparent, and post-quantum secure computation via zk-STARK
Authors: Eli Ben-Sasson* Iddo Bentov[†] Yinon Horesh* Michael Riabzev*
Date: March 6, 2018

Abstract
Human dignity demands that personal information, like medical and forensic data, be hidden from the public. But veils of secrecy designed to preserve privacy may also be abused to cover up lies and deceit by institutions entrusted with Data, unjustly harming citizens and eroding trust in central institutions. Zero knowledge (ZK) proof systems are an ingenious cryptographic solution to this tension between the ideals of personal *privacy* and *computational integrity*, enabling computation in a way that is not compromised for privacy. Public audit demands *non-interactive* proofs from ZK systems, meaning they be set up with no reliance on any trusted party, and have no trapdoors that could be exploited by powerful parties to bear false witness. For ZK systems to be used with Big Data, it is imperative that the public verification process scale sublinearly in data size. Transparent ZK proofs that can be verified *exponentially* faster than data size were first described in the 1990s but early constructions were impractical, and no ZK system realized thus far in code (including that used by crypto-currencies like Zcash™) has achieved both transparency and exponential verification speedup, simultaneously, for general computations. Here we report the first realization of a transparent ZK system (ZK-STARK) in which verification scales exponentially faster than database size, and moreover, this exponential speedup in verification is observed concretely for meaningful and sequential computations, described next. Our system uses several recent advances on interactive oracle proofs (IOP), such as a “fast” (linear time) IOP system for error correcting codes.

Background



STARK

Integrity* via Math
(impractical)

“...a single reliable PC can monitor the operation of a herd of supercomputers with powerful but unreliable software and untested hardware ...”

* Integrity means doing the right thing, even when no one is watching [C.S. Lewis]

1992

Checking Computations in Polylogarithmic Time
László Babai¹
Dept. Comp. Sci.
Univ. of Chicago⁶
Univ. of Budapest⁶
Lance Fortnow²
Dept. Comp. Sci.
Univ. of Chicago⁶
Leonid A. Levin³
Dept. Comp. Sci.
Boston University⁴
Mario Szegedy⁵
Dept. Comp. Sci.
Univ. of Chicago

© 2008 Society for Industrial and Applied Mathematics
SIAM J. COMPUT. Vol. 38, No. 2, pp. 551-607

2004

SHORT PCPS WITH POLYLOG QUERY COMPLEXITY
ELI BEN-SASSON[†] AND MADHU SUDAN[‡]

2015

Interactive Oracle Proofs*
Eli Ben-Sasson¹, Alessandro Chiesa² and Nicholas Spooner³

2018

Fast Reed-Solomon Interactive Oracle Proofs of Proximity
Eli Ben-Sasson* Iddo Bentov[†] Yinon Horesh* Michael Riabzev*
January 12, 2018

Scalable, transparent, and post-quantum secure computation
Eli Ben-Sasson* Iddo Bentov[†] Yinon Horesh* Michael Riabzev*
March 6, 2018

Abstract
Human dignity demands that personal information, like medical and forensic data, be hidden from the public. But veils of secrecy designed to preserve privacy may also be abused to cover up lies and deceit by institutions entrusted with Data, unjustly harming citizens and eroding trust in central institutions. Zero knowledge (ZK) proof systems are an ingenious cryptographic solution to this tension between the ideals of personal privacy and computational integrity, enabling computation in a way that is not compromised by the need for privacy. Public audit demands *non-interactive* proofs from ZK systems, meaning they be set up with no reliance on any trusted party, and have no trapdoors that could be exploited by powerful parties to bear false witness. For ZK systems to be used with Big Data, it is imperative that the public verification process scale sublinearly in data size. Transparent ZK proofs that can be verified exponentially faster than data size were first described in the 1990s but early constructions were impractical, and no ZK system realized thus far in code (including that used by crypto-currencies like Zcash™) has achieved both transparency and exponential verification speedup, simultaneously, for general computations. Here we report the first realization of a transparent ZK system (ZK-STARK) in which verification scales exponentially faster than database size, and moreover, this exponential speedup in verification is observed concretely for meaningful and sequential computations, described next. Our system uses several recent advances on interactive oracle proofs (IOP), such as a “fast” (linear time) IOP system for error correcting codes.

Background

Integrity* via Math

“...a single reliable PC can monitor the operation of a herd of supercomputers with powerful but unreliable software and untested hardware ...”

Bitcoin: A Peer-to-Peer Electronic Cash System



Satoshi Nakamoto
satoshin@gmx.com
www.bitcoin.org

Abstract. A purely peer-to-peer version of electronic cash would allow online payments to be sent directly from one party to another without going through a financial institution. Digital signatures provide part of the solution, but the main benefits are lost if a trusted third party is still required to prevent double-spending. We propose a solution to the double-spending problem using a peer-to-peer network. The network timestamps transactions by hashing them into an ongoing chain of hash-based proof-of-work, forming a record that cannot be changed without redoing the proof-of-work. The longest chain not only serves as proof of the sequence of events witnessed, but proof that it came from the largest pool of CPU power. As long as a majority of CPU power is controlled by nodes that are not cooperating to attack the network, they'll generate the longest chain and outpace attackers. The network itself requires minimal structure. Messages are broadcast on a best effort basis, and nodes can leave and rejoin the network at will, accepting the longest proof-of-work chain as proof of what happened while they were gone.

Background

Integrity* via Math

“...a single reliable PC can monitor the operation of a herd of supercomputers with powerful but unreliable software and untested hardware ...”

* Integrity means doing the right thing, even when no one is watching [C.S. Lewis]

Bitcoin: A Peer-to-Peer Electronic Cash System



Satoshi Nakamoto
satoshin@gmx.com
www.bitcoin.org

Ethereum: A Next-Generation Smart Contract and Decentralized Application Platform.
By Vitalik Buterin (2014).



When Satoshi Nakamoto first set the Bitcoin blockchain into motion in January 2009, it was simultaneously introducing two radical and untested concepts. The first is the "bitcoin", a decentralized peer-to-peer online currency that maintains a value without any backing, intrinsic value or central issuer. So far, the "bitcoin" as a currency unit has taken up the bulk of the public attention, both in terms of the political aspects of a currency without a central bank and its extreme upward and downward volatility in price. However, there is also another, equally important, part to Satoshi's grand experiment: the concept of a proof-of-work-based blockchain to allow for public agreement on the order of transactions. Bitcoin as an application can be described as a first-to-finish system: if one entity has 50 BTC, and simultaneously sends the same 50 BTC to A and to B, only the transaction that gets confirmed first will process. There is no intrinsic way of determining from two transactions which came earlier, and for decades this stymied the development of decentralized digital currency. Satoshi's blockchain was the first credible decentralized solution. And now, attention is rapidly starting to shift toward this second part of Bitcoin's technology, and how the blockchain concept can be used for more than just money.

Background

Integrity* via Math

“...a **single reliable PC** can monitor the operation of a herd of supercomputers with powerful but unreliable software and untested hardware ...”



Blockchain is the “single reliable PC”

Background

Integrity* via Math

2013



“...a **single reliable PC** can monitor the operation of a herd of supercomputers with powerful but unreliable software and untested hardware ...”



Blockchain is the “single reliable PC”

Background

Integrity* via Math

2013



“...a single reliable PC can monitor the operation of a herd of supercomputers with powerful but unreliable software and untested hardware ...”



Blockchain is the “single reliable PC”

2015 - Zcash 1st general ZK for privacy

2018 - StarkWare 1st Proofs for scalability

Cambrian Explosion of Cryptographic Proofs

September 2019

libSTARK

Aurora

Ligero

ZKBoo

genSTARK

STARK

BulletProofs

Halo

Groth16

SONIC

PLONK

Pinocchio

November 2019

libSTARK

Aurora

Hodor

Ligero

ZKBoo

Fractal

Succ. Aurora

openZKP

STARK

genSTARK

BulletProofs

Marlin

SLONK

Halo

Groth16

SONIC

PLONK

Pinocchio

SuperSonic

February 2023

Awesome Stuff!

Awesome Stuff!

Awesome Stuff!

libSTARK

Aurora

Hodor

Ligero

ZKBoo

Fractal

Succ. Aurora

openZKP

STARK

genSTARK

Awesome Stuff!

BulletProofs

Marlin

SLONK

Groth16

SONIC

PLONK

Pinocchio

SuperSonic

Awesome Stuff!

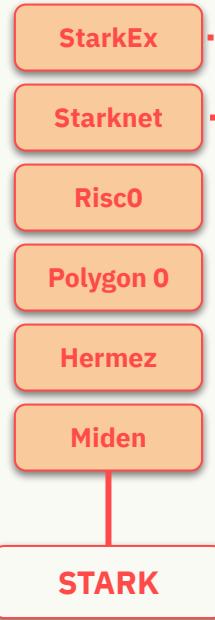
Blockchain Usage February 2023



Blockchain Usage February 2023



Privacy



Scalability

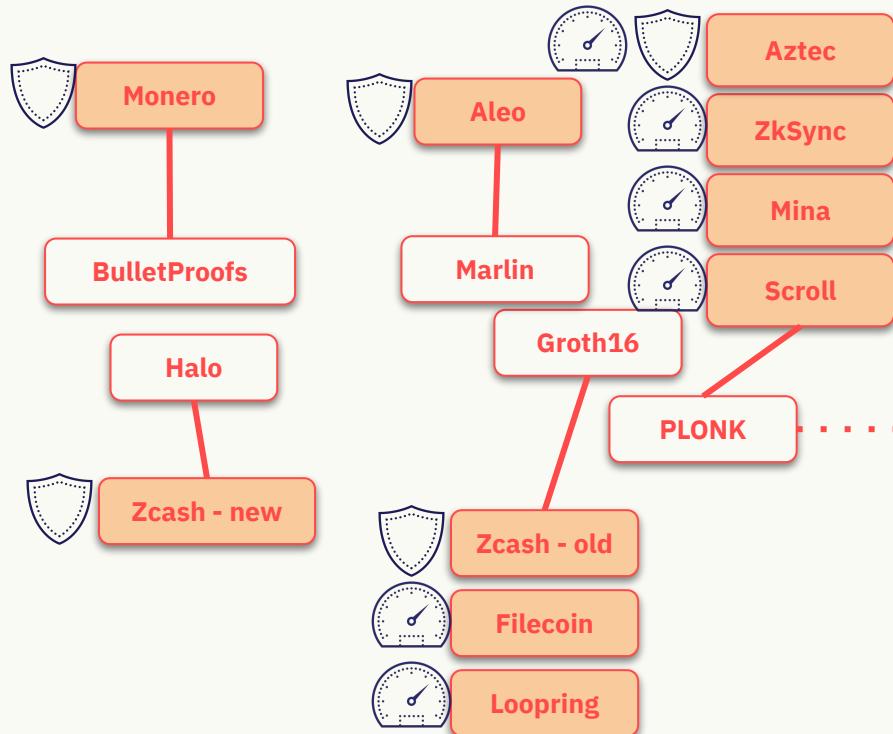
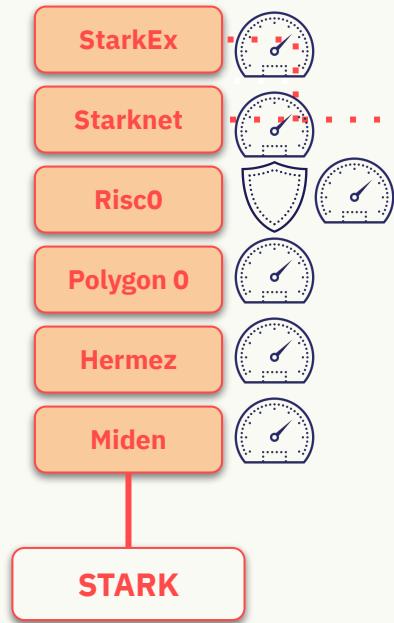


Blockchain Usage
February 2023

Privacy



Scalability



Blockchain Usage February 2023



Privacy

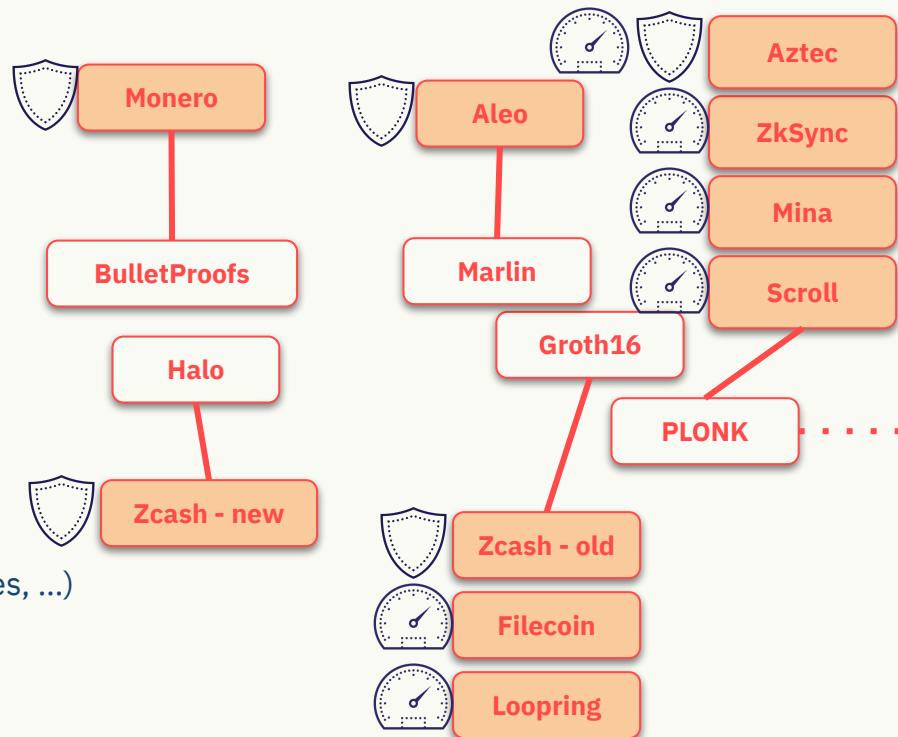


Scalability



Why so few systems in blockchain?

- theory-to-practice takes time
- existing systems good enough for scale
- tech standards (network protocols, programming languages, ...)
- bottleneck is **not** proof/verification efficiency
- bottlenecks: productization, dev tools, integration, ...



Proofs of Computational Integrity (CI)



Privacy (Zero Knowledge, ZK)

Prover's private inputs are shielded



Scalability

Exponentially small verifier running time*

Nearly linear prover running time*



Universality

Applicability to general computation



Transparency

No toxic waste (i.e. no trusted setup)

*With respect to size of computation

Common Ancestors

1. Arithmetization 2. Low degreeness

1) Arithmetization

Arithmetization Converts (“reduces”) Computational Integrity problems to problems about local relations between a bunch of polynomials

Example: For public 256-bit string z , Bob claims knows a SHA2-preimage of z

1) Arithmetization

Arithmetization Converts (“reduces”) Computational Integrity problems to problems about local relations between a bunch of polynomials

Example: For public 256-bit string z , Bob claims knows a SHA2-preimage of z

**Pre-arithmetization
claim**

“*I know y such that
SHA2(y)= z* ”

1) Arithmetization

Arithmetization Converts (“reduces”) Computational Integrity problems to problems about local relations between a bunch of polynomials

Example: For public 256-bit string z , Bob claims knows a SHA2-preimage of z

Pre-arithmetization claim

“I know y such that
 $SHA2(y)=z$ ”

Reduction

produces 2
polynomials:
 $Q(X, Y, T, W)$, $R(X)$ and
degree bound d

1) Arithmetization

Arithmetization Converts (“reduces”) Computational Integrity problems to problems about local relations between a bunch of polynomials

Example: For public 256-bit string z , Bob claims knows a SHA2-preimage of z

Pre-arithmetization claim

“I know y such that
 $SHA2(y)=z$ ”

Reduction

produces 2
polynomials:
 $Q(X, Y, T, W)$, $R(X)$ and
degree bound d

Post-arithmetization claim

I know 4 polynomials
of degree d - $A(x)$, $B(x)$,
 $C(x)$, $D(X)$ - such that:

$Q(X, A(X), B(X+1),$
 $C(2^*X))=D(X) * R(X)$

1) Arithmetization

Arithmetization Converts (“reduces”) Computational Integrity problems to problems about local relations between a bunch of polynomials

Example: For public 256-bit string z , Bob claims knows a SHA2-preimage of z

Pre-arithmetization claim

“I know y such that $SHA2(y)=z$ ”

Reduction

produces 2 polynomials:
 $Q(X, Y, T, W)$, $R(X)$ and degree bound d

Post-arithmetization claim

I know 4 polynomials of degree d - $A(x)$, $B(x)$, $C(x)$, $D(X)$ - such that:

$$Q(X, A(X), B(X+1), C(2^*X)) = D(X) * R(X)$$

Theorem

If A , B , C , D do not satisfy **THIS**,

then nearly all x expose Bob's lie

1) Arithmetization

Assuming Theorem, we get a scalable proof system for Bob's original claim:

1. Apply reduction, ask Bob to provide access to A,B,C,D of degree-d
2. Sample random x and accept Bob's claim iff equality holds for this x

Pre-arithmetization claim

"I know y such that
 $SHA2(y)=z$ "

Reduction

produces 2
polynomials:
 $Q(X, Y, T, W)$, $R(X)$ and
degree bound d

Post-arithmetization claim

I know 4 polynomials
of degree d - $A(x)$, $B(x)$,
 $C(x)$, $D(x)$ - such that:

$$Q(X, A(X), B(X+1),
C(2^*X)) = D(X) * R(X)$$

Theorem

If A , B , C , D do not
satisfy THIS,

then nearly all x
expose Bob's lie

2) Low degreeness

Assuming Theorem, we get a scalable proof system for Bob's original claim:

1. Apply reduction, ask Bob to provide access to A,B,C,D of degree-d
2. Sample random x and accept Bob's claim iff equality holds for this x

New Computational Integrity problem: Force Bob to answer all queries according to some quadruple of degree-d polynomials

Post-arithmetization claim

I know 4 polynomials of degree d - A(x), B(x), C(x), D(X) - such that:

$$Q(X, A(X), B(X+1), C(2^*X)) = D(X) * R(X)$$

Theorem

If A, B, C, D do not satisfy THIS,

then nearly all x expose Bob's lie

Common Ancestors

1. Arithmetization 2. Low degreeness

Differentiating Factors



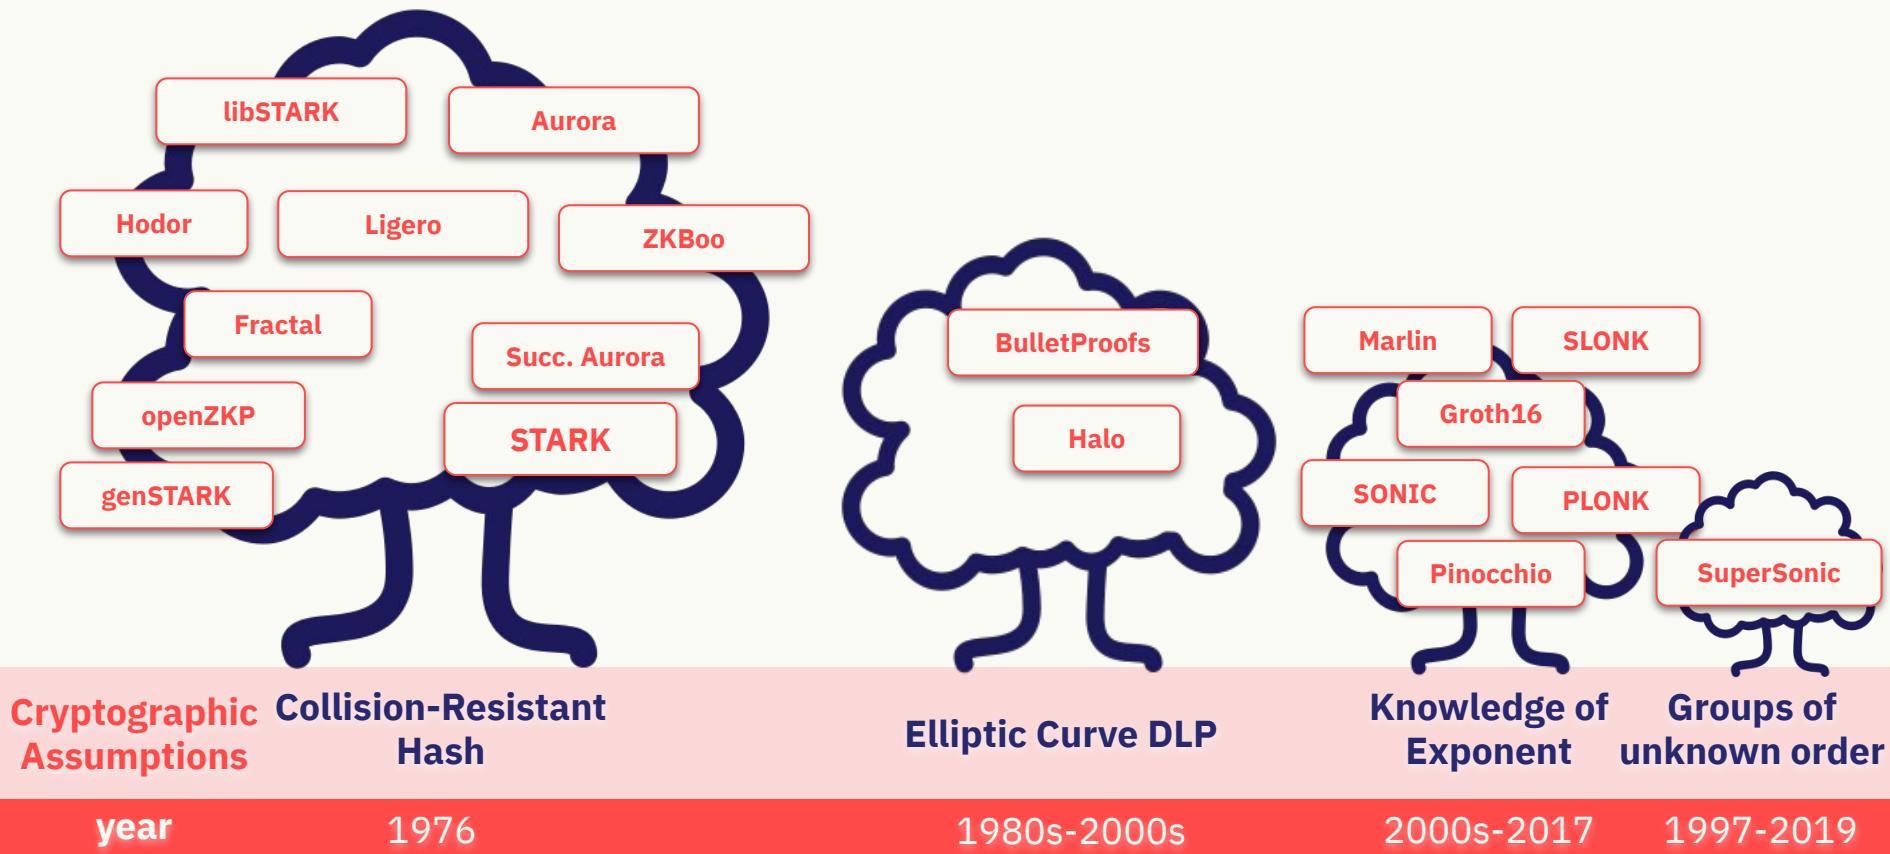
- 1. Arithmetization Method
- 2. Low degreeness enforcement
- 3. Cryptographic assumptions used to get 2



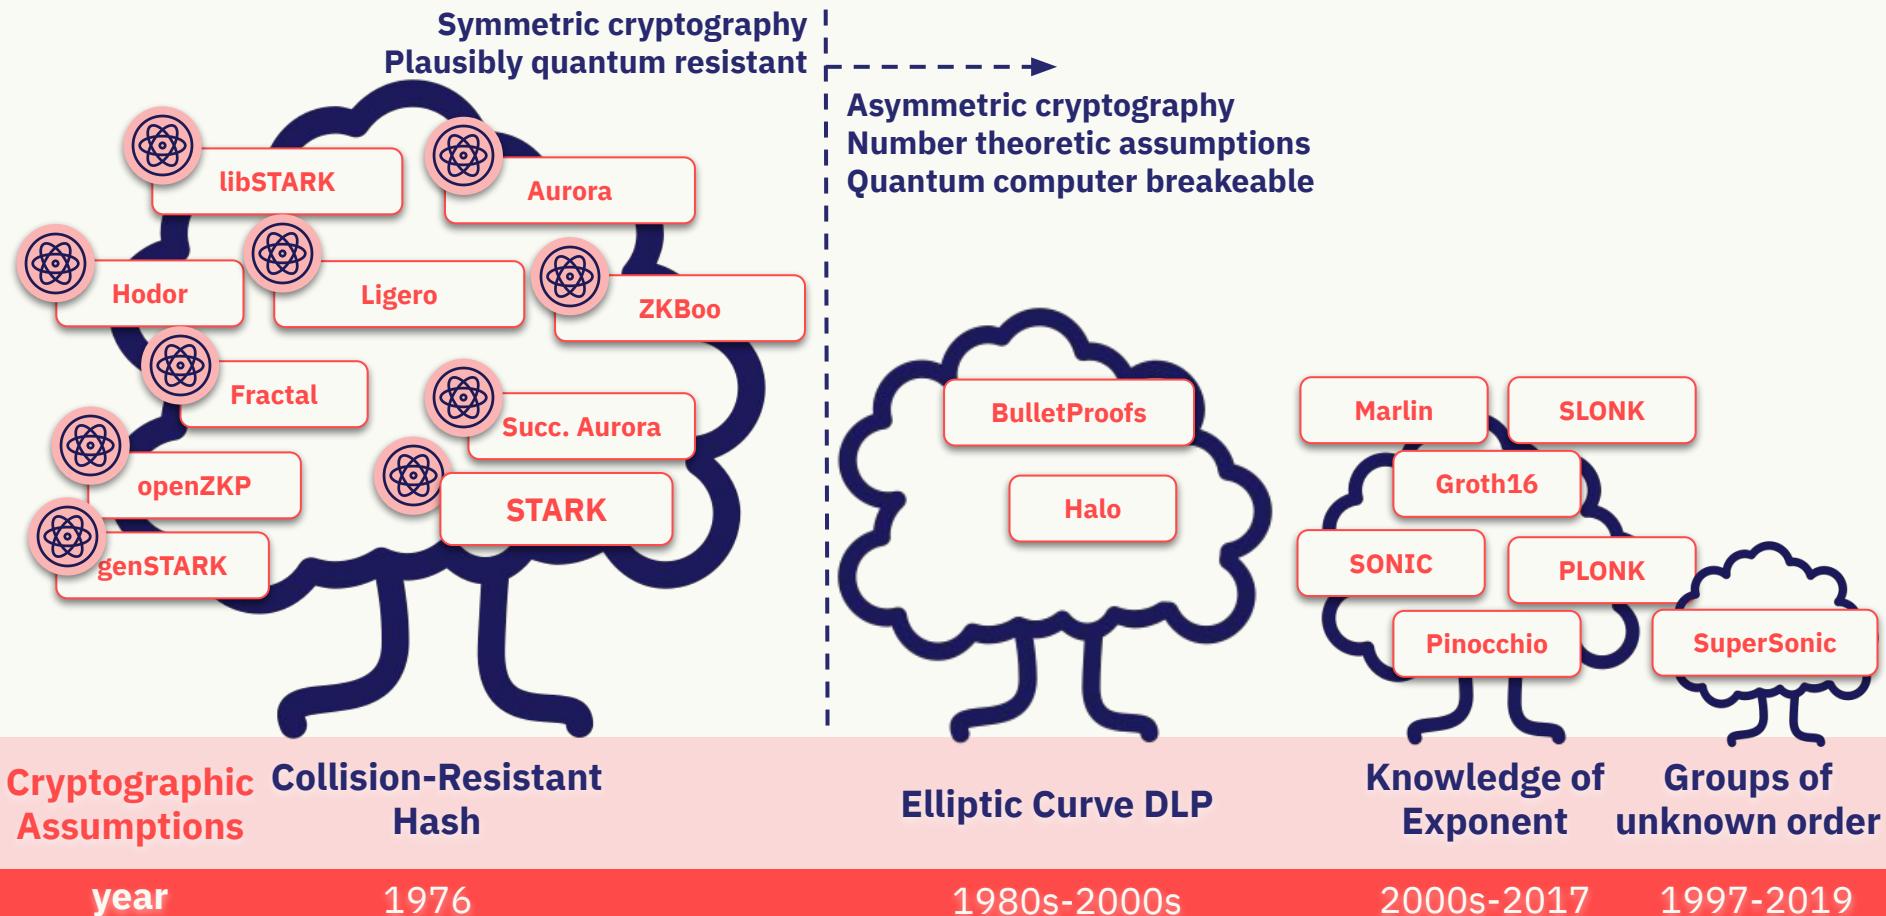
Common Ancestors

1. Arithmetization
2. Low degreeness

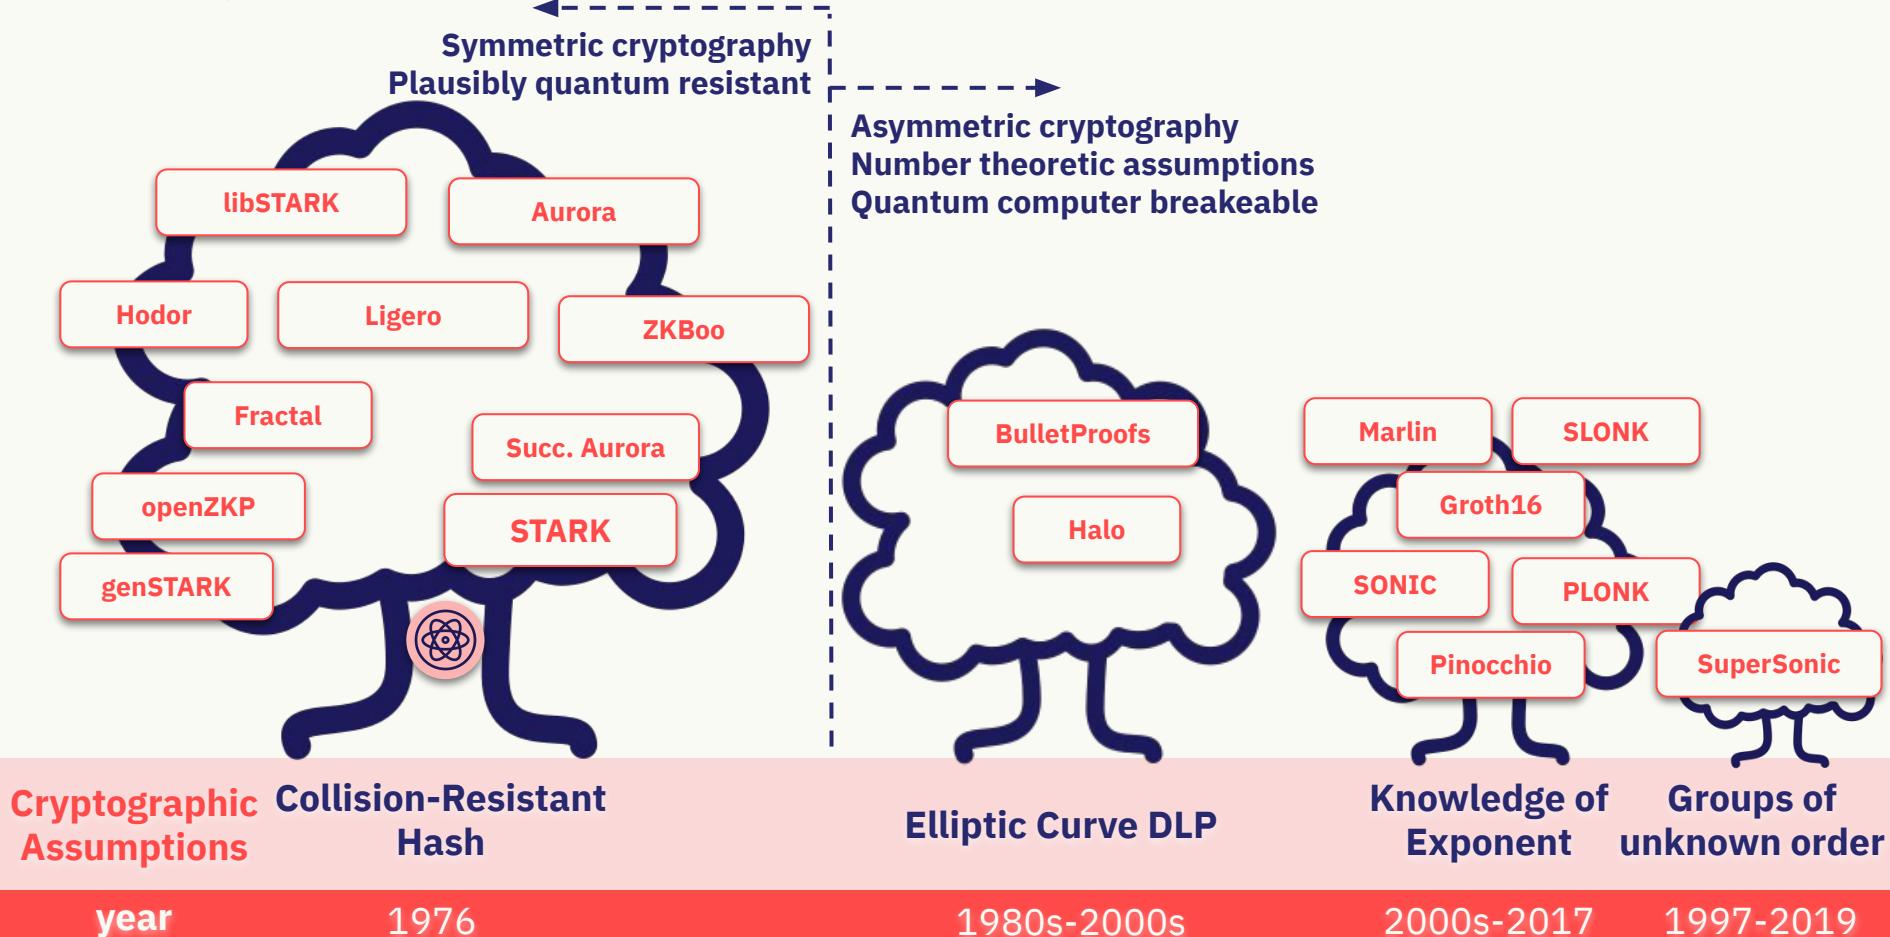
3. Cryptographic Assumptions



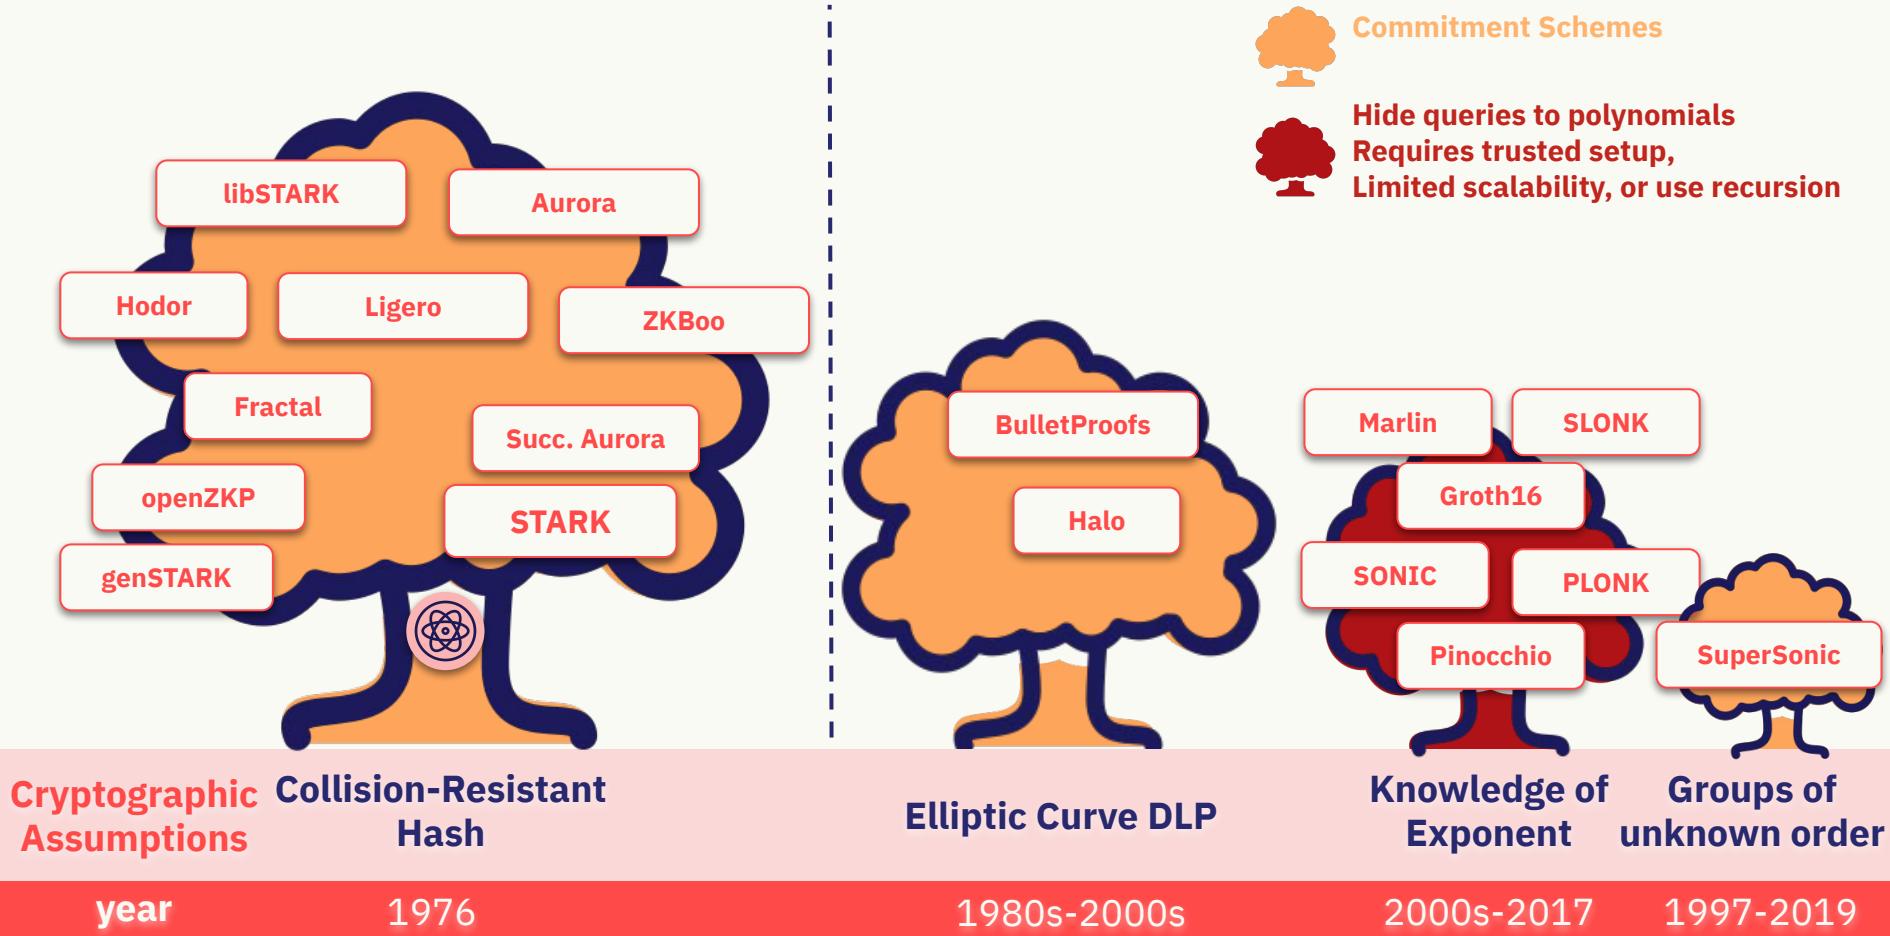
3. Cryptographic Assumptions



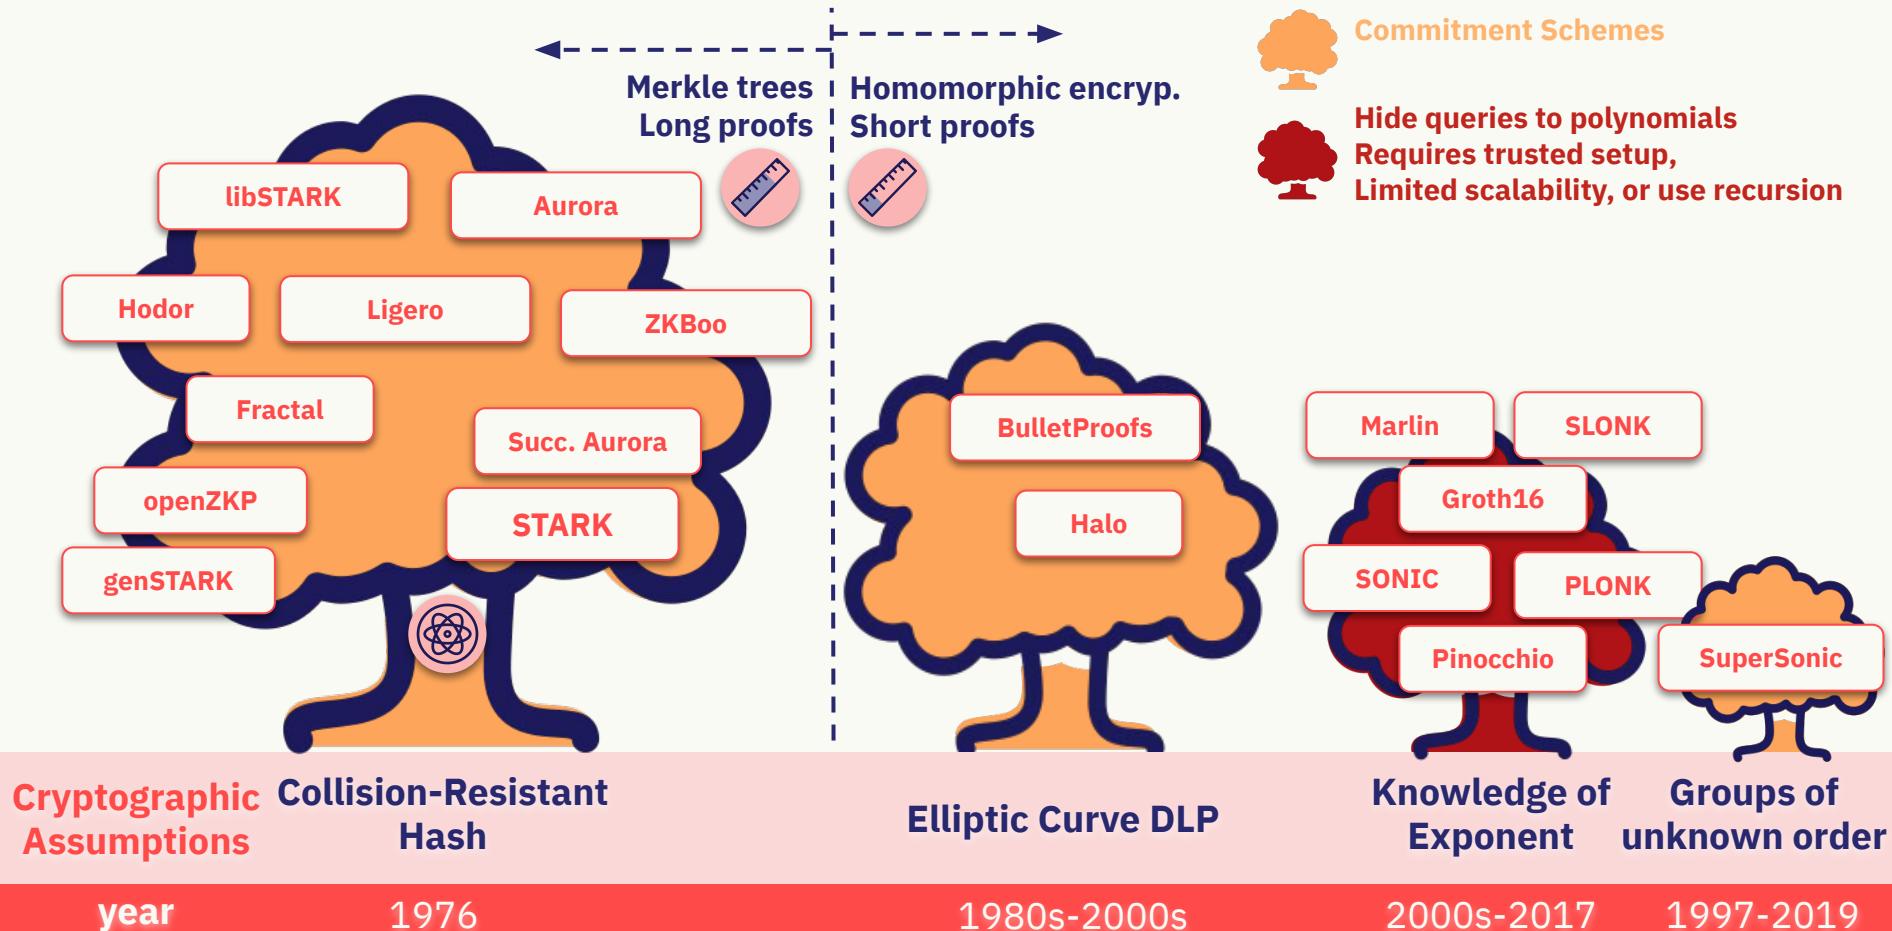
3. Cryptographic Assumptions



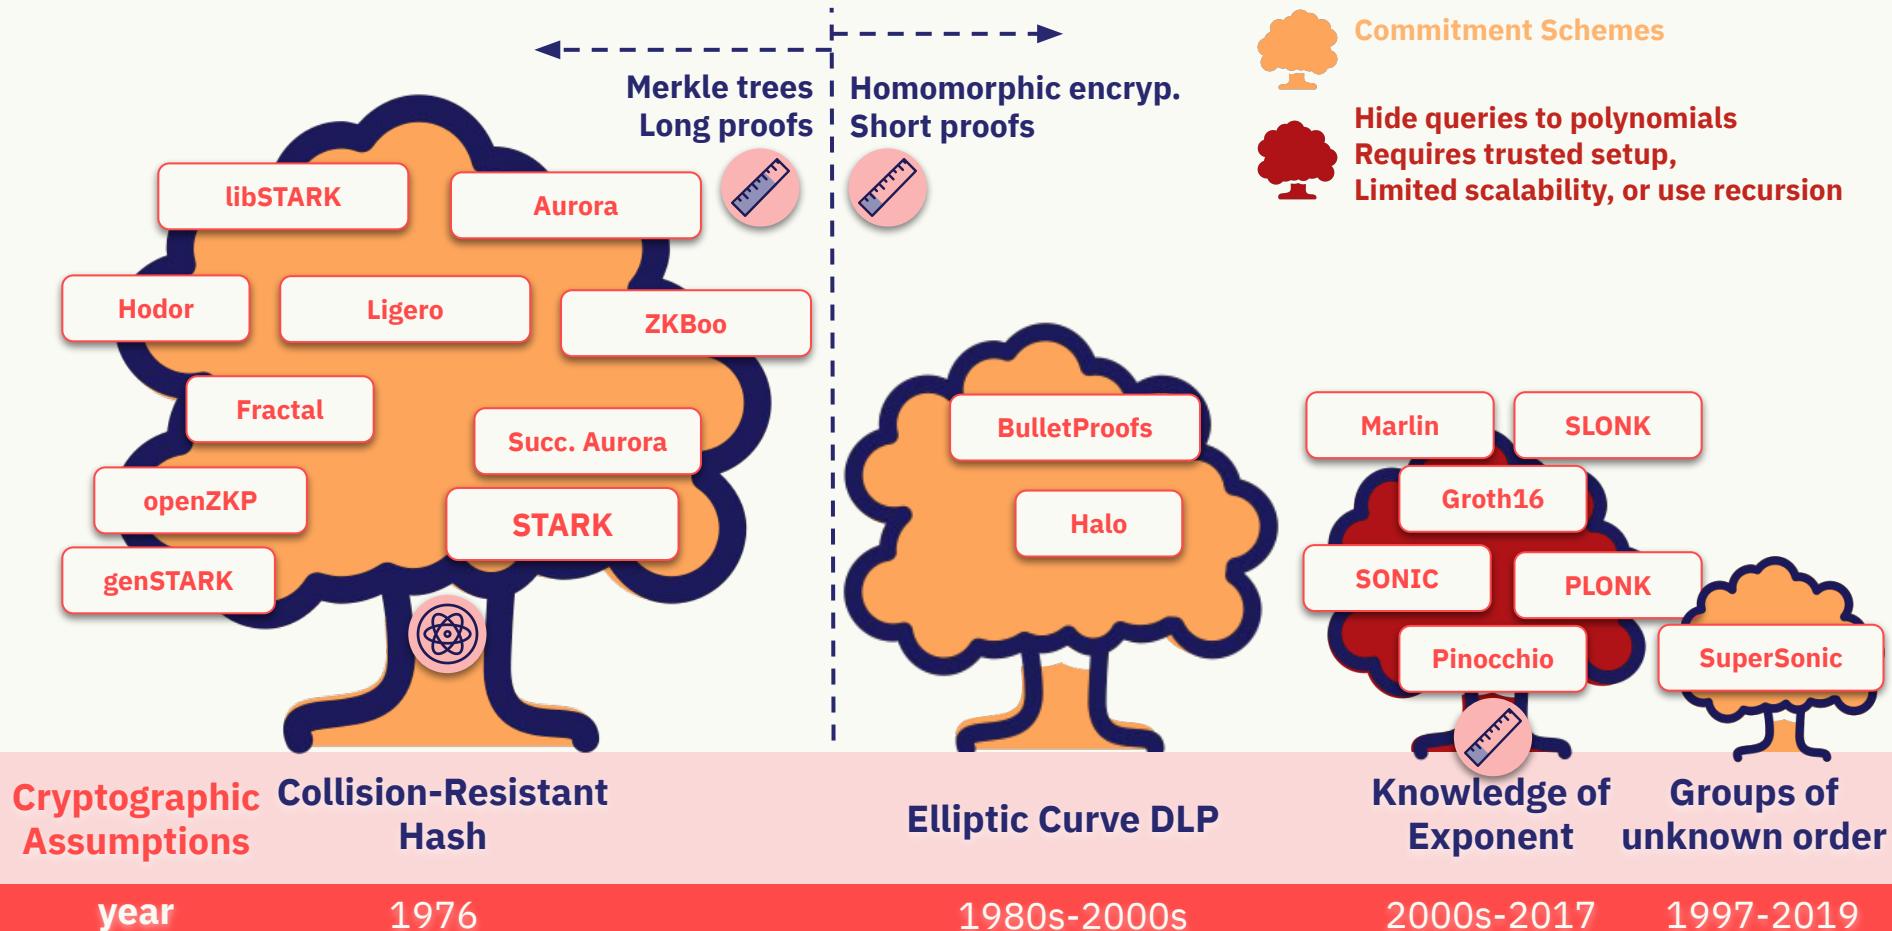
2. Enforcing low-degreeness



2. Enforcing low-degreeness



2. Enforcing low-degreeness



2. Enforcing low-degreeeness

Polynomial Commitment Scheme (PCS) [Field \mathbb{F} , degree d]

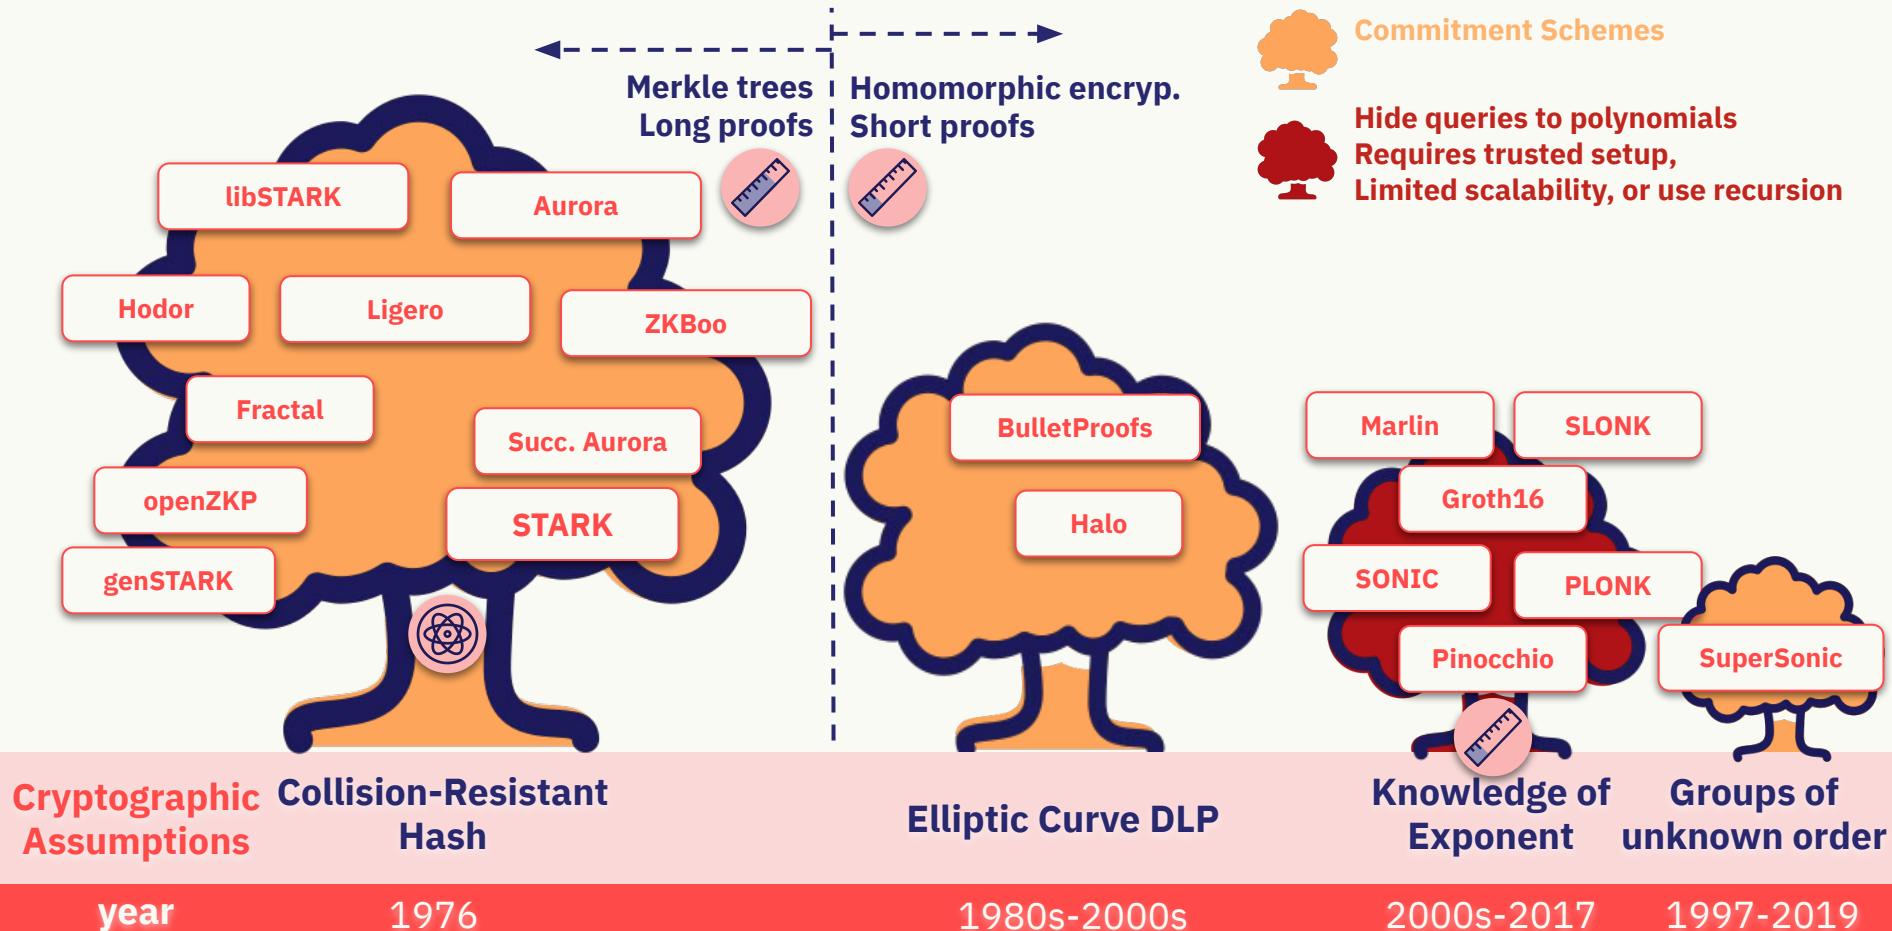
- Prover sends $c = \text{Commit}(P(x))$, $\deg(P) < d$
- Verifier queries $z \in \mathbb{F}$
- Prover answers $a \in \mathbb{F}$, claiming “ $\deg(P) < d$ and $P(z) = a$ ”
- Both parties interact; at end, verifier decides whether to accept/reject claim
- Want
 - **Completeness:** If $P(z) = a$ then Verifier accepts
 - **Soundness:** If $P(z) \neq a$ then whp Verifier rejects

2. Enforcing low-degreeeness

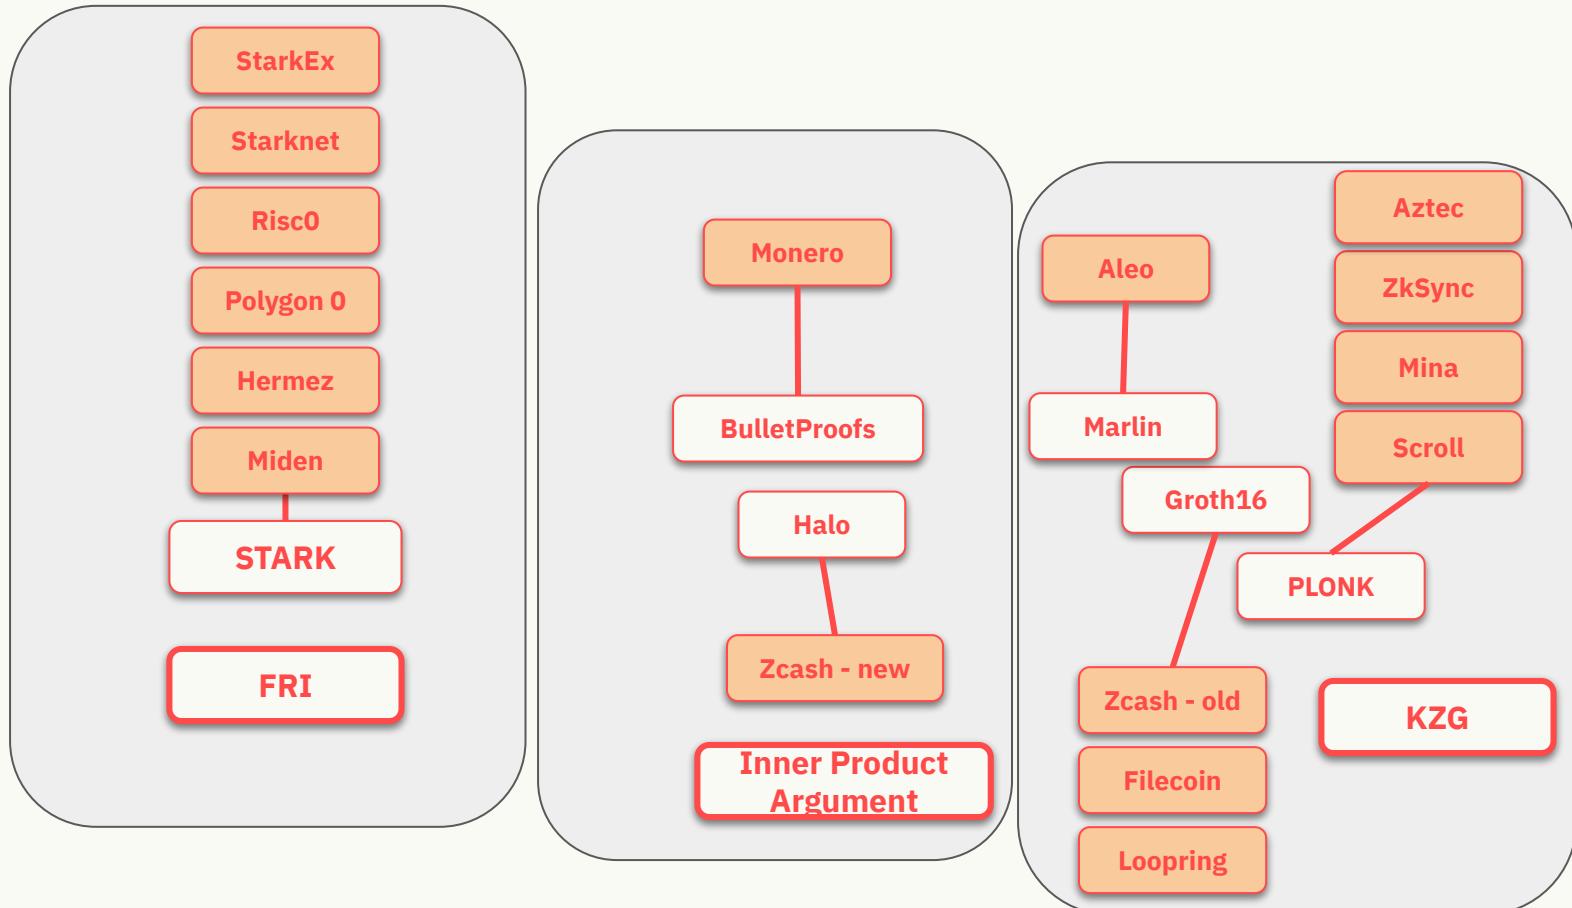
Polynomial Commitment Scheme (PCS) [Field \mathbb{F} , degree d]

- Prover sends $c = \text{Commit}(P(x))$, $\deg(P) < d$
- Verifier queries $z \in \mathbb{F}$
- Prover answers $a \in \mathbb{F}$, claiming “ $\deg(P) < d$ and $P(z) = a$ ”
- Both parties interact; at end, verifier decides whether to accept/reject claim
- Want
 - **Completeness:** If $P(z) = a$ then Verifier accepts
 - **Soundness:** If $P(z) \neq a$ then whp Verifier rejects
 - **Knowledge soundness:** efficient extractor can recover $P(X)$ from good prover
 - **Efficiency:** low proving time, comm complexity, verification time, over all fields, ...
 - **Succinctness:** polylogarithmic verification time (and communication)
 - **Security:** minimal crypto assumptions

2. Enforcing low-degreeness



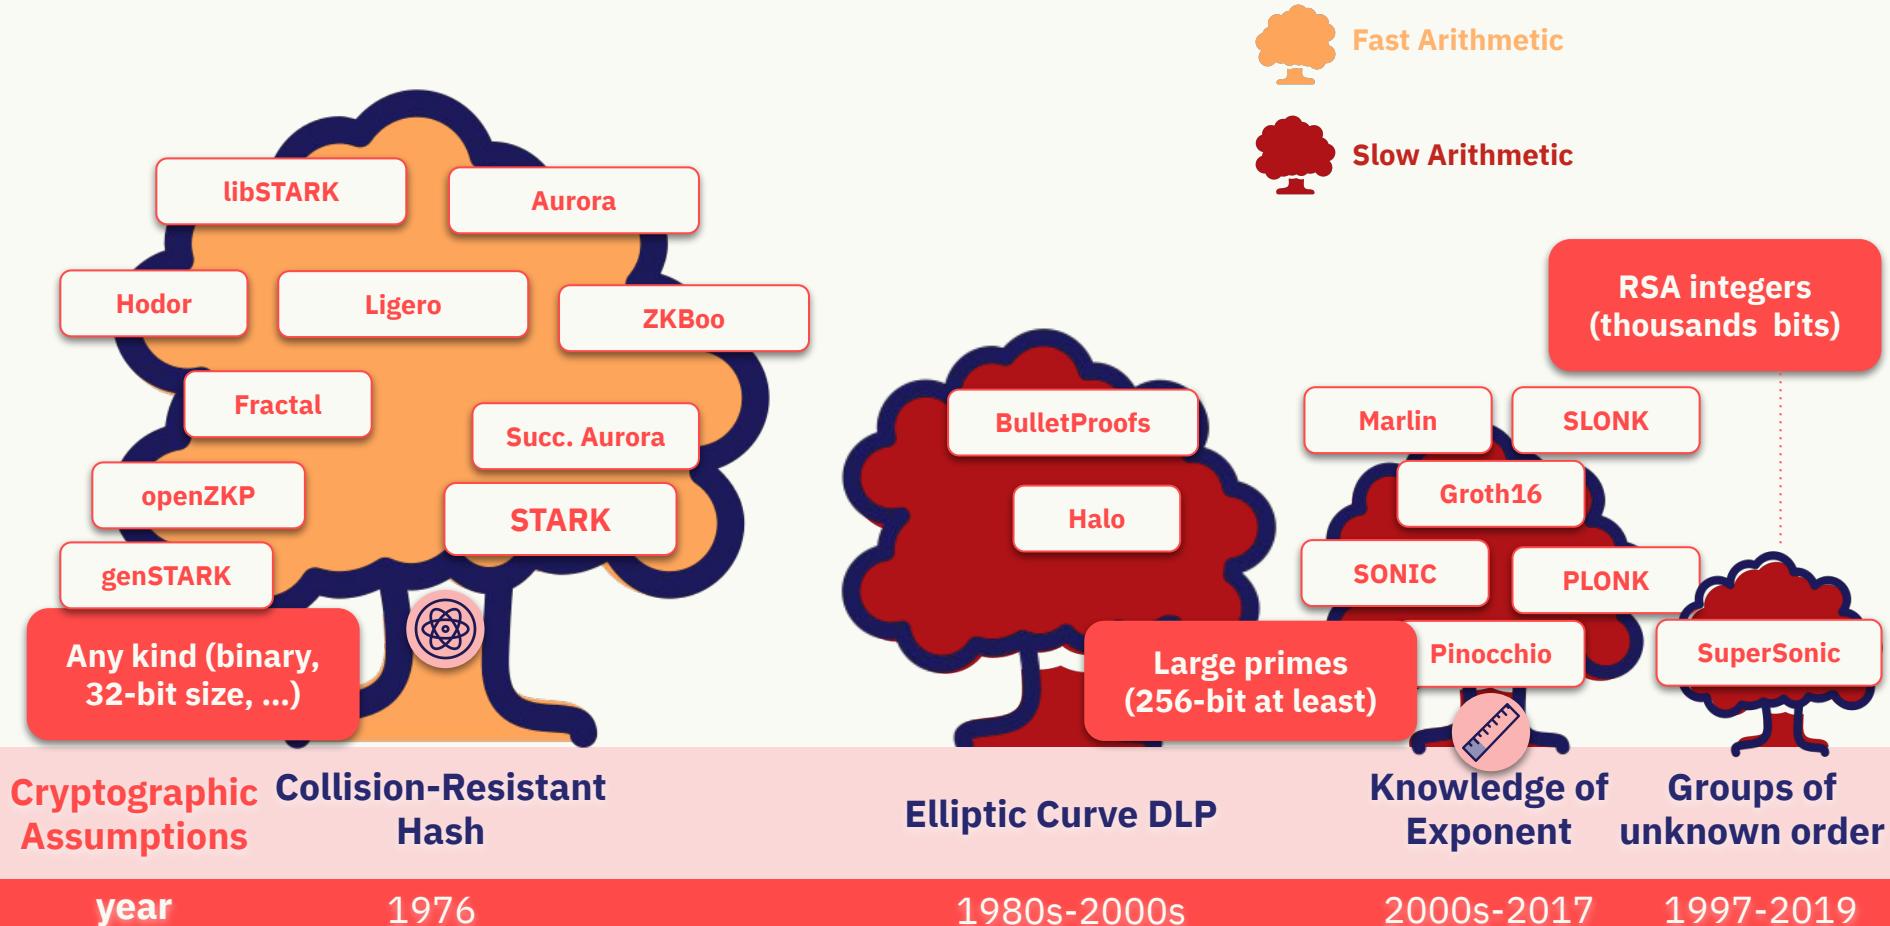
2. Enforcing low-degreeness



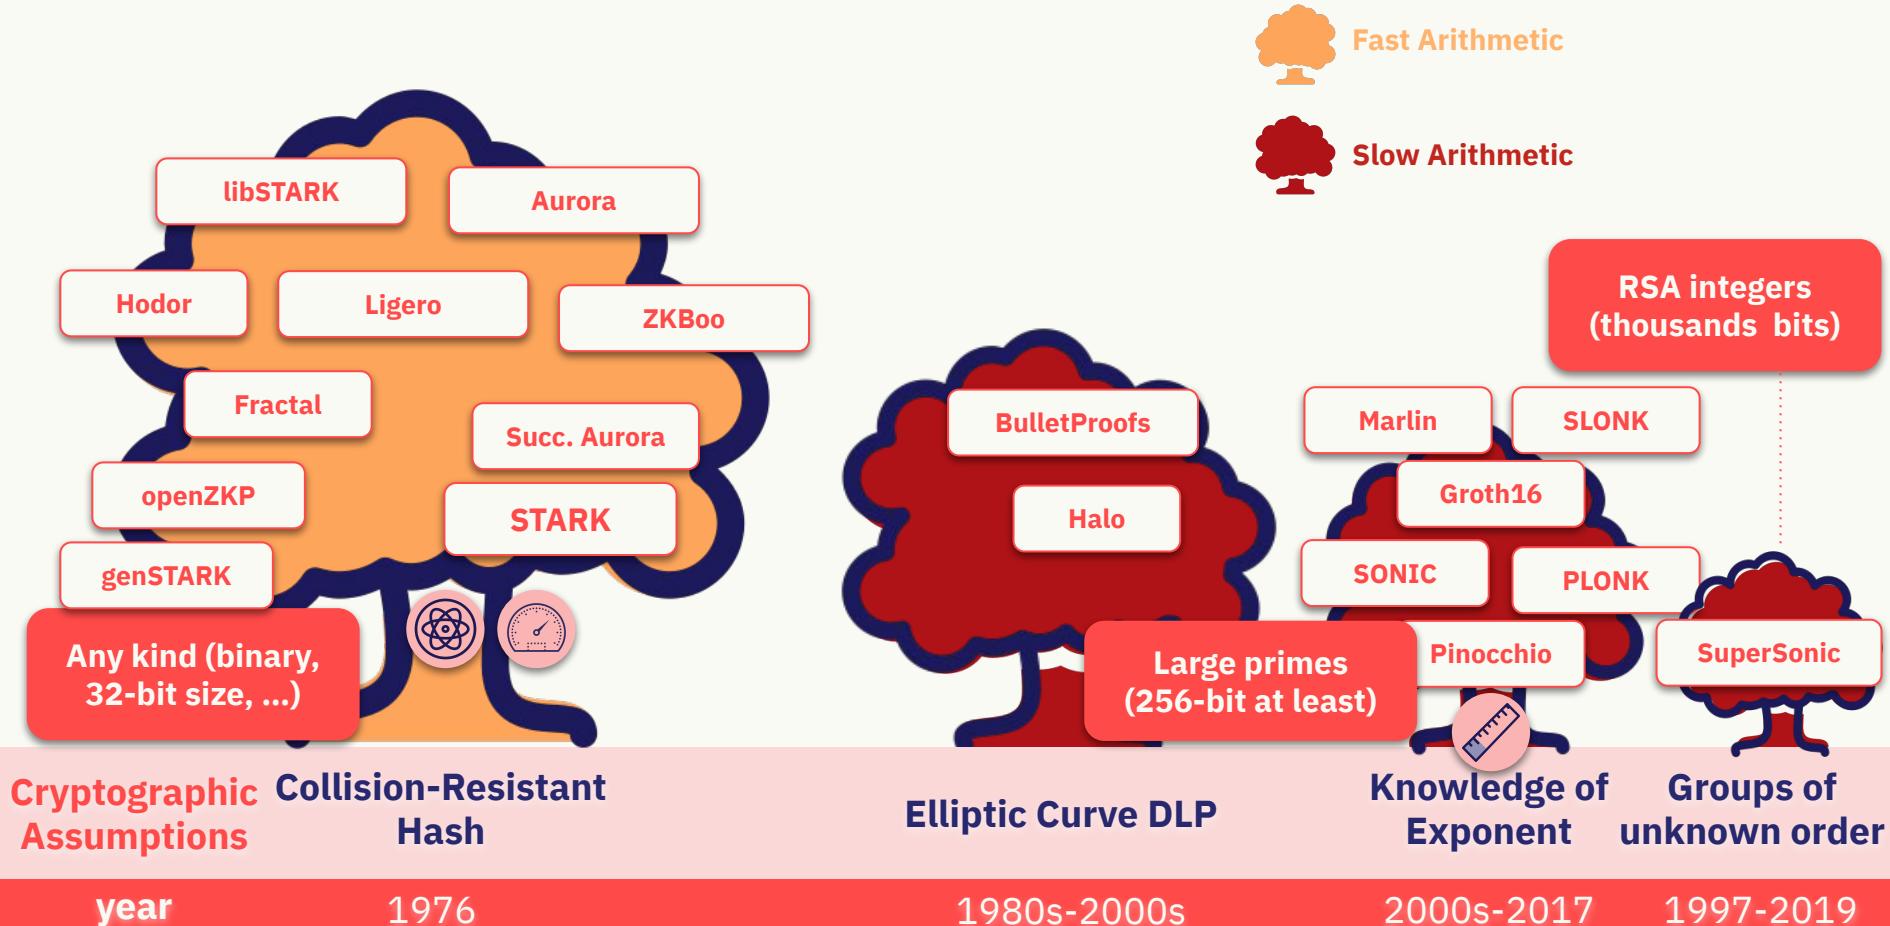
2. Enforcing low-degreeness

	FRI	Inner Product Arguments	KZG
Pros	<ul style="list-style-type: none">- Succinct verification- Succinct setup- Transparent- Post quantum secure- Works over all fields	<ul style="list-style-type: none">- Transparent- Short proofs (KBs)- Additivity	<ul style="list-style-type: none">- Very short pf (<1KB)- Additivity
Cons	<ul style="list-style-type: none">- Long proofs (dozens KB)	<ul style="list-style-type: none">- Linear time verifier- Quantum breakable	<ul style="list-style-type: none">- Trusted setup- Linear size/time setup- Quantum breakable
Assumptions	<ul style="list-style-type: none">- Collision resistant hash	<ul style="list-style-type: none">- Discrete log hardness	<ul style="list-style-type: none">- Knowledge exponent

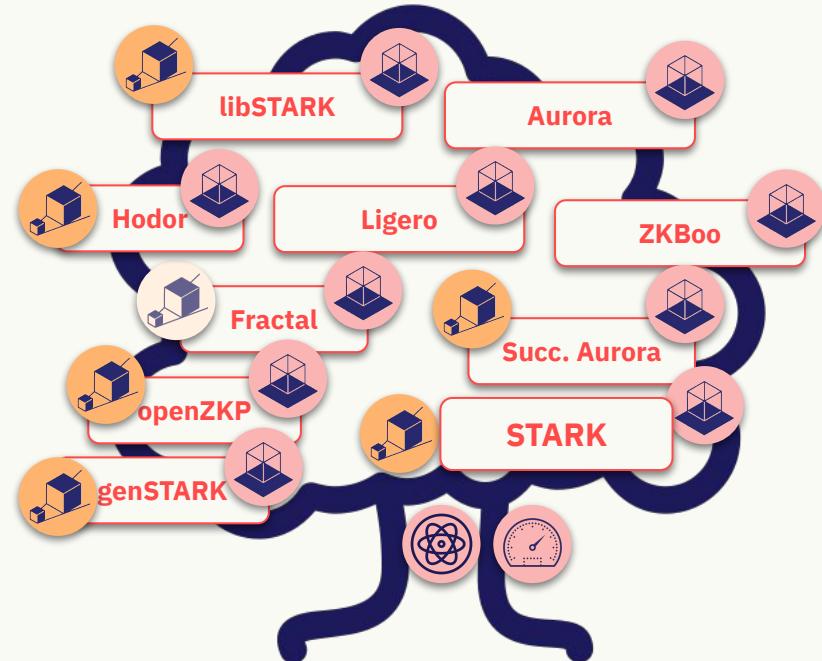
1. Arithmetization - finite field type



1. Arithmetization - finite field type



Scalability and Transparency



Cryptographic Assumptions Collision-Resistant Hash

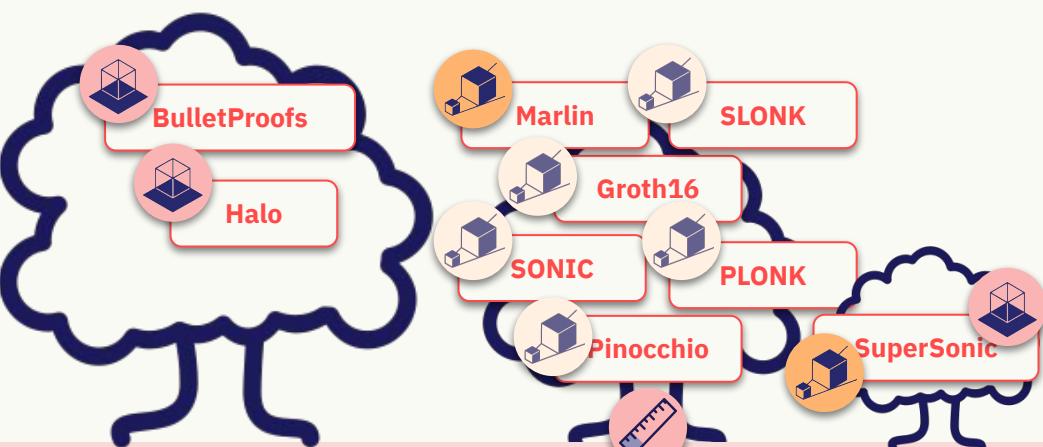
year

1976

1980s-2000s

2000s-2017

1997-2019



Elliptic Curve DLP

Knowledge of Exponent Groups of unknown order



Transparent



Scalable

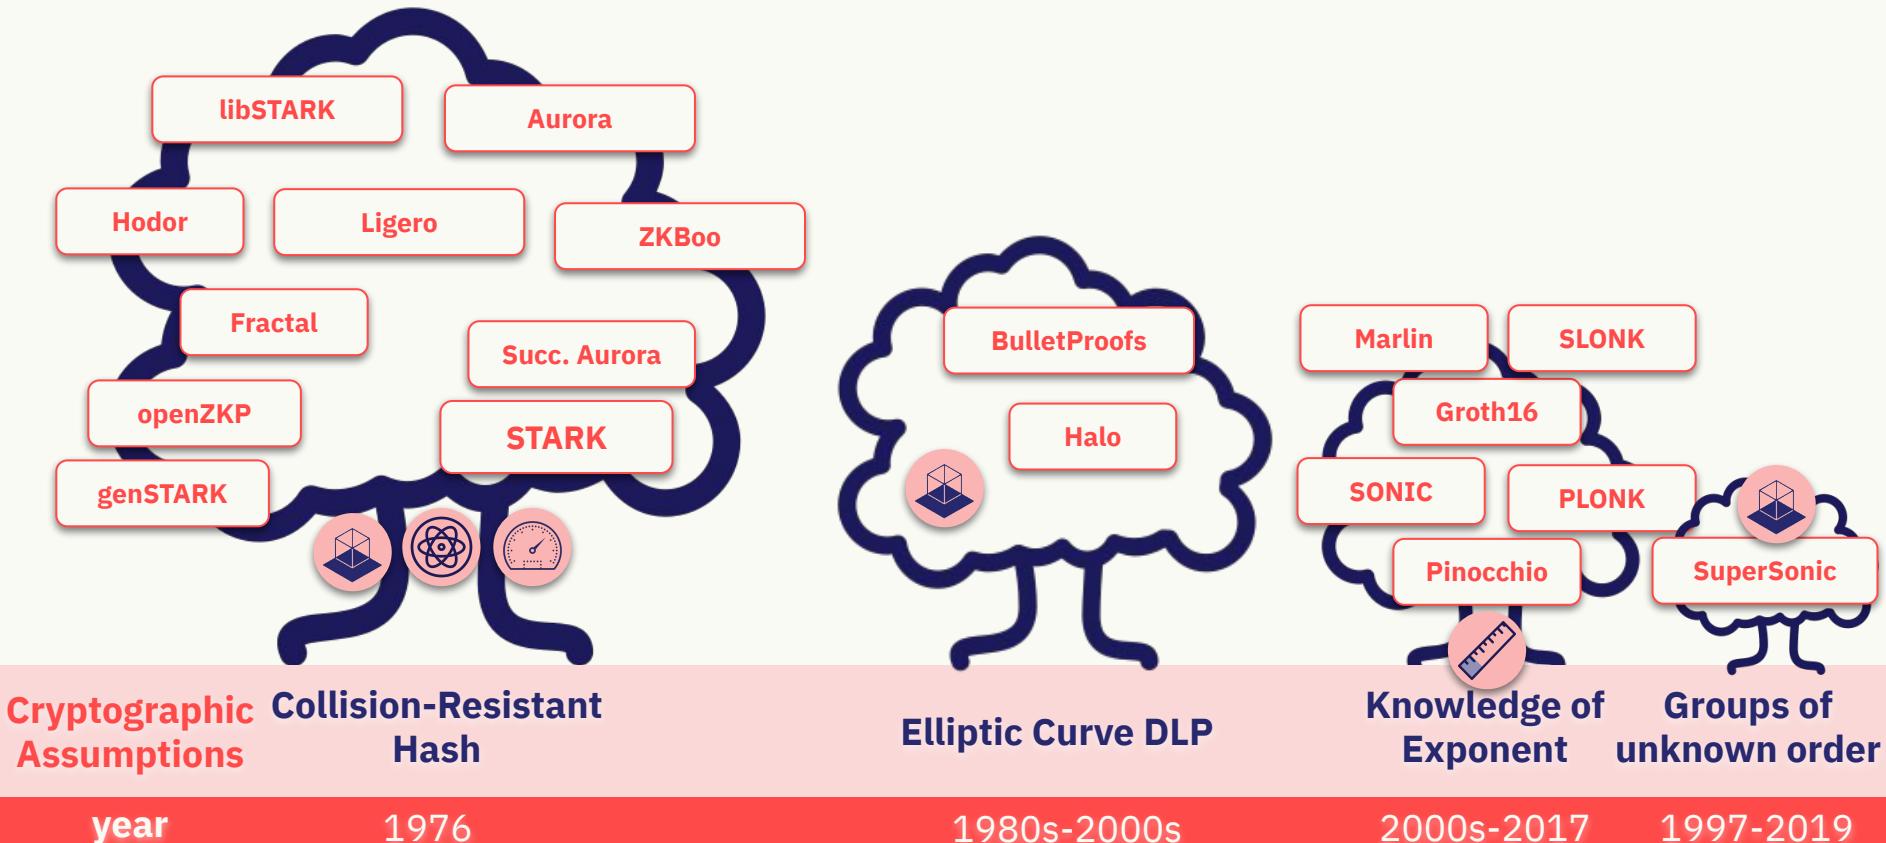


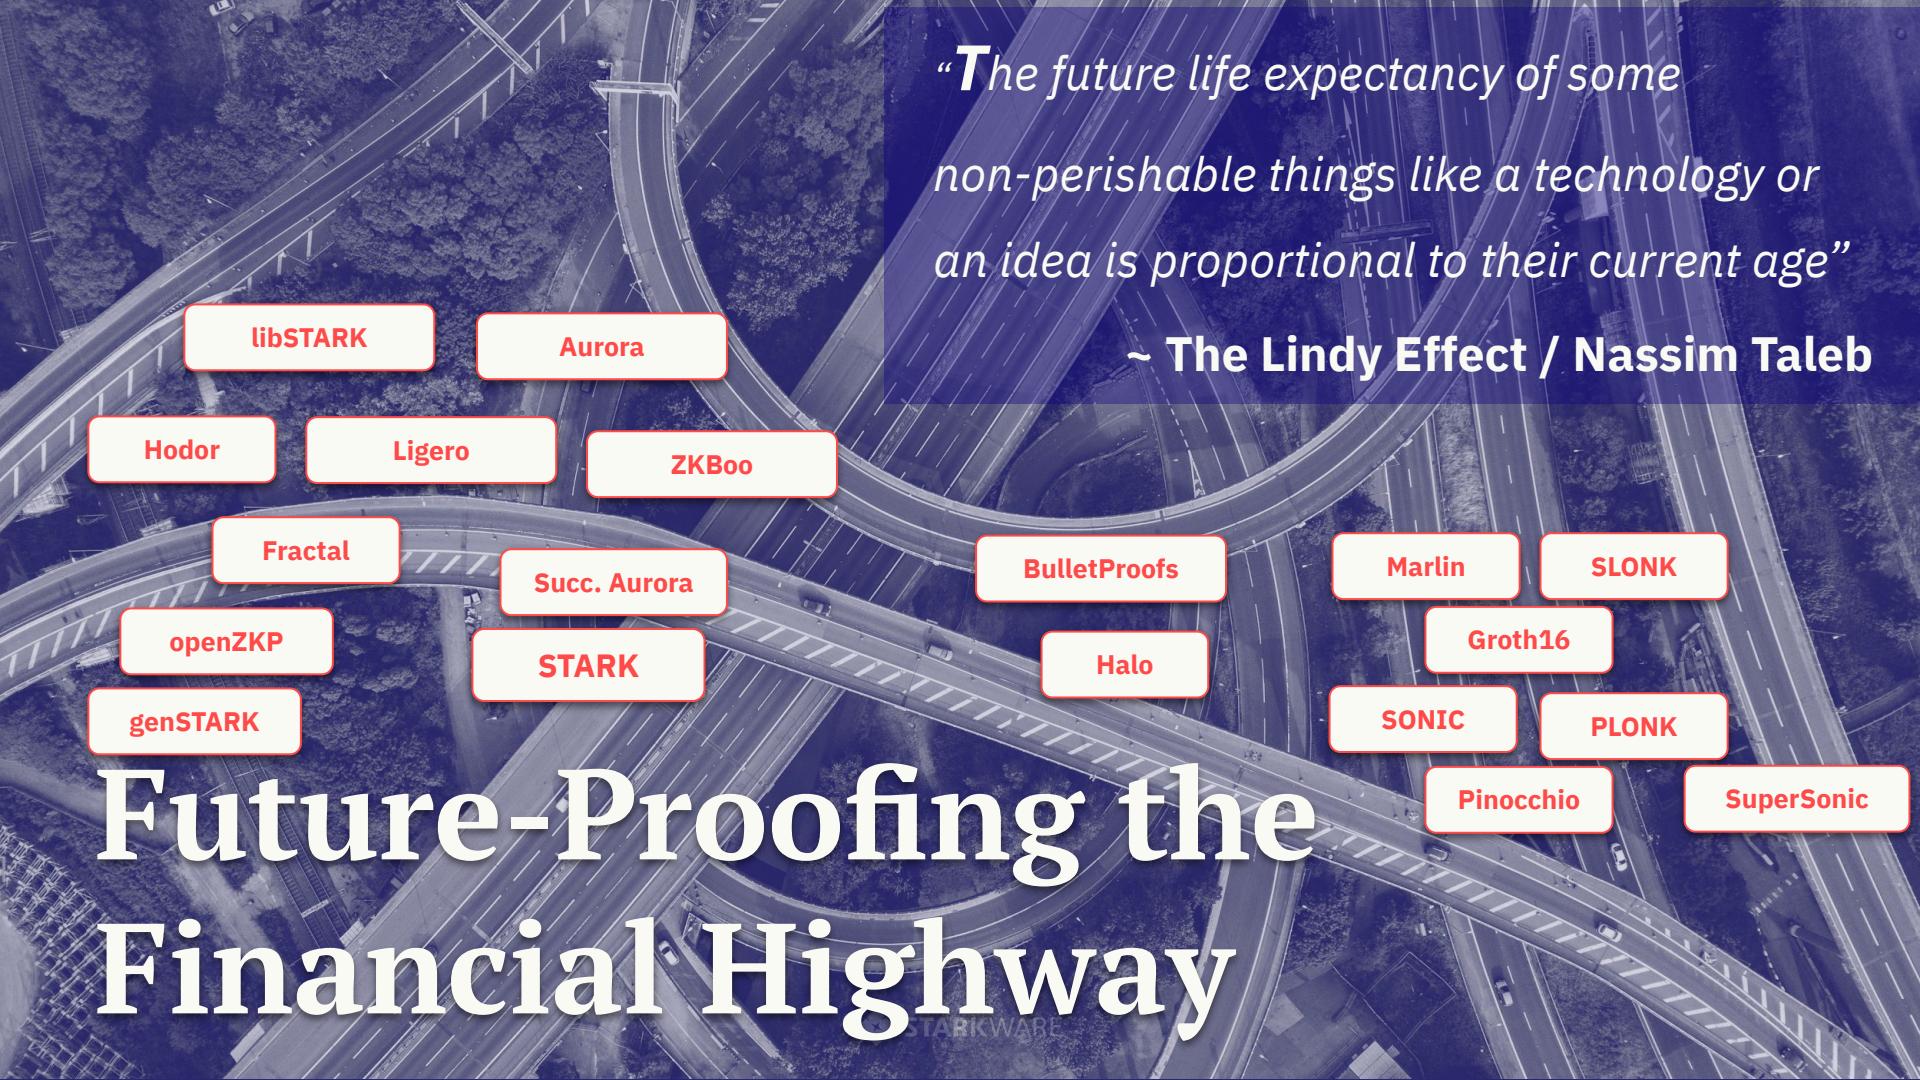
Semi-Scalable
(after linear pre-processing)

Scalability and Transparency



Transparent





“The future life expectancy of some non-perishable things like a technology or an idea is proportional to their current age”

~ The Lindy Effect / Nassim Taleb

libSTARK

Aurora

Hodor

Ligero

ZKBoo

Fractal

Succ. Aurora

openZKP

STARK

genSTARK

BulletProofs

Marlin

SLONK

Halo

Groth16

SONIC

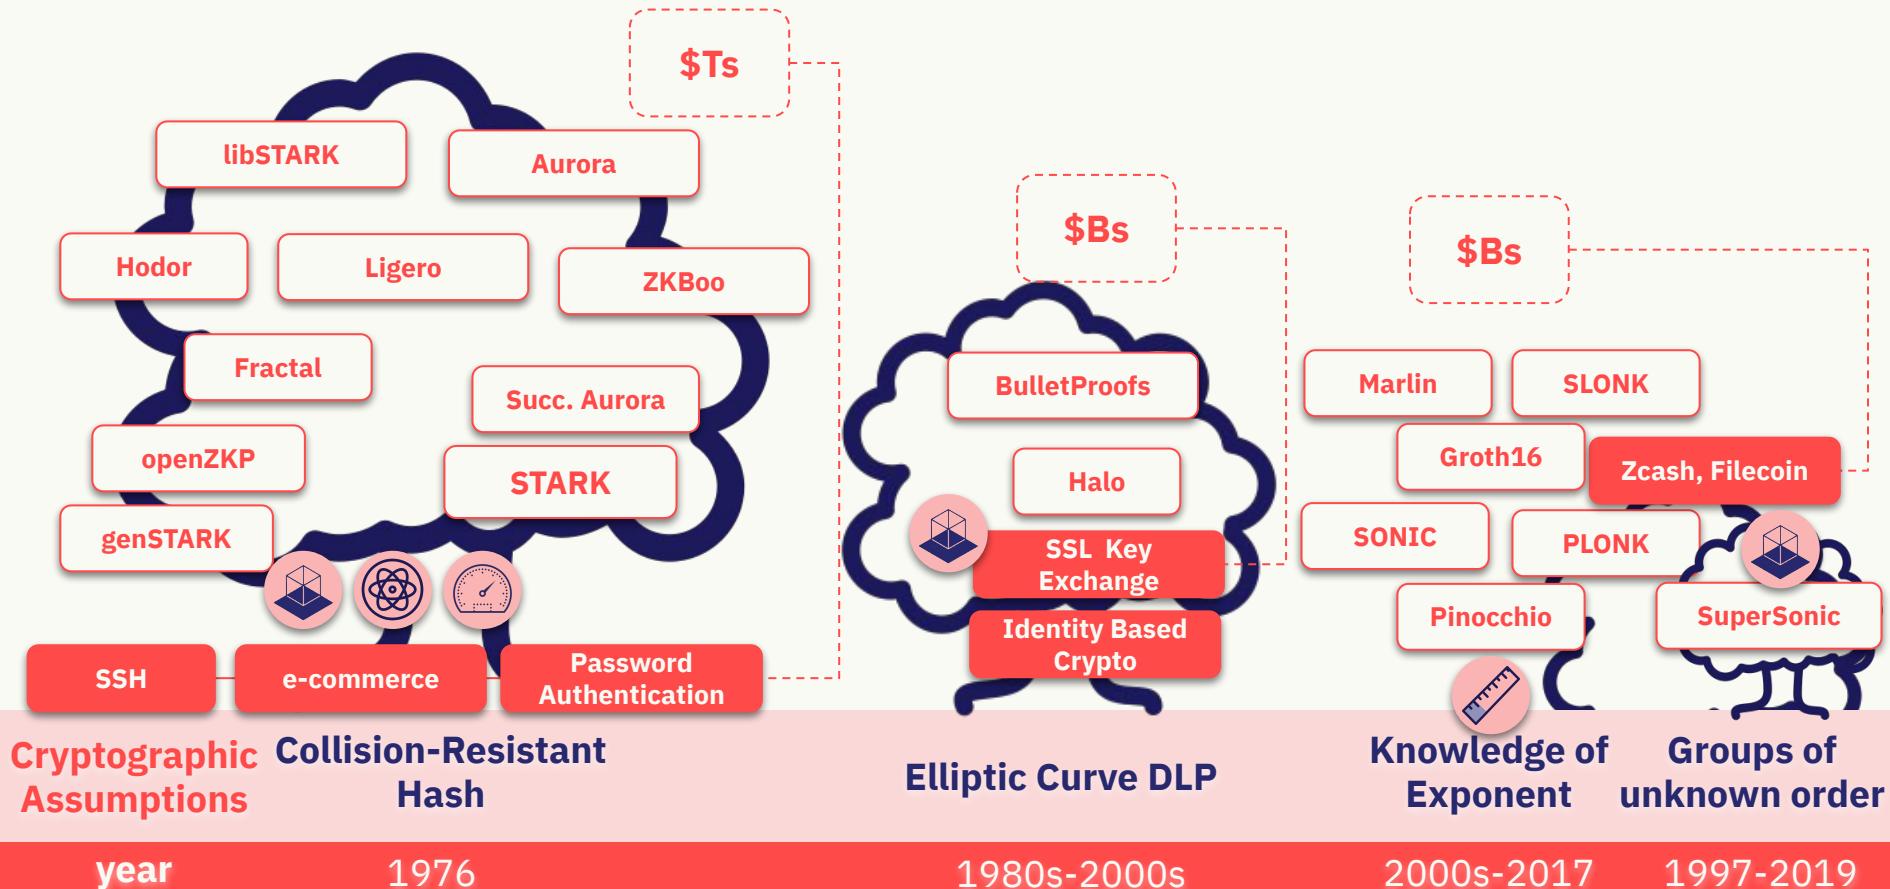
PLONK

Pinocchio

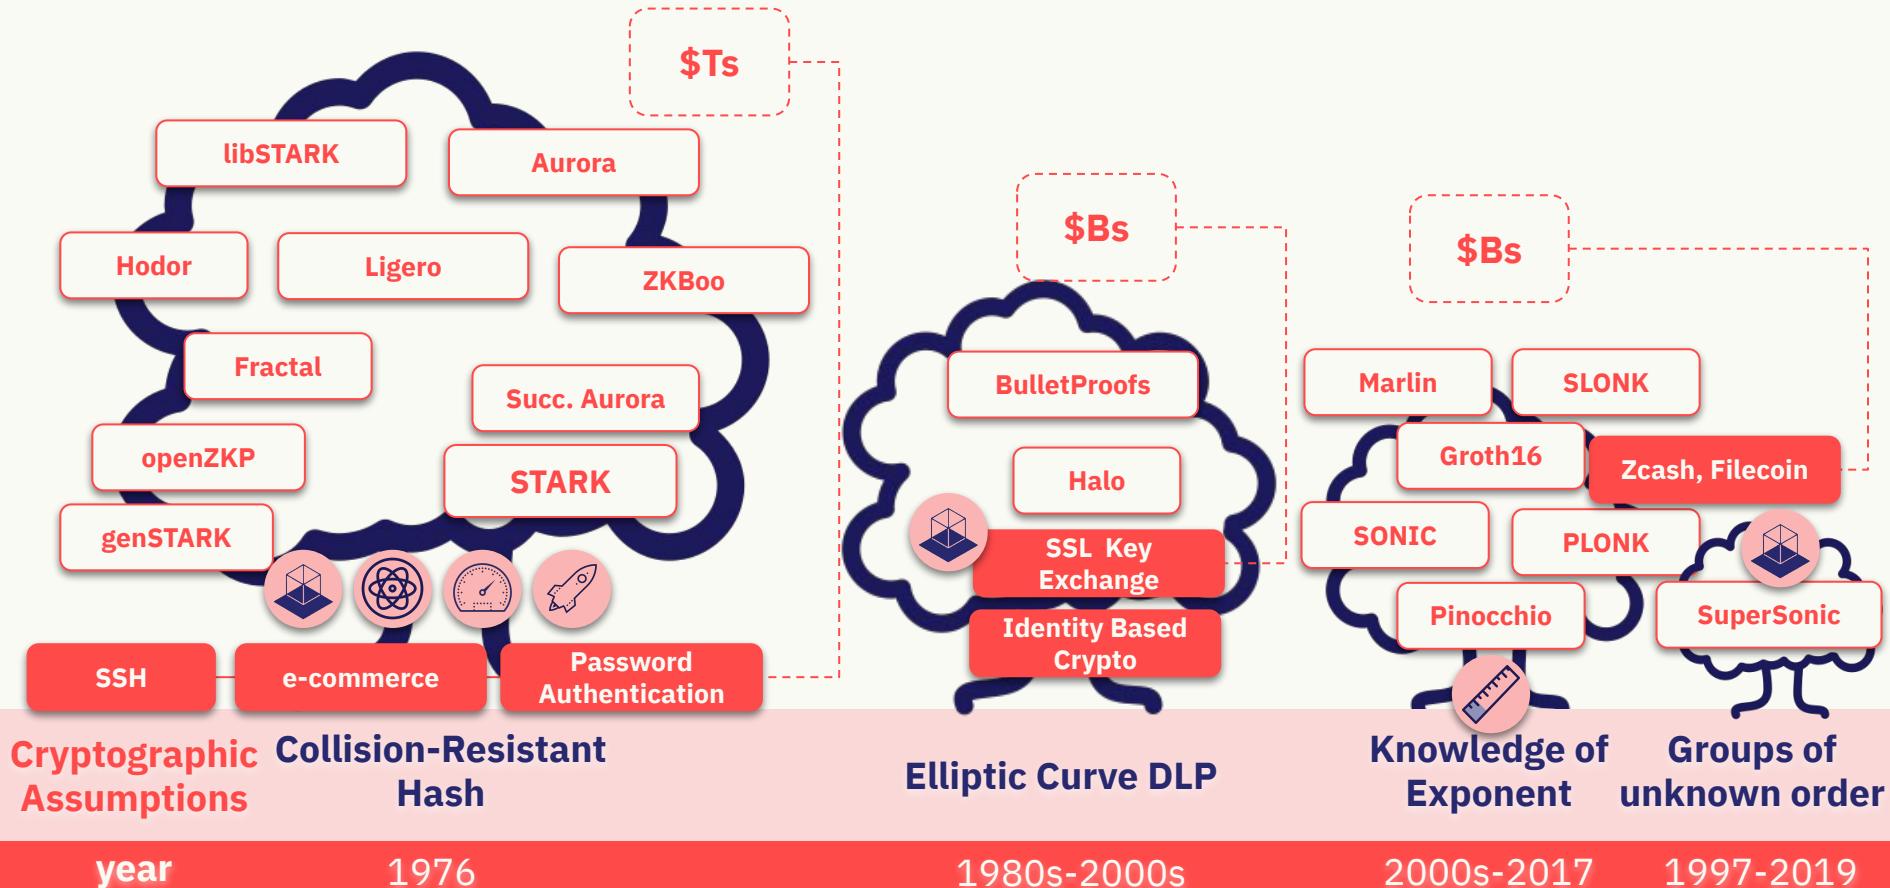
SuperSonic

Future-Proofing the Financial Highway

ZKP Family Trees



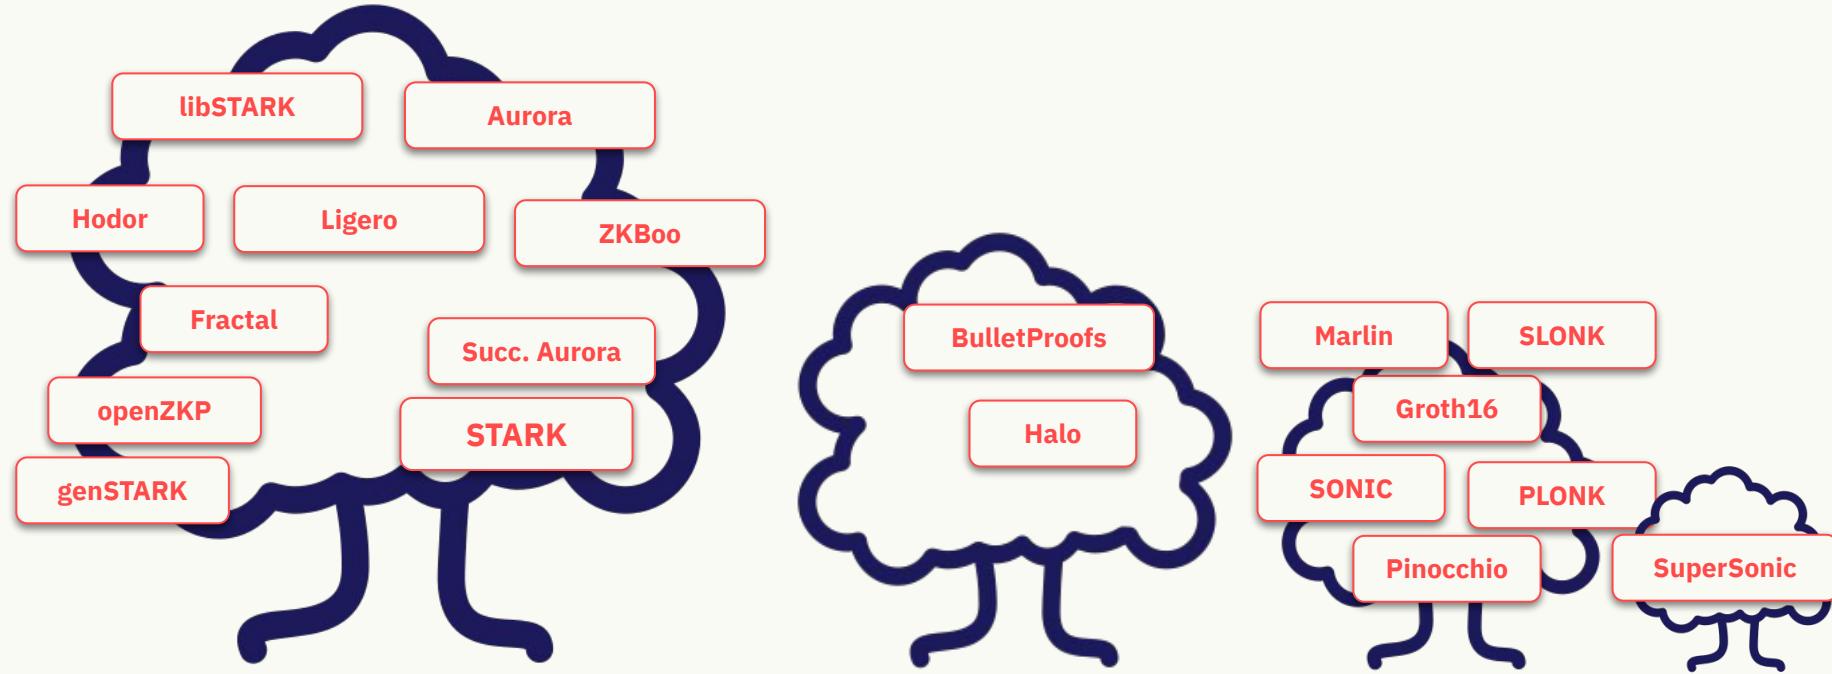
ZKP Family Trees



Summary

ZKP Cambrian explosion ongoing, expect more science!

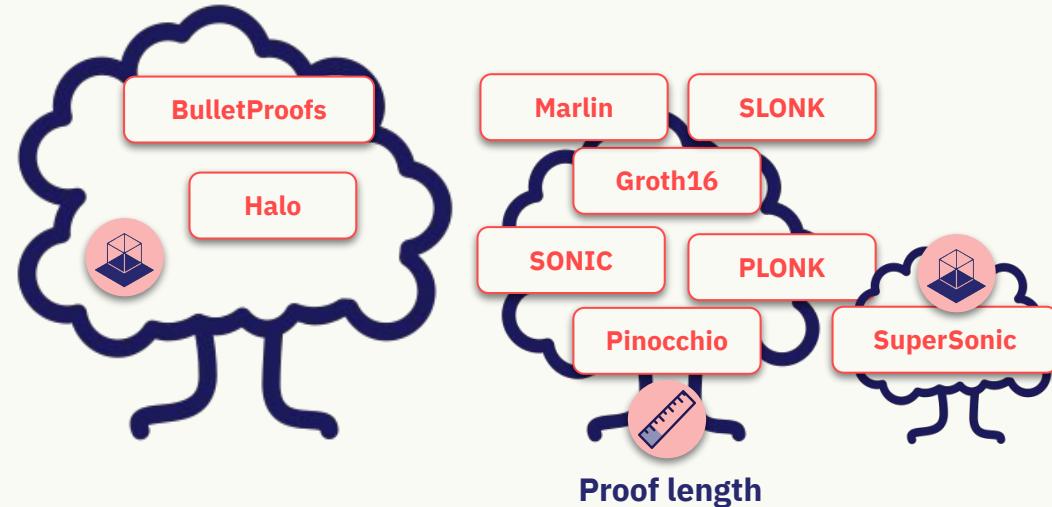
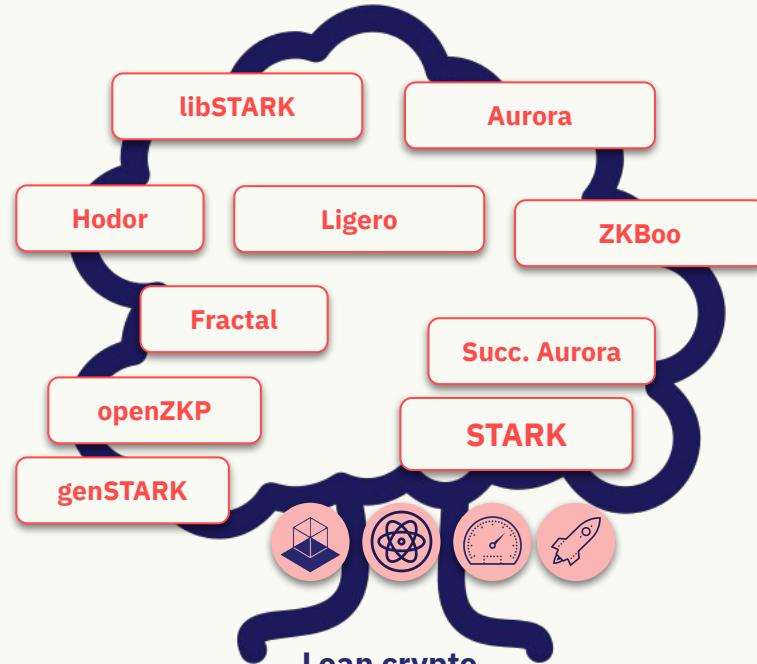
ZKP members differ by (i) arithmetization, (ii) low-degreeness, and (iii) crypto assumptions



Summary

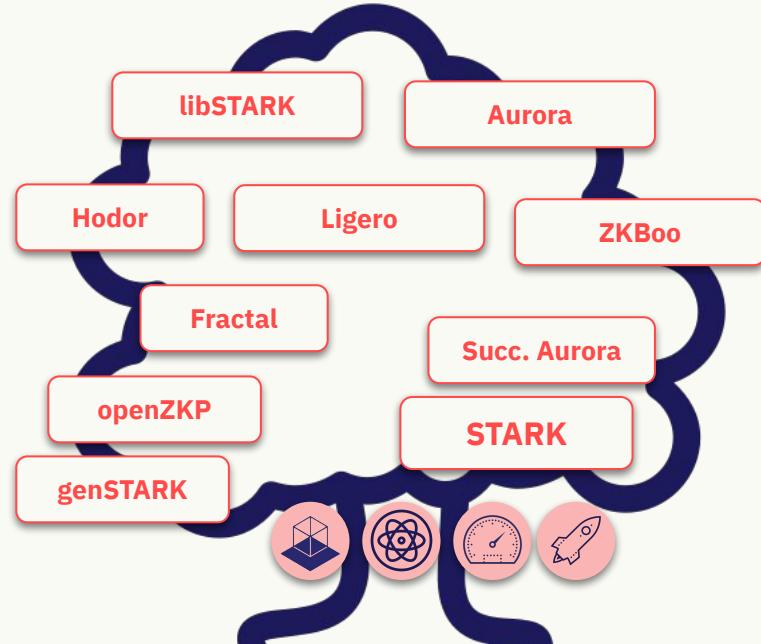
ZKP Cambrian explosion ongoing, expect more science!

ZKP members differ by (i) arithmetization, (ii) low-degreeness, and (iii) crypto assumptions



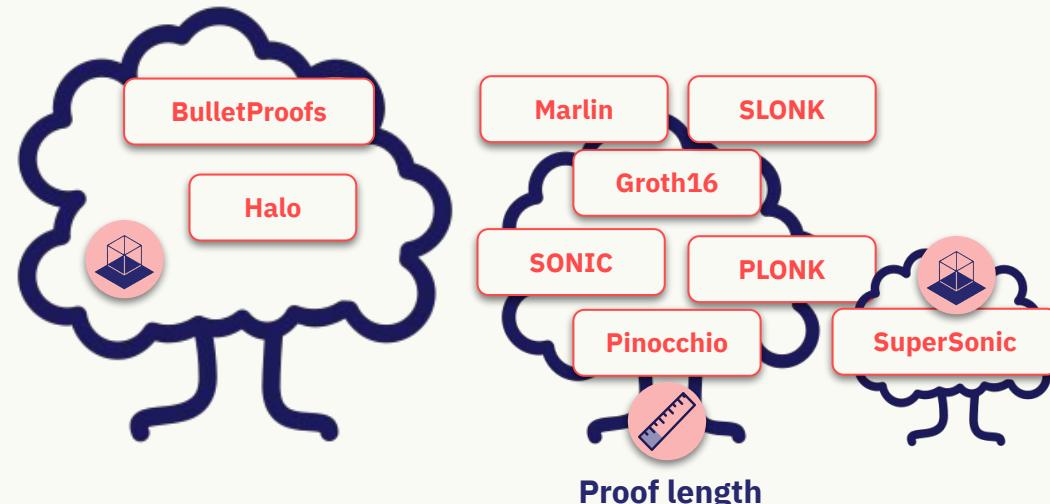
Summary

ZKP Cambrian explosion ongoing, expect more science!



ZKP members differ by (i) arithmetization, (ii) low-degreeness, and (iii) crypto assumptions

For short proofs, use **Groth16 SNARKs**.
For everything else, there's **STARKs!**



The End

November 2019



The Cambrian Explosion of ZKPs

February 2023

The Cambrian Explosion of ZKPs

T

libSTARK

Aurora

T

Hodor

Ligero

ZKBoo

N

Fractal

T

Succ. Aurora

T

openZKP

T

STARK

T

genSTARK

Awesome Stuff!

Awesome Stuff!

N

SNARK

T

STARK

Awesome Stuff!

Awesome Stuff!

BulletProofs

N

Halo

N

Marlin

N

SLONK

N

Groth16

N

SONIC

N

PLONK

N

Pinocchio

T

SuperSonic

Awesome Stuff!

Proofs of Computational Integrity (CI)



Privacy (Zero Knowledge, ZK)

Prover's private inputs are shielded



Scalability

Exponentially small verifier running time*

Nearly linear prover running time*



Universality

Applicability to general computation



Transparency

No toxic waste (i.e. no trusted setup)



Lean & Battle-Hardened Cryptography

e.g. post-quantum secure

*With respect to size of computation

STARK vs. SNARK - emphasizing different aspects

T

STARKs must be

Transparent: no trusted setup

Scalable: logarithmic verifying time **and** nearly-linear proving time

Succinct setup, at most logarithmic time

N

SNARKs must be

Noninteractive: pf is single message (after preprocessing)

Succinct: logarithmic verifying time

Setup can take linear time (and more)

Non-interactive STARKs are SNARKs (transparent ones)

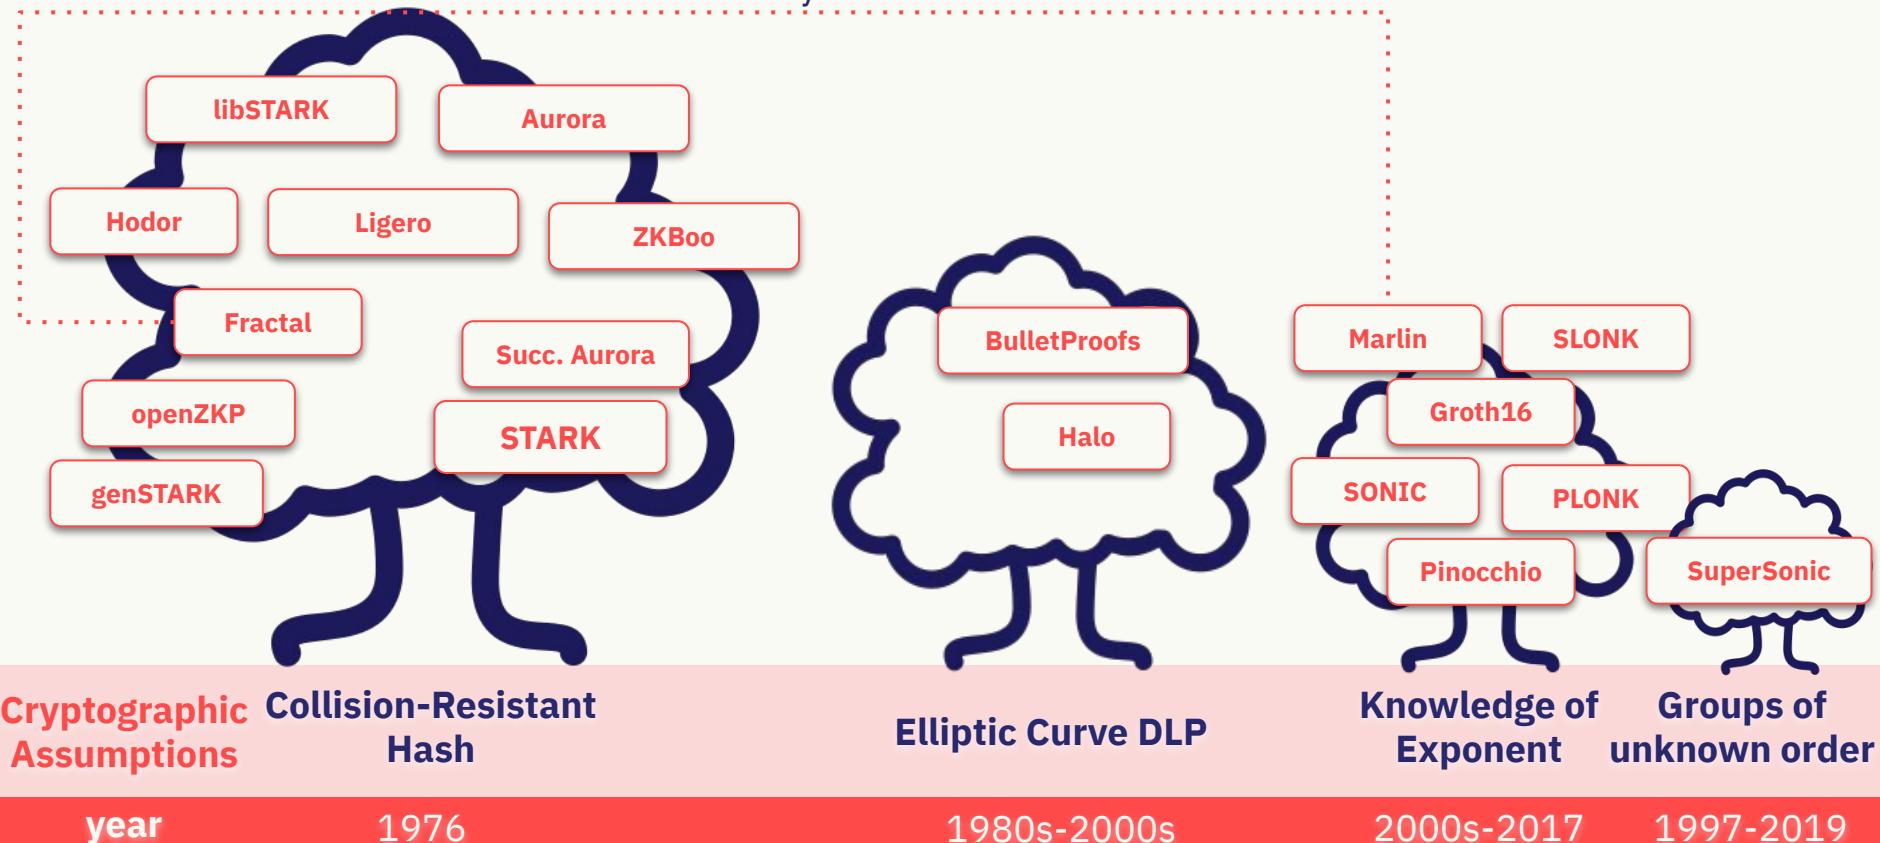
Transparent SNARKs w/ succinct setup are STARKs

The Cambrian Explosion of ZKPs



3. Cryptographic Assumptions

Note: systems can move across trees



STARK efficiency

- Arithmetization over any field
 - Initially over any “FFT-friendly” field, including small binary fields, small primes
 - Recently: over any field, using Elliptic curves [BCKL 2021-2]
- New Computational Model - IOP [RRR 2016; BCS 2016]
- Fast Reed-Solomon IOP of Proximity (FRI) [BBHR 2018]
 - Proving time is $O(n)$, small constants (6 or less)
 - Verification time is $O(\log n)$, small constants (20 or less)
 - Nearly no soundness loss till Johnson bound [BCIKS 2020]
 - Formally: if f is delta-far from RS code, then single query-phase ($\log n$ queries to f and IOP) rejects f w.p. at least $\min(\delta, 1 - \sqrt{\delta})$
 - Proof: relies on the Guruswami-Sudan list decoding algorithm