

From Crypto-Paper to Crypto-Currency: the Cardano Consensus Layer




Real-World Crypto Symposium (RWC) 2021

**Christian
Badertscher**
IOHK

Philipp Kant
IOHK

Duncan Coutts
IOHK
Well-Typed

Joint work with Peter Gaži, Aggelos Kiayias, Alexander Russell



INPUT | OUTPUT

Overview

1

From Theory to Practice

2

**Formal Methods and Implementation
Correctness**

3

What Could Have Gone Wrong

4

Path to Decentralisation

Overview

1

From Theory to Practice

2

Formal Methods and Implementation
Correctness


3

What Could Have Gone Wrong


4

Path to Decentralisation

Distributed Computation over a Network




Distributed Computation over a Network



Protocol execution model:

- Interactive machines

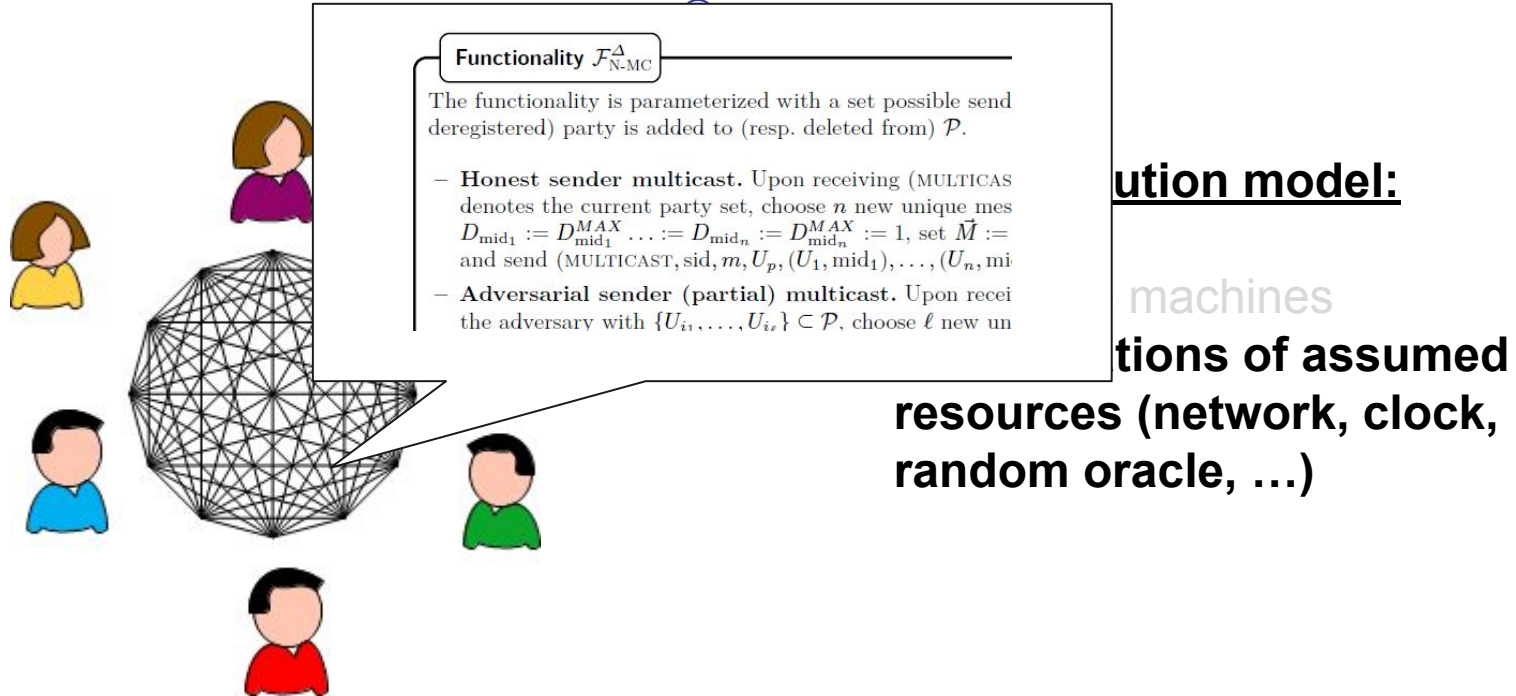
Distributed Computation over a Network



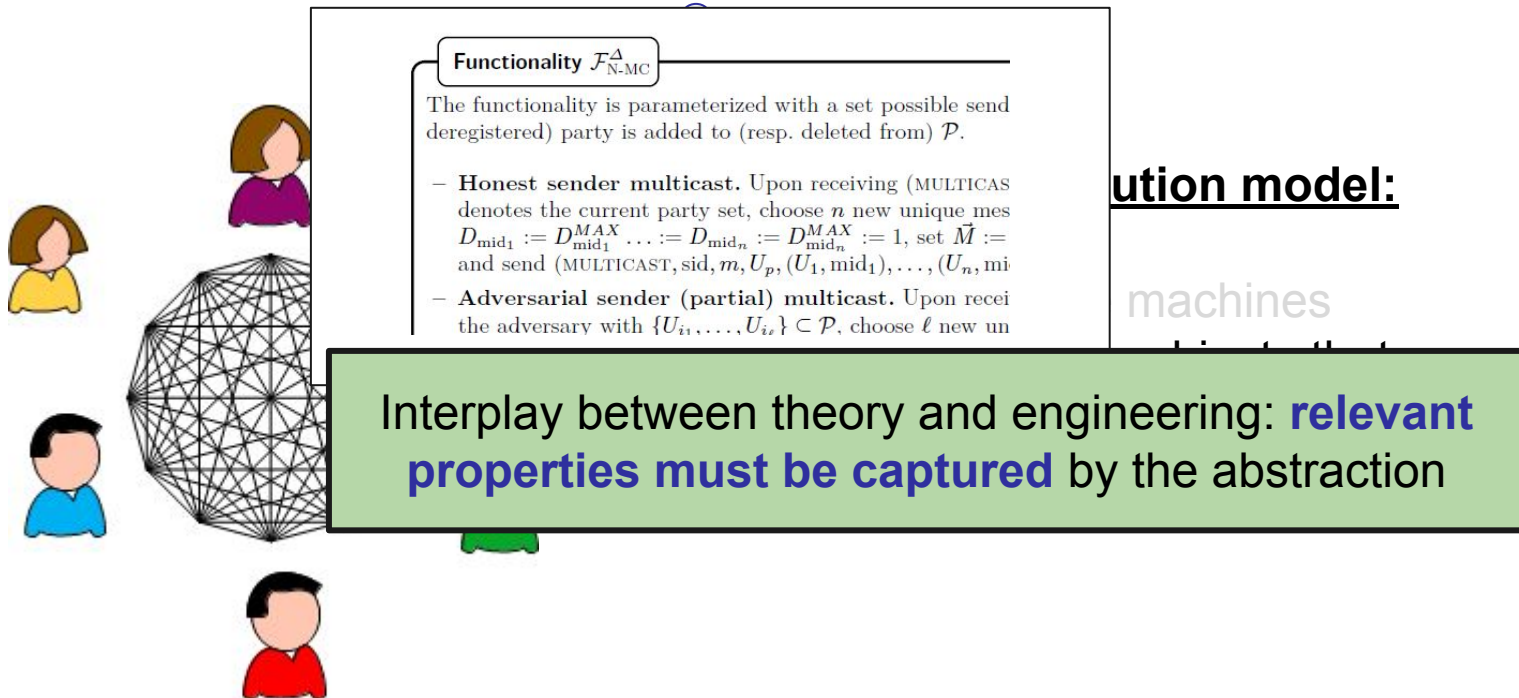
Protocol execution model:

- Interactive machines
- **Formalizations of assumed resources (network, clock, random oracle, ...)**


Distributed Computation over a Network



Distributed Computation over a Network




Distributed Computation over a Network



Protocol execution model:

- Interactive machines
- Formalizations of assumed resources (network, clock, random oracle, ...)
- **Byzantine behavior**


Distributed Computation over a Network



Protocol execution model:

- Interactive machines
- Formalizations of assumed resources (network, clock, random oracle, ...)
- Byzantine behavior
- **Rational behavior**

Distributed Computation over a Network



Protocol execution model:

- Interactive machines
- Formalizations of assumed resources (network, clock, random oracle, ...)
- Byzantine behavior
- Rational behavior
- **Machine failures**

Distributed Computation over a Network



Fine-Grained security model

+ **Security under composition:**

Each external input might depend on the entire view of this and other protocols


⇒ **A security proof is meaningful to practice.**

Protocol execution model:


- Interactive machines
- Formalizations of assumed resources (network, clock, random oracle, ...)
- Byzantine behavior
- Rational behavior
- Machine failures

Consensus Layer of Cardano

Goal:



Realize



Ledger:

- Persistence
- Liveness

(Consensus on ledger state)


Result:

Ouroboros:

A Nakamoto-style **Proof-of-Stake Blockchain Protocol** realizing a ledger

Nakamoto-Style Blockchain:

Goal:



Genesis
Block


Ouroboros:

Result:

A Nakamoto-style **Proof-of-Stake
Blockchain Protocol** realizing a ledger

Nakamoto-Style Blockchain:

Goal:



Genesis
Block


- Lottery on next block proposer(s).
- Proposers elected proportional to owned stake.

Ouroboros:

Result:

A Nakamoto-style **Proof-of-Stake Blockchain Protocol** realizing a ledger

Nakamoto-Style Blockchain:




Security proof:

Establish **common-prefix, chain-growth, chain-quality** properties (under honest majority of stake assumption).

Ouroboros:

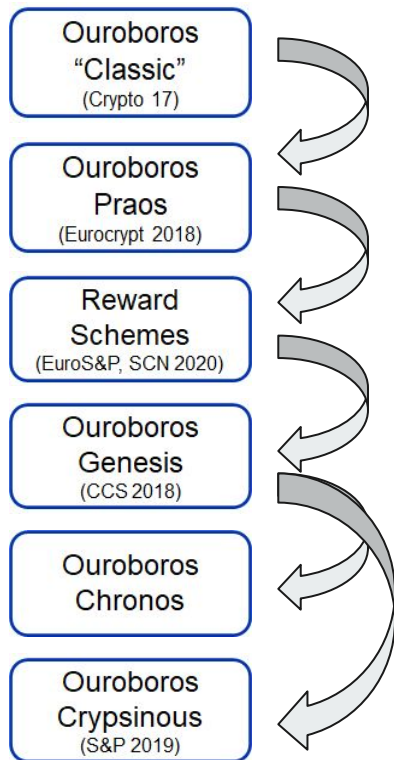
A Nakamoto-style **Proof-of-Stake Blockchain Protocol** realizing a ledger

Goal:



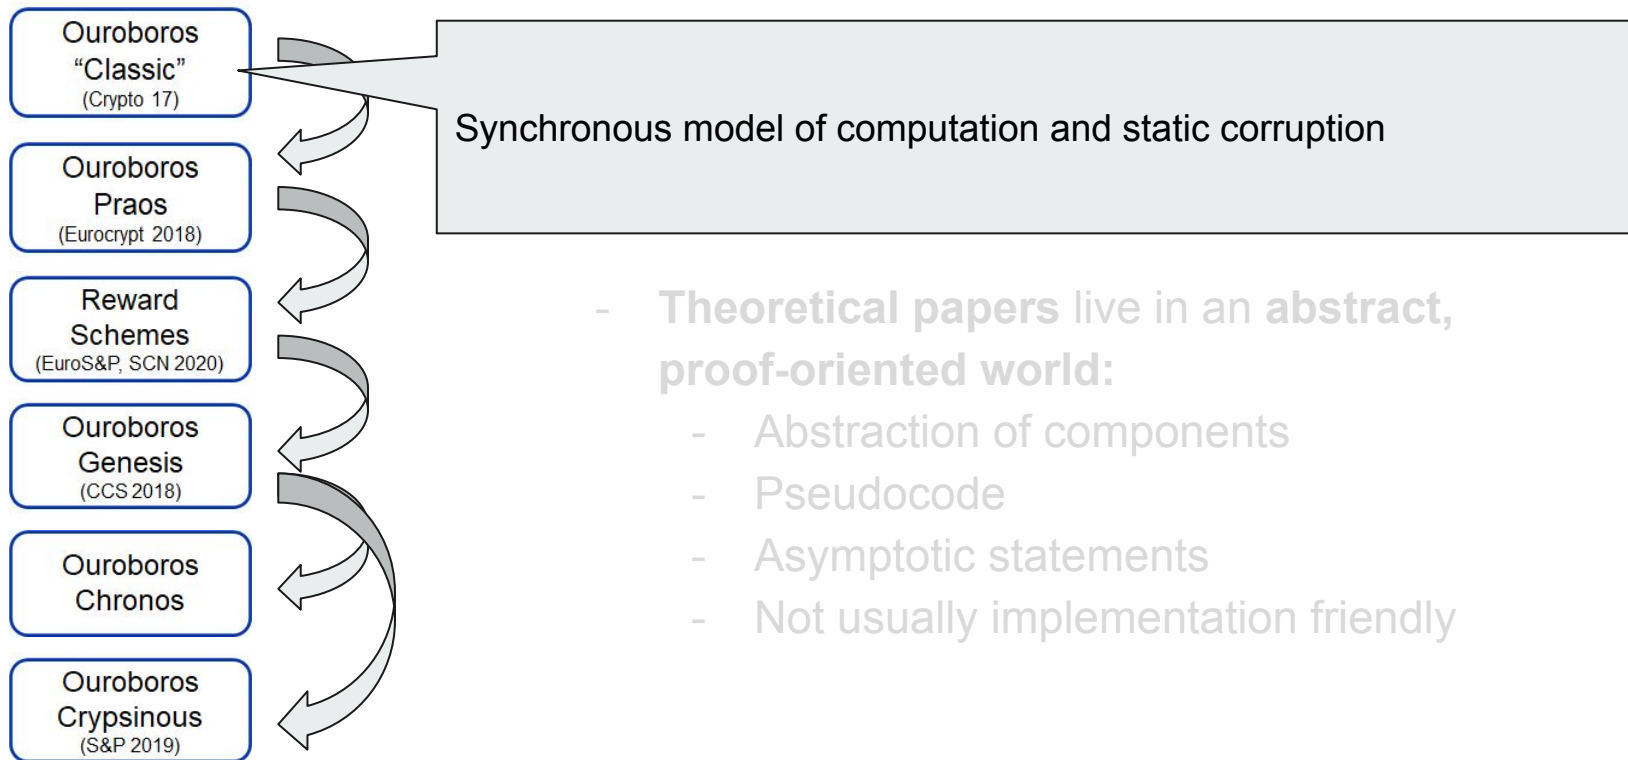
Result:

Theory + Practice Interplay

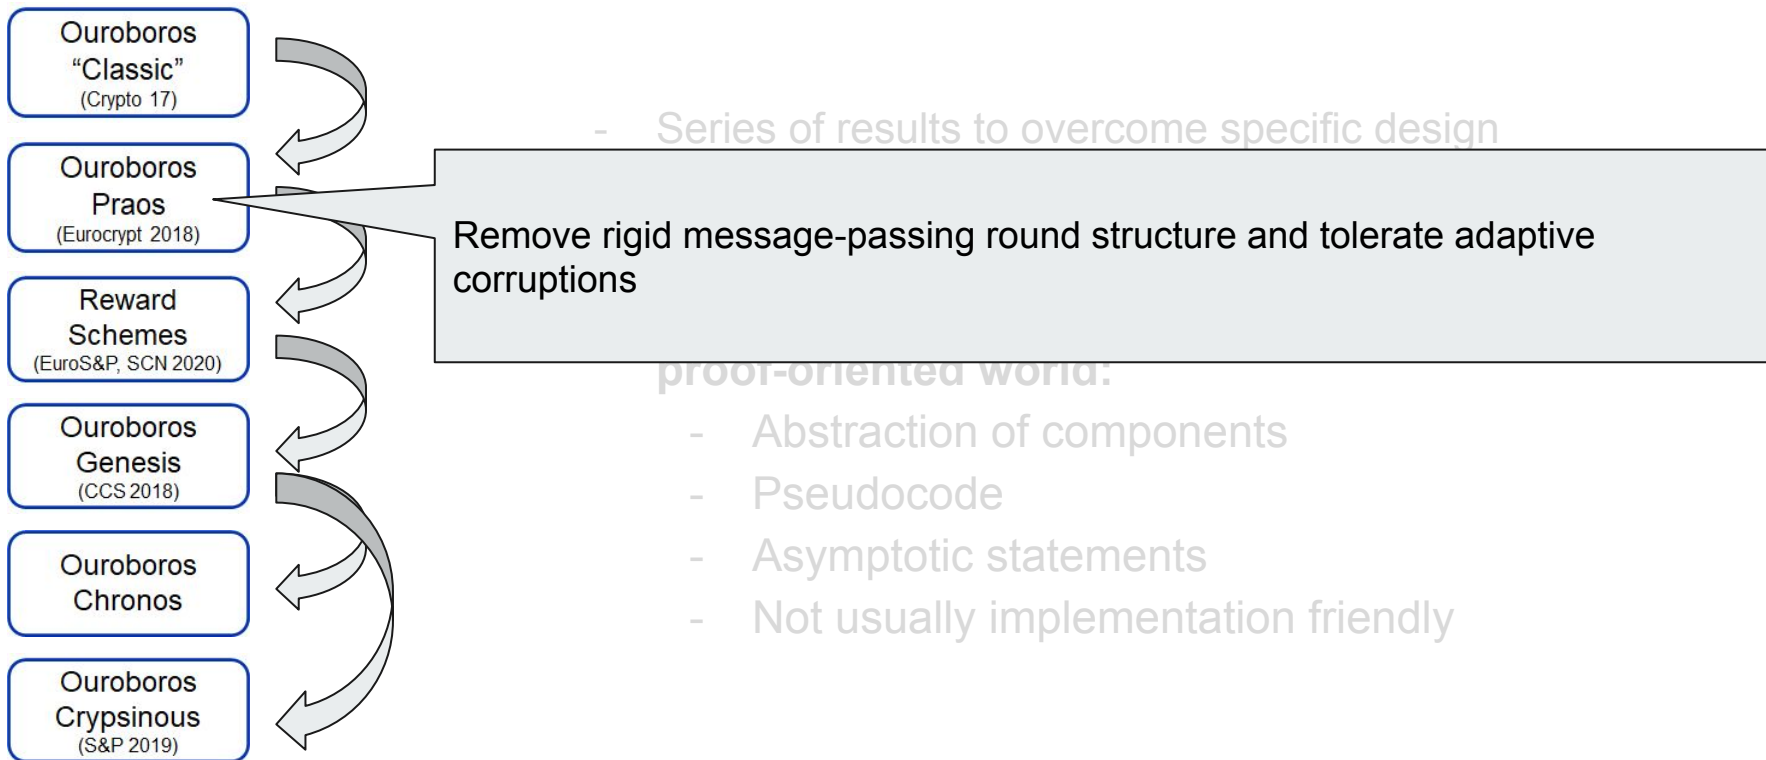


- Series of results to overcome specific design challenges
- **Theoretical papers** live in an **abstract, proof-oriented world**:
 - Abstraction of components
 - Pseudocode
 - Asymptotic statements
 - Not usually implementation friendly

Theory + Practice Interplay



Theory + Practice Interplay



Theory + Practice Interplay

Ouroboros
"Classic"
(Crypto 17)

Ouroboros
Praos
(Eurocrypt 2018)

Reward
Schemes
(EuroS&P, SCN 2020)

Ouroboros
Genesis
(CCS 2018)

Ouroboros
Chronos

Ouroboros
Crypsinous
(S&P 2019)

- Series of results to overcome specific design challenges

Reward sharing scheme for stake pools to establish an equilibrium with a large number of pools with high commitment to the system

- Pseudocode
- Asymptotic statements
- Not usually implementation friendly

Theory + Practice Interplay

Ouroboros
“Classic”
(Crypto 17)

Ouroboros
Praos
(Eurocrypt 2018)

Reward
Schemes
(EuroS&P, SCN 2020)

Ouroboros
Genesis
(CCS 2018)

Ouroboros
Chronos

Ouroboros
Crypsinous
(S&P 2019)

- Series of results to overcome specific design challenges

- **Theoretical papers** live in an **abstract, proof-oriented world:**

Dealing with fluctuating participation, joining after disconnection in a decentralized manner

- Not usually implementation friendly

Theory + Practice Interplay

Ouroboros
"Classic"
(Crypto 17)

Ouroboros
Praos
(Eurocrypt 2018)

Reward
Schemes
(EuroS&P, SCN 2020)

Ouroboros
Genesis
(CCS 2018)

Ouroboros
Chronos

Ouroboros
Crypsinous
(S&P 2019)

- Series of results to overcome specific design challenges

- **Theoretical papers** live in an **abstract, proof-oriented world**:

- Abstraction of components
- Pseudocode

More resilience by less dependency on external timing services

Theory + Practice Interplay

Ouroboros
"Classic"
(Crypto 17)

Ouroboros
Praos
(Eurocrypt 2018)

Reward
Schemes
(EuroS&P, SCN 2020)

Ouroboros
Genesis
(CCS 2018)

Ouroboros
Chronos

Ouroboros
Crypsinous
(S&P 2019)

- Series of results to overcome specific design challenges

- **Theoretical papers** live in an **abstract, proof-oriented world**:

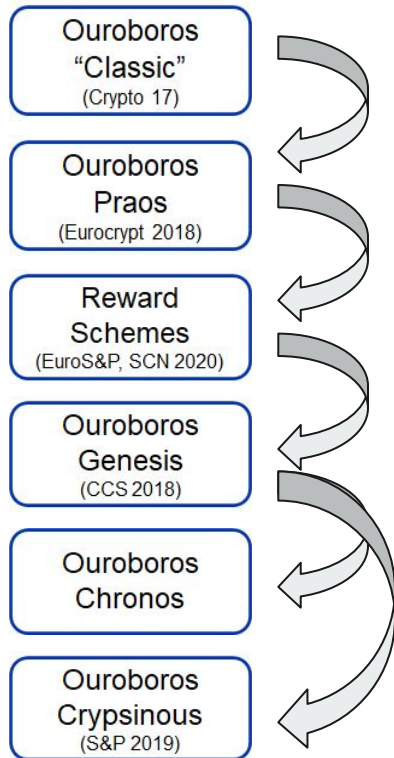
- Abstraction of components

- Pseudocode

Asymptotic statements

More functionality by privacy

Theory + Practice Interplay



Everything so ordered ...

Ouroboros Preprint +
Engineering launches 2016

Theory + Practice Interplay

Ouroboros
"Classic"
(Crypto 17)

Ouroboros
Praos
(Eurocrypt 2018)

Reward
Schemes
(EuroS&P, SCN 2020)

Ouroboros
Genesis
(CCS 2018)

Ouroboros
Chronos

Ouroboros
Crypsinous
(S&P 2019)

September
2017 - first
version live

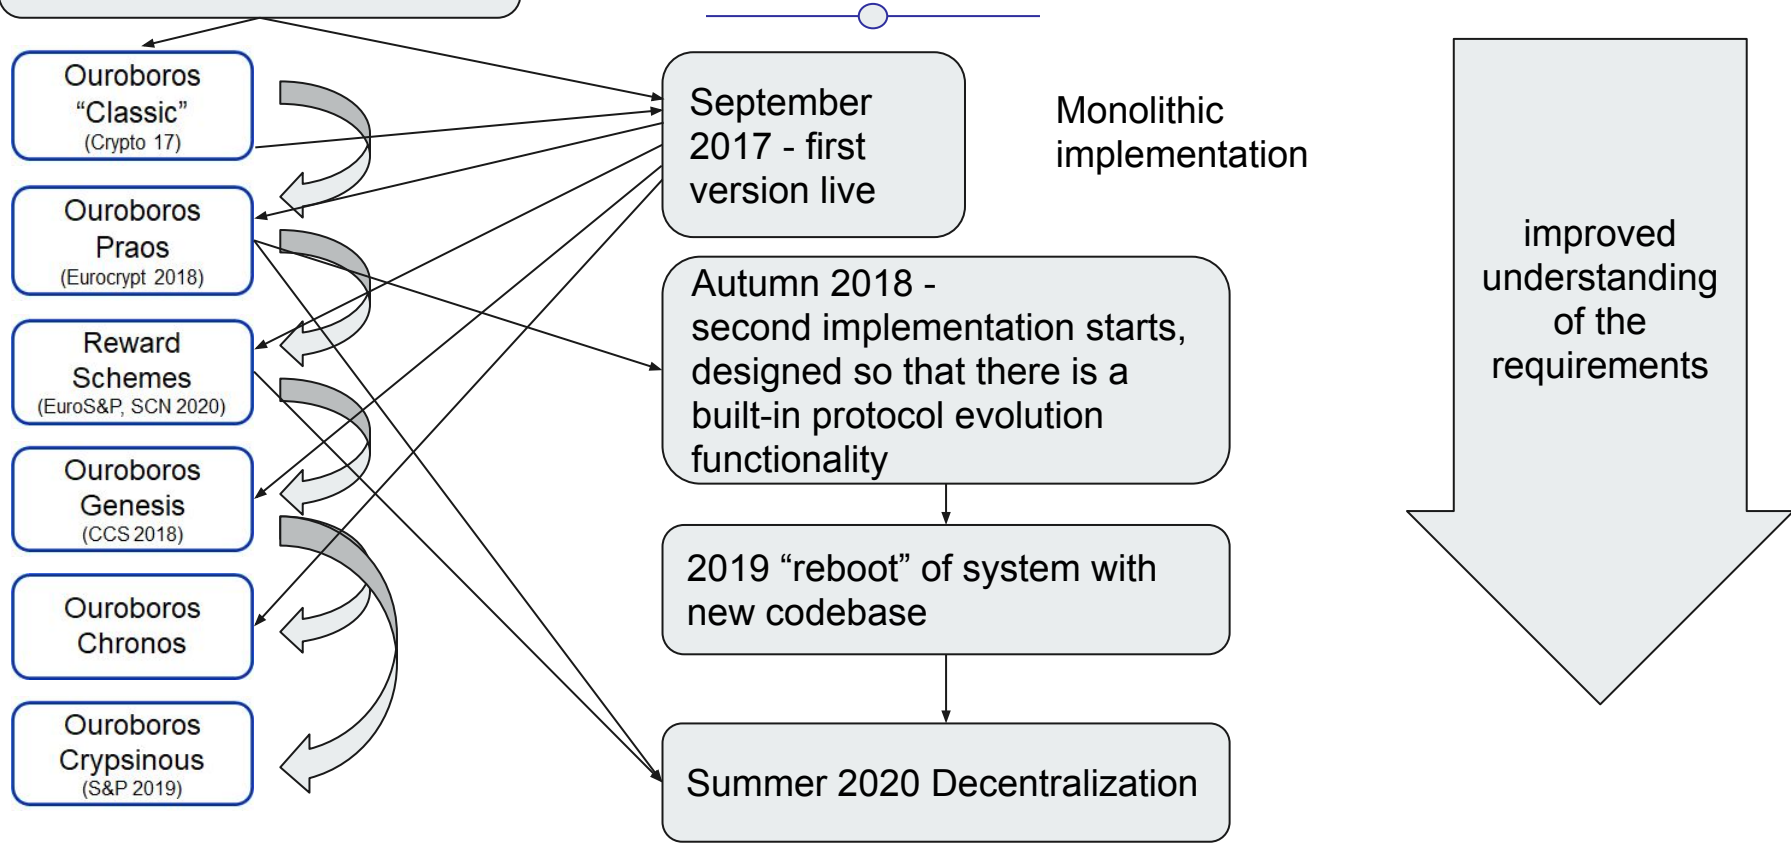
Monolithic
implementation

Autumn 2018 -
second implementation starts,
designed so that there is a
built-in protocol evolution
functionality

2019 "reboot" of system with
new codebase

Summer 2020 Decentralization

improved
understanding
of the
requirements



Overview

1

From Theory to Practice

2

**Formal Methods and Implementation
Correctness**

3

What Could Have Gone Wrong

4


Path to Decentralisation

Implementing ~~Correctly~~ or Quickly?

and

Questions before implementation

- Required level of confidence?
- Time available?
- Requirements fixed?
- Difficulty of Problem?



Correctness Toolbox




Variety of Tools

- Formal specifications
- Proofs (machine verifiable or manual)
- Model checking
- Property-based tests

System Properties

- Determinism?
- Concurrency?
- Finite state space?

Separating the Concerns




Blockchains are messy

- Forks,
eventual consistency
- Distributed system,
concurrency

Ledger can be simple

- Single, linear view
of history



Transaction Ledger



$$\frac{\begin{array}{l} \text{txins } tx \subseteq \text{dom } utxo \quad \text{minfee } pc \, tx \leq \text{txfee } tx \\ \text{balance (txouts } tx) + \text{txfee } tx = \text{balance (txins } tx \triangleleft utxo) \end{array}}{pc \vdash utxo \xrightarrow[UTXO]{tx} (\text{txins } tx \not\triangleleft utxo) \cup \text{txouts } tx}$$

Rule

```
[ Predicate $ λ(−, utxo, tx) →  
  txins tx ‘Set.isSubsetOf’ dom utxo ?! BadInputs  
, Predicate $ λ(pc, −, tx) →  
  pcMinFee pc tx ≤ txfee tx           ?! FeeTooLow  
, Predicate $ λ(−, utxo, tx) →  
  balance (txouts tx) <> txfee tx  
  ≡ balance (txins tx < utxo)       ?! IncreasedTotalBalance  
]
```

(Extension ◦ Transition \$
λ(pc, utxo, tx) → (txins tx $\not\triangleleft$ utxo) \cup txouts tx)

Formal/executable specifications

- Describe **microscopic** behaviour
“small step operational semantics”

Correctness:

- Formulate **macroscopic** properties
“money is conserved”
- Property-based testing
QuickCheck
- Prove them (manually, assisted)

Consensus and Networking



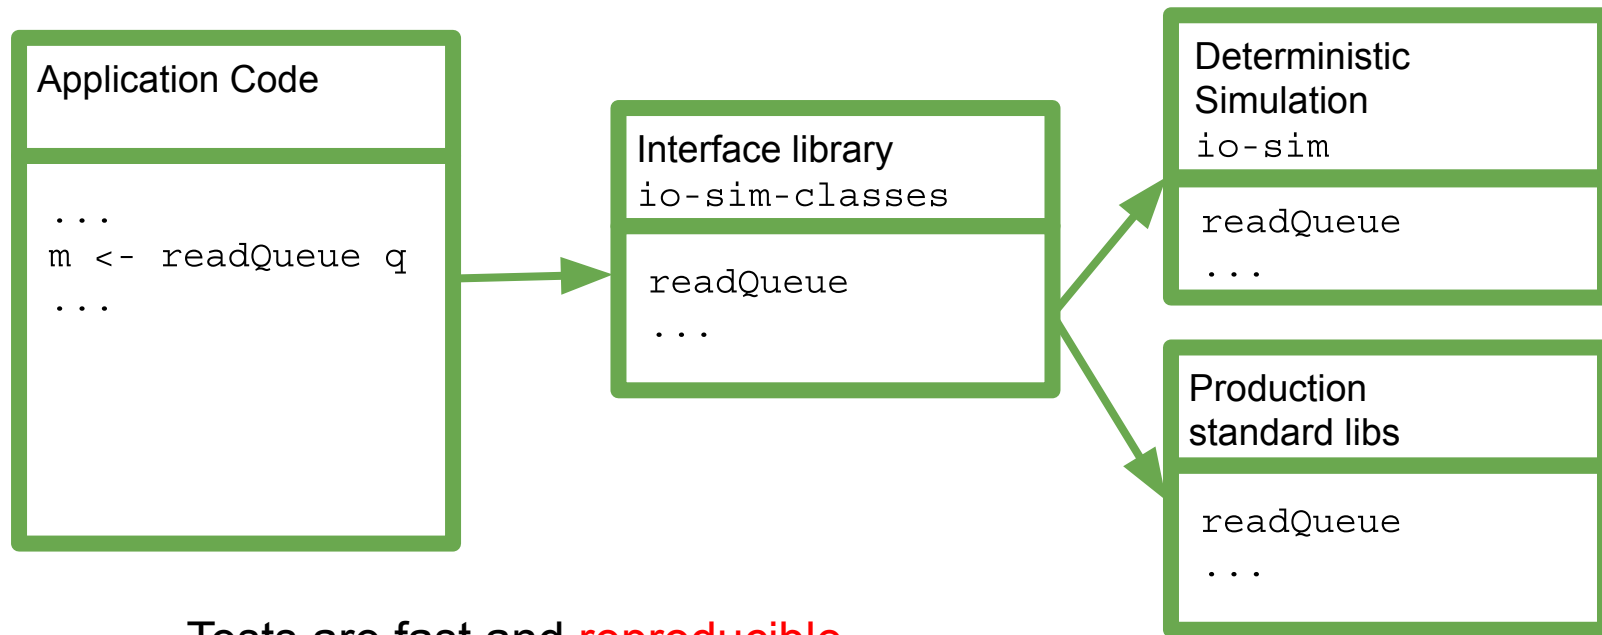
Harder: Concurrency, Distributed Computing

- Language of process calculi
- Not feasible to start from full formal specification within timeframe

Strategy

- Write testable code → **deterministic simulation**
property-based testing
- Use type system for guaranteeing invariants
- “Catch Up” on the formal side

Testing Concurrent Code in Simulation




Tests are fast and **reproducible**

Shrink to minimal **counterexample**

Network Protocols: Session Types

Race Condition Too many are talking


Deadlock Everyone is waiting



Typed Protocols

- Always one sender, one receiver
 - Enforced by Haskell type system
- Protocol type class
- Violations prevented at compile time

Example: Fetching Blocks



Formal Treatment



Goal is to prove equivalence of high and low-level designs

- High level: described exactly as in the papers
- Low level:
 - a practical design;
 - matching how networks work; and
 - operate in bounded resources.
- Modular design: high and low-level related by a series of refinements
- Modular proof: prove equivalence of each refinement and compose
- Machine checked proofs

Formal Treatment




Results so far

- Formalised an asynchronous process calculus
 - Inspired by the π -calculus
 - Using Isabelle theorem prover
 - Proved all the usual properties of a process calculus
 - Framework for proving bisimilarity results
- Formalised high-level paper versions of Ouroboros BFT and Praos
- Proved equivalence of message broadcast vs relay over a network
- Proved equivalence of bulk vs incremental Ouroboros chain selection

Formal Treatment

Message broadcast
(async, with loss and duplication)




Bisimilarity(*) proved within
the process calculus.

The graphical notation has a syntactic
equivalent in the process calculus.

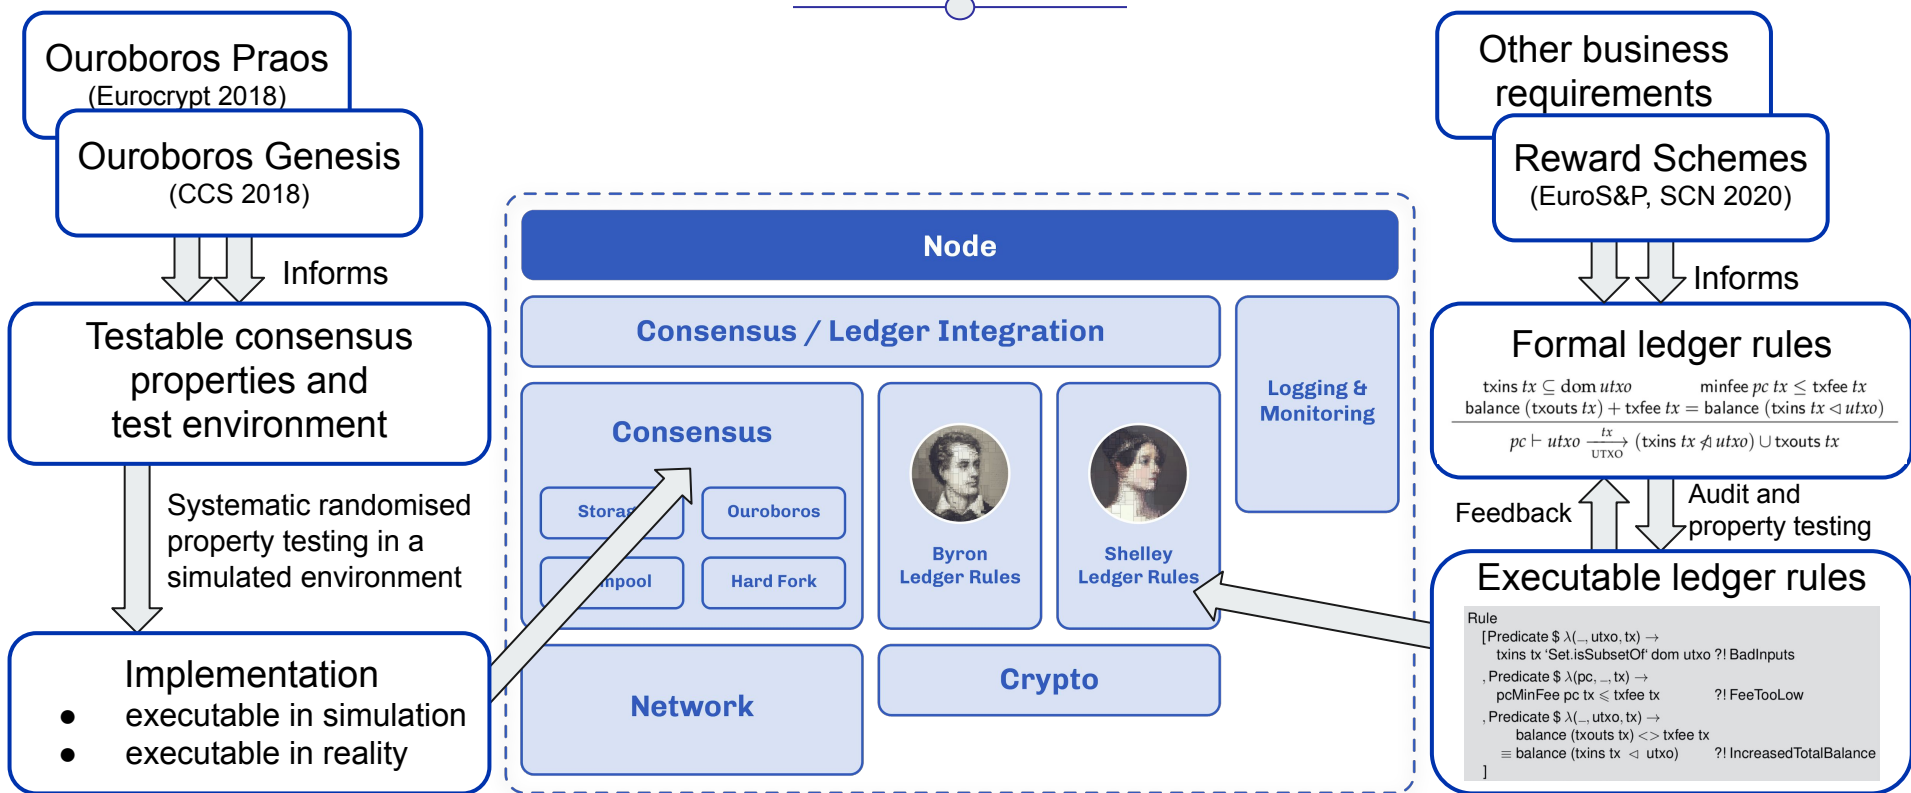
\approx

Relay over a directed graph
(async, with loss and duplication)



(*) Bisimilarity modulo more *structure* to the
message loss in the relay graph:
in relay from $A \Rightarrow B \Rightarrow C$, if a message
arrives at C we can infer it also arrived at
 B , which is not true for broadcast.

Connecting theory with implementation



Overview

1

From Theory to Practice

2

Formal Methods and Implementation
Correctness

3

What Could Have Gone Wrong

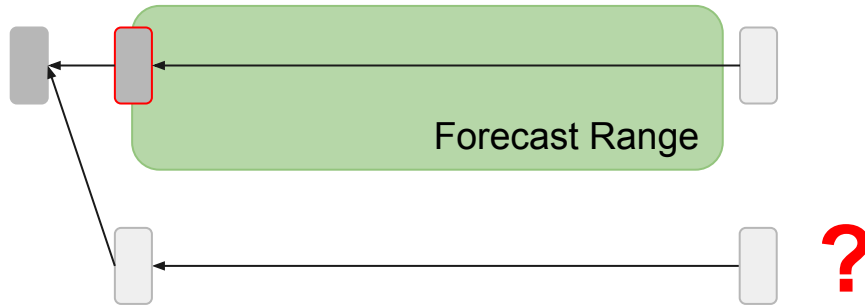
4

Path to Decentralisation

Block-Fetch, Forecasting, Denial of Service

Tension: Open Participation, Limited Resources


- Needs aggressive filtering
- Select chain based on headers before downloading blocks



Block-Fetch, Forecasting, Denial of Service

Edge Cases Require Block Download Before Choosing a Chain

- Proposal: only do that when there is no other way to follow any chain




Implemented: tests show disagreement between nodes

Note: general, randomised test, not unit test

Hard-Fork Combinator

Research: One protocol, starting from scratch

Real-World System: Protocol evolves, but history is immutable



Challenges:

- One codebase needs to understand the whole chain
- Need agreement on when to transition

Hard-Fork Combinator:

- Define new protocol as sequential composition of protocols

De-Risking Decentralisation



Risks in Switching to Decentralised Protocol

- Operators need to gain experience → system not stable
- Little active stake → danger of 51% attack

Gradual Transition


- Fraction \bar{d} of blocks produced by fixed federation
- Gradually hand over control, from $\bar{d}=1$ to $\bar{d}=0$

Real-World Concern: Finite Resources

Open Participation, Limited Resources: worry about DoS

- Proof of Work: producing a chain is much more difficult than validating it
→ obvious, significant advantage for honest nodes
- Proof of Stake: finer balance
→ problem: attacker can cheaply occupy honest resources
- Design needs to analyse worst-case behaviour
excludes many off-the-shelf libraries

e.g. aggressive filtering



Refine Carefully



Papers are Quite Abstract, Implementation Introduces Details

- Interactions of the adversary with the system also more detailed
- Avoid creating new interactions that are detrimental to the security, which are not describable at the high level
- Mind the resource balance between honest and adversarial nodes

Future Work



- Close remaining gaps between theory and implementation
- Formalised version of the security guarantees of Ouroboros
Ph.D. thesis, using Isabelle theorem prover
- Implement newer versions of Ouroboros
- Adding additional ledger functionality
smart contracts, decentralised software updates