

Towards Mechanised Probabilistic Blockchains

Kiran Gopinathan, Ilya Sergey

UCL, NUS

Blockchains & Security

- ▶ Blockchain systems have become commonplace.
- ▶ Hundreds of public Blockchain systems deployed to date.
- ▶ History of bugs and exploits:
 - ▶ 92 billion BTC underflow in 2010.
 - ▶ 5 successful 51% attacks in 2018.




Prior Work

► Formalisations

- Bitcoin Backbone Protocol
- Blockchain in Asynchronous networks

► Mechanisations

- Toychain



Mechanising Blockchain Consensus

George Pirlca
University College London, UK
george.pirlca.15@ucl.ac.uk

Ilya Sergey
University College London, UK
i.sergey@ucl.ac.uk

Abstract

We present the first formalisation of a blockchain-based distributed consensus protocol with a proof of its consistency mechanised in an interactive proof assistant.

Our development includes a reference mechanisation of the *block forest* data structure, necessary for implementing provably correct per-node protocol logic. We also define a model of a network, implementing the protocol in the form of a replicated state-transition system. The protocol's executions are modeled via a small-step operational semantics for asynchronous message passing, in which packages can be rearranged or duplicated.

In this work, we focus on the notion of global system safety, proving a form of eventual consistency. To do so, we provide a library of theorems about a pure functional implementation of block forests, define an inductive system invariant, and show that, in a quiescent system state, it implies a global agreement on the state of per-node transaction ledgers. Our development is parametric w.r.t. implementations of several security primitives, such as hash-functions, a notion of a *proof object*, a *Validator-Acceptance Function*, and a *Fork-Choice Rule*. We precisely characterize the assumptions, made about these components for proving the global system consensus, and discuss their adequacy. All results described in this paper are formalised in Coq.

CCS Concepts • Theory of computation → Program verification; • Networks → Formal specifications;

Keywords blockchain, consensus, protocol verification, Coq

ACM Reference Format
George Pirlca and Ilya Sergey. 2018. Mechanising Blockchain Consensus.

1 Introduction

The notion of decentralised blockchain-based consensus is a tremendous success of the modern science of distributed computing, made possible by the use of basic cryptography, and enabling many applications, including but not limited to cryptocurrencies, smart contracts, application-specific arbitration, voting, etc.

In a nutshell, the idea of a distributed consensus protocol based on *blockchains*, or *transaction ledgers*,¹ is rather simple. In all such protocols, a number of stateful nodes (participants) are communicating with each other in an asynchronous message-passing style. In a message, a node (a) can announce a *transaction*, which typically represents a certain event in the system, depending on the previous state of the node or the entire network (we intentionally leave out the details of what can go into a transaction, as they are application-specific); a node can also (b) create and broadcast a block that contains the encoding of a certain vector of transactions, created locally or received via messages of type (a) from other nodes. Each recipient of a block message should then validate the block (i.e., check the consistency of the transaction sequence included in it), and, in some cases, append it to its local ledger, thus, extending its subjective view of the global sequence of transactions that have taken place in the system to date. The process continues as more messages are emitted and received.

In order to control the number of blocks in the system, distributed ledger protocols rely on certain cryptographic primitives, such as a hash-function *hash* defined both on transactions and blocks, a notion of a *proof object* necessary for defining the validity of a block, and an implementation of a *Validator-Acceptance Function* (VAF) that is used to ensure

Blockchain Protocol

set of
transactions

$\{tx_1, tx_2, tx_3, tx_4, tx_5\}$

Blockchain Protocol

set of
transactions

$\{tx_1, tx_2, tx_3, tx_4, tx_5\}$



blockchain
protocol

global
ordering

$tx_1 \rightarrow tx_2 \rightarrow tx_3 \rightarrow tx_4 \rightarrow tx_5$

Blockchain Protocol

set of
transactions

$\{tx_1, tx_2, tx_3, tx_4, tx_5\}$



$[] \rightarrow [tx_1] \rightarrow [tx_2, tx_3] \rightarrow [tx_4, tx_5]$



global
ordering

$tx_1 \rightarrow tx_2 \rightarrow tx_3 \rightarrow tx_4 \rightarrow tx_5$

Blockchain Protocol

set of
transactions

$\{tx_1, tx_2, tx_3, tx_4, tx_5\}$



$[\] \leftarrow [tx_1] \leftarrow [tx_2, tx_3] \leftarrow [tx_4, tx_5]$




global
ordering

$tx_1 \rightarrow tx_2 \rightarrow tx_3 \rightarrow tx_4 \rightarrow tx_5$

Probabilistic Difficulty

Adversary has a low chance of success


Global Consensus



Probabilistic Difficulty

Adversary has a low chance of success

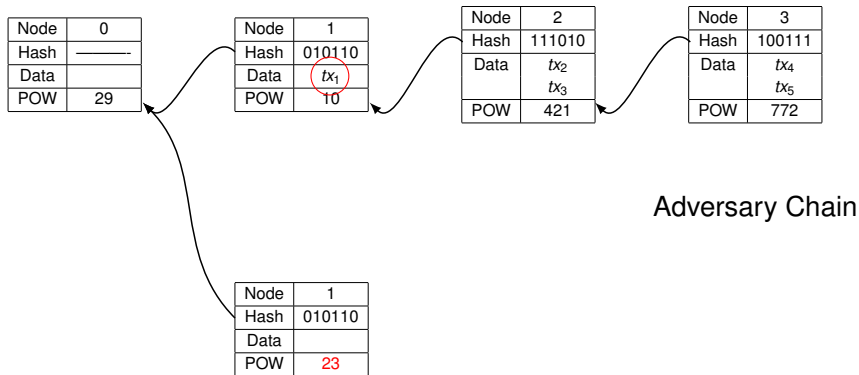
Global Consensus



Probabilistic Difficulty

Adversary has a low chance of success

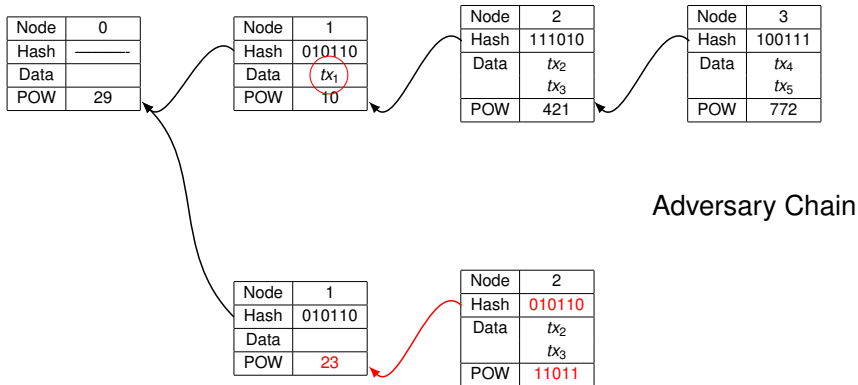
Global Consensus



Probabilistic Difficulty

Adversary has a low chance of success

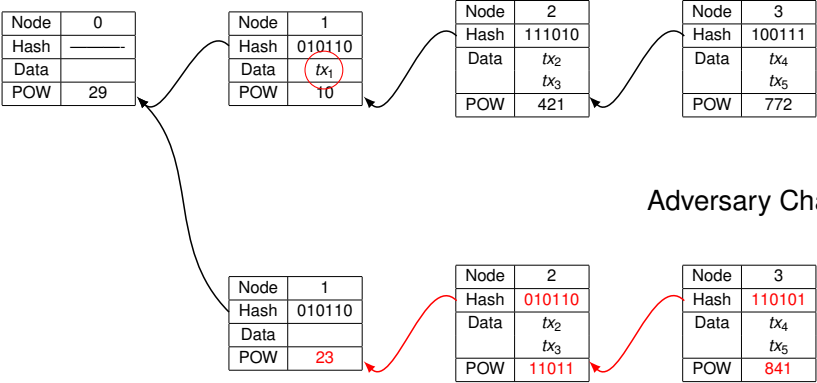
Global Consensus



Probabilistic Difficulty





Adversary has a low chance of success

Global Consensus




Network Model

Network Model



Network Model







Message queue:



Network Model



Round 1







Message queue:



Network Model

Round 1




Message queue:



Network Model


Round 1



Message queue:

Network Model

Round 1




Message queue:



Network Model

Round 1




Message queue:



Network Model

Round 1




Message queue:



Network Model



Round 2




Message queue:



Network Model



Round 2







Message queue:



Network Model

Round 1 + δ




Message queue:



Network Model

Round 1 + δ




[], [tx₁]

Message queue:



Network Model

Round 1 + δ







Message queue:



Network Model

Round 1 + δ




Message queue:



Typical Execution Property

Conditions to exclude edge cases


- ▶ Bounded Successful Rounds - X'
- ▶ Uniquely Bounded Successful Rounds - Y'
- ▶ Number of Adversarial Blocks - Z'



Typical Execution Property

Conditions to exclude edge cases

- ▶ Bounded Successful Rounds - X'
- ▶ Uniquely Bounded Successful Rounds - Y'
- ▶ Number of Adversarial Blocks - Z'



Typical Execution Property

Conditions to exclude edge cases

- ▶ Bounded Successful Rounds - X'
- ▶ Uniquely Bounded Successful Rounds - Y'
- ▶ Number of Adversarial Blocks - Z'

Rounds:

Z'_0	Z'_1	Z'_2	Z'_3	Z'_4	Z'_5	Z'_6	Z'_7	Z'_8
--------	--------	--------	--------	--------	--------	--------	--------	--------

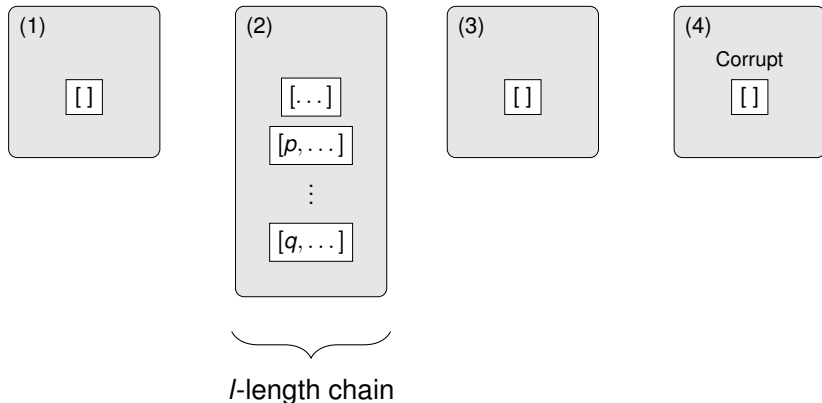
where

$Z'_i \sim$ # blocks hashed by adversary in round i

Chain Growth Property

Adversary can not inhibit chain growth.


Round r



Chain Growth Property

Adversary can not inhibit chain growth.

$$\text{Round } s \geq r + \delta$$




all have chain lengths

$$l' \geq l + \sum_{i=r}^{s-\delta} X'_i$$

Common Prefix Property


Blockchain can achieve global consensus.

At all rounds,




Mechanical Semantics

Round 1



Message queue:

Mechanical Semantics





World

Mechanical Semantics


Round: 1, Active: 1

1	2	3	4
---	---	---	---

Mechanical Semantics




Mechanical Semantics



world_step

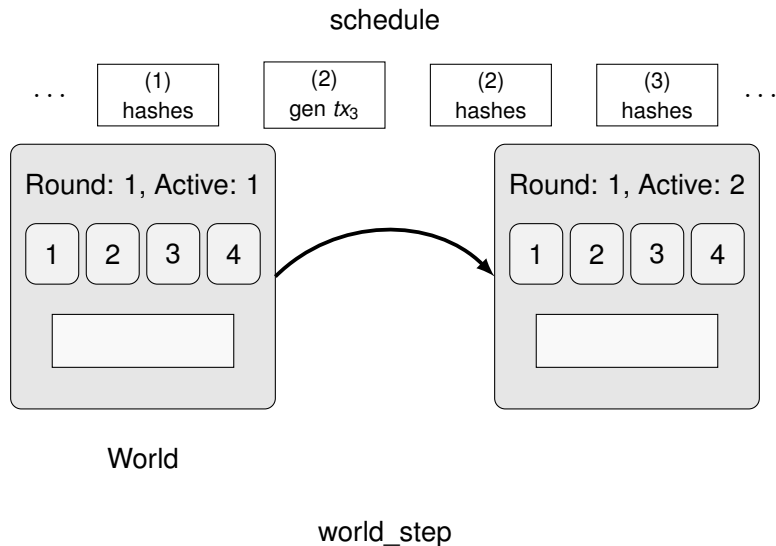
Mechanical Semantics



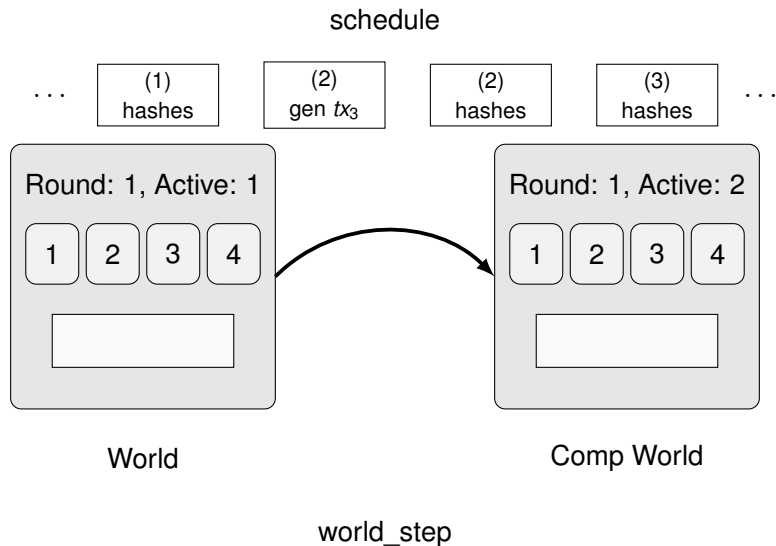
World

world_step

Mechanical Semantics



Mechanical Semantics



Encoding Probability

hash(x) = ???

Encoding Probability

$$\text{hash}(x) = \left\{ \begin{array}{l} \vdots \\ 18 \\ 17 \\ 16 \\ 15 \\ 14 \\ 13 \\ 12 \\ \vdots \end{array} \right.$$

Encoding Probability

$$\text{hash}(x) = \left\{ \begin{array}{l} \vdots \\ 18 \\ 17 \\ 16 \\ 15 : P[\text{hash}(x) = 15] \\ 14 \\ 13 \\ 12 \\ \vdots \end{array} \right.$$

Encoding Probability

$$\text{hash} : A \rightarrow (B \rightarrow \mathbb{R})$$

Encoding Probability

hash : $A \rightarrow \text{dist } B$

Encoding Probability

- ▶ Probability monad defined by Affeldt and Hagiwara.

$\text{bind} : \text{dist } A \rightarrow (A \rightarrow \text{dist } B) \rightarrow \text{dist } B$

$\text{ret} : A \rightarrow \text{dist } A$

Encoding Probability

- ▶ We extend it to probabilistically execute the system.

$\text{eval_dist} : \text{Comp } A \rightarrow \text{dist } A$

Encoding Probability

- ▶ To allow stating properties about probable worlds.

$$\forall sc, \forall w, \text{eval_dist}(\text{world_step } w_0 \text{ } sc) \ w > 0 \implies F \ w$$

$$\forall sc, P[(\text{world_step } w_0 \text{ } sc) \triangleright F] = 1$$

Key lemmas

► Typical Execution Assumption

$$P[\text{world_step } sc \triangleright \text{TEP}_\epsilon \text{ } sc] = 1 - e^{-\Omega(\kappa)}$$

► Chain Growth Property

$$P[\text{world_step } sc \ w_0 \triangleright \text{CGP}] = 1$$

► Common Prefix Property

$$P[\text{world_step } sc \ w_0 \triangleright (\text{CPP}_k \wedge \text{TEP}_\epsilon \text{ } sc)] = \\ P[\text{world_step } sc \ w_0 \triangleright \text{TEP}_\epsilon \text{ } sc]$$

Key lemmas

▶ Typical Execution Assumption

$$P[\text{world_step } sc \triangleright \text{TEP}_\epsilon \text{ } sc] = 1 - e^{-\Omega(\kappa)}$$

▶ Chain Growth Property

$$P[\text{world_step } sc \ w_0 \triangleright \text{CGP}] = 1$$

▶ Common Prefix Property

$$P[\text{world_step } sc \ w_0 \triangleright (\text{CPP}_k \wedge \text{TEP}_\epsilon \text{ } sc)] = \\ P[\text{world_step } sc \ w_0 \triangleright \text{TEP}_\epsilon \text{ } sc]$$

Key lemmas

▶ Typical Execution Assumption

$$P[\text{world_step } sc \triangleright \text{TEP}_\epsilon \text{ } sc] = 1 - e^{-\Omega(\kappa)}$$

▶ Chain Growth Property

$$P[\text{world_step } sc \ w_0 \triangleright \text{CGP}] = 1$$

▶ Common Prefix Property

$$P[\text{world_step } sc \ w_0 \triangleright (\text{CPP}_k \wedge \text{TEP}_\epsilon \text{ } sc)] = \\ P[\text{world_step } sc \ w_0 \triangleright \text{TEP}_\epsilon \text{ } sc]$$

Key lemmas

▶ Typical Execution Assumption

$$P[\text{world_step } sc \triangleright \text{TEP}_\epsilon \text{ } sc] = 1 - e^{-\Omega(\kappa)}$$

▶ Chain Growth Property

$$P[\text{world_step } sc \ w_0 \triangleright \text{CGP}] = 1$$

▶ Common Prefix Property

$$P[\text{world_step } sc \ w_0 \triangleright (\text{CPP}_k \wedge \text{TEP}_\epsilon \text{ } sc)] = \\ P[\text{world_step } sc \ w_0 \triangleright \text{TEP}_\epsilon \text{ } sc]$$

Main Contributions

- ▶ Implemented a mechanised probabilistic blockchain model based on the Bitcoin Backbone Protocol (BBP) by Garay et al.
- ▶ Proved several preliminary lemmas.
- ▶ Formulated the main BBP lemmas within this model.

Future work

- ▶ Completing proofs of the key properties.
- ▶ Elevating the Typical Execution Assumption to a lemma.
- ▶ Extracting the system to an executable implementation.

Take away

- ▶ Blockchain security properties inherently require probabilistic considerations.
- ▶ 2 key properties:
 - ▶ **Chain growth property**
 - ▶ **Common prefix property**
- ▶ Working on a mechanisation of the Bitcoin backbone protocol.
 - ▶ Mechanised protocol model.
 - ▶ Formulated several key lemmas.

<https://github.com/certichain/probchain>

Infotheo Probability Monad

- ▶ **Benefits**

- ▶ SSReflect integration
- ▶ Several utility lemmas

- ▶ **Drawbacks**

- ▶ Fintype requirement.