



# Mechanising Blockchain Consensus

George Pîrlea and Ilya Sergey



# Context

- Hundreds of deployed public blockchains
- ~~\$600~~ ~~625~~ ~~675~~ ~~735~~ ~~755~~ ~~780~~ 820 billion total market cap  
*(7 day progression since Jan 1<sup>st</sup>)*



# This work

- Formalised a blockchain consensus protocol in Coq
- Proved eventual consistency in a clique topology

# Motivation

## 1. Understand blockchain consensus

- **what** it is
- **how** it works: example
- **why** it works: our formalisation

## 2. Lay foundation for *verified* practical implementation

- verified Byzantine-tolerant consensus layer
  - platform for verified smart contracts
- } **Future work**

What it does

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

- transforms a **set** of transactions into a *globally-agreed* **sequence**
- “distributed timestamp server” (Nakamoto2008)

blockchain  
consensus protocol

transactions  
can be *anything*

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

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



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



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

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



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



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



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



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

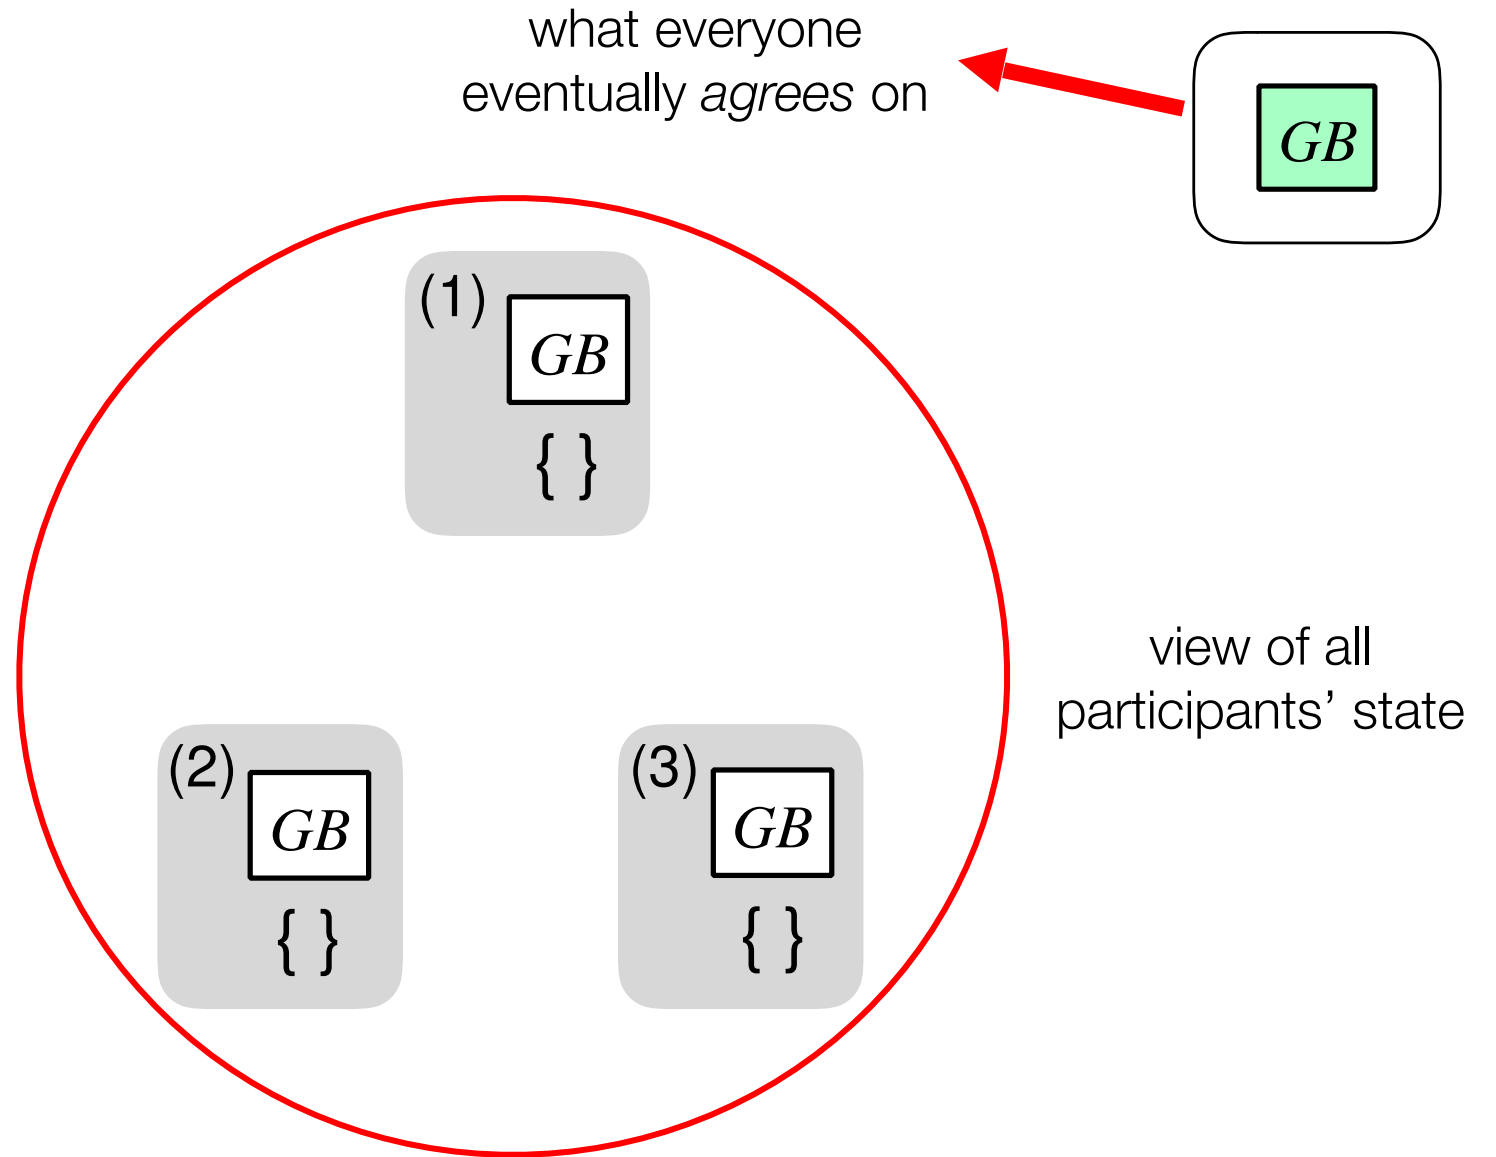
**GB** = genesis block




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

How it works

- **distributed**
  - multiple nodes
- all start with same GB






- **distributed**

- multiple nodes
- message-passing over a network





- all start with same GB



- **distributed**

- multiple nodes
- message-passing over a network


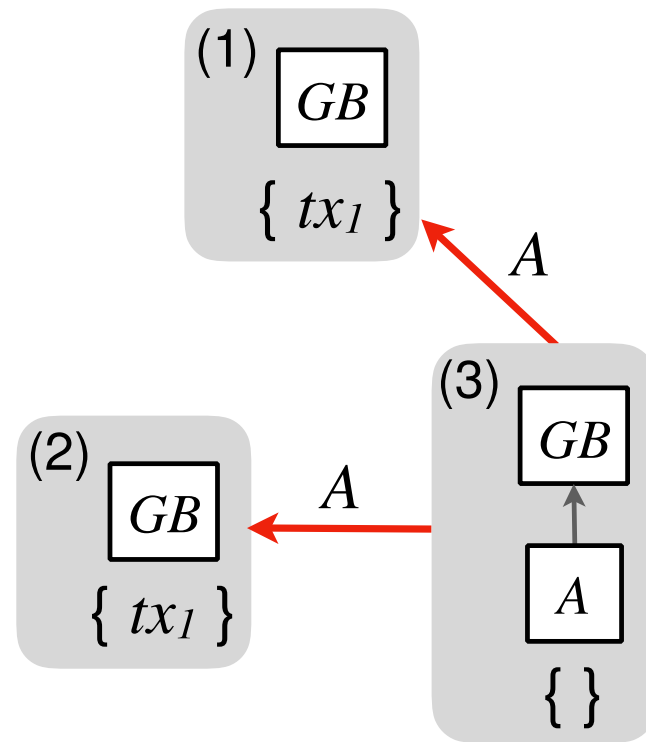
- all start with same GB
- have a transaction pool



- **distributed**

- multiple nodes
- message-passing over a network


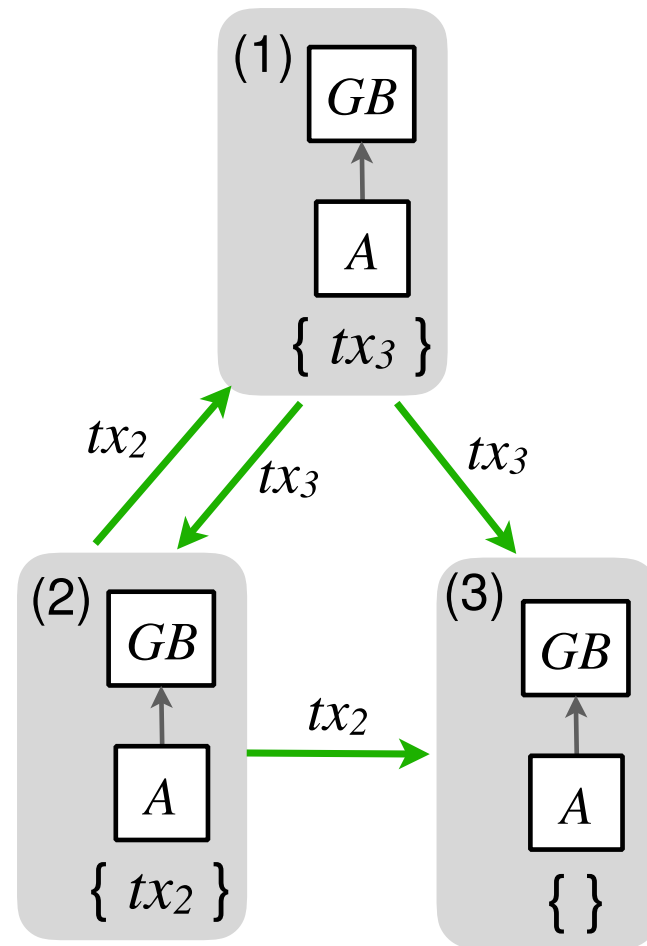
- all start with same GB
- have a transaction pool
- can mint blocks



- **distributed** => concurrent

- multiple nodes
- message-passing over a network



- multiple transactions can be issued and propagated concurrently



- **distributed** =>  
concurrent


- multiple nodes
- message-passing over a network

- blocks can be minted without full knowledge of all transactions

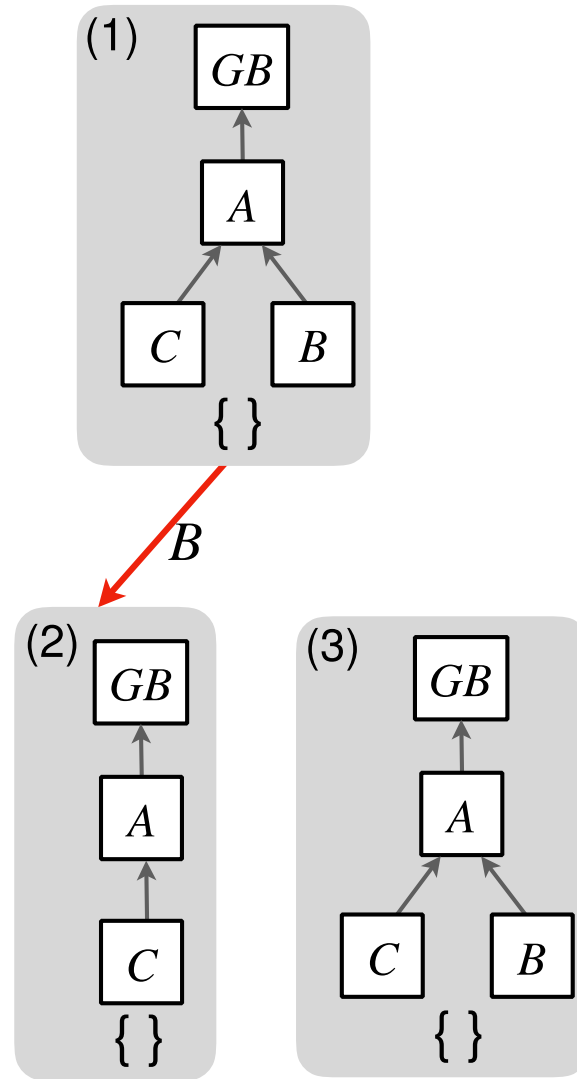




- chain fork has happened, but nodes don't know

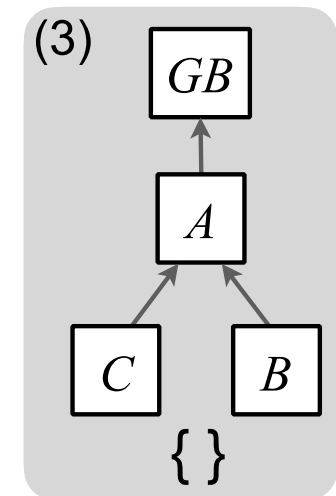



- as block messages propagate, nodes become aware of the fork





# Problem: need to choose

- blockchain “promise” = *one globally-agreed chain*
- each node must choose one chain
- nodes with the same information must choose the same chain





# Problem: need to choose

- blockchain “promise” = *one globally-agreed chain*
- each node must choose one chain
- nodes with the same information must choose the same chain





# Problem: need to choose

- blockchain “promise” = *one globally-agreed chain*
- each node must choose one chain
- nodes with the same information must choose the same chain



# Problem: need to choose

- blockchain “promise” = *one globally-agreed chain*
- each node must choose one chain
- nodes with the same information must choose the same chain



# Solution: fork choice rule

- Fork choice rule (FCR,  $>$ ):
  - given two blockchains, says which one is “heavier”
  - imposes a *strict total order* on all possible blockchains
  - same FCR shared by all nodes
- Nodes adopt “heaviest” chain they know

# FCR ( $\succ$ )


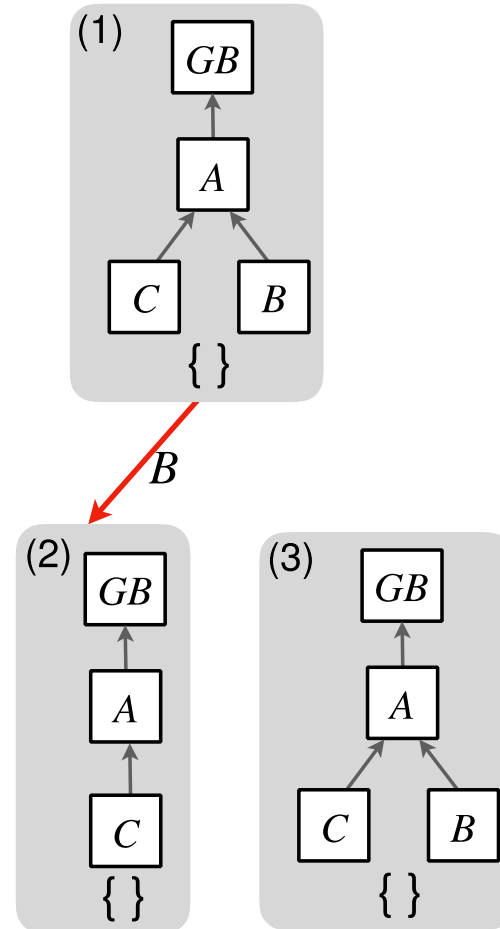
...  $\succ$  [GB, A, C]  $\succ$  ...  $\succ$  [GB, A, B]  $\succ$  ...  $\succ$  [GB, A]  $\succ$  ...  $\succ$  [GB]  $\succ$  ...

Bitcoin: FCR based on “most cumulative work”



# Quiescent consistency

- **distributed**
  - multiple nodes
  - all start with GB
  - message-passing over a network
  - equipped with same FCR
- quiescent consistency: when all block messages have been delivered, everyone agrees



Why it works

## Definitions

- blocks, chains, block forests

## *Parameters and assumptions*

- *hashes* are collision-free
- *FCR* imposes strict total order

## Invariant

- local state + messages “in flight” = global

## Quiescent consistency

- when all block messages are delivered, everyone agrees

# Blocks and chains

links blocks together

$hash_b : \text{Block} \rightarrow \text{Hash}$

$b \in \text{Block} ::= \{ \text{prev} : \text{Hash}; \underline{\text{txs}} : \text{Tx}^*; \text{pf} : \text{Proof} \}$

$c \in \text{Chain} \triangleq \text{Block}^*$

$GB : \text{Block}$

proof-of-work

proof that this block  
was minted in

proof-of-stake

accordance to the  
rules of the protocol

# Minting and verifying

 *try* to generate a proof = “ask the protocol for permission” to mint

*mkProof*: Addr → Chain → option Proof

*VAF*: Proof → Time → Chain → bool

 validate a proof = ensure protocol rules were followed

# Resolving conflict

*FCR* : Chain  $\rightarrow$  Chain  $\rightarrow$  bool

# Assumptions

- Hash functions are collision-free

$$\mathit{hash\_inj} \quad : \quad \forall x \ y, \#x = \#y \implies x = y$$

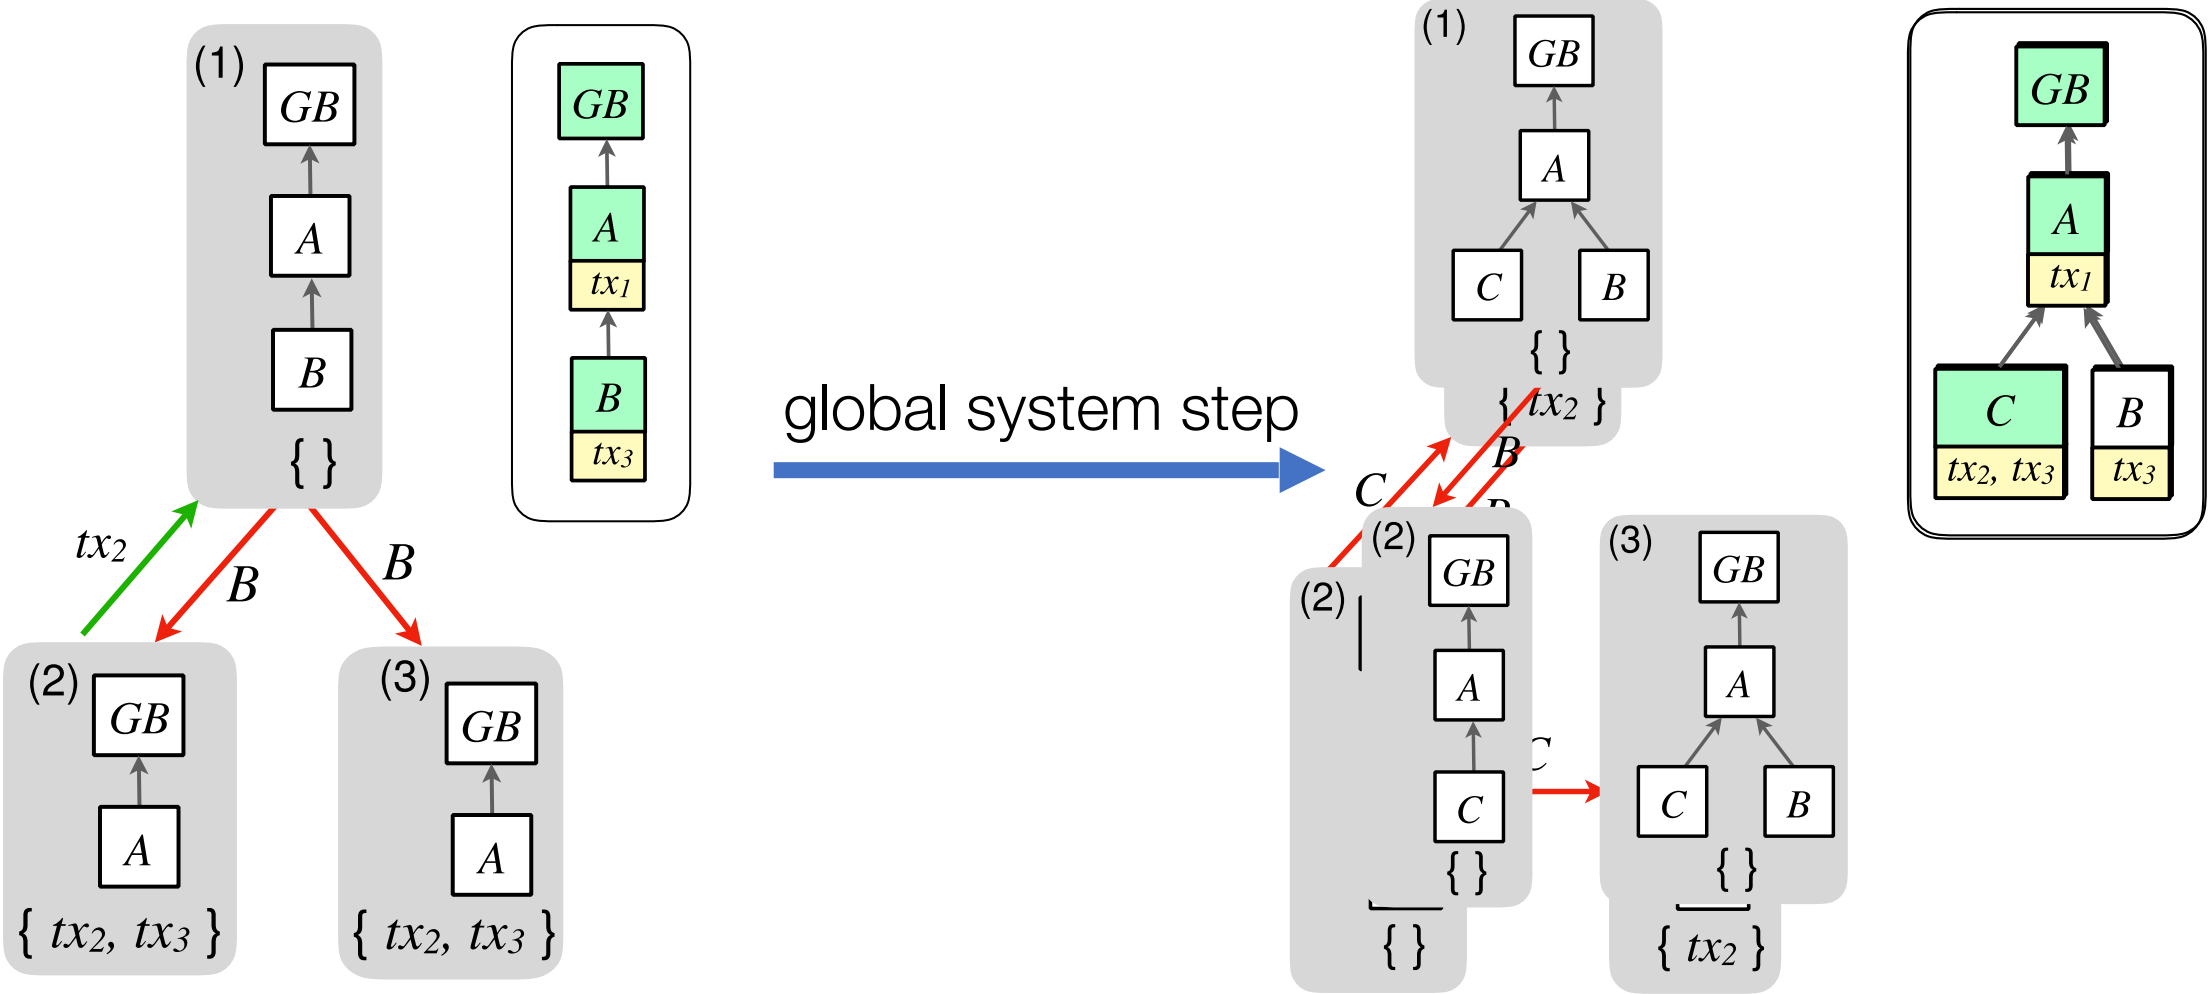
- FCR imposes a *strict total order* on all blockchains

$$\mathit{FCR\_rel} \quad : \quad \forall c_1 \ c_2, c_1 = c_2 \vee c_1 > c_2 \vee c_2 > c_1$$

$$\mathit{FCR\_trans} \quad : \quad \forall c_1 \ c_2 \ c_3, c_1 > c_2 \wedge c_2 > c_3 \implies c_1 > c_3$$

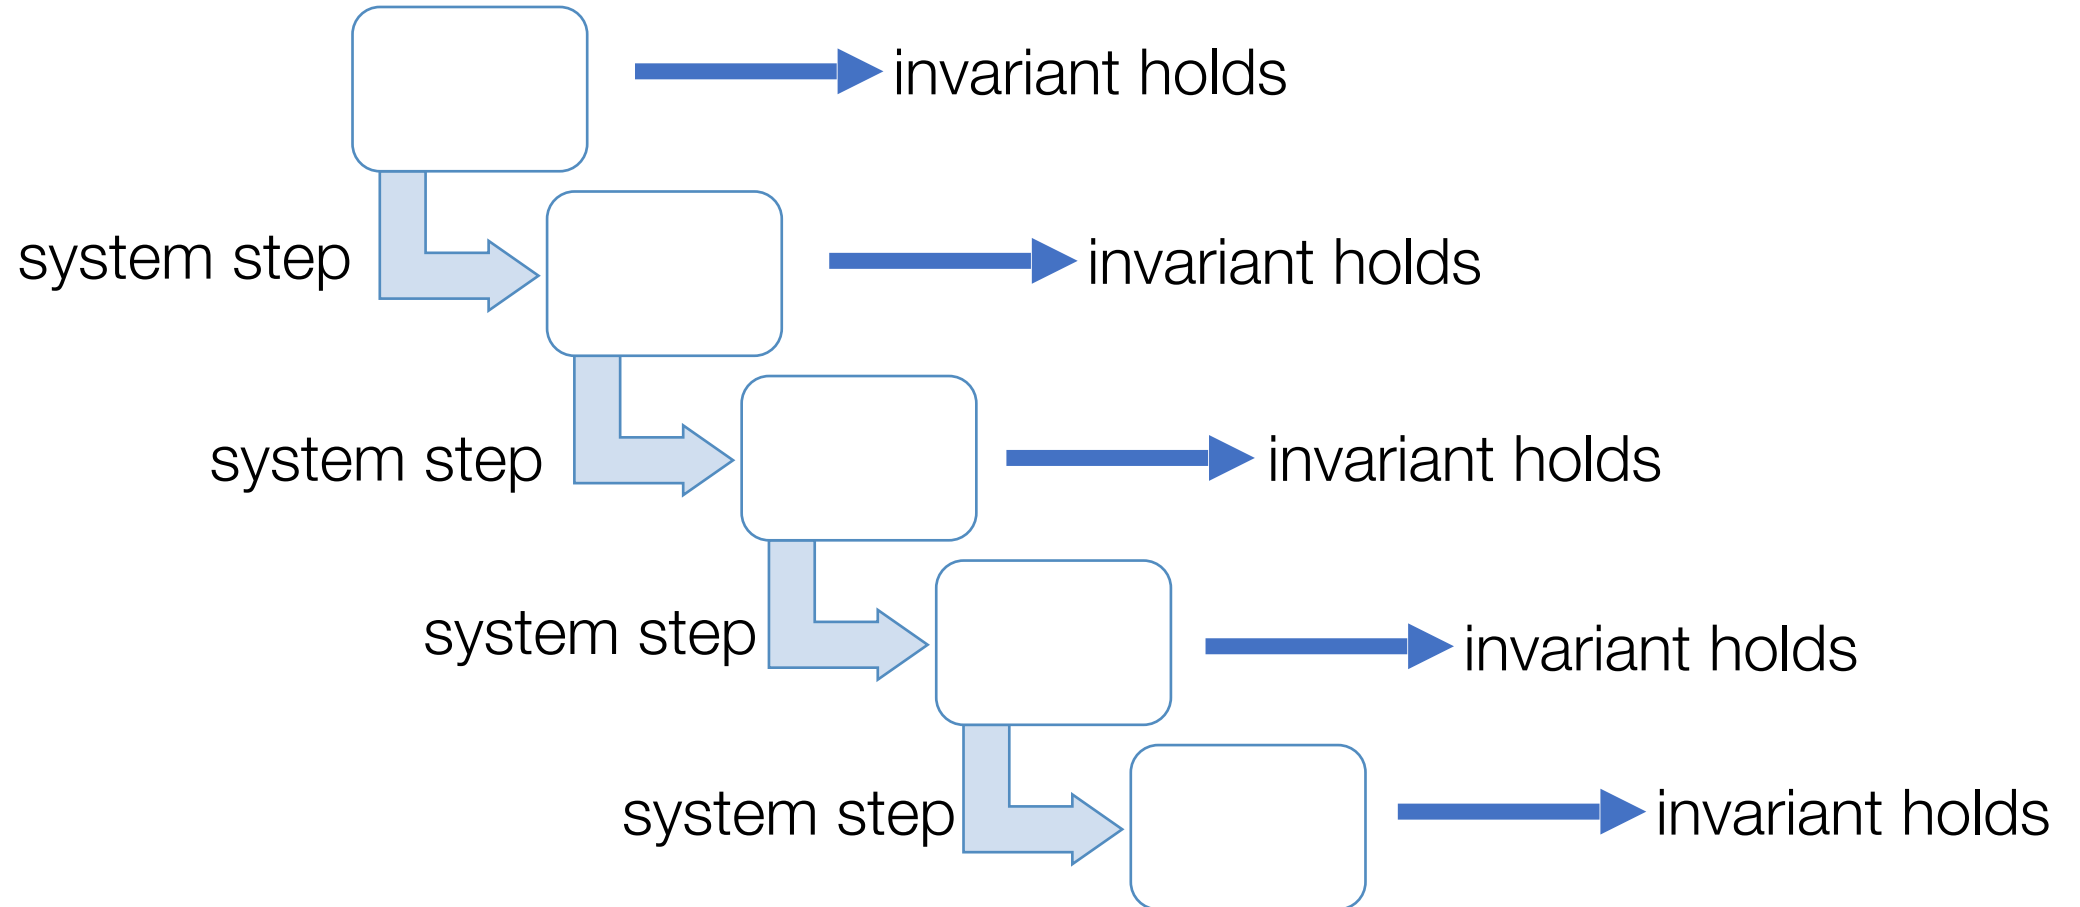
$$\mathit{FCR\_nrefl} \quad : \quad \forall c, c > c \implies \text{False}$$

# Invariant: local state + “in-flight” = global





# Invariant is inductive



# Invariant implies QC

- QC: when all blocks delivered, everyone agrees

How:

- local state + ~~“weight”~~ = global
- use FCR to extract “heaviest” chain out of local state
- since everyone has same state & same FCR
  - consensus

# Reusable components

- Reference implementation of block forests
- Per-node protocol logic
- Network semantics
- Clique invariant, QC property, various theorems

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

# Future work

- Network semantics with nodes joining/leaving at will
- Improved invariants:
  - non-clique topologies
  - network partitions
  - Byzantine faults
- Verified smart contracts platform

# Take away

- Formalisation of a blockchain consensus protocol in Coq:
  - minimal set of required security primitives
  - per-node protocol logic & data structures
  - network semantics
- global eventual consistency in a clique topology

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