# Formally verifying Coco

George Pîrlea

#### About me

Background:

- formal verification
- distributed systems

- currently undergrad at UCL
- looking to do a PhD

Current focus:

• blockchain consensus protocols

#### Research motivation:

• trustworthy censorship-resistant digital infrastructure

#### This work

- Created a formal model of Coco in F\*
  - node-local behaviour of crash-fault-tolerant replication (CFTR)
  - core Coco service: commands, user-defined application, shared log
- Proved some high-level integrity properties
  - any CFTR protocol which satisfies the spec is *correct*
  - Coco, using such a protocol, is **sequentially consistent**

#### Outline

- model break-down
  - focus on highlighting assumptions
- sequential consistency
  - prerequisites
  - visual proof
- conclusions & future work

#### Model

- crash-fault-tolerant consensus
  - described in terms of the local behaviour at each node
- Coco running "on top" of the CFTR protocol
  - each Coco node:
    - participates in the consensus protocol
    - runs the user-defined service (application)
    - can respond to client commands

differs from implementation

### Crash-fault-tolerant replication

Specification of the node-local behaviour of a consensus protocol









#### Extend



#### Extend



#### Truncate

#### Truncate





#### Truncate



can't truncate committed entries



local log















#### CFTR model

- replication system is a network of nodes
- nodes undergo local transitions (extend, truncate, commit)
  - transitions are relations between network states
- users can inject proposals into the system (propose)
- *ghost* global commit log (GCL)

#### CFTR model

- replication system is a network of nodes
- nodes undergo local transitions (extend, truncate, commit)
  - transitions are relations between network states
- users can inject proposals into the system (propose)
- ghost global commit log (GCL)

state required for the proofs; doesn't exist in implementation

#### CFTR model

- replication system is a network of nodes
- nodes undergo local transitions (extend, truncate, commit)
  - transitions are relations between network states
- users can inject proposals into the system (propose)
- *ghost* global commit log (GCL)

#### CFTR correctness

• there are no forks in the committed part of the logs

- a node has the global commit log as its local log
  every other committed log is a prefix
- all entries in the log correspond to a proposal

# $\text{Coco}_{\text{m}}$

Our model of Coco, which (hopefully) over-approximates a subset of Coco's real behaviours

## What is Coco<sub>m</sub>?

• User-defined service running "on top" of CFTR

eval: log -> kvState
service: command -> kvState -> (tx × result)

- assuming commands have some unique identifier
- assuming different commands, when run on the same kvState, produce different transactions (extensions to the log)

- 1. Run a CFTR transition (extend, truncate or commit)
- 2. Process a client command

- 1. Run a CFTR transition (extend, truncate or commit)
- 2. Process a client command
  - a) eval log  $\rightarrow$  kvS

- 1. Run a CFTR transition (extend, truncate or commit)
- 2. Process a client command
  - a) eval log  $\rightarrow$  kvS
  - b) service cmd kvS  $\rightarrow$  (tx, result)

- 1. Run a CFTR transition (extend, truncate or commit)
- 2. Process a client command
  - a) eval log  $\rightarrow$  kvS
  - b) service cmd kvS  $\rightarrow$  (tx, result)
  - c) serialize\_tx tx  $\rightarrow$  propose to CFTR
# What can Coco<sub>m</sub> do?

Any Coco<sub>m</sub> node can:

- 1. Run a CFTR transition (extend, truncate or commit)
- 2. Process a client command
  - a) eval log  $\rightarrow$  kvS
  - b) service cmd kvS  $\rightarrow$  (tx, result)
  - c) serialize\_tx tx  $\rightarrow$  propose to CFTR
  - d) return result to client

# What can Coco<sub>m</sub> do?

Any Coco<sub>m</sub> node can:

- 1. Run a CFTR transition (extend, truncate or commit)
- 2. Process a client command
  - a) eval log  $\rightarrow$  kvS
  - b) service cmd kvS  $\rightarrow$  (tx, result)
  - c) serialize\_tx tx  $\rightarrow$  propose to CFTR
  - d) return result to client

assuming correct transaction serializer & deserializer

assuming proposals uniquely identify the log which they extend assuming one atomic step



## Proof infrastructure

It's proofs all the way down

#### Transitions, invariants

• Transitions are modelled as inductive data types (relations)



• An **invariant** is a property that holds over transitions

### Instant vs. history properties



Network history: sequence of network states linked by valid transitions

- Instant: "there are no inconsistencies in the committed logs"
- **History**: "if a command is committed at some point, it is committed at all later points"

Coco<sub>m</sub> instants

Model:

- network of nodes, each running their own replication node
- ghost state:
  - processed commands
  - mapping from commands to proposals and vice-versa

Proofs:

- the ghost state is internally consistent
- if a command is committed in state *m*, and you undergo a transition, it remains committed in state *m*'
- commands are not created committed

Coco<sub>m</sub> histories

• history = non-empty list of Coco<sub>m</sub> instants



#### Definitions

Instant properties:

• **committed** *cmd* = processed & included in GCL

#### Definitions

Instant properties:

• committed cmd = processed & included in GCL

History properties:

 sequentially committed *cmds* = committed at strictly increasing indices in the history What can Coco., do

#### Definitions

Instant properties:

• committed cmd = processed & included in GCL

History properties:

- sequentially committed cmds = committed at strictly increasing indices in the history
- sequentially consistent *cmds hist* = if *cmds* is sequentially committed, then the corresponding proposals form a subsequence in the GCL of the last instant in *hist*

## Sequential consistency

With pictures!



p <sub>0</sub>	p <sub>1</sub>	p <sub>2</sub>	p <sub>3</sub>	p <sub>4</sub>	р <sub>5</sub>	p <sub>6</sub>	p <sub>7</sub>	p <sub>8</sub>	p <sub>9</sub>	p <sub>10</sub>	p <sub>11</sub>





p <sub>0</sub>	<b>p</b> <sub>1</sub>	p <sub>2</sub>	p <sub>3</sub>	p <sub>4</sub>	<b>p</b> <sub>5</sub>	p <sub>6</sub>	p <sub>7</sub>	p <sub>8</sub>	p <sub>9</sub>	<b>p</b> <sub>10</sub>	p <sub>11</sub>



p <sub>0</sub>	p <sub>1</sub>	p <sub>2</sub>	p <sub>3</sub>	p <sub>4</sub>	p <sub>5</sub>	p <sub>6</sub>	p <sub>7</sub>	p <sub>8</sub>	p <sub>9</sub>	p <sub>10</sub>	p <sub>11</sub>



p <sub>0</sub>	<b>p</b> <sub>1</sub>	p <sub>2</sub>	p <sub>3</sub>	p <sub>4</sub>	р <sub>5</sub>	p <sub>6</sub>	p <sub>7</sub>	p <sub>8</sub>	p <sub>9</sub>	p <sub>10</sub>	p <sub>11</sub>











C <sub>23</sub>	C <sub>52</sub>	C <sub>77</sub>
-----------------	-----------------	-----------------













p <sub>0</sub>	p <sub>1</sub>	p <sub>2</sub>	p <sub>3</sub>	p <sub>4</sub>	<b>p</b> <sub>5</sub>	p <sub>6</sub>	p <sub>7</sub>	p <sub>8</sub>	p <sub>9</sub>	p <sub>10</sub>





p <sub>0</sub>	p <sub>1</sub>	p <sub>2</sub>	p <sub>3</sub>	p <sub>4</sub>	p <sub>5</sub>	p <sub>6</sub>	p <sub>7</sub>	p <sub>8</sub>	p <sub>9</sub>	p <sub>10</sub>



p <sub>0</sub>	p <sub>1</sub>	p <sub>2</sub>	p <sub>3</sub>	p <sub>4</sub>	p <sub>5</sub>	p <sub>6</sub>	p <sub>7</sub>	p <sub>8</sub>	p <sub>9</sub>	p <sub>10</sub>
										1









p <sub>0</sub>	p <sub>1</sub>	p <sub>2</sub>	p <sub>3</sub>	p <sub>4</sub>	p <sub>5</sub>	p <sub>6</sub>	p <sub>7</sub>	p <sub>8</sub>	p <sub>9</sub>	p <sub>10</sub>
										1







p <sub>0</sub>	p <sub>1</sub>	p <sub>2</sub>	p <sub>3</sub>	p <sub>4</sub>	p <sub>5</sub>	p <sub>6</sub>	p <sub>7</sub>	p <sub>8</sub>	p <sub>9</sub>	p <sub>10</sub>
----------------	----------------	----------------	----------------	----------------	----------------	----------------	----------------	----------------	----------------	-----------------





p <sub>0</sub>	p <sub>1</sub>	p <sub>2</sub>	p <sub>3</sub>	p <sub>4</sub>	p <sub>5</sub>	p <sub>6</sub>	p <sub>7</sub>	p <sub>8</sub>	p <sub>9</sub>	p <sub>10</sub>





p <sub>0</sub>	p <sub>1</sub>	p <sub>2</sub>	p <sub>3</sub>	p <sub>4</sub>	<b>p</b> <sub>5</sub>	p <sub>6</sub>	p <sub>7</sub>	p <sub>8</sub>	p <sub>9</sub>	<b>p</b> <sub>10</sub>
										1





p <sub>0</sub>	p <sub>1</sub>	p <sub>2</sub>	p <sub>3</sub>	p <sub>4</sub>	<b>p</b> <sub>5</sub>	p <sub>6</sub>	p <sub>7</sub>	p <sub>8</sub>	p <sub>9</sub>	p <sub>10</sub>





p<sub>0</sub>






## Proof: new commits extend the log



p<sub>0</sub>



## Proof: new commits extend the log







p <sub>0</sub>	p <sub>1</sub>	p <sub>2</sub>	p <sub>3</sub>	p <sub>4</sub>	<b>p</b> <sub>5</sub>	p <sub>6</sub>	p <sub>7</sub>	p <sub>8</sub>	p <sub>9</sub>	p <sub>10</sub>



p <sub>0</sub>	p <sub>1</sub>	p <sub>2</sub>	p <sub>3</sub>	p <sub>4</sub>	p <sub>5</sub>	p <sub>6</sub>	p <sub>7</sub>	p <sub>8</sub>	p <sub>9</sub>	p <sub>10</sub>





p <sub>0</sub>	p <sub>1</sub>	p <sub>2</sub>	p <sub>3</sub>	p <sub>4</sub>	p <sub>5</sub>	p <sub>6</sub>	p <sub>7</sub>	p <sub>8</sub>	p <sub>9</sub>	p <sub>10</sub>



p <sub>0</sub>	p <sub>1</sub>	p <sub>2</sub>	p <sub>3</sub>	p <sub>4</sub>	p <sub>5</sub>	p <sub>6</sub>	p <sub>7</sub>	p <sub>8</sub>	p <sub>9</sub>	p <sub>10</sub>



p <sub>0</sub>	p <sub>1</sub>	p <sub>2</sub>	p <sub>3</sub>	p <sub>4</sub>	p <sub>5</sub>	p <sub>6</sub>	p <sub>7</sub>	p <sub>8</sub>	p <sub>9</sub>	p <sub>10</sub>



p <sub>0</sub>	p <sub>1</sub>	p <sub>2</sub>	p <sub>3</sub>	p <sub>4</sub>	p <sub>5</sub>	p <sub>6</sub>	p <sub>7</sub>	p <sub>8</sub>	p <sub>9</sub>	p <sub>10</sub>







p <sub>0</sub>	p <sub>1</sub>	p <sub>2</sub>	p <sub>3</sub>	p <sub>4</sub>	<b>p</b> <sub>5</sub>	p <sub>6</sub>	p <sub>7</sub>	p <sub>8</sub>	p <sub>9</sub>	p <sub>10</sub>









# Final thoughts

## Implementation effort

I spent a large amount of time proving facts about lists

Majority of the effort related to lemmas for sequential consistency

Component	# of lines	
Common	58	
Lemmas	724	
Replication	140	S W
Сосо	115	ir fe
Replication properties	122	
Coco instant properties	120	
Coco history properties	515	

Specifications were fairly intuitive; took a few hours

## Future work

- Extend the CFTR spec to include:
  - snapshots
  - dynamic membership
  - majorities
- Revise Coco<sub>m</sub> to allow reasoning about:
  - confidentiality
  - governance
  - disaster recovery