

1

Application for Broadcast [Reliable, Uniform, Causal, and Total-Order]

Adi Seredinschi Distributed Programming Laboratory



What do Distributed Systems have in common with Onions?





What do Distributed Systems have in common with Onions?

Answer:

- 1. Layering
- 2. Abstraction
- 3. They make you cry





What do Distributed Systems have in common with Onions?

Answer:

- 1. Layering
- 2. Abstraction
- 3. They make you cry





What do Distributed Systems have in common with Onions?





What do Distributed Systems have in common with Onions?





Links

Network

Question:

What do Distributed Systems have in common with Onions?

Answer: 1. Layering Application 2. Abstraction 3. They make you cry Broadcast Point-to-Point

2



What do Distributed Systems have in common with Onions?

Answer:

- 1. Layering
- 2. Abstraction
- 3. They make you cry





What do Distributed Systems have in common with Onions?

Answer:

- 1. Layering
- 2. Abstraction
- 3. They make you cry



Let's design a simple application





Broadcast (we will investigate different properties..)

CAMIPRO-Bitcoin

- The canonical Bitcoin design
 - Uses gossip (best-effort broadcast)
 - Relies heavily on crypto no time to cover that
- $\boldsymbol{\cdot}$ We will not discuss the canonical design
 - Instead, we will design our own version of Bitcoin
 - Optimized for CAMIPRO



CAMIPRO-Bitcoin





Conceptual goals:

- 1.Replace traditional CAMIPRO
 - Based on CHF

2. Make banks obsolete

CAMIPRO-Bitcoin





Conceptual goals:

- 1.Replace traditional CAMIPRO
 - Based on CHF
- 2.Make banks obsolete





CAMIPRO-Bitcoin





Conceptual goals:

- 1.Replace traditional CAMIPRO
 - Based on CHF

2.Make banks obsolete



LEDGER Alice 20 Bob 15







• Ledger

















• Transaction $\begin{bmatrix} A \rightarrow B \\ 5.0 \text{ BTC} \end{bmatrix}$ "Alice gives some money to Bob"





• Transaction $\begin{bmatrix} A \rightarrow B \\ 5.0 \text{ BTC} \end{bmatrix}$ "Alice gives some money to Bob"

Let's see how they all fit together..





















Blockchain



Recognise any abstractions?





What kind of broadcast?





What kind of broadcast?

PROPERTIES (guarantees)





The Four Dimensions of a cube





The Four Dimensions of a cube





Do we need Reliability?

- Consider the following:
 - User A starts TX1
 - Use best-effort broadcast
 - Validity + !Duplication + !Creation
 - "the burden of reliability is on the sender"
 - Lacks Agreement \Rightarrow Nodes diverge





Do we need Reliability?

- Consider the following:
 - User A starts TX1
 - Use best-effort broadcast
 - Validity + !Duplication + !Creation
 - "the burden of reliability is on the sender"
 - Lacks Agreement \Rightarrow Nodes diverge







Do we need Uniformity?

- Consider the following:
 - User A starts TX1
 - Use regular reliable broadcast
 - Validity + !Duplication + !Creation
 + Agreement for correct nodes
 - Is it OK to deliver and crash?





Do we need Uniformity?

- Consider the following:
 - User A starts TX1
 - Use regular reliable broadcast
 - Validity + !Duplication + !Creation
 + Agreement for correct nodes
 - Is it OK to deliver and crash?



User A

User B





Do we need Uniformity?

- Consider the following:
 - User A starts TX1
 - Use regular reliable broadcast
 - Validity + !Duplication + !Creation
 + Agreement for correct nodes
 - Is it OK to deliver and crash?



What if User B observes TX1 before the nodes crash?

"Uniformity is important if nodes interact with the external world"



Do we need Causality? (partial ordering)

- Consider the following:
 - User A starts TX1 and TX2
 - Use uniform reliable broadcast
 - Validity + !Duplication + !Creation + Uniform Agreement
 ⇒ Applies to all nodes
 - All nodes deliver both TX, but the order may differ





TX2

Node

Blockchair

TX2

Bitcoin network

Blockchain

TX1 TX2

Blockchair

TX1 TX2

Blockchain

TX1 TX2

Blockchain

TX1

Node

Blockchain

PROPERTIES

Do we need Causality? (partial ordering)

- Consider the following:
 - User A starts TX1 and TX2
 - Use uniform reliable broadcast
 - Validity + !Duplication + !Creation + Uniform Agreement
 ⇒ Applies to all nodes

User A

User B

• All nodes deliver both TX, but the order may differ

What if TX2 depends on money from TX1 ? ----



Do we need (Total) Ordering?

- Consider the following:
 - User A starts TX1
 - User B starts TX2 among these two



- Validity + !Duplication + !Creation + Uniform Agreement + Causality
 ⇒ Respect causal dependencies among TXs
- All nodes deliver both TX, but the order may differ

No dependency





Do we need (Total) Ordering?

- Consider the following:
 - User A starts TX1
 - User B starts TX2 among these two



Bitcoin network

- Use causal-order uniform reliable broadcast
 - Validity + !Duplication + !Creation + Uniform Agreement + Causality
 ⇒ Respect causal dependencies among TXs
- All nodes deliver both TX, but the order may differ

No dependency



Commutativity counter-example





Commutativity counter-example



Commutativity counter-example





- Reliability
 - Sender crashes
 - Agreement
- Uniformity
 - Again, crashes
 - Interaction with outside world
- Causality
 - Partial order
 - Dependencies among TXs
- Total order
 - Commutativity





- Reliability
 - Sender crashes
 - Agreement
- Uniformity
 - Again, crashes
 - Interaction with outside world
- Causality
 - Partial order
 - Dependencies among TXs
- Total order
 - Commutativity





- Reliability
 - Sender crashes
 - Agreement
- Uniformity
 - Again, crashes
 - Interaction with outside world
- Causality
 - Partial order
 - Dependencies among TXs
- Total order
 - Commutativity



All properties are desirable



PERFORMANCE

Goals:

Replace traditional CAMIPRO
 Make banks obsolete







PERFORMANCE

Goals:

1.Replace traditional CAMIPRO

2.Make banks obsolete







PERFORMANCE

Goals:

1.Replace traditional CAMIPRO

2.Make banks obsolete







Specification





Module:

Name: CAMIPRO-Bitcoin-Broadcast, instance *cbit*

Properties:

- URB1, URB2, URB3, URB4
- Causal Order (CO)
- Total order (TO)





Specification





Module:

Name: CAMIPRO-Bitcoin-Broadcast, instance cbit

Properties:

- URB1, URB2, URB3, URB4
- Causal Order (CO)
- Total order (TO)



Events:

Request: (*cbit, Start* | TX):

Attempts to commit TX

Indication: (*cbit, Status* | TX, s):

Indicates the status s∈{"Abort","Commit"} of TX





Module:

Name: CAMIPRO-Bitcoin-Broadcast, instance cbit

Properties:

- URB1, URB2, URB3, URB4
- Causal Order (CO)
- Total order (TO)





CAMIPRO-Bitcoin Request Indication

CAMIPRO-Bitcoin-Broadcast

Events:

Request: (*cbit, Start* | TX): Attempts to commit TX

Indication: (*cbit, Status* | TX, s): Indicates the status s∈{"Abort","Commit"} of TX



Causal consistency — one last look

• Given two operations, op1 and op2

op1 → op2 (causally precedes)
 if any of the following three cases applies:

(a) FIFO: A process invokes op1 and then invokes op2

(b) Local: A process invokes the PUT operation op1 and another process invokes the GET operation op2, where op2 observes the written value of op1

(c) Transitivity: There exists an intermediate operation op' such that op1 \rightarrow op' and op' \rightarrow op2.







Causal consistency — one last look





Causal consistency — one last look



