

# Catena: Efficient Non-equivocation via *bitcoin*

**Alin Tomescu**  
[alinush@mit.edu](mailto:alinush@mit.edu)

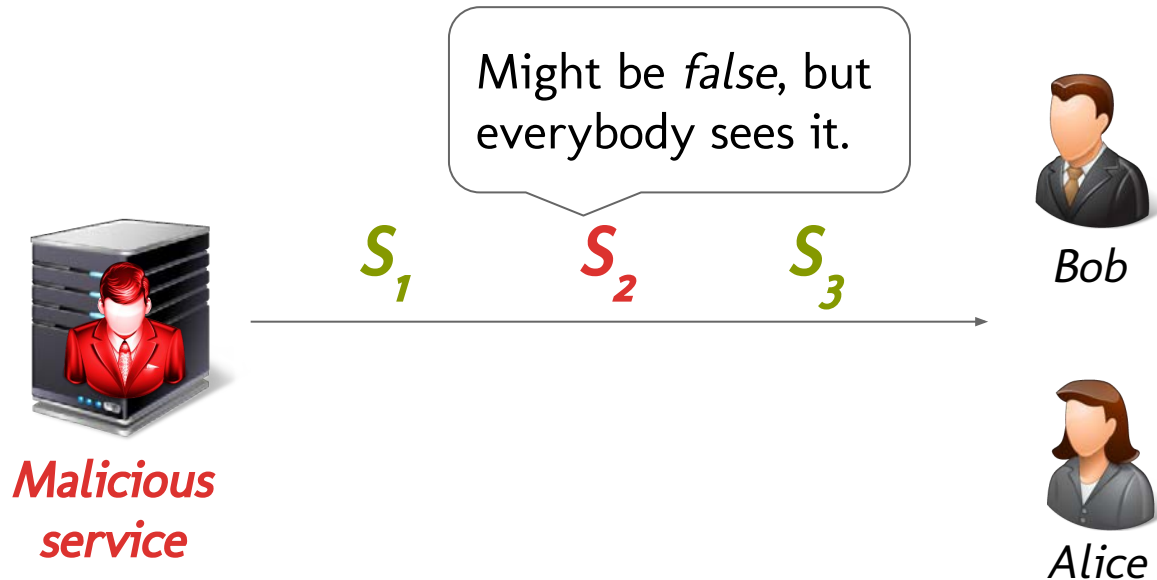
**Srinivas Devadas**  
[devadas@mit.edu](mailto:devadas@mit.edu)

# Overview

1. What?
2. How?
3. Why?

# The equivocation problem


**Non-equivocation:** "Saying the same thing to everybody."



Various clip art images in this document © unknown. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/fairuse>.

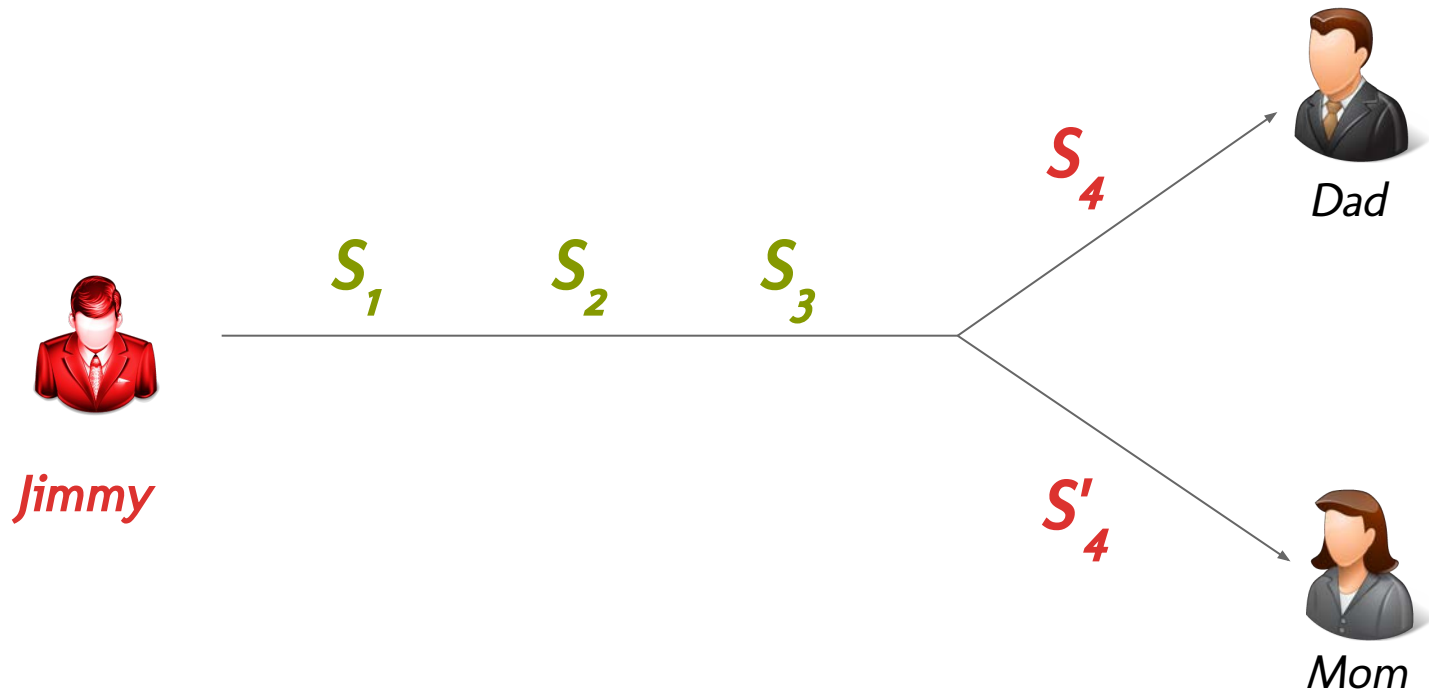
# The equivocation problem

**Equivocation:** *"Saying different things to different people."*



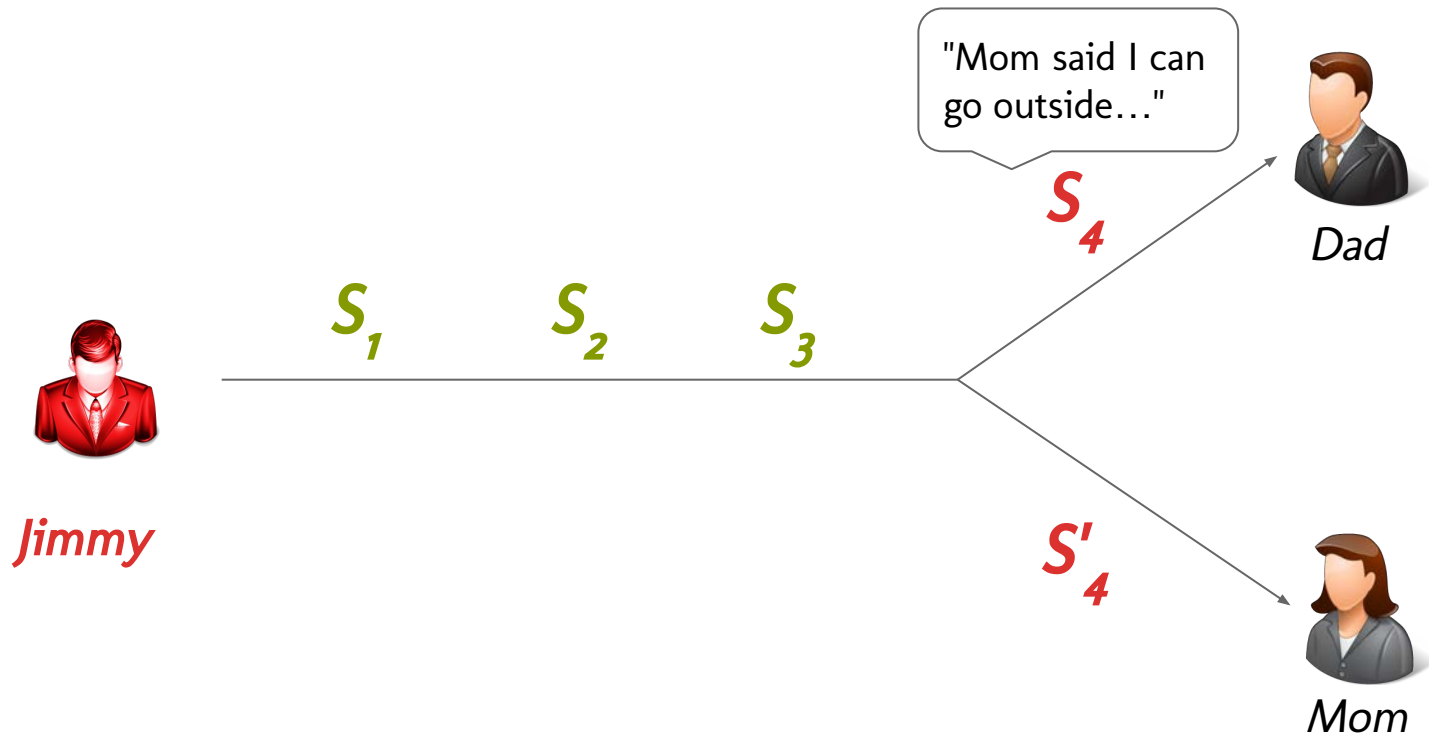
# The equivocation problem

**Equivocation:** *"Saying different things to different people."*



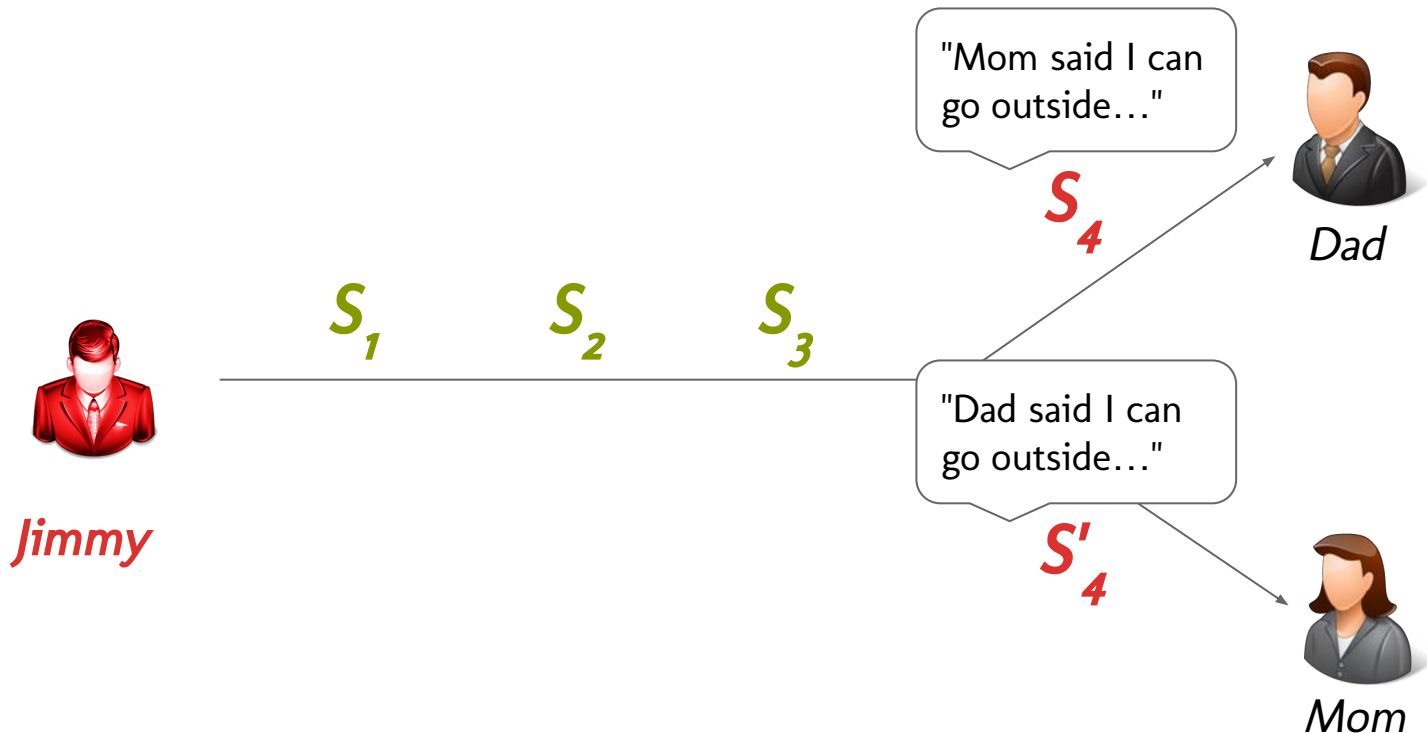
# The equivocation problem

**Equivocation:** *"Saying different things to different people."*




# The equivocation problem

**Equivocation:** *"Saying different things to different people."*



# What is **equivocation**?


- $s_2$ : Leave Alice's key intact, add fake  $PK_B'$  for Bob






# What is **equivocation**?

- $s_2$ : Leave Alice's key intact, add fake  $PK_B'$  for Bob




# What is **equivocation**?

- $s_2'$ : Leave Bob's key intact, add fake  $PK_A'$  for Alice




# What is **equivocation**?

- Alice not impersonated in her view, but Bob is.




# What is **equivocation**?

- Bob not impersonated in his view, but Alice is.




# What is **equivocation**?

- Obtain fake keys for each other  $\Rightarrow$  **MITM**



# What is **equivocation**?

**Bad:** "Stating different things to different people."



# Catena prevents equivocation!

Damn.



*Malicious  
service with  
Catena*

$S_1$

$S_2$

$S_3$




Bob



Alice

# Catena prevents equivocation!





# So what?

## Secure software update

- Attacks on Bitcoin binaries

## Secure messaging

- HTTPS
- *"We assume a PKI."*



### 0.13.0 Binary Safety Warning

17 August 2016

#### Summary

Bitcoin.org has reason to suspect that the binaries for the upcoming Bitcoin Core release will likely be targeted by state sponsored attackers. As a website, Bitcoin.org does not have the technical resources to guarantee that we can defend ourselves from attackers of this calibre. We ask the Bitcoin community, and in particular the Chinese Bitcoin community to be extra vigilant when downloading binaries from our website.



# So what?

## Secure software update

- Attacks on Bitcoin binaries

## Secure messaging

- HTTPS
- *"We assume a PKI."*

## "Blockchain" for X



### 0.13.0 Binary Safety Warning

17 August 2016

#### Summary

Bitcoin.org has reason to suspect that the binaries for the upcoming Bitcoin Core release will likely be targeted by state sponsored attackers. As a website, Bitcoin.org does not have the technical resources to guarantee that we can defend ourselves from attackers of this calibre. We ask the Bitcoin community, and in particular the Chinese Bitcoin community to be extra vigilant when downloading binaries from our website.




# 10,000 feet view

- Bitcoin-based append-only log,
  - Generalizes to other cryptocurrencies
- ...as hard-to-fork as the Bitcoin blockchain
  - Want to fork? Do some work!
- ...but efficiently auditable
  - 600 bytes / statement (but can batch!)
  - 80 bytes / Bitcoin block
- Java implementation (3500 SLOC)
  - <https://github.com/alinush/catena-java>

# Overview


1. What?
2. How?
  - a. **Bitcoin background**
3. Why?

# Bitcoin blockchain




- Hash chain of blocks
  - Arrows are *hash pointers*
- Merkle tree of TXNs in each block
- Proof-of-work (PoW) consensus

# Bitcoin blockchain




- Transactions mint coins

# Bitcoin blockchain



- Transactions mint coins
- Output = # of coins and owner's PK

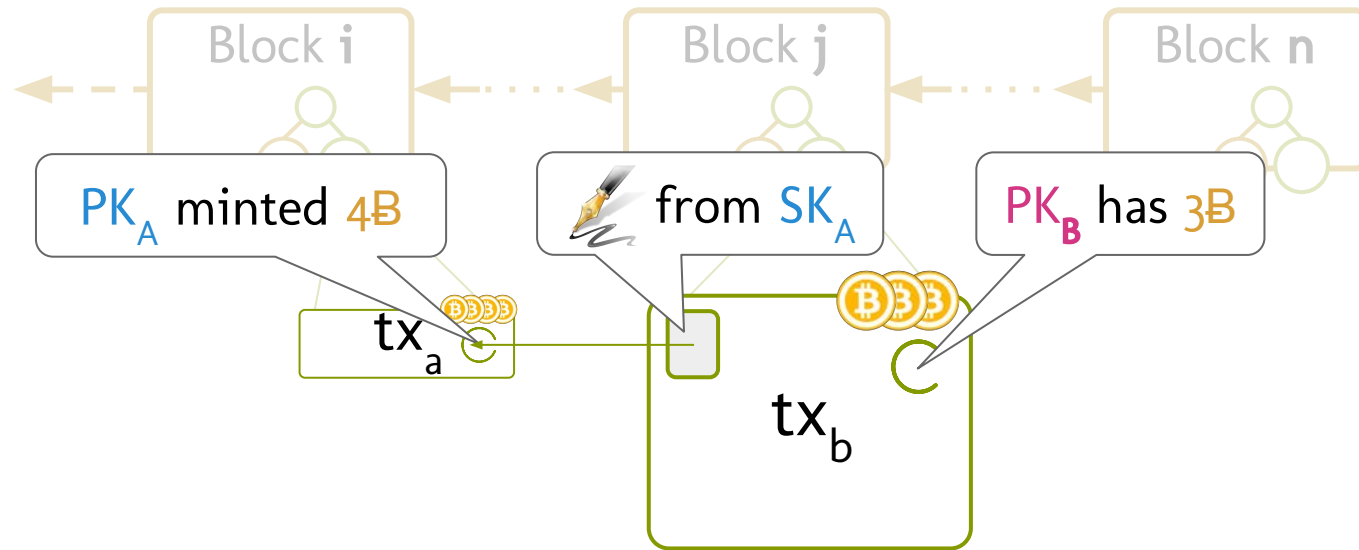
# Bitcoin blockchain



- Transactions mint coins
- Output = # of coins and owner's PK
- Transactions transfer coins (and pay fees)

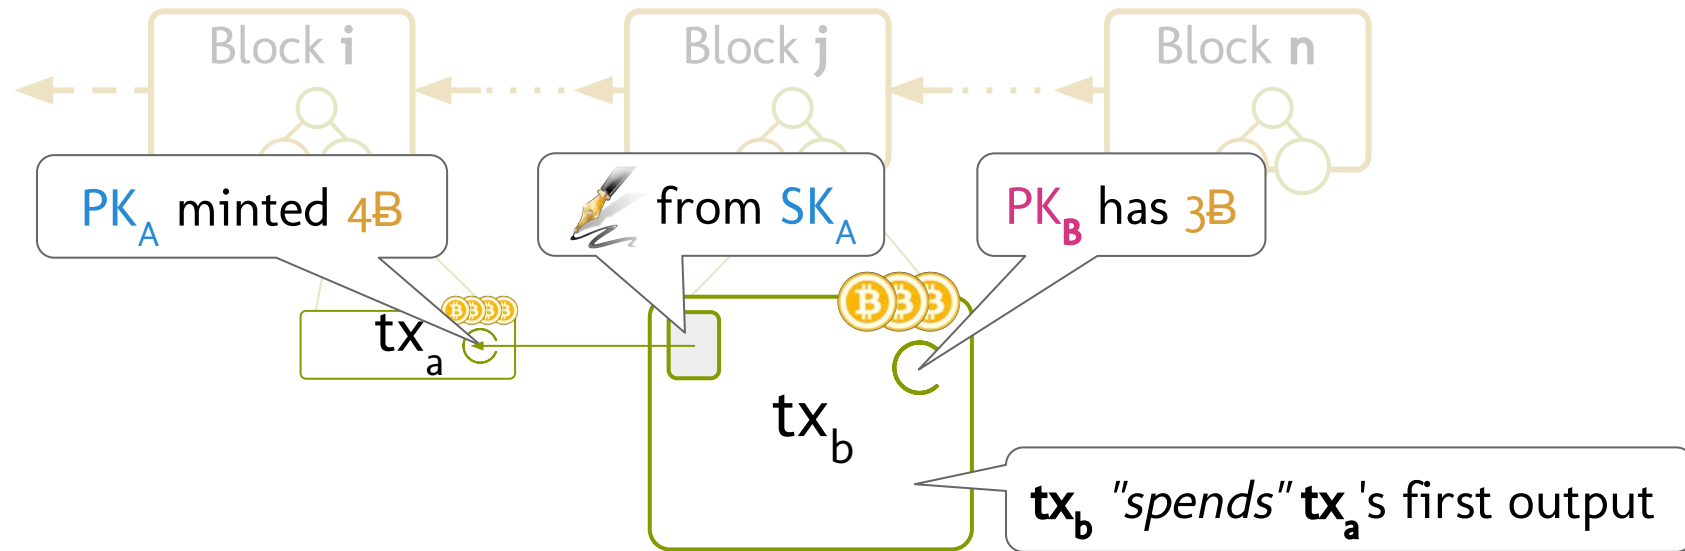


# Bitcoin blockchain



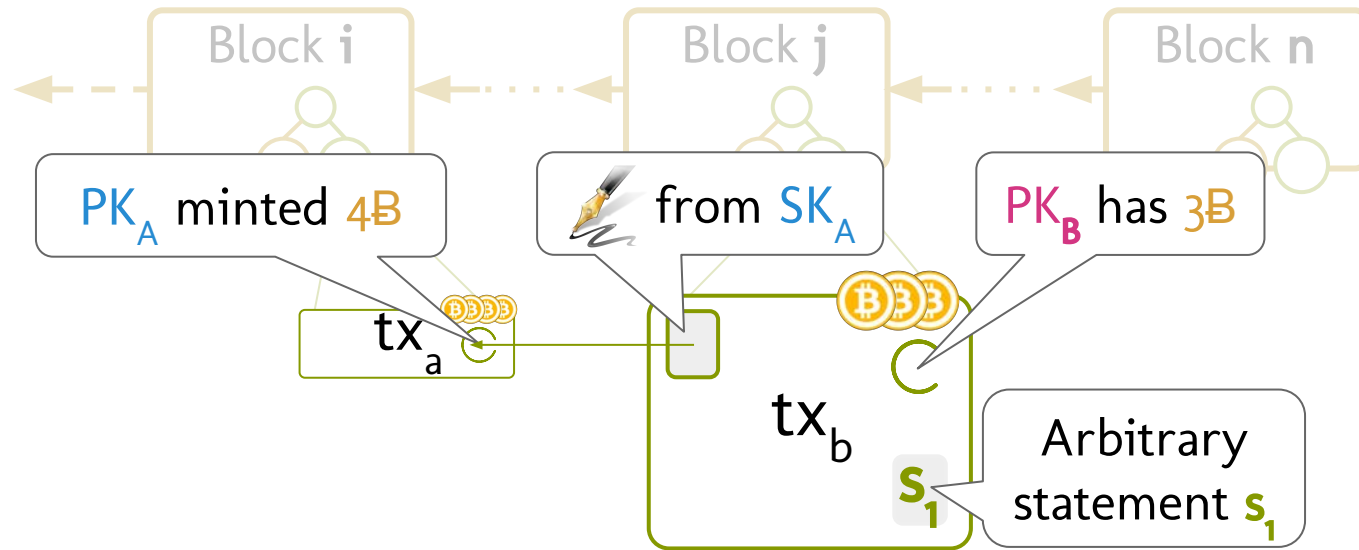
- Transactions mint coins
- Output = # of coins and owner's PK
- Transactions transfer coins (and pay fees)
- Input = hash pointer to output + digital signature

# Bitcoin blockchain



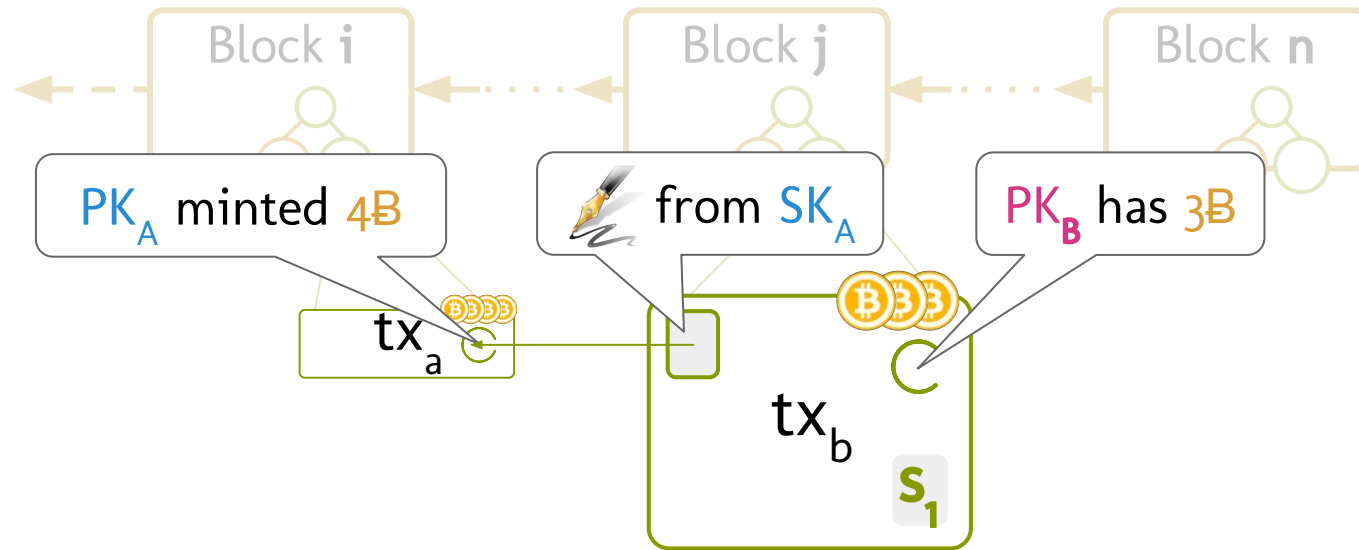
- Transactions mint coins
- Output = # of coins and owner's PK
- Transactions transfer coins (and pay fees)
- Input = hash pointer to output + digital signature

# Bitcoin blockchain



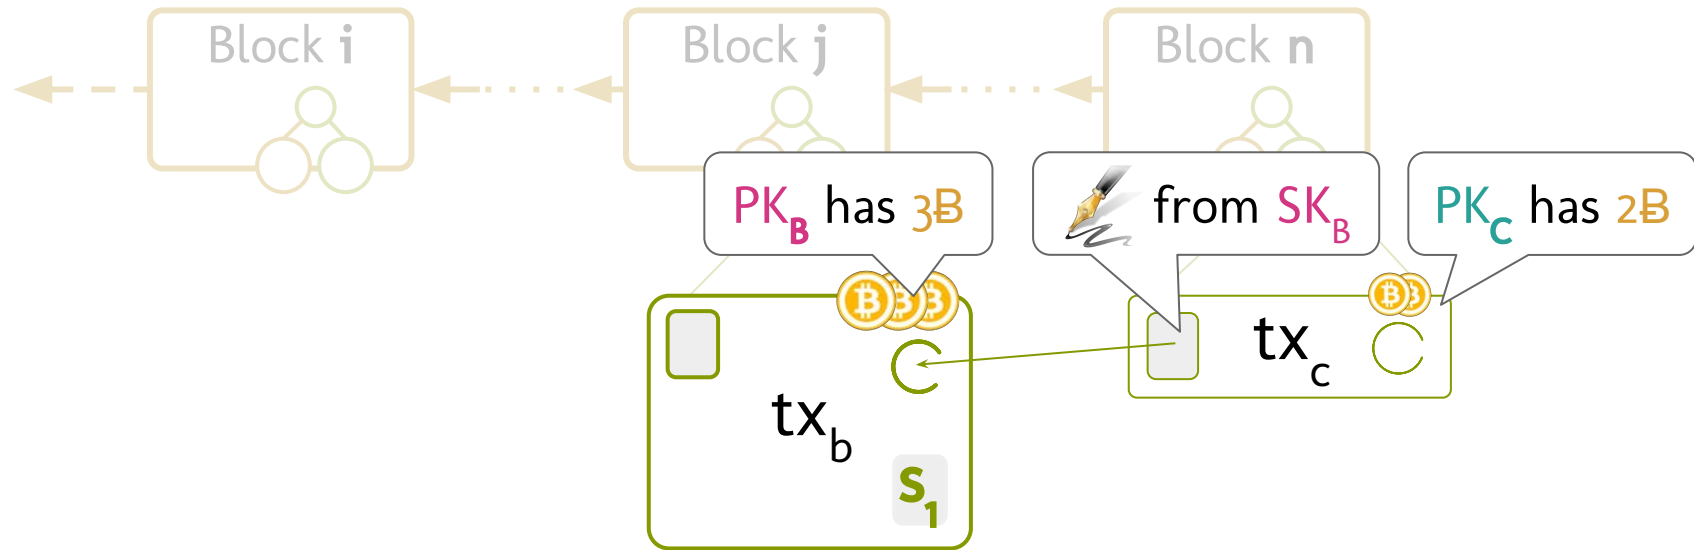
Data can be embedded in TXNs.

# Bitcoin blockchain




Alice gives Bob 3B,  
Bitcoin *miners* collected 1B as a *fee*.

# Bitcoin blockchain



Bob gives Carol 2B,  
Bitcoin *miners* collected another B as a fee.


# Bitcoin blockchain



**No double-spent coins:** A TXN output can only be referred to by a single TXN input.


# Moral of the story

Proof-of-work (PoW) consensus  $\Rightarrow$  No double spends



# Moral of the story

Proof-of-work (PoW) consensus  $\Rightarrow$  No double spends

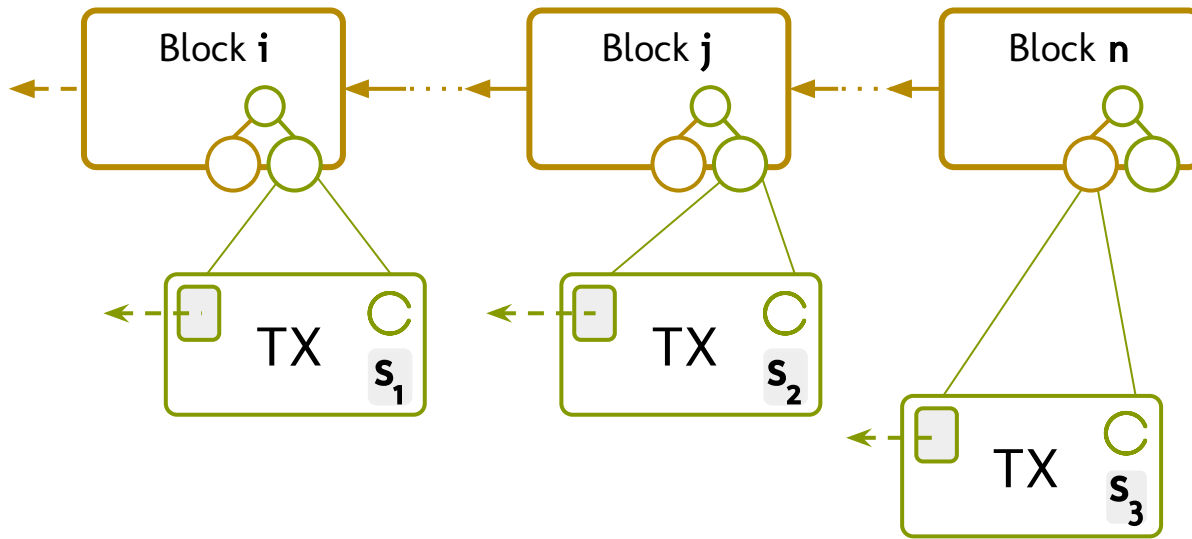




# Overview

1. What?
2. How?
  - a. Bitcoin background
  - b. Previous work**
3. Why?

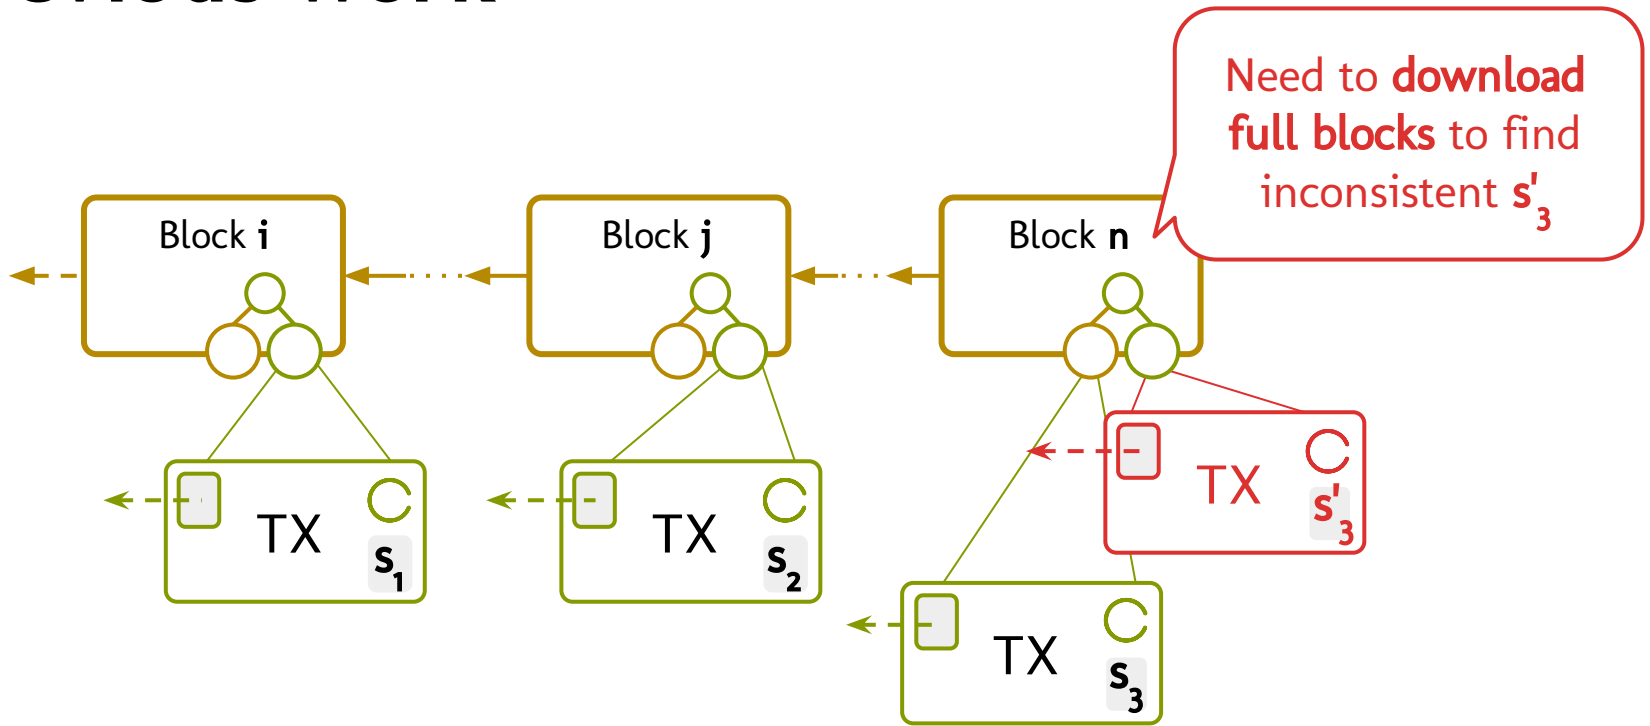
# Previous work



**blockstack**



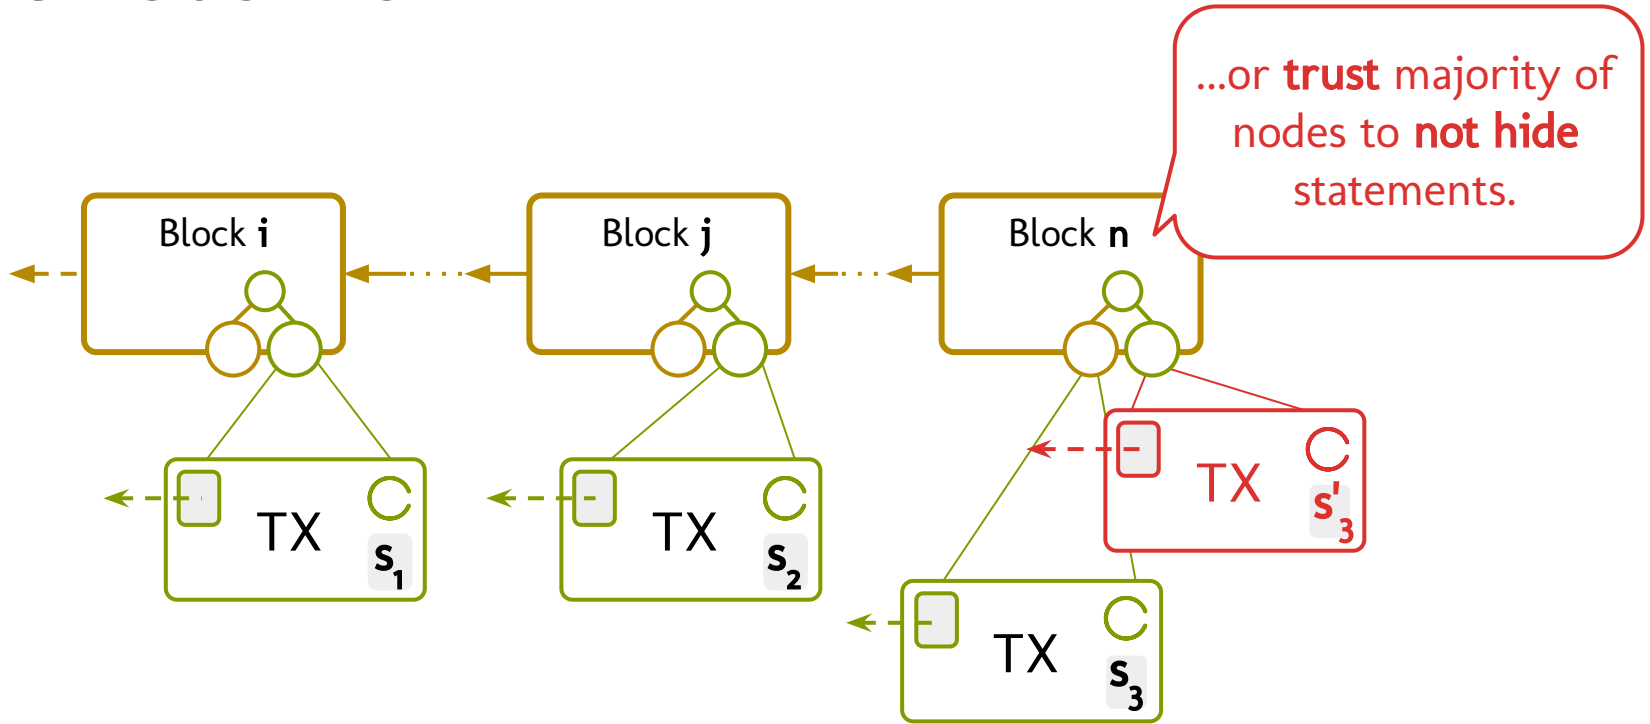
# Previous work



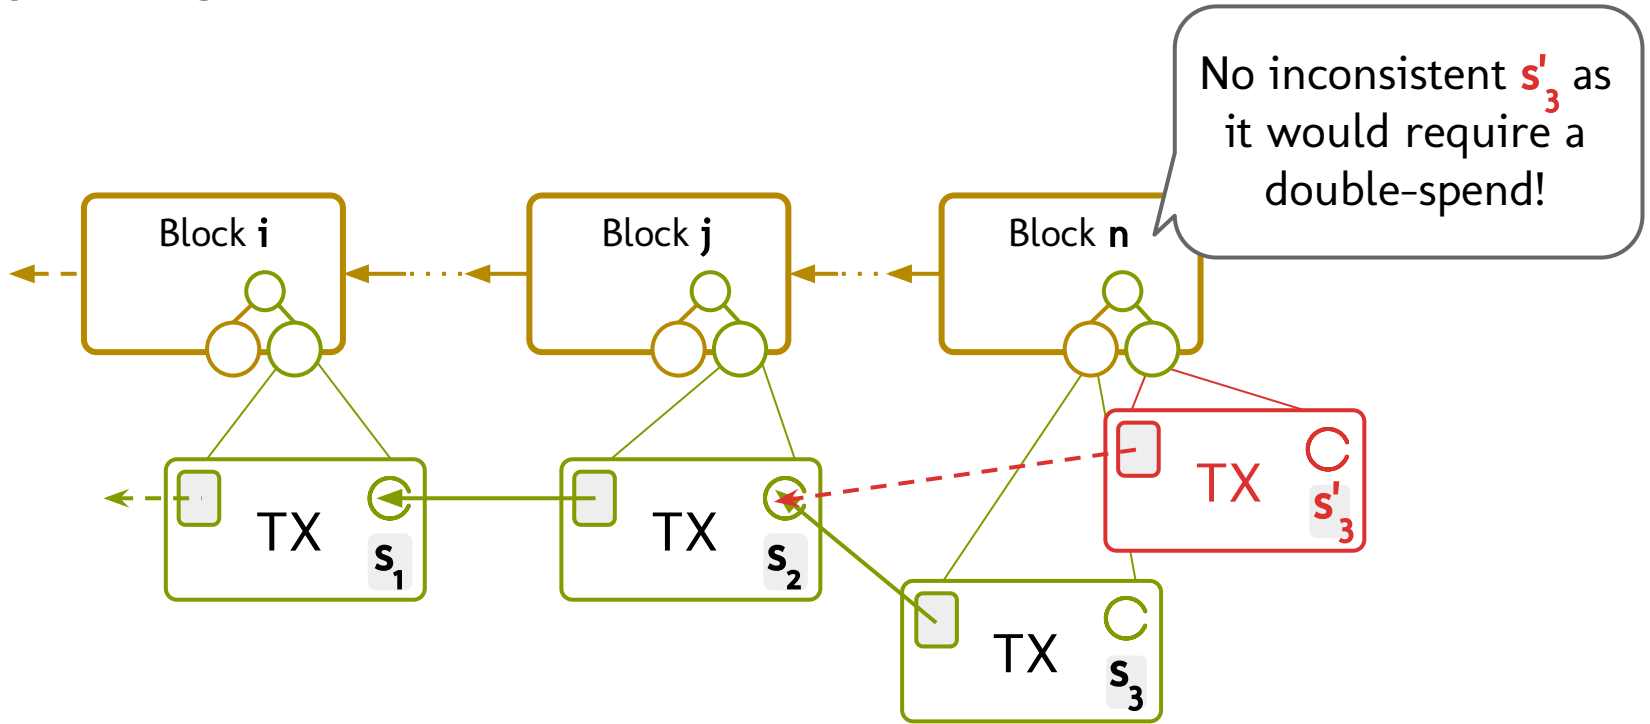
blockstack



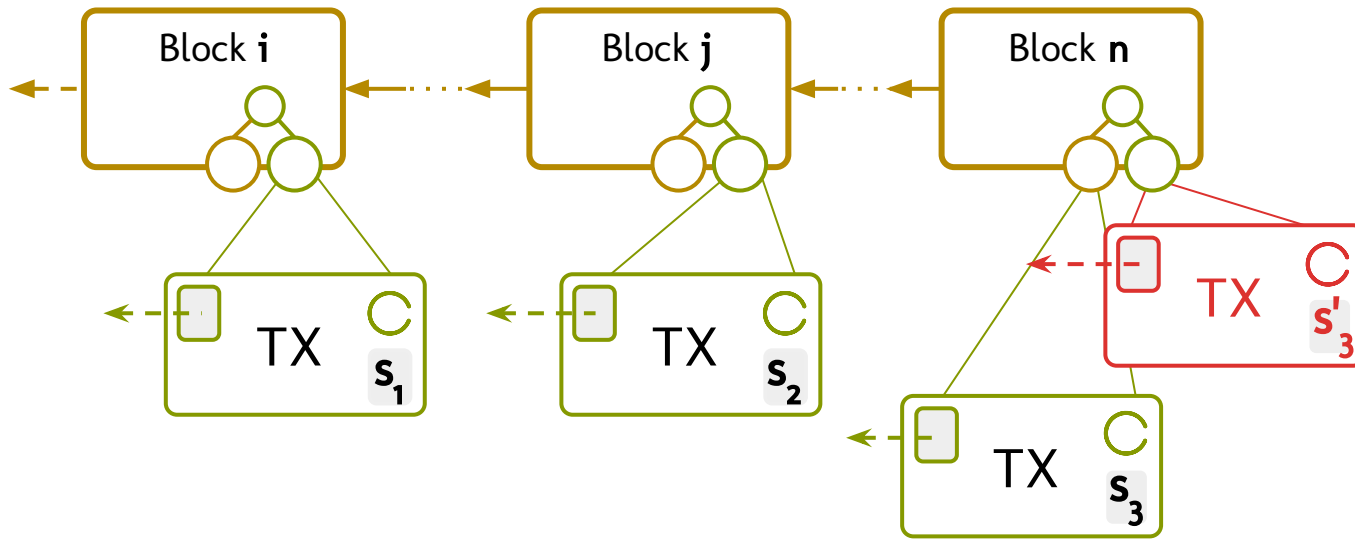
# Previous work



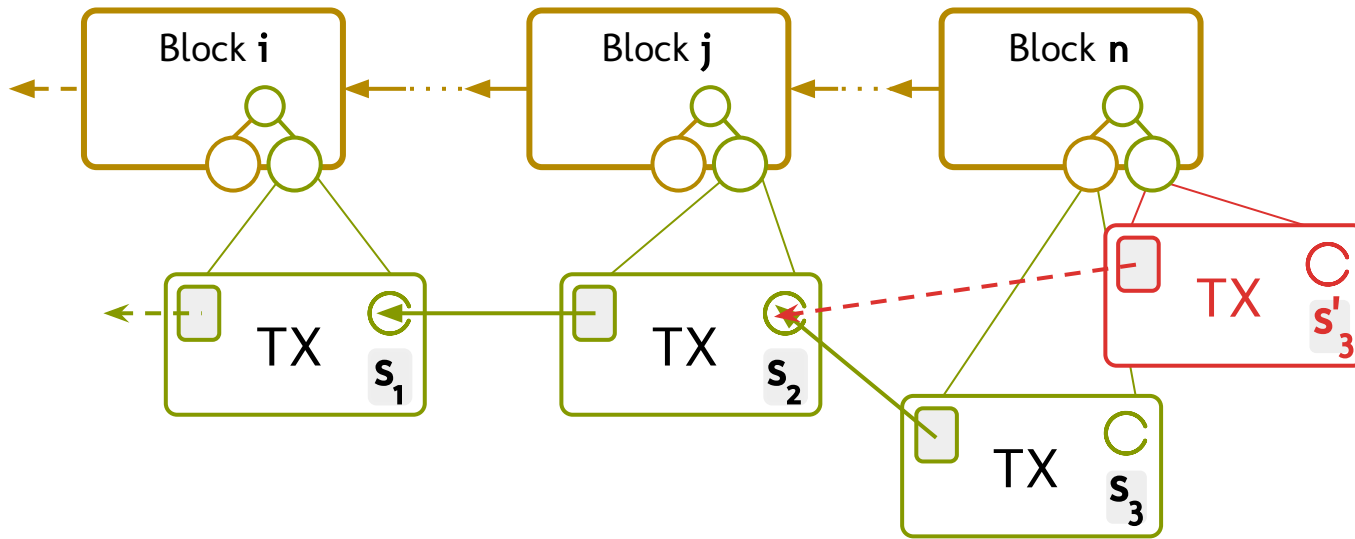
# Our work



# Previous work



# Our work

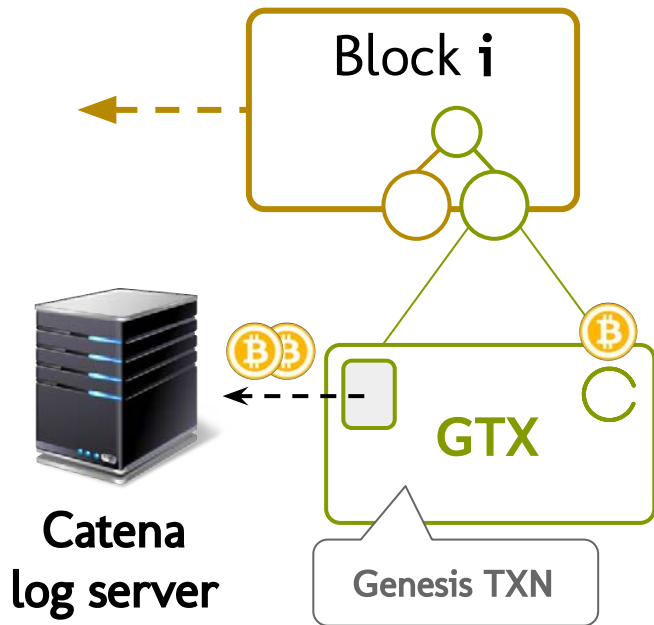


# Overview

1. What?
2. How?
  - a. Bitcoin background
  - b. Previous work
  - c. Design**
3. Why?

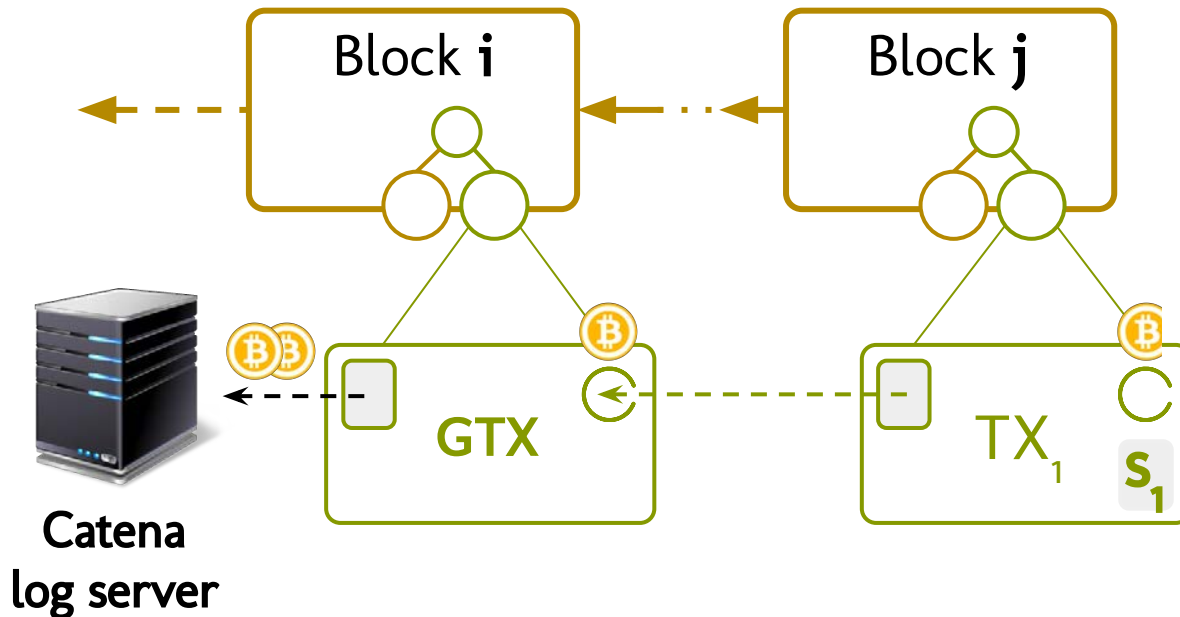


# Starting a Catena log




- *Genesis TXN* (GTX) = log's "public key"
- Coins from server back to server (minus fees)

# Appending to a Catena log




- TX<sub>1</sub> "spends" GTX's output, publishes s<sub>1</sub>
- Coins from server back to server (minus fees)
- Inconsistent s'<sub>1</sub> would require a double-spend

# Appending to a Catena log




- TX<sub>2</sub> "spends" TX<sub>1</sub>'s output, publishes s<sub>2</sub>
- Coins from server back to server (minus fees)
- Inconsistent s<sub>2</sub>' would require a double-spend

# Appending to a Catena log



- Server is compromised, still cannot equivocate.


# Appending to a Catena log



## Advantages:

- (1) Hard to fork
- (2) Efficient to verify

# Appending to a Catena log




## Advantages:

- (1) Hard to fork
- (2) Efficient to verify

## Disadvantages:

- (1) 6-block confirmation delay

# Appending to a Catena log




## Advantages:

- (1) Hard to fork
- (2) Efficient to verify

## Disadvantages:

- (1) 6-block confirmation delay
- (2) 1 statement every 10 minutes

# Appending to a Catena log



## Advantages:


- (1) Hard to fork
- (2) Efficient to verify

## Disadvantages:

- (1) 6-block confirmation delay
- (2) 1 statement every 10 minutes
- (3) Must pay Bitcoin TXN fees



# Appending to a Catena log



## Advantages:

- (1) Hard to fork
- (2) Efficient to verify


## Disadvantages:

- (1) 6-block confirmation delay
- (2) 1 statement every 10 minutes
- (3) Must pay Bitcoin TXN fees
- (4) No freshness guarantee

# Overview

1. Catena: What?
2. Catena: How?
  - a. Bitcoin background
  - b. Previous work
  - c. Design
  - d. Efficient auditing**
3. Potentially-interesting applications

# Efficient auditing

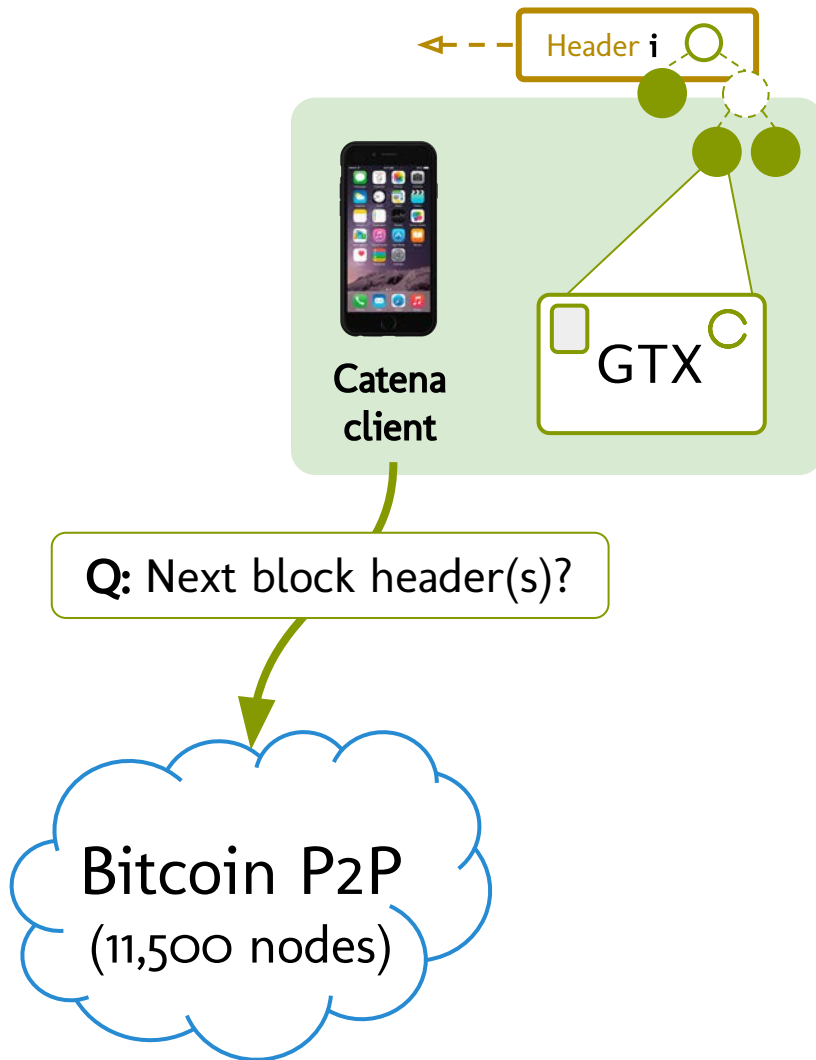


Bitcoin P2P  
(11,500 nodes)

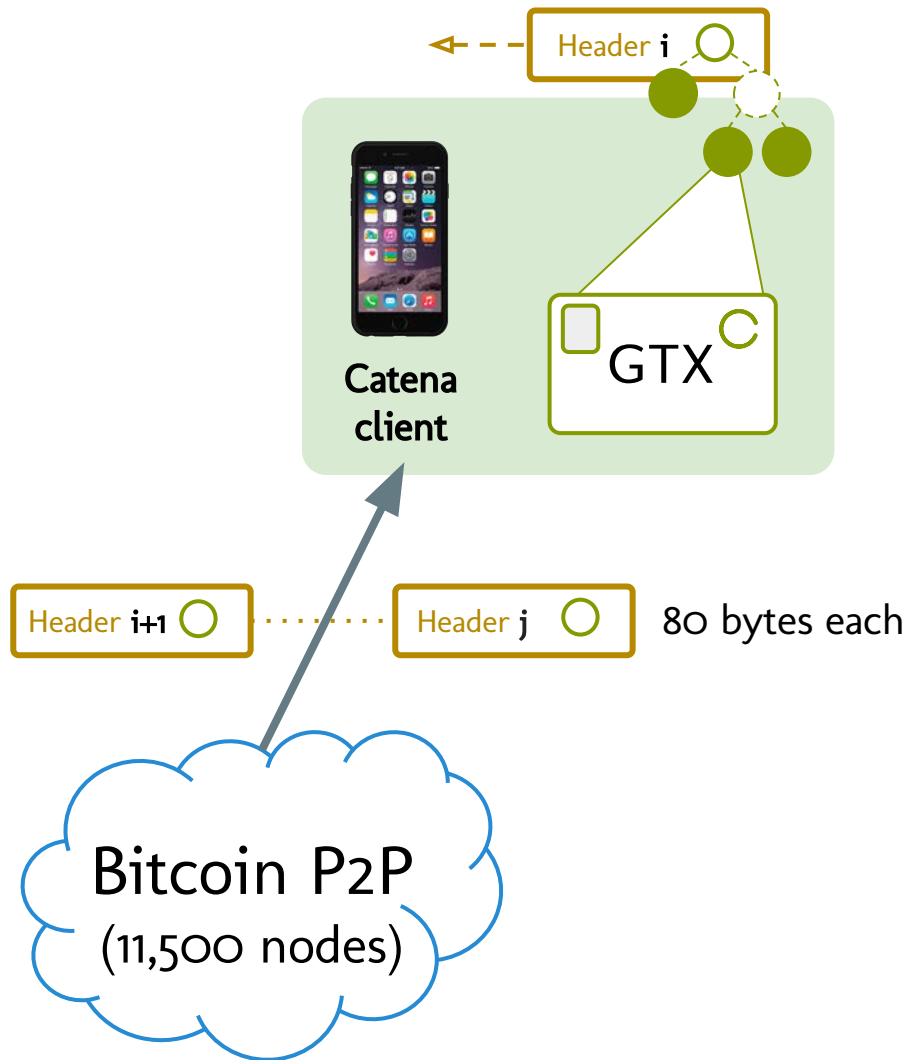


Catena  
log server


# Efficient auditing



# Efficient auditing



# Efficient auditing

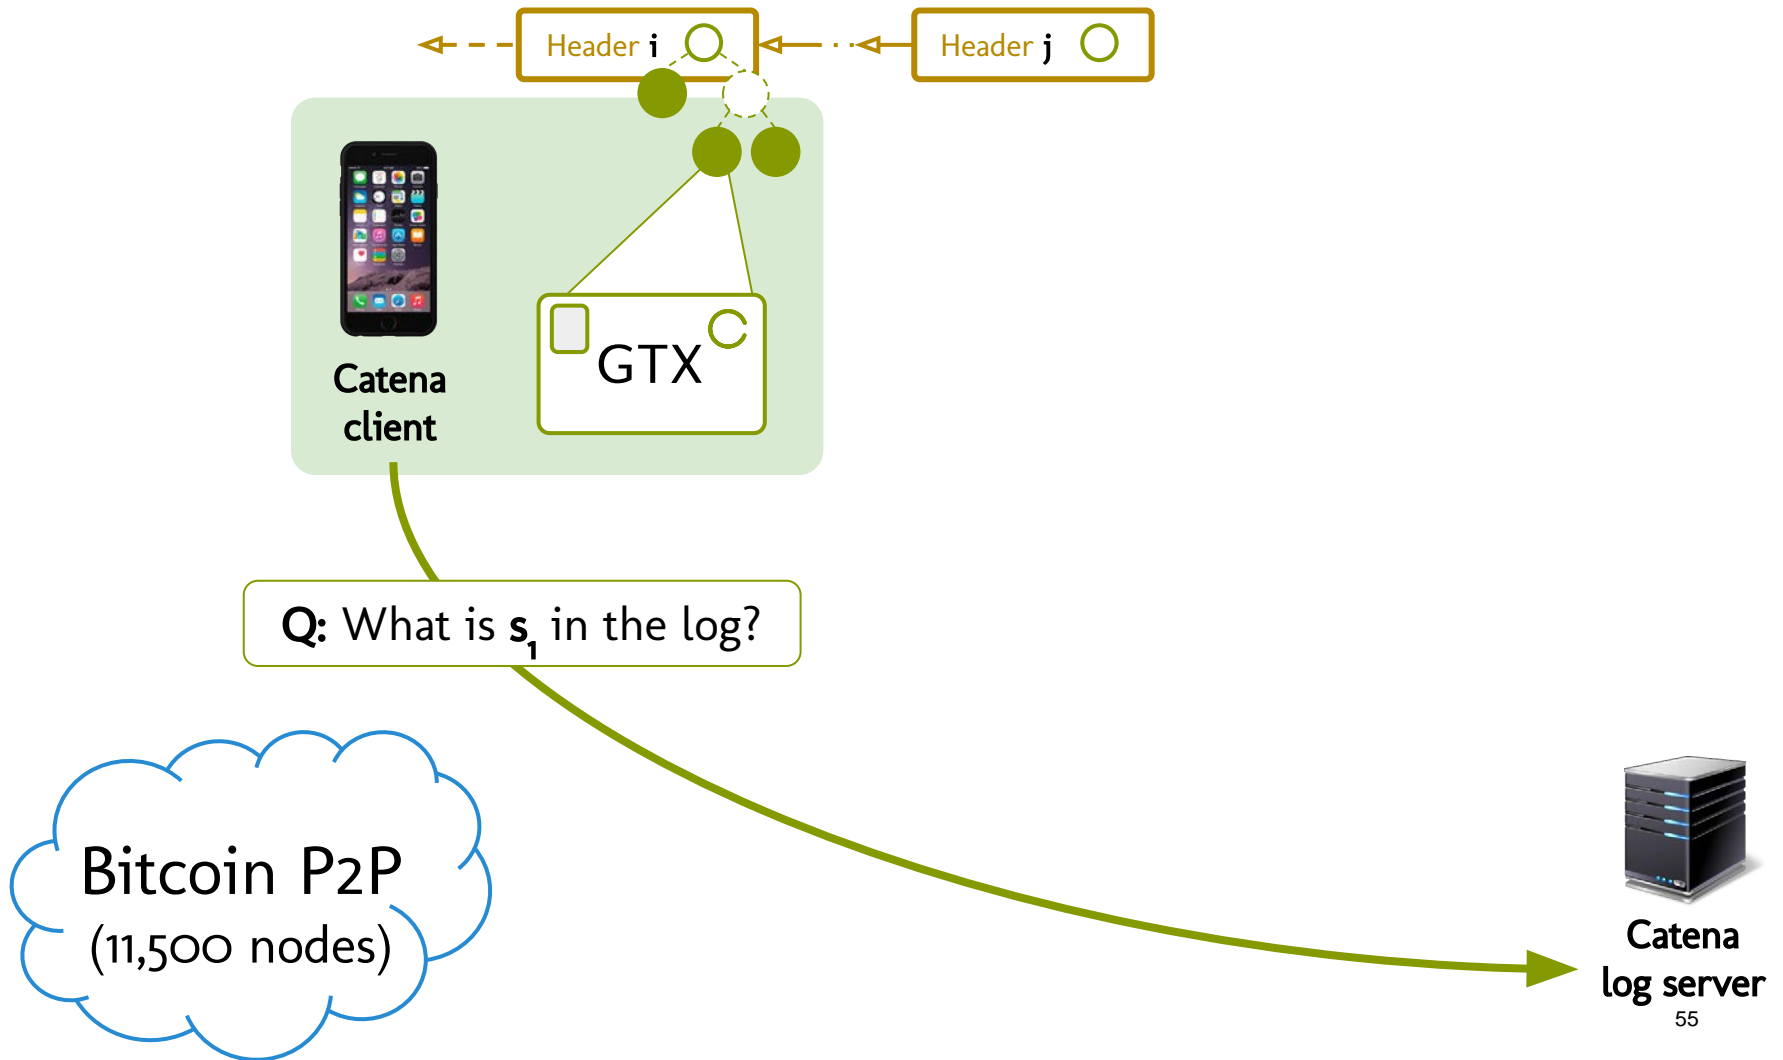


Bitcoin P2P  
(11,500 nodes)

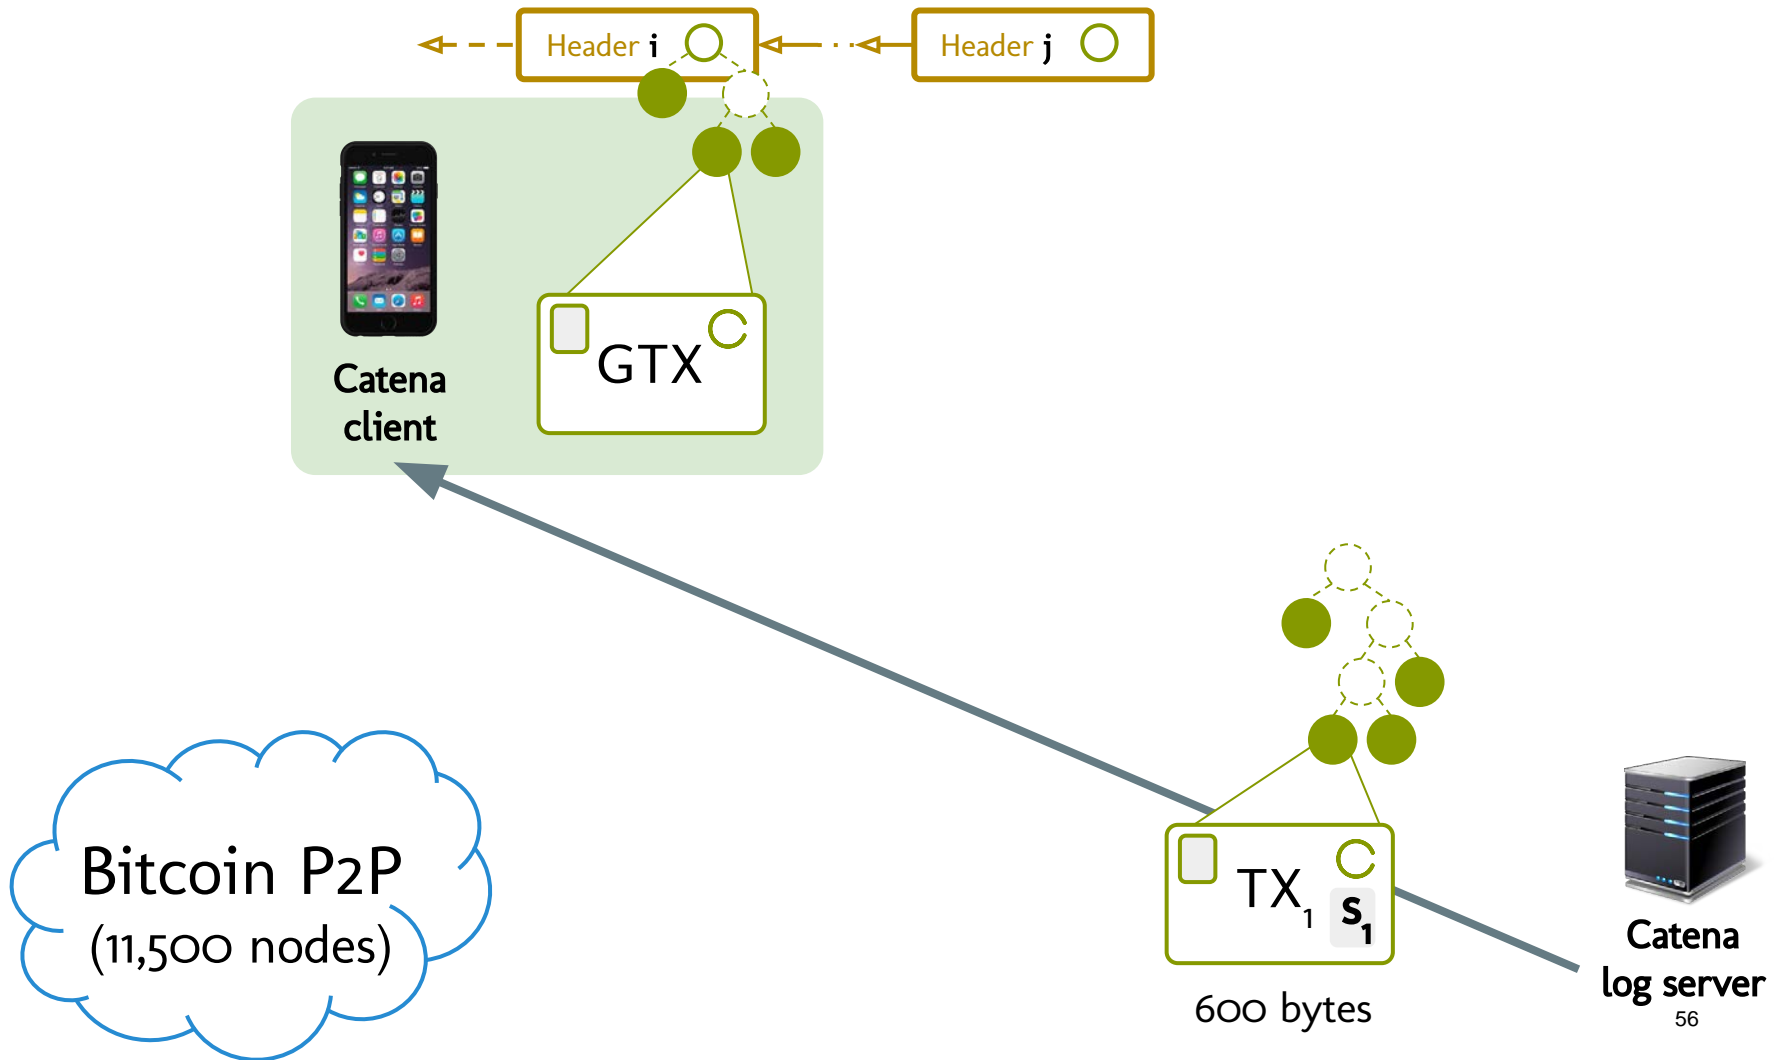


Catena  
log server

# Efficient auditing




# Efficient auditing





# Efficient auditing




Bitcoin P2P  
(11,500 nodes)



Catena  
log server

# Efficient auditing





Q: Next block header(s)?

Bitcoin P2P  
(11,500 nodes)




Catena  
log server

# Efficient auditing



**Catena  
log server**

# Efficient auditing




Bitcoin P2P  
(11,500 nodes)



Catena  
log server

# Efficient auditing

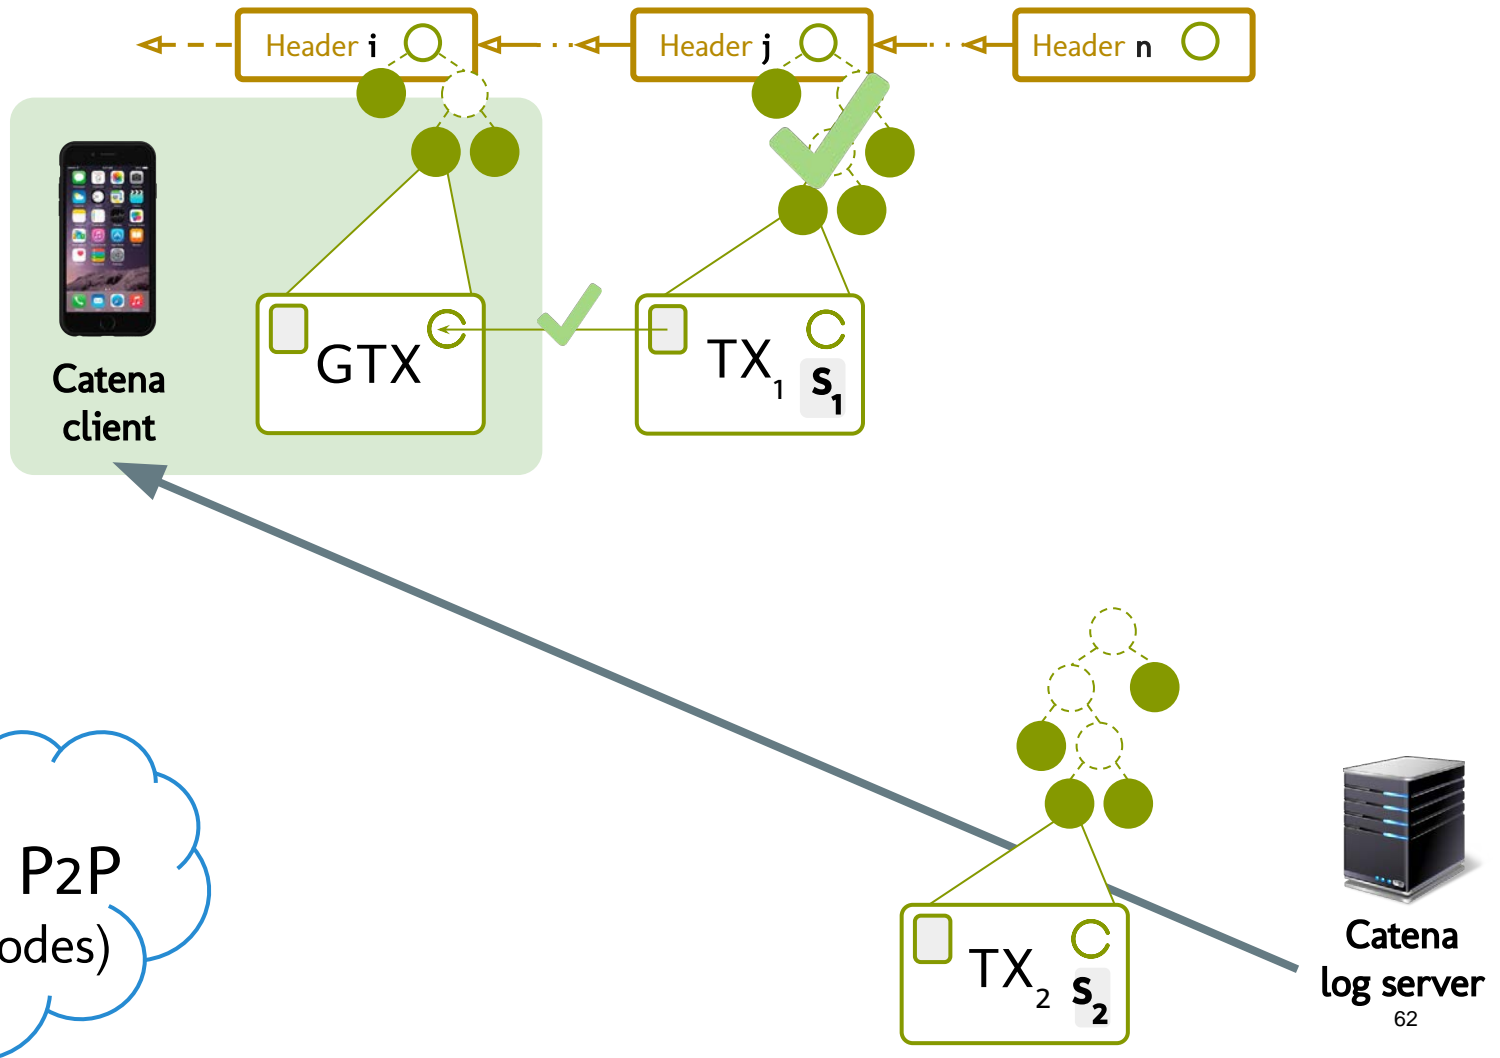


Q: What is  $s_2$  in the log?


Bitcoin P2P  
(11,500 nodes)

Catena  
log server

# Efficient auditing




# Efficient auditing



Bitcoin P2P  
(11,500 nodes)

Catena  
log server

# Auditing bandwidth



*e.g.*, **500K** block headers + **10K** statements = **~46 MB**  
(80 bytes each) (around 600 bytes each)



# Overview

1. What?
2. How?
  - a. Bitcoin background
  - b. Previous work
  - c. Design
  - d. Efficient auditing
  - e. Scalability**
3. Why?

# Catena scalability



Catena client 1



Catena client 2

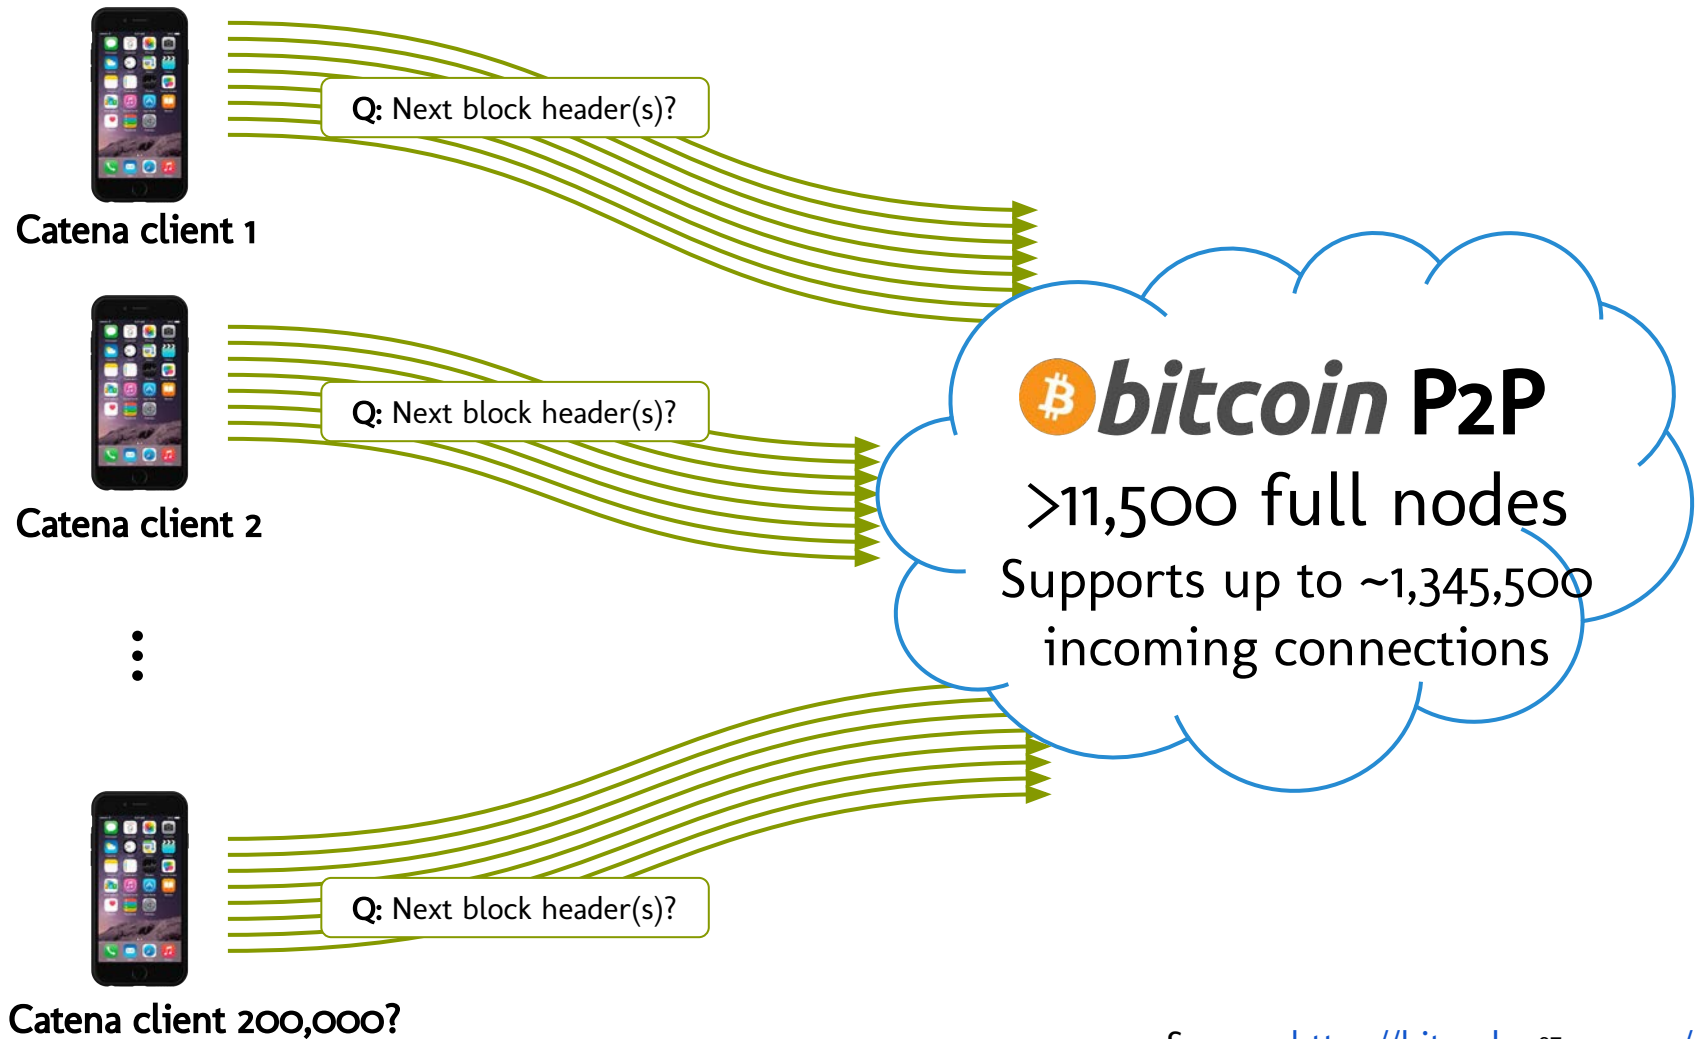


Catena client 200,000?

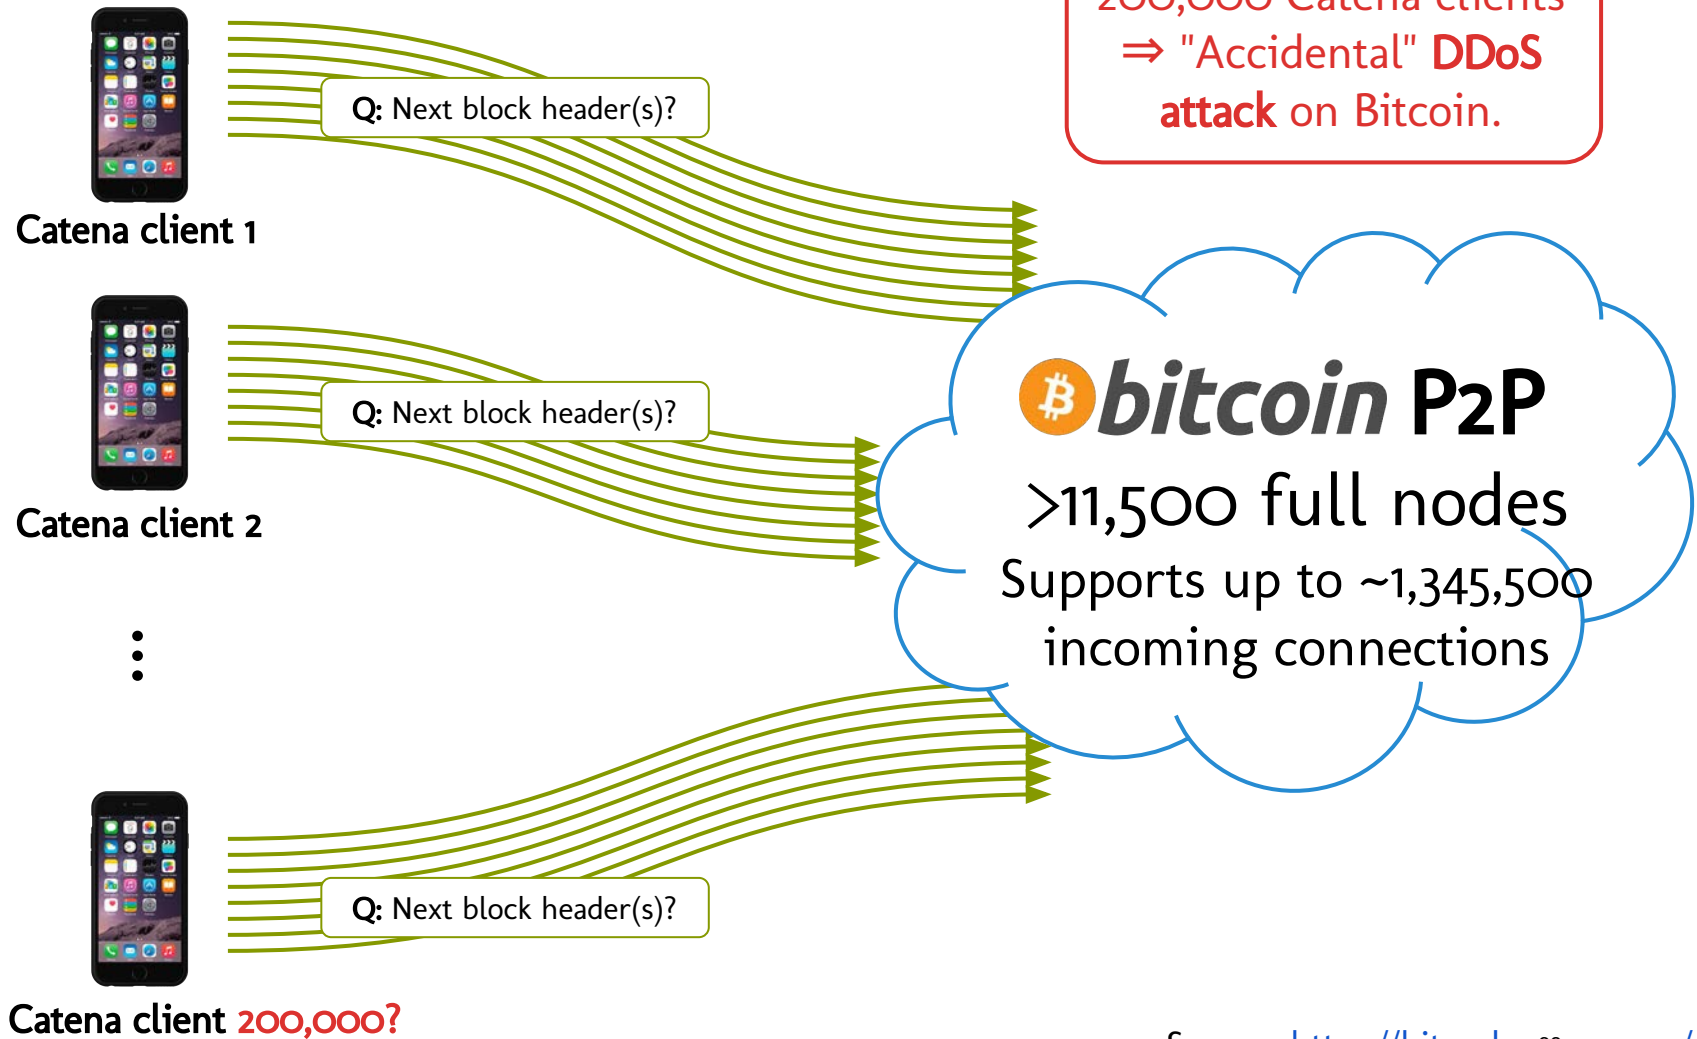
**Bitcoin P2P**  
>11,500 full nodes  
Supports up to ~1,345,500  
incoming connections

A blue-outlined cloud shape containing text about Bitcoin P2P network capacity.

# Catena scalability



# Catena scalability



# Catena scalability



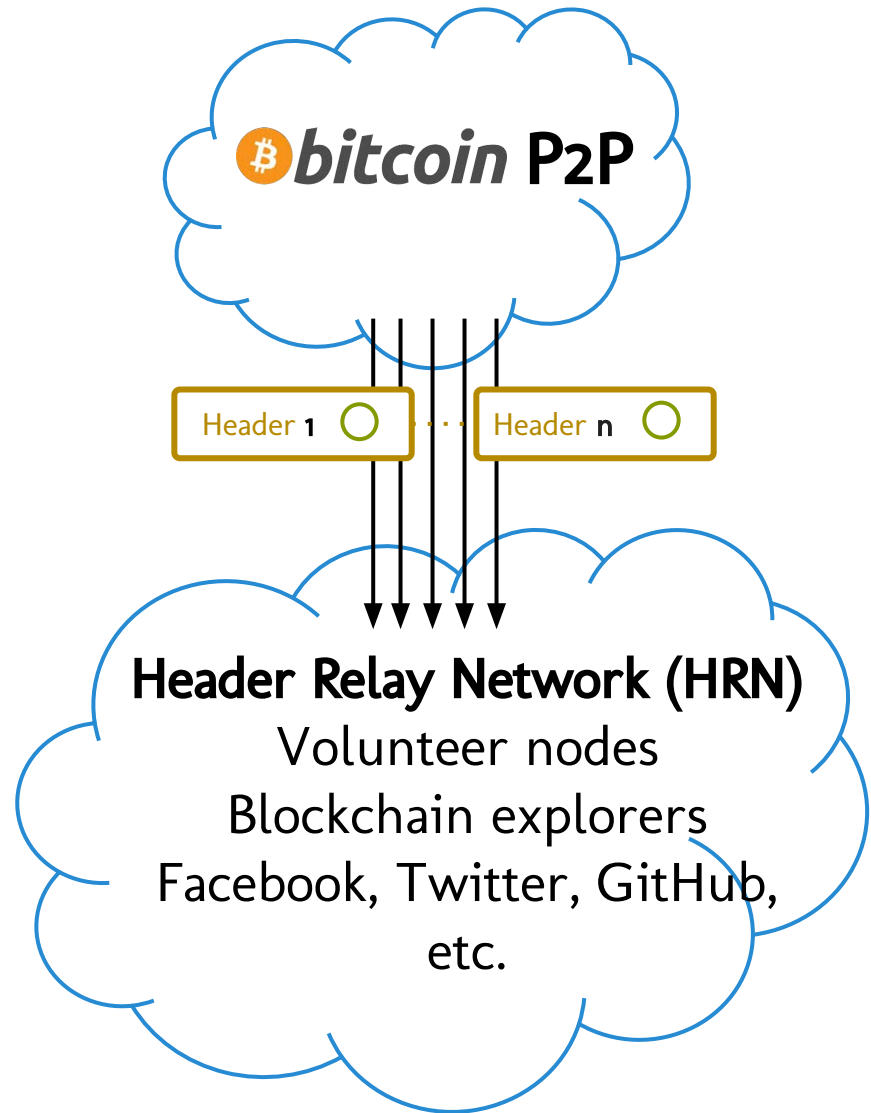
Catena client 1



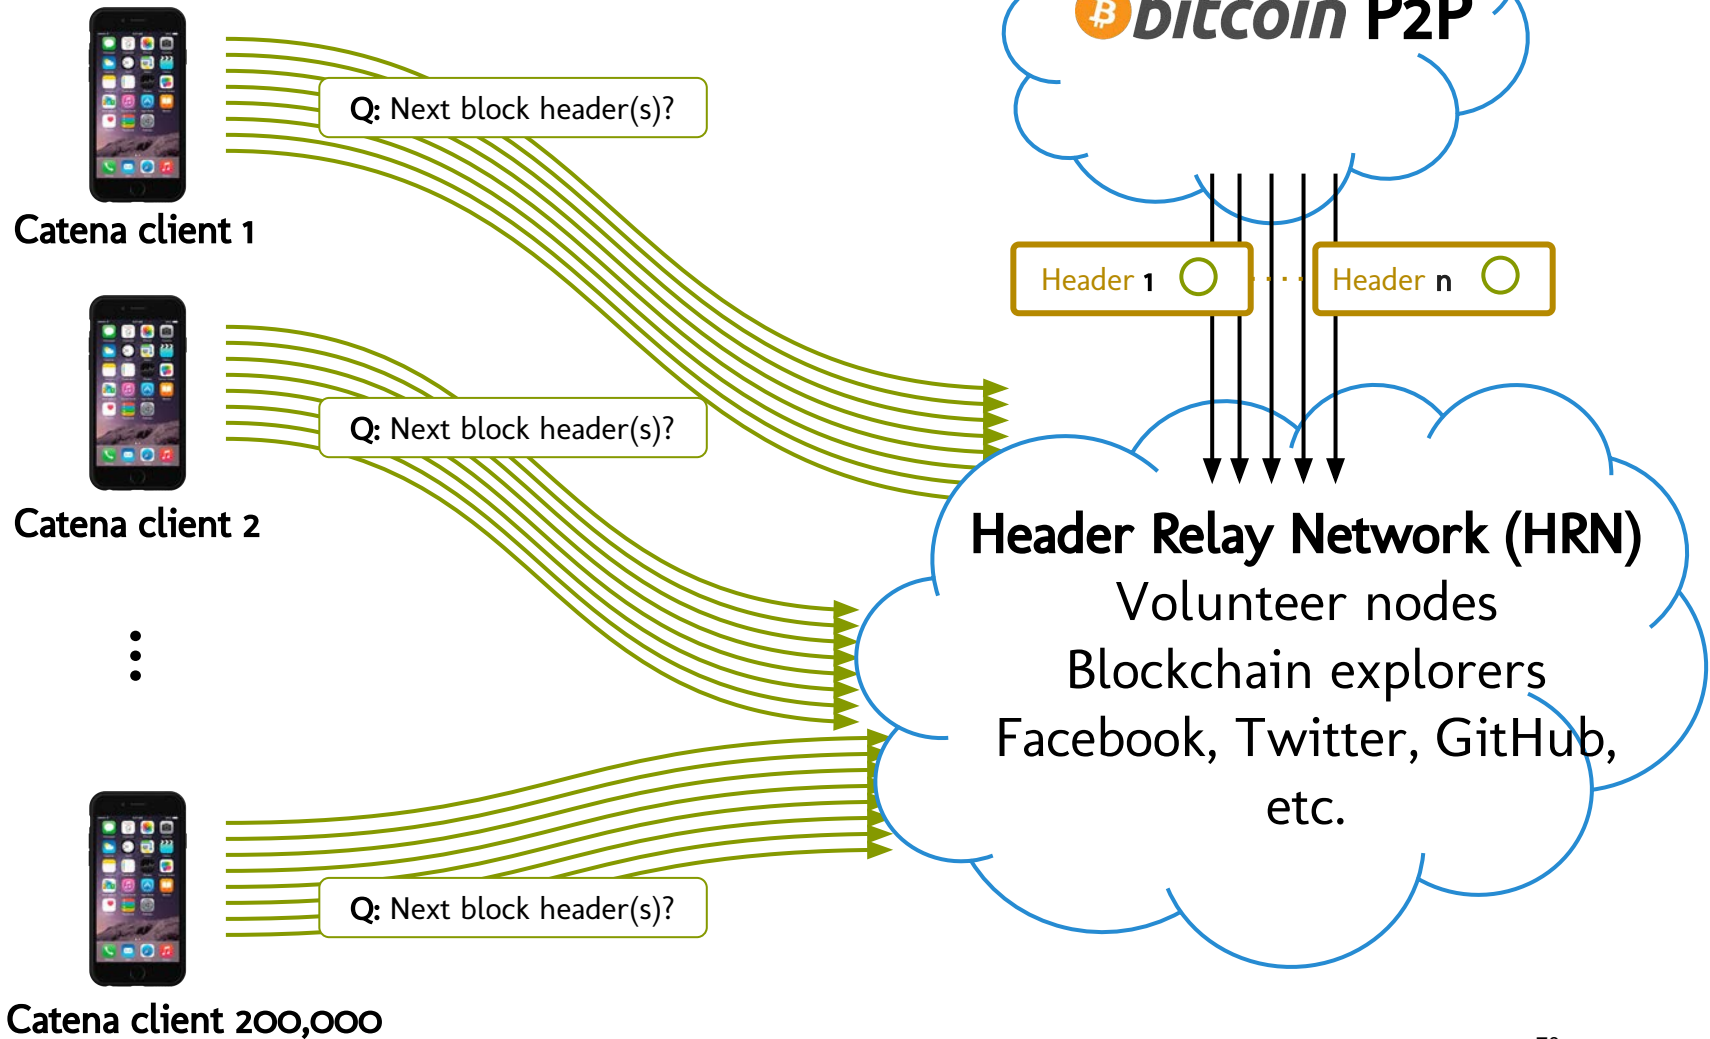
Catena client 2



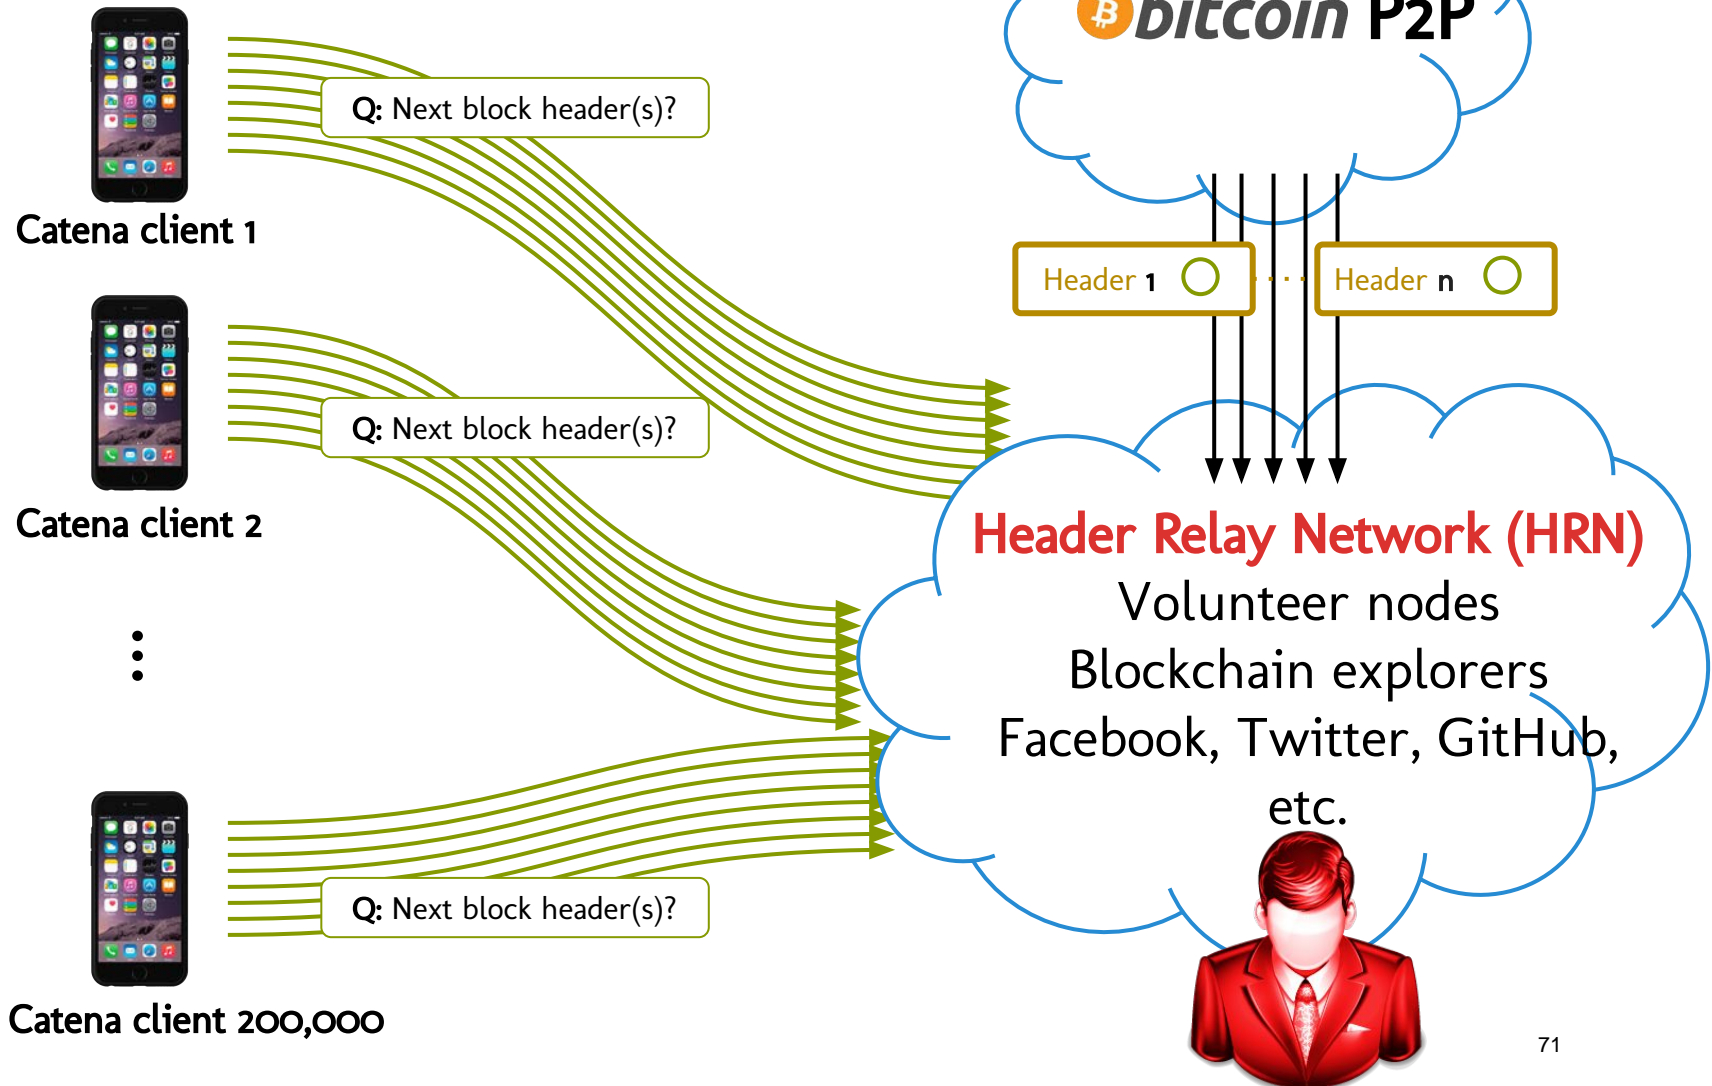
Catena client 200,000



# Catena scalability



# Catena scalability



# The cost of a statement

To append a statement, must issue TXN and pay fee.

**TXN size:** 235-byte

Fee as of Dec 13th, 2017: \$16.24 (10 mins)

Fee as of Feb 28th, 2017: \$0.78 (10 mins)

**PS:** Statements can be "batched" using Merkle trees.



# Overview

1. What?
2. How?
3. Why?
  - a. **Secure software update**

# Secure software update

## Example attack:

- (1) Compromise bitcoin.org (or the network)
- (2) Change the Bitcoin binary to your **malicious** binary
- (3) Wait for people to install your **malicious** Bitcoin binary
- (4) Steal their coins, steal their data, etc.

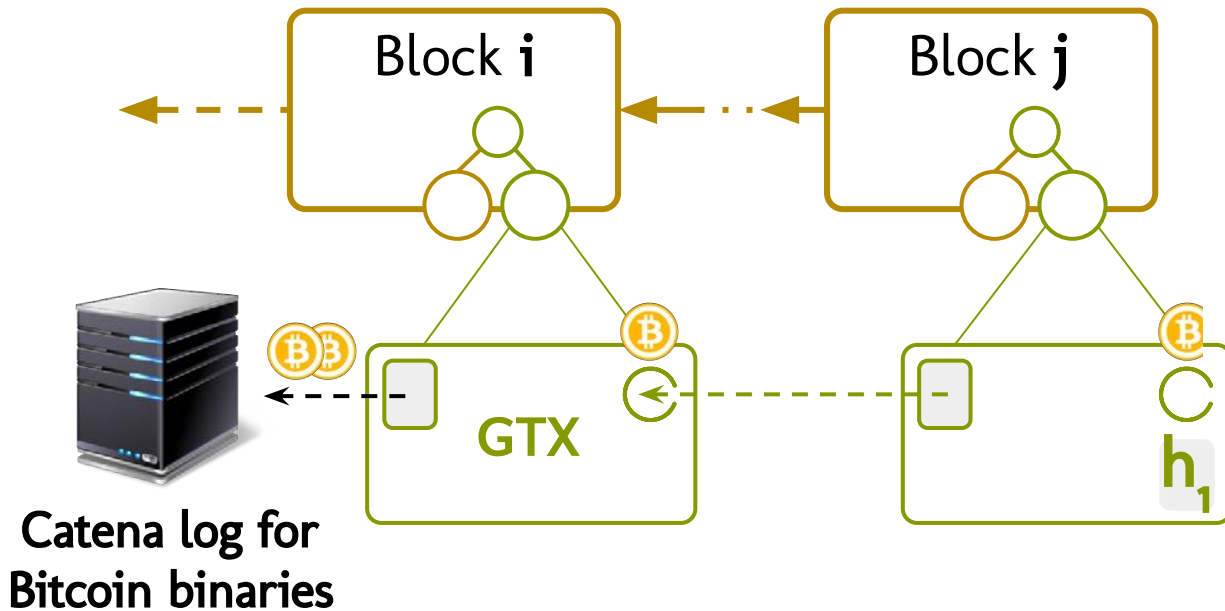
Example: [bitcoin.org, "0.13.0 Binary Safety Warning," August 17th, 2016](https://bitcoin.org/en/0.13.0-binary-safety-warning)

*Typical defense:* Devs sign Bitcoin binaries with **SK** and protect **SK**.

**Problem:** (1) Not everyone checks sig. (2) Hard to detect stolen **SK**.


**Solution:** Publish signatures in a Catena log  $\Rightarrow$  **Can at least detect.**

# Software transparency to the rescue



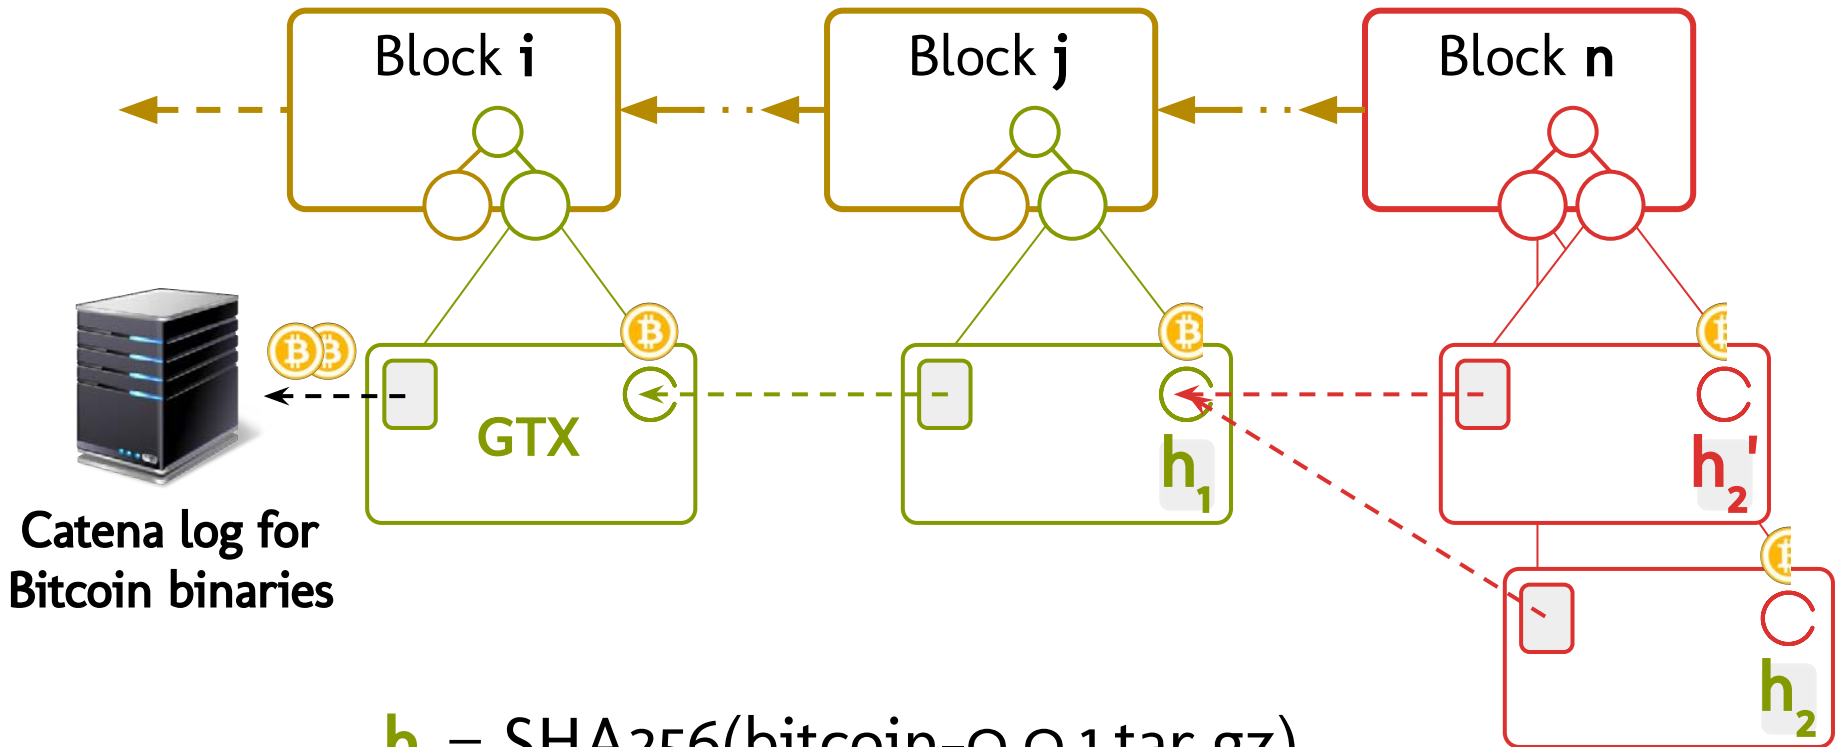
$$h_1 = \text{SHA256}(\text{bitcoin-0.0.1.tar.gz})$$

# Software transparency to the rescue



$$h_1 = \text{SHA256}(\text{bitcoin-0.0.1.tar.gz})$$
$$h_2' = \text{SHA256}(\text{evilcoin-0.0.2.tar.gz})$$

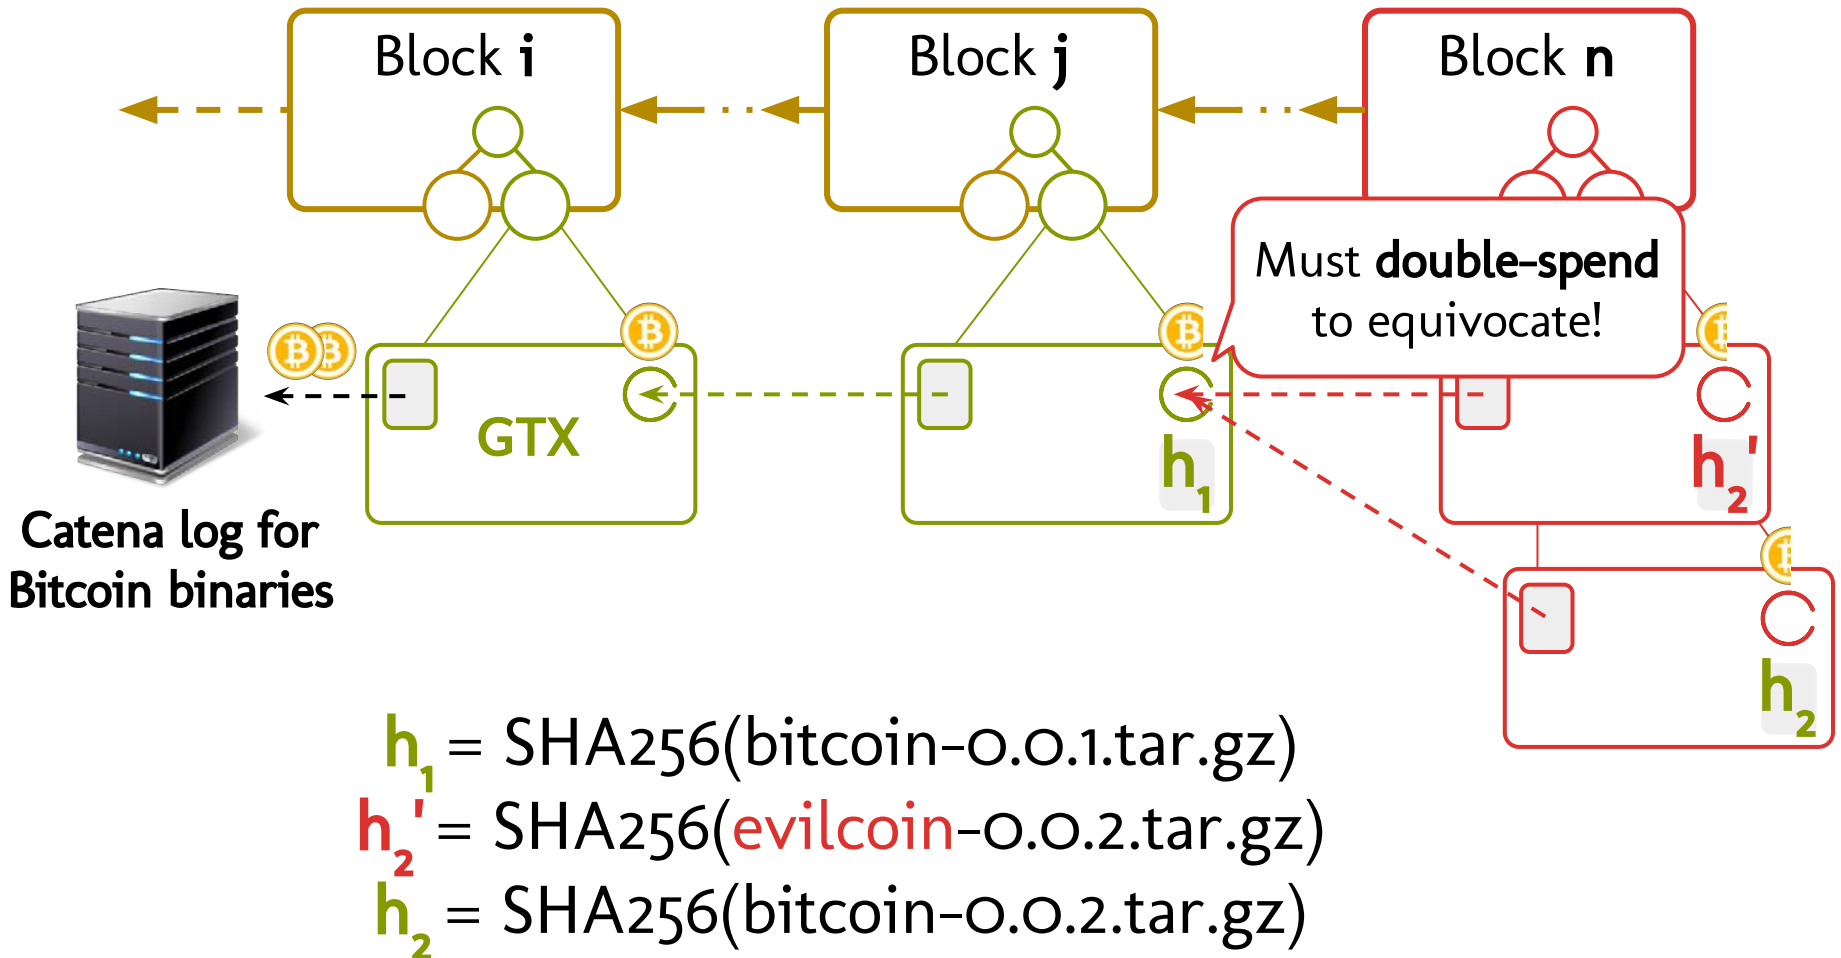
# Software transparency to the rescue



Catena log for  
Bitcoin binaries

$$\begin{aligned} h_1 &= \text{SHA256}(\text{bitcoin-0.0.1.tar.gz}) \\ h_2' &= \text{SHA256}(\text{evilcoin-0.0.2.tar.gz}) \\ h_2 &= \text{SHA256}(\text{bitcoin-0.0.2.tar.gz}) \end{aligned}$$

# Software transparency to the rescue



# Overview

1. What?
2. How?
3. Why?
  - a. Secure software update
  - b. Secure messaging**

# Public-key distribution



$A = \{Alice, PK_A\}, SK_A$

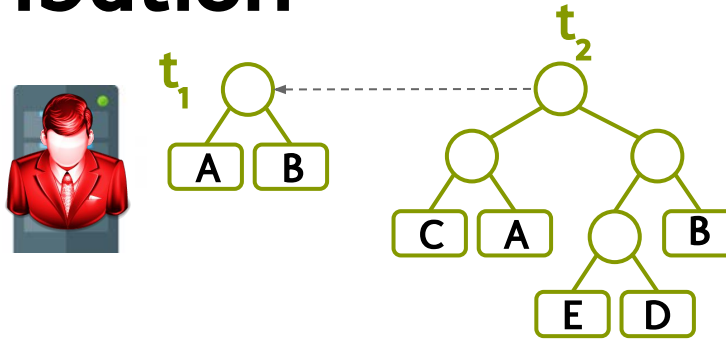


$B = \{Bob, PK_B\}, SK_B$






# Public-key distribution




$A = \{Alice, PK_A\}, SK_A$



$B = \{Bob, PK_B\}, SK_B$



# Public-key distribution



$$A = \{Alice, PK_A\}, SK_A$$




$$A' = \{Alice, PK_M\}, SK_M$$


$$B' = \{Bob, PK_M\}, SK_M$$



$$B = \{Bob, PK_B\}, SK_B$$



# Public-key distribution



I've been impersonated!

$$A = \{Alice, PK_A\}, SK_A$$




$$A' = \{Alice, PK_M\}, SK_M$$

$$B' = \{Bob, PK_M\}, SK_M$$




I've been impersonated!

$$B = \{Bob, PK_B\}, SK_B$$



# Public-key distribution



I've been impersonated!

$$A = \{Alice, PK_A\}, SK_A$$



$$A' = \{Alice, PK_M\}, SK_M$$

$$B' = \{Bob, PK_M\}, SK_M$$

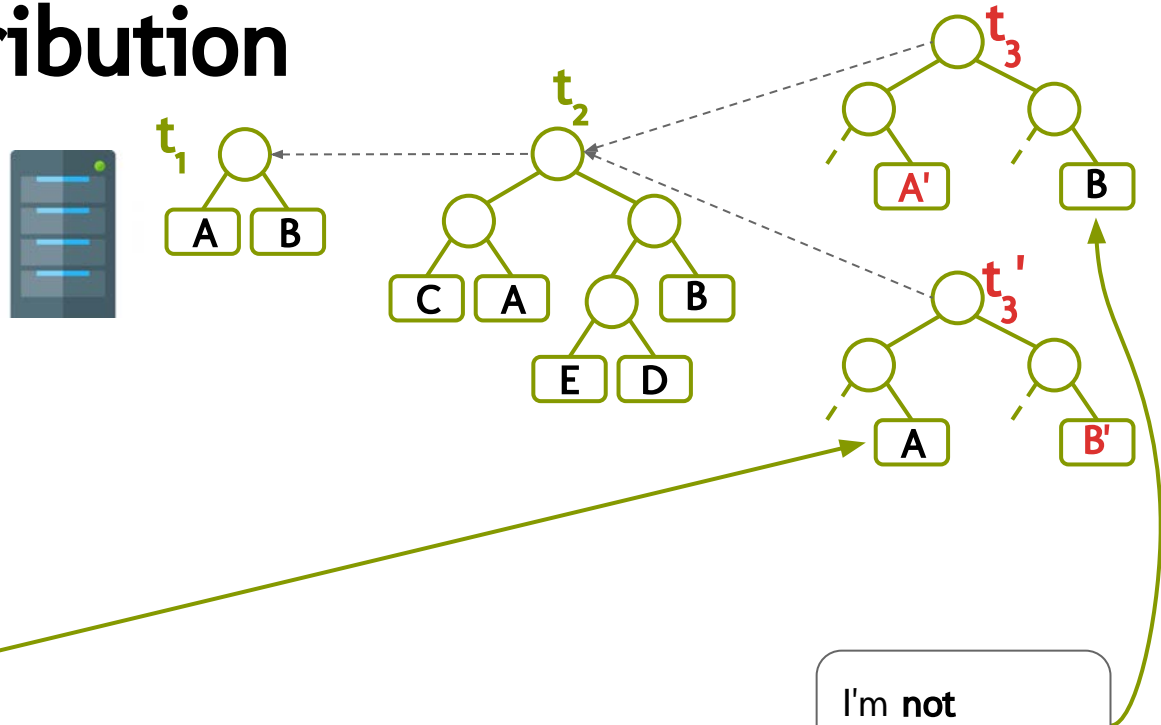



I've been impersonated!

$$B = \{Bob, PK_B\}, SK_B$$



# Public-key distribution



I'm not impersonated.

$$A = \{Alice, PK_A\}, SK_A$$




$$A' = \{Alice, PK_M\}, SK_M$$

$$B' = \{Bob, PK_M\}, SK_M$$




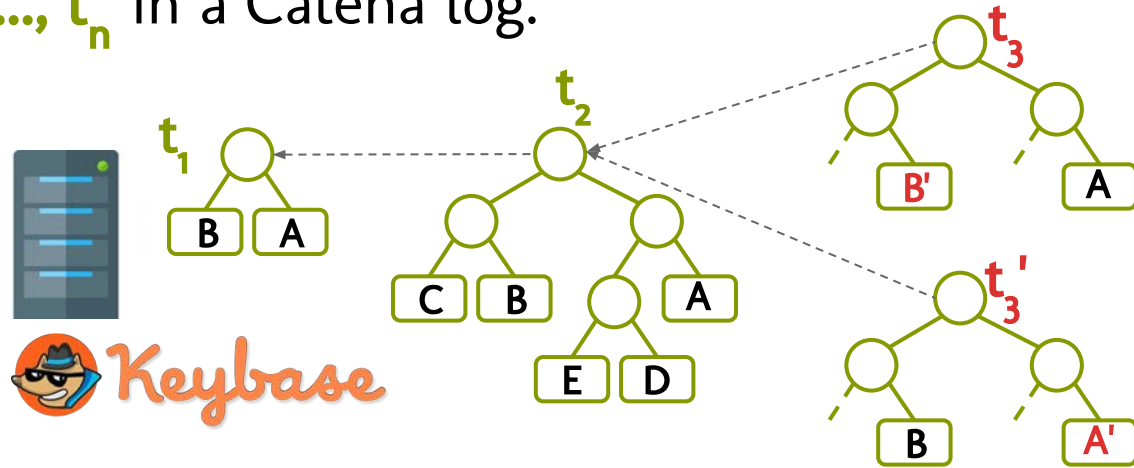
I'm not impersonated.

$$B = \{Bob, PK_B\}, SK_B$$



# KeyChat

Idea: Store  $t_1, t_2, \dots, t_n$  in a Catena log.



# Overview

1. What?
2. How?
3. Why?
  - a. Secure software update
  - b. Secure messaging
  - c. **"Blockchain" for X**

# I need a "blockchain" for ...

IoT? Self-driving cars? Digital identity? Issue diplomas? Health care? Voting?

**"Blockchain"** = Byzantine State Machine Replication (SMR) =  
= Agree on log of ops + Execute ops = Agree on final state.

Permissioned "blockchain" via:

- Your favorite Byzantine SMR algorithm
- Ethereum Smart Contract (pay fees per Ethereum op)
- **Catena +  $2f+1$  replicas (pay fees per op batch)**



# I need a "blockchain" for ...

IoT? Self-driving cars? Digital identity? Issue diplomas? Health care? Voting?

**"Blockchain"** = Byzantine State Machine Replication (SMR) =  
= Agree on log of ops + Execute ops = Agree on final state.

Permissioned "blockchain" via:

- Your favorite Byzantine SMR algorithm
- Ethereum Smart Contract (pay fees per Ethereum op)
- **Catena +  $2f+1$  replicas (pay fees per op batch)**
- Don't need execution? Use Catena directly.

Permissionless "blockchain:" Roll your own. But proceed with caution?

# Conclusions

## *What we did:*

- Enabled applications to efficiently leverage Bitcoin's publicly-verifiable consensus
  - Download transactions selectively rather than full blockchain
  - ~46 MB instead of gigabytes of bandwidth

## *Why it matters:*


- Secure software update schemes
- Public-key directories for HTTPS and secure messaging
- "Blockchain" for X

For more, read [our paper!](#)

# Ask me questions!

<https://people.csail.mit.edu/alinush/>

Previous work



Catena

