

# Blockchains & Distributed Ledgers

Lecture 08

Michele Ciampi



Slide credits: MC, Dimitris Karakostas, Aggelos Kiayias

# Eponymous system

- Each action can be **attributed** to a user's **real-world identity**
- Examples:
  - Facebook - posts/comments are linked with the real-world name of the user who made it
  - Twitter blue check (pre-Musk) - accounts are verified w.r.t. real-world identification documents
  - UK parliament votes - the vote of each MP is (publicly) attributable to each

# Pseudonymous system

- Identities are represented as **tags**
- Each tag is independently assigned to each identity
- An **identity** may be **assigned multiple tags** and vice versa
- Examples:
  - Twitter/Reddit - posts/comments are linked to an (arbitrary) username
  - Email - each message is linked with an email address
  - Graffiti - each piece is signed by a tag/pseudonym (e.g., Banksy)

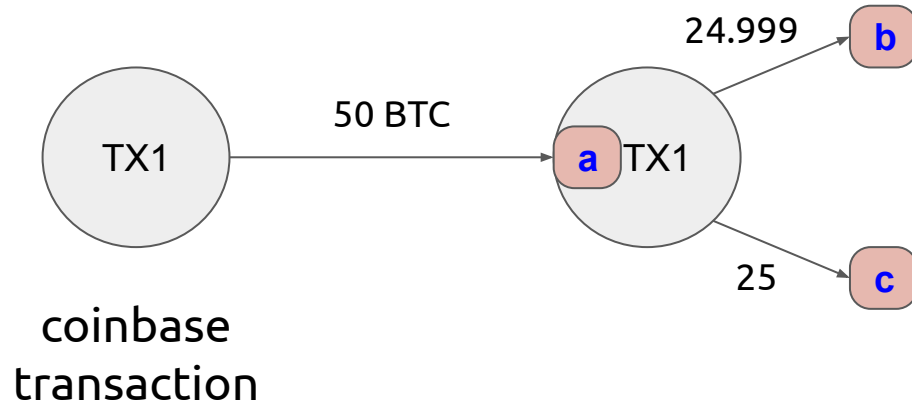
# Anonymous system

- Any performed action is manifested within a set of **indistinguishably-acting participants**
- The set of indistinguishable participants is called **the anonymity set**
  - Hide in public
- Examples:
  - General election voting - e.g., ~14M of 47.6M eligible voters voted Conservatives in 2019
  - Tor browsing - website/hidden service sees only number of Tor connections (not name/IP)

# Privacy in Bitcoin

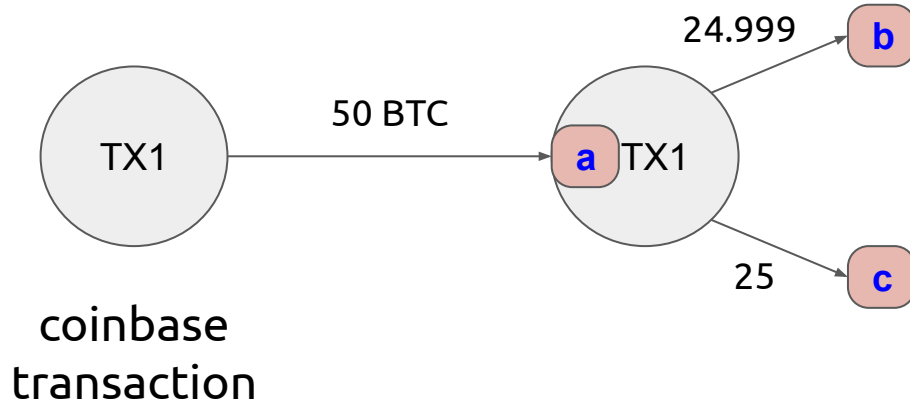
- Users can create multiple accounts/addresses:
  - without cost
  - without association to previous accounts
- Essentially, users can create an **unlimited number of pseudonyms**

# Transaction Graph Analysis




account **a** moves  
50 BTC to  
accounts **b** and **c**  
(minus fees)


# Transaction Graph Analysis




account **a** moves  
50 BTC to  
accounts **b** and **c**  
(minus fees)



# Common Behaviours



peeling chain



star

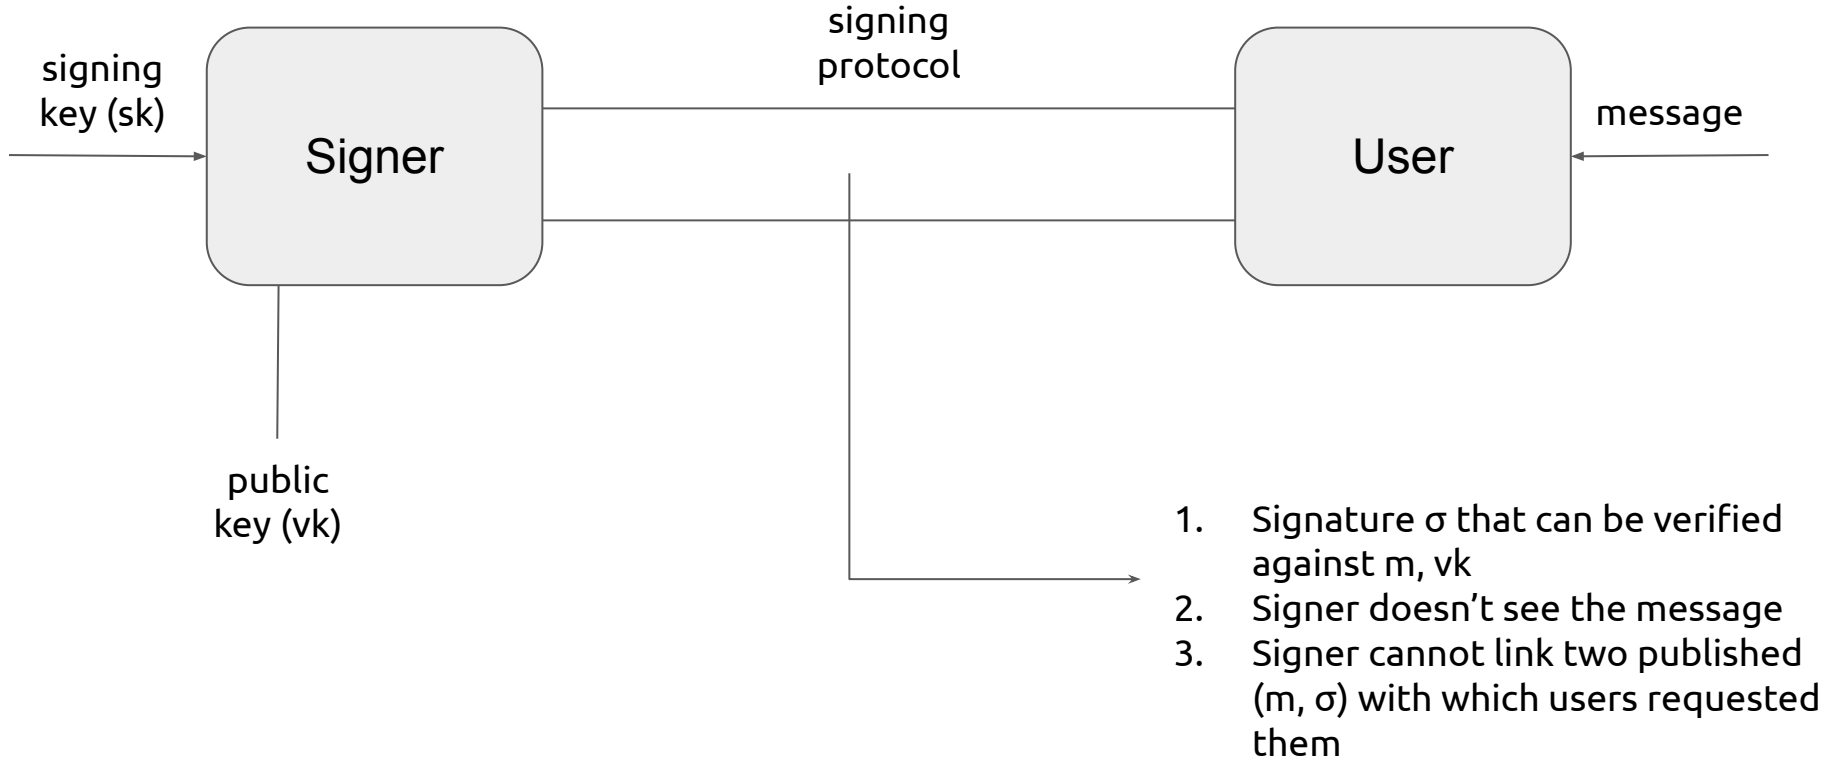


# Fungibility and Privacy

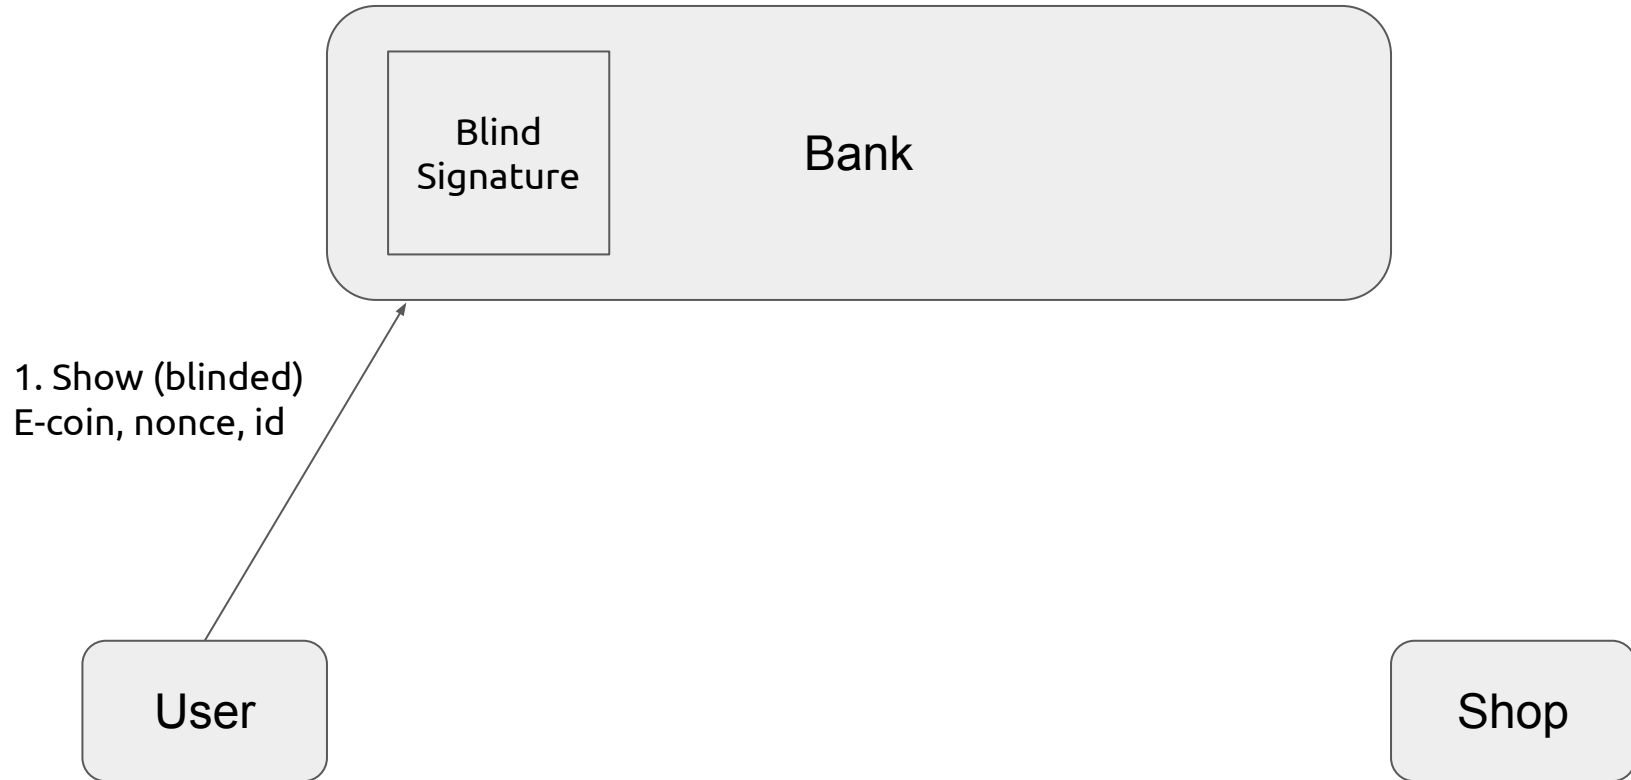
- **Fungibility:** Coins are interchangeable
- However, each “satoshi” has its whole history in the Bitcoin blockchain
  - satoshi fungibility is debatable

# Transaction Anonymization Techniques

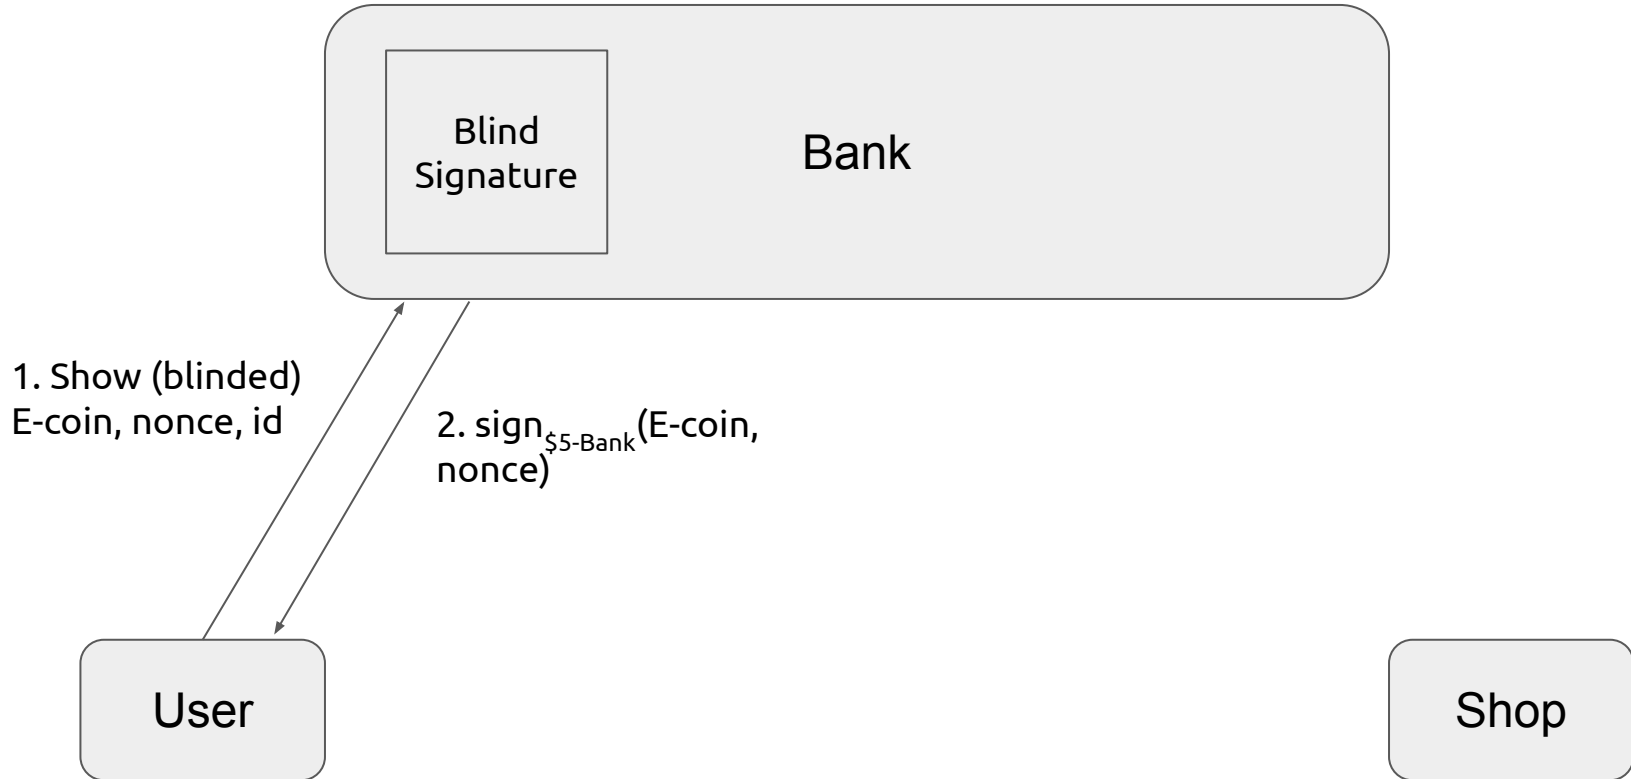
# Blind-Signatures



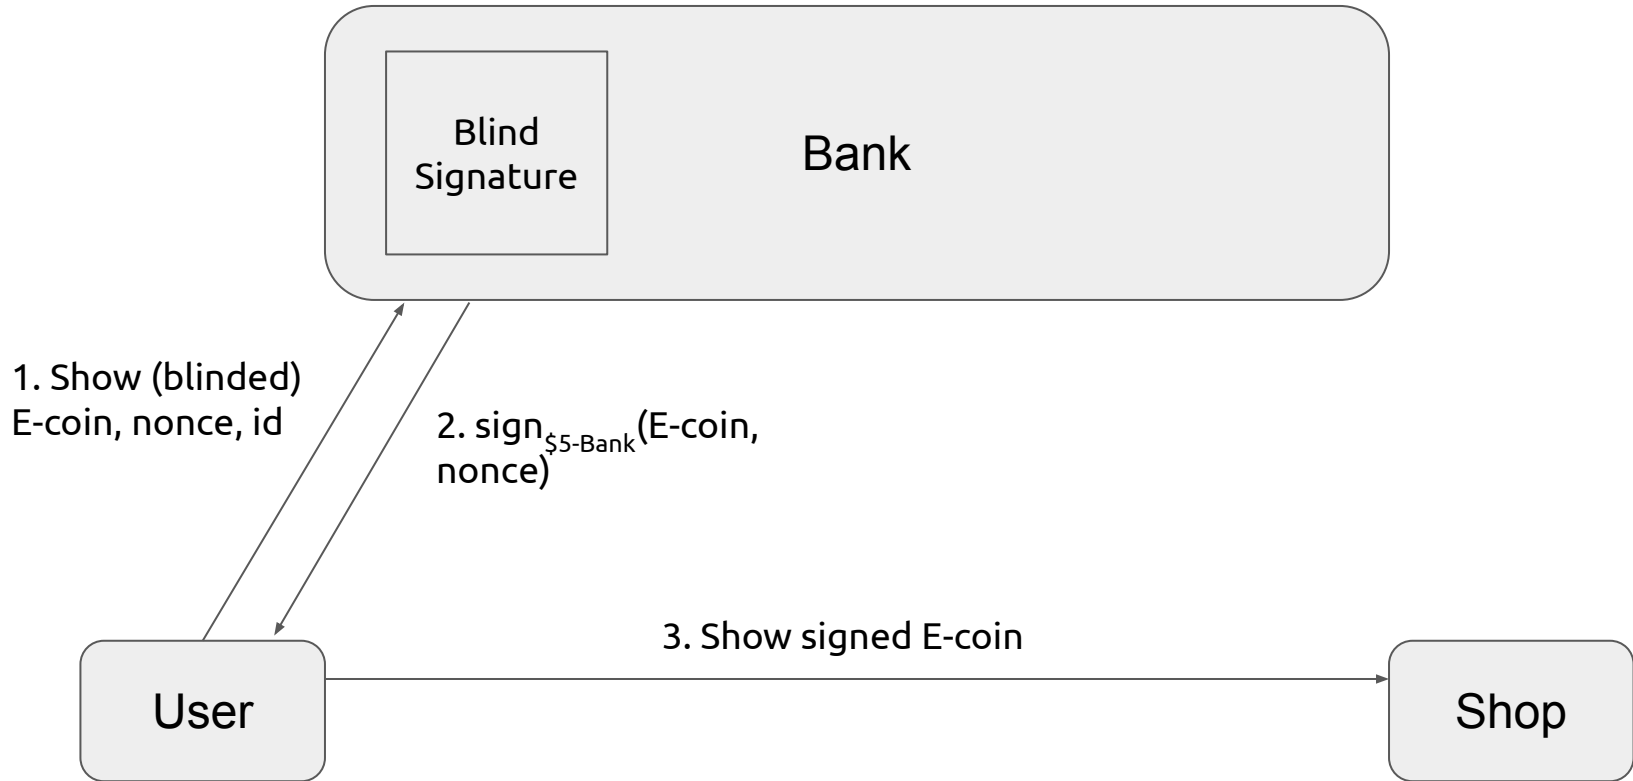
# Chaum's E-cash



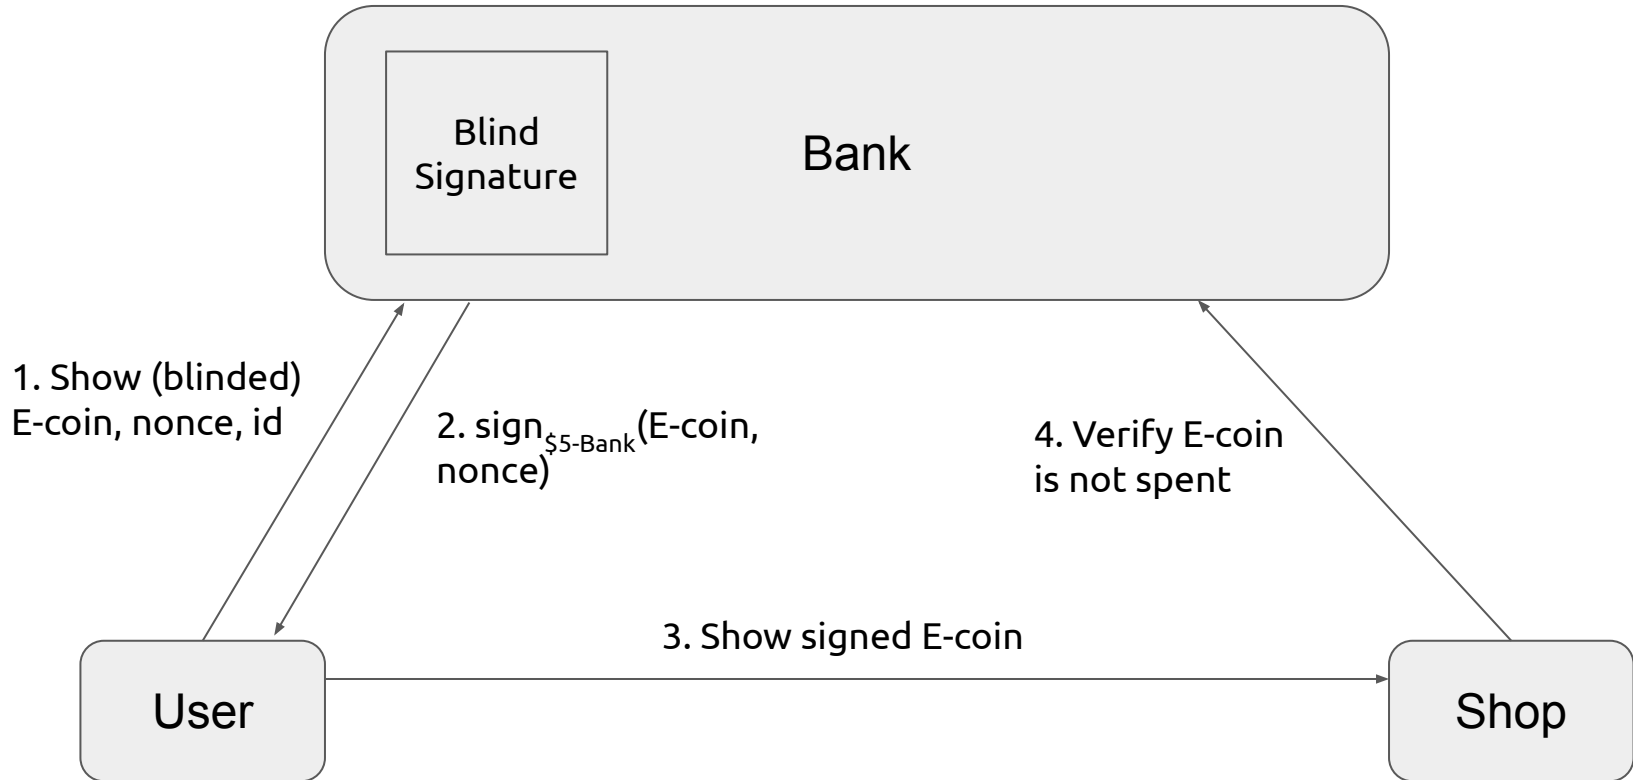
# Chaum's E-cash



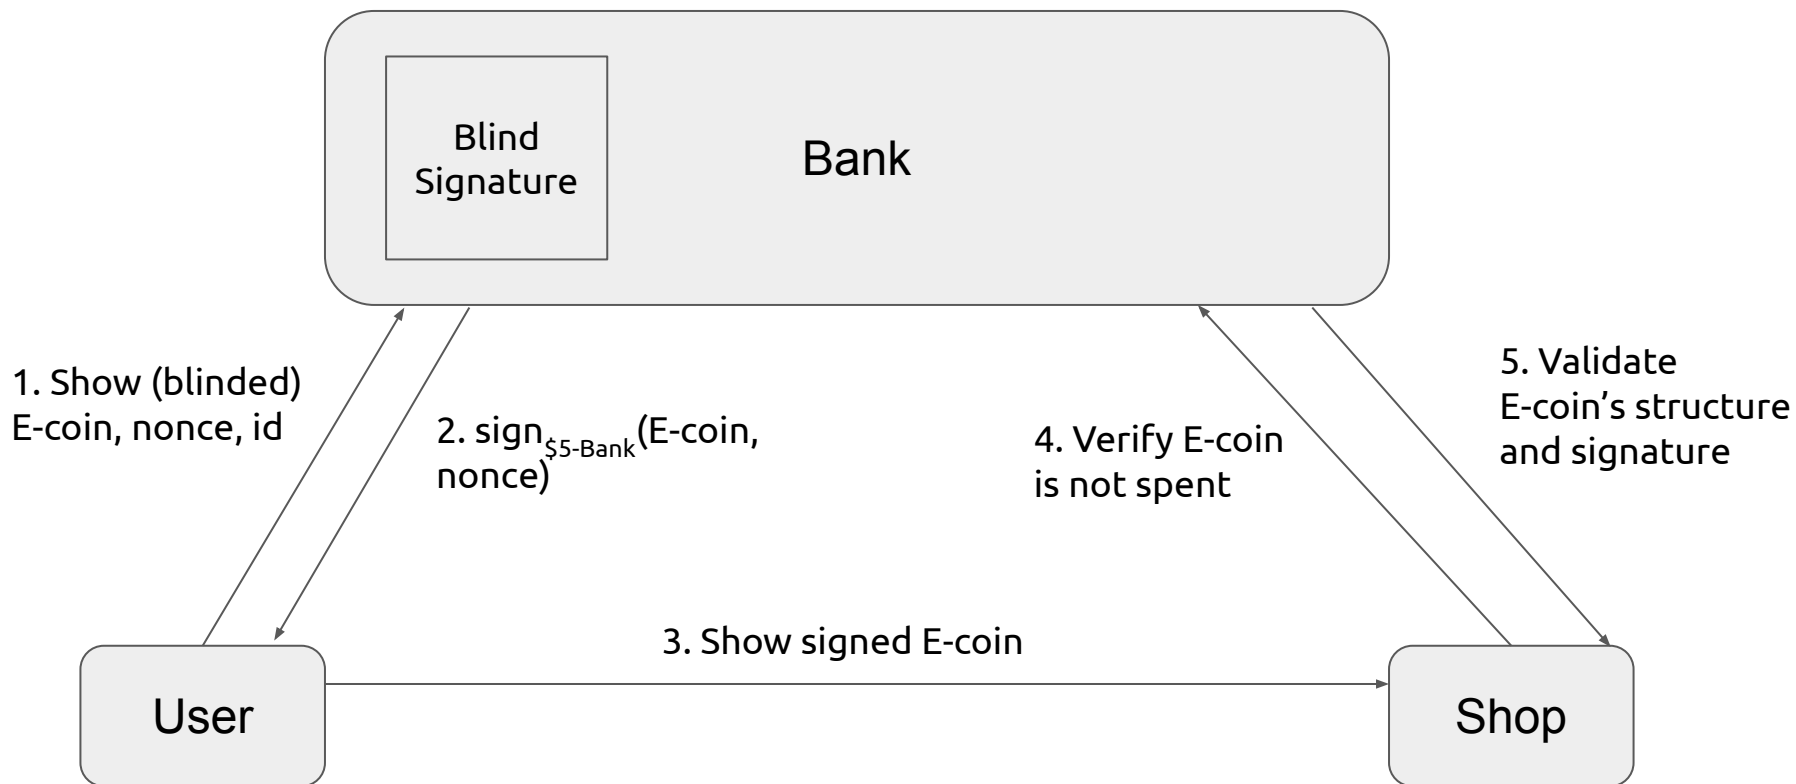
# Chaum's E-cash



# Chaum's E-cash

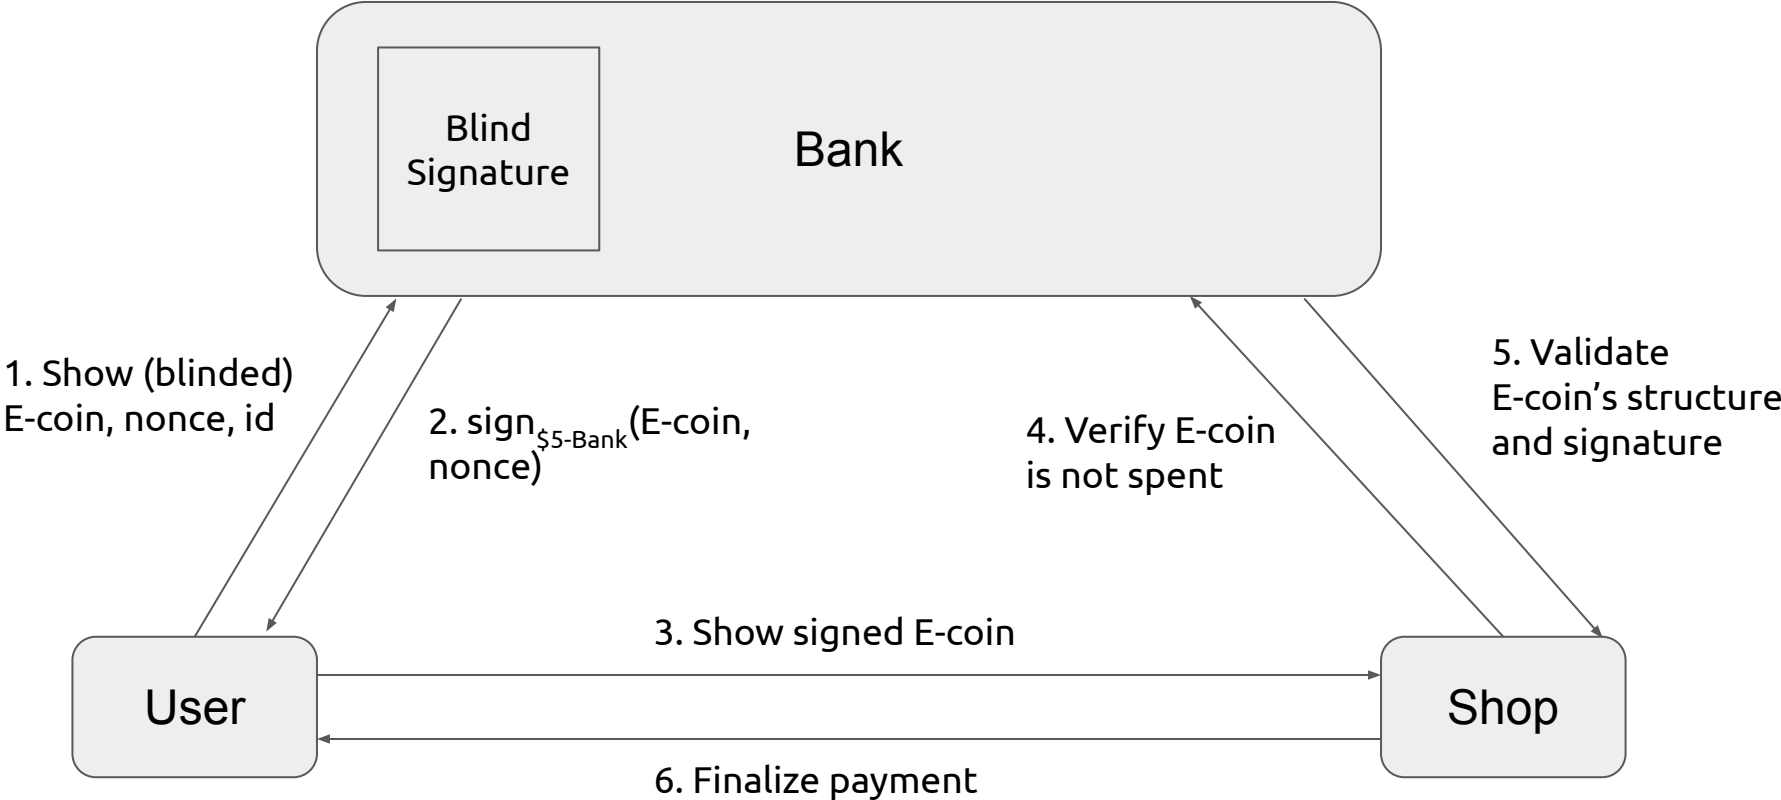


# Chaum's E-cash

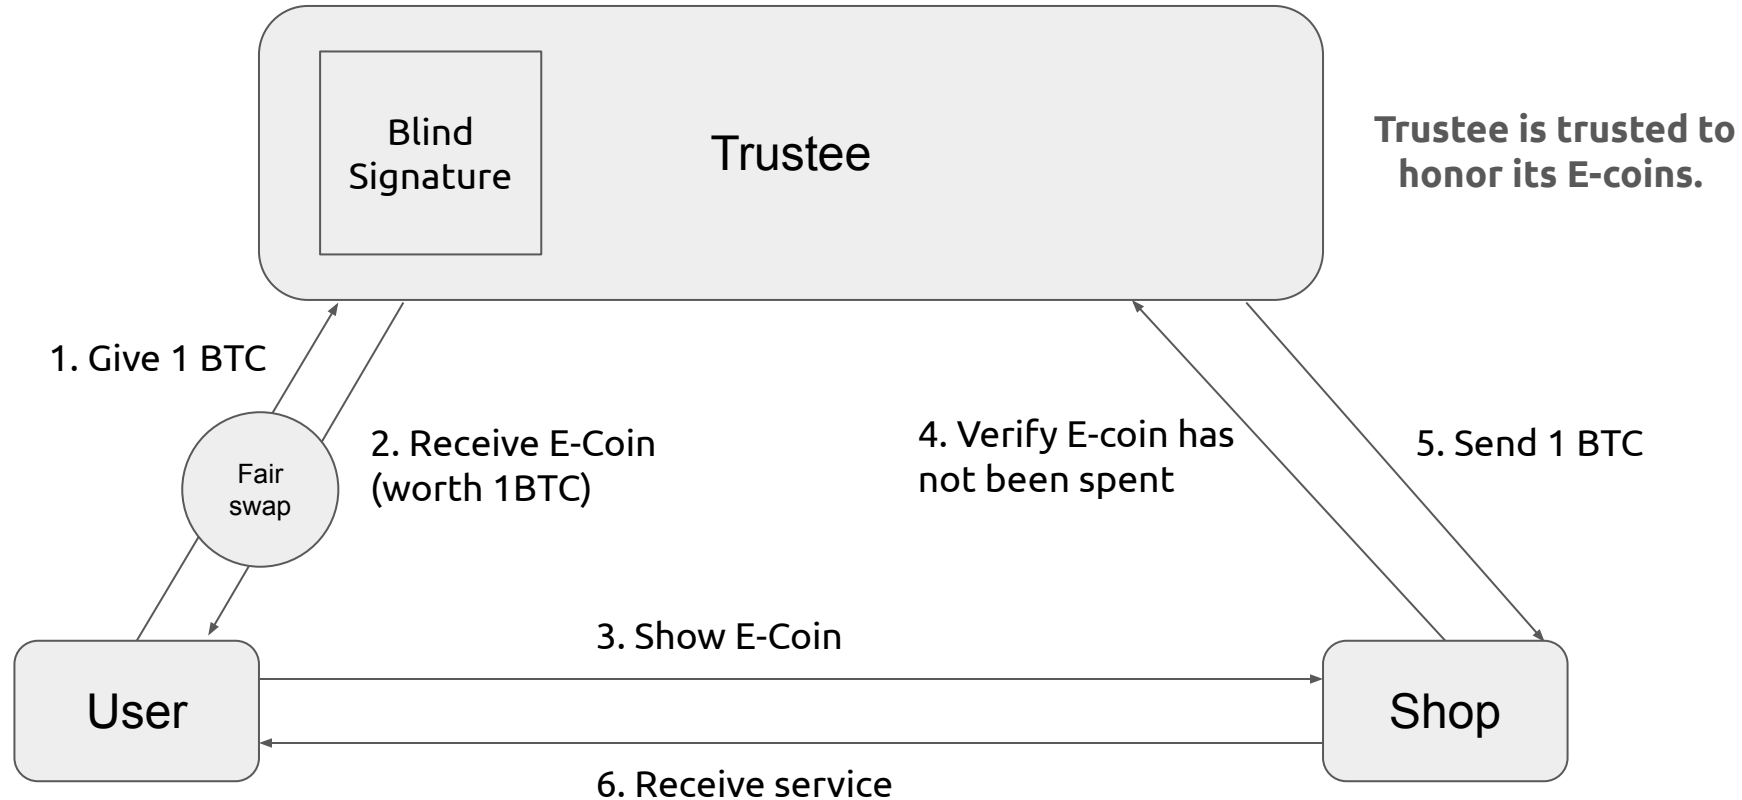




# Chaum's E-cash



# Anonymizing Bitcoin Payments via E-cash



# Fair Swaps


- Alice and Bob would like to exchange secrets s.t.:
  - either none of them gets their output
  - or both do
- Classical problem
- Impossible to solve under standard network assumptions!
- Going around the impossibility:
  - optimistic fair exchange
  - resource-based fair exchange
  - fair swaps with penalties

# Fair Swaps - Construction

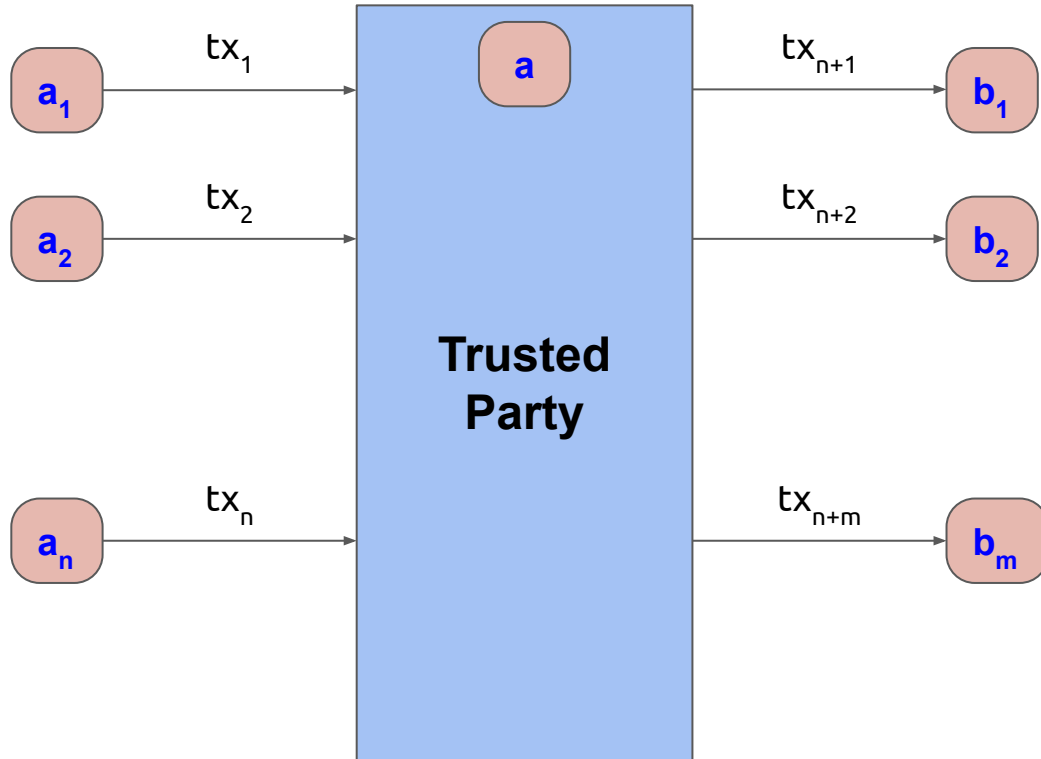
- Using a blockchain that supports **smart contracts**
- A contract that both parties fund to accept their secrets
- The parties are **rational**
  - The security argument will be **game theoretic**
- Key requirements:
  - parties lock up some funds in **deposits**
  - secret submission should be **verifiable** by the contract's code
- Fair swap variation:
  - Either both parties get their output
  - Or the offending party is **penalized financially**

Coinjoin

# Anonymising Transactions - Centralized



# Anonymising Transactions - Centralized




Anonymity set of size  $n$

TP may disappear with the money

# Anonymizing Transactions - CoinJoin


## Without Coinjoin






# Anonymizing Transactions - CoinJoin

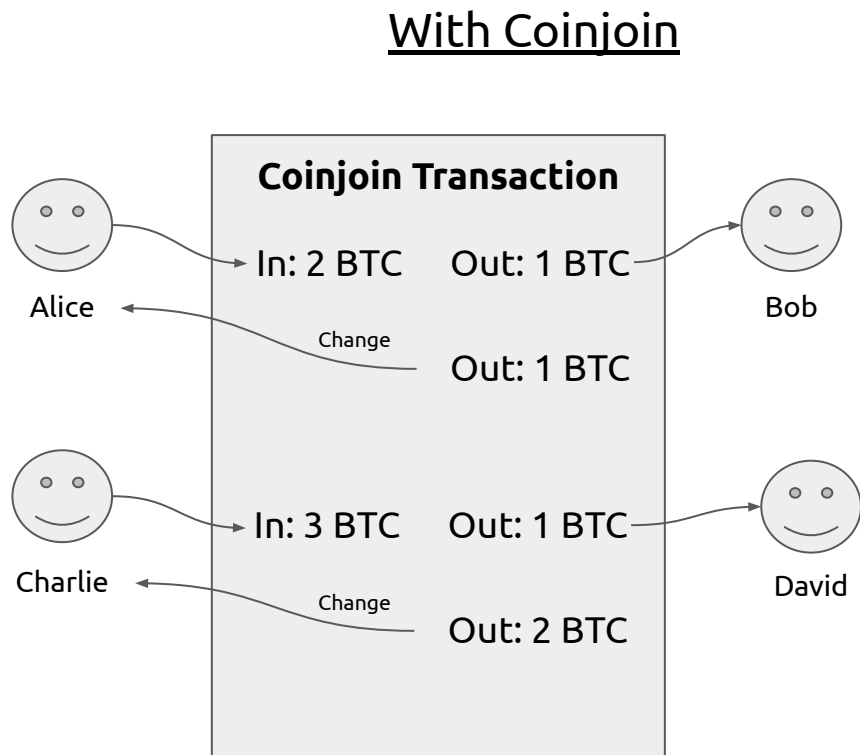
## Without Coinjoin



## With Coinjoin




# Anonymizing Transactions - CoinJoin




# Anonymizing Transactions - CoinJoin

## Without Coinjoin



## With Coinjoin



# Multiple Input Transactions - Setup

- Parties:
  - $n$  participants
  - one designated leader
- The  $i$ -th party sends to the leader:
  - the recipient address  $b_i$
  - the return (change) address  $c_i$
  - the corresponding amounts
- When all  $n$  parties complete this step, the multiple input transaction is formed by the leader and sent to all  $n$  parties

# Multiple Input Transactions - Sign and Publish

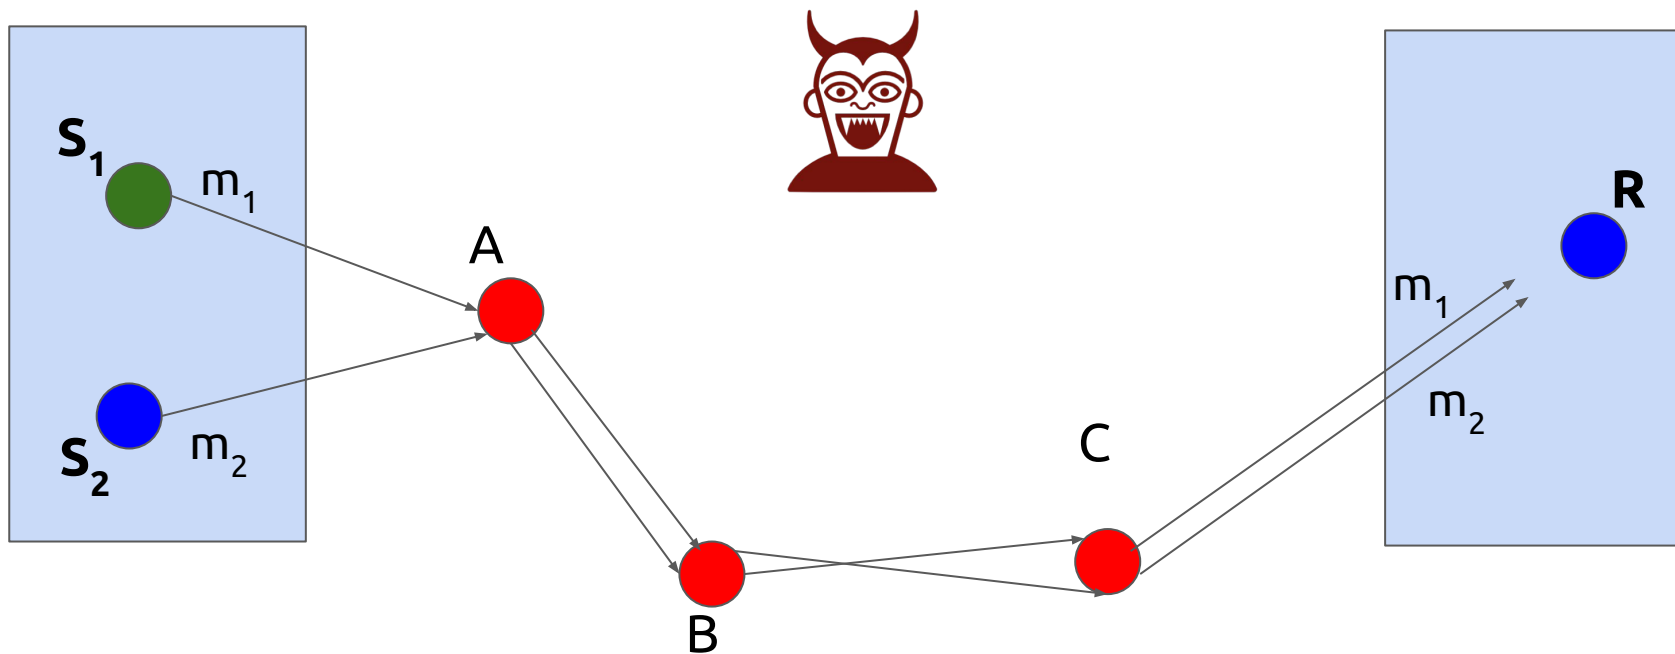
- **Each party** sends a **signature** on the multiple input tx to the leader
- When all  $n$  signatures are received, the multiple input tx is **posted** on the blockchain **by the leader**
- If any of the  $n$  parties **aborts** the protocol, the transaction cannot be validated
- If the **leader is adversarial**, transaction cannot be published/validated

Question: Can we ensure that an adversary does not correlate the IP address of the sender and the receiver?

# Mix-net

- A mix-net facilitates hiding the sender and the receiver of a given message
- Decryption mix-nets
- Re-encryption mix-nets

# Mix-net: simplified scenario (hiding only the sender)



Not possible to relate if  $S_1$  sent  $m_1$  or  $m_2$  (and vice versa for  $S_2$ ) - as long as there is one honest node (even if the adversary can look at what all the nodes receive and output).

# Routing via a Mix-net

Sample 3 encryption keys  
Sym\_key1, Sym\_key2, Sym\_key3

sender



A ( $PK_A$ )



B ( $PK_B$ )



C ( $PK_C$ )



receiver





# Decryption Mix-net

Sample 3 encryption keys  
Sym\_key1, Sym\_key2, Sym\_key3

sender



A ( $PK_A$ )



B ( $PK_B$ )



C ( $PK_C$ )



receiver



Encrypted with sym\_key3

Payload  
destination / info

fixed  
block  
size 2

fixed  
block  
size 1

Encrypted with Public key of C  
*Deliver to R; sym\_key3*



# Decryption Mix-net

Sample 3 encryption keys  
Sym\_key1, Sym\_key2, Sym\_key3

sender



A (PK<sub>A</sub>)



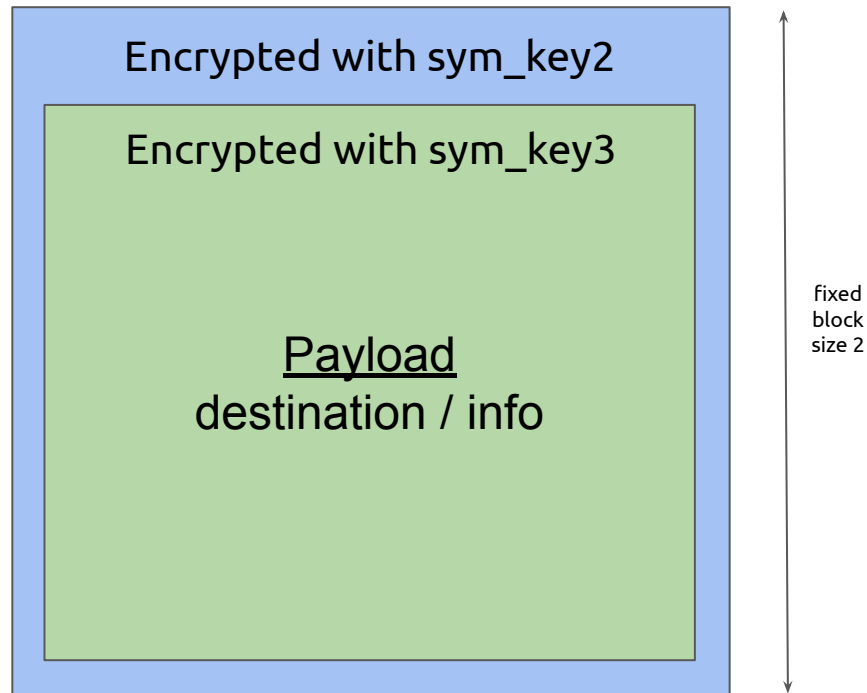
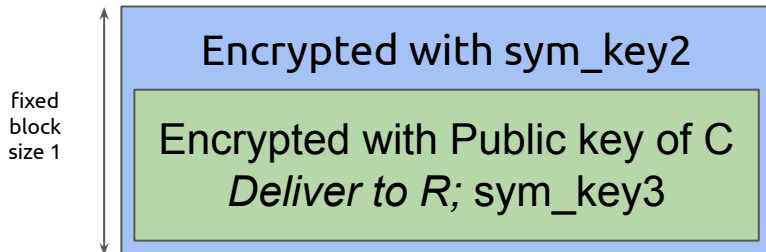
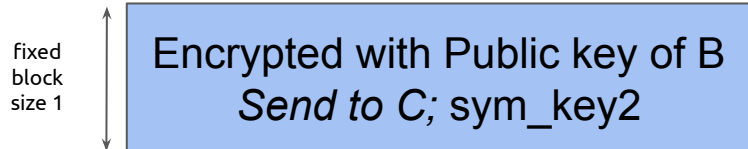
B (PK<sub>B</sub>)



C (PK<sub>C</sub>)

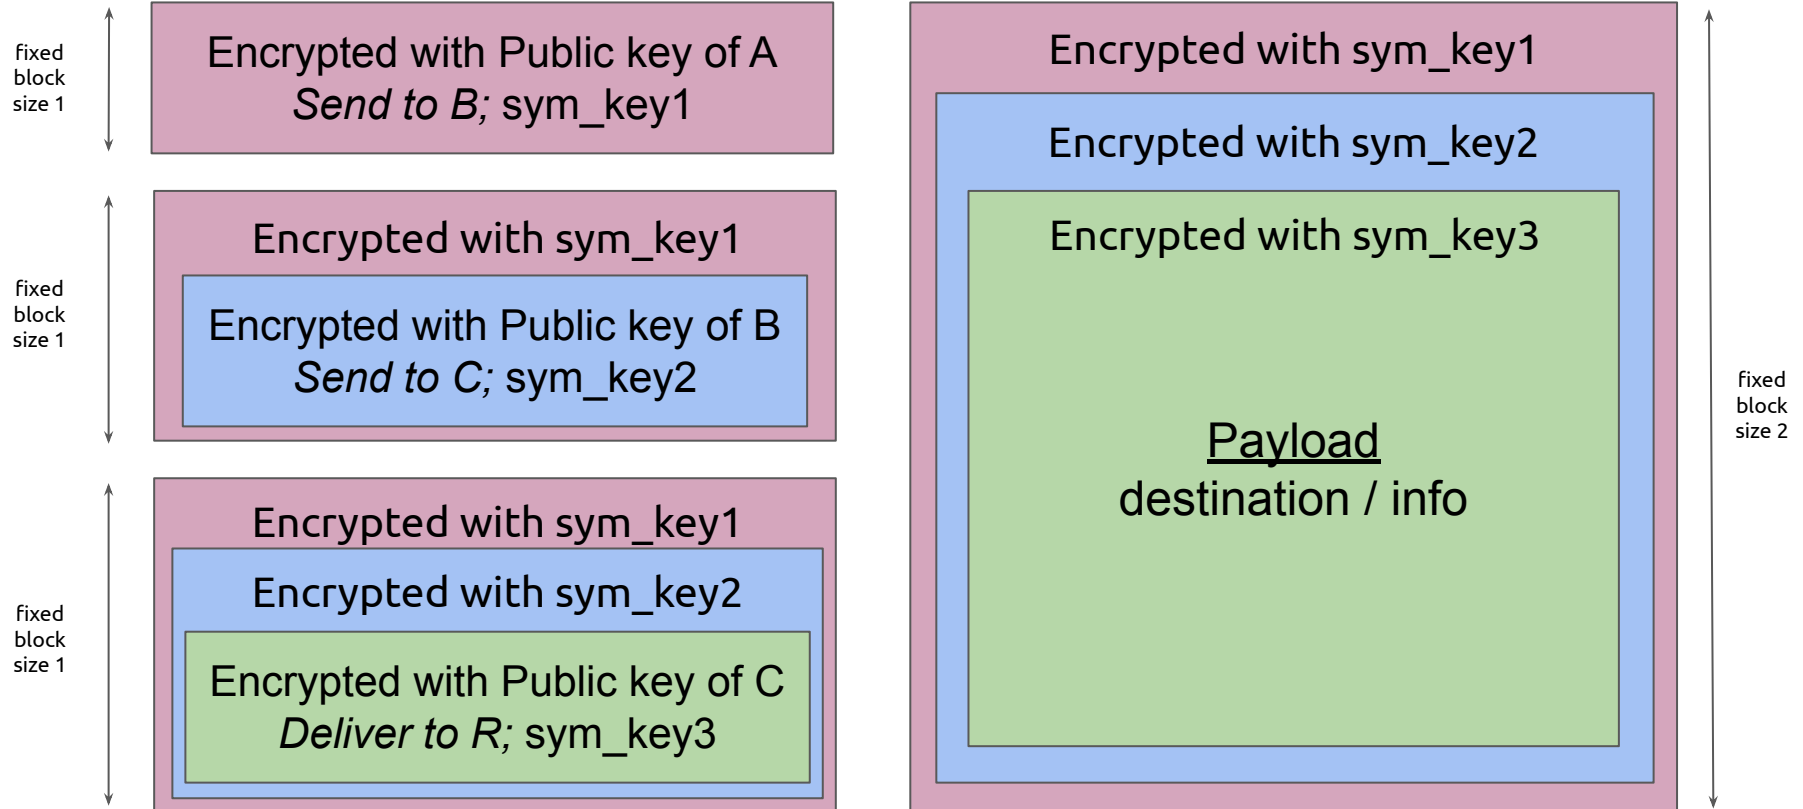


receiver




Sample 3 encryption keys  
Sym\_key1, Sym\_key2, Sym\_key3

# Decryption Mix-net

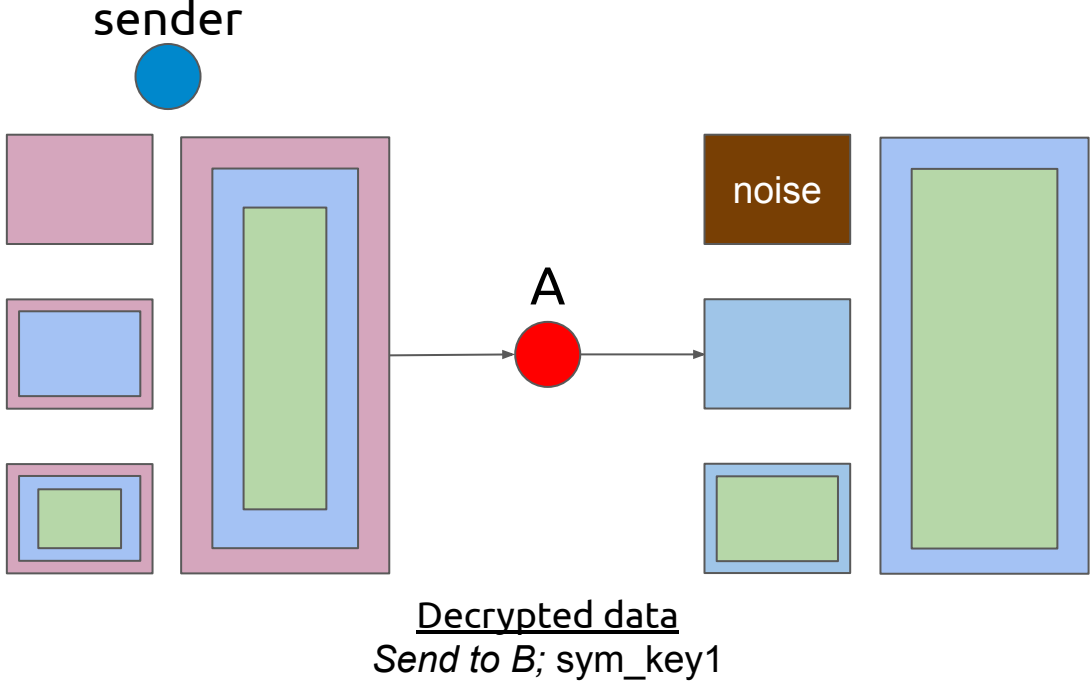


# Routing via a Mix-net

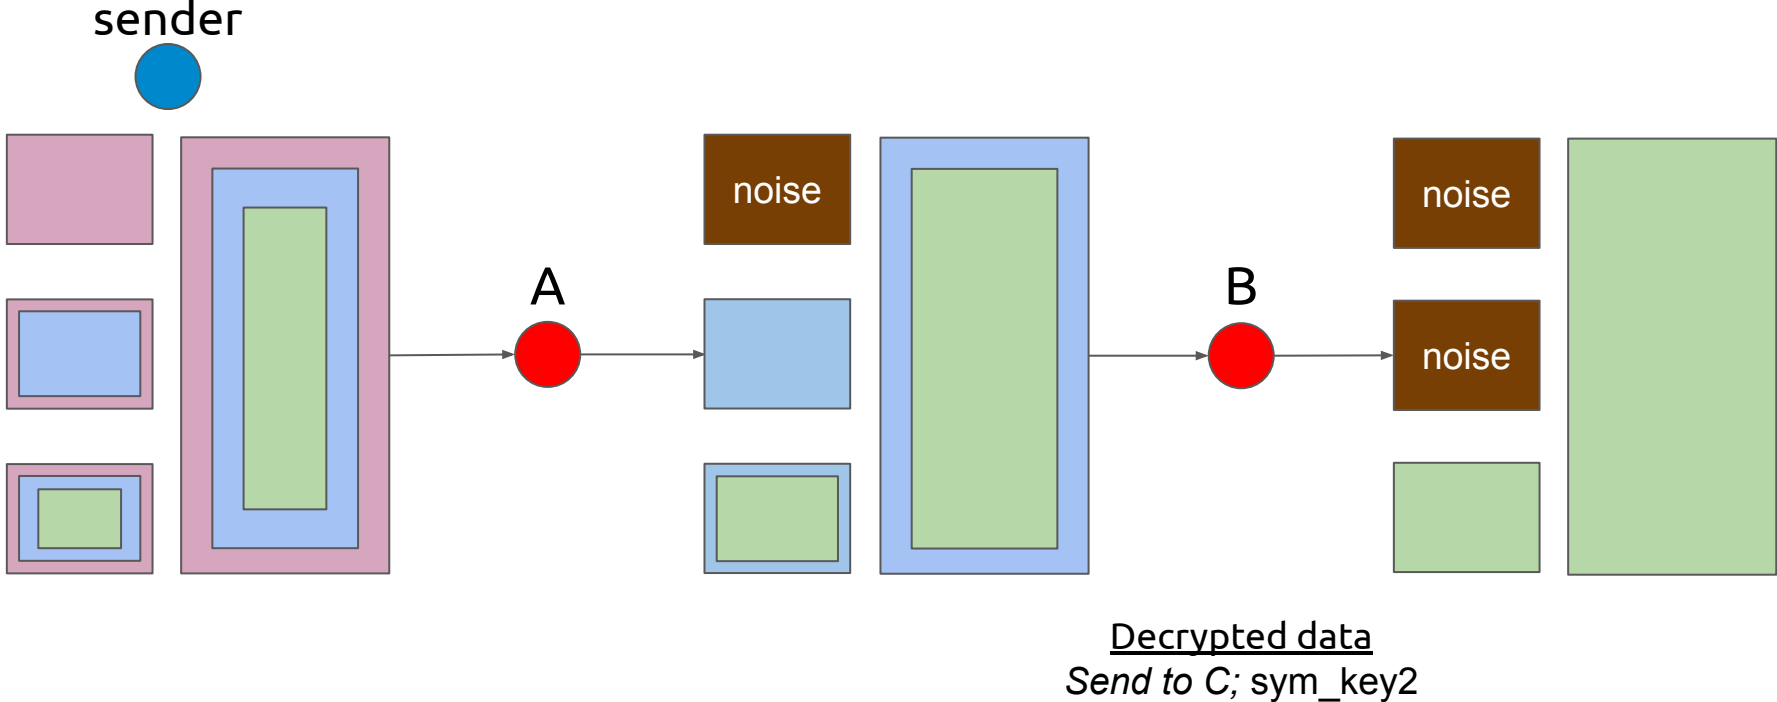
sender



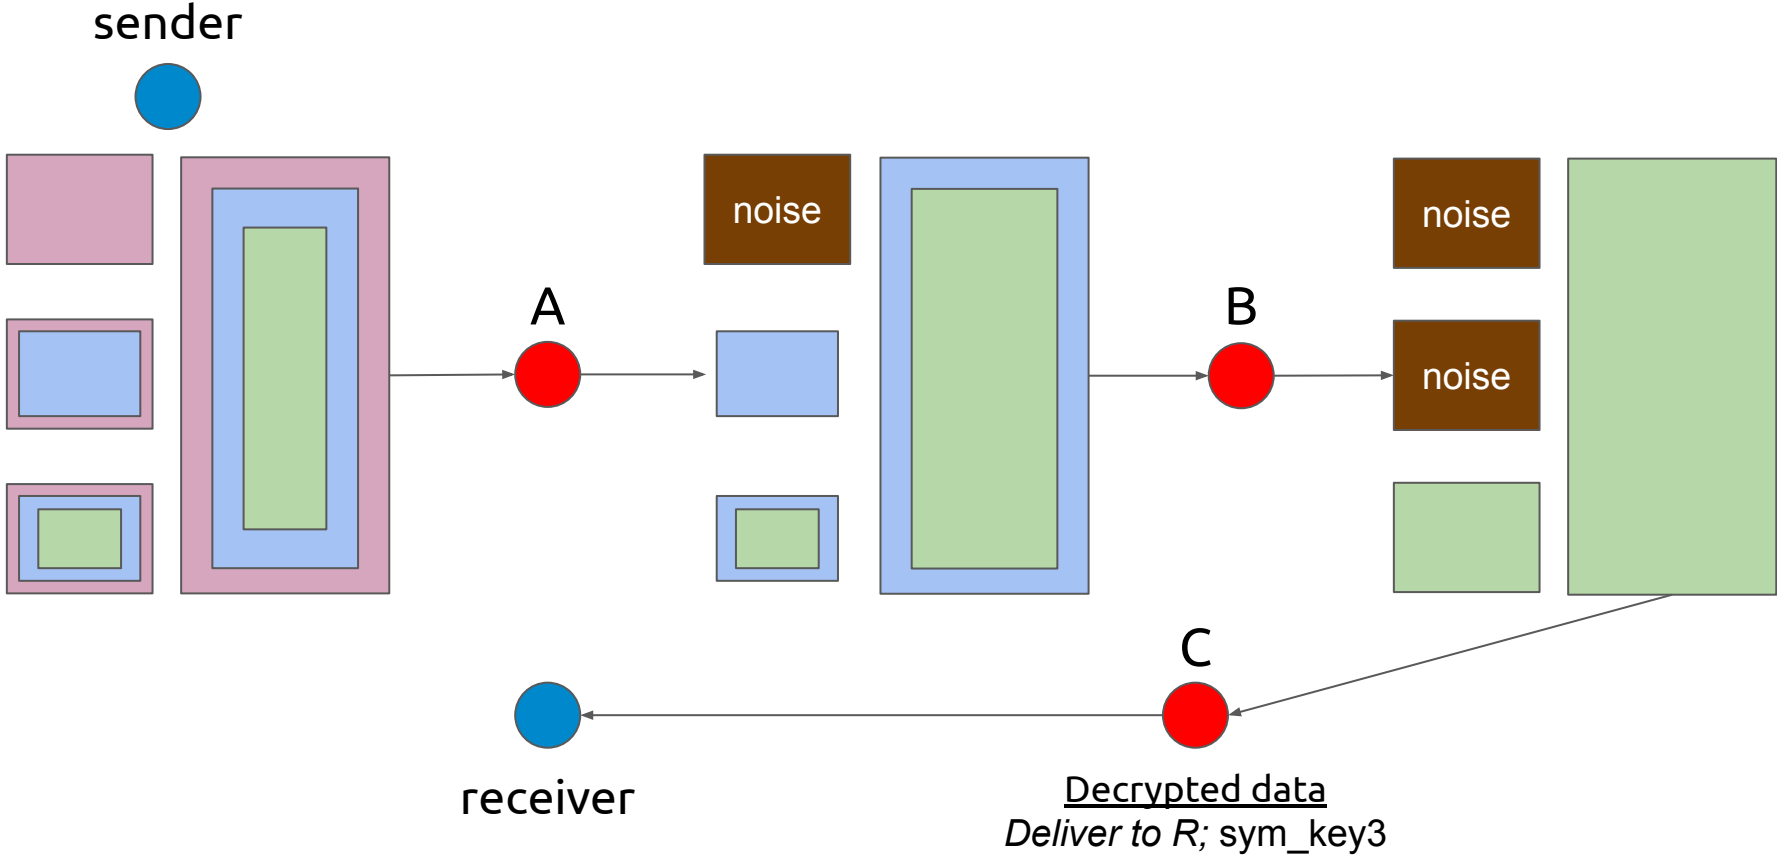
# Routing via a Mix-net



# Routing via a Mix-net



# Routing via a Mix-net



# Coordination

- CoinJoin and similar techniques require:
  - Coordination
  - Message passing between multiple parties
- How do parties find each other?
- How to prevent DoS attacks?
- Is it possible to improve with more advanced cryptographic techniques?

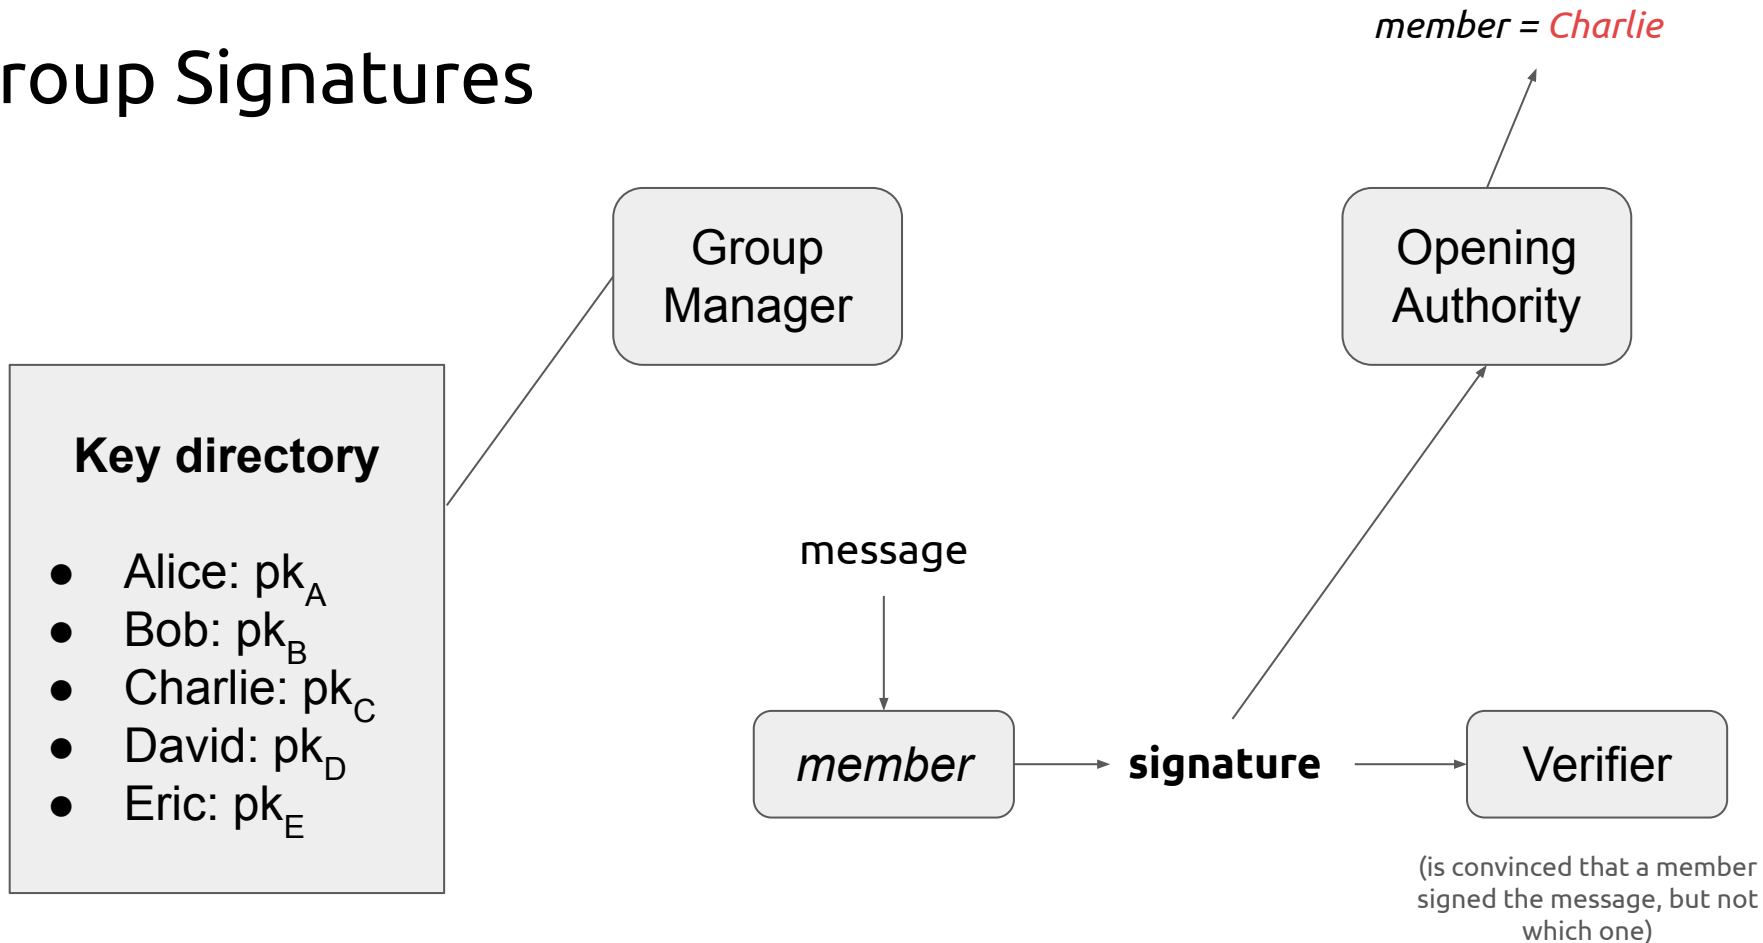


# Anonymity and Digital Signatures

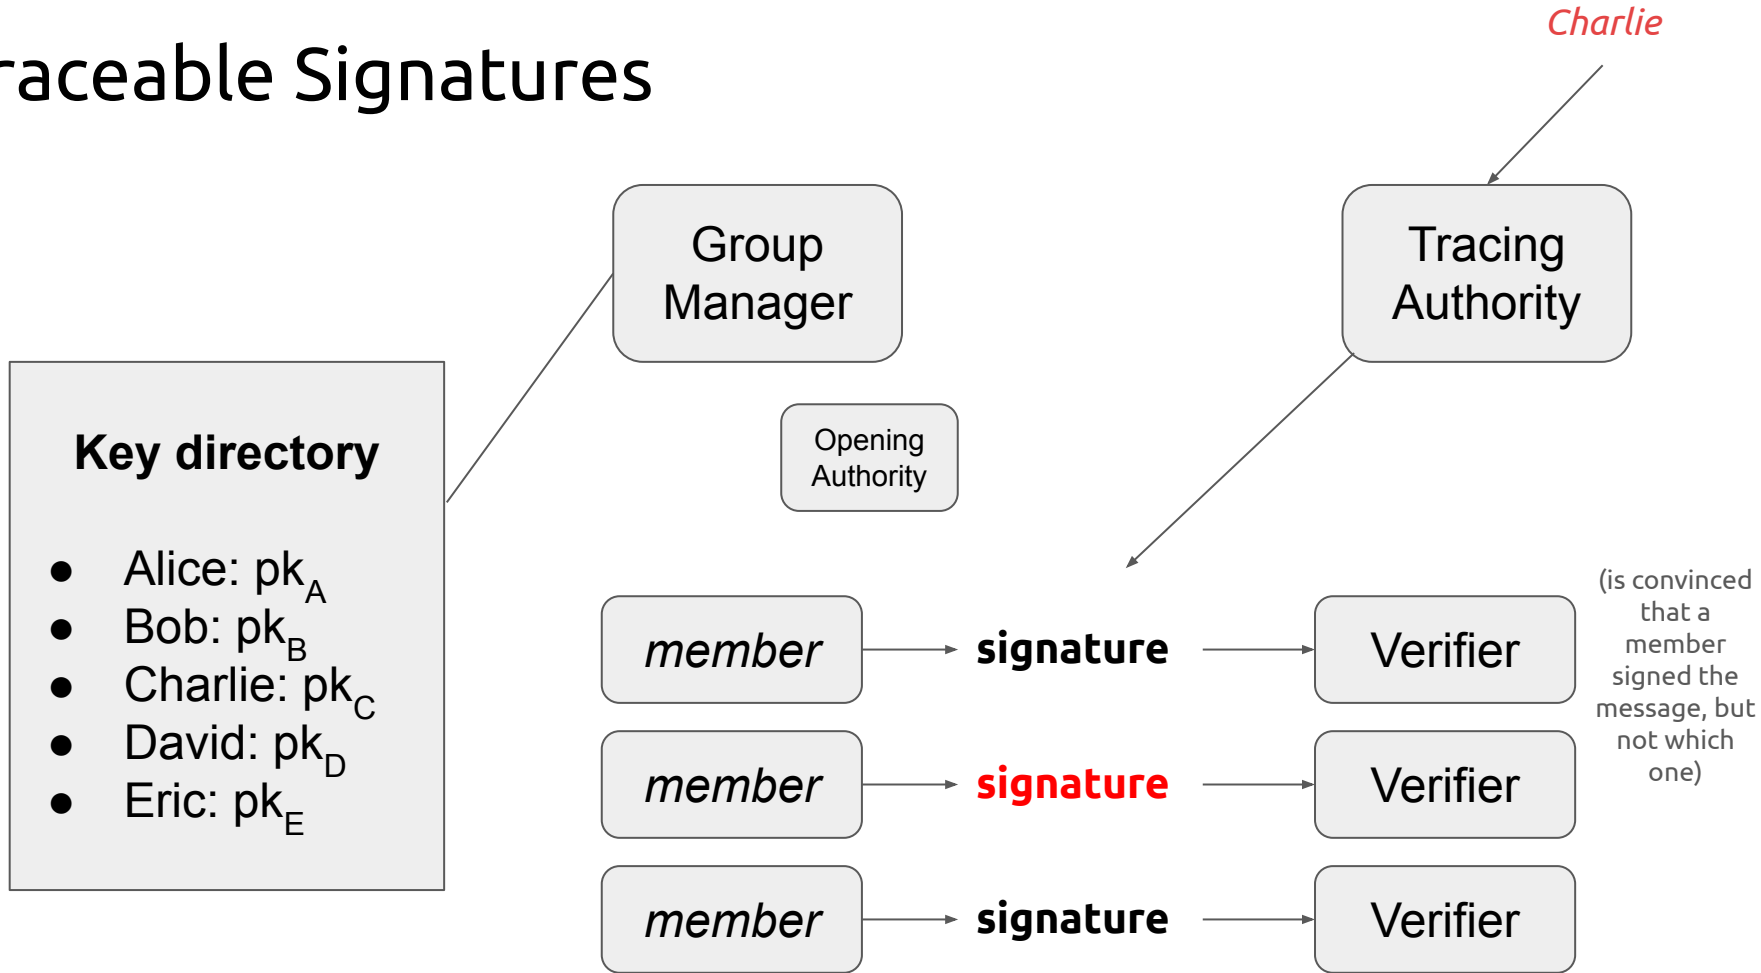
# Anonymity and Digital Signatures

- So far all digital signatures identify the signer
- Is it possible to hide the sender within a group?

# Group Signatures



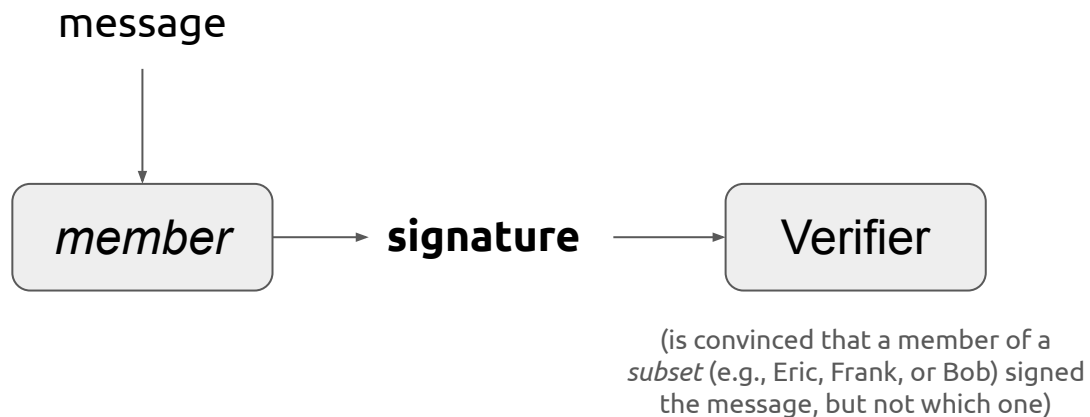
# Traceable Signatures



# Ring Signatures

## Key directory

- Alice:  $pk_A$
- Bob:  $pk_B$
- Charlie:  $pk_C$
- David:  $pk_D$
- Eric:  $pk_E$



# Monero/Cryptonote

- **Linkable ring signatures**
- **“Stealth” addresses**
- For each payment, an anonymity set is selected with accounts of the same monetary value
- A ring signature is issued on behalf of that set:
  - suitably restricted s.t. an account can only be used once
  - if an output is used twice, it is *linkable*
- **Stealth addresses enable:**
  - the sender to create unlinkable addresses for the receiver
  - the receiver to detect said addresses

# Is Monero Anonymous?

- There is potentially more uncertainty in the Monero blockchain compared to a Bitcoin-like blockchain (even with Coinjoin transactions)
- However, it is not obvious how to **quantify the level of anonymization**
- **De-anonymization is feasible** in reasonable real-world threat models
  - e.g., the attacker “sprays” the ledger with transactions s.t. it commands a good number of selected accounts

# The importance of the anonymity set

Dec 18, 2013, 01:46pm EST

## Harvard Student Receives F For Tor Failure While Sending 'Anonymous' Bomb Threat



**Runa A. Sandvik** Former Contributor

Tech

*I cover all things privacy, security and technology.*

Follow

According to the five-page complaint, the student "took steps to disguise his identity" by using Tor, a software which allows users to browse the web anonymously, and Guerrilla Mail, a service which allows users to create free, temporary email addresses.



(Photo credit: joeythibault)

What Kim didn't realize is that Tor, which masks online activity, doesn't hide the fact that you are using the software. In analyzing the headers of the emails sent through the Guerrilla Mail account, authorities were able to determine that the anonymous sender was connected to the anonymity network.

Using that conclusion, they then attempted to discern which students had been using Tor on the Harvard wireless network around the time of the threats. Before firing up Tor, Kim had to log on to the school's

Given how quickly he was found, Kim was likely one of the few—if not the only—individuals on Tor around on Monday morning. According to authorities, he "anonymously" emailed threats including ""bombs



# Increasing and Safeguarding the anonymity set

- A **larger anonymity set** is most **preferable**
- In the techniques seen so far, transaction preparation work **increases linearly** with the anonymity set
- **Goal:** use the set of all possible Unspent Transaction Outputs (UTxOs)

ZeroCash

# ZK-Snarks

- Zero-knowledge succinct arguments of knowledge
- Similar to “zero-knowledge proofs”
- Can prove possession of a witness for any public statement / predicate
- **Zero-Knowledge**
  - Nothing aside the fact that the statement is true is leaked.
- **Computational soundness:**
  - depends on the security of a “common reference string” (a structured cryptographic information that is assumed to be honestly sampled)
- **Succinctness:**
  - the proof size and the verifier’s running time is efficient
  - proportional to the statement only

# Constructing ZK-SNARKs


- There exist a SNARK for any NP-relation  $R$

$NP = \{ L \mid \text{exists } R: x \text{ in } L \text{ iff } (x, w) \text{ in } R; R \text{ is polynomial time} \}$

- The actual proof sizes are small (hundreds of bytes)
- Verification does not depend on the running time of  $R$


# Additional Tools

## Commitment scheme



- Hiding
- Binding  
 $\nexists r', m'$ , with  $m \neq m'$  s.t.  
 $\text{Commit}(m;r) = \text{com}$  and  
 $\text{Commit}(m';r') = \text{com}$


## Pseudo-random functions



- $y$  looks random to someone who does not have  $\text{SK}$

# Anonymous cash

Coin creation




# Anonymous cash

Coin creation



**sn**




# Anonymous cash

Coin creation




$sn$   $r$






# Anonymous cash

Coin creation




# Anonymous cash

Coin creation




# Anonymous cash

Spend the coin




# Anonymous cash

Spend the coin



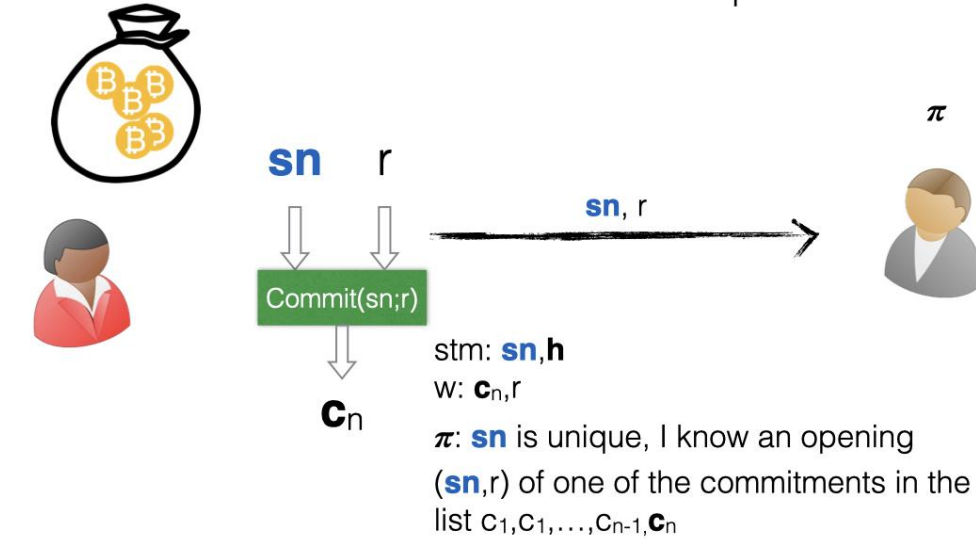
# Anonymous cash

Spend the coin




# Anonymous cash

Spend the coin




# Anonymous cash

Spend the coin



# Anonymous cash

Spend the coin



New hidden coin = new commitment added to the list


stm:  $sn, h$   
 W:  $C_n, r$   
 $\pi$ :  $sn$  is unique, I know an opening  
 ( $sn, r$ ) of one of the commitments in the  
 list  $C_1, C_1, \dots, C_{n-1}, C_n$





# Anonymous cash

Spend the coin



stm:  $sn, h$   
 W:  $C_n, r$   
 $\pi$ :  $sn$  is unique, I know an opening  
 ( $sn, r$ ) of one of the commitments in the  
 list  $C_1, C_1, \dots, C_{n-1}, C_n$


New hidden coin = new commitment  
 added to the list

The size of the blockchain would grow  
 dramatically



# Anonymous cash

Spend the coin



$C_n$

stm:  $sn, h$   
 W:  $C_n, r$   
 $\pi$ :  $sn$  is unique, I know an opening  
 ( $sn, r$ ) of one of the commitments in the  
 list  $C_1, C_1, \dots, C_{n-1}, C_n$

New hidden coin = new commitment  
 added to the list



The size of the blockchain would grow  
 dramatically

Merkle Trees





# Anonymous cash

Spend the coin





# Anonymous cash

Spend the coin




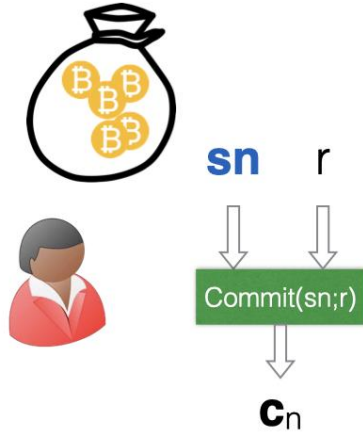
# Anonymous cash

Spend the coin





# Anonymous cash

Spend the coin



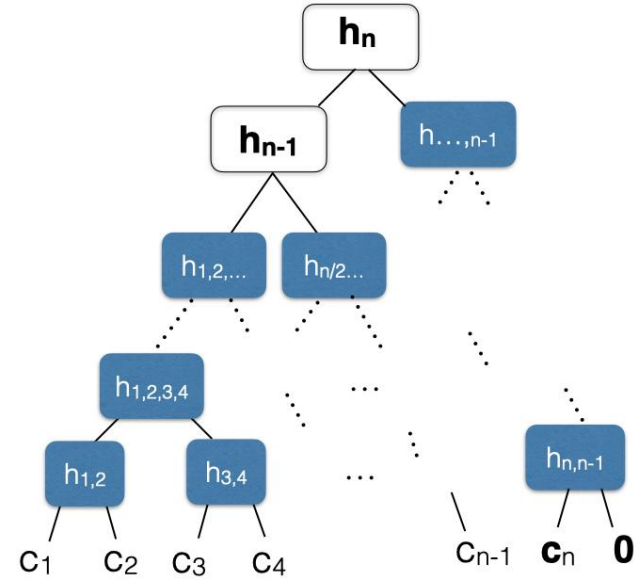

# Anonymous cash

Spend the coin



# Anonymous cash

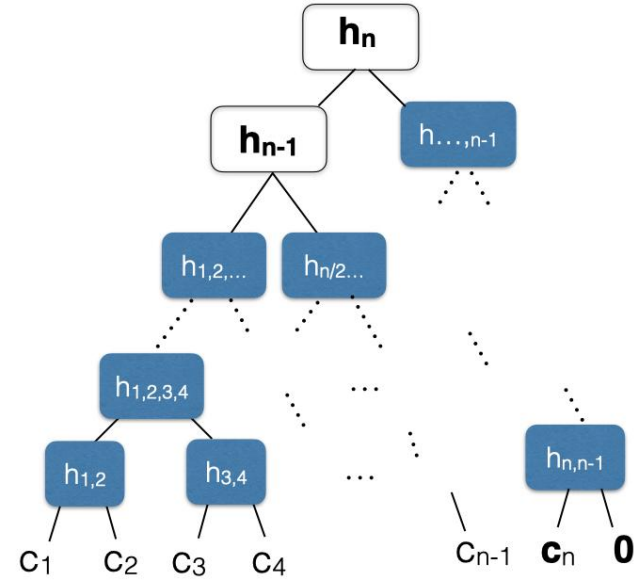
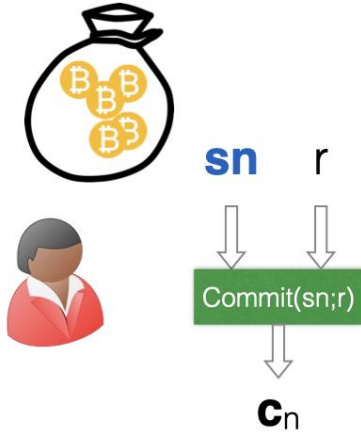
Spend the coin





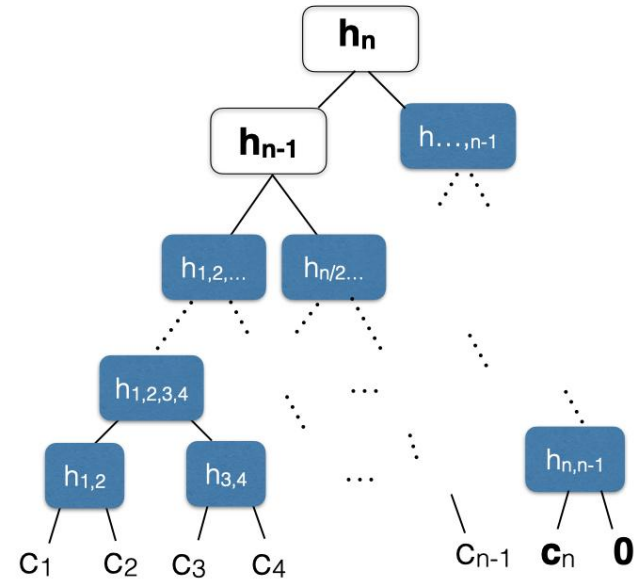

# Anonymous cash

Spend the coin



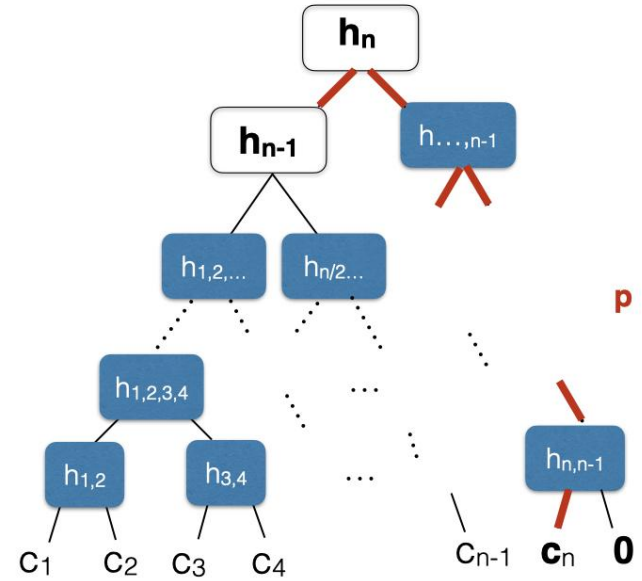

# Anonymous cash

Spend the coin




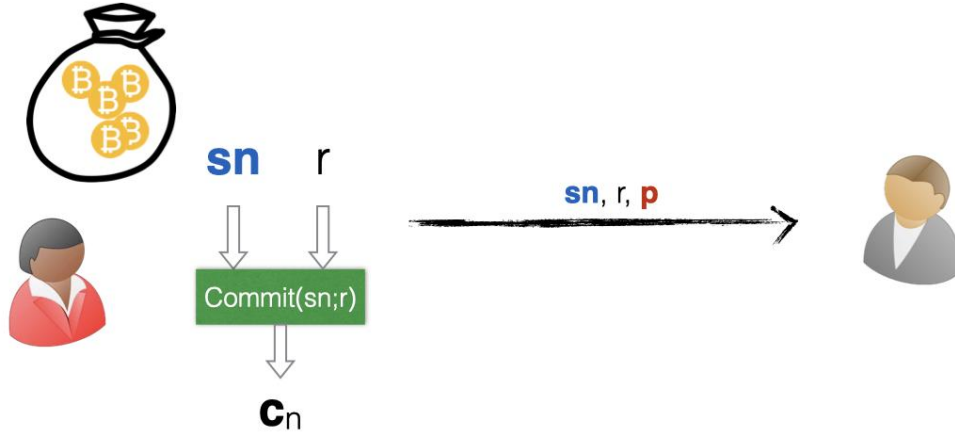
# Anonymous cash

Spend the coin





# Anonymous cash

Spend the coin




# Anonymous cash

Spend the coin



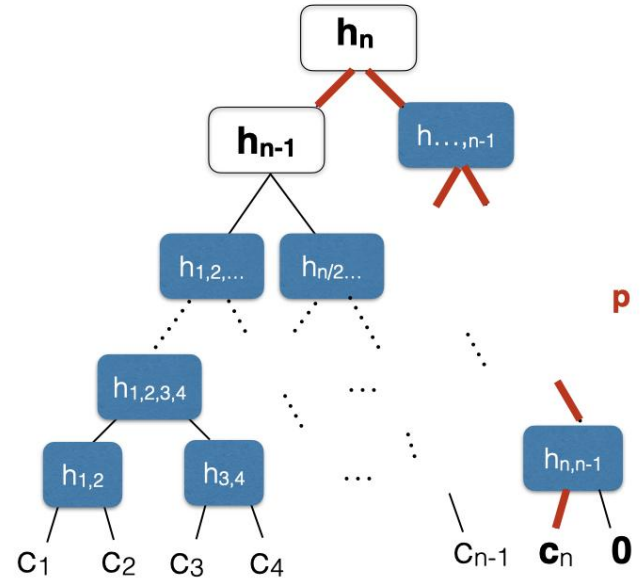
# Anonymous cash

Spend the coin



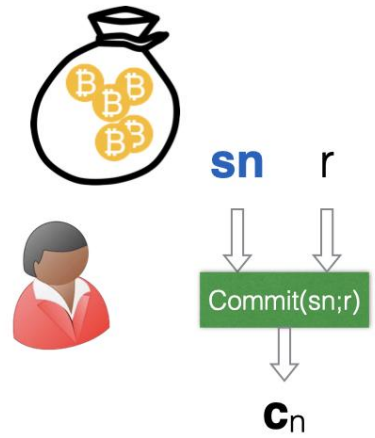
stm: **sn, h**  
 W: C<sub>n</sub>, r, **p**

$\pi$ : **sn** is unique, I know a path **p** in the Merkle tree **h<sub>n</sub>** that contains a commitment for which I know the opening (**sn, r**)



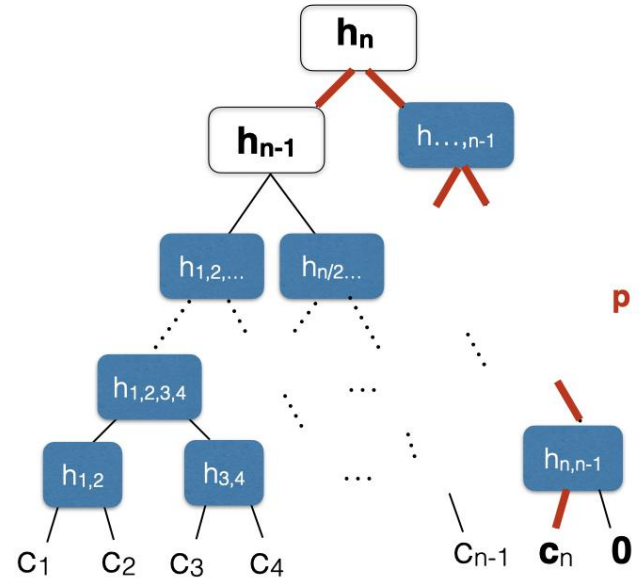
# Anonymous cash

Spend the coin





stm:  $sn, h$   
 W:  $C_n, r, p$

$\pi$ :  $sn$  is unique, I know a path  $p$  in the Merkle tree  $h_n$  that contains a commitment for which I know the opening  $(sn, r)$



# Anonymous cash




Anonymous to anyone








# Anonymous cash

Anonymous to anyone





# Anonymous cash

Anonymous to anyone




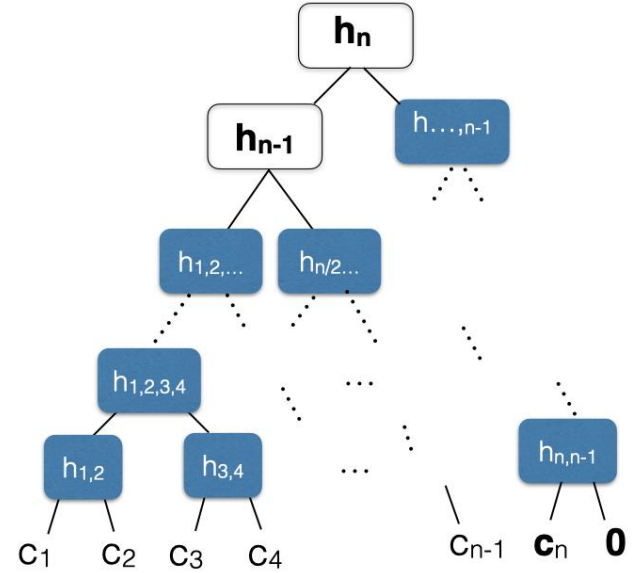


# Anonymous cash

Anonymous to anyone





# Anonymous cash

Anonymous to anyone


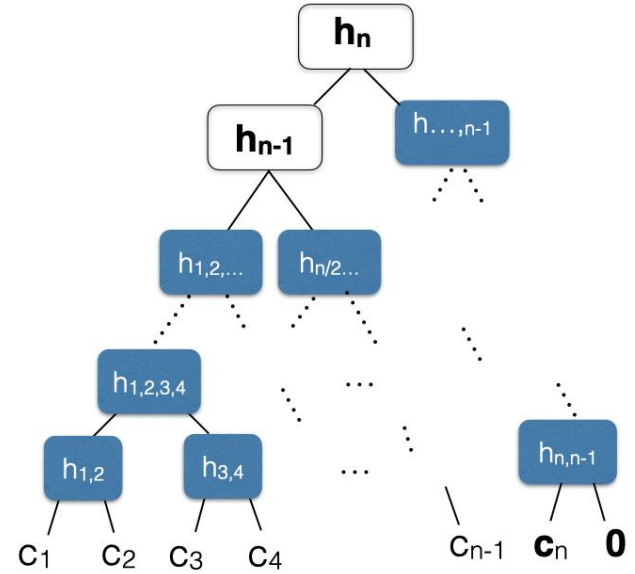


# Anonymous cash

Anonymous to anyone




stm:  $sn, h_n, v$   
 w:  $Int, C_n, r, r', p$



# Anonymous cash

Anonymous to anyone





$sn, v, Int, C_n, r, r', p$




stm:  $sn, h_n, v$   
 w:  $Int, C_n, r, r', p$

$\pi$ :  $sn$  is unique, I know a path  $p$  in the Merkle tree  $h_n$  that contains a commitment for which I know the opening  $((Int, v), r')$ . Moreover, I know the opening of  $Int$  to  $(sn, r)$



# Anonymous cash

Anonymous to anyone




$sn, v, Int, C_n, r, r', p$



stm:  $sn, h_n, v$   
w:  $Int, C_n, r, r', p$

$\pi$ :  $sn$  is unique, I know a path  $p$  in the Merkle tree  $h_n$  that contains a commitment for which I know the opening  $((Int, v), r')$ . Moreover, I know the opening of  $Int$  to  $(sn, r)$



# Anonymous cash

Anonymous to anyone





# Anonymous cash

Anonymous to anyone



PK,SK



# Anonymous cash

Anonymous to anyone

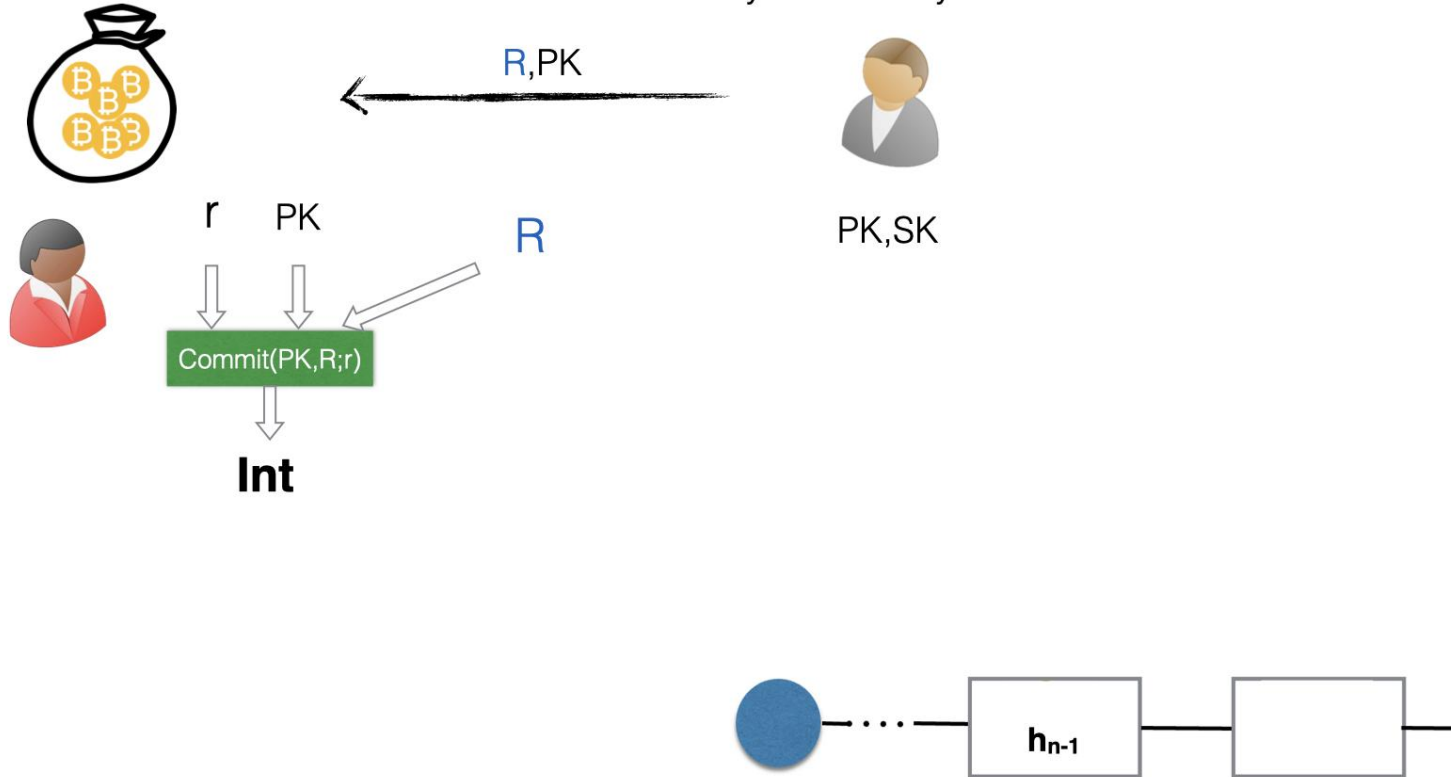


PK, SK



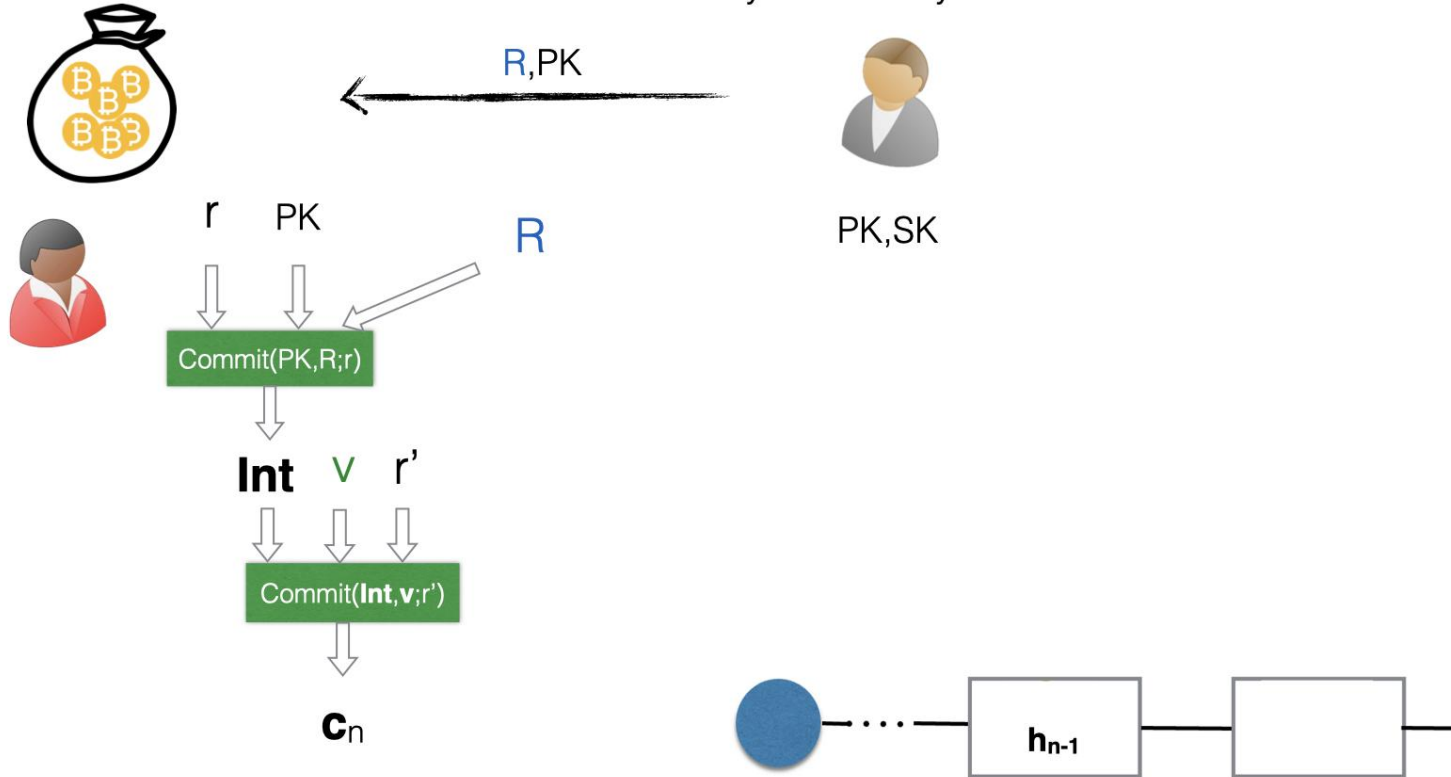
# Anonymous cash

Anonymous to anyone



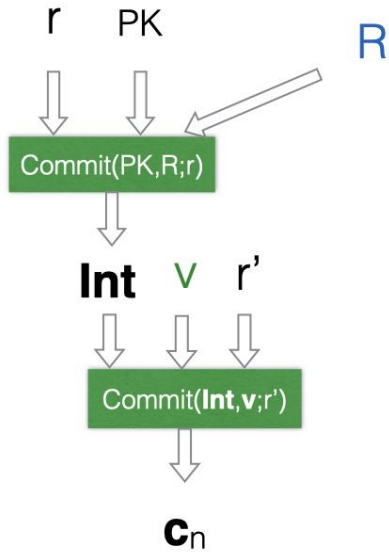
# Anonymous cash

Anonymous to anyone




# Anonymous cash

Anonymous to anyone





$PK, SK$



# Anonymous cash

Anonymous to anyone

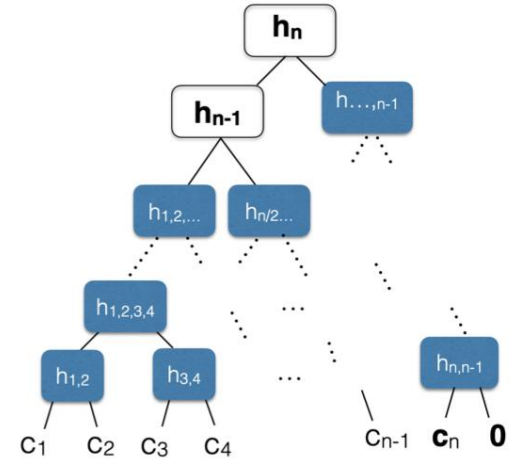


**$C_n$**



$PK, SK$

$R$





# Anonymous cash

Anonymous to anyone




PK,SK



$\text{Int}$   $v$   $r'$



$c_n$

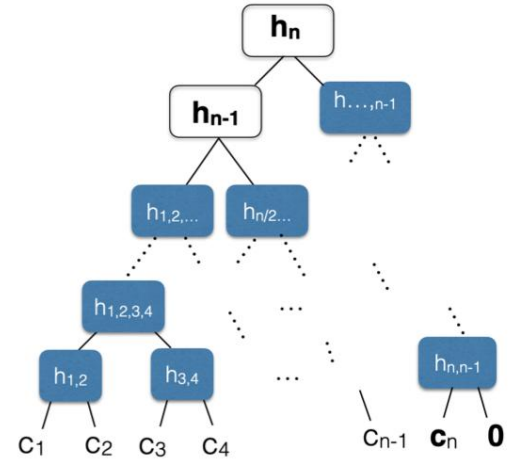




# Anonymous cash

Anonymous to anyone




PK,SK







# Anonymous cash

Anonymous to anyone




PK, SK

stm:  $sn, h_n, v$   
 w:  $\text{Int}, c_n, r, r', R, p$



# Anonymous cash

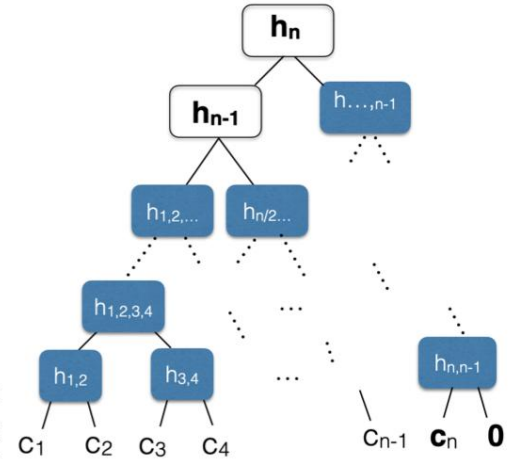
Anonymous to anyone



PK, SK



stm:  $sn, h_n, v$   
 w:  $Int, c_n, r, r', R, p$

$\pi$ :  $sn$  is unique, I know a path  $p$  in the Merkle tree  $h_n$  that contains a commitment for which I know the opening  $((Int, v), r')$ . Moreover, I know the opening of  $Int$  to  $((R, PK), r)$  s.t.  $PRF(SK, R) = sn$  AND  $SK$  is a secret key for  $PK$



# Anonymous cash

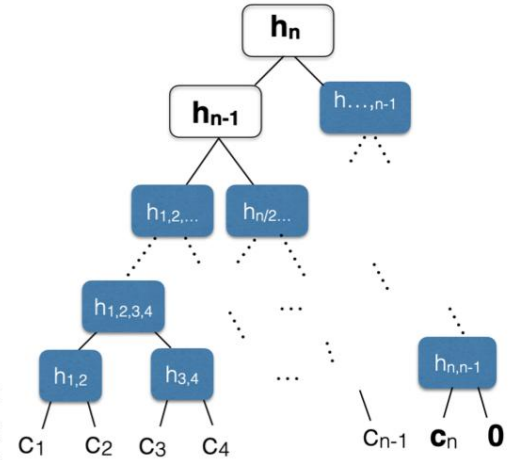
Anonymous to anyone



PK, SK

stm:  $sn, h_n, v$   
 w:  $Int, c_n, r, r', R, p$

$\pi$ :  $sn$  is unique, I know a path  $p$  in the Merkle tree  $h_n$  that contains a commitment for which I know the opening  $((Int, v), r')$ . Moreover, I know the opening of  $Int$  to  $((R, PK), r)$  s.t.  $PRF(SK, R) = sn$  AND  $SK$  is a secret key for  $PK$



Network security

# Overlay Networks


- A reliable network is critical for blockchains and distributed ledger protocols to operate
- Typically they utilize an **overlay network**
  - a network built on top of another network
  - virtual links connect the participating nodes

# Overlay Networks

in a network, we would like nodes to be fully connected

relevant operations :

1. point-to-point communication
2. broadcast



# Network Requirements

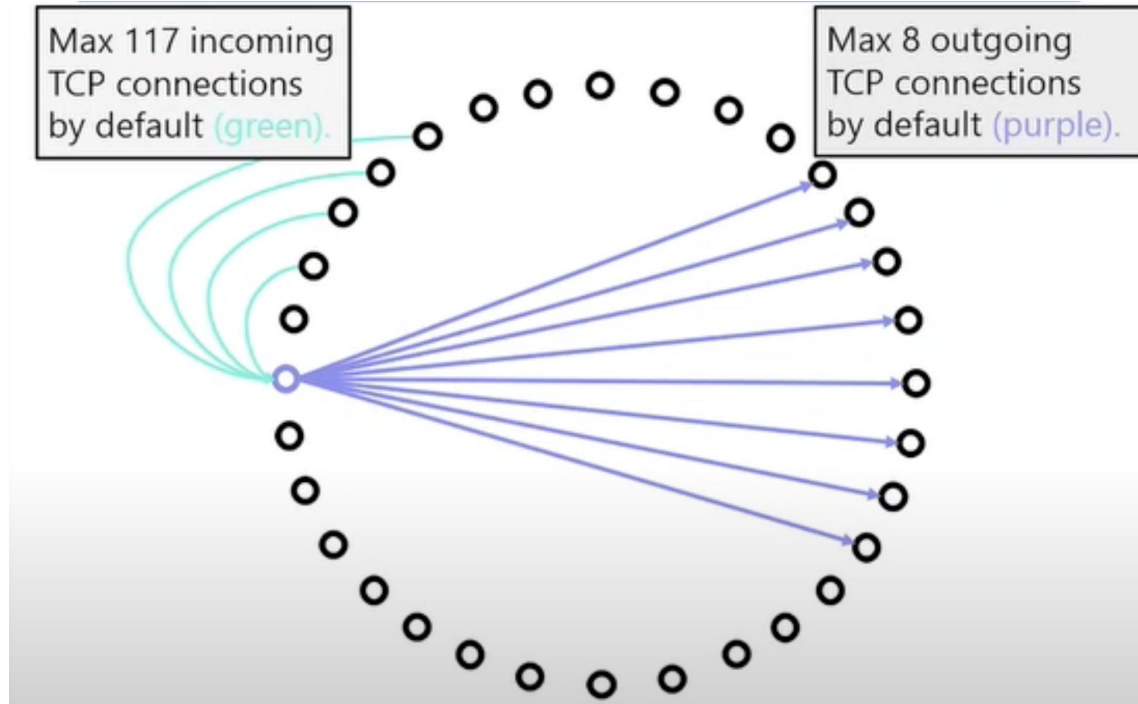
- Synchronicity
- Reliable message transmission
- Reliable Broadcast

# Bitcoin's P2P Network

- A **Peer-to-Peer** (P2P) network over TCP/IP
- Peers are identified by their IP address
- Peers can **diffuse** messages to be propagated to the whole network
- Peers initiate a **small number** of **outgoing** connections
- Peers receive a limited number of **incoming** connections




# Eclipse attack (overview)



# Eclipse attack (overview)


We manipulate the node so all its outgoing connections are to attacker IPs.



# Eclipse attack (overview)


We manipulate the node so all its outgoing connections are to attacker IPs.

1. Fill node's peer tables, with attacker IPs.



# Eclipse attack (overview)

We manipulate the node so all its outgoing connections are to attacker IPs.



1. Fill node's peer tables, with attacker IPs.

2. The node restarts and loses its current outgoing connections.



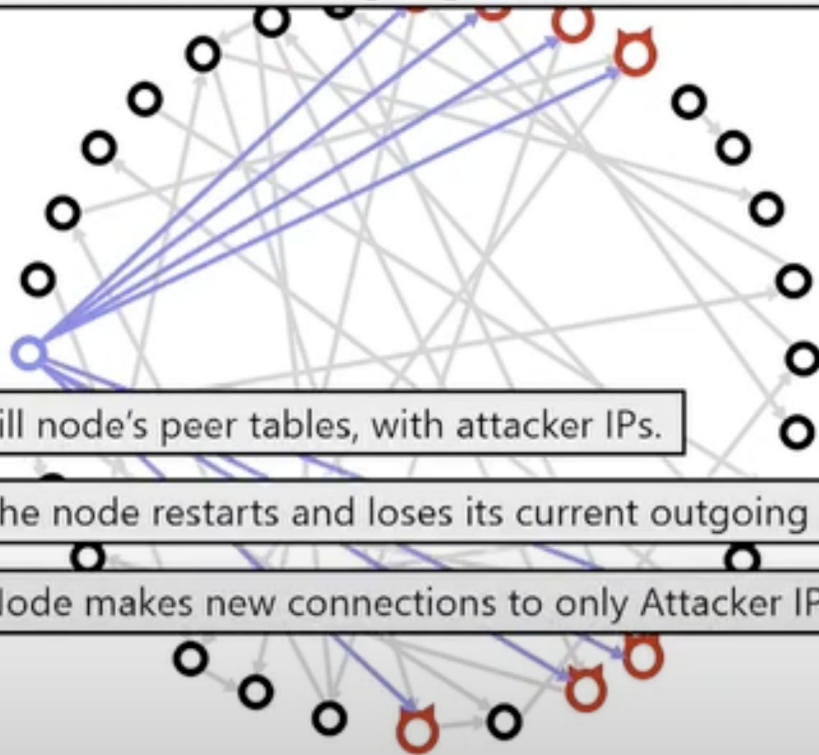
# Eclipse attack (overview)

We manipulate the node so all its outgoing connections are to attacker IPs.

1. Fill node's peer tables, with attacker IPs.

2. The node restarts and loses its current outgoing connections.

3. Node makes new connections to only Attacker IPs.



# Implications of the attack

- Controlling blocks propagation
- Splitting mining power
- Confirmation double spending
  - Send the merchants a tx T, but send T' to the rest of the network.

# (Deeper look into eclipse attack) P2P Networks

- (In the case of Bitcoin) The requesting node contacts a **DNS Seeder**:
  - A node with a public IP address that can serve a list of IP addresses for Bitcoin nodes
  - Obtains those addresses via **crawling**
- If the connection fails, the node has a hardcoded set of IP addresses
- Peers exchange node IP addresses via **ADDR messages** that contain a selection of a peer's address book

# Table maintenance

- Nodes maintain tables of peers that they have learned:
  - Nodes that have proven to be operational
  - Nodes for which the node has been informed about their existence, but they have not been contacted yet
- Tables are updated on a regular basis
- Timestamp information is stored from the last connection attempt



# Attacking the P2P layer - Key Observations

- A node will add an address to the 'tried' table if it receives an incoming connection from another node
- A node will accept unsolicited ADDR messages; these will be added to the 'new' table
- Nodes rarely solicit information from DNS seeders and other nodes

# Eclipse Attack

- Victim is a node with a public IP
- Attacker makes outgoing connection to the node using adversarial nodes
  - 'tried' table gets full with fresh adversarial IP's
- Attacker uses ADDR messages to insert trash IP's into the 'new' table of the victim
- Attacker waits for the victim node to restart (nodes maintain existing outgoing connections)
  - Restarts can happen because of a software update or even deliberately by the attacker via a DOS attack

# Eclipse Attack

- The attacker can repetitively connect to victim node to ensure timestamps of adversarial nodes are fresh
- If a 'new' address is selected:
  - injection of trash IPs ensures that, with some probability, the new node will not be responsive
  - another coin flip will be attempted for the connection, which can result to an adversarial IP

# Eclipse Attack

- Attacker saturates the incoming connections of the victim
  - The protocol allows for the same IP to occupy all 117 incoming TCP/IP connections
- It becomes impossible for other nodes to connect to the victim
- As maximum number of connections is reached, the victim will deny any other incoming connections

# Eclipse Attack

- Once the eclipse takes place, all (incoming/outgoing) communication of the victim is routed via the attacker nodes
  - victim's transactions may be censored
  - victim's blocks can be dropped
  - victim's blockchain could be populated almost entirely by adversarial blocks!
- The rest of the network will eventually completely forget about the victim node
  - a function *isTerrible* is executed periodically on the tables to remove any node that has an over-30-days old timestamp and too many failed connection attempts


# Attack Countermeasures

- Many mitigation techniques can be used:
  - ban unsolicited ADDR messages
  - diversify incoming connections
  - test before evicting addresses from the tried table
- The possibility of an attack cannot be zeroed

# Wallets

# Full nodes

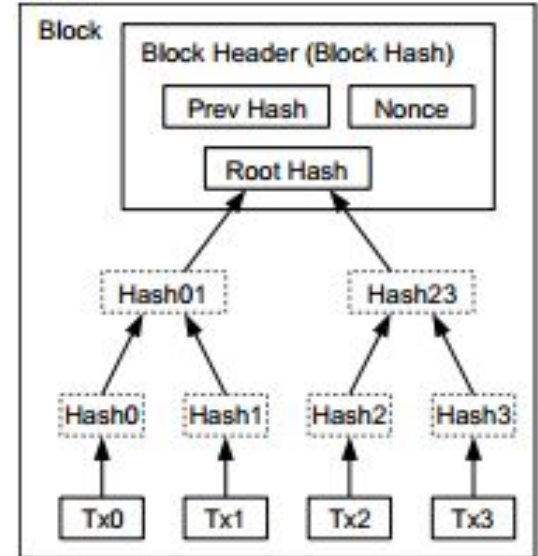
- Some wallets maintain the whole blockchain
- Full nodes:
  - Keep the whole blockchain history
  - Keep the whole UTXO set
  - Verify each tx
  - Verify each block
  - Relay every tx and block





# Recall : Merkle trees of transactions

- Transactions not yet confirmed, but received by a full node are collected into a data structure called the **mempool**
- To build a block, the mempool transactions **are collected** into a Merkle Tree in an (arbitrary, but valid) order defined by the miner
- The application data in the **block header**, for which the Proof-of-Work equation is solved, only contain the **root** of this Merkle Tree: **x**



# Advantages of using a Merkle tree

- Proof-of-Work difficulty **does not depend** on the number of confirmed transactions
  - each miner is incentivized to include all transactions they can, which have a non-zero fee
- The PoW difficulty **only depends on the target T**
  - this allows better control of the mining rate
- It enables **SPV (Simple Payment Verification) wallets!**

# SPV

- Simple Payment Verification
- A different type of wallet
- Useful for mobile, laptops etc.
- Doesn't need to download the whole blockchain
  - Does not download all transactions
  - Much faster than standard (full) node
- Keeps **only the block headers from genesis till today**
- Connects to multiple **untrusted** servers
- Server is a full node which **proves** to the SPV wallet each claim

# SPV

- Wallet sends to the SPV server the bitcoin addresses they have
  - Not the private keys!
  - The SPV server knows which transactions to send to the SPV client
  - The addresses are shared via a Bloom filter
- Wallet **verifies** each block's **PoW** and authenticated ancestry
  - Keeps a longest chain as usual
  - Does not keep transactions
- Wallet verifies **each transaction** it receives
  - Signatures
  - Law of conservation
- Wallet verifies that the transaction belongs to the Merkle Tree root of a block

# SPV Security

- SPV wallets
  - **don't keep** a UTXO
  - **don't verify or receive** transactions they are not interested in
  - **don't verify** coinbase validity
- Have the *same level of security* as a regular full node
  - assuming honest majority
- What can a malicious SPV server achieve?
  - Temporary fork to invalid block (invalid coinbase, transactions, non-existing UTXO, double spending...)

# Wallet seeds and HD wallets

- Hierarchical Deterministic (HD) wallet
- An infinite sequence of wallet private keys can be generated from a single “master private key” (BIP-32)
- A private key can be encoded as a human-readable **seed**
- Seed is sufficient to **recover** all the private keys of a wallet
  - Typically backed up on paper
  - Optionally encrypted with password

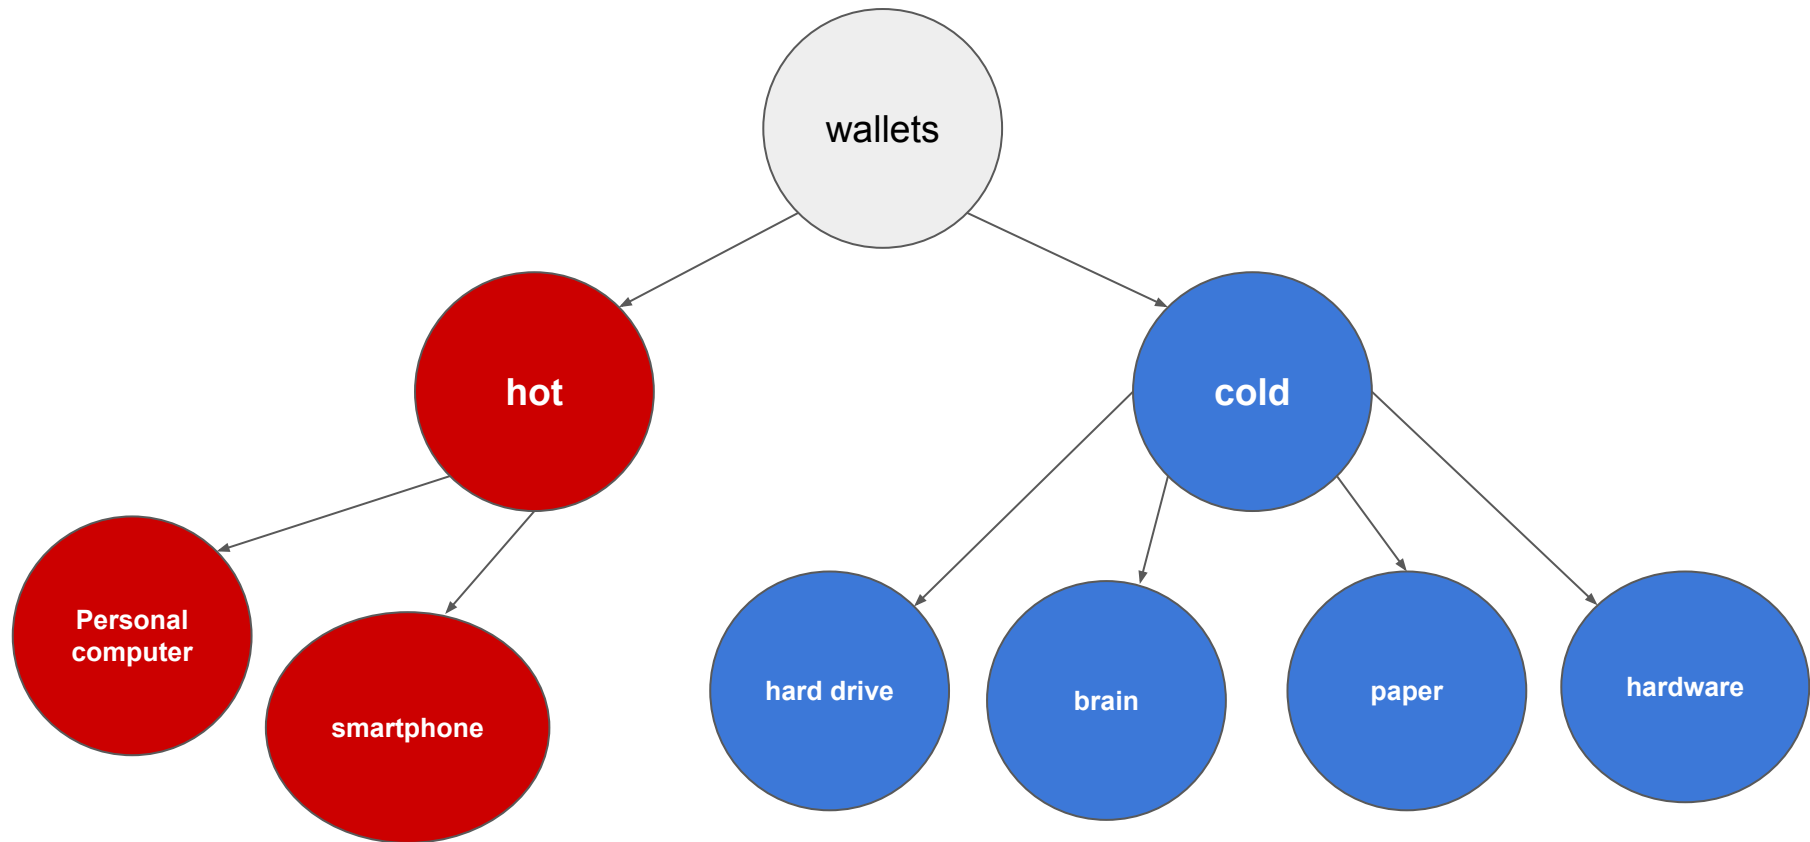
*Seed Example:*

deal smooth awful edit virtual monitor term sign start home shrimp wrestle

# Hot and cold wallets

- Keys on an Internet-connected computer: **Hot** wallet
  - Easy to use
  - Can always spend my money immediately
- Private keys offline: **Cold** wallet
  - Kept on a computer not connected to the Internet or a hard drive
  - Keys cannot easily be stolen
  - Keys can be moved to a hot wallet when needed to spend
  - User can see balance and how much money they have using public keys kept (safely) online

# Wallet classification





# References

- Chaum, David, Amos Fiat, and Moni Naor. "Untraceable electronic cash." Advances in Cryptology—CRYPTO'88: Proceedings 8. Springer New York, 1990. [https://link.springer.com/content/pdf/10.1007/0-387-34799-2\\_25.pdf](https://link.springer.com/content/pdf/10.1007/0-387-34799-2_25.pdf)
- Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, Madars Virza, Zerocash: Decentralized Anonymous Payments from Bitcoin, proceedings of the IEEE Symposium on Security & Privacy (Oakland) 2014, 459-474, IEEE, 2014. <http://zerocash-project.org/media/pdf/zerocash-extended-20140518.pdf>
- Reed, Michael G., Paul F. Syverson, and David M. Goldschlag. "Anonymous connections and onion routing." IEEE Journal on Selected areas in Communications 16.4 (1998): 482-494. <https://www.usenix.org/system/files/conference/usenixsecurity17/sec17-alexopoulos.pdf>
- Bitcoin Network: <https://en.bitcoin.it/wiki/Network>
- Heilman, Ethan, et al. "Eclipse attacks on {Bitcoin's}{peer-to-peer} network." 24th USENIX security symposium (USENIX security 15). 2015. <https://www.usenix.org/system/files/conference/usenixsecurity15/sec15-paper-heilman.pdf>
- For zero-knowledge: References from the book of Goldreich Oded: "Foundations of Cryptography: Volume 1, Basic Tools"
  - Sec. 4.2 until (included) Sec. 4.2.2
  - Sec. 4.3 until (included) Sec. 4.3.2
  - Sec. 4.7 until (included) Definition 4.7.2
- For the notions of commitment schemes, pseudorandom functions, and the Merkle tree construction please check the book "Introduction to Modern Cryptography" by Jonathan Katz Yehuda Lindell.