

# BITCOIN AND BLOCKCHAINS

CS6410

Hakim Weatherspoon

[slide liberally taken from Kevin Seqniqi, Ittay Eyal, Emin Gun Sirer, Robbert van Renesse]

# A Brave New World - The Vision of David



## David Chaum

- PhD CS/Business Adm from Berkeley 1982
- Founded International Association for Cryptologic Research (IACR) same year
- Known for *eCash*, *mix nets*, *voting systems*..

# A Brave New World - The Vision of David Chaum [1983]



## INTRODUCTION

Automation of the way we pay for goods and services is already underway, as can be seen by the variety and growth of electronic banking services available to consumers. The ultimate structure of the new electronic payments system may have a substantial impact on personal privacy as well as on the nature and extent of criminal use of payments. Ideally a new payments system should address both of these seemingly conflicting sets of concerns.

On the one hand, knowledge by a third party of the payee, amount, and time of payment for every transaction made by an individual can reveal a great deal about the individual's whereabouts, associations and lifestyle. For example, consider payments for such things as transportation, hotels, restaurants, movies, theater, lectures, food, pharmaceuticals, alcohol, books, periodicals, dues, religious and political contributions.

On the other hand, an anonymous payments systems like bank notes and coins suffers from lack of controls and security. For example, consider problems such as lack of proof of payment, theft of payments media, and black payments for bribes, tax evasion, and black markets.



# A Brave New World - The Vision of David Chaum [1983]



## INTRODUCTION

Automation of the way we pay for goods and services is already underway, as can be seen by the variety and growth of electronic banking services available to consumers. The ultimate structure of the new electronic payments system may have a substantial impact on personal privacy as well as on the nature and extent of criminal use of payments. Ideally a new payments system should address both of these seemingly conflicting sets of concerns.

On the one hand, knowledge by a third party of the payee, amount, and time of payment for every transaction made by an individual can reveal a great deal about the individual's whereabouts, associations and lifestyle. For example, consider payments for such things as transportation, hotels, restaurants, movies, theater, lectures, food, pharmaceuticals, alcohol, books, periodicals, dues, religious and political contributions.

On the other hand, an anonymous payments systems like bank notes and coins suffers from lack of controls and security. For example, consider problems such as lack of proof of payment, theft of payments media, and black payments for bribes, tax evasion, and black markets.

# A Brave New World - The Vision of David Chaum [1983]



Basically ...

- Electronic payment systems suffer from loss of privacy and cumbersome trust on single entities.
- Privacy protection, however, encounters issues of security and safety of data.

# Nick Szabo [1998]

## Bit gold

A long time ago I hit upon the idea of bit gold. The problem, in a nutshell, is that our money currently depends on **trust in a third party** for its value. As many inflationary and hyperinflationary episodes during the 20th century demonstrated, this is not an ideal state of affairs. Similarly, **private bank note issue**, while it had various advantages as well as disadvantages, similarly depended on a trusted third party.





# Nick Szabo [1998]

## Bit gold

A long time ago I hit upon the idea of bit gold. The problem, in a nutshell, is that our money currently depends on **trust in a third party** for its value. As many inflationary and hyperinflationary episodes during the 20th century demonstrated, this is not an ideal state of affairs. Similarly, **private bank note issue**, while it had various advantages as well as disadvantages, similarly depended on a trusted third party.

Thus, it would be very nice if there were a protocol whereby unforgeably costly bits could be created online with minimal dependence on trusted third parties, and then securely stored, transferred, and assayed with similar minimal trust. Bit gold.



# Nick Szabo [1998]

## Bit gold

A long time ago I hit upon the idea of bit gold. The problem, in a nutshell, is that our money currently depends on **trust in a third party** for its value. As many inflationary and hyperinflationary episodes during the 20th century demonstrated, this is not an ideal state of affairs. Similarly, **private bank note issue**, while it had various advantages as well as disadvantages, similarly depended on a trusted third party.

Thus, it would be very nice if there were a protocol whereby unforgeably costly bits could be created online with minimal dependence on trusted third parties, and then securely stored, transferred, and assayed with similar minimal trust. Bit gold.


My proposal for bit gold is based on computing a string of bits from a string of challenge bits, using functions called variously "client puzzle function," "proof of work function," or "**secure benchmark function**.". The resulting string of bits is the proof of work. Where a **one-way function** is prohibitively difficult to compute backwards, a secure benchmark function ideally comes with a specific cost, measured in compute cycles, to compute backwards.

<http://unenumerated.blogspot.com/2005/12/bit-gold.html>





# Satoshi Nakamoto and the Anon Post [2008]



## Bitcoin open source implementation of P2P currency

Posted by Satoshi Nakamoto on February 11, 2009 at 22:27

 [View Discussions](#)

I've developed a new open source P2P e-cash system called Bitcoin. It's completely decentralized, with no central server or trusted parties, because everything is based on crypto proof instead of trust. Give it a try, or take a look at the screenshots and design paper:

Download Bitcoin v0.1 at <http://www.bitcoin.org>

# Satoshi Nakamoto and the Anon Post [2008]

The root problem with conventional currency is all the trust that's required to make it work. The central bank must be trusted not to debase the currency, but the history of fiat currencies is full of breaches of that trust. Banks must be trusted to hold our money and transfer it electronically, but they lend it out in waves of credit bubbles with barely a fraction in reserve. We have to trust them with our privacy, trust them not to let identity thieves drain our accounts. Their massive overhead costs make micropayments impossible.

# Satoshi Nakamoto and the Anon Post [2008]

A generation ago, multi-user time-sharing computer systems had a similar problem. Before strong encryption, users had to rely on password protection to secure their files, placing trust in the system administrator to keep their information private. Privacy could always be overridden by the admin based on his judgment call weighing the principle of privacy against other concerns, or at the behest of his superiors. Then strong encryption became available to the masses, and trust was no longer required. Data could be secured in a way that was physically impossible for others to access, no matter for what reason, no matter how good the excuse, no matter what.

It's time we had the same thing for money. With e-currency based on cryptographic proof, without the need to trust a third party middleman, money can be secure and transactions effortless.



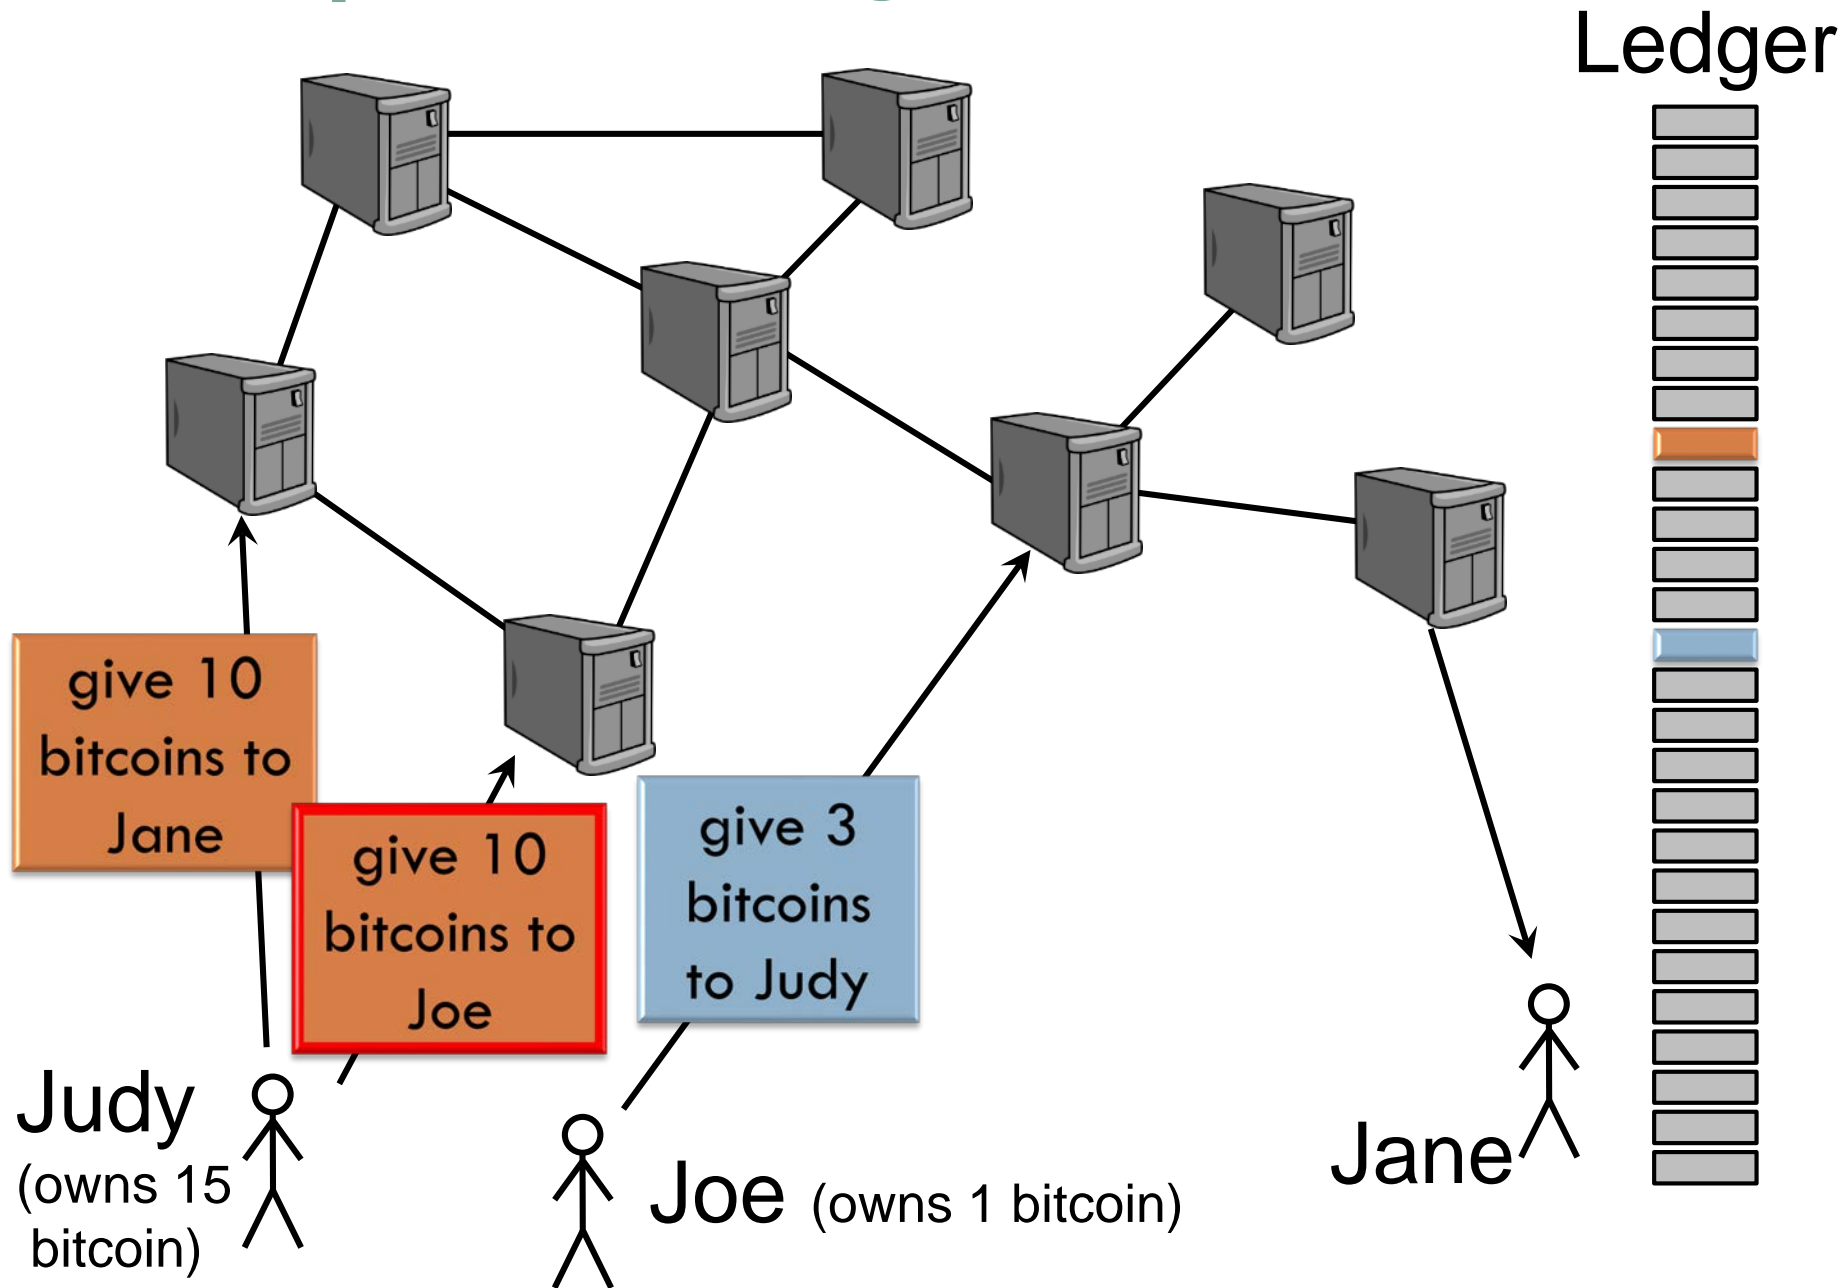
# Goals

An electronic payment system:

- Guarantees safety of transactions, protects against double spends, gives full freedom to owners.
- Yet no central trusted authority, no reliance on quorum since identities are not known.



# A Replicated Ledger of Transactions



# Bitcoin Blockchain



14

- Permissionless, open membership
- Proof-of-Work
- There are thousands of Bitcoin miners
  - ▣ they use ASIC hardware to compute SHA256 hashes
  - ▣ use about more energy than the country of Denmark
- Overall rate is a few transactions per second




# The Blockchain

Ledger





# The Blockchain

Ledger



nonce



$\text{HASH}(\text{Block}) < \textit{target}$  "cryptopuzzle"

# Cryptographic One-Way Hash Function

$\text{hash}(X) = Y$

$\text{HASH}(\text{🗄️}) < \text{target}$

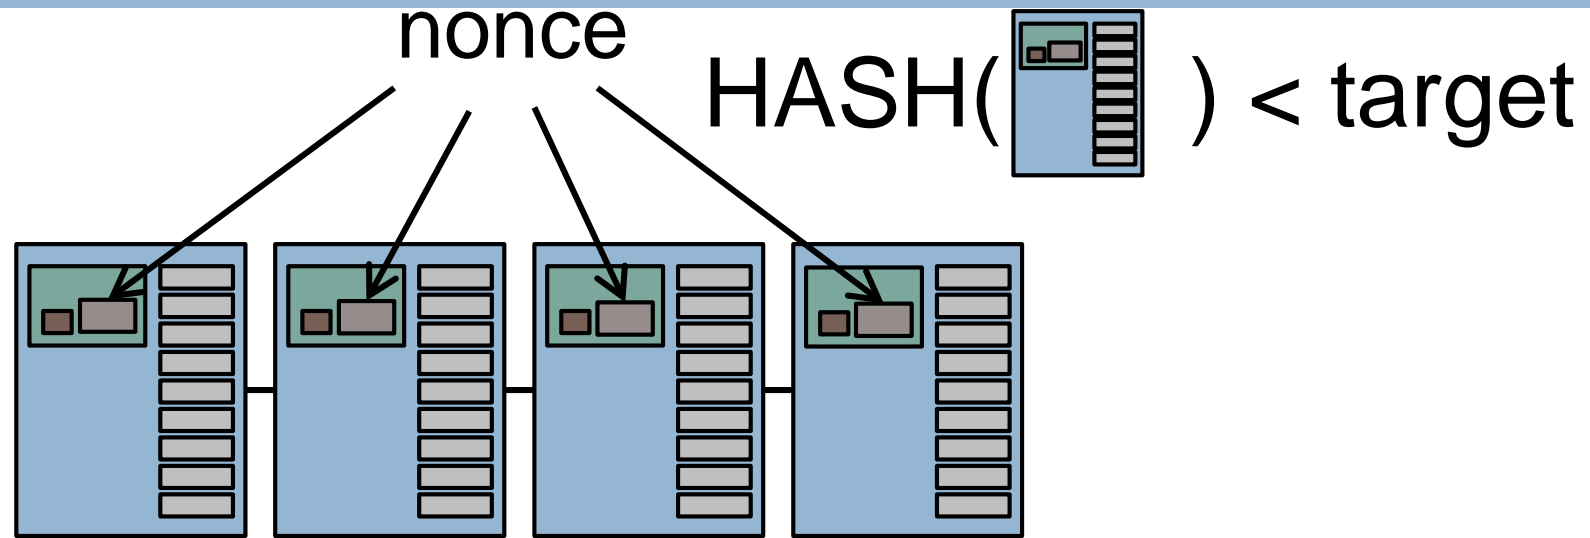
- Given  $X$  it is easy to compute  $Y$  (the *digest*)
- Given  $Y$  it is computationally infeasible to find
  - *unless you already know  $X$ , of course*
- In some sense,  $Y$  identifies  $X$

Examples: ~~MD5~~, SHA-256, SHA-3

*Note: unlike an ordinary hash function where you typically have fewer buckets than objects and thus multiple objects per bucket, with cryptographic hash functions you typically have many more “virtual buckets” than objects, and at most one object in a bucket*




# The Blockchain: Proof-of-work / Mining




- $\text{SHA256}(\text{SHA256}(\text{PrevHash} || \text{Tx} || \text{Tx} || \dots || \text{Nonce})) < \{0\}^k \{0,1\}^*$
- Mining: Find **Nonce** that when hashed with block of transactions results in  $k$  leading 0's.
- Block Identifier: Hash of block identifies the block
- Each hash identifies the entire prefix of the ledger

# The Blockchain




# The Blockchain





# The Blockchain






Exponentially distributed, with  
constant mean interval

***target*** automatically adjusted every  
2016 blocks so that mean interval is **10  
minutes**

# Incentives for Mining


Prize:

- **“Minting”**
- **Transaction Fees**




*Wins proportional to computation power*

# Forks




Two blocks “mined” at approximately the same time  
by two different miners

# Fork Resolution




- **Longest** chain wins
- Transactions on short chain are reverted

# Fork Resolution





A transaction is **confirmed** when  
it is **buried** “deep enough”  
(typically 6 blocks – i.e., one hour)

# Security Threat!







# Security Threat!




Threat: attacker outruns good miners

# Security Threat!



Threat: attacker outruns good miners

→ **Security Assumption:** *good miners own >.5 of the total compute power*






[blockchain.info, April 2015]

# Bitcoin: Network




- 1. New transactions are broadcast to all nodes.
- 2. Each node collects new transactions into a block.
- 3. Each node works on finding a difficult proof-of-work for its block.
- 4. When a node finds a proof-of-work, it broadcasts the block to all nodes.
- 5. Nodes accept the block only if all transactions in it are valid and not already spent.
- 6. Nodes express their acceptance of the block by working on creating the next block in the chain, using the hash of the accepted block as the previous hash.
  
- ***Nodes always consider the longest chain to be the correct one and will keep working on extending it.***

# Bitcoin: Network




TX1 = [Alice → Bob]


TX2 = [Alice → Carol]



# Bitcoin

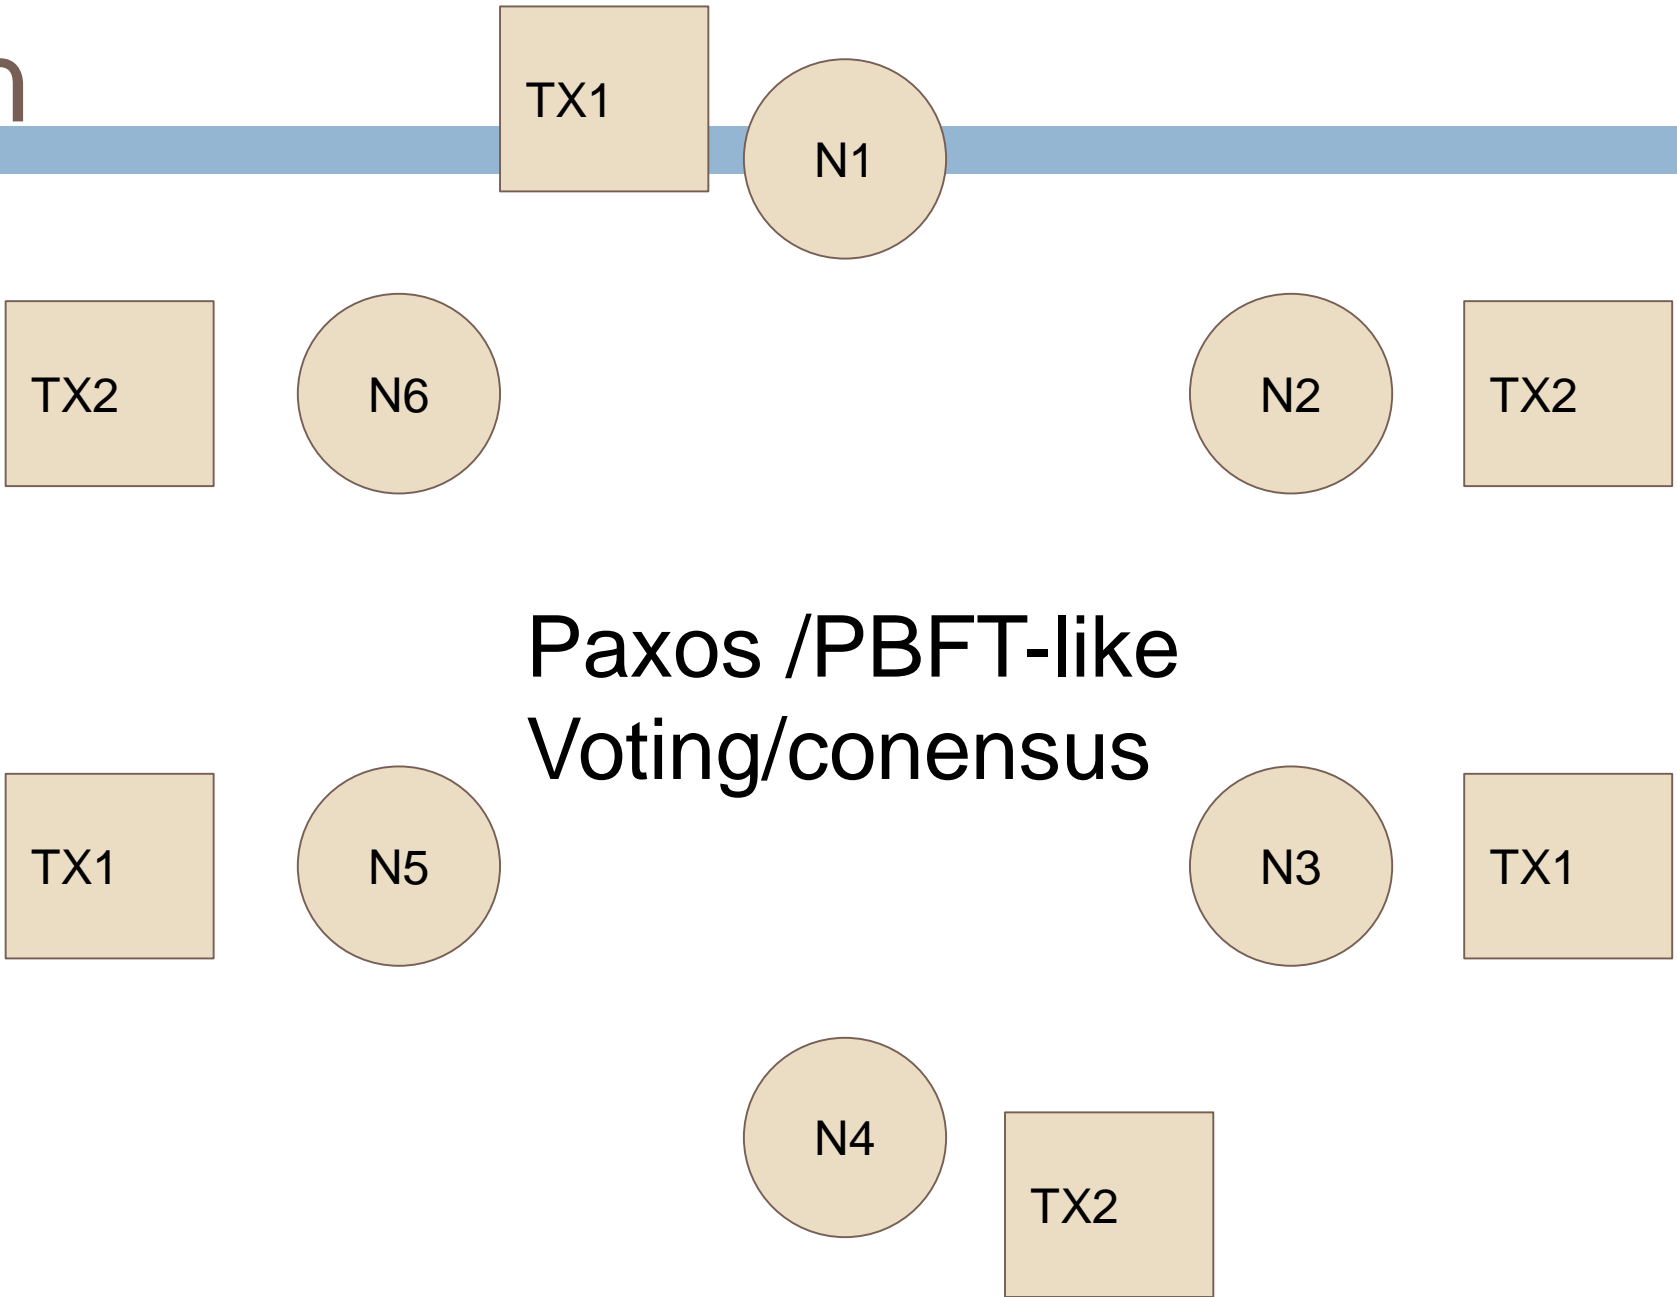


# Bitcoin



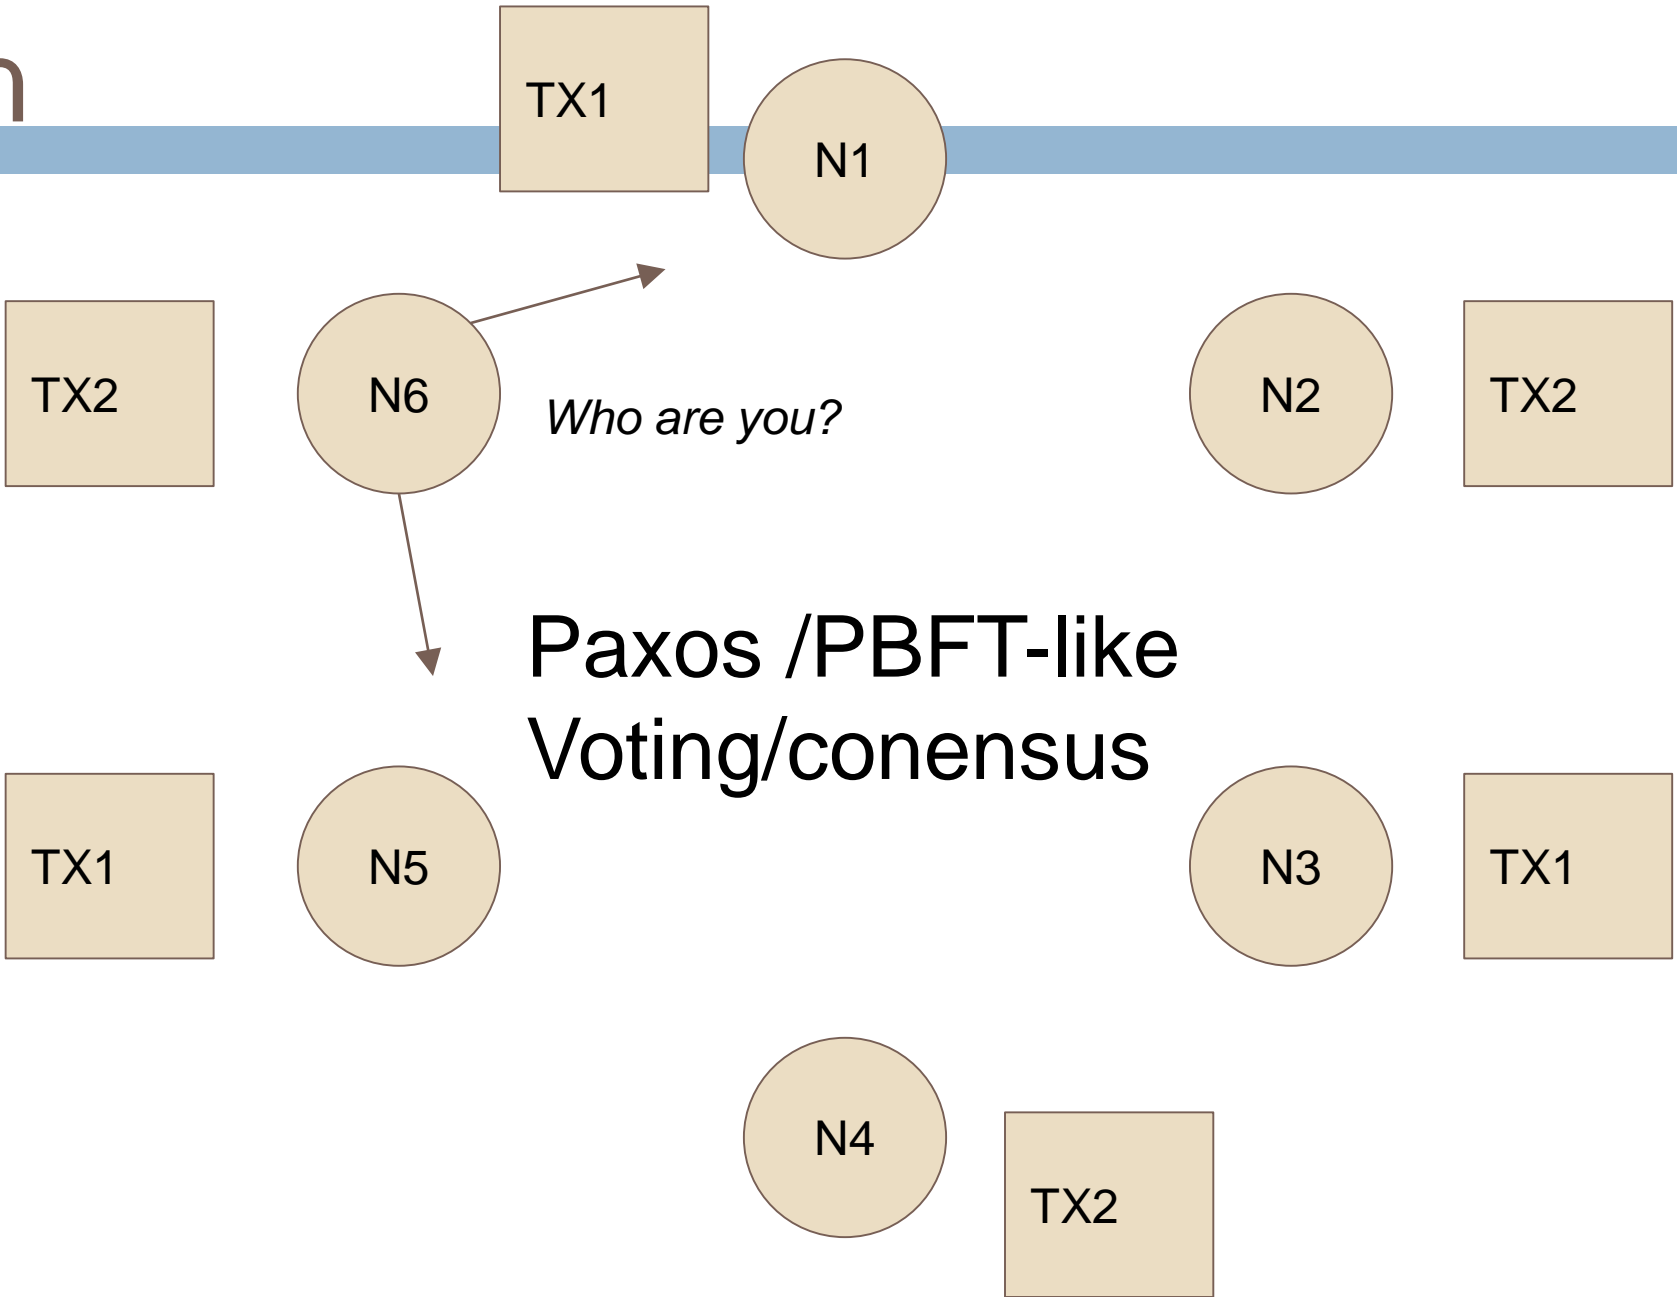


# Bitcoin

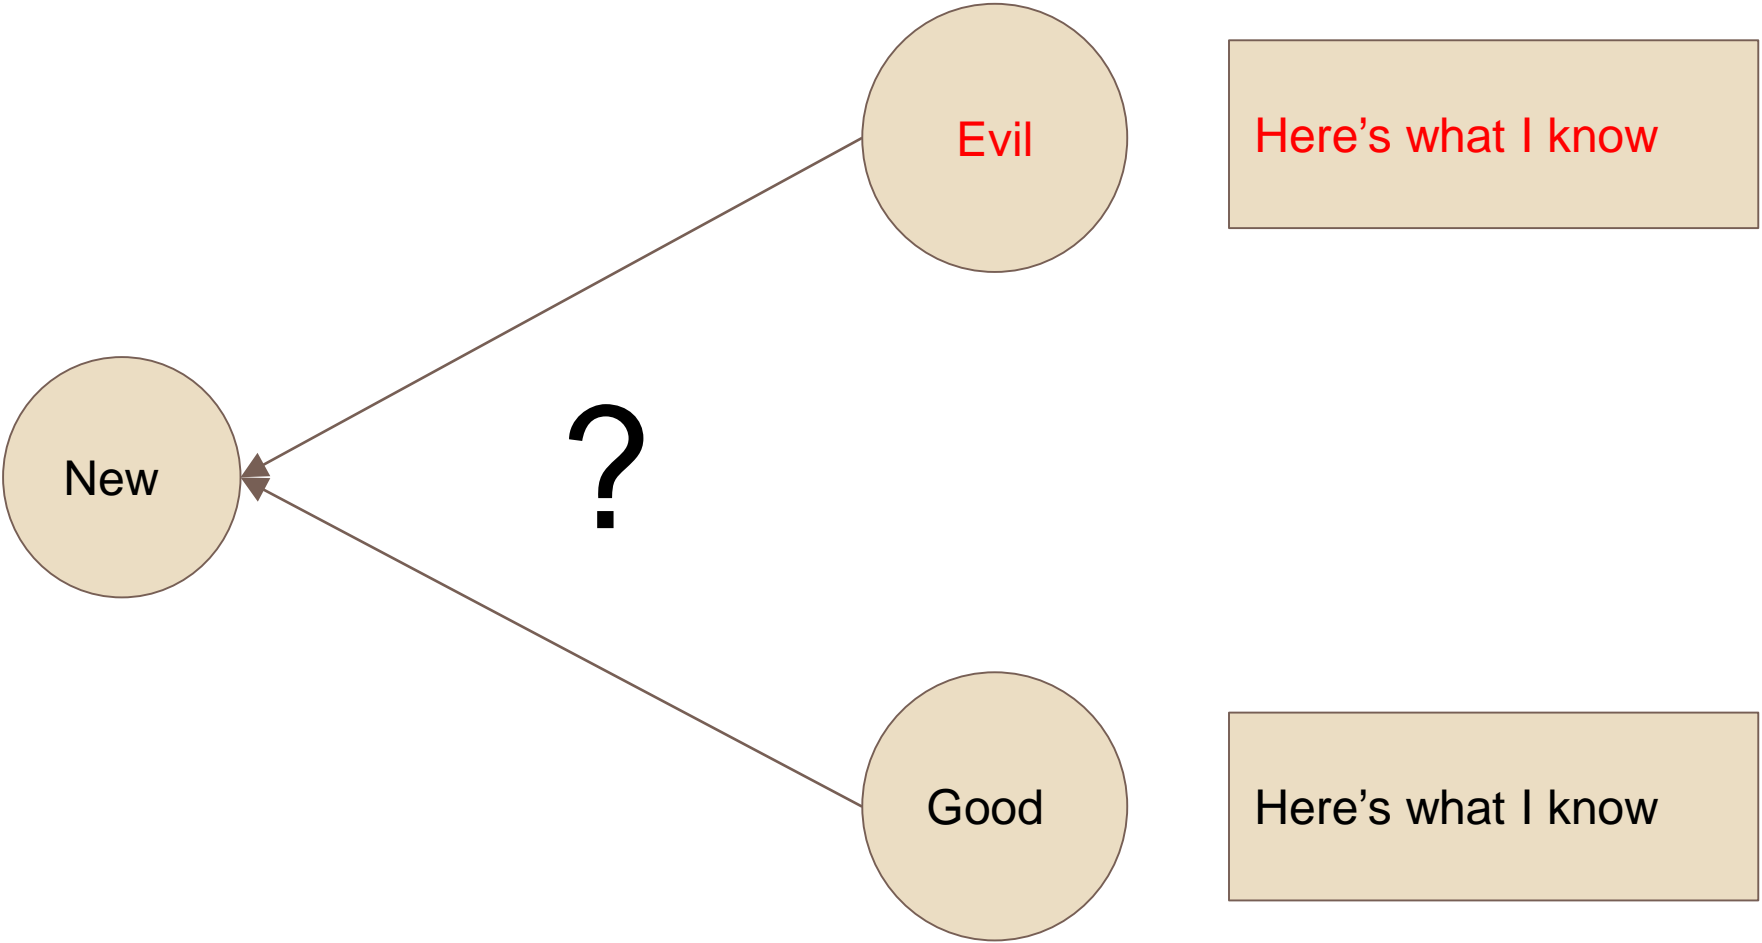


Paxos /PBFT-like  
Voting/conensus

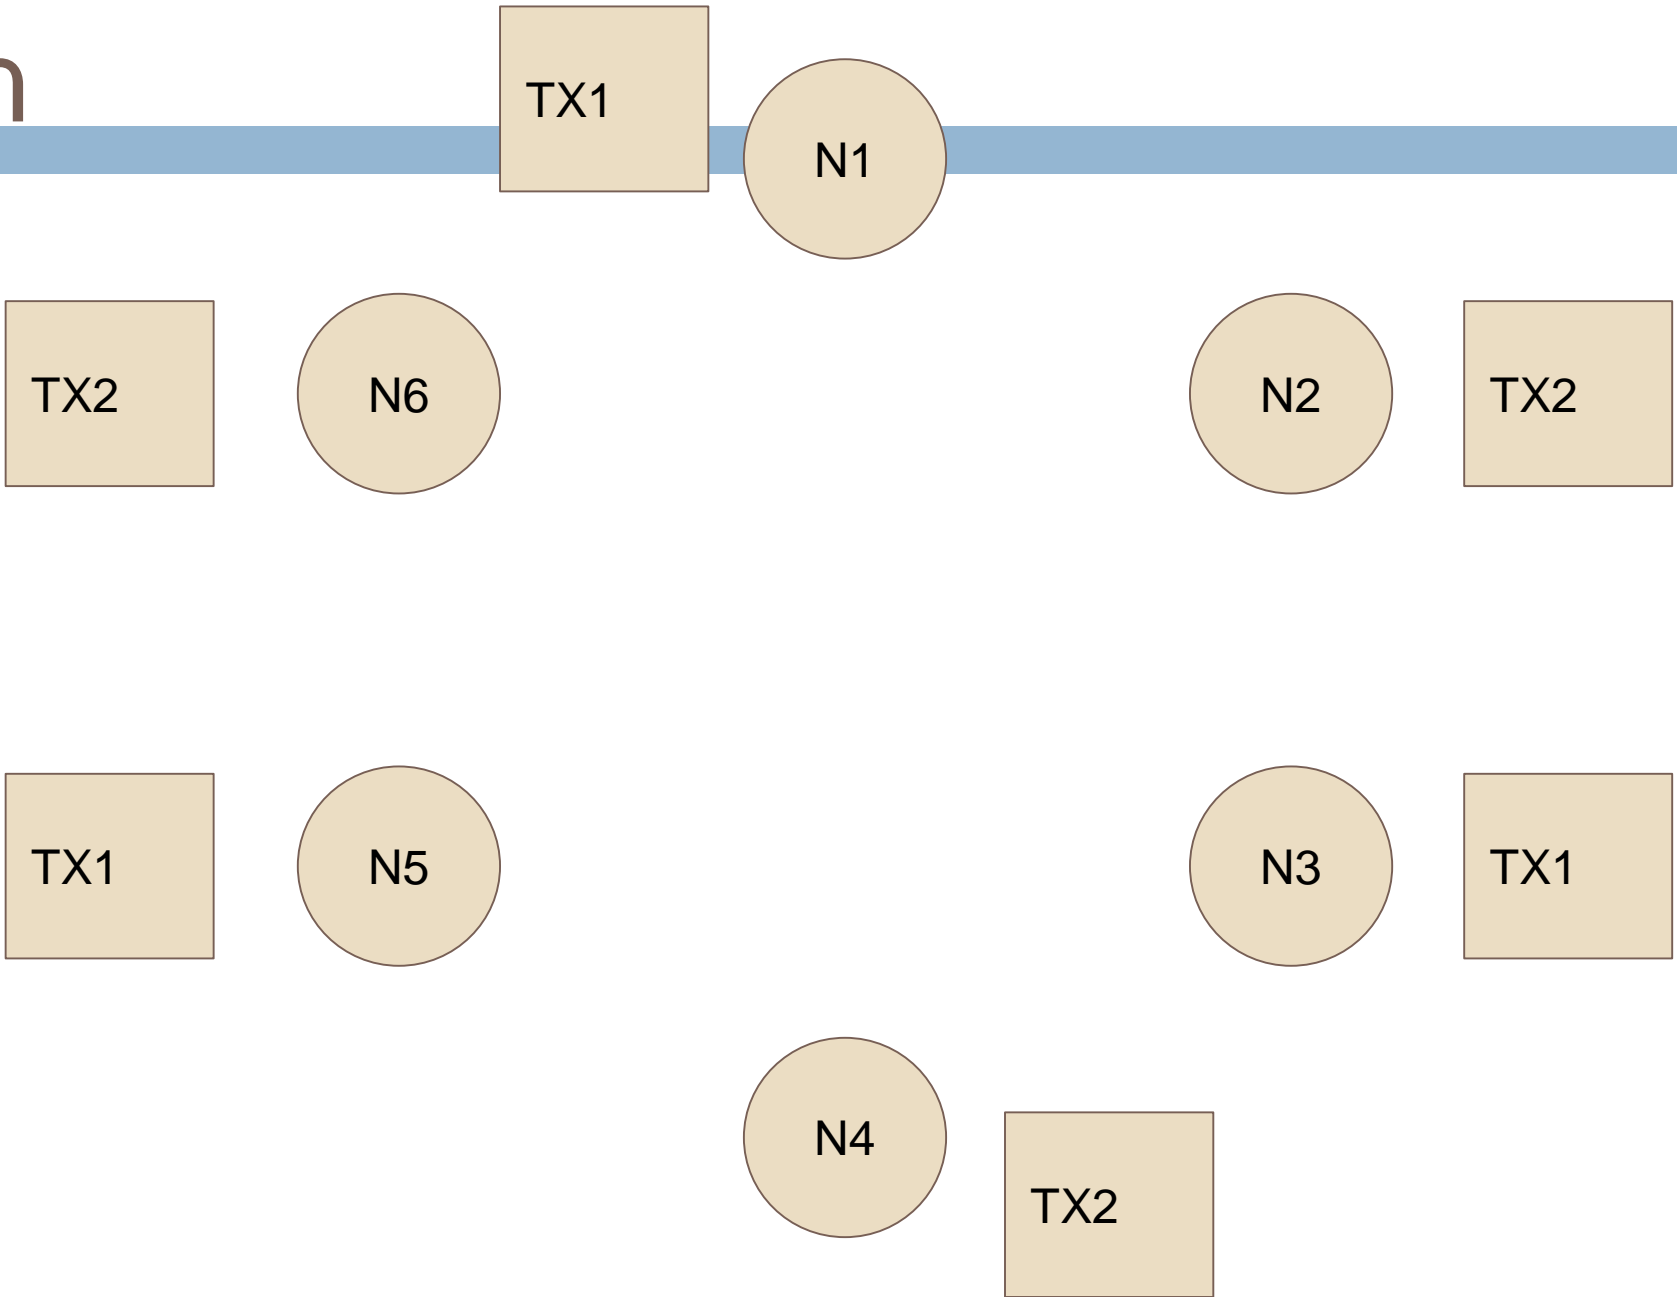
# Bitcoin




# Bitcoin



# Bitcoin



# Bitcoin



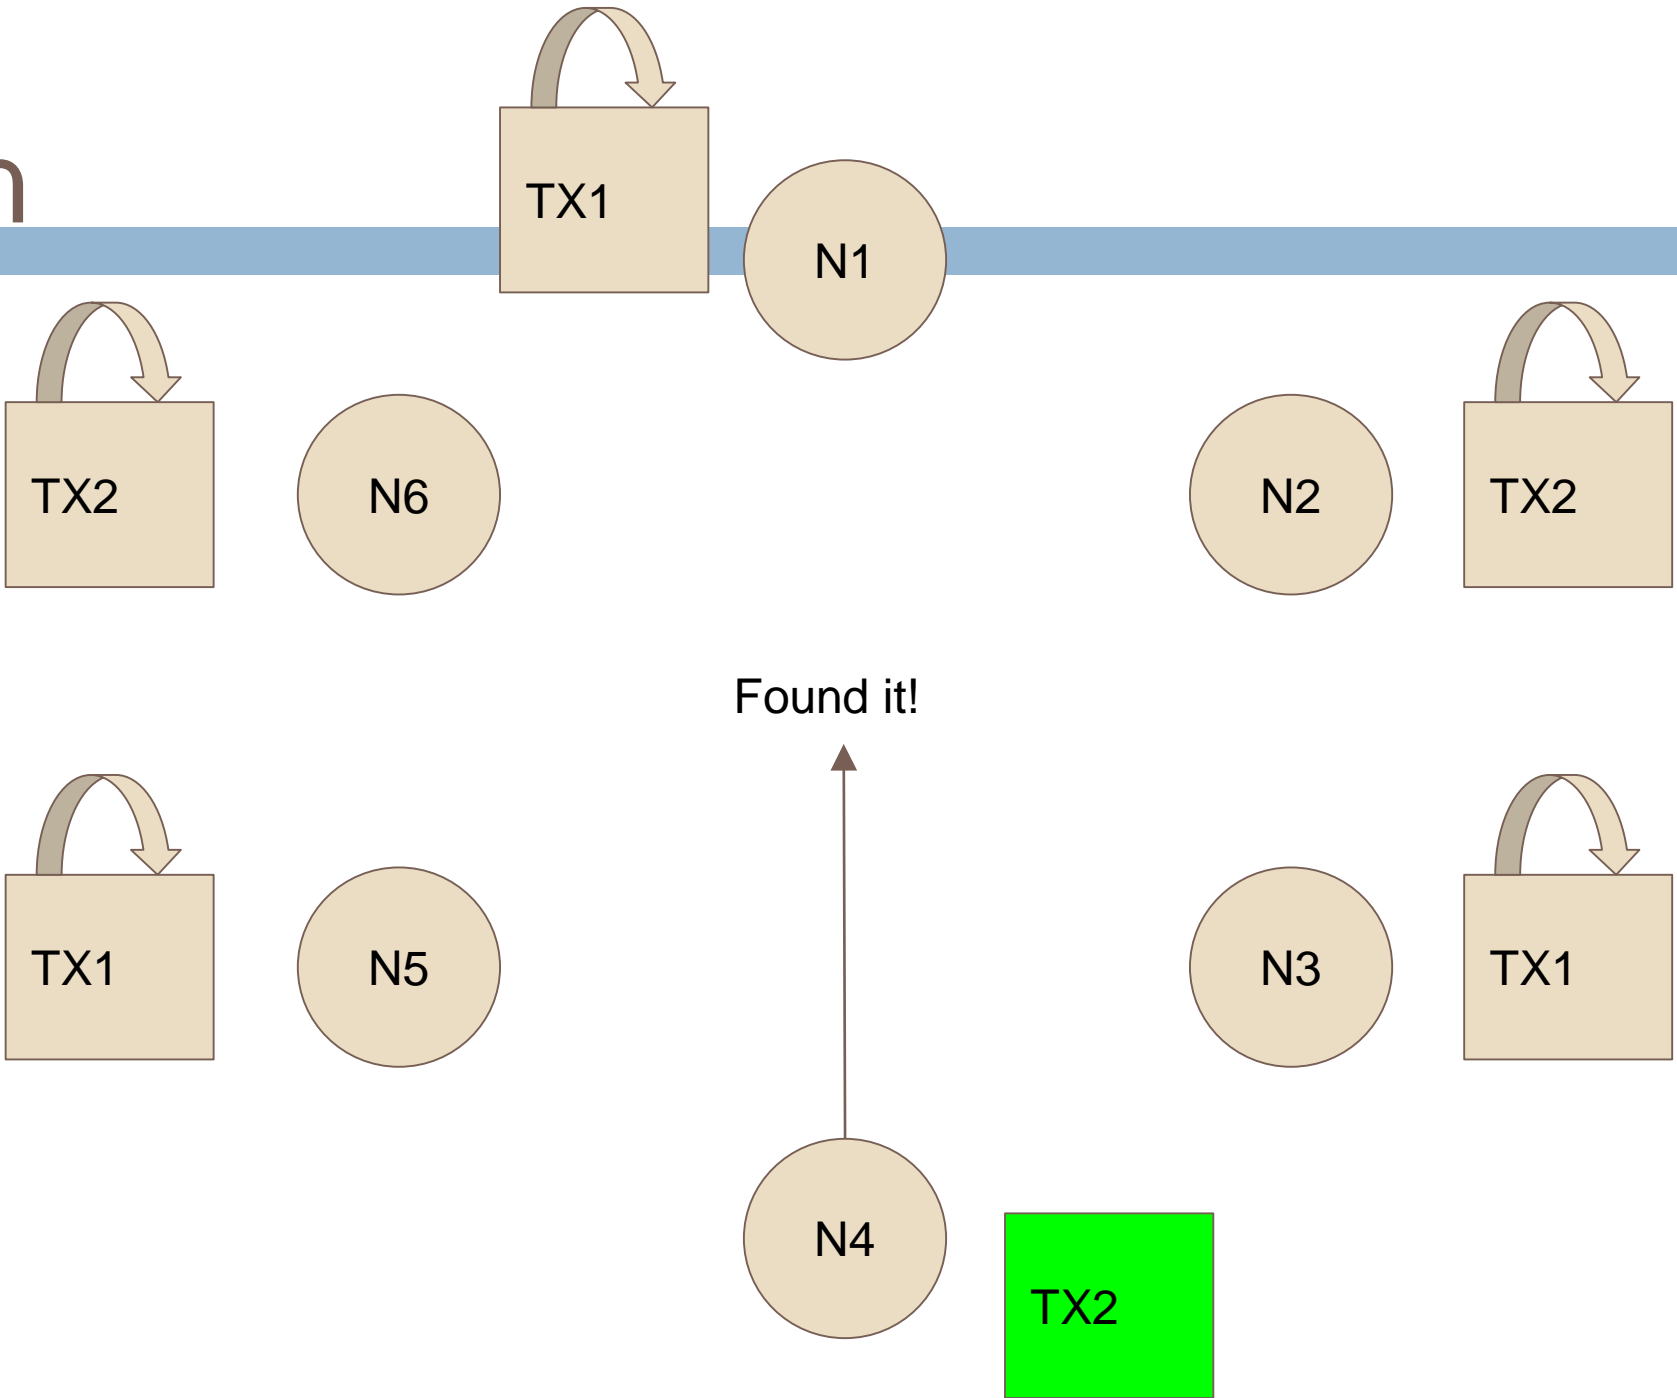
# Bitcoin

---

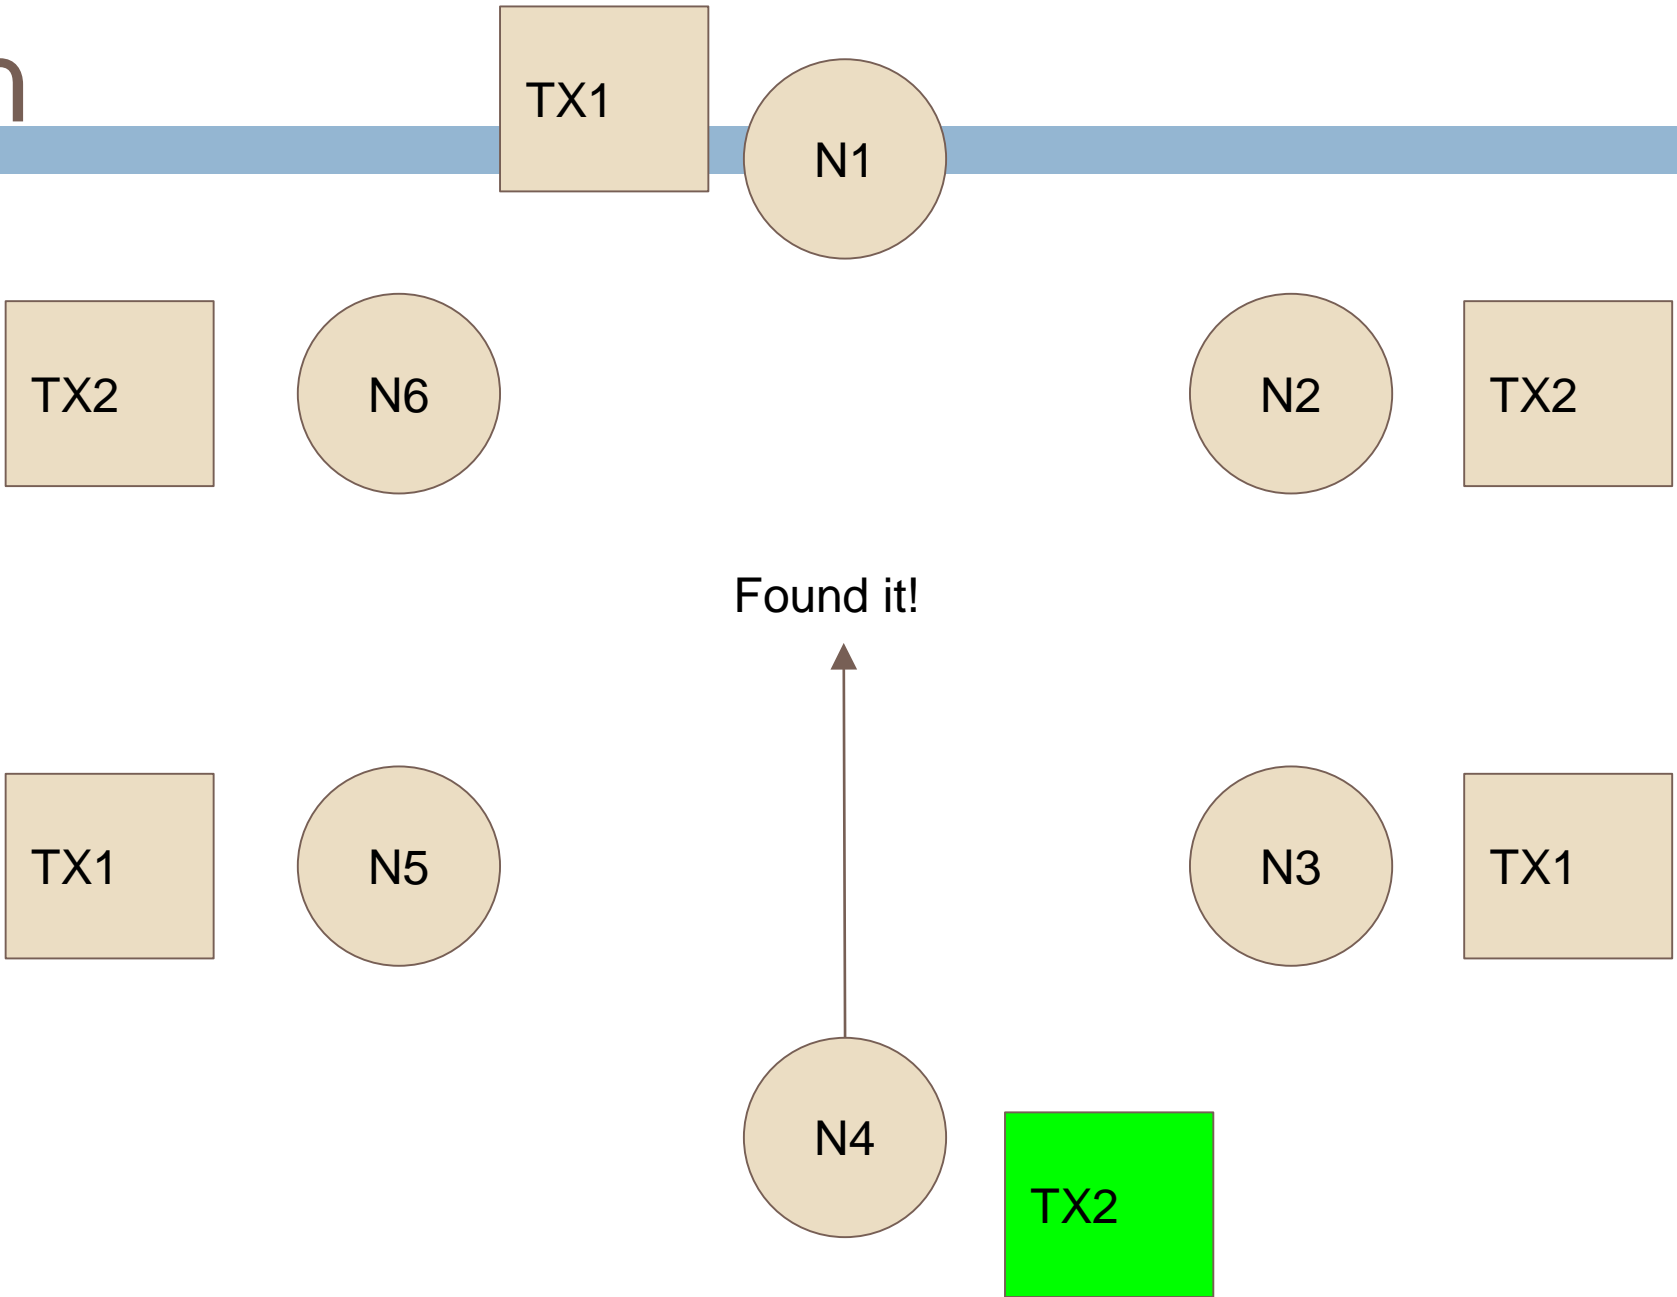
$$\text{SHA256}(\text{SHA256}(\text{TX} \parallel \text{Nonce})) < \{0\}^k \{0,1\}^*$$



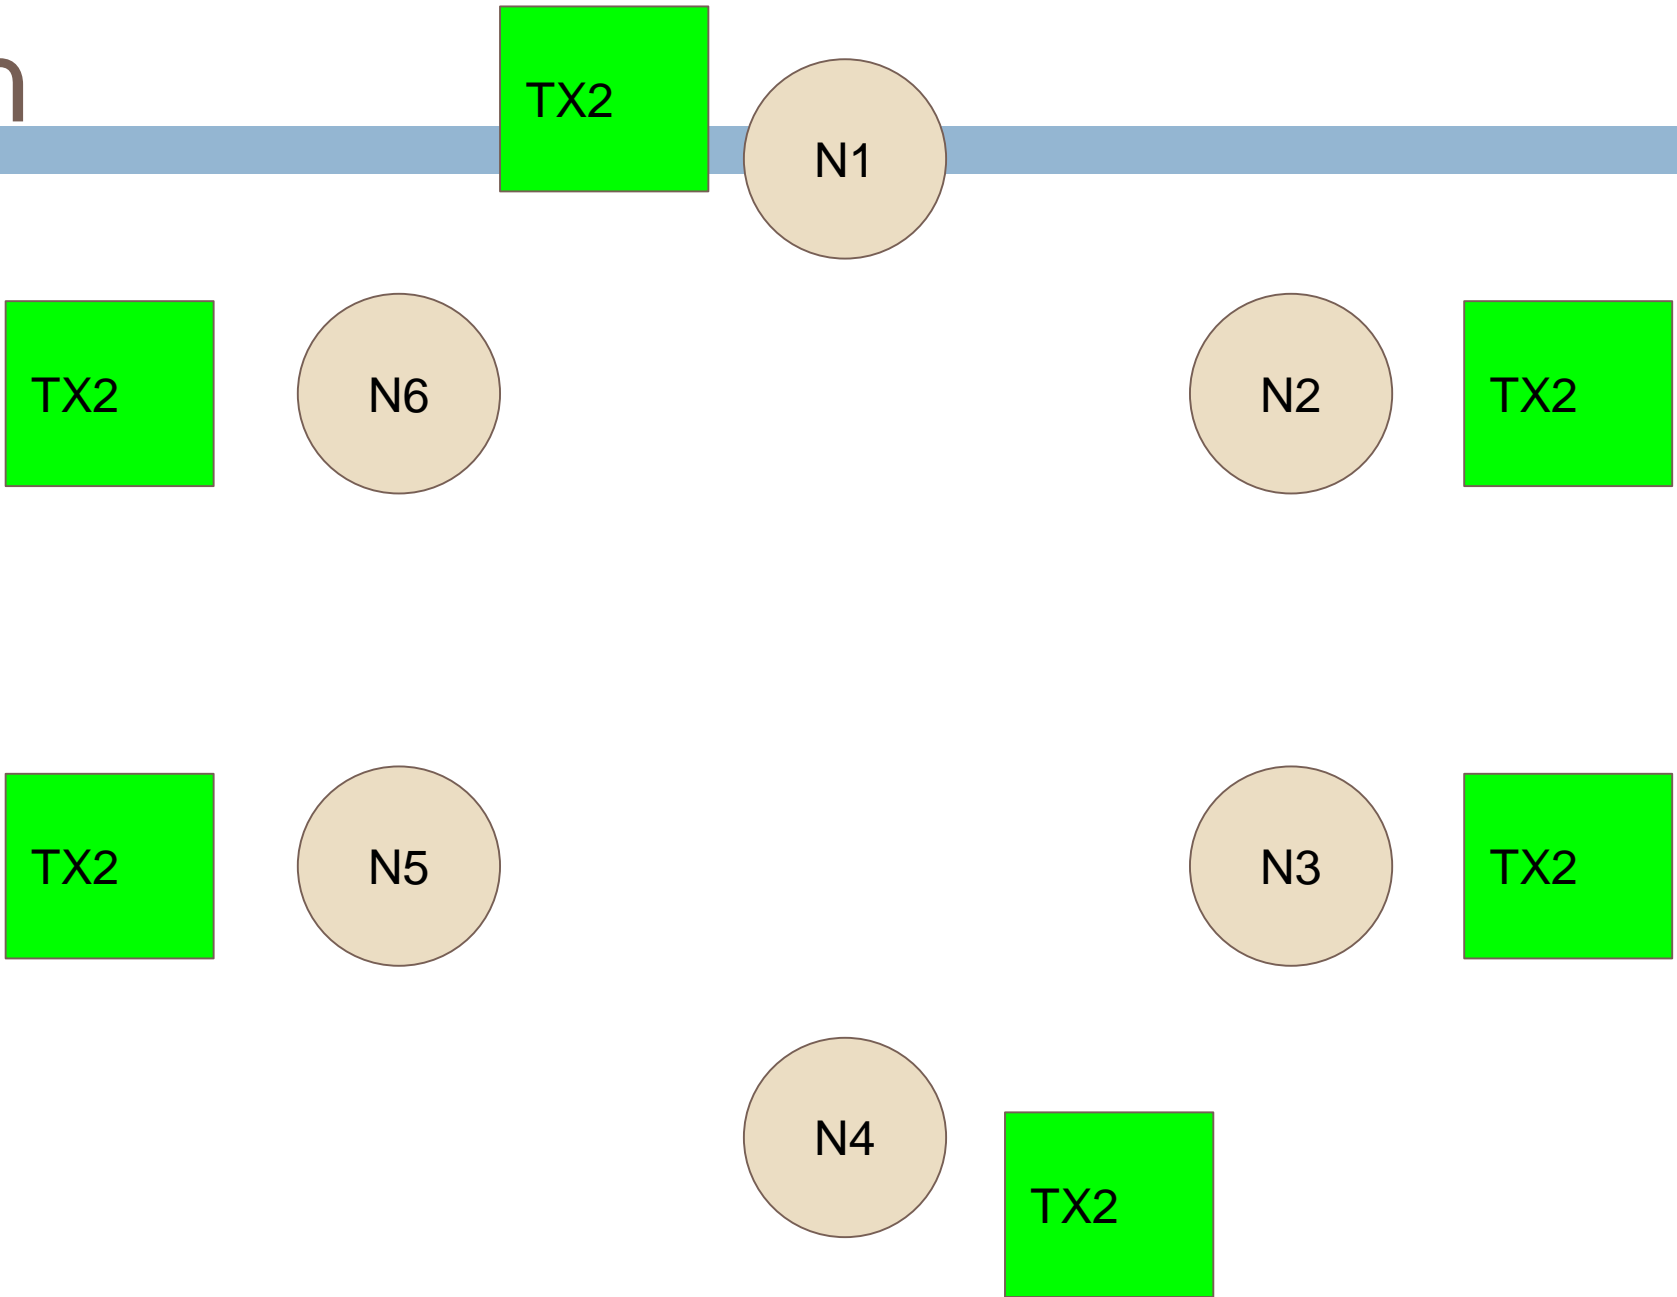
# Bitcoin



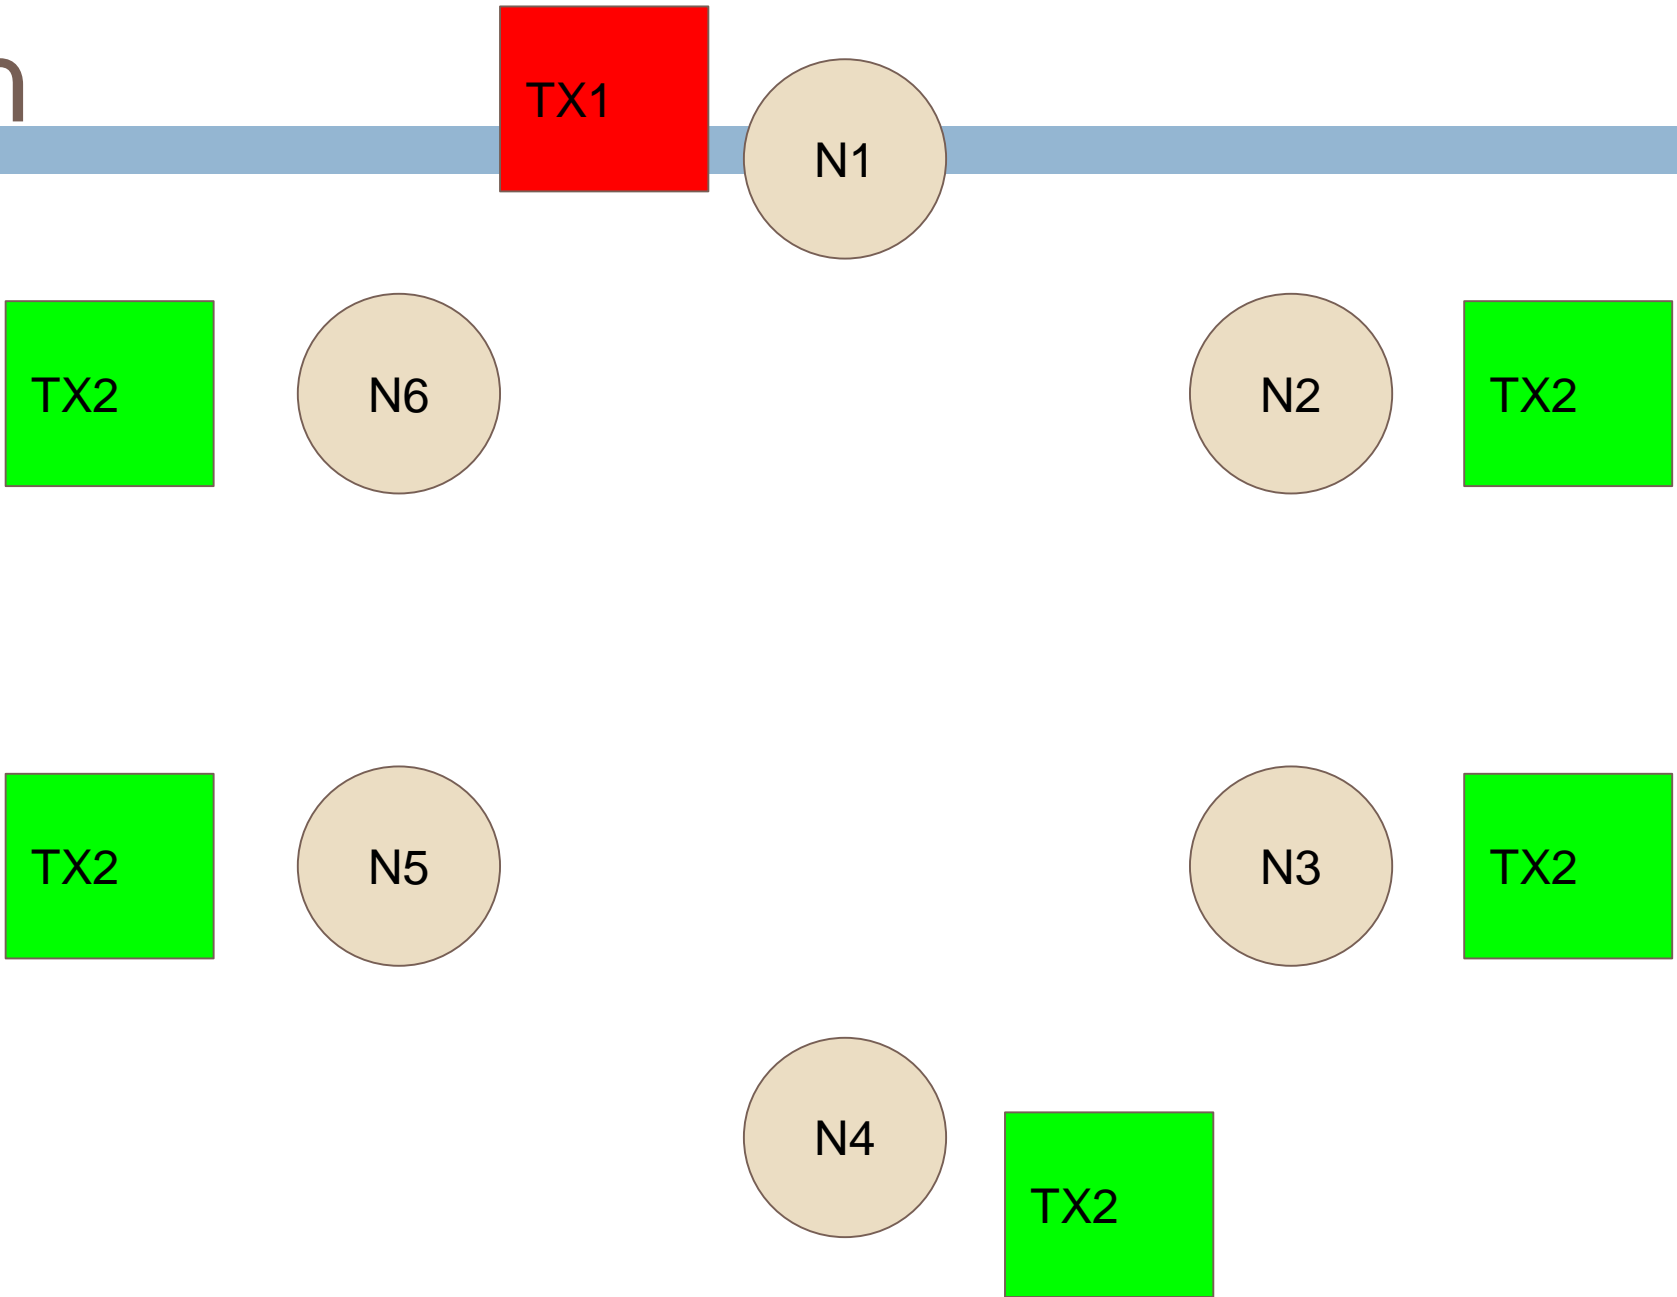
# Bitcoin




# Bitcoin



# Bitcoin

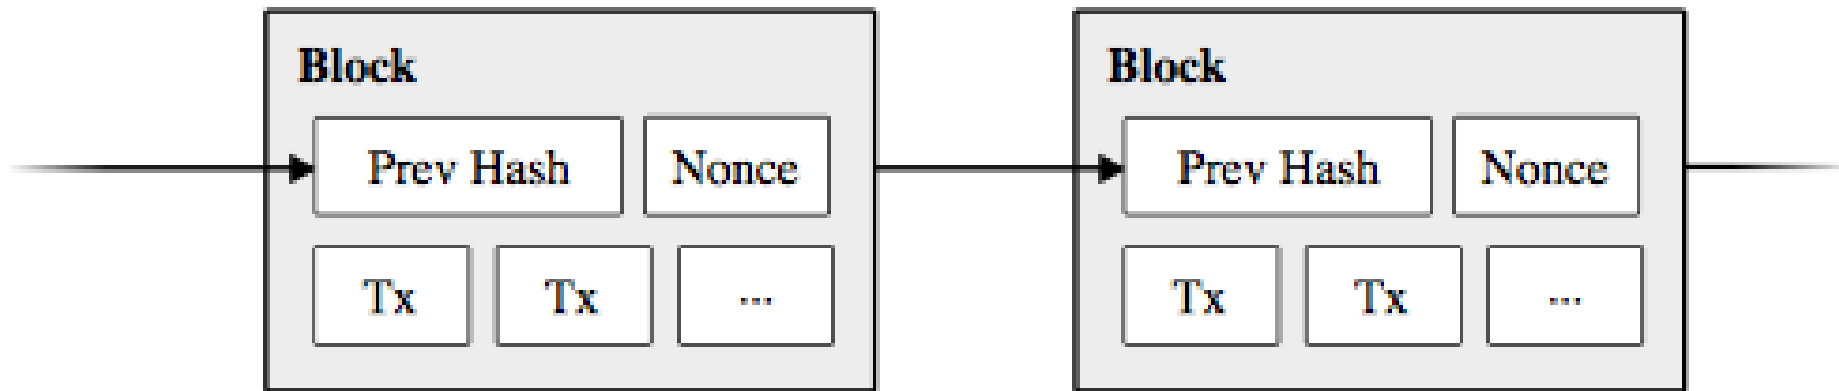


# Bitcoin




I'm confused


# Bitcoin



# Bitcoin

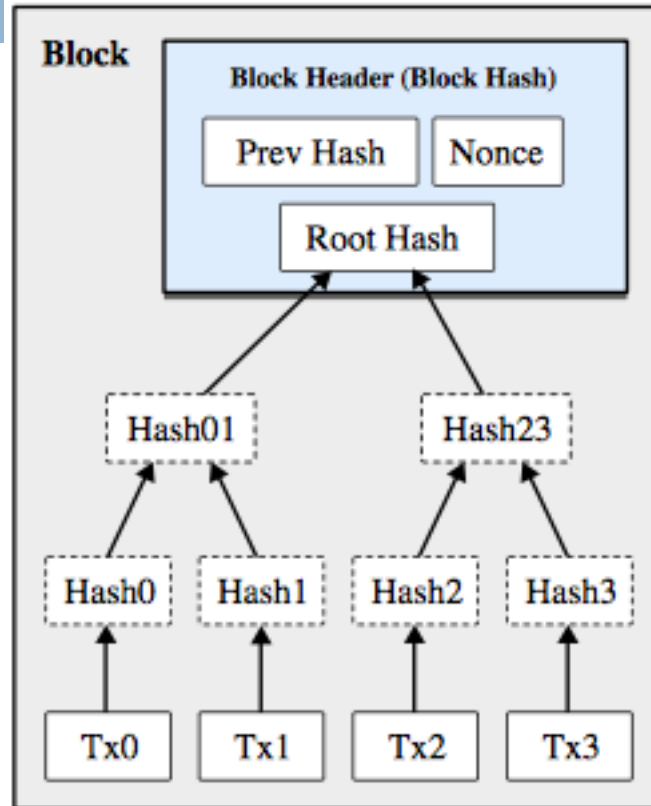


I'm no longer confused






# Bitcoin




Transactions Hashed in a Merkle Tree



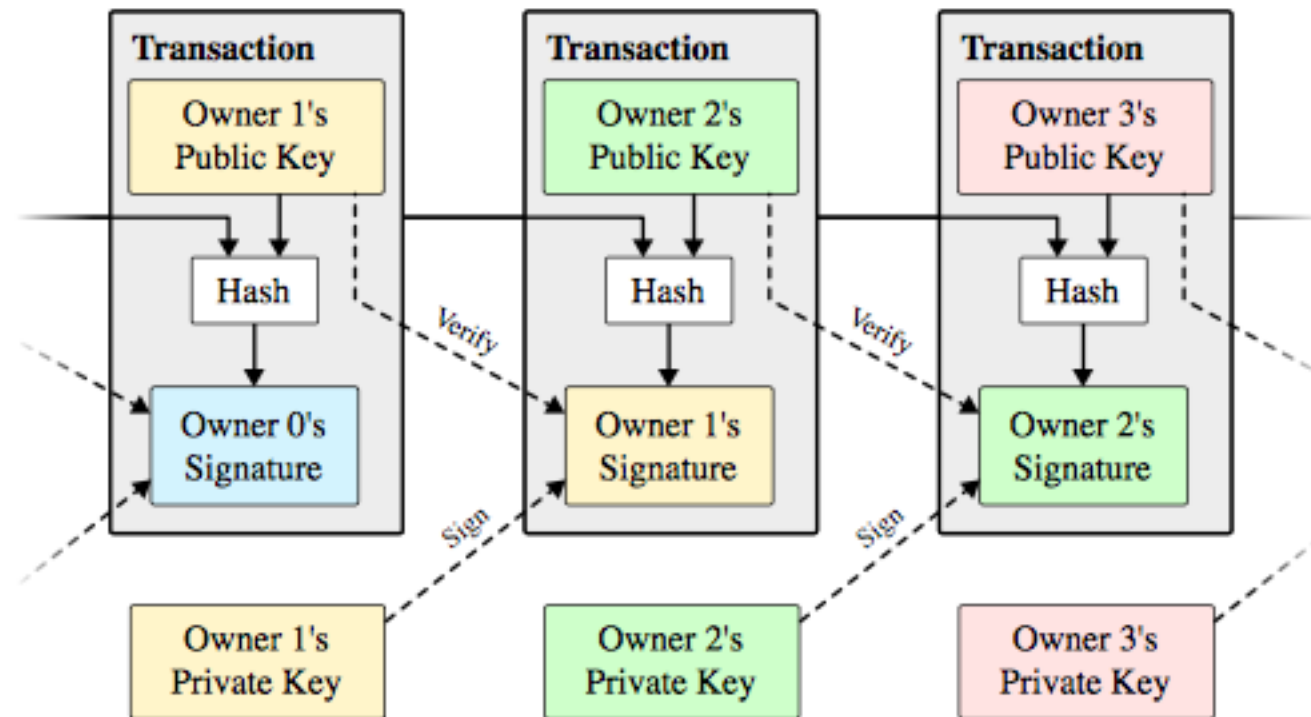
After Pruning Tx0-2 from the Block

# Bitcoin

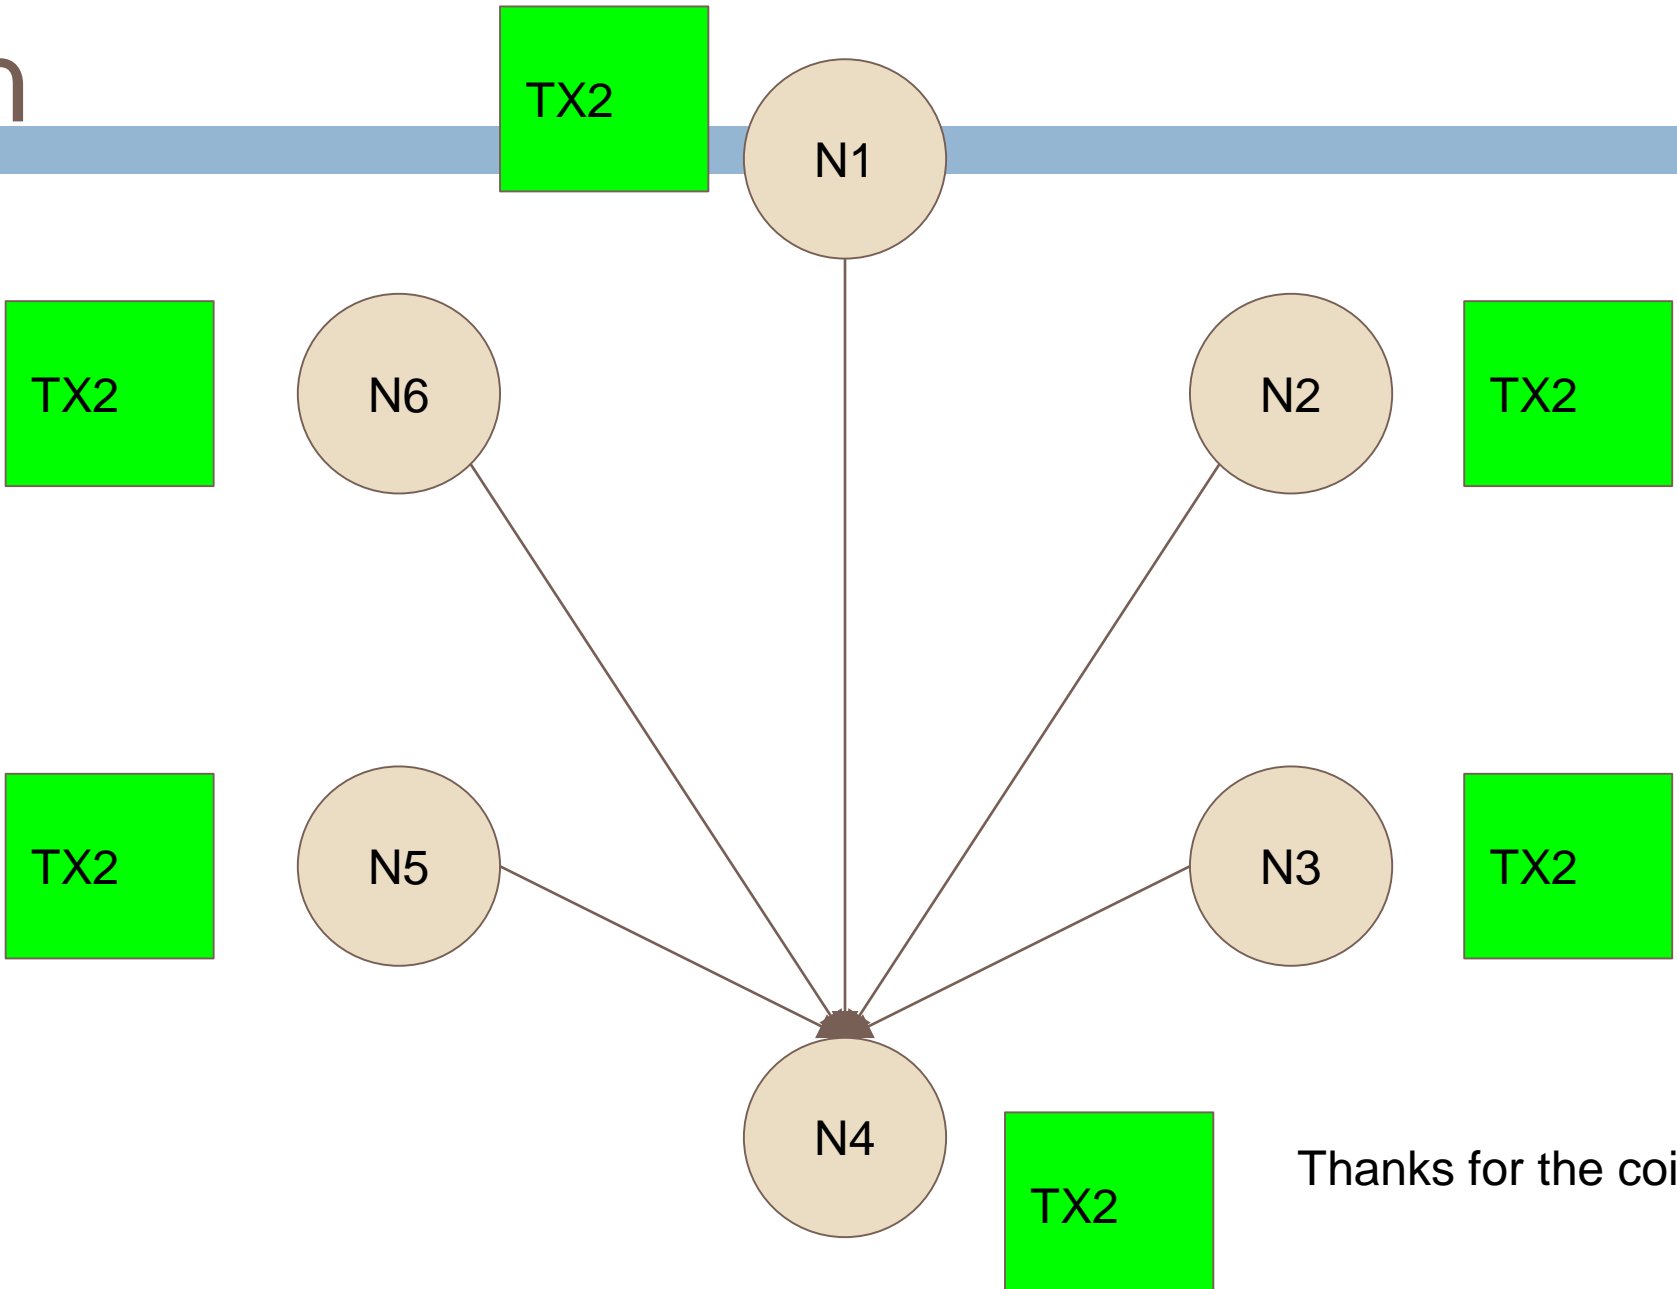


# Bitcoin

## UTXO



# Bitcoin



Thanks for the coins!

# Bitcoin

- Many more subtle details, but core mechanism is computational race.
- Results:
  - Breakthrough consensus mechanism in the permissionless setting
    - Everybody *agrees* on what is on the blockchain
    - Always available for reading and appending
    - Fair
    - Tamperproof (can't change or truncate blockchain)
    - No Single Administrative Domain
    - Open membership
- Challenges:
  - Electrical usage of a medium-sized country.
  - Very slow confirmation times.
  - 3 tx/second.



Power Consumption...

# Bitcoin

- Three years ago, when Kevin Sekniqi gave talk in December 2016, 1 BTC == \$1100  
Today at \$10,261
- Crypto market cap high of \$180B last week, today \$126B
- 21 million total possible Bitcoins. As supplies dwindle, price skyrockets.

# Bitcoin: Tradeoffs

	Permissionless	Permissioned
Approach	Competitive	Cooperative
Basic technique	Proof-of-Resource	Byzantine Consensus
Trust requirements	Crypto (+ peers...)	Peers (+ crypto...)
Membership	Open	Closed
Energy-efficiency	Often terrible	Excellent
Transaction rate	At best hundreds / sec	Many thousands per second
Txn latency	As high as many minutes	Less than a second

# Majority is Not Enough: Bitcoin Mining is Vulnerable



## Ittay Eyal



- Postdoc @ Cornell, faculty @ Technion
- Major contributor to Bitcoin community





## Emin Gun Sirer

- Also major contributor to the Bitcoin community





49%



51%

# Majority is Not Enough: Bitcoin Mining is Vulnerable

---

## Algorithm 1: Selfish-Mine

---


```
1 on Init
2   public chain  $\leftarrow$  publicly known blocks
3   private chain  $\leftarrow$  publicly known blocks
4   privateBranchLen  $\leftarrow$  0
5   Mine at the head of the private chain.

6 on My pool found a block
7    $\Delta_{prev} \leftarrow$  length(private chain)  $-$  length(public chain)
8   append new block to private chain
9   privateBranchLen  $\leftarrow$  privateBranchLen + 1
10  if  $\Delta_{prev} = 0$  and privateBranchLen = 2 then           (Was tie with branch of 1)
11    publish all of the private chain                       (Pool wins due to the lead of 1)
12    privateBranchLen  $\leftarrow$  0
13  Mine at the new head of the private chain.


14 on Others found a block
15    $\Delta_{prev} \leftarrow$  length(private chain)  $-$  length(public chain)
16   append new block to public chain
17   if  $\Delta_{prev} = 0$  then
18     private chain  $\leftarrow$  public chain                   (they win)
19     privateBranchLen  $\leftarrow$  0
20   else if  $\Delta_{prev} = 1$  then
21     publish last block of the private chain               (Now same length. Try our luck)
22   else if  $\Delta_{prev} = 2$  then
23     publish all of the private chain                       (Pool wins due to the lead of 1)
24     privateBranchLen  $\leftarrow$  0
25   else                                                     ( $\Delta_{prev} > 2$ )
26     publish first unpublished block in private block.
27   Mine at the head of the private chain.
```

---


# Majority is Not Enough: Bitcoin Mining is Vulnerable




# Majority is Not Enough: Bitcoin Mining is Vulnerable




# Majority is Not Enough: Bitcoin Mining is Vulnerable




# Majority is Not Enough: Bitcoin Mining is Vulnerable



# Majority is Not Enough: Bitcoin Mining is Vulnerable




# Majority is Not Enough: Bitcoin Mining is Vulnerable







# Majority is Not Enough: Bitcoin Mining is Vulnerable



# Majority is Not Enough: Bitcoin Mining is Vulnerable



# Majority is Not Enough: Bitcoin Mining is Vulnerable





# Perspective

- Exciting application possibilities
- Interesting research problems in distributed systems and practical cryptography
- Bitcoin “works”, but is extremely inefficient

	Permissionless	Permissioned
Approach	Competitive	Cooperative
Basic technique	Proof-of-Resource	Byzantine Consensus
Trust requirements	Crypto (+ peers...)	Peers (+ crypto...)
Membership	Open	Closed ( <i>but stay tuned</i> )
Energy-efficiency	Often terrible	Excellent
Transaction rate	At best hundreds / sec	Many thousands per second
Txn latency	As high as many minutes	Less than a second

# Next Time

- Read and write review for Thursday, September 6:
  - ▣ **Required** The Design and Implementation of a Log-Structured File System, Mendel Rosenblum and Ousterhout. Proceedings of the thirteenth ACM symposium on Operating systems principles, October 1991, pages 1--15. On the duality of operating system structures, H. C. Lauer and R. M. Needham. ACM SIGOPS Operating Systems Review Volume 12, Issue 2 (April 1979), pages 3--19.
  - ▣ **Optional:** A Fast File System for UNIX. Marshall K. McKusick, William N. Joy, Samuel J. Leffler, Robert S. Fabry. ACM TOCS 2(3), Aug 1984, pages 181-197.

# Next Time

- Read and write review:
- MP1.a – due next Monday, September 10
- Project Proposal due next Tuesday, September 11
  - ▣ talk to me and other faculty and email and talk to me
- Check website for updated schedule