

Blockchain

IF306 2018-2019


Bitcoin

- Classic 2008 Nakamoto paper « Bitcoin: A Peer-to-Peer Electronic Cash System »
- Motivation: replace credit card for internet payments


What is needed is an electronic payment system based on cryptographic proof instead of trust, allowing any two willing parties to transact directly with each other without the need for a trusted third party. Transactions that are computationally impractical to reverse would protect sellers from fraud, and routine escrow mechanisms could easily be implemented to protect buyers. In

- No trust, no central authority
- Irreversible transactions : seller point of view

Bitcoin today



Bitcoin Energy Consumption Relative to Several Countries



Bitcoin network versus VISA network average consumption

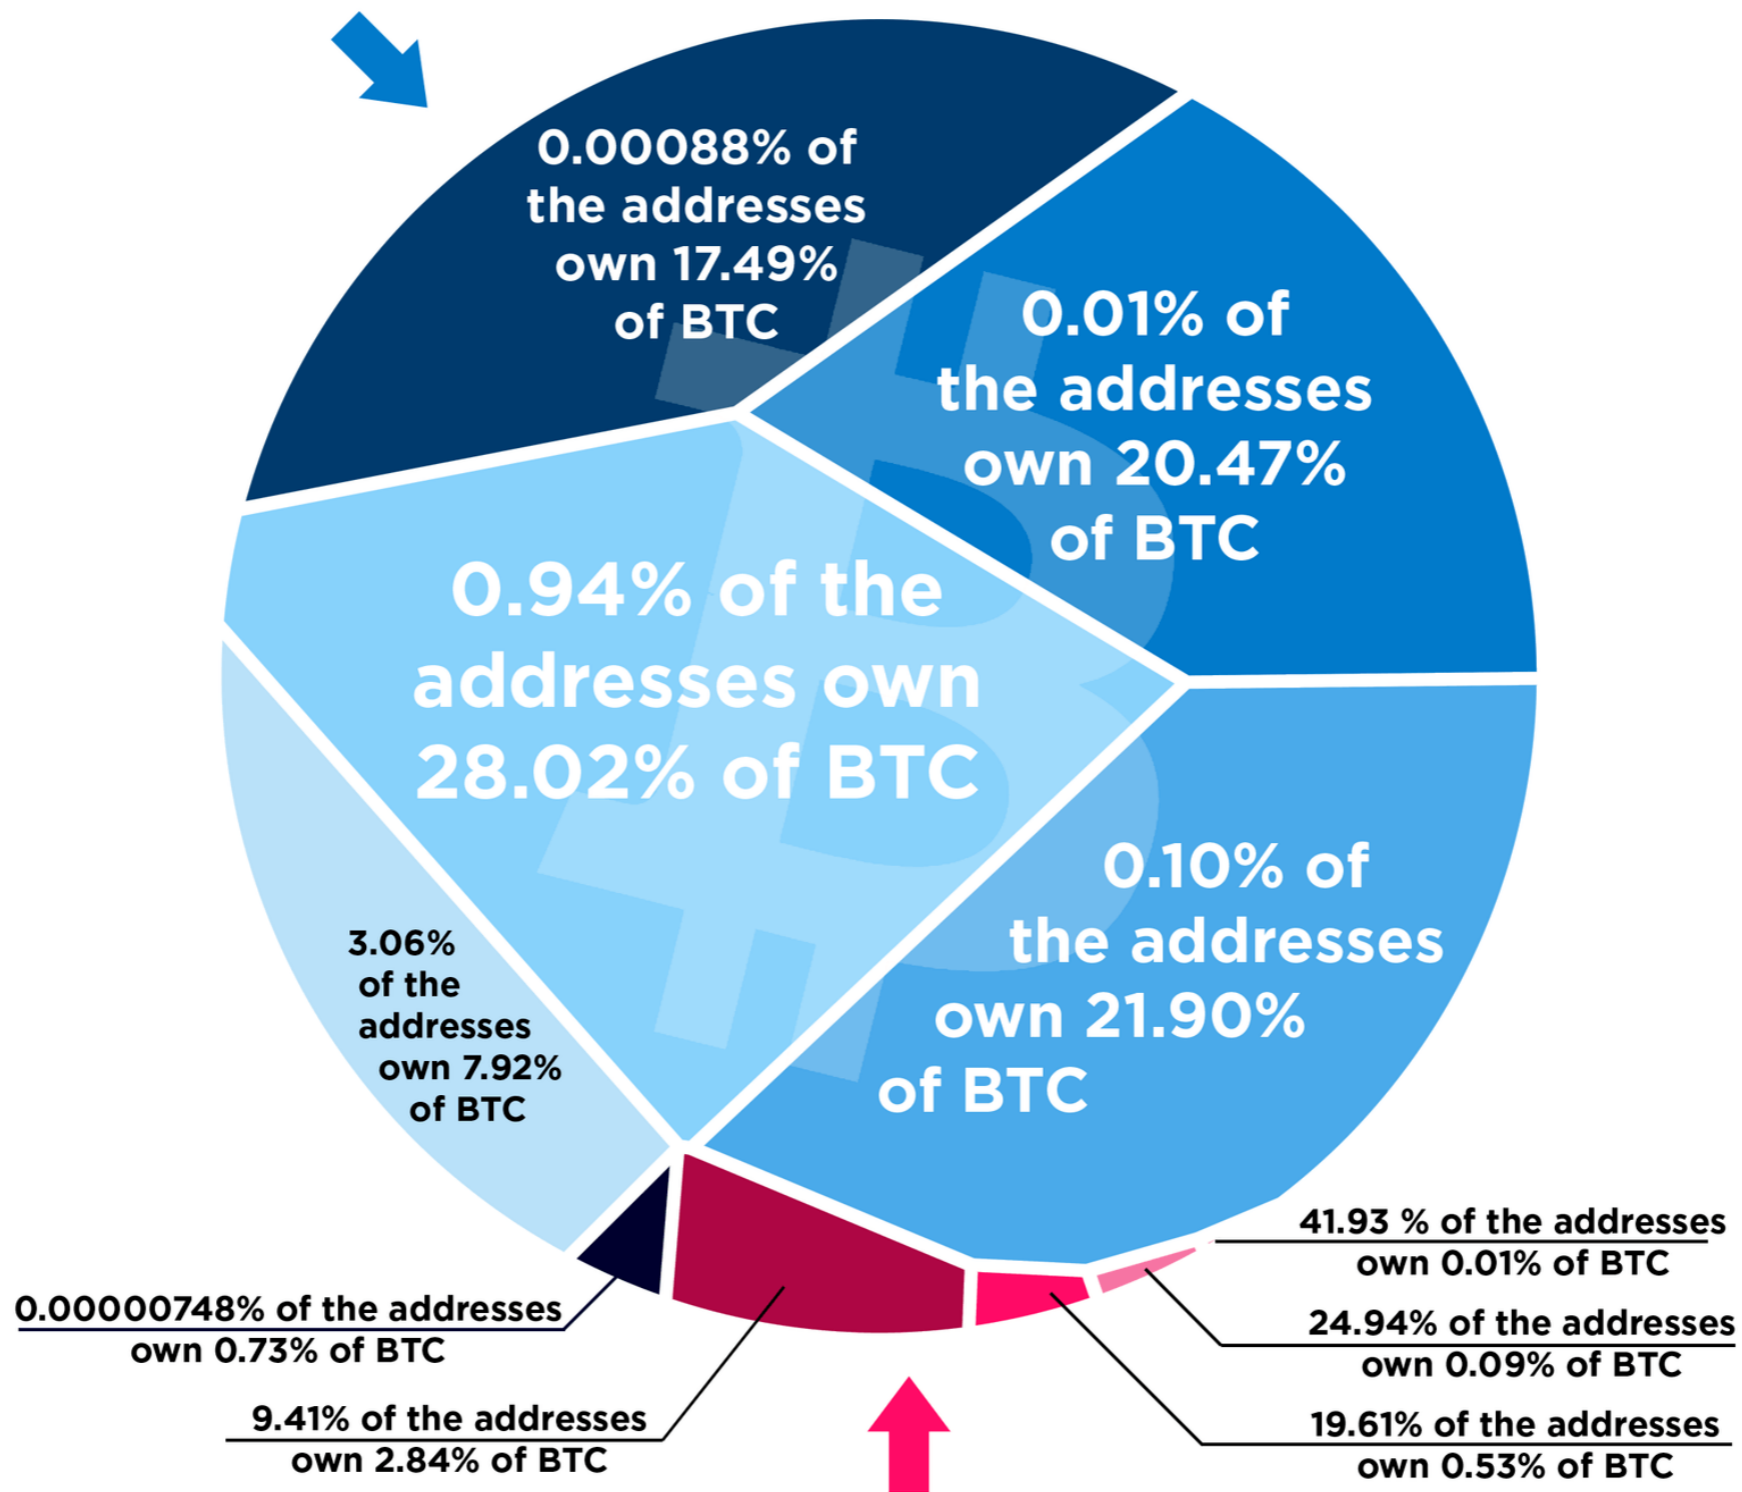


Deflation ?

- BTC supply is bounded (~21 millions)
- Krugman's co-op baby-sitting story:
 - 500 coupons 1h babysitting
 - Soon people were preferring to save rather to spend

The bitcoin Wealth Distribution

4.11% OF ADDRESSES OWN 96.53% OF BTC*




95.89% OF ADDRESSES OWN 3.47% OF BTC*

Scandals



the DAO hack

 The DAO

CREATION



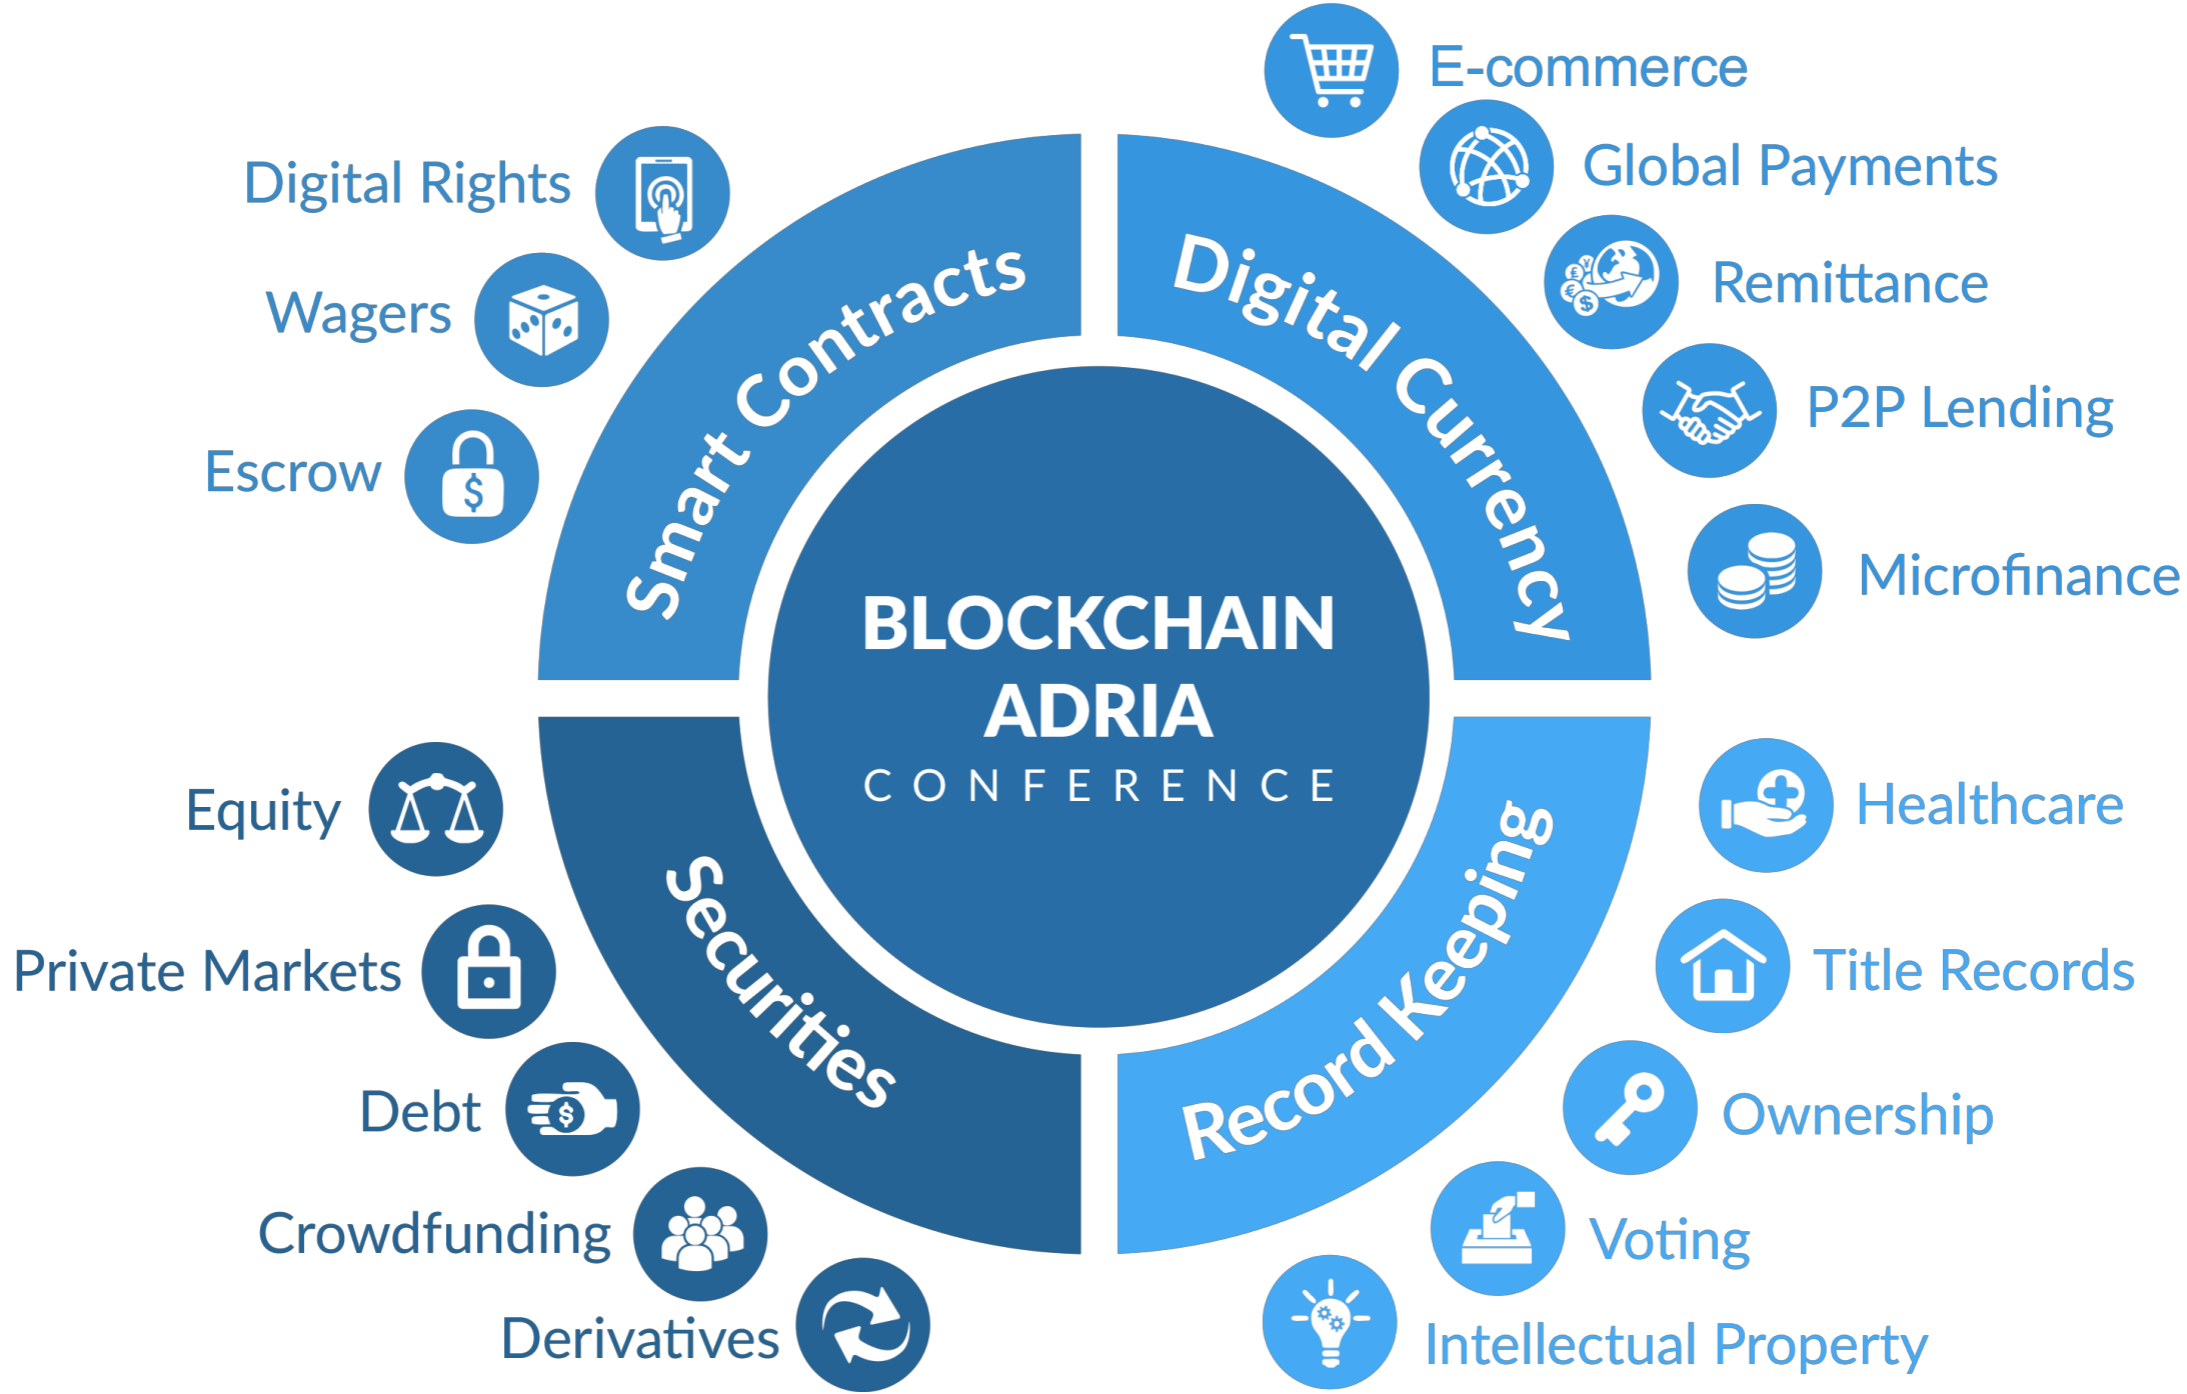
**The DAO Attacked: Code Issue Leads to
\$60 Million Ether Theft**

**Still every current blockchain technology
originated from Nakamoto's paper**

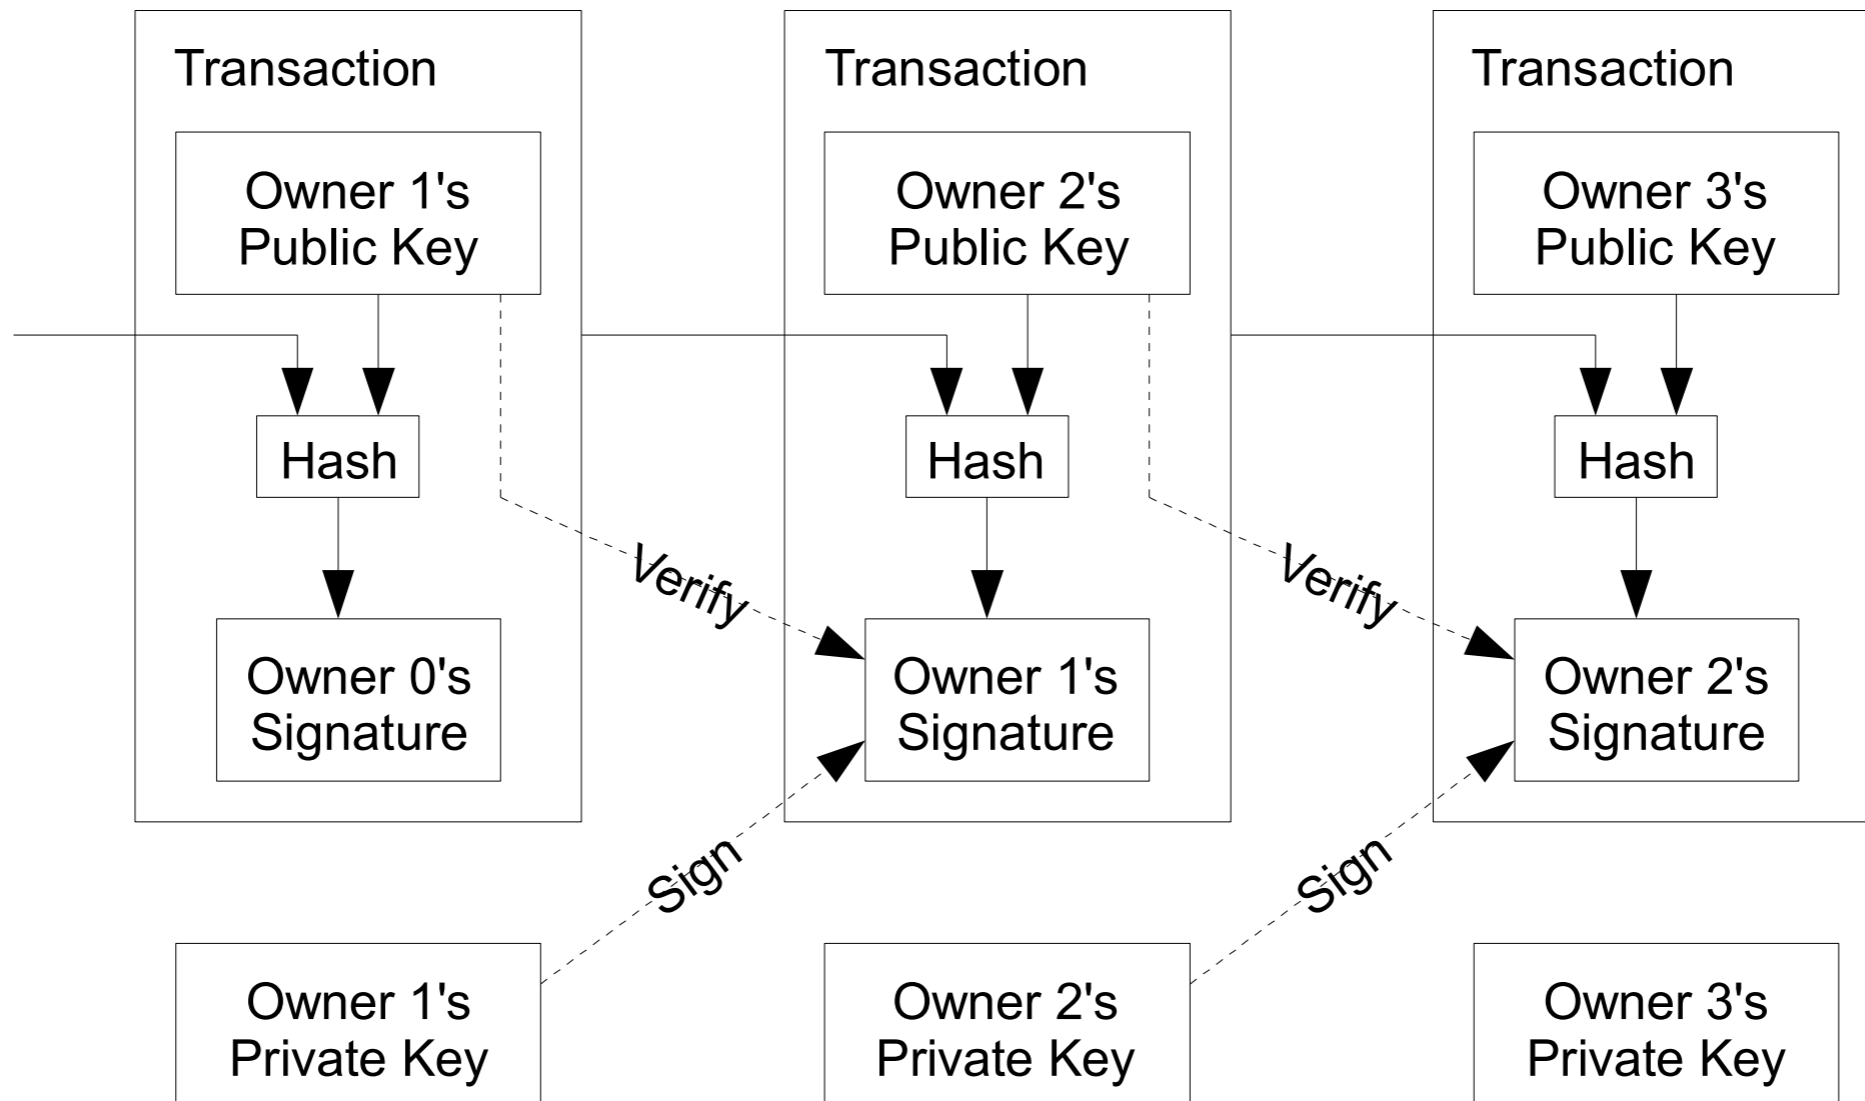


Cryptocurrencies



Blockchain promises



How does it work?



Cryptographic Hash



Preimage resistance



- Given $H(x)$, **computationally hard** to find x

Cryptographic Hash

given



hard to find



Cryptographic Hash

- **2nd preimage resistance** Given $H(x)$, **computationally hard** to find x' such that $H(x) = H(x')$
- **Collision resistance** **computationally hard** to find any x, x' such that $H(x) = H(x')$

Signature

Know me broken by my master
Teach thee on child of love hereafter

Into the flood again
Same old trip it was back then
So I made a big mistake
Try to see it once my way

Drifting body it's sole desertion
Flying not yet quite the notion

Into the flood again
Same old trip it was back then
So I made a big mistake
Try to see it once my way

Into the flood again
Same old trip it was back then
So I made a big mistake
Try to see it once my way

Am I wrong?
Have I run too far to get home?



Know me broken by my master
Teach thee on child of love hereafter


Into the flood again
Same old trip it was back then
So I made a big mistake
Try to see it once my way

Drifting body it's sole desertion
Flying not yet quite the notion

Into the flood again
Same old trip it was back then
So I made a big mistake
Try to see it once my way

Into the flood again
Same old trip it was back then
So I made a big mistake
Try to see it once my way

Am I wrong?
Have I run too far to get home?



- Only Alice can sign
- Everybody knowing Alice's public key can verify

Blockchain Abstraction: Distributed Ledger

Cash				
Date	Description	Increase	Decrease	Balance
Jan. 1, 20X3	Balance forward			\$ 50,000
Jan. 2, 20X3	Cash sale	10,000		60,000
Jan. 3, 20X3	Cash sale	5,000		65,000
Jan. 5, 20X3	Paid rent		2,000	63,000
Jan. 7, 20X3	Paid bills		3,000	60,000
Jan. 8, 20X3	Cash sale	4,000		64,000
Jan. 8, 20X3	Paid bills		2,000	62,000
Jan. 10, 20X3	Paid tax		1,000	61,000
Jan. 12, 20X3	Collected receivable	7,000		68,000

Append-only list of events

Not just financial

Everyone agrees on content

Tamper-proof

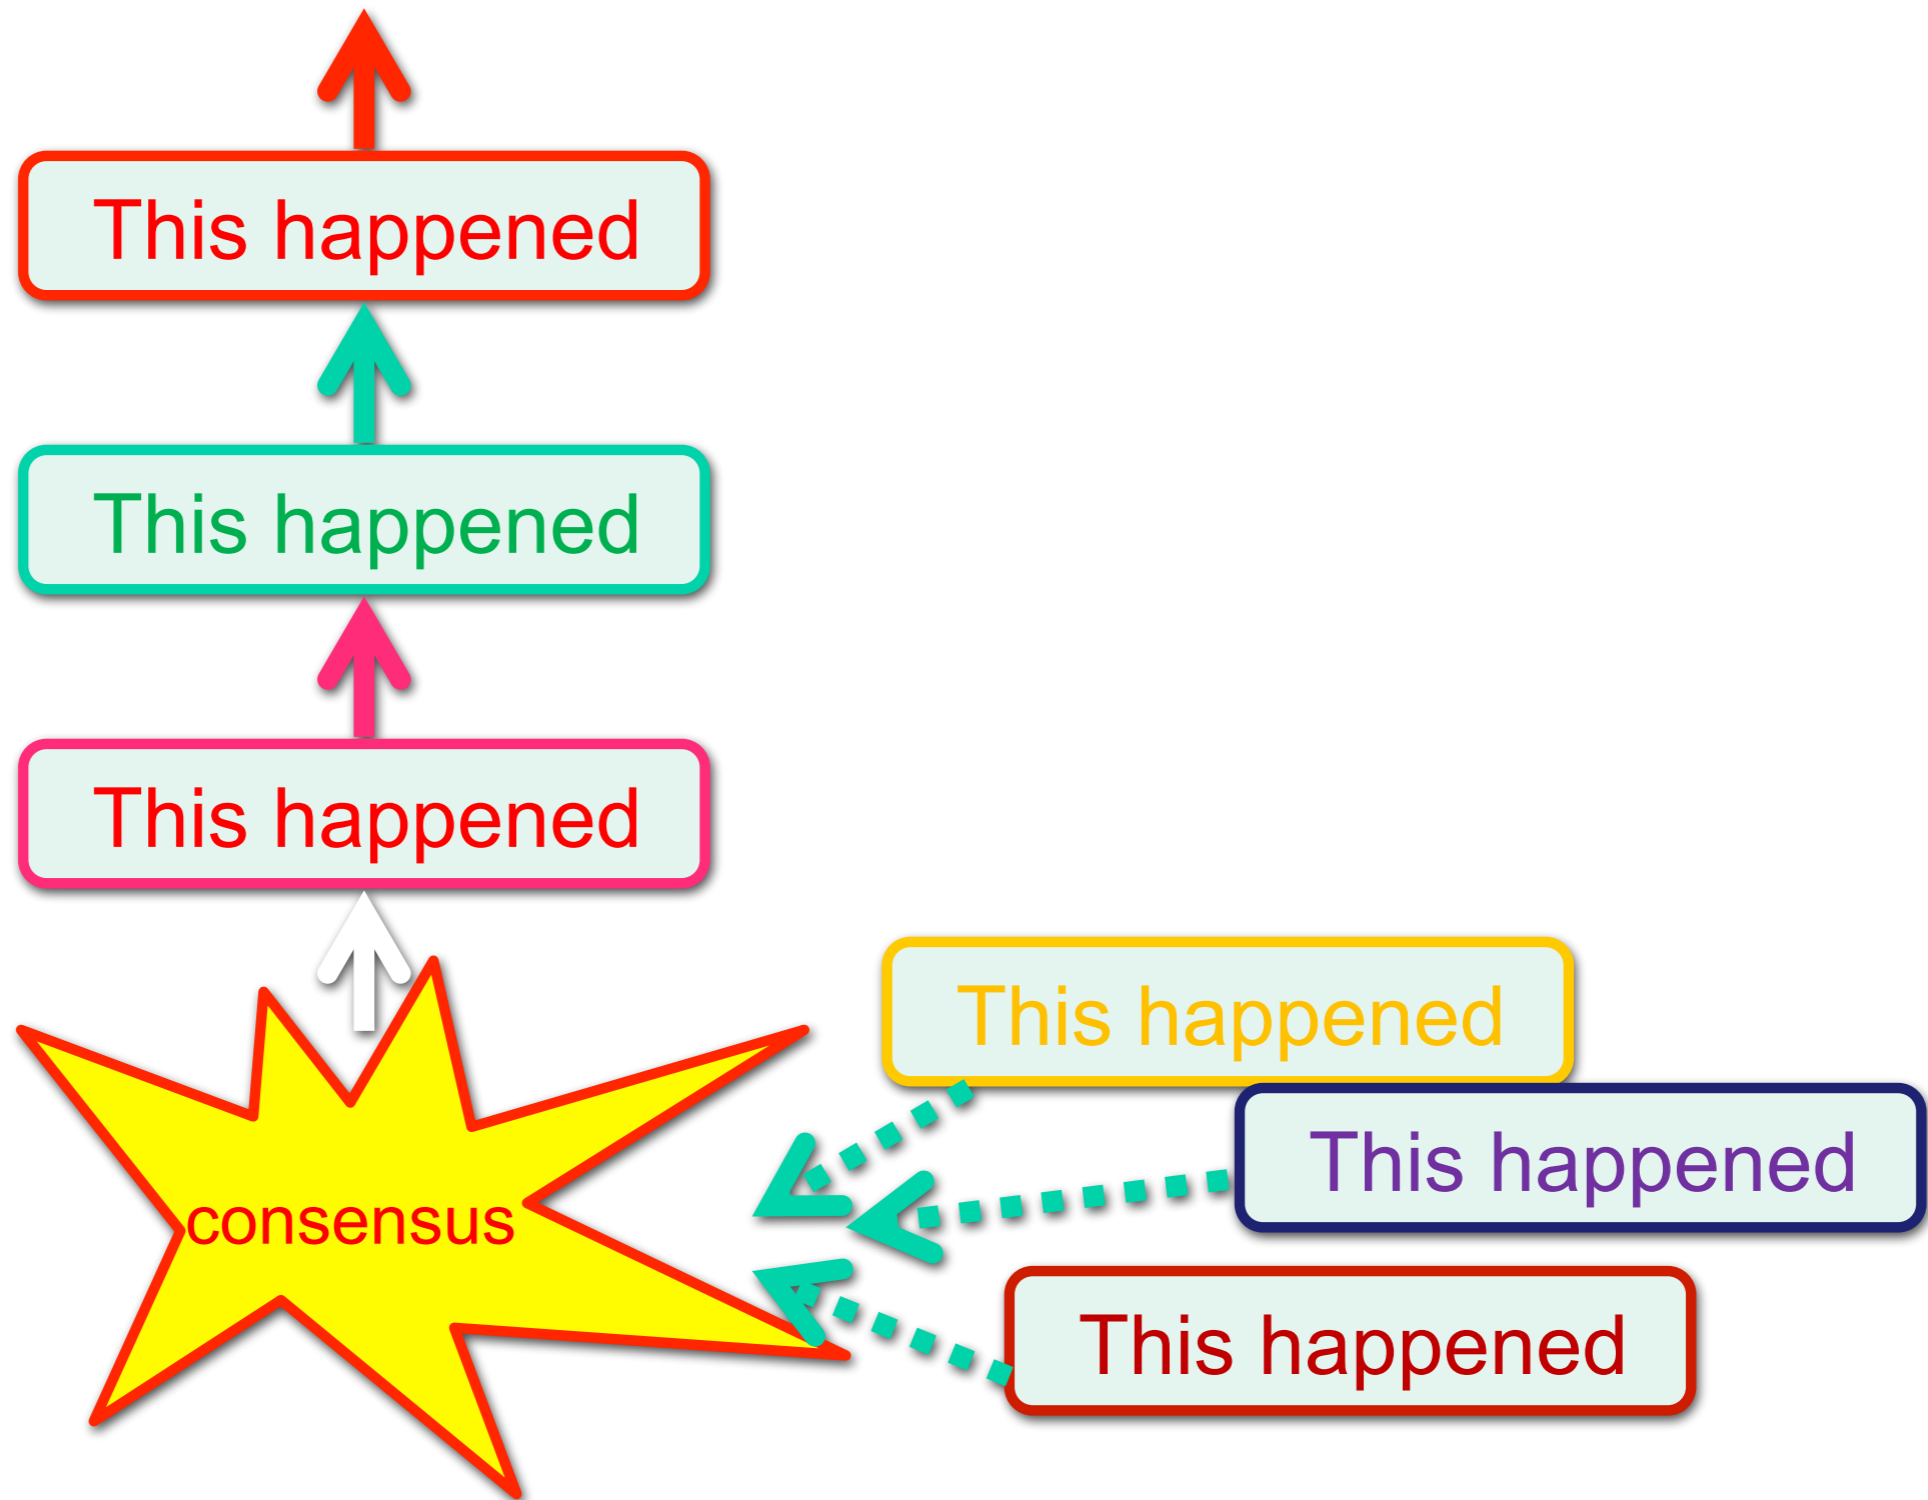
Everyone agrees on content?

Consensus!

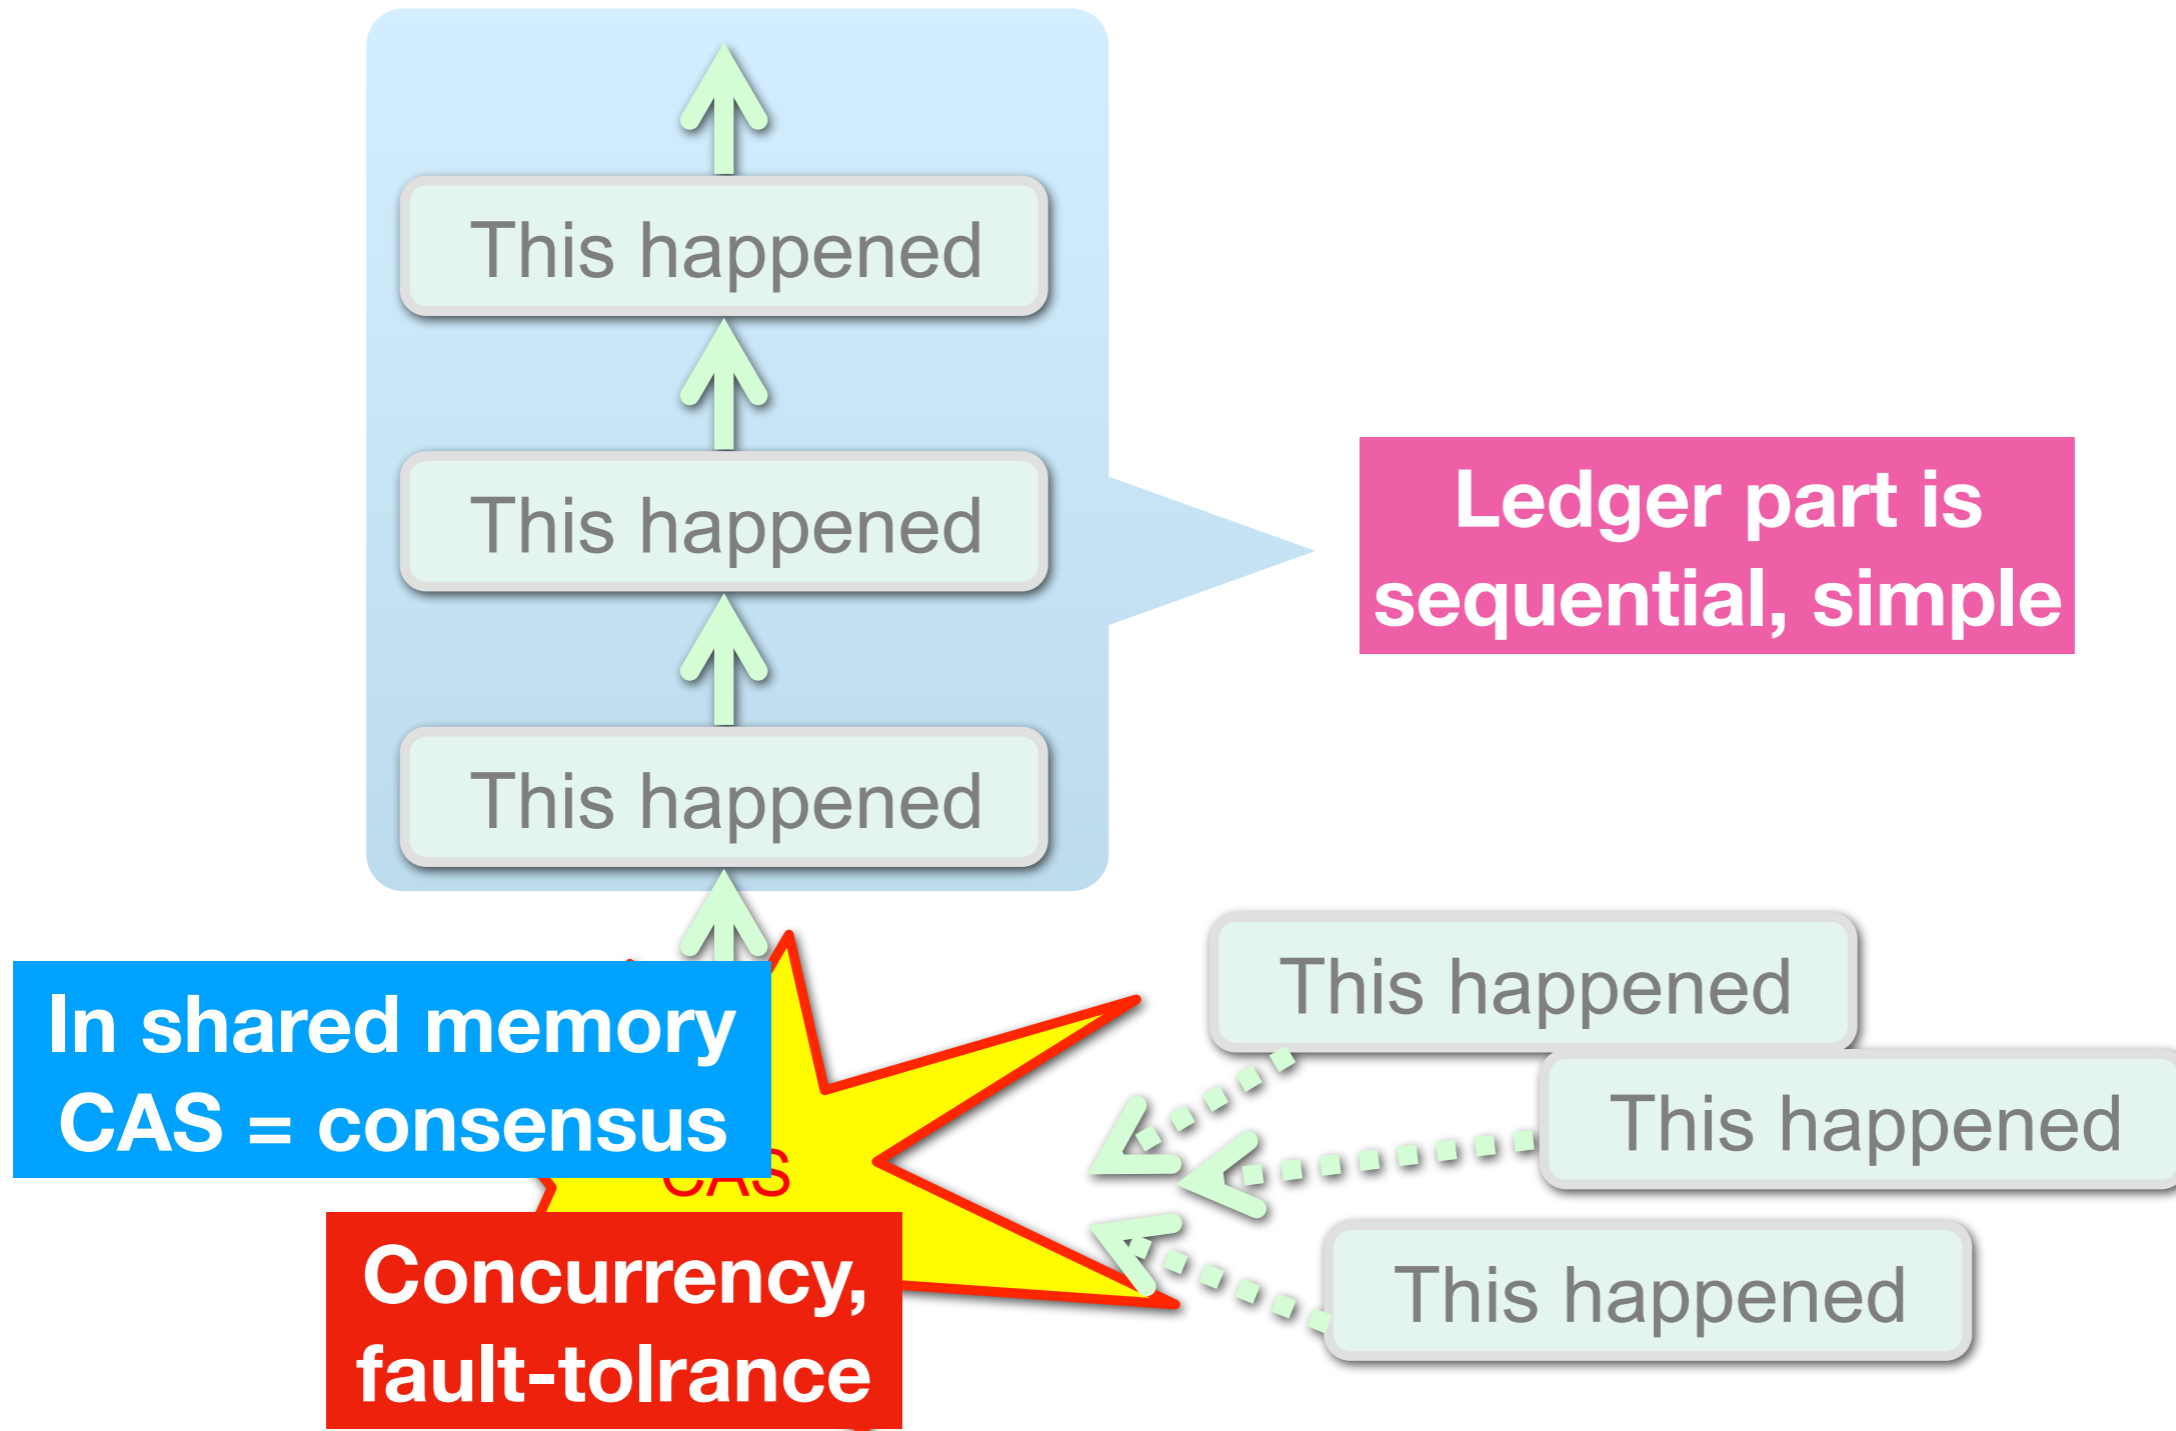
Each thread has a private input and must decide a value

- **Agreement** : they decide the same value
- **Validity** : decision is one of the proposal
- **Agreement** : non-faulty process decide


Universal Construction



Universal Construction: Shared Memory



Atomic Broadcast



Parallel Universes

Traditionnal DC

Blockchain

Consensus, Universal
construction, Atomic
Broadcast

Distributed Ledger

Ids

Pseudonymous

Paxos, PBFT,
zzyyvva, and hundred
more

Nakamoto
consensus, PoS,
PoA

Huge peer-reviewed
academic literature

White papers

Chubby, Raft,
Zookeeper

Many flaws, Bugs,
Hacks

Bitcoin Transaction



123,456 btc



In

1782352eab45

refs to previous *unspent* TXs whose recipient is alice

Out

123,456 4b35147fc

Amount

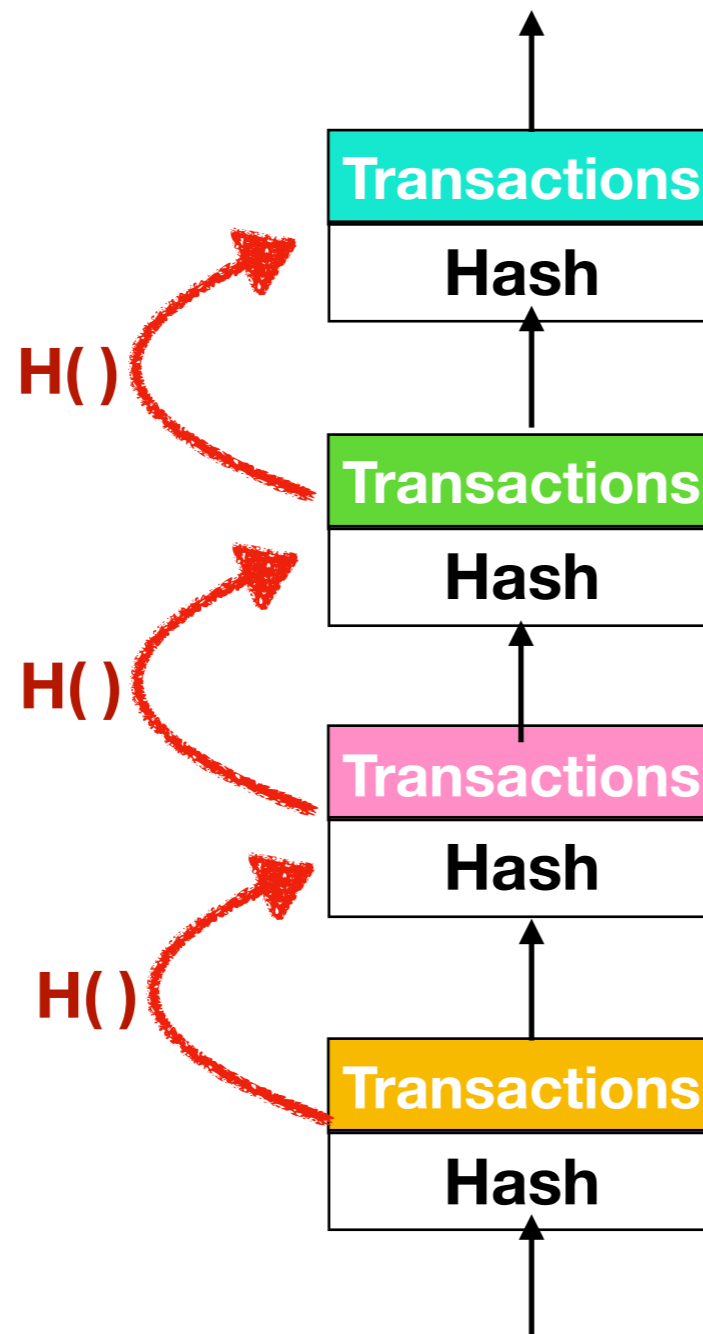
Bob's public key

Sig

3050122eaa90

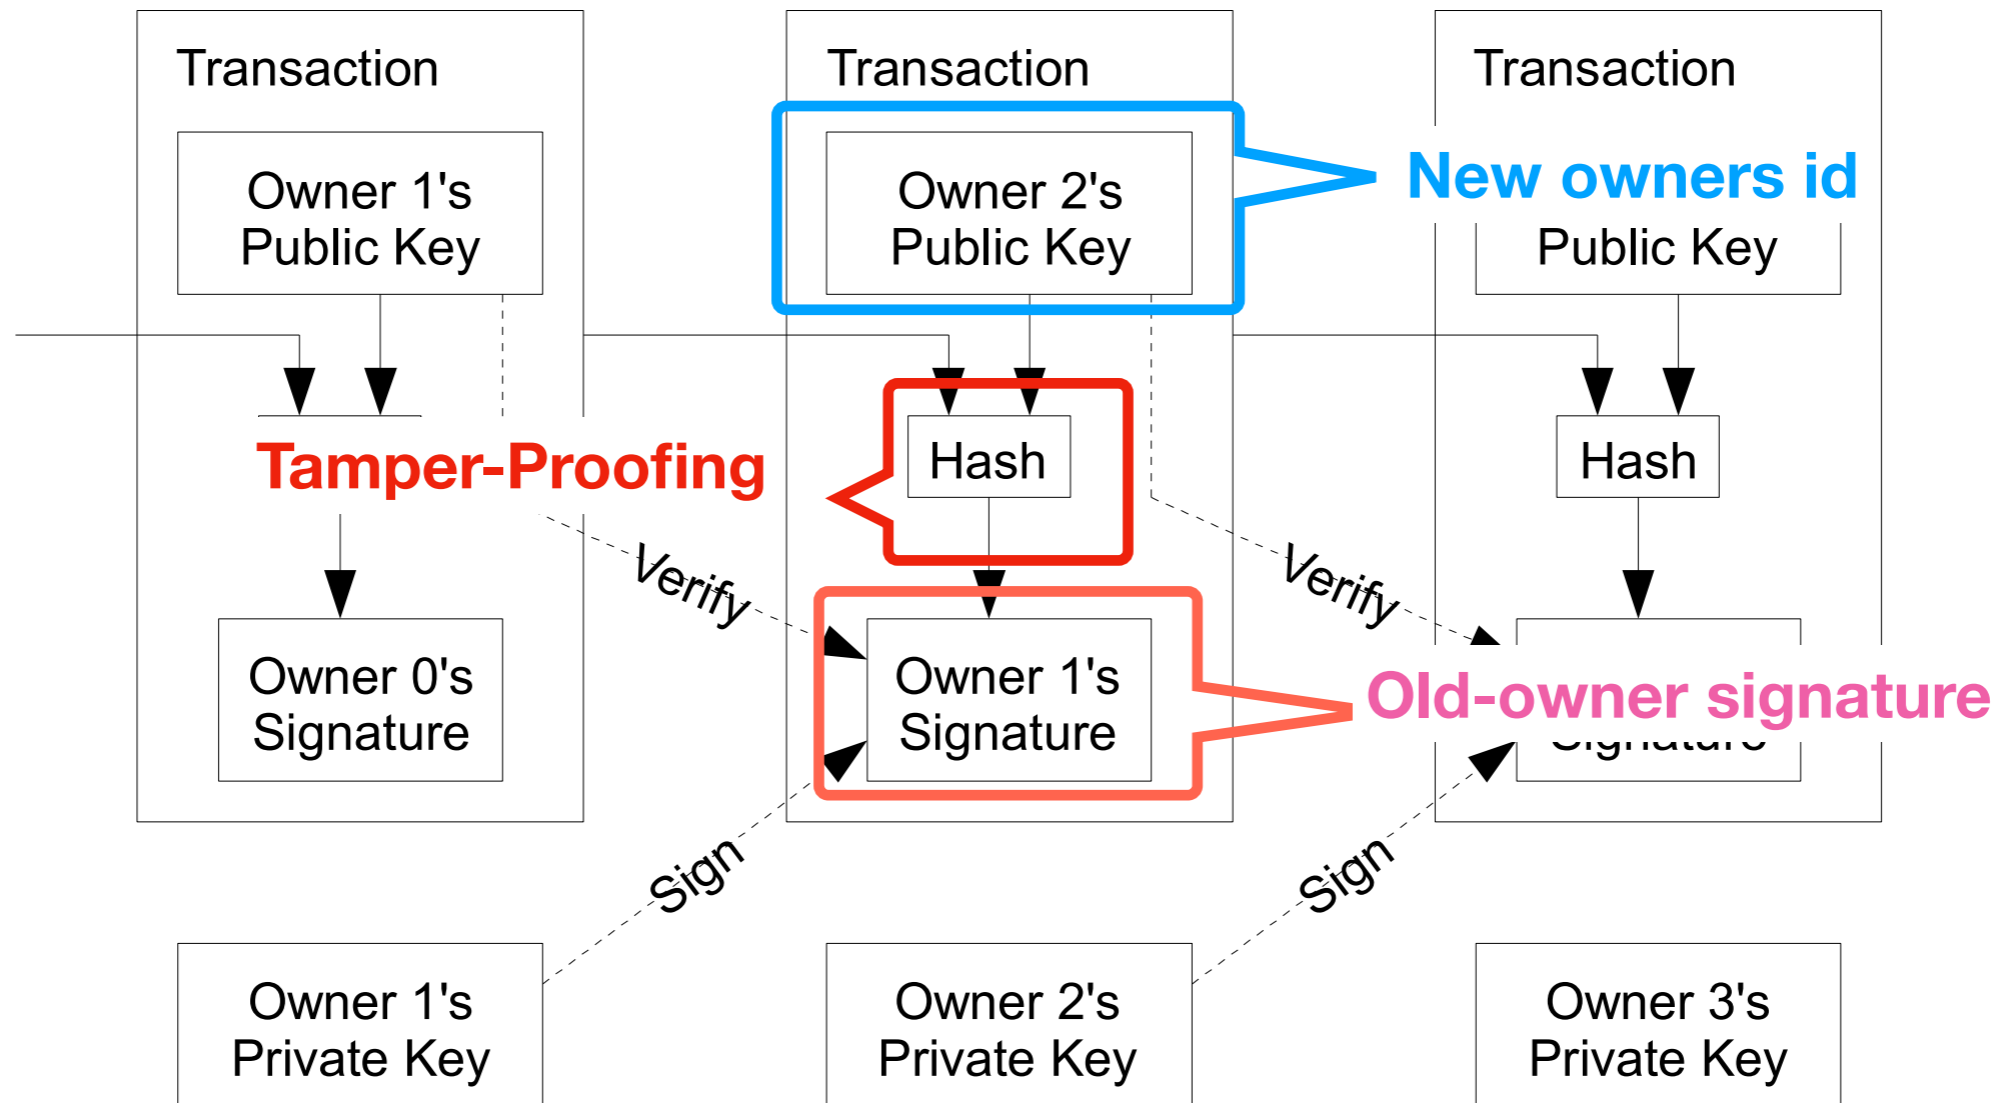
Alice's signature

Tamper-proof




Each TXs block contain the hash of the previous block

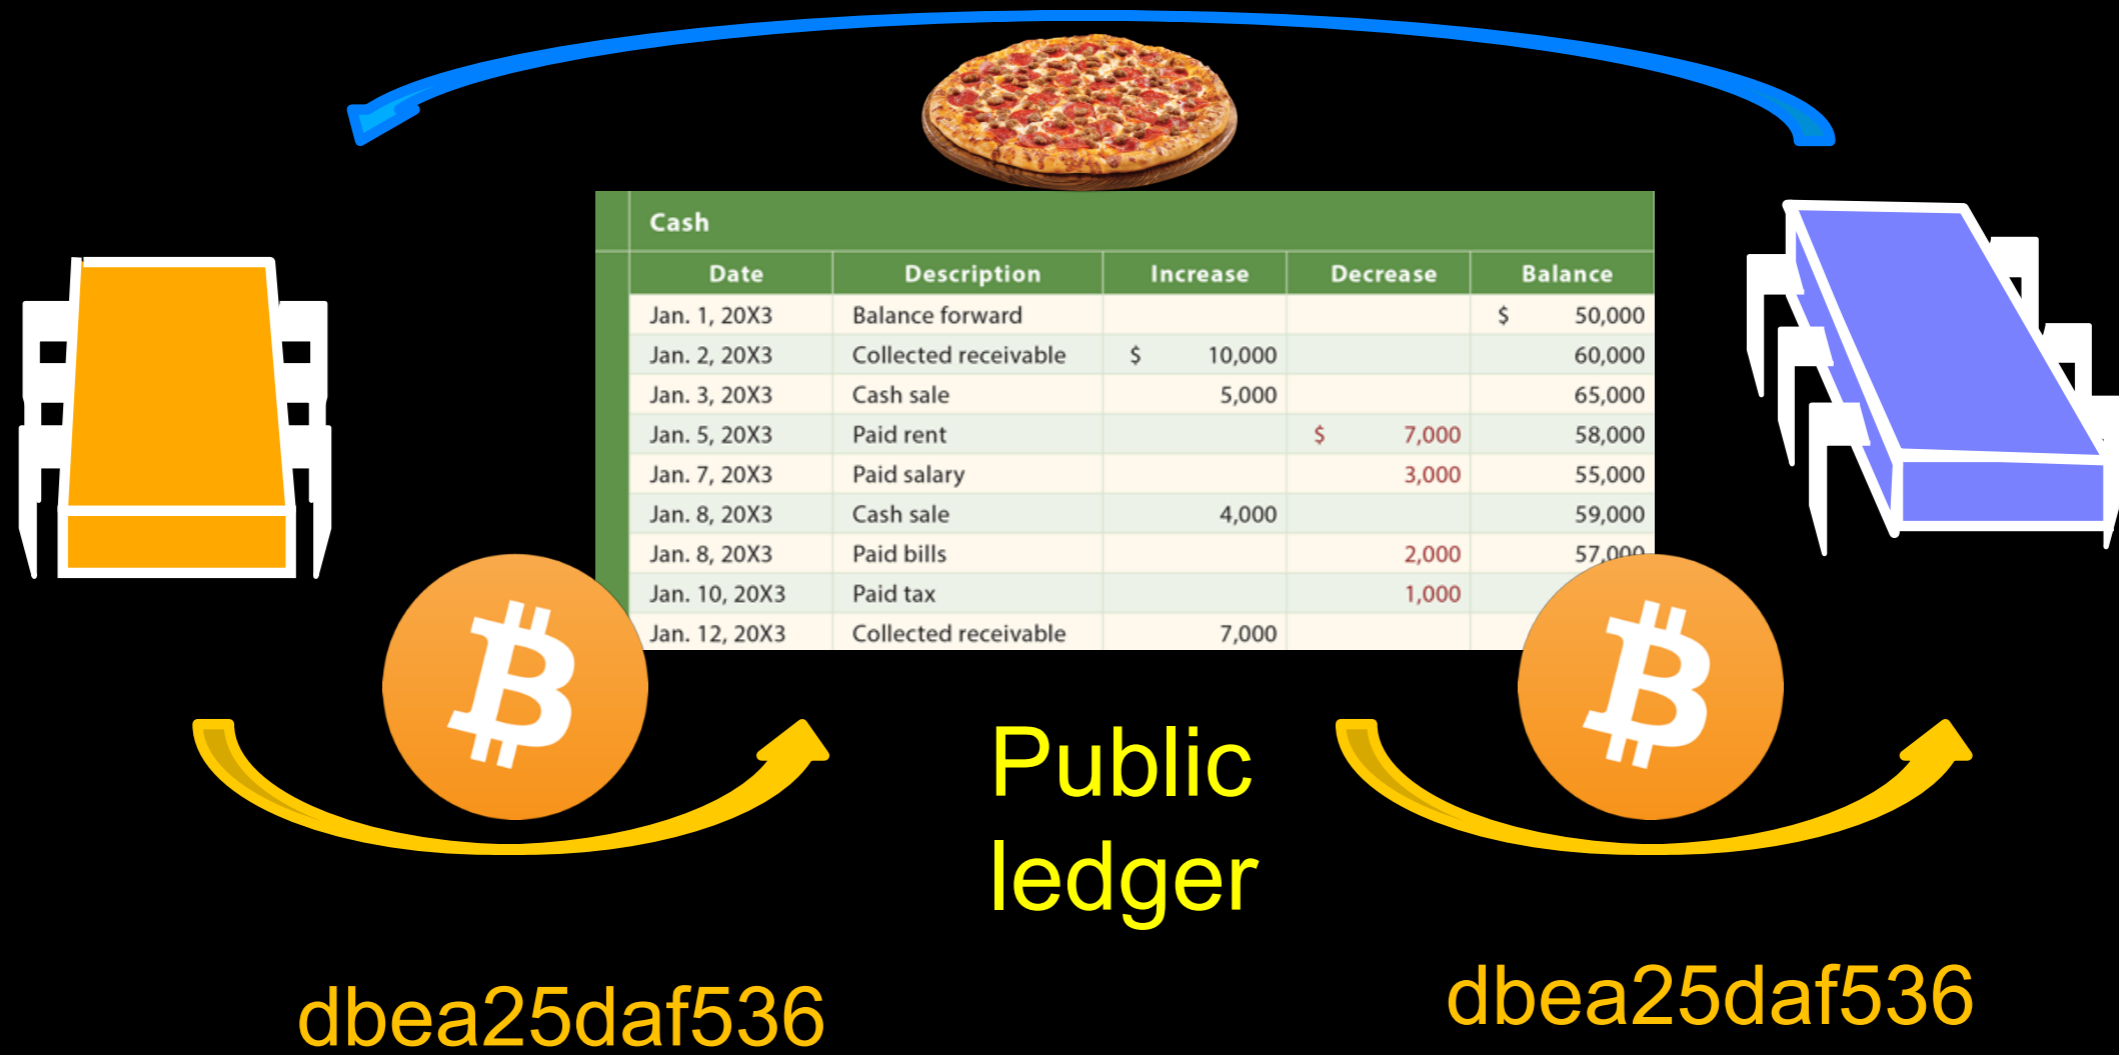
How does it work?



Problem: Double Spending



Nakamoto Solution



Cash				
Date	Description	Increase	Decrease	Balance
Jan. 1, 20X3	Balance forward			\$ 50,000
				60,000
				65,000
				58,000
Jan. 5, 20X3	Paid rent		7,000	55,000
Jan. 7, 20X3	Paid salary		3,000	52,000
				50,000
Jan. 10, 20X3	Paid tax		1,000	49,000
		7,000		56,000
				63,000

Every node keeps a copy of every transaction

Widely considered reckless at the time

Still a scalability issue

Traditionnnaal DC Consensus

A Common Design Pattern

Phase 1 : Conciliation

Select a (block of) proposal

Phase 2 : Conciliation

Adopt/Commit a proposal

**Leader Collect and
Chose a proposal**

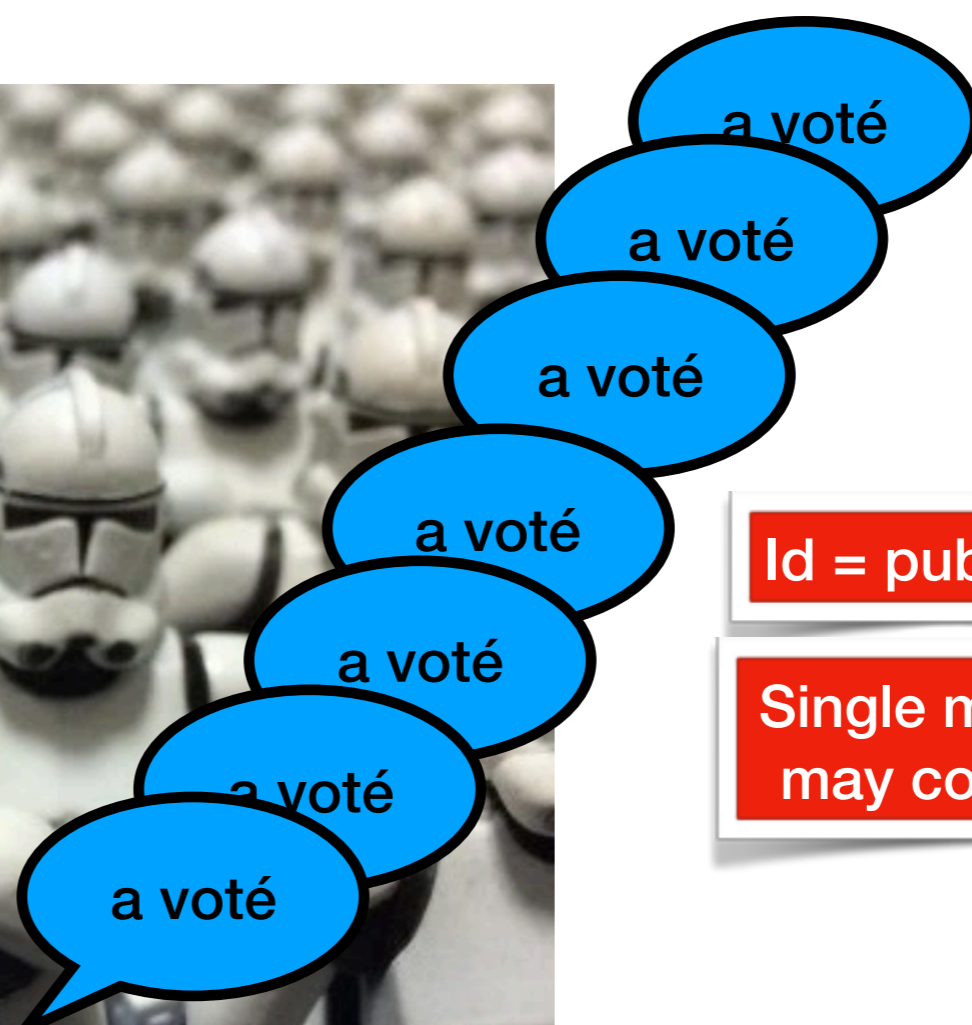

Vote

Iterate if do not succeed

Let's vote



Sybil Attack



Id = public key

Single malicious player may control many ids

Proof of Work



Dwork and Naor 1993

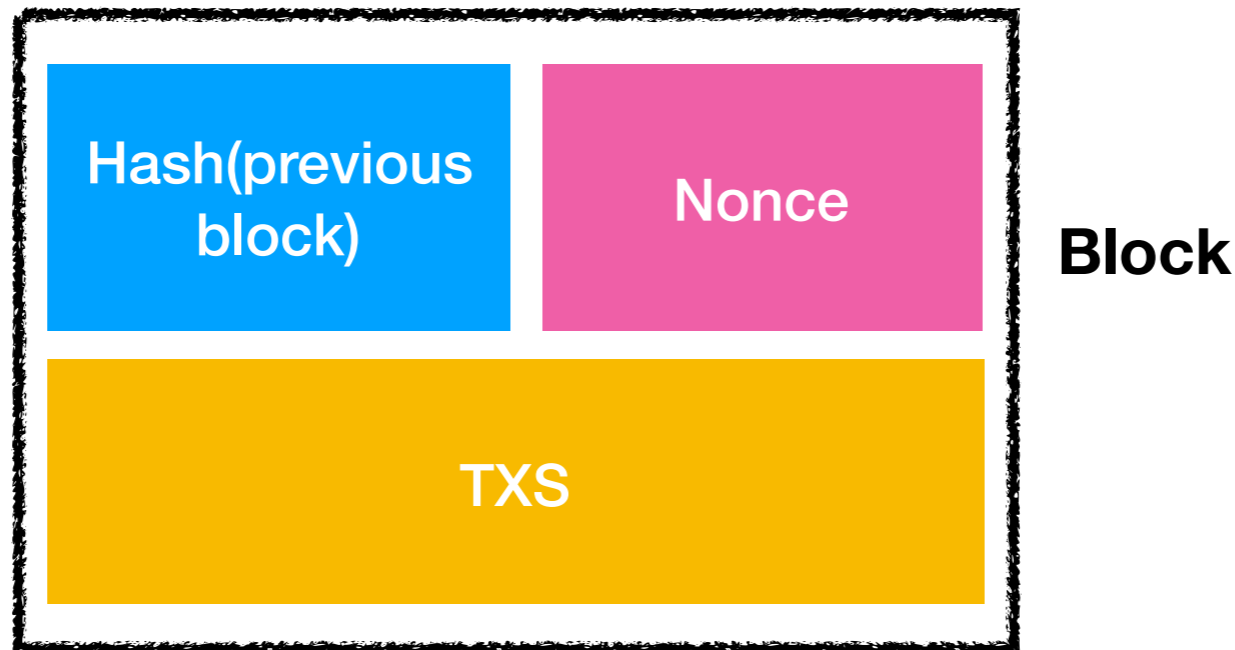
Expensive to fake

Adapted to PoW consensus

PoW Consensus

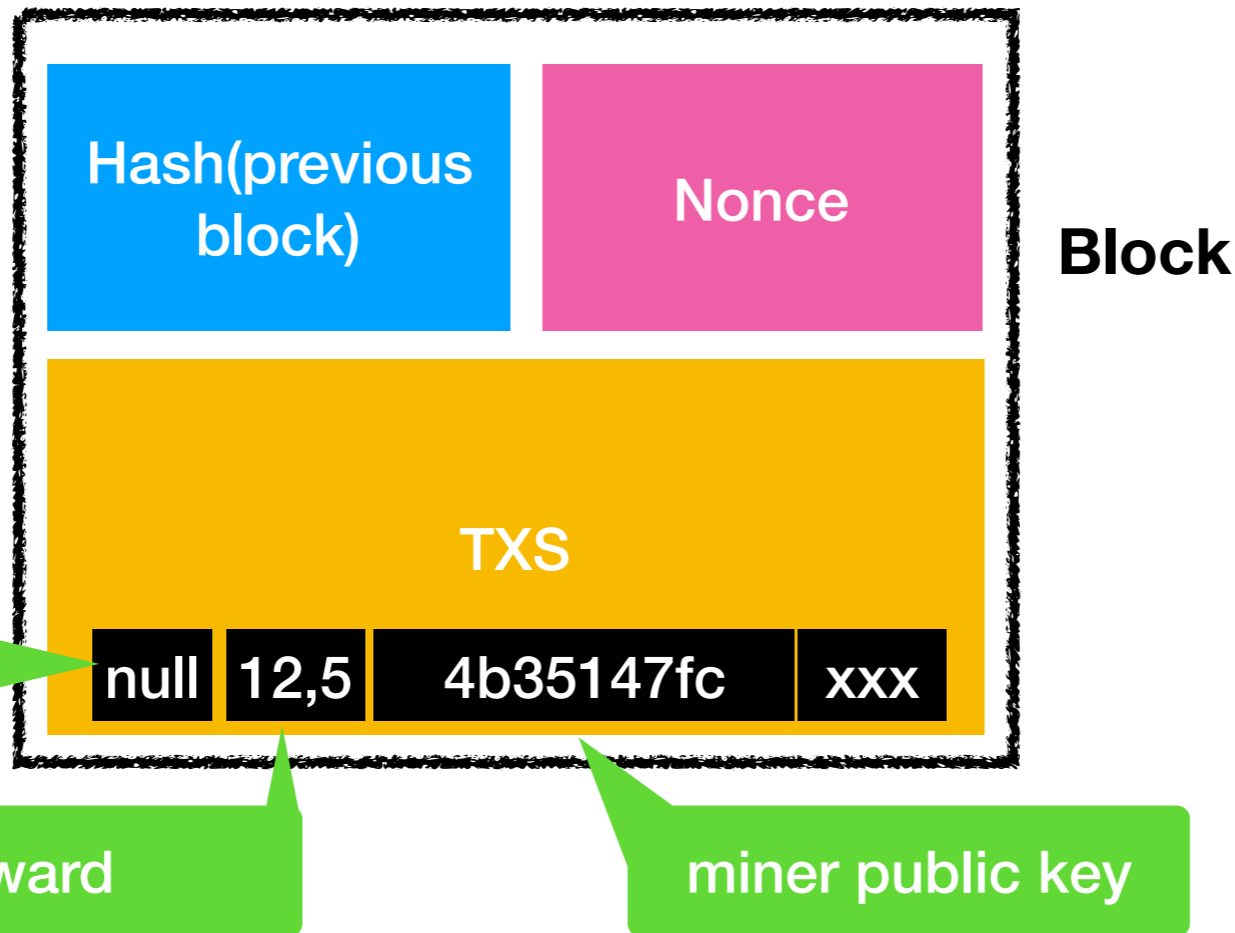
- Miners compete to append block to the chain
- Entry ticket is expensive
- Multiple winners possible

PoW



- Find **Nonce** such that Hash(**Block**) has k leading 0's
- Randomized leader election !
- Chance of winning ~ hashing power

Reward & Incentive to behave



- Reward: newly minted coins
- Winner also collects TXs fees

Multiple Winners ?

- Multiple near simultaneous winners create « forks »
- Infrequent but does happen
- Subsequent winners decide which fork wins
- Differs from classical consensus

Honest Majority Hypothesis

in it. If a majority of CPU power is controlled by honest nodes, the honest chain will grow the fastest and outpace any competing chains. To modify a past block, an attacker would have to redo the proof of work of the block and all blocks after it and then catch up with and surpass the

**Honest miners build on
longest chain ...**

**... longest chain reflects will of
honest miners**

**Dishonest miners would have to out-
compute all honest miners**

Limited Throughput is Feature, not Bug

To compensate for increasing hardware speed and varying interest in running nodes over time, the proof-of-work difficulty is determined by a moving average targeting an average number of blocks per hour. If they're generated too fast, the difficulty increases.

Number of blocks/time kept approximately constant

By varying PoW difficulty

This will become a problem as Bitcoin becomes successful

Parallel Universes

Classical Consensus

Unique winner

Once a decision is reached,
it is final

Permissioned
number of threads fixed
No cheating on Ids

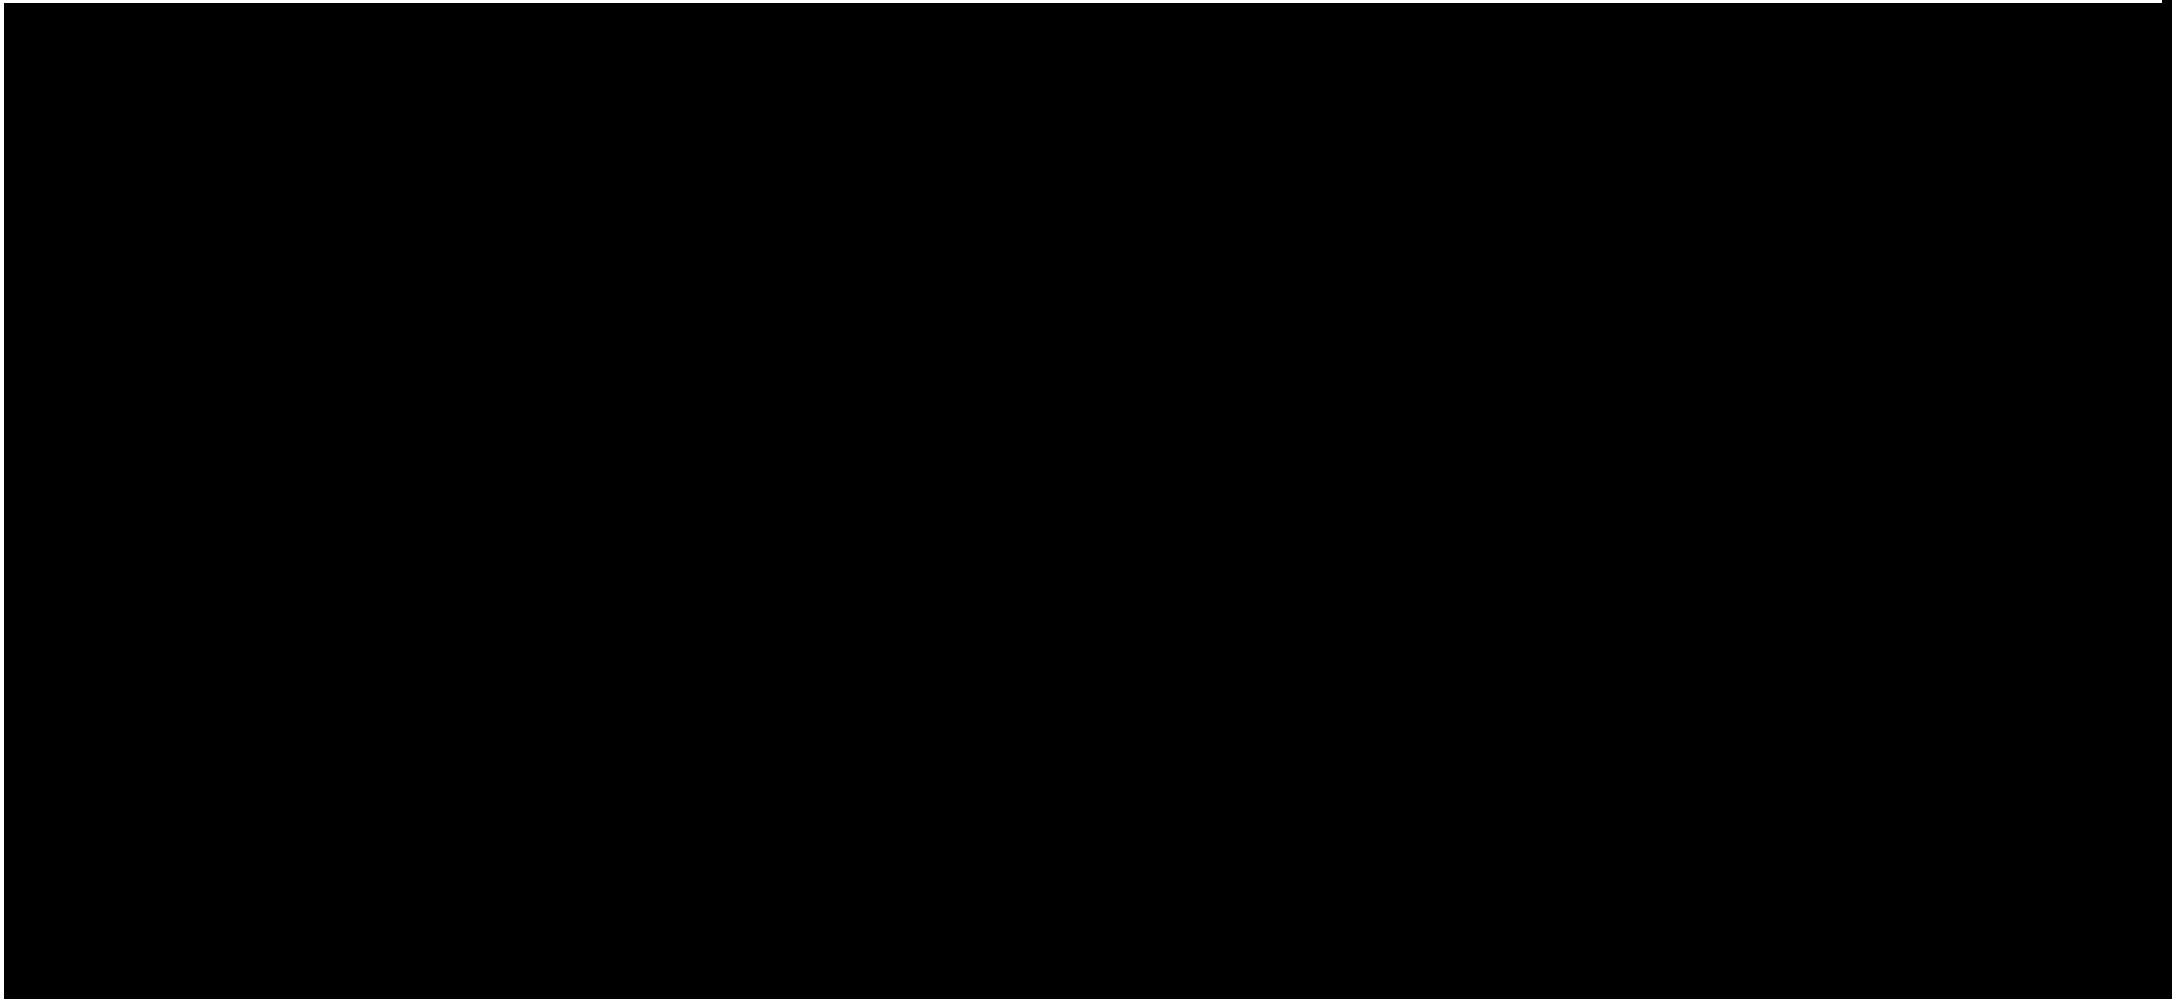
PoW Consensus

Multiple winners possible

Agreement emerges
over time

Permissionless
Anyone can participate
Faking id is cheap

The steps to run the network are as follows:

- 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.
- 

The steps to run the network are as follows:

- 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.

Clients send transactions to miners

Does anyone ever talk about the Bitcoin P2P layer?

Rumor: mining cartels use faster side-channels

Information Propagation in the Bitcoin Network

Christian Decker,* Roger Wattenhofer†

*ETH Zurich, Switzerland cdecker@tik.ee.ethz.ch

†Microsoft Research wattenhofer@microsoft.com

Abstract—Bitcoin is a digital currency that unlike traditional currencies does not rely on a centralized authority. Instead Bitcoin relies on a network of volunteers that collectively implement a replicated ledger and verify transactions. In this paper we analyze how Bitcoin uses a multi-hop broadcast to propagate transactions and blocks through the network to all replicas. We then use the gathered information to propose a conjecture that the propagation of information in Bitcoin is a primary cause for the observed inconsistency in the network. We avoid this inconsistency by pushing information to the replicas.

Empirical Study of Bitcoin P2P network as of 2013

Bitcoin is a truly decentralized global currency system. Like any other currency, its main purpose is to facilitate the exchange of goods and services by offering a commonly accepted good. Unlike traditional currencies however, it is not issued by a state or even a single authority. Since its inception in late 2008, Bitcoin has enjoyed a rapid growth, both in value and in the number of transactions. Its success is mostly due to innovative use of a peer-to-peer network to implement all aspects of a currencies lifecycle, from creation to its transfer between users. This is the fundamental difference from previous research, which concentrated on building systems that rely on either a centralized issuer [5], [16], [18] or creating inter-user credit [9]. These systems required users to trust the original issuer, which was still used to eventually clear transactions. Bitcoin has often been compared to a near-instantaneous and non-trust-based system beyond the scope of the current paper.

Furthermore, an inconsistency about the validity of a replica state, any inconsistency may jeopardize the security of the consensus itself. This may facilitate an attacker that attempts to rewrite transaction history. In this work we analyze Bitcoin from a networking perspective, i.e., how information is disseminated or propagated in the Bitcoin network, we identify key weaknesses as well as the resulting problems. In particular, we analyze the synchronization mechanism which fails to synchronize the information stored at the ledger with a non-negligible probability. This problem not only causes a prolonged inconsistency that goes unnoticed by a large number of nodes, but also weakens the system's defenses against attackers. We then propose some changes to the current protocol that, while not a solution to the intrinsic problems of the communication model used by Bitcoin, mitigate them.

In this section we describe the communication model used by Bitcoin.

The steps to run the network are as follows:

- 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.

**Miners assemble transactions
into blocks**

**Economy of scale: single
transaction too expensive**

**Block size becomes major
headache later on!**

The steps to run the network are as follows:

- 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.

Miners race to do Proof of Work

Today, consumes lots of energy

Cartels with access to cheap power and ASICs control most of hashing power

The steps to run the network are as follows:

- 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.

If multiple winners at the same time ...

the blockchain forks ...

**Result: high latency because need to wait until
your transaction deep enough in chain**

The steps to run the network are as follows:

- 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.

Sanity check: malformed txns rejected

Incentive for miners to behave ...

Double spending filter

The steps to run the network are as follows:

- 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.

Successors build on recent well-formed blocks

Pick longest chain if there is a fork

Break ties arbitrarily

Crime doesn't Pay

The incentive may help encourage nodes to stay honest. If a greedy attacker is able to assemble more CPU power than all the honest nodes, he would have to choose between using it to defraud people by stealing back his payments, or using it to generate new coins. He ought to

**Suppose dishonest party acquires
lots of hashing power ...**

Unlimited double spending?

Or collect all the rewards?

Vandalism destroys coin values!

Calculation

p = probability an honest node finds the next block

Back of the envelope calculation

q_z = probability the attacker will ever catch up from z blocks behind

$q_z = \left\{ \begin{array}{l} \dots \\ (q \dots) \end{array} \right.$ **How likely dishonest miner can overtake honest miner to reverse transaction?**

Exponentially small in gap size

Calculation naïve but probably mostly right

More Precise Calculations

- Garay et. al The Bitcoin Protocol: Analysis and Applications
- R Pass and E Shi. [The Sleepy Model of Consensus](#)

Bitcoin Today Problems

- Long and unpredictable confirmation times
- Block size limit and technical complexity
- High transaction fees

Fail as a medium of exchange

Bombastic success for investment/speculation

Research Directions

Permissionless consensus protocol

- Eventual consensus: agreement on a prefix of the blockchain
- Exponential convergence: probability of fork of depth k is $1/2^k$
- Liveness : new block as a reasonable rate
- Correctness : blocks in the correct chain are valid
- Fairness: miner success rate proportional to hash power

Power of the adversary

- Honest majority assumption
- But what if collusion of miner somewhat control network delay ?
- Selfish mining strategy / Mining cartel

Bitcoin interface

- Bitcoin wallet (lot of attacks)
- Swap with other (crypto)currency

Privacy

- Transaction are public
- User = Public key
- Analysis of transaction network leaks private data
- Cash

Alternative to PoW consensus

- PoW is bad for the planet
- Alternatives PoS, PoA, Proof of Space/Time, ASICs resistant, Useful computation