


Secure Decentralization for a Global-Scale Trustworthy Infrastructure

Vassilis Zikas

9.1.2023


NTUA

The Decentralization Paradigm




Classical Examples:

- Database
- Voting System
- Auctions System
- Authentication Service



The Decentralization Paradigm



Why?

- Trust and Power
- Reliability
- Single point of failure (privacy, DoS ..)
- Load (scalability)




Blockchain as a Scientist/Engineer

Blockchains and Decentralized Ledger Technology (DLT)




Game-changer: Promises to disrupt/improve basic infrastructure:

- Financial Transactions (Cryptocurrencies)
- Trustworthy online computing (Smart contracts)
- Supply chain coordination and tracking
- Fair exchange of digital goods
- Coordinated IoTs
- Biomedical data sharing and access



Blockchain as an Layman



ios	Newsletters	Podcasts	CoinDesk
s	Companies	Policy	Technology
	Web3	Learn	Layer 2
+1.06%	Ethereum ▲ \$1,123.18 +1.91%	Binance Coin ▲ \$262.72 +3.50%	XRP ▲ \$0.37258014 +4.43%
	Binance USD ▲ \$0.9997072		Crypto Pri

We need a principled approach to take the technology to its next step!

Exclusive: At least \$1 billion of client funds missing at failed crypto firm FTX

By Angus Berwick


By Emily Nicolle

24 May 2022, 04:48 GMT-4

Roadmap

- **Part 1: A Principled Design and Analysis of Decentralization with Blockchain**
 - Proof of Work, Proof of Stake, and beyond
- **Part 2 (in passing...): Economics and blockchain**


The Decentralization Paradigm



Methodology

1. Specification
2. System Model
3. Distributed protocol
4. Assumptions and security analysis

The Decentralization Paradigm: Bitcoin



Methodology

1. Specification
2. System Model
3. Distributed protocol
4. Assumptions and security analysis

“Reverse Engineering” Bitcoin

Bitcoin: A Peer-to-Peer Electronic Cash System

Satoshi Nakamoto
satoshi@gmx.com
www.bitcoin.org

Abstract. A purely peer-to-peer version of electronic cash would allow online payments to be sent directly from one party to another without going through a financial institution. Digital signatures provide part of the solution, but the main benefits are lost if a trusted third party is still required to prevent double-spending. We propose a solution to the double-spending problem using a peer-to-peer network. The network timestamps transactions by hashing them into an ongoing chain of hash-based proof-of-work, forming a record that cannot be changed without redoing the proof-of-work. The longest chain not only serves as proof of the sequence of events witnessed, but proof that it came from the largest pool of CPU power. As long as a majority of CPU power is controlled by nodes that are not cooperating to attack the network, they'll generate the longest chain and outpace attackers. The network itself requires minimal structure. Messages are broadcast on a best effort basis, and nodes can leave and rejoin the network at will, accepting the longest proof-of-work chain as proof of what happened while they were gone.

1. Introduction


Commerce on the Internet has come to rely almost exclusively on financial institutions serving as trusted third parties to process electronic payments. While the system works well enough for most transactions, it still suffers from the inherent weaknesses of the trust-based system. Completely non-reversible transactions are not really possible, since a trusted third party must be involved to avoid mediating disputes. The cost of mediation increases transaction size and cutting off the possibility of reversal. With the possibility of reversal, the need for a minimum practical transaction size and cutting off the possibility of reversal. With the possibility of reversal, the need for a minimum practical transaction size and cutting off the possibility of reversal. With the possibility of reversal, the need for a minimum practical transaction size and cutting off the possibility of reversal.

What is needed is an electronic payment system based on cryptography, 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 provide protection from fraud, and routine escrow mechanisms could easily be implemented. In this paper, we propose a solution to the double-spending problem using a peer-to-peer network. The network timestamps transactions by hashing them into an ongoing chain of hash-based proof-of-work, forming a record that cannot be changed without redoing the proof-of-work. The longest chain not only serves as proof of the sequence of events witnessed, but proof that it came from the largest pool of CPU power. As long as a majority of CPU power is controlled by nodes that are not cooperating to attack the network, they'll generate the longest chain and outpace attackers. The network itself requires minimal structure. Messages are broadcast on a best effort basis, and nodes can leave and rejoin the network at will, accepting the longest proof-of-work chain as proof of what happened while they were gone.



10. Privacy

The traditional banking model achieves a level of privacy by limiting access to information to the parties involved and the trusted third party. The necessity to announce all transactions publicly precludes this method, but privacy can still be maintained by breaking the flow of information in another place: by keeping public keys anonymous. The public can see that someone is sending an amount to someone else, but without information linking the transaction to anyone. This is similar to the level of information released by stock exchanges, where the time and size of individual trades, the "tape", is made public, but without telling who the parties were.



As an additional firewall, a new key pair should be used for each transaction to keep them from being linked to a common owner. Some linking is still unavoidable with multi-input transactions, which necessarily reveal that their inputs were owned by the same owner. The risk is that if the owner of a key is revealed, linking could reveal other transactions that belonged to the same owner.

11. Calculations

We consider the scenario of an attacker trying to generate an alternate chain faster than the honest chain. Even if this is accomplished, it does not throw the system open to arbitrary changes, such as changing a transaction or taking money that never belonged to the attacker. Nodes are required to wait for a certain number of blocks before accepting a new block as payment, and honest nodes will never accept a block with a previous block's hash. An attacker can only try to change one of his own transactions to take back

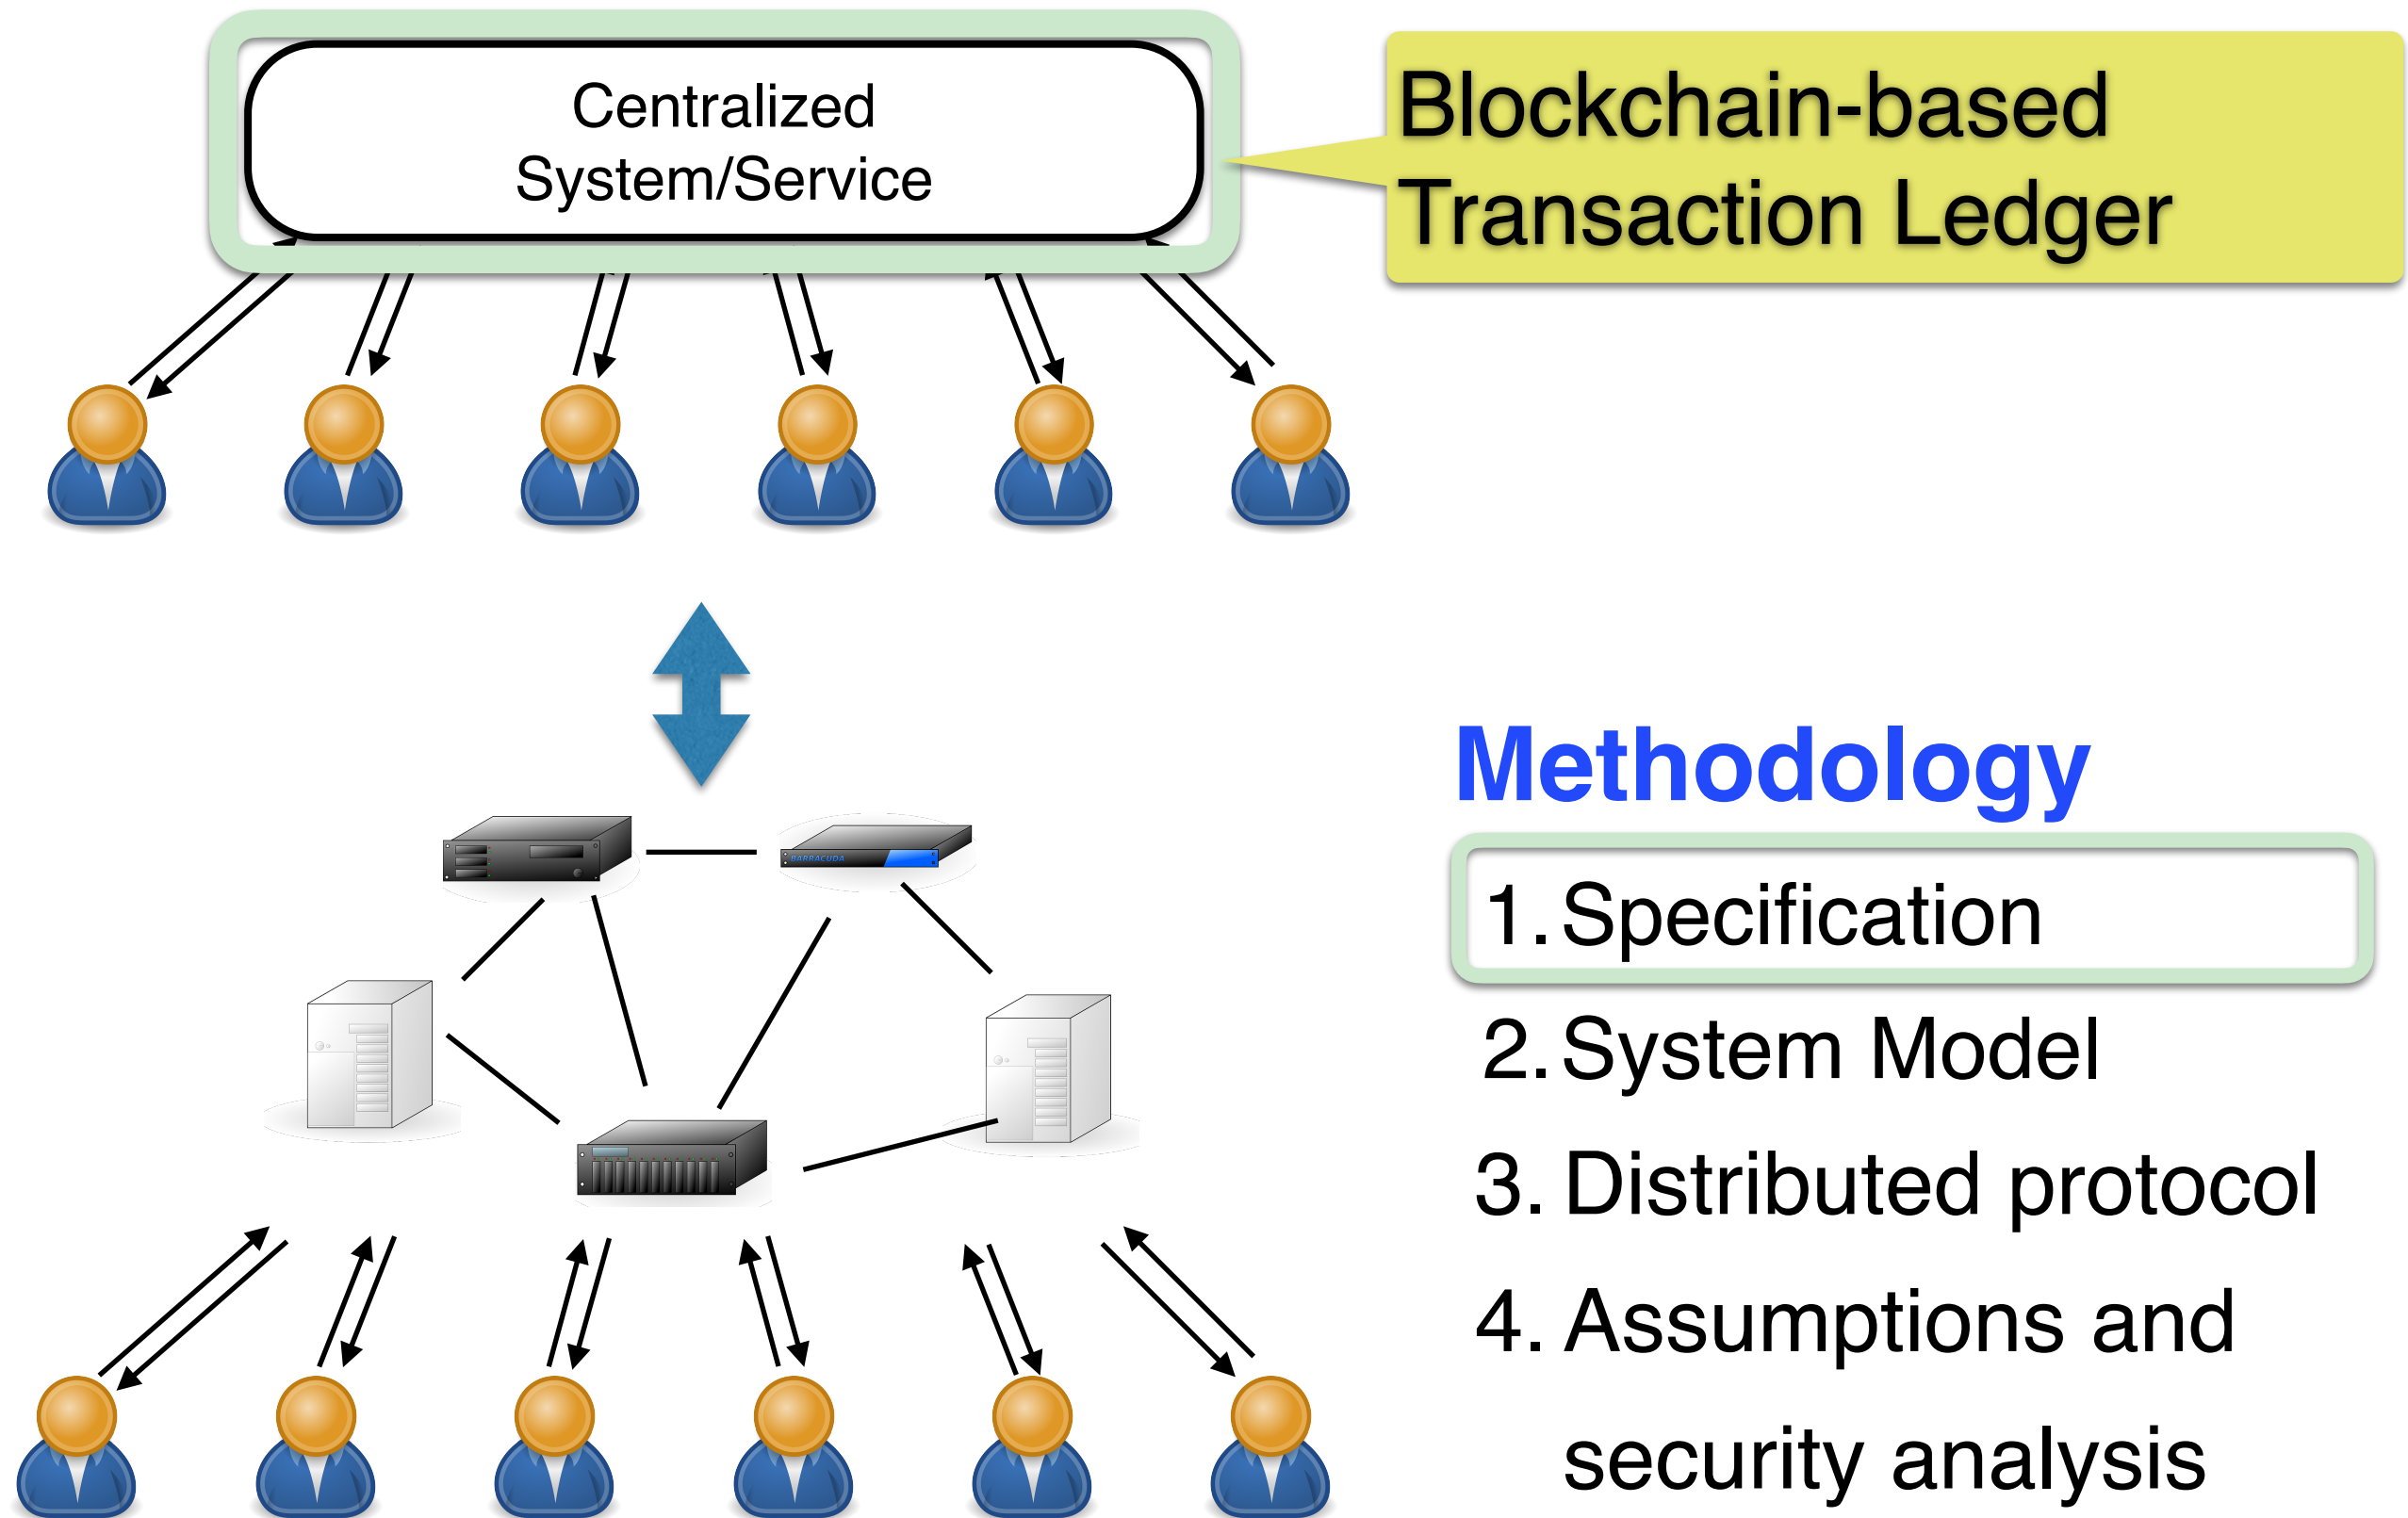
the honest chain and an attacker chain can be characterized as a Binomial process. The honest chain being extended by one block, increasing its length by one. The attacker's chain being extended by one block, reducing the deficit by one.

An attacker catching up from a given deficit is analogous to a Gambler's Ruin problem. A gambler with unlimited credit starts at a deficit and plays potentially an infinite number of times to try to reach breakeven. We can calculate the probability he ever catches up with the honest chain, as follows [8]:

Let q be the probability that the honest node finds the next block before the attacker. Then the probability that the attacker will ever catch up from z blocks behind is:

$$q_z = \begin{cases} 1 & \text{if } p \leq q \\ (q/p)^z & \text{if } p > q \end{cases}$$



The Decentralization Paradigm: Bitcoin



The Decentralization Paradigm: Bitcoin

Blockchain-based Transaction Ledger

Background: Two common tools from crypto

- **Hash function $H()$** as a random n-bit function
 - $H(x) =$  $01011...1101 = y$
 - $H(x') =$  $01000...0101$
 - $H(x) = y$
- **Digital signature $\text{sig}[\cdot]$**
 - Alice has a secret key sk_{Alice}
 - Only Alice can produce signatures $\text{sig}_{\text{sk}_{\text{Alice}}}[\cdot]$
 - Anyone with Alices public key pk_{Alice} can verify

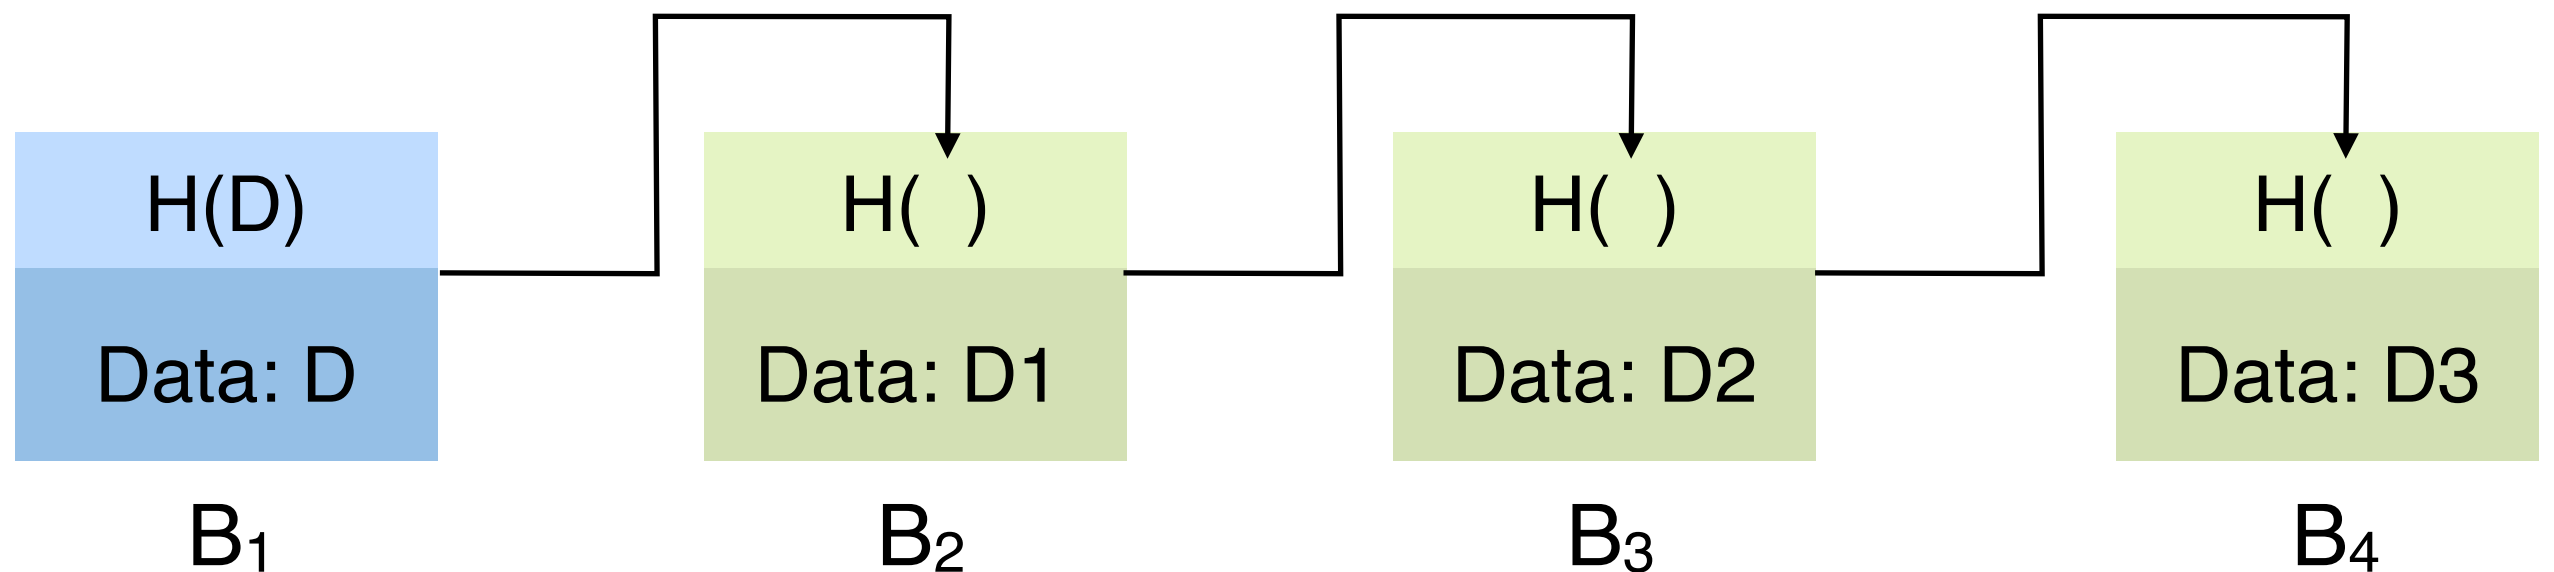
The Decentralization Paradigm: Bitcoin

Blockchain-based Transaction Ledger

The Decentralization Paradigm: Bitcoin

Blockchain-based Transaction Ledger

A hash-pointer-based Data Structure

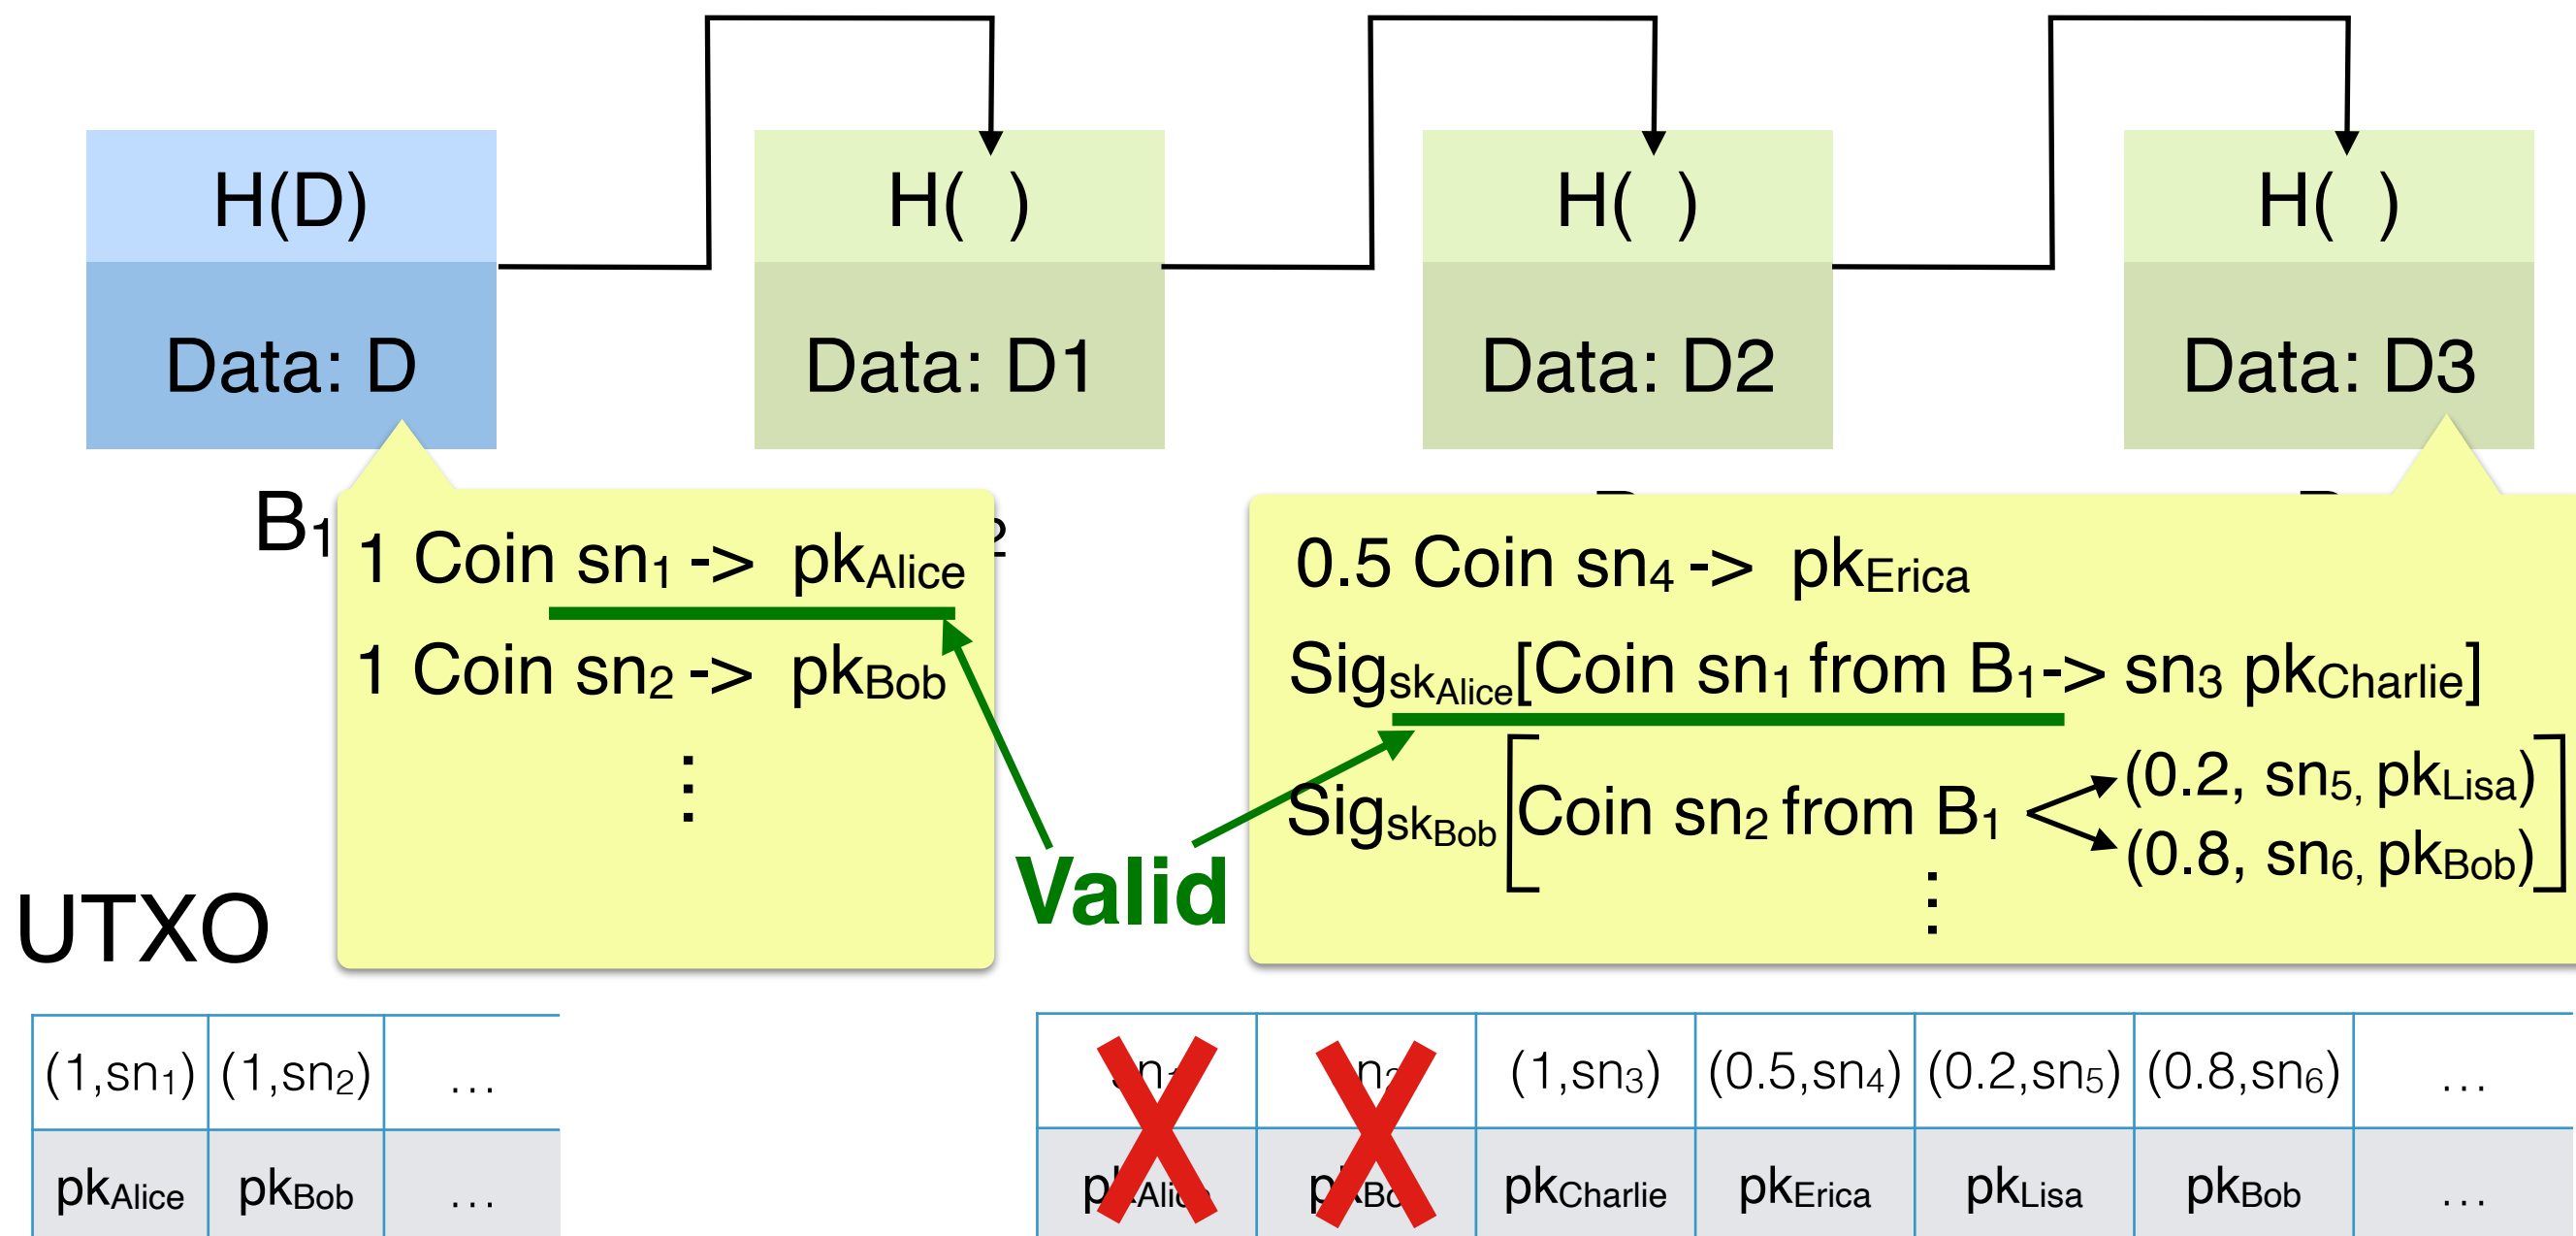


- **Causality (aka “time-stamping”)**: You cannot create block B_i without knowing block B_{i-1}
- **Immutability**: You cannot modify a block anywhere without changing every follow-up block

The Decentralization Paradigm: Bitcoin

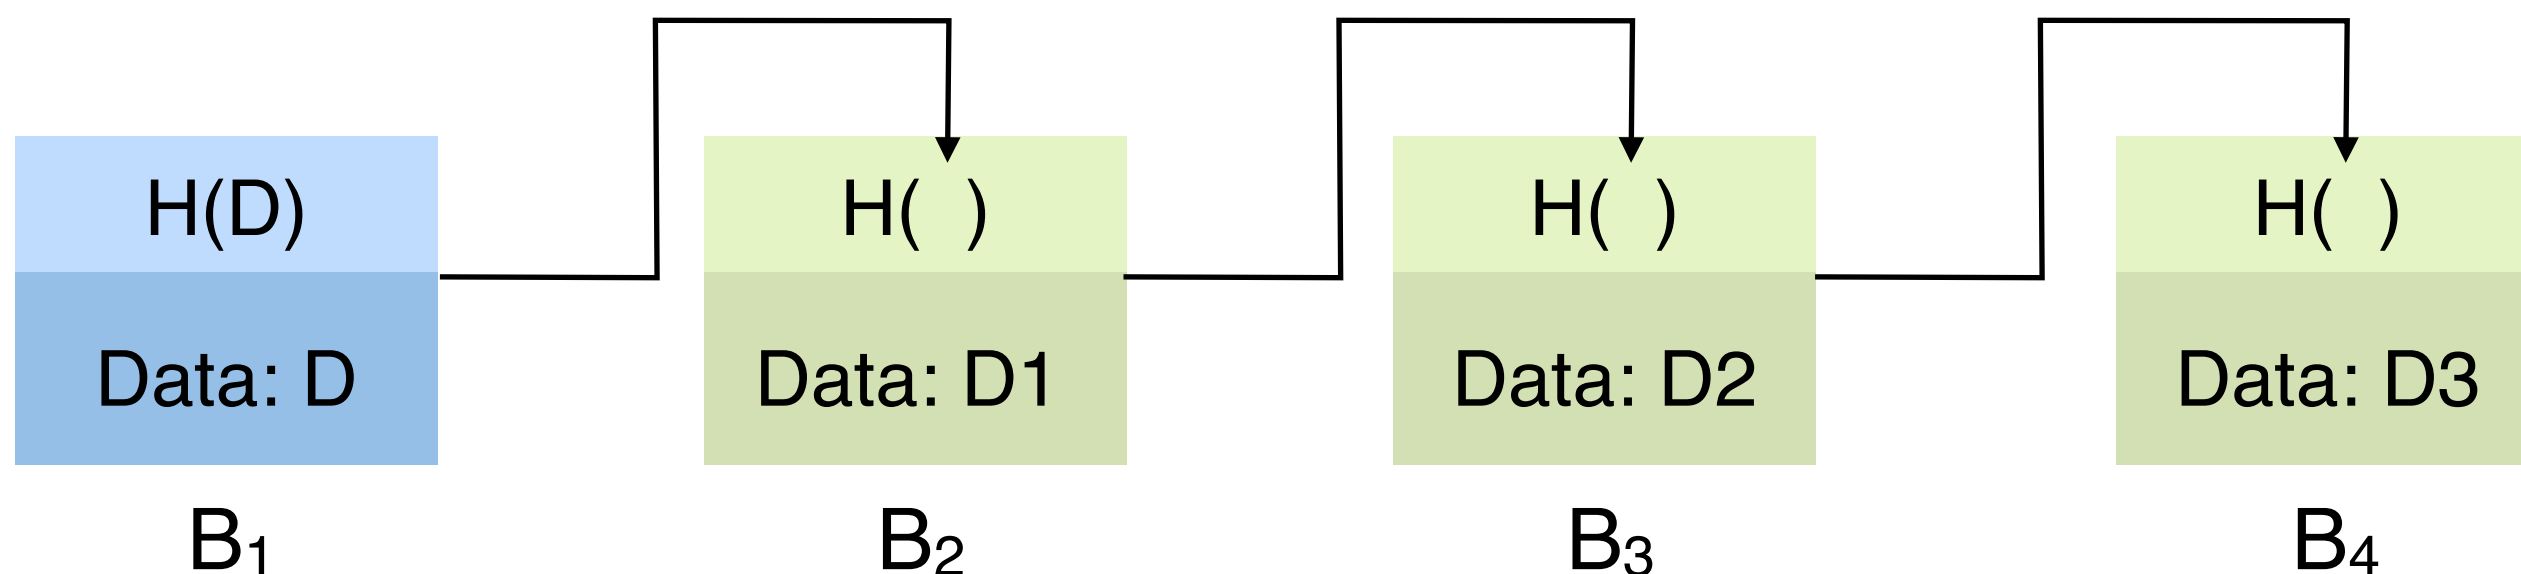
Blockchain-based Transaction Ledger

Putting semantics on the blockchain data



The Decentralization Paradigm: Bitcoin

Blockchain-based Transaction Ledger



The Decentralization Paradigm: Bitcoin

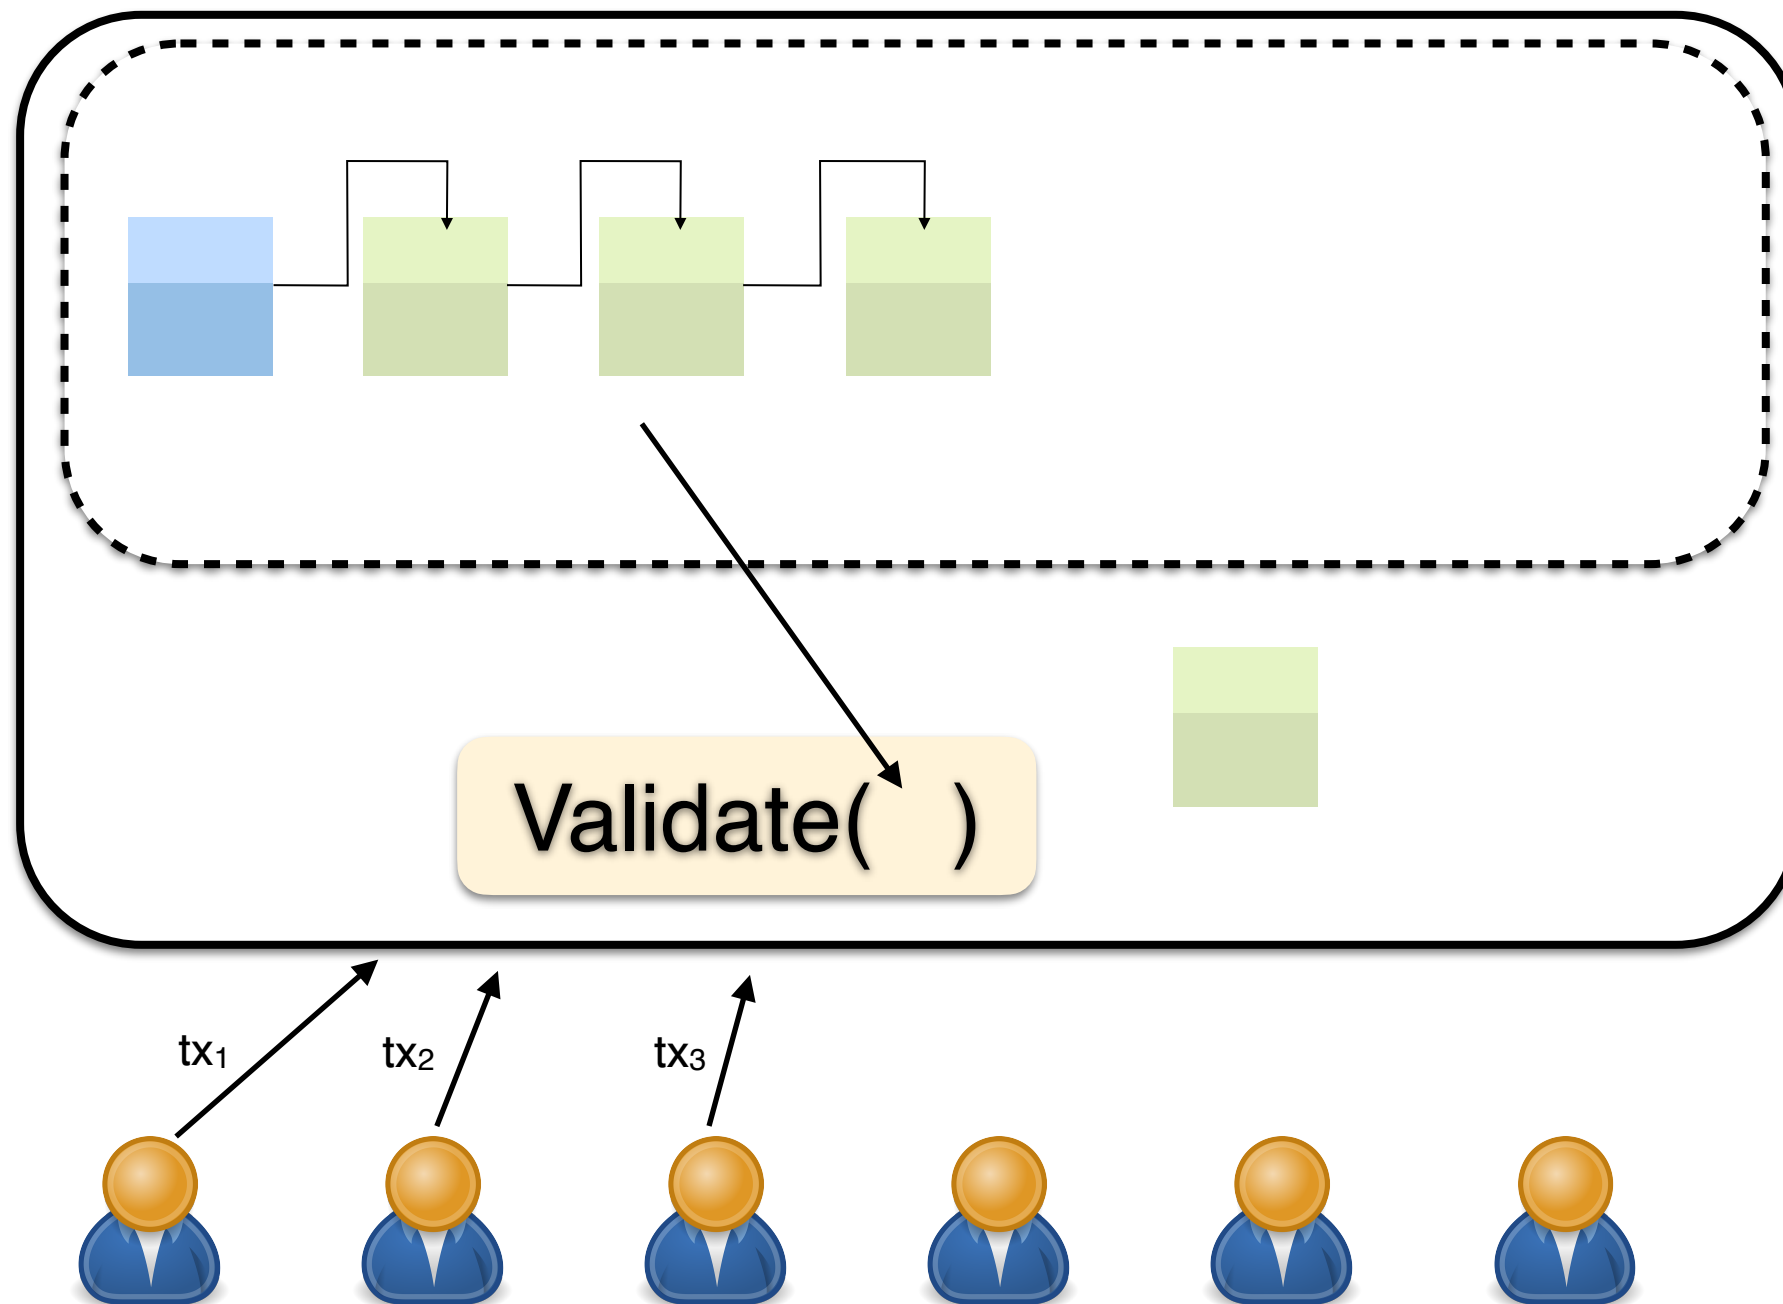
Blockchain-based Transaction Ledger

The Decentralization Paradigm: Bitcoin

Blockchain Ledger


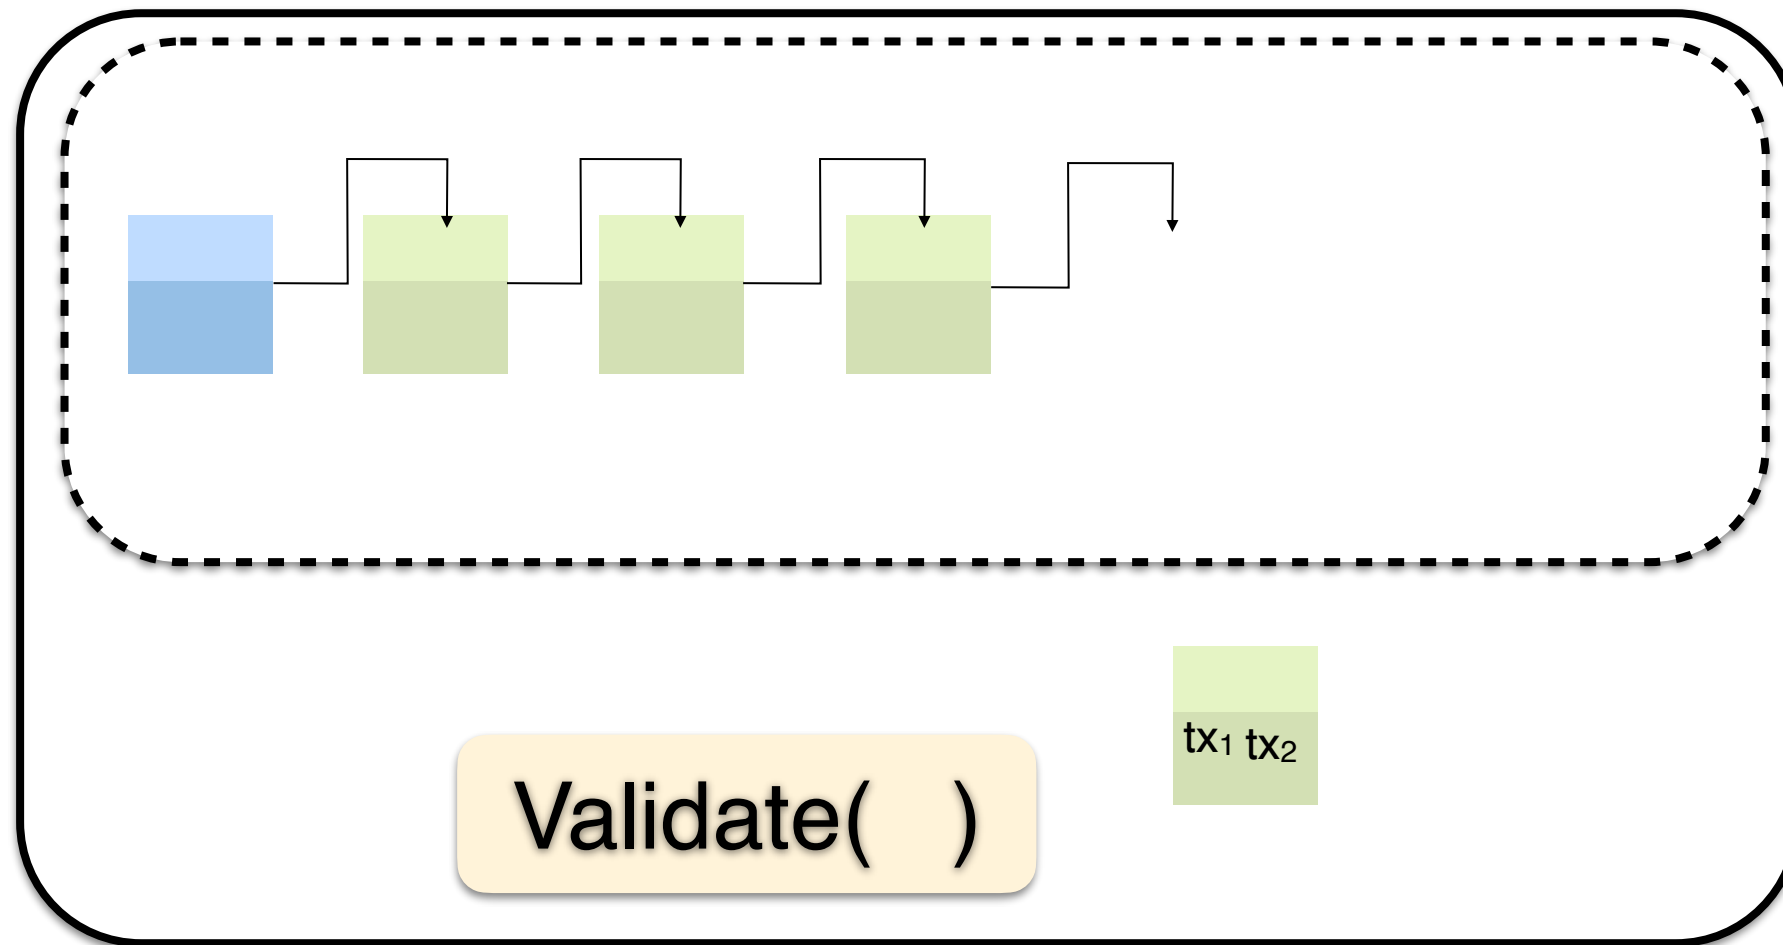
The Decentralization Paradigm: Bitcoin

Blockchain Ledger




The Decentralization Paradigm: Bitcoin

Blockchain Ledger




The Decentralization Paradigm: Bitcoin

Blockchain Ledger




[Badertscher-Maurer-Tshudi-**Z**: CRYPTO17]: Detailed Specification of Bitcoin Ledger:

- Outputs might be prefixes of one another
- A transaction might not make it in the immediately next block




The Decentralization Paradigm: Bitcoin



Methodology

1. Specification
2. System Model
3. Distributed protocol
4. Assumptions and security analysis

The Decentralization Paradigm: Bitcoin



Methodology

1. Specification
2. System Model
3. Distributed protocol
4. Assumptions and security analysis

The Decentralization Paradigm: Bitcoin

Computing Infrastructure:


- Simple devices that can hash, do simple computation, store, communicate

Communication:

- **Challenges:**
 - Unknown topology (ad hoc): no direct all to all links
 - *Dynamic Participation*
- **Solution:** Gossiping over the Internet



The Decentralization Paradigm: Bitcoin



Methodology

1. Specification
2. System Model
3. Distributed protocol
4. Assumptions and security analysis


The Decentralization Paradigm: Bitcoin




The Decentralization Paradigm: Bitcoin




The Decentralization Paradigm: Bitcoin



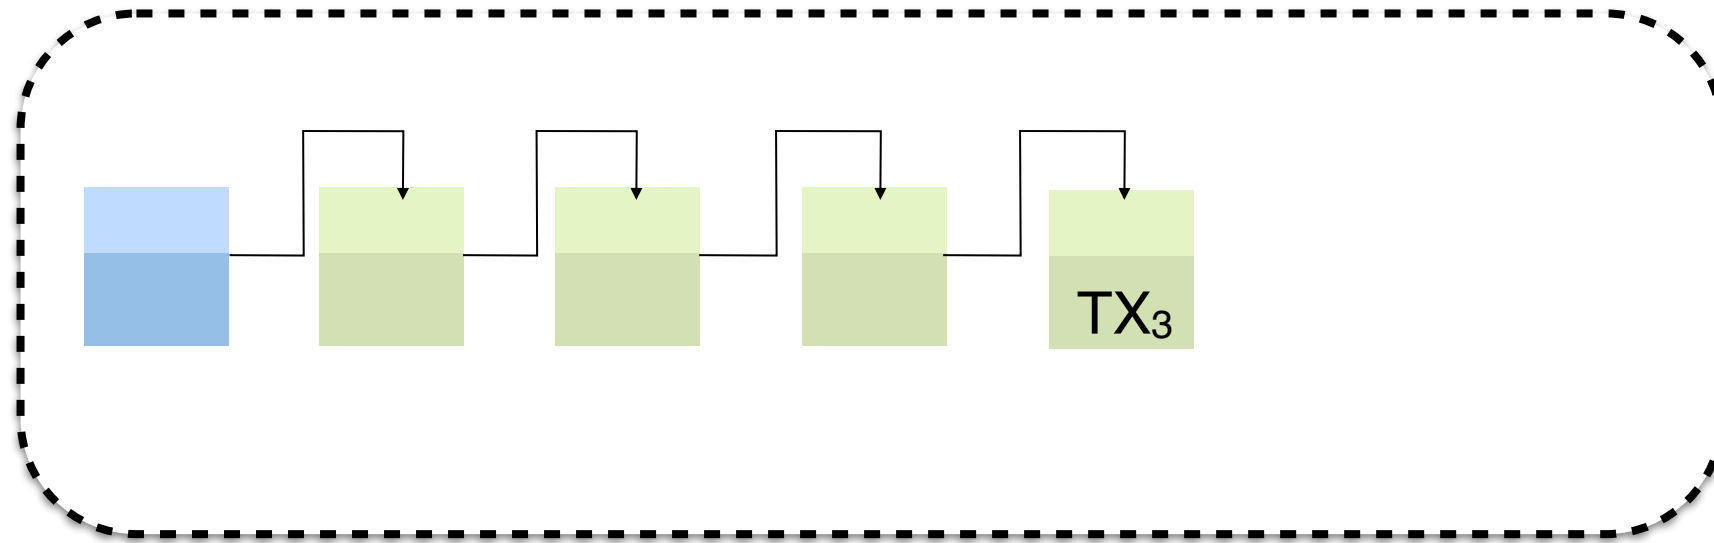
The Decentralization Paradigm: Bitcoin



The Decentralization Paradigm: Bitcoin



The Decentralization Paradigm: Bitcoin



The Decentralization Paradigm: Bitcoin

Decentralizing a Blockchain Ledger

Key Challenges:

- Who can propose the next block?



- *Let's choose randomly*
 - From what population?

- How to agree on the blockchain-ledger state?

The Decentralization Paradigm: Bitcoin

Decentralizing a Blockchain Ledger

Key Challenges:

- Who can propose the next block? ←



- *Let's choose randomly from the unspent coin keys*
 - Sybil attack!


sn ₁	sn ₂	(1,sn ₃)	(0.5,sn ₄)	(0.2,sn ₅)	(0.8,sn ₆)	...
pk _{Alice}	pk _{Bob}	pk _{Charlie}	pk _{Erica}	pk _{Lisa}	pk _{Bob}	...

- How to agree on the blockchain-ledger state?


The Decentralization Paradigm: Bitcoin

Decentralizing a Blockchain Ledger

Key Challenges:

- **Who can propose the next block?** 
 - Whoever solves a hash puzzle (proof of work)
- **How to agree on the blockchain-ledger state?**

The Decentralization Paradigm: Bitcoin



Hash(Prev. Block, nonce)

01001...0101

01101...0101







0001...0101

0001...0101


1001...0101

1001...0001

1001...0001



The Decentralization Paradigm: Bitcoin



PoW: Chances to find a solution proportional to the number of hashes

$$\text{Hash}(\text{Prev. Block}, \text{nonce}) < T$$

01011...1101

00101...0101








0011...0000

0100...0101


1001...0011

1011...0011


1001...1011



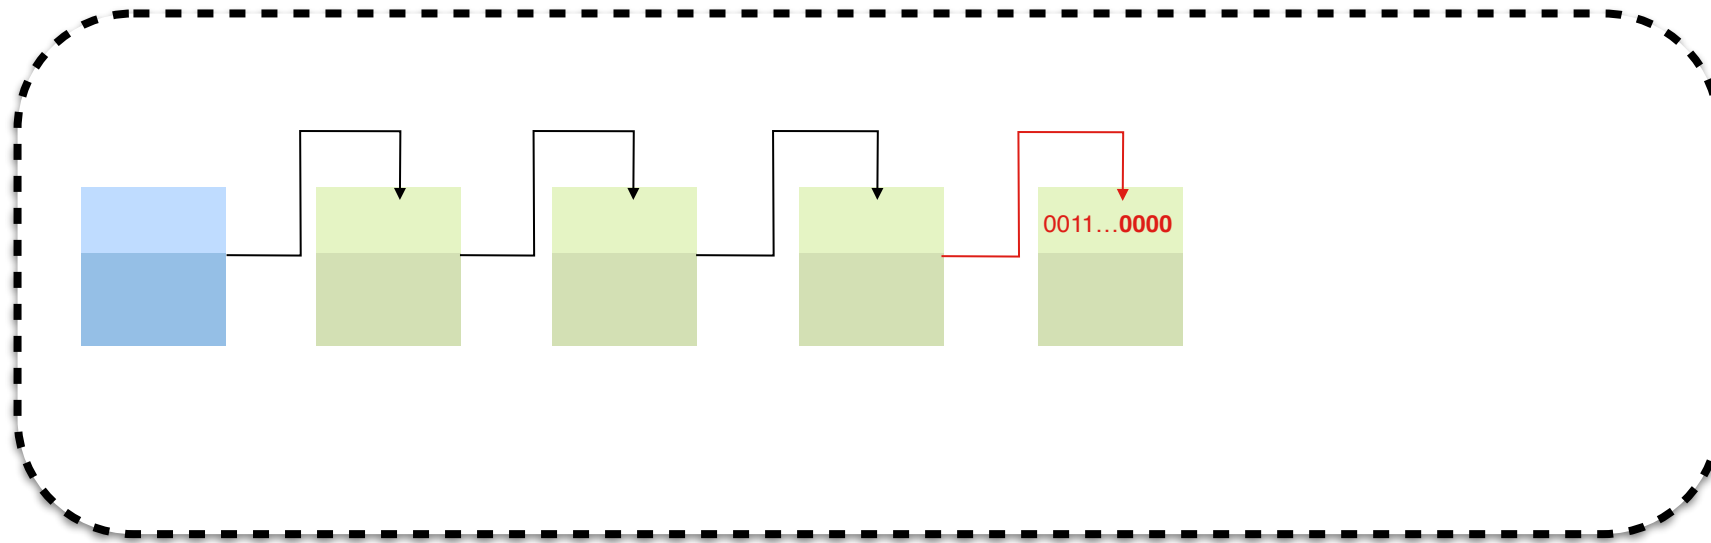
The Decentralization Paradigm: Bitcoin



The Decentralization Paradigm: Bitcoin




The Decentralization Paradigm: Bitcoin



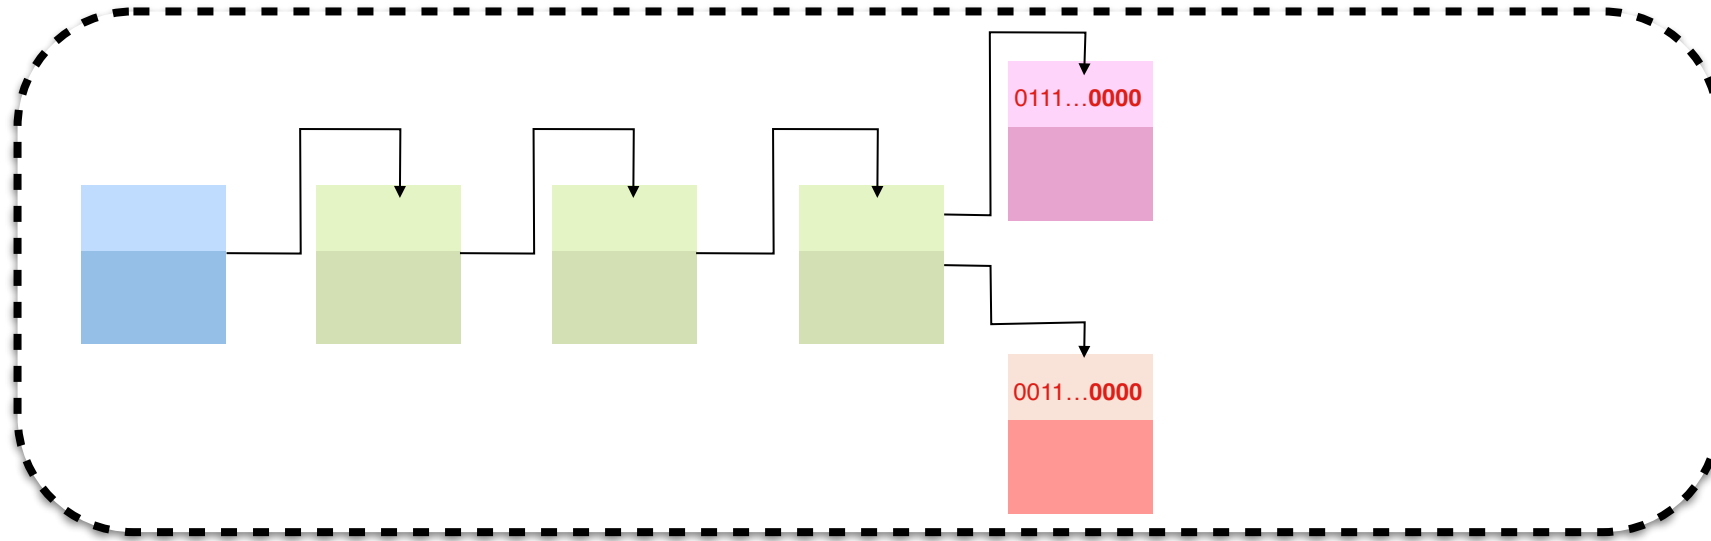
The Decentralization Paradigm: Bitcoin

Decentralizing a Blockchain Ledger

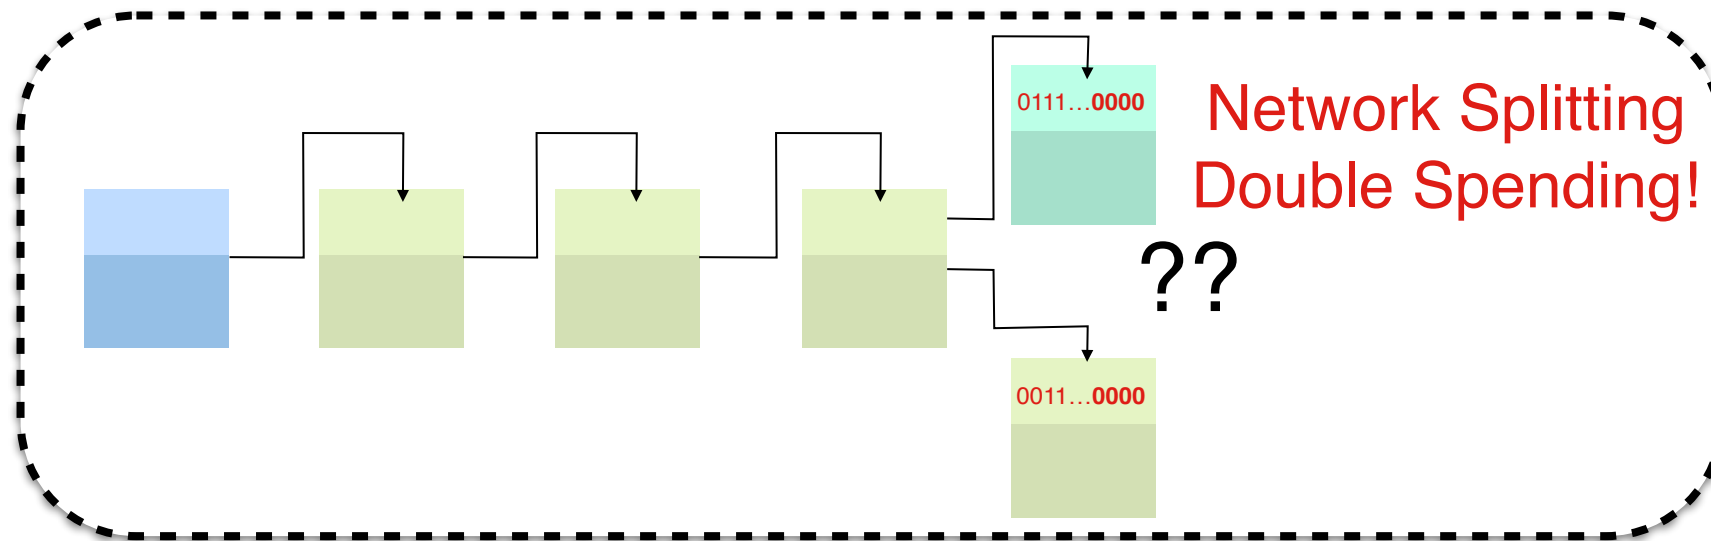
Key Challenges:

- **Who can propose the next block?**
 - Whoever solves a hash puzzle (proof of work)
 - Chances to win proportional to # of attempts (work)
- **How to agree on the blockchain-ledger state?** 

The Decentralization Paradigm: Bitcoin




The Decentralization Paradigm: Bitcoin



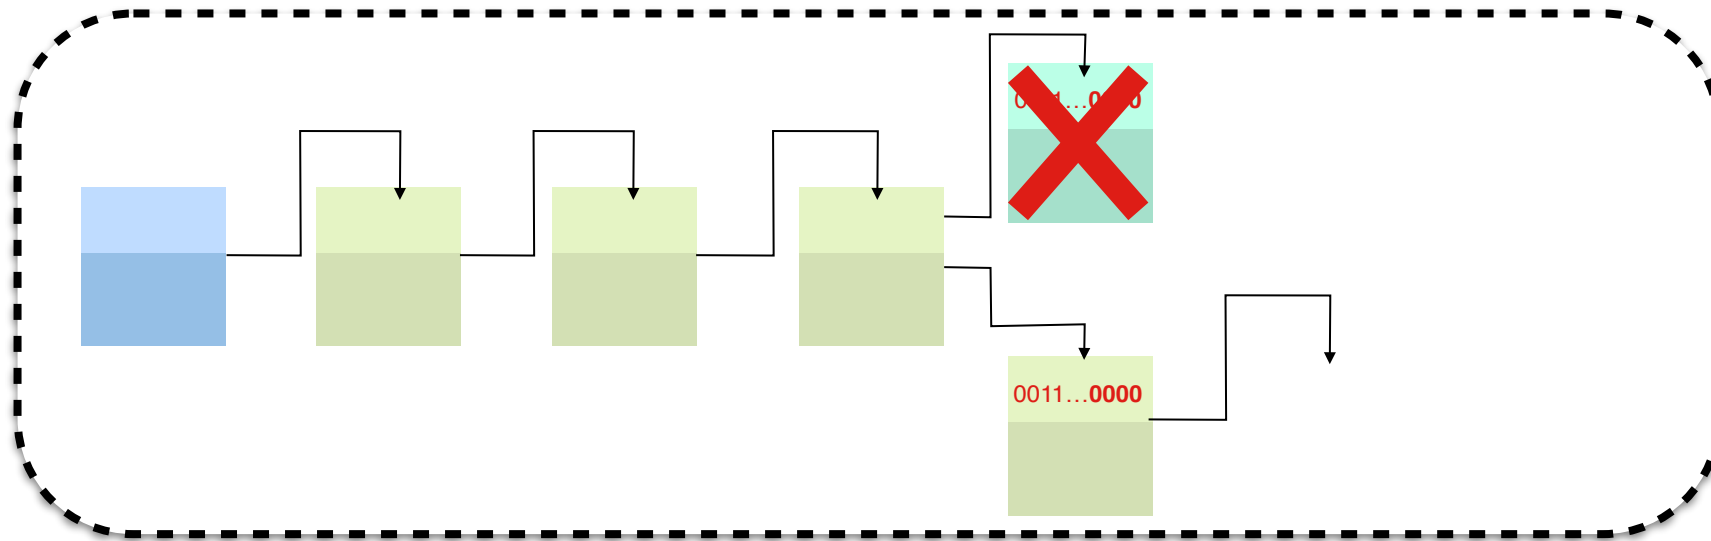
The Decentralization Paradigm: Bitcoin

Decentralizing a Blockchain Ledger


Key Challenges:

- **Who can propose the next block?**
 - Whoever solves a hash puzzle (proof of work)
 - Chances to win proportional to # of attempts (work)
- **How to agree on the blockchain-ledger state?** 
 - Whenever in doubt between two chains adopt the longer (the one with most work)

The Decentralization Paradigm: Bitcoin




The Decentralization Paradigm: Bitcoin





I don't like the last block




The Decentralization Paradigm: Bitcoin



The Decentralization Paradigm: Bitcoin



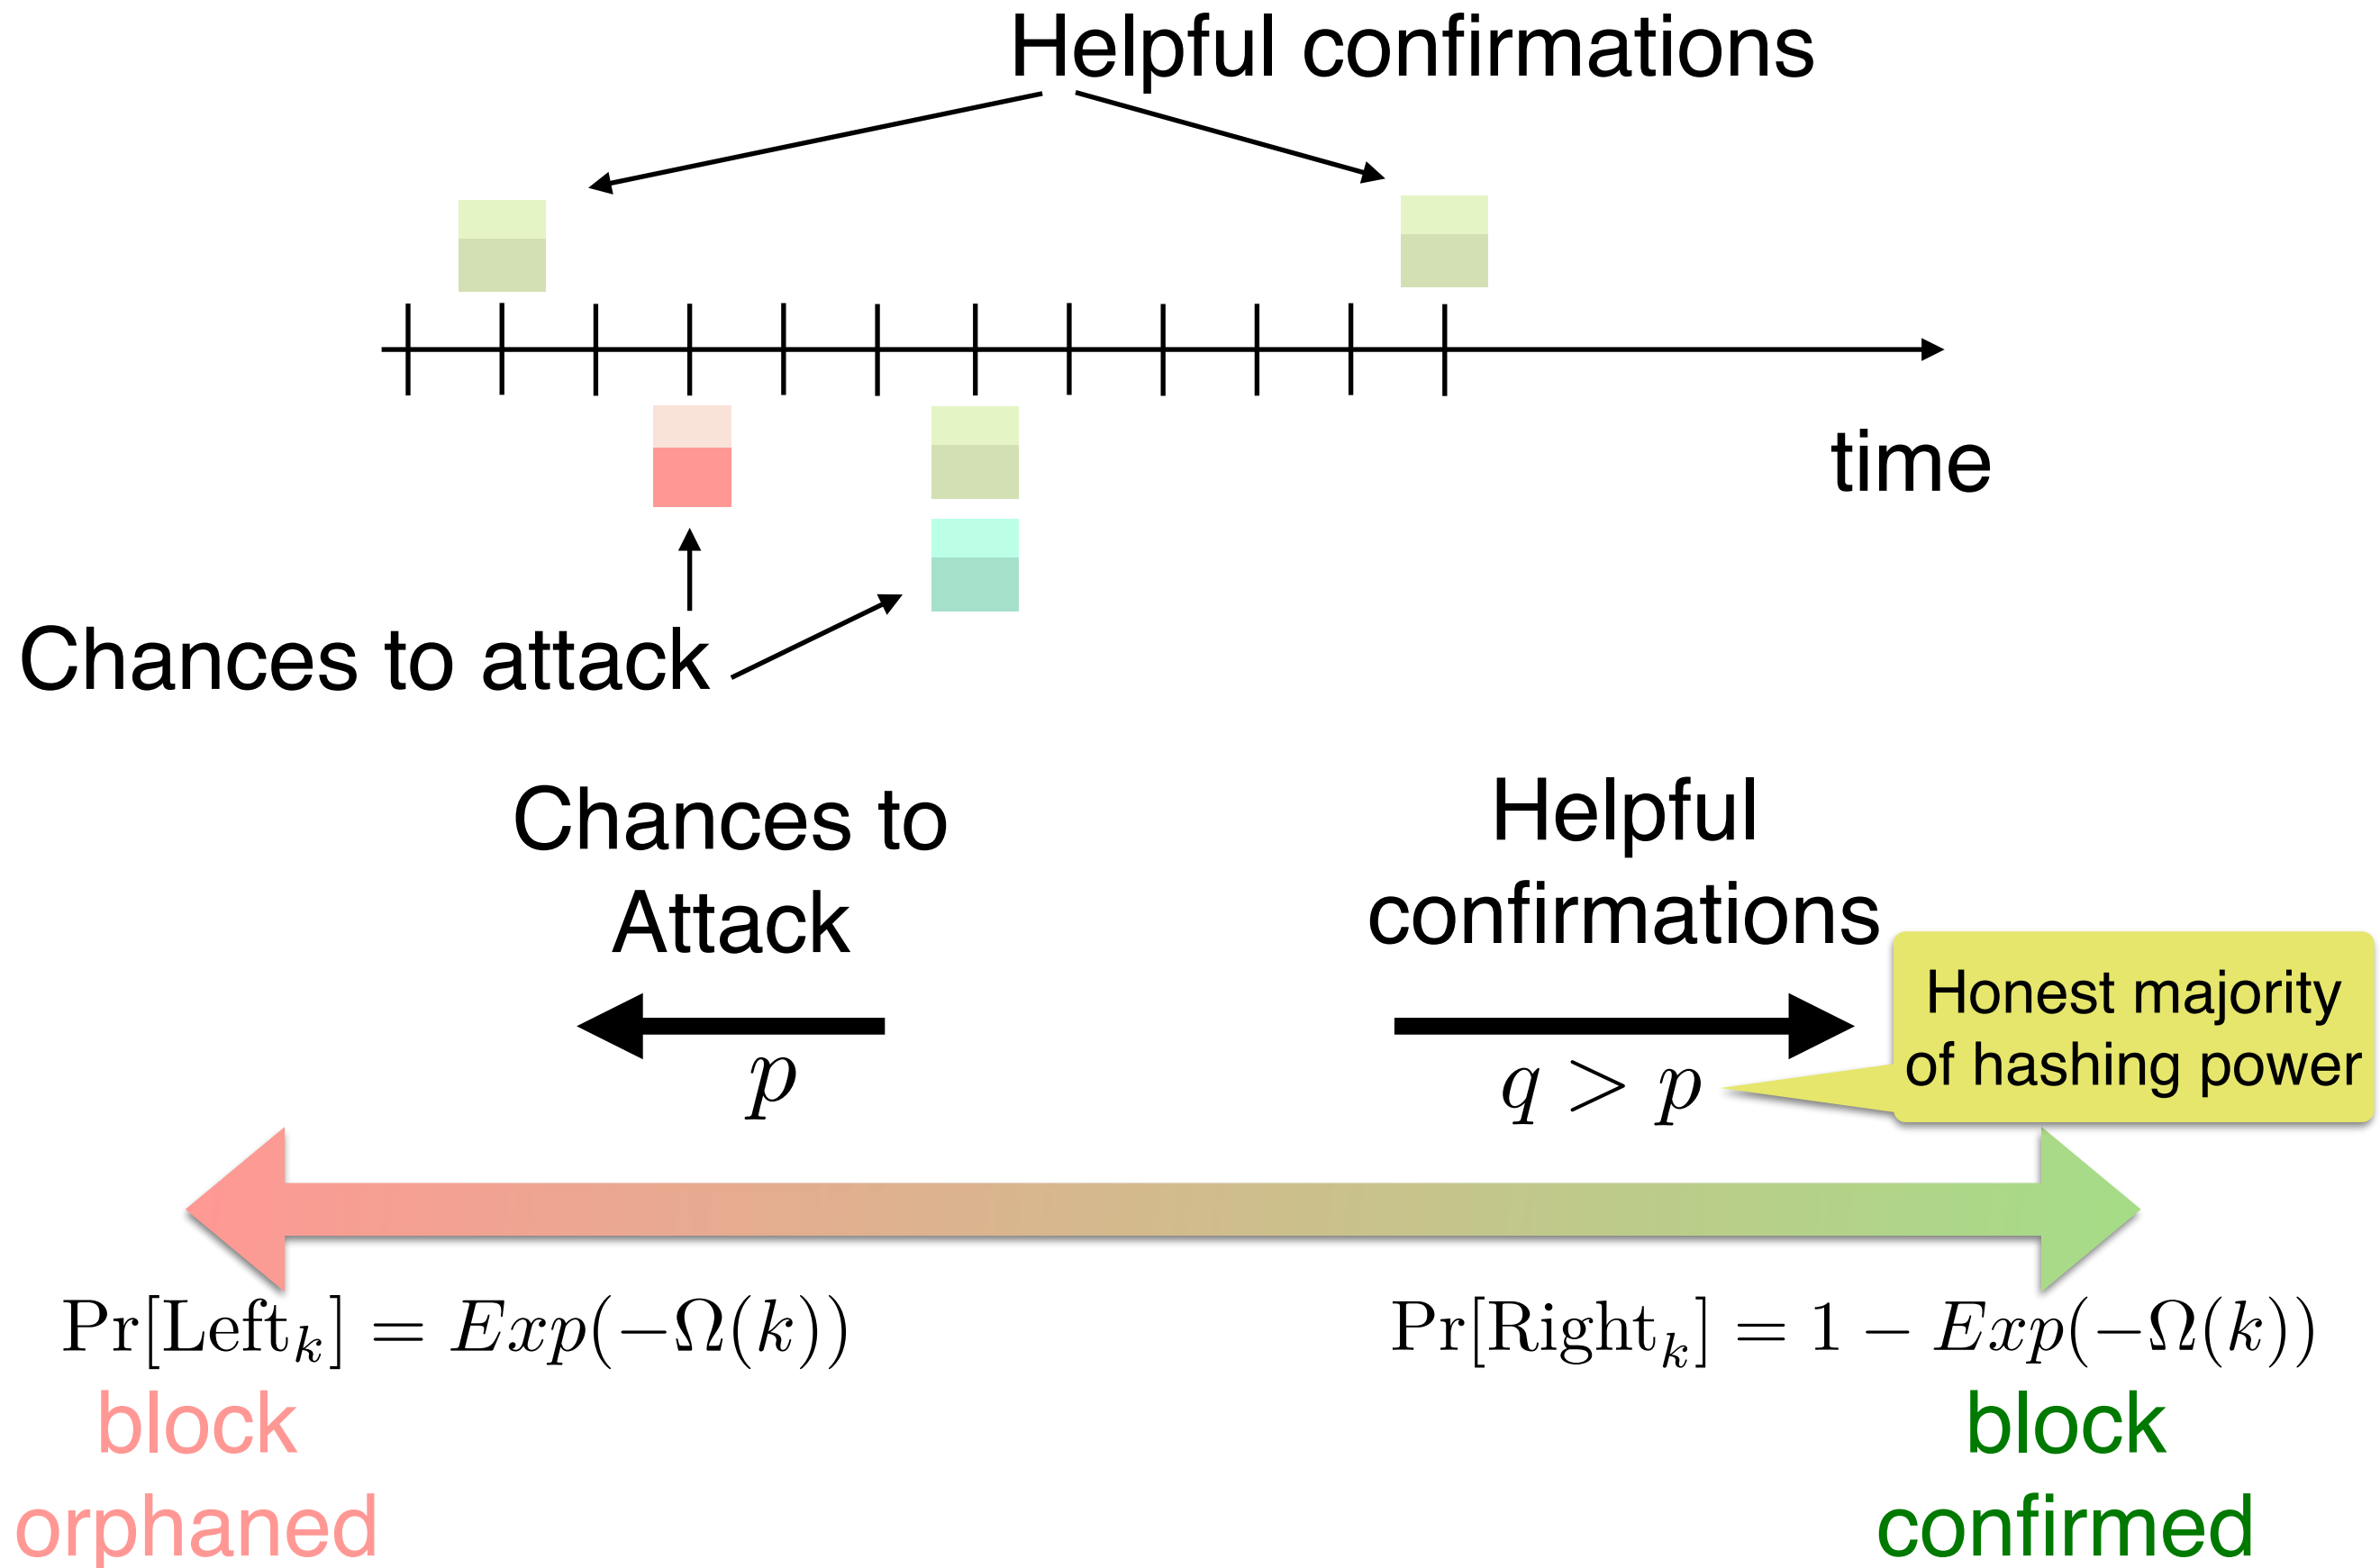
The Decentralization Paradigm: Bitcoin



Honest majority of hashing power =>
Deep enough blocks cannot be orphaned



The Decentralization Paradigm: Bitcoin



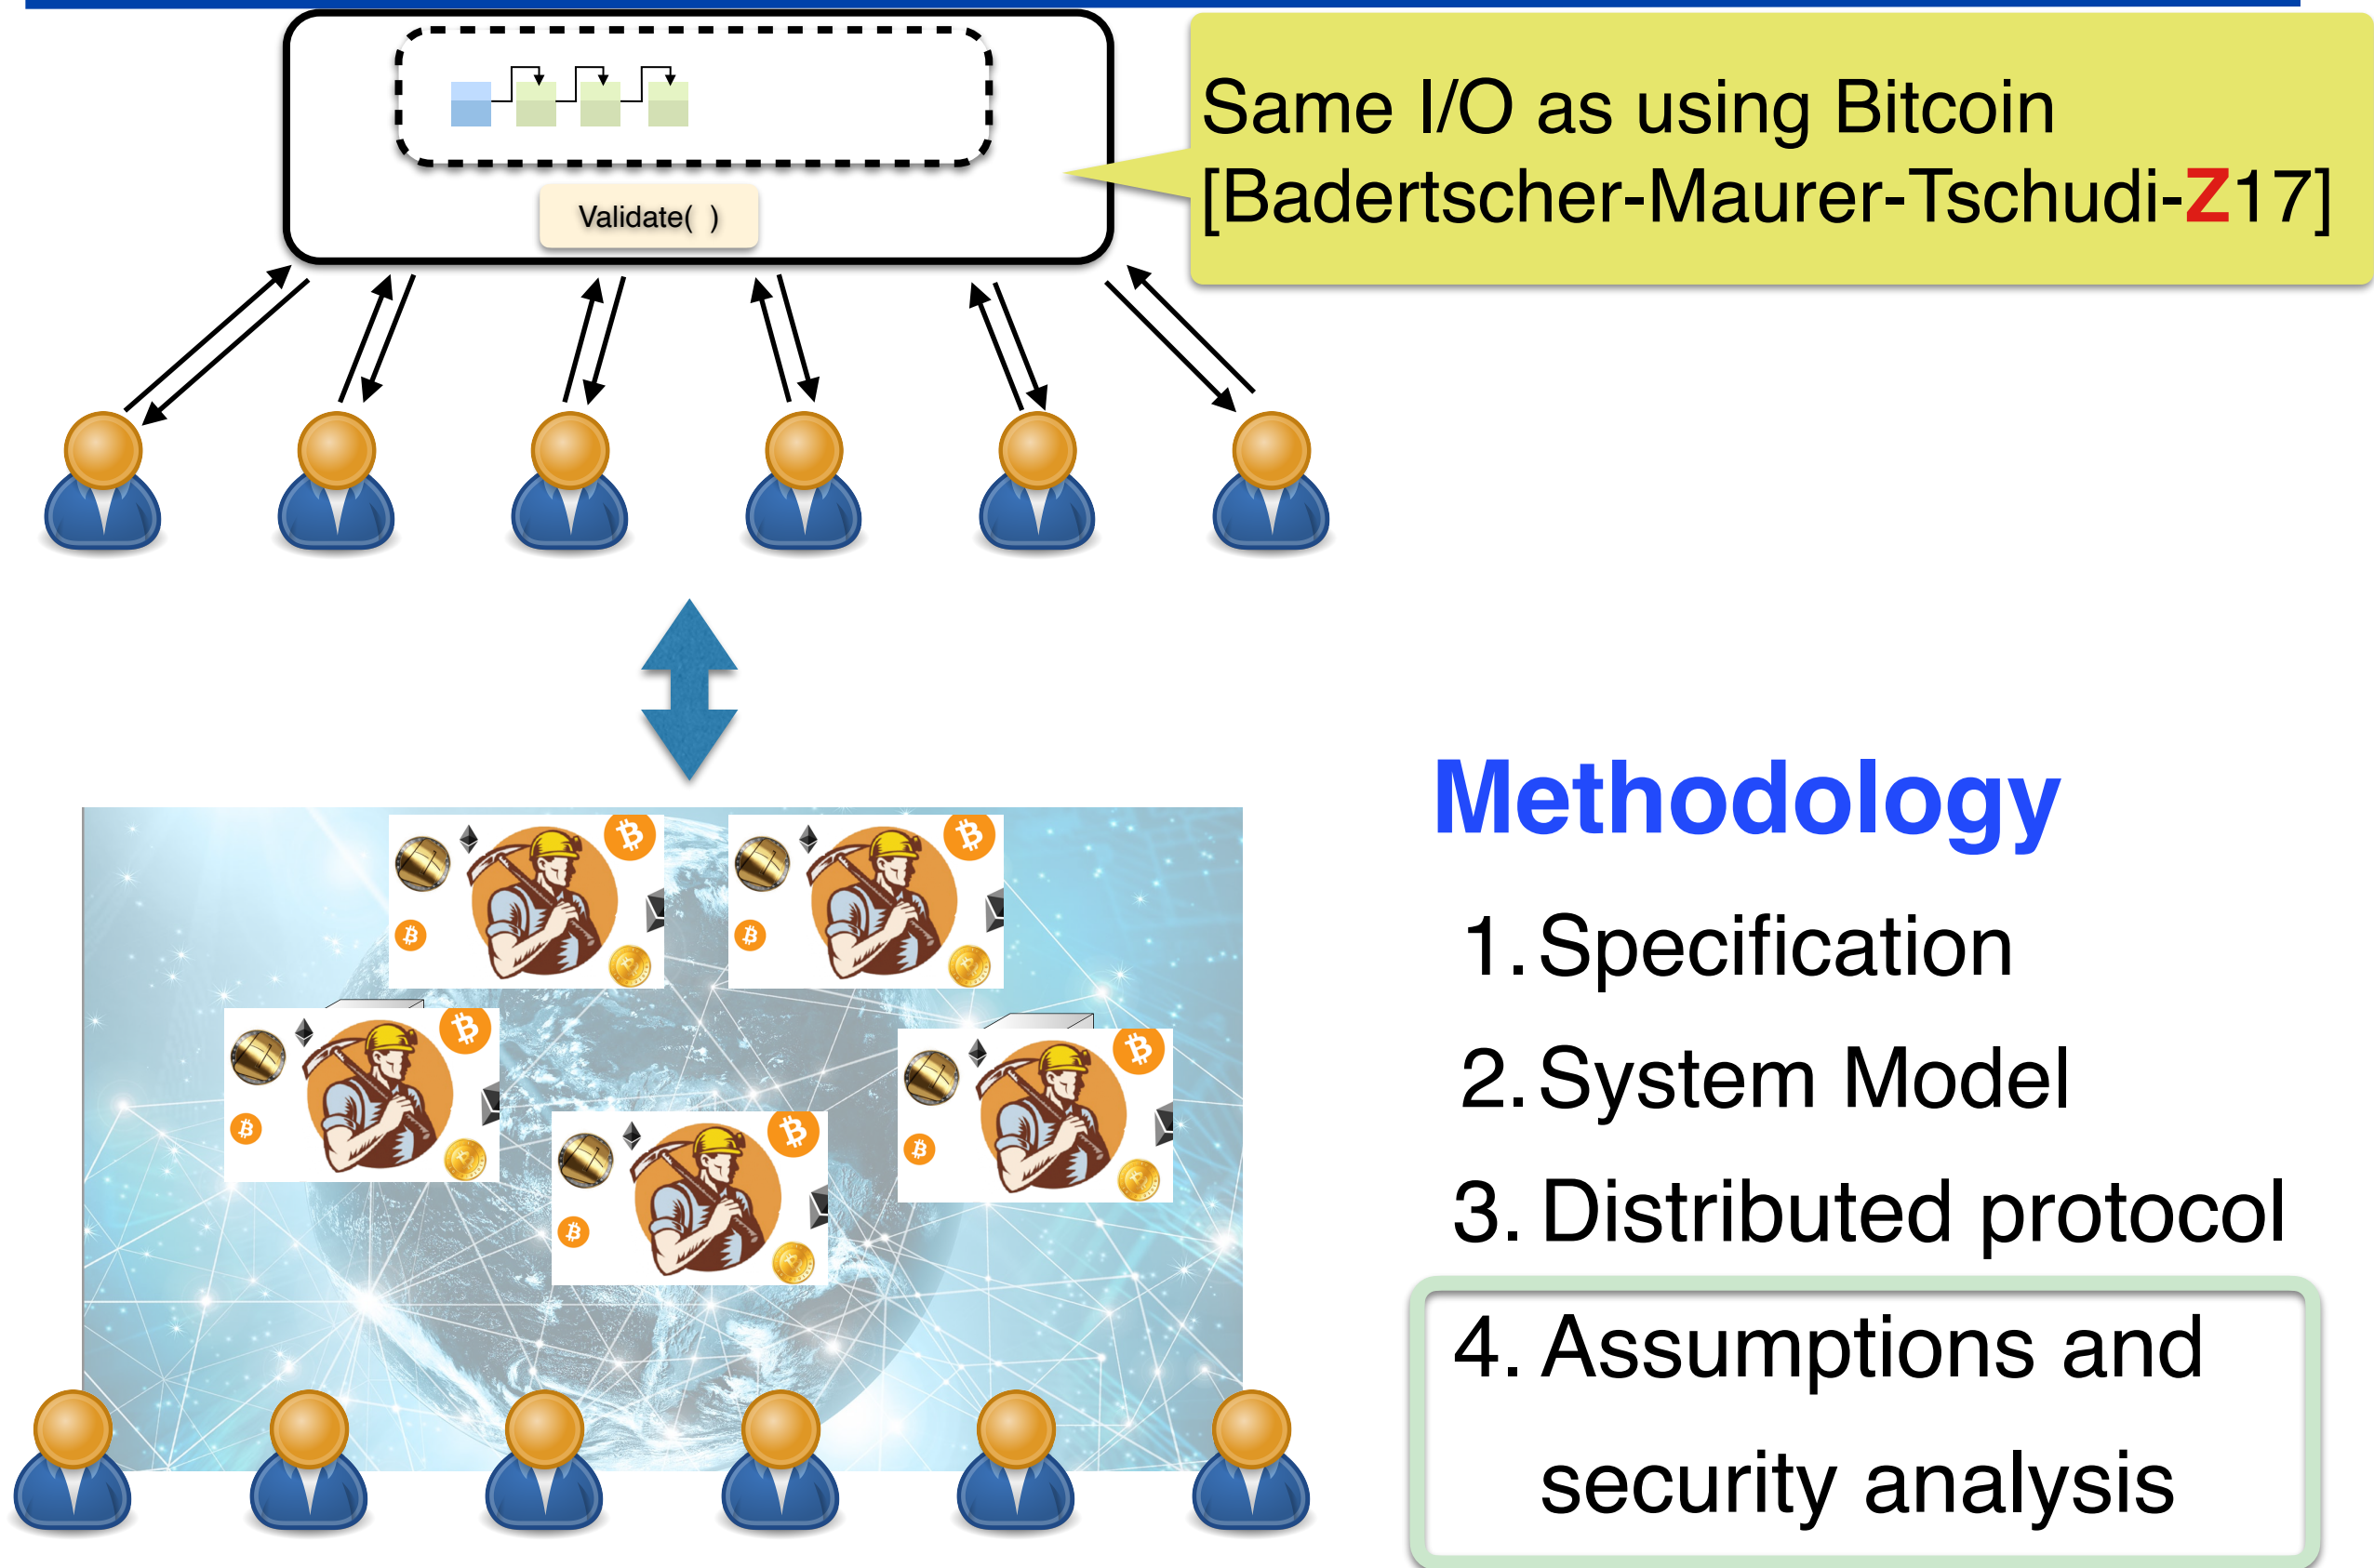
The Decentralization Paradigm: Bitcoin

That was over-simplified ...

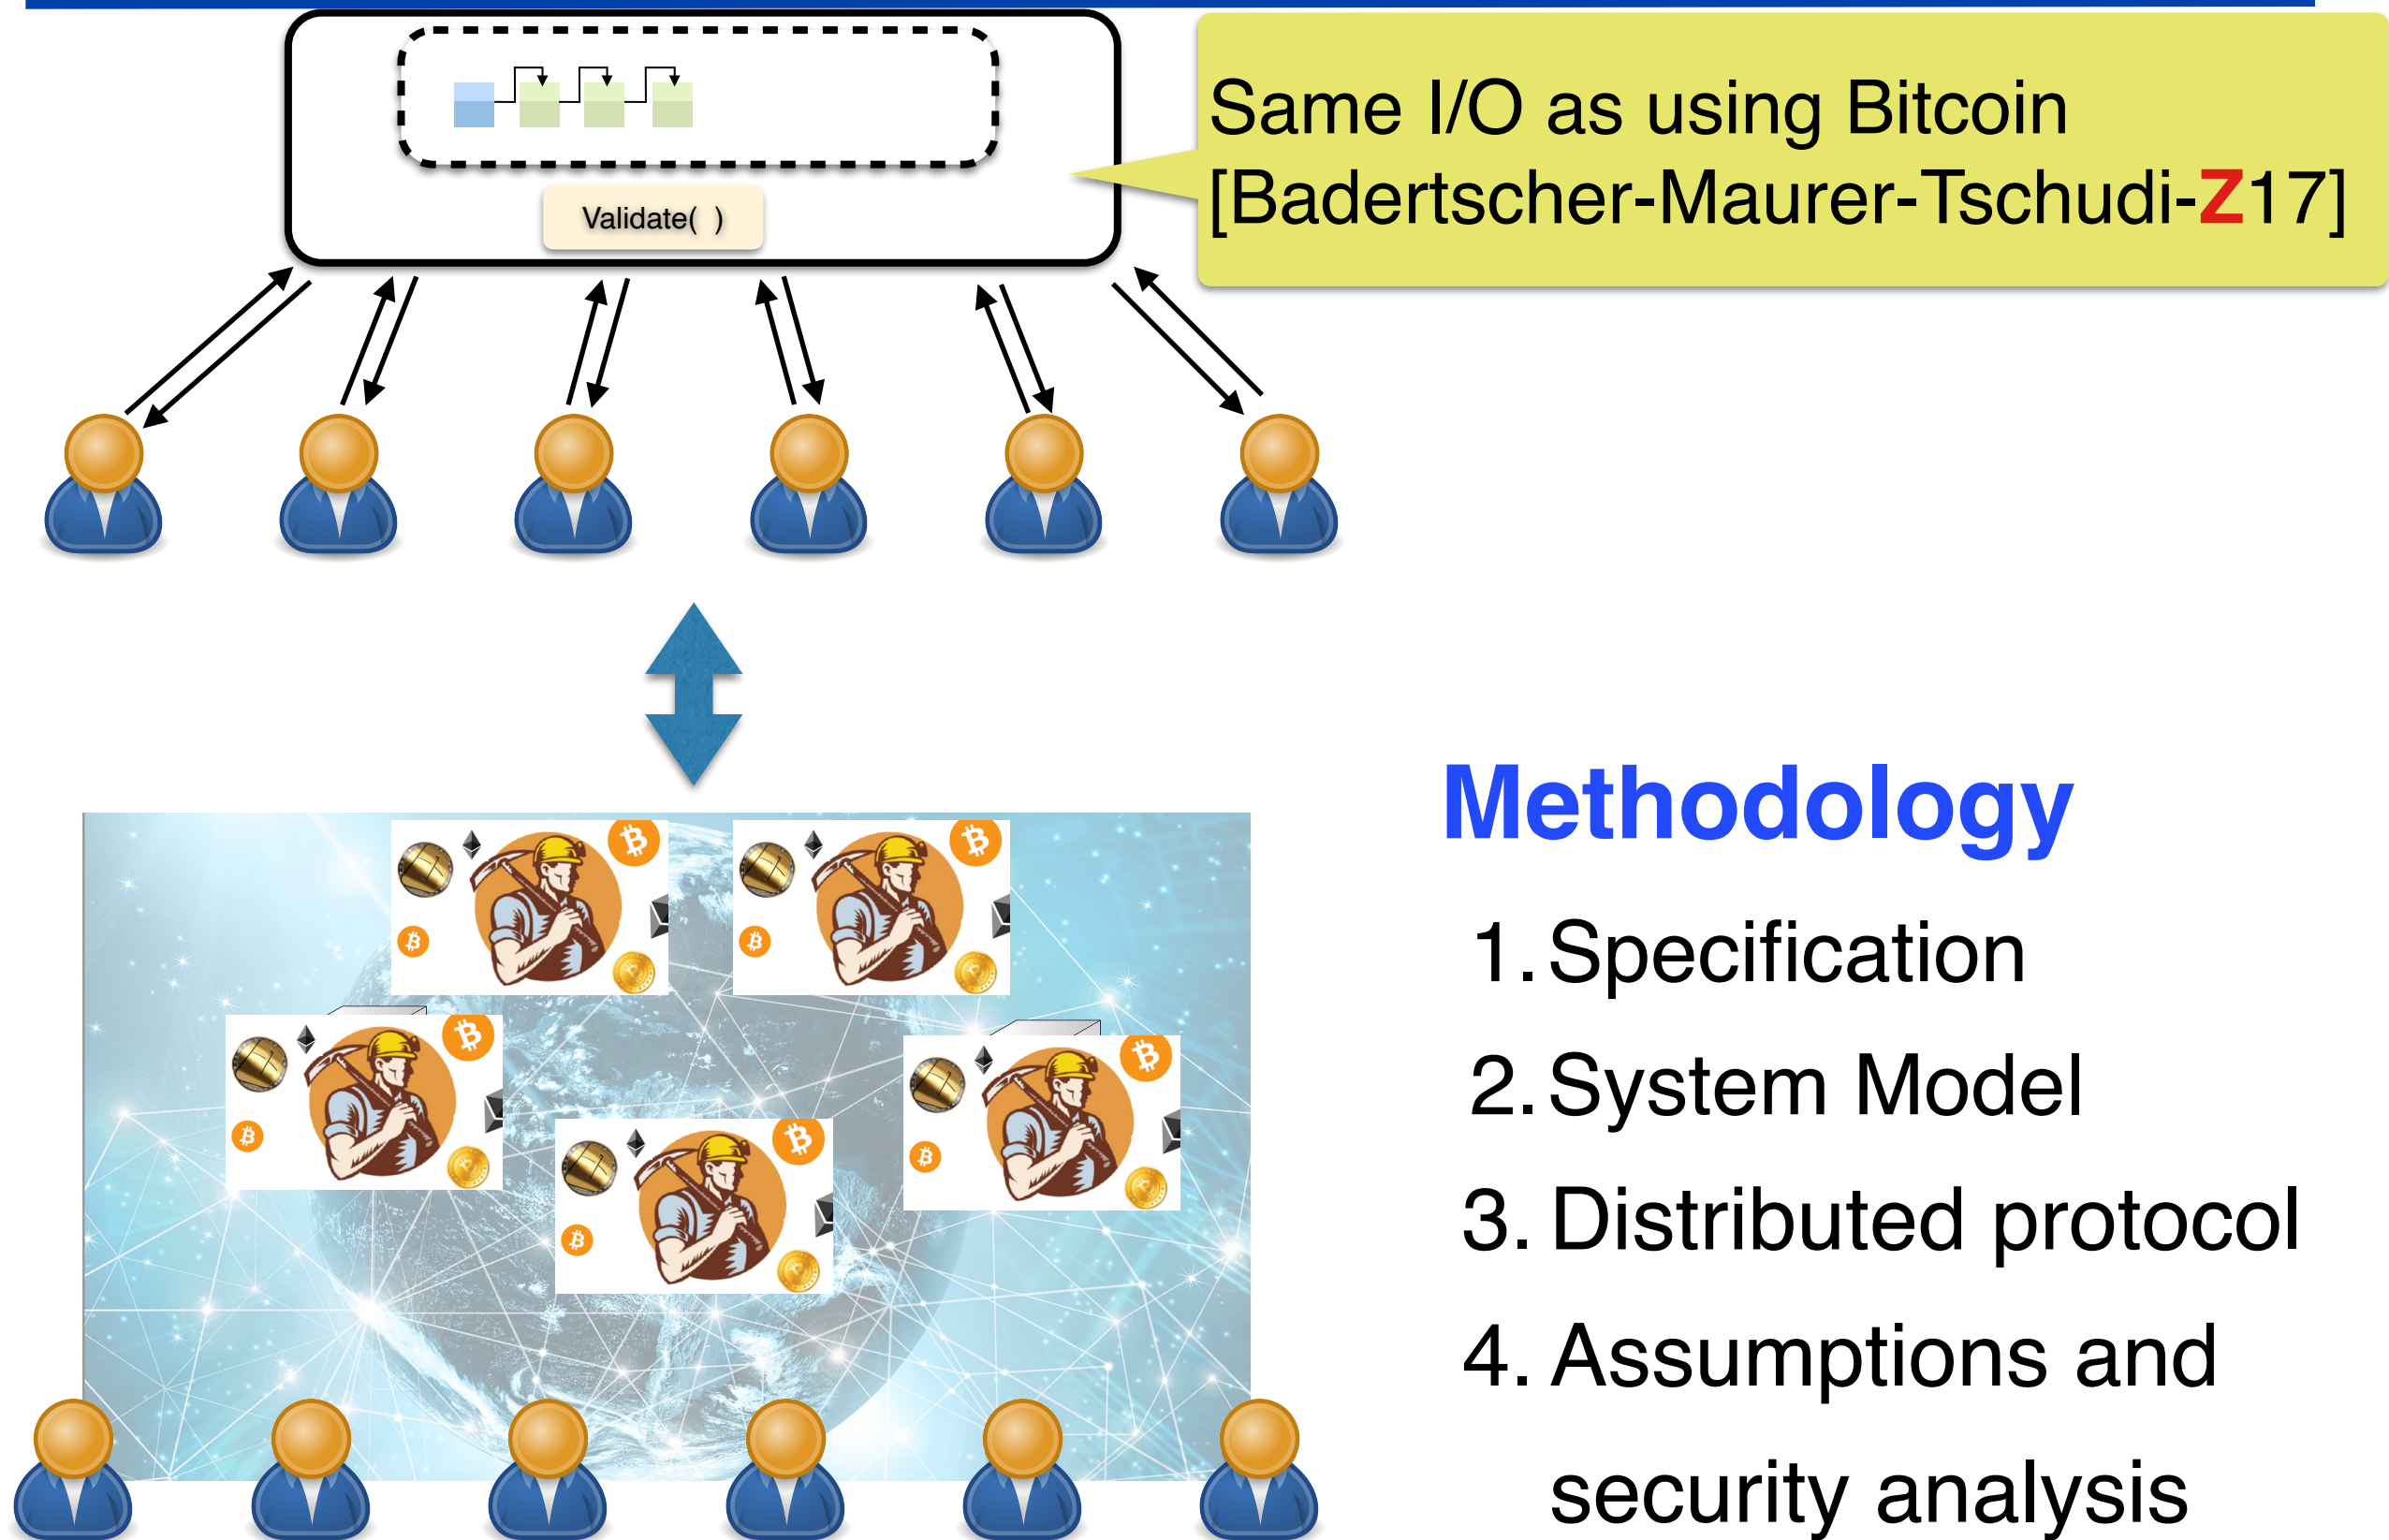
Actual security analysis is very delicate:

- Block withholding (e.g., selfish mining) reduces honest miners' "effective" hashing
 - When is it optimal to announce?
- Network delay provide opportunities to attacker
 - Buy time to work on a block
 - Temporarily split the network
- Difficulty Readjustment

The Decentralization Paradigm




The Decentralization Paradigm



Methodology

1. Specification
2. System Model
3. Distributed protocol
4. Assumptions and security analysis


The Decentralization Paradigm



Science: Cross-disciplinary research



- Economics
 - Rational analysis of crypto [Badertscher-Lu-**Z**: EUROCRYPT'20]
 - Incentive-compatible online-poker [Ciampi-Lu-**Z**: CSF'20]
- AI/ML: Trusted model outsourcing
- Law: Regulation of Cryptocurrencies

The Decentralization Paradigm



Systems Engineering:

- API to blockchain ("*ledger.h*")
- Benchmarks specs for different systems



Proof-of-Work Blockchains (Bitcoin)



Economics and Blockchain

Rational Analysis of Blockchains

The world (Reality)

vs.

Crypto(graphy)



Is Bitcoin (blockchain) secure?

So ... what if the adversary gets majority?



It is secure!!!!!! ... **assuming** honest majority of hashing power

not secure ... (or “less secure” if adversarial majority is temporary [BGKR^Z20])

Economic Robustness (A new type of security statements):

Attacking the assumption and/or the underlying system is irrational

- A good fallback of standard cryptographic security
- Makes modern cryptocurrency DLTs more than “just” a ledger!
- Can improve security/understanding of blockchain-aided protocols

Economics and Blockchain

Rational Analysis of Blockchains

...building on Game-Theoretic Security/ Cryptography

[CRYPTO'12, ICALP'12, FOCS'13, PODC15, DISC'15]

- **Framework for Economic Robustness of PoW Blockchains**

[Badertsher-Garay-Maurer-Garay-Tschudi-**Z**: EUROCRYPT 2018]

- Rational Crypto Analysis of the bitcoin backbone protocol

- **Analysis of 51% Attacks on PoW Blockchains**

[Badertsher-Lu-**Z**: CRYPTO 2020]

- Susceptibility and patch

Predicted attack times and cost

- Matches forensics result on the ETC attack(s)

r _{cost} (USD)	\bar{t} (days)	Cost/day (USD)
\$0.0001	24.0	\$78,084
\$0.0002	10.5	\$156,167
\$0.0003	4.3	\$234,251
\$0.0004	3.2	\$312,334
\$0.0005	2.6	\$390,418
\$0.0006	2.1	\$468,501

- **DeFi Application: FairMM: Front-running resistant crypto exchanges**

[Ciampi-Ishaq-MagdonIsmail-Ostrovsky-**Z**: CSCML 2022]

- A front-running resistant DEX

Feature	FairMM	Uniswap
Front Running Resilience	Yes	No
Gas Price Auctions	No	Yes
Miner Influence	No	Yes
Trade Execution (seconds)	≈ 0.30	≥ 15
Average Trade Cost (K)	≈ 101	$\approx 141^*$
Max Trade Cost (K)	≈ 101	$\approx 1,316^\dagger$
Max Throughput \ddagger	≈ 475	≈ 340

The Purdue Blockchain Lab (The Pub)



Blockchain Lab


[Home](#) [Research](#) [Publications](#) [People](#) [News](#) [Events](#) [Contact](#)

Purdue Blockchain

The PuB's mission is to investigate and extend the foundations and applications of blockchains and decentralized ledger technology. The lab's approach is a practice-driven methodology to solving pressing problems in the blockchain ecosystem, by means of interdisciplinary research combining ideas from cryptography, distributed computing and systems, and game theory.



The Purdue Blockchain Lab (The Pub)



People

Lab Director



Vassilis Zikas

Associate Professor of Computer Science

[Vassilis' website](#)

<https://www.cs.purdue.edu/blockchain>



<https://twitter.com/PBlockchainLab>

Affiliated Faculty:



Jeremiah Blocki

Assistant Professor of Computer Science

[Jeremiah's website](#)



Aniket Kate

Associate Professor of Computer Science

[Aniket's website](#)



Christina Garman

Assistant Professor of Computer Science

[Christina's website](#)



Alex Psomas

Assistant Professor of Computer Science

[Alex's website](#)

Thank you!

MEGA-ACE: Purdue Partners with NTUA

