

CS251 Fall 2021

(<https://cs251.stanford.edu>)



Cryptocurrencies and Blockchain Technologies

Dan Boneh

Benedikt Bünz

Stanford University

[videos on canvas, discussions on edstem, homework on gradescope]

What is a blockchain?

Abstract answer: a blockchain provides

- coordination between many parties,
- when there is no single trusted party

if trusted party exists \Rightarrow no need for a blockchain

[financial systems: often no trusted party]

What is all the excitement about?

- (1) Basic application: a digital currency (stored value)
- Current largest: Bitcoin (2009), Ethereum (2015)
 - Global: accessible to anyone with an Internet connection

Opinion The New York Times

Bitcoin Has Saved My Family

“Borderless money” is more than a buzzword when you live in a collapsing economy and a collapsing dictatorship.

By Carlos Hernández
Mr. Hernández is a [Venezuelan economist](#).

Feb. 23, 2019

What is all the excitement about?

(2) Beyond stored value: **decentralized applications (DAPPs)**

- **DeFi**: financial instruments managed by public programs
 - examples: stablecoins, lending, exchanges,
- **Asset management (NFTs)**: art, game assets, domain names.
- **Decentralized organizations (DAOs)**: (decentralized governance)
 - DAOs for investment, for donations, for collecting art, etc.

(3) New programming model: writing decentralized programs

Assets managed by DAPPs

Total Value Locked (USD) in DeFi

TVL (USD)




Sep. 2021



Transaction volume

24h volume

Sep. 2021

 Bitcoin • BTC	\$30.6B
 Ethereum • ETH	\$19.2B
 Cardano • ADA	\$2.3B

Central Bank Digital Currency (CBDC)

China Moves Forward With National Digital Currency

by [Sam Klebanov](#) — September 3, 2021

What is a blockchain?

user facing tools (cloud servers)

applications (DAPPs, smart contracts)

compute layer (blockchain computer)

consensus layer

Consensus layer (informal)

A public append-only data structure:

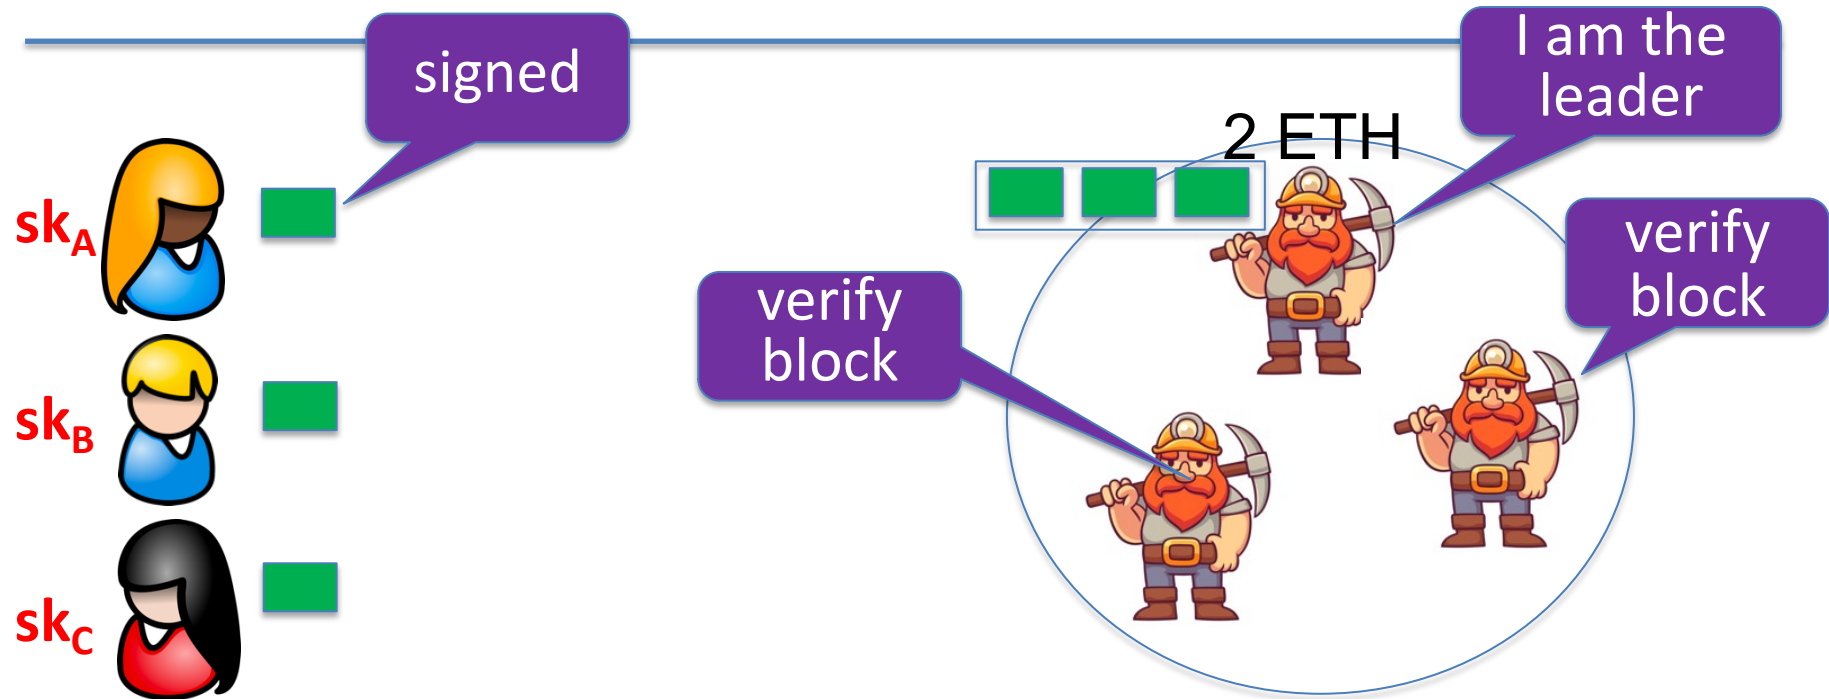
achieved by replication

- **Persistence:** once added, data can never be removed*
- **Safety:** all honest participants have the same data**
- **Liveness:** honest participants can add new transactions
- **Open(?):** anyone can add data (no authentication)

consensus layer

How are blocks added to chain?

blockchain

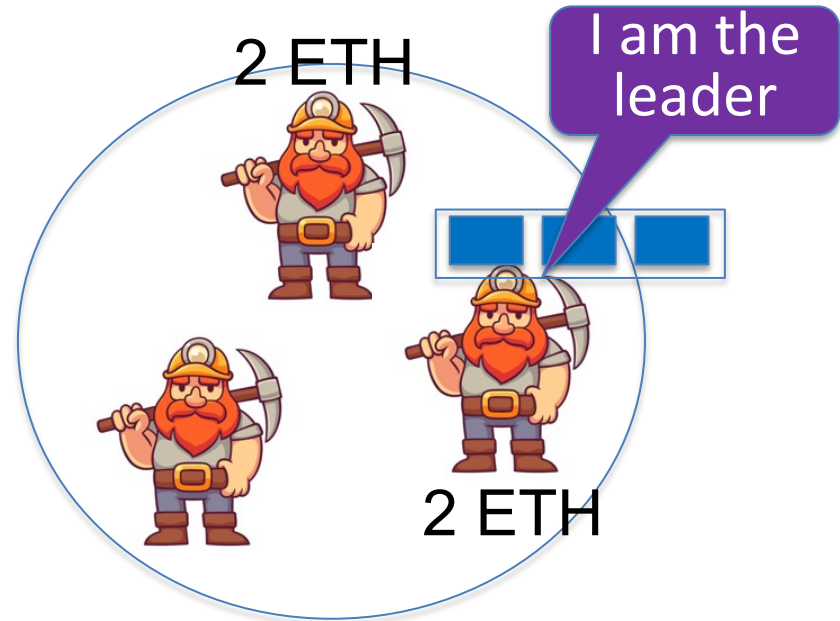
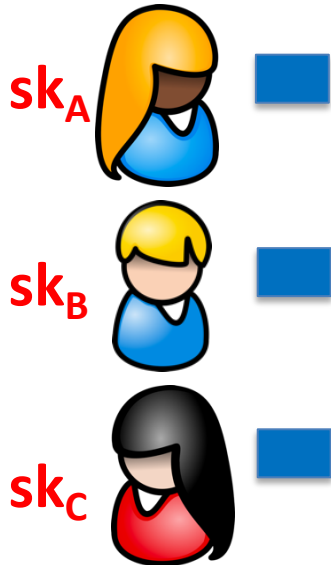


How are blocks added to chain?

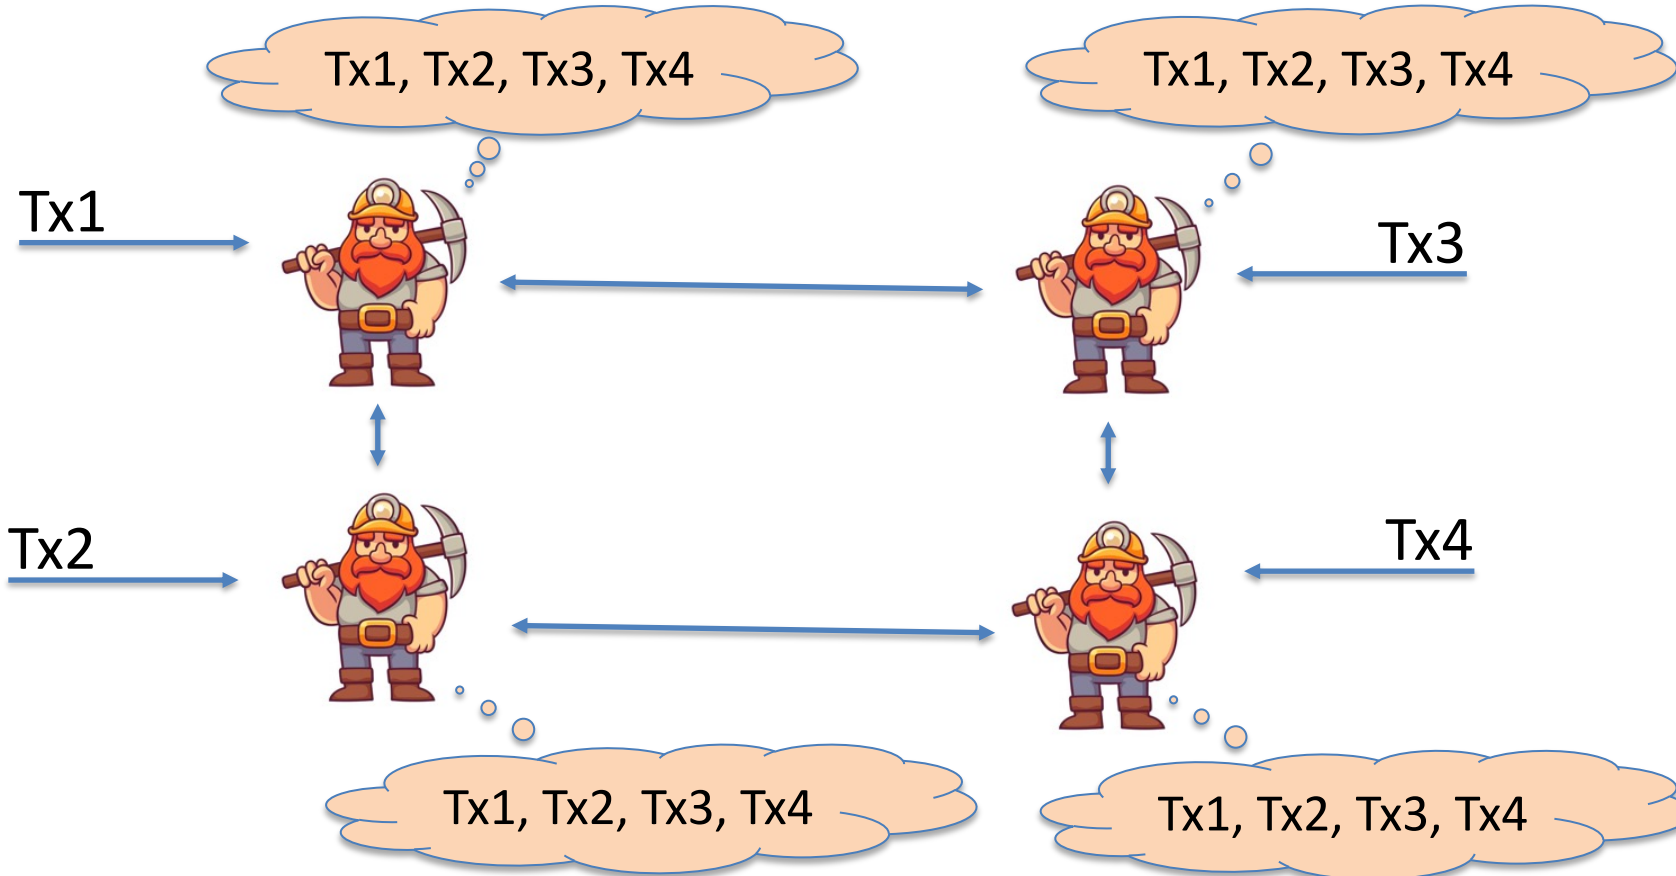
blockchain



...




Why is consensus a hard problem?



Why is consensus a hard problem?

Problems:

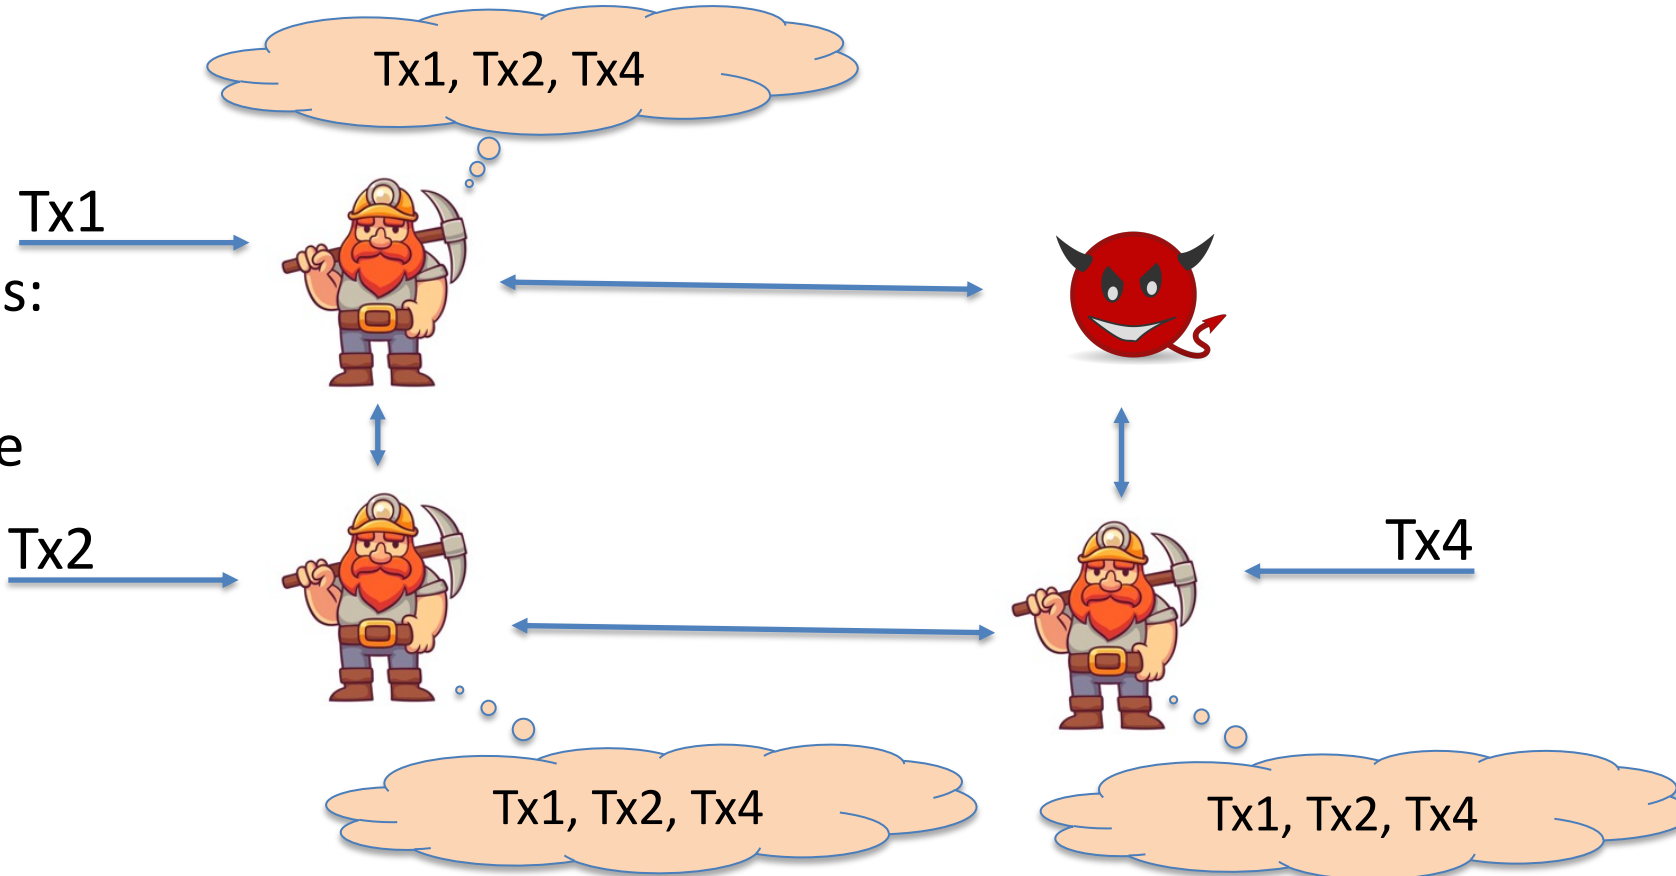
- Network delay
- Network partition



Why is consensus a hard problem?

Problems:

- crash
- malice



The blockchain computer

DAPP logic is encoded in a program that runs on blockchain

- Rules are enforced by a public program (public source code)
⇒ **transparency**: no single trusted 3rd party
- The DAPP program is executed by parties who create new blocks
⇒ **public verifiability**: everyone can verify state transitions

compute layer

consensus layer

Decentralized applications (DAPPS)

Run on
blockchain
computer



applications (DAPPs, smart contracts)

blockchain computer


consensus layer

Common DAPP architecture

Top layer: user facing servers



end user



Ethereum's DeFi



Payments

request network
∞ X Protocol
Dai Card

OPEN PLATFORM

xDai Chain
Groundhog
DAIDEN

Custodial Services

MyEtherWallet
ZERION
argent TRUST WALLET
METAMASK
Balance
MyCrypto

Infrastructure

connext 0xcert
SETTLE GITCOIN
DutchX Ethlance
0x FOAM BOUNTIES NETWORK

Exchanges & Liquidity

Uniswap
AIRSWAP
ForkDelta
slow.trade
RADAR
TOTE hydro LOOPRING
PARADEX Bancor Ren

Investing

Set 22X SWARM
HARBOR FETCH
MELONPORT
Brickblock SPICE bskt MERIDIO
BETOKEN SLICE
SCIENCE BLOCKCHAIN

KYC & Identity

SELFKEY JOLOCOM
sovrin uport
civic Bloom

Derivatives

MARKETPROTOCOL
expo UMA
veil LENDROID
DY/SX DAXIA
b2x VARIABL

Marketplaces

Rare Bits
district0x
ORIGIN
OpenSea

Stablecoins

DAI
USD Coin
GEMINI dollar
StableUnit
PAXOS STANDARD
TrueUSD
CARBON
Reserve
Terra
Ampleforth

Prediction Markets

Guesser augur
Bodhi
STOX veil
GNOSIS

Insurance

ETHERISC
Nexus Mutual
iXledger
VouchForMe
ai gang

Credit & Lending

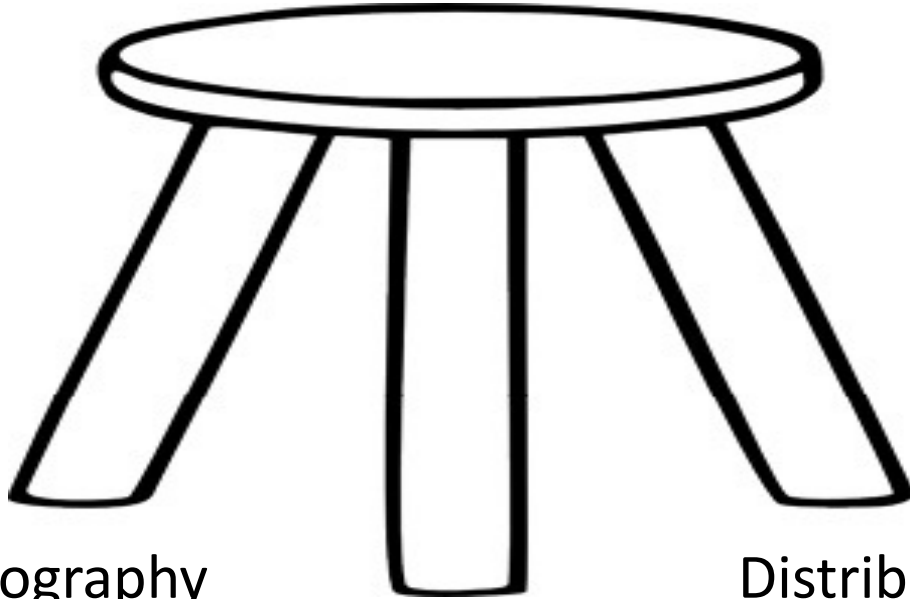
LENDROID
Lendbit
Compound
Ripio Credit Network
Dharma MAKER
ETHLend
nūo InstaDApp
Marble. SALT
BLOQBOARD
COLENDI

[source: the Block Genesis]

lots of experiments ...

				<u>locked</u>
1.	Aave	Multichain	Lending	\$16.00B
2.	Maker	Ethereum	Lending	\$13.32B
3.	Curve Finance	Multichain	DEXes	\$12.73B
4.	InstaDApp	Ethereum	Lending	\$12.53B
5.	Compound	Ethereum	Lending	\$10.91B
6.	Uniswap	Ethereum	DEXes	\$6.54B
7.	Convex Finance	Ethereum	Assets	\$6.51B

This course



Cryptography

Economics

Distributed systems

Course organization

1. The starting point: Bitcoin mechanics
2. Consensus protocols
3. Ethereum and decentralized applications
4. Economics of decentralized applications
5. Scaling the blockchain: 10K Tx/sec
6. Private transactions on a public blockchain
(SNARKs and zero knowledge proofs)
7. Interoperability among chains: bridges and wrapped coins

Course organization

cs251.stanford.edu

- Three homework problems, four projects, final exam(?)
- Optional weekly sections on Friday

Please tell us how we can improve ...
Don't wait until the end of the quarter


Let's get started ...

Cryptography Background

(1) cryptographic hash functions

An efficiently computable function $H: M \rightarrow T$

where $|M| \gg |T|$



Collision resistance

Def: a **collision** for $H: M \rightarrow T$ is pair $x \neq y \in M$ s.t. $H(x) = H(y)$

$|M| \gg |T|$ implies that many collisions exist

Def: a function $H: M \rightarrow T$ is **collision resistant** if it is “hard” to find even a single collision for H (we say H is a CRHF)

Example: **SHA256:** $\{x : \text{len}(x) < 2^{64} \text{ bytes}\} \rightarrow \{0,1\}^{256}$

details in CS255

Application: committing to data on a blockchain

Alice has a large file m . She posts $h = H(m)$ (32 bytes)

Bob reads h . Later he learns m' s.t. $H(m') = h$

H is a CRHF \Rightarrow Bob is convinced that $m' = m$
(otherwise, m and m' are a collision for H)

We say that $h = H(m)$ is a **binding commitment** to m

(note: not hiding, h may leak information about m)

Committing to a list

(of transactions)

Alice has $S = (m_1, m_2, \dots, m_n)$

32 bytes



Goal:

- Alice posts a short binding commitment to S , $h = \text{commit}(S)$
- Bob reads h . Given $(m_i, \text{proof } \pi_i)$ can check that $S[i] = m_i$
Bob runs $\text{verify}(h, i, m_i, \pi_i) \rightarrow \text{accept/reject}$

security: adv. cannot find (S, i, m, π) s.t. $m \neq S[i]$ and
 $\text{verify}(h, i, m, \pi) = \text{accept}$ where $h = \text{commit}(S)$

Merkle tree

(Merkle 1989)

commitment



h

Merkle tree
commitment

m_1 m_2 m_3 m_4 m_5 m_6 m_7 m_8



list of values S

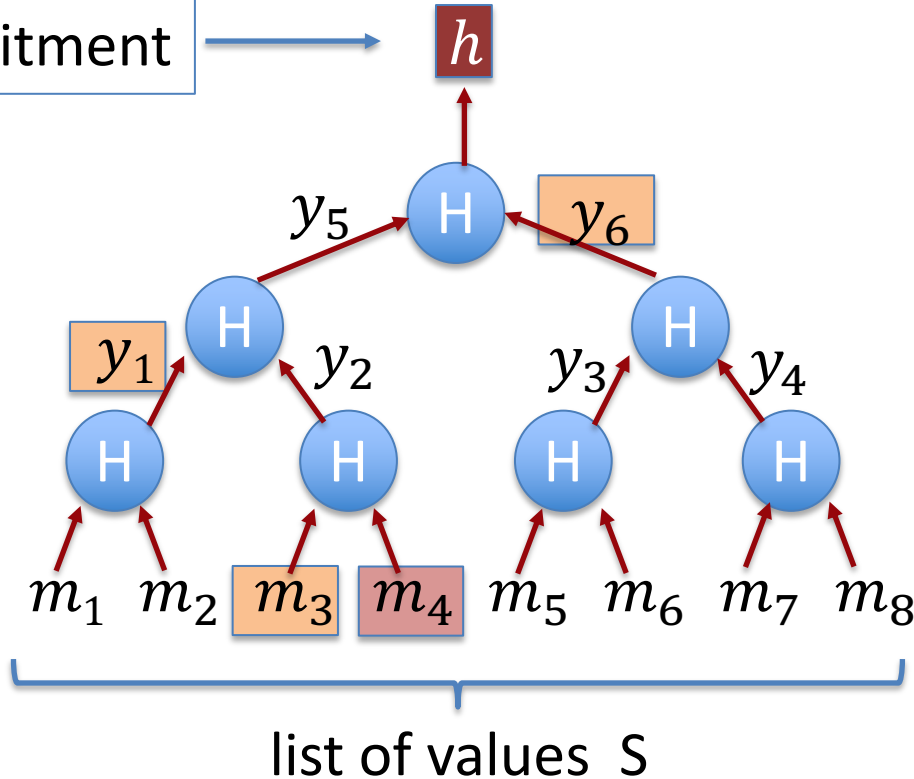
Goal:

- commit to list S of size n
- Later prove $S[i] = m_i$

Merkle tree

(Merkle 1989)

commitment



Goal:

- commit to list S of size n
- Later prove $S[i] = m_i$

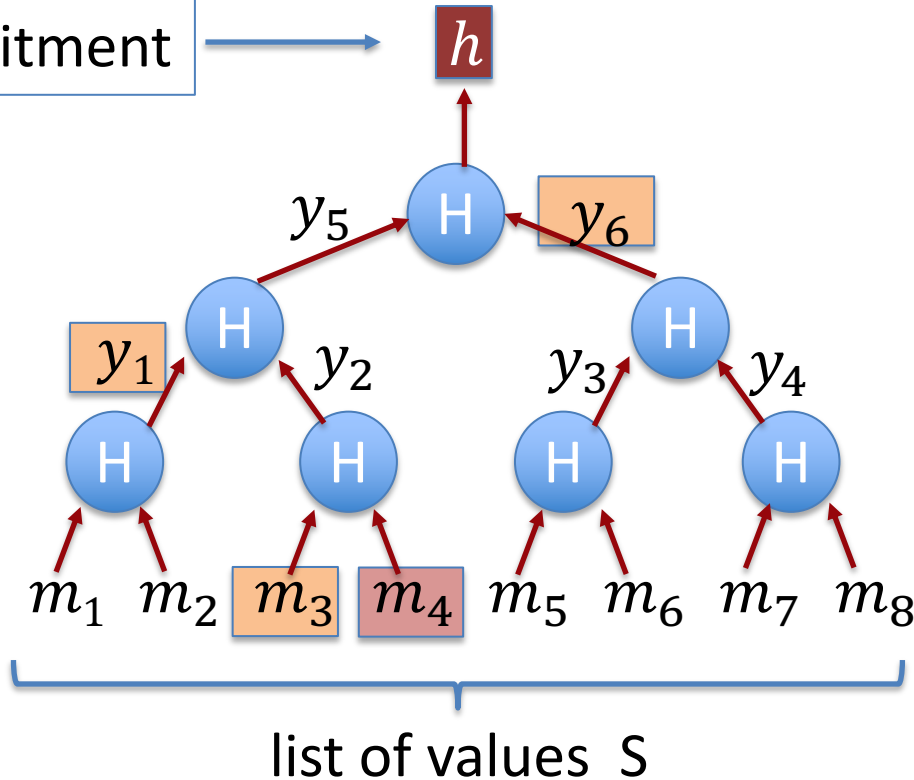
To prove $S[4] = m_4$,
proof $\pi = (m_3, y_1, y_6)$

length of proof: $\log_2 n$

Merkle tree

(Merkle 1989)

commitment



To prove $S[4] = m_4$,
proof $\pi = (m_3, y_1, y_6)$

Bob does:

$$y_2 \leftarrow H(m_3, m_4)$$

$$y_5 \leftarrow H(y_1, y_2)$$

$$h' \leftarrow H(y_5, y_6)$$

accept if $h = h'$

Merkle tree

(Merkle 1989)

Thm: H CRHF \Rightarrow adv. cannot find (S, i, m, π) s.t. $m \neq S[i]$ and
 $\text{verify}(h, i, m, \pi) = \text{accept}$ where $h = \text{commit}(S)$

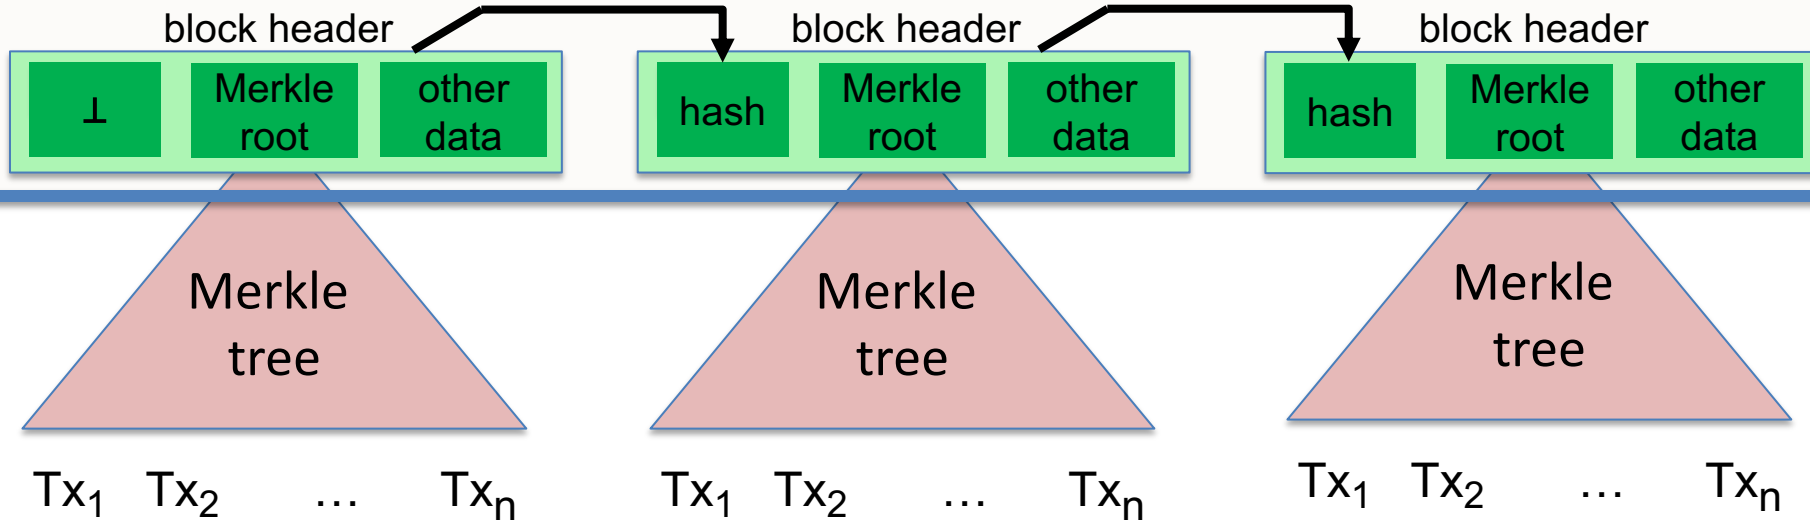
(to prove, prove the contra-positive)

How is this useful? Super useful. Example

- When writing a block of transactions S to the blockchain, suffices to write $\text{commit}(S)$ to chain. Keep chain small.
- Later, can prove contents of every Tx.

Abstract block chain

blockchain



Merkle proofs are used to prove that a Tx is “on the block chain”

Another application: proof of work

Goal: computational problem that

- takes time $\Omega(D)$ to solve, but
- solution takes time $O(1)$ to verify

(D is called the **difficulty**)

How? $H: X \times Y \rightarrow \{0, 1, 2, \dots, 2^n - 1\}$ e.g. $n = 256$

- puzzle: input $x \in X$, output $y \in Y$ s.t. $H(x, y) < 2^n / D$
- verify(x, y): accept if $H(x, y) < 2^n / D$

Another application: proof of work

Thm: if H is a “random function” then the best algorithm requires D evaluations of H in expectation.

Note: this is a parallel algorithm

⇒ the more machines I have, the faster I solve the puzzle.

Proof of work is used in some consensus protocols (e.g., Bitcoin)



Bitcoin uses $H(x, y) = \text{SHA256}(\text{SHA256}(x.y))$

Cryptography background: Digital Signatures

How to authorize a transaction

Signatures

Physical signatures: bind transaction to author

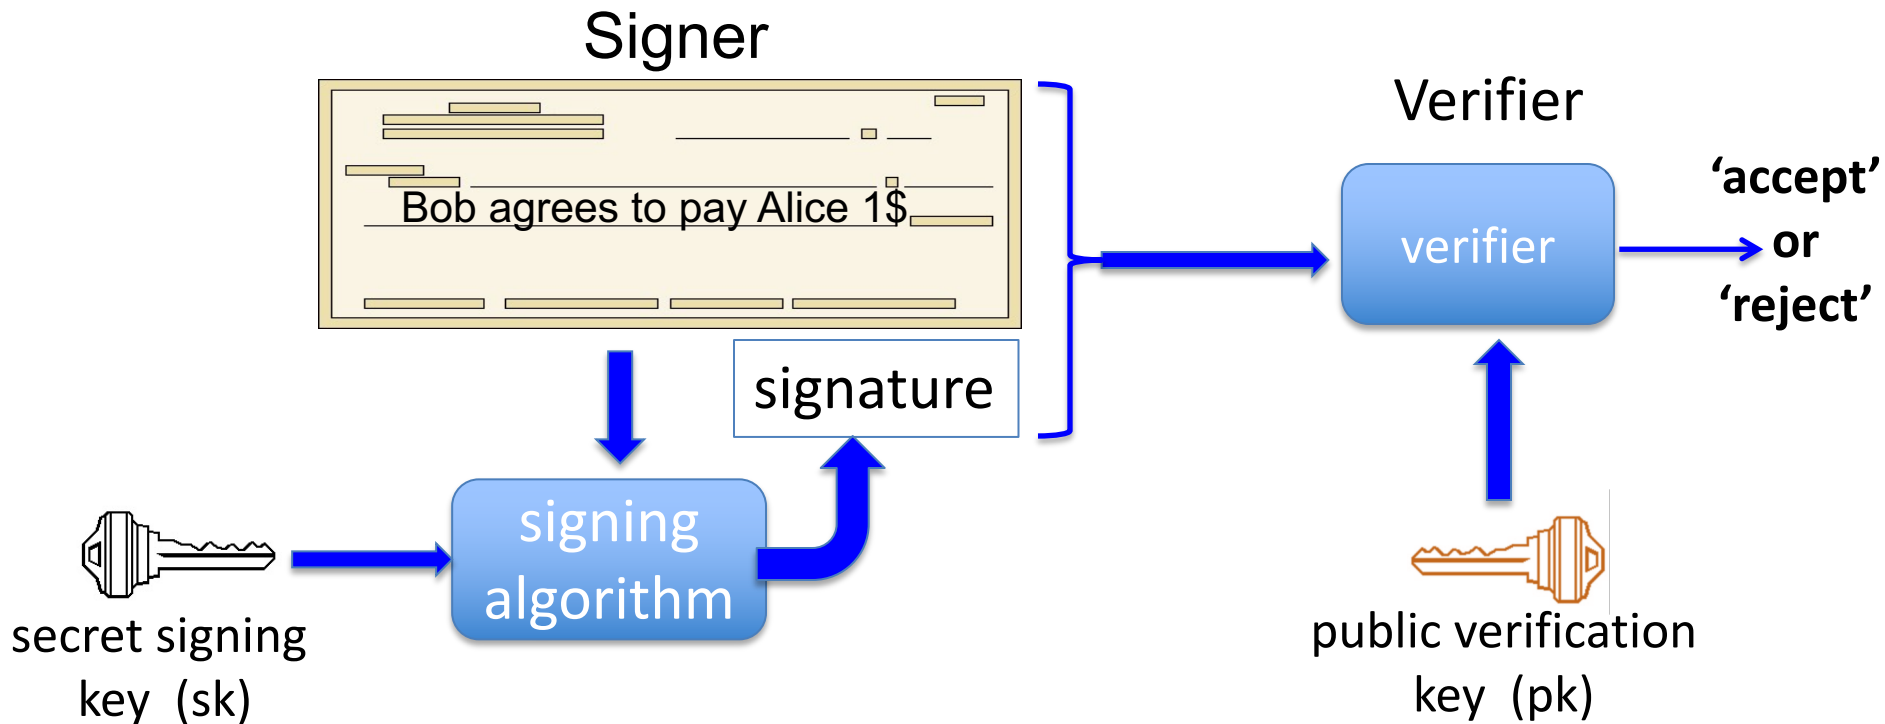


Problem in the digital world:

anyone can copy Bob's signature from one doc to another

Digital signatures

Solution: make signature depend on document



Digital signatures: syntax

Def: a signature scheme is a triple of algorithms:

- **Gen()**: outputs a key pair (pk, sk)
- **Sign**(sk, msg) outputs sig. σ
- **Verify**(pk, msg, σ) outputs 'accept' or 'reject'

Secure signatures: (informal)

Adversary who sees signatures on many messages of his choice, cannot forge a signature on a new message.

Families of signature schemes


1. RSA signatures (old ... not used in blockchains):
 - long sigs and public keys (≥ 256 bytes), fast to verify
2. Discrete-log signatures: Schnorr and ECDSA (Bitcoin, Ethereum)
 - short sigs (48 or 64 bytes) and public key (32 bytes)
3. BLS signatures: 48 bytes, aggregatable, easy threshold
(Ethereum 2.0, Chia, Dfinity)
4. Post-quantum signatures: long (≥ 768 bytes)

details in CS255

Signatures on the blockchain

Signatures are used everywhere:

- ensure Tx authorization,
- governance votes,
- consensus protocol votes.



END OF LECTURE

Next lecture: the Bitcoin blockchain