CS261: Security in Computer Systems

Bitcoin

Prof. Raluca Ada Popa

Sept 9, 2018

What is Bitcoin?

- Bitcoin is **a cryptocurrency**: a digital currency whose rules are enforced by cryptography and not by a trusted party (e.g., bank)
- **Core ideal:** avoid trust in institutions (e.g., banks, governments)
  - Reasons: Ideological, financial (avoid fees), peudo-anonymity
- Created by Satoshi Nakamoto, an anonymous identity, in 2009
- Bitcoin is also a **ledger.** Its protocol is built on a technique called a blockchain, which has applications beyond Bitcoin

Replacing banks

"IN BANKS WE DISTRUST"

Basic notions a bank provides:

- Identity management
- Transactions
- Prevents double spending

How can we enforce these properties cryptographically?

Two components

1. **Ledger:**
    1. publicly-visible,
    2. append-only, and
    3. immutable,

    log
2. **Cryptographic transactions**

Cryptographic transactions

- assume the existence of a trusted ledger

Identity

Q: How can we give a person a cryptographic identity?

- Each user has a PK and SK
- User referred to by PK

Transactions

Q: How can Alice transfer 10 ฿ (bitcoins) to Bob in a secure way?

- Idea: Alice signs transaction using her SK_A
- $sign_{SK_A}$("PK_A transfers *10 ฿ to PK_B*")
- Anyone can check Alice intended transaction

Q: Problems?

- Alice can spend more money than she has. She can sign as much as she wants.

Q: Ideas how to solve this still assuming a "trusted ledger owner"?

Include only correct transactions in the public ledger

- **For now only**: assume a trustworthy ledger owner, assume initial budgets for each PK

Q: how would you prevent double spending?

- Assume all signatures/transactions are sorted in order of creation; include previous transaction where money came from

$TX = (PK_{sender} \rightarrow PK_{receiver}; X$ ฿;
list of transactions L where money came from)

time

Initial budgets: PK_A has 10 ฿	$TX_1 = (PK_A \rightarrow PK_B; 10$ ฿; *from initial budgets*) $sign_{SK_A}(TX_1)$	$TX_2 = (PK_B \rightarrow PK_C; 5$ ฿; *from* TX_1*)* $sign_{SK_B}(TX_2)$

How does the ledger owner check a transaction?

1. The signature on TX verifies with the PK of the sender
2. The transactions in L have PK of sender as “to”
3. Checks sender had X bitcoins.
   - If all lists have only one transaction. Check the outcome of the transaction minus sum of all amounts spent from this transaction prior is at least X.
   - If Lists have more than one transaction, compute the amount unspent for each transaction in L and see if the total is at least X.

Bitcoin’s ledger

Solution: blockchain + consensus via proof of work

Blockchain

- Chain transactions using their hashes => hashchain
- Each transaction contains hash of previous transaction **(which contains the hash of its own previous transaction, and so on)**

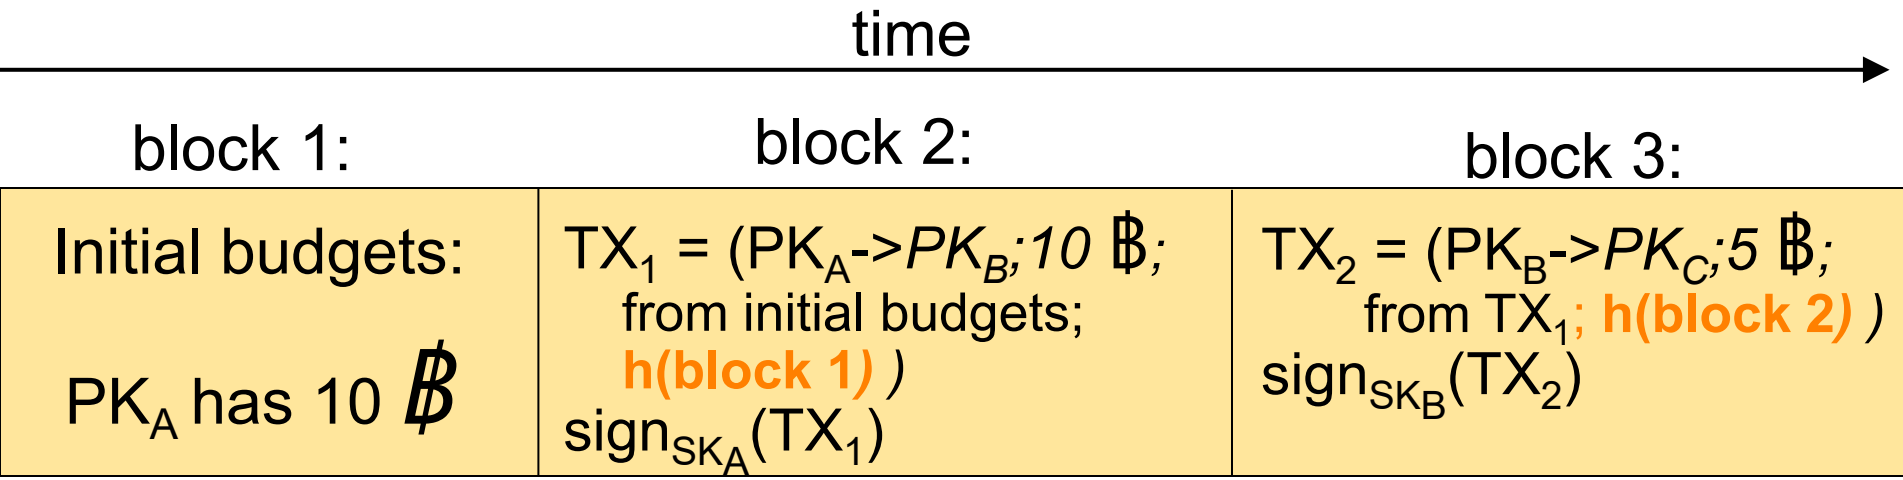

block i refers to the entire block (transaction description and signature), so the hash is over all of this

Properties of the hashchain

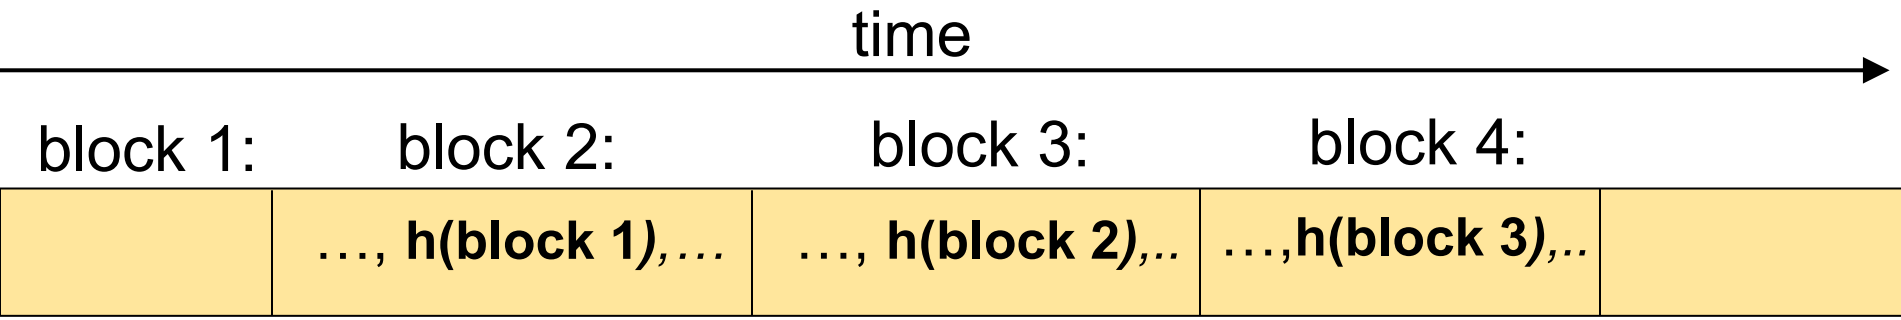

Given h(block i) from a trusted source and all the blocks *1 … i* from an untrusted source, Alice can verify that blocks 1 … i are not compromised using h(block i)

Q: How?

A: Alice recomputes the hashes of each block, checks it matches the hash in the next block, and so on, until the last block, which she checks it matches the hash from the trusted source

Why can't attacker cheat?

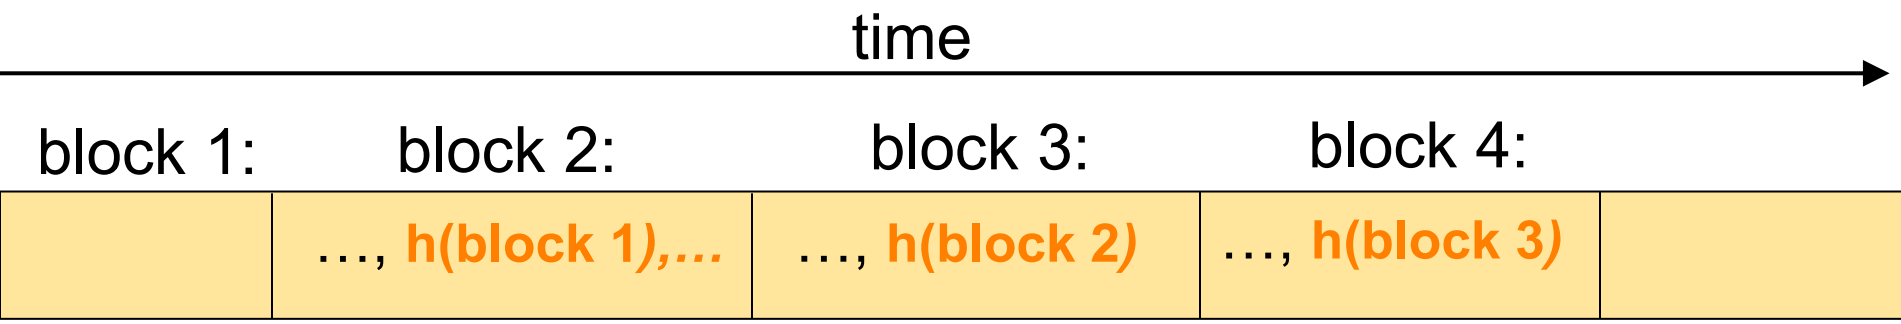

Say Alice obtains h(block 4) from somewhere **trusted**

She fetches the entire blockchain from a compromised server.

Q: Why can't the attacker give Alice an incorrect chain?

Say block 2 is incorrect.

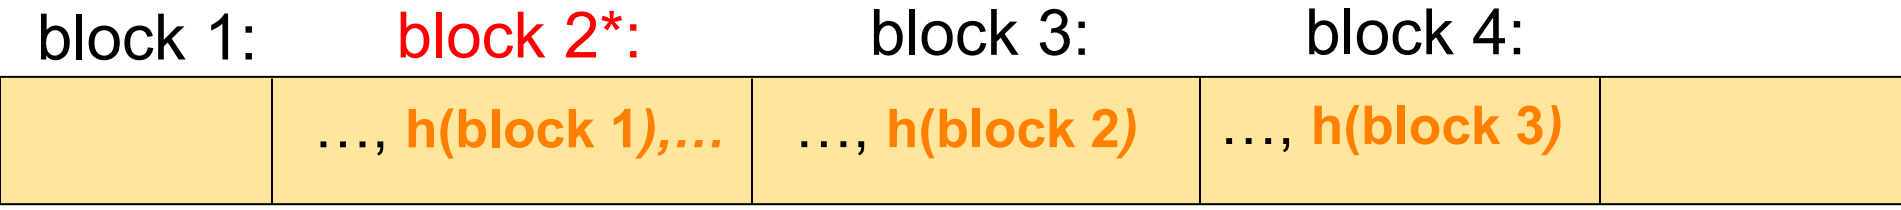

A: because the hash is collision resistant

She fetches the entire blockchain from a compromised server.

Q: Why can't the attacker give Alice an incorrect chain?

Say block 2 is incorrect.

block 1:	block 2:	block 3:	block 4:	
	…, **h(block 1),…**	…, **h(block 2)**	…, **h(block 3)**	

- If block 2* is incorrect, then hash(block 2*) ≠ hash(block 2)
- Then the third block is block 3* ≠ block 3 because it includes hash(block 2*)
- So hash(block 3*) ≠ hash(block 3)
- Then the fourth block is block 4* ≠ block 4 because it includes hash(block 3*)
- So hash(block 4*) ≠ hash(block 4)
- Hence, the hash of the block chain from the server will not match the trusted hash, detecting misbehavior
- If the hash does match, the attacker supplied the correct block chain

Back to building the trustworthy ledger

- Every participant in Bitcoin stores a copy of the entire blockchain
- When someone wants to create a new transaction, they broadcast the transaction to everyone
- Every node checks the transaction, and if it is correct, it creates a new block including this transaction and adds it to its local blockchain

- **Q: Problem?**
- A: People can choose to truncate blockchain or not include certain transactions

Problem: Consensus

- Problem: Mallory can fork the hash chain
- Say she buys Bob's house from him for $500K in Bitcoins. Then, she goes back in time and, starting from the block chain just before this transaction was added to it, she starts appending new entries from there.

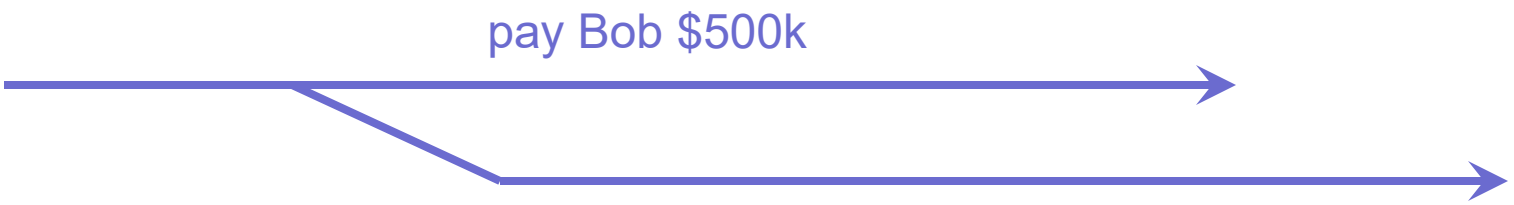

Q: How can we reach consensus on the correct chain?

Proof of Work and Mining

- Not everyone is allowed to add blocks to the blockchain, but only certain people, called miners
- All miners try to solve a proof of work: the hash of the new block (which includes the hash of the blocks so far) must start with 33 zero bits
  - Can include a random number in the block and increment that so the hash changes until the proof of work is solved
- Once a miner solves a proof of work, it includes all transactions it heard about after checking they are correct

Consensus

- Consensus: **longest correct** chain wins
- Everyone checks all blocks and all transactions. If a miner appends a block with some incorrect transaction, the block is ignored
- Assumes most miners are honest

“Longest chain” wins

- **Q: What if two different parts of network have different hash chains?**

- Solution: Whichever is “longer” wins; the other is discarded

How can we convince people to mine?

- A: Give a reward to anyone who successfully appends – they receive a free coin
  - Essentially they may include a transaction from no one to their PK having a coin

Consensus

- **Q: Can Mallory fork the block chain?**
- A: No. Assume Bob waits three blocks till he deems the transaction complete (and gives Mallory his house). She cannot fork the chain unless she has ≥51% of the computing power in the world. Longest chain wins, and her forked one will be shorter (unless she can mine new entries faster than aggregate mining power of everyone else in the world).

Let's chew on consensus

- **Q: What happens if Miner A and Miner B at the same time solve a proof of work and append two different blocks thus forking the network?**
- A: The next miner that appends onto one of these chains, invalidates the other chain. Longest chain wins.
- **Q: What happens if Miner Mallory discards the last few blocks in the block chain and miners from there?**
- A: Unless Miner Mallory has more than 50% of the computation power in the world, she will not be successful because the longest chain will keep being appended
- **Q: If a miner included your transaction in the latest block created, are you guaranteed that your transaction is forever in the blockchain?**
- A: No, there could have been another miner appending a different block at the same time and that chain might be winning. So wait for a few blocks, e.g. 3 until your transaction is committed with high probability

Let’s chew on consensus

- **Q: What happens if a miner who just mined a block refuses to include my transaction?**
- A: Hopefully the next miner will not refuse this. Each transaction also includes a fee which goes to the miner, so a miner would want to include as many transactions as possible

Proof of work can be adapted

- Mining frequency is ~15 mins
- If it takes too long to mine on average, make the proof of work easier (less zeros), else make it harder (more zeros)
- **Q: what is the economic insight?**
- A: if mining is rare, it means few machines in the network, give more incentives to join the network

Watch the blockchain live

- https://blockchain.info/

Mining pools

- It used to be easy to mine in early days, but now it is too hard for a regular person to mine, they need too much compute
- But you can contribute your cycles to a mining pool, which is a group of many machines with good success of mining on average
- Receive a more predictable income based on the average mining of the group and how many cycles you contribute

Top mining countries

1. China
2. Georgia
3. Sweden

(the ranking is influenced by price of electricity)

First few blocks were mined by Satoshi Nakamoto

- Wrote beautiful white paper on Bitcoin, in the syllabus
- No one knows who he is, online presence only
- Name stands for clear/wise medium; most likely not Japanese, but pseudonym
- He is very rich! [But hasn’t changed yet]

Bitcoin

- Public, distributed, peer-to-peer, hash-chained audit log of all transactions (“block chain”).
- Mining: Each entry in block chain must come with a proof of work (its hash value starts with *k* zeros). Thus, appending takes computation.
- Lottery: First to successfully append to block chain gets a small reward (if append is accepted by others). This creates new money. Each block contains a list of transactions, and identity of miner (who receives the reward).
- Consensus: If there are multiple versions of the block chain, longest one wins.

Bitcoin

- Transactions: If Alice wants to give $10 to Bob, she signs this transaction. She gives the signed transaction to all miners and asks them to include it in the block they're trying to append to the chain.
- Honest miners check integrity of block chain entries and try to append to the latest, longest valid version of block chain.
- Bob knows he has received $10 once this transaction appears in the consensus block chain.

Is Bitcoin anonymous?

Quantitative Analysis of the Full Bitcoin Transaction Graph

Dorit Ron and Adi Shamir

Department of Computer Science and Applied Mathematics,
The Weizmann Institute of Science, Israel
{dorit.ron,adi.shamir}@weizmann.ac.il

Abstract. The Bitcoin scheme is a rare example of a large scale global payment system in which all the transactions are publicly accessible (but in an anonymous way). We downloaded the full history of this scheme, and analyzed many statistical properties of its associated transaction graph. In this paper we answer for the first time a variety of interest-

It might look anonymous because you only use your PK and not your name as at a bank. But all your transactions can be tied to your PK. People can identify you from transactions you make: parking fee near your work, people you transact with, etc.

They can even see how wealthy you are

Mitigations: use multiple PKs

Solution: Zcash, anonymous version of Bitcoin

Many other cryptocurrencies

"The number of cryptocurrencies available over the internet as of 19 August 2018 is over 1600 and growing." [Wikipedia]

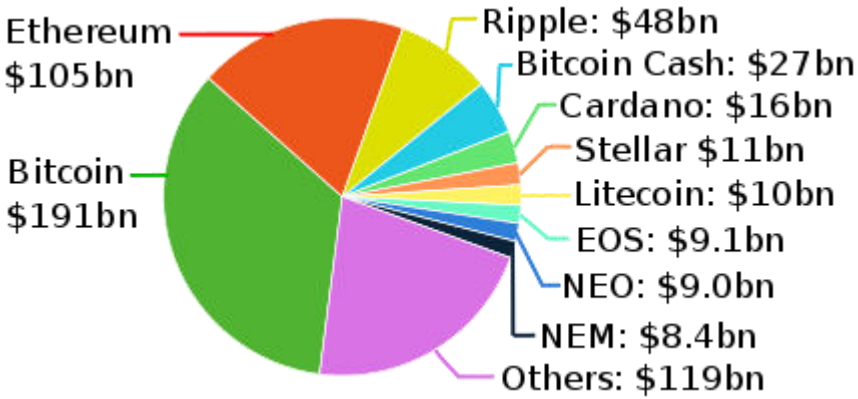

[January 27,2018]

Other cryptocurrencies we will study

2nd largest. Introduces the powerful idea of ”smart contracts”, running code in the blockchain.

Privacy preserving via zero-knowledge proofs

Very new concepts, no proof of work, much more scalable

Q&A on blockchain/cryptocurrencies

- How can Alice turn dollars into bitcoins, or vice versa?
- Why is Bitcoin popular?
- Should I think of Bitcoin as a short-term currency or as a long-term investment?
- Is it ethical to build a system that relies upon wasting CPU cycles (and thus energy)?