

Bitcoin

Lecture 16

Software Security Engineering

Winter 2023
Thompson Rivers University

What is Bitcoin?

The
Economist



Bitcoin: explained

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), pseudoanonymity
- Bitcoin is also a **ledger**
 - Its protocol is built on a technique called a **blockchain**, which has applications beyond Bitcoin
- Created by Satoshi Nakamoto
 - an anonymous identity, in 2009



Satoshi Nakamoto

- Wrote beautiful white paper on Bitcoin (in the moodle)
- 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]


Cryptocurrencies have
supporters and opposers

Supporters say...

- No need to trust or depend on banks or the government
- Low transaction fees when transferring
- Helps disadvantaged areas (due to the fee but also because you don't need a bank account or official identity)
- Anybody can partake
- Cross nation trading easier
- Transparent
- Anonymity
- “Digital gold”

Digital gold

- Used to be thought to replace cash but that seems less likely today
- People associate it to digital gold, a way to invest in a currency that does not belong to one government



Critics say...

- A major criticism is that it brings waste (e.g., proof of work, many large copies of blockchain)
- Not as decentralized as we wished
- Not scalable
- Market fluctuations
- Anonymity
- No security in case of loss
- Helps criminals
- Not a replacement for cash, not all vendors accept it

Bitcoin technical design

Bitcoin technical design

Let's work it out together!

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

- Ledger:
 - publicly-visible
 - append-only
 - immutable, log
- Cryptographic transactions

Cryptographic Transactions

For now, assume the existence of a trusted ledger (append-only, immutable, everyone can see what is on it)

Identity

- How can we give a person a cryptographic identity?
- Each user has a PK and SK
- User referred to by PK

Transactions

- How can Alice transfer 10 ₿ (bitcoins) to Bob in a secure way?
 - Idea: Alice signs transaction using her SK_A
 - $\text{sign}_{SK_A}(\text{"PK}_A \text{ transfers 10 ₿ to PK}_B\text{"})$
 - Anyone can check Alice intended the transaction
- what can go wrong?!

Transactions


- How can Alice transfer 10 ₿ (bitcoins) to Bob in a secure way?
 - Idea: Alice signs transaction using her SK_A
 - $\text{sign}_{SK_A}(\text{"PK}_A \text{ transfers 10 ₿ to PK}_B \text{"})$
 - Anyone can check Alice intended the transaction
- what can go wrong?!
 - Alice can spend more money than she has.
 - She can sign as much as she wants.
- How to solve this still assuming a “trusted ledger owner”?

Include only correct transactions in the public ledger

- For now only: assume there is a **trustworthy ledger owner**, assume **initial budgets** for each PK
- how would you prevent double spending?
 - Assume all signatures/transactions are sorted in order of creation
 - include previous transaction where money came from
- $TX = (PK_{\text{sender}} \rightarrow PK_{\text{receiver}}; X_{\mathbb{B}}; PK_{\text{sender}} \rightarrow PK_{\text{sender}}; R_{\mathbb{B}}; L)$
 - $X_{\mathbb{B}}$ - amount transferred
 - $R_{\mathbb{B}}$ - amount remained
 - L - list of transactions L where money came from

Public Ledger

Initial budgets: PK _A has 10 ₿	TX ₁ = (PK _A → PK _B ; 10₿; PK _A → PK _A ; 0₿; from initial budgets) sign _{SKA} (TX ₁)	TX ₂ = (PK _B → PK _C ; 6₿; PK _B → PK _B ; 4₿; TX ₁) sign _{SKB} (TX ₂)
--	---	--



How does the ledger owner check a transaction?

Verify TX

- The signature on TX verifies with the PK of the sender
- The transactions in L have PK of sender as their recipient
 - (that is, the sender receives Bitcoins in the transactions in L)
- The transactions in L have not been spent before by sender
 - (each transaction A->B can only be spent once by B, and once by A if there were remaining bitcoins in it)
- Sender had X+R Bitcoins in L: the sum of the amounts received in the transactions in L total to X+R.

$$TX_n = (PK_A \rightarrow PK_B; X\text{฿};$$

$$PK_A \rightarrow PK_A; R\text{฿};$$

L)

$$\text{sign}_{SK_A}(TX_n)$$

Two Main Components

- Cryptographic transactions
- Ledger:
 - publicly-visible
 - append-only
 - immutable log

Bitcoin's ledger

- Blockchain
- Consensus via proof of work

What is Blockchain?

What is Blockchain?

You know it!

What is Blockchain?

You know it!

Blockchain is Hashchian

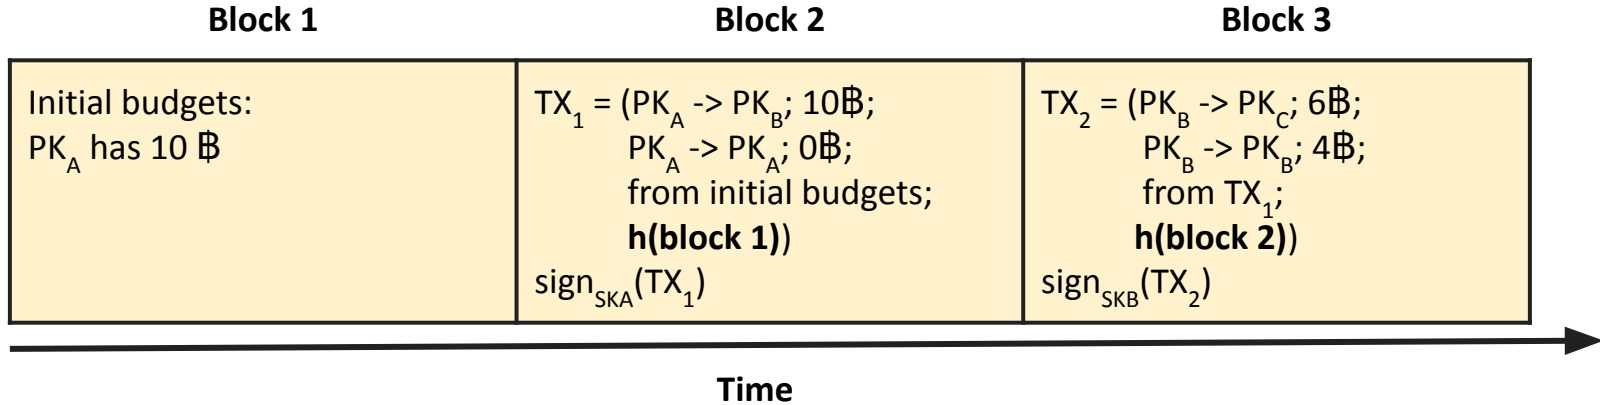
Hash of transactions

$H(H(H(\dots(H(\text{Transactions}))\dots)))$


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

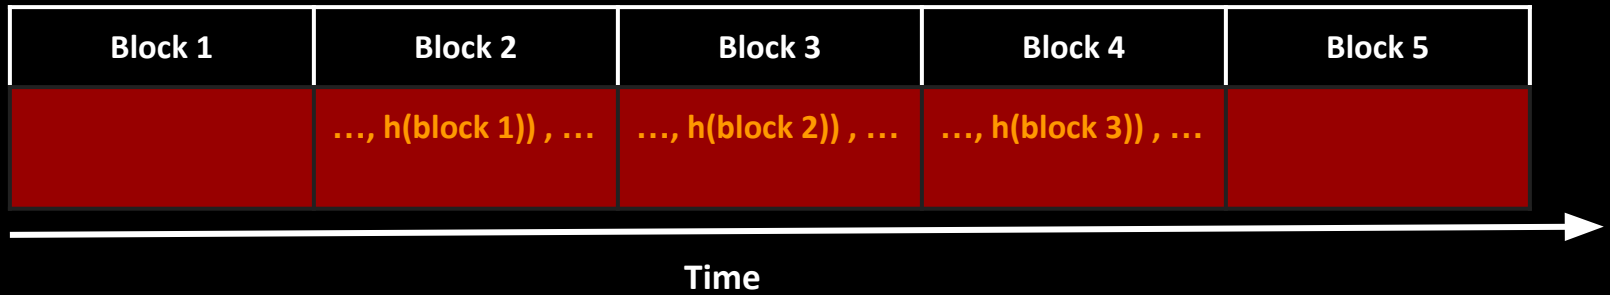
Public Ledger



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

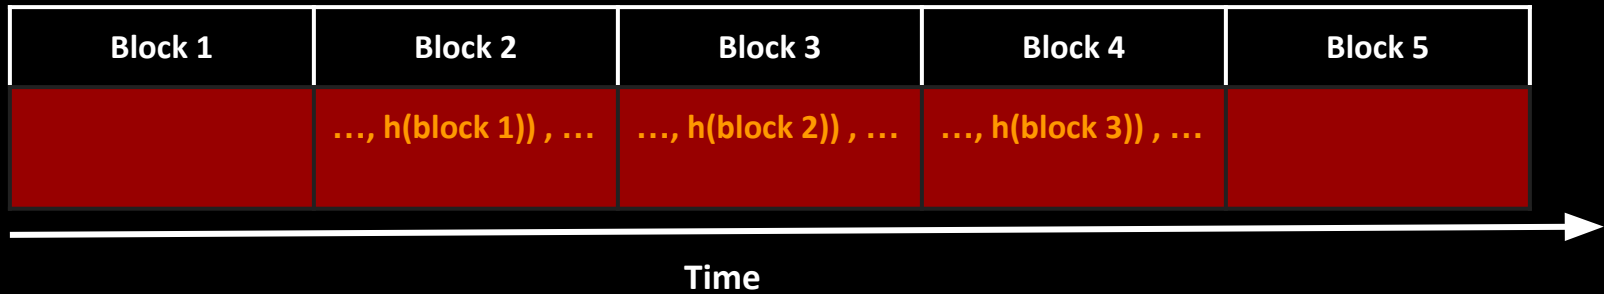


Given $h(\text{block } i)$ from a *trusted source* and
all the **blocks 1 ... i** from an *untrusted source*,
Alice can verify that **blocks 1 ... i** are not compromised
using $h(\text{block } i)$



Given $h(\text{block } i)$ 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 $h(\text{block } i)$

How?



Given $h(\text{block } i)$ 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 $h(\text{block } i)$

How?

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?

Block 1	Block 2	Block 3	Block 4	Block 5
	..., h(block 1) ,, h(block 2)) ,, h(block 3)) , ...	
Block 1	Block 2*	Block 3	Block 4	Block 5
	..., h(block 1)) ,, h(block 2)) ,, h(block 3)) , ...	

Alice obtains $h(\text{block 4})$ from somewhere **trusted**

Block 1	Block 2	Block 3	Block 4	Block 5
	..., h(block 1) ,, h(block 2)) ,, h(block 3)) , ...	

Block 1	Block 2*	Block 3	Block 4	Block 5
	..., h(block 1)) ,, h(block 2)) ,, h(block 3)) , ...	

Say Alice obtains $h(\text{block 4})$ from somewhere **trusted**
She fetches the entire blockchain from a **compromised** server.

Block 1	Block 2	Block 3	Block 4	Block 5
	..., h(block 1) ,, h(block 2)) ,, h(block 3)) , ...	
Block 1	Block 2*	Block 3	Block 4	Block 5
	..., h(block 1)) ,, h(block 2)) ,, h(block 3)) , ...	

Say Alice obtains $h(\text{block 4})$ from somewhere **trusted**
 She fetches the entire blockchain from a **compromised** server.
 Why can't the attacker give Alice an incorrect chain?
 Say block 2 is incorrect.

Block 1	Block 2	Block 3	Block 4	Block 5
	..., h(block 1) ,, h(block 2)) ,, h(block 3)) , ...	
Block 1	Block 2*	Block 3	Block 4	Block 5
	..., h(block 1)) ,, h(block 2)) ,, h(block 3)) , ...	

Say Alice obtains $h(\text{block 4})$ from somewhere **trusted**
She fetches the entire blockchain from a **compromised** server.

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

Say block 2 is incorrect.

because the hash is **collision resistant**

Block 1	Block 2	Block 3	Block 4	Block 5
	..., h(block 1) ,, h(block 2) ,, h(block 3) , ...	
Block 1	Block 2*	Block 3	Block 4	Block 5
	..., h(block 1) ,, h(block 2) ,, h(block 3) , ...	

If **Block 2*** is incorrect, then $\text{hash}(\text{block } 2^*) \neq \text{hash}(\text{block } 2)$

Then the third block is $\text{block } 3^* \neq \text{block } 3$ because it includes $\text{hash}(\text{block } 2^*)$

So $\text{hash}(\text{block } 3^*) \neq \text{hash}(\text{block } 3)$

$\text{hash}(\text{block } 4^*) \neq \text{hash}(\text{block } 4)$

It will not match the trusted hash, detecting misbehavior

In Bitcoin:

- Every participant stores the blockchain
- There is no central party storing it
- 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
- Some participants can be **malicious**
- The majority are assumed to be **honest**

What can go wrong?

What can go wrong?

People can choose to truncate blockchain or not
include certain transactions

What can go wrong?

People can choose to truncate blockchain or not
include certain transactions

So we need a way for everyone to agree on the
content of the blockchain

What can go wrong?

People can choose to truncate blockchain or not
include certain transactions

So we need a way for everyone to agree on the
content of the blockchain

consensus

Example

- 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
- Can she get others to accept this forked chain, so she gets her \$500K back?
- Yes.



This is an example of
fork attack, double spending

How do users agree on
the same history?

How do users agree on
the same history?

Consensus via proof of work

Proof of work / Mining

- Not everyone is allowed to add blocks to the blockchain
 - but only certain people, called **miners**
- An honest miner will include all transactions it hears about after checking them
- 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 N (e.g. 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
 - Eg: Hash(block || random_number) = 000...0000453a48b244
 - Currently someone in the world solves the proof of work every 10-20 mins

Propagating blocks

- Miners broadcast blocks with proof of work
- All (honest) Bitcoin nodes
 - listen for such blocks
 - check the blocks for correctness, and
 - accept the longest correct chain
- If a miner appends a block with some incorrect transaction, the block is ignored

Consensus: longest correct chain wins

- Everyone will always prefer the longer correct chain

Example

- An honest miner M1 stores current blockchain: b1->b2->b3
- M1 hears about transactions T
- M1 tries to mine for block b4 to include T
- Another miner M2 mines first b4 and broadcasts b4, with b3->b4
- M1 checks b4, accepts b4, and starts mining for block 5

Example (cont'd)

- M1 now has blockchain
b1->b2->b3->b4
- M1 hears that some miners are broadcasting
b1->b2->b3->b4'->b5'
- M1 checks this new chain, and then accepts
- this new chain, essentially discarding b4

Assumption

- Assumes more than half of the computing power is in the hands of honest miners
- So honest miners will always have an advantage to mine the longest chain

Can Mallory fork the block chain?

Can Mallory fork the block chain?

No, not unless she has $\geq 51\%$ of the computing power in the world.

Can Mallory fork the block chain?

No, not unless she has $\geq 51\%$ 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).

“Longest chain” wins

- Problem: What if two different parts of network have different hash chains?
- Solution: Whichever is “longer” wins; the other is discarded

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

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
- Q: What happens to a miner's reward if his block was removed because an alternate longer chain appears?
- A: The miner lost their reward. Only the transactions and rewards on the longest chain “exist”.

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?

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?
The next miner that appends onto one of these chains,
invalidates the other chain. Longest chain wins.

If a miner included your transaction in the latest block created, are you guaranteed that your transaction is forever in the blockchain?

If a miner included your transaction in the latest block created, are you guaranteed that your transaction is forever in the blockchain?

No, there could have been another miner appending a different block at the same time and that chain might be winning.

If a miner included your transaction in the latest block created, are you guaranteed that your transaction is forever in the blockchain?

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, though you can never be sure.

What happens if a miner who just mined a block
refuses to include my transaction?

What happens if a miner who just mined a block
refuses to include my transaction?

Hopefully the next miner will not refuse this.

What happens if a miner who just mined a block
refuses to include my transaction?

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


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

Distribution of Bitcoin mining hashrate from September 2019 to January 2022, by country



Source

Various sources (BTC.com, Poolin, and ViaBTC)
© Statista 2023

Additional Information:

Worldwide; Cambridge Centre for Alternative Finance; September 2019 to January 2022; The country names underneath so to remove certain countries, or get to a particular country of interest

How to become a miner?

How to become a miner?

Not a good idea!

reduces its your PC's lifespan

requires significant computing power these days

consumes huge electricity power

How to become a miner?

Open a crypto wallet in websites like coinbase.com

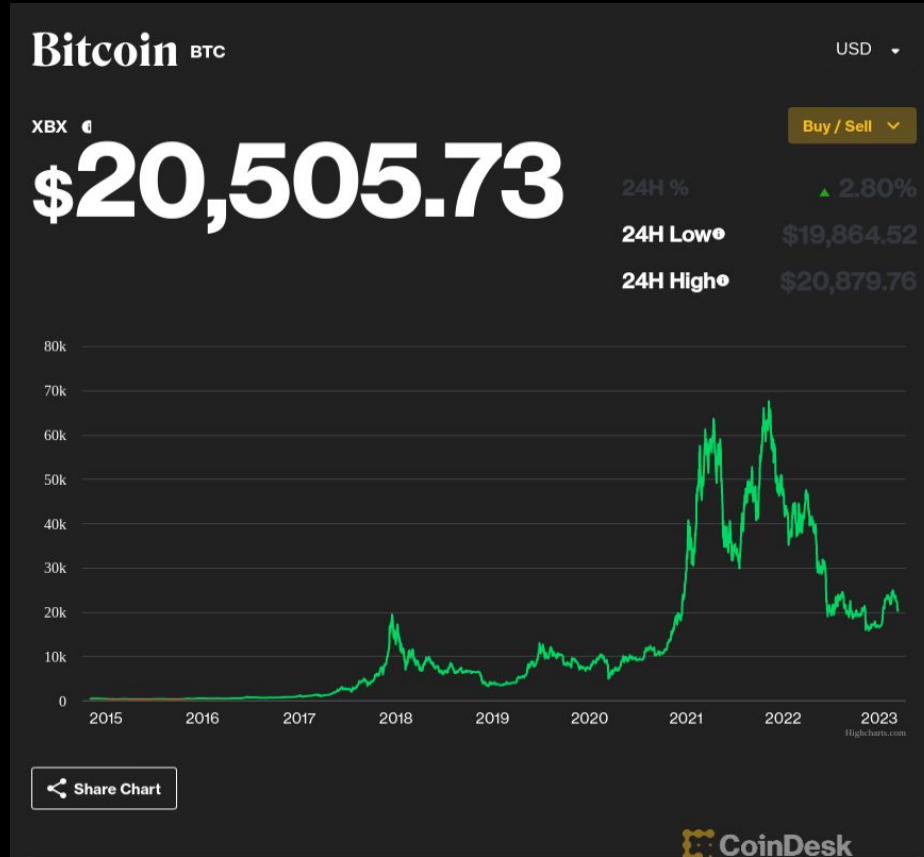
Join a mining pool like Luxor

install apps and set configurations







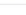


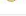









Is Bitcoin anonymous?

- 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

Value fluctuations



Many other cryptocurrencies (21,844 as March 2023)

Rank	Name	Symbol	Market Cap	Price	Circulating Supply	Volume(24h)	% 1h	% 24h	% 7d
1	 Bitcoin	BTC	\$394,976,128,809	\$20,449.50	19,314,706 BTC	\$30,921,366,540	-0.51%	2.60%	-8.06%
2	 Ethereum	ETH	\$179,867,434,231	\$1,469.82	122,373,866 ETH *	\$15,210,726,426	-0.27%	4.07%	-5.61%
3	 Tether	USDT	\$72,604,966,600	\$1.01	72,074,498,292 USDT *	\$66,995,533,534	-0.29%	0.68%	0.72%
4	 BNB	BNB	\$43,520,099,172	\$275.63	157,892,860 BNB *	\$522,687,672	-0.23%	-0.06%	-4.12%
5	 USD Coin	USDC	\$39,239,697,839	\$0.9603	40,861,915,901 USDC *	\$26,826,823,167	-2.23%	-3.96%	-3.96%
6	 XRP	XRP	\$18,595,909,528	\$0.365	50,950,912,949 XRP *	\$1,068,418,326	-0.62%	-0.98%	-1.81%
7	 Cardano	ADA	\$10,622,870,918	\$0.3062	34,695,282,189 ADA *	\$450,852,627	-0.70%	-1.37%	-8.29%
8	 Polygon	MATIC	\$9,139,506,408	\$1.05	8,734,317,475 MATIC *	\$755,323,738	-0.51%	1.28%	-5.56%
9	 Dogecoin	DOGE	\$8,761,562,060	\$0.06604	132,670,764,300 DOGE	\$458,256,416	-0.49%	1.54%	-10.47%
10	 Binance USD	BUSD	\$8,418,885,515	\$1.00	8,396,420,505 BUSD *	\$9,000,698,588	0.03%	0.27%	0.25%
11	 Solana	SOL	\$6,953,456,789	\$18.16	382,960,145 SOL *	\$728,854,716	-0.75%	3.37%	-11.54%
12	 Polkadot	DOT	\$6,346,393,214	\$5.45	1,165,235,464 DOT *	\$314,559,951	0.01%	-0.29%	-6.57%
13	 Shiba Inu	SHIB	\$5,634,851,726	\$0.00001026	549,063,278,876,302 SHIB *	\$277,670,434	-0.39%	0.93%	-6.82%
14	 Dai	DAI	\$5,467,331,646	\$0.971	5,630,656,035 DAI *	\$4,265,290,520	-1.56%	-2.87%	-2.84%
15	 TRON	TRX	\$5,363,672,254	\$0.05876	91,287,989,703 TRX *	\$423,814,651	-0.58%	3.64%	-11.51%
16	 Litecoin	LTC	\$4,999,554,182	\$69.00	72,460,927 LTC	\$1,071,278,468	-1.20%	-2.73%	-21.97%
17	 Avalanche	AVAX	\$4,683,713,607	\$14.39	325,526,001 AVAX *	\$250,466,530	-0.33%	-2.05%	-10.39%
18	 Uniswap	UNI	\$4,152,323,929	\$5.45	762,209,327 UNI *	\$146,273,338	-0.40%	-2.41%	-10.40%
19	 UNUS SED LEO	LEO	\$3,348,326,849	\$3.51	953,954,130 LEO *	\$1,068,280	-0.50%	2.62%	3.89%
20	Chainlink	LINK	\$3,185,180,988	\$6.16	517,099,970 LINK *	\$341,862,235	-0.41%	-0.24%	-9.45%
21	Wrapped Bitcoin	WBTC	\$3,138,637,353	\$20,478.63	153,264 WBTC *	\$763,696,786	-1.37%	2.62%	-7.80%

Blockchain

- Usage of blockchain goes beyond cryptocurrencies.
- The idea is a ledger storing information in an immutable way that can be accessed cross organizations.
 - Financial usages (e.g., ledgers for bank transactions)
 - Healthcare (e.g., personal health records encrypted in the blockchain so only certain insurance and medical providers can access them)
 - Key distribution
 - Certificate Transparency

Another usage of a blockchain



Another usage of a blockchain

Love letter embedded in the blockchain

3505443530030ccfb8275d37e2db1cbd9368247c0842c7eac23d2cc5ad1966e8

2017-01-14 03:12:39

1DearSPQ51n2CKgSLQwMXrEFjWKmfuaoA6



1DayahDover111111111111111112JYRq2	0.00314159 BTC
1YourPersona1ity1sUnmatched43YzMv	0.00314159 BTC
1YourInte11igenceJustShines4B7QFA	0.00314159 BTC
1YouCanDoThingsFewPeop1eCan1G6NPV	0.00314159 BTC
1AndYoureA1waysJustGorgeous2x1SyG	0.00314159 BTC
1YouAreRea11yMyEntireWor1d116eypT	0.00314159 BTC
1GivingMyLifeMeaningAndFun13pcr5P	0.00314159 BTC
1Dayah7Px1kbs5x5cQbQMhtMm9wnUWJYTG	0.00314159 BTC
11LoveYou11111111111111111111GPc4r	0.00314159 BTC
1Forever111111111111111111113RMwCB	0.00314159 BTC

0.0314159 BTC

Another usage of a blockchain

Love letter embedded in the blockchain

It stays **forever!**

General problem with blockchain:

cannot erase information

Consider private information about you or
your organization leaking, the power of law used to
be able to remove it

Bitcoin (summary)

- 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 N 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.