


CS 419: Computer Security

Week 5: Hash Pointers, Blockchain, & Bitcoin

Paul Krzyzanowski



© 2024 Paul Krzyzanowski. No part of this content may be reproduced or reposted in whole or in part in any manner without the permission of the copyright owner.

Hash Pointers

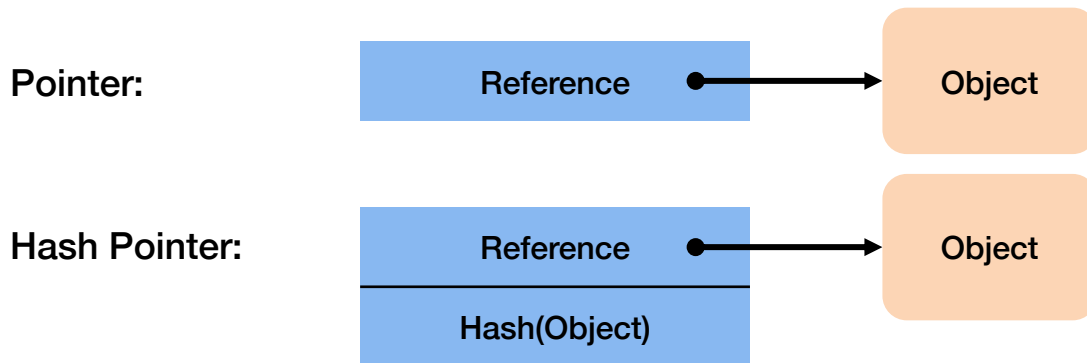
Hash Pointers

Extension of pointers in data structures


Hash pointer = { pointer, hash(data) }

pointer = reference that identifies where the object is:
memory location, file name, object ID, server/object, ...

hash(data) = hash function applied to the data being pointed to

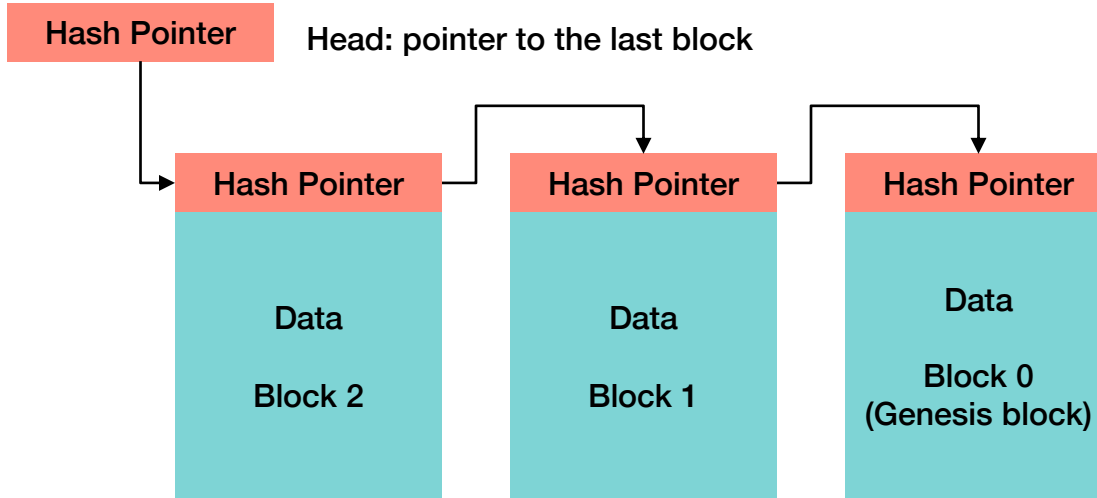


Tamper Detection With A Hash Pointer



- If an attacker modifies data, $hash(data) \neq hash$ in pointer
- This allows us to verify that the information we're pointing to has not changed
 - Before using that data, do a $hash(data)$ and see if it matches the hash in the hash pointer

Hash Pointers: Linked Lists = **Blockchain**

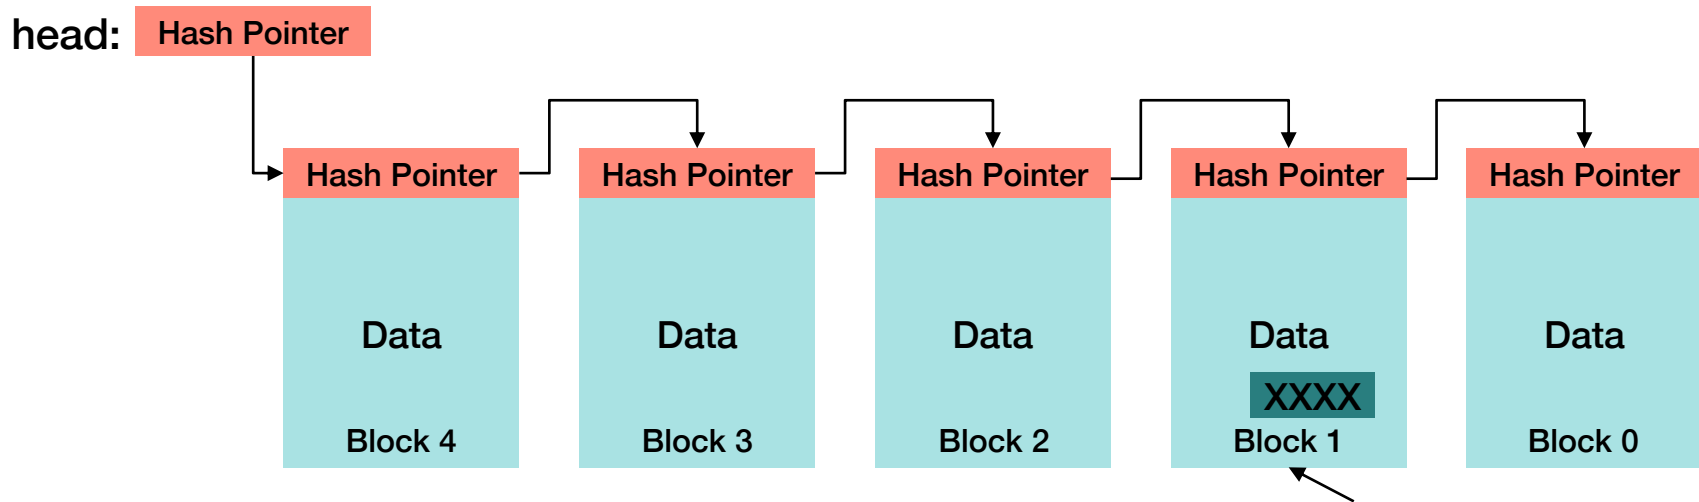


Add new data blocks to the list

- Each hash pointer contains a pointer & a hash of the entire data structure to which it is pointing: the application data and the hash pointer in that structure

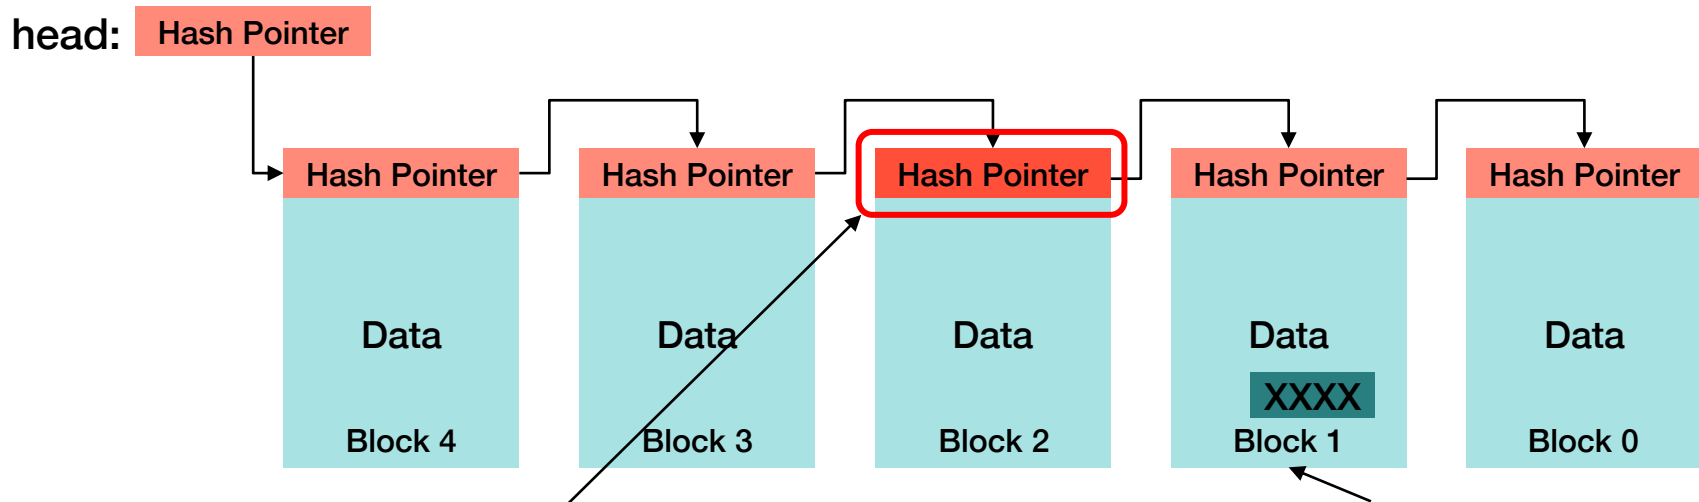
Tamper Evident Log = Blockchain

Tamper detection



Suppose an attacker changes data in this block

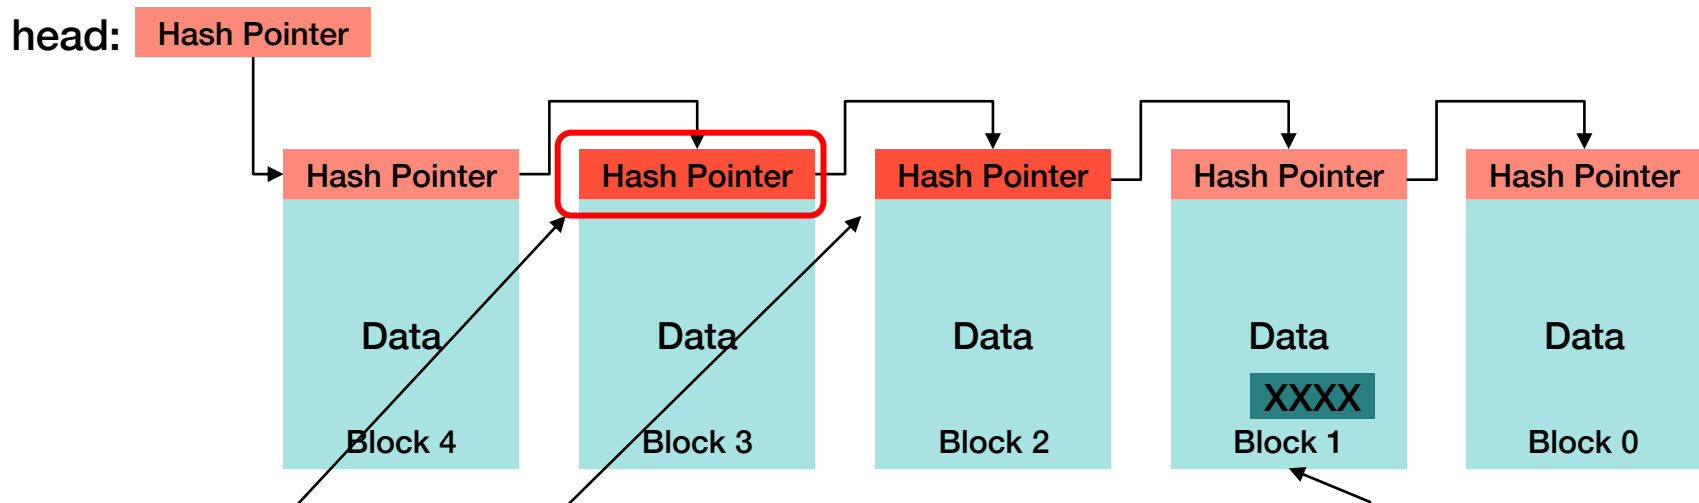
Tamper detection



Suppose an attacker changes data in this block

Then this hash pointer needs to be changed
The attacker needs to update the hash in the pointer to match the hash of the modified block

Tamper detection

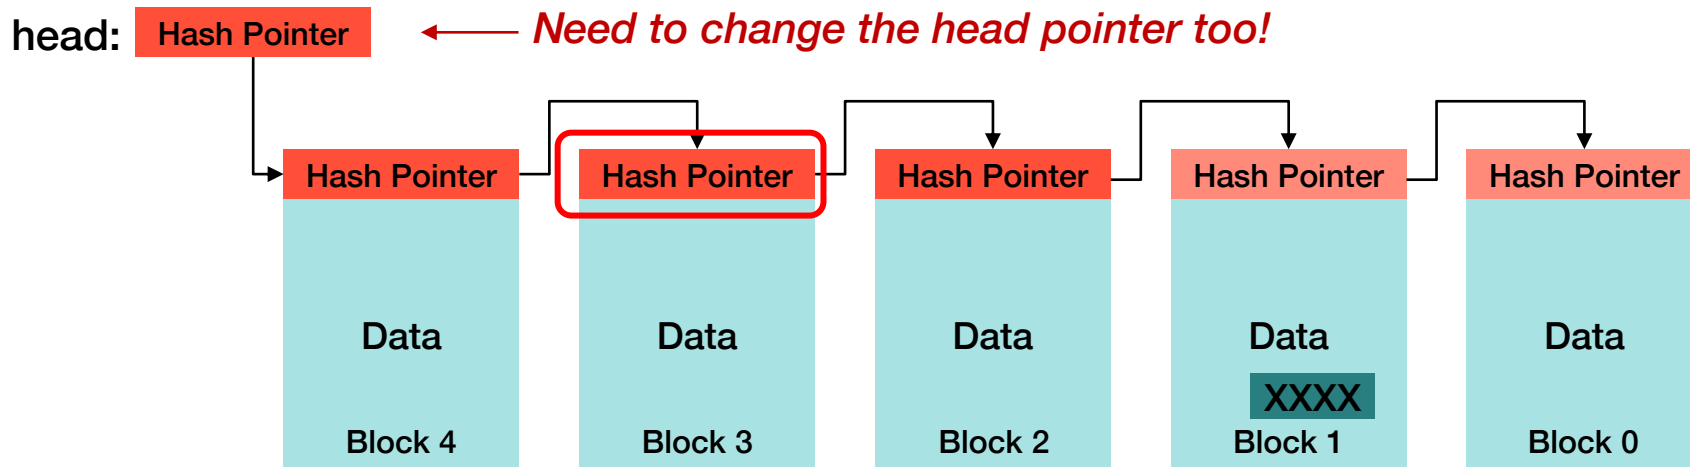


Suppose an attacker changes data in this block

Then this hash pointer needs to be changed
The attacker needs to update the hash in the pointer to match the hash of the modified block

The hash in this pointer is now invalid, so it needs to be updated

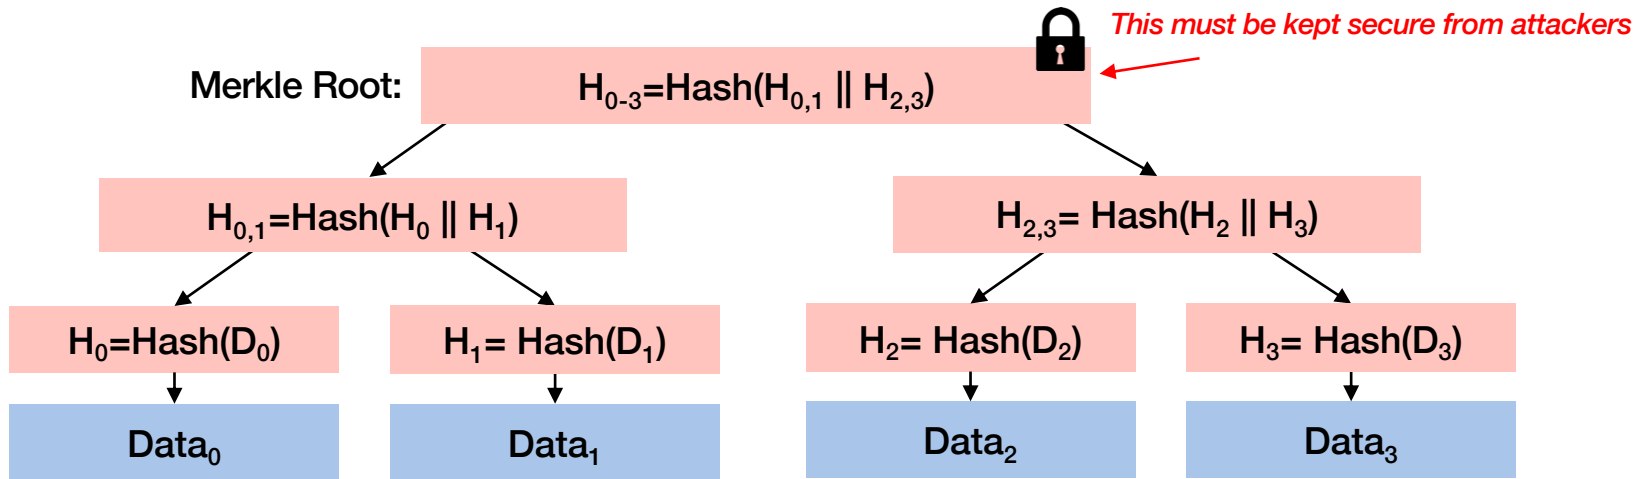
Tamper detection



- The attacker will have to change all hash pointers back to the head
- If we can keep the head of the list safe so an attacker cannot modify it, then we can always detect tampering

Merkle Trees: Binary trees with hash pointers

Merkle Tree hash pointer = { left_subtree, right_subtree, $hash(left \parallel right)$ }

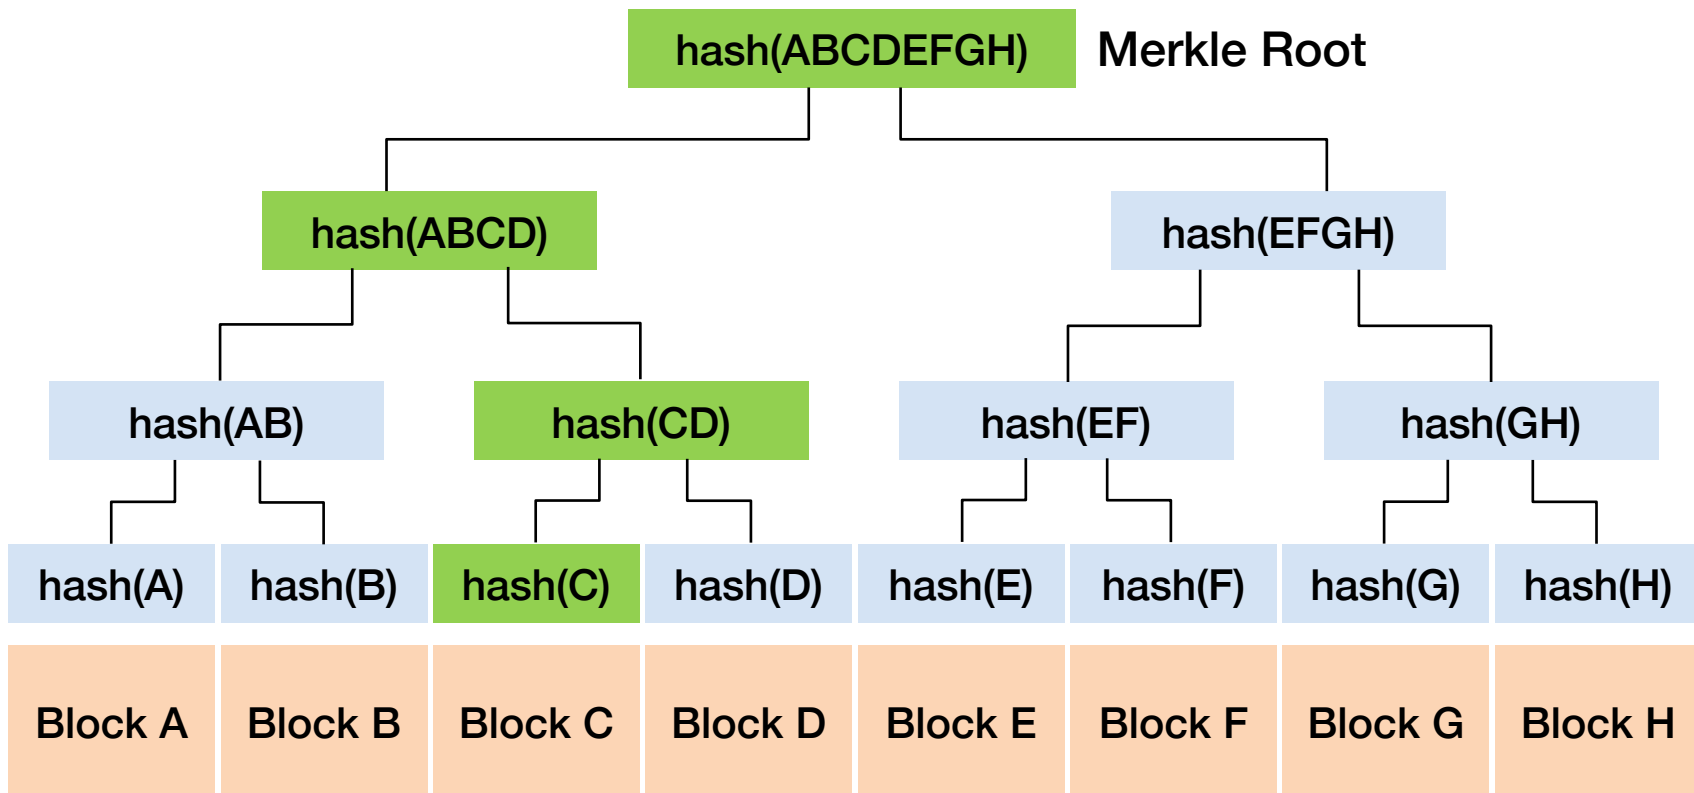


Tamper evident tree

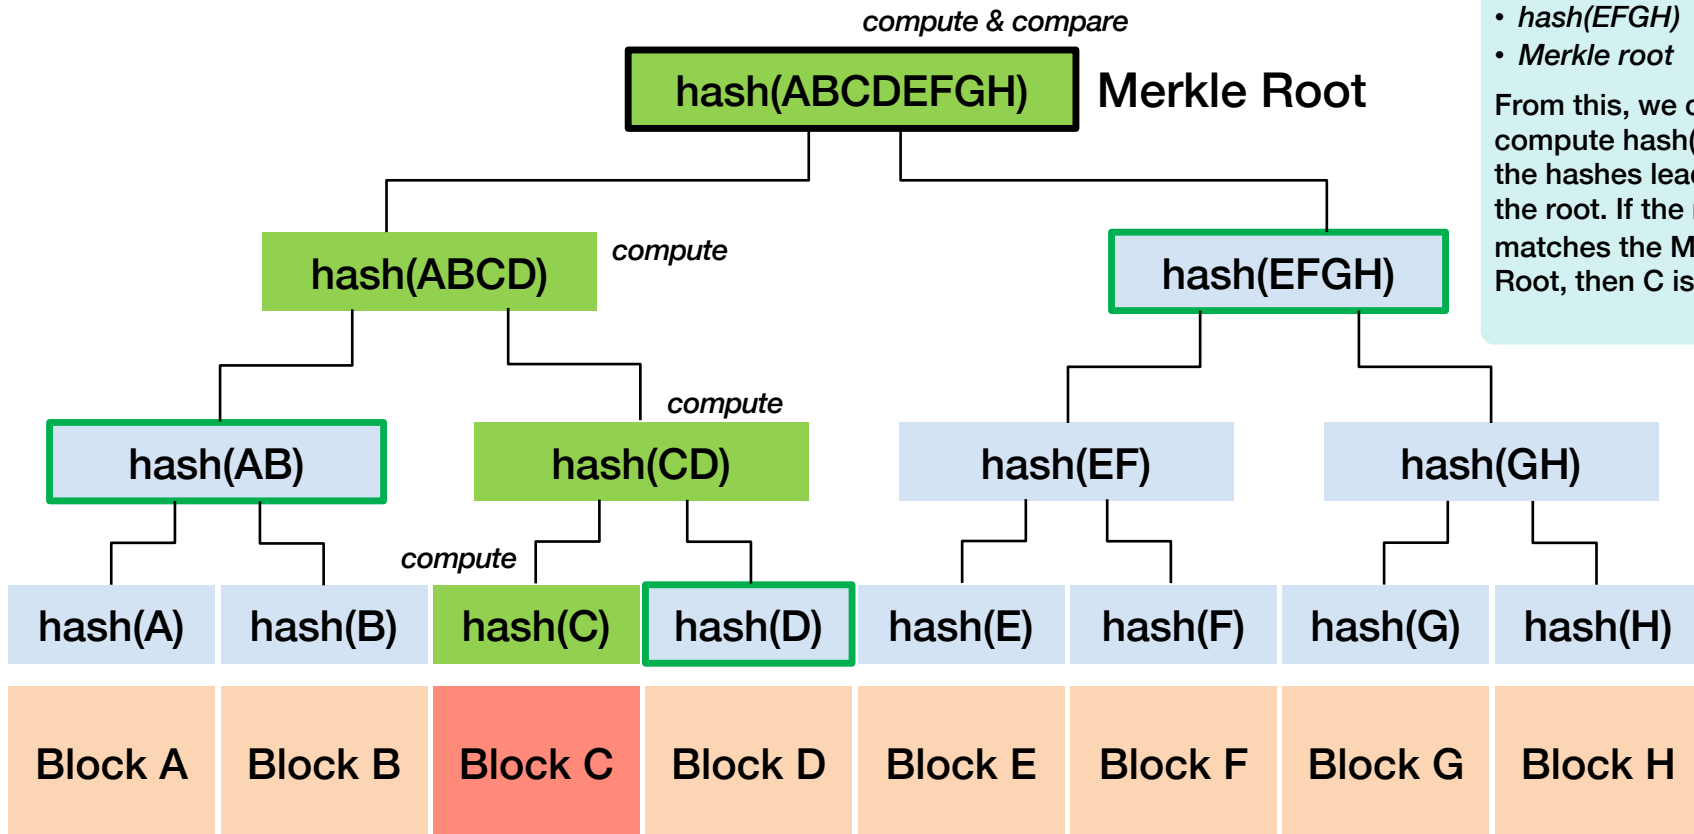
Only need to examine $O(\log_2 n)$ hashes to validate a data block belongs to the tree

$a \parallel b$ means a concatenated with b

Does Block C belong in the file?



Does Block C belong in the file?



We need to have:

- $hash(D)$
- $hash(AB)$
- $hash(EFGH)$
- Merkle root

From this, we can compute $hash(C)$ and the hashes leading to the root. If the root matches the Merkle Root, then C is valid.

Merkle Trees (Hash Trees): Uses

- **Commonly used in peer-to-peer data updates**
- **You receive updated content from an untrusted peer**
 - Validate that the data blocks have not been damaged or modified
 - Don't need to wait for all content to be downloaded
 - **Root hash must be obtained from a trusted place (or signed)**
- **Used in**
 - Version control systems: Git, Mercurial
 - File systems (to detect data damage): ZFS, IPFS
 - Distributed databases: Cassandra, Dynamo, Riak
 - Backup systems: Tahoe-LAFS
 - Decentralized websites: ZeroNet
 - Cryptocurrency: Bitcoin & Ethereum (maybe others)

Bitcoin

Bitcoin Cryptocurrency

- Introduced in 2009
 - anonymously by Satoshi Nakamoto
- First blockchain
- Designed to be public
 - Anyone can participate in the system & use it
 - Users are anonymous
- Currency that is totally separate from any sovereign government
 - Anyone can create money!

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.

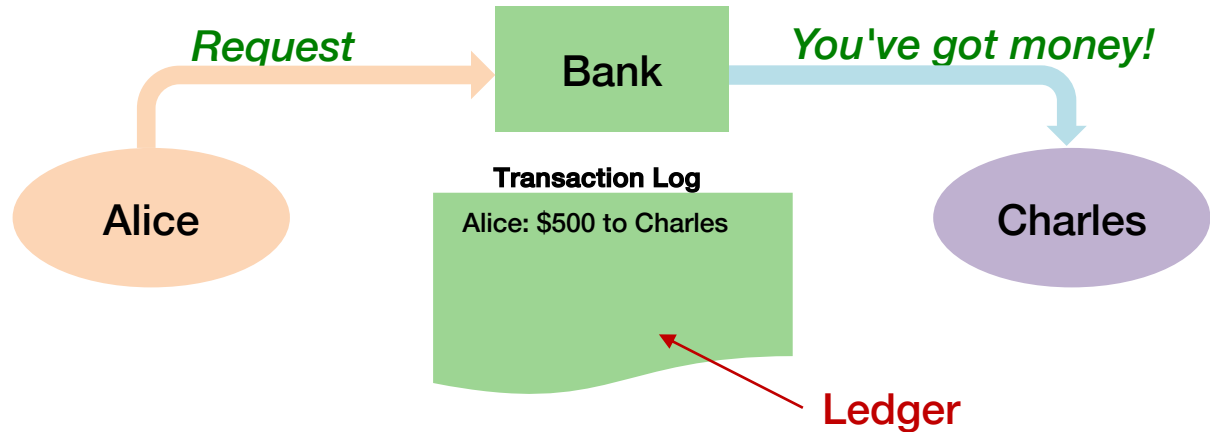
I. Introduction

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 model. Completely non-reversible transactions are not really possible, since financial institutions cannot avoid mediating disputes. The cost of mediation increases transaction costs, limiting the minimum practical transaction size and cutting off the possibility for small casual transactions, and there is a broader cost in the loss of ability to make non-reversible payments for non-reversible services. With the possibility of reversal, the need for trust spreads. Merchants must be wary of their customers, hassling them for more information than they would otherwise need. A certain percentage of fraud is accepted as unavoidable. These costs and payment uncertainties over a communications channel without a trusted party.

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 this paper, we propose a solution to the double-spending problem using a peer-to-peer distributed timestamp server to generate computational proof of the chronological order of transactions. The system is secure as long as honest nodes collectively control more CPU power than any cooperating group of attacker nodes.

Traditional Payments

- **Suppose Alice wants to pay Charles**
 - Send a message to the bank: *transfer \$500 from Alice to Charles*
- **Bank is a trusted third party**
 - Owns register of activity & account balances
 - Only the bank can manipulate the data
 - Also – banks control supply of money



Centralized systems

Transactions are simply modifications to the bank's database

- **We can simply**
 - Subtract \$500 from Alice's account
 - Add \$500 to Charles' account
- **Having a log (ledger) is nice for auditing but not necessary**

Problems?

This is a centralized system

We trust the bank – it is a **trusted third party**

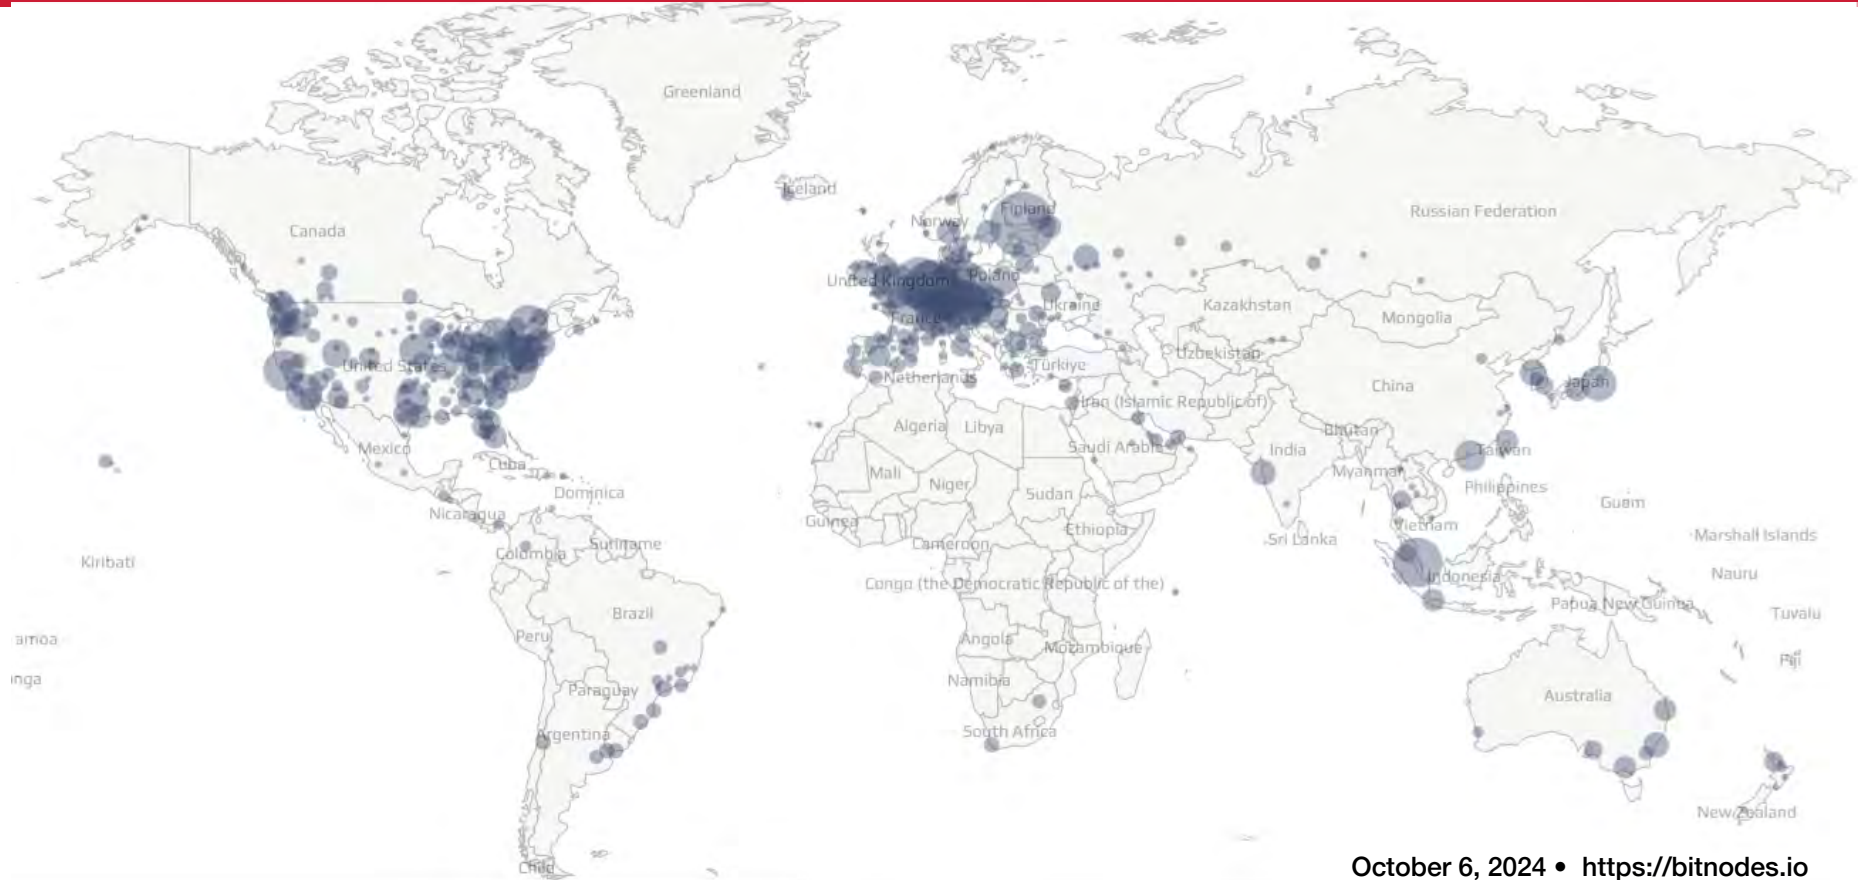
- **What if the bank disappears?**
- **What if the banker makes a mistake?**
- **What if the banker is corrupt?**

Decentralized Solution – Bitcoin

- **Blockchain = ledger = complete list of ALL transactions**
 - Since Bitcoin was started in January 2009
 - 1.09B transactions, 605.9 GB as of Oct 5, 2024
(See <https://www.blockchain.com/en/charts/blocks-size> for the current size of the ledgers)
- **Complete copies of the ledger are replicated around the world**
 - 18,984 nodes (Oct 7, 2024) – all peers – running identical software (See <https://bitnodes.io>)
 - There is **no master node** or master version of the data
- **New systems can do a DNS query for well known peers**
 - Names hardcoded in source (DNS seeds)
 - Return list of IP addresses of bitcoin nodes
 - Then use peer discovery process to find others




Global Bitcoin Nodes



October 6, 2024 • <https://bitnodes.io>

Global Bitcoin Nodes – compute power

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



<https://www.statista.com/statistics/1200477/bitcoin-mining-by-country/>

Identities

- **User creates a {public, private} key pair that defines her wallet**
 - 256-bit Elliptic Curve Digital Signature Algorithm (ECDSA) used
 - **The wallet is just local place for users to store these keys**
 - Wallets may store a transaction list but that's just for user records – the bitcoin network doesn't care
- **Anonymity: bitcoins are associated with keys, not users**
 - Users are anonymous; a user's public key has no association to name
 - The user's identity is called their **address** — directly derived from their public key
 - Users may have multiple keys & multiple addresses
- **Every transaction is signed with the creator's private key**
 - Transaction identifies the user by the public key and can be verified
 - We know only the person with the corresponding private key could have created the request

Nobody to call if you lose your private key!

Bitcoin address = hash(public_key)

Bitcoin uses ECDSA: Elliptic Curve Digital Signature Algorithm

A user creates one or more identities = { private, public } key pairs

You can create an identity (address) for each new transaction

How Bitcoin addresses are created*

1. Generate an ECDSA public, private key pair
2. Create a SHA-256 hash of the public key
3. Perform a RIPEMD-160 hash on that
4. Add a version byte in front of the result
5. Perform a SHA-256 hash on the result ... and a SHA-256 hash on that
6. First 4 bytes of the result = address checksum
7. Add 4 bytes from [6] to the end of the RIPEMD-160 hash from [4]
8. Convert the byte to a base-58 string using Base58Check encoding.
This produces a 20-byte *address*

Why addresses?

A user's bitcoin address is used by other users to identify to send money – as the output of a transaction.

A user's public key could be used as the output in place of the address in transactions.

It is derived directly from the user's public key (hash of the key)

It is shorter than a public key and contains a checksum to detect typing errors.

*You don't have to know this.

Addresses vs. keys

- **Spending: Bob wants to send Alice 5 bitcoins**
 - Bob creates a transaction with a **digital signature** using his private key
 - Presents his public key along with the transaction
 - Any receiving node can validate that the transaction was signed by someone with the corresponding private key
 - The destination of the money is Alice's **address**
- **Addresses are not accounts**
 - They only receive funds
 - You can use those funds if you prove you know the private key that corresponds to the address
 - If Alice wants to use the coins she received
 - She creates a new transaction with her public key & a digital signature
 - Any node can validate that the address belongs to her
 - No node can figure out Alice's public key just by looking at the address

<https://learnmeabitcoin.com/guide/public-key-hash160>

Transactions: Inputs

If Alice wants to send some bitcoin to Charles

- She creates a **transaction** and sends it to one or more bitcoin nodes
- A node tells its peers about the transaction
- Within ~5 sec. every peer on the network has it
- The transaction is currently **unconfirmed**

A blockchain is **NOT** a database – it's a list of transactions

- There are no accounts to query
- Alice needs to provide links to previous transactions that will add up to at least the required amount – these are **inputs**

A node verifies inputs


- Make sure they have not been used by another transaction—this would be **double spending**
- Make sure there is sufficient money in the inputs

Sending

Bob's transaction: send 5 bitcoins to Alice

Hash pointers to previous blocks showing where Bob got paid

Transaction



Alice's transaction: send 5 bitcoins to Charles

The input is a hash pointer to a transaction where Alice received at least 5 bitcoins

Transaction



Addresses vs. keys – Inputs and Outputs

If Alice wants to use the coins she received:

- She creates a new transaction with her public key & a digital signature
- Any node can validate the signature using her public key

Transaction 10732

Output: 1PMycacnJaSqwwJqjawXBEnLsZ7RkXUAs *Alice's address*
Amount: ₿ 0.1
...

Transaction 71991

Output: PWJ2sc9aV72kknbi3R9sjcXVcMXpkdh9Le5
Address to whom she's sending the money

Input:

Source: 10732 *pointer to transaction where Alice got the money*
Public key: 0250863ad64a87ae8a2fe83c1af1a...dad8a04887e5b2352 *Alice's public key*
Signature: a3bb7c5f22079c.... *Alice's signature (using her private key)*

The transaction can be validated by validating each input:

1. Validate the signature using Alice's public key (in the transaction). This proves that whoever created the signature has the private key corresponding to the public key.
2. Hash the public key in the transaction to create the address – see if it matches the address in the referenced transaction
3. No node can figure out Alice's public key just by looking at the address

<https://learnmeabitcoin.com/guide/public-key-hash160>

Transactions: Inputs & Outputs

A transaction contains:

1. One or more **inputs**: transaction IDs & address where coin comes from
 - Contains signature & public key
 - An input is a reference to the output from a previous transaction
2. **Output**: whom the money goes to – destination address & amount
3. **Change**: owner's address & amount
 - Every input must be completely spent
 - Any excess **change** can be generated as another output to the owner of the transaction
4. **Transaction fee** (fluctuates over time: \$ 0.99 on October 6, 2024, \$0.61 on August 17, 2024 ~\$6.72 on Feb 15, 2024; ~ \$14 on Feb 3, 2024)
 - Default transaction fee: 0.0001 bitcoin per 1000 bytes
 - There's a limited amount of space (1 MB) in a block. A transaction is about 250 bytes. To get your transaction processed quickly, you need to outbid others.

The amount of bitcoin you own is the set of transactions in the system that are outputs to your address but have NOT been used as inputs in any transaction

Blocks


Transactions are grouped into blocks

- Each block holds ~2,220 transactions @250 bytes and is ~1.25 MB in size

Bitcoin ledger = linked list of chronologically-ordered blocks


Approximately every 10 minutes, a new block of transactions is added to the blockchain

Genesis block



Each block has

- A link to the previous block
- SHA-256 hash of the previous block

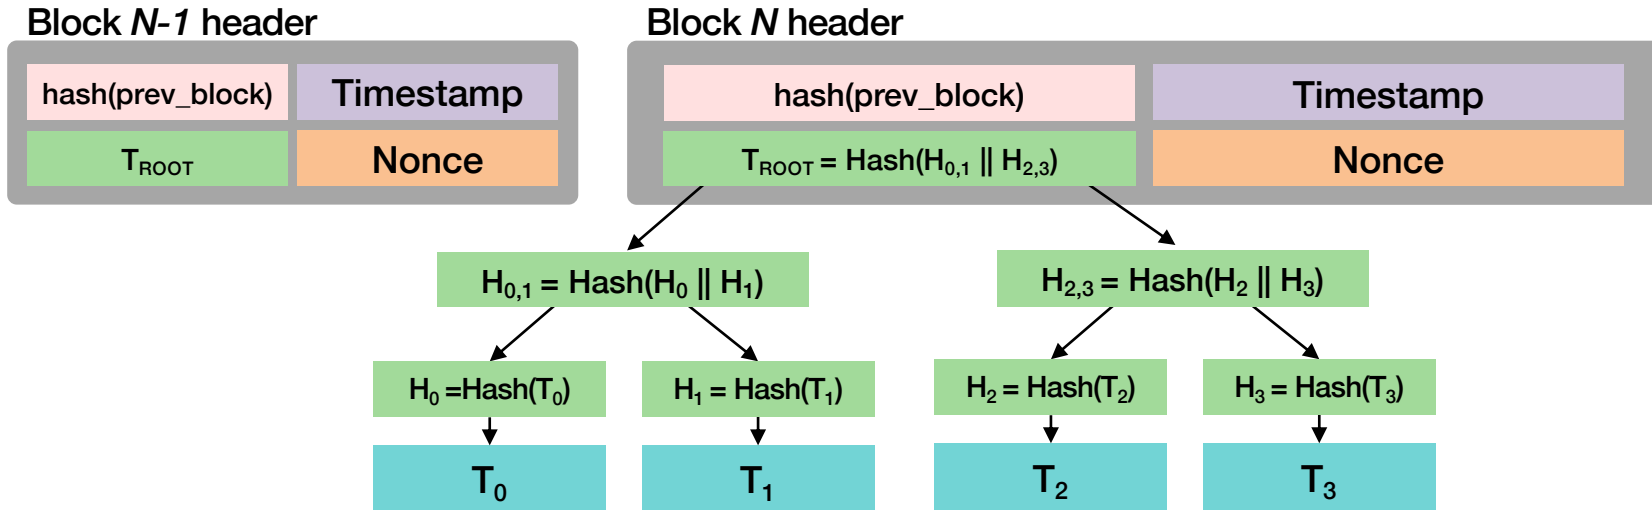


This creates the **blockchain**:
hash pointers – a tamper-evident log

Transactions in blocks: Merkle trees

Transactions within each block are stored in a **Merkle tree**

- A tree makes it easy to find any one of thousands of transactions
- A Merkle tree makes it easy to check if the transaction belongs in the block
- Makes it easy to find and check each input of a transaction to confirm they are valid



Agreement & adding blocks

Each node groups transactions into a block & can propose it as the next block in the blockchain

- Transactions can reach nodes in a different order
- We want all nodes **to agree on the sequence** of blocks in the blockchain

A linked list of hash pointers (blockchain) is a tamper-proof structure

- If the contents of any block are modified, then the hash pointer that points to the block will not be valid (the hash in the pointer won't match)
- ***But can't anyone change the hash pointers?***
We might want to use signed pointers but there's no central authority (no trusted party), so that won't work

Let's create a system where:

- (a) Everyone can agree on the sequence of blocks**
- (b) That sequence cannot be modified**

To add a block to the chain, the hash of the block must meet a certain requirement

Make block addition difficult: *create a puzzle*

If it's sufficiently difficult, there will be a very low chance that multiple nodes will add a block at the same time – and virtually no chance they will be able to do so repeatedly

Suppose we want a hash value to have a specific property:

- Example: the hash should start with with "0000"

There is no algorithmic way to do this

Must try lots of variations of the input

But once found –

it is easy for anyone to verify that the data hashes to the result

Just hash the data and see if the hash starts with "0000"

Solving this "puzzle" is called **mining**

- A block has a 32-bit field in the block where we can try different numbers
- Try to get the block to hash to a desired output
- The resulting number is called the **Proof of Work** – *difficult to generate but easy to verify*

We demonstrate that work has been put into figuring out what the value should be to create the desired hash

Everyone in the network can participate in this

- The first system that finds it **announces the block to everyone else** in the network
- Upon receiving an announcement of a new block:
 1. Each system validates the Proof of Work number against the block
 2. A majority of systems must grant approval
 3. If they do, the block (with the Proof of Work) is made part of the **blockchain**

What's the puzzle?

Bitcoin uses a version of **hashcash** (created in 1997)

– Search for a SHA-256 hash:

$$\text{hash}(\text{block_header}) < \text{target hash}$$

– Choice of *target_hash* sets the difficulty of the problem

Target hash:

**You need to find a proof of work # so that
hash (block) < target hash**

- *Roughly* – the number of leading zero bits in the hash
- Target is a # with a very large # of leading 0s – currently ~17

We change this # and compute hash(block).
If $\text{hash}(\text{block_header}) \geq \text{target hash}$, we try another value of N

Block header

- Version number
- Hash pointer to the previous block
- Merkle root hash (hash of all transactions in the block)
- Timestamp when mining started on the block
- Difficulty **target value** (compact format – 4 bytes)
- **N – nonce:** a 32-bit number

Difficulty Adjustment Algorithm (DAA)

- Bitcoin self-tunes so a new block is added approximately every 10 minutes
- Adjusts for changes in the hashing power of all the miners
- DAA adjusts the Proof of Work **target hash**
 - Selects a predetermined reference block.
 - Compares the timestamp between the reference block and the current block
 - Uses the interval to adjust the target to predict a target that will keep the hash rate at 1 block/10 min

See: <https://chainbulletin.com/proof-of-work-explained-in-simple-terms/>
<https://reference.cash/protocol/blockchain/proof-of-work/difficulty-adjustment-algorithm>

Current Target Hash

You can find the current difficulty by checking the ***bits*** field of the most recent block at <https://whatsonchain.com>

- **Example: The block at 2024-10-07 03:34:21**

Transactions: 78 Size: 33.55KB Nonce: 478580058 Bits: **180ca427**

- **Bits are translated to a Target Hash via a formula:**

See <https://en.bitcoin.it/wiki/Difficulty>

$$0x0ca427 \times 256^{(0 \times 18 - 3)} =$$

0x00000000000000000000000000ca42700

- **If we find a block hash smaller than this (17 leading 0s), we win!**

**We're not solving a difficult problem: we're just computing a hash
It's like a lottery. If our hash is below a certain # then we win**

Isn't a 32-bit nonce too small?

A 32-bit nonce field might not be sufficient


- There might be NO value of the nonce that will produce the right hash

The node then needs to modify other data in the block header & try again

- Make small changes to the timestamp
- New transactions may be added
- If nonce overflows, increment **extraNonce** & reset nonce
 - **extraNonce** is 2-100 bytes
 - Changing **extraNonce** alters the Merkle root hash
 - That needs to be recomputed

How much work is going on?

Oct 6, 2024 hash rate $\approx 665.1 \times 10^{18}$ hashes per second



See <https://bitinfocharts.com/comparison/bitcoin-hashrate.html#3y>

Bitcoin mining

Computing the hash = mining

If you come up with the right answer, you win!

1. You send your block, with the nonce in it, to the whole network
2. Others validate it
3. When a computer validates your block, it adds it to its ledger
4. You get a reward for solving the puzzle! Currently 6.25 BTC; 3.125 in April 2024
5. You also get paid transaction fees in the transactions you put into the block
6. The block (not transactions!) is **confirmed** and you get paid
7. Individual transactions may require the confirmation of multiple subsequent blocks
... More on this later....

Bitcoin mining

The more hashes you can try, the better your chances of winning

- You're competing with every other miner
- **People moved from CPU-based mining to GPU-based**
 - GPU power approximately = 30 CPUs
 - Then FPGA mining: approximately 3-100x faster than GPUs
 - ASIC mining (application-specific integrated circuit):
 - Special hardware built for hash computation: faster & more power-efficient
- **Mining pools = group miners together & share rewards**
 - There are over a dozen large pools for Bitcoin

Mining Hardware

CPU → GPU → FPGA → ASIC



Example:

Antminer AL1 Pro

Computes SHA-256 hashes at 16.6 Th/s
~\$9,900

Consumes 3730 watts

Estimated income: **\$22,781.65/year**

Estimated electricity: **-\$3,920.98/year**

Estimated profit: +\$19,177.09/year

<https://www.asicminervalue.com/miners/ibelink/bm-ks-max>

Mining Hardware

CPU → GPU → FPGA → ASIC



Example:

iBeLink BM-KS Max

Computes SHA-256 hashes at 10.5 Th/s
~\$1,799

Consumes 3400 watts

Estimated income: **\$3,305.89/year**

Estimated electricity: **-\$3,574.08/year**

Estimated profit: -\$222.28/year

<https://www.asicminervalue.com/miners/ibelink/bm-ks-max>

It takes an estimated seven nuclear plants to power our bitcoin mining

That's 21.8 million solar panels worth of production.

Andrew Tarantola – August 26, 2020

Turns out that plugging a bunch of computers into our electrical grid that do nothing but draw current and hash through algorithms has had some negative environmental impacts. Recent studies suggest that Bitcoin-related power consumption has reached record highs this year — with more than seven gigawatts of power being pulled in the pursuit of the suspect digital currency.




A study from the Cambridge Center for Alternative Finance released on Monday estimates that the global bitcoin mining industry uses 7.46 GW, equivalent to around 63.32 terawatt-hours of energy consumption. The study also notes that miners are paying around \$0.03 to \$0.05 per kWh this year. Given that a March estimate put the cost to mine a full bitcoin is around \$7,500, the average miner still stands to make over \$4,000 in profit from the operation.

<https://www.engadget.com/it-takes-an-estimated-seven-nuclear-plants-to-power-our-bitcoin-mining-212441059.html>

Bitcoin Energy Consumption

Bitcoin Energy Consumption



Carbon Footprint

96.08 Mt CO₂



Comparable to the carbon footprint of **Uzbekistan**.

Electrical Energy

172.26 TWh



Comparable to the power consumption of **Poland**.

Electronic Waste

35.15 kt



Comparable to the small IT equipment waste of **the Netherlands**.

Fresh Water Consumption

2,715 GL




Comparable to the total water use of **Switzerland**.



Tom Gara ✓ @tomgara · Oct 20



"Bitcoin and Ethereum are now using up the same amount of electricity as the whole of Austria. Carrying out a payment with Visa requires about 0.002 kilowatt-hours; the same payment with bitcoin uses up 906 kWh, more than half a million times as much"



Blockchain, the amazing solution for almost nothing


Blockchain technology is going to change everything: the shipping industry, the financial system, government ... in fact, what won't it ...

[the correspondent.com](https://thecorrespondent.com)

Competing chains

What if a malicious participant wants to modify an old transaction?

- It will need to modify an old block
- And recompute the Proof of Work (which takes a lot of effort) for the block and *each successive block* (tons of work)
- **The participant will be creating another competing chain in the blockchain**



Competing chains

BUT:

- One malicious participant will not be able to catch up with the cumulative work of all the others
 - The original chain continues to grow as new transactions come in and new blocks are added
- It is expected that some nodes will occasionally have different versions
- **Length of chain = score**

If we two versions of the blockchain,
we select the one that was the hardest to generate (= longest chain)

Blockchain rules state that

The longest chain in the network is the correct one

If a participant receives a higher-scoring version, it overwrites its blockchain with the better data & transmits updates to peers

Producing a longer ledger than the current one requires computing power that competes with the rest of the entire network

51% Attack

If a participant has the majority of the hash rate, the protocol will fail

Blockchain works only because of the assumption that the majority of participants are honest

To double-spend a bitcoin:

- You would need to rewrite the blockchain (change past transactions)
- An attacker would need to control more than 50% of computing capacity
 - **This is a lot:** in 12/17, The Economist estimated *"bitcoin miners now have 13,000 times more combined number-crunching power than the world's 500 biggest supercomputers"*
 - Even if someone tried to do this attack, they'd likely only modify transactions in the past few blocks
- **Keeping history of all transactions among all participants allows anyone to check for double spending**

We started with the usual framework of coins made from digital signatures, which provides strong control of ownership, but is incomplete without a way to prevent double-spending. To solve this, we proposed a peer-to-peer network using proof-of-work to record a public history of transactions that quickly becomes computationally impractical for an attacker to change if honest nodes control a majority of CPU power.

— Satoshi Nakamoto

Confirming transactions

A transaction is **confirmed** after N number of additional blocks are added to the blockchain

- Large values of N are recommended for high-value transactions

The more blocks are added after a transaction, the more difficult it is to modify it

Higher values of N mean that an attacker will need to recompute $N+1$ Proof of Work values to modify the blockchain

- Computationally not feasible

Bitcoin Confirmation Recommendations

- 1: Small payments <\$1,000
- 3: Deposits and payments of \$1,000-\$10,000
- 6: Large payments \$10k-\$1M
- 60: Payments >\$1M

<https://www.buybitcoinworldwide.com/confirmations/>

Computing the Proof of Work takes a lot of work – why do it?

To get earn bitcoin:

- First participant to compute the Proof of Work gets rewarded with bitcoin
- BUT ... only after another 99 blocks have been added to the ledger
- This gives miners an incentive to participate & validate transactions

The reward decreases over time

(assumption: bitcoins will be more valuable)

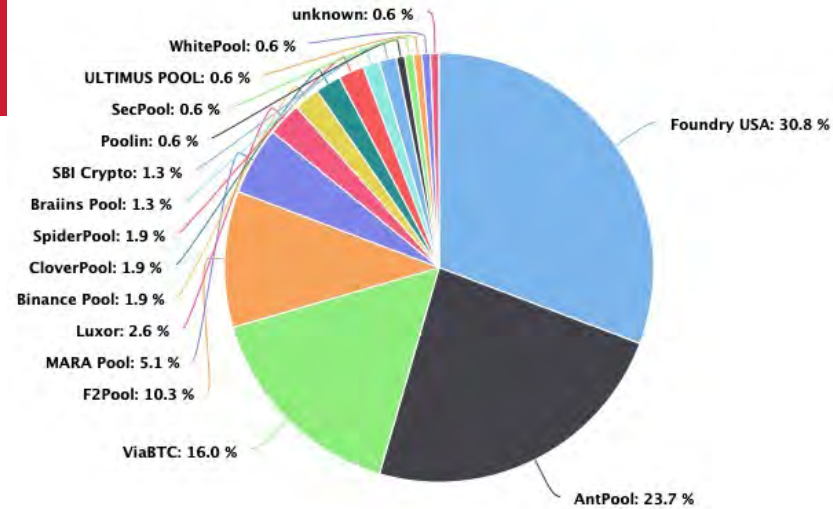
Eventually there will be a maximum of ~21 million bitcoins

Reward	Block #	Date
₿ 50	0	January 3, 2009
₿ 25	210,000	November 28, 2012
₿ 12.5	420,000	July 9, 2016
₿ 6.25	630,000	May 24, 2020
₿ 3.125	840,000	April 18, 2024
₿ 1.625	1,050,000	2028

There are also transaction fees even if the block reward = 0

Centralization

- **Anyone can run a bitcoin node**
 - Requires a good chunk of disk space but is accessible to anyone
- **Mining**
 - Anyone can mine but requires a lot of computing power
 - Not as decentralized as we'd like
- **Software development/support**
 - Open but there's a core set of trusted developers
 - As of 2024, there are only 5 maintainers for Bitcoin Core
 - Only 17 people had the ability to change the code since its launch in 2009
 - Bugs may be fixed ... but transactions cannot be undone
- **Access**
 - People don't run their own Bitcoin nodes
 - Most users trade via large exchanges (Binance, Coinbase, FTX, Kraken, ...)
 - These follow rules on reporting, sanctioning
- **In theory**
 - Teams of sneaky developers may be able to mount an attack
 - Mining pools may try to mount a 51% attack
 - Both scenarios are highly unlikely today



October 2024:
Top 2 pools control 54.5% of the hashrate
Top 5 pools control 85.91% of the hashrate

Source: https://btc.com/stats/pool?pool_mode=day

Decentralization: key findings (2022)

- Now way to optimally distribute a blockchain without a central trusted third party
- A dense subnetwork of bitcoin nodes is responsible for reaching consensus; most nodes do not contribute to the health of the network
- Bitcoin traffic is unencrypted – messages can be observed and dropped
- 60% of bitcoin traffic traverses just three ISPs
- Tor routes traffic for about ½ of bitcoin's nodes. A malicious exit node can modify or drop traffic
- 21% of bitcoin's nodes were running an old version of the bitcoin core client that is vulnerable
- 90% of recently deployed Ethereum smart contracts are at least 56% similar to each other



Decentralization: key findings (2022)

- Trust that programmers won't introduce a bug
- Off-chain code must share the same trust as centralized
- Privileged entities that can modify the semantics of the blockchain to change past transactions
- Most Bitcoin nodes do not participate in mining
- No penalty for dishonesty
- Protocol for coordination within mining pools is unencrypted & unauthenticated
- When any nodes have an incorrect view of the network, it lowers the % of the hash rate needed for a 51% attack



https://assets-global.website-files.com/5fd11235b3950c2c1a3b6df4/62af6c641a672b3329b9a480_Unintended_Centralities_in_Distributed_Ledgers.pdf

51% attack: difficult, not impossible

MIT Technology Review

Once hailed as unhackable, blockchains are now getting hacked

More and more security holes are appearing in cryptocurrency and smart contract platforms, and some are fundamental to the way they were built.

By Mike Orcutt February 19, 2019

Early last month, the security team at Coinbase noticed something strange going on in Ethereum Classic, one of the cryptocurrencies people can buy and sell using Coinbase's popular exchange platform. Its blockchain, the history of all its transactions, was under attack.

An attacker had somehow gained control of more than half of the network's computing power and was using it to rewrite the transaction history. That made it possible to spend the same cryptocurrency more than once—known as “double spends.” The attacker was spotted pulling this off to the tune of \$1.1 million.

<https://www.technologyreview.com/s/612974/once-hailed-as-unhackable-blockchains-are-now-getting-hacked/>

Two Miners Purportedly Execute 51 % Attack on Bitcoin Cash Blockchain

Max Boddy May 25, 2019

Two miners have reportedly executed a 51% attack on the bitcoin cash (BCH) blockchain, according to tweets by Cryptoconomy Podcast host Guy Swann on May 24.

A 51% attack occurs when someone controls the majority of mining power on a Proof-of-Work blockchain network. This means that the majority block verifier can prevent other users from mining and reverse transactions.

While many have assumed that a 51% attack would be carried out with malicious intent, the above case happened as **the two mining pools attempted to prevent an unidentified party from taking some coins that — due to a code update — were essentially “up for grabs.”**

<https://cointelegraph.com/news/two-miners-purportedly-execute-51-attack-on-bitcoin-cash-blockchain>

Achieving anonymity is difficult

CNN

BUSINESS

Markets Tech Media Success Perspectives Video

How Mueller used Bitcoin to catch Russia



By Donie O'Sullivan, CNN Business

Updated 6:17 PM ET, Fri April 19, 2019

The blockchain contains no personally identifiable information. But once someone figures out a user is responsible for one transaction, they can track the entire Bitcoin history.

New York (CNN Business) – Russian operatives used cryptocurrency at almost every stage in their online efforts to interfere in the 2016 U.S. presidential election, according to Special Counsel Robert Mueller's [final report on his investigation](#).

Systems used in the hacking of the Democratic Party were paid for using Bitcoin, as were online hosting services that supported websites which published hacked materials and were used in the targeting of disinformation at American voters. The hacking and disinformation campaigns accounted for the vast majority of Russia's online efforts to influence the 2016 election.

All Bitcoin transactions are posted to an immutable public ledger, known as a blockchain. While the blockchain doesn't contain obvious identifying information about the person behind a transaction, once someone figures out a user is responsible for one transaction it can be possible to track their entire Bitcoin history.

<https://www.cnn.com/2019/04/19/tech/bitcoin-mueller-russia/index.html>

Pipeline Investigation Upends Idea That Bitcoin Is Untraceable

The New York Times

The F.B.I.'s recovery of Bitcoins paid in the Colonial Pipeline ransomware attack showed cryptocurrencies are not as hard to track as it might seem

Nicole Perloth, Erin Griffith and Katie Benner • June 9, 2021

When Bitcoin burst onto the scene in 2009, fans heralded the cryptocurrency as a secure, decentralized and anonymous way to conduct transactions outside the traditional financial system.

...

But this week's revelation that federal officials had recovered most of the Bitcoin ransom paid in the recent Colonial Pipeline ransomware attack exposed a fundamental misconception about cryptocurrencies: They are not as hard to track as cybercriminals think.

On Monday, the Justice Department announced it had traced 63.7 of the 75 Bitcoins — some \$2.3 million of the \$4.3 million — that Colonial Pipeline had paid to the hackers as the ransomware attack shut down the company's computer systems, prompting fuel shortages and a spike in gasoline prices. Officials have since declined to provide more details about how exactly they recouped the Bitcoin.

Yet for the growing community of cryptocurrency enthusiasts and investors, the fact that federal investigators had tracked the ransom as it moved through at least 23 different electronic accounts belonging to DarkSide, the hacking collective, before accessing one account showed that law enforcement was growing along with the industry.

<https://www.nytimes.com/2021/06/09/technology/bitcoin-untraceable-pipeline-ransomware.html>

Feds Arrest an Alleged \$336M Bitcoin-Laundering Kingpin

The alleged administrator of Bitcoin Fog kept the dark web service running for 10 years before the IRS caught up with him.

Andy Greenberg • April 27, 2021

FOR A DECADE, Bitcoin Fog has offered to obscure the source and destination of its customers' cryptocurrency, making it one of the most venerable institutions in the dark web economy. Now the IRS says it has finally identified the Russian-Swedish administrator behind that long-running anonymizing system and charged him with laundering hundreds of millions of dollars worth of bitcoins, much of which was sent to or from dark web drug markets. **What gave him away? The trail of his own decade-old digital transactions.**

US authorities on Tuesday arrested Roman Sterlingov in Los Angeles, according to court records, and charged him with **laundering more than 1.2 million bitcoins**—worth \$336 million at the times of the payments—over the 10 years that he allegedly ran Bitcoin Fog. According to the IRS criminal investigations division, Sterlingov, a citizen of Russia and Sweden, **allowed users to blend their transactions with those of others to prevent anyone examining the Bitcoin blockchain from tracing any individual's payments.** He took commissions on those transactions of 2 to 2.5 percent.

...

Most remarkable, however, is the **IRS's account of tracking down Sterlingov using the very same sort of blockchain analysis that his own service was meant to defeat.** The complaint outlines how Sterlingov allegedly paid for the server hosting of Bitcoin Fog at one point in 2011 using the now-defunct digital currency Liberty Reserve.

<https://www.wired.com/story/bitcoin-fog-dark-web-cryptocurrency-arrest/>

SafeDollar ‘stablecoin’ drops to \$0 following \$248,000 DeFi exploit on Polygon

SafeDollar (SDO) stablecoin proved to be anything but stable today as its price collapsed on the heels of an alleged cyberattack.

Liam Frost • June 28, 2021

The price of SafeDollar (SDO), an algorithmic decentralized finance (DeFi) stablecoin based on the Polygon (MATIC) blockchain, has plummeted to literally zero as a result of what appears to be an exploit today.

While details are yet scarce, block explorer Polygonscan shows that 202,000 USDC and 46,000 USDT stablecoins were suddenly drained from SDO’s smart contract today—worth around \$248,000 in total.

As a result, SafeDollar’s price—which was supposed to always be equal to \$1 since it’s a stablecoin—has plummeted to zero, according to the protocol’s own website.

Stablecoins are a special type of cryptocurrency tokens that are pegged to certain fiat currencies, usually the U.S. dollar. They are designed to always retain the value of their corresponding assets and—in theory—should always be tradeable or redeemable in a one-to-one ratio.

<https://cryptoslate.com/safedollar-stablecoin-drops-to-0-following-248000-defi-exploit-on-polygon/>



A single anonymous market manipulator caused bitcoin to top \$20,000 two years ago, study shows

Michael Sheetz
November 4, 2019

A forensic study on bitcoin's 2017 boom has found that nearly the entire rise of the digital currency at the time is attributable to "one large player," although the market manipulator remains unidentified.

Finance professors John Griffin and Amin Shams – instructors at University of Texas and the Ohio State University, respectively – analyzed over 200 gigabytes of data for the transaction history between bitcoin and tether, another digital currency. Tether is an asset known as a "stablecoin," which has its trading value connected to the dollar.

The professors' study found that tethers being traded for bitcoins revealed a pattern.

One of the SEC's top worries is that crypto is subject to manipulation

A forensic study found that tethers, a digital currency, being traded for bitcoins, revealed a pattern of manipulation during the 2017 cryptocurrency boom.

"Almost the entire price impact can be attributed to this one large player," finance professors John Griffin and Amin Shams wrote.

<https://www.cnbc.com/2019/11/04/study-single-anonymous-market-manipulator-pushed-bitcoin-to-20000.html>

Where are we heading?

- **There are currently over 10,000 cryptocurrencies**
 - Most will fail (over 14,000 already failed)
- **Some are tied to real currency**
 - Tether, Binance USD, Pax Dollar, USD Coin, ... are stablecoins backed by U.S. \$
 - Designed to park funds during times of high volatility
 - But Terra & Luna collapsed, wiping almost \$500B in May 2022
 - USDC lost its peg to the \$ after customers withdrew billions
 - 2021: 2.9% of Tether backed by cash; >65% backed by commercial paper (e.g., T-bills)
- **Privacy tokens (e.g., Monero, Dash)**
 - Obscure origins & destination: focus on *anonymity* and *untracability*
 - **Stealth addresses**: new address for each transaction
 - **Coin mixing** (CoinJoin): merge transactions from different individuals into one transaction and disburse them to users at new addresses
 - **Zk-SNARKs** (Zero-Knowledge Succinct Non-Interactive Argument of Knowledge): prove transaction validity without divulging identifying information

Targeted by anti money laundering regulators

South African Brothers Vanish, and So Does \$3.6 Billion in Bitcoin

Bloomberg

Lawyers hired by investors piece together funds' disappearance
Africrypt loss would rank among biggest-ever crypto heists

Roxanne Henderson and Loni Prinsloo • June 23, 2021

A pair of South African brothers have vanished, along with Bitcoin worth \$3.6 billion from their cryptocurrency investment platform.

A Cape Town law firm hired by investors says they can't locate the brothers and has reported the matter to the Hawks, an elite unit of the national police force. It's also told crypto exchanges across the globe should any attempt be made to convert the digital coins.

Following a surge in Bitcoin's value in the past year, the disappearance of about 69,000 coins -- worth more than \$4 billion at their April peak -- would represent the biggest-ever dollar loss in a cryptocurrency scam. The incident could spur regulators' efforts to impose order on the market amid rising cases of fraud.

The first signs of trouble came in April, as Bitcoin was rocketing to a record. Africrypt Chief Operating Officer Ameer Cajee, the elder brother, informed clients that the company was the victim of a hack. He asked them not to report the incident to lawyers and authorities, as it would slow down the recovery process of the missing funds.

...

The firm's investigation found Africrypt's pooled funds were transferred from its South African accounts and client wallets, and the coins went through tumblers and mixers -- or to other large pools of bitcoin -- to make them essentially untraceable.

<https://www.bloomberg.com/news/articles/2021-06-23/s-african-brothers-vanish-and-so-does-3-6-billion-in-bitcoin>

What's Blockchain Actually Good for, Anyway? For Now, Not Much

Not long ago, blockchain technology was touted as a way to track tuna, bypass banks, and preserve property records. Reality has proved a much tougher challenge.



Story of the day

21 August 2020 · Reading time 14 · 09 minutes · *News analysis letter*

Blockchain technology is going to change everything: the shipping industry, the financial system, government ... in fact, what won't it change? But enthusiasm for it mainly stems from a lack of knowledge and understanding. The blockchain is a solution in search of a problem.

Blockchain, the amazing solution for almost nothing

<https://thecorrespondent.com/655/blockchain-the-amazing-solution-for-almost-nothing/86714927310-8f431cae>

Report: North Korea steals cryptocurrency to fund nuclear program



Karen Hoffman • February 7, 2022

Truth is stranger than fiction.

In the 2014 movie "The Interview," tabloid TV journalists (played by Seth Rogen and James Franco) travel to interview Kim Jong-un and learn that lots of the weirdness in the life of the leader of North Korea is not just accurate, it is less shocking than reality.

North Korea has stolen roughly \$50 million in cryptocurrency between 2020 and mid-2021, according to a report prepared for the United Nations Security Council's North Korea sanctions committee, Reuters reported on Saturday. That amount is only a small fraction of the \$400 million stolen by North Korean hackers. according to a January report by tech consultancy Chainalysis, or the \$300 million that was assessed in a U.N. report in October.

“According to a member state, DPRK [Democratic People's Republic of Korea] cyberactors stole more than \$50 million between 2020 and mid-2021 from at least three cryptocurrency exchanges in North America, Europe and Asia,” the Feb. 5 U.N. report stated, according to Reuters.

North Korean hackers have been targeting U.S. and foreign exchanges and bridge sites, which convert crypto from one denomination to another, realizing that these payment sites offer a rich opportunity to deliver money to suspect programs like the country's missile development. According to widespread reports, North Korea employs at least 7,000 hackers to collect money and data and attack other countries' financial and critical infrastructure systems, like power plants and electrical systems (witness Colonial Pipeline).

<https://www.scmagazine.com/news/cryptocurrency/report-north-korea-steals-cryptocurrency-to-fund-nuclear-program>

Hackers rob thousands of Coinbase customers using MFA flaw

Lawrence Abrams • October 1, 2021

Crypto exchange Coinbase disclosed that a threat actor stole cryptocurrency from 6,000 customers after using a vulnerability to bypass the company's SMS multi-factor authentication security feature.

...

To conduct the attack, Coinbase says the attackers needed to know the customer's email address, password, and phone number associated with their Coinbase account and have access to the victim's email account.

...

However, Coinbase states a vulnerability existed in their SMS account recovery process, allowing the hackers to gain the SMS two-factor authentication token needed to access a secured account.

"Even with the information described above, additional authentication is required in order to access your Coinbase account," explained a Coinbase notification to customers seen by BleepingComputer.

"However, in this incident, for customers who use SMS texts for two-factor authentication, the third party took advantage of a flaw in Coinbase's SMS Account Recovery process in order to receive an SMS two-factor authentication token and gain access to your account."

<https://www.bleepingcomputer.com/news/security/hackers-rob-thousands-of-coinbase-customers-using-mfa-flaw/>

Crypto exchange BitMart loses \$196 million to hackers

The theft might be difficult to track.

J. Fingas • December 5, 2021

Large-scale cryptocurrency heists remain a significant headache. According to *Coindesk*, the crypto exchange BitMart has lost the equivalent of \$196 million (originally estimated at \$150 million) to a hack. The intruder breached Ethereum and Binance wallets with a flood of transfers starting around 2:30PM Eastern on December 4th, followed by an exodus of tokens two hours later that included Shiba and USDC.

Founder Sheldon Xia said only a "small percentage" of BitMart's assets were at risk. Even so, the company has frozen withdrawals "until further notice" and is reviewing security.

It's not clear who was responsible, but the culprit may have been knowledgeable. The stolen funds have been sent to an Ethereum mixing service that could make it difficult to trace the funds. Crypto thieves aren't always that astute. The Poly Network attacker, for instance, offered to "surrender" and wound up returning all their loot. They claimed they were contributing to Poly's security, but that might also have been an attempt to avoid repercussions after researchers obtained potentially identifying data.

<https://www.engadget.com/crypto-exchange-bitmart-hack-theft-222846482.html>

Crypto exchanges keep getting hacked, and there's little anyone can do



One of the biggest heists happened this month, when the crypto trading platform Bitmart said hackers stole almost \$200 million after they broke into a company account.

Kevin Collier • December 17, 2021

It's not just lucky investors getting rich from crypto.

Hackers have made off with billions of dollars in virtual assets in the past year by compromising some of the cryptocurrency exchanges that have emerged during the bitcoin boom.

There have been more than 20 hacks this year where a digital robber stole at least \$10 million in digital currencies from a crypto exchange or project. In at least six cases, hackers stole more than \$100 million, according to data compiled by NBC News. By comparison, bank robberies netted perpetrators an average of less than \$5,000 per heist last year, according to the FBI's annual crime statistics.

Despite the large dollar amounts associated with these thefts, they often lack the drama or attention of traditional bank robberies. But cryptocurrency experts say they offer a warning to would-be crypto investors: Exchanges are now lucrative targets for hackers.

<https://www.nbcnews.com/tech/security/bitcoin-crypto-exchange-hacks-little-anyone-can-do-rcna7870>

Ukraine legalizes cryptocurrency as it receives millions in crypto donations

The Ukrainian government has been actively soliciting crypto donations

James Vincent • March 17, 2022

Ukrainian President Volodymyr Zelenskyy has signed into law a bill that effectively legalizes the cryptocurrency sector in the country. The decision comes as Ukraine has received cryptocurrency donations worth tens of millions of dollars from individuals and groups hoping to help the country's war effort against Russia.

The bill signed by Zelenskyy was approved by Ukraine's parliament last month and "creates conditions for the launch of a legal market for virtual assets in Ukraine." It allows Ukrainian banks to open accounts for crypto firms; appoints the National Bank of Ukraine and the National Commission on Securities and Stock Market as financial watchdogs for the sector; and, as reported by CoinTelegraph, means crypto exchanges and companies that handle other virtual assets will have to register with the government. The state says it will protect citizens' cryptocurrency holdings with the same legal force as its fiat currency, the hryvnia.





<https://www.theverge.com/2022/3/17/22982608/ukraine-cryptocurrency-sector-legalized-zelenskyy-signs-bill>

Opinion

Bitcoin: Gold 2.0? Try Reserve Asset 3.0

The conflict between Russia and Ukraine is beginning to send ripples through the global economy that could lead to a new monetary system.

By **George Kaloudis**  CoinDesk Insights

 Mar 13, 2022 at 10:50 a.m. EDT

Updated Mar 17, 2022 at 10:30 a.m. EDT



<https://www.coindesk.com/markets/2022/03/13/bitcoin-gold-20-try-reserve-asset-30/>

Beeple sold an NFT for \$69 million

THE VERGE


Through a first-of-its-kind auction at Christie's

By Jacob Kastrenakes • March 11, 2021

Until October, the most Mike Winkelmann — the digital artist known as Beeple — had ever sold a print for was \$100.

Today, an NFT of his work sold for \$69 million at Christie's. The sale positions him “among the top three most valuable living artists,” according to the auction house.

The record-smashing NFT sale comes after months of increasingly valuable auctions. In October, Winkelmann sold his first series of NFTs, with a pair going for \$66,666.66 each. In December, he sold a series of works for \$3.5 million total. And last month, one of the NFTs that originally sold for \$66,666.66 was resold for \$6.6 million.



Beeple's collage, Everydays: The First 5000 Days

<https://www.theverge.com/2021/3/11/22325054/beeple-christies-nft-sale-cost-everydays-69-million>

GOVTECH BIZ

Vermont City, Real Estate Startup Try Out Blockchain for Recording Property Transactions

Examples of local government using blockchain are pretty hard to come by, but one city is testing the technology on property transactions.

• Ben Miller

Among the many possible uses of blockchain, techno-enthusiasts have long thought of property record tracking as a prime example of how government could use the technology.

Now, a company is working with a Vermont city to test the concept. Since blockchain is still very much an emerging technology, it represents an early example of a city government trying the technology out.

<https://www.govtech.com/biz/vermont-city-real-estate-startup-try-out-blockchain-for-recording-property-transactions.html>

Proof of Stake vs. Proof of Work

- **Primary goal: Better energy efficiency (reduced computation)**
- **Users stake their coins to become **validators****
 - Validators do the same thing as miners in proof-of-work
 - If chosen: Create and order transactions in a new block
 - Otherwise: validate new blocks
- **Staking = putting aside coins to act as a validator**
 - Validators chosen at random as $f(\text{size of stake})$
 - Incentivizes good behavior: you lose stake for attesting to malicious blocks
 - Get rewarded for creating and for attesting (validating) proposed blocks
- **Ethereum requires a committee of at least 128 validators to attest a new block – at least 2/3 of validators must agree for a block to be finalized**
 - Validators will lose their stake if they try to revert this later via a 51% attack

Smart Contracts

- **Executable programs in a blockchain transaction**
 - Executed by a quorum of blockchain nodes
 - Updates the state of the blockchain if execution is successful
- **All honest nodes must produce the same result**
 - Same input values (the blockchain), same execution logic
- **Code may be stored on the blockchain or outside with a hash pointer to the code in the blockchain**
 - Ethereum stores code on the blockchain; Hyperledger Fabric & Corda outside
- **Functions that update state include**
 - Transfer, approve, transfer_from, mint, burn
 - These generate transactions that are processed as new transactions on the blockchain

The End