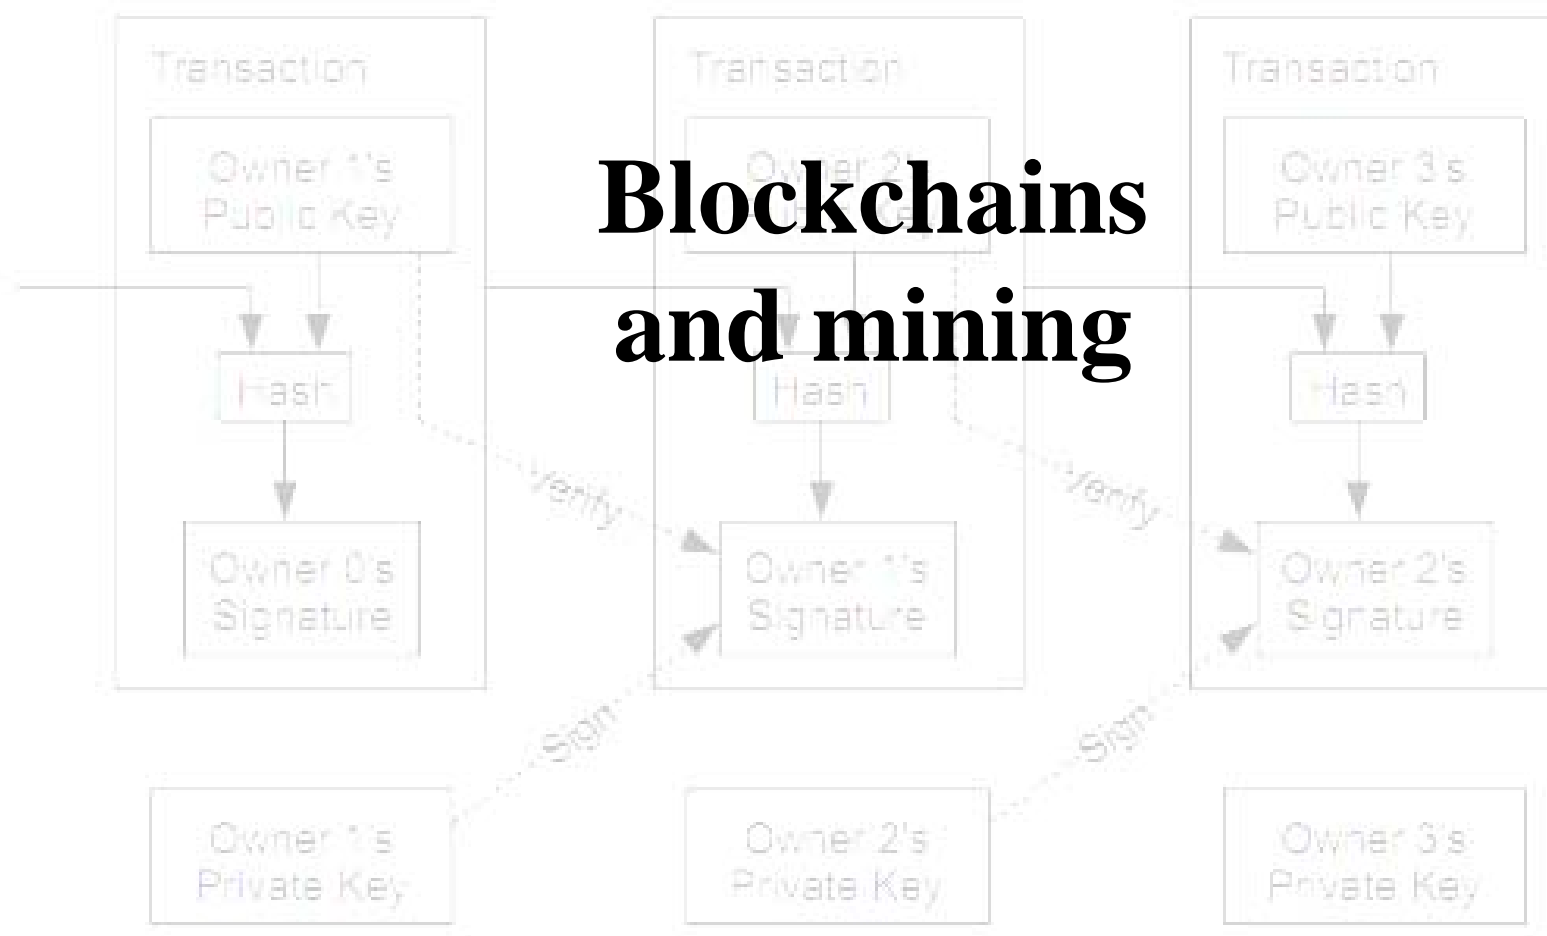


Blockchains and mining




The four major innovations that Nakamoto designed into Bitcoin


- Using double-key cryptography to make secure transfers of assets from one wallet to another (British intelligence, 1970s)
- Recording new data sequentially in a write-only, indelible ledger, the “blockchain” (IBM, 1976)
- Decentralizing the ledger to provide transparency of data to all users and interested third parties (Bell Labs, 1991)
- Validating new data by cryptographic “consensus” proof, in recurring 10-minute open competitions, instead of relying on a trusted third party (Nakamoto, 2008)

“Message verification and transmission error detection by block chaining”

U.S. patent granted to IBM scientists in 1976

Publication number	US4074066 A
Publication type	Grant
Application number	US 05/680,404
Publication date	Feb 14, 1978
Filing date	Apr 26, 1976
Priority date 	Apr 26, 1976
Also published as	CA1100588A , CA1100588A1 , DE2715631A1 , DE2715631C2
Inventors	William F. Ehram , Carl H. W. Meyer , John L. Smith , Walter L. Tuchman
Original Assignee	International Business Machines Corporation
Export Citation	BiBTeX , EndNote , RefMan
Patent Citations (5), Referenced by (52), Classifications (10)	
External Links: USPTO , USPTO Assignment , Espacenet	

Logic of a blockchain



Source: Ethereum white paper

- Each transaction n is *encrypted* into $Tx(n)$.
- Each new block n includes:
 - The new transaction, $Tx(n)$
 - An encryption of the previous block, $S(n - 1)$.
- Two implications of this structure:
 - Even if $Tx(1) = Tx(2)$, we will have $S(1) \neq S(2)$, making it impossible to recover the raw data
 - If $Tx(n)$ is changed, **every block $n, n+1, n+2, \dots$, will also change**

Encrypting data with hash functions

Developed at IBM in early 1950s

Input

“NYU Stern School of Business”

“N.Y.U. Stern School of Business”

“N.Y.U.”

Output (SHA3-256)

d17deb9093208f1d2dc201d5d2455f68
860e089896d1f43b4b7239755ca8d708

de52a08cf9418d8facac03eb21e2b683a
8ce53a03f051a900efb0210da34892a

27716e313c1944ddd72f4561fe82955f
083bc304819059e2fa2ca8a41fed5047

What is a hash function?

- A “digital fingerprint”
 - A person can prove their identity by matching their fingerprint with one stored in a database
 - Someone who breaks into the database and steals the fingerprint cannot use it to re-create the person, or even tell what the person looks like

Input to a hash function

- Anything that can be stored in digital form
 - Text
 - Data
 - Video, music, photographs
 - Fingerprints, irises
 - Etc.
- Limit: 2.09 exabytes
 - *So large it would take 220 years just to read in*

What does a hash function do with the input?

- Input is converted to a hexadecimal “hash” by scrambling it in a way that is impractical to invert
 - For instance, “Take every third digit in the file, multiply that number by 7, add the digits together and divide the total by every fourth number in the file. Append every number not used in the previous calculation to the number you have, etc. . .”



Output from a hash function

- Fixed length, generally 64 or 128 characters
 - Tells you nothing about the length of the input
- 16 possible characters in each space
 - Digits 0 through 9, letters a through f
 - Number of possible outputs = 16^{64} or approximately 1.16×10^{77}
- Small changes in input will drastically change the output
 - Cannot use patterns of characters in the output as a roadmap for recovering the input, *even if you know the hash function that generated the output*
 - Only trial-and-error decryption will work

Test this yourself

- <https://emn178.github.io/online-tools/sha256.html>
(select different hash functions on the right side of page)

Input: Lincoln's *Gettysburg Address*

Output: SHA 256:
73203f72f551bd1326fc42e3acc03d98bfdbad5f5eb460d1796bd93101a0f250

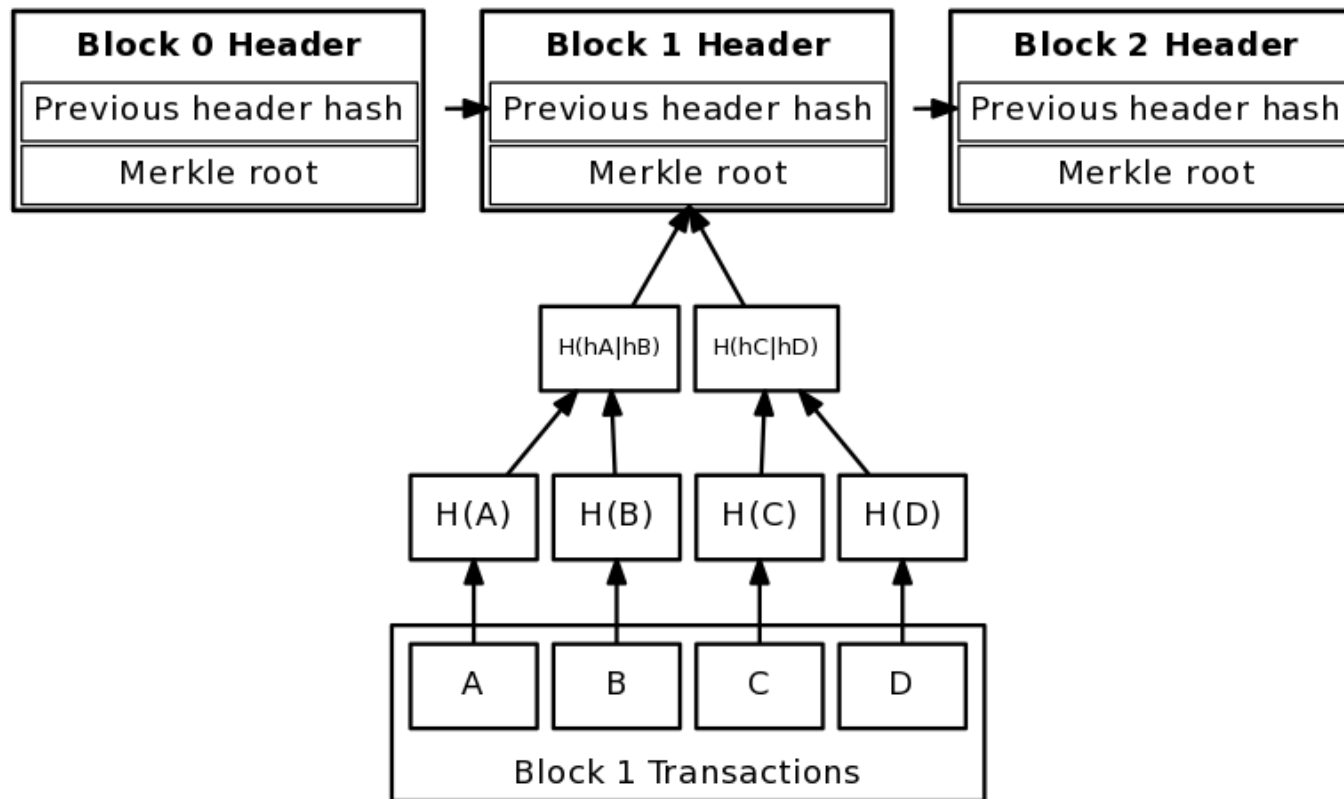
 SHA3-256:
b5d396a63292a26c2e2a4dacfa30f368227b35f69d4804757d1b415e3174b3f1

 SHA3-512:
7a1db3e95f6d25c440e328109a6c185fd4a6eac5db79cfe52f979fe2b9f7b375
2678ce7d35094f9b23aca730946b9f10f0d550344cf89cdca985fedd10a19e84

What are hash functions good for?

- Authentication
 - A password-protected website would store the hash function of your password, rather than storing the password itself
 - A “digital passport” could be a hash function of your fingerprint stored by the Dept. of Homeland Security
- Security
 - If somebody steals the hash function, they cannot invert it to recreate the input data

Storing large amounts of data: bundle it into blocks using Merkle trees



Merkle tree connecting block transactions to block header merkle root

Records in a block of the Bitcoin blockchain

Time	Digital Signature(s) used in current transaction:	Source Address (controlled by current signatory)	Reference to prior transaction	Recipient Address	Data	Bitcoins at source address prior to transaction	Bitcoins Sent to Recipient	Fee to Verif Agent	Signature(s) required for next transaction:
2:59:38 PM	<i>Tammy Tone</i>	1Zefew	←---	1estgE	[a secret]	0.050	0.020	0.015	Person A or B
2:53:31 PM	<i>John Smith</i>	1wEfet	←	1ewYUe	null	25.000	6.000	0.010	Frank Xao
2:52:37 PM	<i>Joe Bookie</i>	1Nuyts	←	1wEfet	[bet winner]	87.500	25.000	0.020	John Smith
2:52:25 PM	<i>John Smith</i>	1EWseg	←	1Nuyts	[sports bet]	12.515	12.500	0.015	Joe Bookie
2:51:04 PM	<i>Frank Heinz</i>	1Wefvs	←---	1EWseg	null	18.000	12.515	0.015	John Smith

Links to addresses further down in the blockchain

Not all entries are required at all times but some must always be included (examples: signatures, references to prior updates)

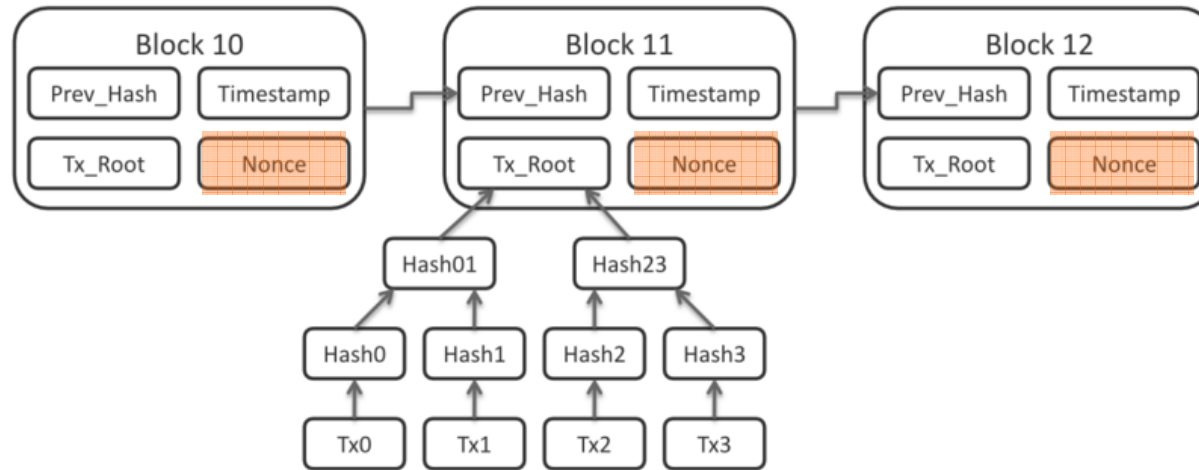
Source: SolidX Partners Inc.

Who updates the blockchain?

- Haber and Stornetta (1991)
 - A *trusted third party* takes responsibility for coding blocks
 - The chain is posted publicly, becoming a *distributed ledger* that can be verified by anyone
- Nakamoto's (2008) crowd-sourcing solution
 - *Network members compete* to create new blocks
 - Anyone can join the network and take part
 - A reward goes to the fastest
(seigniorage of new coins)

A blockchain with “proof of work”

Nakamoto (2008)




- A valid “nonce” must be discovered by trial-and-error, such that the hash function for the entire block begins falls below a critical value
 - *Creates a high cost for hackers*
 - *Proves that the miner has invested resources in the security of the network.*

Block size

7-day moving average

Maximum block size was 1.000 MB,
which is about 2,000 transactions.
Introduced secretly by Nakamoto in
2010 to deter industrial-scale mining:
<https://cointelegraph.com/news/satoshis-best-kept-secret-why-is-there-a-1-mb-limit-to-bitcoin-block-size>




Segregated Witness activated by miners' consensus
on 24 August 2017 to circumvent block size limit

<https://blockchain.info/charts/avg-block-size>

Mempool size


7-day moving average, May 2016 – Feb 2019



<https://blockchain.info/charts/mempool-count>

Median transaction confirmation time


7-day moving average, transactions with mining fees only



<https://blockchain.info/charts/median-confirmation-time>

User fees and confirmation speed

February 24, 2019




The fastest and cheapest transaction fee is currently **32 satoshis/byte**, shown in green at the top. For the median transaction size of **257 bytes**, this results in a fee of **8,224 satoshis**.

$$257 \times 32 \times (\$5,300 \times 10^{-8}) = \$0.45 \text{ fee for a typical transaction}$$

<https://bitcoinfees.com/>


Recalibrating the time required to mine a new block

[illegible]

Recalibrating time to mine each block

- As the mining difficulty changes during a 2,016-block cycle (blue series), difficulty is re-set for the next 2,016 blocks (red series) such that the new expected block time equals 600 seconds. It then could drift upward or downward.

Goal



<https://bitcoinwisdom.com/bitcoin/difficulty>

Mining technology

- Central Processing Units (CPUs) of a high-end laptop or desktop, circa 2009
 - **14 million** SHA-256 hashes per second
- Graphics Processing Units (GPUs), circa 2010
 - **850 million** per second, 40 times faster than CPUs
 - Designers began building PCs that had up to seven GPUs. These required improvised housing in milk crates, etc.
 - Interesting that this occurred barely after “Bitcoin Pizza Day” (May 22, 2010) and before active markets existed for trading Bitcoin



Satoshi Nakamoto on mining technology

2009 posting

- “We should have a gentleman's agreement to postpone the GPU arms race as long as we can for the good of the network.
- “It's much [easier] to get new users up to speed if they don't have to worry about GPU drivers and compatibility.
- “It's nice how anyone with just a CPU can compete fairly equally right now.”



Mining technology

- Purpose-built mining hardware, circa 2011
 - Fivefold reduction in power consumption
- Specialized application specific integrated circuit (ASIC) chips
 - Avalon 1 (2013), **66 GH** per second and tenfold reduction in power consumption
 - This created a huge initial advantage for early adopters
 - Mining revenue of approximately *\$300 per chip per day* initially
 - Competed away by market entry within three months
 - Vast technical improvements continuously since then

Source: Joshua Cain

Mining hardware

Useful life \approx 1 to 2 years

Bitcoin Mining Hardware Comparison

Currently, based on (1) price per hash and (2) electrical efficiency the best Bitcoin miner options are:

AntMiner S7



Advertised Capacity:
4.73 Th/s

Power Efficiency:
0.25 W/Gh

Weight:
8.8 pounds

Guide:
Yes

Price:
\$479.95



Appx. BTC Earned Per Month:
0.1645

AntMiner S9



Advertised Capacity:
13.5 Th/s

Power Efficiency:
0.098 W/Gh

Weight:
8.1 pounds

Guide:
Yes

Price:
\$2,279.95



Appx. BTC Earned Per Month:
0.3603

Avalon6



Advertised Capacity:
3.5 Th/s

Power Efficiency:
0.29 W/Gh

Weight:
9.5 pounds

Guide:
No

Price:
\$499.95




Appx. BTC Earned Per Month:
0.1232

Manuf.	Model	Delivery Date	Shutoff Date
Avalon	Avalon1	3/8/13	5/12/14
Bitmain	AntMiner S1	12/6/13	9/28/14
Spondoolies	SP10	3/29/14	1/13/15
Bitmain	AntMiner S2	4/3/14	1/13/15
Bitmain	AntMiner S3	7/18/14	1/14/15
Spondoolies	SP30	8/15/14	1/15/16
Bitmain	AntMiner S4	10/1/14	8/24/15
Spondoolies	SP20	11/19/14	9/21/15
Avalon	Avalon4	12/10/14	8/24/15
Bitmain	AntMiner S5	12/29/14	2/7/16
Bitmain	AntMiner S4+	4/27/15	1/16/16
Bitmain	AntMiner S7	9/23/15	8/1/16
Avalon	Avalon6	11/21/15	7/9/16
Bitmain	AntMiner S9	6/13/16	
Bitmain	AntMiner R4	9/2/16	
Avalon	721	11/23/16	
Avalon	741	1/20/17	

<https://bitcoinmining.com>

Source: Joshua Cain

Hash rate per dollar cost of ASIC chips



Source: Joshua Cain

Mining farms

Some can switch between currencies
or mine multiple currencies simultaneously



Life Inside a Secret Chinese Bitcoin Mine:

<https://www.youtube.com/watch?v=K8kua5B5K3I>



Miners: successors to accountants

“Competitive bookkeeping”

- Mining is computationally intensive, with supercomputers specially configured to look for nonces at very high “hash rates”
- Generally located in bunkers where electric power is cheap
 - Iceland
 - Inner Mongolia
 - Tibet
 - Venezuela




Icelandic Bitcoin mine

The New York Times


Hash rate of network

Trillions of hashes per second



<https://blockchain.info/charts/hash-rate>

Competition among miners creates the indelibility of data on a blockchain




- Fraud = rewriting old transactions
- Prohibitively difficult to recreate a block; nonces must be found *for all subsequent blocks* before honest miners code the next block
- Implication: transactions are *indelible*, but also *irreversible*
- Note potential conflict if two miners each discover a valid block without knowing of one another

Source: Mark Montgomery / IEEE Spectrum

Mining strategies


How to overcome huge barriers to entry due to capital cost?

- Mining pools
 - Syndicates of miners who work together and agree to share the rewards
 - Similar to lottery pools involving many people who purchase tickets and share the prizes equally
- Cloud mining
 - Using background computing power that would otherwise be idle




Hashrate shares of mining pools

as of February 24, 2019



Botnet miners



2017 data is for January – August period only

Cloud mining



DMM is an eCommerce company with \$700 million in revenue - www.dmm.com

The company's press release

September 8, 2017

■Overview of the virtual currency mining business, "DMM Mining Farm"

1.Operation of a mass-scale mining farm

"DMM Mining Farm" will operate a mass-scale, made-in-Japan quality, mining farm whose operating size will be unmatched by any of the domestic operators. In the future, DMM plans operation that ranks in top three of the world's mining farm companies in terms of scale. After completion of successful trials in October 2017, "DMM POOL" will be released for world-wide use within the year 2017. During the year 2018, DMM will be one of the 10 largest mining farms in the world.

2 Providing a safe and secure service, "DMM Cloud Mining" (currently in development)


DMM offers "DMM Cloud Mining", a safe and secure service that enables general public to easily join "DMM Mining Farm" at any time. As one of its advantages over other existing cloud mining services, which are often made overseas, "DMM Cloud Mining" will be operated by DMM, which offers superiority in its reliability and services for which DMM's distinct reputation attests.

■About virtual currency mining

Virtual currencies such as Bitcoin and Ethereum are expected to be the currencies of the next generation as they are not bound by country-specific or regional-level restrictions. There are already over 800 types of virtual currencies in existence in the world. Their market size has grown to the scale of \$160 billion. Data of these virtual currencies are maintained by being recorded and updated in a distributed ledger called Blockchain. The recording and updating work of this distributed ledger is called mining. In other words, the mining process contributes to the maintenance of the virtual currency networks being used around the world. To execute the action of this mining process consumes computing resources and electricity, for which there is a mechanism to pay for these resources via virtual currencies as a compensation. Due to its nature for having compatible legal currency and their exchangeability, the virtual currency mining market has been growing rapidly.


Cost per transaction (\$USD)

7 day moving average, block rewards plus mining fees



<https://blockchain.info/charts/cost-per-transaction>

Mining revenue / Value processed 7 day moving average



<https://blockchain.info/charts/cost-per-transaction-percent>