

11. Bitcoin, incentives and the future

Modelling Social Interaction in Information Systems

www.davidhales.com/msiis

David Hales, University of Szeged

Bitcoin

- Bitcoin system is quite complex in detail
- A whole course is devoted to it *
- Here we focus on:
 - the basic architecture (blockchain)
 - the incentive system(s)
 - future of blockchain approach (briefly)
- Hence I'm giving a *highly simplified* overview

* For detail see this very good free course, includes draft chapters of a soon to be published book: <https://www.coursera.org/course/bitcointech>

* Also see bitcoin wiki: https://en.bitcoin.it/wiki/Main_Page

Bitcoin

- Created by Satoshi Nakamoto (2009)
- Aimed at creating a secure, P2P payment system over the internet
- Without the need for banks, existing currencies or other trusted 3rd parties
- anyone should be able to join, do what they need to do, and leave (open system)
- Open protocol including incentives

Satoshi Nakamoto (2009) Bitcoin: A Peer-to-Peer Electronic Cash System.
<https://bitcoin.org/bitcoin.pdf>

Bitcoin

- In a nutshell bitcoin is a P2P decentralised protocol that supports a secure transactions ledger (blockchain) that all nodes agree on
- It does this through consensus, cryptography and incentives
- Firstly let us take a step back...
- What do we mean by a transaction ledger?

A centralised ledger

- Suppose we had a centralised trusted ledger
- stored a balance against each person in a population
- Each person could authorise to transfer some amount of their balance to some other person
- Each transaction would be verified & recorded in the ledger
- We could call this money, but:
 - everyone would have to trust the ledger
 - everyone would have to have access to the ledger
 - everyone would need verifiable and unique identities
 - initial balances would need to be decided somehow
 - why would anyone ascribe value to the balances?

Distributed ledger

- Imagine a simple P2P system where
 - Each peer stores it's own copy of the ledger
 - Broadcasts transactions it wishes to enact
 - Receives transactions from other peers
 - Updates it's ledger rejecting transactions that are invalid (i.e. don't have funds to transfer)
- Solves problem of trusting central ledger but none of the other problems
- In fact it is worse because now everyone must trust everybody
- Suppose I send two transaction at the same time to two different nodes spending same value twice (double spending problem)

Distributed ledger with digital signatures

- Digital signatures rely on public / private key crypto approach:
 - Given a private secret key and an associated public key
 - A peer can sign a message such that others can verify that it was sent by the pub. key holder
 - And that the message has not been altered
- A signature is a hash value derived from the (message, private key) by the sender
- The receiver can verify the message by applying a function to (message, signature, public key)

Aside: Hash function

- A hash function $H()$ takes some string of bits (data) of any length and produces a fixed length string of bits (hash value)
 - $H(\text{data}) \Rightarrow \text{hash value}$
- In general (with very high probability):
 - each unique data has a unique hash value
 - similar data has very different hash values
 - easy to compute forwards (find hash value from data)
 - hard to compute backwards (find data that makes a given hash value)

Distributed ledger with digital signatures

- If we ascribe identity to any public key
- Any peer can generate pub / priv. key pairs at any time
- Then we can represent transactions as digitally signed transfers between public keys
- These can be broadcast and verified by each peer in the network:
 - Check the signature
 - Check enough funds are in the source to cover the transaction to the destination by looking back at all the previously stored transactions

Distributed ledger with dig. sigs.

- At a simple level this is how bitcoin works:
 - All transactions are broadcast to all peers
 - They constitute transfers between public keys *
 - They are digitally signed by the owner of the private key associated with the source public key
 - Any peer can create pub. / priv. key pair anytime
 - Each node checks the transactions to make sure they are valid before adding them to their ledger

* Transactions are a little more complex and flexible than this in bitcoin

Aside: Bitcoin Transactions

- The structure of BTC transactions (detail):
 - Unique ID
 - A set of inputs (links to previous trans outputs)
 - A set of outputs (sent to public keys)
 - Value of inputs \geq value of outputs
 - Flexible: Script that translates how inputs mapped to outputs (a simple language – no looping)
 - Flexible: can implement multiple signatures etc.
 - Efficient: inputs of the transaction are specified as outputs from previous transactions (secure linked data structure using hash pointers)

Smallest subunit of bitcoin is 100,000,000th and called a Satoshi

Distributed ledger with dig. sigs.

- Several problems with this approach:
 - (**creation**) Where do initial balances come from associated with public keys?
 - (**consensus**) How do we ensure all the peers have the same ledger (agree on ordering of transactions):
 - What happens if a peer sends two incompatible transactions quickly to the network such that two peers incorporate them into their ledger (double spend)?
 - (**incentives**) what incentive do peers have to run the system (correctly) in the first place or ascribe value to balances?

Blockchain

- Bitcoin elegantly solves these problems using a data structure called the blockchain
 - linked list (chain) of blocks (groups) of transactions
- New blocks are added to the end of the blockchain
- Each block points to previous block using a “hash pointer”
 - A hash pointer is a unique hash value derived from the contents of the thing being pointed to
 - This means that neither the pointer NOR the contents of what it points to can be changed without detection
 - Hence hash pointers are used to maintain the integrity of data structures

Adding a block to the blockchain

- Any node can work on creating a new block
 - collect new transactions received
 - compose into a candidate block
 - solve a special problem (see next)
 - broadcast block to other nodes so they can add them to their blockchain
- Nodes will reject any block that:
 - contains invalid transactions
 - does not point to previous valid block at end of chain
 - does not contain the solution to the *special problem*

Note: you can view blocks and transactions in real-time at: <https://blockexplorer.com>

Blockchain – proof of work


- Each valid block must contain the correct solution to a special problem
- To solve this problem requires significant brute force computation (it is hard)
- Find a number (n) that when combined with block data (BD) produces a hash value (using a known hash function) < some threshold (T)
 - $\text{Hash}(\text{BD} \mid n) < T$

Blockchain – proof of work

- Since each different n produces a completely different hash value:
 - only way to find suitable n is to search by repeatedly computing hashes for different n
 - The lower the value of T the harder this search becomes
- T adapts overtime such that on average one new block is discovered every 10mins

Since n is only used once this is called a “nonce” in crypto terminology

Blockchain



Tx = transaction

Note: transactions are actually stored in a “Merkle Tree” data structure which is a binary tree of hash values over the transaction data. It provides efficiency and integrity properties.

From: <https://bitcoin.org/bitcoin.pdf>

Blockchain - incentives

- The blockchain provides a way to agree on the order of transactions (**consensus**, see later) but..
- Why would a peer do all this work to create a new block?
- Two incentive systems (**incentives**):
 - Block reward: new bitcoins from “thin air” awarded to the creator of the block
 - Transaction fees: those producing transactions can add an optional small fee which can be kept by the creator of the block

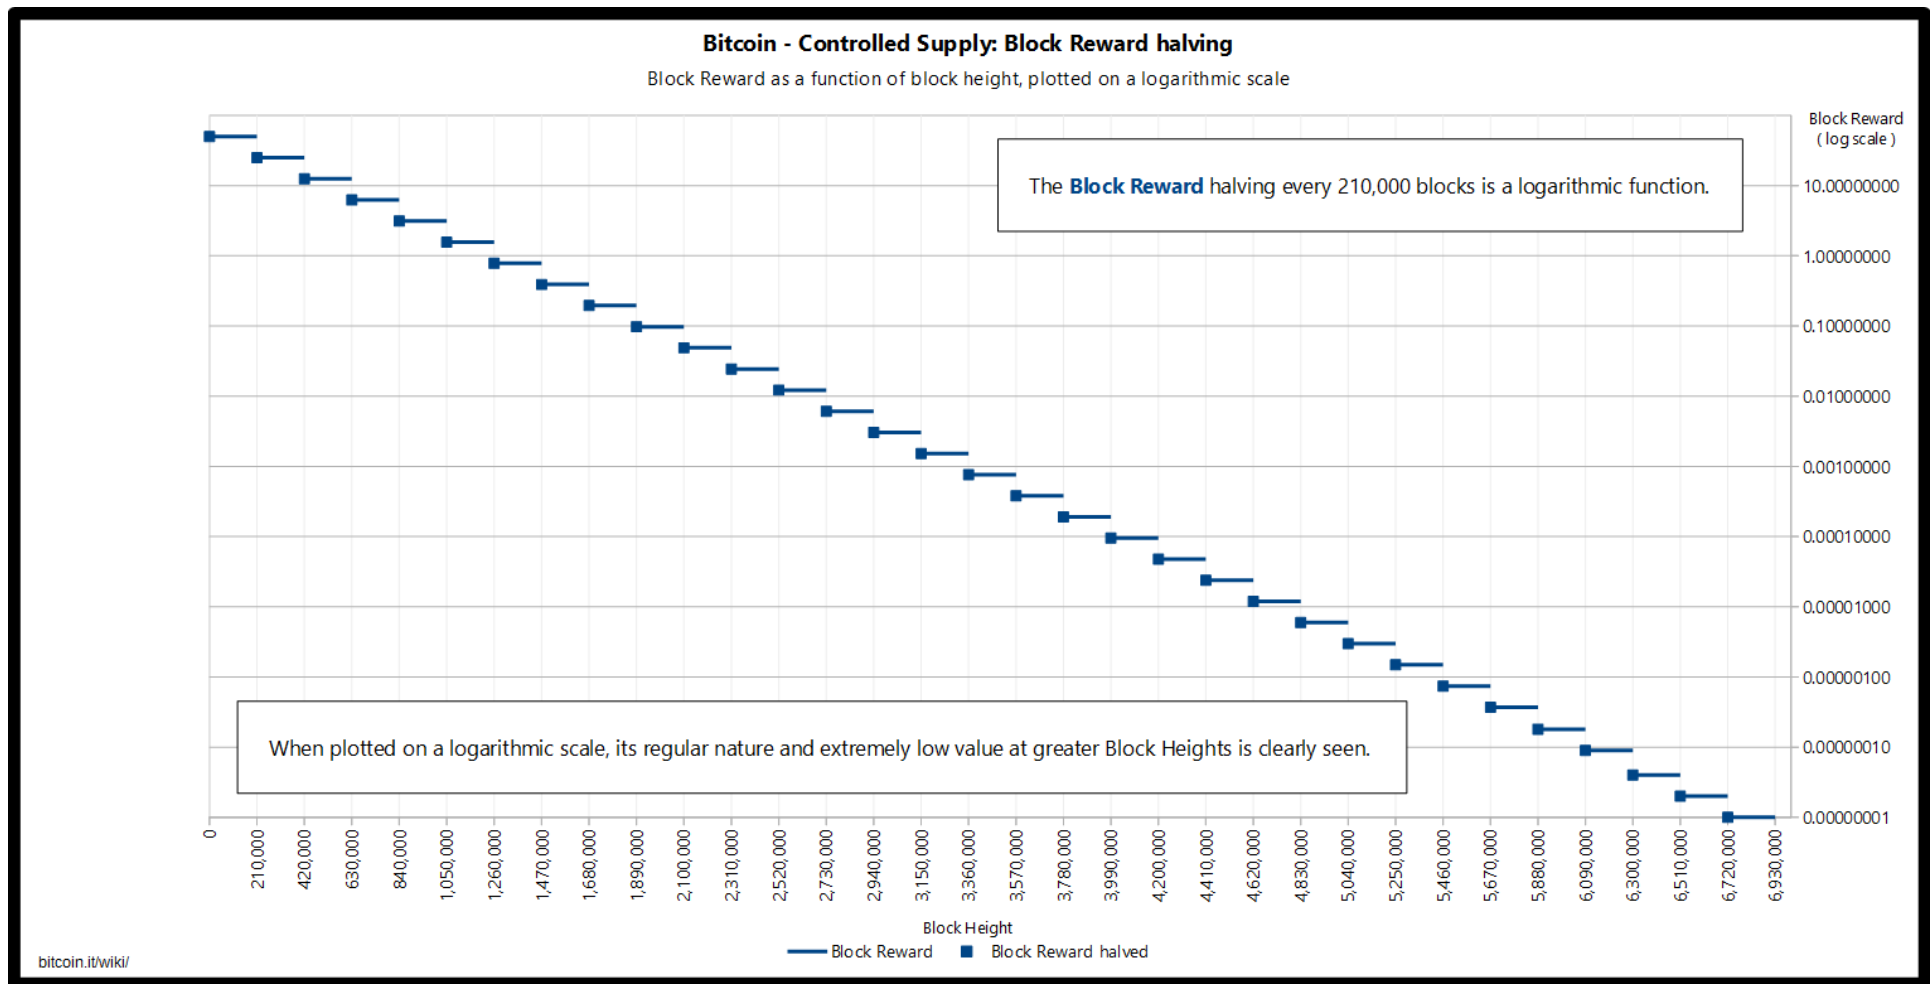
Blockchain – new coins

- Each new block contains a special transaction: a block reward (or coinbase) transaction:
 - A new bitcoin balance (say 25) is created and allocated to a public key (chosen by the block creator)
 - Hence it is a transaction without a source only a destination key (bitcoin **creation**)
 - The block reward is adapted system wide (halving every 210,000 blocks)

Blockchain - mining


- Hence new bitcoins are created and allocated when new blocks are created (**creation**)
- Attempting to create new blocks is called “mining”:
 - Solving the special problem is hard work
 - But you get bitcoins if successful
 - It is also risky (you can only do a random search)
- miners are in competition because if someone else creates a block before you then you have to start over
- Remember that block data contains a hash pointer to the previous block at the end of the blockchain

Block reward schedule



From: https://en.bitcoin.it/wiki/Controlled_supply

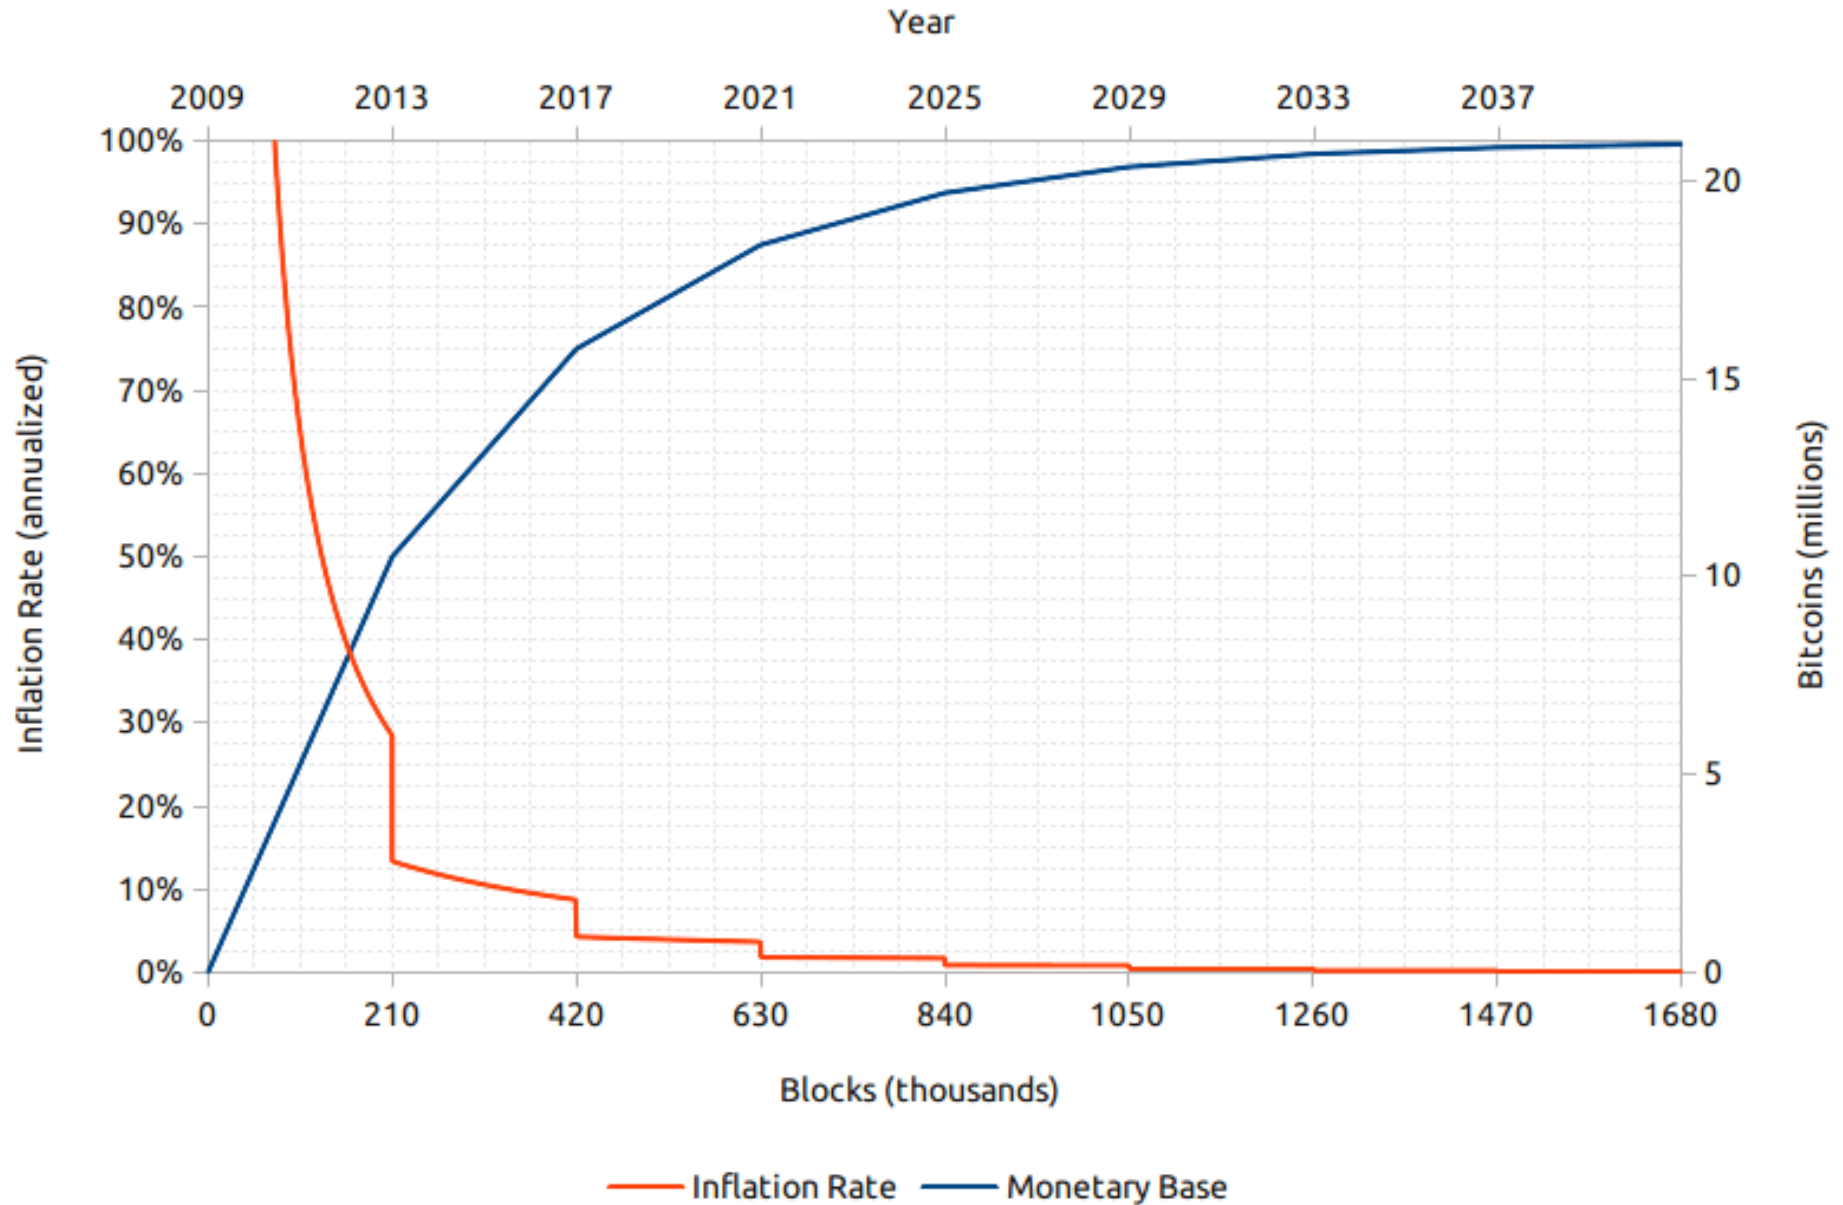
BTC increase



From: https://en.bitcoin.it/wiki/Controlled_supply


Bitcoin Inflation vs. Time

From: <https://bitcointalk.org/index.php?topic=130619.0>



Hash Rate

Source: blockchain.info



From: blockchain.info

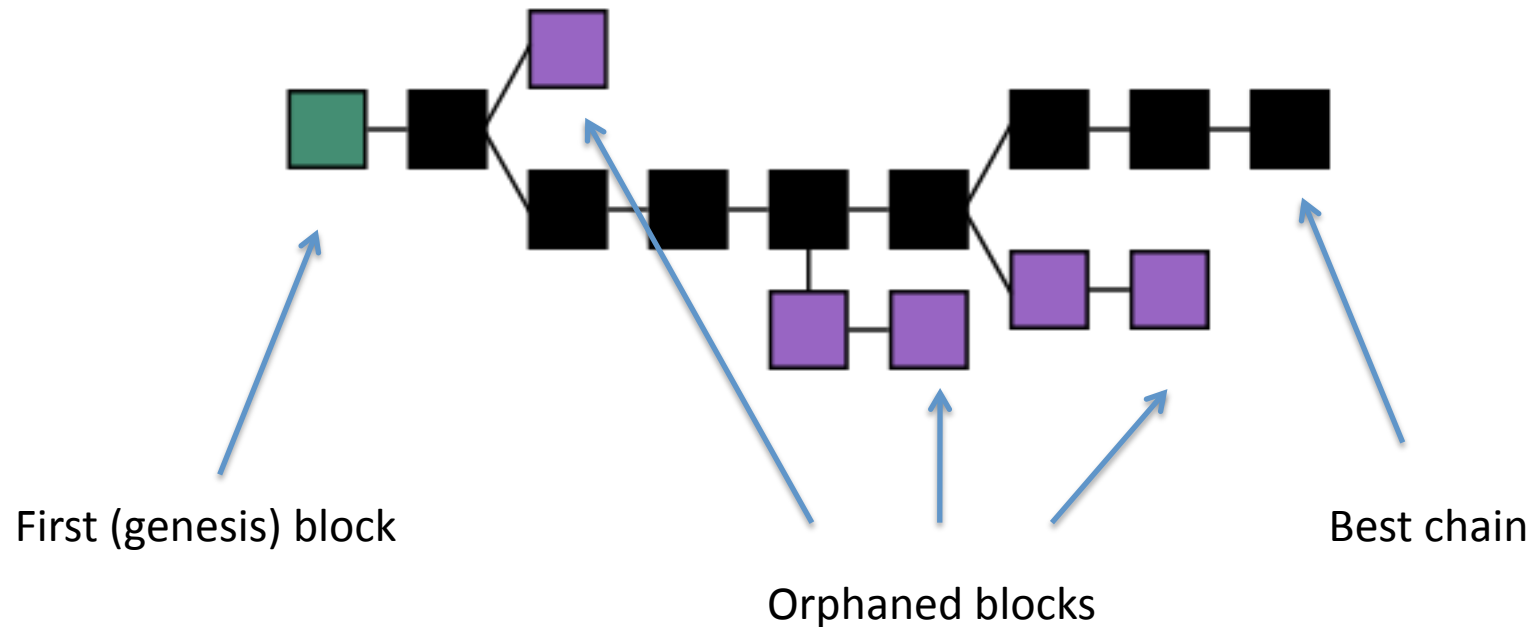
Blockchain - consensus

- What happens if two new valid blocks (A & B) are created at (close to) the same time
 - Both are valid
 - Some nodes will get A first others B first
 - This creates a fork in the blockchain where two alternative paths exist
 - This means two alternative transaction histories exist which is highly undesirable

Blockchain - consensus

- Nodes store competing blocks but eventually discard those that end-up in the shorter chain
 - Transactions in shorter chain are lost (orphan block)
 - Does not occur so often due to low rate of new blocks (~10 mins) and randomness of solving special problem
- Hence transactions are not considered fully “validated” until they are some depth into the blockchain (say 7)
- Consensus is achieved through computational effort (longest chain)
- Hence **incentive** for miners to work on longest current chain (also block reward locked until ~100 depth)

Blockchain orphan blocks



Aside: An orphan chain could get quite long. Longest known is about 50 blocks due to a software error in a new release causing a potential “permanent fork in the block chain” until it was fixed by developers and cooperation of mining pools (see later)

Picture from: https://en.wikipedia.org/wiki/Bitcoin_network

Incentives (infrastructure)

- Process of creation and consensus support, and are supported by, incentives:
 - Creating new blocks (block reward)
 - Including transactions in blocks (trans. fees)
 - Building a longest consensus chain (keeping block reward and trans. fees)
- Hence assumption miners will provide trusted infrastructure due to incentives of obtaining bitcoin

Incentives (long-term)

- The staged rewards and increase in bitcoins overtime
 - Early adopters got lots of BTC with little effort
 - There is a limit to how many BTC can be produced
 - The value of BTC depends on more users
 - Hence social incentive for early adopters to continue to use and spread usage (create publicity)
 - Even when all BTC have been released (~21m) trans. fees will provide incentives to provide infrastructure

Incentives and value

- Bitcoins can be traded on exchanges for other fiat currencies (euro, dollar etc)
- Why do they have value? (various theories):
 - Similar to our discussion on other currencies
 - Useful for pseudonymous money transfer
 - traders accepting them
 - Cypherpunk Ideology and Publicity
 - Speculation
 - Liquid markets that trade for fiat currencies

Incentives and value

- Cycle of incentives:
 - More hash power => more secure (true?)
 - More secure => more users
 - More users => more value (more publicity)
 - More value => more hash power (incentive)
- However, diminishing / risky returns:
 - More hash power => more expensive to mine
 - Equipment + electricity
 - Speculation => price volatility

Other social incentives

- Bitcoin emerged from the Cypherpunk community who generally believe in:
 - anonymity
 - Freedom from state / corporate interference
 - Free association, value creation
 - The use of cryptography to achieve this
 - Decentralisation, eliminating trusted 3rd parties
 - Open source transparency / free software
- Bitcoin adheres to these principles

One could think of this as a kind of cypherpunk “social contract”

Limitations of Bitcoin


- Centralisation (how P2P is it really?)
- Scalability (can it really scale up?)
- Waste of energy (green objection)
- Privacy (transactions are public)

For a brief and clear discussion on the limitations of bitcoin design see post by Victor Grishchenko, 12 May 2011, "Bitcoin?"
<http://www.pds.ewi.tudelft.nl/~victor/bitcoin.html>

Centralisation (mining)

- Some concern about a form of centralisation in the system:
 - Competitive nature of mining means ever more powerful dedicated systems used
 - Standard computer can no longer compete
 - Mining centralised into “mining pools” shared resources / risk
 - Miners decide which transactions get put into blocks

Hash rate (%) over mining pools (Jun'15)



From: <http://cryptomining-blog.com/tag/bitcoin-mining-pools/>

Aside: Mining pool incentives

- Mining pools and individuals within pools have complex incentives
 - “block withholding attacks”
 - One pool can attack another through infiltration
- Has been formulated into a game
 - The miner’s dilemma (a Prisoner's Dilemma)
 - Predicts potential breakdown of large open pools

Ittay Eyal (2015) Bitcoin – The Miner’s Dilemma. SWIFT institute working paper no. 2014-006.

See discussion of this by Ittay Eyal:

<http://hackingdistributed.com/2014/12/03/the-miners-dilemma/>

Centralisation (thin clients)

- Many bitcoin clients are not full nodes and do not store the entire blockchain or check transitions
 - Dependent on other nodes (miners or others storing whole blockchain)
 - Power becomes focused into those nodes
 - “web wallets” and other 3rd party solutions no better (or worse) than using standard banks
- One way to view this: bitcoin starts to resemble an alternative “base money” mediated not by banks but by “bitcoin service providers”

Centralisation (developers)

- Open source project
- Core set of developers (unclear governance?)
- Responsible for bug fixes and updates
- But new versions must be accepted / downloaded by majority of (CPU of) network
- Major players (mining pools, payment operators)
- Recent disagreements on increasing block size above 1MB (see next slide) led to a controversial fork from the core code (Bitcoin-XT)

See: <https://medium.com/faith-and-future/why-is-bitcoin-forking-d647312d22c1>

Centralisation (Exchanges)

- Currently to buy / sell bitcoin (in fiat e.g. \$)
 - Need to use centralised exchanges (matching engines)
 - Or find people and physically meet them
 - In both cases require trust with those entities
 - No current fully decentralised P2P matching engine in operation
 - A early popular exchange MTGOX went bankrupt nobody seems sure where the bitcoin went

Bitcoin scalability

- Full nodes use efficient ways to store and retrieve blockchain
- However, it does continue to grow
- Also current parameters mean < 10 transactions processed per second (on ave. globally)
- Could be solved by future revisions to the core code assuming mining pools choose to run it
- But if not then transaction prices could rise to stop congestion (i.e. miners could throw away tx with low fees)

See: Gavin Andresen: On The Blocksize And Bitcoin's Governance
<https://youtu.be/B8l11q9hsJM>

Waste of energy?



Picture from: https://en.wikipedia.org/wiki/Bitcoin_network










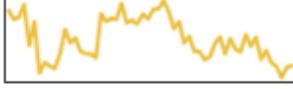






Other blockchain systems

- Altcoins
- Smart contracts (developing / future)
- What is emerging?

Altcoins

- There are many variants of coins based on a BTC like blockchain called altcoins
 - In general each is self-contained (but not all)
 - simplest just clone BTC code but change some parameters (eg. Block rate, total number of coin)
 - Some use other techniques than proof-of-work for mining
 - Some targeted to specific application (namecoin)
 - Many tradable on centralised market exchanges
 - In some sense these compete with each other for users / comp. power / publicity

Various altcoins (nov.'15)

All ▾	Currencies ▾	Assets ▾	USD ▾	Next 100 →	View All		
▲ #	Name	Market Cap	Price	Available Supply	Volume (24h)	% Change (24h)	Price Graph (7d)
1	 Bitcoin	\$ 4,782,957,374	\$ 322.59	14,826,600 BTC	\$ 135,940,000	-13.72 %	
2	 Ripple	\$ 143,141,334	\$ 0.004317	33,156,211,683 XRP *	\$ 618,601	-5.68 %	
3	 Litecoin	\$ 129,569,036	\$ 3.00	43,120,685 LTC	\$ 5,196,080	-7.83 %	
4	 Ethereum	\$ 66,614,868	\$ 0.892607	74,629,560 ETH	\$ 916,381	-11.96 %	
5	 Dash	\$ 14,055,263	\$ 2.35	5,980,429 DASH	\$ 89,502	-9.11 %	
6	 Dogecoin	\$ 12,800,031	\$ 0.000126	101,771,707,292 DOGE	\$ 177,363	-4.47 %	
7	 Stellar	\$ 9,921,273	\$ 0.002051	4,837,356,606 STR *	\$ 7,979	-4.41 %	
8	 Peercoin	\$ 8,487,420	\$ 0.373065	22,750,513 PPC	\$ 62,195	-11.91 %	

From: <http://coinmarketcap.com/>

More general blockchain

- Transaction scripts extended to full computer language:
 - Transactions represent contracts
 - Public & computationally verifiable
 - Could code rules of institutions such as
 - Voting rights
 - Ownership transfers of real objects
 - called “Smart Contracts”, “smart property”

Ethereum white paper: <https://github.com/ethereum/wiki/wiki/White-Paper>

See: <https://en.wikipedia.org/wiki/Ethereum>

What is emerging?

- A new programming paradigm in which:
 - Data, code and incentives are packaged together
 - Publically maintained in a distributed way
 - Cryptographically secured
 - Self-funding through value creation mechanisms
- Applying ideas from:
 - Comp. Sci. & Crypto
 - Economics, game theory, social interaction
 - law, political philosophy

What is emerging?

- The Block chain could be as significant as the invention of double entry bookkeeping
- Or the joint stock company
- Or maybe it's all a fad and will die off!
- Or live on in a “hackers dream” niche?
- Next 5 years will be interesting

Readings and Questions

- Readings:
 - Chap. 1,2 & 3 from draft textbook at: <https://www.coursera.org/course/bitcointech>
 - Nick Szabo's work: <http://szabo.best.vwh.net>
- Questions:
 - In the first ever (genesis) block Nakamoto included that day's date and the headline from the Times of London – why might he have done this?
 - If one person controlled > 50% of the bitcoin hash power could they destroy the value of bitcoin? would they?
 - Can you think of other uses for a blockchain-like data structure?