Lecture #21: Nakamoto Consensus: Difficulty Adjustment, Block Rewards, and Selfish Mining

> COMS 4995-001: The Science of Blockchains URL: https://timroughgarden.org/s25/

> > Tim Roughgarden

Goals for Lecture #21

- 1. Difficulty adjustment in Nakamoto consensus.
 - tuning the difficulty threshold to achieve a target block rate
- 2. Limitations of proof-of-work.
 - drawbacks of Nakamoto consensus hold for PoW protocols generally

3. Why cryptocurrencies?

- among other reasons, bootstrap a PoW blockchain protocol
- 4. Block rewards and selfish mining.
 - incentivizes block production; does it incentive other behavior even more?

Recall: Nakamoto Consensus

Nakamoto consensus: longest-chain consensus + PoW leader selection.

- selection probability proportional to hashrate
- random length of a view = time for someone to solve puzzle



Recall: Nakamoto Consensus

Nakamoto consensus: longest-chain consensus + PoW leader selection.

- selection probability proportional to hashrate
- random length of a view = time for someone to solve puzzle

Puzzle format: find x with $h(x) \le t$ where:

h = cryptographic hash fn, t = difficulty parameter (both protocol-defined)

B₀

- x = block header of the form < tx Merkle root II pred II pk II nonce >
 - point of the nonce: "grind" through possibilities until find a solution

C = longest chain

В

Canonical PoW puzzle: for a threshold t, find a valid x with $h(x) \le t$.

Question: how to set t? [controls puzzle difficulty]

Canonical PoW puzzle: for a threshold t, find a valid x with $h(x) \le t$.

Question: how to set t? [controls puzzle difficulty]

Answer: to achieve a target rate of block production.

Canonical PoW puzzle: for a threshold t, find a valid x with $h(x) \le t$.

Question: how to set t? [controls puzzle difficulty]

Answer: to achieve a target rate of block production.

• to balance the pros of faster block production (lower latency)

Canonical PoW puzzle: for a threshold t, find a valid x with $h(x) \le t$.

Question: how to set t? [controls puzzle difficulty]

Answer: to achieve a target rate of block production.

• to balance the pros of faster block production (lower latency) with the cons (more wasted work due to inadvertent forks)



Canonical PoW puzzle: for a threshold t, find a valid x with $h(x) \le t$.

Question: how to set t? [controls puzzle difficulty]

Answer: to achieve a target rate of block production.

• to balance the pros of faster block production (lower latency) with the cons (more wasted work due to inadvertent forks)

Consequence: threshold most decrease (resp., increase) as total amount of hashrate increases (resp., decreases).

- difficulty adjustment algorithm programmatically makes these updates

• target rate = 1 block/10 minutes (\rightarrow 144 blocks/day)

- target rate = 1 block/10 minutes (\rightarrow 144 blocks/day)
- update t every epoch := 2016 blocks

- target rate = 1 block/10 minutes (\rightarrow 144 blocks/day)
- update t every epoch := 2016 blocks
- if elapsed time in epoch = $\beta \cdot (14 \text{ days})$, reset t := $\beta \cdot t$
 - if next epoch like previous one, expect to see 1 block/10 minutes

- target rate = 1 block/10 minutes (\rightarrow 144 blocks/day)
- update t every epoch := 2016 blocks
- if elapsed time in epoch = $\beta \cdot (14 \text{ days})$, reset t := $\beta \cdot t$
 - if next epoch like previous one, expect to see 1 block/10 minutes

Question: how is time measured?

- target rate = 1 block/10 minutes (\rightarrow 144 blocks/day)
- update t every epoch := 2016 blocks
- if elapsed time in epoch = $\beta \cdot (14 \text{ days})$, reset t := $\beta \cdot t$
 - if next epoch like previous one, expect to see 1 block/10 minutes

Question: how is time measured?

Answer: timestamps.

recorded in block headers



- target rate = 1 block/10 minutes (\rightarrow 144 blocks/day)
- update t every epoch := 2016 blocks
- if elapsed time in epoch = $\beta \cdot (14 \text{ days})$, reset t := $\beta \cdot t$
 - if next epoch like previous one, expect to see 1 block/10 minutes

Question: how is time measured?

Answer: timestamps.

- recorded in block headers
 - rules to limit timestamp manipulation by Byzantine validators (see HW7)



- target rate = 1 block/10 minutes (\rightarrow 144 blocks/day)
- update t every epoch := 2016 blocks
- if elapsed time in epoch = $\beta \cdot (14 \text{ days})$, reset t := $\beta \cdot t$
 - if next epoch like previous one, expect to see 1 block/10 minutes
- time measured via timestamps in block headers

Also: need to redefine Nakamoto consensus so that (honest) leaders extend the *heaviest* chain (rather than the longest chain).

- weight of block with threshold t := 2^{256} /t [expected # of attempts to obtain]
- weight of chain = sum of weights of blocks in the chain

Recall: Guarantees for Nakamoto Consensus

Assumptions: (all necessary)

- 1. Synchronous network (i.e., max message delay $\leq \Delta$).
- 2. < 50% Byzantine *hashrate at all times*.
- 3. k large enough that, in every interval of $\ge 2k+2$ views, a strict majority of the leaders are honest with high probability.
- 4. difficulty threshold t small enough that avg view length $>> \Delta$

Guarantee: under these assumptions, Nakamoto consensus is consistent and live *with high probability*.

Drawbacks of Nakamoto consensus:

Drawbacks of Nakamoto consensus:

· loses consistency in the partially synchronous setting

- true already for permissioned version



Drawbacks of Nakamoto consensus:

- loses consistency in the partially synchronous setting
 - true already for permissioned version
- even in synchrony, consistency + liveness only probabilistic
 - wasn't a problem is permissioned case (deterministic round-robin)

Drawbacks of Nakamoto consensus:

- · loses consistency in the partially synchronous setting
 - true already for permissioned version
- even in synchrony, consistency + liveness only probabilistic
 - wasn't a problem is permissioned case (deterministic round-robin)

Questions:

- tweak Nakamoto consensus so that one/both problems fixed?
- combine PoW with Tendermint rather than longest-chain?

Facts: [Lewis-Pye/Roughgarden, 2020-3]

Facts: [Lewis-Pye/Roughgarden, 2020-3]

1. no PoW protocol is consistent in partial synchrony

Facts: [Lewis-Pye/Roughgarden, 2020-3]

- 1. no PoW protocol is consistent in partial synchrony
 - assuming protocol is live in synchrony with no Byzantine validators
 - rules out combining PoW with Tendermint in any straightforward way

Facts: [Lewis-Pye/Roughgarden, 2020-3]

- 1. no PoW protocol is consistent in partial synchrony
 - assuming protocol is live in synchrony with no Byzantine validators
 - rules out combining PoW with Tendermint in any straightforward way
- even in synchrony, no POW protocol always guarantees (deterministic) consistency + liveness

Facts: [Lewis-Pye/Roughgarden, 2020-3]

- 1. no PoW protocol is consistent in partial synchrony
 - assuming protocol is live in synchrony with no Byzantine validators
 - rules out combining PoW with Tendermint in any straightforward way
- even in synchrony, no POW protocol always guarantees (deterministic) consistency + liveness

Upshot: drawbacks of Nakamoto consensus fundamental to all PoW protocols.

- can be overcome (under extra assumptions) with proof-of-stake protocols

- 1. no PoW protocol is consistent in partial synchrony
- 2. even in synchrony, no POW protocol guarantees (deterministic) consistency + liveness

Intuition for (1): catch-22 a la "CAP principle" argument.

- 1. no PoW protocol is consistent in partial synchrony
- 2. even in synchrony, no POW protocol guarantees (deterministic) consistency + liveness

Intuition for (1): catch-22 a la "CAP principle" argument. If validator hears no messages for a long time, can't distinguish between:

• (i) in synchrony, other validators turned off their machines

- 1. no PoW protocol is consistent in partial synchrony
- 2. even in synchrony, no POW protocol guarantees (deterministic) consistency + liveness

Intuition for (1): catch-22 a la "CAP principle" argument. If validator hears no messages for a long time, can't distinguish between:

- (i) in synchrony, other validators turned off their machines
- (ii) in partial synchrony + pre-GST, all messages delayed

- 1. no PoW protocol is consistent in partial synchrony
- 2. even in synchrony, no POW protocol guarantees (deterministic) consistency + liveness

Intuition for (1): catch-22 a la "CAP principle" argument. If validator hears no messages for a long time, can't distinguish between:

- (i) in synchrony, other validators turned off their machines
- (ii) in partial synchrony + pre-GST, all messages delayed

Should the validator ever finalize any additional txs?

- 1. no PoW protocol is consistent in partial synchrony
- 2. even in synchrony, no POW protocol guarantees (deterministic) consistency + liveness

Intuition for (1): catch-22 a la "CAP principle" argument. If validator hears no messages for a long time, can't distinguish between:

- (i) in synchrony, other validators turned off their machines
- (ii) in partial synchrony + pre-GST, all messages delayed

Should the validator ever finalize any additional txs?

yes → might be in scenario (ii), cause a consistency violation

- no PoW protocol is consistent in partial synchrony 1.
- even in synchrony, no POW protocol guarantees (deterministic) consistency + liveness 2.

Intuition for (1): catch-22 a la "CAP principle" argument. If validator hears no messages for a long time, can't distinguish between:

- (i) in synchrony, other validators turned off their machines
- (ii) in partial synchrony + pre-GST, all messages delayed

Should the validator ever finalize any additional txs?

- yes \rightarrow might be in scenario (ii), cause a consistency violation
- no \rightarrow might be in scenario (i), liveness violation (in synchrony) $_{32}$

- 1. no PoW protocol is consistent in partial synchrony
- 2. even in synchrony, no POW protocol guarantees (deterministic) consistency + liveness

Intuition for (1): catch-22 a la "CAP principle" argument. If validator hears no messages for a long time, can't distinguish between:

- (i) in synchrony, other validators turned off their machines
- (ii) in partial synchrony + pre-GST, all messages delayed

Should the validator ever finalize any additional txs?

- yes \rightarrow might be in scenario (ii), cause a consistency violation
- no \rightarrow might be in scenario (i), liveness violation (in synchrony)

Proof of (2): similar to proof of FLP Impossibility Theorem.

– churning validators can substitute for unbounded message delays
33

Incentivizing Validators

Question: why run a validator (e.g., for a PoW protocol)?

Why Cryptocurrencies?

Cryptocurrecy: currency native to a blockchain protocol (i.e., minted/burned/tracked by the protocol).

Why Cryptocurrencies?

Cryptocurrecy: currency native to a blockchain protocol (i.e., minted/burned/tracked by the protocol). Uses (incomplete list):

1. interesting in its own right (e.g., Bitcoin)
- 1. interesting in its own right (e.g., Bitcoin)
- 2. charge for usage (i.e., transaction fees)

- 1. interesting in its own right (e.g., Bitcoin)
- 2. charge for usage (i.e., transaction fees)
- 3. rewards contributions to protocol (e.g., PoW validators)

- 1. interesting in its own right (e.g., Bitcoin)
- 2. charge for usage (i.e., transaction fees)
- 3. rewards contributions to protocol (e.g., PoW validators)
- 4. proof-of-stake sybil-resistance

- 1. interesting in its own right (e.g., Bitcoin)
- 2. charge for usage (i.e., transaction fees)
- 3. rewards contributions to protocol (e.g., PoW validators)
- 4. proof-of-stake sybil-resistance
- 5. punish protocol deviators (a.k.a. "slashing")

Cryptocurrecy: currency native to a blockchain protocol (i.e., minted/burned/tracked by the protocol). Uses (incomplete list):

- 1. interesting in its own right (e.g., Bitcoin)
- 2. charge for usage (i.e., transaction fees)
- 3. rewards contributions to protocol (e.g., PoW validators)
- 4. proof-of-stake sybil-resistance
- 5. punish protocol deviators (a.k.a. "slashing")

Questions: (i) how does cryptocurrency get distributed initially?

Cryptocurrecy: currency native to a blockchain protocol (i.e., minted/burned/tracked by the protocol). Uses (incomplete list):

- 1. interesting in its own right (e.g., Bitcoin)
- 2. charge for usage (i.e., transaction fees)
- 3. rewards contributions to protocol (e.g., PoW validators)
- 4. proof-of-stake sybil-resistance
- 5. punish protocol deviators (a.k.a. "slashing")

Questions: (i) how does cryptocurrency get distributed initially? (ii) why a validator be "honest" (vs. profit-maximizing)?

Question: why run a validator (e.g., for a PoW protocol)?

Question: why run a validator (e.g., for a PoW protocol)?

Answer in Nakamoto consensus: block rewards in native currency.

- newly minted coins paid to the "miner" of each block on longest chain
- in Bitcoin: currently 3.125 BTC per block

Question: why run a validator (e.g., for a PoW protocol)?

Answer in Nakamoto consensus: block rewards in native currency.

- newly minted coins paid to the "miner" of each block on longest chain
- in Bitcoin: currently 3.125 BTC per block

Worry: does incentivizing block production incentivize any deviations from Nakamoto consensus?

- hope: validators maximize rewards by following the protocol

Question: why run a validator (e.g., for a PoW protocol)?

Answer in Nakamoto consensus: block rewards in native currency.

- newly minted coins paid to the "miner" of each block on longest chain
- in Bitcoin: currently 3.125 BTC per block

Worry: does incentivizing block production incentivize any deviations from Nakamoto consensus?

- hope: validators maximize rewards by following the protocol
- paid for getting blocks on the longest chain, not for being honest per se
- e.g., could validators have their blocks orphaned at different rates?

Recall: difficulty adjustment \rightarrow fixed average rate of growth of longest chain (e.g., 2016 new blocks on longest chain per 14 days).

- says nothing about the number of orphaned blocks during this time

Recall: difficulty adjustment \rightarrow fixed average rate of growth of longest chain (e.g., 2016 new blocks on longest chain per 14 days).

- says nothing about the number of orphaned blocks during this time
- fixed rate of block rewards (e.g., 6300 BTC per 14 days)
 - \rightarrow to maximize rewards, maximize share of blocks on longest chain

Recall: difficulty adjustment \rightarrow fixed average rate of growth of longest chain (e.g., 2016 new blocks on longest chain per 14 days).

- says nothing about the number of orphaned blocks during this time
- fixed rate of block rewards (e.g., 6300 BTC per 14 days)
 - \rightarrow to maximize rewards, maximize share of blocks on longest chain
- all validators honest → share rewards proportional to hashrate

Recall: difficulty adjustment \rightarrow fixed average rate of growth of longest chain (e.g., 2016 new blocks on longest chain per 14 days).

- says nothing about the number of orphaned blocks during this time
- fixed rate of block rewards (e.g., 6300 BTC per 14 days)
 - \rightarrow to maximize rewards, maximize share of blocks on longest chain
- all validators honest \rightarrow share rewards proportional to hashrate

Question: can a validator with an α fraction of the hash rate get > α fraction of overall rewards by deviation from the protocol?

Recall: difficulty adjustment \rightarrow fixed average rate of growth of longest chain (e.g., 2016 new blocks on longest chain per 14 days).

- says nothing about the number of orphaned blocks during this time
- fixed rate of block rewards (e.g., 6300 BTC per 14 days)
 - \rightarrow to maximize rewards, maximize share of blocks on longest chain
- all validators honest \rightarrow share rewards proportional to hashrate

Question: can a validator with an α fraction of the hash rate get > α fraction of overall rewards by deviation from the protocol?

• in general, yes! [Eyal/Sirer 14]

A Profitable Deviation

Setup: adversary controls $\alpha < \frac{1}{2}$ of the overall hashrate.

- other 1α fraction obediently follows the protocol
- synchronous model with $\Delta=0$ (all msgs delivered instantly)
- assume ties broken in favor of adversary (can relax w/more work)

- goal (i): get as many "A-blocks" on longest chain as possible (ideally, all of them)
- goal (ii): orphan as many "H-blocks" as possible, to maximize share of A-blocks on the longest chain























B

Profitable deviation from Nakamoto consensus:

- let h = max height of any block produced thus far
- case 1: if there is an A-block at height h,
 ^Y^{B'3}
 try to extend it [successful → delay announcement]
- case 2: if only an H-block at height h, try to orphan it [successful → announce immediately]

delay announcement

orphaned (break ties in favor of adversary)

 B_0

Profitable deviation from Nakamoto consensus:

- let h = max height of any block produced thus far
- case 1: if there is an A-block at height h,
 try to extend it [successful → delay announcement]
- case 2: if only an H-block at height h, try to orphan it [successful → announce immediately]
- throughout: announce an A-block only once (+ immediately after) there in an H-block at the same height

delay announcement

orphaned (break ties in favor of adversarv)

 B_0

Profitable deviation from Nakamoto consensus:

- let h = max height of any block produced thus far
- case 1: if there is an A-block at height h,
 try to extend it [successful → delay announcement]
- case 2: if only an H-block at height h, try to orphan it [successful → announce immediately]
- throughout: announce an A-block only once (+ immediately after) there in an H-block at the same height
 - → every A-block deployed to knock an H-block off of the longest chain

delay announcement

orphaned (break ties in favor of adversarv)

Consider sequence of N rounds (e.g., N in the 1000s)

1. expect $\approx \alpha N$ A-blocks, $\approx (1 - \alpha) N$ H-blocks to be produced

- 1. expect $\approx \alpha N$ A-blocks, $\approx (1 \alpha) N$ H-blocks to be produced
- 2. every A-block gets on the longest chain

- 1. expect $\approx \alpha N$ A-blocks, $\approx (1 \alpha) N$ H-blocks to be produced
- 2. every A-block gets on the longest chain
- 3. every A-block orphans a distinct H-block

- 1. expect $\approx \alpha N$ A-blocks, $\approx (1 \alpha) N$ H-blocks to be produced
- 2. every A-block gets on the longest chain
- 3. every A-block orphans a distinct H-block
- \Rightarrow # of H-blocks on longest chain = $\approx (1 \alpha)N \alpha N = (1 2\alpha)N$
Consider sequence of N rounds (e.g., N in the 1000s)

- 1. expect $\approx \alpha N$ A-blocks, $\approx (1 \alpha) N$ H-blocks to be produced
- 2. every A-block gets on the longest chain
- 3. every A-block orphans a distinct H-block
- \Rightarrow # of H-blocks on longest chain = $\approx (1 \alpha)N \alpha N = (1 2\alpha)N$
- # of A-blocks on longest chain $\approx \alpha N$

Consider sequence of N rounds (e.g., N in the 1000s)

- 1. expect $\approx \alpha N$ A-blocks, $\approx (1 \alpha) N$ H-blocks to be produced
- 2. every A-block gets on the longest chain
- 3. every A-block orphans a distinct H-block
- \Rightarrow # of H-blocks on longest chain = $\approx (1 \alpha)N \alpha N = (1 2\alpha)N$
- # of A-blocks on longest chain $\approx \alpha N$
- \Rightarrow A's share of blocks n longest chain $\approx \alpha N/[\alpha N + (1 2\alpha)N]$

Consider sequence of N rounds (e.g., N in the 1000s)

- 1. expect $\approx \alpha N$ A-blocks, $\approx (1 \alpha) N$ H-blocks to be produced
- 2. every A-block gets on the longest chain
- 3. every A-block orphans a distinct H-block
- \Rightarrow # of H-blocks on longest chain = $\approx (1 \alpha)N \alpha N = (1 2\alpha)N$
- # of A-blocks on longest chain $\approx \alpha N$
- A's share of blocks n longest chain $\approx \alpha N/[\alpha N + (1 2\alpha)N]$ = $\alpha/(1 - \alpha) > \alpha$

Consider sequence of N rounds (e.g., N in the 1000s)

- 1. expect $\approx \alpha N$ A-blocks, $\approx (1 \alpha) N$ H-blocks to be produced
- 2. every A-block gets on the longest chain
- 3. every A-block orphans a distinct H-block
- \rightarrow # of H-blocks on longest chain = $\approx (1 \alpha)N \alpha N = (1 2\alpha)N$
- # of A-blocks on longest chain $\approx \alpha N$
- \rightarrow A's share of blocks n longest chain $\approx \alpha N/[\alpha N + (1 2\alpha)N]$ $= \alpha/(1-\alpha) > \alpha$

Upshot: can boost rewards by deviating from intended behavior! 76