

# Ethereum Blocks

Saravanan Vijayakumaran  
sarva@ee.iitb.ac.in

Department of Electrical Engineering  
Indian Institute of Technology Bombay

August 31, 2018

# Ethereum Block Header

Block = (Header, Transactions, Uncle Headers)

## Block Header

|                  |           |       |
|------------------|-----------|-------|
| parentHash       | 32        | bytes |
| ommersHash       | 32        | bytes |
| beneficiary      | 20        | bytes |
| stateRoot        | 32        | bytes |
| transactionsRoot | 32        | bytes |
| receiptsRoot     | 32        | bytes |
| logsBloom        | 256       | bytes |
| difficulty       | $\geq 1$  | byte  |
| number           | $\geq 1$  | byte  |
| gasLimit         | $\geq 1$  | byte  |
| gasUsed          | $\geq 1$  | byte  |
| timestamp        | $\leq 32$ | bytes |
| extraData        | $\leq 32$ | bytes |
| mixHash          | 32        | bytes |
| nonce            | 8         | bytes |

# Simple Fields in Block Header

|                         |           |       |
|-------------------------|-----------|-------|
| <b>parentHash</b>       | 32        | bytes |
| ommersHash              | 32        | bytes |
| <b>beneficiary</b>      | 20        | bytes |
| <b>stateRoot</b>        | 32        | bytes |
| <b>transactionsRoot</b> | 32        | bytes |
| receiptsRoot            | 32        | bytes |
| logsBloom               | 256       | bytes |
| difficulty              | $\geq 1$  | byte  |
| <b>number</b>           | $\geq 1$  | byte  |
| gasLimit                | $\geq 1$  | byte  |
| gasUsed                 | $\geq 1$  | byte  |
| <b>timestamp</b>        | $\leq 32$ | bytes |
| <b>extraData</b>        | $\leq 32$ | bytes |
| mixHash                 | 32        | bytes |
| nonce                   | 8         | bytes |

- parentHash = Keccak-256 hash of parent block header
- beneficiary = Destination address of block reward and transaction fees
- stateRoot = Root hash of world state trie after all transactions are applied
- transactionsRoot = Root hash of trie populated with all transactions in the block
- number = Number of ancestor blocks
- timestamp = Unix time at block creation
- extraData = Arbitrary data; Miners identify themselves in this field

## gasLimit and gasUsed

|                  |           |       |
|------------------|-----------|-------|
| parentHash       | 32        | bytes |
| ommersHash       | 32        | bytes |
| beneficiary      | 20        | bytes |
| stateRoot        | 32        | bytes |
| transactionsRoot | 32        | bytes |
| receiptsRoot     | 32        | bytes |
| logsBloom        | 256       | bytes |
| difficulty       | $\geq 1$  | byte  |
| number           | $\geq 1$  | byte  |
| <b>gasLimit</b>  | $\geq 1$  | byte  |
| <b>gasUsed</b>   | $\geq 1$  | byte  |
| timestamp        | $\leq 32$ | bytes |
| extraData        | $\leq 32$ | bytes |
| mixHash          | 32        | bytes |
| nonce            | 8         | bytes |

- gasUsed is the total gas used by all transactions in the block
- gasLimit is the maximum gas which can be used
- $|\text{gasLimit} - \text{parent.gasLimit}| \leq \frac{\text{parent.gasLimit}}{1024}$
- Miner can choose to increase or decrease the gasLimit

# logsBloom and receiptsRoot

|                     |           |       |
|---------------------|-----------|-------|
| parentHash          | 32        | bytes |
| ommersHash          | 32        | bytes |
| beneficiary         | 20        | bytes |
| stateRoot           | 32        | bytes |
| transactionsRoot    | 32        | bytes |
| <b>receiptsRoot</b> | 32        | bytes |
| <b>logsBloom</b>    | 256       | bytes |
| difficulty          | $\geq 1$  | byte  |
| number              | $\geq 1$  | byte  |
| gasLimit            | $\geq 1$  | byte  |
| gasUsed             | $\geq 1$  | byte  |
| timestamp           | $\leq 32$ | bytes |
| extraData           | $\leq 32$ | bytes |
| mixHash             | 32        | bytes |
| nonce               | 8         | bytes |

- Bloom filter = Probabilistic data structure for set
  - **Query:** Is  $x$  in the set? **Response:** “Maybe” or “No”
- receiptsRoot is the root hash of transaction receipts trie
  - Each transaction receipt contains Bloom filter of addresses and “topics”
- logBloom is the OR of all transaction receipt Bloom filters
- Light clients can efficiently retrieve only transactions of interest

Mining

# Ethash Mining Algorithm

- An epoch lasts 30,000 blocks
- Epoch index  $EI = \text{block\_number} / 30000$
- At an epoch beginning
  - A list called cache of size  $\approx 2^{24} + EI \times 2^{17}$  bytes is created
  - A list called dataset of size  $\approx 2^{30} + EI \times 2^{23}$  bytes is created
- The dataset is also called the DAG (directed acyclic graph)

| Block Number | Epoch | Cache Size | DAG Size | Start Date   |
|--------------|-------|------------|----------|--------------|
| 30000        | 1     | 16 MB      | 1 GB     | 17 Oct, 2015 |
| 3840000      | 128   | 32 MB      | 2 GB     | 21 Jul, 2017 |
| 7680000      | 256   | 48 MB      | 3 GB     | 30 Apr, 2019 |
| 192000000    | 640   | 96 MB      | 6 GB     | 25 Aug, 2024 |

Source: [https://investoon.com/tools/dag\\_size](https://investoon.com/tools/dag_size)

- Mining nodes need to store full dataset (ASIC resistance)
- Light nodes store cache and recalculate specific dataset items

# Ethash Mining Algorithm

|                  |           |       |
|------------------|-----------|-------|
| parentHash       | 32        | bytes |
| ommersHash       | 32        | bytes |
| beneficiary      | 20        | bytes |
| stateRoot        | 32        | bytes |
| transactionsRoot | 32        | bytes |
| receiptsRoot     | 32        | bytes |
| logsBloom        | 256       | bytes |
| difficulty       | $\geq 1$  | byte  |
| number           | $\geq 1$  | byte  |
| gasLimit         | $\geq 1$  | byte  |
| gasUsed          | $\geq 1$  | byte  |
| timestamp        | $\leq 32$ | bytes |
| extraData        | $\leq 32$ | bytes |
| <b>mixHash</b>   | 32        | bytes |
| <b>nonce</b>     | 8         | bytes |

- Cache calculation involves hashing previous cache elements pseudorandomly
- Every dataset element involves hashing 256 pseudorandom cache elements
- Mining loop takes partial header hash, `nonce`, and dataset as input
- 128 dataset elements are used to create 256-bit `mixHash`

Mining output = Keccak256 (Keccak512(HdrHash||nonce)||mixHash)



# Mining Difficulty

|                   |           |       |
|-------------------|-----------|-------|
| parentHash        | 32        | bytes |
| ommersHash        | 32        | bytes |
| beneficiary       | 20        | bytes |
| stateRoot         | 32        | bytes |
| transactionsRoot  | 32        | bytes |
| receiptsRoot      | 32        | bytes |
| logsBloom         | 256       | bytes |
| <b>difficulty</b> | $\geq 1$  | byte  |
| number            | $\geq 1$  | byte  |
| gasLimit          | $\geq 1$  | byte  |
| gasUsed           | $\geq 1$  | byte  |
| timestamp         | $\leq 32$ | bytes |
| extraData         | $\leq 32$ | bytes |
| <b>mixHash</b>    | 32        | bytes |
| <b>nonce</b>      | 8         | bytes |

- Proof of work is valid if `mixhash` and `nonce` lead to

$$\text{Keccak256}(\text{Keccak512}(\text{HdrHash}||\text{nonce})||\text{mixHash}) \leq \frac{2^{256}}{\text{difficulty}}$$


- Partial validation of PoW in block can be done without DAG or cache
- Difficulty adjustment algorithm explained after discussing uncles

# Uncle Incentivization

# Uncle Blocks


- Block = (Block Header, Transactions List, Uncle Header List)
- `ommersHash` in block header is hash of uncle header list
- **Problem:** Low inter-block time leads to high stale rate
  - Stale blocks do not contribute to network security
  - High stale rate may lead to mining centralization
- **Solution:** Reward stale block miners and also miners who include stale block headers
- Rewarded stale blocks are called uncles or ommers
  - Transactions in uncle blocks are invalid
  - Only a fraction of block reward goes to uncle creator; no transaction fees
- Greedy Heaviest Observed Subtree (GHOST) protocol proposed by Sompolinsky and Zohar in December 2013
- Ethereum uses a simpler version of GHOST

# GHOST Protocol




- A policy for choosing the main chain in case of forks
- Given a block tree  $T$ , the protocol specifies  $\text{GHOST}(T)$  as the block representing the main chain
- Mining nodes calculate  $\text{GHOST}(T)$  locally and mine on top of it
- Heaviest subtree rooted at fork is chosen

# GHOST Protocol



```
function CHILDRENT(B)  
    return Set of blocks with B as immediate parent  
end function  
function SUBTREET(B)  
    return Subtree rooted at B  
end function  
function GHOST(T)  
    B ← Genesis Block  
    while True do  
        if CHILDRENT(B) = ∅ then return B and exit  
        else B ← argmaxC ∈ CHILDRENT(B) |SUBTREET(C)|  
        end if  
    end while  
end function
```

# GHOST Protocol Example



- Suppose an attacker secretly constructs the chain 1A, 2A, ..., 6A
- All other blocks are mined by honest miners
- Honest miners' efforts are spread over multiple forks
- Longest chain rule gives 0, 1B, 2D, 3F, 4C, 5B as main chain
  - Shorter than attacker's chain
- GHOST rule gives 0, 1B, 2C, 3D, 4B as main chain

# Main Chain Selection and Uncle Rewards

- Chain with maximum total difficulty is chosen
  - Total difficulty is sum of block difficulty values
- Uncles contribute to difficulty since Oct 2017 (Byzantium)
- A uncle block of a given block satisfies the following
  - Cannot be a direct ancestor of given block
  - Cannot already be included as an uncle block in the past
  - Has to be the child of given block's ancestor at depth 2 to 7
- Mining reward
  - Block reward = 3 ETH, Nephew reward =  $\frac{3}{32}$  ETH
  - Total reward to block miner is

$$\text{Block reward} + \text{NumUncles} \times \text{Nephew reward}$$

- NumUncles can be at most 2
- Uncle miner gets

$$\text{Block reward} \times \frac{(8 + \text{UncleHeight} - \text{BlockHeight})}{8}$$

# Difficulty Adjustment



# Difficulty Adjustment Algorithm Evolution

Frontier Release, July 2015

```
1  MIN_DIFF = 131072
2
3  def calc_difficulty(parent, timestamp):
4      offset = parent.difficulty // 2048
5      sign = 1 if timestamp - parent.timestamp < 13 else -1
6      return int(max(parent.difficulty + offset * sign, MIN_DIFF))
```

- If difference between current timestamp and parent's timestamp is less than 13 seconds, difficulty is increased
- Otherwise, difficulty is decreased
- Quantum of change is  $\frac{1}{2048}$  of parent block's difficulty
- Difficulty is not allowed to go below a fixed minimum

# Difficulty Adjustment Algorithm Evolution

Patch to Frontier Release, August 2015

```
1 MIN_DIFF = 131072
2 EXPDIFF_PERIOD = 100000
3 EXPDIFF_FREE_PERIODS = 2
4
5 def calc_difficulty(parent, timestamp):
6     offset = parent.difficulty // 2048
7     sign = 1 if timestamp - parent.timestamp < 13 else -1
8     o = int(max(parent.difficulty + offset * sign, MIN_DIFF))
9     period_count = (parent.number + 1) // EXPDIFF_PERIOD
10    if period_count >= EXPDIFF_FREE_PERIODS:
11        o = max(o + 2** (period_count - EXPDIFF_FREE_PERIODS),
12               MIN_DIFF)
13    return o
```

- Difficulty time bomb was added to force move to proof-of-stake
- Bomb term added to every block's difficulty double every 100,000 blocks
- Ice age = Blocks too difficult to find

# Difficulty Adjustment Algorithm Evolution

Homestead Release, March 2016, Block 1150000

```
1 def calc_difficulty(parent, timestamp):
2     <snip>
3     time_diff = timestamp - parent.timestamp
4     sign = max(1 - time_diff // 10, -99)
5     <snip>
```

- Protocol requires `timestamp > parent.timestamp`
- `time_diff` in range 1, 2, ..., 9  $\implies$  `sign = 1`
- `time_diff` in range 10, 11, ..., 19  $\implies$  `sign = 0`
- `time_diff` in range 20, 21, ..., 29  $\implies$  `sign = -1`
- `time_diff  $\geq$  1010  $\implies$  sign = -99`
- Rationale
  - Previous algorithm targeted a median block time of 13 seconds
  - New algorithm targets a mean block time of 15 seconds
    - See vague justification in EIP 2

# Difficulty Adjustment Algorithm Evolution

Byzantium Release, October 2017, Block 4370000

```
1 def calc_difficulty(parent, timestamp):
2     EXPDIFF_PERIOD = 100000
3     <snip>
4     time_diff = timestamp - parent.timestamp
5     uncle_factor = 2 if len(parent.uncles) else 1
6     sign = max(uncle_factor - time_diff // 9, -99)
7     <snip>
8     period_count = (parent.number + 1) // EXPDIFF_PERIOD
9     period_count = period_count - 30
10    if period_count >= 2:
11        o = max(o + 2**(period_count - 2), MIN_DIFF)
12    return o
```

- Take uncles into account while adjusting difficulty
  - <https://github.com/ethereum/EIPs/issues/100>
- Delays ice age by approximately 42 million seconds to account for PoS delays
- Block reward reduced from 5 ETH to 3 ETH

# References

- **Yellow paper** <https://ethereum.github.io/yellowpaper/paper.pdf>
- **Light client protocol**  
<https://github.com/ethereum/wiki/wiki/Light-client-protocol>
- **Ethash** <https://github.com/ethereum/wiki/wiki/Ethash>
- **GHOST paper** <https://eprint.iacr.org/2013/881>
- **Uncle calculations** <https://github.com/ethereum/pyethereum/blob/develop/ethereum/pow/consensus.py>
- **Homestead difficulty adjustment**  
<https://ethereum.stackexchange.com/questions/5913/how-does-the-ethereum-homestead-difficulty-adjustment-algorithm-work>
- **Rationale for Homestead difficulty adjustment**  
<https://github.com/ethereum/EIPs/blob/master/EIPS/eip-2.md>
- **Byzantium difficulty adjustment** <https://blog.ethereum.org/2017/10/12/byzantium-hf-announcement/>