L02 — Blockchain Fundamentals

Cheatsheet

Block Structure

Formal Definition

$\text{Block}_n = (\text{Hash}_{n-1},\ \text{Timestamp},\ \text{Nonce},\ \text{MerkleRoot},\ \text{Txns})$

Header (80 bytes in Bitcoin):

Field	Size	Purpose
Version	4 B	Protocol rules in effect
Prev. Hash	32 B	Links to previous block
Merkle Root	32 B	Fingerprint of all txns
Timestamp	4 B	Unix time of block creation
Bits (Target)	4 B	Current difficulty
Nonce	4 B	Mining solution counter

Body: ordered list of transactions; first is always the *coinbase* (miner reward, no inputs).

Hash Chain & Immutability

Each block includes the SHA-256 hash of the *previous* block header, creating a cryptographic chain.

Tamper cascade: modifying Block k changes its hash ⇒ Block k+1's "Prev Hash" is now invalid ⇒ must redo PoW for all subsequent blocks.

Practical security: attacker needs $> 50\%$ of global hash rate to outpace the honest chain—economically prohibitive.

Genesis block (Block 0): Previous Hash $= \underbrace{00\cdots0}_{64}$;

Satoshi embedded the message *"Chancellor on brink of second bailout for banks"* (3 Jan 2009).

Transaction Lifecycle

1. **Create**: User constructs transaction (inputs, outputs, amount)
2. **Sign**: User signs with private key (ECDSA proof of ownership)
3. **Broadcast**: Sent to connected peers via gossip protocol
4. **Mempool**: Unconfirmed transactions wait; miner selects by fee
5. **Mine**: Miner includes txn in block, solves PoW puzzle
6. **Propagate**: New block gossips across network (∼seconds)
7. **Confirm**: Each subsequent block = 1 more confirmation; 6 confirmations ≈ finality for Bitcoin

Replace-By-Fee (RBF): user can rebroadcast stuck txn with higher fee.

Merkle Trees

Structure

Binary hash tree: leaves $= H(\text{Tx}_i)$; parents $= H(\text{left}\|\text{right})$. Root stored in block header summarises all transactions.

Root formula (4 transactions):

$$\text{Root} = H(H(H(T_1)\|H(T_2)) \parallel H(H(T_3)\|H(T_4)))$$

SPV proof size: $O(\log_2 n)$ hashes needed to prove inclusion.
Example: 1 000 txns → only 10 hashes; 1 000 000 txns → only 20 hashes.

Light client (SPV): downloads only block headers (∼75 MB vs ∼720 GB for full node); verifies via Merkle proof that a transaction is included without trusting a third party.

Network Node Types

Node Type	Role & Resources
Full node	Stores entire chain; validates all txns; enforces consensus rules
Light (SPV)	Stores headers only; relies on Merkle proofs; mobile-friendly
Mining node	Full node + creates blocks; competes for block reward
Archival node	Full node preserving all historical state (ETH: older pruned data)

Bitcoin network: ≈25 000 reachable nodes; ≈50 000+ total estimated.

Fork Types

Fork	Meaning & Examples
Soft fork	New rules are a *subset* of old rules; backward compatible. Old nodes accept new blocks. Majority upgrade sufficient. *Examples*: SegWit (2017), Taproot (2021)
Hard fork	Makes previously invalid blocks valid; *not* backward compatible. Old nodes reject new blocks; permanent chain split if not unanimous. *Examples*: Bitcoin Cash (2017, 8 MB blocks), Ethereum Classic (2016, DAO rollback), Bitcoin SV (2018)

Scalability Trilemma (Buterin)

Any blockchain can fully achieve **at most two** of:

Property	Meaning
Scalability	High TPS; fast settlement
Decentralisation	Many independent validators
Security	Resistant to $> 50\%$ attack

Throughput comparison:

- Bitcoin (BTC): ≈7 TPS

- Ethereum (ETH): ≈15–30 TPS
- Visa: ≈24 000 TPS (peak ∼65 000)
- Solana: ∼3 000–65 000 TPS (less decentralised)

Layer 2 Solutions

L2 inherits L1 security while dramatically increasing throughput.

State channels (e.g., Lightning Network): off-chain payment channel; only open/close txns settle on L1; instant, near-zero fees; best for repeated micro-payments.

Optimistic rollups (Arbitrum, Optimism): bundle txns off-chain, post compressed data + state root to L1; assume valid, dispute window ∼7 days; 10×–100× throughput gain.

ZK-rollups (zkSync, StarkNet): generate cryptographic validity proof for each batch; no dispute period; higher proving cost but faster finality.

Sidechains (Polygon PoS): independent chain with own consensus; bridge to L1; weaker security guarantees than rollups.

Key Formulas

Block reward halving:

$$R_n = \frac{50}{2^{\lfloor n/210{,}000 \rfloor}} \text{ BTC}$$

where n = block height. Last BTC mined ≈ year 2140; total supply = 21M BTC.

Confirmation security (attacker hash fraction q, honest $p = 1 - q$):

$$P(\text{catch up after } z \text{ blocks}) \approx \left(\frac{q}{p}\right)^z$$

For $q = 0.1$, $z = 6$: $P < 0.1\%$.

Merkle proof size: $O(\log_2 n)$ hashes for n transactions.

Mining target: valid block requires $H(\text{header}) <$ target; difficulty adjusts every 2 016 blocks (≈2 weeks) to maintain 10 min/block.

Terminology Quick Reference

Term	One-Line Definition
Merkle root	Single hash summarising all block transactions
Nonce	Counter varied by miner to satisfy PoW target
Mempool	Pool of unconfirmed transactions at each node
Coinbase txn	First block transaction; generates new BTC reward
UTXO	Unspent Transaction Output — Bitcoin's balance model
Orphan block	Valid block not on canonical chain (shorter branch)
Reorg	Network switches to longer valid chain, abandons tip
SPV	Simplified Payment Verification (light client)
Gossip	Peer-to-peer epidemic broadcast of new data
Confirmation	Each block added on top of a transaction's block
51% attack	Majority hash rate enables chain rewrite
SegWit	Soft fork separating signature data; raises effective block size