

Theory — Fractal Persistence in detail

Contents

Formal Specification of Blockchain Systems (v2.0)	2
0. Preliminaries: Cryptographic Primitives	2
Definition 0.1 (Digital Signature Scheme)	2
Definition 0.2 (Cryptographic Hash Function)	3
Definition 0.3 (Verifiable Delay Function, VDF)	4
I. Layer 1: Application & State Machine	5
Definition 1.1 (State Space)	5
Definition 1.2 (Event)	6
Definition 1.3 (State Transition Function)	7
Definition 1.4 (State Commitment)	8
II. Layer 2: Distributed Chronometry	9
Definition 2.1 (Epochs)	9
Definition 2.2 (Temporal Authority Oracle)	10
Definition 2.3 (Epoch Transition Function)	11
III. Layer 3: Network Topology & Message Propagation	12
Definition 3.1 (The Overlay Graph)	12
Definition 3.2 (Local Mempool)	13
Definition 3.3 (Event Origination)	14
Definition 3.4 (Gossip Protocol)	15
Definition 3.5 (Block Propagation)	16
Definition 3.6 (Network Adversary)	17
Definition 3.7 (Eclipse Resistance)	18
IV. Layer 4: Consensus & Fork Choice	19
Definition 4.1 (Block Structure)	19
Definition 4.2 (Network-Adjusted Stability Condition)	20
Definition 4.3 (The Blockchain)	21
Definition 4.3a (Slot-Skip Recovery for Liveness)	22
Definition 4.4 (Fork Choice Rule)	23
Definition 4.5 (Latency Arbitrage)	24
V. Layer 5: Cryptoeconomic Security	25
Definition 5.1 (Resource Commitment)	25
Definition 5.2 (Slashing Function)	26
Definition 5.3 (Incentive Compatibility with Network)	27
VI. System Properties	28
Definition 6.1 (Safety)	28
Definition 6.2 (Liveness)	29
Definition 6.3 (Censorship Resistance)	30
VII. Critical Theorems	31
Theorem 7.1 (Network-Safety Interdependence)	31
Theorem 7.3 (Propagation-Liveness Trade-off)	32
Appendix: Complete Instantiation Examples	33
Summary of Changes (v2.0)	34
Summary of Dependencies	35

Formal Specification of Blockchain Systems (v2.0)

0. Preliminaries: Cryptographic Primitives

Blockchain systems rely on cryptographic primitives to ensure security, authenticity, and verifiability. These primitives are modeled as idealized functions with well-defined security properties.

Definition 0.1 (Digital Signature Scheme)

A **digital signature scheme** is a tuple of probabilistic polynomial-time (PPT) algorithms:

$$\mathcal{S} = (\text{KeyGen}, \text{Sign}, \text{Verify})$$

where: - $\text{KeyGen}(1^\lambda) \rightarrow (sk, pk)$: Generates a secret key sk and public key pk for security parameter λ .
- $\text{Sign}(sk, m) \rightarrow \sigma$: Produces a signature σ for message m . - $\text{Verify}(pk, m, \sigma) \rightarrow \{0, 1\}$: Returns 1 if σ is valid for m under pk .

Security Property (EUF-CMA): An adversary, given access to a signing oracle, cannot produce a valid signature σ^* on a message m^* that was not previously signed, except with negligible probability.

Intuition: Signatures ensure that transactions and blocks are authored by legitimate participants. In blockchain systems, they prevent forgery and enable non-repudiation.

Definition 0.2 (Cryptographic Hash Function)

A **cryptographic hash function** is a function:

$$\mathcal{H} : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$$

modeled as a **random oracle** (an idealized function that outputs uniformly random values for new inputs).

Security Properties: 1. **Collision Resistance:** It is computationally infeasible to find $x \neq y$ such that $\mathcal{H}(x) = \mathcal{H}(y)$. 2. **Pre-image Resistance:** Given $y = \mathcal{H}(x)$, it is computationally infeasible to recover x . 3. **Second Pre-image Resistance:** Given x , it is computationally infeasible to find $x' \neq x$ such that $\mathcal{H}(x) = \mathcal{H}(x')$.

Intuition: Hash functions are used to: - Commit to data (e.g., Merkle roots for state commitments). - Link blocks in a chain (via block hashes). - Generate pseudorandom values (e.g., for leader election).

Definition 0.3 (Verifiable Delay Function, VDF)

A **VDF** is a triple of algorithms:

$$\mathcal{VD}\mathcal{F} = (\text{Setup}, \text{Eval}, \text{Verify})$$

where: - $\text{Setup}(1^\lambda, \Delta) \rightarrow pp$: Generates public parameters pp for a delay parameter Δ . - $\text{Eval}(pp, x) \rightarrow (y, \pi)$: Computes output y and proof π after Δ sequential steps. - $\text{Verify}(pp, x, y, \pi) \rightarrow \{0, 1\}$: Verifies y in $O(\log \Delta)$ time.

Intuition: VDFs introduce **sequential computation** that cannot be parallelized, ensuring that: - Block production requires real time (e.g., in Solana's Proof-of-History). - Adversaries cannot speed up computation via parallelism.

I. Layer 1: Application & State Machine

This layer defines the logical state, the programs residing within it, and the generic “Events” that trigger state evolution.

Definition 1.1 (State Space)

The **state space** Σ is defined as a tuple comprising data storage and program logic:

$$\Sigma = (\mathcal{D}, \Omega)$$

Where: - $\mathcal{D} : \mathcal{A} \rightarrow \mathbb{N}$ is the **Data Layer** (e.g., mapping addresses to balances/UTXOs). - $\Omega : \mathcal{A} \rightarrow \{0, 1\}^*$ is the **Program Space** (Smart Contracts). $\Omega(a)$ stores the immutable bytecode at address a . - The **genesis state** $\sigma_0 = (\mathcal{D}_0, \emptyset)$ initializes the data and contains no deployed programs.

Definition 1.2 (Event)

An **Event** ϵ is the atomic unit of state transition. It generalizes the concept of a transaction to include simple transfers, smart contract deployments, and contract executions.

$$\epsilon = (\text{type}, \text{payload}, \pi_{\text{auth}})$$

Where: - **type** $\in \{\text{Transfer}, \text{Deploy}, \text{Invoke}\}$: Discriminator defining the event's nature. - **payload**: The data payload, structure varies by type: - If **type** = **Transfer**: Contains sender, recipient, value. - If **type** = **Deploy**: Contains sender, init-code (bytecode to store). - If **type** = **Invoke**: Contains sender, target address a , function selector, and input arguments. - π_{auth} : Authentication proof (e.g., digital signature) authorizing the event.

Validity Predicate:

$$\mathcal{V} : \Sigma \times \mathcal{E} \rightarrow \{0, 1\}$$

$\mathcal{V}(\sigma, \epsilon) = 1$ if the signature is valid and the sender has sufficient resources (gas/balance).

Definition 1.3 (State Transition Function)

The **state transition function** δ updates the state based on the event type. It is a partial function:

$$\delta : \Sigma \times \mathcal{E} \rightarrow \Sigma \cup \{\perp\}$$

The transition logic is defined by cases on **type**:

1. **Simple Transfer:**

$$\delta(\sigma, \epsilon) = \sigma' \text{ where } \mathcal{D}' = \mathcal{D} \text{ updated with value transfer}$$

2. **Smart Contract Deployment (Deploy):**

$$\delta(\sigma, \epsilon) = \sigma' \text{ where } \Omega'(a_{\text{new}}) = \text{init_code}$$

This event adds new logic to the state.

3. **Smart Contract Execution (Invoke):**

$$\delta(\sigma, \epsilon) = \text{Execute}(\sigma, \Omega(a_{\text{target}}), \text{payload})$$

This loads the bytecode from Ω at the target address and executes it. The execution may read/write \mathcal{D} and is bounded by a resource limit (Gas) Γ :

$$\text{If } \Gamma_{\text{consumed}} > \Gamma_{\text{limit}} \implies \text{Revert}(\sigma)$$

Intuition: In traditional blockchains, transactions are simple database updates. In smart contract platforms, an **Event** triggers the execution of the **Program** stored in the state, allowing the state to mutate itself autonomously based on predefined logic.

Definition 1.4 (State Commitment)

A **state commitment** remains a succinct cryptographic binding to the entire state:

$$\mathcal{M} : \Sigma \rightarrow \{0, 1\}^\lambda$$

Note: For Ethereum, $\mathcal{M}(\sigma)$ is the **State Root**, which commits to both the account balances (\mathcal{D}) and the contract storage/code (Ω) via a Modified Merkle Patricia Trie.

II. Layer 2: Distributed Chronometry

This layer defines **how time is measured** in a decentralized system, where nodes may have differing local clocks.

Definition 2.1 (Epochs)

An **epoch** is a discrete time interval during which a subset of validators is authorized to propose blocks. The sequence of epochs is totally ordered:

$$\varepsilon_0, \varepsilon_1, \varepsilon_2, \dots$$

where: - Each epoch ε_i has duration $\delta_i \in \mathbb{R}^+$. - δ_i may be fixed (e.g., 12s in Ethereum) or adaptive (e.g., Bitcoin's difficulty adjustment).

Intuition: Epochs provide a **synchronization mechanism** for validators, ensuring that: - Only one validator proposes a block per epoch (in leader-based protocols). - Validators agree on the order of blocks.

Definition 2.2 (Temporal Authority Oracle)

The **temporal authority oracle** \mathcal{O} determines which validator is authorized to propose a block in epoch i :

$$\mathcal{O} : \mathbb{N} \times \mathcal{V} \times \Theta \times \Sigma \rightarrow \{0, 1\}$$

where: - \mathcal{V} is the set of validators. - Θ represents temporal parameters (e.g., RANDAO seed, VRF output). - Σ is the current state (e.g., stake distribution).

Output:

$$\mathcal{O}(i, v, \theta_i, \sigma_i) = 1 \iff v \text{ is authorized to propose in epoch } i.$$

Examples: 1. **Proof-of-Work (Bitcoin):** - $\mathcal{O}(i, v, \theta_i, \sigma_i) = 1$ if v finds a nonce n such that $\mathcal{H}(B_{i-1} \| n) < D$, where D is the difficulty target. 2. **Proof-of-Stake (Ethereum):** - $\mathcal{O}(i, v, \theta_i, \sigma_i) = 1$ if v is selected via RANDAO + VRF based on stake. 3. **Proof-of-History (Solana):** - $\mathcal{O}(i, v, \theta_i, \sigma_i) = 1$ if v is the next validator in a VDF-generated sequence.

Definition 2.3 (Epoch Transition Function)

The **epoch transition function** Ψ dynamically adjusts the epoch duration based on network conditions:

$$\delta_{i+1} = \Psi(\sigma_i, \gamma_i, \delta_{\text{base}})$$

where: - σ_i is the state at the end of epoch i . - $\gamma_i \in \Gamma$ represents **network observations** (e.g., block propagation time, congestion metrics). - δ_{base} is a baseline duration (e.g., 12s in Ethereum).

Examples: 1. **Bitcoin:** - Ψ adjusts δ_{base} every 2016 blocks to target a 10-minute block time. - γ_i includes the actual time taken to mine the last 2016 blocks. 2. **Ethereum (Post-Merge):** - $\delta_{\text{base}} = 12s$ (fixed). - Ψ does not adjust δ_i (slot time is constant). 3. **Solana:** - δ_i is adaptive based on network load. - γ_i includes metrics like transaction throughput and leader rotation speed.

III. Layer 3: Network Topology & Message Propagation

This layer grounds the abstract “mempool” in the **physical network**, modeling how transactions and blocks propagate across nodes.

Definition 3.1 (The Overlay Graph)

The **overlay graph** $G_t = (N_t, E_t, \omega_t)$ is a time-varying directed weighted graph representing the P2P network at time t : - **Nodes** N_t : - \mathcal{V}_t : **Validators** (consensus participants, e.g., miners or stakers). - \mathcal{C}_t : **Clients** (transaction originators, full/light nodes). - \mathcal{R}_t : **Relay infrastructure** (e.g., sentries, block builders, MEV relayers). - **Edges** $E_t \subseteq N_t \times N_t$: - Persistent P2P connections (e.g., TCP links). - May be **asymmetric** (e.g., a validator may connect to a relay but not vice versa). - **Latency Weight** $\omega_t : E_t \rightarrow \mathbb{R}^+$: - $\omega_t(u, v)$ is the **propagation delay** for a message from u to v . - Depends on: - Physical distance (geographic latency). - Bandwidth constraints. - Network congestion.

Intuition: The overlay graph captures the **real-world constraints** of blockchain networks: - Nodes are not fully connected (unlike idealized models). - Latency varies significantly (e.g., 10ms within a data center vs. 200ms intercontinental). - Adversaries can manipulate G_t (e.g., eclipse attacks).

Definition 3.2 (Local Mempool)

The **mempool** is the set of pending Events. Each node u maintains:

$$\mathcal{M}_u(t) = \{\epsilon \mid \epsilon \text{ received by } u \text{ before } t\} \setminus \{\epsilon \mid \epsilon \in B_j\}$$

where: - B_j is a confirmed block. - $\mathcal{M}_u(t)$ is **node-specific** and may differ from $\mathcal{M}_v(t)$ for $u \neq v$.

Properties: 1. **Non-global:** Two nodes may have different mempools due to network latency or censorship. 2. **Dynamic:** Transactions are added when received and removed when included in a block. 3. **Priority-based:** Transactions may be ordered by fee, arrival time, or other heuristics.

Example (Ethereum): - $\mathcal{M}_u(t)$ is a priority queue ordered by gas price. - Nodes may drop low-fee transactions if the mempool is full.

Definition 3.3 (Event Origination)

An event ϵ originates at a source client s and propagates via the gossip protocol. The **propagation time** to a validator $v \in \mathcal{V}_t$ is:

$$\Delta_{s,v}(t) = \min_{p \in \text{Paths}_{G_t}(s,v)} \sum_{(u,w) \in p} \omega_t(u,w)$$

where: - $\text{Paths}_{G_t}(s,v)$ is the set of all paths from s to v in G_t . - $\Delta_{s,v}(t)$ is the **shortest-path latency** from s to v .

Intuition: - Events do not appear instantly in all mempools. - Validators closer to s (in network terms) see ϵ earlier. - This creates **latency arbitrage opportunities** (see Definition 4.5).

Definition 3.4 (Gossip Protocol)

The **gossip protocol** \mathcal{G} determines how transactions are forwarded:

$$\mathcal{G} : \mathcal{M}_u \times 2^{N_t} \rightarrow 2^{\mathcal{T}}$$

where $\mathcal{G}(\mathcal{M}_u, \text{Neigh}(u))$ returns the set of transactions to send to neighbors $\text{Neigh}(u)$.

Variants: 1. **Flooding (Naive Gossip):** - $\mathcal{G}(\mathcal{M}_u, \text{Neigh}(u)) = \mathcal{M}_u \setminus \mathcal{M}_u^{\text{sent}}$. - Simple but bandwidth-inefficient (redundant transmissions). 2. **Diffusion (Erlay):** - Uses **Invertible Bloom Lookup Tables (IBLTs)** to reconcile mempools with neighbors. - Reduces bandwidth by only sending differences. 3. **Structured Gossip (e.g., Ethereum's Gossipsub):** - Nodes form a **mesh** and forward messages along the mesh. - More efficient but requires topology maintenance.

Intuition: Gossip protocols balance: - **Bandwidth efficiency** (minimize redundant transmissions). - **Propagation speed** (ensure fast inclusion). - **Censorship resistance** (prevent adversaries from blocking transactions).

Definition 3.5 (Block Propagation)

When a validator v produces a block B_i , it propagates the block via:

$$\mathcal{G}_B : B_i \times N_t \rightarrow \{\text{ack}, \text{nack}\}^{\text{Neigh}(v)}$$

where: - \mathcal{G}_B sends B_i to neighbors and waits for acknowledgments. - The **block propagation diameter** D_t is the worst-case time for B_i to reach all honest validators:

$$D_t = \max_{v, u \in \mathcal{V}_{\text{honest}}} \Delta_{v,u}(t)$$

Intuition: - D_t is a **critical parameter** for consensus safety and liveness. - If D_t is too large, validators may propose conflicting blocks (forks). - Protocols like **Compact Blocks (Bitcoin)** or **Weak Subjectivity (Ethereum)** reduce D_t by sending only block headers first.

Definition 3.6 (Network Adversary)

The **adversary** \mathcal{A} controls a subset of nodes $N_{\text{corr}} \subset N_t$ and can: 1. **Delay Messages:** - Add δ_{adv} to ω_t on edges incident to N_{corr} . - Example: A malicious ISP throttling traffic to a validator. 2. **Drop Messages:** - Remove edges (selective censorship). - Example: A relay node refusing to forward transactions. 3. **Eclipse Attacks:** - Partition the graph such that all paths between honest validators pass through N_{corr} . - Example: An adversary controlling all peers of a victim validator.

Intuition: The adversary's goal is to: - **Break safety** (create forks by delaying blocks). - **Break liveness** (censor transactions). - **Extract MEV** (exploit latency advantages).

Definition 3.7 (Eclipse Resistance)

A protocol has (ϵ, δ) -**Eclipse Resistance** if for any honest validator $v \in \mathcal{V}$:

$$\Pr[\mathcal{A} \text{ can eclipse } v] \leq \epsilon$$

provided: 1. $\deg_{G_t}(v) \geq \delta$ (sufficiently many peers). 2. Fewer than $\frac{1}{3}$ of v 's neighbors are adversarial (for BFT protocols).

Intuition: Eclipse resistance ensures that: - Validators cannot be isolated by the adversary. - The network remains connected even under attack.

Example (Ethereum): - Validators are required to maintain ≥ 50 peers. - The adversary needs to control ≥ 17 of a validator's peers to eclipse it (assuming $\frac{1}{3}$ threshold).

IV. Layer 4: Consensus & Fork Choice

This layer defines how validators agree on a single chain of blocks, despite network delays and adversarial behavior.

Definition 4.1 (Block Structure)

A block B_i is a tuple:

$$B_i = (H_{i-1}, \sigma_{\text{root}}^{(i)}, \vec{e}_i, \pi_{\text{consensus}}, \eta_i)$$

where: - H_{i-1} : Hash of the parent block. - $\sigma_{\text{root}}^{(i)}$: State commitment after applying the events. - $\vec{e}_i \subseteq \mathcal{M}_v(t_i^{\text{start}})$: **An ordered sequence of Events** (Transfers, Deploys, and Invocations) included by the proposer. - $\pi_{\text{consensus}}$: Proof of authority (PoW/PoS). - η_i : Additional metadata (e.g., timestamp, gas limit).

Intuition: A block is a batch of Events. Validating a block requires re-executing the sequence \vec{e}_i starting from the previous state σ_{i-1} to verify that the resulting state matches the claimed $\sigma_{\text{root}}^{(i)}$.

Definition 4.2 (Network-Adjusted Stability Condition)

The **epoch duration** δ_i must satisfy:

$$\delta_i > D_t + \Delta_{\text{process}}(\vec{c}_i) + \Delta_{\text{verify}}(\vec{c}_i)$$

Note: Δ_{process} and Δ_{verify} now depend on the computational complexity of the smart contracts invoked in \vec{c}_i (bounded by the block gas limit).

Intuition: This condition ensures that: - All honest validators observe B_{i-1} before proposing B_i . - No validator proposes B_i before seeing all pending events.

Example (Ethereum): - $\delta_i = 12s$ (fixed). - $D_t \approx 2s$ (for well-connected validators). - Δ_{process} and Δ_{verify} scale with block gas usage.

Definition 4.3 (The Blockchain)

A **blockchain** \mathcal{C}_n is a sequence of blocks. Validity requires that for every block B_i : 1. **Integrity:** Parent hashes link correctly. 2. **State Validity:** $\delta(\sigma_{i-1}, \vec{\epsilon}_i) = \sigma_i$ (The events transition the state correctly). 3. **Temporal Authority:** The proposer was authorized.

Definition 4.3a (Slot-Skip Recovery for Liveness)

When the elected leader for slot s fails to produce a block within the slot duration δ_i (e.g., node offline, network partition), the protocol must **recover** to preserve liveness. Define the **current physical slot**:

$$s_{\text{phys}}(t) = \left\lfloor \frac{t - t_{\text{genesis}}}{\delta_i} \right\rfloor$$

When $s_{\text{phys}}(t) \geq |\mathcal{C}|$ (chain height), the leader for slot $s_{\text{phys}}(t)$ may produce the next block, even if $s_{\text{phys}}(t) > |\mathcal{C}|$ (missed slots). The block's timestamp must fall within slot $s_{\text{phys}}(t)$'s window. This ensures the chain progresses when nodes come and go.

Implementation (PoS2/TokenX): `can_produce_block` uses $\max(|\mathcal{C}|, s_{\text{phys}}(t))$ for leader election. Consensus derives slot from block timestamp when validating.

Definition 4.4 (Fork Choice Rule)

The **fork choice rule** \mathcal{F} selects the canonical chain from the set of observed chains \mathcal{G} :

$$\mathcal{F} : \mathcal{G} \rightarrow \mathcal{C}$$

where $\mathcal{C}^* = \mathcal{F}(\mathcal{G})$ is the longest chain (in PoW) or the heaviest justified chain (in PoS).

Examples: 1. **Nakamoto Consensus (Bitcoin):** - $\mathcal{F}(\mathcal{G})$ selects the chain with the most accumulated work:

$$\mathcal{C}^* = \arg \max_{\mathcal{C} \in \mathcal{G}} \sum_{B \in \mathcal{C}} \text{Work}(B)$$

- $\text{Work}(B) = \frac{2^{256}}{D}$ (where D is the difficulty target). 2. **Gasper (Ethereum PoS):** - $\mathcal{F}(\mathcal{G})$ selects the chain with the highest **justification score**, where: - A block is **justified** if $\geq \frac{2}{3}$ of validators attest to it. - A block is **finalized** if it is justified and its child is also justified. 3. **Avalanche:** - $\mathcal{F}(\mathcal{G})$ uses **repeated sub-sampling** to converge on a chain.

Intuition: The fork choice rule must: - Be **deterministic** (all nodes agree on \mathcal{C}^*). - Be **resilient to forks** (prevent long-range attacks). - Incentivize **honest behavior** (e.g., following the longest chain).

Definition 4.5 (Latency Arbitrage)

A validator v with a low-latency view of the network can order Events within \vec{e}_i to maximize MEV:

$$\vec{e}_i^* = \arg \max_{\vec{e} \subset \mathcal{M}_v} \text{MEV}(\vec{e}, \prec_v)$$

Intuition: This accounts for the fact that the *order* of Events (especially *Invoke* events) drastically affects the final state σ_i .

- Validators closer to high-value event sources (e.g., exchanges) see events earlier.
- They can **front-run** or **sandwich** transactions before other validators see them.
- This creates a **centralizing force** (validators with better network positions earn more).

Mitigations: 1. **Fair Sequencing (e.g., Flashbots):** - Transactions are ordered by a trusted relay rather than the proposer. 2. **Time-Based Ordering (e.g., FSS):** - Transactions are ordered by arrival time at a reference node. 3. **Commit-Reveal Schemes:** - Transactions are committed before being revealed to prevent front-running.

V. Layer 5: Cryptoeconomic Security

This layer defines the **economic incentives** that ensure validators behave honestly.

Definition 5.1 (Resource Commitment)

Each validator v commits an **external resource** $\alpha_v \in \mathbb{R}^+$ to participate in consensus: - **Proof-of-Work (Bitcoin)**: α_v is hashrate (computational power). - **Proof-of-Stake (Ethereum)**: α_v is staked ETH. - **Proof-of-Space (Chia)**: α_v is allocated disk space.

Intuition: - α_v determines v 's influence in consensus (e.g., probability of proposing a block). - The total resource commitment $\sum_v \alpha_v$ defines the **security budget** of the protocol.

Definition 5.2 (Slashing Function)

The **slashing function** \mathcal{S} imposes economic penalties for misbehavior:

$$\mathcal{S} : \mathcal{V} \times \Sigma \rightarrow \mathbb{R}^+$$

where $\mathcal{S}(v, \sigma)$ is the amount slashed from v for provable misbehavior.

Examples: 1. **Safety Violations:** Signing conflicting blocks. 2. **Availability Violations:** Failing to propose blocks. 3. **Invalid State Transitions:** Proposing a block B_i where $\delta(\sigma_{i-1}, \vec{e}_i) \neq \sigma_i$ (i.e., claiming an invalid state root).

(Concrete instantiations: Double-signing in Ethereum, invalid state transition in Cosmos, liveness faults in Tendermint.)

Intuition: Slashing ensures that: - Misbehavior is **costly** (deters attacks). - Honest validators are **rewarded** (via block rewards and fees).

Definition 5.3 (Incentive Compatibility with Network)

A protocol is **Network-Incentive Compatible** if for all rational validators v :

$$U_v(\Pi) \geq U_v(\Pi^*) + \text{negl}(\lambda)$$

where: - $U_v(\Pi)$ is v 's utility under the honest protocol Π . - $U_v(\Pi^*)$ is v 's utility under a deviating strategy Π^* that includes **network-level attacks** (e.g., eclipse attacks, latency arbitrage).

Intuition: This requires that: - The **cost of controlling high-centrality positions** in G_t (e.g., running a relay node near an exchange) exceeds the **MEV extracted** from latency advantages. - Validators have no incentive to **deviate from the protocol** even when they can exploit network asymmetries.

Example (Ethereum): - Running a relay node near an exchange costs \approx \$10,000/month. - The MEV extracted from front-running may be \approx \$5,000/month. - Thus, the protocol is **not** network-incentive compatible in this case. - **Solution:** Use **MEV-boost** to auction block space, ensuring that MEV is distributed fairly.

VI. System Properties

These properties define the **security and performance** of the blockchain system.

Definition 6.1 (Safety)

For any two honest nodes u, v with chains $\mathcal{C}_u, \mathcal{C}_v$ at times $t_u, t_v > \text{GST}$ (Global Stabilization Time):

$$\mathcal{C}_u \preceq \mathcal{C}_v \vee \mathcal{C}_v \preceq \mathcal{C}_u$$

where \preceq denotes the **prefix relation** (i.e., one chain is a prefix of the other).

Intuition: - **No conflicting finalized blocks:** Honest nodes never disagree on the canonical chain. -

Common prefix property: If two nodes see different chains, one is a prefix of the other.

Example (Bitcoin): - Safety is **probabilistic:** The probability of a fork decreases exponentially with the number of confirmations. - After 6 confirmations, the probability of a reorg is negligible.

Definition 6.2 (Liveness)

For any valid transaction τ submitted to at least one honest node $s \in \mathcal{C}_t$ at time t , and assuming \mathcal{A} does not control the min-cut between s and $\mathcal{V}_{\text{honest}}$:

$$\exists i : \tau \in B_i \in \mathcal{C}(t+T) \text{ where } T \leq \frac{D_t \cdot k}{1-f}$$

where: - k is the expected number of epochs for inclusion. - f is the adversarial stake fraction.

Intuition: - **Transactions are eventually included** in the blockchain. - The **inclusion time** T depends on: - Network latency (D_t). - Adversarial power (f). - Protocol parameters (k).

Example (Ethereum): - $D_t \approx 2s$. - $k \approx 2$ (transactions are usually included within 2 epochs). - $f < \frac{1}{3}$.
- Thus, $T \approx 6s$ (in practice, often faster due to parallel block proposals).

Definition 6.3 (Censorship Resistance)

A protocol is $(\epsilon, \Delta_{\max})$ -**Censorship Resistant** if for any transaction τ with fee $\geq f_{\min}$, the probability that τ remains unconfirmed after time Δ_{\max} is $\leq \epsilon$, provided \mathcal{A} does not control a vertex cut separating the originator from $\frac{2}{3}$ of \mathcal{V} by weight.

Intuition: - **Transactions with sufficient fees are eventually included.** - The adversary cannot **permanently censor** transactions unless they control a large fraction of the network.

Example (Ethereum): - $\epsilon = 0.01$ (1% chance of censorship after Δ_{\max}). - $\Delta_{\max} = 120s$ (2 minutes). - $f_{\min} = 1$ gwei (minimum fee to avoid censorship).

Mitigations for Censorship: 1. **Inclusion Lists (Ethereum):** - Proposers must include transactions from a predefined list. 2. **Decentralized Relayers:** - Transactions are submitted to multiple relayers to avoid single points of censorship. 3. **Time-Based Inclusion:** - Transactions are included after a fixed delay, regardless of fees.

VII. Critical Theorems

These theorems formalize the **interdependence** between network topology, consensus, and security.

Theorem 7.1 (Network-Safety Interdependence)

In a partially synchronous BFT system tolerating $f < \frac{1}{3}$ Byzantine validators, if the underlying graph G_t has **vertex connectivity** $\kappa(G_t) \leq 3f$, then \mathcal{A} can partition the network to create conflicting finalized blocks without violating local BFT assumptions, thus breaking Safety (Definition 6.1).

Proof Sketch: 1. If $\kappa(G_t) \leq 3f$, the adversary can partition \mathcal{V} into three sets V_1, V_2, F where: - $|F| \leq 3f$ (adversarial nodes). - F controls all paths between V_1 and V_2 . 2. The adversary delivers B_i to V_1 and B'_i to V_2 , where B_i and B'_i are conflicting blocks. 3. Each partition sees only one block and finalizes it locally. 4. When the partition heals, the network has two conflicting finalized blocks, violating safety.

Corollary 7.2: For BFT consensus, a **necessary condition for safety** is:

$$\kappa(G_t) > 3f$$

Intuition: - **High connectivity** is required to prevent partitions. - In practice, this means: - Validators must have **many diverse peers** (e.g., geographically distributed). - The network must be **resilient to node failures**.

Theorem 7.3 (Propagation-Liveness Trade-off)

Under the Network-Adjusted Stability Condition (Definition 4.2), if the gossip protocol \mathcal{G} has propagation complexity $O(|N_t| \cdot \log |N_t|)$, then liveness (Definition 6.2) holds with $T \in O(\delta_i \cdot \log |\mathcal{V}|)$.

Proof Sketch: 1. The gossip protocol ensures that a transaction τ propagates to all honest validators in $O(\log |\mathcal{V}|)$ rounds. 2. Each round takes $O(D_t)$ time (block propagation diameter). 3. Thus, τ is included in a block within $O(\delta_i \cdot \log |\mathcal{V}|)$ time.

Intuition: - **Efficient gossip** ensures fast transaction propagation. - **Liveness depends on:** - The **scalability** of the gossip protocol. - The **epoch duration** δ_i .

Appendix: Complete Instantiation Examples

Parameter	Bitcoin (PoW)	Ethereum PoS	Solana	Avalanche
\mathcal{O}	$\mathcal{H}(B) < D$	RANDAO + VRF	PoH (VDF-sequence)	Sub-sampling
Ψ	2016 blocks retarget	Fixed 12s	Adaptive (400ms target)	Async (no slots)
G_t	Random graph (8 conn)	Discv5 DHT + mesh	Turbine (tree)	Kademlia
\mathcal{G}	Erlay	Gossipsub 1.1	Turbine (FEC)	Gossip
Safety	Probabilistic	$\kappa > \frac{1}{3} \mathcal{V} $	Leader rotation	Snowball finality
Censorship Resistance	Topology-agnostic	Relayer-dependent (MEV-boost)	Leader-dominant	Subnet-dependent

Summary of Changes (v2.0)

1. **Terminology:** “Transaction” (τ) is generalized to “**Event**” (ϵ).
2. **State Definition** (Σ): Explicitly split into **Data** (\mathcal{D}) and **Programs** (Ω) to formally recognize smart contracts as part of the state.
3. **Transition Function** (δ): Expanded to handle three types of events (Transfer, Deploy, Invoke), including the execution of stored bytecode.
4. **Block Validity:** Updated to require re-execution of the event sequence to verify state commitments.

Summary of Dependencies

The layers of a blockchain system are **interdependent**: 1. **Layer 1 (State)** defines the application logic and state transitions (δ). 2. **Layer 2 (Time)** defines epochs (ε_i) and proposer selection (\mathcal{O}). 3. **Layer 3 (Network)** defines how events and blocks propagate (G_t, \mathcal{M}_u). 4. **Layer 4 (Consensus)** defines block structure (B_i) and fork choice (\mathcal{F}), constrained by network latency (D_t). 5. **Layer 5 (Economics)** defines incentives (\mathcal{S}), which must account for network-level attacks (e.g., latency arbitrage).

Key Insight: The **Network Topology layer (Layer 3)** is the **physical constraint** that binds the logical consensus (Layer 4) to real-world latency and geographic distribution. Ignoring network effects leads to **security vulnerabilities** (e.g., eclipse attacks, MEV extraction) or **performance bottlenecks** (e.g., slow block propagation).