AIONET Finality Comparison & Proof of Memory (PoM) Consensus Model

AIONET Research

August 18, 2025

Overview

Finality is the time after which a transaction is practically irreversible and confirmed by the network. Traditional blockchains rely on energy (PoW) or economic stake (PoS), creating latency and complex game dynamics. **AIONET** introduces **Proof of Memory (PoM)**—a real-time, bandwidth-driven consensus layer—with **Proof of Drift (PoD)** providing continuous entropy/drift scoring for anomaly detection.

Model: Finality Decomposition

We separate finality into a compute/bandwidth component and a network/coordination component, plus a small PoD scoring overhead:

$$T_{\text{final}} \approx \max\left(T_{\text{compute}} + \delta_{\text{PoD}},\, T_{\text{network}}\right). \tag{1}$$

Compute/bandwidth term. Let

- B = sustained memory bandwidth per validator (bytes/s),
- N = number of concurrent memory lanes/channels,
- P = parallelization factor (vector width / HBM bit-level concurrency),
- $\eta \in (0, 1]$ = effective utilization (contention, cache/micro-ops, scheduler efficiency),
- $V \in (0, 1]$ = validator health factor from AI scoring,
- D = bytes of data that must be observed/checked for a block (or decision window).

Then the effective memory throughput for validation is

$$\Theta_{\text{mem}} = B \cdot N \cdot P \cdot \eta \cdot V,$$

and the compute-bound time is

$$T_{\text{compute}} = \frac{D}{\Theta_{\text{mem}}}.$$

Network/coordination term. Let

- R = number of confirmation rounds/commit steps,
- RTT = median round-trip network latency among validators,
- Δ = clock skew / propagation slack.

Then

$$T_{\text{network}} = R \cdot \text{RTT} + \Delta.$$

PoD overhead. PoD entropy/drift scoring is typically lightweight; we model it as an additive δ_{PoD} (tunable, implementation-dependent).

Throughput view. For a transaction size (or validation budget) of d_{tx} bytes per transaction, an upper bound on throughput is:

$$\text{TPS}_{\max} \approx \min\left(\frac{\Theta_{\text{mem}}}{d_{\text{tx}}}, \frac{\kappa}{R \cdot \text{RTT} + \Delta}\right),$$

where κ represents the number of transactions the protocol can commit per round (committee parallelism and batching limits).

Notation (summary)

Symbol	Meaning
B	Sustained memory bandwidth per validator (bytes/s)
N	Concurrent memory lanes/channels
P	Parallelization factor (vector/bit-level concurrency)
η	Utilization/efficiency (0–1)
V	Validator health factor from PoD/AI scoring (0–1)
D	Bytes to be validated per decision window/block
R	Commit rounds
RTT	Median round-trip network latency
Δ	Slack: propagation & clock skew
δ_{PoD}	PoD overhead

Illustrative Comparison (High-Level)

Network	Consensus	Block time	Finality (typical)	Limitation signal
Bitcoin	PoW	∼10 min	∼60 min (6 blocks)	energy/latency tradeoff
Ethereum	PoS	∼12 s	∼60 s to 900 s	complex validator game theory
AIONET (target)	PoM + PoD	<1 s	∼1 s to 2 s	bounded by bandwidth + network

HBM Scaling (Illustrative)

HBM Gen	Bandwidth/stack	Channels (N)	Est. PoM Finality (compute-bound)
HBM3e	1.2 TB/s to 1.4 TB/s	8–16	∼1.5 s to 2.0 s
HBM4	2 TB/s to 3 TB/s	16–32	∼0.8 s to 1.2 s
HBM8*	≥4 TB/s (proj.)	32–64+	$<$ 0.3 s (theoretical)

Notes. Values are illustrative and depend on D, η, V, and board/SoC limits. Network term can dominate in poor connectivity; finality is max of compute vs network per Eq. (1).

Sensitivity (Example)

For a fixed D, halving RTT or doubling η produces similar first-order gains. Example parameters:

$$\eta = 0.6,\ V = 0.9,\ R = 2,\ \mathrm{RTT} = 150\,\mathrm{ms},\ \Delta = 50\,\mathrm{ms},\ \delta_{\mathrm{PoD}} = 20\,\mathrm{ms}.$$

On HBM4 ($B{=}2.5$ TB/s$, N{=}24, P{=}1$), if D requires 1 GB of observed memory behavior, then

$$T_{\mathrm{compute}} \approx \frac{1 \times 10^{9}\,\mathrm{B}}{2.5{\times}10^{12} \cdot 24 \cdot 1 \cdot 0.6 \cdot 0.9} \approx 0.03\,\mathrm{s}.$$

Network side: $T_{\mathrm{network}} = 2 \times 0.15 + 0.05 = 0.35\,\mathrm{s}$. Hence $T_{\mathrm{final}} \approx \max(0.03 + 0.02,\ 0.35) = 0.35\,\mathrm{s}$.

Methodology & Limitations

This simplified model abstracts protocol details and treats PoM as bandwidth-bounded observation with PoD-based validator health. Actual deployments depend on committee sizes, batching, leader selection, adversarial behavior, and implementation factors (*e.g.*, NUMA, DMA engines, cache coherence). Numbers in the tables are indicative targets; **testnet metrics** will supersede estimates. See the accompanying *Finality Analysis PDF* for edge cases, parameter sweeps, and reproducibility instructions.

References

- KAIST HBM roadmaps (capacity/bandwidth trends).
- AIONET Finality Analysis (detailed derivations & sims).