AIONET Science & Vision

Measuring Trust with Entropy and Memory Behavior

AIONET Research

August 18, 2025

Overview

AIONET is an AI-native blockchain that replaces energy or stake with *observable behavior*. Validators operate under **Proof of Memory (PoM)**, where real-time memory bandwidth and access patterns bound validation throughput, and **Proof of Drift (PoD)**, which continuously scores entropy and signal drift to detect anomalies. The result is fast finality with physics-anchored auditability.

Core Ideas

PoM (Proof of Memory). Validation is treated as an observation task over memory behavior. Let B be sustained bandwidth per validator (bytes/s), N memory channels, P a parallelization factor, $\eta \in (0, 1]$ utilization, $V \in (0, 1]$ health (from PoD), and D bytes to observe per decision window. Effective throughput is

$$\Theta_{\text{mem}} = B \cdot N \cdot P \cdot \eta \cdot V, \quad T_{\text{compute}} = \frac{D}{\Theta_{\text{mem}}}.$$

PoD (Proof of Drift). A lightweight scoring pipeline estimates validator health V by tracking entropy and drift in time-series features (e.g., access timing, refresh jitter, bank/row activity). A simple formulation:

$$V = \sigma\big(w_0 + \sum_k w_k \, \phi_k(\Delta s_k, H_k, \text{SNR}_k)\big), \quad V \in (0, 1],$$

where ϕ_k are calibrated features (entropy H_k, drift Δs_k, signal-to-noise SNR_k), and σ is a squashing function. Low V can quarantine or de-weight a validator.

Finality at a Glance

We decompose finality into compute/bandwidth and network/coordination, with a small PoD overhead δ_{PoD}:

$$T_{\text{final}} \approx \max\big(T_{\text{compute}} + \delta_{\text{PoD}},\ T_{\text{network}}\big), \qquad T_{\text{network}} = R \cdot \text{RTT} + \Delta, \tag{1}$$

where R is the number of commit rounds, RTT median round-trip latency, and Δ slack for propagation/clock skew.

Why It Scales

- **Physics-bounded:** As HBM bandwidth grows, Θ_{mem} increases; compute-bound time shrinks.
- **Security via behavior:** PoD ties validator weight to stable, reproducible signals; drift triggers investigation or de-weighting.
- **Composable:** PoM/PoD can sit under standard mempool/consensus plumbing with committee parallelism and batching.

HBM Trend (Illustrative)

HBM Gen	Bandwidth/stack	Channels (N)	PoM compute time (indicative)
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)

Method Highlights

- Effective throughput: $\Theta_{\text{mem}} = BNP\eta V$.
- Finality: $T_{\text{final}} \approx \max(T_{\text{compute}} + \delta_{\text{PoD}}, T_{\text{network}})$ (Eq. 1).
- Throughput ceiling: $\text{TPS}_{\max} \approx \min\big(\Theta_{\text{mem}}/d_{\text{tx}},\ \kappa/(R \cdot \text{RTT} + \Delta)\big)$ for tx budget d_{tx} and per-round capacity κ.

Reproducibility & Open Questions

- Benchmarks across HBM generations; parameter sweeps for η, V, R, RTT.
- Robustness of PoD features under load, temperature, and adversarial patterns.
- Committee sizing, batching, and leader selection impact on R and κ.

Further Reading

- **Finality Analysis (PDF):** full derivations, sensitivity, and edge cases.
  https://aionet.tech/Full_Finality_Analysis.pdf
- **KAIST HBM Roadmap (PDF):** bandwidth/capacity trends.
  https://aionet.tech/KAIST_HBM_Roadmap.pdf
- **Researcher View (web):** tabs with artifacts and links.
  https://aionet.tech/

Disclaimer. Figures are indicative targets; real-world results depend on network conditions, validator set size, implementation details, and adversarial behavior. Testnet data will supersede estimates.