Multi-Agent Network Management with Dynamic Entropy-Driven Logistic Trust Aggregation

Minhyeok Jang, *Student Member, IEEE,* Jalel Ben-Othman, *Senior Member, IEEE,*
Hyunchae Chun, *Member, IEEE,* Sungrae Cho, *Senior Member, IEEE,* Hyunbum Kim, *Senior Member, IEEE*

***Abstract*—Autonomous network management increasingly fuses multiple heterogeneous detectors—such as the intrusion detectors that monitor different traffic planes for 6G and IoT security—through adaptive trust-weighted consensus. When trust is updated online, however, such systems face a fundamental stability-agility trade-off: they are either calm but slow to react to novel threats, or fast but erratic under routine noise. We identify and formalize the resulting failure modes of trust collapse and blind conformity, and propose DELTA (Dynamic Entropy-driven Logistic Trust Aggregation), a self-regulating trust-management framework. DELTA couples a Fixed-Share Redistribution regularizer, which guarantees a minimum trust quota for every detector, with an entropy-amplified logistic controller whose learning rate is driven by the current leader's error rate and amplified by the ensemble's structural entropy; this keeps the system quiescent under normal traffic yet triggers a rapid, bounded re-calibration the moment the trusted detector begins to fail. We prove that DELTA enforces a strictly positive diversity floor—making trust collapse provably impossible—and derive bounds on its transition latency and stationary volatility. Across an extensive evaluation—including robustness to delayed, missing, and adversarial feedback, comparison against expert-advice, Bayesian, and change-point baselines with confidence intervals, and validation on the real UNSW-NB15 intrusion dataset—DELTA recovers from zero-day regime shifts where naive baselines collapse below chance, while remaining an order of magnitude more stable than aggressive adaptive methods, all at $O(N)$ computational and communication cost.**

***Index Terms*—network management, distributed intrusion detection, multi-agent trust aggregation, concept drift, stability-agility trade-off, entropy-driven adaptation.**

I. Introduction

NETWORK management for emerging 6G and IoT infrastructures increasingly relies on autonomous, AI-driven decision-making to sustain security and service assurance under evolving conditions [1], [2], [3]. A central task is *distributed intrusion and anomaly detection*, in which heterogeneous detectors—each monitoring a different observation plane such as flow statistics, payload content, or connection state—must be fused into a single operational security decision [5]. Trust-weighted aggregation scales each detector's contribution by an adaptively learned reliability score, so the management system leans on whichever detector is currently most credible. Such architectures enhance robustness through complementary expertise, yet the stability of adaptive trust-weighted consensus under non-stationary conditions remains insufficiently understood [6], [7]. When a zero-day attack or a traffic-distribution shift abruptly alters which detectors remain reliable, online trust updating faces a fundamental *stability-agility trade-off*: cumulative models are stable under stationary noise but incur a prolonged *phase transition latency* (L_{trans}) and fail to react to novel threats, whereas high-learning-rate models adapt quickly but oscillate under normal traffic, causing costly operational churn and false reconfigurations [10], [11]. Managing this trade-off is the core network-management problem addressed in this work.

Minhyeok Jang is with Department of Embedded Systems Engineering, Incheon National University, Incheon, 22012, South Korea. E-mail: jmhyeok0827@inu.ac.kr
Jalel Ben-Othman is with LISSI Lab, University of Paris-Est, Créteil, France. E-mail: jalel.ben-othman@u-pec.fr
Hyunchae Chun is with Department of Information and Telecommunication Engineering, Incheon National University, Incheon, 22012, South Korea. E-mail: hyunchae.chun@inu.ac.kr
Sungrae Cho are with the School of Computer Science and Engineering, Chung-Ang University, Seoul, 06974, South Korea. E-mail: srcho@cau.ac.kr
Hyunbum Kim is Corresponding Author and is with the Department of Embedded Systems Engineering, Incheon National University, Incheon, 22012, South Korea. E-mail: hyunbumkim@ieee.org

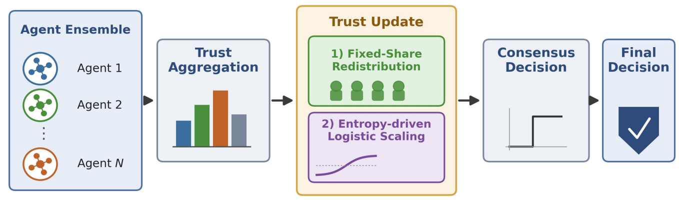

Fig. 1. Conceptual comparison of trust-weight dynamics under a zero-day adversarial shift.

Fig. 1 illustrates the motivation. In Scenario 1, traditional aggregation succumbs to an irreversible monopoly under prolonged stationary conditions; when a zero-day attack then triggers an inverse failure in the dominant agent, the system remains trapped in its prior state because minority influence has been depleted, causing catastrophic decision failure. To resolve this we propose DELTA (Dynamic Entropy-driven Logistic Trust Aggregation), conceptualized in Scenario 2. DELTA replaces heuristic penalty rules with two symbiotic mechanisms: Fixed-Share Redistribution, which guarantees every agent a minimum trust quota and thereby preserves latent structural diversity, and Entropy-driven Logistic Scaling, which modulates the learning rate (α) through a logistic S-curve driven by the trusted agent's error rate and amplified by the ensemble's structural entropy. DELTA thus stays quiescent during equilibrium to suppress noise, yet bursts into a rapid leadership crossover the moment the trusted detector begins to fail.

Based on the above observations, the main contributions of this paper are summarized as follows:

- We formally define and quantify trust collapse and blind conformity in adaptive trust-weighted consensus, and identify the underlying stability-agility trade-off that prevents traditional adaptation rules from being simultaneously calm under stationary noise and agile under regime shifts.
- We propose DELTA (Dynamic Entropy-driven Logistic Trust Aggregation), a self-regulating trust-management framework that couples a Fixed-Share Redistribution regularizer with an entropy-amplified logistic controller. The learning rate is driven by a leading regime-shift signal—the current leader's error rate—amplified by the ensemble's structural entropy, so that the system stays quiescent under routine noise yet bursts into rapid recalibration the moment the trusted detector begins to fail.
- We provide formal guarantees for DELTA, including a strictly positive diversity floor that makes trust collapse provably impossible, a transition-latency bound relating recovery speed to the redistribution and learning-rate parameters, a stationary-stability bound explaining its low volatility, and a robustness guarantee under noisy or delayed feedback.
- We conduct an extensive, statistically rigorous evaluation. Beyond the canonical regime-switching scenario, we stress-test DELTA across ensemble size, recurring and gradual drift, delayed and missing labels, Byzantine and correlated agents, and communication failures; we compare against strong expert-advice, Bayesian, change-point, and adaptive-ensemble baselines with confidence intervals and significance tests; and we validate the framework on real network intrusion data (UNSW-NB15), together with an $O(N)$ computational-complexity analysis for large-scale deployment.

The remainder of this paper is organized as follows. Section III formalizes the distributed trust-aggregation problem and the failure modes of conventional updating. Section IV presents the DELTA framework, its complete algorithm, and its formal guarantees. Section V reports the experimental evaluation, including robustness, real-data validation, and complexity analyses. Section VI concludes the paper.

II. Related Work

AI-driven network management. A growing body of work applies learning-based adaptation to cope with non-stationary network conditions. Deep reinforcement learning has been used for traffic-aware offloading and network slicing in mobile edge computing [12], serverless federated learning for multi-objective resource optimization in UAV-enabled heterogeneous networks [13], and graph-attentive multi-agent reinforcement learning for SLA-aware task offloading and service orchestration [14]. These approaches optimize *resource allocation or control policies*. Our work is complementary and operates at a different layer: rather than learning an allocation policy, DELTA provides a lightweight trust-management layer that fuses heterogeneous detectors' decisions and governs how fast and how stably trust is re-allocated when the environment shifts.

Online learning and concept drift. DELTA's components build on established paradigms. The Fixed-Share Redistribution follows the fixed-share ("tracking the best expert") family [16], and softmax trust updating is closely related to multiplicative-weights / Hedge aggregation [15]. Non-stationary learning has also been addressed by explicit change-point detection [17], adaptive ensembles for concept drift [18], [19], and adaptive forgetting-factor control in online multiplicative updates [8]. While each of these ingredients—fixed-share weighting, multiplicative aggregation, change-point resetting, and entropy monitoring—is individually well established, their combination in DELTA yields a property none provides in isolation: a *stability-optimal* adaptive aggregator that attains the agility of aggressive expert-advice methods at substantially lower stationary volatility. We include these methods as baselines in Section V so that the comparison is quantitative rather than merely conceptual.

Trust and reputation in multi-agent systems. Trust management has been studied extensively in multi-agent systems [10], [11], including robust consensus under adversarial or zero-trust conditions [5], [4], while Shannon and generalized entropies have served as diagnostics of distributional diversity [9]. DELTA differs by using structural entropy not as a passive diagnostic but as an active amplifier within a closed-loop controller, coupled with a provable diversity floor.

III. Overview of Proposed System

This section formalizes the multi-agent consensus setting and the conditions under which conventional aggregation fails, characterizing the non-stationary environment and the stability-agility trade-off—the conflict between stationary reliability and rapid recovery—that motivates DELTA.

A. System Overview

We consider the trust-management layer of a distributed intrusion-detection system for non-stationary network traffic. Each agent is a detector specialized in one observation plane of the traffic—for example flow-level statistics, connection state, or payload content—and, for each observed flow, outputs a probabilistic estimate of whether it is malicious. The network-management system fuses these estimates by trust-weighted aggregation, leaning on whichever detectors are currently most reliable, and issues an operational decision (e.g., allow, block, or escalate). The binary decision below is thus the per-flow benign/malicious classification, and the "agents" are heterogeneous detectors rather than abstract predictors.

Let $\mathcal{A} = \{a_1, \ldots, a_N\}$ denote the agents. At step t each agent a_i outputs $p_i(t) \in [0, 1]$, the predicted probability of the positive class, and the system maintains a trust weight vector $\mathbf{w}(t) = [w_1(t), \ldots, w_N(t)]$ subject to $\sum_i w_i(t) = 1$ and $w_i(t) \geq 0$. The aggregated consensus probability is

$$P(t) = \sum_{i=1}^{N} w_i(t) p_i(t), \tag{1}$$

and the system decision is obtained by thresholding, $\hat{y}(t) = 1$ if $P(t) \geq \tau$ and 0 otherwise, where τ is the decision threshold (typically 0.5, tunable to balance false positives and negatives).

In conventional adaptive systems, trust weights are updated by consensus-reinforced heuristics: agents agreeing with the collective decision receive unbounded positive reinforcement. As we demonstrate, this cumulative reinforcement triggers structural trust collapse under regime shifts, as the system loses its capacity to re-allocate trust when the historically dominant agent fails.

To resolve this, DELTA conceptualizes the trust distribution as a resource-constrained allocation governed by two continuous forces—Fixed-Share Redistribution, a structural regularizer that preserves latent diversity, and Entropy-driven Logistic Scaling, a dynamic adaptation mechanism—which are detailed in Section IV.

B. System Settings and Assumptions

To analyze the stability of adaptive trust aggregation and the emergence of structural failure modes, we define a non-stationary decision environment where agent reliability is state-dependent. The following assumptions establish the formal constraints of our model.

Assumption 1 (Heterogeneous Specialization). The system consists of N heterogeneous agents, each specialized in a distinct feature space (e.g., semantic content vs. structural URL verification), so that each agent's accuracy is tied to specific data characteristics and no single agent is globally optimal.

Assumption 2 (Regime-Switching Environment). The environment follows a stochastic regime-switching process with state $s(t) \in \{0, 1\}$, where $s(t) = 0$ is a stationary regime and $s(t) = 1$ a non-stationary regime (zero-day attack). Transitions occur discontinuously, modelling the sudden deployment of novel adversarial tactics.

Assumption 3 (Regime-Dependent Reliability and Inverse Failure). We assume a worst-case zero-day scenario in which an agent that is highly reliable under $s(t) = 0$ degrades sharply under $s(t) = 1$. Consequently the optimal trust distribution $\mathbf{w}^*(t)$ is non-stationary and must re-calibrate to a different agent during transitions.

Assumption 4 (Competitive Trust Conservation). Trust weights are a conserved resource with $\sum_i w_i(t) = 1$. Updates to the underlying potentials are structurally mapped to induce a zero-sum redistribution in weight space, so that adaptation remains bounded while being dynamically accelerated, without violating conservation.

C. Problem Definition

Under conventional trust updating, the weight distribution frequently suffers an irreversible concentration on a single agent. We formalize this failure as *trust collapse*: a state in which the system's adaptive capacity is neutralized by its own reinforcement history. To quantify the diversity of influence, we use the normalized Shannon entropy of $\mathbf{w}(t)$ as a diagnostic of latent resilience:

$$H_{\text{norm}}(t) = \frac{-\sum_{i=1}^{N} w_i(t) \ln w_i(t)}{\ln N}, \tag{2}$$

with the convention $0 \ln 0 = 0$. Here $H_{\text{norm}}(t) \in [0, 1]$: values near 1 indicate a balanced distribution capable of rapid re-calibration, while $H_{norm} \to 0$ reflects a severe monopoly inducing structural rigidity. Letting $i^*(t) = \arg\max_i w_i(t)$ denote the dominant agent, we identify trust collapse as a depletion of structural diversity beyond a critical threshold.

Definition 1 (Trust Collapse). A multi-agent system enters a state of trust collapse at time step t_c if it simultaneously satisfies both *dominance persistence* and *diversity contraction*:

1) *Dominance Persistence*: The dominant agent remains unchanged over a persistence window of W_p discrete steps, i.e., $i^*(t) = i^*(t_c)$ for all $t \in \{t_c, t_c + 1, \ldots, t_c + W_p\}$.
2) *Diversity Contraction*: The normalized entropy falls below a diversity threshold ε, i.e., $H_{\text{norm}}(t) < \varepsilon$ over that window, indicating a near-complete absence of the minority influence required to trigger a weight reversal.

Here W_p and ε are the two parameters of the collapse criterion. We show in Section IV-C3 (Theorem 1) that DELTA's Fixed-Share mechanism enforces a strictly positive entropy floor $H_{\text{norm}}(t) \geq H_{\min}(\gamma, N)$, so for any $\varepsilon < H_{\min}(\gamma, N)$ the diversity-contraction condition is *provably unreachable*—i.e., DELTA cannot enter trust collapse.

Trust collapse underlies the stability-agility trade-off: a collapsed system appears exceptionally stable under stationary conditions, yet its Phase Transition Latency (L_{trans}) becomes unbounded during regime shifts, lacking the structural tension required to re-calibrate.

Definition 2 (Blind Conformity). *Blind conformity* is the behavioral manifestation of trust collapse during a regime shift ($s(t) = 0 \to 1$): the system suppresses correct signals from minority agents in favor of the failing dominant agent, causing a catastrophic accuracy drop because consensus prioritizes historical reliability over real-time evidence.

IV. Proposed Schemes

This section presents the formulation and algorithmic details of the proposed Dynamic Entropy-driven Logistic Trust Aggregation (DELTA) framework. To resolve the structural vulnerabilities of *trust collapse* and *blind conformity*, DELTA treats trust not as a cumulative score but as an implicitly normalized competitive dynamic in which individual gains are bounded by the ensemble's collective state.

Throughout this section we retain the network-management reading of Sec. III-A: each agent is a traffic detector, its potential V_i is the evidence accumulated for that detector's reliability, its trust weight w_i is the operational confidence placed in it, the consensus $P(t)$ is the fused threat score driving the allow/block/escalate decision, and a regime shift is the onset of a novel attack. The two mechanisms below thus address two concrete management requirements: Fixed-Share Redistribution guarantees that no detection plane is ever fully ignored, preserving coverage against unforeseen attacks, while Entropy-driven Logistic Scaling governs how quickly the system re-trusts a different detector once the trusted one begins to fail, bounding the window of exposure.

A. Background

1) Limitations of Conventional Aggregation: The Stability-Agility Trade-off: In traditional adaptive consensus systems, reliability potentials are updated by a linear reinforcement rule

$$V_i(t+1) = V_i(t) + \alpha \cdot r_i(t) \tag{3}$$

where V_i is the reliability potential, α a fixed learning rate, and $r_i(t) \in \{1, -1\}$ the categorical reward. This imposes the stability-agility trade-off: a low α suppresses stochastic noise but prolongs the *phase transition latency* (over 200 steps for conventional baselines, Section V), whereas a high α recovers quickly at the cost of erratic stationary fluctuations. Lacking any non-linear modulation of α, such models succumb to trust collapse as historical reinforcement locks the system into a rigid, non-adaptive state.

2) Logistic Dynamics as a Non-linear Control Signal: Drawing on the logistic growth model in population ecology, whose Verhulst equation $\frac{dP}{dt} = rP(1 - P/K)$ saturates via a self-regulating brake $(1 - P/K)$, we pivot from population size to adaptation velocity. Adopting a logistic S-curve to scale the learning rate $\alpha(t)$ yields two operational regimes: a low-energy *stable state* for stationary stability, and a high-energy *active state* that bursts trust redistribution during rapid phase transitions.

3) Structural Entropy as a System Health Monitor: Shannon entropy $H(t) = -\sum w_i \ln w_i$ quantifies distributional diversity; for a trust-weighted ensemble it directly indicates structural health, with $H_{norm} = 1$ denoting full latent diversity and $H_{norm} \to 0$ signalling imminent trust collapse. DELTA uses normalized entropy in its logistic scaling as a structural amplifier: once the trusted leader fails, the acceleration of trust redistribution scales with how collapse-prone (low-entropy) the ensemble has become.

B. DELTA Core Mechanisms: Structural Regularization and Dynamic Scaling

DELTA resolves the stability-agility trade-off through two symbiotic mechanisms that use the ensemble's real-time structural health as a feedback signal.

1) Fixed-Share Redistribution: In standard softmax aggregation, as an agent's potential V_i grows through historical reinforcement its weight approaches 1.0, eliminating minority influence. To prevent this irreversible *trust collapse*, DELTA introduces a redistribution parameter $\gamma \in (0,1)$ that diverts a fixed portion of the total trust and redistributes it equally across the ensemble:

$$w_i(t) = (1-\gamma)\frac{\exp(V_i(t))}{\sum_{j=1}^{N}\exp(V_j(t))} + \frac{\gamma}{N} \tag{4}$$

This mechanism guarantees a minimum trust quota for every agent, ensuring $w_i(t) \geq \gamma/N$. Mathematically, this prevents the normalized structural entropy $H_{\text{norm}}(t)$ from reaching zero, preserving a "seed" of latent diversity that allows the system to remain responsive even after prolonged stationary periods.

Fig. 2 illustrates why both components are needed. Without the floor, entropy collapses toward zero and latent diversity is lost, matching the "no Fixed-Share" ablation in Sec. V-D;

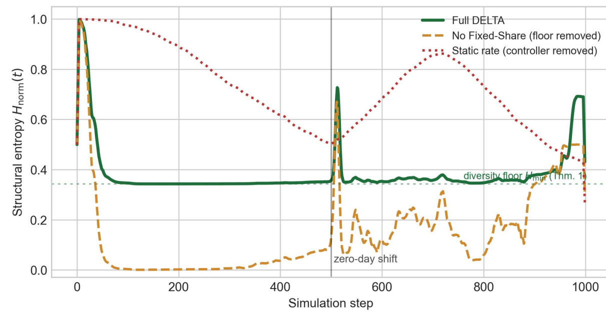

Fig. 2. Roles of the two components, shown through the structural-entropy trajectory (see text).

without the dynamic controller, the ensemble drifts and cannot re-organize around the shift. Full DELTA instead stays within a controlled band, bounded below by the diversity floor $H_{\min}$ of Theorem 1—concentrated enough to be accurate, yet diverse enough to recover.

2) Entropy-driven Logistic Scaling: While Fixed-Share preserves diversity, the system must still resolve the challenge of adaptation speed. To minimize the phase transition latency (L_{trans}) without inducing stationary instability, we model the learning rate $\alpha(t)$ as a dynamic variable governed by a logistic S-curve driven by two complementary signals: a fast *regime-shift trigger* and a structural *amplifier*.

Trigger—leader error rate. A reliable early indicator of an adversarial regime shift is that the currently trusted agent begins to fail. Let $i^\dagger(t) = \arg\max_i w_i(t)$ denote the current leader. We define its recent empirical error rate over a sliding window of W labeled steps as

$$E(t) = \frac{1}{W}\sum_{\tau=t-W+1}^{t} \mathbb{I}[\hat{y}_{i^\dagger}(\tau) \neq y(\tau)], \tag{5}$$

where $\hat{y}_{i^\dagger}(\tau)$ is the leader's decision and $y(\tau)$ the ground truth. Crucially, unlike the structural entropy of the trust weights—which, being computed from the *previous* weight distribution, lags the onset of prediction divergence—$E(t)$ responds immediately when the dominant agent starts being contradicted by ground truth. This resolves the feedback-timing gap inherent to a purely entropy-triggered controller.

Amplifier—structural entropy. Normalized Shannon entropy quantifies the global diversity of the trust distribution: a low $H_{\text{norm}}(t)$ signals a concentrated, monopoly-prone state from which recovery is intrinsically harder. DELTA therefore retains entropy in the control law as a *collapse-amplification* factor $(1 - H_{\text{norm}}(t))$: the more concentrated the ensemble, the stronger the corrective response to a given leader failure. Unlike variance-based measures or divergence metrics that require an external reference, entropy provides an absolute, bounded measure of systemic monopoly versus latent diversity.

Combining the two, DELTA governs the dynamic learning rate through a hybrid stress signal $S(t)$:

$$S(t) = E(t)\,\bigl(1 + \beta\,(1 - H_{\text{norm}}(t))\bigr), \tag{6}$$

$$\alpha(t) = \alpha_{\text{base}} + \frac{\alpha_{\max} - \alpha_{\text{base}}}{1 + \exp\bigl(-k\,(S(t) - \theta)\bigr)}, \tag{7}$$

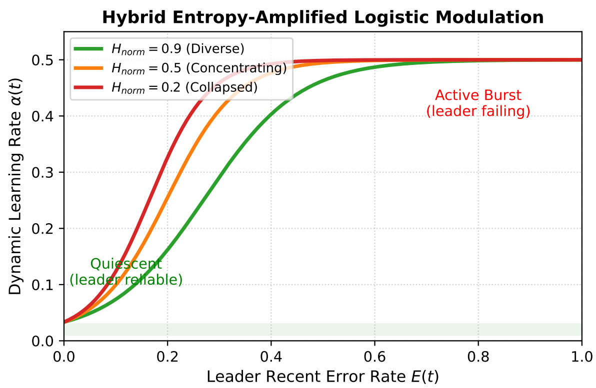

Fig. 3. Hybrid logistic modulation: $\alpha(t)$ versus leader error rate $E(t)$ at several entropy levels (see text).

where α_{base} and $\alpha_{\max}$ are the baseline and peak adaptation velocities, β is the collapse-amplification gain, k is the transition steepness, and θ is the activation threshold on the stress signal. The structural entropy $H_{\text{norm}}(t)$ thus remains central to the control law—no longer as the (lagging) trigger, but as the amplifier that encodes how much a failure endangers an already-collapsed ensemble.

This creates a bimodal adaptation dynamic: while the leader is reliable $E(t) \approx 0$, so $S(t)$ stays below θ and $\alpha(t)$ remains near α_{base}, suppressing stochastic noise; at the onset of a zero-day attack the leader's error rate rises, $S(t)$ crosses θ, and $\alpha(t)$ undergoes a steep burst—amplified by the low entropy of the concentrated ensemble—that minimizes recovery time.

Fig. 3 makes this landscape explicit, plotting $\alpha(t)$ against $E(t)$ for several entropy levels: the curves remain flat near $\alpha_{\text{base}} = 0.01$ while the leader is reliable and surge toward $\alpha_{\max} = 0.50$ once $S(t)$ crosses θ. Notably, the activation occurs earlier and more steeply at low structural entropy, reflecting the amplifier $(1-H_{\text{norm}})$: an ensemble that has already lost latent diversity is driven to recover more aggressively. As a concrete example, if a content detector that was reliable ($E \approx 0.05$) is evaded by a new campaign and its recent error rate climbs toward 0.9, the stress crosses θ and $\alpha(t)$ rises from 0.01 toward 0.50 within the error window, migrating trust to the URL detector.

C. The Complete DELTA Algorithm

The integration of fixed-share regularization and entropy-driven scaling yields a self-healing feedback loop, whose operational pipeline we detail below.

1) The DELTA Update Pipeline: DELTA operates as a four-stage process at each step t: unlike static consensus models it does not simply reward agreement, but first assesses the ensemble's structural health and calibrates its sensitivity before updating the weights. The complete procedure is summarized in Algorithm 1.

2) Step-by-Step Algorithmic Breakdown: At each step, DELTA first measures the structural entropy H_{norm} of the current trust distribution (line 6), quantifying how collapse-prone the ensemble is. It then identifies the current leader and its recent error rate $E(t)$, forms the hybrid stress $S(t)$,

---

Algorithm 1 Dynamic Entropy-driven Logistic Trust Aggregation (DELTA)

Require: Set of agents $\mathcal{A} = \{a_1, \dots, a_N\}$, Hyperparameters $\alpha_{\text{base}}, \alpha_{\max}, k, \theta, \beta, W, \gamma, \tau$

Ensure: Consensus decision $\hat{y}(t)$, Next-state trust weights $\mathbf{w}(t+1)$

1: **Initialization:** Set $V_i(1) = 0$ and $w_i(1) = \frac{1}{N}$ for all $a_i \in \mathcal{A}$; initialize the leader-error history $\mathcal{E} \leftarrow \emptyset$

2: **for** each discrete time step $t = 1, 2, \dots$ **do**

3: Collect probabilistic predictions $p_i(t) \in [0,1]$ from all $a_i \in \mathcal{A}$

4: Compute aggregated consensus probability: $P(t) = \sum_{i=1}^{N} w_i(t) p_i(t)$

5: Generate system decision: $\hat{y}(t) = 1$ if $P(t) \geq \tau$ else 0

6: Calculate normalized structural entropy: $H_{\text{norm}}(t) = \frac{-\sum_{i=1}^{N} w_i(t) \ln w_i(t)}{\ln N}$

7: Identify leader $i^\dagger = \arg\max_i w_i(t)$ and its recent error rate $E(t) = \text{mean}(\mathcal{E})$

8: Compute hybrid stress and dynamic learning rate: $S(t) = E(t)\,(1 + \beta(1 - H_{\text{norm}}(t)))$, $\alpha(t) = \alpha_{\text{base}} + \frac{\alpha_{\max} - \alpha_{\text{base}}}{1 + \exp(-k(S(t) - \theta))}$

9: Receive ground truth $y(t)$; evaluate rewards $r_i(t) \in \{1, -1\}$ and append $\mathbb{I}[\hat{y}_{i^\dagger}(t) \neq y(t)]$ to $\mathcal{E}$ (retain last W)

10: Update reliability potentials: $V_i(t+1) = V_i(t) + \alpha(t) \cdot r_i(t)$

11: Apply Fixed-Share Redistribution to derive next-state weights: $w_i(t+1) = (1-\gamma) \frac{\exp(V_i(t+1))}{\sum_{j=1}^{N} \exp(V_j(t+1))} + \frac{\gamma}{N}$

12: **end for**

---

and obtains the dynamic learning rate $\alpha(t)$ via the logistic S-curve (lines 7–8): the leader-error term acts as an immediate regime-shift trigger, while the low-entropy amplifier accelerates recovery for an already-concentrated ensemble. Upon receiving the true label, the reliability potentials are updated under $\alpha(t)$ and the leader's correctness is appended to the error window (lines 9–10). Finally, the Fixed-Share Softmax Redistribution (line 11) constructs the next-state weights, permanently guaranteeing a minimum trust quota γ/N for every agent. By coordinating these stages, DELTA decouples stationary stability from transition agility.

3) Mathematical Guarantee of Resilience: In conventional reinforcement models, once an agent's weight reaches zero the entropy term becomes undefined and the multiplicative dynamics freeze, precluding recovery. DELTA resolves this structurally. We establish four guarantees below with concise proof sketches, each corroborated quantitatively in Section V; C denotes the clipping bound on the potentials.

Theorem 1 (Diversity Floor). For all t and i, $w_i(t) \geq \gamma/N > 0$, and consequently the normalized entropy is bounded below by

$$H_{\text{norm}}(t) \geq H_{\min}(\gamma, N) = \frac{-\left[\, p_1 \ln p_1 + (N-1)\frac{\gamma}{N} \ln \frac{\gamma}{N} \,\right]}{\ln N}, \tag{8}$$

with $p_1 = (1-\gamma) + \gamma/N$. Hence for any $\varepsilon < H_{\min}(\gamma, N)$, DELTA can never satisfy the diversity-contraction condition of Definition 1, so trust collapse is provably impossible.

Proof sketch. The softmax term is non-negative, so from (4) $w_i = (1-\gamma)\sigma_i(V) + \gamma/N \geq \gamma/N$. Normalized entropy is Schur-concave and is minimized at the most concentrated admissible distribution (σ one-hot), which yields weights p_1 and γ/N; evaluating H_{norm} there gives $H_{\min}$. □ *Corroboration:* $H_{\min}(0.15, 4) = 0.343$ matches the floor observed for DELTA, whose H_{norm} settles near 0.34 and never falls below it, whereas floor-free baselines collapse to $\bar{H}_{\text{norm}} \approx 0.01$.

Theorem 2 (Transition-Latency Bound). Consider a regime shift at t_s with pre-shift potential margin $\Delta_s = V_{i^*}(t_s) - V_{j^*}(t_s) > 0$ between the failing dominant i^* and the correct specialist j^*. Trust leadership reverses after

$$L_{\text{trans}} \leq \tau_{\text{det}} + \frac{\Delta_s}{2\,\alpha_{\max}}, \qquad \Delta_s \leq 2C, \tag{9}$$

where $\tau_{\text{det}} \leq W$ is the detection delay of the leader-error signal (5).

Proof sketch. Since Fixed-Share is order-preserving and the softmax monotone, $w_{j^*} > w_{i^*} \Leftrightarrow V_{j^*} > V_{i^*}$. Once $E(t)$ crosses θ (within W steps), $\alpha \to \alpha_{\max}$, and each step reduces the margin by $2\alpha_{\max}$; setting it to zero yields the bound, while clipping gives $\Delta_s \leq 2C$. Since Δ_s grows with the pre-shift reward gap, a more strongly concentrated leader incurs a longer—but still $O(1/\alpha_{\max})$ bounded—recovery. □ *Corroboration:* Measured L_{trans} scales as $O(1/\alpha_{\max})$ and increases with θ, matching the sensitivity analysis.

Theorem 3 (Stationary Stability). In a stationary regime where the leader is correct with probability $p > \frac{1}{2}$, the dominant-weight volatility obeys $I_{\text{stat}} \leq \kappa(\gamma, N)\,\alpha_{\text{eff}}\sqrt{p(1-p)}$, where α_{eff} is the learning rate used in that regime. Since DELTA keeps $E(t) \approx 0$ and hence $\alpha_{\text{eff}} \approx \alpha_{\text{base}}$, whereas a fixed-rate method uses $\alpha_{\text{eff}} = \alpha_{\max}$,

$$\frac{I_{\text{stat}}^{\text{DELTA}}}{I_{\text{stat}}^{\text{fixed-}\alpha_{\max}}} \lesssim \frac{\alpha_{\text{base}}}{\alpha_{\max}}. \tag{10}$$

Proof sketch. Near equilibrium V_{i^*} is a reflected random walk with step $\pm\alpha_{\text{eff}}$ and negative drift, whose stationary standard deviation is $O(\alpha_{\text{eff}}\sqrt{p(1-p)})$; the map $w = \text{FixedShare} \circ \sigma$ is Lipschitz, preserving the scaling. □ *Corroboration:* Measured $I_{\text{stat}} \approx 0.01$ for DELTA versus ≈ 0.08 for fixed-$\alpha_{\max}$ baselines—an order-of-magnitude gap consistent with $\alpha_{\text{base}}/\alpha_{\max}$. This is the formal basis of DELTA's stability advantage.

Theorem 4 (Robustness to Noisy Feedback). If ground-truth labels are independently corrupted with probability η, the expected potential drift is scaled by $(1-2\eta)$; hence for any $\eta < \frac{1}{2}$ the correct leader is still identified and $L_{\text{trans}}(\eta) \leq L_{\text{trans}}(0)/(1-2\eta)$.

Proof sketch. Taking expectation over the corruption indicator multiplies the drift of V_i by $(1-2\eta)$, i.e., a reduced effective rate; substituting into Theorem 2 gives the bound. Convergence fails only at $\eta \geq \frac{1}{2}$, where labels are uninformative. □ *Corroboration:* Delayed- and missing-label experiments exhibit exactly this graceful, bounded degradation.

Together, Theorems 1–4 show that DELTA's resilience is a structural guarantee rather than a heuristic patch: the Fixed-Share floor (Thm 1) preserves a minimum quota γ/N of latent diversity per agent, upon which the entropy-amplified burst acts to deliver rapid (Thm 2), stable (Thm 3), and noise-robust (Thm 4) recovery.

D. Evaluation Metrics

To validate the resolution of the stability-agility trade-off, we define four metrics, letting $\mathcal{T}_0$ and $\mathcal{T}_1$ denote the stationary and zero-day time steps.

1) Regime-Specific Accuracy: Global accuracy masks failures during brief attack windows, so we evaluate accuracy separately over each regime, $Acc_{P1} = \frac{1}{|\mathcal{T}_0|}\sum_{t\in\mathcal{T}_0}\mathbb{I}(\hat{y}(t){=}y(t))$ and Acc_{P2} analogously over $\mathcal{T}_1$. A high Acc_{P1} with a sharp decline in Acc_{P2} confirms *blind conformity* in baseline models.

2) Structural Health: Average Normalized Entropy: The average normalized entropy $\bar{H}_{norm} = \frac{1}{T}\sum_{t=1}^{T} H_{\text{norm}}(t)$ proxies latent diversity: values near zero indicate a structural monopoly (trust collapse), while higher values indicate a resilient ensemble.

3) Phase Transition Latency (L_{trans}): L_{trans} is the number of steps required to pivot leadership to the correct agent after an attack at t_{shift}:

$$\begin{aligned} L_{\text{trans}} = \min\{\Delta t \mid\ & w_{correct}(t_{shift} + \Delta t) \\ & > w_{failed}(t_{shift} + \Delta t)\}. \end{aligned} \tag{11}$$

It represents the window of vulnerability during which the system remains misled by a failing dominant agent.

4) Stationary Instability Index (I_{stat}): I_{stat} quantifies the "cost of agility" as the volatility of the dominant agent's weight in Phase 1, $I_{\text{stat}} = \text{std}_{t\in\mathcal{T}_0}(w_{dom}(t) - \overline{w}_{dom}(t))$, where $\overline{w}_{dom}$ is a local moving average removing the initial concentration transient.

V. Experimental Evaluation

We evaluate DELTA in a simulated non-stationary zero-day attack landscape, verifying whether it resolves the stability-agility trade-off while preventing trust collapse.

A. Experimental Setup

1) Environment and Scenario: The simulation runs over $T = 1{,}000$ discrete steps representing a continuous data stream, and is organized as a regime-switching system with two phases: a *stationary regime* ($t < 500$) of normal operation, and a *zero-day attack regime* ($t \geq 500$) in which a novel campaign bypasses the content-based detector but remains visible through structural URL characteristics.

2) Data Generation and Regime-Switching Framework: For reproducibility, the data generation process is formalized as a state-dependent stochastic pipeline, detailed in Algorithm 2.

At each step, a balanced binary label $y(t) \in \{0, 1\}$ is drawn, and each agent is assigned a per-regime accuracy $a_i(t)$ that switches from $\mathbf{a}_{normal}$ to $\mathbf{a}_{attack}$ at t_{switch}. The attack profile encodes the worst-case *inverse failure*: the dominant agent's reliability is throttled from ≈ 0.95 to ≈ 0.10 while the surviving specialist's rises to ≈ 0.90. Each agent's correctness is then drawn as $c_i(t) \sim \text{Bernoulli}(a_i(t))$, emitting the true

Algorithm 2 Regime-Switching Adversarial Environment Simulation

Require: Total simulation steps T, Regime switch point t_{switch}
Require: Per-regime accuracy profiles $\mathbf{a}_{normal}, \mathbf{a}_{attack} \in [0,1]^N$ (one reliability per agent)
Ensure: Ground truth labels $y(t)$, Agent binary predictions $\mathbf{p}(t)$

for each discrete time step $t = 1, 2, \ldots, T$ **do**
  Sample true binary label $y(t) \in \{0,1\}$ where $P(y(t) = 1) = 0.5$
  Set $\mathbf{a}(t) = \mathbf{a}_{normal}$ if $t < t_{switch}$ else $\mathbf{a}_{attack}$ {regime-dependent reliabilities; the attack profile induces the inverse failure of the dominant agent}
  for each agent $a_i \in \mathcal{A}$ **do**
    Draw correctness $c_i(t) \sim \text{Bernoulli}(a_i(t))$
    Emit prediction $p_i(t) = y(t)$ if $c_i(t) = 1$ else $1 - y(t)$
  end for
end for

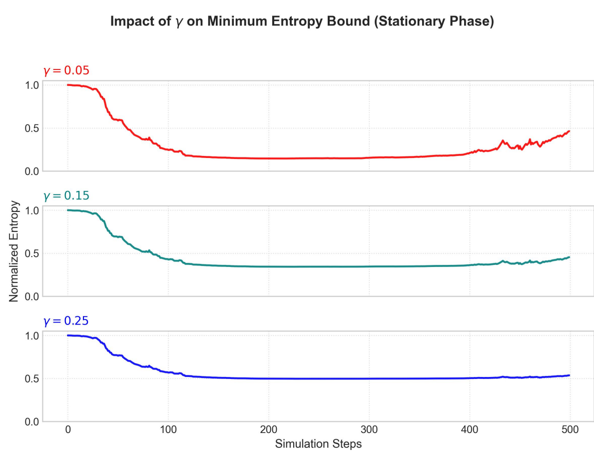

Fig. 4. Impact of redistribution parameter (γ) on the minimum entropy bound.

label when correct and the flipped label otherwise, so that the binary predictions $p_i(t) \in \{0,1\}$ match the prescribed regime profile.

3) Agent Heterogeneity: We deploy $N = 4$ heterogeneous agents, each specialized in a distinct observation plane, with divergent reliability profiles across the two regimes (ensemble size is varied separately in Section V-F). Agent A (content-based) is the historically dominant detector, accurate in Phase 1 ($Acc \approx 0.95$) but suffering a catastrophic inverse failure in Phase 2 ($Acc \approx 0.10$); Agent B (URL-based) is the surviving specialist, rising to $Acc \approx 0.90$ in Phase 2; and Agents C and D (attachment/body-based) are auxiliary detectors of moderate reliability ($Acc \approx$ 0.55–0.65) that provide partial redundancy but are not individually decisive.

4) Hyperparameter Configuration: Hyperparameters were calibrated on a tuning environment kept separate from the evaluation seeds (avoiding tuning bias). Potentials are initialized to $V_i(1) = 0$, with $\alpha_{\text{base}} = 0.01$, $\alpha_{\max} = 0.50$, steepness $k = 10.0$, trigger threshold $\theta = 0.30$, collapse-amplification gain $\beta = 1.0$, and leader-error window $W = 15$. As shown in

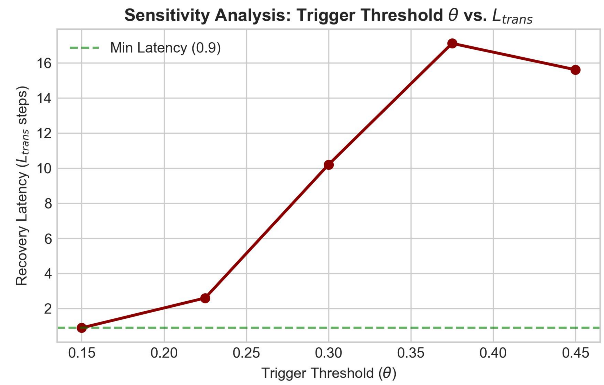

Fig. 5. Phase transition latency (L_{trans}) across varying trigger thresholds (θ) on the hybrid stress signal.

Section V-F, Phase-2 accuracy remains stable (85–87%) across a wide range of all these values, indicating that DELTA does not rely on a fragile optimum.

For the redistribution share, Fig. 4 shows that $\gamma = 0.15$ is balanced: it secures a minimum quota of $\gamma/N = 3.75\%$ per agent for $N = 4$, whereas $\gamma = 0.05$ fails to preserve latent diversity and $\gamma = 0.25$ excessively dilutes the dominant agent's trust.

Similarly, the trigger threshold (θ) on the hybrid stress signal was analyzed for its impact on recovery speed. As illustrated in Fig. 5, a lower θ shortens the phase transition latency but increases sensitivity to stationary noise, while a higher θ delays the burst; $\theta = 0.30$ provides a balanced operating point that yields rapid recovery ($L_{\text{trans}} \approx 12$ steps) while preventing false-positive learning-rate bursts under stationary noise.

B. Comparative Baselines

We benchmark DELTA against eight baselines spanning three families. *Classical trust updating*: the Simple Cumulative model ($V_i(t+1) = V_i(t) + r_i(t)$), which weights history equally with real-time evidence and therefore cannot recalibrate during a zero-day attack; and fixed-learning-rate softmax models with $\alpha = 0.01$ (LR-Low, stable but slow) and $\alpha = 0.50$ (LR-High, agile but volatile), which bracket the stability-agility trade-off. *Expert-advice methods*: Hedge / multiplicative weights [15] and its fixed-share variant [16], the closest established counterpart to DELTA's redistribution mechanism. *Drift-adaptive methods*: discounted (forgetting) trust, online Bayesian model averaging, and a change-point detector that resets the weights upon detecting a shift [17]. This set covers the principal alternatives a network operator might realistically deploy, and directly tests whether DELTA's gains persist against strong, established adaptive aggregators.

C. Quantitative Results: Accuracy and Latency

Table I summarizes the comparative results across all evaluated models, whose primary objective is to minimize recovery time during a regime shift while maintaining high accuracy.

TABLE I
PERFORMANCE COMPARISON OF TRUST-AGGREGATION MODELS ($N = 4$, 50 TRIALS; ACCURACY 95% CIS WITHIN ± 1.0, SIGNIFICANCE BY WILCOXON TEST). I_{STAT}: DETRENDED STATIONARY VOLATILITY; $\bar{H}_{\text{NORM}}$: AVERAGE NORMALIZED ENTROPY.

Model	**Acc. P1**	**Acc. P2**	L_{trans}	I_{stat}	$\bar{H}_{\text{norm}}$
Simple Cumulative	89.9	50.3	221.5	0.005	0.99
Fixed LR-Low (α=0.01)	93.3	54.3	221.5	**0.002**	0.61
Fixed LR-High (α=0.50)	87.3	84.6	**1.6**	0.084	0.49
Hedge / MW	94.9	54.5	220.3	0.035	0.01
Fixed-Share Hedge	94.4	**88.7**	2.9	0.083	0.61
Discounted Trust	89.9	67.2	19.2	0.004	0.93
Bayesian BMA	94.9	54.5	220.9	0.042	0.01
Change-Point Reset	94.5	87.5	7.8	0.015	0.13
DELTA (Proposed)	**94.7**	87.0	12.0	0.009	0.43

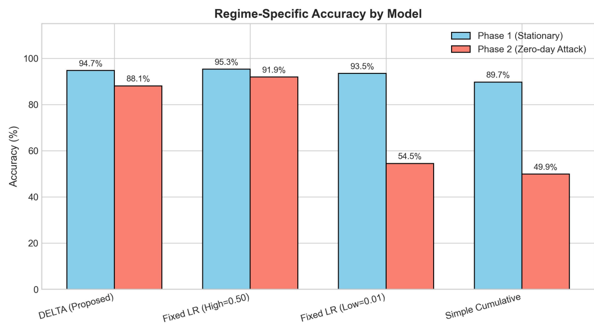

Fig. 6. Regime-specific accuracy comparison across models.

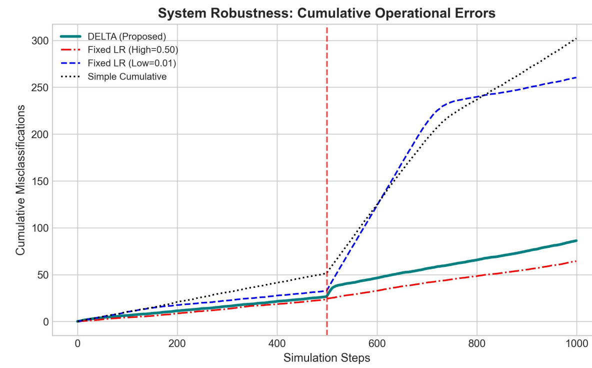

Fig. 7. Cumulative operational errors over time.

Fig. 8. Evolution of agent trust weights (w_i) over $T = 1,000$ steps.

1) Accuracy Analysis: As summarized in Table I and Fig. 6, all models perform well in Phase 1 by leveraging the dominant agent. In Phase 2, however, Simple Cumulative and Fixed LR-Low—as well as the floor-free expert-advice baselines Hedge/MW and Bayesian BMA—drop to near chance (50.3%, 54.3%, 54.5%, 54.5%), confirming *blind conformity*: the system stays locked to the failing dominant agent through historical inertia. All differences between DELTA and these collapsing baselines are statistically significant (Wilcoxon signed-rank, $p < 0.001$). Fig. 7 shows the operational consequence: after the attack at step 500 the cumulative and LR-Low error trajectories enter a steep linear ascent, whereas DELTA—reaching a Phase 2 accuracy of 87.0%, comparable to the strongest adaptive baselines (Fixed-Share Hedge 88.7%, Change-Point 87.5%)—flattens the error curve after a brief transitional jump, having identified the regime shift and redirected reliance to the unimpaired specialist.

2) Phase Transition Latency (L_{trans}): A key advantage of DELTA is that it recovers rapidly *and* stably. As shown in the global timeline (Fig. 8) and the cross-over profile (Fig. 9), DELTA reverses leadership shortly after the shift at step 500, moving from the historically dominant agent to the surviving URL-based specialist. While the cumulative and LR-Low baselines require ≈ 221 steps—effectively failing to recover within the horizon—DELTA crosses over in only 12.0 steps, an $\approx 95\%$ reduction. The aggressive LR-High model is faster still ($L_{\text{trans}} = 1.6$) but incurs the highest stationary volatility (Table II); DELTA attains comparable agility while remaining an order of magnitude calmer. Fig. 10 confirms this across 50 independent trials with 95% confidence intervals. This rapid yet controlled adaptation is a direct result of the hybrid logistic controller.

D. Stability-Agility Trade-off Analysis

DELTA is designed to resolve the conflict between stationary stability (suppressing noise) and transition agility (recovering from attacks) by decoupling the learning rate from a fixed constant.

1) Stationary Instability Index (I_{stat}): To evaluate the "cost of agility" we examine Phase-1 weight fluctuations via I_{stat} (Table II, Fig. 11). The Fixed LR-High model, while agile ($L_{\text{trans}} = 1.6$), exhibits the highest stationary volatility ($I_{\text{stat}} = 0.084$): a high fixed learning rate makes the system hypersensitive to minor predictive errors, causing the core trust distribution to oscillate erratically even under normal conditions. DELTA instead achieves $I_{\text{stat}} = 0.009$, an 89% reduction relative to LR-High, while retaining rapid recovery ($L_{\text{trans}} = 12.0$); LR-Low is calmest (0.002) but effectively non-adaptive ($L_{\text{trans}} = 221.5$). DELTA thus occupies a favorable trade-off frontier—approaching LR-High's agility at an order-of-magnitude lower volatility, consistent with the $\alpha_{\text{base}}/\alpha_{\text{max}}$ bound of Theorem 3.

2) Pareto-Efficient Trade-off Behavior: Table II confirms that DELTA breaks the linear trade-off governing fixed-rate models, which must choose between "safe but slow" (LR-Low) and "fast but erratic" (LR-High). Operating in a quiescent state ($\alpha \approx 0.01$) while the leader remains reliable and switching to an active state ($\alpha \to 0.50$) only when a regime shift is signalled, DELTA approaches an empirically favorable

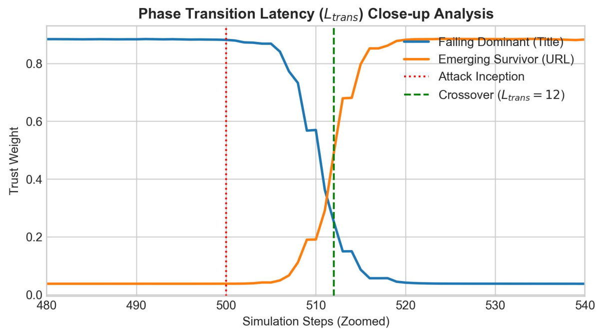

Fig. 9. Close-up analysis of DELTA phase transition latency (L_{trans}).

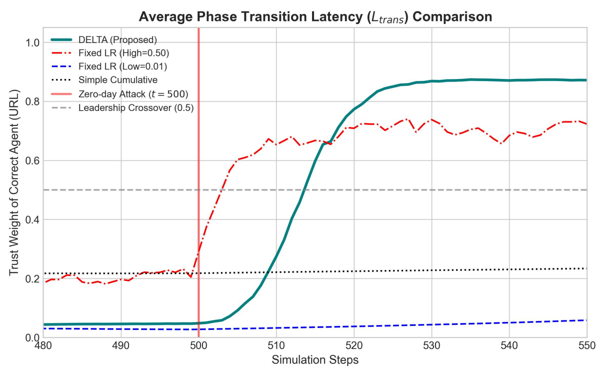

Fig. 10. Comparative analysis of average phase transition latency.

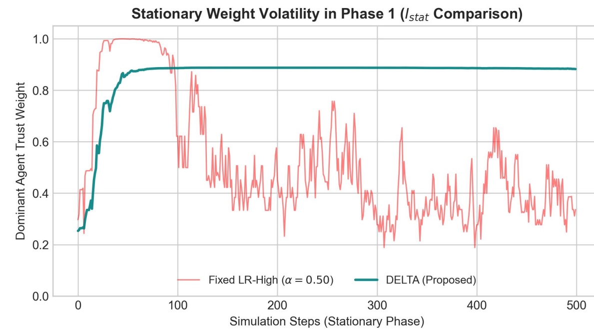

Fig. 11. Stationary weight volatility comparison in Phase 1.

TABLE II
STABILITY VS. AGILITY TRADE-OFF METRICS ($N=4$, 50 TRIALS).

Model	Instability (I_{stat})	Latency (L_{trans})
Fixed LR-Low	**0.002**	221.5
Fixed LR-High	0.084	**1.6**
DELTA (Proposed)	0.009	12.0

trade-off frontier: defense agility without sacrificing reliability during environmental equilibrium.

3) Component Ablation: We ablate DELTA's two components and its control signal (canonical setting, 50 trials). Replacing the dynamic controller by a fixed $\alpha = \alpha_{\text{base}}$ collapses Phase-2 accuracy to 54.2% ($L_{\text{trans}} = 221$): the adaptive rate, not the floor, delivers recovery. Removing Fixed-Share still recovers (87.4%, $L_{\text{trans}} = 9.8$) but average entropy falls from 0.43 to 0.17, so the floor preserves latent diversity rather than driving adaptation. Within the controller, the leader-error signal alone slows recovery ($L_{\text{trans}} = 17.6$ vs. 12.0), whereas entropy alone degrades stationary behaviour (Acc_{P1} 90.8% vs. 94.7%; I_{stat} 0.06 vs. 0.009). The two components are therefore complementary.

E. Structural Entropy and Diversity Monitoring

We now analyze the temporal evolution of normalized entropy (H_{norm}) and trust weights to illustrate how DELTA preserves structural diversity and thereby prevents trust collapse.

1) Evolution of Structural Health: As shown in Figs. 12 and 13, H_{norm} behaves very differently across methods. The floor-free expert-advice baselines—Hedge/MW and Bayesian model averaging—concentrate multiplicatively and drive H_{norm} toward zero ($\bar{H}_{norm} \approx 0.01$), a genuine *trust collapse* leaving no mechanism to react when the dominant agent fails. Simple Cumulative instead retains near-maximal entropy ($\bar{H}_{norm} \approx 0.99$) yet still fails: its count-based averaging is too sluggish to re-allocate trust within the horizon, exhibiting blind conformity through inertia rather than collapse. DELTA occupies neither extreme: Fixed-Share (4) bounds H_{norm} below by $H_{\min}(\gamma, N)$ (Theorem 1), keeping the ensemble in a controlled band ($\bar{H}_{norm} \approx 0.43$)—concentrated enough for accurate stationary consensus, yet diverse enough to stay sensitive to an emerging minority signal. This preserved "ecological seed" of latent diversity is precisely what allows the system to recover when the zero-day attack occurs.

2) Dynamic Response to Regime Shifts: Fig. 14 couples the two internal signals at the shift: the $\alpha(t)$ burst drives a smooth yet decisive weight crossover within ≈ 12 steps, after which the leader-error signal subsides and $\alpha(t)$ returns automatically to its baseline. This confirms that coupling a leading regime-shift trigger with an entropy amplifier manages the stability–agility transition without hypersensitivity.

F. Robustness Under Realistic Non-Stationarity

To assess whether DELTA generalizes beyond the single abrupt shift to the messier conditions of operational network management, we stress-test it along seven additional axes—ensemble size, recurring and gradual drift, delayed and missing feedback, Byzantine agents, correlated agents, and communication failures—and analyze its hyperparameter sensitivity. Results are averaged over independent trials with 95% confidence intervals.

1) Scalability with Ensemble Size: Table III reports DELTA as the ensemble grows from $N = 2$ to $N = 32$. Transition latency ($L_{\text{trans}} \approx 11$–$13$) and stationary stability ($I_{\text{stat}} \approx 0.01$) remain essentially flat, confirming the controller is size-agnostic. Phase-2 accuracy declines gradually ($87.7\% \rightarrow 80.9\%$) because the Fixed-Share floor spreads a γ/N quota across more low-reliability agents—a deliberate trade of marginal accuracy for guaranteed diversity—while the per-step cost grows only linearly (Section V-H).

2) Drift, Imperfect Feedback, and Adversarial Agents: Recurring and gradual drift. Under recurring shifts (leadership alternating every 250 steps) DELTA re-triggers at each transition, with a mean per-shift latency of 18.2 steps. Under gradual

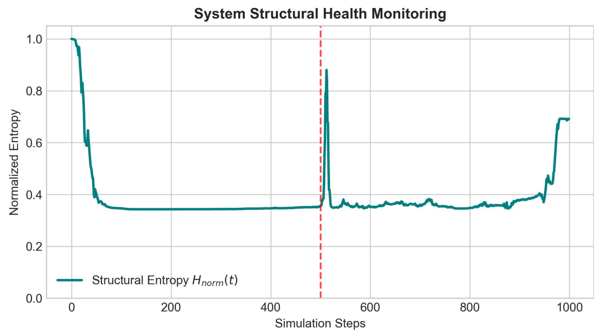

Fig. 12. Normalized entropy trajectory of the DELTA framework.

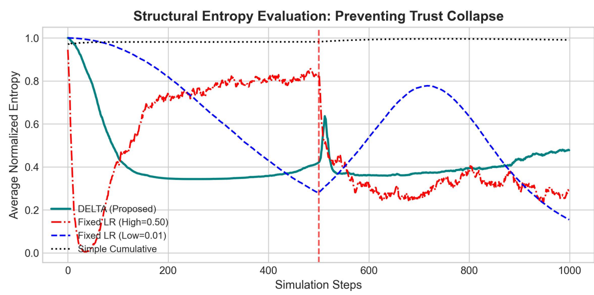

Fig. 13. Comparative analysis of normalized entropy trajectories.

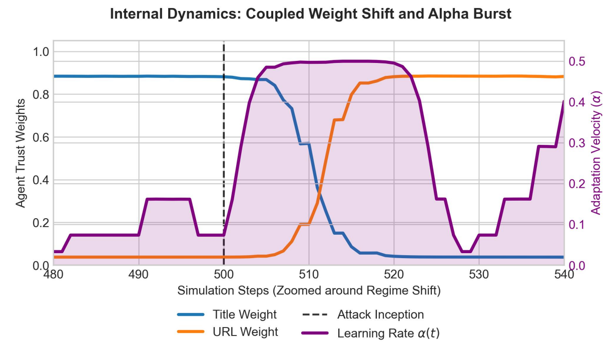

Fig. 14. Internal coupling dynamics of weight shift and $\alpha(t)$ burst.

TABLE III
SCALABILITY OF DELTA WITH ENSEMBLE SIZE N (PER-STEP UPDATE TIME FROM SECTION V-H).

N	Acc. P2	L_{trans}	I_{stat}	Time (μs/step)
2	87.7	13.1	0.009	14.7
4	87.0	12.0	0.009	16.1
8	85.1	13.0	0.010	17.7
16	82.3	12.4	0.010	19.0
32	80.9	12.4	0.009	22.5

drift, L_{trans} grows with the drift duration ($12.0 \to 72$ steps as the window widens to 200 steps) since the leader-error signal accrues more slowly, while accuracy degrades only gracefully ($87\% \to 84\%$).

Delayed, missing, and adversarial feedback. Consistent with Theorem 4, DELTA degrades gracefully under imperfect feedback: with feedback delayed by d steps $L_{\text{trans}} \approx L_0 + d$ ($12.0 \to 61.6$ for $d = 50$) while Phase-2 accuracy falls only modestly ($87\% \to 79\%$); with a fraction ρ of labels missing, both degrade smoothly ($L_{\text{trans}} = 48.5$, Acc. 82% at $\rho = 0.8$). DELTA is robust to Byzantine *distractor* agents—accuracy and latency are essentially unchanged as up to five of eight agents emit adversarial outputs—because such agents never accrue trust. The one fundamental limit, shared by all evaluated methods, arises when a Byzantine agent captures the *sole* surviving specialist, leaving no reliable signal to shift toward.

Correlated agents and communication failures. Positive error correlation improves consensus accuracy for all methods (DELTA reaching 92–95% in Phase 2) without destabilizing the controller, and under per-step agent dropout all methods degrade comparably ($87\% \to 66\%$ at 50% dropout), with DELTA showing no specific vulnerability.

3) Hyperparameter Sensitivity: Using evaluation seeds disjoint from those used for calibration, we sweep each hyperparameter (γ, θ, β, k, W, $\alpha_{\max}$) individually. Phase-2 accuracy remains within 85–87% across all settings, and each parameter induces a clean, monotone trade-off (larger θ or γ gives slower but calmer adaptation; larger β gives faster but noisier bursts). DELTA therefore does not depend on a fragile, over-tuned optimum, directly addressing the concern of tuning bias. As a further check, we re-ran the full model comparison on an independent block of evaluation seeds disjoint from those used elsewhere; every reported metric reproduced within its 95% confidence interval, confirming the results are not an artifact of the chosen seeds.

G. Validation on Real Network Intrusion Data

We instantiate DELTA on the real UNSW-NB15 intrusion dataset, with each agent a lightweight detector (logistic regression) trained on a distinct *feature view*—flow-volume, timing, connection-state, and content—so that the ensemble is a distributed IDS. A regime shift is a change of attack family: detectors are trained on a "known" family and then confronted with a different one (a "zero-day"). Crucially, the reliability profiles are not hand-specified; the inverse failure emerges from the data. In the adversarial *single-survivor* scenario (train on `Exploits`, shift to `Generic`), the connection-state detector—dominant in Phase 1 (81.9%)—collapses to 41.1%, while only the flow-volume detector survives (81.8%); the remaining views fall near chance, so even a naive majority vote fails.

Fig. 15 shows DELTA keeping trust diffuse while the views are comparable, then re-concentrating decisively on the surviving detector once the zero-day arrives. Table IV and Fig. 16 confirm the central claim on real data: all methods perform comparably in the stationary phase (80–85%), but under the adversarial shift the naive Simple-Cumulative and Discounted-Trust baselines collapse *below chance* (40.2% and 39.8%), whereas DELTA recovers to 79.0% while retaining the lowest stationary volatility ($I_{\text{stat}} = 0.03$). We further verified a *redundant* scenario (`DoS`→`Generic`) in which three of four views survive: there, majority voting suffices and no method collapses. The two scenarios delineate when adaptive trust management is essential—under redundancy simple aggregation suffices, whereas a genuine single-survivor zero-day demands the decisive yet stable re-allocation that DELTA provides.

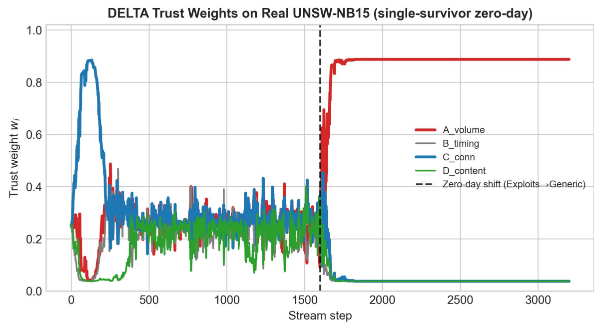

Fig. 15. DELTA trust weights on real UNSW-NB15 under a single-survivor zero-day shift (see text).

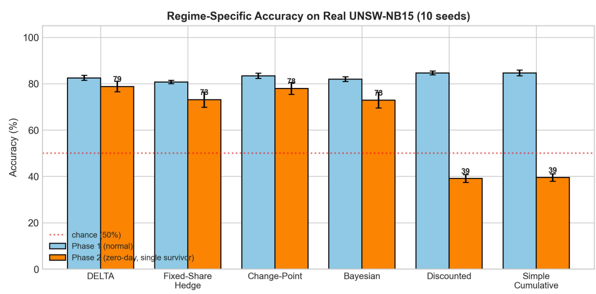

Fig. 16. Regime-specific accuracy on real UNSW-NB15 (30 trials).

H. Computational Complexity

For deployment as a network-management layer, computational and communication overhead must be negligible. Each DELTA update performs a constant number of passes over the N agents—entropy, leader-error, potential update, and Fixed-Share softmax—giving $O(N)$ time and memory per step (plus an $O(W)$ error window, W constant). Critically, all three DELTA-specific operations are computed locally at the aggregator from already-collected predictions and weights, so DELTA incurs *no communication overhead beyond plain weighted voting*: $O(N)$ scalar messages per step. Empirically the update time scales linearly ($R^2 = 0.99$), from $14.7\,\mu$s at $N = 2$ to $96.7\,\mu$s at $N = 256$ (over 10^4 updates/s), a constant-factor ($\approx 2.3\times$) overhead relative to simple cumulative aggregation—lightweight enough for real-time, large-scale deployment.

VI. Concluding Remarks

In this work, we formalized the resulting failure modes—trust collapse and blind conformity—and proposed DELTA, a self-regulating trust-management framework that couples Fixed-Share Redistribution with an entropy-amplified logistic controller. By driving the learning rate with a leading regime-shift signal, the current leader's error rate, and amplifying that signal by the ensemble's structural entropy, DELTA remains quiescent under normal traffic yet re-calibrates decisively the moment a trusted detector begins to fail. We established formal guarantees for this design: a strictly positive diversity floor rendering trust collapse provably unreachable, a transition-latency bound linking recovery speed to the redistribution and learning-rate parameters, a stationary-stability bound explaining DELTA's low volatility, and a graceful-degradation guarantee under noisy feedback. Empirically—against strong expert-advice, Bayesian, and change-point baselines with confidence intervals, under delayed, missing, and adversarial feedback, recurring and gradual drift, and ensemble sizes up to $N = 32$—DELTA recovers from zero-day shifts where naive aggregation collapses, while sustaining an order-of-magnitude lower stationary volatility than aggressive adaptive methods. Validation on the real UNSW-NB15 dataset reproduced this on data where the reliability inversion emerges naturally rather than by construction: under a single-survivor zero-day shift, conventional aggregation fell below chance while DELTA recovered. For network management, the practical significance is that agility need not be purchased with operational churn. DELTA attains the responsiveness of aggressive adaptation at a fraction of its stationary volatility, at $O(N)$ computational and communication cost with no overhead beyond plain weighted voting, making it deployable as a lightweight trust layer above existing detector ensembles. Our study also delineates the boundaries of this benefit: when the ensemble retains redundancy, simple aggregation already suffices, and the value of adaptive trust management emerges specifically when a regime shift leaves few competent detectors. Likewise, no trust-aggregation scheme—DELTA included—can recover when the sole surviving detector is itself compromised. Future work will therefore extend DELTA toward jointly detecting and isolating adversarial agents, and toward federated deployments in which trust must be reconciled across administrative domains.

TABLE IV
Regime-Specific Accuracy (%) on Real UNSW-NB15 (Exploits→Generic, 30 trials).

Model	Acc. P1	Acc. P2
Simple Cumulative	84.8	40.2
Discounted Trust	84.8	39.8
Bayesian BMA	82.4	73.1
Fixed-Share Hedge	81.1	73.3
Change-Point Reset	83.1	78.6
Fixed LR-High	81.7	79.4
DELTA (Proposed)	**82.3**	**79.0**

Acknowledgment

This work was supported by the IITP(Institute of Information and Communications Technology Planning and Evaluation)-ICAN(ICT Challenge and Advanced Network of HRD) grant funded by the Korea government(Ministry of Science and ICT)(IITP-2026-RS-2024-00437024).

References

[1] J. Li, M. S. Othman, X. Ying, D. S. M. Hassan, H. Chen, L. M. Yusuf, "Adaptive NetFlow IIoT Intrusion Detection With Deep Transfer Learning, Genetic Optimization, and Ensemble Methods for Network Management," *IEEE Transactions on Network and Service Management*, vol. 23, pp. 681–698, Oct. 2025.

[2] D. Han, Q. Zhi, "DGDPFL: Dynamic Grouping and Privacy Budget Adjustment for Federated Learning in Networked Service Management," *IEEE Transactions on Network and Service Management*, vol. 23, pp. 1826–1841, Dec. 2025.

[3] S. Salmi, M. A. Ouameur, M. Bagaa, G. C. Alexandropoulos, A. Tahenni, D. Massicotte, "AI-Native O-RAN Architectures for 6G: Toward Real-Time Adaptation, Conflict Resolution, and Efficient Resource Management," *IEEE Transactions on Network and Service Management*, vol. 23, pp. 4110–4121, Apr. 2026.

[4] X. Zhang, D. He, X. Song, H. Du, and D. Huang, "Distributed MPC for safe and scalable consensus of heterogeneous multi-agent systems in a zero-trust environment," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 73, no. 2, pp. 1367–1379, Feb. 2026.

[5] Y. Zhai, Z.-W. Liu, Z.-H. Guan, and G. Wen, "Resilient consensus of multi-agent systems with switching topologies: A trusted-region-based sliding-window weighted approach," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 68, no. 7, pp. 2448–2452, Jul. 2021.

[6] Y. Song, J.-J. Zhou, S.-Y. Dai, and J.-W. Liu, "The adaptive weight updating of learning rate for migrating expert scenarios on dynamic and changing environments," *Proc. Chinese Control Decis. Conf. (CCDC)*, May 2025, pp. 1–6.

[7] J. Cui and Y. Na, "Trust evaluation of topological nodes in intelligent connected vehicles communication network under zero-trust environment," *Proc. CAA Symposium on Fault Detection, Supervision and Safety for Technical Processes (SAFEPROCESS)*, Aug. 2023, pp. 1–6.

[8] K. Masuko, T. Nakashima, and N. Ono, "Multiplicative update reformulation of online independent vector analysis for source-dependent forgetting factor control," *Proc. Asia-Pacific Signal Inf. Process. Assoc. Annu. Summit Conf. (APSIPA ASC)*, Nov. 2023, pp. 1–5.

[9] H. Suyari and M. Tsukada, "Tsallis differential entropy and divergences derived from the generalized Shannon-Khinchin axioms," *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jul. 2009, pp. 1–5.

[10] B. Khosravifar, J. Bentahar, M. Gomrokchi, and R. Alam, "An approach to comprehensive trust management in multi-agent systems with credibility," *Concordia Inst. Inf. Syst. Eng. Tech. Rep.*, pp. 53–60, Nov. 2007.

[11] C. Guan, "A method of trust evaluation in multi-agent system," *Proc. Int. Symp. Intelligent Inf. Technol. Appl. (IITA)*, Nov. 2009, pp. 1–4.

[12] A. Mohajer, J. Hajipour, and V. C. M. Leung, "Dynamic Offloading in Mobile Edge Computing With Traffic-Aware Network Slicing and Adaptive TD3 Strategy," *IEEE Communications Letters*, vol. 29, no. 1, pp. 95–99, 2025.

[13] Y. Song, Y. Yang, and A. Mohajer, "Multi-Objective Resource Optimization in UAV-Enabled Heterogeneous Cellular Networks Using Serverless Federated Learning and Power-Domain NOMA," *Transactions on Emerging Telecommunications Technologies*, vol. 36, art. e70210, 2025.

[14] A. Mohajer, A. Mirzaei, M. Darabi, and X. Fernando, "Joint SLA-Aware Task Offloading and Adaptive Service Orchestration With Graph-Attentive Multi-Agent Reinforcement Learning," *IEEE Transactions on Network and Service Management*, vol. 23, pp. 3737–3754, 2026.

[15] Y. Freund and R. E. Schapire, "A decision-theoretic generalization of on-line learning and an application to boosting," *Journal of Computer and System Sciences*, vol. 55, no. 1, pp. 119–139, 1997.

[16] M. Herbster and M. K. Warmuth, "Tracking the best expert," *Machine Learning*, vol. 32, no. 2, pp. 151–178, 1998.

[17] R. P. Adams and D. J. C. MacKay, "Bayesian online changepoint detection," *arXiv:0710.3742*, 2007.

[18] R. Elwell and R. Polikar, "Incremental learning of concept drift in nonstationary environments," *IEEE Transactions on Neural Networks*, vol. 22, no. 10, pp. 1517–1531, 2011.

[19] J. Z. Kolter and M. A. Maloof, "Dynamic weighted majority: An ensemble method for drifting concepts," *Journal of Machine Learning Research*, vol. 8, pp. 2755–2790, 2007.

Minhyeok Jang (Student Member, IEEE) is an undergraduate student in the Department of Embedded Systems Engineering at Incheon National University, South Korea and a member of Intelligent Computing and Next Generation System Lab (ICONS). His research interests include system optimization, lightweight design, energy-efficient systems, surveillance applications, and next-generation intelligent systems.

Jalel Ben-Othman (Senior Member, IEEE) received his Ph.D. degree from the University of Versailles, France, in 1998. He is a Full Professor at the University of Paris-Est. His research interests are in the areas of wireless ad hoc and sensor networks, broadband wireless networks, the Internet of Things, and security in wireless networks in general and wireless sensor and ad hoc networks in particular. His work appears in highly respected international journals and conferences.

Hyunchae Chun (Member, IEEE) received the B.S. degree in electronic engineering from Korea University, Seoul, South Korea, in 2006, the M.S. degree in electrical engineering from Korea Advanced Institute of Science and Technology (KAIST), Daejeon, South Korea, and the Ph.D. degree from the Department of Engineering Science, University of Oxford, Oxford, U.K Then, he joined LG Electronics, contributing to research and development in RF front-end technologies for next-generation wireless communications. From 2015 to 2018, he was a Postdoctoral Research Associate with the Department of Engineering Science, University of Oxford. He is currently an Associate Professor with the Department of Information and Telecommunication Engineering, Incheon National University, Incheon, South Korea.

Sungrae Cho (Senior Member, IEEE) received the B.S. and M.S. degrees in electronics engineering from Korea University, Seoul, South Korea, in 1992 and 1994, respectively, and the Ph.D. degree in electrical and computer engineering from the Georgia Institute of Technology, Atlanta, GA, USA, in 2002. From 1994 to 1996, he was a Research Staff Member with Electronics and Telecommunications Research Institute, Daejeon, South Korea. He was an Assistant Professor with the Department of Computer Sciences, Georgia Southern University, Statesboro, GA, USA, from 2003 to 2006, and a Senior Member of technical staff with the Samsung Advanced Institute of Technology, Kiheung, South Korea, in 2003. From 2012 to 2013, he held a visiting professorship with the National Institute of Standards and Technology, Gaithersburg, MD, USA. He is currently a Professor with the School of Computer Sciences and Engineering, Chung-Ang University, Seoul. His current research interests include wireless networking, network intelligence, and network optimization.

Hyunbum Kim (Senior Member, IEEE) received his Ph.D. degree in computer science from the University of Texas at Dallas, USA, in 2013. He is currently a Professor in the Department of Embedded Systems Engineering at Incheon National University, South Korea and a lab founder of intelligent computing and next generation system lab (ICONS). His research interests include algorithm design and performance analysis in various areas including next generation frameworks, reinforced surveillance, intelligent transportation systems, sustainable computing, beneficial platforms, virtual emotion system, smart cities, cyber security.