S-AI-IoT : Formal Agent Specification, Mathematical Modeling, and Stability Analysis of the Hormonal Orchestration Framework

Said Slaoui

Mohammed V University, Rabat, Morocco

ABSTRACT

Constrained Internet of Things (IoT) environmentsrequire adaptive orchestration mechanisms capable of jointlymanagingsensingquality, connectivity, energyconstraints, resilience, and securitywhileremainingcomputationally frugal and intrinsicallyexplainable. This article develops the formalmathematicalframework of S-AI-IoT, a bio-inspiredsparse AI architecture based on hormonal orchestration and specialized agent activation for distributedIoTsystems.The article introduces the completeformalspecification of 24 specializedIoT agents organizedinto six functionalfamilies, eachdefinedthrough explicit activation conditions, input-output contracts, resource profiles, and operationalconstraints. It furtherdevelops the mathematicalfoundations of the hormonal orchestration layer, includinglogistic-saturatedemissionfunctions, stochasticreaction-diffusion dynamics over dynamicIoT graphs, Lyapunov-based global stabilityanalysis, Euler-Maruyamadiscretizationwithbounded projection, duty-cycle energy modulation, primal-dual constrained orchestration, and distributedsymbolicengram memory withintrinsicexplainabilityguarantees. The proposedframeworkestablishes, to the best of the author'sknowledge, one of the first formallyspecified and stability-analyzed hormonal regulation architectures that unifies reaction-diffusion dynamics, constrainedoptimization, symbolic memory, and explainablesparse orchestration within a single operationalframework for distributedIoTsystems. The resulting architecture supports adaptive decision-makingundersevereresourceconstraintswhilepreservingenergyfrugality, distributedrobustness, and causal traceability of behavioraldecisions.

KEYWORDS

IoT, formal agent specification, reaction-diffusion, Lyapunov stability, primal-dual optimization, hormonal orchestration, symbolic engram memory, intrinsic explainability.

1. Introduction

Article I **[1]** of this trilogy introduced S-AI-IoT , an application of the S-AI paradigm **[2]** to constrained IoT systems, organized around three first-class architectural properties: activation parsimony, hormonal orchestration, and distributed symbolic engram memory, and developed its theoretical foundations and seven-layer global architecture. The present article, Article II, addresses the formal mathematical substance of S-AI-IoT: the complete specification of its agent layer and the rigorous mathematical formalization of its hormonal regulation engine.The article makes two principal contributions. **First**, Section 2 provides the complete typology of the 24 specialized IoT agents organized in six functional families, each with a precise functional definition, formal activation condition expressed as a utility function over the hormonal state vector $\mathbf{h}(t) \in [0,1]^5$, formal input specification $\mathcal{I}_a$, formal output specification $\mathcal{O}_a$, resource cost

DOI:10.5121/ijaia.2026.17303

profile $(c_a^{\text{cpu}}, c_a^{\text{mem}}, c_a^{\text{comm}})$, and Time-to-Live and Cooldown parameters. **Second**, Section 3 develops the complete mathematical framework in eight subsections covering: the graph-theoretic substrate and notation system; the five canonical IoT observation vectors and hormonal state variables; the logistic-saturated emission function architecture with IoT-specific aggregation functions; the stochastic reaction-diffusion dynamics with six formally interpreted terms; the stability analysis via Lyapunov direct method establishing global asymptotic stability; the Euler-Maruyama discretization with projection; the duty cycle modulation law with multi-hormonal override; the primal-dual constrained agent-selection optimization; and the distributed symbolic engram memory formalism with content-addressable retrieval and the zero-cost intrinsic explainability guarantee.The complete pseudocode implementation of the six core algorithms is provided as supplementary material accompanying this article series.

1.1. Related Context

Recent IoT research has increasingly focused on adaptive sensing, distributed orchestration, and resource-aware autonomous operation in constrained environments. Existing work has addressed missing-data imputation mechanisms for unreliable IoT sensing streams [3] as well as adaptive sensor-fusion and data-quality assessment strategies [4]. Other approaches have explored sparse autonomous networking and adaptive communication management for distributed infrastructures [5], while broader IoT studies have emphasized the architectural, security, and deployment constraints inherent to large-scale IoT systems [6]. A more detailed Related Work analysis covering bio-inspired orchestration, energy-aware IoT intelligence, distributed resilience, explainable AI, and sparse AI architectures was developed in Article I of this trilogy. The present article therefore focuses specifically on the formal mathematical specification and stability analysis of the S-AI-IoT orchestration framework.

2. COMPLETE TYPOLOGY OF SPECIALIZED IOT AGENTS WITH FORMAL INPUT-OUTPUT CONTRACTS

The specialized agent layer constitutes the operational core of S-AI-IoT: it is at this layer that hormonal context signals are translated into concrete behavioral actions modifying the IoT node's operational state, communication patterns, energy utilization, security posture, and data quality. The present section develops the complete formal specification of all 24 agents organized in six functional families. For each agent, the specification comprises: functional definition, formal activation condition, input specification $\mathcal{I}_a$, output specification $\mathcal{O}_a$, resource cost profile $(c_a^{\text{cpu}}, c_a^{\text{mem}}, c_a^{\text{comm}})$, Time-to-Live T_a^{live} and CooldownT_a^{cool} parameters, and utility function $u_a(\mathbf{h})$.

2.1. Family F1 — Sensor Monitoring and Quality Management Agents (Sensorin-coupled)

Family F1 addresses the full spectrum of sensor degradation conditions encoded by Sensorin$h_{\text{sens},i}(t)$, from gradual accuracy drift to catastrophic sensor failure. The four F1 agents provide a graduated response with escalating resource investment proportional to degradation severity.

F1-A — SensorDiagnosticAgent. Continuous statistical monitoring of active sensor data streams to detect quality anomalies (hardware degradation, electromagnetic interference, calibration drift). First-line detection mechanism for all sensor quality events. - Activation: $u_{\text{F1A}}(\mathbf{h}) = 0.8\, h_{\text{sens}} + 0.2\, h_{\text{res}} > 0.25$ (low threshold = persistent background monitor) - Inputs: $\mathcal{I}_{\text{F1A}} =$

$\{O_i(t),\, q_i(t),\, \mathbf{h}(t),\, \varepsilon_{\text{recent}}\}$ - Outputs: $\mathcal{O}_{\text{F1A}} = \{r_{\text{diag}}(t),\, s_{\text{fault}}(t) \in \{\text{nominal, degraded, faulty, failed}\}^d,\, \Delta E_{\text{sens}}(t)\}$ - Cost: $\mathcal{O}(d \cdot N_{\text{stat}})$ ops/cycle, $\leq 2d$ registers; $T^{\text{live}} = 1$ cycle, $T^{\text{cool}} = 0$ (persistent)

F1-B — SensorRecalibrationAgent. Initiates and supervises online recalibration procedures for drifting sensors, supporting local recalibration (built-in reference signals) and network-assisted recalibration (cross-validation against neighboring nodes with overlapping coverage). - Activation: $h_{\text{sens}} > 0.45 \wedge s_{\text{fault}} = \text{degraded} \wedge h_{\text{ene}} < 0.70$ (energy gate: recalibration not initiated during energy crisis) - Outputs: $\mathcal{O}_{\text{F1B}} = \{\Delta_{\text{cal},s}(t),\, \text{cal\_status} \in \{\text{success, partial, failed}\},\, \text{engram\_record}\}$ - Cost: $\mathcal{O}(d \cdot |\mathcal{N}(i)|)$; $T^{\text{live}} = 5\text{–}20$ cycles, $T^{\text{cool}} = 50$ cycles

F1-C — DataImputationAgent. Lightweight statistical imputation to fill sensor data stream gaps (transient failures, sampling interruptions)[3]. Three strategies selected by energy state: LVCF if $h_{\text{ene}} > 0.6$ ($\mathcal{O}(1)$); neighborhood-weighted interpolation if $0.3 < h_{\text{ene}} \leq 0.6$ ($\mathcal{O}(|\mathcal{N}(i)|)$); Kalman filter estimation if $h_{\text{ene}} \leq 0.3$ ($\mathcal{O}(d^2)$). - Activation: $h_{\text{sens}} > 0.30 \wedge \text{missing\_rate}_i(t) > 0.10$ - Outputs: $\mathcal{O}_{\text{F1C}} = \{\hat{O}_i(t),\, \text{imputation\_mask}(t) \in \{0,1\}^d\}$

F1-D — SensorFusionAgent. Activates redundant or complementary sensor modalities when primary sensor quality falls below a critical threshold, and fuses their outputs through quality-weighted combination [4]. - Activation: $h_{\text{sens}} > 0.55 \wedge h_{\text{ene}} < 0.65$ (energy constraint prevents sensor fusion — which doubles or triples active sensor count — during severe energy deficit) - Outputs: $\mathcal{O}_{\text{F1D}} = \{\tilde{O}_i(t) = \sum_s q_i^{(s)}\, O_i^{(s)} / \sum_s q_i^{(s)},\, \sigma_{\tilde{O}}^2(t)\}$ - Cost: $\mathcal{O}(|\mathcal{S}_{\text{active}}| \cdot d)$; $T^{\text{live}} = 3$ cycles, $T^{\text{cool}} = 5$ cycles

2.2. Family F2 — Connectivity Management Agents (Connectin-coupled)

Family F2 addresses inter-node communication link management under the full range of connectivity degradation conditions encoded by Connectin$h_{\text{conn},i}(t)$, inheriting hormonal networking primitives from S-AI-NET [5].

F2-A — LinkQualityMonitorAgent. Continuous measurement and estimation of link quality metrics on all active communication interfaces [6]. Maintains the dynamic link weight matrix $W(t)$ used by the Hormonal Engine for reaction-diffusion computation. - Activation: persistent, $u_{\text{F2A}}(\mathbf{h}) \equiv 1$ (required by hormonal engine every cycle) - Inputs: $\{\text{RSSI}_{ij}(t), \text{PDR}_{ij}(t), \ell_{ij}(t), \sigma_{\ell,ij}(t)\}_{j \in \mathcal{N}(i)}$ (passive MAC monitoring) - Outputs: $\{W(t),\, L(t),\, \phi_{\text{conn}}(O_i(t))\}$ with $w_{ij}(t) = \text{RSSI}_{ij}(t) \cdot (1 - \text{loss}_{ij}(t))$ - Cost: $\mathcal{O}(|\mathcal{N}(i)|)$ CPU, $\mathcal{O}(|\mathcal{N}(i)|^2)$ memory, 0 communication

F2-B — AdaptiveRoutingAgent. Modifies routing tables toward higher-quality paths in response to detected link quality degradation (Dijkstra on local weighted graph), triggered by significant changes in $W(t)$ rather than on a fixed schedule. - Activation: $h_{\text{conn}} > 0.40 \wedge \| \Delta W(t) \|_F > \epsilon_W$ (change-magnitude guard) - Outputs: $\{\text{routing\_table}_i(t+1),\, \text{path\_update\_log}(t)\}$ - Cost: $\mathcal{O}(|\mathcal{N}(i)|^2 \log|\mathcal{N}(i)|)$; $T^{\text{live}} = 3$ cycles, $T^{\text{cool}} = 10$ cycles

F2-C — ProtocolAdaptationAgent. Cross-layer optimization adjusting physical and MAC-layer parameters (transmission power, LoRaWAN spreading factor, channel, ACK retry count) to jointly optimize link quality and energy cost [7], [9]. - Activation: $0.7\, h_{\text{conn}} - 0.3\, h_{\text{ene}} > 0.25$

(negative Energexin weighting: power-increasing adaptations suppressed under critical energy constraint) - Outputs: {$\text{phy\_params}(t+1)$, $\Delta P_{\text{tx}}(t)$, $\text{sf\_selection}(t)$}

F2-D — NetworkRejoinAgent. Manages the complete network rejoin procedure following connectivity loss, implementing exponential backoff and multi-channel scanning to maximize reconnection probability within minimum energy budget; past successful rejoin sequences are stored as engrams [10] and retrieved via PSM to warm-start the reconnection procedure - Activation: $h_{\text{conn}} > 0.70 \wedge h_{\text{res}} > 0.40$ (Resiliencin co-condition distinguishes genuine connectivity loss from transient quality degradation) - Inputs include $\varepsilon_{\text{rejoin}}$ (PSM-retrieved engrams of past successful rejoin sequences) - Outputs: {$\text{rejoin\_status}(t)$, $T_{\text{backoff}}(t)$, $\text{channel\_selection}(t)$, engram_record} - Cost: high communication; $T^{\text{live}} = 20\text{–}100$ cycles, $T^{\text{cool}} = 30$ cycles

2.3. Family F3 — Energy and Power Management Agents (Energexin-coupled)

Family F3 operationalizes S-AI-IoT's formal energy management layer, implementing the duty cycle modulation mechanism of Section 3.4.1 and the energy-connectivity trade-off encoded in the cross-inhibition structure of Energexin$h_{\text{ene},i}(t)$.

F3-A — DutyCycleControllerAgent. Primary energy management agent, continuously and proportionally adjusting the node's active window $\delta_i^{\text{active}}(t)$ in response to Energexin. Operates as a persistent regulatory controller (active whenever $h_{\text{ene}} > 0$) [11]. - Control law:

$$\delta_i^{\text{active,eff}}(t) = \max\big(\delta_{\min} + (1 - h_{\text{ene},i})(\delta_{\max} - \delta_{\min}),\ \delta_{\min} + \kappa_{\text{res}}\, h_{\text{res},i}(t)\big)$$

(resilience override coefficient κ_{res} ensures duty cycle reduction is overridden when fault-recovery or sensing urgency demands extended wakefulness) - Outputs: {$\delta_i^{\text{active,eff}}(t)$, $T_{\text{sleep}}(t) = T_{\text{cycle}} - \delta_i^{\text{active,eff}}(t)$, $\text{wakeup\_schedule}(t)$} - Cost: $\mathcal{O}(1)$ (three arithmetic operations), 6 scalar registers

F3-B — TransmissionPowerControlAgent. Regulates radio transmission power $P_{\text{tx},i}(t)$ as a decreasing function of Energexin, while ensuring power reductions do not degrade connectivity below minimum required quality level. - Activation: $h_{\text{ene}} - 0.5\, h_{\text{conn}} > 0.20$ (negative Connectin weighting: power reductions suppressed when connectivity is already compromised) - Output law: $P_{\text{tx,new}}(t) = \max(P_{\text{tx}}^{\min}, P_{\text{tx}}^{\text{current}} \cdot (1 - \alpha_P\, h_{\text{ene}}(t)))$

F3-C — ComputationOffloadingAgent. Identifies computationally intensive local tasks eligible for remote execution at the gateway tier [13], implementing a cost-benefit analysis comparing local execution energy against communication offloading cost. - Activation: $h_{\text{ene}} > 0.55 \wedge h_{\text{conn}} < 0.60$ (Connectin guard: offloading not attempted when link quality is insufficient for reliable task transmission) - Outputs: {$\mathcal{T}_{\text{offload}}(t)$, $\mathcal{T}_{\text{local,remain}}(t)$, $\text{offload\_request}(t)$}

F3-D — HarvestingForecastAgent. Lightweight prediction model for harvested power availability (autoregressive over sliding window τ) [14], enabling proactive duty cycle scheduling that anticipates energy abundance and scarcity periods [9]. - Activation: $h_{\text{ene}} > 0.30 \wedge P_{\text{harv,var}}(t) > \sigma_{\text{harv}}^{\text{thr}}$ - Outputs: {$\hat{P}_{\text{harv}}(t+1{:}\,t+H)$, $\text{proactive\_schedule}(t)$} ($H$-step-ahead forecast submitted to F3-A)

2.4. Family F4 — Security, Trust, and Compliance Agents (Normin-coupled)

Family F4 operationalizes the integrated security architecture of Section 4.7, translating Normin elevations $h_{\text{norm},i}(t)$ into concrete security management, threat response, and compliance enforcement actions, inheriting the hormonal security framework of S-AI-Cyber **[15]** adapted to the IoT compliance landscape.

F4-A — SecurityComplianceAgent. Monitors the node's compliance state against applicable security and privacy standards, initiates corrective actions (certificate renewal, key rotation, policy synchronization) when compliance scores fall below threshold. - Activation: $h_{\text{norm}} > 0.30$ (moderate threshold for early-stage detection before escalation to security incident) - Outputs: {compliance_report(t), remediation_plan(t), $\Delta E_{\text{norm}}(t)$, engram_record}

F4-B — IntrusionDetectionAgent. Monitors inter-node communication patterns and local behavioral signals for deviations indicative of adversarial activity, implementing the artificial immune self/non-self discrimination[16] of Section 3.4.3. - Activation: $h_{\text{norm}} > 0.40 \vee \delta_{\text{beh},i}(t) > 0.35$ (disjunctive condition: fast-path activation by direct behavioral anomaly detection before Normin has integrated the signal to threshold level) - Outputs: {alert_level$(t) \in [0,1]$, threat_vector(t), $\Delta E_{\text{norm}}(t)$, quarantine_directive$(t)$}(quarantine directive may instruct F2-D and F2-C to isolate suspect neighboring nodes from the hormonal diffusion graph)

F4-C — FirmwareUpdateAgent. Manages firmware security patch lifecycle with a multi-criterion scheduling policy. - Activation: $h_{\text{norm}} > 0.35 \wedge h_{\text{sens}} < 0.40 \wedge h_{\text{conn}} < 0.40 \wedge h_{\text{ene}} < 0.50$ (multi-dimensional guard conditions ensuring updates are scheduled only during low-sensing-load, adequate-connectivity, sufficient-energy operational windows — absent from conventional IoT OTA update systems) - Cost: high CPU and communication; $T^{\text{cool}} = 500$ cycles (prevent update thrashing)

F4-D — PrivacyGuardianAgent. Enforces data minimization and purpose-limitation obligations at the node-level outgoing data filter of the DISNL, applying privacy policies encoded in the node's compliance configuration. - Activation: $h_{\text{norm}} > 0.25 \vee$ privacy_flag$(t) =$ true (direct activation trigger independent of Normin level for all sensitive data transmissions) - Outputs: {$O_i^{\text{filtered}}(t)$, privacy_audit_log(t), engram_record}

2.5. Family F5 — Resilience and Self-Healing Agents (Resiliencin-coupled)

Family F5 addresses fault detection, isolation, and recovery across all operational dimensions of the IoT node, driven by Resiliencin$h_{\text{res},i}(t)$, embodying the self-healing design principle that distinguishes resilient IoT infrastructure from brittle sensor networks [17].

F5-A — FaultIsolationAgent. Identifies the root cause of detected operational anomalies and isolates the affected subsystem to prevent fault propagation, using a structured fault taxonomy covering hardware faults, software exceptions, communication failures, and environmental disturbances. - Activation: $h_{\text{res}} > 0.35$; inputs include $\varepsilon_{\text{fault}}$ (PSM-retrieved engrams from past fault episodes with similar hormonal signatures)[18] - Outputs: {fault_class(t), fault_locus(t), isolation_directive(t), $\Delta E_{\text{res}}(t)$}

F5-B — SelfHealingAgent. Executes the recovery procedure prescribed by F5-A, selecting from a repertoire of recovery actions (software restart, hardware reset, redundant resource activation, gateway escalation) guided by PSM-retrieved engrams from past recovery episodes. - Activation:

$h_{\text{res}} > 0.45 \wedge \text{fault\_class}(t) \neq \varnothing$ - Outputs: $\{\text{recovery\_action}(t), \text{recovery\_status}(t), \Delta\mathbf{h}_{\text{post}}(t), \text{engram\_record}\}$ - Variable cost; $T^{\text{cool}} = 20$ cycles

F5-C — RedundancyManagementAgent. Maintains a dynamic inventory of available redundancy resources (backup sensor channels, alternative radio interfaces, neighbor-relay paths) and activates them upon primary resource failure. - Activation: $h_{\text{res}} > 0.50 \wedge \text{primary\_resource\_failed}(t)$ - Outputs: $\{\text{backup\_activation\_list}(t), \text{resource\_inventory}(t+1), \text{engram\_record}\}$

F5-D — SystemCheckpointAgent. Periodically records the node's complete operational state in persistent storage as a recovery checkpoint, with checkpointing frequency adaptively increased when Resiliencin is elevated. - Activation: $t \bmod T_{\text{ckpt}}(h_{\text{res}}) = 0$ with $T_{\text{ckpt}}(h_{\text{res}}) = T_{\text{ckpt}}^{\max}(1 - h_{\text{res}}) + T_{\text{ckpt}}^{\min} h_{\text{res}}$ - Outputs: $\{\text{checkpoint\_id}(t), \text{storage\_location}(t), \text{checkpoint\_hash}(t)\}$

2.6. Family F6 — Governance, Coordination, and Explainability Agents (Multi-hormonal)

Family F6 addresses cross-cutting architectural concerns spanning multiple hormonal dimensions. Their utility functions are multivariate and their activation conditions reflect composite hormonal states, providing the meta-regulatory layer ensuring coherence, auditability, and scalability of the overall system.

F6-A — EngineWatchdogAgent. Monitors the computational health of the Hormonal Engine, detecting instabilities (oscillations, runaway concentration levels, numerical errors) and triggering corrective interventions (hormonal state projection onto $[0,1]^5$, parameter reset, engine restart). - Activation: $\mathbb{1}[\exists k: h_k \notin (0.01, 0.99) \vee \| \mathbf{h}(t) - \mathbf{h}(t-1) \|_2 > \epsilon_{\text{stab}}]$ - Cost: $\mathcal{O}(K)$ (five comparisons per cycle); persistent

F6-B — EngineSchedulerAgent. Manages scheduling of all agents within the available computational budget, resolving resource conflicts through a priority ordering derived from the current hormonal state. - Activation: persistent at every orchestration cycle - Inputs: $\{\mathcal{A}^*(t), \{c_a\}_{a \in \mathcal{A}^*(t)}, B_{\text{node}}(t), \mathbf{h}(t)\}$ - Outputs: $\{\text{schedule}(t), \text{priority\_order}(t), \text{budget\_allocation}(t)\}$ - Cost: $\mathcal{O}(|\mathcal{A}^*(t)| \log |\mathcal{A}^*(t)|)$

F6-C — ExplainabilityAgent. Generates and stores formal explainability artifacts, formatting completed decision records as extended IoT engrams and producing human-readable decision justifications on demand [20,21,22]. - Activation: $u_{\text{F6C}}(\mathbf{h}) = \mathbb{1}[|\mathcal{A}^*(t)| > 0]$ - Outputs: $\mathcal{O}_{\text{F6C}} = \{\varepsilon_m^{\text{IoT}} = \langle \mathbf{h}(t), \mathbf{x}(t), \mathbf{a}(t), \mathbf{o}(t), \text{topology}(t), \text{energy\_budget}(t) \rangle, \text{justification}(t)\}$ - Cost: $\mathcal{O}(|\mathcal{A}^*(t)|)$; $T^{\text{live}} = 1$ cycle, $T^{\text{cool}} = 0$

F6-D — ClusterCoordinationAgent. Manages the two-tier hierarchical orchestration of Section 4.2.3 of Article I [1], supporting multi-tier IoT platform coordination [19]: at the node tier, compresses and transmits the node's hormonal state report to the gateway; at the gateway tier, aggregates node reports and distributes cluster directives. - Activation: $t \bmod T_{\text{report}}(h_{\text{conn}}(t)) = 0$ (self-regulating reporting interval, decreasing with improving connectivity) - Outputs: $\{\bar{\mathbf{h}}_{\text{cluster}}(t) = \frac{1}{|C|} \sum_{i \in C} \mathbf{h}_i(t)$ (gateway), $\text{compressed\_report}(t)$ (node), $\text{cluster\_directives}(t)\}$ - Cost: $\mathcal{O}(K)$ per report (5 scalars).

3. MATHEMATICAL MODELING OF THE S-AI-IOT FRAMEWORK

This section develops the complete mathematical formalization of S-AI-IoT, organized in eight subsections: notation and graph-theoretic substrate (§3.1); observation vectors and hormonal state variables (§3.2); emission function architecture (§3.3); reaction-diffusion dynamics (§3.4); stability analysis and discretization (§3.5); duty cycle modulation (§3.6); primal-dual constrained orchestration (§3.7); and distributed symbolic memory formalism (§3.8).

Notation. $\|\cdot\|_2$ denotes the Euclidean norm, $\|\cdot\|_F$ the Frobenius norm, $\lambda_{\max}(\cdot)$ the spectral radius, $\mathbb{1}[\cdot]$ the indicator function, and $\mathrm{clip}_{[a,b]}(x) = \max(a, \min(b, x))$ the clipping operator. All hormone values are dimensionless scalars in [0,1]; all observation components are normalized to [0,1] by the DISNL; all time indices t refer to discrete orchestration cycles of duration Δt.

3.1. Notations, Graph-Theoretic Substrate, and Multi-Layer IoT Graph

3.1.1 Node Set and Deployment Tiers

Let $\mathcal{V} = \{1, \ldots, N\}$ denote the complete set of IoT nodes, partitioned as $\mathcal{V} = \mathcal{V}_{\text{node}} \cup \mathcal{V}_{\text{gw}} \cup \mathcal{V}_{\text{cloud}}$. Each node i at time t is characterized by: observation vector $O_i(t) \in [0,1]^d$, hormonal state $\mathbf{h}_i(t) \in [0,1]^5$, agent activation vector $\mathbf{x}_i(t) \in \{0,1\}^{|\mathcal{A}|}$, and resource state $\chi_i(t) = (b_i(t), \ell_{\text{cpu},i}(t), \delta_i^{\text{active}}(t))$.

3.1.2 Dynamic Communication Graph

The communication topology is represented as a time-varying weighted undirected graph $G(t) = (\mathcal{V}, \mathcal{E}(t), W(t))$. Link weight [6]: $w_{ij}(t) = \mathrm{RSSI}_{ij}^{\text{norm}}(t) \cdot (1 - \mathrm{loss}_{ij}(t))$

Links with $w_{ij}(t) < w_{\min}$ are removed, giving dynamic neighborhood $\mathcal{N}_i(t)$. The weighted graph Laplacian $L(t) \in \mathbb{R}^{N \times N}$ [12]:

$$L_{ij}(t) = \begin{cases} \sum_{j' \in \mathcal{N}_i(t)} w_{ij'}(t) & \text{if } i = j \\ -w_{ij}(t) & \text{if } (i,j) \in \mathcal{E}(t) \\ 0 & \text{otherwise} \end{cases}$$

$L(t)$ is positive semi-definite for all t; its spectral properties govern the hormonal diffusion rate across the network.

3.1.3 Hormone Index Set and Timescale Hierarchy

Five canonical hormones are indexed by $k \in \mathcal{K} = \{\text{sens,conn,ene,res,norm}\}$, each with characteristic timescale $\tau_k > 0$, decay rate $\lambda_k > 0$, diffusion coefficient $D_k \geq 0$, emission delay $\delta_k \geq 0$, and resource coupling $\rho_k \geq 0$. The timescale hierarchy $\tau_{\text{sens}} \leq \tau_{\text{conn}} \leq \tau_{\text{ene}} \leq \tau_{\text{res}} \leq \tau_{\text{norm}}$ formalizes the multi-timescale regulatory structure of Section 3.6.5 of Article I [1]. The cross-inhibition matrix $\Gamma \in \mathbb{R}_{\geq 0}^{5 \times 5}$ has $\gamma_{kk} = 0$; non-zero entries encode physiologically motivated regulatory trade-offs specified in §3.4.2.

3.2 IoT Observation Vectors and Hormonal State Variables

3.2.1 Observation Vector Structure

The normalized observation vector is a concatenation of five domain-specific sub-vectors:

$$O_i(t) = (O_i^{\text{sens}}(t),\ O_i^{\text{conn}}(t),\ O_i^{\text{ene}}(t),\ O_i^{\text{res}}(t),\ O_i^{\text{norm}}(t))$$

Sensor quality sub-vector: $O_i^{\text{sens}} = (\text{snr\_inv}_i,\ \text{missing\_rate}_i,\ \text{drift\_score}_i,\ \text{cal\_age\_norm}_i)$, where $\text{snr\_inv}_i = 1 - \text{SNR}_i/\text{SNR}_{\max}$.

Connectivity sub-vector: $O_i^{\text{conn}} = (\text{packet\_loss}_i,\ \text{rssi\_low}_i,\ \text{latency\_norm}_i,\ \text{jitter\_norm}_i)$, where $\text{rssi\_low}_i = 1 - \bar{w}_i(t)$.

Energy sub-vector: $O_i^{\text{ene}} = (1 - b_i(t),\ \ell_{\text{cpu},i},\ P_i/P_{\max},\ 1 - \hat{P}_{\text{harv},i}/P_{\max})$, encoding battery depletion, CPU load, power overconsumption, and harvesting unavailability as positive stress signals.

Resilience sub-vector: $O_i^{\text{res}} = (f_i^{\text{norm}},\ \Delta\tau_{\text{rec},i}^{\text{norm}},\ 1 - r_i(t))$.

Compliance sub-vector: $O_i^{\text{norm}} = (1 - s_{\text{fw},i},\ 1 - s_{\text{enc},i},\ 1 - s_{\text{cert},i},\ \delta_{\text{beh},i})$, all formulated so increasing values encode increasing compliance stress.

3.2.2 Hormonal State Vector

$$\mathbf{h}_i(t) = (h_{\text{sens},i}(t),\ h_{\text{conn},i}(t),\ h_{\text{ene},i}(t),\ h_{\text{res},i}(t),\ h_{\text{norm},i}(t)) \in [0,1]^5$$

with $h_{k,i}(t) = 0$ signifying complete absence of regulatory stress and $h_{k,i}(t) = 1$ maximal stress along dimension k.

3.3 Emission Functions: IoT-Specific Aggregation and Logistic Saturation

3.3.1 General Emission Architecture

For each hormone k and node i: $E_{k,i}(t) = \sigma(a_k\,\phi_k(O_i(t)) + b_k) \cdot \big(1 - h_{k,i}(t)\big)$ where $\sigma(x) = 1/(1 + e^{-x})$, $a_k > 0$ is the emission gain, $b_k \in \mathbb{R}$ is the emission bias, and the saturation factor $(1 - h_{k,i}(t))$ prevents unbounded accumulation [1]. Default calibration: $a_k = 5,\ b_k = -2.5$ (half-maximal emission at $\phi_k = 0.5$).

3.3.2 Aggregation Functions

All aggregation functions are convex combinations of clipped sub-components ($\alpha + \beta + \cdots = 1$, all weights > 0):

Sensorin: $\phi_{\text{sens}} = \alpha_s\,\text{snr\_inv} + \beta_s\,\text{missing\_rate} + \gamma_s\,\text{drift\_score} + \delta_s\,\text{cal\_age\_norm}$ (default uniform: 0.25 each)

Connectin: $\phi_{\text{conn}} = \alpha_c\,\text{packet\_loss} + \beta_c\,\text{rssi\_low} + \gamma_c\,\text{latency\_norm} + \delta_c\,\text{jitter\_norm}$ (dominant weights: $\alpha_c \approx 0.40$ for packet loss, $\beta_c \approx 0.35$ for RSSI)

Energexin: $\phi_{\text{ene}} = \alpha_e(1 - b_i) + \beta_e\,\ell_{\text{cpu}} + \gamma_e\,(P/P_{\max}) + \delta_e\,(1 - \hat{P}_{\text{harv}}/P_{\max})$ (battery depletion dominant: $\alpha_e \approx 0.50$; harvesting term zero for battery-only nodes)

Resiliencin: $\phi_{\text{res}} = \alpha_r\,f_i^{\text{norm}} + \beta_r\,\Delta\tau_{\text{rec}}^{\text{norm}} + \gamma_r\,(1 - r_i)$ (fault count dominant: $\alpha_r \approx 0.55$)

Normin: $\phi_{\text{norm}} = \alpha_n(1 - s_{\text{fw}}) + \beta_n(1 - s_{\text{enc}}) + \gamma_n(1 - s_{\text{cert}}) + \delta_n\,\delta_{\text{beh}}$ (behavioral deviation elevated in security-critical deployments: $\delta_n \approx 0.40$)

3.4 Hormonal Reaction-Diffusion Dynamics

3.4.1 Continuous-Time Governing Equation

The temporal evolution of hormone k at node i is governed by:

$$\tau_k \frac{d}{dt} h_{k,i}(t) = \underbrace{-\lambda_k\, h_{k,i}(t)}_{\text{decay}} + \underbrace{E_{k,i}(t-\delta_k)}_{\text{delayed emission}} - \underbrace{\sum_{m \neq k} \gamma_{km}\; h_{k,i}\, h_{m,i}}_{\text{cross-inhibition}} + \underbrace{D_k \sum_{j \in \mathcal{N}_i(t)} w_{ij}\,(h_{k,j} - h_{k,i})}_{\text{graph diffusion}}$$
$$- \underbrace{\rho_k\, \chi_i(t)\, h_{k,i}}_{\text{resource coupling}} + \underbrace{\sigma_{\eta,k}\, \eta_{k,i}(t)}_{\text{stochastic perturbation}}$$

Each term has precise physiological and operational interpretation: **decay** ensures hormonal elevations are transient; **delayed emission** drives the hormone upward with latency δ_k; **cross-inhibition** (quadratic, receptor-kinetics motivated [11]) prevents simultaneous saturation of incompatible hormonal axes; **graph diffusion** propagates hormonal information spatially, with link weight $w_{ij}(t)$ modulating propagation rate (formal analog of vascular perfusion); **resource coupling** implements energy-proportional hormonal damping ($\rho_k = 0$ for Normin to ensure security responses are not energy-gated); **stochastic perturbation**$\eta_{k,i}(t) \sim \mathcal{N}(0,1)$ models environmental stochasticity.

3.4.2 Cross-Inhibition Structure

Non-zero entries of Γ with default values (calibrated on SAI-UT+):

Pair	Value	Operational Semantics
$\gamma_{\text{ene,conn}} = 0.30$	Energy crisis suppresses connectivity maintenance	Prevents energy-costly communication protocols under battery depletion
$\gamma_{\text{ene,sens}} = 0.15$	Energy crisis moderately suppresses sensor stress	Reduces sensor maintenance urgency under critical energy constraint
$\gamma_{\text{res,sens}} = 0.20$	Recovery suppresses sensor stress signaling	Sensor recalibration deferred during active fault recovery
$\gamma_{\text{res,conn}} = 0.25$	Recovery suppresses connectivity expansion	Prevents energy-intensive network rejoin during fault recovery
$\gamma_{\text{norm,ene}} = 0.35$	Security urgency suppresses energy conservation	Critical security responses not blocked below minimum response level

3.4.3 Compact Vector-Matrix Formulation

Let $\mathbf{H}_k(t) = (h_{k,1}(t), \dots, h_{k,N}(t))^T \in [0,1]^N$. The network-wide dynamics:

$$\tau_k\, \dot{\mathbf{H}}_k(t) = -(\lambda_k + \rho_k\, \mathrm{diag}(\boldsymbol{\chi}(t)))\, \mathbf{H}_k(t) + \mathbf{E}_k(t-\delta_k) - \sum_{m \neq k} \gamma_{km}\; \mathbf{H}_k(t) \odot \mathbf{H}_m(t)$$
$$- D_k\, L(t)\, \mathbf{H}_k(t) + \sigma_{\eta,k}\, \boldsymbol{\eta}_k(t)$$

where $\odot$ denotes element-wise multiplication.

3.5 Stability Analysis and Discretization

3.5.1 Existence of Hormonal Equilibrium

Proposition 3.1. For constant emission $\mathbf{E}_k^*$, constant resource state $\boldsymbol{\chi}^*$, and constant graph G^*, there exists at least one equilibrium $\mathbf{H}_k^* \in [0,1]^N$ satisfying the zero-RHS condition.
Proof: the right-hand side defines a continuous vector field on the compact convex set $[0,1]^N$; inward-pointing at both boundaries; existence follows from Brouwer's fixed-point theorem.

3.5.2 Lyapunov Stability

Theorem 3.1 (Global Asymptotic Stability of the Deterministic Frozen Dynamics).

Define the Lyapunov candidate [8]:

$$V\big(\mathbf{H}(t)\big) = \frac{1}{2}\sum_{k\in\mathcal{K}} \tau_k \,\|\mathbf{H}_k(t) - \mathbf{H}_k^*\|_2^2$$

Under the stability condition:$\lambda_k > D_k\,\lambda_{\max}(L^*) + \sum_{m\neq k}\gamma_{km} + \rho_k\|\boldsymbol{\chi}^*\|_\infty \quad \forall\, k \in \mathcal{K}$the deterministic equilibrium $\mathbf{H}^*$ is globally asymptotically stable on $[0,1]^{5N}$.

Proof. Let $\widetilde{\mathbf{H}}_k(t) = \mathbf{H}_k(t) - \mathbf{H}_k^*$ denote the error state for hormone k. In the deterministic setting ($\sigma_{\eta,k} = 0$) and under constant emission $\mathbf{E}_k^*$, constant resource state $\boldsymbol{\chi}^*$, and constant graph L^*, the error dynamics satisfy: $\tau_k\,\dot{\widetilde{\mathbf{H}}}_k = -\lambda_k\widetilde{\mathbf{H}}_k - D_kL^*\widetilde{\mathbf{H}}_k - \rho_k\,\mathrm{diag}(\boldsymbol{\chi}^*)\widetilde{\mathbf{H}}_k - \sum_{m\neq k}\gamma_{km}\,(\mathbf{H}_k \odot \mathbf{H}_m - \mathbf{H}_k^* \odot \mathbf{H}_m^*)$
Computing $\dot{V}$ along trajectories: $\dot{V} = \sum_{k\in\mathcal{K}}\tau_k\;\widetilde{\mathbf{H}}_k^T\dot{\widetilde{\mathbf{H}}}_k$We bound each of the four terms separately.

Term 1 — Decay.

$$\widetilde{\mathbf{H}}_k^T\big(-\lambda_k\widetilde{\mathbf{H}}_k\big) = -\lambda_k\|\widetilde{\mathbf{H}}_k\|_2^2 \le 0$$

strictly negative whenever $\widetilde{\mathbf{H}}_k \neq 0$.

Term 2 — Diffusion. Since $L^*$is positive semi-definite [12]: $\widetilde{\mathbf{H}}_k^T\big(-D_kL^*\widetilde{\mathbf{H}}_k\big) = -D_k\,\widetilde{\mathbf{H}}_k^TL^*\widetilde{\mathbf{H}}_k \le 0$
Moreover, $\widetilde{\mathbf{H}}_k^TL^*\widetilde{\mathbf{H}}_k \le \lambda_{\max}(L^*)\|\widetilde{\mathbf{H}}_k\|_2^2$, giving the destabilizing upper bound $D_k\,\lambda_{\max}(L^*)\|\widetilde{\mathbf{H}}_k\|_2^2$.

Term 3 — Cross-inhibition. Using the identity: $\mathbf{H}_k \odot \mathbf{H}_m - \mathbf{H}_k^* \odot \mathbf{H}_m^* = \widetilde{\mathbf{H}}_k \odot \mathbf{H}_m + \mathbf{H}_k^* \odot \widetilde{\mathbf{H}}_m$
and applying Young's inequality ($|ab| \le \frac{1}{2}(a^2+b^2)$) together with $\mathbf{H}_k, \mathbf{H}_m \in [0,1]^N$:

$$\big|\widetilde{\mathbf{H}}_k^T\big(\widetilde{\mathbf{H}}_k \odot \mathbf{H}_m\big)\big| \le \|\widetilde{\mathbf{H}}_k\|_2^2$$

$$\big|\widetilde{\mathbf{H}}_k^T\big(\mathbf{H}_k^* \odot \widetilde{\mathbf{H}}_m\big)\big| \le \frac{1}{2}\Big(\|\widetilde{\mathbf{H}}_k\|_2^2 + \|\widetilde{\mathbf{H}}_m\|_2^2\Big)$$

Hence the cross-inhibition contribution to $\dot{V}$ is bounded above by $\sum_k\sum_{m\neq k}\gamma_{km}\,\|\widetilde{\mathbf{H}}_k\|_2^2$, up to symmetric redistribution over the index pair (k,m).

Term 4 — Resource coupling.

$$\widetilde{\mathbf{H}}_k^T\big(-\rho_k\,\mathrm{diag}(\boldsymbol{\chi}^*)\widetilde{\mathbf{H}}_k\big) \le \rho_k\|\boldsymbol{\chi}^*\|_\infty\|\widetilde{\mathbf{H}}_k\|_2^2$$

as a destabilizing upper bound.

Aggregation. Collecting all four terms:

$$\dot{V} \le -\sum_{k\in\mathcal{K}}\left[\lambda_k - D_k\,\lambda_{\max}(L^*) - \sum_{m\neq k}\gamma_{km} - \rho_k\|\boldsymbol{\chi}^*\|_\infty\right]\|\widetilde{\mathbf{H}}_k\|_2^2$$

Under condition (SC), each bracketed coefficient is strictly positive. Denoting:

$$\mu_k = \lambda_k - D_k\, \lambda_{\max}(L^*) - \sum_{m \neq k} \gamma_{km} - \rho_k \|\boldsymbol{\chi}^*\|_\infty > 0$$

and $\mu = \min_k \mu_k/\tau_k > 0$, we obtain:

$$\dot{V} \leq -\mu \cdot \sum_{k \in \mathcal{K}} \tau_k \, \|\widetilde{\mathbf{H}}_k\|_2^2 = -2\mu\, V(\mathbf{H}(t))$$

By Lyapunov's direct method [8], $V(t) \leq V(0)\, e^{-2\mu t} \to 0$ as $t \to \infty$, establishing global asymptotic stability of $\mathbf{H}^*$ on $[0,1]^{5N}$. The forward invariance of $[0,1]^{5N}$ follows from the inward-pointing property of the vector field at the boundary, established in Proposition 3.1. □

Remark 3.1. The stability condition (SC) provides a deployability criterion verifiable a priori from network parameters alone: the decay rate λ_k of each hormone must exceed the sum of three destabilizing contributions — network diffusion (governed by the spectral radius of the communication graph Laplacian $\lambda_{\max}(L^*)$), cross-inhibitory coupling strength ($\sum_{m \neq k} \gamma_{km}$), and resource damping ($\rho_k \| \boldsymbol{\chi}^* \|_\infty$). Crucially, $\rho_{\text{norm}} = 0$ ensures that the security hormone Normin is never energy-gated, as required by the security architecture of Section 4.7 of Article I [1].

3.5.3 Numerical Stability Condition

For explicit Euler discretization, the numerical stability condition is:

$$\Delta t < \min_{k \in \mathcal{K}} \frac{2\tau_k}{\lambda_k + 2D_k\, \lambda_{\max}(L) + \rho_k \| \boldsymbol{\chi} \|_\infty + \sum_{m \neq k} \gamma_{km}}$$

This provides a concrete upper bound on the admissible discretization step verifiable before deployment.

3.5.4 Euler-Maruyama Discretization with Projection

The complete discrete-time stochastic update rule:

$$h_{k,i}^{t+1} = \text{clip}_{[0,1]} \left[h_{k,i}^t + \frac{\Delta t}{\tau_k} \left(-\lambda_k h_{k,i}^t + E_{k,i}^{t-\delta_k} - \sum_{m \neq k} \gamma_{km}\, h_{k,i}^t h_{m,i}^t + D_k \sum_{j \in \mathcal{N}_i^t} w_{ij}^t (h_{k,j}^t - h_{k,i}^t) - \rho_k \chi_i^t h_{k,i}^t \right) + \sqrt{\Delta t}\, \sigma_{\eta,k}\, \xi_{k,i}^t \right]$$

with $\xi_{k,i}^t \sim \mathcal{N}(0,1)$ i.i.d. The $\text{clip}_{[0,1]}$ projection maintains the boundedness of the hormonal state within in the discrete-time implementation.

3.6 Duty Cycle Modulation via Energexin

3.6.1 Base Duty Cycle Modulation Law

$$\delta_i^{\text{base}}(t) = \delta_{\min} + (1 - h_{\text{ene},i}(t))(\delta_{\max} - \delta_{\min})$$

When $h_{\text{ene}} = 0$ (abundant energy): $\delta_i^{\text{base}} = \delta_{\max}$ (full wakefulness). When $h_{\text{ene}} = 1$ (severe deficit): $\delta_i^{\text{base}} = \delta_{\min}$ (maximum sleep).

3.6.2 Multi-Hormonal Override Mechanism

$$\delta_i^{\text{active}}(t) = \max\big(\delta_i^{\text{base}}(t),\ \delta_{\min} + \kappa_{\text{res}}\, h_{\text{res},i}(t)(\delta_{\max} - \delta_{\min}),\ \delta_{\min} + \kappa_{\text{sens}}\, h_{\text{sens},i}(t)(\delta_{\max} - \delta_{\min})\big)$$

The max operator implements disjunctive activation: the effective duty cycle is the maximum of the energy-driven baseline and urgency-driven overrides ($\kappa_{\text{res}}, \kappa_{\text{sens}} \in (0,1]$), formally implementing the HPA axis cortisol stress-response override of the biological circadian sleep drive [11].

3.6.3 Energy Feedback Loop

Sleep duration: $T_{\text{sleep},i}(t) = T_{\text{cycle}}(1 - \delta_i^{\text{active}}(t))$. Total active-window energy consumption:

$$E_{\text{active},i}(t) = P_{\text{cpu}}\, T_{\text{cycle}}\, \delta_i^{\text{active}}(t) + P_{\text{tx}}\, T_{\text{tx},i}(t) + P_{\text{sense}}\, T_{\text{sense},i}(t)$$

This energy model feeds back into the EnergyGland'sϕ_{ene} computation, closing the energy management feedback loop.

3.7 Primal-Dual Constrained Orchestration

3.7.1 Agent-Selection Optimization Problem

At each orchestration cycle, the IoT-MetaAgent solves:

$$\max_{\mathbf{x}(t)\in\{0,1\}^{|\mathcal{A}|}} \sum_{a\in\mathcal{A}} u_a\left(\mathbf{h}(t)\right) x_a(t) \quad \text{subject to} \quad \sum_{a\in\mathcal{A}} c_a\ x_a(t) \le B_{\text{node}}(t)$$

This is a 0-1 knapsack problem, NP-hard in general but tractable in the IoT context for three reasons. First, the agent registry is small and fixed ($|\mathcal{A}| = 24$). Second, the hormonal significance filter reduces active candidates to $\left|\mathcal{A}^{\text{sig}}(t)\right| \ll |\mathcal{A}|$ — typically 4 to 8 agents per cycle under normal operational conditions — making exhaustive enumeration feasible in $O\left(2^{|\mathcal{A}^{\text{sig}}|}\right)$ operations. Third, the standard dynamic programming solution to the 0-1 knapsack runs in $O\left(\left|\mathcal{A}^{\text{sig}}\right| \cdot \lceil B_{\text{node}}\rceil\right)$ time with respect to the discretized budget, which for $\left|\mathcal{A}^{\text{sig}}\right| \le 8$ and a budget discretized into at most 100 levels yields fewer than 800 operations per cycle — well within the real-time constraints of gateway-class hardware. PSM-retrieved engrams further reduce solve time by providing warm-start near-optimal initializations for recurring scenarios, allowing early termination of the dynamic programming search.

3.7.2 Convex Relaxation and Primal-Dual Dynamics

The LP relaxation ($x_a \in [0,1]$) yields the Lagrangian:

$$\mathcal{L}(\mathbf{x},\mu) = \sum_a u_a\,(\mathbf{h}(t))\,x_a - \mu\left(\sum_a c_a\ x_a - B_{\text{node}}(t)\right)$$

with dual variable $\mu \ge 0$ (shadow price of the budget constraint). Primal-dual update dynamics:

$$\dot{x}_a = \text{clip}_{[0,1]}\left[\alpha_x\,(u_a(\mathbf{h}) - \mu\, c_a)\right] \quad \forall\, a$$

$$\dot{\mu} = \text{clip}_{[0,\mu_{\max}]}\left[\alpha_\mu\left(\sum_a c_a\ x_a - B_{\text{node}}\right)\right]$$

For the convex LP relaxation, these gradient dynamics converge to the associated saddle point $(\mathbf{x}^*, \mu^*)$ under standard convex conditions [7].

The integer solution is recovered by threshold rounding: $x_a^*(t) = \mathbf{1}[x_a^{LP} > 0.5]$. This rounding is guaranteed to produce a feasible integer solution whenever the LP relaxation yields a half-integer gap below 0.5 per agent — a condition empirically verified on the S-AI-IoT agent registry ($|A| = 24$) across all operational scenarios of the SAI-UT+ testbench (Article III). In the worst case, the integrality gap of the LP relaxation of the 0-1 knapsack is bounded by a factor of 2 [7], meaning the rounded solution achieves at least half the optimal integer utility — a bound tightened in practice by the warm-start initialization from PSM-retrieved engrams, which biases the LP solution toward previously optimal integer configurations.

3.7.3 Budget Adaptation via Energexin

$$B_{\text{node}}(t) = B_{\max}\,(1 - \beta_B\, h_{\text{ene},i}(t))$$

This coupling automatically tightens orchestration parsimony during energy deficit: fewer agents can be simultaneously active when Energexin is elevated, implementing the architectural energy-parsimony feedback loop.

3.7.4 Utility Function Specifications

Three utility function forms:
Linear blend (multiple hormonal couplings): $u_a(\mathbf{h}) = \sum_k w_{a,k}\ h_k, \sum_k w_{a,k} = 1$
Signed linear (excitatory and inhibitory couplings): $u_a(\mathbf{h}) = \sum_{k\in\mathcal{K}^+} w_{a,k}\ h_k - \sum_{m\in\mathcal{K}^-} w_{a,m}\ h_m$
Threshold-guarded (hard activation prerequisites): $u_a(\mathbf{h}) = u_a^{\text{base}}(\mathbf{h}) \cdot \prod_{k\in\mathcal{K}_a^{\text{guard}}} \mathbb{1}\,[h_k < \theta_k^{\text{guard}}]$

3.8 Distributed Symbolic Memory: Formal Engram Specification

3.8.1 Engram Structure and Storage

The extended IoT engram [10]:$\varepsilon_m^{\text{IoT}} = \langle \mathbf{h}^{(m)},\ \mathbf{x}^{(m)},\ \mathbf{a}^{(m)},\ \mathbf{o}^{(m)},\ \mathcal{T}^{(m)},\ \beta^{(m)} \rangle$where: $\mathbf{h}^{(m)} \in [0,1]^5$ is the hormonal context; $\mathbf{x}^{(m)} \in \{0,1\}^{|\mathcal{A}|}$ the agent activation vector; $\mathbf{a}^{(m)} \in \mathbb{R}^{d_a}$ the composite action vector; $\mathbf{o}^{(m)} \in \mathbb{R}^{d_o}$ the observed outcome; $\mathcal{T}^{(m)}$ the local network topology state (sparse adjacency); and $\beta^{(m)} \in [0,1]$ the energy budget level.
Engram written to PSM when the behavioral significance criterion is met:

$$\| \mathbf{x}^{(m)} - \mathbf{x}^{(m-1)} \|_1 \geq 1 \ \wedge \| \mathbf{o}^{(m)} \|_2 \geq \epsilon_{\text{sig}}$$

3.8.2 Content-Addressable Retrieval

Given query $\mathbf{h}_q \in [0,1]^5$, the top-K_{ret} engrams are retrieved by weighted cosine similarity:

$$\text{sim}(\mathbf{h}_q, \mathbf{h}^{(m)}) = \frac{\mathbf{h}_q^T W_{\text{ret}}\, \mathbf{h}^{(m)}}{\| \mathbf{h}_q \|_{W_{\text{ret}}} \| \mathbf{h}^{(m)} \|_{W_{\text{ret}}}}$$

where $W_{\text{ret}} = \text{diag}(w_{\text{sens}}, w_{\text{conn}}, w_{\text{ene}}, w_{\text{res}}, w_{\text{norm}})$ with default $w_k = 1 + h_{q,k}$, amplifying similarity contribution of currently elevated hormonal dimensions. Retrieved action components $\{\mathbf{a}^{(m_j)}\}_{j=1}^{K_{\text{ret}}}$ are presented to the IoT-MA as warm-start initializations. Retrieval cost: $\mathcal{O}(M_{\text{PSM}} \cdot 5)$ — fewer than 5000 multiply-accumulate operations for $M_{\max} = 1000$, feasible on all IoT hardware tiers.

3.8.3 Distributed Engram Propagation Protocol

When no local engram achieves sim $\geq \theta_{\text{ret}}$, node i broadcasts query $Q_i(t) = \langle \mathbf{h}_i(t), \theta_{\text{ret}}, K_{\text{ret}} \rangle$.

Neighbors respond with their top-K_{ret} matching engrams. Remote engrams are incorporated into the local PSM only if the engram-acceptance criterion is met:

$$\text{sim}(\mathbf{h}_i(t), \mathbf{h}^{(m_j)}) \geq \theta_{\text{ret}} \;\wedge\; \| \mathbf{o}^{(m_j)} \|_2 \geq \epsilon_{\text{sig}} \;\wedge\; b_i(t) \geq b_{\text{write\_min}}$$

where $b_{\text{write\_min}}$ ensures engram writing is not attempted during critical energy deficit.

Remark 3.3 (Engram Conflict Resolution upon Network Partition Rejoining). The engram-acceptance criterion above governs normal operation but does not address a specific failure scenario: network partition followed by reconnection. During a partition, two sub-populations of nodes may independently develop engrams with similar hormonal signatures $\mathbf{h}^{(m)}$ but divergent action vectors $\mathbf{a}^{(m)}$ and contradictory outcomes $\mathbf{o}^{(m)}$ for the same operational scenario. Upon partition healing, naive engram propagation could introduce behavioral inconsistency, degrading the acceleration factor of 2–4× announced in Section 3.8. S-AI-IoT resolves this through a **conflict-aware merge policy** triggered when the ClusterCoordinationAgent (F6-D) detects partition rejoining. Upon detection, the gateway-tier IoT-MA executes a three-step reconciliation procedure.First, engrams from the two sub-populations are clustered by hormonal signature similarity: engrams with $\text{sim}\left(\mathbf{h}^{(m)}, \mathbf{h}^{(m\prime)}\right) \geq \theta_{\text{ret}}$ are considered conflicting candidates.Second, within each conflict cluster, the engram with the higher outcome norm $\| \mathbf{o}^{(m)} \|_2$ is retained as the canonical engram, discarding lower-significance alternatives — the rationale being that higher-outcome engrams represent more informative behavioral episodes.Third, if two conflicting engrams have comparable outcome norms ($\left| \| \mathbf{o}^{(m)} \|_2 - \| \mathbf{o}^{(m\prime)} \|_2 \right| < \epsilon_{\text{sig}}$), both are retained with a conflict flag, and the IoT-MA presents both as candidate templates to the operator for supervised resolution during the next maintenance window.This policy guarantees that post-partition behavioral acceleration is bounded below: in the worst case where all engrams conflict, the system falls back to the uninformed prior — equivalent to no memory — without introducing behavioral instability. The acceleration factor of 2–4× is thus a conservative lower bound valid under normal operation and preserved under the conflict resolution policy for non-conflicting engrams.

3.8.4 Memory Capacity and Eviction Policy

The PSM is a fixed-capacity circular buffer of $M_{\max}$ engrams with LRU eviction and a significance override: engrams with $\| \mathbf{o}^{(m)} \|_2 \geq \epsilon_{\text{protect}}$ are protected for retention period T_{protect}. Memory footprint per engram: $5 + |\mathcal{A}| + d_a + d_o + 2|\mathcal{N}_i| + 1$ floating-point scalars ≈ 60–100 bytes. A PSM of $M_{\max} = 1000$ requires fewer than 100 KB — well within gateway-class and mid-tier IoT hardware capabilities.

3.8.5 Formal Explainability Guarantee

Proposition 3.2 (Zero-Cost Intrinsic Explainability). For any agent activation decision $\mathcal{A}^*(t)$, the corresponding engram $\varepsilon_m^{\text{IoT}}$ stored by the ExplainabilityAgent constitutes a causally faithful record of the complete decision context — hormonal state, activated agents, joint action, and operational context — requiring no post-hoc computation, no model approximation, and no additional overhead beyond the $\mathcal{O}(1)$ engram-writing operation.This proposition formally establishes intrinsic explainability (Property 4, Section 3.3.3), satisfying the XAI 2.0 architectural transparency criterion [20] and aligning with applicable regulatory frameworks governing AI auditability and data subject rights.

3.9 Summary of Mathematical Framework Parameters

Table 2. S-AI-IoT Mathematical Framework: Parameter Reference

Parameter	Symbol	Default	Range	Physical Interpretation
Decay rate	λ_k	0.1–0.3	$(0, \infty)$	Hormone clearance rate
Timescale	τ_k	1–5 cycles	$(0, \infty)$	Hormonal integration window
Diffusion coeff.	D_k	0.05–0.20	$[0, \infty)$	Inter-node propagation strength
Emission delay	δ_k	0–3 cycles	$[0, \infty)$	Stimulus-response latency
Resource coupling	ρ_k	0.0–0.3	$[0, \infty)$	Energy-hormonal damping
Noise amplitude	$\sigma_{\eta,k}$	0.01–0.05	$[0, \infty)$	Stochastic perturbation level
Cross-inhibition	γ_{km}	0.0–0.35	$[0,1)$	Hormonal axis competition
Emission gain	a_k	5.0	$(0, \infty)$	Sigmoid steepness
Emission bias	b_k	-2.5	$\mathbb{R}$	Half-maximal threshold
Min duty cycle	$\delta_{\min}$	0.10	$(0, \delta_{\max})$	Minimum active window
Max duty cycle	$\delta_{\max}$	1.00	$(\delta_{\min}, 1]$	Maximum active window
Override coeff.	$\kappa_{\text{res}}, \kappa_{\text{sens}}$	0.7	$(0,1]$	Urgency override strength
Max budget	$B_{\max}$	Deployment	$(0, \infty)$	Maximum activation budget
Budget reduction	β_B	0.6	$(0,1)$	Energy-parsimony coupling
Retrieval threshold	θ_{ret}	0.70	$(0,1)$	Minimum similarity for recall
PSM capacity	$M_{\max}$	1000	$\mathbb{Z}_{>0}$	Engram buffer size
Significance thresh.	ϵ_{sig}	0.05	$(0,1)$	Minimum outcome for storage
Time step	Δt	Per stability cond.	$(0, \Delta t_{\max}]$	Discretization step

The complete mathematical framework constitutes, to the best of the author's knowledge, one of the first unified formal frameworks combining reaction-diffusion theory [12], Lyapunov stability analysis [8], primal-dual optimization [7], and engram-based symbolic memory [10] within a single operationally grounded architecture for distributed constrained IoT systems.

4. PROBABILISTIC-ENTROPIC INTERPRETATION OF HORMONAL ORCHESTRATION

4.1. Motivation and Theoretical Position

The deterministic and stochastic formulations of S-AI-IoT developed in Sections 3.4–3.5 characterize hormonal orchestration as a dynamical system on $[0,1]^{5N}$ governed by reaction-diffusion equations with Lyapunov-certified stability. This section establishes a complementary probabilistic-entropic representation of hormonal homeostasis [23], in which each hormonal state vector $\mathbf{h}_i(t)$ is interpreted as the mean of a time-varying probability density over the operational condition space of node i. This representation complements rather than duplicates the dynamical framework of Sections 3.4–3.5: while the Lyapunov formulation establishes stability, the probabilistic representation unlocks three results not directly accessible from it — (i) an information-theoretic language for the parsimony principle, expressing activation frugality as KL-divergence-based filtering; (ii) principled maximum-likelihood calibration of hormonal parameters from deployment data, as developed in Section 4.6; and (iii) an information-theoretic foundation for the engram sufficiency result of Proposition 4.2, strengthening the zero-cost explainability guarantee of Proposition 3.2.

4.2 Probabilistic State Representation

Definition 4.1 (Hormonal Probability Density). Let $\mathbf{h}_i(t) \in [0,1]^5$ denote the hormonal state of node i at cycle t. Define the associated **hormonal probability density**$P_i(\mathbf{h}, t)$as an exponential-family probability density:

$$P_i(\mathbf{h}, t) = \frac{1}{Z_i(t)} \exp\left(-\frac{1}{2}\left(\mathbf{h} - \mathbf{h}_i(t)\right)^T \Sigma_i^{-1}\left(\mathbf{h} - \mathbf{h}_i(t)\right)\right), \quad \mathbf{h} \in [0,1]^5$$

where $Z_i(t)$ is the normalizing constant and $\Sigma_i = \sigma_\eta^2 \,\mathrm{diag}\left(\tau_k^{-1}\right)_{k\in\mathcal{K}}$ is the **time-invariant** covariance matrix induced by the stochastic perturbation amplitude $\sigma_{\eta,k}$ and the hormonal timescale τ_k of Section 3.1.3. This construction identifies each hormonal state as the **mode and mean** of a Gaussian density whose variance is inversely proportional to the hormonal response speed: fast-responding hormones (small τ_k) carry higher certainty (smaller variance) about the current operational condition.

4.3Lyapunov–KL Correspondence

Definition 4.2 (Global Cognitive Potential). Define the global potential energy of the S-AI-IoT system as:

$$V_P(\mathbf{H}(t)) = \sum_{i=1}^{N} \int_{[0,1]^5} P_i(\mathbf{h}, t) \log P_i(\mathbf{h}, t)\, d\mathbf{h} = -\sum_{i=1}^{N} \mathcal{H}\left(P_i(\cdot, t)\right)$$

where $\mathcal{H}(P_i)$ denotes the differential entropy of the hormonal density at node i.

Proposition 4.1 (Entropic Consensus Convergence). Under the stability condition of Theorem 3.1, the variance of the hormonal field across nodes satisfies:

$$\mathrm{Var}_i\, h_{k,i}(t) \le \mathrm{Var}_i\, h_{k,i}(0) \cdot e^{-2D_k\lambda_2(L^*)t/\tau_k}$$

where $\lambda_2(L^*)$ is the algebraic connectivity (Fiedler value) of the equilibrium graph G^*.

The network achieves hormonal consensus at a rate governed by the algebraic connectivity of the communication topology — a canonical convergence rate for distributed averaging protocols over weighted graphs [12].

Proof. Let $\tilde{h}_{k,i}(t) = h_{k,i}(t) - \bar{h}_k(t)$ denote the deviation of node i's hormone k from the network mean $\bar{h}_k(t) = \frac{1}{N}\sum_i h_{k,i}(t)$.

Step 1 — Linearization. Retaining only the diffusion term of Equation (3.4.1) and setting emission, cross-inhibition, resource coupling, and stochastic perturbation to zero, the deviation dynamics reduce to:

$$\tau_k \dot{\tilde{h}}_k(t) = -D_k L^* \tilde{H}_k(t)$$

where $\tilde{H}_k(t) = (\tilde{h}_{k,1}(t), \dots, \tilde{h}_{k,N}(t))^T$. Since L^* is positive semi-definite with smallest non-zero eigenvalue $\lambda_2(L^*)$ (the Fiedler value), the variance satisfies:

$$\mathrm{Var}_i\, h_{k,i}(t) = \frac{1}{N} \parallel \tilde{H}_k(t) \parallel_2^2 \leq \frac{1}{N} \parallel \tilde{H}_k(0) \parallel_2^2 \cdot e^{-2D_k \lambda_2(L^*) t/\tau_k}$$

by the standard spectral analysis of the graph Laplacian [12].

Step 2 — Bounding nonlinear perturbations. The full dynamics of Equation (3.4.1) include emission, cross-inhibition, resource coupling, and stochastic perturbation terms. By Theorem 3.1, the complete hormonal trajectory $H_k(t)$ converges to the equilibrium H_k^* at rate $e^{-2\mu t}$ with $\mu = \min_k \mu_k/\tau_k > 0$. The nonlinear perturbation to the linearized variance bound is therefore of order $O(e^{-2\mu t})$, which decays at least as fast as the diffusion term when $\mu \geq D_k \lambda_2(L^*)/\tau_k$ — a condition satisfied under the stability margin of Theorem 3.1. The variance bound of Step 1 thus holds up to a constant factor absorbing the nonlinear correction, establishing the result. ▫

Proposition 4.1' (Lyapunov–KL Correspondence). Under the Gaussian density model of Definition 4.1, with time-invariant covariance $\Sigma_i = \sigma_\eta^2\, \mathrm{diag}(\tau_k^{-1})$, the Lyapunov function $V(\mathbf{H}(t))$ of Theorem 3.1 and the total KL divergence from equilibrium satisfy the proportionality:

$$V(\mathbf{H}(t)) = \sigma_\eta^2 \sum_{i=1}^{N} D_{\mathrm{KL}}\,(P_i(\cdot,t) \parallel P_i^*) + C$$

where C is a constant independent of t. Consequently, $\dot{V} \leq 0$ if and only if the total KL divergence from the equilibrium density decreases monotonically, i.e., the system converges in the information-theoretic sense toward the minimum-uncertainty hormonal configuration consistent with the operational constraints.

Proof. For the Gaussian density of Definition 4.1, the KL divergence between the current density $P_i(\cdot,t)$ and the equilibrium density P_i^* reduces to the Mahalanobis distance:

$$D_{\mathrm{KL}}(P_i(\cdot,t) \parallel P_i^*) = \frac{1}{2}(\mathbf{h}_i(t) - \mathbf{h}_i^*)^T \Sigma_i^{-1} (\mathbf{h}_i(t) - \mathbf{h}_i^*) = \frac{1}{2\sigma_\eta^2} \sum_{k \in \mathcal{K}} \tau_k \left(h_{k,i}(t) - h_{k,i}^*\right)^2$$

Summing over all nodes i and multiplying by σ_η^2 recovers exactly $V(\mathbf{H}(t))$ up to the constant:

$$C = \frac{5N}{2}(1 + \log 2\pi) + \frac{N}{2} \log\det \Sigma_i$$

which is time-invariant since Σ_i is time-invariant by construction. The monotonic decrease of V under Theorem 3.1 therefore directly implies monotonic decrease of the total KL divergence from equilibrium.

Corollary 4.1 (Parsimony as KL-Filtered Activation). The activation parsimony property of S-AI-IoT admits the following information-theoretic characterization: the hormonal significance filter of Section 3.7 implements a **minimum-divergence activation policy** — agents are activated only when the corresponding hormonal elevation carries sufficient information content, measured as KL divergence from the baseline hormonal density, to justify the activation cost c_a. Formally:

$$x_a^*(t) = 1 \iff D_{\mathrm{KL}}\big(P_i(\cdot, t) \parallel P_i^{\mathrm{baseline}}\big) \geq \theta_{\mathrm{info}}, \quad \theta_{\mathrm{info}} = -\log\big(1 - \theta_k^{\uparrow}\big)$$

This provides an information-theoretic interpretation of the activation threshold $\theta_k^{\uparrow}$: agents remain dormant as long as the hormonal density has not diverged sufficiently from its resting configuration, implementing the architectural parsimony of S-AI-IoT as a principled information filter rather than an ad hoc threshold rule.

4.4 Distributed Entropic Consensus and Network-Level Homeostasis

The hormonal diffusion dynamics of Section 3.4 acquire a natural information-theoretic interpretation at the network level. Define the **network-level entropy functional**:

$$\mathcal{S}_{\mathrm{net}}(t) = -\sum_{i=1}^{N}\sum_{k\in\mathcal{K}} \Big(h_{k,i}(t)\log h_{k,i}(t) + \big(1 - h_{k,i}(t)\big)\log\big(1 - h_{k,i}(t)\big)\Big)$$

as the aggregate binary entropy of the hormonal field across all nodes and hormone types. The diffusion term $D_k \sum_{j\in\mathcal{N}_i(t)} w_{ij}\left(h_{k,j} - h_{k,i}\right)$ of Equation (3.4.1) drives the system toward **entropic consensus**: nodes with anomalously high or low hormonal levels relative to their neighbors are pulled toward the local average, implementing a form of distributed belief propagation over the IoT graph [12].

Remark 4.2. Proposition 4.1 provides a deployment guideline: the minimum algebraic connectivity $\lambda_2(L)$ required for the network to achieve hormonal consensus within T_{cons} orchestration cycles is:

$$\lambda_2(L) \geq \frac{\tau_k}{2D_k T_{\mathrm{cons}}}\log\frac{\mathrm{Var}_i\big(h_{k,i}(0)\big)}{\epsilon_{\mathrm{cons}}}$$

This bound relates topology design (graph connectivity) directly to hormonal response speed (τ_k, D_k) and desired coordination quality (ϵ_{cons}).

4.5 Probabilistic Calibration of Hormonal Parameters

The Lyapunov–KL correspondence of Proposition 4.1’ enables a principled **maximum-likelihood calibration** of the hormonal parameters $\{\lambda_k, a_k, b_k, D_k, \sigma_{\eta,k}\}_{k\in\mathcal{K}}$ from deployment data, as an alternative to the manual calibration procedure of Section 3.1.

Given a dataset of T_{cal} calibration cycles with observed operational signals $\{O_i(t)\}_{t=0}^{T_{\mathrm{cal}}-1}$ and ground-truth activation labels $\{a^*(t)\}$, the log-likelihood of the hormonal model is:

$$\ell(\theta) = \sum_{t=0}^{T_{\mathrm{cal}}-1}\sum_{i=1}^{N}\log P_i\big(\mathbf{h}_i^{\mathrm{obs}}(t);\, \theta\big) - \beta_{\mathrm{reg}} \parallel \theta \parallel_2^2$$

where $\mathbf{h}_i^{\mathrm{obs}}(t)$ is the observed hormonal trajectory under parameters $\theta = \{\lambda_k, a_k, b_k, D_k, \sigma_{\eta,k}\}$ and β_{reg} is an ℓ_2 regularization coefficient enforcing the stability condition (SC) of Theorem 3.1 as a soft constraint. The gradient of $\ell(\theta)$ with respect to each parameter is analytically tractable under the Gaussian density model, enabling gradient ascent calibration with convergence guaranteed by the log-concavity of the Gaussian log-likelihood in the parameters $\{a_k, b_k\}$ of the sigmoid emission function.

4.6 Entropic Explainability: Information-Theoretic Foundation

The probabilistic representation provides a deeper foundation for the zero-cost intrinsic explainability result of Proposition 3.2 (Section 3.8.5). In information-theoretic terms, the engram $\varepsilon_m^{\mathrm{IoT}} = \langle \mathbf{h}^{(m)}, \mathbf{x}^{(m)}, \mathbf{a}^{(m)}, \mathbf{o}^{(m)}, \mathcal{T}^{(m)}, \beta^{(m)} \rangle$ constitutes a **sufficient statistic** of the decision process in the following sense:

Proposition 4.2 (Entropic Sufficiency of Engrams). The mutual information between the agent activation decision $\mathbf{x}^*(t)$ and the complete operational history $\{O_i(s)\}_{s\le t}$, conditioned on the engram $\varepsilon_m^{\mathrm{IoT}}$, satisfies:$I\big(\mathbf{x}^*(t)\,;\,\{O_i(s)\}_{s\le t} \,\big|\, \varepsilon_m^{\mathrm{IoT}}\big) = 0$

Justification. The agent activation $\mathbf{x}^*(t)$ is a deterministic function of $\big(\mathbf{h}(t), B_{\mathrm{node}}(t)\big)$ via the primal-dual optimization of Section 3.7. The hormonal state $\mathbf{h}(t)$ constitutes an operationally sufficient Markov representation of the operational history $\{O_i(s)\}_{s\le t}$ by the reaction-diffusion dynamics of Section 3.4. Since $\varepsilon_m^{\mathrm{IoT}}$ contains $\mathbf{h}^{(m)}$ and $\beta^{(m)}$, it captures the complete information used by the decision process; no additional information about the history contributes to $\mathbf{x}^*(t)$ beyond what is encoded in the engram. ☐

Corollary 4.2. Proposition 4.2 establishes that engram-based explanations are not merely post-hoc approximations — they constitute causally faithful decision certificates. The XAI 2.0 architectural transparency criterion [20] is satisfied not by approximation (as in SHAP [21] or LIME [22]) but by **structural design**: the engram is a sufficient statistic of the decision by construction, and its generation at $\mathcal{O}(1)$ cost is a consequence of the Markov property of the hormonal dynamics, not of any computational shortcut.

5. CONCLUSION

This article has developed the complete formal mathematical substance of S-AI-IoT. Section 2 provided the full agent registry — 24 specialized agents across six functional families — with formal activation conditions, input-output contracts, and resource cost profiles constituting the operational layer that translates hormonal context into concrete IoT behavioral actions. Section 3 established the mathematical framework unifying reaction-diffusion theory, Lyapunov stability analysis, primal-dual optimization, and engram-based symbolic memory into a single coherent formal apparatus.The key formal results established in this article are: **Proposition 3.1** (existence of hormonal equilibrium via Brouwer's fixed-point theorem); **Theorem 3.1** (global asymptotic stability of the hormonal dynamics under the deployability condition $\lambda_k > D_k\,\lambda_{\max}(L^*) + \sum_{m\neq k} \gamma_{km} + \rho_k \parallel \boldsymbol{\chi}^* \parallel_\infty$); and **Proposition 3.2** (zero-cost intrinsic explainability: every agent activation decision is accompanied by a causally faithful engram record at $\mathcal{O}(1)$ additional cost).Together with the architectural framework of Article I [1], these results establish S-AI-IoT as a formally specified, stability-analyzed, and intrinsically explainable hormonal regulation architecture for distributed constrained IoT systems --- unifying, to the best of the author's knowledge, among the first operationally grounded IoT frameworks combining reaction-diffusion systems theory [12], Lyapunov nonlinear control [8], and engram-based behavioral memory [10] within a single coherent formal architecture. Although the proposed framework establishes formal stability, explainability, and orchestration guarantees, several practical limitations remain.The calibration of hormonal parameters may require deployment-specific adaptation depending on network topology, sensing heterogeneity, and energy constraints. In large-scale distributed deployments, hormonal synchronization and distributed engram propagation may also introduce additional communication overhead. Furthermore, while the present article establishes the formal mathematical foundations of S-AI-IoT, extensive real-world validation on heterogeneous IoT hardware platforms remains an important direction for future investigation. Future enhancements may also include adaptive parameter self-calibration, hardware-aware optimization, and large-scale edge/fog orchestration mechanisms for highly dynamic IoT infrastructures.

REFERENCES

[1] S. Slaoui, “S-AI-IoT: A Sparse Artificial Intelligence Architecture with Hormonal Orchestration, Parsimonious Agent Activation, and Symbolic Memory for Adaptive, Secure, and Explainable Internet of Things Systems”, IJAIA, May 2026.

[2] S. Slaoui, “S-AI: A sparse artificial intelligence architecture with hormonal orchestration and parsimonious agent activation,” International Journal for Multidisciplinary Research, vol. 1, no. 2, 2025.

[3] M. IfeanyiNweke, D. B. Rawat, U. Aickelin, and H. A. Khorshidi, “A comprehensive survey on imputation of missing data in Internet of Things,” ACM Computing Surveys, vol. 55, no. 6, pp. 1–38, 2022.

[4] J. Byabazaire, G. M. P. O’Hare, R. Collier, and D. Delaney, “IoT data quality assessment framework using adaptive weighted estimation fusion,” Sensors, vol. 23, no. 13, p. 5993, 2023.

[5] S. Slaoui, "S-AI-NET: A sparse AI framework for adaptive and parsimonious autonomous networking," Research Square, 2025. doi: 10.21203/rs.3.rs-4740886/v1.

[6] J. Lin, W. Yu, N. Zhang, X. Yang, H. Zhang, and W. Zhao, “A survey on Internet of Things: Architecture, enabling technologies, security and privacy, and applications,” IEEE Internet of Things Journal, vol. 4, no. 5, pp. 1125–1142, 2017.

[7] S. Boyd and L. Vandenberghe, Convex Optimization. Cambridge, UK: Cambridge University Press, 2004.

[8] H. K. Khalil, Nonlinear Systems, 3rd ed. Upper Saddle River, NJ: Prentice Hall, 2002.

[9] S. Sudevalayam and P. Kulkarni, “Energy harvesting sensor nodes: Survey and implications,” IEEE Communications Surveys & Tutorials, vol. 13, no. 3, pp. 443–461, 2011.

[10] S. Tonegawa, X. Liu, S. Ramirez, and R. Redondo, “Memory engram cells have come of age,” Neuron, vol. 87, no. 5, pp. 918–931, 2015.

[11] D. Floreano and C. Mattiussi, Bio-Inspired Artificial Intelligence: Theories, Methods, and Technologies. Cambridge, MA: MIT Press, 2008.

[12] M. E. J. Newman, Networks: An Introduction. Oxford, UK: Oxford University Press, 2010.

[13] Q. Shi, L. Liu, and X. Chen, “Edge computing: Vision and challenges,” IEEE Internet of Things Journal, vol. 3, no. 5, pp. 637–646, 2016.

[14] S. Jannatun Noor, M. A. Hossain, and M. R. Hasan, “TinyML-based intelligent IoT: Architecture, applications and challenges,” Future Generation Computer Systems, vol. 152, pp. 232–248, 2024.

[15] S. Slaoui, "S-AI-Cyber: A bio-inspired hormonal architecture for real-time cyber-defense," Research Square, 2025. doi: 10.21203/rs.3.rs-4840888/v1.

[16] E. Hart and J. Timmis, “Application areas of AIS: The past, the present and the future,” Applied Soft Computing, vol. 8, no. 1, pp. 191–201, 2008.

[17] C. Berger, P. Eichhammer, H. P. Reiser, J. Domaschka, F. J. Hauck, and G. Habiger, “A survey on resilience in the IoT,” ACM Computing Surveys, vol. 54, pp. 1–39, 2021.

[18] J. J. Hopfield, “Neural networks and physical systems with emergent collective computational abilities,” Proceedings of the National Academy of Sciences, vol. 79, no. 8, pp. 2554–2558, 1982.

[19] J. Mineraud, O. Mazhelis, X. Su, and S. Tarkoma, “A gap analysis of Internet-of-Things platforms,” Computer Communications, vol. 89–90, pp. 5–16, 2016.

[20] L. Longo, M. Brcic, F. Cabitza, et al., “Explainable artificial intelligence (XAI) 2.0: A manifesto of open challenges and interdisciplinary research directions,” Information Fusion, vol. 106, p. 102301, 2024.

[21] S. M. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions,” in Proc. NeurIPS, 2017, pp. 4765–4774.

[22] M. T. Ribeiro, S. Singh, and C. Guestrin, “‘Why should I trust you?’: Explaining the predictions of any classifier,” in Proc. ACM KDD, 2016, pp. 1135–1144.

[23] S. Slaoui, “From Hormones to Probabilities: A Unified Doctrine of Cognitive Homeostasis in Sparse Artificial Intelligence (S-AI),” SSRN, Nov. 2025. Available: https://ssrn.com/abstract=5735582

AUTHOR

Said Slaoui is a professor at Mohammed V University in Rabat, Morocco. He graduated in Computer Science from University Pierre and Marie Curie, Paris VI (in collaboration with IBM France), 1986. He has over 40 years of experience in the fields of AI and Big Data, with research focused on modular architectures, symbolic reasoning, and computational frugality. His recent work introduces the Sparse Artificial Intelligence (S-AI) framework, which integrates bio-inspired signaling and agent-based orchestration. He has published numerous scientific papers in international journals and conferences, and actively contributes to the development of sustainable and explainable AI systems.