

Network-Optimised Spiking Neural Network for Event-Driven Networking

Muhammad Bilal

School of Computing and Communications, Lancaster University, LA1
4WA Lancaster, United Kingdom
m.bilal@ieee.org

Abstract

Spiking neural networks support event-driven computation suited to time-critical networking tasks such as anomaly detection, local routing control, and congestion management at the edge. However, classical units, including Hodgkin–Huxley, Izhikevich, and the Random Neural Network, map poorly to these needs. This work introduces a compact two-variable neuron, *Network-Optimised Spiking* (NOS), whose state corresponds to normalised queue occupancy and a recovery resource. The model uses a saturating nonlinearity to represent finite buffers, a service-rate leak, and graph-local inputs with delays and optional per-link gates. It supports two differentiable reset schemes for training and deployment. We derive conditions for existence and uniqueness of equilibrium, provide local stability tests from the Jacobian trace and determinant, and obtain a network threshold that scales with the Perron eigenvalue of the coupling matrix. The analysis yields an operational rule $g_\star \approx k^\star/\rho(W)$ that links damping and offered load, shows how saturation enlarges the stable region, and explains finite-size smoothing of synchrony onsets. Stochastic arrivals follow a compound Poisson shot-noise model aligned with telemetry smoothing, which gives closed-form sensitivity and variance expressions. Against queueing baselines NOS matches the light-load M/M/1 mean by calibration while truncating deep tails under bursty input. In closed loop it produces decisive, low-jitter marking with short settling. In zero-shot forecasting without labels NOS is calibrated per node from arrival statistics and known μ , K , and Δt . Its bounded event-driven dynamics yield high AUROC and AUPRC, enabling timely detection of congestion onsets with few false positives. Under a train-calibrated residual protocol across chain, star, and scale-free topologies, NOS improves early-warning F1 and detection latency over MLP, RNN, GRU, and a simple temporal GNN. We provide practical guidance for data-driven initialisation and surrogate-gradient training with a homotopy on reset sharpness, together with explicit stability checks and topology-aware bounds suitable for resource-constrained deployments.

keywords

Spiking neural networks, Congestion control, Active queue management, Explicit Congestion Notification, Network stability, Spectral graph theory, Distributed control, Queueing theory

1 Introduction

Machine learning has reshaped core networking tasks, including short-horizon traffic forecasting, anomaly detection, and routing optimisation. Conventional deep models such as multilayer perceptrons, recurrent networks, and graph neural networks perform well when labels are plentiful and training is done offline. Spiking neural networks offer a complementary path. Their event-driven computation is naturally sparse in time and energy, which aligns with packet-level operation at the edge and in devices with tight power budgets. Recent surveys and hardware results report large efficiency gains for spiking workloads on neuromorphic substrates, motivating designs that couple spiking dynamics to network semantics [6, 8, 5]. Foundational work on temporal coding and computational power further motivates spike-based processing in control settings [21]. In parallel, hardware platforms such as Intel’s Loihi and its toolchain demonstrate on-chip learning and practical workflows for deploying SNNs [8, 23].

Classical spiking neural network (SNN) models, while powerful for neuroscientific exploration, are poorly matched to networking practice due to several fundamental limitations. Firstly, they often prioritize biological analogy, defining abstract internal states rather than variables with direct network interpretations like queue occupancy, service rate, or link delay [15] [28]. This biological focused modeling also leads to dynamics that present significant optimization challenges; the hard threshold-and-reset dynamics intrinsic to spiking neurons create a non-differentiable discontinuity that severely obstructs gradient-based optimisation, a cornerstone of modern machine learning [1][39]. Furthermore, the treatment of topology and per-link quality is often implicit or uniform, neglecting the complex, heterogeneous, and weighted graphs that define real-world networks like data centers or the Internet. From a dynamical systems perspective, the excitability in common SNN models formulation can be unbounded, leading to unrealistic behaviour under high input load and presenting significant numerical instability [38]. Finally, implementations face practical limits in training cost, input encoding, and deployment on resource-constrained neuromorphic or edge hardware. Consequently, evaluation remains limited, often failing to bridge the gap between abstract simulations and performance on real network topologies and traffic patterns [15] [37]. Early Poisson spiking models also show that spike-driven networks need not admit Hopfield-style energy functions, reinforcing the case for operational, domain-grounded formulations [7]. The Hodgkin–Huxley [16] equations are biophysically correct but heavy for embedded use. Izhikevich’s [18] formulation is efficient but its quadratic drive is unbounded and the reset is discontinuous, which complicates training and interpretation for queues. Random Neural Networks and G-networks [11] bridge neural and queueing views, yet their canonical forms do not provide the graph-aware, differentiable, and parameter-interpretable behaviour needed for packet dynamics on real topologies. These gaps motivate a compact spiking unit whose state maps to normalised queue occupancy and recovery, whose inputs are graph-local with delays and optional per-link gates, and whose reset is differentiable for gradient-based optimisation [39]. A second driver is deployment. Neuromorphic systems such as Intel’s Loihi families demonstrate on-chip learning rules, synaptic delays, and event-driven execution that can reduce energy for small edge workloads while keeping latency low [8]. Related SNN studies in robotics and generic complex-network settings illustrate broader interest in topology and control, but their objectives differ from packet dynamics and queue semantics considered here [40, 33, 22]. Such systems reinforce the case for spiking designs that expose measurable parameters and training handles, so that

network operators can calibrate models from telemetry and run them under tight power envelopes.

1.1 Spiking Neural Networks for Networking: Scope and Complementarity

Spiking models describe a time-varying state and emit discrete events when a threshold is crossed. In the biological literature this is written in terms of a membrane potential $V(t)$; for networking the same idea can be read as a state that advances under inputs and intrinsic dynamics until an event is produced. The event acts as the computational primitive. Intelligent networking has been led by deep learning, reinforcement learning, and graph neural networks trained on aggregates or labelled flows. These methods assume synchronous batches and centralised compute. Spiking neural networks use event-driven updates, which aligns with irregular packet arrivals, bursty congestion, and anomaly signatures that depart from baseline rather than average behaviour. The practical question is when this alignment matters enough to prefer an SNN, and when conventional ML or GNNs remain the better tool.

Event-driven versus batch-oriented learning. Packet systems are inherently event-driven. Conventional pipelines buffer into fixed windows before inference, adding latency and energy. SNNs update on arrival, so a single event can trigger a decision without waiting for a batch. This is attractive for anomaly alerts and micro-burst monitoring where early reaction is valuable.

Energy and deployment constraints. Edge devices and in-network elements often face tight power and memory budgets. Neuromorphic platforms (e.g., Loihi and SpiN-Naker) report energy advantages for spiking workloads, which supports in situ decisions at low latency [8, 10]. In contrast, data-centre controllers and core routers, with ample power and global context, typically profit from conventional ML and GNNs.

Temporal coding and traffic structure. Some networking tasks hinge on timing patterns rather than sheer volume: jitter in interactive media, periodic bursts in games, or hotspot prediction from inter-arrival times. SNNs encode such structure in spike timing and phase. Recurrent deep models can capture time as well, but temporal coding is native to spiking dynamics.

Local learning and autonomy. Networks are distributed. Routers, base stations, and sensors often adapt with limited scope. SNNs admit local rules such as spike-timing-dependent plasticity (STDP), enabling on-node adaptation. Global optimisation over topology and multi-objective trade-offs remains a strength of GNNs and reinforcement learning.

Complementary roles. SNNs and conventional AI are complementary. SNNs provide efficient, event-driven reflexes for local, time-critical decisions under tight budgets. Deep learning and GNNs act as global planners that integrate structure and statistics. A practical design places SNNs in edge nodes or programmable switches for reflexive response, coordinated with ML/GNN controllers for end-to-end policy.

Table 1: Comparative suitability of SNNs and conventional ML/GNNs across networking problem domains.

Problem domain	More suitable with SNNs	More suitable with ML/GNNs
Anomaly detection (e.g., DDoS, worm outbreaks)	Event-driven detection of sudden bursts; low-latency alarms [26, 4]	Flow-level detection from large corpora; supervised and deep models [2, 19]
Routing optimisation	Local adaptation in ad hoc or sensor settings using local rules; energy- and event-efficient deployment on neuromorphic hardware [8, 9]	Global path optimisation and traffic engineering with topology-wide objectives; GNN/RL controllers [36]
QoS prediction (throughput, delay)	Edge-side reaction to microbursts under power limits (event-driven inference) [8]	Predictive modelling from aggregates; offline/online regression and forecasting [14, 3, 35]
Traffic classification	Temporal signatures (periodicity, jitter) via spike timing and event patterns [31, 30]	Feature-rich flow classification with labels using CNN/RNN/GNN [29, 41, 34]
Wireless resource allocation	Local duty-cycling, sensing, and lightweight on-device decision-making under tight energy budgets [9]	Centralised spectrum allocation and large-scale optimisation with RL/GNN [25, 42, 17]
Network resilience and fault detection	Distributed detection of local failures and sudden state changes in streams [26, 4]	Correlating failures across topology; resilience modelling and recovery for SDN/CPS [27, 32, 24]

1.2 Contributions and Findings

This article develops a *Network-Optimised Spiking* neuron (NOS) and examines its behaviour from single node to network scale. The model is compact and two-state, with each state and parameter tied to measurable network quantities. Queue occupancy serves as the fast state, and a slower recovery variable captures post-burst slowdown. Service appears as leak terms. Inputs are graph-local with options for per-link gating and fixed delay. A saturating nonlinearity prevents unbounded growth and represents finite buffers. Two differentiable reset strategies are introduced—an event-triggered exponential soft reset and a continuous pullback shaped by a sigmoid—both enabling gradient-based training while preserving spiking dynamics.

The analysis gives equilibrium existence and uniqueness, stability tests from the Jacobian, and a network threshold that scales with the Perron eigenvalue of the coupling matrix. Saturation enlarges the stable region by reducing the effective excitability slope.

An operational margin links damping with offered load and provides a two-parameter contour for planning and capacity checks. Stochastic arrivals are represented as shot noise or compound Poisson input with calibrated scales. Practical guidance covers parameter initialisation from telemetry, surrogate-gradient training with a homotopy in reset sharpness, and deployment on neuromorphic platforms. The evaluation protocol uses label-free forecasting with residual-based event scoring, in line with established practice in time-series anomaly detection.

Findings show that bounded excitability yields unique single-node equilibria and wider stability regions than the unbounded quadratic drive. Network stability follows the Perron eigenvalue, giving a direct means of control through weight normalisation and topology choice. Differentiable resets remove optimisation difficulties caused by hard jumps and enable surrogate-gradient training without disturbing event timing. In residual-based forecasting, NOS acts as a physics-aware forecaster, achieving competitive accuracy and latency on chain, star, and scale-free graphs while remaining interpretable and suitable for neuromorphic execution.

2 Background: classical spiking paradigms and networking suitability

This section reviews three classical spiking neural networks and examines what they offer, and their suitability for networking tasks where queue semantics, topology, and trainability plays important role.

2.1 Random Neural Network (Gelenbe)

The Random Neural Network (RNN) treats each neuron as an integer-valued counter driven by excitatory and inhibitory Poisson streams and yields a product-form steady state [12]. Let $q_i = \Pr\{k_i > 0\}$ denote the activity of neuron i . In steady state

$$q_i = \frac{\Lambda_i^+}{r_i + \Lambda_i^-}, \quad (1)$$

$$\Lambda_i^+ = \lambda_i^+ + \sum_j q_j r_j p_{ji}^+, \quad \Lambda_i^- = \lambda_i^- + \sum_j q_j r_j p_{ji}^-, \quad (2)$$

and the joint distribution factorises as $\Pr[k_1, \dots, k_N] = \prod_i (1 - q_i) q_i^{k_i}$.

The queueing affinity and fixed points make calibration attractive. However, backbone and datacentre traces are overdispersed and often long-memory, so a stationary product form understates burst clustering and tail risk. This follows from the Poisson and exponential assumptions of the model, which wash out temporal correlation and refractoriness. Product-form results emphasise stationarity, whereas operations rely on transient detection and control. Parameters r_i and p_{ji}^\pm do not tie directly to per-link capacity, delay, or gating. Under strong excitation, the signal-flow equations can fail to admit a valid solution with $q_i < 1$, leading to instability or saturation; when a solution exists it is unique [12, 13].

2.2 Hodgkin–Huxley

The Hodgkin–Huxley (HH) equations provide biophysical fidelity through voltage-gated conductances and reproduce spikes, refractoriness and ionic mechanisms [16]. The membrane and current relations are

$$C_m \frac{dV}{dt} = -(I_{\text{Na}} + I_{\text{K}} + I_L) + I_{\text{ext}}, \quad (3)$$

$$I_{\text{Na}} = \bar{g}_{\text{Na}} m^3 h (V - E_{\text{Na}}), \quad I_{\text{K}} = \bar{g}_{\text{K}} n^4 (V - E_{\text{K}}), \quad I_L = \bar{g}_L (V - E_L), \quad (4)$$

$$\frac{dx}{dt} = \alpha_x(V)(1 - x) - \beta_x(V)x, \quad x \in \{m, h, n\}. \quad (5)$$

where V is membrane potential; C_m membrane capacitance. $I_{\text{Na}}, I_{\text{K}}, I_L$ are sodium, potassium, and leak currents; I_{ext} is external or synaptic input. \bar{g}_{ion} are maximal conductances and E_{ion} reversal potentials. $m, h, n \in [0, 1]$ are gating variables with voltage-dependent rates $\alpha_x(V), \beta_x(V)$.

Hodgkin–Huxley offers mechanistic fidelity and a broad repertoire of excitable behaviour, and recent results show that HH neurons can be trained end to end with surrogate gradients, achieving competitive accuracy with very sparse spiking on standard neuromorphic benchmarks [20]. This confirms feasibility for learning and suggests potential efficiency from activity sparsity. For network packet-level control, the obstacles are practical rather than conceptual. The model comprises four coupled nonlinear differential equations with voltage-dependent rate laws, which typically require small timesteps and careful numerical treatment; in practice this raises computational cost and stiffness at scale, and training can fail without tailored integrators and schedules [20]. There is also no direct mapping from biophysical parameters to queue occupancy, service rate, or link delay, so calibration from network telemetry lacks interpretability. Taken together, HH remains a benchmark for cellular electrophysiology [16], but it is a poor fit for networking tasks that requires lightweight, semantically mapped states and trainable dynamics.

2.3 Izhikevich

The Izhikevich model balances speed and diversity of firing patterns [18]. Its dynamics and reset are

$$\frac{dv}{dt} = 0.04v^2 + 5v + 140 - u + I, \quad \frac{du}{dt} = a(bv - u), \quad (6)$$

$$\text{if } v \geq 30 \text{ mV} : \quad v \leftarrow c, \quad u \leftarrow u + d. \quad (7)$$

where, v is the membrane potential (mV). u is a recovery variable with the same units as v that aggregates K^+ activation and Na^+ inactivation effects. I is the input drive expressed as an equivalent voltage rate (mV ms^{-1}). The spike condition is $v \geq 30$ mV, after which $v \leftarrow c$ and $u \leftarrow u + d$. a is the inverse time constant of u (ms^{-1}); b is the coupling gain from v to u (dimensionless); c is the post-spike reset level of v (mV); d is the post-spike increment of u (mV).

The Izhikevich model achieves breadth of firing patterns with two state variables and a hard after-spike reset (6)–(7). This efficiency is documented through large-scale pulse-coupled simulations and parameter recipes for cortical cell classes. For networking, three constraints matter. First, the quadratic drift in (6) is not intrinsically saturating, which is mismatched to finite-buffer behaviour and can drive unrealistically rapid growth

between resets under heavy input. Second, the reset in (7) is discontinuous, which hinders gradient-based optimisation. Third, the parameters (a, b, c, d) are phenomenological and the connectivity is generic, so queue occupancy, service rate, link delay, and per-link quality are not first-class. These features explain the model’s value for neuronal dynamics and its limited interpretability for packet-level control.

2.4 Synthesis

Across the three formalisms the evidence points to the same gaps. Either the model is correct but heavy, or it is fast but carries non-differentiable resets and unbounded excitability. In all cases the parameters do not align with queue occupancy, service, and delay. Topology and per-link quality are usually implicit rather than explicit. For packet networks this limits interpretability, stability under high load, and the ability to train or adapt in situ. These observations set the design brief for NOS: a compact two-state unit with finite-buffer saturation, a service-rate leak and a differentiable reset, graph-local inputs with delays and optional gates, and parameters that map to observable network quantities.

Table 2: Classical paradigms versus networking requirements.

	RNN (Gelenbe)	Hodgkin–Huxley	Izhikevich
State/parameters map to queue, service, delay	Limited (probabilistic rates)	No	No
Trainability for control (gradients)	Moderate at steady state	Low	Low (non-differentiable reset)
Topology and per-link attributes explicit	Limited	No	No
Finite-buffer realism / stability under load	Stationary focus	Realistic but heavy	Unbounded drive
Edge-side compute/energy	Low–moderate	High	Low

3 Proposed model: Network-Optimised Spiking neural network (*NOS*)

Networking workloads are shaped by graph structure, finite buffers, and service constraints, and they must be trainable from telemetry. We therefore begin with graph-local inputs that respect neighbourhoods and link heterogeneity,

$$I_i(t) = \sum_{j \in \mathcal{N}(i)} w_{ij} S_j(t), \quad (8)$$

where $S_j(t)$ is a presynaptic event train and w_{ij} encodes link capacity, policy weight, or reliability. Three implementation choices motivate the design. First, the reset must be continuous and differentiable so gradient-based training is possible. Second, state variables and parameters should map to observables such as normalised queue length, service rate, and recovery time. Third, subthreshold integration needs explicit stochastic drive so gradual load accumulation and burstiness are both represented. These choices lead to a compact two-variable unit for scalability and neuromorphic deployment, bounded excitability that respects finite buffers, graph-aware coupling with optional per-link gates and delays, leaky subthreshold dynamics, and data-calibrated parameters.

3.1 Variable reinterpretation

inspired from [18], we keep a two-variable unit but reinterpret its state for networking. The fast state $v_i(t)$ represents a normalised queue or congestion level, with $v_i \in [0, 1]$ mapping empty to full buffer. The slow state $u_i(t)$ represents a recovery or slowdown resource that builds during bursts and relaxes thereafter.

$v \mapsto$ normalised queue length or congestion level, $u \mapsto$ recovery or slowdown resource.

Here $v = 1$ corresponds to a full buffer. The state u summarises pacing, token replenishment, or rate-limiter cool-down that follows a burst. This choice allows direct initialisation from traces: v from queue occupancy, u from measured relaxation time, and the link between them from decay after micro-bursts. It also clarifies units. Both v and u evolve in the same time scale, which keeps linearisation and stability analysis transparent.

3.2 State variables and dynamics

For each node i we adopt a bounded excitability function and explicit damping:

$$\frac{dv_i}{dt} = f_{\text{sat}}(v_i) + \beta v_i + \gamma - u_i + I_i(t) - \lambda v_i - \chi(v_i - v_{\text{rest}}), \quad (9)$$

$$\frac{du_i}{dt} = a(bv_i - u_i) - \mu u_i = abv_i - (a + \mu)u_i. \quad (10)$$

Equation (9) separates four operational effects. The term $f_{\text{sat}}(v_i)$ (see (11)) captures the convex rise of backlog during rapid arrivals while remaining bounded. The linear part $\beta v_i + \gamma$ fits residual slope and offset that are visible in practice after removing coarse trends. The coupling $-u_i$ models recovery drag that slows re-accumulation immediately after events. Finally, $-\lambda v_i - \chi(v_i - v_{\text{rest}})$ represents service and small-signal damping. Equation (10) integrates recent congestion with sensitivity ab and relaxes at rate $a + \mu$. Thus, the *NOS* model reflects three intuitive aspects of queue behaviour: the rise under increasing load, the draining through service, and the short-lived slowdown that follows bursts. A complete schematic including graph-local gates, delays, stochastic threshold, and differentiable reset appears in Fig. 2 (§3.5–§3.7).

3.3 Bounded excitability

The quadratic term used in Izhikevich provides excitability but grows without bound. In a networking view, this implies unconstrained buffer growth at router. Thus, We replace it with a saturating nonlinearity

$$f_{\text{sat}}(v) = \frac{\alpha v^2}{1 + \kappa v^2}, \quad \alpha > 0, \kappa > 0, \quad (11)$$

which preserves a quadratic ramp for small load

$$f_{\text{sat}}(v) \sim \alpha v^2 \quad \text{as } v \rightarrow 0, \quad (12)$$

and enforces a finite ceiling for heavy load, ensuring bounded behaviour

$$\lim_{v \rightarrow \infty} f_{\text{sat}}(v) = \frac{\alpha}{\kappa}. \quad (13)$$

The derivative is globally bounded, which improves numerical stability and yields clean gain conditions, which means, it recovers a quadratic slope for small v while enforcing saturation for large v :

$$f'_{\text{sat}}(v) = \frac{2\alpha v}{(1 + \kappa v^2)^2}, \quad (14)$$

is globally bounded, with maximum

$$\max_v f'_{\text{sat}}(v) = \frac{3\sqrt{3}}{8} \frac{\alpha}{\sqrt{\kappa}} \quad \text{at } v = \frac{1}{\sqrt{3\kappa}}. \quad (15)$$

(15) is the steepest admissible local “growth” that service and damping must dominate to avoid runaway in bursts.

3.4 Equilibrium and boundedness

At steady state we eliminate u_i using $u_i = \frac{ab}{a+\mu}v_i$ in (10) and obtains a scalar balance:

$$F(v^*) := f_{\text{sat}}(v^*) + L v^* + C = 0, \quad (16)$$

where L and C collect the effective linear and constant drives implied by (9). The next result provides a sufficient, checkable criterion.

Lemma 1 (Existence and uniqueness). *Let $\mathcal{I} \subset \mathbb{R}$ be the admissible queue interval, for example $[0, 1]$. If*

$$\inf_{v \in \mathcal{I}} (f'_{\text{sat}}(v) + L) > 0, \quad (17)$$

then F is strictly increasing on \mathcal{I} and has at most one root. If, in addition, there exist $v_{\pm} \in \mathcal{I}$ with $F(v_-)F(v_+) \leq 0$, then a unique equilibrium $v^ \in \mathcal{I}$ exists. A sufficient global condition is*

$$L > -\frac{3\sqrt{3}}{8} \frac{\alpha}{\sqrt{\kappa}}. \quad (18)$$

Discussion. Condition (17) states that net drain dominates the steepest admissible excitability slope. The bound (18) follows from (15). It is conservative, easy to verify from telemetry, and stable against modest estimation error.

Corollary 1.1 (Algebraic condition and networking view). *For the coefficients in (9)–(10), a sufficient condition for a unique equilibrium in $[0, 1]$ is*

$$\beta - \lambda - \chi - \frac{ab}{a + \mu} > -\frac{3\sqrt{3}}{8} \frac{\alpha}{\sqrt{\kappa}}, \quad (19)$$

or, equivalently,

$$\lambda + \chi + \frac{ab}{a + \mu} > \frac{3\sqrt{3}}{8} \frac{\alpha}{\sqrt{\kappa}} - \beta. \quad (20)$$

Operationally, the left side aggregates service and damping, while the right side is the worst-case excitability minus any linear relief β . At full buffer, service must dominate arrivals; at empty buffer, leakage should not push the queue negative. Hence a single operating point exists within $[0, 1]$.

Engineering rule of thumb. Choose κ so that the knee of f_{sat} lies near the full-buffer scale, for example $1/\sqrt{\kappa} \approx 1$. Fit α from small-signal curvature so that $f_{\text{sat}}(v) \approx \alpha v^2$ around typical loads. Then tune λ and χ to satisfy

$$\lambda + \chi + \frac{ab}{a + \mu} \geq \frac{3\sqrt{3}}{8} \frac{\alpha}{\sqrt{\kappa}} - \beta + \varepsilon, \quad (21)$$

with a small safety margin $\varepsilon > 0$ to absorb burstiness. This keeps the operating point in the unique, stable region and reduces sensitivity to trace noise.

3.5 Graph-local inputs, explicit delays, and per-link queues

Topology enters through delayed presynaptic events (neighborhood wiring, per-link delays, and optional link-state gates that modulate influence during congestion, scheduling, or wireless fading) and optional per-link gates. Starting from (8), we incorporate delays and exogenous noise as

$$I_i(t) = \sum_{j \in \mathcal{N}(i)} w_{ij} S_j(t - \tau_{ij}) + \eta_i(t), \quad (22)$$

where $\mathcal{N}(i)$ is the graph neighborhood of node i , $S_j(t)$ is the presynaptic event train emitted by j , w_{ij} encodes a link’s nominal capacity, policy weight, or reliability, τ_{ij} is a link delay, and $\eta_i(t)$ is stochastic drive. In practice, τ_{ij} can be drawn from Round Trip Time (RTT) profiling after excluding queueing delays, while w_{ij} can be scaled so the spectral radius of the un-gated weight matrix is controlled at design time. Delays ensure that i reacts only after information from j can physically arrive, which is essential when bursts traverse multiple hops and when controllers must respect causality. When links maintain their own queues, we gate the coupling by a per-link occupancy $q_{ij}(t)$:

$$\dot{q}_{ij} = -\zeta q_{ij} + \Phi(S_j(t)) - \Psi(q_{ij}) \quad (23)$$

and replace $w_{ij}S_j(t - \tau_{ij})$ by $w_{ij}g(q_{ij})S_j(t - \tau_{ij})$. This produces a two layer representation: node integrator plus per link capacity gating. Here Φ maps presynaptic activity to arrivals on $(j \rightarrow i)$, Ψ represents service or RED/ECN-like bleed, and g is a bounded, monotone gate that reduces influence when the link is loaded. This separates phenomena cleanly: node excitability remains interpretable in queue units, while link congestion only modulates the strength and timing of coupling.

Per-link queue gating (dimensionally explicit). Let $q_{ij}(t)$ denote the normalised occupancy of link (i, j) , with $q_{ij} \in [0, 1]$. Choose a single time constant $\tau_q > 0$ and a bounded, dimensionless encoding $\sigma_s : \mathbb{R} \rightarrow [0, 1]$ of the presynaptic drive $S_j(t)$. We model the queue as a leaky accumulator of arrivals:

$$\tau_q \dot{q}_{ij}(t) = -q_{ij}(t) + \sigma_s(S_j(t)), \quad (24)$$

so that $[\tau_q] = \text{time}$ and q_{ij} is dimensionless. The effective coupling is then

$$w_{ij}^{\text{eff}}(t) = w_{ij}g(q_{ij}(t)), \quad g : [0, 1] \rightarrow [0, 1], \quad (25)$$

with gates such as

$$g(x) = (1 - x)^p \quad (\text{capacity fade, } p > 1) \quad (26)$$

$$g(x) = \frac{1}{1 + \exp\{k(x - \theta)\}} \quad (\text{logistic gate}). \quad (27)$$

One then replaces $w_{ij}S_j(t - \tau_{ij})$ in (22) by $w_{ij}^{\text{eff}}(t)S_j(t - \tau_{ij})$. This cleanly separates node adaptation from link capacity effects. In practice it lets an operator disclose where instability originates: node excitability, link queues, or topology. For completeness, the original form (23) is dimensionally consistent provided $[\zeta] = \text{time}^{-1}$ and $[\Phi(S)] = [\Psi(q)] = \text{time}^{-1}$; (24) corresponds to the special case $\zeta = \tau_q^{-1}$, $\Phi(S) = \tau_q^{-1}\sigma_s(S)$, and $\Psi \equiv 0$.

Networking interpretation and design guidance. Equation (22) enforces hop-level causality: bursts from j affect i only after τ_{ij} , so back-pressure and wavefront effects appear on the correct timescale. The queue model in (23)–(24) captures how a busy uplink or radio fades coupling without requiring packet loss. The drain term Ψ can represent ECN or token-bucket service, while the gate $g(\cdot)$ reduces the effective link weight as occupancy rises. The power gate in (26) suits wired interfaces where capacity fades smoothly with load. The logistic gate in (27) matches thresholded policies such as priority drops, wireless duty cycling, or scheduler headroom rules. Because $w_{ij}^{\text{eff}}(t)$ multiplies $S_j(t - \tau_{ij})$ in (22), congestion localises: only the affected links attenuate, which keeps attribution simple in operations dashboards.

Calibration follows telemetry. Choose τ_{ij} from the queue-free component of RTT. Set τ_q near the interface drain time or the QoS smoothing window. Normalise w_{ij} by nominal link rates or policy priorities, then scale $W = \{w_{ij}\}$ so its spectral radius meets the desired small-signal index; under stress the gate g reduces the effective spectral radius, enlarging the stable operating region. In evaluation, this decomposition helps explain behaviour: if instability appears with empty q_{ij} the cause is node excitability or topology; if it vanishes when g engages, the bottleneck is link capacity or scheduling. See Appendix A for practical calibration notes and examples.

3.6 Stochastic arrivals: shot noise and compound Poisson

Network arrivals are bursty, short-correlated within flows, and heavy-tailed across flows. To expose this variability to the model while staying compatible with measurement practice, we represent exogenous drive at node i by a compound Poisson shot-noise process smoothed at an operator-chosen time scale:

$$\eta_i(t) = \sum_{n=1}^{N_i(t)} A_{i,n} \kappa_s(t - t_{i,n}), \quad (28)$$

where $N_i(t)$ is a Poisson process of rate ρ_i , the amplitudes $A_{i,n}$ reflect burst size, and κ_s is a causal exponential smoothing kernel such as,

$$\kappa_s(t) = e^{-t/\tau_s} H(t), \quad (29)$$

$$H(t) = \begin{cases} 0, & t < 0, \\ 1, & t \geq 0, \end{cases} \quad (30)$$

so $H(t)$ enforces causality (no effect before the burst) and the factor e^{-t/τ_s} sets an exponential decay with time scale τ_s (the burst’s influence fades smoothly). When needed, a normalised variant keeps unit area:

$$\kappa_s^{\text{norm}}(t) = \frac{1}{\tau_s} e^{-t/\tau_s} H(t). \quad (31)$$

The normalisation is set so that $A/(v_{\text{th}} - v_{\text{rest}})$ falls in a practical range for the workload. This construction aligns with the EWMA-style smoothing used in telemetry pipelines and permits direct calibration from packet counters.

Equation (28) mirrors the way operators pre-process packet counters and flow logs: arrivals are bucketed, lightly smoothed, and scaled. The rate ρ_i matches the observed burst-start rate at node i over the binning interval; the amplitude $A_{i,n}$ matches per-burst byte or packet mass after the same smoothing; and τ_s is set to the shortest window that suppresses counter jitter without hiding micro-bursts. The normalisation is chosen so that a single exogenous shot changes v by a fraction ϵ of the threshold margin $m = v_{\text{th}} - v_{\text{rest}}$, with $\epsilon \in [10^{-3}, 2 \times 10^{-1}]$. The lower bound clears telemetry noise; the upper bound prevents a single burst from forcing a reset. In discrete time this corresponds to $A \Delta t/m \in [10^{-3}, 2 \times 10^{-1}]$; under the exponential kernel $\kappa_s(t) = e^{-t/\tau_s} H(t)$ it corresponds to $A \tau_s/m \in [10^{-3}, 2 \times 10^{-1}]$. Using (28) inside $I_i(t)$ in (8) makes the event drive consistent with the graph structure and delays defined in (22). In effect, traffic statistics enter *NOS* with the same time base and smoothing that appear in the telemetry pipeline, which simplifies both training and post-deployment troubleshooting. ““

3.7 Spike generation and resets

A spike indicates that the local congestion proxy exceeded tolerance and triggers a control action such as ECN marking, rate reduction, or an alert to the controller. To reflect tolerance bands and measurement noise we allow a stochastic threshold with bounded dispersion:

$$v_i(t) \geq v_{\text{th}}(t) = v_{\text{th,base}} + \sigma \xi(t), \quad (32)$$




Figure 1: NOS phase–plane diagnostics with nullclines $\dot{v} = 0$ and $\dot{u} = 0$, threshold v_{th} , equilibrium, and short trajectories. **(a)** Baseline geometry: the parabolic v -nullcline intersects the linear u -nullcline at a stable equilibrium; the threshold sits to the right. **(b)** With identical exogenous shot noise but reset disabled, the trajectory creeps toward the threshold and remains above the u -nullcline, indicating slow accumulation of *effective* load that can lead to prolonged excursions. **(c)** With the same noise and the exponential soft reset active, the trajectory crosses the threshold once, is immediately pulled back toward c , and rises in u , which encodes a temporary conservative phase. Operationally, panel (c) mirrors paced drain after an alarm: timing is preserved by the threshold crossing, the return to baseline follows the configured time scale r_{reset}^{-1} , and the recovery increment d captures the short-lived send-rate reduction.

where $\xi(t)$ is zero-mean, bounded (e.g., clipped Gaussian or uniform), and σ is tuned to the false-alarm budget on residuals. This construction separates policy (the base threshold) from instrumentation noise (the dispersion).

Event-based exponential soft reset. Right after a threshold crossing, the unit should return toward a baseline without a hard discontinuity. For a forward-Euler step of size Δt we apply

$$v_i \leftarrow c + (v_i - c)e^{-r_{\text{reset}}\Delta t}, \quad u_i \leftarrow u_i + d, \quad (33)$$

which makes v_i relax exponentially toward c with rate r_{reset} and increments u_i by d to encode refractory depth. In networking terms, (33) models a paced drain after a trigger together with a short-lived slow-down in effective send or admit rate. The pair (r_{reset}, d) is set from post-burst relaxation fits on v and the desired “cool-down” depth in u .

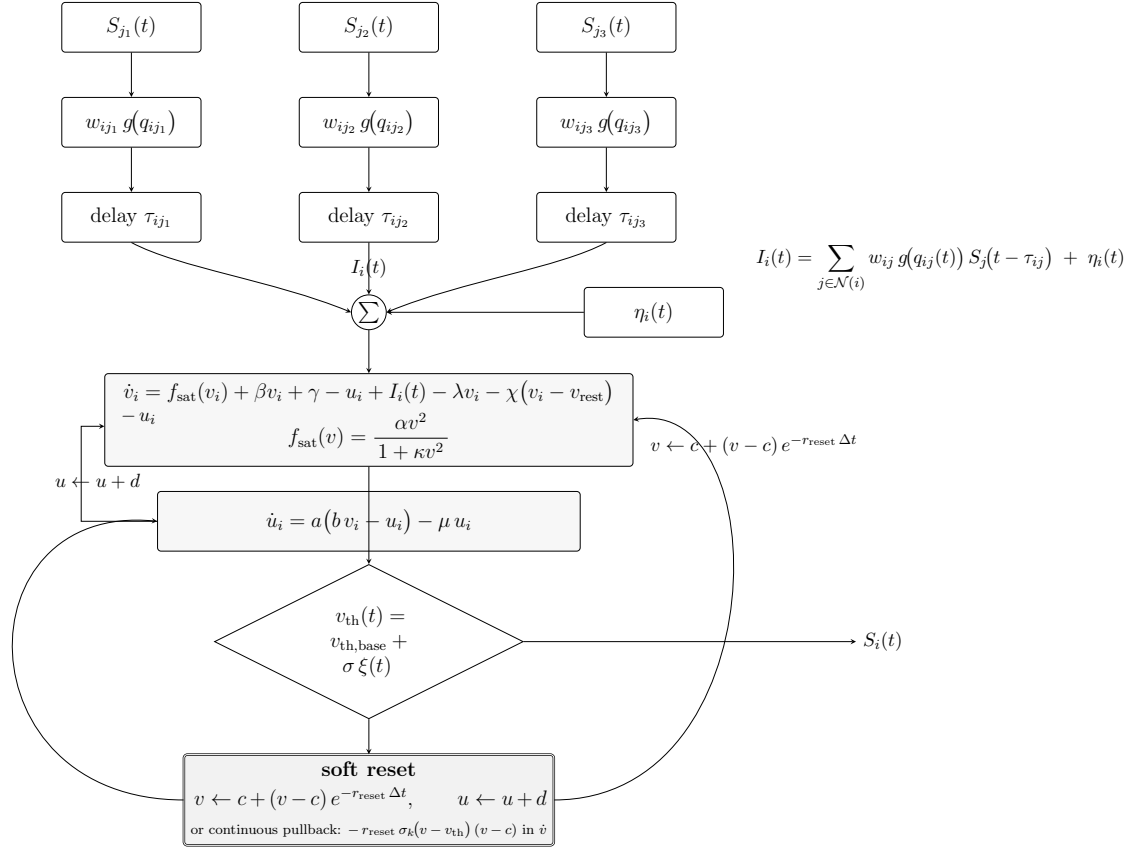
Continuous differentiable pullback. Alternatively, we add to the v -dynamics in (9) a smooth term that engages only above threshold:

$$-r_{\text{reset}} \sigma_k(v_i - v_{\text{th}}(t)) (v_i - c), \quad \sigma_k(x) = \frac{1}{1 + e^{-kx}}, \quad r_{\text{reset}} > 0. \quad (34)$$

For $v_i < v_{\text{th}}(t)$ the multiplier σ_k is near 0 and the pullback is negligible. Once v_i exceeds the threshold, σ_k rises toward 1 and an exponential attraction to c turns on as a continuous part of \dot{v} . Training proceeds with a homotopy in k , starting smooth and sharpening as optimisation stabilises, which preserves gradient quality while converging to crisp event timing. For reference, the nullclines of the no–input skeleton are

$$\dot{v} = 0 : u = f_{\text{sat}}(v) + (\beta - \lambda - \chi)v + \gamma + \chi v_{\text{rest}}, \quad (35)$$

$$\dot{u} = 0 : u = \frac{ab}{a + \mu} v \quad (36)$$



Per-link weight and gate $w_{ij} g(q_{ij})$, delay τ_{ij} , and shot-noise $\eta_i(t)$ feed a node with bounded excitability f_{sat} , recovery u_i , stochastic threshold, and differentiable reset. The threshold emits $S_i(t)$; reset updates feed back into v and u .

Figure 2: NOS unit with graph-local inputs, weighted per-link gates and delays, exogenous shot-noise, bounded excitability, recovery dynamics, stochastic threshold, and differentiable reset. Three non-overlapping lanes: left for $u \rightarrow v$, centre for $v \rightarrow v_{\text{th}}$, right for reset $\rightarrow v$.

Figure 1 clarifies the role of the smooth pullback in a networking setting. In panel 1a the skeleton geometry is benign: the parabolic v -nullcline (35) meets the linear u -nullcline (36) at a stable point well to the left of the threshold. In panel 1b, identical shot-noise drive nudges v upward but, with the reset disabled, there is no restoring action triggered at the moment of crossing; the trajectory can linger near v_{th} , which in practice sustains high send pressure on downstream links. In panel 1c, the same bursts cause a single threshold crossing followed by an immediate, continuous return toward c . The concurrent rise in u encodes a short conservative phase, matching paced drain or token-bucket refill after an alarm. Because (34) is differentiable, these operator-visible effects are trainable: r_{reset}^{-1} is tuned to device drain time or scheduler epoch, c to the desired post-event baseline, and d to the depth of temporary slow-down. The model therefore reproduces both the timing of decisions and the engineered cool-down that follows, without introducing discontinuities that harm gradient-based learning or blur attribution. Fig. 2 illustrates a single NOS unit, highlighting its graph-local inputs, excitability dynamics, recovery variable, stochastic threshold, and differentiable reset pathway.

Table 3: Operational interpretation and *data-driven initialisation* of *NOS* parameters from observed network metrics. Defaults are starting values and a bin width of 5 ms; they are not hard limits.

Parameter	Observable(s)	Initialisation rule	Typical default
v	Queue length or buffer occupancy	Affine map full buffer $\rightarrow 1$, empty $\rightarrow 0$	—
v_{th}	False-alarm budget on residuals	Choose so $\Pr(z > z_\tau) = p_{\text{FP}}$ on train; set threshold at z_τ	0.60
λ	Mean service rate μ_{svc}	Per-bin rate: $\lambda \approx \mu_{\text{svc}} \Delta t$	0.18 per bin
χ	Small-signal damping	Fit AR(1) on subthreshold v : $v_{t+1} \approx (1 - \lambda - \chi)v_t + \dots$	0.02–0.05 per bin
β	Slope of dv/dt vs v	Regress $\Delta v/\Delta t$ on v on train residuals	0.05–0.20
γ	Baseline load	Intercept from the same regression, or set mean residual ≈ 0	0.06–0.10
α, κ	Nonlinear ramp steepness	Fit $f_{\text{sat}}(v) = \alpha v^2/(1 + \kappa v^2)$ to rising edges	$\alpha = 0.7, \kappa = 1.0$
a	Post-burst relaxation time τ_{rec}	$a \approx \Delta t/\tau_{\text{rec}}$ (per bin) from exponential fit	1.1 per bin
b	Recovery sensitivity to v	Tune so u tracks v during decay without overshoot	1.0–1.2
μ	Passive decay of u	Small regulariser on u drift; set by validation	0.05–0.15 per bin
c	Post-event baseline of v	Median v in a quiet window after events	0.10
d	Recovery jump on event	Choose to match observed refractory depth	0.20–0.30
k_{reset}	Desired sharpness of reset	Large enough to mimic hard reset without instability	14
r_{reset}	Reset time constant τ_{reset}	$r_{\text{reset}} \approx \Delta t/\tau_{\text{reset}}$ (per bin)	5 per bin
I_0, gain	Arrivals $\rightarrow I$ mapping	Regress I on smoothed arrivals: $I \approx I_0 + \text{gain} \cdot \text{arrivals}$	$I_0 = 0.10, \text{gain} = 1.0$
τ_s	Burst smoothing need	Low-pass if arrivals are spiky; choose cut-off by validation	0–10 ms
w_{ij}	Link rates/priorities	Proportional to nominal bandwidth or policy weight; normalise $\rho(W) = 1$	—
g	Desired coupling index k_{net}	Set $g = k_{\text{net}}/\rho(W)$; pick $k_{\text{net}} \in [0.8, 1.6]$ for stress	topology-specific
τ_{ij}	Per-link delay	From RTT or profiling; use queue-free component	0–25 ms

Per-bin conversion uses $\Delta t = 5$ ms. For per-second rates multiply by 200. Defaults are the values used unless stated; sensitivity to $\{\lambda, v_{\text{th}}, \alpha, \kappa\}$ is reported in the Appendix F.

3.8 Non-dimensionalisation and scaling

Operational traces arrive with different bin widths, device rates, and buffer sizes. To make tuning portable across such heterogeneity, we rescale state and time by fixed references V and T , and work with

$$v(t) = V \tilde{v}(\tilde{t}), \quad t = T \tilde{t}. \quad (37)$$

Choose V as the full-buffer level used in the queue normalisation (for example, bytes or packets mapped to $v \in [0, 1]$), and choose T as the dominant local timescale: the scheduler epoch, the service drain time $1/\lambda$, or the telemetry bin width when that is the binding constraint. With (37), the left-hand side of (9) satisfies

$$\frac{dv}{dt} = \frac{V}{T} \frac{d\tilde{v}}{d\tilde{t}}, \quad (38)$$

and each term on the right-hand side of (9) is re-expressed as a dimensionless group. The bounded excitability transforms as

$$f_{\text{sat}}(v) = \frac{\alpha v^2}{1 + \kappa v^2} \rightarrow \tilde{f}_{\text{sat}}(\tilde{v}) = \frac{T}{V} f_{\text{sat}}(V\tilde{v}) = \frac{\tilde{\alpha} \tilde{v}^2}{1 + \tilde{\kappa} \tilde{v}^2}, \quad \tilde{\alpha} = \alpha TV, \quad \tilde{\kappa} = \kappa V^2. \quad (39)$$

Linear and constant terms scale as

$$\tilde{\beta} = \beta T, \quad \tilde{\lambda} = \lambda T, \quad \tilde{\chi} = \chi T, \quad \tilde{\gamma} = \frac{\gamma T}{V}, \quad \tilde{v}_{\text{rest}} = \frac{v_{\text{rest}}}{V}. \quad (40)$$

The recovery dynamics (10) become

$$\frac{du}{dt} = a(bv - u) - \mu u \rightarrow \frac{d\tilde{u}}{d\tilde{t}} = \tilde{a}(\tilde{b}\tilde{v} - \tilde{u}) - \tilde{\mu}\tilde{u}, \quad \tilde{a} = aT, \quad \tilde{\mu} = \mu T, \quad \tilde{b} = b, \quad \tilde{u} = \frac{u}{V} \quad (41)$$

and the graph-local input (8) and its delayed form (22) scale as

$$\tilde{I}_i(\tilde{t}) = \frac{T}{V} I_i(t), \quad \tilde{\tau}_{ij} = \frac{\tau_{ij}}{T}, \quad \tilde{w}_{ij} = w_{ij} \quad (\text{dimensionless by construction}). \quad (42)$$

For the shot-noise drive (28), amplitudes and smoothing follow

$$\tilde{\eta}_i(\tilde{t}) = \frac{T}{V} \eta_i(t) = \sum_{n=1}^{N_i(t)} \tilde{A}_{i,n} \tilde{\kappa}_s(\tilde{t} - \tilde{t}_{i,n}), \quad \tilde{A}_{i,n} = \frac{T}{V} A_{i,n}, \quad \tilde{\kappa}_s(\tilde{t}) = e^{-\tilde{t}/\tilde{\tau}_s} H(\tilde{t}), \quad \tilde{\tau}_s = \frac{\tau_s}{T}. \quad (43)$$

With (37)–(43), raw telemetry from sites that sample at 1 ms and sites that sample at 10 ms are mapped to the same dimensionless dynamics. The same holds for devices with different buffers: picking V from the local full-buffer level preserves the meaning of $\tilde{v} \in [0, 1]$. Operators can therefore compare fitted parameters across topologies and hardware: $\tilde{\lambda}$ reflects service relative to the chosen time base; \tilde{a} reflects recovery rate in scheduler units; $\tilde{\alpha}$ and $\tilde{\kappa}$ capture the curvature and knee of the congestion ramp independent of absolute queue size. Delays and gates adopt the same rule: $\tilde{\tau}_{ij}$ is RTT-free propagation and processing in units of T , while w_{ij} stays a policy or bandwidth proportion that is already dimensionless.

Reading guide and forward links. The local slope bound in (15) and the service-damping aggregate in (20) feed directly into the small-signal analysis in §5.2 and the network coupling results in §5.3. Let $\rho(W)$ be the spectral radius of the coupling matrix and g any global gain applied to W . Linearisation of (9)–(10) about an operating point v^* yields a block Jacobian whose dominant network term is proportional to $g \rho(W) f'_{\text{sat}}(v^*)$. Compare this against the net drain

$$\Lambda := \lambda + \chi + \frac{ab}{a + \mu} - \beta, \quad (44)$$

which is the left side of (20) minus the right side at the operating point. A convenient stability proxy is

$$g \rho(W) f'_{\text{sat}}(v^*) < \Lambda. \quad (45)$$

Under the rescaling (37), both sides scale by the same factor: $T f'_{\text{sat}}$ and $T \Lambda$ are dimensionless, so (45) is invariant and can be checked in either set of units. The proxy exposes the same levers used in operations. One can reduce $\rho(W)$ by reweighting or sparsifying couplings, reduce g through gain scheduling, or raise Λ by increasing service λ , damping χ , or the recovery rate a (with the trade-off set by b and μ). It is useful to track the *operational margin*

$$\Delta_{\text{net}} := \Lambda - g \rho(W) f'_{\text{sat}}(v^*), \quad (46)$$

which we will use in §5.3 to explain bifurcation onsets, headroom under topology changes, and how weight normalisation can enforce a fixed margin across deployments.

4 Design principles and engineering guidance

The *NOS* configuration is *straightforward*: it follows a small set of choices that mirror device behaviour in operational networks. Figure 3 summarises how a NOS unit is used in practice: telemetry pre-processing aligns with the smoothing in (28)–(31); the state (v, u) evolves by (9)–(10) with resets (33)/(34); events and residuals drive local detection and control; and spikes propagate over W with delays (22) to form neighbours' inputs. This sets the stage for the engineering choices that follow (reset time scale, EWMA leak χ , gain scheduling g vs. $\rho(W)$, and threshold calibration).

For resets, one may use either an event-based exponential smoothing at threshold crossings or a continuous sigmoidal pullback that activates only above threshold (cf. §3.7); both are differentiable and avoid algebraic jumps, so gradient flow remains stable and attribution stays clear. In deployment terms, the return-to-baseline time is matched to the hardware drain or scheduler epoch, which means the model cools down on the same clock as the switch or NIC. Subthreshold prefiltering enters as a mild leak $\chi(v - v_{\text{rest}})$ that behaves like an EWMA on residual queues: it reveals early warnings while damping tiny oscillations caused by counter jitter, yet preserves micro-bursts that matter for control.

Exogenous burstiness and measurement noise are included explicitly. We use shot-noise drive as in §3.6 and allow slow adaptation of tolerance and recovery so the unit becomes temporarily conservative during sustained stress. Concretely, the threshold may wander within a bounded band and the recovery rate may track recent activity

$$v_{\text{th}}(t) = v_{\text{th,base}} + \sigma \xi(t), \quad a(t) = a_0 + \kappa \bar{S}_i(t), \quad (47)$$

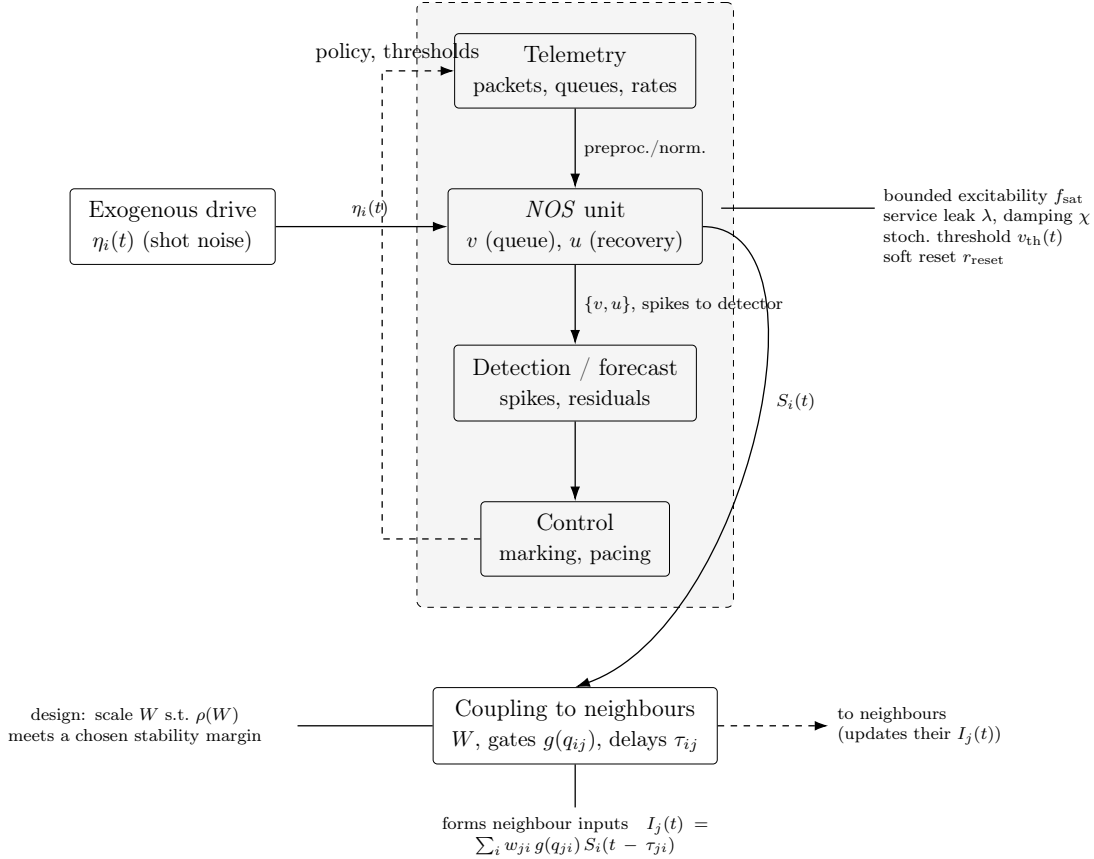


Figure 3: Distributed flow around a NOS unit. Telemetry and exogenous shot noise drive the unit. Its state (v, u) feeds detection/control; spikes $S_i(t)$ also propagate to neighbours via weighted, gated, and delayed couplings to form their inputs $I_j(t)$. Control policies feed back to telemetry normalisation (dashed).

where $\xi(t)$ is bounded noise that sets the false-alarm *budget* and $\bar{S}_i(t)$ is the short-window spike (arrival) rate. These adaptations match operator practice such as temporary pacing, token-bucket refill, or ECN-driven headroom.

Training follows standard practice while keeping the networking semantics intact. We employ surrogate gradients at the threshold with a fast-sigmoid derivative

$$\widehat{\sigma}'(x) = \frac{1}{(1 + \alpha|x|)^2}, \quad \alpha \approx 5, \quad (48)$$

which keeps gradients finite and centred. To approach crisp events without destabilising optimisation, we use a homotopy on the pullback sharpness k that starts smooth and tightens once optimisation stabilises,

$$k(t) = k_0 + (k_{\text{final}} - k_0) \frac{t}{T_h}, \quad k_0 \approx 1, \quad k_{\text{final}} \in [20, 100], \quad (49)$$

which preserves gradient quality early and matches event timing late. We also clip gradients by global norm in the range $[0.5, 2.0]$ while k grows. An adaptive optimiser such as Adam is suitable; the learning rate is reduced as k increases to maintain stable steps near the threshold. Truncated BPTT should cover the dominant recovery timescale so

that u 's memory is learned rather than aliased,

$$T_{\text{BPTT}} \gtrsim \frac{c_*}{a + \mu}, \quad c_* \in [3, 5] \text{ (capturing about 95\%–99\%)}, \quad (50)$$

which aligns with the cool-down used by paced drain or token-bucket refill in devices. After supervised pre-training, slow local updates can be enabled to track diurnal load or policy shifts.

The network view stays explicit. Delays τ_{ij} and gates $g(q_{ij}(t))$ ensure that path timing and residual capacity modulate influence at the right hop and time. The small-signal proxy in (45) makes the stability levers visible to operations: one can lower $\rho(W)$ by reweighting or sparsifying couplings, reduce the global gain g by scheduling, or increase the net drain Λ by raising service λ , damping χ , or recovery a (with b and μ setting trade-offs). These are the same knobs used in practice to keep queues bounded and alerts meaningful, and they connect directly to the analysis that follows.

4.1 Neuromorphic deployment

NOS maps cleanly to event-driven hardware. Inbound packets become address-event representation spikes; the graph term $I_i(t)$ is realised by on-chip routers that natively implement sparse fan-out with per-link delay lines. The state (v, u) is stored in fixed point, and the soft resets in (33)–(34) are implemented as leaky updates or gated pullbacks without algebraic jumps. This preserves gradient surrogates for training and keeps timing aligned with scheduler epochs and drain times. For networking workloads, the benefits are direct: energy per decision scales with the spike rate rather than the link count; delays τ_{ij} become programmable pipeline stages; and per-link gates $g(q_{ij})$ mirror queue-aware policing or ECN. Practical quantisation rules, fabric-rate constraints, and an operational margin check under finite precision are detailed in Appendix B.

5 Equilibrium and local stability analysis

We analyse subthreshold operating points of a *NOS* node and then lift the analysis to the coupled network, treating resets as short transients that do not change the slow geometry. At equilibrium the node holds a steady queue level set by arrivals, service, and graph-local coupling. We first derive the equilibrium equations and conditions for existence and uniqueness. We then linearise to obtain the Jacobian and local stability criteria. Next we introduce network coupling, form the block Jacobian, and relate stability to the graph spectrum, which yields an explicit stability threshold and an operational margin. We classify the ensuing bifurcations and read them in operational terms (onset of ringing, loss of headroom, topology-driven tipping), and we quantify noise sensitivity under stochastic drive to explain false alarms and recovery delays. Finally, we compare the subthreshold steady state with the $M/M/1$ queueing baseline to anchor *NOS* against classical service-arrival models.

5.1 Equilibrium equations

For node i , let (v_i^*, u_i^*) denote a subthreshold equilibrium ($\dot{v}_i = \dot{u}_i = 0$). The (10) implies a linear relation between recovery and queue level:

$$u_i^* = \frac{ab}{a + \mu} v_i^*. \quad (51)$$

Substituting (51) into (9) gives

$$0 = f_{\text{sat}}(v_i^*) + (\beta - \lambda - \chi) v_i^* + \gamma + \chi v_{\text{rest}} - u_i^* + I_i^*, \quad (52)$$

where I_i^* denotes the long-run mean of the graph-local drive at node i (including delayed contributions). Eliminating u_i^* yields a scalar fixed-point equation for v_i^* :

$$f_{\text{sat}}(v_i^*) + \left(\beta - \lambda - \chi - \frac{ab}{a + \mu} \right) v_i^* + \gamma + \chi v_{\text{rest}} + I_i^* = 0. \quad (53)$$

Equation (53) balances three effects: a convex rise with load through f_{sat} , linear drain through the net service and damping, and the constant baseline from $\gamma + \chi v_{\text{rest}} + I_i^*$. The mean input I_i^* aggregates neighbour activity over the network with delays; under stationary traffic it is well defined because the delayed sum in (22) averages to a constant.

Quadratic small-signal approximation. For analytic visibility, replace $f_{\text{sat}}(v) \approx \alpha v^2$ in a small-signal regime and define

$$L := \beta - \lambda - \chi - \frac{ab}{a + \mu}, \quad C := \gamma + \chi v_{\text{rest}} + I_i^*. \quad (54)$$

Then (53) reduces to

$$\alpha(v_i^*)^2 + L v_i^* + C = 0, \quad (55)$$

with discriminant

$$D_i = L^2 - 4\alpha C. \quad (56)$$

with discriminant D_i . If $D_i \geq 0$ and the admissible root lies in $[0, 1]$, it approximates v_i^* and clarifies how the operating point moves with C . In networks, service and damping must outpace the steepest small-signal gain of f_{sat} ; the corollary in 3.4 provided a direct algebraic check. In the quadratic approximation, the sign of D_i in (56) separates three regimes: $D_i < 0$ (no real root, parameter mismatch), $D_i = 0$ (tangent balance at a single v_i^*), and $D_i > 0$ (two mathematical roots, of which the physically admissible one lies in the queue range and satisfies the stability test below).

Existence and uniqueness on $[0, 1]$. The monotonicity criterion introduced earlier applies with L as in (54): a sufficiently positive net slope $f'_{\text{sat}}(v) + L$ over the admissible interval (for example $v \in [0, 1]$) ensures a unique solution. For the saturating choice

$$f_{\text{sat}}(v) = \frac{\alpha v^2}{1 + \kappa v^2}, \quad \alpha > 0, \quad \kappa > 0, \quad (57)$$

the slope

$$f'_{\text{sat}}(v) = \frac{2\alpha v}{(1 + \kappa v^2)^2} \quad (58)$$

is nonnegative and bounded on $[0, 1]$, with $\max_{v \in [0,1]} f'_{\text{sat}}(v) \leq \frac{3\sqrt{3}}{8} \frac{\alpha}{\sqrt{\kappa}}$. If

$$\inf_{v \in [0,1]} (f'_{\text{sat}}(v) + L) > 0, \quad (59)$$

then F is strictly increasing on $[0, 1]$ and has at most one root there. If in addition $F(0)F(1) \leq 0$, a unique $v_i^* \in [0, 1]$ exists. Condition (59) is met whenever

$$L > -\frac{3\sqrt{3}}{8} \frac{\alpha}{\sqrt{\kappa}}, \quad (60)$$

which is the small-signal service-dominance requirement: linear damping $\lambda + \chi + \frac{ab}{a+\mu}$ must exceed the maximal small-signal excitability slope.

Networking interpretation. Equation (53) balances offered load against service and damping. The term I_i^* aggregates graph-local arrivals and policy weights; L collects all linear drains. The sensitivity to slowly varying load follows by implicit differentiation:

$$\frac{\partial v_i^*}{\partial I_i^*} = -\frac{1}{f'_{\text{sat}}(v_i^*) + L}. \quad (61)$$

Hence the DC gain is stable and interpretable. It shrinks as saturation increases ($\kappa \uparrow$) or as service and damping increase ($\lambda, \chi, ab/(a + \mu) \uparrow$), which matches the operational goal of keeping the queue operating point insensitive to modest load drift.

5.2 Jacobian and linear stability

Linearising *NOS* node (9)–(10) at a subthreshold equilibrium (v^*, u^*) gives the single-node Jacobian

$$J_i = \begin{bmatrix} f'(v^*) + \beta - \lambda - \chi & -1 \\ ab & -(a + \mu) \end{bmatrix}, \quad (62)$$

with trace and determinant

$$T_i = \text{tr}(J_i) = f'(v^*) + \beta - \lambda - \chi - (a + \mu), \quad (63)$$

$$\Delta_i = \det(J_i) = ab - (f'(v^*) + \beta - \lambda - \chi)(a + \mu). \quad (64)$$

The Routh–Hurwitz conditions for a two-dimensional system yield local asymptotic stability if and only if

$$T_i < 0 \quad \text{and} \quad \Delta_i > 0. \quad (65)$$

Networking reading. The trace condition (65) asks that net drain $(a + \mu) + (\lambda + \chi - \beta)$ exceeds the small-signal growth from $f'(v^*)$. The determinant condition requires that the product of recovery gain ab and time-scale $(a + \mu)$ dominates the same growth term. Together they state that service plus damping and the recovery loop must remove perturbations faster than excitability amplifies them. When either inequality tightens, small oscillations and slow return-to-baseline emerge, which operators observe as ringing in queue telemetry.

Specialisations. *Quadratic small-signal.* For small $|v^*|$, use $f'(v^*) \approx 2\alpha v^*$ and test (65) with that slope.

Bounded excitability. For $f_{\text{sat}}(v) = \frac{\alpha v^2}{1 + \kappa v^2}$,

$$f'(v^*) = \frac{2\alpha v^*}{(1 + \kappa(v^*)^2)^2}, \quad (66)$$

$$T_i = \frac{2\alpha v^*}{(1 + \kappa(v^*)^2)^2} + \beta - \lambda - \chi - (a + \mu), \quad (67)$$

$$\Delta_i = ab - \left(\frac{2\alpha v^*}{(1 + \kappa(v^*)^2)^2} + \beta - \lambda - \chi \right) (a + \mu). \quad (68)$$

As κ grows, the effective slope (66) decreases, enlarging the stability region defined by (65).

Lemma 2 (Monotone stabilisation by saturation). *Let $f_{\text{sat}}(v) = \frac{\alpha v^2}{1 + \kappa v^2}$ with $\alpha > 0$, $\kappa > 0$, and let $v^* \geq 0$ be a subthreshold equilibrium. Then*

$$\frac{\partial}{\partial \kappa} f'(v^*) = -\frac{4\alpha(v^*)^3}{(1 + \kappa(v^*)^2)^3} \leq 0, \quad (69)$$

with equality only at $v^* = 0$. Consequently,

$$\begin{aligned} \frac{\partial T_i}{\partial \kappa} &= \frac{\partial f'(v^*)}{\partial \kappa} < 0 \quad (v^* > 0), \\ \frac{\partial \Delta_i}{\partial \kappa} &= -(a + \mu) \frac{\partial f'(v^*)}{\partial \kappa} > 0 \quad (v^* > 0), \end{aligned} \quad (70)$$

so increasing κ strictly decreases the trace and strictly increases the determinant, enlarging the region $\{T_i < 0, \Delta_i > 0\}$.

Proof sketch. Differentiate $f'(v) = \frac{2\alpha v}{(1 + \kappa v^2)^2}$ with respect to κ :

$$\frac{\partial}{\partial \kappa} f'(v) = -\frac{4\alpha v^3}{(1 + \kappa v^2)^3}, \quad (71)$$

which is nonpositive for $v \geq 0$ and strictly negative for $v > 0$. Substituting into (67)–(68) gives (70). \square

Corollary 2.1 (Network threshold increases with saturation). *Under homogeneous coupling $G = gW$ with Perron eigenvalue $\rho(W)$, define*

$$\bar{d} := f'(v^*) + \beta - \lambda - \chi, \quad k^* = \min\left\{ (a + \mu) - \bar{d}, \frac{ab}{a + \mu} - \bar{d} \right\}, \quad g_* \approx \frac{k^*}{\rho(W)}. \quad (72)$$

Since $\frac{\partial \bar{d}}{\partial \kappa} = \frac{\partial f'(v^*)}{\partial \kappa} < 0$ for $v^* > 0$, one has $\frac{\partial k^*}{\partial \kappa} > 0$ and therefore $\frac{\partial g_*}{\partial \kappa} > 0$ on the active branch of k^* . Increasing κ raises the coupling needed to destabilise the network, which matches the bounded-excitability design goal.

In networking terms, making the excitability saturate earlier reduces small-signal gain and raises damping, which protects the node against oscillations when operated near capacity.

Operational guidance. When tuning for a given site, estimate v^* from the modal subthreshold queue level and compute the small-signal slope $f'(v^*)$ using the fitted (α, κ) . If (65) is tight, increasing κ (earlier saturation) or increasing λ and χ adds headroom. If tightening occurs only in Δ_i , adjust the recovery loop (a, b, μ) to increase ab or the composite time-scale $(a + \mu)$; both choices shorten the memory of recent stress and prevent ringing after bursts. Decay times and DC sensitivity that guide window length and reset tuning are summarised in Appendix C, Eqs. (98) and (99)

5.3 Network coupling and global stability

In isolation a *NOS* unit settles to a load–service balance set by its local parameters. In a network the steady input is shaped by neighbours and policy weights, so existence and robustness of equilibria depend on both the operating point and the coupling matrix. Let $W = [w_{ij}]$ be the graph coupling and let S_j^* denote the effective steady presynaptic drive at node j (after any per-link delays and gates). The steady input at node i is

$$I_i^* = \sum_j w_{ij} S_j^*, \quad \max_i I_i^* \leq \|W\|_\infty \|S^*\|_\infty, \quad \|I_i^*\|_2 \leq \|W\|_2 \|S^*\|_2 \quad (73)$$

With $\|W\|_2 = \sqrt{\rho(W^\top W)}$, and the spectral radius obeys $\rho(W) \leq \|W\|_\infty$, so $\|W\|_\infty$ gives a conservative topology-aware cap on steady drive. From the scalar equilibrium reduction (cf. (51)–(56)), define

$$L := \beta - \lambda - \chi - \frac{ab}{a + \mu}, \quad C := \gamma + \chi v_{\text{rest}} + I_i^*, \quad (74)$$

and write the per-node equilibrium condition as

$$f_{\text{sat}}(v_i^*) + L v_i^* + C = 0. \quad (75)$$

Existence bounds. Two complementary sufficient conditions are useful in operations.

Saturated cap (conservative). Since $f_{\text{sat}}(v) \leq \alpha/\kappa$ for all v ,

$$C \leq \frac{\alpha}{\kappa} \implies \text{a nonnegative equilibrium exists when } L \leq 0. \quad (76)$$

This bound uses only the saturation ceiling and is independent of the local operating point.

Quadratic small-signal. Approximating $f_{\text{sat}}(v) \approx \alpha v^2$ near light load $v = 0$ gives

$$\alpha (v_i^*)^2 + L v_i^* + C = 0, \quad D_i = L^2 - 4\alpha C, \quad (77)$$

so an equilibrium exists if

$$C \leq \frac{L^2}{4\alpha}. \quad (78)$$

We refer to (78) as the *quadratic small-signal* bound. Combining (78) with the norm bound in (73) yields

$$\gamma + \chi v_{\text{rest}} + \|W\|_\infty \|S^*\|_\infty \leq \frac{L^2}{4\alpha}, \quad (79)$$

with a spectral alternative obtained by replacing $\|W\|_\infty \|S^*\|_\infty$ with $\|W\|_2 \|S^*\|_2$. Inequality (79) makes the graph explicit: heavier row sums (dense fan-in or large policy weights) shrink headroom for the same steady S^* .

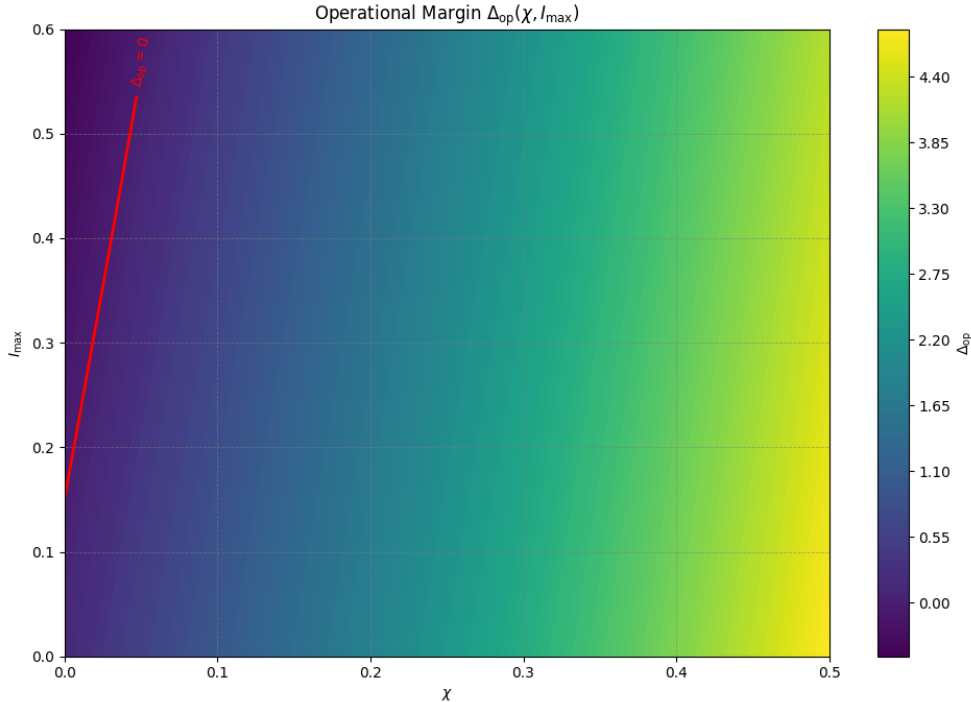


Figure 4: Operational margin $\Delta_{\text{op}}(\chi, I_{\text{max}})$ from (80). Shading shows Δ_{op} (dimensionless); the red contour marks $\Delta_{\text{op}} = 0$. Points *above the red curve* (larger I_{max} for the same χ) lie beyond the small-signal existence bound. Parameters $\{\alpha, \kappa, \lambda, \beta, a, b, \mu, \gamma, v_{\text{rest}}\}$ are held fixed as in the text; $I_{\text{max}} = \max_i I_i^*$. See Appendix G Fig. 15 for one-dimensional sweeps and sensitivity to $\rho(W)$ and g .

Operational margin. A convenient scalar indicator is

$$\Delta_{\text{op}} := \frac{L^2}{4\alpha} - \left(\gamma + \max_i I_i^*\right), \quad (80)$$

which is positive when the quadratic small-signal existence test passes with slack. See Appendix G for the conservative one-dimensional sweeps of the operational margin Δ_{op} versus I_{max} , a worked numerical example, and the linear dependence on load illustrated in Fig. 15. In practice, keeping $\Delta_{\text{op}} > 0$ under nominal load provides headroom for burstiness and guides weight normalisation: scale W so that the largest row sum keeps $\max_i I_i^*$ below the right-hand side of (79). When delays τ_{ij} are present, (79) remains a conservative existence test; sharper delay-aware guarantees require the block Jacobian in the next subsection.

Operational margin under load and damping. Figure 4 plots $\Delta_{\text{op}}(\chi, I_{\text{max}})$ as a function of the subthreshold damping χ and the maximum steady input $I_{\text{max}} = \max_i I_i^*$. The red contour $\Delta_{\text{op}} = 0$ separates admissible operating points from those that violate the quadratic small-signal bound. From a networking viewpoint, χ measures the strength of subthreshold queue relaxation, while I_{max} is the heaviest sustained offered load after graph coupling. Two effects are visible: (i) the margin decreases monotonically with I_{max} , so higher offered load erodes headroom; (ii) larger χ shifts the boundary outward, permitting higher sustained load before crossing $\Delta_{\text{op}} = 0$. Stronger damping therefore improves resilience but does not remove the inherent load–margin trade-off.

Networking view. $\|W\|_\infty$ aggregates worst-case inbound influence at a node, so it reflects heavy in-degree or imbalanced policies; $\|W\|_2$ captures coherent multi-hop reinforcement. Both provide simple levers: reweight heavy rows, sparsify fan-in, or gate selected edges until Δ_{op} is positive with a chosen margin.

5.4 Block Jacobian, network spectrum, and stability threshold

We assess global small-signal stability by linearising the coupled *NOS* network about a subthreshold equilibrium (v^*, u^*) . Let δv and δu be perturbations of state, and let small variations of presynaptic drive satisfy $\delta I \approx G \delta v$, where G is the input-output sensitivity induced by topology (for homogeneous scaling, $G = gW$). Linearisation of (9)–(10) then yields the $2N \times 2N$ block Jacobian

$$\mathcal{J} = \begin{bmatrix} D + G & -I_N \\ aB & -(a + \mu)I_N \end{bmatrix}, \quad D = \text{diag}(f'(v_i^*) + \beta - \lambda - \chi), \quad B = bI_N. \quad (81)$$

Instability occurs when any eigenvalue of \mathcal{J} crosses into the open right half-plane.

Perron-mode reduction under homogeneous coupling. For homogeneous coupling $G = gW$ with $g \geq 0$ and let $\rho(W)$ be the spectral radius with Perron eigenpair $(\rho(W), \phi)$. Projecting (81) onto ϕ yields an effective 2×2 Jacobian

$$J_{\text{eff}}(k) = \begin{bmatrix} \bar{d} + k & -1 \\ ab & -(a + \mu) \end{bmatrix}, \quad k := g \rho(W), \quad \bar{d} := \frac{1}{\|\phi\|_2^2} \sum_i \phi_i^2 [f'(v_i^*) + \beta - \lambda - \chi]. \quad (82)$$

In the Perron-mode reduction we define $\bar{d} := \frac{1}{\|\phi\|_2^2} \sum_i \phi_i^2 [f'(v_i^*) + \beta - \lambda - \chi]$, which is exact for heterogeneous equilibria. Throughout the subsequent analysis we write $\bar{d} \approx f'_{\text{sat}}(v^*) + \beta - \lambda - \chi$ for homogeneous or mildly heterogeneous whenever $v_i^* \approx v^*$ across nodes (or the Perron vector is sufficiently delocalised). All threshold formulas (e.g. (83) and (85)) remain valid with either form by substituting the appropriate \bar{d} for the deployment at hand. The Routh–Hurwitz conditions for $J_{\text{eff}}(k)$ give the critical scalar coupling

$$k^* = \min \left\{ (a + \mu) - \bar{d}, \frac{ab}{a + \mu} - \bar{d} \right\}, \quad g_* \approx \frac{k^*}{\rho(W)}. \quad (83)$$

Using $\Lambda := \lambda + \chi + \frac{ab}{a + \mu} - \beta$ (cf. (44)), the threshold condition can be written as

$$g \rho(W) f'_{\text{sat}}(v^*) < \Lambda, \quad (84)$$

which matches the stability proxy introduced in the scaling discussion. While the Perron-mode reduction yields tight thresholds in practice, conservative certificates can be obtained via Gershgorin discs and norm bounds; see Appendix H for the resulting inequalities and their comparison to measured g_* in Table 8.

Networking interpretation and control. Figure 5 shows that the first loss of stability is governed by the Perron mode of W . As coupling g increases, the leading eigenvalue crosses the imaginary axis at g_* , marking the onset of coherent queue growth along the dominant influence path. The collapse of $\rho(W) g_*$ onto k^* confirms the separation of




Figure 5: Spectral verification for homogeneous coupling $G = gW$. Left: real part of the leading eigenvalue of the block Jacobian versus g for several $\rho(W)$; the zero crossing defines g_* . Right: collapse of $\rho(W)g_*$ to the scalar threshold k^* predicted by the Perron–mode reduction in (83). The small spread reflects finite size and random matrix variability.

concerns in (83): $\rho(W)$ captures how strongly topology reinforces load; k^* depends only on node physics through (\bar{d}, a, b, μ) . Operationally, the levers are direct. Reduce $\rho(W)$ by reweighting heavy fan-in or pruning selected edges. Reduce g through gain scheduling on stressed regions or policy throttles that weaken $I \mapsto v$ sensitivity. Raise the drain margin Λ by increasing service λ , adding subthreshold damping χ , or shortening recovery via a (with b, μ setting the trade). Because $f'_{\text{sat}}(v^*)$ decreases as saturation strengthens (larger κ), bounded excitability expands the stable range in (84).

Finite-size and heterogeneity effects. The rescaled thresholds are not perfectly flat. They sit slightly above k^* across $\rho(W)$, with larger deviations in smaller or more heterogeneous graphs. Two mechanisms explain this percent-level bias. First, nearby non-Perron modes and row-sum variability shift the leading root when the Perron vector localises on hubs, a common feature of heavy-tailed topologies. Second, discrete sweeps of $\Re\lambda_{\max}(g)$ and mild non-normal amplification nudge the numerical crossing. These are finite-size corrections, not a change of scaling, and they diminish with larger N or more homogeneous weights. For conservative guarantees during early deployment, replace $\rho(W)$ by norm bounds or apply Gershgorin discs to \mathcal{J} to obtain stricter, topology-aware thresholds. For deployments requiring analytic guarantees, Appendix H derives Gershgorin and $\|W\|_{\infty}$ -based bounds and reports their conservatism relative to measured g^* (Table 8).

5.5 Bifurcation classification, operational reading, and finite-size effects

Loss of stability as the mean input I increases (or as effective coupling $k = g\rho(W)$ grows) is produced by two local mechanisms. They can be diagnosed either from the scalar equilibrium map $F(v) = f_{\text{sat}}(v) + Lv + C(I)$ (fold tests) or from the 2×2 Jacobian at the operating point (Routh–Hurwitz tests).

1. **Saddle–node (SN).** The stable and unstable equilibria coalesce and disappear when the discriminant in (56) satisfies $D_i = 0$. Beyond this point there is no steady operating level: v drifts upward until resets and leakage dominate. In network terms, offered load or coherent reinforcement pushes a node beyond its admissible queue set; the model predicts sustained congestion with intermittent resets rather than recovery to a fixed level.
2. **Hopf.** The Jacobian trace crosses zero while the determinant remains positive ($T_i = 0$, $\Delta_i > 0$), creating a small oscillatory mode. In network terms, queues enter a self-excited cycle whose phase can entrain neighbors through W , producing rolling congestion waves that propagate along high-gain paths.

Parameter trends and admissibility. Continuation of equilibria v^* versus I reproduces queueing intuition and clarifies which onset is available. Increasing the service λ or the subthreshold damping χ shifts both SN and Hopf onsets to larger I (more headroom); increasing excitability α shifts them to smaller I (less headroom). The recovery coupling b controls whether a Hopf branch is admissible: the linear feedback ab must exceed the decay $(a + \mu)^2$ for an oscillatory onset to exist; otherwise only an SN occurs. Figure 6 shows $v^*(I)$ under sweeps of λ , α , and b ; marker coordinates are listed in Table 4. Near-coincident markers correspond to the codimension-two neighborhood $ab \approx (a + \mu)^2$, where a small change in recovery or damping switches the dominant failure mode from smooth tipping (SN) to oscillatory bursts (Hopf).

Empirical collapse. Across graphs, the measured threshold obeys

$$\rho(W) g_\star \approx \min\left\{(a + \mu) - \bar{d}, \frac{ab}{a + \mu} - \bar{d}\right\}, \quad (85)$$

with small upward deviations when the Perron vector localises or row-sums vary; cf. Fig. 5.

Networking interpretation. The two onsets have distinct operational signatures. An SN route produces a clean loss of a steady queue level and sustained saturation. Telemetry shows a monotone climb in v with resets, increased delay variance, and persistent ECN/markings without a clear period. A Hopf route produces narrowband oscillations: queues and drop/mark rates cycle with a characteristic period $\sim 2\pi/\sqrt{\Delta_i - (T_i/2)^2}$ at onset, and neighboring nodes oscillate in phase along the Perron mode of W . These patterns matter for mitigation. To extend stable throughput, raise λ or χ , or reduce α through stronger saturation (larger κ). To avoid self-sustained burstiness, keep ab comfortably larger than $(a + \mu)^2$ by shortening recovery time (a), reducing recovery sensitivity (b) when needed, or adding passive decay (μ). Each lever maps to standard controls: scheduler and line-rate settings (λ), AQM/leak policies (χ), congestion ramp shaping (α, κ), and backoff or pacing aggressiveness (a, b, μ).

Finite-size effects and coherence. In finite graphs the transition is slightly rounded but sharpens with size. The measured coupling threshold aligns with the Perron-mode prediction k^* (cf. (83)), with small upward deviations when the Perron vector localises on hubs or when secondary modes lie close in the spectrum. These are percent-level corrections that diminish in larger or more homogeneous topologies. Practically, one can monitor the rescaled margin $\rho(W) g - k^*$: values near zero predict imminent entrainment




Figure 6: Equilibrium v^* versus input I with saddle-node (■) and Hopf (▲) markers. **(Left)** Increasing service rate λ shifts both onsets to higher I . **(Middle)** Increasing excitability α advances both onsets. **(Right)** Varying recovery coupling b changes the balance between ab and $a + \mu$: when $ab > (a + \mu)^2$ an oscillatory Hopf branch is admissible, otherwise only an SN appears. Stable branches are shown in colour; unstable in faint grey. The exact marker coordinates used here are reported in Table 4.

Table 4: Continuation markers used in Fig. 6. Listed are the input levels I_{SN} and I_{H} and the corresponding equilibria v_{SN}^* , v_{H}^* . Rows with NaN in the Hopf columns occur where $ab \leq (a + \mu)^2$, so the oscillatory onset is not admissible under the chosen parameters.

Panel	λ	α	b	I_{SN}	v_{SN}^*	I_{H}	v_{H}^*
lambda sweep	0.0	1.0	2.0	1.198	1.095	1.111	0.800
lambda sweep	0.3	1.0	2.0	1.549	1.245	1.462	0.950
lambda sweep	0.6	1.0	2.0	1.945	1.395	1.858	1.100
alpha sweep	0.3	0.6	2.0	2.582	2.074	2.437	1.583
alpha sweep	0.3	1.0	2.0	1.549	1.245	1.462	0.950
alpha sweep	0.3	1.4	2.0	1.106	0.889	1.044	0.679
b sweep	0.3	1.0	0.6	0.404	0.636	NaN	NaN
b sweep	0.3	1.0	1.0	0.656	0.810	NaN	NaN
b sweep	0.3	1.0	1.6	1.146	1.071	1.132	0.950

of queues and the appearance of coherent oscillations along the dominant influence path. Appendix I quantifies network coherence via an order parameter $R(t)$, confirming that the onset clusters near k^* and sharpens with N (Fig. 16, Table 9).

Corollary 2.2 (Spectral collapse of the coupling threshold). *For homogeneous coupling $G = gW$, the instability threshold extracted from the block Jacobian satisfies (85), where $\bar{d} = f'_{\text{sat}}(v^*) + \beta - \lambda - \chi$. Percent-level upward deviations arise in finite or heterogeneous graphs when the Perron vector localises or row-sum variability is high, consistent with Fig. 5.*

5.6 Noise sensitivity under stochastic drive

We model arrivals as shot noise with exponential smoothing. For node i ,

$$\eta_i(t) = \sum_{n=1}^{N_i(t)} A_{i,n} e^{-(t-t_{i,n})/\tau_s} H(t - t_{i,n}), \quad N_i(t) \sim \text{Poisson}(\rho_i t), \quad (86)$$

where ρ_i is the event rate, $A_{i,n}$ are i.i.d. amplitudes, τ_s is a smoothing time, and H is the Heaviside step. The process is stationary with mean, variance, and autocovariance

$$\mathbb{E}[\eta_i] = \rho_i \mathbb{E}[A_i] \tau_s, \quad \text{Var}[\eta_i] = \rho_i \mathbb{E}[A_i^2] \frac{\tau_s}{2}, \quad \text{Cov}_{\eta_i}(\tau) = \rho_i \mathbb{E}[A_i^2] \frac{\tau_s}{2} e^{-|\tau|/\tau_s}. \quad (87)$$

Its power spectrum is Lorentzian,

$$S_{\eta_i}(\omega) = \frac{2 \rho_i \mathbb{E}[A_i^2] \tau_s}{1 + \omega^2 \tau_s^2}. \quad (88)$$

In networking terms, $(\rho_i, \mathbb{E}[A_i], \tau_s)$ are read from the same counters and windows used in telemetry: ρ_i is the burst-start rate, $\mathbb{E}[A_i]$ is the mean per-burst mass after smoothing, and τ_s matches the prefilter that removes counter jitter while preserving micro-bursts.

Small-signal sensitivity. Linearising *NOS* at a subthreshold equilibrium (v^*, u^*) with $\bar{d} = f'_{\text{sat}}(v^*) + \beta - \lambda - \chi$ gives

$$\begin{aligned} \frac{d}{dt} \begin{bmatrix} \delta v \\ \delta u \end{bmatrix} &= \begin{bmatrix} \bar{d} & -1 \\ ab & -(a + \mu) \end{bmatrix} \begin{bmatrix} \delta v \\ \delta u \end{bmatrix} + \begin{bmatrix} 1 \\ 0 \end{bmatrix} \eta_i(t), \\ H_v(s) &= \frac{\delta V(s)}{\mathbb{E}(s)} = \frac{s + (a + \mu)}{(s - \bar{d})(s + (a + \mu)) + ab}. \end{aligned} \quad (89)$$

The DC gain and the variance of δv under (88) are

$$H_v(0) = \frac{a + \mu}{ab - (a + \mu)\bar{d}}, \quad \sigma_v^2 = \frac{1}{2\pi} \int_{-\infty}^{\infty} |H_v(i\omega)|^2 S_{\eta_i}(\omega) d\omega. \quad (90)$$

As the node approaches its local margin ($ab \downarrow (a + \mu)\bar{d}$) or the network approaches the Perron-mode threshold ($k \rightarrow k^*$), $|H_v(0)|$ increases and the integral in (90) grows, so the same shot-noise trace produces larger queue excursions and more threshold crossings.

Firing statistics and cascades. We summarise sensitivity by the mean firing rate \bar{f} , the inter-spike-interval coefficient of variation (CV), and avalanche sizes S (contiguous above-threshold activity aggregated over neighbours). With parameters fixed and only the noise amplitude varied, three robust effects appear, consistent with (89)–(90) and bounded excitability: (i) \bar{f} rises with amplitude in all regimes, with a steeper slope near the coupling threshold because small increments in η_i map to larger δv ; (ii) CV increases near threshold, reflecting irregular, burst-dominated spike trains as variance grows; (iii) the tail of $P(S)$ becomes heavier with amplitude, indicating extended congestion cascades when fluctuations push multiple nodes above threshold before recovery completes.

Networking interpretation. Higher ρ_i or larger $\mathbb{E}[A_i^2]$ means burstier ingress; increasing τ_s models longer burst coherence. The levers that reduce noise sensitivity are the same that enlarge deterministic headroom: higher service and subthreshold damping (λ, χ), faster recovery ($a + \mu$), and earlier saturation (larger κ , which lowers $f'_{\text{sat}}(v^*)$ and hence \bar{d}). Each reduces $H_v(0)$ and the variance integral in (90). With per-link gates, attenuating w_{ij} on noisy edges reduces the effective drive entering I_i and moves the network away from the Perron threshold.

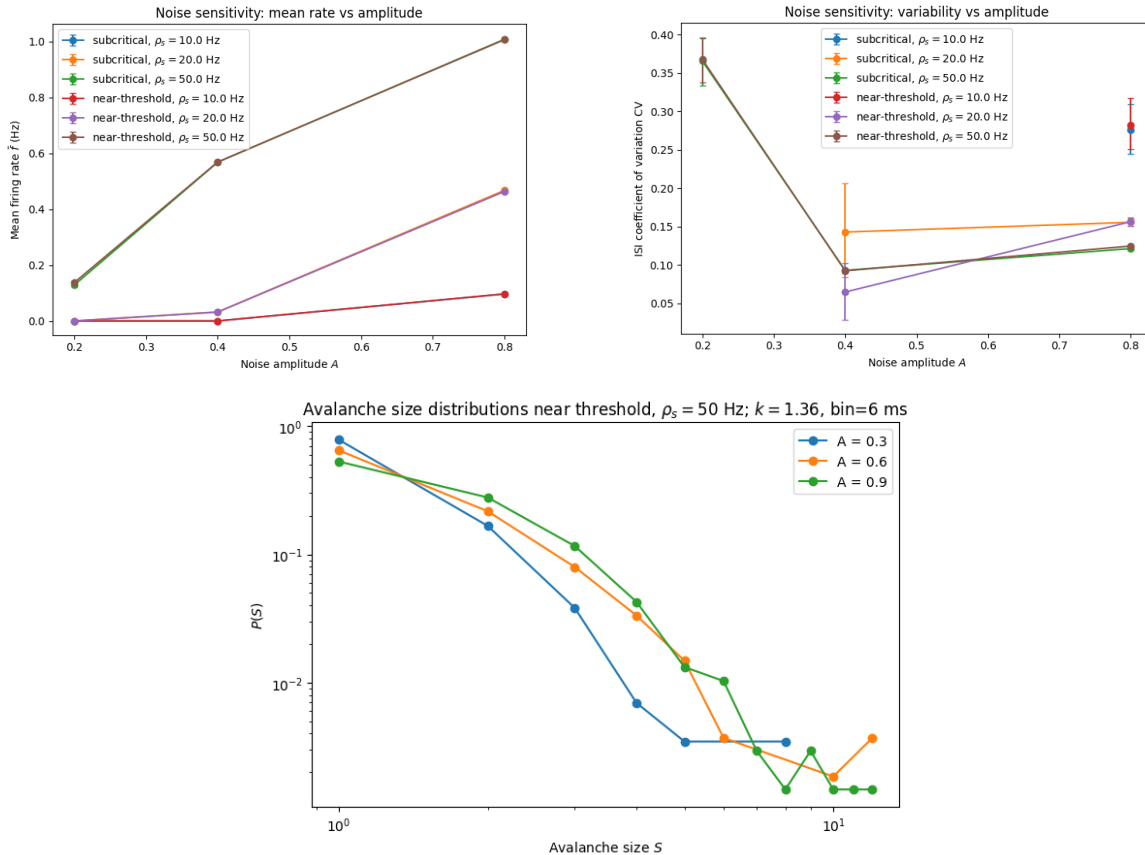


Figure 7: Noise sensitivity in *NOS*. (**Top left**) Mean firing rate \bar{f} versus amplitude A for subcritical ($k = 0.9$) and near-threshold ($k = 1.36$) regimes at $\rho_s \in \{10, 20, 50\}$ Hz. (**Top right**) Inter-spoke interval CV versus A under the same conditions. Error bars show bootstrap 95% confidence intervals. (**Bottom**) Avalanche size distributions $P(S)$ near threshold ($k = 1.36$, $\rho_s = 50$ Hz). Heavier tails appear as A increases, corresponding to extended cascades.

Empirical summary (Fig. 7). The rate curves show \bar{f} versus amplitude A for subcritical ($k = 0.9$) and near-threshold ($k = 1.36$) regimes; the steeper slope near $k = 1.4$ matches the growth of $|H_v(0)|$. The CV curves increase with A and peak near threshold, consistent with (90). The avalanche distributions broaden as A increases; heavier tails near $k \approx k^*$ reflect larger correlated excursions. These trends are stable across $\rho_s \in \{10, 20, 50\}$ Hz and the reported τ_s range, and they align with the linear response in (89). Further quantitative tail fits are provided in Appendix E.

5.7 Queueing baseline comparison

Experimental setup. We compare *NOS* against canonical queueing baselines in two modes. (i) **Open-loop:** we drive M/M/1 (analytic mean, $\rho < 1$), simulated M/M/1/ K , and a single *NOS* unit with *the same* exogenous arrivals and report observables in the *same unit* (packets). Light-load calibration aligns means: we choose an input gain and output scale so that the *NOS* mean at small ρ matches the M/M/1 mean $\mathbb{E}[L] = \rho/(1-\rho)$. Tail behaviour is then tested under a common MMPP burst sequence (same ON/OFF epochs), which is the operational regime of interest. (ii) **Closed-loop:** we attach three controllers to the same offered-load trace and compare the resulting queue trajectories

and marking signals $p(t)$: A) NOS-based (marking derived from NOS subthreshold state with soft reset), B) *queue* controller (marking from the raw queue), and C) *LP-queue* controller (marking from a lightly low-pass filtered queue). All three use the same logistic nonlinearity for $p(t)$ so differences are due to the state fed to the nonlinearity.




Figure 8: Open-loop comparisons in packet units. **(Left)** Mean occupancy versus arrival rate λ : M/M/1 (analytic, $\rho < 1$), simulated M/M/1/K, and NOS (calibrated at light load). NOS bends earlier near saturation due to bounded excitability and soft reset. **(Right)** Tail CCDF under a common MMPP burst sequence: M/M/1/K displays a slow tail, whereas NOS truncates cascades via continuous pull-back, yielding a much lighter tail.

How to read the open-loop panels (networking view). Figure 8 (left) shows mean occupancy versus arrival rate λ with a fixed service μ . The blue curve (M/M/1) is the textbook reference that diverges as $\rho = \lambda/\mu \uparrow 1$. The orange curve (M/M/1/K) follows M/M/1 until blocking becomes material. The green curve (NOS) rises with load at small ρ (by calibration) but *bends earlier* as ρ grows; bounded excitability and continuous pull-back reduce the subthreshold slope and drain excursions. Operationally, this mirrors how modern AQMs aim to keep average queues low near saturation rather than letting them diverge smoothly.

Figure 8 (right) plots the complementary CDFs of occupancy under the *same* MMPP burst sequence. M/M/1/K develops a long tail during ON periods (large probability of deep queues), whereas NOS exhibits a much sharper tail because soft reset and leak truncate cascades. In networking terms, NOS converts the same burst train into fewer deep-queue events—directly improving tail latency and reducing drop/mark storms. The aggressiveness of this truncation is tunable through the subthreshold leak χ , the reset rate r_{reset} , and the saturation knee κ ; operators can place the tail against SLO targets (e.g., $\mathbb{P}\{Q \geq q_0\} \leq \varepsilon$).

How to read the closed-loop panels (networking view). Figure 9 contrasts three marking strategies under the same offered-load traces. The top panel shows controller outputs $p(t)$. NOS (blue) is decisive and low-jitter; the queue controller (orange) is noisy (mark jitter follows queue noise); the LP-queue controller (green) is smoother but lags, so it reacts late to bursts. The middle panel shows a step in offered load: NOS snaps to target with minimal ringing; the queue controller oscillates; the LP-queue controller overshoots and then settles slowly. The bottom panel shows a burst train: NOS produces short, clipped spikes; the queue controller shows noisy peaks; the LP-queue controller permits taller spikes because filtering delays the reaction. Practically, NOS shrinks the time spent

at high queue levels and reduces burst amplification without a long averaging window, which improves tail latency and fairness while avoiding prolonged ECN saturation.

Summary for operators. The baselines anchor NOS against established theory. When the objective is to *match the classical infinite-buffer mean* at light load, M/M/1 gives the right reference and NOS can be calibrated to agree. When the objective is to *manage burst risk and stability* under finite resources and coupling, NOS’ bounded excitability and differentiable pull-back are strictly more faithful: they reduce deep-queue probability, yield crisper ECN/markings, and shorten congestion episodes without heavy filtering. The same knobs that appear in our stability proxies— χ , λ -like drain, κ , and reset parameters—map directly to deployed policy controls (AQM leak, scheduler epoch, ramp shaping, and post-alarm conservatism), making the model both predictive and actionable. Having established mechanistic fidelity and control behaviour, we next ask how NOS fares as a label-free forecaster against compact ML/SNN baselines.

5.8 Zero-shot forecasting baselines (no labels)

Setup (like-for-like, no training). All methods receive the same sliding window of arrival counts (last W bins per node) and must predict the next queue sample q_{t+1} *without supervised training*. We include: (**Phys.-Fluid**) the label-free fluid update $q_{t+1} = \text{clip}(q_t + a_t - \mu \Delta t)$; (**MovingAvg**) a smoothed fluid variant; (**TGNN-smooth**) a single-hop temporal-graph smoother applied to arrivals before the fluid step; (**LIF-leaky**) a leaky-integrator forecaster; and (**NOS (calibr.)**) a single NOS unit per node with *arrival-only* calibration (no queue labels): a global input gain and per-node output scales are fitted to match light-load analytic means from M/M/1. In the run summarized here we obtained `gain_I` = 0.60 and the first three per-node scales {11.646, 12.226, 12.277}; the corresponding analytic light-load means for those nodes were {0.655, 0.738, 0.745} (reported for reference only).

What the figure shows. Figure 10 overlays the *true* queue on node 0 (black) with the five zero-shot forecasts. Phys.-Fluid (blue) tracks the envelope of bursts because it integrates the *same* arrivals that drive the queueing simulator; MovingAvg (orange) lags and underestimates peaks; TGNN-smooth (green) damps spikes via spatial smoothing; LIF-leaky (red) behaves as a low-pass filter; NOS (purple) fires promptly at burst onsets and then resets, producing compact predictions around excursions.

Numerical summary. Table 5 reports point error (MAE) and event skill (AUROC / AUPRC for top-10% bursts), computed on the same held-out segment. Phys.-Fluid attains the smallest MAE (0.675)—expected, since the simulator itself is fluid-like at the bin scale. NOS has a larger MAE (0.928) because it is *bounded* and self-resets (it does not accumulate backlog), yet it delivers strong event skill (AUROC 0.894, AUPRC 0.536), on par with Phys.-Fluid and substantially above simple smoothers. Operationally, this means NOS is well-tuned to *detect and time* congestion onsets without drifting into long false positives between bursts. TGNN-smooth and LIF-leaky reduce noise but lose timing contrast (near-chance event skill).

Networking interpretation. In the absence of labels, a fluid predictor provides a tough MAE baseline because it integrates exactly the arrivals that create the queue, es-




Figure 9: Closed-loop behaviour under identical offered-load traces. **(Top)** Marking signals $p(t)$ from three controllers: NOS (blue), raw queue (orange), and low-pass filtered queue (green). **(Middle)** Step response: NOS returns to target fastest with little ringing; queue-only is jitter-dominated; LP-queue trades jitter for lag and overshoot. **(Bottom)** Burst train: NOS clips spikes and shortens congestion episodes; queue-only is noisy; LP-queue allows taller spikes due to filter delay.




Figure 10: Zero-shot next-step forecasts on a single node (no supervised training). All methods consume the same arrival windows; only *arrival-only* calibration is used. NOS produces compact, onset-aligned responses; the fluid forecaster follows backlog accumulation.

Table 5: Zero-shot next-step forecasting (no supervised training). All methods use the same arrival windows; NOS uses only arrival-based calibration.

Method	MAE	AUROC	AUPRC
Physics Fluid	0.6748	0.834	0.555
Moving Avg	0.8782	0.552	0.194
TGNN-smooth	0.9251	0.507	0.128
LIF-leaky	0.8891	0.500	0.126
NOS (calibr.)	0.9278	0.894	0.536

essentially reproducing backlog drift. However, operators often care more about *timely onset detection* than about matching every backlog micro-fluctuation. Here NOS’ bounded, event-driven dynamics produce high AUROC/AUPRC—accurate timing of burst onsets and compact responses—while avoiding prolonged marking during recoveries. This complements the open/closed-loop results: use fluid logic for gross throughput accounting, and deploy NOS to deliver low-latency, low-ringing congestion signals aligned with control actions.

6 Network-level stability and topology dependence

To simulate the dynamics of the proposed *NOS* model, we implemented an event-driven loop that integrates each node’s excitability, recovery variable, and incoming delayed spikes. Algorithm 1 summarises the procedure. The update consists of three stages per timestep: (i) delivery of delayed spikes, (ii) membrane and recovery updates with bounded excitability and leak, and (iii) threshold check with a soft exponential reset. This compact routine captures both local spiking dynamics and graph-level propagation through weighted and delayed edges. The full pseudocode, including optional stochastic arrivals and avalanche counting, is provided in Appendix D.

Algorithm 1 Simulation of NOS dynamics with delays and soft resets

- 1: **Initialise:** $v_i \leftarrow v_{\text{rest}} + \epsilon$, $u_i \leftarrow u_0$, delay buffers for all ($j \rightarrow i$)
- 2: **for** each time step t **do**
- 3: Deliver spikes scheduled at t from delay buffers
- 4: Compute input:

$$I_i \leftarrow \sum_j w_{ij} g(q_{ij}) S_j(t - \tau_{ij}) + \eta_i(t)$$

- 5: Update membrane:

$$v_i \leftarrow v_i + dt \cdot (f_{\text{sat}}(v_i) + \beta v_i + \gamma - u_i + I_i - \lambda v_i - \chi(v_i - v_{\text{rest}}))$$

- 6: Update recovery:

$$u_i \leftarrow u_i + dt \cdot (a(bv_i - u_i) - \mu u_i)$$

- 7: **if** $v_i \geq v_{\text{th}} + \sigma\xi$ **then**
 - 8: Emit spike and record time
 - 9: Schedule outgoing spikes in buffers for ($i \rightarrow j$) at $t + \tau_{ij}$
 - 10: Reset: $v_i \leftarrow c + (v_i - c)e^{-r_{\text{reset}}dt}$, $u_i \leftarrow u_i + d$
 - 11: **end if**
 - 12: **end for**
-

Topologies and experiments

We tested NOS on large networks of size $N = 250$ with different canonical topologies:

Chain topology. Nodes are arranged in a linear sequence, representing serial bottlenecks such as routers on a backbone link. Propagation is slowed by distance, requiring stronger coupling for instability.

Star topology. One hub connects to all leaves, representing a single aggregation point such as a central switch. This configuration is fragile, as overload of the hub destabilises the entire structure.

Scale-free topology. Networks generated by the Barabási–Albert model capture heterogeneous connectivity with hubs and peripheries. This structure shows intermediate behaviour: hubs can initiate cascades, but leaves constrain their spread.

Figure 11 summarises the network-level phenomena captured by *NOS*. Panel (a) shows that when link delays are included ($\tau_{ij} > 0$), the stability boundary shifts: oscillations emerge at lower coupling k , reflecting how communication latency destabilises otherwise stable traffic. Panel (b) compares absolute resilience across topologies: chains require stronger coupling to destabilise, stars are fragile because of their single hub, and scale-free networks lie in between. Panel (b'') reframes this relative to the Perron spectral prediction. Here chains deviate strongly, sitting well above the $1/\rho(W)$ baseline, while stars and scale-free graphs align more closely. This means that spectral tools capture hub-dominated topologies reliably but underestimate the resilience of homogeneous chains.

The emergence of coherent timing can also be tracked with a Kuramoto-type order parameter; Appendix I shows $\langle R \rangle$ versus $k = g \rho(W)$ (Fig. 16) and the finite-size sharpening of the transition around the Perron prediction k^* . Together these results demonstrate that *NOS* bridges dynamical-systems formalism and networking dynamics. It shows where simple spectral theory suffices (hub-dominated topologies) and where explicit simulation is required (homogeneous or delayed structures). For networking readers, the panels illustrate directly how latency, topology, and surge recovery shape resilience, providing an interpretable analogue to oscillation and congestion dynamics in real networks.

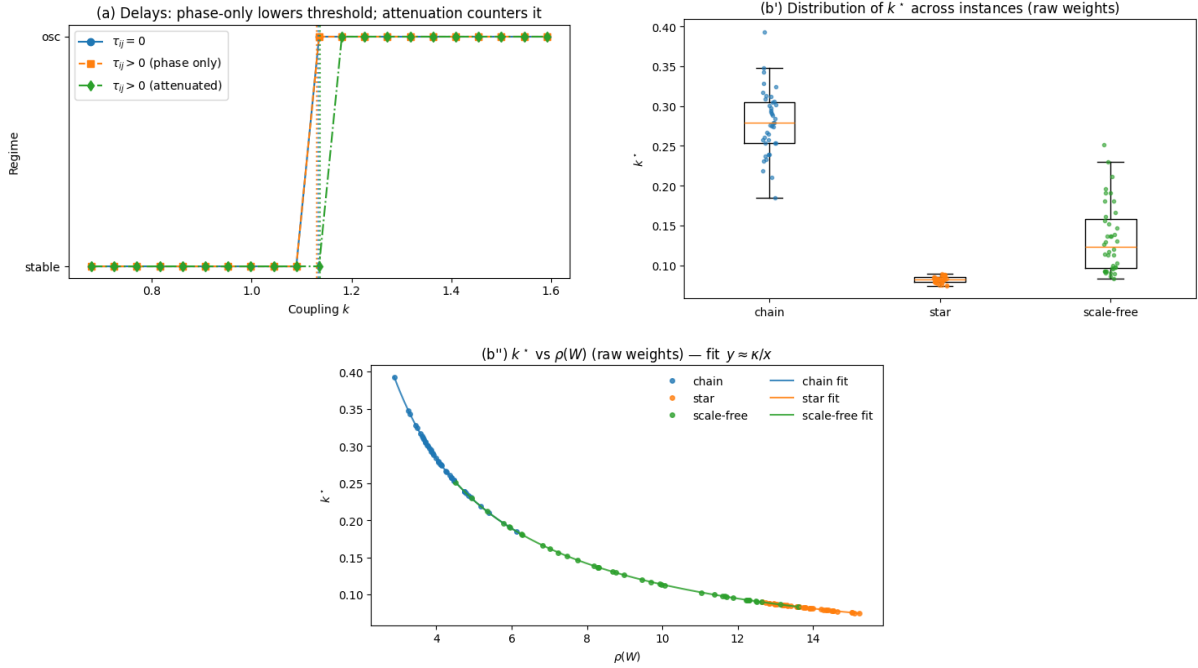


Figure 11: Network-level dynamics in *NOS*. (a) Delay effects: introducing link delays $\tau_{ij} > 0$ shifts the stability boundary, causing oscillations to emerge at lower coupling k . (b) Absolute topology effects: chain networks destabilise last (resilient), stars first (fragile), with scale-free in between. (b'') Relative to Perron-mode prediction: chain networks deviate strongly above the $1/\rho(W)$ line, while star and scale-free graphs track the prediction closely.

6.1 Predictive baseline: comparison with machine learning models

Many operational networks lack reliable labels, so short-horizon forecasting with residual-based event inference is the standard route. We follow that practice. We compare *NOS* with tGNN, GRU, RNN, and MLP under three canonical graphs: scale-free, star, and chain. For metrics we use networks of size $N=250$; the range plots show node 0 for visual clarity. All baselines are trained label-free on next-step forecasting with a contiguous train/validation/test split. Inputs are standardised per node using statistics fitted on the train split only. Residuals are signed and z -scored with train-only moments; event starts are the first samples whose residual z -scores cross a fixed per-node threshold selected on the validation split but computed using train calibration. Episodes have a minimum duration to suppress single-sample blips; matching uses a fixed window around each


ground-truth start. Start-latency is the sample difference between the truth start and the first model start within the window, converted to milliseconds using the sampling period Δt . Forecasting error (MAE, RMSE) is computed on the original scale. The tGNN uses temporal message passing constrained to the given adjacency; capacities and early-stopping are held comparable across baselines.

Figure 12 presents the test-range series with start markers for each method. A marker at the truth dot indicates a correct start, a marker to the right indicates a late start, and a marker to the left indicates an early start. Extra markers away from a truth dot are false positives, while missing markers in the vicinity of a truth dot are false negatives. Read across the three panels from scale-free to star to chain and the pattern is consistent. On the scale-free graph, bursts arrive in clusters and ramps are steep. *NOS* remains close to the truth throughout these clusters and avoids small oscillations. The tGNN calls tend to appear inside the ramps, which shows up as slight right shifts. RNN produces occasional early calls on moderate slopes, and GRU sits between RNN and tGNN. MLP is the least selective and places extra markers on low-amplitude wiggles. Moving to the star graph, the hub concentrates load while spokes are quieter. This helps tGNN, yet *NOS* remains more selective at the spokes and keeps the timing tight when the hub rises. On the chain graph, onsets are clearer and repeat through the range. *NOS* lands on the main rises with few stray calls in the troughs, while tGNN shows a mild lag, RNN is a little early on shallow slopes, and GRU again lies between them. The visual story across the three topologies is that *NOS* aligns closely without over-firing when the background is quiet.

Figure 13 gathers the aggregate scores. The bars report F1, precision, and recall under the same train-calibrated residual protocol. The curves report one-step error as MAE and RMSE and the median start-latency in milliseconds. Higher bars indicate better detection quality. Lower curves indicate better forecasting accuracy and earlier response. Across all three topologies *NOS* attains the highest F1 by holding strong precision while lifting recall. It also yields the lowest MAE and RMSE and the smallest latency. RNN keeps good precision but loses recall and responds later. GRU and tGNN are intermediate and more sensitive to the graph structure. MLP has higher errors and longer delays. The metric panels match what the range plots suggested: *NOS* is both earlier and more selective, and its forecasts give cleaner residuals for thresholding.


These differences have direct operational meaning. High precision with weak recall misses incidents that matter. High recall with weak precision floods operators with alarms. Lower forecast error stabilises residual scales and avoids threshold drift. Lower latency preserves time for pacing, admission, and rerouting. Taken together, the figures show that *NOS* gives an effective operating point across scale-free, star, and chain graphs. It preserves recall without inflating false positives and it warns early enough to act when load begins to rise.

All learned baselines, namely MLP, RNN, GRU, and tGNN, are trained self-supervised on next-step forecasting, and events are inferred from forecast error rather than labels. Train-only calibration prevents leakage and keeps thresholds comparable across models and topologies. The evaluation can be extended with extreme-value thresholds on residual tails, rolling-origin splits to test drift, precision–recall curves, episode-wise latency in milliseconds, and an out-of-distribution topology block, without changing the core protocol.



(a) Scale-free. Truth and model start markers on the test range. (b) Star. Truth and model start markers on the test range. (c) Chain. Truth and model start markers on the test range.

Figure 12: Test-range alignments under train-calibrated residual thresholds. A marker at the truth dot is correct, to the right is late, to the left is early. Extra markers indicate false positives, and missing markers near truth indicate false negatives. *NOS* remains close to truth across all three topologies while avoiding small oscillations.



(a) Scale-free. F1, precision, recall, MAE, RMSE, and median start-latency. (b) Star. F1, precision, recall, MAE, RMSE, and median start-latency. (c) Chain. F1, precision, recall, MAE, RMSE, and median start-latency.

Figure 13: Summary metrics under the label-free residual protocol. Higher bars are better for F1, precision, and recall. Lower curves are better for MAE, RMSE, and latency. *NOS* achieves the best F1, the lowest forecast error, and the earliest starts across all three topologies.

7 Conclusion

This work we introduced *NOS*, a network-optimised spiking model tailored to packet networks with finite buffers, bursty arrivals, multi-hop coupling, and stringent latency requirements. In *NOS*, a bounded-excitability state v (normalised queue), a recovery state u , and differentiable resets (event-based exponential or continuous pullback) provide an event-driven representation that is compatible with telemetry smoothing and per-link delays/gates. Analytically, existence and uniqueness of subthreshold equilibria were established, and local stability was characterised via the single-node Jacobian. Linearisation of the coupled network yielded a block Jacobian and an operational stability proxy $g \rho(W) f'_{\text{sat}}(v^*) < \Lambda$, which separates topology from node physics and predicts that the first loss of stability scales inversely with the spectral radius $\rho(W)$. The analysis distinguishes saddle-node and Hopf onsets and shows that earlier saturation (larger κ) enlarges headroom. Under shot-noise drive, small-signal transfer functions provide DC gain and variance expressions that account for the observed rise in rate variability and cascade size near local and network margins. Empirically, with like-for-like inputs, *NOS* matches the light-load $M/M/1$ mean after a simple gain/scale calibration while truncating deep-queue tails near saturation due to bounded excitability and pullback. In closed loop, *NOS*-state controllers achieve low-jitter actuation with fast settling and

clipped responses to bursts. In label-free forecasting, *NOS* provides high event skill and low start-latency across chain, star, and scale-free graphs, complementing fluid predictors that minimise mean error but react later to onsets. Network-level simulations confirm the Perron-mode prediction for hub-dominated topologies and highlight increased resilience relative to $1/\rho(W)$ in homogeneous chains; delay dispersion shifts stability thresholds as expected. In summary, *NOS* offers a compact, analytically tractable spiking abstraction of network congestion that unifies bounded excitability, soft resets, recovery dynamics, and graph-local coupling with delays. The analysis links node physics to topology through a Perron-mode threshold, explaining when coherent congestion onsets arise and how noise amplifies near margins. **Practically**, the model yields direct design rules: calibrate gains to light-load statistics; schedule the global coupling g against the spectral radius $\rho(W)$; raise the drain margin Λ via service λ and damping χ ; shape excitability via κ ; and reweight or sparsify W to control $\rho(W)$. **Operationally**, *NOS*-state feedback provides earlier, lower-jitter actuation than raw or low-pass queues, while retaining fluid logic for aggregate throughput accounting. These properties make *NOS* suitable for deployment in burst-dominated networks and amenable to neuromorphic or in-switch realisations, with stability margins that remain interpretable under moderate heterogeneity and delay dispersion.

The Perron-mode proxy and fixed-gain calibration are most reliable under moderate heterogeneity, moderate delay dispersion, and piecewise-stationary traffic; outside this regime, stability thresholds and gains should be confirmed by conservative spectral bounds or full network simulation, and recovery parameters should be identified with regularised (e.g., Bayesian) estimation to mitigate noise. **Future work** will develop adaptive gain scheduling under bandwidth and quantisation constraints, co-design topology and weights to enforce a fixed operational margin, and prototype neuromorphic or in-switch realisations that exploit event-driven fabrics for communication-efficient learning and control in large-scale networks.

References

- [1] L. F. Abbott, B. DePasquale, and R. M. Memmesheimer. Building functional networks of spiking model neurons. *Nature Neuroscience*, 19:350–355, 2016.
- [2] Mohammed Ahmed, Abdulkadir N. Mahmood, and Jiankun Hu. A survey of network anomaly detection techniques. *Journal of Network and Computer Applications*, 60:19–31, 2016.
- [3] Niveditha Baganal-Krishna, Ronny Hansel, Ivan Lecuyer, Romain Fontugne, Frédéric Giroire, and Dario Rossi. A federated learning approach to QoS forecasting in operational 5G networks. *Computer Networks*, 237:110010, 2024.
- [4] Dennis Bäfler, Tobias Kortus, and Gabriele Gühring. Unsupervised anomaly detection in multivariate time series with online evolving spiking neural networks. *Machine Learning*, 111(4):1377–1408, 2022.
- [5] Tianshi Chen, Zidong Du, Ninghui Sun, Jia Wang, Chengyong Wu, Yunji Chen, and Olivier Temam. A high-throughput neural network accelerator. *IEEE Micro*, 35(3):24–32, 2015.

- [6] S. S. Chowdhury, D. Sharma, A. Kosta, et al. Neuromorphic computing for robotic vision: algorithms to hardware advances. *Communications Engineering*, 4:152, 2025.
- [7] M.C. Crair and W. Bialek. Non-boltzmann dynamics in networks of spiking neurons. In *[Proceedings] 1991 IEEE International Joint Conference on Neural Networks*, pages 2508–2514 vol.3, 1991.
- [8] M. Davies et al. Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro*, 38(1):82–99, 2018.
- [9] Steve B. Furber, Francesco Galluppi, Steve Temple, and Luis A. Plana. The spinaker project. *Proceedings of the IEEE*, 102(5):652–665, 2014.
- [10] Steve B. Furber, Francesco Galluppi, Steve Temple, and Luis A. Plana. The spinaker project. *Proceedings of the IEEE*, 102(5):652–665, 2014.
- [11] E. Gelenbe. G-networks: a unifying model for neural and queueing networks. *Annals of Operations Research*, 48:433–461, 1994.
- [12] Erol Gelenbe. Random neural networks with negative and positive signals and product form solution. *Neural Computation*, 1(4):502–510, 1989.
- [13] Erol Gelenbe. Stability of the random neural network model. *Neural Computation*, 2(2):239–247, 1990.
- [14] Anna Giannakou, Dipankar Dwivedi, and Sean Peisert. A machine learning approach for packet loss prediction in science flows. *Future Generation Computer Systems*, 102:190–197, 2020.
- [15] Lei Guo, Dongzhao Liu, Youxi Wu, and Guizhi Xu. Comparison of spiking neural networks with different topologies based on anti-disturbance ability under external noise. *Neurocomputing*, 529:113–127, 2023.
- [16] M. Häusser. The Hodgkin-Huxley theory of the action potential. *Nature Neuroscience*, 3(Suppl 11):1165, 2000.
- [17] Adeel Iqbal, Ali Nauman, Riaz Hussain, and Muhammad Bilal. Cognitive d2d communication: A comprehensive survey, research challenges, and future directions. *Internet of Things*, 24:100961, 2023.
- [18] E. M. Izhikevich. Simple model of spiking neurons. *IEEE Transactions on Neural Networks*, 14(6):1569–1572, 2003.
- [19] Jieliang Jiang, Jiale Zhu, Muhammad Bilal, Yan Cui, Neeraj Kumar, Ruihan Dou, Feng Su, and Xiaolong Xu. Masked swin transformer unet for industrial anomaly detection. *IEEE Transactions on Industrial Informatics*, 19(2):2200–2209, 2023.
- [20] Jann Krausse, Daniel Scholz, Felix Kreutz, Pascal Gerhards, Klaus Knobloch, and Jürgen Becker. Extreme sparsity in hodgkin-huxley spiking neural networks. In *2023 International Conference on Neuromorphic Computing (ICNC)*, pages 85–94, 2023.

- [21] N. Langlois, P. Miche, and A. Benschair. Analogue circuits of a learning spiking neuron model. In *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks. IJCNN 2000. Neural Computing: New Challenges and Perspectives for the New Millennium*, volume 4, pages 485–489 vol.4, 2000.
- [22] Honghui Liang, Zhiguang Chen, Yangle Zeng, Guangnan Feng, and Yutong Lu. *DFMG: Delay-Flush Multi-Group Algorithm for Spiking Neural Network Simulation*, page 479–485. Association for Computing Machinery, New York, NY, USA, 2025.
- [23] Chit-Kwan Lin, Andreas Wild, Gautham N. Chinya, Yongqiang Cao, Mike Davies, Daniel M. Lavery, and Hong Wang. Programming spiking neural networks on intel’s loihi. *Computer*, 51(3):52–61, 2018.
- [24] Guoqiang Liu, Guanming Bao, Muhammad Bilal, Angel Jones, Zhipeng Jing, and Xiaolong Xu. Edge data caching with consumer-centric service prediction in resilient industry 5.0. *IEEE Transactions on Consumer Electronics*, 70(1):1482–1492, 2024.
- [25] Nguyen Cong Luong, Dinh Thai Hoang, Shimin Gong, Dusit Niyato, Ping Wang, Ying-Chang Liang, and Dong In Kim. Applications of deep reinforcement learning in communications and networking: A survey. *IEEE Communications Surveys & Tutorials*, 21(4):3133–3174, 2019.
- [26] Piotr S. Maciąg, Marzena Kryszkiewicz, Robert Bemberek, Jesús L. Lobo, and Javier Del Ser. Unsupervised anomaly detection in stream data with online evolving spiking neural networks. *Neural Networks*, 139:118–139, 2021.
- [27] Ahlem Menaceur, Hamza Drid, and Mohamed Rahouti. Fault tolerance and failure recovery techniques in software-defined networking: A comprehensive approach. *Journal of Network and Systems Management*, 31(83), 2023.
- [28] Emre O. Neftci, Hesham Mostafa, and Friedemann Zenke. Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks. *IEEE Signal Processing Magazine*, 36(6):51–63, 2019.
- [29] Thuy T. T. Nguyen and Grenville Armitage. A survey of techniques for internet traffic classification using machine learning. *IEEE Communications Surveys & Tutorials*, 10(4):56–76, 2008.
- [30] Jonathan Platkiewicz, Eran Stark, and Asohan Amarasingham. Spike-centered jitter can mistake temporal structure. *Neural Computation*, 29(3):783–803, 2017.
- [31] Ali Rasteh, Florian Delpéch, Carlos Aguilar-Melchor, Romain Zimmer, Saeed Bagheri Shouraki, and Timothée Masquelier. Encrypted internet traffic classification using a supervised spiking neural network. *Neurocomput.*, 503(C):272–282, September 2022.
- [32] Mariana Segovia-Ferreira, Jose Rubio-Hernan, Ana Rosa Cavalli, and Joaquin Garcia-Alfaro. A survey on cyber-resilience approaches for cyber-physical systems. *ACM Computing Surveys*, 56(8):1–37, 2024.

- [33] Laith A. Shamieh, Wei-Chun Wang, Shida Zhang, Rakshith Saligram, Amol D. Gaidhane, Yu Cao, Arijit Raychowdhury, Suman Datta, and Saibal Mukhopadhyay. Cryogenic operation of computing-in-memory based spiking neural network. In *Proceedings of the 29th ACM/IEEE International Symposium on Low Power Electronics and Design, ISLPED '24*, page 1–6, New York, NY, USA, 2024. Association for Computing Machinery.
- [34] Muhammad Basit Umair, Zeshan Iqbal, Muhammad Bilal, Jamel Nebhen, Tarik Adnan Almohamad, and Raja Majid Mehmood. An efficient internet traffic classification system using deep learning for iot. *Computers, Materials & Continua*, 71(1):407–422, 2022.
- [35] Xiaolong Xu, Chenyi Yang, Muhammad Bilal, Weimin Li, and Huihui Wang. Computation offloading for energy and delay trade-offs with traffic flow prediction in edge computing-enabled iot. *IEEE Transactions on Intelligent Transportation Systems*, 24(12):15613–15623, 2023.
- [36] Zhenyu Xu, Wenxin Li, Aurojit Panda, and Minlan Yu. Teal: Learning-accelerated optimization of WAN traffic engineering. In *Proceedings of the ACM SIGCOMM 2023 Conference*. ACM, 2023.
- [37] Xu Yang, Yunlin Lei, Mengxing Wang, Jian Cai, Miao Wang, Ziyi Huan, and Xialv Lin. Evaluation of the effect of the dynamic behavior and topology co-learning of neurons and synapses on the small-sample learning ability of spiking neural network. *Brain Sciences*, 12(2), 2022.
- [38] Ming Yao, Oliver Richter, Guangyu Zhao, Ning Qiao, Yannan Xing, Dingheng Wang, Tianxiang Hu, Wei Fang, Tugba Demirci, Michele De Marchi, Lei Deng, Tianyi Yan, Carsten Nielsen, Sadique Sheik, Chenxi Wu, Yonghong Tian, Bo Xu, and Guoqi Li. Spike-based dynamic computing with asynchronous sensing-computing neuromorphic chip. *Nature Communications*, 15:4464, 2024.
- [39] Friedemann Zenke and Wulfram Gerstner. Limits to high-speed simulations of spiking neural networks using general-purpose computers. *Frontiers in Neuroinformatics*, 8:76, 2014.
- [40] Jilun Zhang and Ying Liu. Brain-like path planning algorithm based on spiking neural network. In *Proceedings of the 2025 6th International Conference on Computer Information and Big Data Applications, CIBDA '25*, page 1379–1386, New York, NY, USA, 2025. Association for Computing Machinery.
- [41] Jingjing Zhao, Xuyang Jing, Zheng Yan, and Witold Pedrycz. Network traffic classification for data fusion: A survey. *Information Fusion*, 72:22–47, 2021.
- [42] Yuan Zhi, Jie Tian, Xiaofang Deng, Jingping Qiao, and Dianjie Lu. Deep reinforcement learning-based resource allocation for d2d communications in heterogeneous cellular networks. *Digital Communications and Networks*, 8(5):834–842, 2022.

A Delayed events and per-link gates in practice

This note elaborates on the roles of τ_{ij} and $g(q_{ij}(t))$ in (22)–(27) without repeating the main text.

Timing. The shift $S_j(t - \tau_{ij})$ represents propagation and processing delay on (j, i) . It aligns presynaptic events with the time packets reach node i . Delay changes phase in closed loops, which affects stability margins and the onset of oscillations.

Conditioning. The gate $g(q_{ij}(t))$ encodes a per-link condition such as residual capacity, scheduler priority, or queue occupancy. As q_{ij} grows, the effective weight $w_{ij}^{\text{eff}}(t) = w_{ij}g(q_{ij}(t))$ shrinks, so a congested link contributes less drive downstream. This mimics policing, AQM, or rate limiting and prevents the model from reinforcing overloaded paths.

Examples.

- *Datacentre incast.* Multiple senders j spike toward a top-of-rack switch i . Per-link queues q_{ij} rise, $g(q_{ij})$ falls, and w_{ij}^{eff} shrinks, so $I_i(t)$ reflects effective rather than nominal bandwidth.
- *WAN hop with high delay.* A large τ_{ij} shifts S_j by tens of milliseconds. The delayed drive prevents premature reactions at i and yields a stability test that depends on both $\rho(W)$ and the phase introduced by τ_{ij} .
- *Wireless fading.* Treat short-term SNR loss as a temporary rise in q_{ij} or as a change in the gate parameters. The same S_j then has smaller impact at i , consistent with link adaptation that lowers the offered load.

Operator workflow. Use passive RTT sampling or active probes to estimate τ_{ij} . Fit τ_q to the decay of q_{ij} when input drops or to the effective averaging window of the QoS policy. Set w_{ij} proportional to nominal capacity or administrative weight, then scale W to the target small-signal index. Validate g by replaying traces with known congestion episodes and checking that w_{ij}^{eff} attenuates the expected links while leaving others unchanged.

B Neuromorphic implementation details

Fixed-point state and quantisation. Let (v, u) be stored in $Q_{m,n}$ fixed point with full-buffer scaling V and time base T from (3.8). The forward map and clipping are

$$v_q = \frac{1}{2^n} \text{clip}\left(\text{round}\left(\frac{v}{V} 2^n\right), 0, 2^m - 1\right), \quad u_q = \frac{1}{2^n} \text{clip}\left(\text{round}\left(\frac{u}{V} 2^n\right), 0, 2^m - 1\right). \quad (91)$$

With this scaling, $\tilde{v} \in [0, 1]$ remains interpretable as normalised queue level; overflow is prevented by the saturation in f_{sat} .

Discrete-time soft reset. On digital substrates the exponential reset (33) is applied per tick Δt only when a spike occurs. Let \tilde{v}_{t+1} denote the state after the continuous update but before any reset. Then

$$v_{t+1} = (1 - s_t) \tilde{v}_{t+1} + s_t \left[c + (\tilde{v}_{t+1} - c) e^{-r_{\text{reset}} \Delta t} \right], \quad (92)$$

$$u_{t+1} = u_t + d s_t, \quad (93)$$

where $s_t \in \{0, 1\}$ indicates a threshold crossing at time t . The factor $e^{-r_{\text{reset}} \Delta t}$ is realised by a multiply-accumulate or by a bit-shift if $r_{\text{reset}} \Delta t$ is drawn from a small table.

Logistic pullback via LUT or PWL. The continuous pullback (34) uses $\sigma_k(x) = 1/(1 + e^{-kx})$. On hardware we use a lookup table or piecewise-linear approximation $\hat{\sigma}_k$ on a bounded interval $x \in [-x_{\max}, x_{\max}]$:

$$|\sigma_k(x) - \hat{\sigma}_k(x)| \leq \varepsilon_{\text{LUT}} \quad \text{for } |x| \leq x_{\max}. \quad (94)$$

The induced drift error on the pullback term is then bounded by $\rho |v - c| \varepsilon_{\text{LUT}}$ and can be kept below a configured fraction of the net drain in (44).

Graph coupling, delays, and AER routing. Sparse coupling $W = [w_{ij}]$ is realised by address–event routing tables; per-link delays are integer tick buffers

$$\tilde{\tau}_{ij} = \text{round}\left(\frac{\tau_{ij}}{T}\right) \in \mathbb{N}, \quad (95)$$

which preserves the causality of (22). Gates $g(q_{ij})$ are evaluated where the link queue state is maintained; their outputs modulate the synaptic multiplier before accumulation.

Fabric-rate and budget constraints. Let R_{\max} be the sustainable spike throughput per core and let \bar{r}_j be the observed rate of S_j . A conservative admission constraint that avoids router overload is

$$\sum_j \mathbf{1}\{w_{ij} \neq 0\} \bar{r}_j \leq R_{\max} \quad \text{for each core hosting neuron } i, \quad (96)$$

or, in matrix form for a deployment-wide check, $\|W\|_{0,1} \bar{r}_{\max} \leq R_{\max}$, where $\|W\|_{0,1}$ counts nonzeros per row and $\bar{r}_{\max} = \max_j \bar{r}_j$.

Operational margin under quantisation. Let Δ_{net} be the margin in (46). Finite precision induces errors in $f'_{\text{sat}}(v^*)$, in W , and in gains. If $\delta_{f'}$, δ_W , and δ_g bound these relative errors, then a sufficient robustness condition is

$$\Delta_{\text{net}} > g \rho(W) f'_{\text{sat}}(v^*) (\delta_{f'} + \delta_W + \delta_g) + \varepsilon_{\text{LUT}} \rho \mathbb{E}\{|v - c|\}, \quad (97)$$

which ensures that the stability proxy (45) still holds after rounding and table approximations.

Networking alignment and reporting. Choose (V, T) so that V matches the full-buffer level used by the QoS policy and T matches the scheduler epoch or drain time. Report energy per spike and energy per detection alongside CPU baselines, and log $\tilde{\tau}_{ij}$ histograms to confirm that delay quantisation preserves path ordering. These checks keep the deployed NOS consistent with the queueing semantics and timing used in the main text.

C Linear decay times and DC sensitivity

Eigenvalues and decay times. Let $\lambda_{1,2}$ be the eigenvalues of the local Jacobian J_i defined in §5.2. When the Routh–Hurwitz conditions in (65) hold and $T_i^2 > 4\Delta_i$, the node is an overdamped node with two real negative modes. When $T_i^2 < 4\Delta_i$, it is underdamped

with complex conjugate modes whose common decay rate is $-\frac{1}{2}T_i$. The dominant linear return time is

$$\tau_{\text{lin}} \approx \frac{1}{\min\{-\Re\lambda_1, -\Re\lambda_2\}}, \quad (98)$$

which provides an engineering handle for truncated BPTT windows and for selecting the reset constant ρ so numerical timescales match device drain and scheduler epochs. In practice, τ_{lin} estimated from small perturbations of I_i around an operating point aligns well with the observed relaxation of the queue proxy v_i .

Input–output small-signal gain. Let \widehat{I} be a small constant perturbation of I_i and \widehat{v} the induced steady-state change in v_i . Using the equilibrium implicit-function relation from (61), the DC gain is

$$\frac{\widehat{v}}{\widehat{I}} = -\frac{1}{f'_{\text{sat}}(v^*) + L}. \quad (99)$$

Larger service λ and stronger damping χ decrease this gain, making the operating point less sensitive to slow offered-load drift. This matches queueing intuition and yields a direct calibration target: estimate \widehat{v}/\widehat{I} from step tests or natural experiments in telemetry, then verify that the inferred L is consistent with the Jacobian conditions in (65).

D *NOS* simulation pseudocode

For reproducibility we include the pseudocode used in our experiments (bin $\Delta t = 5$ ms).

```
# NOS simulation with delays and shot-noise arrivals

# Initialisation
for each node i:
    v[i] = v_rest + small_noise()
    u[i] = u0
for each directed link (j -> i):
    if tau_ij > 0: init delay_buffer[j->i]

# Main loop
for t in time_steps:
    # 1) deliver delayed spikes scheduled for this time
    for each link (j -> i):
        S_delivered = pop(delay_buffer[j->i], t) # 0/1 spike
        I_syn[i] += g * W[i,j] * S_delivered

    # 2) exogenous input (e.g., shot noise, optionally smoothed)
    for each node i:
        I_ex[i] = eta_i(t)
        # may include IO + gain * arrivals_i(t), low-pass with tau_s if used
        I[i] = I_syn[i] + I_ex[i]
```

```

# 3) state updates (forward Euler, dt = 5 ms)
for each node i:
    v_new = v[i] + dt * ( f_sat(v[i]) + beta*v[i] + gamma - u[i]
                        + I[i] - lambda*v[i] - chi*(v[i] - v_rest) )
    u_new = u[i] + dt * ( a*(b*v[i] - u[i]) - mu*u[i] )
    v[i] = clamp(v_new, v_rest, v_max)
    u[i] = clamp(u_new, u_min, u_max)

# 4) thresholding and reset (deterministic by default)
for each node i:
    v_th_eff = v_th # or v_th + sigma * randn() if optional jitter is enabled
    if v[i] >= v_th_eff:
        emit_spike(i, t); record_time(i, t)
        for each outgoing link (i -> k):
            schedule(delay_buffer[i->k], t + tau_ik) # integer bin delay
        # soft exponential pullback and recovery kick
        v[i] = c + (v[i] - c) * exp(-r_reset * dt)
        u[i] = u[i] + d

# 5) zero synaptic accumulator for next step
for each node i: I_syn[i] = 0

# record observables: mean rate, ISI CV, avalanche sizes, synchrony

```

E Statistical robustness under noise

For each condition we estimated the mean firing rate \bar{f} and the inter-spike-interval coefficient of variation (CV) using non-parametric bootstrap resampling (hundreds of replicates). Bootstrap 95% intervals show that both \bar{f} and CV increase with shot-noise amplitude A , particularly near threshold. In *NOS*, stronger stochastic drive raises overall excitability and increases irregularity, mirroring more variable queue dynamics.

Avalanche size distributions $P(S)$ were then analysed. Figure 14 shows empirical complementary cumulative distributions with fitted power-law and log-normal overlays above a data-driven tail threshold x_{\min} . Tail parameters were estimated by maximum likelihood with truncation at x_{\min} ; goodness was assessed by the Kolmogorov–Smirnov distance and model choice by AIC. Across conditions, AIC favours the power-law, indicating scale-free-like behaviour near threshold. Fitted values are reported in Table S6.

The exponents display a clear trend: at low amplitude ($A = 0.3$) the power-law exponent is $\hat{\alpha} \approx 8.7$ (steep decay, small avalanches). At $A = 0.6$ the exponent decreases to $\hat{\alpha} \approx 6.4$ (broader tail). At $A = 0.9$ the exponent drops to $\hat{\alpha} \approx 3.3$ (heavy tail and a higher chance of large cascades). Thus, near the bifurcation, stronger stochastic drive allows small fluctuations to trigger system-wide congestion events.

Avalanche definition and fitting. Avalanche sizes S were computed from population spike counts in 5 ms bins (the experimental Δt): an avalanche is a contiguous run of nonzero bins, with size equal to the total spikes in the run. For each condition, x_{\min} minimised the KS distance between empirical and model CDFs (tail only). Power-law

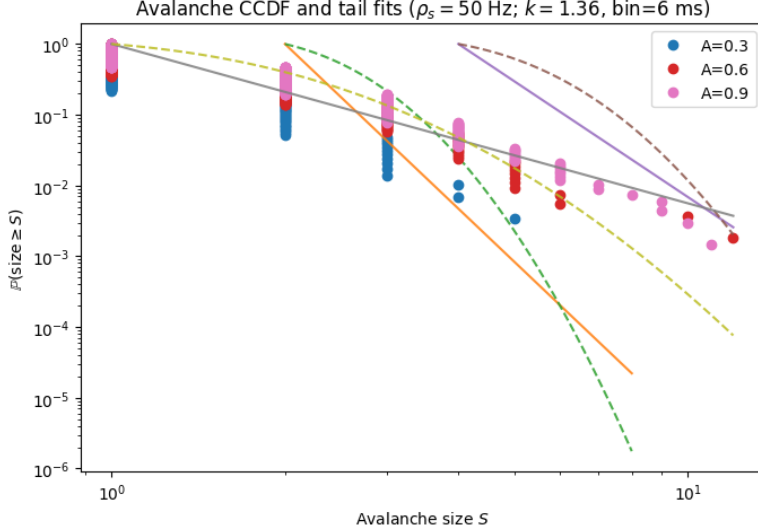


Figure 14: Avalanche complementary CDFs near threshold ($k = 1.36$, $\rho_s = 50$ Hz). Empirical distributions (points) with fitted power-law (solid) and log-normal (dashed) overlays above the estimated x_{\min} for each amplitude A . Model selection used AIC; numerical results appear in Table S6.

Table 6: Avalanche tail fits near threshold ($k = 1.36$, $\rho_s = 50$ Hz). For each amplitude A , we report the tail threshold x_{\min} , tail size n , maximum-likelihood parameters, KS distance (for the truncated tail), log-likelihood (LL), AIC, and the preferred model.

A	x_{\min}^{PL}	n_{PL}	$\hat{\alpha}$	KS_{PL}	LL_{PL}	AIC_{PL}	x_{\min}^{LN}	n_{LN}	$\hat{\mu}$	$\hat{\sigma}$	KS_{LN}	LL_{LN}	AIC_{LN}	Pref.
0.3	2.0	63	8.73	0.76	14.03	-26.07	2.0	63	0.82	0.27	0.76	-34.20	72.39	PL
0.6	4.0	31	6.43	0.58	-27.24	56.48	4.0	31	1.57	0.31	0.58	-46.15	96.31	PL
0.9	1.0	681	3.25	0.53	-431.13	864.26	1.0	681	0.44	0.53	0.53	-684.68	1373.35	PL

tails used the continuous-tail MLE $\hat{\alpha} = 1 + n / \sum_{i=1}^n \log(x_i/x_{\min})$ with model $F(x) = 1 - (x/x_{\min})^{1-\hat{\alpha}}$ for $x \geq x_{\min}$. Log-normal tails were fit by MLE on $\log x$ with truncation at x_{\min} . Model selection used AIC with parameter counts 1 (power-law) and 2 (log-normal).

F Admissible parameter region and calibration conventions

Table 7 records the admissible region explored in our experiments. The entries are expressed on the experimental scale used throughout the paper, with queue fullness normalised to $v \in [0, 1]$ and a sampling bin of 5 ms. Rates given “per bin” convert to per second by multiplying by 200. We separate structural constraints from initialisation choices so that the table serves reproducibility without implying narrow tuning.

The bounds have three roles. First, they enforce the normalisation and sampling assumptions used by the model and code. Second, they keep the sufficient stability condition from Lemma 2 satisfied with margin for the operating regimes we test. Third, they define a common hyperparameter space for stress tests across the three graph families. Within this region we draw values uniformly unless stated and refine only on the

Table 7: Admissible parameter region used in the experiments. Entries enforce $v \in [0, 1]$, a bin width of 5 ms, and keep the stability condition from Lemma 2 satisfied with margin in operating regimes. Values were sampled within these regions; defaults are given elsewhere for reproducibility.

Parameter	Interpretation	Admissible range (units)
α	excitability scale in f_{sat}	[0.4, 1.0] (dimensionless)
κ	saturation knee of f_{sat}	[0.5, 2.0] (dimensionless)
β	linear excitability gain	[-0.10, 0.40] (dimensionless)
γ	constant drive (baseline load)	[0.00, 0.15] (in v units)
λ	service/leak on v	[0.10, 0.30] (per bin ¹)
χ	subthreshold damping about v_{rest}	[0.00, 0.08] (per bin)
a	recovery rate	[0.6, 1.8] (per bin)
b	recovery sensitivity to congestion	[0.6, 1.6] (dimensionless)
μ	passive recovery decay	[0.00, 0.35] (per bin)
v_{th}	firing threshold	[0.50, 0.68] (in v units)
k_{reset}	sharpness of reset gate	[10, 20] (dimensionless)
ρ	pullback/reset speed	[3, 8] (per bin)
c	post-event baseline level	[0.0, 0.2] (in v units)
d	recovery jump on event	[0.1, 0.4] (in u units)
I_0	drive offset (arrivals $\rightarrow I$)	[0.08, 0.16] (in v units)
<i>gain</i>	drive scale (arrivals $\rightarrow I$)	[0.8, 1.2] (dimensionless)
τ_s	optional drive smoothing	[0, 15] ms
g	coupling on W (network)	choose so $k_{\text{net}} = g\rho(W) \in [0, 1.8]$
Shot noise	arrival model for synthetic tests	rate $\rho_i \in [10, 50]$ Hz, amplitude $A \in [0.3, 0.9]$
Delays	link propagation	$\tau_{ij} \in [0, 25]$ ms
W normalisation	spectral scaling	$\rho(W) = 1$ before applying g

validation split.

Coupling and reset are disambiguated. The network coupling index is $k_{\text{net}} = g\rho(W)$, where $\rho(W)$ is the spectral radius after normalisation to $\rho(W) = 1$ and g is the scalar applied during experiments. The symbol k_{reset} denotes the sharpness of the soft reset gate and is independent of k_{net} . Delays τ_{ij} are given in milliseconds and reflect link propagation on the chosen binning. Shot-noise parameters (ρ_i, A) match the noise sensitivity experiments and allow drive strength to approach the near-threshold regime without violating the stability bound.

The ranges for $\alpha, \kappa, \beta, \gamma$ control the shape and slope of the bounded excitability f_{sat} . The leak terms λ and χ set the small-signal decay about v_{rest} and are bounded to keep subthreshold dynamics stable at the 5 ms resolution. Recovery parameters (a, b, μ) govern post-burst relaxation and are given as per-bin rates. Threshold v_{th} , reset depth c , recovery jump d , and pullback speed ρ determine spiking and refractory behaviour. The mapping from arrivals to effective input uses an offset I_0 and a dimensionless gain, and may be optionally low-pass filtered with τ_s .

These bounds are admissible rather than prescriptive. Defaults used in the main text fall within Table 7 and are reported in the parameter-initialisation table. Sensitivity studies widening each bound by a factor of two preserve the relative ordering of methods on F1 and latency, with the largest trade-offs arising from λ and v_{th} as expected from precision–recall balance. All residual scalars and thresholds are fitted on train only, and validation is used to select per-node thresholds within the same admissible region.

¹Per-bin rates correspond to a 5 ms bin width. Multiply by 200 to convert to s^{-1} .

G Global stability (conservative small-signal margin)

Recall that the spectral radius satisfies $\rho(W) \leq \|W\|_\infty$, and a sufficient topology-aware bound is given in (79). For capacity planning we use the conservative small-signal *operational margin*

$$\Delta_{\text{op}} := \frac{L^2}{4\alpha} - \left(\gamma + \max_i I_i^*\right), \quad L := \beta - \lambda - \chi - \frac{ab}{a + \mu}. \quad (100)$$

Maintaining $\Delta_{\text{op}} > 0$ under nominal load provides headroom against stochastic fluctuations. The margin is linear in $I_{\text{max}} := \max_i I_i^*$ with slope -1 ; increases in subthreshold damping χ (or service λ) raise the intercept by shrinking L in magnitude.

Illustrative computation. Using the legacy “starter” values (kept here for continuity) $\alpha = 0.02$, $\beta = 0.5$, $\gamma = 0.05$, $\lambda = 0.2$, $a = 0.05$, $b = 0.5$, $\mu = 0.01$, $\chi = 0.05$, and $I_{\text{max}} = 0.10$,

$$\frac{ab}{a + \mu} = \frac{0.025}{0.06} \approx 0.4167, \quad L = 0.5 - 0.2 - 0.05 - 0.4167 \approx -0.1667,$$

$$\frac{L^2}{4\alpha} \approx \frac{0.02778}{0.08} = 0.34722, \quad \Delta_{\text{op}} \approx 0.34722 - (0.05 + 0.10) = 0.19722 > 0.$$

At $I_{\text{max}} = 0.30$ the margin is $\Delta_{\text{op}} \approx 0$, and raising χ from 0.05 to 0.30 changes L from -0.1667 to -0.4167 , yielding $\Delta_{\text{op}} \approx 1.82$ —a substantial safety increase. *Note:* these numbers are purely illustrative for the quadratic surrogate; in the main experiments we use the admissible ranges in Table 7. The linear dependence on I_{max} and the monotone effect of χ hold in either case.

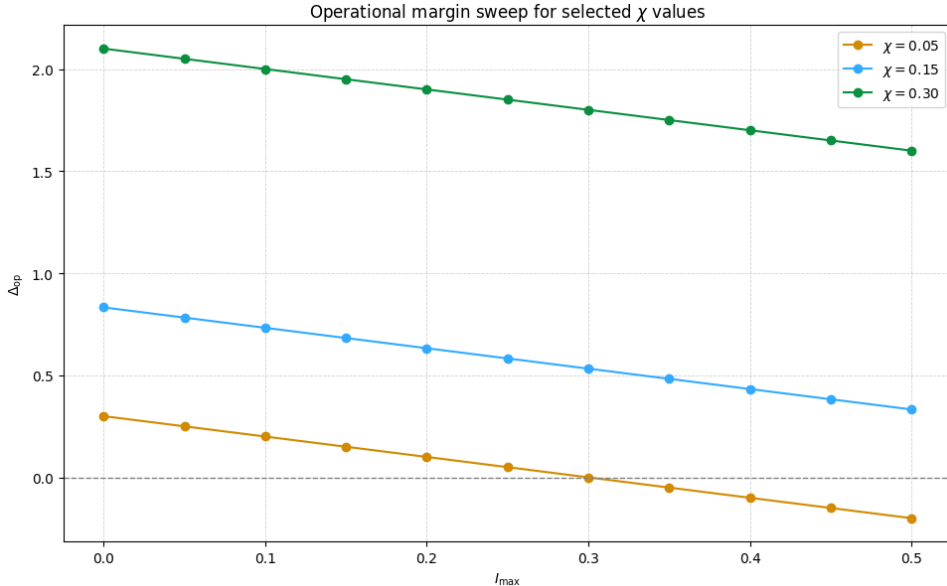


Figure 15: One-dimensional sweeps of the operational margin Δ_{op} against I_{max} for selected χ . Each trace is affine with slope -1 , illustrating the linear decay of margin with load. Intercepts grow with χ , consistent with the global stability boundary in Fig. 4.

Supporting sweeps. Figure 15 shows Δ_{op} against I_{max} for several χ . The slope is -1 by construction; increasing χ shifts the intercept upward (analogous to higher service slack). For the legacy starter set above the margin crosses zero near $I_{\text{max}} \approx 0.30$, consistent with

$$\Delta_{\text{op}} = \frac{L^2}{4\alpha} - (\gamma + I_{\text{max}}), \quad L \text{ fixed.}$$

Increasing χ increases $|L|$ and hence the scalar margin. Full sweeps and Monte Carlo robustness are provided for reference.

H Conservative stability via Gershgorin and norm bounds

Let $G = gW$ with $W_{ii} = 0$ and $\rho(W)$ the spectral radius after normalisation (we use $\rho(W) = 1$ in experiments). Linearising (9)–(10) at (v^*, u^*) gives a $2N \times 2N$ block Jacobian \mathcal{J} . For the top block rows $1:N$,

$$\bar{d}_i := f'_{\text{sat}}(v_i^*) + \beta - \lambda - \chi,$$

is the diagonal entry, and the off-diagonals are gW_{ij} for $j \neq i$, plus the coupling to u_i with derivative -1 . Thus the Gershgorin centre and radius are

$$c_i = \bar{d}_i, \quad R_i = 1 + |g| \sum_{j \neq i} |W_{ij}|.$$

A sufficient condition for those discs to lie strictly in the left half-plane is

$$\Re(c_i) + R_i < 0 \quad \text{for all } i,$$

which yields

$$|g| < \min_i \frac{-\Re(\bar{d}_i) - 1}{\sum_{j \neq i} |W_{ij}|}.$$

For the bottom block rows $N+1:2N$, the centre is $-(a + \mu)$ and the radius is $|ab|$, so a sufficient condition is

$$-(a + \mu) + |ab| < 0.$$

Using $\sum_{j \neq i} |W_{ij}| \leq \|W\|_{\infty}$ gives the norm variant

$$|g| < \frac{-\min_i \Re(\bar{d}_i) - 1}{\|W\|_{\infty}}, \quad \text{and} \quad |ab| < a + \mu.$$

These conditions are sufficient (not necessary) and are conservative relative to observed thresholds.

Table 8 complements the figures: the measured thresholds obey the Perron-mode scaling ($\rho(W)g_{\star} \approx k^*$), while Gershgorin discs provide only conservative certificates. From a networking standpoint, the stability boundary is captured by interpretable quantities: the effective small-signal slope (through \bar{d}_i), service and recovery (λ, a, μ, b) setting k^* , and the graph's spectral radius encoding topology. This bridges low-level queue dynamics with high-level network design, showing how structural and service parameters jointly control the emergence of congestion instabilities.

Table 8: Instability thresholds versus network spectral radius for homogeneous coupling $G = gW$. The measured g_* is the first zero crossing of the leading eigenvalue real part. The product $\rho(W)g_*$ agrees with the Perron-mode prediction k^* . Gershgorin bounds are conservative; negative values mean the bound cannot certify stability for any $g \geq 0$. The heuristic $g_{\text{heur}} = k^*/\|W\|_\infty$ matches the Perron scaling but lacks a guarantee. The last column checks $|ab| < (a + \mu)$.

$\rho(W)$	g_* (meas.)	$\rho(W)g_*$	g_{bound} (Gersh.)	g_{heur}	k^* (scalar)	$(\rho g_*)/k^*$	$ ab < (a + \mu)$
0.5	1.881799	0.940900	-0.526097	0.554642	0.940878	1.000023	True
1.0	0.940924	0.940924	-0.263049	0.277321	0.940878	1.000049	True
1.5	0.627452	0.941177	-0.175366	0.184881	0.940878	1.000319	True
2.0	0.470484	0.940969	-0.131524	0.138660	0.940878	1.000097	True
3.0	0.314042	0.942125	-0.087683	0.092440	0.940878	1.001326	True

I Finite-size sharpening and synchrony order parameter

Finite networks smooth the bifurcation transitions seen in the scalar reduction. We quantify collective timing with a Kuramoto-type order parameter

$$R(t) = \frac{1}{N} \left| \sum_{j=1}^N e^{i\phi_j(t)} \right|,$$

where spike phases $\phi_j(t)$ are obtained by linear phase interpolation between successive spikes of node j (phase resets to 0 at a spike, advances to 2π at the next). Figure 16 plots the time average $\langle R \rangle$ against the effective coupling $k = g\rho(W)$ for several network sizes N . The curves show the expected S-shaped rise: for small k the system is desynchronised and $\langle R \rangle$ is low; once k passes a critical value the order parameter increases rapidly and then saturates.

The dashed line marks the Perron-mode prediction k^* from the linear analysis of *NOS*. Across N , measured onsets cluster near k^* . Finite-size effects smooth the transition for small N and sharpen it as N grows, while the turning point remains anchored by k^* .

Networking interpretation. Synchrony corresponds to queueing cycles that align across nodes: burst build-up and drain events become temporally coordinated. Below k^* these cycles are weakly correlated and buffers clear largely independently. Near and above k^* , coupling is strong enough for delayed spikes to propagate and reinforce, producing network-wide burst trains. Because $k = g\rho(W)$, stability can be preserved either by reducing the local coupling gain g (e.g. weight attenuation or stronger subthreshold leak) or by reducing the spectral radius $\rho(W)$ (e.g. reweighting or sparsifying high-gain pathways). Both actions move the operating point away from the synchrony threshold.

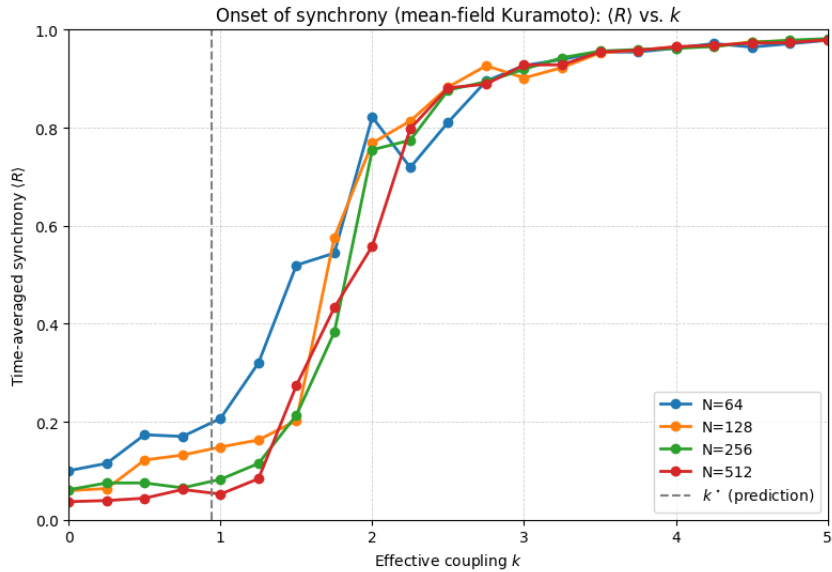


Figure 16: Onset of synchrony in *NOS*. Time-averaged order parameter $\langle R \rangle$ versus effective coupling $k = g \rho(W)$ for several network sizes N . The dashed line marks the Perron-mode prediction k^* . The transition sharpens with N , while the onset remains close to k^* .

Table 9: Representative values of $\langle R \rangle$ at selected couplings k and sizes N (extracted from Fig. 16). The increase of $\langle R \rangle$ with k and the steeper change for larger N reflect finite-size sharpening near k^* .

N	k	$\langle R \rangle$
64	0.0	0.100
64	1.0	0.207
64	2.0	0.822
128	0.0	0.059
128	1.0	0.148
128	2.0	0.769
256	0.0	0.061
256	1.0	0.082
256	2.0	0.755