RESEARCH-ARTICLE

# EdgeSynapse: Towards Leaky-Spike Transmission for Sustainable Edge Sensing

ATAKAN ARAL, University of Vienna, Vienna, Vienna, Austria

Open Access Support provided by:

University of Vienna



Published: 03 December 2025



SEC '25: Tenth ACM/IEEE Symposium on Edge Computing
*December 3 - 6, 2025*
*VA, Arlington, USA*

Conference Sponsors:
SIGMOBILE

SEC '25: Proceedings of the Tenth ACM/IEEE Symposium on Edge Computing (December 2025)
https://doi.org/10.1145/3769102.3774704
ISBN: 9798400722387

# EdgeSynapse: Towards Leaky-Spike Transmission for Sustainable Edge Sensing

Atakan Aral
atakan.aral@univie.ac.at
University of Vienna
Vienna, Austria

## Abstract

EdgeSynapse is a mechanism for energy-efficient sensing in remote, energy-constrained settings. It leverages neuromorphic principles by having sensor nodes and cluster heads generate discrete *spikes* when local integrator states cross thresholds. Specifically, each sensor uses a leaky integrate-and-fire (LIF) model to accumulate observations. When the membrane potential exceeds a threshold, the node emits an *excitatory spike* via LoRaWAN uplink and resets. Cluster heads similarly integrate incoming spikes and forward a higher-level spike when their threshold is reached, optionally using LoRaWAN downlink slots for *inhibitory spikes* to moderate traffic. This hierarchical spiking approach mimics biological signaling and reduces unnecessary transmissions. We provide a mathematical model of the LIF integration at sensors and cluster heads, along with the transmission algorithm tailored to LoRaWAN Class A communication. A preliminary simulation using real sensor data reveals favorable energy–accuracy trade-offs and a significant reduction in transmissions relative to send-on-Delta, alongside improved stability under synchronized bursts.

## CCS Concepts

• **Computer systems organization** → **Sensor networks**; **Neural networks**; • **Networks** → *Network performance analysis*; • **Hardware** → *Power and energy*.

## Keywords

Edge AI, Neuromorphic Sensing, Spiking Neural Networks, Low-Power IoT, LoRaWAN

**ACM Reference Format:**
Atakan Aral. 2025. EdgeSynapse: Towards Leaky-Spike Transmission for Sustainable Edge Sensing. In *The Tenth ACM/IEEE Symposium on Edge Computing (SEC '25), December 3–6, 2025, Arlington, VA, USA*. ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3769102.3774704

## 1 Introduction

Sustainable monitoring of remote environments (forests, rivers, agricultural fields, infrastructure, etc.) poses significant challenges. Deploying a large number of sensors over wide areas is costly and energy-constrained, especially where power grid and reliable communication are lacking [1]. High installation and maintenance costs, combined with obstacles (trees, terrain) and long distances, often yield poor link quality and many retransmissions, which drain node energy [17]. In such rural settings, resilience and low-power operation are critical. However, most existing systems rely on periodic sampling or conventional communication, which may waste energy when sensed values change slowly or unpredictably.

In this work, we draw inspiration from spiking neural networks such that sensor nodes only transmit when they detect significant changes, rather than on a fixed schedule. We use a *leaky integrate-and-fire* (LIF) model at each node. The node integrates incoming measurements into a membrane potential that decays (*leaks*) over time. When the potential exceeds a threshold, the node *fires* a spike and resets. Excitatory spikes raise the potential of downstream neurons, while inhibitory spikes temporarily raise the effective firing threshold. Neuromorphic processors operate this way to enable event-driven, parallel computation, keeping hardware inactive when no spikes occur and yielding large energy savings [5, 11]. Here, we repurpose these ideas for the network layer, where sensors send spikes, and cluster heads integrate them.

In summary, the contributions of this work-in-progress paper include: (i) EdgeSynapse system model that mimics biological spiking and LIF neurons at sensor nodes and cluster heads (Section III); (ii) a protocol design for implementing this model over LoRaWAN Class A (Section IV); (iii) initial evaluation on real sensor data showing significant reductions in transmissions (Section V); and (iv) a discussion of how this approach aligns with sustainability goals and directions for future work (Section VI).

## 2 Related Work

Event-triggered sensing and transmission have a long history in control and sampling theory. Send-on-Delta (SoD), level-crossing, and prediction-error sampling reduce traffic by transmitting only when the signal deviates sufficiently from its last report or predictor [13]. Similar ideas appear in LoRaWAN systems that send only residuals or prediction errors to cut airtime [14]. This provides a strong baseline for EdgeSynapse's SoD comparisons. Wireless spiking networks have been explored in other contexts for low-power edge AI [18]. EdgeSynapse is novel in applying hierarchical spiking to a LoRaWAN-based IoT monitoring network. Unlike prior works that treat neuromorphic transmission for high-speed systems, our focus is on low-data-rate LoRa links with strict duty-cycle limits.

Neuromorphic communication studies advocate spike-based, asynchronous signaling to save energy and exploit sparsity end-to-end (e.g., Address-Event Representation, AER). Prior work explores neuromorphic wireless cognition and distributed spiking systems [4], and hierarchical spike routing in processor fabrics (e.g., HiAER) [16]. However, these efforts typically remain within neuromorphic hardware stacks rather than mapping spike gating

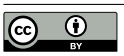


*SEC '25, Arlington, VA, USA*

ACM ISBN 979-8-4007-2238-7/2025/12
https://doi.org/10.1145/3769102.3774704

onto *hierarchical edge* over LPWANs. In parallel, the edge/IoT literature outlines the promise of event-driven and neuromorphic methods for environmental monitoring and the broader edge AI roadmap [6, 9, 12].

Finally, LoRaWAN-specific mechanisms to curb airtime, including downlink scarcity in Class A and the potential of wake-up radios for rare control messages, are relevant when designing *inhibitory* back-pressure [3]. To our knowledge, no prior work combines: (i) a LIF-style leaky accumulator that controls transmissions at *each tier* (sensor→cluster head→cloud), and (ii) explicit inhibitory control spikes that temporarily raise thresholds (or increase leak) to suppress bursts. This cross-tier, neuromorphic control-plane perspective distinguishes EdgeSynapse from SoD/prediction- and scheduler-only approaches.

## 3 System Model

We consider a network of sensor nodes grouped into clusters. Each node continuously observes a time-varying signal and encodes it into spikes using a LIF model. Specifically, let each sensor $i$ maintain a membrane potential $u_i(t)$ that evolves in discrete time steps $\Delta t$. On each step, the sensor updates

$$u_i(t) = \lambda u_i(t - \Delta t) + I_i(t)\,, \tag{1}$$

where $0 < \lambda < 1$ is a leak factor and $I_i(t)$ is the node–local, dimensionless input drive (e.g., rescaled and aggregated sensor readings). If $u_i(t)$ exceeds a threshold $\Theta_i$, the node emits an excitatory spike $s_i(t) = 1$ and resets $u_i(t) \leftarrow 0$; otherwise $s_i(t) = 0$. In this way, the sensor only communicates when its accumulated signal crosses the threshold. This is analogous to a biological neuron, where spikes are asynchronous impulses, and the leak causes older inputs to gradually decay, thereby preventing the accumulation of stale information, which reduces accuracy [2]. Spikes are sent to an upstream aggregator (e.g., cluster head). The aggregator likewise has a membrane potential $v(t)$ driven by incoming spikes:

$$v(t) = \mu v(t - \Delta t) + \sum_{i \in C} w_i\, s_i(t)\,, \tag{2}$$

where $0 < \mu < 1$ is a cluster leak factor, $C$ is the set of sensors in the cluster, and $w_i$ are weights. When $v(t) \geq \Theta_C$, the aggregator fires its own spike $S(t) = 1$ to the next tier (e.g., network server) and resets $v(t) \leftarrow 0$. In effect, the aggregator filters and compresses the sensor spikes. Multiple sensor spikes may accumulate into one cluster spike, further reducing transmissions. Both sensor and aggregator membrane equations implement a *leaky integrate-and-fire* (LIF) neuron model, which is widely used in neuromorphic computing for its simplicity [8]. The aggregator could also send an inhibitory spike to sensors to temporarily raise their thresholds, for example, when congestion/airtime limits approach.

A key advantage of this model is energy efficiency. Since components do nothing when inputs are below threshold, energy is used only for occasional transmissions. Previous work notes that neuromorphic hardware remains largely idle when inputs are sparse, yielding immense energy savings [5, 11]. In our network, spikes are sparse events, resulting in reduced communication. We stress that all terminology is defined in an application-agnostic way: a "spike" is simply a short message or packet, and an "inhibitory spike" is a message that causes downstream nodes to increase their thresholds.

## 4 Proposed Mechanism

EdgeSynapse implements an event-driven uplink using a LIF accumulator. Each sensor node and aggregator runs the following algorithm.

---

**Sensor $i$:**

(1) Continuously sense environment, obtaining input $I_i(t)$.
(2) Update membrane: $u_i \leftarrow \lambda u_i + I_i$.
(3) If $r_i > 0$, then $r_i \leftarrow r_i - 1$
(4) Else If $u_i \geq \Theta_i$, then:
- Transmit a spike (LoRaWAN uplink message) carrying minimal information (e.g., node ID).
- Reset $u_i \leftarrow 0$.
- Set $r_i \leftarrow \rho$
(5) Else do nothing (stay in low-power state).

---

**Aggregator (cluster head):**

(1) Receive spikes $s_i$ from sensors in the cluster.
(2) Update aggregate potential: $v \leftarrow \mu v + \sum_i w_i s_i$.
(3) If $v \geq \Theta_C$, then:
- Transmit a cluster-level spike (single LoRaWAN uplink) to server.
- Reset $v \leftarrow 0$.
(4) If congestion/airtime limits approach, then: Use downlink windows to send an inhibitory signal.

---

The above pseudo-code can be summarized as the EdgeSynapse protocol. Sensors mimic neurons by accumulating (integrating) their readings and firing (transmitting) when active. The aggregator acts like a higher-layer neuron, summing sensor spikes. This hierarchical structure means that small sensor fluctuations only cause local spikes, while significant cluster-wide changes trigger messages up the chain. In this way, communication is driven by events rather than periodic timing. The protocol also includes two control mechanisms to control spike frequency and suppress synchronized storms without frequent downlinks.

(1) **Refractory period (local, offline):** after a node fires a spike, it self-suppresses for a fixed window. The refractory counter, $r_i$, blocks further spikes for $\rho$ steps, acting as a local debounce.
(2) **Inhibitory spike (network back-pressure):** the aggregator monitors (i) a rolling count of child spikes and (ii) its own airtime utilization with respect to the legal duty-cycle. If a burst is detected or airtime margin is low, it sends a short *broadcast* inhibitory message that temporarily increases effective thresholds at children.

In LoRaWAN implementation (Class A devices), each sensor is normally in sleep mode. When $u_i$ crosses the threshold, the node sends an uplink spike. Immediately after, the node opens two receive windows (RX1 and RX2) as per LoRaWAN Class A operation [10]. The network server could exploit these windows to send an *inhibitory spike* (a short downlink packet) when aggregate burstiness or airtime pressure is detected. This spike instructs the cluster or node to temporarily raise the effective threshold $\Theta_{\text{eff}}(t) = \beta\,\Theta$ for $T_{\text{inh}}$ minutes. In practice, we focus on uplink spikes due to strict regional duty-cycle limits [19]. To that end, inhibitors are

rate-limited by a token bucket at the server, prefer multicast or short control payloads, and are only sent to recent talkers or nodes within the hot subcluster. If an inhibitor is missed, nodes simply continue with their local LIF policy. This mapping maintains a sparse control plane compatible with Class-A constraints.

## 5 Preliminary Evaluation

We evaluate whether a simple leaky integrate-and-fire (LIF) transmission policy with short refractory debounce and sparse broadcast inhibition can (i) reduce *energy* (airtime) at comparable accuracy and (ii) remain stable during multi-node bursts, compared to Periodic and Send-on-Delta (SoD) baselines.

### 5.1 Experimental Setup

We use real air quality traces from a multi-sensor node that integrates an *alpha-spectrometry radon sensor* (passive diffusion chamber), an *NDIR $CO_2$ sensor*, a *metal-oxide (MOX) gas sensor* for VOCs, a *solid-state temperature sensor*, a *digital relative-humidity sensor*, a *MEMS barometric-pressure sensor*, and a *laser-scattering optical particle counter* for $PM_1/PM_{2.5}$.

We consider three sampling windows: *Period-1* (Jan-Apr 2025), *Period-2* (Jul-Aug 2025), and *Period-3* (17-24 Oct 2025). All streams were resampled to a 5-minute grid with forward fill. We report total Energy [J], Number of Transmissions (TX), Number of Collisions, Airtime [s], and reconstruction error (RMSE). Energy is computed from LoRa airtime (SF7, 125 kHz, CR 4/5, preamble 8, 12 B payload) with constant TX current; MCU idle, RX-window energy, and downlink costs are not included (a limitation we discuss).

We compare the following policies.

- **Periodic:** one bundled uplink on a fixed interval (1 hour).
- **SoD:** transmit when the deviation from the last sent value exceeds a threshold $\delta$.
- **LIF:** proposed approach without inhibitory messages.
- **LIF+Inhib:** the cluster head monitors burstiness/airtime and, when needed, issues *broadcast* inhibitory messages that temporarily raise $\Theta$ for a group of nodes.

Unless stated otherwise, we use $\lambda = 0.99$ (discrete leak), $\Theta \in [15, 25]$ depending on period, refractory $\rho = 2$–$3$ sampling intervals, inhibitor gain $\beta = 2$, and $T_{\text{inh}} = 60$ min. The SoD threshold $\delta$ is tuned to match pre-storm activity for a fair comparison. The storm experiment uses a pure-ALOHA channel (no capture/SF diversity), per-packet random jitter within the 5 min slot, and identical traffic replicated to $N = 50$ nodes with a 1 h synchronous disturbance.

### 5.2 Energy–Accuracy Trade-offs

Figure 1 plots *Energy* vs. *RMSE* for the four policies in each period. Period-3 is the shortest window and therefore naturally sits toward lower energy but higher RMSE. Within a color (same period), LIF consistently improves the balance relative to Periodic and SoD, moving leftward (less energy) at similar or lower RMSE. Adding inhibition keeps energy low while preserving accuracy, especially when activity is bursty. The detailed metrics for Period-3 are presented in Table 1. Details of periods 1–2 are omitted for brevity, and follow the same trend.

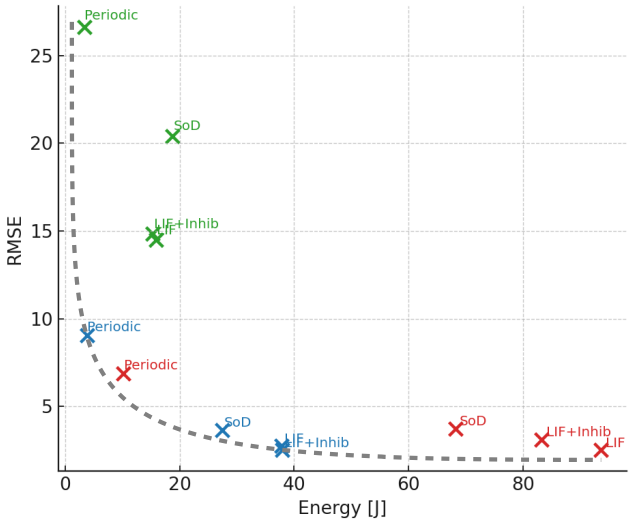


**Figure 1: Pareto comparison of energy consumption and reconstruction error across three sampling periods (red: Period-1, blue: Period-2, green: Period-3). Pareto front (dashed gray line) is interpolated for visualization purposes.**

**Table 1: Metrics for Period-3 (Oct 17–24, 2025)**

| Policy | Transmissions | Airtime [s] | Energy [J] | RMSE |
|---|---|---|---|---|
| Periodic | 179 | 7.378 | 2.922 | 26.850 |
| SoD | 1138 | 46.904 | 18.574 | 20.517 |
| LIF | 998 | 41.134 | 16.289 | 14.327 |
| LIF+Inhib | 939 | 38.702 | 15.326 | 14.680 |

### 5.3 Timing Behavior

Figure 2 illustrates how LIF debounces micro-bursts and how a short inhibitory window lifts the effective threshold. The plot shows the fused drive $I_i(t)$ (rescaled, arbitrary units) as well as LIF and LIF+Inhib state traces: $u_i(t)$ under the baseline threshold $\Theta$ and under the effective threshold $\Theta_{\text{eff}}$. When aggregate chatter or airtime pressure is detected, the *network server* issues a rare group downlink (forwarded by gateways) that, upon reception in the next RX window, instructs the node to raise its local threshold. This temporary elevation defers subsequent spikes while $u_i(t)$ continues to leak, breaking burst trains and lowering channel load. When the window ends, $\Theta_{\text{eff}}$ returns to $\Theta$ and normal spiking resumes. Note that, if $u_i(t)$ already exceeds $\Theta_{\text{eff}}$ at entry, a spike may still occur once refractory clears.

### 5.4 Storm Robustness (Multi-node)

To evaluate network-level effects, we simulate 50 edge nodes experiencing a synchronized disturbance ("event storm") under a pure-ALOHA LoRaWAN channel. Figure 3 compares the collision rates of SoD, LIF, and LIF + Inhib policies. While SoD nodes produce sharp bursts of simultaneous packets (0.24 collision rate), the LIF accumulator naturally desynchronizes activity through its leaky integration and short refractory period. Adding a sparse, broadcast

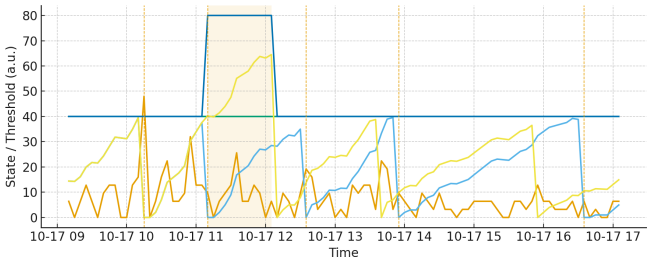


**Figure 2: Timing sketch from Oct 17, 2025. Light blue: LIF state $u_i(t)$ under fixed threshold, $\Theta$ (green); yellow: LIF+Inhib state $u_i(t)$ under effective threshold, $\Theta_{\text{eff}}$ (dark blue); orange: fused input drive $I_i(t)$ (scaled, arbitary units). The shaded region indicates inhibition.**

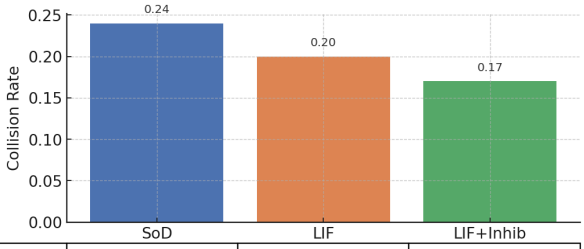


| | SoD | LIF | LIF+Inhib |
|---|---|---|---|
| Attempted_TX | 100 | 150 | 100 |
| Successful_TX | 76 | 120 | 83 |
| Collisions | 24 | 30 | 17 |

**Figure 3: Collision rates under a 1-hour event storm (50 nodes). SoD suffers the highest rate due to synchronized bursts, while LIF and LIF+Inhib reduce collisions by up to 29% through leak, refractory, and sparse inhibitory control.**

inhibitory signal further suppresses correlated transmissions, reducing the collision rate to 0.17 without significant throughput loss. These results confirm that hierarchical inhibitory feedback provides stability advantages during collective node activation events.

## 5.5 Takeaways

The experiments across three periods and a multi-node storm scenario yield several practical lessons about LIF telemetry with refractory debounce and sparse, broadcast inhibition.

- **Energy–accuracy balance.** Across all periods, LIF shifts the operating point leftward on the Energy [J]–RMSE plane relative to Periodic and SoD, offering lower energy at comparable (often better) RMSE.
- **Local stability via debounce.** A short refractory $\tau_{\text{ref}}$ acts as a practical debounce, suppressing micro-bursts from single disturbances without materially harming accuracy.
- **Network stability via inhibition.** Sparse, broadcast inhibitory spikes that temporarily raise $\Theta$ reduce correlated bursts and collision rates in multi-node "storm" scenarios, improving successful delivery.

# 6 Limitations and Future Work

The proposed spiking scheme shows promise for sustainable edge sensing, but practical challenges remain. The choice of leak rates and thresholds affects both responsiveness and transmission frequency. Too high a threshold saves energy but risks missing events, while too low yields many spikes. Additionally, hardware support for low-power integration is needed. Simple microcontrollers can implement a LIF update, but care is needed to minimize sleep current, as LoRaWAN devices have non-negligible idle draw [15].

From a networking standpoint, LoRaWAN Class A restricts downlinks, so true inhibitory signaling is limited. Future work could explore LoRaWAN Class B or C, or even mesh/routed LoRa, to allow more control messages. Scalability is also a concern: collisions become likelier if many nodes spike simultaneously. EdgeSynapse mitigates this via the cluster-level threshold, but real networks will require careful channel and duty-cycle management. Another practical limitation of our current evaluation is that the energy model accounts only for TX airtime with constant-current draw. MCU idle, RX1/RX2 windows, and downlink costs are excluded. Incorporating these terms will clarify the net energy picture and could further favor policies that suppress storms, such as EdgeSynapse, which has a dual mechanism against those. Finally, the very first packet of a burst at each active node cannot be preempted by design, because inhibitors are issued reactively and, in Class A, can only be delivered after a node opens RX1/RX2.

This work fits into broader visions of edge AI. For example, Ding *et al.* (2022) and Meuser *et al.* (2024) present roadmaps where novel ML methods optimize distributed sensor systems [6, 12]. Ferdowsi and Aral (2025) also highlight the potential of distributed neuromorphic edge computing for environmental monitoring [7]. EdgeSynapse contributes a specific communication mechanism in this space. In future work, we plan to prototype EdgeSynapse on real LoRaWAN sensors and cluster heads, measure energy savings, and integrate local inference (e.g., anomaly detection) into the spiking process. Extending the model to multi-hop networks and exploring bio-inspired synaptic plasticity are also promising directions.

# 7 Conclusion

We present EdgeSynapse, a preliminary design for a neuromorphic-inspired, hierarchical spiking protocol for remote sensing over LoRaWAN. By transmitting only when local LIF integrators fire, sensors and aggregators could dramatically reduce network traffic. Our initial evaluation on real data shows a significant decrease in messages and collisions, suggesting improved energy efficiency. While the Periodic update policy remains the absolute low-energy baseline in our traces, EdgeSynapse consistently improves the energy–accuracy trade-off over Send-on-Delta and reduces collision rates under synchronized bursts, suggesting a pragmatic path to more stable low-power sensing with sparse control. This approach bridges neuromorphic edge computing and low-power IoT, which we plan to further refine and validate.

## Acknowledgments

This research was funded in part by CHIST-ERA-22-SPiDDS-07 (TROCI Project) and by the Austrian Science Fund (FWF) [DOI: 10.55776/I6647].

## References


[1] Atakan Aral. 2024. The promise of neuromorphic edge AI for rural environmental monitoring. *Environmental Data Science* 3 (2024), e34.

[2] Atakan Aral, Melike Erol-Kantarci, and Ivona Brandić. 2020. Staleness control for edge data analytics. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 4, 2 (2020), 1–24.

[3] Lluís Casals, Bernat Mir, Rafael Vidal, and Carles Gomez. 2017. Modeling the energy performance of LoRaWAN. *Sensors* 17, 10 (2017), 2364.

[4] Jiechen Chen, Nicolas Skatchkovsky, and Osvaldo Simeone. 2023. Neuromorphic wireless cognition: Event-driven semantic communications for remote inference. *IEEE Transactions on Cognitive Communications and Networking* 9, 2 (2023), 252–265.

[5] Mike Davies, Narayan Srinivasa, Tsung-Han Lin, Gautham Chinya, Yongqiang Cao, Sri Harsha Choday, Georgios Dimou, Prasad Joshi, Nabil Imam, Shweta Jain, et al. 2018. Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro* 38, 1 (2018), 82–99.

[6] Aaron Yi Ding, Ella Peltonen, Tobias Meuser, Atakan Aral, Christian Becker, Schahram Dustdar, Thomas Hiessl, Dieter Kranzlmüller, Madhusanka Liyanage, Setareh Maghsudi, Nitinder Mohan, Joerg Ott, Jan Rellermeyer, Stefan Schulte, Henning Schulzrinne, Gurkan Solmaz, Sasu Tarkoma, Blesson Varghese, and Lars Wolf. 2022. Roadmap for edge AI: A Dagstuhl perspective. *ACM SIGCOMM Computer Communication Review* 52, 1 (2022), 28–33.

[7] Arman Ferdowsi and Atakan Aral. 2025. Distributed Neuromorphic Edge Computing: Theory and Applications in Environmental Monitoring. In *IEEE International Conference on Edge Computing and Communications (EDGE)*. IEEE, Piscataway, NJ, USA, 203–212.

[8] Wulfram Gerstner and Werner M. Kistler. 2002. *Spiking Neuron Models: Single Neurons, Populations, Plasticity.* Cambridge University Press, Cambridge, UK.

[9] Dragi Kimovski, Nishant Saurabh, Matthijs Jansen, Atakan Aral, Auday Al-Dulaimy, André B Bondi, Antonino Galletta, Alessandro V Papadopoulos, Alexandru Iosup, and Radu Prodan. 2023. Beyond von neumann in the computing continuum: Architectures, applications, and future directions. *IEEE Internet Computing* 28, 3 (2023), 6–16.

[10] LoRa Alliance. 2018. LoRaWAN Specification v1.0.3. https://lora-alliance.org/wp-content/uploads/2020/11/lorawan1.0.3.pdf. Accessed Oct. 24, 2025.

[11] Paul A Merolla, John V Arthur, Rodrigo Alvarez-Icaza, Andrew S Cassidy, Jun Sawada, Filipp Akopyan, Bryan L Jackson, Nabil Imam, Chen Guo, Yutaka Nakamura, et al. 2014. A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science* 345, 6197 (2014), 668–673.

[12] Tobias Meuser, Lauri Lovén, Monowar Bhuyan, Shishir G Patil, Schahram Dustdar, Atakan Aral, Suzan Bayhan, Christian Becker, Eyal De Lara, Aaron Yi Ding, and Gurkan Solmaz. 2024. Revisiting Edge AI: Opportunities and challenges. *IEEE Internet Computing* 28, 4 (2024), 49–59.

[13] Marek Miskowicz. 2006. Send-on-delta concept: An event-based data reporting strategy. *Sensors* 6, 1 (2006), 49–63.

[14] Mariusz Nowak, Rafał Różycki, Grzegorz Waligóra, Joanna Szewczyk, Adrian Sobiesierski, and Grzegorz Sot. 2022. Data Processing with Predictions in LoRaWAN. *Energies* 16, 1 (2022), 411.

[15] Solomon Ould and Nick S Bennett. 2021. Energy performance analysis and modelling of LoRa prototyping boards. *Sensors* 21, 23 (2021), 7992.

[16] Jongkil Park, Theodore Yu, Siddharth Joshi, Christoph Maier, and Gert Cauwenberghs. 2016. Hierarchical address event routing for reconfigurable large-scale neuromorphic systems. *IEEE transactions on neural networks and learning systems* 28, 10 (2016), 2408–2422.

[17] Margaret Richardson Ansah, Robert A Sowah, Joan Melià-Seguí, Ferdinand A Katsriku, Xavier Vilajosana, and Wiafe Owusu Banahene. 2020. Characterising foliage influence on LoRaWAN pathloss in a tropical vegetative environment. *IET Wireless Sensor Systems* 10, 5 (2020), 198–207.

[18] Pietro Savazzi, Anna Vizziello, and Fabio Dell'Acqua. 2024. Mutual Information Analysis of Neuromorphic Coding for Distributed Wireless Spiking Neural Networks. In *IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE)*. IEEE, Piscataway, NJ, USA, 160–165.

[19] The Things Network. 2025. *Duty Cycle.* https://www.thethingsnetwork.org/docs/lorawan/duty-cycle/ Accessed Oct. 24, 2025.