A Threshold-Fire-and-Refractory Reservoir: Event-Driven Dynamics Behind a Hidden Membrane

A Preprint

Caitlyn Meeks
AuxiLab Tenerife — research@auxi.cafe

August 12, 2026

registered 2026-08-12 | internal run log #31, #32 | 2 seeds × 15 cells, zero reversals | replicated before publication | promoted 2026-08-12

Abstract

This study originated in an open-ended design brief. It introduces the program's first event-driven reservoir: a threshold-fire-and-refractory network built from integrate-and-fire units rather than leaky integrators. Each of 1024 units silently accumulates drive into a hidden membrane potential; crossing a threshold θ emits an event of fixed amplitude, resets the membrane, and opens a brief refractory window; between emissions the emitted trace decays at rate `decay`. Only the emitted trace is ever fed back or read out — the membrane is structurally invisible, a hard information bottleneck with no learned component. Sweeping θ across five values and decay across three at the program's standard budget, on two independently drawn reservoirs (the second seed hardware-drawn): **ridge validation bits per character (bpc) has a strict interior optimum at θ=0.80**, both grid ends worse in both seeds, **and rises monotonically as the emitted-trace decay slows** — 10/10 adjacent-pair comparisons across both seeds, zero reversals. That direction is the opposite of how persistence parameters (spectral radius, leak rate) behave in every other reservoir family measured in this program. Linear decode depth agrees with bpc throughout — no decodability–usability split. The best cell reproduces within 0.0026 bits across seeds (3.5799 vs 3.5825), with active fraction 7.93% and 7.51%, inside the registered 3–15% band. Every cell trails the reference sparse-tanh reservoir by 0.43–0.44 bits: a new substrate family with a reproducible structural law, not an improved predictor.

Keywords reservoir computing · integrate-and-fire · event-driven sparse coding · spiking dynamics · emitted-trace decay · pre-registration

---

§1 Introduction

Every reservoir previously studied in this program updates every unit at every step through a smooth, continuous nonlinearity. An alternative is to compute with threshold events: a unit stays silent until its accumulated drive crosses a threshold, then emits a discrete event and falls silent again. That idea founds the liquid-state tradition [2]; this study builds the smallest version of it that remains directly comparable to the program's reference reservoir, and asks two registered questions. Is there an interior optimum in how readily the units fire? And does a slower-decaying emitted trace — greater state persistence, the direction that helps in every leaky-integrator reservoir measured here — help or hurt?

---

§2 Method

The state of each unit is a triple (emitted trace g, membrane potential m, refractory counter r):

```
drive = W @ g + Win @ e_c + b     # input one-hot e_c
# units not refractory integrate silently:
m ← 0.7·m + drive
if m crosses θ:                   # EMIT
    g ← 1.0                       # emit at fixed amplitude
    m ← 0                         # membrane resets
    r ← 3                         # refractory window opens
else:
    g ← decay · g                 # emitted trace decays
# ONLY g feeds back through W and reaches the readout;
# m is structurally invisible outside the unit
```

W, Win, and the bias are constructed exactly as in the program's reference reservoir [1] (sparse random, fan-in 10, spectral radius $\rho = 0.95$, N=1024), so the comparison against that reference sparse-tanh reservoir isolates the change of nonlinearity alone. The threshold grid $\theta \in \{0.35, 0.55, 0.80, 1.15, 1.70\}$ was anchored to the reservoir's own dynamic range by a disclosed never-firing probe over real text; decay $\in \{0.50, 0.80, 0.95\}$. Standard text8 splits [3], 2M training characters; instruments: ridge and logistic validation bpc, and U3 — a linear decode-depth probe, in the program's internal instrument numbering. The second seed (3818644351) was drawn from a hardware random-number generator, with the raw draw recorded in the registration before running.

---

§3 Pre-registered predictions

A 100k-character pilot run over the full 5×3 grid was carried out first and disclosed: it showed the interior θ optimum *and* reversed the a-priori expectation on decay (slower decay — greater persistence — had been expected to yield more memory; the pilot run showed the opposite). Predictions were finalized after it and labeled as pilot-informed:

- **J1** (headline): at decay 0.50, ridge validation bpc is non-monotone in θ with a strictly interior minimum. Named loss: the full-budget optimum sits at a grid endpoint.
- **J2** (direction): at every θ, ridge validation bpc rises monotonically with decay — 5/5, zero reversals allowed.
- **J3**: decode depth shows the same qualitative pattern as J1 and J2 — a decodability–usability split would be reported as a named alternative, not folded into a miss.
- **J4** (manipulation check): all 15 cells worse than the reference-reservoir anchor (3.1472) on ridge validation bpc. Named alternative, rated a major surprise: any cell beats it.

- **JQ1/JQ2** (bands): the best cell's active fraction lands in [3%, 15%]; the best cell trails the reference reservoir by more than 0.15 bits.

The seed-2 replication re-registered all six as R-J1–R-J4, R-JQ1–R-JQ2, with promotion pre-committed to R-J1 and R-J2 both being confirmed cleanly.

---

§4 Results

Table 1. Ridge validation bpc, both seeds (internal run log entries #31 / #32; each cell: seed 0 / seed 3818644351). Best cell in bold. Reference sparse-tanh reservoir anchor, cited: 3.1472. No logistic readout was voided by the program's fit-invalidating divergence rule.

θ \ decay	0.50	0.80	0.95
0.35	3.6468 / 3.6515	3.7097 / 3.7098	3.7694 / 3.7688
0.55	3.6078 / 3.6134	3.6880 / 3.6920	3.7555 / 3.7543
0.80	**3.5799 / 3.5825**	**3.6767 / 3.6855**	**3.7545 / 3.7581**
1.15	3.5877 / 3.6060	3.7093 / 3.7312	3.7763 / 3.8015
1.70	3.7067 / 3.7186	3.8405 / 3.8713	3.8985 / 3.9450

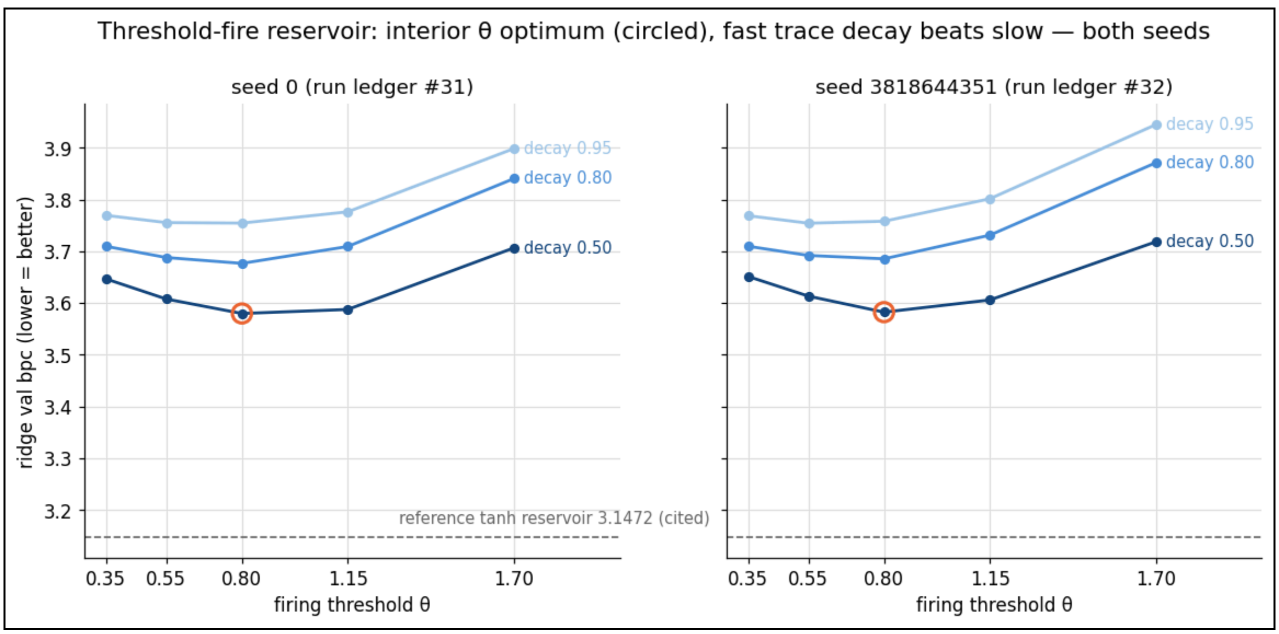

Figure 1. Ridge validation bpc against firing threshold θ for the three emitted-trace decay rates, both seeds (internal run log entries #31, #32). The circled point is the interior optimum (θ=0.80, decay=0.50), the same location in both seeds. The dashed line is the cited reference sparse-tanh reservoir anchor (3.1472) — every cell of this family sits well above it.

All six registered predictions were confirmed at both seeds, with zero reversals anywhere. The interior optimum sits at θ=0.80 in both seeds, with both grid ends worse; the best cell lands within 0.0026 bits of itself across two fully independent wiring draws. The decay ordering 0.50 < 0.80 < 0.95 holds at every θ in both seeds — 10/10 adjacent pairs. Decode depth mirrors bpc throughout (interior maximum at θ=0.80, monotone worsening as decay slows) — no decodability–usability split, unlike several other families in this program. Active fraction at the best cell: 7.93% / 7.51%, inside the registered band.

Comparison against the reference reservoir. Every cell in both seeds — 30 of 30 — is worse than the reference sparse-tanh anchor, by 0.4327 bits at the best cell (seed 0) and 0.4353 (second seed). The registered framing of a performance band rather than an improvement held exactly: event-driven sparse coding is not competitive with continuous mixing at this budget.

A diagnostic limitation, found and retained. At the sparsest-firing corner (θ=1.70, all three decay values), the U3 lag-1 diagnostic gate fails at full budget in both seeds, as it did at pilot scale — so both runs score 12/15 cells on the gated depth instrument, and the depth reading at that corner sits outside the gate (its direction is consistent with the rest of the grid). No bpc conclusion depends on that gate. The first run's committed conclusion text misrecorded this count as 0/15 gate failures; the replication caught it, and the correction is flagged in the program's records for its audit pass.

---

§5 Discussion

A persistence parameter with the opposite sign. In every leaky-integrator reservoir measured in this program, greater state persistence (spectral radius, leak rate) buys more memory until it buys instability. The emitted-trace decay rate is also a persistence parameter — and yet slower decay is monotonically worse, everywhere, at both seeds. The mechanistic reading registered with the finding: a slowly decaying trace keeps re-triggering the units it is wired to through W — measured active fraction rises with decay — and that recurrent re-excitation adds activity that behaves like noise, not memory. The information the readout uses lives in *which units have just fired*, and fast-decaying traces keep that signal clean.

An interior optimum, not a monotone trend. A threshold that is too low saturates the network with activity (12.33% of unit-steps at peak at the densest tested extreme, seed 0); one that is too high leaves it nearly silent (1.80% at the sparsest). The optimum sits between, at a reproducible threshold and a reproducible ~7.5–7.9% activity level. What this study does not claim: that any setting of this family is competitive with the reference sparse-tanh reservoir (none tested is), or anything beyond two seeds — uncertainty is reported as a min–max range across two draws, and the program's independent audit of this entry is still pending.

---

§6 References

1. H. Jaeger, "The 'echo state' approach to analysing and training recurrent neural networks," GMD Report 148, German National Research Center for Information Technology, 2001.
2. W. Maass, T. Natschläger, H. Markram, "Real-time computing without stable states: a new framework for neural computation based on perturbations," *Neural Computation* 14(11), 2002.
3. M. Mahoney, "Large text compression benchmark" (text8), mattmahoney.net/dc/textdata.

---

§7 Provenance

- **fb338dda** — the originating design brief; 15 cells, seed 0 (internal run log #31). Probe and pilot run disclosed before registration; all six predictions confirmed ⇒ provisional, replication automatically queued.

- **208425c9** — seed 3818644351, hardware-drawn with the raw draw recorded before running (internal run log #32). All six replication predictions confirmed, zero reversals ⇒ promoted to the program's confirmed-findings record per the pre-stated rule; the seed-0 gate-count error found and flagged during this pass.

All claims judged strictly against the registered wording; 2 seeds; replicated before publication; the program's independent audit of this entry is still pending. Internal designation: jellyfish swarm.
