BLOCK-DIAGONAL FRAGMENTATION OF A RESERVOIR IS FREE OR BETTER AT MATCHED WIDTH

A PREPRINT

CAITLYN MEEKS
AuxiLab Tenerife — research@auxi.cafe

August 11, 2026 (amended August 12, 2026)

registered 2026-08-11 | internal run log #12–#15 | 10 seeds, four registered passes
replicated before publication | promoted 2026-08-11 | amended 2026-08-12 after audit

Amendment notice, 2026 August 12. This preprint was independently audited on 2026-08-12. The decode-depth finding was an outright pass. The bits-per-character finding passed with a recorded **concern**, and two qualifications now attach to how it is stated. First, "never costs ridge bpc, strictly negative in all 10 seeds" is a *validation-split* result; on the untouched test split it is 9 of 10, and the trained-readout result falls from 8 of 10 to 7 of 10. Second, the two arms' spectral radii were set by different estimators: the 64 blocks land on ρ=0.950000 exactly, while the intact reservoir they are compared against runs slightly hot in 4 of the 10 seeds — an asymmetry that runs in the partition's favour. Neither qualification overturns the finding: an audit re-implementation with both arms scaled exactly, on four fresh seeds and two budgets, favours the partition 5 of 5 on validation *and* 5 of 5 on test. §4.1 reports all of this in full. No number in Table 1 changed.

ABSTRACT

An open-ended design brief on reservoir computing became a controlled fragmentation study: the mixing matrix of the program's reference 1024-unit sparse tanh echo state network was made block-diagonal in H equal blocks, each block's spectral radius ρ matched to the intact reservoir's *requested* ρ=0.95 (the realised radii differ; see §4.1), with the readout always seeing the full 1024-wide concatenated state. We registered a fragmentation *tax*; the pre-named alternative fired instead and, in its final ridge-first wording, hit two consecutive registered passes — ten seeds across four registered passes are on record. Two findings. (1) **Fragmentation is free or better, ridge-first**: the 64-block partition never costs ridge validation bits per character (bpc) — strictly negative in all 10 seeds — and wins outright on the trained logistic readout in 8 of 10, the two misses being dead ties (+0.0004, +0.0006), not reversals. An independent audit has since qualified how that first clause should be read (§4.1): "never" is a validation-split statement, holding 9 of 10 on the untouched test split, and the two arms' spectral radii were set by different estimators, an asymmetry favouring the partition. An audit re-implementation with both arms scaled exactly finds the effect intact, 5 of 5 on validation and test alike. The reference reservoir at ρ=0.95 sits measurably past its own bpc optimum; block-diagonalizing the mixing matrix moves the network to an effectively lower-ρ operating point at no cost. (2) **Memory survives**: the partitioned ensemble's decode depth never falls more than 0.20 characters short of the intact reservoir's, 10/10 seeds, despite each block spanning only 16 of the 1024 units. A third, stronger-sounding claim — that a mean-field coupled variant beats the intact reservoir outright — stands at 9 wins in 10 seeds and is deliberately *not* promoted: it missed its registered all-seeds bar, and unpromoted claims are reported here as clearly labelled context only.

Keywords reservoir computing · echo state networks · block-diagonal recurrence · sub-reservoir ensembles · spectral radius · pre-registration

§1 The question

The program's reference machine is a single large well-mixed reservoir: a sparse random recurrent matrix, spectral radius ρ=0.95, every unit potentially coupled to every other within a few steps [1]. The brief that prompted this study asks the opposite: what if the same 1024 units were instead 64 *separate* sub-reservoirs of 16 units each, driven by the same text stream side by side and never exchanging state? At matched total readout width, is fragmentation a tax (less mixing, less computation), a free change of coordinates, or — the outcome no registered prediction favored — an improvement?

§2 Method

The program's reference sparse tanh reservoir (N=1024, fan-in 10, standard input map and bias, leak 1.0), with its mixing matrix W made block-diagonal: H equal contiguous blocks of n = 1024/H units, each block's own spectral radius set to 0.95 — the intact reservoir's *requested* value — so per-block gain is nominally matched rather than diluted. The two arms' radii were set by different estimators and the intact arm's realised radius runs slightly high in four of ten seeds; §4.1 quantifies this and reports an exactly-scaled re-implementation. The readout always sees the full 1024-wide concatenated state, eliminating the feature-width confound measured separately in a companion study. Three instruments carry the paper: ridge validation bpc (optimizer-free), the trained logistic readout's validation bpc, and U3 decode depth — the linear decode-depth probe, so designated in the program's internal instrument numbering, which reports how many characters of recent history a ridge probe can reconstruct — all linear-readout instrumentation in the reservoir-computing tradition [3]. Standard text8 splits [4], 2M training characters.

The first pass (seed 0) swept H ∈ {1, 4, 16, 64}, plus two variants: a *mean-field coupled* ensemble, in which every block additionally receives its input map applied to the ensemble's mean state — a weak mean-state feedback channel, the coupling scaled to spectral radius 0.2 and the whole matrix rescaled so that total gain matches the uncoupled twin (topology changes, gain does not) — and a *heterogeneous-ρ* ensemble at H=16 (per-block ρ log-spaced 0.80 → 1.128, geometric mean 0.95). A reproduction gate ran first: the H=1 cell through the copied pipeline reproduced the archived reference reservoir of the addressed-memory study to Δ = 0.0000 on all three instruments, so every partitioned cell is directly comparable to the archive.

§3 Pre-registered predictions

Seed 0 registered five ordinal predictions with named alternatives. The registered tax prediction missed onto its pre-named alternative, and the registered depth band missed from above — in the favorable direction, both times — and those two misses are the paper:

- **R1** (fragmentation tax): bpc strictly increases along H = 1 → 4 → 16 → 64. Named alternative **A1** (reduced-mixing win): fragmentation *decreases* bpc toward the optimum of the program's ρ=0.8 spectral-radius ladder. (A1 fired.)
- **R2** (mean-field rescue): the mean-state feedback channel helps at H=16 and H=64. Named alternatives: A2, a readout-suffices null (|Δ| ≤ 0.005 everywhere); A3, over-synchronization (coupling hurts). (A2 fired at seed 0 — and was itself refuted by the reruns; see §4.1.)
- **R3** (heterogeneity no-win): the heterogeneous-ρ ensemble does not beat the uniform-ρ ensemble. (Hit; the heterogeneous arm cost +0.0070.)
- **R4** (depth falls): decode depth strictly decreases with H. Named alternative **A5** (redundant-encoding rescue): depth(H=64) lands within 1.0 character of intact. (The ordinal hit — and the magnitude belonged to A5: total loss 0.20 characters, against a registered band of [4, 8] characters of *remaining* depth that the partitioned ensemble beat from above.)
- **R5** (heterogeneous blocks decode deeper): the heterogeneous-ρ ensemble out-decodes the uniform-ρ ensemble at H=16. (Hit, thinly: +0.07 characters.)

Three registered replication passes followed — seeds 1–3, 4–6, and 7–9 — each written before running. The middle pass bundled its claims conjunctively and paid for it: its coupling sub-claim missed while the composite, the reduced-mixing win, and the depth rescue each hit clean, and the registered conjunctive promotion rule blocked *all* promotion that pass. The bundling was recorded as the registration's own design error, and the surviving claims were re-registered *unbundled* for seeds 7–9, each with its own promotion rule: T1, the mean-field coupled composite (all 3 seeds beyond −0.005); T2, the reduced-mixing win in ridge-first wording (ridge win in all 3 seeds AND logistic beyond −0.005 in ≥2 of 3); T3, the depth rescue (within 1.0 character in all 3 seeds).

---

§4 Results

Seed 0's ladder set the shape: logistic validation bpc 2.6727 → 2.6742 → 2.6627 → 2.6586 along H = 1/4/16/64 (the H=4 step inside the ±0.005 tie band; H=16 and H=64 clear wins), with the optimizer-free ridge instrument agreeing in direction and magnitude (3.1472 → 3.1303/3.1321 at H=16/64). The registered tax band ([+0.01, +0.10] bits) was missed with the opposite sign: the registered tax is a measured gain. Ten seeds later, the pattern of Table 1 is the finding.

Table 1. Intact reservoir (H=1) versus the 64-block independent partition, all ten seeds. Δ = partitioned − intact (negative favors the partition); depth in characters. Runs: ledger #12 (seed 0), #13 (1–3), #14 (4–6), #15 (7–9). No cell was voided by the program's divergence tripwire, a fit-invalidating divergence rule, in 48/48 cells across the family. All Δ are validation-split; §4.1 gives the test-split figures (ridge 9/10, logistic 7/10) and the spectral-radius asymmetry affecting seeds 0, 2, 6 and 9.

seed	Δ ridge val	Δ logistic val	depth H1	depth H64	depth loss
0	−0.0151	−0.0141	10.11	9.91	0.20
1	−0.0140	−0.0115	10.10	9.94	0.16
2	−0.0209	−0.0124	10.12	10.03	0.09
3	−0.0059	+0.0004	10.07	9.91	0.16
4	−0.0181	−0.0055	10.07	9.95	0.12
5	−0.0080	−0.0084	10.06	10.06	0.00
6	−0.0223	−0.0197	10.11	9.94	0.17

7	−0.0140	−0.0084	10.14	10.03	0.11
8	−0.0023	+0.0006	10.05	10.04	0.01
9	−0.0217	−0.0180	10.10	10.01	0.09

The ridge column is negative ten times out of ten. The logistic column is beyond the −0.005 bar eight times out of ten, and the two misses (seeds 3 and 8) are dead ties a fraction of a thousandth of a bit wide — on the validation split the partitioned ensemble never loses, it occasionally only draws; §4.1 gives the test-split figures, where seed 8 reverses. The depth column's worst case is 0.20 characters, for blocks that each span 16 units where the intact reservoir's probe sees 1024. T2 and T3 each hit their unbundled registered bars on seeds 7–9 and were promoted on their own merits (T2's second ridge-first pass; T3's third pass).

---

§4.1 Amendment: two qualifications found by the independent audit

This study was independently audited on 2026-08-12 and the audit recorded a **concern** against the bpc finding. Every promoted number re-derived exactly from the raw result files, the reproduction gate against the archived reference reservoir reproduced it to the last digit, and all 20 cells cleared the lag-1 diagnostic gate with no divergence events. Two qualifications attach to how the result was stated, and both belong in this paper rather than only in the program's files.

(a) The two arms were not scaled by the same instrument. "Matched per-block spectral radius" is true of the blocks and not of the reservoir they are compared against. The auditor rebuilt every matrix and measured its *realised* spectral radius with dense eigenvalues: all 640 blocks of the partitioned arm land on ρ = 0.950000 to within 6×10^{-15}, but the intact arm does not, because the sparse iterative eigenvalue routine used to scale it under-estimates on a 1024 × 1024 draw. The intact reservoir therefore runs *hotter* than requested in 4 of the 10 seeds (by +0.0015, +0.0051, +0.0030 and +0.0070; never cooler), and this program's own reference anchor is in fact a ρ = 0.9515 reservoir rather than a ρ = 0.95 one. Since higher ρ is worse bpc over this range, the asymmetry runs in the partition's favour wherever it bites: those 4 seeds carry 1.93× the mean ridge win of the 6 clean ones (−0.0200 against −0.0104). The program's own ρ ladder accounts for only about 12% of that group gap, leaving most of it attributable to chance at n = 10, but the direction of the bias is not in doubt.

(b) "Never" is a statement about the validation split. That is what was registered, and ridge's temperature parameter is fitted in-sample on validation. On the untouched test split the ridge result is 9 of 10 rather than 10 of 10 — seed 8, whose validation margin was already the thinnest at −0.0023, flips to +0.0038 — and the logistic result falls to 7 of 10 beyond the −0.005 bar. The effect keeps about 85% of its size (test mean −0.0123 against validation −0.0142).

The finding survives its own strongest attack. The auditor re-implemented the comparison end to end — own states, own ridge solve, own temperature fit, own bpc — with *both* arms scaled to ρ = 0.95 by exact dense eigenvalues, removing qualification (a) by construction. Across four fresh hardware-drawn seeds and two training budgets, the partitioned arm wins 5 of 5 on ridge validation *and* 5 of 5 on ridge test, at Δ from −0.0008 to −0.0285 on validation and −0.0060 to −0.0322 on test. The effect is real and is not an artifact of the scaling asymmetry. It should be quoted with that spread rather than as a single number, and the ten-seed "never" should be quoted as a validation-split result. A repair run is registered in the program's queue, and the radius estimator itself has been filed as an apparatus defect rather than a scientific one.

The companion decode-depth finding was audited in the same session and **passed**. Decode depth was re-implemented from the registration's wording and recomputed from each result file's raw per-lag accuracy ladder rather than read from its stored summary: 20 of 20 cells agree exactly, at a maximum difference of 0.0000. The per-seed losses are 0.20 / 0.16 / 0.09 / 0.16 / 0.12 / 0.00 / 0.17 / 0.11 / 0.01 / 0.09 — maximum 0.20 and median 0.115, exactly as stated, with no seed decoding deeper than the intact reservoir. Qualification (a) runs *conservative* here: decode depth rises with ρ over this range, so an accidentally hot intact reservoir decodes deeper and makes the measured loss larger, which makes the ≤ 0.20 bound harder to satisfy rather than easier. Depth is a closed-form ridge ladder with no learning rate, no early stopping and no divergence rule, so it carries none of the optimizer exposure the program's trained-readout claims do. The auditor's independent re-measurement, both arms exactly scaled, gives depth losses of 0.00 / 0.19 / 0.07 / 0.06 / 0.01 — 5 of 5 inside the promoted bound.

What did not survive, reported with equal weight. The mean-field coupled composite — H=64 with the mean-state feedback channel beating the *intact* reservoir — hit its bar on seeds 4–6 (−0.0132 / −0.0166 / −0.0152) but missed on seeds 7–9: −0.0095 / −0.0024 / −0.0277, where seed 8's −0.0024 sits inside the tie band. Across all seeds the composite stands at 9/10, the sole miss a tie, and it is *not* in the program's confirmed-findings record — the registered rule required all three, and a 9/10 regularity that failed its stated bar stays unpromoted. Likewise the coupling sub-benefit (mean-field coupled versus its own uncoupled twin): favored in 11 of 14 comparisons across seeds 0–6, mean ≈ −0.003 bits, sign-unstable — retired to measured-context wording, deliberately not re-registered. And the seed-0 registered claims R3/R5 (heterogeneity costs bpc +0.0070; heterogeneous blocks decode +0.07 deeper) remain single-pass results, not confirmed findings.

---

§5 Discussion

The reference reservoir operates above its own optimum. The intact reservoir at a requested ρ=0.95 — realised at roughly 0.9515 in four of the ten seeds (§4.1) — sits measurably past its own bpc optimum (the program's spectral-radius ladder places that optimum near ρ=0.8). Cutting W into 64 blocks — while explicitly holding per-block spectral radius at 0.95 — mixes less, and behaves as an effectively *lower-ρ* operating point at matched width and matched per-block gain. That account is partial rather than complete: the partition recovers only about 47% of the ρ-ladder gap between ρ=0.95 and ρ=0.8, so "effectively lower ρ" names a direction the mechanism moves in, not the whole of it. The result is not an artifact of the trained readout's optimizer: the optimizer-free ridge instrument shows it in all ten seeds on validation, and nine of ten on the untouched test split (§4.1), which is why the promoted wording is ridge-first. Decomposing a reservoir into decoupled sub-reservoirs has precedent as an architecture proposal [2]; what this study adds is the matched-gain, matched-width control and the pre-registered accounting that turns the absence of a fragmentation cost into a measured statement rather than a design preference.

Redundant partial views collectively remember. Sixty-four 16-dimensional partial views of the same character stream are almost as ridge-decodable, taken together, as one undivided 1024-dimensional encoding — worst loss a fifth of a character across ten seeds. Depth, unlike bpc, was expected to pay heavily for fragmentation (the registered band put the partition's remaining depth at just [4, 8] characters); it barely paid at all.

A methodological note. The strongest-sounding result in this family — that the mean-field coupled ensemble beat the intact reservoir every time — held for eight consecutive seeds, then drew on the ninth; nine wins in ten is a weaker statement than the seven-for-seven the claim stood at when its final registered pass was written. The program's rule against promoting on evidence accumulated outside the registered bar is exactly what kept it out of this paper's confirmed findings, and it is recorded here as context, clearly labelled, because that is what this section exists for.

---

§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. Y. Xue, L. Yang, S. Haykin, "Decoupled echo state networks with lateral inhibition," *Neural Networks* 20(3), 2007.
3. M. Lukoševičius, H. Jaeger, "Reservoir computing approaches to recurrent neural network training," *Computer Science Review* 3(3), 2009.
4. M. Mahoney, "Large text compression benchmark" (text8), mattmahoney.net/dc/textdata.

---

§7 Provenance

- **77cd2338** — the initiating design brief; 9 cells, seed 0 (internal run log #12). Registered before running; exact-reproduction gate against the archived reference reservoir passed at Δ = 0.0000.
- **8f5bdd07** — 15 cells, seeds 1–3 (internal run log #13). The replication split: reduced-mixing win 2/3 on logistic, 3/3 on ridge; the coupling null refuted on this pass, later destabilised (§4.1).
- **cc447b8d** — 15 cells, seeds 4–6 (internal run log #14). All but one registered claim hit; the conjunctive promotion rule blocked promotion, recorded as a registration-design error.
- **d3881e88** — 9 cells, seeds 7–9 (internal run log #15). Claims unbundled; T2 (reduced-mixing win, ridge-first) and T3 (depth rescue) promoted on their own rules; T1 (mean-field coupled composite) hit 2 of 3 seeds, short of its registered all-3 bar, and stays out.

Audit status: independently audited 2026-08-12. The bpc finding drew a recorded CONCERN — the entry stays promoted; the decode-depth finding was an outright AUDIT PASS. The concern, stated in one line: the intact and partitioned arms had their spectral radii set by different estimators, so "matched per-block gain" holds of the blocks but not of the reservoir they are measured against, and the promoted "never" is a validation-split result that stands at 9 of 10 on the untouched test split. Both qualifications are reported in §4.1, together with the audit's own re-implementation — both arms exactly scaled — in which the effect survives 5 of 5 on validation and test. A repair run is registered in the program's queue.

All claims judged strictly against the registered wording, misses included above; 10 seeds across four registered passes; replicated before publication. Internal designation: goat herd.
