BEYOND FIXED POINTS: SUPERPOLYNOMIAL CAPACITY OF ASYMMETRIC HOPFIELD NETWORKS

Anonymous authors
Paper under double-blind review

ABSTRACT

Classical Hopfield networks are limited to static patterns due to symmetric weights, whereas asymmetric networks can encode temporal sequences via limit-cycle attractors. Achieving high-capacity storage of long sequences in classical synchronous asymmetric networks, however, has remained a challenge. We present a simple and robust construction within the classical asymmetric Hopfield model with binary neurons and synchronous updates, that allows n neurons to support $\exp\big(\Omega(n/(\log n)^2)\big)$ distinct limit-cycle attractors, each with period $\exp\big(\Omega(\sqrt{n}/\log n)\big)$ and robust to random noise with flip probability up to $\frac{1}{2}-o(1)$, so that both the number and the length of stored sequences are superpolynomial in n. This is the first demonstration of such capacity for asymmetric Hopfield networks, which we obtain by combining results from combinatorics, number theory and the analysis of opinion dynamics. Our findings show that synchronous asymmetric Hopfield networks possess a sequence-memory capacity that is larger and more robust than previously recognized, suggesting that, in both biological and artificial neural systems, robust sequence representation can be achieved through coarse architectural motifs rather than complex nonlinearities.

1 INTRODUCTION

Associative memory, the ability to retrieve complete patterns from partial or noisy cues, is a fundamental concept in both neuroscience and artificial intelligence. The Hopfield network, a recurrent neural network model, stands as a canonical example of how collective neuronal dynamics can give rise to associative memory (Hopfield, 1982). In its classic formulation, the network is characterized by a symmetric weight matrix. This symmetry guarantees the existence of a Lyapunov energy function, ensuring that the network's state always converges to a stable fixed point.[1] These fixed points act as stored memories, making the symmetric Hopfield network an effective model for static pattern recognition.

However, the constraint of symmetry, while ensuring stability, fundamentally limits the network's dynamical repertoire to point attractors (see Section 3). Many cognitive processes, such as motor control, central pattern generation, and sequential reasoning, are inherently dynamic and involve recalling and generating temporal sequences rather than static patterns. By allowing for asymmetric connections, asymmetric Hopfield networks break the limitations imposed by a global energy function. This asymmetry opens the door to a richer set of dynamics, including the emergence of limit cycles, i.e., periodic sequences of states that can serve as attractors. These limit-cycle attractors, formalized in Section 3 (Definition 2), provide a natural mechanism for storing and retrieving temporal sequences. For instance, in biological systems, high-capacity limit cycles can act as robust Central Pattern Generators (CPGs) for complex rhythmic motor control, while in artificial intelligence, they can serve as high-dimensional "clocks" for temporal reasoning and sequence timing. While the potential for asymmetric (or directed) networks to store sequences is known, and various construction methods have been proposed (Zhang et al., 2013; Muscinelli et al., 2017), the question of their storage capacity remains a critical area of investigation. Muscinelli et al. (2017) demonstrated that asymmetric Hopfield networks with n neurons can support limit-cycle attractors of length 2^n, thereby realizing dynamics that traverse the entire state space. However, this result relies on a highly specialized

[1] If the update rule is synchronous, the network may exhibit limit cycles of period two; however, limit cycles of higher period cannot occur.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

construction in which the weights are carefully engineered on a case-by-case basis. As a consequence, while theoretically striking, the model is unlikely to provide a realistic account of memory storage in biological or practical settings. A further challenge for such constructions is robustness. In practical settings, one would expect the dynamics to reliably return to the same limit-cycle attractor after perturbations induced by noise. However, if the dynamics traverse the entire state space, robustness is fundamentally unattainable. Recent advances have pushed the capacity of Hopfield-type networks by modifying the standard model with nonlinear interaction terms or continuous states (Krotov & Hopfield, 2021; Ramsauer et al., 2021; Chaudhry et al., 2023). This raises a key question:

How can one construct a model that

- *stays within the classical synchronous asymmetric Hopfield framework,*
- *encodes temporal sequences as limit-cycle attractors,*
- *stores a large number of such sequences, each of them long,*
- *remains robust to noise, and*
- *depends only on coarse architectural constraints rather than on finely tuned weights?*

This paper addresses this challenge directly. We introduce a novel construction for asymmetric Hopfield networks based on a very simple network topology, presented in Section 4. We show that this architecture supports a superpolynomial number of distinct limit-cycle attractors (with respect to the number of neurons), with periods that can themselves be superpolynomial in the network size. These periods come not from one block cycle, whose period is at most its length, but from many cycles of different lengths running in parallel: their lcm is superpolynomial in n although each length is $O(\sqrt{n})$ (Figure 1). Moreover, the proposed construction is inherently robust to even extreme levels of noise, admitting non-trivial basins of attraction around each cycle. Throughout, we work with the parallel update rule introduced by Little (Little, 1974); the asynchronous single-neuron regime is left open (see the Limitations). This should be viewed as a classical variant of the same binary-threshold architecture rather than as a modern modification of the model: Little's rule uses the same linear local fields and sign activation, differing only in the update schedule. Since our goal is to study deterministic sequence storage through periodic attractors, the synchronous regime is also the natural classical setting in which those limit cycles appear as first-class dynamical objects. We now present an informal statement of our main results.

Informal Theorem (Informal statement of Theorems 1 and 2)**.** *There exists an asymmetric Hopfield network with n neurons that supports $\exp\big(\Omega\big(\frac{n}{(\log n)^2}\big)\big)$ distinct limit-cycle attractors, each with period $\exp\big(\Omega(\sqrt{n}/\log n)\big)$. Moreover, these limit cycles remain robust under independent flips with probability up to $\frac{1}{2} - \Omega\big(\frac{1}{\sqrt{\log n}}\big)$, and hence up to $\frac{1}{2} - o(1)$,* (Theorem 1)*. Furthermore, this robustness extends to highly sparsified network topologies, with only a modest reduction in capacity and period length, both of which remain superpolynomial in n* (Theorem 2)*; the corresponding quantitative noise thresholds are made explicit in the proofs.*

We further show in Section 5 (Theorem 3) that only a sparse backbone of $3n$ edges needs to be arranged in advance: an adversary may control a vanishing fraction of the neurons, spread evenly across the network, giving them arbitrarily many incoming edges, and the dynamics still return to the limit cycle up to a vanishing fraction of neurons. In Appendix E.5, we also record complementary global convergence statements: in the dense construction every initial configuration rapidly reaches some limit cycle deterministically, in the sparse construction the same holds with high probability, and under adversarial perturbations the trajectory reaches the corresponding weak limit cycle. We also provide empirical support to our results through experiments in Section 6. Our main contributions are summarized below.

- **Classical Synchronous Asymmetric Hopfield Model:** We demonstrate that it is possible to achieve very large sequence memory capacity without altering the classical asymmetric Hopfield network model. Our construction uses standard binary neurons and the synchronous update rule (Little, 1974), in contrast to recent works that rely on nonlinear interactions or continuous states. Our block-based architectural motif suggests that enforcing coarse block-cyclic topologies in modern ML sequence models (e.g., structured state-space models or sparse RNNs) could inherently boost robust sequence capacity.
- **Superpolynomial Capacity and Period Length:** We provide a constructive proof that our network architecture can store a superpolynomially large number of distinct limit-cycle attractors, where the period of these attractors is also superpolynomially long in the number of neurons.

- **High Simplicity and Inherent Robustness:** Up to coarse architectural constraints, the result is preserved under highly sparsified connectivity and addition of adversarial connections, enhancing biological plausibility relative to highly fine-tuned alternatives. Moreover, our construction is robust, ensuring stable retrieval even under extreme perturbations by random flips with probability up to $\frac{1}{2} - \Omega(1/\sqrt{\log n})$ in the dense case, and more generally up to $\frac{1}{2} - o(1)$ once $d = \omega(\log n)$.

2 Related Work

Within Hopfield-type networks, the move from symmetric to asymmetric weight matrices is what enabled temporal dynamics to be modeled, and exactly solvable asymmetric models were analyzed by Derrida et al. (1987) and, at arbitrary finite size, by Fasoli et al. (2018), who also give a learning rule storing point and cyclic attractors under noise. Sequences can be embedded with temporally asymmetric Hebbian rules, for which Gillett et al. (2020) characterize retrieval and compute the storage capacity analytically.

Classic results on the capacity of asymmetric Hopfield models include Bastolla & Parisi (1997), who showed that the dynamics can exhibit $O(n)$ limit cycles with periods exponential in n. This picture was refined by Hwang et al. (2019), who (informally) characterized the distribution of these long cycles. In a different regime, Hwang et al. (2020) demonstrated that an exponential number of limit cycles of constant length $O(1)$ can arise.

More targeted construction methods have been developed to store specific cycles. For instance, Zhang et al. (2013) studied when a prescribed cycle is admissible, constructed corresponding connectivity matrices via the pseudoinverse learning rule (Personnaz et al., 1986), and analyzed the topology of the resulting networks. Closer to our work, Muscinelli et al. (2017) demonstrated a method to construct a standard asymmetric Hopfield network that produces a stable orbit of maximal length (2^n), covering the entire state space. While this highlights the potential for long sequences, our approach differs by focusing on the combinatorial generation of a large number of distinct, long limit-cycle attractors. Unlike the single maximal-length orbit, which requires highly specific weight assignments, our construction relies on coarse architectural constraints rather than precise tuning, making it more appealing from the standpoint of biological plausibility.

Separately, a recent revolution in associative memory has come from “Modern Hopfield networks" or “Dense Associative Memories" (Krotov & Hopfield, 2021; Ramsauer et al., 2021). These models modify the energy function of the classical Hopfield network model, often by introducing higher-order or non-polynomial interactions. This modification allows for a storage capacity for static patterns that is exponential in the number of neurons, a dramatic improvement over the linear capacity of the original model. These modern Hopfield networks have also been shown to be equivalent to the attention mechanism in Transformers (Ramsauer et al., 2021). A related line studies bipartite, two-block architectures in which the high capacity comes from the block representation itself (Krotov & Hopfield, 2021; Chandra et al., 2025; Shafiei Kafraj et al., 2026). Those models are continuous and symmetric where ours is binary and asymmetric, but the reliance on block structure to obtain capacity is a point of contact with the construction developed here.

Naturally, an effort has been made to extend the benefits of these high-capacity models to the domain of sequence memory. Chaudhry et al. (2023) adapted the Dense Associative Memory framework to store and retrieve long sequences, achieving a significantly improved sequence capacity by leveraging a nonlinear interaction term. A key distinction of our work is that we do not modify the fundamental neuron model or its update rule. We remain within the classical framework of the synchronous asymmetric Hopfield network with binary neurons and majority vote update rule. Our contribution is to show that, through the realization of certain families of graph topologies, it is possible to achieve a large number of long periodic sequences without resorting to the nonlinearities or continuous states that characterize modern Hopfield networks. A detailed side-by-side comparison is deferred to Table 4 in Appendix G.

Beyond Hopfield-type models, cyclic structure is also what makes attractor counting tractable in networks of coupled oscillators, where ring topologies have long been used for this purpose. Closest to our setting, Guo et al. (2025) build a Kuramoto associative memory on a honeycomb topology, a chain of cycles sharing nodes, encoding memories as stable phase-locked configurations: they obtain $(2\lceil n_c/4 \rceil - 1)^m$ of them for m cycles of n_c oscillators each, with no spurious memories, and an extended version adds a guaranteed minimum basin size independent of network scale (Ogranovich

et al., 2026). Their count is exponential in the network size, hence larger than ours, and we do not claim to improve on it. The two results count different objects: their memories are stable *equilibria*, that is static patterns, whereas each of ours is a limit cycle of superpolynomial period. The model classes differ too, continuous-phase oscillators against binary-threshold synchronous updates.

3 Setup: Asymmetric Hopfield Networks

We begin by introducing the basic framework of asymmetric Hopfield networks and the dynamical notions relevant to our analysis. For convenience, Table 1 in Appendix G collects the recurring notation used throughout the paper.

An asymmetric Hopfield network consists of n binary neurons with state vector $x(t) \in \{-1,+1\}^n$ evolving in discrete time. The network is specified by a (not necessarily symmetric) weight matrix $W \in \mathbb{R}^{n\times n}$, where w_{ij} represents the influence of neuron j on neuron i, and we assume $w_{ii} = 0$ for all i. The input to neuron i at time t is $h_i(t) = \sum_{j=1}^{n} w_{ij}x_j(t)$. We consider synchronous majority-rule dynamics, where all neurons update simultaneously according to $x_i(t+1) = \mathrm{sign}(h_i(t))$, with the convention that $\mathrm{sign}(0) = x_i(t)$, so that neurons maintain their state in the case of a tie. This defines a deterministic update map $F : \{-1,+1\}^n \to \{-1,+1\}^n$.

Definition 1 (Orbit)**.** *Let $F : \{-1,+1\}^n \to \{-1,+1\}^n$ denote the synchronous update map of the network. The* orbit *or* trajectory *starting from an initial state $x \in \{-1,+1\}^n$ is the sequence $x,\ F(x),\ F^2(x),\ F^3(x),\dots$*

Definition 2 (Limit cycle and basin of attraction)**.** *A* limit cycle *of period $P \geq 1$ is a set of distinct states $\mathcal{C} = \{x^{(0)}, x^{(1)}, \dots, x^{(P-1)}\}$ such that, for the synchronous update map $F : \{-1,1\}^n \to \{-1,1\}^n$, we have $F(x^{(i)}) = x^{(i+1)}$ for $i = 0,\dots,P-2$ and $F(x^{(P-1)}) = x^{(0)}$. The special case $P = 1$ is a fixed point. The* basin of attraction *of $\mathcal{C}$ is $\mathcal{B}(\mathcal{C}) = \{y \in \{-1,+1\}^n : \exists T \geq 0$ such that $F^T(y) \in \mathcal{C}\}$. Clearly, $\mathcal{C} \subseteq \mathcal{B}(\mathcal{C})$; we call $\mathcal{C}$ a* limit-cycle attractor *if $\mathcal{C} \subsetneq \mathcal{B}(\mathcal{C})$.*

Since the state space is finite, every trajectory eventually becomes periodic. The long-term behavior is therefore described by fixed points and limit cycles.

In contrast to the symmetric case, asymmetric Hopfield networks generally do not admit a global energy function, and their dynamics can exhibit nontrivial periodic behavior. The state space is thus partitioned into disjoint basins corresponding to different attractors. To capture robustness under partial recovery, we introduce a weaker notion of limit cycles that allows a bounded number of neuron mismatches. We start with the notion of Block-max Hamming distance.

Definition 3 (Block-max Hamming distance)**.** *Let Σ be an alphabet, let $r, d \in \mathbb{N}$, and let $\mathcal{S}$ be a partition of $\{1,\dots,rd\}$ into r parts of size d. After fixing an ordering of the parts in $\mathcal{S}$, any two strings $u, v \in \Sigma^{rd}$ decompose as $u = (u_1, u_2, \dots, u_r)$ and $v = (v_1, v_2, \dots, v_r)$, where $u_j, v_j \in \Sigma^d$ are the jth blocks, for $j = 1,\dots,r$. Let d_H denote the Hamming distance on Σ^d. The block-max Hamming distance between u and v with respect to $\mathcal{S}$ is defined as $d_{\mathrm{BM}}^{\mathcal{S}}(u,v) = \max_{1\leq j\leq r} d_H\big(u_j, v_j\big)$.*

In our block-cyclic architecture constructions, $\mathcal{S}$ is the canonical partition into blocks of size d. We therefore write $d_{\mathrm{BM}}^{(d)}$ for the corresponding block-max Hamming distance.

Remark 1. The block-max Hamming distance is an ℓ_∞-type block metric: it controls the worst number of mismatches inside any single block, rather than the total number of mismatches across the network. This viewpoint is reminiscent of locality notions in coding theory, such as locally recoverable codes, where the value of a coordinate can be recovered by querying a small prescribed set of other coordinates (Gopalan et al., 2012).

Definition 4 ((k,d)-weak limit cycle)**.** *Let $F : \{-1,+1\}^n \to \{-1,+1\}^n$ be a synchronous update map, and let $\mathcal{C} = (c_0, c_1, \dots, c_{P-1})$ be a cyclic reference sequence of distinct states. The sequence need not itself be an exact limit cycle of F. Assume that the neurons are partitioned into r disjoint blocks $S_1,\dots,S_r$, each of size d, so that $n = rd$. After fixing an ordering of the neurons inside each block, this partition induces a decomposition of any state $x \in \{-1,+1\}^n$ as $x = (x_1,\dots,x_r)$, where $x_j \in \{-1,+1\}^d$. Let $d_{\mathrm{BM}}^{(d)}$ be the corresponding block-max Hamming distance.*

*We say that $\mathcal{C}$ is a (k,d)-*weak limit cycle *for F if for every timestep of the reference sequence $i \in \{0,\dots,P-1\}$ there exists a neighborhood $\mathcal{N}_i \subseteq \{-1,+1\}^n$ with $c_i \in \mathcal{N}_i$, such that for every*

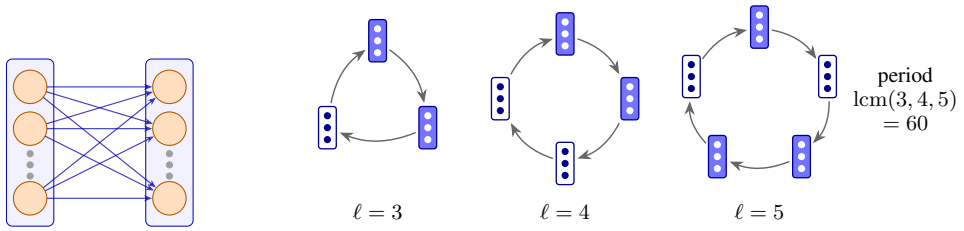

(a) Two connected blocks.

(b) Block cycles of lengths 3, 4, 5 with $d = 3$: filled blocks are in state $+1$, white ones in state -1, and each arrow is a connection as in (a).

Figure 1: Where the long periods come from. At each synchronous update every block copies the sign of its predecessor, so each labeling rotates by one block, and a block cycle of length ℓ with an aperiodic labeling has period ℓ. Cycles on disjoint neurons run in parallel, so the network's period is the lcm of their lengths: 60 here, though no cycle is longer than 5.

initial state $\tilde{x}(0) \in \mathcal{N}_i$, there exist integers $T \geq 0$ and $s \in \{0, \ldots, P-1\}$ satisfying

$$d_{\mathrm{BM}}^{(d)}\big(F^{T+t}(\tilde{x}(0)),\, c_{s+t \bmod P}\big) \leq k \quad \textit{for all } t \geq 0.$$

When the reference sequence is an exact limit cycle of F, this reduces to the usual cycle notion with blockwise error tolerance k.

A central object in this work is that of noise-robust limit cycles.

Definition 5 (Noise robustness)**.** *Let $\mathcal{C} = \{x^{(0)}, \ldots, x^{(P-1)}\}$ be a limit cycle of an asymmetric Hopfield network with update map F. We say that $\mathcal{C}$ is (p, ρ)-noise robust if, for $\tilde{x}(0)$ obtained from any $x^{(i)} \in \mathcal{C}$ by independently flipping each neuron (from 1 to -1 or from -1 to 1) with probability p, $\Pr\big[\exists\, T \geq 0$ such that $F^T(\tilde{x}(0)) \in \mathcal{C}\big] \geq 1 - \rho$, where the probability is over the noise (and any randomness in the network, if applicable).*

4 Super-Polynomial Capacity of Asymmetric Hopfield Networks

In this section, we introduce the construction underlying our theoretical results. These results concern capacity rather than learning: as in classical capacity results, we exhibit weights under which superpolynomially many long limit cycles coexist, and we do not give a rule that finds weights for prescribed sequences. We then present our main results along with proof sketches; full proofs are deferred to Appendix E. In Section 5, we discuss several relaxations of this construction.

Definition 6 (Blocks and block cycles)**.** *A* block *of size d is a collection of d neurons with no internal connections between them. Given two blocks S_1 and S_2, we write $S_1 \to S_2$ if every neuron in S_1 is connected to every neuron in S_2, with all edges directed from S_1 to S_2, and each such edge assigned weight 1 (see Figure 1a). A* block cycle *of length ℓ is a collection of blocks $S_0, S_1, \ldots, S_{\ell-1}$ (depicted in Figure 1b) such that*

$$S_0 \to S_1 \to \cdots \to S_{\ell-1} \to S_0.$$

A block cycle thus has period at most ℓ; long periods arise only by combining cycles (Figure 1).

Definition 7 (Dense block-cyclic architecture)**.** *Let m and d be positive integers. A* block-cyclic architecture with parameters d and m, *denoted* $\mathrm{BCA}(d, m)$, *is an asymmetric Hopfield network constructed as follows.*

1. *Let A be a multiset of $z = \lceil 4m \rceil$ integers chosen independently uniformly at random from the set $\{1, 2, \ldots, m\}$.*
2. *For each $\ell \in A$, create a* block cycle *of length ℓ of blocks of size d.*
3. *Let $N_{\mathrm{blocks}} = \sum_{\ell \in A} \ell$ be the total number of blocks, and denote by $n = dN_{\mathrm{blocks}}$ the total number of neurons in the* block-cyclic architecture.

The total number of blocks is $N_{\mathrm{blocks}} = \Theta(m^2)$ with high probability, and hence the total number of neurons in $\mathrm{BCA}(d, m)$ is $n = dN_{\mathrm{blocks}} = \Theta\left(m^2 d\right)$ with high probability, as shown in Lemma 2 in

Appendix D. Also note that, after removing repeated samples from A, the resulting set contains an independently sampled subset of $\{1, \dots, m\}$ with constant inclusion probability with high probability; see Lemma 3 in Appendix D. We now state our main result.

Theorem 1. *Let $m \in \mathbb{Z}^+$ be sufficiently large. Fix a sufficiently large constant K, set $d = \lceil K(\log m)^2 \rceil$. Let n denote the number of neurons in $\mathrm{BCA}(d, m)$. Then there exist constants $c_1, c_2 > 0$ such that, with probability $1 - o(1)$, $\mathrm{BCA}(d, m)$ contains a family $\mathcal{A}$ of limit-cycle attractors satisfying the following:*

1. *(**Number of limit cycles**) Writing $N_{\mathrm{cycles}} := |\mathcal{A}|$, we have*
$$N_{\mathrm{cycles}} \geq \exp\Big(c_1 \frac{n}{(\log n)^2}\Big). \tag{1}$$
2. *(**Periods of the constructed attractors**) Every attractor $\mathcal{C} \in \mathcal{A}$ has period at least*
$$\exp\left(c_2 \frac{\sqrt{n}}{\log n}\right). \tag{2}$$
3. *(**Noise robustness**) For n large enough, fix any $\mathcal{C} \in \mathcal{A}$ and any state $x \in \mathcal{C}$. If each neuron of x is independently flipped with probability p_n at time $t = 0$, where $p_n \leq \frac{1}{2} - \sqrt{\log n/d}$, then with probability $1 - o(1)$ the perturbed trajectory returns to $\mathcal{C}$.*

Remark 2. The choice $d = \lceil K(\log m)^2 \rceil$ is made for concreteness. More generally, the same proof shows that if the block size $d = d(n)$ satisfies $d = \omega(\log n)$, then the dense construction remains noise robust for flip probabilities $p_n \leq \frac{1}{2} - \sqrt{\log n/d} = \frac{1}{2} - o(1)$, while the bounds in (1) and (2) become $N_{\mathrm{cycles}} \geq \exp(\Omega(n/d))$ and $P \geq \exp(\Omega(\sqrt{n/d}))$. Appendix G.4 evaluates this trade-off exactly at finite size, together with the spread of the available periods.

We next provide a proof sketch of Theorem 1. A rigorous proof thereof is deferred to Appendix E.

Proof Sketch. *1.* Consider a block cycle in the architecture. Assign signs to the neurons such that neurons in the same block have all the same sign and such that the resulting binary sequence of blocks is aperiodic. Such a configuration forms a limit cycle, where blocks simply transfer their state to the succeeding one, making a rotating loop. The desired lower bound (1) then follows by bounding the number of such configurations using results of Appendix A.1.

2. For every attractor in the constructed family, the full block-cyclic architecture evolves as the product of rotations on its block cycles. Since the binary string on each block cycle is aperiodic, the period equals the lcm of the block-cycle lengths, which can be bounded as in (2) via Appendix A.2.

3. The choice of parameters guarantees that if each neuron is independently flipped with probability at most $\frac{1}{2} - \sqrt{\log n/d}$, then the majority of neurons within each block is preserved with probability $1 - o(1)$ by Hoeffding's bound, meaning that random perturbations are corrected at the next update. □

Remark 3 (Worst-case robustness). The noise robustness of Theorem 1 has a deterministic counterpart. Since every neuron takes the strict majority sign of the preceding block, any state that differs from a state of some $\mathcal{C} \in \mathcal{A}$ in fewer than $d/2$ neurons of every block is mapped in one synchronous update to the next state of $\mathcal{C}$ (Step 1 of the proof in Appendix E). Only the number of flipped neurons per block matters, not which ones, so this covers arbitrary perturbations. It is specific to the dense case; with h-sparse connections correction takes $O(\log d)$ rounds (Theorems 2 and 3).

Remark 4. Appendix E also records a global, fast convergence statement for arbitrary initial configurations. For the dense construction of Theorem 1, every trajectory reaches some limit cycle after at most one turn around the longest block cycle; sparse and adversarial analogues are in Theorem 8.

Remark 5. Even if the block-cycle lengths are chosen more carefully so as to be pairwise co-prime, the period does not improve drastically. For justification, see Remark 9 in Appendix E.2.

5 Sparsity, Randomness and Biological Plausibility

Two empirical observations motivate what follows. First, asymmetry itself is grounded in Dale's law, the observation that the synapses of a neuron are either all excitatory or all inhibitory. The

synaptic matrix then has columns of uniform sign, so once both neuron types are present it cannot be symmetric, an independent reason to study asymmetric networks. Second, van Es et al. (2025) report that large-scale cortical functional networks are visited in a robust cyclical order, on a timescale of 300 to 1000 ms and reproduced across five magnetoencephalography datasets. This concerns the cyclical ordering of *activity*, not block-cyclic anatomical wiring, so it is not direct evidence for the topology we assume; what it does show is that robust cyclic traversal of a repertoire of states is a real feature of cortical dynamics, which is the phenomenon our limit cycles model.

In this section, we present several results that highlight the biological plausibility of our theory. The adjacent blocks in the block cycles of Theorem 1 are complete bipartite graphs. We first show that this assumption can be relaxed: instead of all-to-all connectivity between adjacent blocks, each neuron may receive only h incoming edges from the preceding block, for any odd $h \geq 3$.

Definition 8 (h-sparse connection)**.** *Let S_1 and S_2 be two blocks. We say that S_1 is* h-sparsely connected *to S_2 if every neuron in S_2 is connected to exactly h neurons in S_1 chosen uniformly at random with replacement, with all edges directed from S_1 to S_2, and each edge assigned weight equal to its multiplicity, i.e., if a neuron in S_1 is chosen v times, the edge has weight v.*

Definition 9 (Sparse block-cyclic architecture)**.** *Let $m, b, h, d \in \mathbb{N}$ be fixed parameters. Suppose $3 \leq h < d$ is an odd integer and $b < m$. An* h-sparse block-cyclic architecture with parameters h, d, b, m, *denoted h-sparse* $\mathrm{BCA}(d, b, m)$, *is an asymmetric Hopfield network constructed as follows.*

1. *Let A be a multiset of $z = \lceil 4(m - b + 1) \rceil$ integers chosen independently uniformly at random from the set $\{b, b+1, \ldots, m\}$.*
2. *For each $\ell \in A$, create a* block cycle *of length ℓ of blocks of size d where adjacent blocks in these block cycles are h-sparsely connected.*
3. *Let $N_{\text{blocks}} = \sum_{\ell \in A} \ell$ be the total number of blocks, and denote by $n = dN_{\text{blocks}}$ the total number of neurons in the* sparse block-cyclic architecture.

We now state the h-sparse version of the result. To prove this result, we have used the classic theory on consensus reaching in opinion dynamics, which is reviewed in Appendix F.

Theorem 2. *Let $\varepsilon \in (0, 1)$ satisfy $\varepsilon(1 + \xi) > 1$, let $h \in \mathbb{N}$ be odd with $h \geq 3$, and let $m \in \mathbb{N}$ be sufficiently large.*[2] *Set $d = \lceil m^{2\varepsilon/(1-\varepsilon)} \rceil$ and choose $b = b(m)$ such that $b \geq c_3 \log m$ for a sufficiently large constant c_3 and $b/m \to 0$. Let n denote the number of neurons in the resulting h-sparse $\mathrm{BCA}(d, b, m)$. Then there exist constants $c_1, c_2 > 0$ such that, with probability $1 - o(1)$, the resulting h-sparse block-cyclic architecture contains a family $\mathcal{A}$ of limit-cycle attractors satisfying the following:*

1. *(**Superpolynomial number of limit cycles**) Writing $N_{\text{cycles}} := |\mathcal{A}|$, we have $N_{\text{cycles}} \geq \exp\big(c_1 n^{1-\varepsilon}\big)$.*
2. *(**Periods of the constructed attractors**) Every attractor $\mathcal{C} \in \mathcal{A}$ has period at least $\exp\big(c_2 (n^{1-\varepsilon})^{\frac{1}{2}}\big)$.*
3. *(**Noise robustness**) For n large enough, fix any $\mathcal{C} \in \mathcal{A}$ and any state $x \in \mathcal{C}$. If each neuron of x is independently flipped at time $t = 0$ with probability p, where $p \leq \frac{1}{2} - C_p \sqrt{\log n / d}$ for a sufficiently large constant C_p, then with probability $1 - o(1)$ (as $n \to \infty$) the perturbed trajectory returns to $\mathcal{C}$.*

Proof Sketch. 1. Item 1 follows from the counting argument for Item 1 of Theorem 1, with $d = \lceil K \log^2 m \rceil$ replaced by $d = \lceil m^{\frac{2\varepsilon}{1-\varepsilon}} \rceil$. The same quotient by the global orbit period appears, but its logarithm is $O(m) = o(n/d)$, so it does not affect the stated lower bound. Let $\mathcal{A}$ denote the corresponding family of attractors obtained from the aperiodic assignments.

2. Claim 2 is proved by applying Lemma 3 on the interval $\{b, \ldots, m\}$ and then using Theorem 5.

3. Because each neuron receives exactly h inputs from the preceding block and h is odd, the Hopfield update on a block coincides with synchronous h-majority dynamics. If each neuron is flipped

[2] ξ denotes the exponent in the h-majority recovery bound under the bias induced by the noise level p: starting from that bias, the process returns to the correct consensus in $O(\log d)$ rounds with probability $1 - 1/d^{\xi}$. See Appendix F for details.

with probability $p \le \frac{1}{2} - C_p\sqrt{\log n/d}$, then each neuron remains unflipped with probability at least $\frac{1}{2} + C_p\sqrt{\log n/d}$; a Hoeffding bound therefore shows that every block still has bias at least $\lambda_1\sqrt{d\log d}$ toward its pre-noise state with probability $1 - o(1)$. The large-bias recovery estimate for h-majority reviewed in Appendix F then implies that each block returns to its original state within $O(\log d)$ rounds with failure probability at most $d^{-\xi}$, where ξ is the exponent furnished by that recovery bound for this noise level. A union bound over all blocks yields overall recovery with probability $1 - o(1)$. See Appendix E for the formal proof. □

We now state our result with added adversarial connections. In this result, we use the weak notion of limit cycles (Definition 4). Appendix G.3 includes spectral embeddings of representative sparse block-cyclic architectures with adversarial connections. Let $\mathcal{A}$ be the class of limit cycles on the sparse block-cyclic architecture obtained from the aperiodic monochromatic assignments used in Theorem 2. We show that if a sufficiently small multiple of $\sqrt{d}/\log d$ neurons in each block receive adversarial incoming edges from elsewhere in the block-cyclic architecture, and each neuron is independently flipped with probability $p \in (0, 1/2)$, then the trajectory tracks the original cycle within block-max distance $O(\sqrt{d}/\log d)$ with high probability.

Definition 10 (q-adversarial block-cyclic architecture)**.** *Let $\mathcal{H}$ be an h-sparse* BCA (d, b, m) *and let $q \in \mathbb{N}$. A* q-adversarial version *of $\mathcal{H}$ is obtained by designating, in every block, an arbitrary set of at most q neurons as* adversarial*; the remaining neurons are* regular*. Regular neurons keep exactly the incoming edges prescribed by $\mathcal{H}$ and update by the synchronous majority rule on them. An adversarial neuron may in addition receive incoming edges from an arbitrary set of neurons of the architecture, of arbitrary cardinality, possibly lying in other blocks or in other block cycles, subject only to consistency with the direction of the construction; its state may evolve arbitrarily, and adaptively in the whole history of the process. Every neuron, adversarial or regular, keeps the outgoing edges prescribed by $\mathcal{H}$.*

Remark 6. Only the number of neurons receiving adversarial edges is limited, to q per block; the number of those edges is not. Since an adversarial state may depend on the entire history, the adversary is in particular free to know the construction. What must be arranged in advance is therefore only the h-sparse backbone, that is hn edges in total for $h = 3$, and the guarantee below holds whatever is added on top of it at q neurons per block.

Theorem 3 (Robustness under adversarial perturbations)**.** *Let $\varepsilon \in (0, 1)$ satisfy $\varepsilon(1 + \xi) > 1$, where $\xi > 0$ is the polynomial-decay exponent in the per-block recovery bound from Appendix F, and let $m \in \mathbb{N}$ be sufficiently large. Set $d = \lceil m^{2\varepsilon/(1-\varepsilon)} \rceil$ and choose $b = b(m)$ such that $b \ge c_4(\log m)^2$ for a sufficiently large constant c_4 and $b/m \to 0$. Let $\mathcal{H}$ be a 3-sparse* BCA (d, b, m) *with n neurons, and let $\mathcal{A}$ denote the set of its limit-cycle states obtained from aperiodic monochromatic assignments. Fix any initial state $c \in \mathcal{A}$. Choose a sufficiently small constant $\gamma > 0$ and a sufficiently large constant $C_{\mathrm{adv}} > 0$, and set*

$$q = \left\lfloor \gamma \frac{\sqrt{d}}{\log d} \right\rfloor, \qquad k = \left\lceil C_{\mathrm{adv}} \frac{\sqrt{d}}{\log d} \right\rceil.$$

Let $\mathcal{H}'$ be any q-adversarial version of $\mathcal{H}$ in the sense of Definition 10. If, at time $t = 0$, each neuron of $\mathcal{H}'$ is independently flipped with probability $p \in (0, 1/2)$ from the initial state c, then with probability $1 - o(1)$ as $n \to \infty$, the resulting trajectory returns to the (k, d)-weak limit cycle associated with the original monochromatic orbit through c, for the canonical block partition.

Proof Sketch. Theorem 11 asserts that the majority process converges even in the presence of a sufficiently small multiple of $\sqrt{d}/\log d$ adversarial agents, and that the resulting almost-consensus has at most a constant multiple of $\sqrt{d}/\log d$ exceptions. Reusing the same union bound over the n/d blocks as in Theorem 2, the condition $\varepsilon(1 + \xi) > 1$ upgrades this per-block high-probability estimate to an overall $1 - o(1)$ bound. See Appendix E for details. □

Remark 7. Appendix E also contains global convergence statements for arbitrary initial configurations. In the dense construction of Theorem 1, every trajectory reaches a limit cycle deterministically after at most one full turn around the longest block cycle. Under the hypotheses of Theorem 2, any fixed initial configuration quickly reaches some limit cycle with high probability. Under the adversarial perturbations of Theorem 3, the same argument yields fast convergence to a (k, d)-weak limit cycle; see Theorem 8.

6 EXPERIMENTS

In this section, we present finite-size simulations illustrating the recovery mechanisms behind our theoretical results. We begin with a single sparse block cycle of length $\ell = 200$ with $h = 3$ inputs per neuron. The initial state is assigned as in Step 1 of our construction (Theorem 2), all block signs being taken equal to 1, so that the experiment isolates the denoising mechanism rather than the period-length effect of nonconstant block sequences. Each neuron is then flipped independently with probability p and the dynamics are run to see whether the initial state is recovered before a full cycle completes. Over 50 trials this estimates the revival probability for varying d and p; see Figure 2(a).

Next, we consider a 3-sparse block-cyclic architecture with two block cycles of lengths $\ell_1 = 150$ and $\ell_2 = 200$ and block size $d = 10000$, adding random extra incoming connections to k neurons per block from elsewhere in the network. These are not adversarial: they are sampled randomly, and their targets are ignored in the weak recovery criterion. Increasing k therefore need not hurt revival monotonically, and at high noise the extra inputs can even help, by supplying further samples from a population already biased toward the correct sign, as in the highest-noise curve of Figure 2(b). The protocol is otherwise that of the first experiment. Appendix G.2 reports the implementation details, together with a stress test motivated by Theorem 3 in which genuinely adversarial nodes update against the majority rule and the initial state still revives, and an experiment for Theorem 8 showing that a randomly initialized block cycle converges to a blockwise monochromatic configuration.

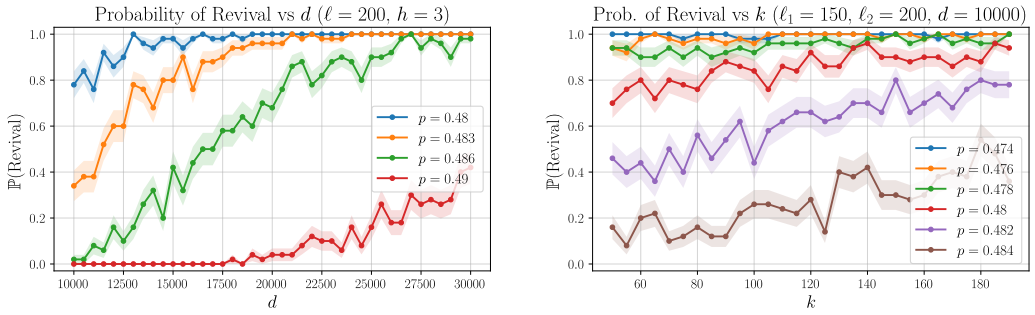

(a) (b)

Figure 2: Revival probability $\pm$ standard error over 50 trials. (a) Single sparse block cycle, $\ell = 200$ and $h = 3$: the x-axis varies the block size d, and curves correspond to the initial flip probability p. (b) Two block cycles of lengths 150 and 200, with $d = 10000$ and $h = 3$: the x-axis varies the number k of neurons per block receiving random extra incoming connections, curves again corresponding to p.

7 CONCLUSION

We showed that classical synchronous asymmetric Hopfield networks can support large, robust sequence memory without nonlinear interactions or continuous states: n neurons can store $\exp\big(\Omega(n/(\log n)^2)\big)$ distinct limit-cycle attractors of period $\exp\big(\Omega(\sqrt{n}/\log n)\big)$, robust to independent flips up to $\frac{1}{2} - o(1)$ and to sparse and adversarial perturbations. Block-cyclic motifs thus offer a route to sequence memory through coarse structure rather than fine weight tuning. Natural next steps are learning such motifs and storing deliberately correlated memories.

Limitations. Our results establish capacity, not learning. The topology is prescribed in advance, we give no local rule that forms it from unstructured connectivity, and the construction does not store arbitrary prescribed sequences. In the dense case the construction is exactly a disjoint product of directed cyclic shift registers, each block acting as a repetition code and the period being the lcm of the register lengths; this reading stops being exact under sparsification to in-degree 3 and under adversarial connections and neurons (Section 5), which is where the substance of our results lies. Our guarantees also assume synchronous Little-style updates; whether they persist under asynchronous single-neuron dynamics remains open. The experiments are synthetic simulations of the theoretical mechanisms, not evaluations on biological recordings or task-level benchmarks.

Statement on the Use of Large Language Models

Large language models were used only as writing aids, to polish the wording and grammar of the text. They were not used for research ideation, proofs, experiments, or analysis.

References

U. Bastolla and G. Parisi. Attractors in fully asymmetric neural networks. *Journal of Physics A: Mathematical and General*, 30(16):5613, August 1997. ISSN 0305-4470. doi: 10.1088/0305-4470/30/16/007.

L. Becchetti, A. Clementi, E. Natale, F. Pasquale, and L. Trevisan. Stabilizing consensus with many opinions. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 620–635. Society for Industrial and Applied Mathematics, December 2016. doi: 10.1137/1.9781611974331.ch46. URL http://dx.doi.org/10.1137/1.9781611974331.ch46.

Petra Berenbrink, Amin Coja-Oghlan, Oliver Gebhard, Max Hahn-Klimroth, Dominik Kaaser, and Malin Rau. On the hierarchy of distributed majority protocols. In Eshcar Hillel, Roberto Palmieri, and Etienne Rivière (eds.), *26th International Conference on Principles of Distributed Systems (OPODIS 2022)*, volume 253 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pp. 23:1–23:19, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi: 10.4230/LIPICS.OPODIS.2022.23. URL https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.OPODIS.2022.23.

Sarthak Chandra, Sugandha Sharma, Rishidev Chaudhuri, and Ila Fiete. Episodic and associative memory from spatial scaffolds in the hippocampus. *Nature*, 638:739–751, 2025. doi: 10.1038/s41586-024-08392-y.

Hamza Tahir Chaudhry, Jacob A Zavatone-Veth, Dmitry Krotov, and Cengiz Pehlevan. Long sequence hopfield memory. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=Tz2uONpgpy.

Javier Cilleruelo, Juanjo Rué, Paulius Šarka, and Ana Zumalacárregui. The least common multiple of random sets of positive integers. *Journal of Number Theory*, 144:92–104, November 2014. ISSN 0022-314X. doi: 10.1016/j.jnt.2014.04.011. URL http://dx.doi.org/10.1016/j.jnt.2014.04.011.

Andrea Clementi, Mohsen Ghaffari, Luciano Gualà, Emanuele Natale, Francesco Pasquale, and Giacomo Scornavacca. A tight analysis of the parallel undecided-state dynamics with two colors. In *International Symposium on Mathematical Foundations of Computer Science*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi: 10.4230/LIPICS.MFCS.2018.28. URL https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.MFCS.2018.28.

Bernard Derrida, Elizabeth Gardner, and Annette Zippelius. An exactly solvable asymmetric neural network model. *Europhysics Letters*, 4(2):167–173, 1987. doi: 10.1209/0295-5075/4/2/007.

Diego Fasoli, Anna Cattani, and Stefano Panzeri. Pattern Storage, Bifurcations, and Groupwise Correlation Structure of an Exactly Solvable Asymmetric Neural Network Model. *Neural Computation*, 30(5):1258–1295, May 2018. ISSN 0899-7667. doi: 10.1162/neco_a_01069.

Maxwell Gillett, Ulises Pereira, and Nicolas Brunel. Characteristics of sequential activity in networks with temporally asymmetric Hebbian learning. *Proceedings of the National Academy of Sciences*, 117(47):29948–29958, November 2020. doi: 10.1073/pnas.1918674117.

Parikshit Gopalan, Cheng Huang, Huseyin Simitci, and Sergey Yekhanin. On the locality of codeword symbols. *IEEE Transactions on Information Theory*, 58(11):6925–6934, November 2012. doi: 10.1109/TIT.2012.2208937.

Taosha Guo, Arie Ogranovich, Arvind R. Venkatakrishnan, Madelyn R. Shapiro, Francesco Bullo, and Fabio Pasqualetti. Oscillatory associative memory with exponential capacity. In *IEEE Conference on Decision and Control (CDC)*, 2025. Preprint arXiv:2504.03102.

J J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558, April 1982. ISSN 1091-6490. doi: 10.1073/pnas.79.8.2554. URL http://dx.doi.org/10.1073/pnas.79.8.2554.

Sungmin Hwang, Viola Folli, Enrico Lanza, Giorgio Parisi, Giancarlo Ruocco, and Francesco Zamponi. On the number of limit cycles in asymmetric neural networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(5):053402, May 2019. ISSN 1742-5468. doi: 10.1088/1742-5468/ab11e3.

Sungmin Hwang, Enrico Lanza, Giorgio Parisi, Jacopo Rocchi, Giancarlo Ruocco, and Francesco Zamponi. On the number of limit cycles in diluted neural networks. *Journal of Statistical Physics*, 181(6):2304–2321, December 2020. ISSN 0022-4715, 1572-9613. doi: 10.1007/s10955-020-02664-3.

Dmitry Krotov and John J. Hopfield. Large associative memory problem in neurobiology and machine learning. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=X4y_10OX-hX.

W. A. Little. The existence of persistent states in the brain. *Mathematical Biosciences*, 19(1–2):101–120, February 1974. ISSN 0025-5564. doi: 10.1016/0025-5564(74)90031-5. URL https://doi.org/10.1016/0025-5564(74)90031-5.

Samuel P. Muscinelli, Wulfram Gerstner, and Johanni Brea. Exponentially long orbits in hopfield neural networks. *Neural Computation*, 29(2):458–484, February 2017. ISSN 1530-888X. doi: 10.1162/neco_a_00919. URL http://dx.doi.org/10.1162/NECO_a_00919.

Arie Ogranovich, Taosha Guo, Arvind R. Venkatakrishnan, Madelyn Shapiro, Francesco Bullo, and Fabio Pasqualetti. Oscillator-based associative memory with exponential capacity: Theory, algorithms, and hardware implementation. *arXiv preprint arXiv:2604.01469*, 2026.

L. Personnaz, I. Guyon, and G. Dreyfus. Collective computational properties of neural networks: New learning mechanisms. *Physical Review A*, 34(5):4217–4228, November 1986. doi: 10.1103/PhysRevA.34.4217.

Hubert Ramsauer, Bernhard Schäfl, Johannes Lehner, Philipp Seidl, Michael Widrich, Lukas Gruber, Markus Holzleitner, Thomas Adler, David Kreil, Michael K Kopp, Günter Klambauer, Johannes Brandstetter, and Sepp Hochreiter. Hopfield networks is all you need. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=tL89RnzIiCd.

Mohadeseh Shafiei Kafraj, Dmitry Krotov, and Peter E. Latham. A biologically plausible dense associative memory with exponential capacity. In *International Conference on Learning Representations (ICLR)*, 2026.

Mats WJ van Es, Cameron Higgins, Chetan Gohil, Andrew J Quinn, Diego Vidaurre, and Mark W Woolrich. Large-scale cortical functional networks are organized in structured cycles. *Nature Neuroscience*, 28(10):2118–2128, 2025.

Chuan Zhang, Gerhard Dangelmayr, and Iuliana Oprea. Storing cycles in hopfield-type networks with pseudoinverse learning rule: Admissibility and network topology. *Neural Networks*, 46:283–298, October 2013. ISSN 0893-6080. doi: 10.1016/j.neunet.2013.06.008. URL http://dx.doi.org/10.1016/j.neunet.2013.06.008.

A Analytical Tools

In this section, we present several preliminary results that will be used in the proof of our main theorem. In Subsection A.1, we recall results concerning the counting of aperiodic binary sequences, while in Subsection A.2, we review results related to the growth of the least common multiple (lcm) of random integers; further details can be found in Appendices B and C, respectively.

A.1 Aperiodic Binary Sequence Counting

We briefly recall standard asymptotic estimates related to the counting of aperiodic binary sequences, that will be used later in the argument. Binary sequences arise naturally in several combinatorial problems and the distinction between periodic and aperiodic necklaces plays a key role in enumeration.

Definition 11 (Aperiodic binary sequence)**.** *Let $\ell \geq 1$ and let $x \in \{-1,+1\}^\ell$. For $k \in \{0,1,\ldots,\ell-1\}$, define the cyclic shift $\sigma^k(x)$ by $\left(\sigma^k(x)\right)_i = x_{(i+k) \bmod \ell}$ for $i = 0,1,\ldots,\ell-1$. The sequence x is said to be* aperiodic *if $\sigma^k(x) \neq x$ for all $1 \leq k < \ell$, equivalently, if all its cyclic shifts $\{\sigma^k(x) : 0 \leq k < \ell\}$ are pairwise distinct.*

We now state an asymptotic estimate for the number of aperiodic binary sequences.

Lemma 1. *The number $\mathcal{M}_2(\ell)$ of aperiodic binary sequences of length ℓ is given by $\mathcal{M}_2(\ell) = 2^\ell\,(1+o(1))$ as $\ell \to \infty$.*

For a proof of Lemma 1, see Appendix B.

Remark 8. A randomly chosen binary sequence of length ℓ is aperiodic with probability tending to 1 as $\ell \to \infty$.

A.2 LCM of Random Sets of Positive Integers

The function $\psi(m) = \log \operatorname{lcm}\{1 \leq a \leq m\}$ was introduced by Chebyshev in his study of the distribution of prime numbers. A well-known consequence of the Prime Number Theorem is that $\psi(m) \sim m$ as $m \to \infty$, where by $a(m) \sim b(m)$ we mean that $\lim_{m\to\infty} a(m)/b(m) = 1$.

Cilleruelo et al. (2014) extended this analysis to random subsets of $\{1,2,\ldots,m\}$ as follows. For a subset $A \subseteq \{1,2,\ldots,m\}$, define $\psi(A) = \log \operatorname{lcm}\{a : a \in A\}$. Then, the following result holds for random subsets of $\{1,2,\ldots,m\}$.

Theorem 4. *Let A be the set constructed by choosing each element from $\{1,2,\ldots,m\}$ independently with probability $\delta = \delta(m)$. If $\delta = \delta(m) < 1$ and $\delta \cdot m \to \infty$, then*

$$\psi(A) \sim m\frac{\delta \log(\delta^{-1})}{1-\delta} \tag{3}$$

asymptotically almost surely as $m \to \infty$.

Using the analysis of Cilleruelo et al. (2014), we show that the same asymptotic result holds when each element is chosen independently from $\{b,\ldots,m\}$ with probability δ, where $b < m$ satisfies $b/m = o(1)$.

Theorem 5. *Let A be the set constructed by choosing each element from $\{b,\ldots,m\}$ independently with a constant probability δ, for some integer $b = b(m)$ satisfying $b/m = o(1)$. Then, (3) holds asymptotically almost surely as $m \to \infty$.*

In particular, in both cases, $\log \operatorname{lcm}(A)$ grows linearly in m asymptotically almost surely. For further details, see Appendix C. Note that the above results imply that in both cases, $\operatorname{lcm}(A) = \exp(\Omega(m))$.

B Aperiodic Binary Sequences

Definition 12 (Möbius function)**.** *The* Möbius function $\mu : \mathbb{N} \to \{-1,0,1\}$ *is defined by*

$$\mu(n) \;=\; \begin{cases} 1 & \text{if } n = 1, \\ 0 & \text{if } n \text{ is divisible by the square of a prime,} \\ (-1)^r & \text{if } n \text{ is the product of } r \text{ distinct primes.} \end{cases}$$

Theorem 6. *Let $k \geq 1$ and $\ell \geq 1$ be integers. The number $\mathcal{M}_k(\ell)$ of aperiodic k-ary sequences of length ℓ is*

$$\mathcal{M}_k(\ell) \;=\; \sum_{j \mid \ell} \mu(j)\, k^{\,\ell/j},$$

where μ denotes the Möbius function.

Proof of Lemma 1. By Theorem 6,

$$\mathcal{M}_2(\ell) = \sum_{j|\ell} \mu(j)\, 2^{\ell/j} = 2^\ell + \sum_{\substack{j|\ell \\ j>1}} \mu(j)\, 2^{\ell/j}.$$

For every proper divisor $j > 1$ of ℓ, we have $\ell/j \le \ell/2$, so

$$\left| \sum_{\substack{j|\ell \\ j>1}} \mu(j)\, 2^{\ell/j} \right| \le \sum_{\substack{j|\ell \\ j>1}} 2^{\ell/j} \le \tau(\ell)\, 2^{\ell/2},$$

where $\tau(\ell)$ is the number of divisors of ℓ. Hence

$$\left|\mathcal{M}_2(\ell) - 2^\ell\right| \le \tau(\ell)\, 2^{\ell/2} \implies \mathcal{M}_2(\ell) = 2^\ell(1 + o(1)),$$

since $\tau(\ell) = o(2^{\ell/2})$. □

C LCM OF RANDOM SUBSETS

In this section, we prove Theorem 5. We restate it for the convenience of the reader.

Theorem 7. *Let A be the set constructed by choosing each element from $\{b, \dots, m\}$ independently with a constant probability δ, for some integer $b = b(m)$ satisfying $b/m = o(1)$. Then*

$$\psi(A) \sim m\frac{\delta \log(\delta^{-1})}{1-\delta},$$

asymptotically almost surely as $m \to \infty$.

Proof. Let $R_m \subseteq \{1, 2, \dots, m\}$ be the random set obtained by including each integer independently with the same fixed probability δ. Write

$$A_m := R_m \cap \{b, b+1, \dots, m\} \qquad \text{and} \qquad D_m := R_m \cap \{1, 2, \dots, b-1\}.$$

Then A_m has the distribution appearing in the statement of the theorem, and $R_m = A_m \cup D_m$ is the corresponding random subset of $\{1, \dots, m\}$.

Using the elementary inequality

$$\operatorname{lcm}(S \cup T) \le \operatorname{lcm}(S) \operatorname{lcm}(T) \qquad (S, T \subset \mathbb{N} \text{ finite}),$$

we obtain

$$\psi(R_m) = \log \operatorname{lcm}(A_m \cup D_m) \le \psi(A_m) + \psi(D_m).$$

Since $D_m \subseteq \{1, \dots, b-1\}$, we have

$$\psi(D_m) \le \psi(\{1, \dots, b-1\}) = \log \operatorname{lcm}(1, 2, \dots, b-1) \sim b,$$

and therefore $\psi(D_m) = o(m)$ because $b/m \to 0$.

Consequently,

$$\psi(A_m) \ge \psi(R_m) - o(m).$$

On the other hand, since $A_m \subseteq R_m$, we trivially have

$$\psi(A_m) \le \psi(R_m).$$

Hence

$$\psi(A_m) = \psi(R_m) + o(m).$$

Now Theorem 4 applied to R_m gives

$$\psi(R_m) \sim m\,\frac{\delta \log(\delta^{-1})}{1-\delta} \qquad \text{a.a.s.}$$

Combining this with the previous display yields

$$\psi(A_m) \sim m\,\frac{\delta \log(\delta^{-1})}{1-\delta} \qquad \text{a.a.s.}$$

as claimed. □

D SAMPLING OF BLOCK CYCLE LENGTHS

Lemma 2. *Let $M \in \mathbb{N}$ and let $z = z(M)$ satisfy $z \to \infty$. Let $\{Y_i\}_{i=1}^{z}$ be a sequence of i.i.d. random variables, each uniformly distributed over $\{1, 2, \ldots, M\}$, and define*

$$S \triangleq \sum_{i=1}^{z} Y_i.$$

Then, with probability at least $1 - \exp\big(-\Theta(z)\big)$,

$$S = \Theta(zM).$$

Proof. The mean of S equals

$$\mu \triangleq \mathbb{E}\left[S\right] = \mathbb{E}\left[\sum_{i=1}^{z} Y_i\right] = \sum_{i=1}^{z} \mathbb{E}\left[Y_i\right] = z\frac{M+1}{2}.$$

Since $\mathbb{E}\left[Y_i\right] = (M+1)/2$, we have $\mu = \Theta(zM)$.

Let $\varepsilon > 0$. Since each $Y_i \in \{1, 2, \ldots, M\}$, the summands in S are bounded in $[1, M]$. Hence, by Hoeffding's inequality,

$$\begin{aligned}\mathbb{P}\left(|S - \mu| \geq \varepsilon\mu\right) &\leq 2\exp\left(\frac{-2\mu^2\varepsilon^2}{\sum_{i=1}^{z}(M-1)^2}\right) \\ &= 2\exp\left(\frac{-2\mu^2\varepsilon^2}{z(M-1)^2}\right).\end{aligned}$$

Using $\mu = \Theta(zM)$, the exponent is $-\Theta(z)$. Therefore,

$$\mathbb{P}\left(|S - \mu| \geq \varepsilon\mu\right) = \exp\big(-\Theta(z)\big).$$

Consequently, with probability at least $1 - \exp\big(-\Theta(z)\big)$,

$$S = (1 \pm \varepsilon)\mu = \Theta(zM).$$

□

Corollary 1. *Let $\{Y_i\}_{i=1}^{z}$ be i.i.d. uniform on $\{b, \ldots, m\}$, and assume $b/m = o(1)$. Let*

$$M \triangleq m - b + 1, \qquad z \triangleq \lceil 4M \rceil.$$

Then, with probability at least $1 - \exp\big(-\Theta(M)\big)$,

$$\sum_{i=1}^{z} Y_i = \Theta(m^2).$$

Proof. Let

$$X_i \triangleq Y_i - b + 1, \qquad i = 1, \ldots, z.$$

Then $\{X_i\}_{i=1}^{z}$ are i.i.d. uniform on $\{1, 2, \ldots, M\}$.

Hence, by Lemma 2, with probability at least $1 - \exp\{\Theta(-M)\}$,

$$\sum_{i=1}^{z} X_i = \Theta(M^2).$$

Moreover,

$$\sum_{i=1}^{z} Y_i = \sum_{i=1}^{z} \{X_i + b - 1\} = \sum_{i=1}^{z} X_i + z(b-1).$$

Since $b/m = o(1)$ and $M = m - b + 1$, we have $b = o(M)$, while $z = \Theta(M)$. Therefore,

$$z(b-1) = \Theta(M) \cdot o(M) = o(M^2).$$

Consequently,

$$\sum_{i=1}^{z} Y_i = \Theta(M^2).$$

Finally, because $M = m - b + 1 = m(1 - o(1))$, we have $M = \Theta(m)$ and thus

$$\sum_{i=1}^{z} Y_i = \Theta(m^2).$$

□

Lemma 3. *Let I be an interval of integers with $|I| = M$, and let $S_{4M}^{(\mathrm{rep})}$ be the subset of elements of I observed at least once when taking $\lceil 4M \rceil$ samples from I uniformly at random with replacement. There is a coupling with a subset $S_{\delta}^{(\mathrm{ind})} \subseteq I$, obtained by including each element of I independently with probability $\delta = 1 - e^{-2}$, such that*

$$\mathbb{P}\Big(S_{\delta}^{(\mathrm{ind})} \subseteq S_{4M}^{(\mathrm{rep})}\Big) \geq 1 - \exp(-\Theta(M)).$$

Proof. For each $a \in I$, let X_a be an independent Poisson random variable with mean 2, and set

$$S_{\delta}^{(\mathrm{ind})} = \{a \in I : X_a \geq 1\}.$$

Then $S_{\delta}^{(\mathrm{ind})}$ includes each element independently with probability $\delta = 1 - e^{-2}$. Let $T = \sum_{a \in I} X_a$. Then T is Poisson with mean $2M$, and hence

$$\mathbb{P}(T > \lceil 4M \rceil) \leq \exp(-\Theta(M)).$$

Conditional on $T = t$, the vector $(X_a)_{a \in I}$ has the multinomial distribution obtained from t independent uniform samples from I. Therefore, on the event $T \leq \lceil 4M \rceil$, we can realize these Poisson counts as the first T samples in a sequence of $\lceil 4M \rceil$ independent uniform samples from I, and fill the remaining positions with independent uniform samples. The set of elements seen in the full sequence has the distribution of $S_{4M}^{(\mathrm{rep})}$ and contains $S_{\delta}^{(\mathrm{ind})}$ whenever $T \leq \lceil 4M \rceil$. This proves the claim. □

E Proof of Main Results

E.1 Proof of Theorem 1

Proof. We prove statements (1)–(3) for block-cyclic architectures as defined in the statement.

Step 1: Constructing limit cycles. Consider the following state of the block-cyclic architecture: arbitrary signs are given to all neurons with one constraint, namely that all neurons in a given block must have the same sign. We claim that such configurations lead to a limit-cycle attractor. To see this, note that all neurons in the same block receive identical inputs at every synchronous update because their incoming signals are identical (and all weights are $+1$). Hence, if every block is initially monochromatic (all d neurons share the same sign), then the next update again produces a monochromatic state for each block, which is just a one-step rotation of the previous state. The effective state of each block cycle is then a binary string of length ℓ_i, which is rotated synchronously at each time step. Therefore, the block-cyclic architecture exhibits a limit cycle, where each block cycle is an independent rotating binary string. Having shown that our configurations lead to a limit cycle, we now impose another constraint on the configurations: all binary strings on the block cycles are aperiodic, meaning no nontrivial rotation fixes them. In this case, the overall period of the block-cyclic architecture is given by the lcm of the individual periods, i.e., overall period $= \mathrm{lcm}\{\ell_i\}_{i=1}^{z}$, where $|A| = z$. Moreover, any state whose blockwise strict majority signs agree with one of these monochromatic cycle states is mapped in one synchronous update to the next state of the cycle. Since $d \geq 2$ for all sufficiently large m, this gives states outside the cycle in its basin, so the cycle is a limit-cycle attractor. Let $\mathcal{A}$ be the family of limit-cycle attractors obtained from these aperiodic assignments.

Step 2: Counting limit cycles (proof of (1)). We start by counting how many such distinct limit-cycle attractors are there. For a single block cycle of length ℓ, the number of aperiodic binary labelings equals the number $\mathcal{M}_2(\ell)$ of binary aperiodic sequences. By Lemma 1, we have

$$\mathcal{M}_2(\ell) = 2^\ell(1 + o(1)).$$

Since the block cycles are disjoint, there are $\prod_{i=1}^{z} \mathcal{M}_2(\ell_i)$ aperiodic monochromatic states of this form. These are states rather than global orbits: the same limit cycle is counted once for each point on its orbit. Since the block strings are aperiodic, every such global orbit has period $P = \mathrm{lcm}(\ell_1, \ldots, \ell_z)$. Therefore,

$$N_{\text{cycles}} := |\mathcal{A}| = \frac{1}{P} \prod_{i=1}^{z} \mathcal{M}_2(\ell_i).$$

Hence we have

$$N_{\text{cycles}} = \frac{2^{N_{\text{blocks}}}(1 + o(1))^z}{P}, \qquad N_{\text{blocks}} := \sum_{i=1}^{z} \ell_i.$$

Taking $\log$ on both sides and using $N_{\text{blocks}} = n/d$,

$$\log N_{\text{cycles}} \geq (\log 2)\frac{n}{d} + z \log(1 + o(1)) - \log P.$$

Since each $\ell_i \leq m$, we have $P \leq \mathrm{lcm}(1, \ldots, m)$, and hence $\log P \leq \psi(m) = O(m)$. By Lemma 2, we have $m = \Theta((n/d)^{\frac{1}{2}})$. Since $d = \lceil K(\log m)^2 \rceil$ and $n = \Theta(m^2 d)$, it follows that $\log m = \Theta(\log n)$ and hence $d = \Theta((\log n)^2)$. Also noting that $z = \Theta(m) = \Theta((n/d)^{\frac{1}{2}})$ and $N_{\text{blocks}} = n/d = \Theta(n/(\log n)^2)$, the dominant term is $(\ln 2)n/d$, and the latter two terms are negligible. Hence there exists $c_1 > 0$ such that

$$N_{\text{cycles}} \geq \exp\Big(c_1 \frac{n}{(\log n)^2}\Big),$$

proving (1).

Step 3: Lower bound on the period (proof of (2)). Let

$$P := \mathrm{lcm}(\ell_1, \ldots, \ell_z)$$

be the least common multiple of the block-cycle lengths. Recall that the multiset $\{\ell_i\}_{i=1}^{z}$ is chosen uniformly at random with replacement from $\{1, 2, \ldots, m\}$ with $z = \lceil 4m \rceil$. Consider the set Q formed by ignoring repeated elements in the multiset $\{\ell_i\}_{i=1}^{z}$. By Lemma 3, this construction contains, with probability $1 - \exp(-\Theta(m))$, a set that includes each integer in $\{1, \ldots, m\}$ independently with probability $\delta = 1 - e^{-2}$. Since adding elements to a set cannot decrease its lcm, Theorem 4 implies that, with high probability,

$$\log P \geq c'm$$

for some $c' > 0$ depending only on δ. With $m = \Theta((n/d)^{\frac{1}{2}})$ (Lemma 2), this yields

$$P \geq \exp\Big(c_2 \Big(\frac{n}{d}\Big)^{1/2}\Big) = \exp\Big(c_2 \frac{\sqrt{n}}{\log n}\Big),$$

proving (2).

Step 4: Noise robustness (proof of (3)). Fix any $\mathcal{C} \in \mathcal{A}$ and any state $x \in \mathcal{C}$. Note that due to the majority updating rule, a block only gives a different signal to its successor block in a block cycle if more than half of its neurons are flipped. Assume that each neuron is flipped independently with probability $p_n \leq \frac{1}{2} - \sqrt{\frac{\log n}{d}}$, and compute the probability that there exists no block with half or more of its neurons flipped. Let $p_n \leq \frac{1}{2} - \sqrt{\frac{\log n}{d}}$, and let $Y \sim \mathrm{Binomial}(d, p_n)$ be the number of flipped neurons in a given block. For n, and equivalently d large enough, a Hoeffding's bound gives

$$\mathbb{P}(Y \geq d/2) \leq \exp\left(-2\frac{(d/2 - dp_n)^2}{d}\right) \leq e^{-2\log n} = n^{-2},$$

There are $N_{\text{blocks}} = n/d$ blocks, so by a union bound the probability that any block has a majority flipped is at most

$$\frac{n}{d} \cdot n^{-2} = \frac{1}{dn} = o(1).$$

Thus, with probability $1 - o(1)$ no block changes its majority sign, and the next synchronous update maps the perturbed state exactly to $F(x) \in \mathcal{C}$. The network therefore returns to $\mathcal{C}$, establishing (3) and completing the proof. □

E.2 Co-prime block-cycle lengths

In our construction in Theorem 1, the block-cycle lengths $\{\ell_1, \ell_2, \ldots, \ell_z\}$ are chosen uniformly at random from $\{1, 2, \ldots, m\}$. In this subsection, we show that even if these lengths are chosen more carefully so as to be pairwise co-prime, we do not gain a huge improvement.

Remark 9. If the block-cycle lengths $\{\ell_1, \ell_2, \ldots, \ell_z\}$, with each $\ell_i \in [m]$, are chosen to be pairwise co-prime, then the lower bound on the maximum period improves from

$$\exp\left(\Omega\left(\frac{\sqrt{n}}{\log n}\right)\right) \quad \text{to} \quad \exp\left(\Omega\left(\sqrt{\frac{n}{\log n}}\right)\right).$$

To see this, observe that the largest possible size of such a set is at most $\pi(m) + 1$, where $\pi(m)$ denotes the number of primes up to m. By the Prime Number Theorem, the sum $S(m)$ of the elements in this set satisfies

$$S(m) \sim \frac{m^2}{\log m}.$$

Since the number of neurons is

$$n \sim \frac{m^2 d}{\log m},$$

it follows that

$$m \sim \sqrt{\frac{n}{d} \log\left(\frac{n}{d}\right)}.$$

Finally, since the logarithm of the maximum period P scales as $\log P \sim m$ and $d \sim \log^2 n$, we obtain

$$\log P \sim \sqrt{\frac{n}{d} \log\left(\frac{n}{d}\right)} \sim \sqrt{\frac{n}{\log n}}.$$

E.3 Proof of Theorem 2

Proof. To prove Theorem 2, we notice its similarity to Theorem 1. We will see that the first two parts follow by a very similar argument to Theorem 1. Then we prove the third part.

Step 1: Constructing limit cycles. The aperiodic monochromatic assignments used in the dense proof still define limit cycles for the sparse block-cyclic architecture, because the sparse edges transmit the same block value whenever the preceding block is monochromatic. We assume aperiodicity here as well, and let $\mathcal{A}$ denote the corresponding family of attractors. By Corollary 1, with probability $1-o(1)$ the total number of blocks satisfies $N_{\text{blocks}} = \Theta(m^2)$. Since $d = \lceil m^{2\varepsilon/(1-\varepsilon)} \rceil$ and $n = dN_{\text{blocks}}$, on this event we have $d = \Theta(n^{\varepsilon})$ and $n/d = \Theta(n^{1-\varepsilon})$. We work on this event throughout the proof.

Step 2: Counting limit cycles (proof of (1)). As in Step 2 of the proof of Theorem 1, the count of aperiodic monochromatic states must be divided by the global orbit period $P = \operatorname{lcm}(\ell_1, \ldots, \ell_z)$. Since $P \le \operatorname{lcm}(1, \ldots, m)$, we have $\log P = O(m) = o(n/d)$. Thus the logarithm of $N_{\text{cycles}} := |\mathcal{A}|$ is still dominated by n/d. Since here we have $d = \Theta(n^{\varepsilon})$, there exists $c_1 > 0$ such that

$$N_{\text{cycles}} \ge \exp\left(c_1 n^{1-\varepsilon}\right),$$

proving (1).

Step 3: Lower bound on the period (proof of (2)). Let $I = \{b, \ldots, m\}$ and $M = |I| = m - b + 1$. The multiset of block-cycle lengths is formed by taking $\lceil 4M \rceil$ uniform samples from I

with replacement. By Lemma 3, after ignoring repeats, the sampled lengths contain an independently sampled subset of I with constant inclusion probability with high probability. Since adding elements to a set cannot decrease its lcm, Theorem 5 implies that

$$P \geq \exp(c_2 m).$$

On the high-probability event from Step 1, $m = \Theta((n/d)^{1/2})$. Since here we have $d = \Theta(n^{\varepsilon})$, we get

$$P \geq \exp(c_2 (n^{1-\varepsilon})^{1/2}),$$

proving (2).

Step 4: Noise robustness (proof of (3)). For d neurons in any block of any given block cycle consider d "imaginary agents", undergoing a majority dynamics with 2 opinions. For instance, consider a block with some distribution of states (opinions). At the next unit of time, these opinions will be transmitted to the next block, as the Hopfield update rule is a similar majority update rule, and this will be the new state of the "imaginary agents". Hence we need whichever opinion is popular among these d agents to reach consensus before completing the cycle. The only difference between Hopfield dynamics and majority dynamics is that in case of a tie, majority dynamics chooses an opinion uniformly at random, whereas Hopfield dynamics preserves the old state. Since h is odd by assumption, they are exactly the same here. Also note that in consensus dynamics, the agents are chosen with replacement, which we handle by setting the weights equal to a node's multiplicity (see Definition 8).

According to Theorems 9 and 10, the agents should reach consensus (the state before they were flipped) in at most $C \log d$ units of time (this is why we keep the cycles longer than $C \log d$) given that the initially popular opinion has at least $d/2 + \lambda_1 \sqrt{d \log d}$ supporters. Since this happens for each block, the entire block cycle must come back to its original state (before flipping) in at most $C_1 \log d$ time. We will now calculate the probability of the event that each original block opinion has at least $d/2 + \lambda_1 \sqrt{d \log d}$ supporters after the random flips.

Let X_i be the number of neurons that remain unflipped in block i. Each neuron stays unflipped with probability $1 - p = \frac{1}{2} + \eta$, so

$$X_i \sim \mathrm{Bin}\left(d, \frac{1}{2} + \eta\right).$$

Hence the expectation is

$$\mu := \mathbb{E}[X_i] = d\left(\frac{1}{2} + \eta\right) = \frac{d}{2} + \eta d.$$

Now take $\eta = C_p \sqrt{\frac{\log n}{d}}$. Then

$$\mu = \frac{d}{2} + C_p \sqrt{d \log n}.$$

Since $d = \Theta(n^{\varepsilon})$, there is a constant $C_\lambda > 0$, depending only on λ_1 and ε, such that for all sufficiently large n,

$$\lambda_1 \sqrt{d \log d} \leq C_\lambda \sqrt{d \log n}.$$

Thus the gap is at least

$$\mu - \left(\frac{d}{2} + \lambda_1 \sqrt{d \log d}\right) \geq (C_p - C_\lambda) \sqrt{d \log n}.$$

Choose $C_p > C_\lambda$, as allowed by the statement. Let

$$a := (C_p - C_\lambda) \sqrt{d \log n}.$$

By Hoeffding's inequality,

$$\mathbb{P}\left(X_i \leq \mu - a\right) \leq \exp\left(-\frac{2a^2}{d}\right).$$

Substituting a,

$$\mathbb{P}\left(X_i \leq \mu - a\right) \leq \exp\left(-2\left(C_p - C_\lambda\right)^2 \log n\right) = n^{-2(C_p - C_\lambda)^2}.$$

Since $\frac{n}{d} = \Theta(n^{1-\varepsilon})$, there is a constant $C_N > 0$ such that the union bound gives

$$\mathbb{P}\left(\exists i : X_i \le \frac{d}{2} + \lambda_1\sqrt{d\log d}\right) \le C_N n^{1-\varepsilon} \cdot n^{-2(C_p - C_\lambda)^2}.$$

For sufficiently large C_p, this probability tends to 0. Hence with high probability, for all blocks i,

$$X_i > \frac{d}{2} + \lambda_1\sqrt{d\log d}.$$

Hence we have ensured the assumptions of Theorems 9 and 10. Suppose an h majority process converges in at most $c\log d$ time with probability say $1 - 1/d^\xi$. A union bound over all blocks gives the failure probability

$$\frac{1}{d^\xi} \cdot \frac{n}{d} = O\left(n^{1-\varepsilon(1+\xi)}\right) = o(1),$$

since $\varepsilon(1+\xi) > 1$. Now choose c_3 such that $c_3 \log m \ge c\log d$, which ensures that even for the smallest possible cycle, there is enough time to return to consensus, which proves Theorem 2. □

E.4 Proof of Theorem 3

Proof. By the same calculation as in the proof of Theorem 2, with probability $1 - o(1)$, for each block the initial perturbation preserves a bias of at least $\lambda_1\sqrt{d\log d}$ towards the original block opinion. Let $q = \lfloor \gamma\sqrt{d}/\log d \rfloor$ and $k = \lceil C_{\mathrm{adv}}\sqrt{d}/\log d \rceil$, as in the theorem statement. Since only at most q neurons in each block receive adversarial incoming edges, choosing γ sufficiently small ensures that the adversarial budget is within the range allowed by Theorem 11.

We may therefore apply Theorem 11 to each block. Under this bias condition and with at most q adversarially affected neurons, the majority dynamics converges back to the original majority state up to at most k exceptions within $O((\log d)^2)$ rounds, after choosing C_{adv} large enough to dominate the hidden constant in Theorem 11. The failure probability is at most $d^{-\xi}$ for some constant $\xi > 0$. Reusing the same union bound as above over the $n/d = \Theta(n^{1-\varepsilon})$ blocks, the total failure probability is at most

$$\frac{1}{d^\xi} \cdot \frac{n}{d} = O\left(n^{1-\varepsilon(1+\xi)}\right) = o(1),$$

by the hypothesis $\varepsilon(1+\xi) > 1$. Thus the entire network returns to the (k, d)-weak limit cycle associated with the original monochromatic orbit through c with probability $1 - o(1)$. The convergence time required is say $c'(\log d)^2$. Choose c_4 such that $c_4(\log m)^2 \ge c'(\log d)^2$, which ensures that even for the smallest possible cycle, there is enough time to return to consensus. □

E.5 Global convergence from arbitrary initial configurations

Theorem 8. *Fix any initial configuration $x(0) \in \{-1, +1\}^n$.*

For every realization of the dense block-cyclic architecture in Theorem 1, the synchronous Hopfield trajectory starting from $x(0)$ reaches a limit cycle within at most m rounds.

Under the assumptions of Theorem 2, with probability $1 - o(1)$ over the random construction of the h-sparse block-cyclic architecture, the synchronous Hopfield trajectory starting from $x(0)$ reaches a limit cycle within $O(\log d)$ rounds.

Under the assumptions of Theorem 3, with at most $q = \lfloor \gamma\sqrt{d}/\log d \rfloor$ adversarially perturbed neurons per block, with probability $1 - o(1)$ the synchronous Hopfield trajectory starting from $x(0)$ reaches a (k, d)-weak limit cycle within $O((\log d)^2)$ rounds, where $k = \lceil C_{\mathrm{adv}}\sqrt{d}/\log d \rceil$.

Proof. We begin with the dense block-cyclic architecture. Fix one dense block cycle $S_0, S_1, \ldots, S_{\ell-1}$ of length ℓ. If the sum of the neuron states in S_j at time t is nonzero, then every neuron in S_{j+1} receives the same nonzero input at time $t+1$, so S_{j+1} becomes monochromatic at time $t+1$. If instead the sum of the neuron states in S_j is zero, then every neuron in S_{j+1} receives input zero and therefore keeps its previous state by the convention $\mathrm{sign}(0) = x_i(t)$. Consequently, if every block in the cycle is balanced at time 0, then the entire block cycle is already a fixed point. Otherwise choose

one block whose neuron states have nonzero sum at time 0. Its successor becomes monochromatic after one round, the next block becomes monochromatic one round later, and so on. After at most ℓ rounds every block in that dense block cycle is monochromatic. From that moment onward the block cycle evolves by cyclically shifting these monochromatic signs, and hence it has entered a limit cycle. Since the dense block-cyclic architecture is a disjoint union of such block cycles and every length is at most m, the full dense network reaches a limit cycle within at most m rounds.

We first consider the non-adversarial sparse block-cyclic architecture. Fix one block cycle $S_0, S_1, \ldots, S_{\ell-1}$ of length ℓ, and for $0 \le t < \ell$ let $Y_t \in \{-1,+1\}^d$ denote the state of S_t at global time t. Because adjacent blocks are connected by independent h-sparse random bipartite graphs and h is odd, the evolution $Y_0, Y_1, \ldots$ agrees with the h-majority gossip process on d agents for as long as we do not complete a full turn around the block cycle. By choosing the constant c_3 in Theorem 2 sufficiently large and using $\log d = \Theta(\log m)$, we may assume that every block cycle has length at least $b \ge c \log d$, where c is the convergence constant from Theorem 9. Hence, before any block cycle repeats an edge layer, Theorem 9 implies that Y_t becomes monochromatic within $T = O(\log d)$ rounds with failure probability at most $d^{-\xi}$.

Applying the same argument to each of the $n/d = \Theta(n^{1-\varepsilon})$ blocks and taking a union bound, the total failure probability is at most

$$\frac{n}{d} \cdot d^{-\xi} = O\left(n^{1-\varepsilon(1+\xi)}\right) = o(1).$$

Therefore, with probability $1 - o(1)$, after $T = O(\log d)$ rounds every block is monochromatic. From that time onward each block simply copies the common sign of its predecessor, so the network evolves by cyclically shifting the resulting binary labels on the block cycles. Hence the trajectory has entered a limit cycle.

For the adversarial extension, the same reduction yields on each block cycle an h-majority process on d agents with at most $q = \lfloor \gamma\sqrt{d}/\log d \rfloor$ adversarially perturbed agents. By choosing c_4 sufficiently large in Theorem 3, we may assume that every block cycle has length at least $b \ge c'(\log d)^2$, where c' is the convergence constant from Theorem 11. Hence Theorem 11 implies that within $T' = O((\log d)^2)$ rounds each block reaches consensus on some sign up to at most $k = \lceil C_{\mathrm{adv}}\sqrt{d}/\log d \rceil$ exceptional neurons, with failure probability at most $d^{-\xi}$. The same union bound over the n/d blocks again gives total failure probability $o(1)$. Once every block is monochromatic up to those k neurons, the regular part of the network again evolves by cyclically shifting the resulting binary labels. Equivalently, the full trajectory stays within block-max Hamming distance at most k from the corresponding monochromatic orbit for all subsequent times, and therefore it has entered a (k, d)-weak limit cycle. □

F Background on Opinion Dynamics

Opinion dynamics protocols are fundamental models for information spreading and consensus formation in distributed systems. They capture how agents interact locally yet produce global convergence through repeated, randomized communication. We describe the widely studied Gossip model (Berenbrink et al., 2023; Clementi et al., 2018; Becchetti et al., 2016).

Let $N \in \mathbb{N}$ be the number of agents. The system is represented by a complete graph on N anonymous and identical agents. Time proceeds in discrete, synchronous rounds indexed by $t \in \mathbb{N}_0$. At each time t, the global configuration is a vector $C_t \in \{a, b\}^N$ assigning to every agent an opinion from a finite set (here $\{a, b\}$).

In every round, all agents are activated simultaneously. Each agent u independently samples a multiset of j agents $v_1, \ldots, v_j$ uniformly at random with replacement from the population. Self-sampling is allowed, i.e., $v_i = u$ is possible. The sampling of different agents and different rounds is mutually independent.

Given the sampled opinions, agent u computes its new opinion according to a deterministic update rule based solely on the multiset $\{C_t(v_1), \ldots, C_t(v_j)\}$. In the j-majority dynamics, the update rule is defined as follows: agent u adopts the majority opinion among the sampled opinions. In case of a tie (which may occur when j is even), the agent breaks the tie uniformly at random.

All agents perform this sampling and decision step using the configuration C_t. After all agents have determined their new opinions, the system performs a synchronous update, producing the next configuration C_{t+1}.

Thus, the evolution $(C_t)_{t\geq 0}$ defines a discrete-time stochastic process on $\{a,b\}^N$, where transitions are governed by independent random sampling and simultaneous opinion updates across all agents.

In the study of opinion dynamics, a quantity of interest is the convergence time, i.e., how many rounds are required before all agents reach the same state. We now state a result by Berenbrink et al. (2023) on the convergence time of j majority.

Theorem 9 (Convergence of j-Majority in the Gossip Model). *Let $j \geq 3$ be a fixed integer, and let $(C_t)_{t\geq 0}$ be the stochastic process defined above. Then there exists a constant $c = c(j) > 0$ such that, for any initial configuration $C_0 \in \{a,b\}^N$, the process reaches consensus within $O(\log N)$ rounds with high probability. More precisely, if T denotes the (random) convergence time, then*

$$\mathbb{P}[T \leq c \log N] \geq 1 - N^{-\xi},$$

where $\xi = \Omega(1)$.

We now state a result by Clementi et al. (2018), which shows that if the initial configuration has a sufficiently large bias towards one of the opinions, then the initially popular opinion wins at the end with high probability.

Theorem 10 (Large Bias Implies Correct Consensus). *Let $j \geq 3$ be fixed, and let $(C_t)_{t\geq 0}$ be the j-majority gossip process. Let X_t denote the number of agents holding opinion a at time t, and assume without loss of generality that $X_0 \geq N/2$.*

There exists a constant $\gamma > 0$ such that if the initial bias satisfies

$$X_0 - (N - X_0) \geq \gamma\sqrt{N \log N},$$

then the process converges to the consensus configuration in which all agents hold opinion a within $O(\log N)$ rounds with high probability. In particular, if T denotes the convergence time, then

$$\mathbb{P}[C_T = (a, \ldots, a)] \geq 1 - N^{-\Omega(1)}.$$

A *dynamic adversary* with budget F is an adaptive process that, at each round t, after observing the configuration C_t, may arbitrarily change the opinions of up to F agents (possibly depending on the entire history of the process). The following result by Becchetti et al. (2016) shows robustness under dynamic adversaries.

Theorem 11 (Robust Convergence under Dynamic Adversaries). *Let $(C_t)_{t\geq 0}$ be the j-majority gossip process with $j \geq 3$. Then there exists a constant $\beta > 0$ such that if $F \leq \beta\sqrt{N}/\log N$, the following holds: for any initial configuration $C_0 \in \{a,b\}^N$, the process reaches an almost-consensus within $O(\log^2 N)$ rounds with high probability. More precisely, there exists a round $T = O((\log N)^2)$ such that*

$$|\{u : C_T(u) = x\}| \geq N - O(\sqrt{N}/\log N)$$

for some opinion $x \in \{a,b\}$, and all but at most $O(\sqrt{N}/\log N)$ agents keep opinion x for any polynomial number of subsequent rounds, despite the adversarial actions. Equivalently, this event holds with probability at least $1 - N^{-\xi}$ for some constant $\xi = \Omega(1)$.

Furthermore, if the initial bias satisfies

$$|\{u : C_0(u) = a\}| - |\{u : C_0(u) = b\}| \geq K\sqrt{N \log N}$$

for a sufficiently large constant $K > 0$, then with high probability the opinion initially held by the majority of agents is the opinion x appearing above.

For later use, fix $\lambda_1 > 0$ large enough that an opinion count of at least $N/2 + \lambda_1\sqrt{N \log N}$ satisfies the bias hypotheses in Theorems 10 and 11.

G Additional Experimental Details and Related-Work Comparison

G.1 Notation Summary

Table 1 collects the recurring notation used throughout the paper.

Table 1: Summary of the recurring notation used throughout the paper.

Notation	Meaning
n	total number of neurons in the network.
$x(t) \in \{-1, +1\}^n$	network state at time t.
$W = (w_{ij})$	asymmetric weight matrix; w_{ij} is the edge weight from neuron j to neuron i.
F	synchronous update map induced by the majority-rule dynamics.
$\mathcal{C}$, P	a limit cycle and its period.
$\mathcal{A}$	family of constructed limit-cycle attractors.
BCA (d, m)	dense block-cyclic architecture with block size d and scale parameter m.
$A = \{\ell_1, \ldots, \ell_z\}$	sampled multiset of block-cycle lengths.
ℓ, ℓ_i	length of a block cycle.
N_{blocks}	total number of blocks in the architecture.
d	number of neurons in each block.
m	scale parameter controlling the range of sampled block-cycle lengths.
h	odd in-degree used in the h-sparse construction.
p	independent flip probability in the noise model.
k	maximum number of adversarially perturbed neurons per block.
$d_{\text{BM}}^{(d)}$	block-max Hamming distance with respect to the canonical block partition.
h-sparse BCA (d, b, m)	sparse block-cyclic architecture in which each neuron receives exactly h inputs from the preceding block.
$\psi(m)$	$\log \operatorname{lcm}\{1 \le a \le m\}$
N_{cycles}	total number of limit cycles in a given directed Hopfield network.
d_H	hamming distance
$\mathcal{B}(\mathcal{C})$	basin of attraction of a limit cycle $\mathcal{C}$.

G.2 Additional Experimental Details

Details of Experiments in Section 6. In the first set of experiments, we consider a block cycle of length $\ell = 200$, where each block has size d (varied across the experiment), and each neuron receives $h = 3$ inputs, corresponding to the sparsest network topology allowed by our theoretical results. The initial configuration is assigned as in Step 1 of our construction (Theorem 2), so that all neurons within a given block share the same sign. For simplicity we keep all signs +1 while initializing for all further experiments, unless stated; this isolates the block-level denoising mechanism rather than the period-length effect of nonconstant block sequences. We then independently flip each neuron with probability p (also varied across the experiment), after which the Hopfield dynamics are run to test whether the system returns to the initial configuration. This procedure is repeated 50 times. In each trial, the network topology is resampled: for every neuron, its 3 neighbors are selected anew. The retrieval probability is then estimated empirically. The results are presented in Figure 2(a). This experiment illustrates reliable recovery of noisy block-constant states while varying network sparsity, noise, and random topology.

Next, we run a random extra-connection experiment motivated by the weak-recovery setting of Theorem 3. Specifically, we consider a 3-sparse block-cyclic architecture consisting of two block cycles of lengths $\ell_1 = 150$ and $\ell_2 = 200$, each with block size $d = 10000$. We randomly select k neurons in each block and connect them to randomly chosen neurons in the Hopfield network, with edges directed toward these k neurons. These added edges are not adversarial: they are sampled randomly, and the selected target neurons are ignored when evaluating recovery. The value of k is varied across the experiment. The initial configuration is assigned as in the previous experiment, i.e., all neurons in the (+1) state. Each neuron is then independently flipped with probability p,

after which the Hopfield dynamics are run to determine whether the system recovers the initial state before completing a full cycle, ignoring the neurons that receive additional inputs. This procedure is repeated 50 times to estimate the revival probability. In each trial, the network topology is resampled, including the random interconnections. The experiments are conducted over a range of values for k and p. The results are shown in Figure 2(b). Since the added inputs are random rather than worst-case, they can sometimes help: when the population is already biased toward the correct sign, extra random samples may reinforce that bias, which explains why the highest-noise curve improves for larger k.

Further Experiments. As a further stress test motivated by Theorem 3, we study the robustness of the dynamics in the presence of extreme adversaries, namely, nodes whose updates are always opposite to the majority rule. We consider a single 3-sparse block cycle of length $\ell = 200$ with block size $d = 5000$. We randomly select k neurons in the cycle and modify their update rule so that each chosen neuron updates to the opposite of the state prescribed by the standard Hopfield dynamics. The network is initialized with all neurons in the $+1$ state, after which each neuron is independently flipped with probability $p = 0.45$. We then run the Hopfield dynamics and check whether the system returns to the original limit cycle, disregarding the adversarial neurons when evaluating recovery. The probability of successful revival is estimated over 50 trials, and the results are shown in Figure 3(a). These experiments illustrate that the network can retain recovery behavior even in the presence of highly adversarial nodes.

Next, we probe the length condition behind Theorem 8 (Appendix E.5) by asking whether a randomly initialized block cycle converges to a configuration in which all neurons within each block share the same sign. We consider a single 3-sparse block cycle with fixed length $\ell = 20$ and block size d, where d varies throughout the experiment. Starting from a uniformly random initial configuration, we run the dynamics and measure the probability that the system converges to a state in which every neuron within a given block has identical sign before one full rotation. The probability is estimated over 50 trials for each value of d. The results are presented in Figure 3(b). The slight downward trend as d grows is expected because the theorem requires enough time for the h-majority process to reach consensus, namely a cycle length growing at least logarithmically with d. Thus this fixed-ℓ experiment is best read as a finite-time diagnostic: even with a cycle length that does not scale with d, success remains high, while the deterioration is consistent with the predicted need for $\ell = \Omega(\log d)$.

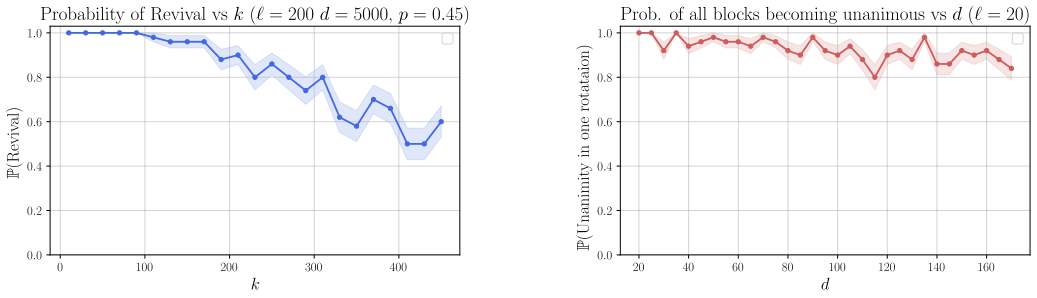

(a) (b)

Figure 3: (a) Recovery probability on a sparse block cycle versus number of adversarial neurons, showing robustness. (b) Probability of convergence to a state where each neuron in any given block has the same sign before one full rotation, starting from a random configuration with fixed cycle length $\ell = 20$. The mild decline with d reflects that the theorem requires the cycle length to grow at least logarithmically with d.

G.3 Spectral Embedding Plots

To provide a geometric view on our architecture, we present spectral embeddings of representative sparse block-cyclic architectures augmented with adversarial connections, capturing the network's connectivity structure. Nodes are embedded via the leading eigenvectors of the graph Laplacian, which expose the model's underlying organization. As shown in Figure 4, this coarse structure remains largely preserved even after introducing adversarial edges.

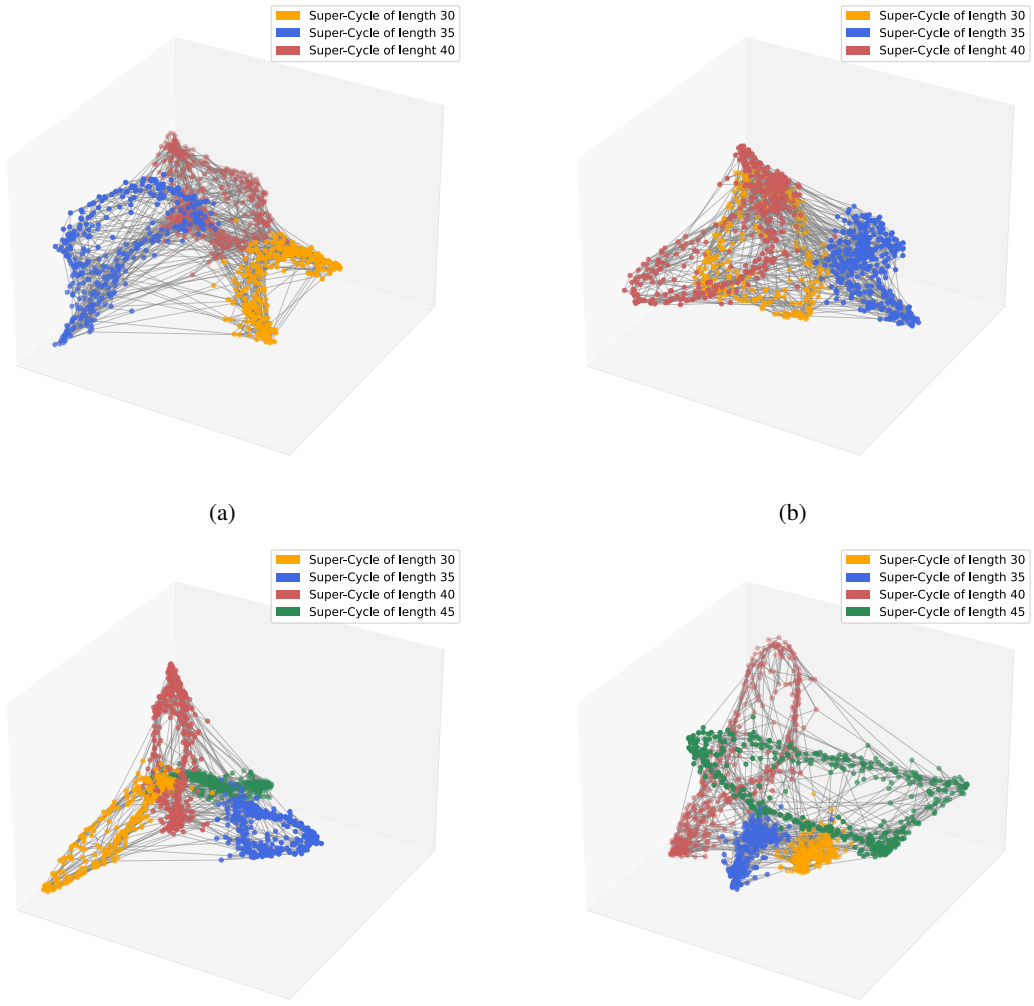

Figure 4: Spectral embeddings of two sparse block-cyclic architectures with adversarial connections: (a) & (b) block-cycle lengths 30, 35, and 40; (c) & (d) lengths 30, 35, 40, and 45. In all panels, each block has size $d = 10$ and each neuron receives $h = 3$ incoming edges from the preceding block.

G.4 FINITE-SIZE BEHAVIOUR OF THE CONSTRUCTION

Trade-off between block size and number of blocks. For a fixed neuron budget n, the block size d can be traded against the number of blocks, and this moves the capacity, the period and the noise tolerance together. Table 2 evaluates the quantities in the proof of Theorem 1 exactly at $n = 10^8$, with no asymptotic approximation. For each d we take m so that $dN_{\text{blocks}} \approx n$, draw the multiset A of $z = \lceil 4m \rceil$ lengths, and compute the period $P = \operatorname{lcm}(A)$ and the number of limit cycles $\prod_{\ell \in A} \mathcal{M}_2(\ell)/P$ from the exact counts of aperiodic binary strings. The column $p^\star$ is the largest flip probability at which, with probability at least 0.99, every block keeps a strict majority, so that by Remark 3 the state returns exactly to its cycle; since blocks are flipped independently this probability is $\mathbb{P}[\mathrm{Bin}(d, p) < d/2]^{N_{\text{blocks}}}$. Entries are averages over 8 independent draws of the architecture; the two logarithmic columns vary by at most 7% between draws, and $p^\star$ only in the fourth decimal.

Table 2: Finite-size trade-off at $n = 10^8$ neurons.

d	N_{blocks}	$\log_{10} N_{\text{cycles}}$	$\log_{10} P$	$p^\star$
50	2.0×10^6	6.0×10^5	429	0.148
100	1.0×10^6	3.0×10^5	300	0.237
200	5.0×10^5	1.5×10^5	218	0.311
400	2.5×10^5	7.5×10^4	152	0.367
800	1.3×10^5	3.8×10^4	106	0.408

Capacity and period decrease as d grows while the tolerated noise rises towards $\frac{1}{2}$, and over the whole range both the number of cycles and the period remain very large. The same trade-off holds at smaller sizes: at $n = 10^6$ and $d = 800$ the construction still gives about 10^{375} cycles of period about $10^{10.3}$, with $p^\star = 0.42$.

Spread of the available periods. In the family $\mathcal{A}$ of Theorem 1 every block cycle carries an aperiodic string, so all its attractors share the period $\operatorname{lcm}(A)$. The same network, however, also contains limit cycles in which only a set B of the block cycles carries aperiodic strings and the others are constant; such a cycle has period $\operatorname{lcm}(B)$, and Remark 3 applies to it unchanged, since it does not use aperiodicity. Table 3 reports the distribution of $\ln \operatorname{lcm}(B)$ over 10^4 uniformly random subsets $B \subseteq \{1, \ldots, m\}$ (with $z = \lceil 4m \rceil$ samples, A misses only about a fraction e^{-4} of these lengths). The log-period is spread over a wide band rather than concentrated at one value. Its mean is close to $m \ln 2$ (34.7, 69.3 and 138.6), the value Theorem 4 predicts for inclusion probability $\frac{1}{2}$. This diversity matters: a family of limit cycles sharing a single period could be re-expressed as fixed points in a co-rotating frame, whereas no single frame can do so for cycles of many different periods under the same update map.

Table 3: Distribution of $\ln \operatorname{lcm}(B)$ for uniformly random $B \subseteq \{1, \ldots, m\}$, over 10^4 samples.

m	mean	std	1%	median	99%
50	33.5	5.3	21.1	33.5	45.1
100	65.9	7.8	47.2	66.0	83.3
200	139.2	13.1	109.1	139.2	169.5

G.5 Related-Work Comparison

Table 4 compares our results with related sequence-memory results in Hopfield-type networks, in terms of the model, the number and length of the cycles, and their robustness.

Table 4: Comparison with related sequence-memory results in Hopfield-type networks.

Work	Model / setting	Number of cycles	Cycle length	Robustness	Main distinction
Bastolla and Parisi (Bastolla & Parisi, 1997)	Fully asymmetric random neural networks	Mean total number grows linearly in n	Typical length grows exponentially in n	Basin statistics studied, but no designed retrieval robustness	Long cycles arise statistically in random asymmetric dynamics rather than from a constructive sequence-memory architecture.
Hwang et al. (Hwang et al., 2019)	Asymmetric neural networks	Computes mean number n_L of cycles of fixed length L	Finite L, with an exponential finite-size cutoff for long cycles	Basins not analyzed in the counting formalism	Gives a statistical-mechanics characterization of cycle counts and the chaotic long-cycle regime.
Hwang et al. (Hwang et al., 2020)	Diluted neural networks	Exponentially many finite-length cycles when the complexity is positive	Fixed finite length, mainly $L = 4$ in the fully asymmetric diluted case	Basins discussed as motivation, not retrieval robustness	Shows how dilution changes the exponential complexity of short temporal cycles.
Zhang et al. (Zhang et al., 2013)	Hopfield-type networks with pseudoinverse learning	Prescribed admissible cycles	Depends on prescribed cycle	Retrieval robustness not analyzed	Gives an admissibility criterion for prescribed cycles and a pseudoinverse construction, with emphasis on network topology.
Muscinelli et al. (Muscinelli et al., 2017)	Standard asymmetric Hopfield network	Explicitly constructs a maximal orbit	2^n	No nontrivial basin	Constructs a maximal-length orbit covering the full state space; the orbit cannot be attractive.
Chaudhry et al. (Chaudhry et al., 2023)	Dense associative memory / modern Hopfield-type model	Not framed as counting multiple limit-cycle attractors	Sequence capacity is polynomial or exponential, depending on the nonlinearity	Probabilistic sequence recall below capacity	Achieves long prescribed sequence recall using nonlinear interactions beyond the synchronous asymmetric Hopfield model.
This work	synchronous asymmetric Hopfield network with binary neurons and majority updates	$\exp(\Omega(n/(\log n)^2))$	$\exp(\Omega(\sqrt{n}/\log n))$	Robust to random flips up to $1/2-o(1)$; persists under sparsification and adversarial connections	Simultaneously obtains many long, robust limit-cycle attractors using a simple block / block-cycle architecture.