SIAM J. Applied Dynamical Systems
Vol. 23, No. 1, pp. 855–884

Diversity of Emergent Dynamics in Competitive Threshold-Linear Networks*

Katherine Morrison†, Anda Degeratu‡, Vladimir Itskov§, and Carina Curto§

Abstract. Threshold-linear networks consist of simple units interacting in the presence of a threshold nonlinearity. Competitive threshold-linear networks have long been known to exhibit multistability, where the activity of the network settles into one of potentially many steady states. In this work, we find conditions that guarantee the *absence* of steady states, while maintaining bounded activity. These conditions lead us to define a combinatorial family of competitive threshold-linear networks, parametrized by a simple directed graph. By exploring this family, we discover that threshold-linear networks are capable of displaying a surprisingly rich variety of nonlinear dynamics, including limit cycles, quasi-periodic attractors, and chaos. In particular, several types of nonlinear behaviors can co-exist in the same network. Our mathematical results also enable us to engineer networks with multiple dynamic patterns. Taken together, these theoretical and computational findings suggest that threshold-linear networks may be a valuable tool for understanding the relationship between network connectivity and emergent dynamics.

Key words. attractors, networks, threshold-linear networks, competition, index theorem

MSC code. 92

DOI. 10.1137/22M1541666

1. Introduction.

Networks are complex dynamical systems that consist of nodes and their interactions. They are commonly used models in fields as disparate as ecology, economics, and neuroscience. Even when the building blocks are simple, networks can display rich *emergent dynamics,* whose complexity cannot be reduced to a sum of constituent parts. Moreover, the most interesting dynamic phenomena that arise are fundamentally nonlinear behaviors, such as multistability, periodic attractors, and chaos.

Despite this, network dynamics are often approximated using linear models, namely linear systems of ordinary differential equations (ODEs). This is because the accompanying mathematical theory is well developed. Indeed, one might say that networks with linear interactions are the complex systems we already understand. While reducing more complicated models

*Received by the editors December 16, 2022; accepted for publication (in revised form) by B. Rink November 9, 2023; published electronically March 13, 2024.
https://doi.org/10.1137/22M1541666
Funding: The first author's research was supported by NIH R01 EB022862 and NSF DMS-1951599. The third author's research was supported by NSFIOS-155925 and the NSF Next Generation Networks for Neuroscience Program (award 2014217). The fourth author's research was supported by NIH R01 EB022862, NSF DMS-1951165, and NSF DMS-1516881. The research of the first, third, and fourth authors was also supported by the Statistical and Applied Mathematical Sciences Institute under grant NSF DMS-1127914.
†School of Mathematical Sciences, University of Northern Colorado, Greeley, CO 80639 USA (katherine.morrison@unco.edu).
‡Fachbereich Mathematik, Universität Stuttgart, Stuttgart, 70569 Germany (anda.degeratu@mathematik.uni-stuttgart.de).
§Department of Mathematics, The Pennsylvania State University, University Park, PA 16802 USA (vladimir.itskov@psu.edu, ccurto@psu.edu).

to linear approximations can be a useful approach, this strategy also poses severe limitations. Phenomena such as multistability, chaos, and robust periodic attractors (limit cycles) simply do not occur in linear models. Can we replace approximation by linear systems of ODEs with something "almost" linear, simple enough that a useful mathematical theory can be developed and yet be capable of capturing the full variety of nonlinear behavior?

Motivated by this question, we study the dynamics of threshold-linear networks (TLNs). Emergent dynamics in these networks are not inherited from intrinsically oscillating nodes or a fluctuating external input; instead, they can be attributed solely to the structure of connectivity, given by a matrix W. The nonlinear behavior stems entirely from the presence of a simple threshold at each node, which guarantees that the activity of individual units cannot go negative. This nonnegativity is natural in any system where the dynamic variables represent fundamentally nonnegative quantities, such as the size of a population, a chemical concentration, or the firing rate of a neuron. Though TLNs look essentially linear, the presence of the threshold changes everything. With it, the entire repertoire of complex nonlinear behavior comes into play (multistability, limit cycles, and quasi-periodic behavior), and even deterministic chaos emerges.

Historically, TLNs have been studied with an eye towards *stable fixed points,* as these are the traditional attractors considered in the Hopfield model and related neural network literature [1, 2]. This led to an early emphasis on the case of symmetric W, where convergence to stable fixed points can be guaranteed [3, 4]. In the early 2000s, the study of "permitted" and "forbidden" sets also provided simple conditions under which symmetric (inhibitory) TLNs give rise to multistability and showed that the collection of all permitted sets of a symmetric network satisfies the structure of a simplicial complex [3]. Notably, a permitted set is not a fixed point but a subset of neurons that can in principle be co-activated at a stable fixed point for at least one external input; this simplifies the problem to one of analyzing the spectral properties of principal submatrices of W, while ignoring questions about which fixed points can arise for a given external input.

The present authors continued the study of stable fixed points and permitted sets of TLNs in subsequent work [5, 6, 7], also focusing primarily on the symmetric case. We then shifted our attention to nonsymmetric W, leading to the work presented here. A preliminary (unpublished) version of the current article was our first serious foray into the dynamics of nonsymmetric TLNs [8]; those results became the seeds of this article and inspired related developments [9, 10, 11, 12, 13]. In addition to biological realism, our primary motivation for this shift was a desire to explore a wider range of nonlinear behavior, including more complex attractors. There is also, however, a body of related work in the nonsymmetric case that maintains its focus on stable fixed points [14, 15].

In order to identify networks with complex attractors, we initially sought conditions that would guarantee the *absence* of stable fixed points, so that network activity would be forced to be oscillatory or chaotic. This resulted in Theorem 3.3, whose conditions on pairwise interactions can be framed in terms of a directed graph. These conditions then led us to define a combinatorial family of competitive TLNs that are parametrized by a directed graph; we call this family the *CTLN model* ("C" for "combinatorial"). By exploring CTLNs, we have discovered that TLNs are capable of displaying a surprisingly rich variety of nonlinear dynamics, which we illustrate in section 4.

Although they exhibit high-dimensional nonlinear dynamics, TLNs and CTLNs are also surprisingly tractable. We have been able to prove a series of mathematical results, Theorems 3.1, 3.3, 3.4, and 3.5, which provide valuable information about the stable and unstable fixed points of these networks. Theorem 3.1 is an index theorem, which constrains the set of fixed points that can co-exist in a given TLN with fixed input. Specifically, for a TLN with connectivity matrix W, each fixed point x^* has an *index,* $\mathrm{idx}(\sigma) = \mathrm{sgn}\det(I - W_\sigma) \in \{\pm 1\}$, where $\sigma \subseteq [n]$ is the set of active neurons of x^*. Theorem 3.1 states that these indices must sum to $+1$. In particular, this implies the total number of fixed points is odd. Since the index of a stable fixed point is always $+1$, we also obtain an upper bound on the total number of stable fixed points.

As previously discussed, Theorem 3.3 provides conditions guaranteeing the *absence* of stable fixed points in TLNs with bounded activity. Because the conditions are graph-theoretic, and independent of the precise choice of weights W, it is easy to specialize this result to the case of CTLNs, yielding Theorem 3.4. Finally, Theorem 3.5 revisits the topic of stable fixed points, identifying graph structures that give rise to stable equilibria in CTLNs.

The proofs of Theorems 3.1, 3.3, 3.4, and 3.5 required novel approaches and techniques different from what was previously used in the symmetric case. In particular, in section 5 we consider TLNs as a patchwork of linear systems and develop methods to carefully analyze the relationship between fixed points of the component linear systems and those of the full nonlinear TLN. The proof of Theorem 3.1 involves a novel application of the Poincaré–Hopf theorem; to our knowledge, this has not previously been used in the TLN literature. We also prove a number of new lemmas, specifically for the nonsymmetric case. For example, a key result in proving Theorem 3.3 is Lemma 5.8, which relates the stability of a nonsymmetric matrix to that of its 2×2 principal submatrices.

What is perhaps most striking about our CTLN theorems is how they have enabled us to engineer networks exhibiting multiple dynamic attractors, with distinct attractors corresponding to different initial conditions. In section 4, we illustrate the diversity of emergent dynamics via simulations of example networks that are guaranteed to have no stable fixed points (by Theorem 3.4). We also use Theorem 3.5 to generate networks with transiently active "cell assemblies." We conclude this section with some "engineered" networks exhibiting multiple prescribed dynamic attractors. In each of these cases, the networks are built using constraints and intuition relating directly to the graph of connectivity, without any need for parameter tuning. Taken together, our theoretical and computational results suggest that TLNs are a valuable tool for studying the relationship between emergent dynamics and network connectivity.

While CTLNs are a subfamily of TLNs, it should be noted that TLNs are themselves a special case of a much larger class of (nonsmooth) switching dynamical systems. These systems are piecewise-linear and have been extensively studied in the context of stability and control (see, e.g., [16, 17, 18, 19, 20]). Interestingly, by studying the constrained subfamily of CTLNs we have been able to obtain a richer variety of dynamics, as well as more detailed theoretical results, than what is readily available in the more general piecewise-linear settings. Nevertheless, it would be valuable to better understand how the results presented here relate to the broader switching dynamical systems literature.

2. Preliminaries. A *threshold-linear network* (TLN) is a rate model consisting of n nodes, with dynamics governed by the system of ODEs:

$$\frac{dx_i}{dt} = -x_i + \left[\sum_{j=1}^{n} W_{ij}x_j + b_i\right]_+, \quad i = 1,\dots,n. \tag{1}$$

The dynamic variables $x_1,\dots,x_n$ give the activity levels[1] of nodes $1,\dots,n$. The matrix entries W_{ij} are directed connection strengths between pairs of nodes, the vector $b = (b_1,\dots,b_n) \in \mathbb{R}^n$ represents the external drive to each node, and the threshold nonlinearity $[\cdot]_+$ is given by $[y]_+ \stackrel{\text{def}}{=} \max\{y, 0\}$. We refer to the TLN with matrix W and vector b as (W, b). Note that although TLNs are nonlinear dynamical systems, they are piecewise linear: the threshold nonlinearity decomposes the state space into chambers, within which the dynamics are linear, of the form $dx/dt = Ax + c$ (see section 5.1 for more details).

A *fixed point* of (1) is a vector $x^* \in \mathbb{R}^n$ that satisfies $\frac{dx_i}{dt}\big|_{x=x^*} = 0$ for each $i \in \{1,\dots,n\}$. We will be interested in both stable and unstable fixed points of TLNs. The *support* of a fixed point x^* is the subset of active nodes,

$$\operatorname{supp} x^* \stackrel{\text{def}}{=} \{i \mid x_i^* > 0\}. \tag{2}$$

Throughout this paper, we restrict ourselves to considering TLNs that are *nondegenerate* (a precise definition is given in Definition 5.1). Briefly, nondegeneracy requires that certain determinants are nonzero, such that each of the linear systems of the TLN is nondegenerate and thus has a unique fixed point. Another requirement is that at least one $b_i > 0$, which guarantees that 0 is never a fixed point of the TLN. A key feature of nondegenerate TLNs is that they are guaranteed to have at most one fixed point per support (see Corollary 5.3). Note that almost all networks of the form (1) are nondegenerate, since having a zero determinant is a highly fine-tuned condition.

Although these networks have been around for decades in the neural networks community, the mathematical theory is still a work in progress. It began in earnest about 20 years ago, with work by Hahnloser, Seung, and others [3, 4, 21]. Theirs was the first serious attempt to develop a mathematical theory of threshold-linear networks to rival that of Hopfield networks [1]. Not surprisingly, the initial results were confined to the case where W is a symmetric matrix. In [3], precise conditions were found to guarantee that network activity always converges to a stable fixed point, and a characterization was given of symmetric threshold-linear networks exhibiting multistability. In [4], the role of lateral inhibition was highlighted as playing an especially important role, enabling selective competition between groups of neurons.

2.1. Competitive TLNs. A *competitive* TLN is a special case of (1) with the additional restrictions: $W_{ij} \le 0$, $W_{ii} = 0$, and $b_i \ge 0$ for all $i, j = 1,\dots,n$. *Unless otherwise specified, we will assume all TLNs are both* competitive *and* nondegenerate. Note that nondegeneracy implies that not all b_i can be 0, and so despite the competitive dynamics, trajectories cannot decay to the origin.

[1] If the nodes are neurons, the activity level is typically called a "firing rate."

An important property of competitive networks is that they are guaranteed to have bounded activity. In fact, we can prove that the asymptotic dynamics are confined to the box

$$\mathcal{B} \stackrel{\text{def}}{=} \prod_{i=1}^{n} [0, b_i].$$

Lemma 2.1. *The box $\mathcal{B} = \prod_{i=1}^{n}[0,b_i]$ is a globally attracting set of the competitive TLN dynamics (1). That is, if $x(0) \in \mathcal{B}$, then $x(t) \in \mathcal{B}$ for all $t > 0$. Moreover, if $x(0) \notin \mathcal{B}$, the trajectory $x(t)$ approaches or enters $\mathcal{B}$ as $t \to \infty$.*

Proof. To show that trajectories that start inside $\mathcal{B}$ stay there, it suffices to show that for each x_i, the derivative dx_i/dt on the boundary $\partial\mathcal{B}$ is either 0 or it points inside the box. That is, for $x \in \mathcal{B}$ we must have $dx_i/dt \geq 0$ whenever $x_i = 0$ and $dx_i/dt \leq 0$ whenever $x_i = b_i$. Now observe that $x \in \mathcal{B}$ implies $x_i \geq 0$ for each i; putting this together with the condition that $W_{ij} \leq 0$, we have

$$\left[\sum_{j=1}^{n} W_{ij}x_j + b_i\right]_+ \leq b_i.$$

From here we immediately see that for all $x \in \mathcal{B}$ (in fact, for all $x \geq 0$),

$$-x_i \leq \frac{dx_i}{dt} \leq -x_i + b_i.$$

Evaluating the derivatives on the boundary, we see that $dx_i/dt \geq 0$ at $x_i = 0$, and $dx_i/dt \leq 0$ at $x_i = b_i$.

Next, consider what happens for $x(0) \notin \mathcal{B}$. If $x_i \notin [0,b_i]$ for some i, then either $x_i < 0$ and thus $dx_i/dt > 0$, or $x_i > b_i$ and thus $dx_i/dt < 0$. If follows that any trajectory $x(t)$ initialized outside of $\mathcal{B}$ will either approach or enter $\mathcal{B}$ as $t \to \infty$. ∎

In particular, Lemma 2.1 guarantees that $\mathcal{B}$ contains all the fixed points.

Corollary 2.2. *All fixed points of a competitive TLN (W,b) lie inside $\mathcal{B} = \prod_{i=1}^{n}[0,b_i]$.*

This fact is also easy to see directly: at a fixed point, we must have $x_i^* = [\sum_{j=1}^{n} W_{ij}x_j^* + b_i]_+ \geq 0$ for each i, which immediately implies $0 \leq x_i^* \leq b_i$, since $W_{ij} \leq 0$.

2.2. The set of all fixed points FP(W,b). Recall that nondegenerate TLNs have at most one fixed point per support (for the proof, see Corollary 5.3). We can thus label all the fixed points of a given network by their supports. We denote this as

$$\text{FP}(W,b) \stackrel{\text{def}}{=} \{\sigma \subseteq [n] \mid \sigma = \operatorname{supp} x^* \text{ for some fixed point } x^* \text{ of the associated TLN}\},$$

where

$$[n] \stackrel{\text{def}}{=} \{1, \ldots, n\}.$$

Note that for each support $\sigma \in \text{FP}(W,b)$, the fixed point itself is easily recovered. Outside the support, $x_i^* = 0$ for all $i \notin \sigma$. Within the support, x^* is given by

$$x^*_\sigma = (I - W_\sigma)^{-1} b_\sigma,$$

where x^*_σ and b_σ are the column vectors obtained by restricting x^* and b to the indices in σ, and W_σ is the induced principal submatrix obtained by restricting rows and columns of W to σ. Note that $x^* \in \mathcal{B}$ (see Corollary 2.2), though it is not obvious from this above formula.

In [3], the authors studied the collection of *stable* fixed points of these networks through the lens of *permitted sets*; these are subsets $\sigma \subseteq [n]$ for which the network (W, b) supports a stable fixed point for at least one $b \in \mathbb{R}^n$. The theory of permitted sets was further developed in [5, 6]. In particular, it was shown that a fixed point with support σ is stable if and only if all eigenvalues of $(-I + W)_\sigma$ have negative real part (this was first shown for symmetric matrices in [3] and then generalized in [6]). Finally, in [7], attention was shifted to the study of stable fixed points that can be simultaneously realized for a single uniform external input b. In this context, there is additional mathematical tractability that allows us to generalize to asymmetric networks. Thus, we will pay special attention to the case of uniform inputs, where $b_i = \theta > 0$ for all i. In order to further isolate the role of network connectivity in shaping dynamics, we will also consider a special family of TLNs, known as CTLNs. These are networks where the matrix W has binary synapses, corresponding to strong and weak inhibition, dictated by an underlying connectivity graph.

2.3. The CTLN model. The self-inhibition (or decay rate) of a single node, which has been normalized to -1, provides a natural scale for the strength of inhibition. If $W_{ij} > -1$, then node j inhibits node i *less* than it inhibits itself. If, on the other hand, $W_{ij} < -1$, the inhibition from j to i is stronger than the self-inhibition of j. This distinction ends up playing an important role in shaping the dynamics.

To simplify our study, and further isolate the role of connectivity, we specialize to competitive networks with only two values for the connection strengths: one value for $W_{ij} < -1$ and another for $W_{ij} > -1$. To any simple directed graph G on n vertices, and for any $0 < \varepsilon < 1$ and $\delta > 0$, we can associate a corresponding $n \times n$ connectivity matrix $W = W(G, \varepsilon, \delta)$ as follows:

$$W_{ij} = \begin{cases} 0 & \text{if } i = j, \\ -1 + \varepsilon & \text{if } i \leftarrow j \text{ in } G, \\ -1 - \delta & \text{if } i \not\leftarrow j \text{ in } G, \end{cases} \tag{3}$$

where $i \leftarrow j$ indicates that G has an edge from j to i, and $i \not\leftarrow j$ indicates that there is no such edge. Clearly, W satisfies the conditions of a competitive network.

We refer to TLNs of the form (1), with $b_i = \theta > 0$ and $W = W(G, \varepsilon, \delta)$ as in (3), as the *combinatorial threshold-linear network* (CTLN) model. Note that this model is completely specified by the choice of directed graph, G, along with three positive real parameters: ε, δ, and θ (see Figures 1A–C).

In order to ensure that the graph G is most meaningful for the dynamics, it is useful to require that any pair of nodes i, j with a single unidirectional edge $i \to j$ (but $j \not\to i$) cannot support a stable fixed point, so that activity will flow along directed edges. This is guaranteed

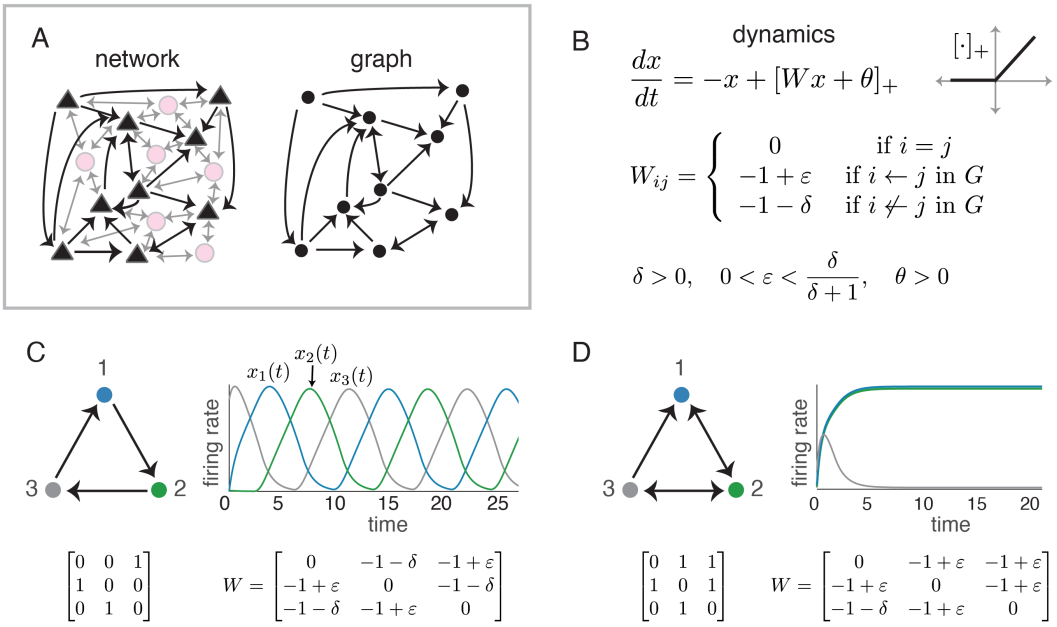

Figure 1. *Combinatorial threshold-linear networks with strong and weak inhibition.* (A) *(Left) A neural network with excitatory pyramidal neurons (triangles) and a background network of inhibitory interneurons (pink circles) that produce a global inhibition. (Right) The corresponding graph retains only the excitatory neurons and their connections.* (B) *The equations for a CTLN network.* (C) *(Left) An oriented graph on three nodes, with the corresponding adjacency matrix below. (Right) Network activity follows the arrows in the graph, with peak activity occurring sequentially in the cyclic order* 123. (D) *Cliques correspond to stable fixed points but only if they are* target-free *cliques (see Theorem* 3.5*). The clique* 12 *is target-free, but* 23 *is not. Unless otherwise noted, all simulations have parameters* $\varepsilon = 0.25$, $\delta = 0.5$, *and* $\theta = 1$.

if the 2×2 matrix $(-I + W)_{\{i,j\}}$ for the directed edge $i \to j$ is always unstable. Specifically, the eigenvalues of the matrix

$$(-I + W)_{\{i,j\}} = \begin{pmatrix} -1 & -1-\delta \\ -1+\varepsilon & -1 \end{pmatrix}$$

should not both have strictly negative real part. Since the trace is negative, this occurs precisely when $\det(-I + W)_{\{i,j\}} = \varepsilon(1+\delta) - \delta < 0$. We thus define the *legal range* for a CTLN as follows.

Definition 2.3. *We say that the CTLN parameters* $\varepsilon, \delta, \theta$ *are in the* legal range *if* $\varepsilon, \delta, \theta > 0$ *and* $\varepsilon < \delta/(\delta+1)$.

We interpret a CTLN as modeling a network of n excitatory neurons in the presence of strong background inhibition (Figure 1A). When $j \not\to i$, we say j *strongly inhibits* i; when $j \to i$, we say j *weakly inhibits* i. The strong inhibition is just the global background inhibition, while the weak inhibition can be thought of as the sum of an excitatory connection and global inhibition. These differences in inhibition promote the flow of activity along the arrows of the graph.

We have seen that competitive TLNs have bounded activity that remains confined to a box, $\prod_{i=1}^{n}[0,b_i]$ (Lemma 2.1). For CTLNs, the corresponding box is $[0,\theta]^n$. What can we say about the dynamics inside this box? In the simplest cases, where all interactions are equal so that $W_{ij}=w$ for all $i\neq j$, we have two extremes: $w=-1-\delta$ or $w=-1+\varepsilon$. If inhibition is strong, so that $w=-1-\delta<-1$, we obtain a classical winner-take-all (WTA) network. Such a network has n stable fixed points, one corresponding to each node, and no other attractors. The activity always converges to the fixed point of the "winning" neuron. If, on the other hand, inhibition is weak, so that $w=-1+\varepsilon>-1$, then the network synchronizes and activity always converges to a single stable fixed point in which all nodes have equal activity. Neither case is particularly exciting.

When a competitive network has *both* strong and weak inhibition, things get more interesting, particularly when the interactions are asymmetric. Such networks can exhibit a variety of stable and unstable fixed points, and dynamic attractors appear to correspond to a key subset of these unstable fixed points [13]. Thus, we are motivated to understand the full collection of fixed points of a CTLN (or a TLN), not only those that are stable. Moreover, we are particularly interested in networks with *no* stable fixed points, as these will be guaranteed to have dynamic attractors.

3. Mathematical results.

We begin with some results that hold for the more general class of nondegenerate competitive TLNs. Theorem 3.1 is an index theorem for the fixed points of a TLN, which has a number of important consequences. In particular, it implies that the total number of fixed points is always odd and provides an upper bound on the number of stable fixed points. Theorem 3.3 gives conditions on the connectivity matrix W that guarantee that the TLN has no stable fixed points. Specializing to CTLNs, we obtain Theorem 3.4, which tells us that oriented graphs with no sinks yield networks with no stable fixed points. At the other extreme, Theorem 3.5 shows that *cliques* correspond to stable fixed points (provided that the embedding is *target-free*). Moreover, we conjecture that these are the only subgraphs that correspond to stable fixed points. These results are illustrated with numerous examples in section 4. The proofs are postponed to section 5.

3.1. Index and parity.

For each fixed point of a competitive TLN (W,b), labeled by its support $\sigma\in\mathrm{FP}(W,b)$, we define the *index* as

$$\mathrm{idx}(\sigma)\overset{\mathrm{def}}{=}\mathrm{sgn}\det(I-W_\sigma).$$

Since we assume our TLNs are nondegenerate, we have $\det(I-W_\sigma)\neq 0$ and thus $\mathrm{idx}(\sigma)\in\{\pm 1\}$. Note that if σ is the support of a *stable* fixed point, then the eigenvalues of $-I+W_\sigma$ must all have negative real part, and so those of $I-W_\sigma$ all have positive real part. This implies that $\mathrm{idx}(\sigma)=+1$ for all stable fixed points.

The following theorem indicates that fixed points with index $+1$ and -1 are almost perfectly balanced. It also tells us that the parity of the total number of fixed points is always odd.

Theorem 3.1 (parity). *Let (W,b) be a competitive nondegenerate TLN on n nodes, with $b_i>0$ for all $i\in[n]$. Then*

$$\sum_{\sigma \in \mathrm{FP}(W,b)} \mathrm{idx}(\sigma) = +1.$$

In particular, the total number of fixed points $|\mathrm{FP}(W,b)|$ *is always odd.*

As an immediate corollary, we obtain an upper bound on the number of stable fixed points. Note that $|\mathrm{FP}(W,b)| \leq 2^n - 1$, since nondegenerate TLNs have at most one fixed point per support (Corollary 5.3) and the origin (empty support) is never a fixed point. Because stable fixed points all have index $+1$, their upper bound is at most half of $2^n - 1$.

Corollary 3.2. *The number of stable fixed points of a competitive nondegenerate TLN on* n *nodes, with* $b_i > 0$ *for all* $i \in [n]$*, is at most* 2^{n-1}.

3.2. Competitive networks with no stable fixed points. Historically, much of the mathematical theory of TLNs has been developed in the symmetric case with a focus on networks with guaranteed convergence to stable fixed points [3, 4, 21, 5, 6, 7]. These networks can serve as models of associative memory storage and retrieval, similar to Hopfield networks, with static memory patterns encoded as stable fixed points. But the brain also exhibits *dynamic* memory patterns, such as rhythms and sequences. This type of activity is more naturally modeled with dynamic attractors, such as limit cycles and higher-dimensional attractors, not stable fixed points. Can we find conditions that guarantee that a TLN only exhibits dynamic attractors?

Our next result establishes sufficient conditions such that a competitive network has no stable fixed points (i.e., no steady states). The statement of the theorem makes use of a simple[2] directed graph G_W on n vertices that is defined from the $n \times n$ connectivity matrix W as follows:

$$G_W \text{ has an edge from } j \to i \Leftrightarrow W_{ij} > -1 \text{ (for } i \neq j). \tag{4}$$

Note that W_{ij} represents the influence of node j on node i. The edges of G_W correspond to inhibitory interactions that are weaker than the self-inhibition of each node. If W is the matrix of a CTLN with graph G, then $G_W = G$.

The next theorem uses the following graph-theoretic terminology: a graph is *oriented* if there are no bidirectional edges, and a *sink* is a vertex with no outgoing edges. The proof is given in section 5.

Theorem 3.3. *Consider a competitive nondegenerate TLN with connectivity matrix* W*, associated graph* G_W*, and uniform inputs* $b_i = \theta$*. Suppose that*

(i) G_W *is an oriented graph with no sinks, and*

(ii) *whenever* $j \to i$ *in* G_W*,* $W_{ij} < \dfrac{1}{W_{ji}}$.

Then the network (1) *has no stable fixed points. Moreover, the network activity is bounded.*

[2]A graph is *simple* if it has no multiple edges and no self-loops. Simple graphs have binary adjacency matrices with zeros on the diagonal.

In the case of a CTLN, it is easy to see that condition (ii) of Theorem 3.3 is always satisfied, provided that the parameters are within the *legal range,* so that $\varepsilon < \delta/(1+\delta)$. In fact, this was one of our motivations for the definition of legal range. We thus have the following result, obtained by specializing Theorem 3.3 to the CTLN case.

Theorem 3.4. *Let G be an oriented graph with no sinks, and consider an associated nondegenerate CTLN with $W = W(G, \varepsilon, \delta)$ for ε and δ in the legal range. Then the network has bounded activity and no stable fixed points.*

Figure 1C displays the smallest oriented graph with no sinks, together with the activity of the corresponding CTLN model. The solutions to (1) for this W always converge to the same perfectly periodic trajectory, irrespective of the initial conditions. It is surprisingly nontrivial to explicitly prove that this limit cycle exists. A recent proof was given in [12] for the existence of a unique limit cycle in CTLNs whose graph is a 3-cycle, as in Figure 1C, as well as a k-cycle more generally.

3.3. Stable fixed points in the CTLN model. Theorem 3.4 gave conditions on the graph G of a CTLN that guaranteed the absence of stable fixed points. In this section, we present a theorem and a conjecture relating certain graph structures to the existence of stable fixed points.

To state the next theorem, we need a few graph-theoretic definitions. A subset of vertices σ is a *clique* of G if the nodes in σ are all-to-all bidirectionally connected, i.e., $i \leftrightarrow j$ for all pairs $i, j \in \sigma$. We say that a vertex k is a *target* of σ if $k \notin \sigma$ and $i \to k$ for each $i \in \sigma$. If the clique σ has no targets, we say that it is *target-free.* A clique is *maximal* if it is not contained in any larger clique of G. Note that all target-free cliques are necessarily maximal, but maximal cliques need not be target-free. For example, the graph in Figure 1D has two maximal cliques, but only one of them is target-free. It turns out that the only cliques that can support stable fixed points are target-free cliques.

Theorem 3.5. *Let G be a simple directed graph, and consider an associated nondegenerate CTLN with $W = W(G, \varepsilon, \delta)$ for any choice of the parameters $\varepsilon, \delta, \theta > 0$ with $\varepsilon < 1$. If σ is a clique of G, then there exists a stable fixed point with support σ if and only if σ is target-free.*

The proof is given in section 5. Note that the result is valid for parameters beyond the legal range, including $\delta/(\delta+1) \leq \varepsilon < 1$.

While Theorem 3.5 identifies precisely which cliques support stable fixed points, there may still exist additional stable fixed points that do not correspond to cliques. We have not, however, been able to find any such example. This leads us to the following conjecture.

Conjecture 3.6. *Consider a nondegenerate CTLN model $W = W(G, \varepsilon, \delta)$, where G is a simple directed graph and ε and δ are within the legal range. There exists a stable fixed point with support σ if and only if σ is a target-free clique of G.*

The conjecture was previously proven to hold in the case of symmetric W, i.e., when all the edges in G are bidirectional [7]. Note that in that context, a clique is target-free precisely when it is maximal. Outside of the symmetric case, the conjecture has been proven in a few other settings, such as for graphs on $n \leq 4$ nodes or within a particular region of the legal range of parameters ε and δ. These results and all other current evidence in support of the conjecture are collected in [22].

4. Simulations. The CTLN model captures a surprisingly rich diversity of nonlinear dynamics. In this section, we provide a variety of examples to illustrate both this diversity and the theorems from the previous section. Unless otherwise noted, all simulations have parameters $\varepsilon = 0.25$, $\delta = 0.5$, and $\theta = 1$. These are referred to as the *standard parameters*. When the parameters are identical, differences in network dynamics arise solely as a result of differences in the underlying graph G.

The MATLAB package CTLN Basic 2.0 can be used to reproduce the simulations in Figures 1–13. In particular, graphs and initial conditions are provided for each of the figures, with the exception of Figure 9. The package is available online from https://github.com/nebneuron/CTLN-Basic-2.0.

4.1. Dynamic diversity from network connectivity. Figure 2 displays adjacency matrices for three different graphs on $n = 25$ nodes, along with two-dimensional projections of solutions that are periodic (A), chaotic (B), and quasi-periodic (C). Since each of these graphs is oriented with no sinks, Theorem 3.4 guarantees that the networks have bounded activity but no stable fixed points. Note, however, the large variety of dynamics that can arise by varying the graph structure. For example, the quasi-periodic behavior in Figure 2C is shaped by a highly structured graph, and this graph differs markedly from the ones in Figures 2A, B.

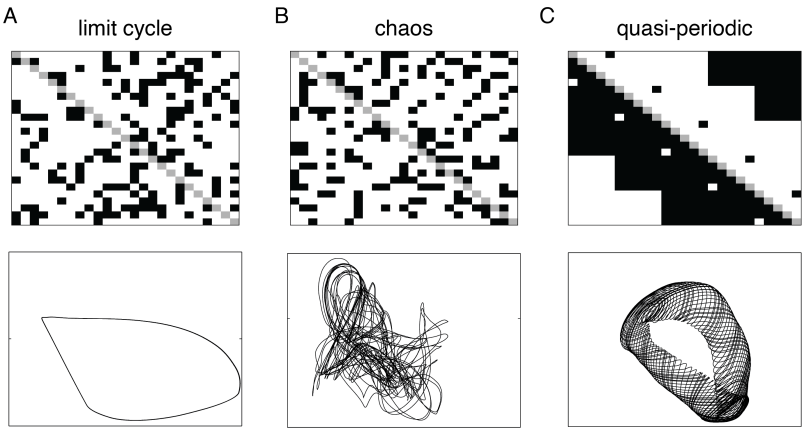

Figure 2. *Nonlinear dynamics of the CTLN model. Top panels show adjacency matrices for three oriented graphs on $n = 25$ nodes with no sinks, satisfying the conditions of Theorem 3.4. (Black = 1, white = 0, and gray = diagonal elements, which are ignored.) The corresponding CTLN model networks produce* (A) *a limit cycle,* (B) *a chaotic attractor, and* (C) *quasi-periodic behavior. Bottom panels show random two-dimensional projections of the 25-dimensional trajectories.*

What aspects of the connectivity graph G determine the emergent dynamics of the network? Are local properties of the connectivity matrix sufficient to predict the behavior that will emerge? One possibility is that the degree profile, that is, the list of in-degrees and out-degrees of each node, is predictive.

Figure 3 makes it clear that the degree profile is insufficient to predict the patterns of activity that emerge. The four networks shown in panels A–D have graphs on $n = 5$ nodes with exactly the same degree profile: $\{(2,2),(2,1),(2,1),(1,2),(1,2)\}$, where (a,b) indicates in-degree a and out-degree b of a single node. Despite having identical local properties,

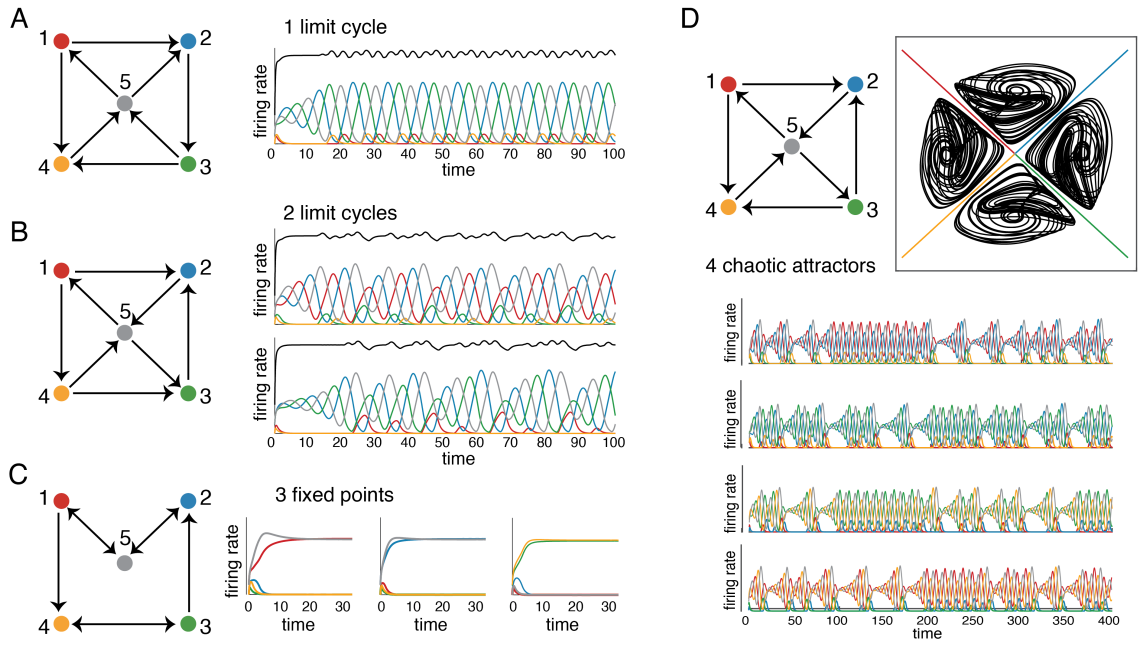

Figure 3. *Degree-matched networks with different dynamics.* (A–D) *Four graphs are shown together with the attractors that emerge from their corresponding CTLNs. Each activity plot shows a solution arising from a different initial condition, where the colors match the nodes in the graph. Networks* A–D *have precisely the same degree profile but significantly different dynamics: limit cycles* (A, B)*, stable fixed points* (C)*, and strange/chaotic attractors* (D)*. A two-dimensional projection of the four strange/chaotic attractors is shown in* D.

dramatically different dynamics arise: limit cycles in networks A and B, stable fixed points in network C, and strange/chaotic attractors in network D. These CTLNs thus exhibit *emergent dynamics* in the strongest sense, as differences in activity depend only on the global structure of the network and not on local properties of individual nodes (which are identical across networks A–D). Note that the graphs in panels A, B, and D are all oriented, and so the absence of stable fixed points is guaranteed by Theorem 3.4. In contrast, the graph in C is not oriented; its stable fixed points are fully predicted by Theorem 3.5.

The networks in Figures 3B–D also exhibit multiple co-existing attractors, with different initial conditions leading to different patterns of activity. For example, network B has two distinct limit cycles, each of which can be accessed by changing initial conditions, without changing any network parameters. Similarly, network C has three stable fixed points, and network D has four strange/chaotic attractors. While each of these networks possesses multiple attractors, all attractors within a given network are of the same type. Is it possible for different types of attractors to coexist within a single network?

The network in Figure 4 shows that different types of attractors can co-exist: two stable fixed points, a limit cycle, and a chaotic attractor. Once again, the selected attractor depends only on the choice of initial conditions, without any change in network parameters. And the stable fixed points can again be predicted using Theorem 3.5.

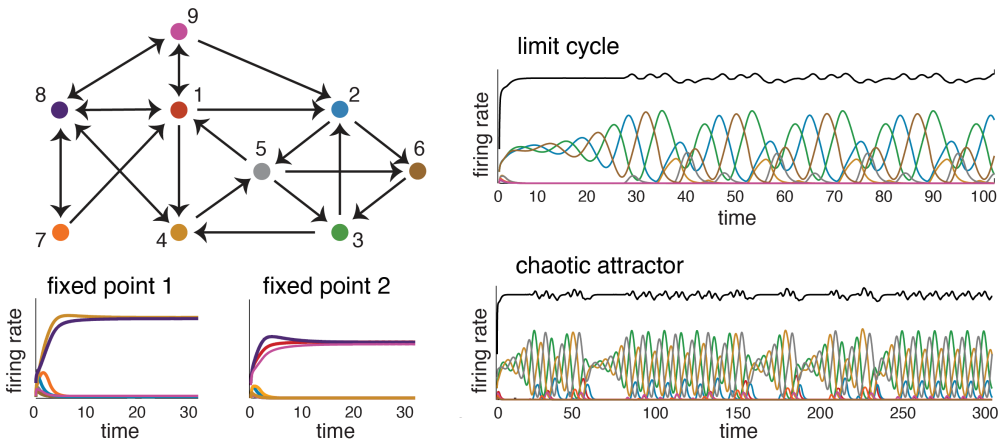

Figure 4. *Variety of emergent dynamics in a single network with* $\boldsymbol{n = 9}$ *nodes. Solutions corresponding to different initial conditions are shown. The network has four attractors: two stable fixed points, one limit cycle, and one chaotic attractor. The equations for the dynamics are identical in each case; only the initial conditions differ.*

4.2. Emergent sequences. In Figure 1C, Figures 3A–B, and Figure 4, we saw limit cycles in which nodes were activated in a regular sequence. Sequential patterns of activity are in fact quite common in CTLNs, as competition between nodes results in a tendency for neurons to "take turns" in reaching their peak firing. The emergent sequences, however, are often irregular and surprising. Figure 5 displays a network with $n = 7$ nodes and a single emergent sequence (irrespective of initial conditions). Although the graph has many cycles, the sequential activity follows only one cycle in the graph. Moreover, the activity of node 2 always decays to zero, even though there are other nodes (3 and 5) with smaller in-degree whose activity persists in the attractor. This provides another example where local properties of the graph are not sufficient to predict emergent dynamics; the resulting sequence has been shaped by the structure of the graph as a whole.

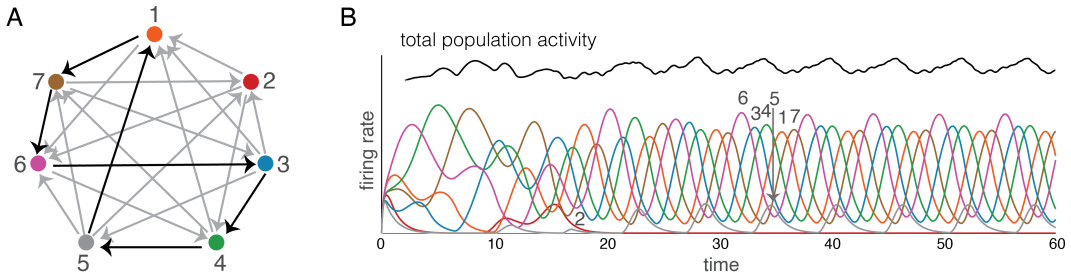

Figure 5. *Emergence of an irregular sequence.* (A) *A graph on seven nodes with a particular cycle highlighted in black. The CTLN model makes no distinction between black and gray edges, but black edges are highlighted here, because they correspond to the emergent sequence of activation.* (B) *Node 2 decays to zero after a short period of transient activity, while the remaining nodes settle into a limit cycle of sequential activation with ordering* 634517*. The same sequence emerges irrespective of initial conditions and is robust to small perturbations of the matrix* W*.*

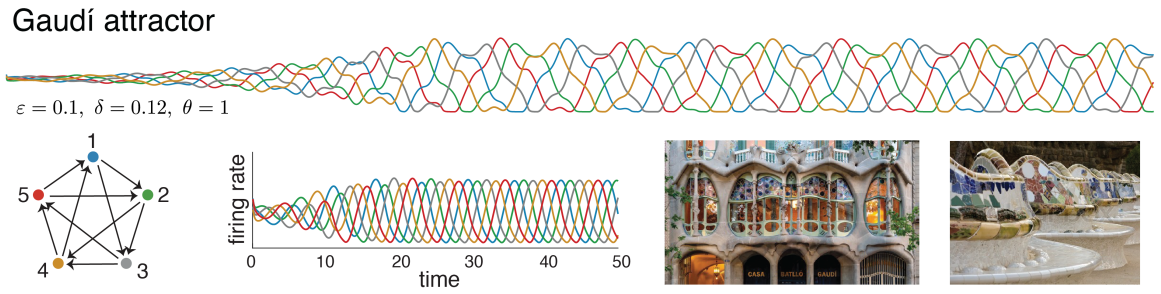

Figure 6. *The Gaudi attractor. A CTLN for a cyclically symmetric tournament on $n = 5$ nodes produces two distinct attractors, depending on initial conditions. We call the top one the Gaudi attractor, because the undulating curves are reminiscent of the architect from Barcelona. Note the different parameters for this CTLN: $\varepsilon = 0.1$ and $\delta = 0.12$.*

Figure 6 shows the dynamic attractors of a CTLN whose graph is a cyclically symmetric tournament on $n = 5$ nodes. For some initial conditions, the dynamics converge to a somewhat boring limit cycle with the firing rates $x_1(t), \ldots, x_5(t)$ all peaking in the expected sequence, 12345 (bottom middle). However, for a different set of initial conditions, the solution converges to the beautiful but unusual attractor displayed at the top.

Similar behavior emerges from a CTLN whose graph is a generalization of this cyclic tournament structure to $n = 7$ nodes (see Figure 7). For some initial conditions, a simple limit cycle emerges with the expected sequence of peaks 1234567 (middle left). For different initial conditions, another attractor emerges (middle right), which is quasi-periodic with a torus-like trajectory shown on the right. Interestingly, this CTLN has a unique fixed point, which is unstable and has all nodes active, and both of these attractors are accessible from perturbations of that fixed point. This same phenomenon holds for the CTLN in Figure 6 as well.

The dynamics in Figure 7 are reminiscent of traveling wave solutions in bump attractor networks [23, 24, 25, 26]. The traveling wave interpretation is natural if we consider the nodes as spatially organized as in a ring model. In classic ring models, the connectivity strength typically drops off with distance and the connectivity matrix is symmetric. In contrast, the

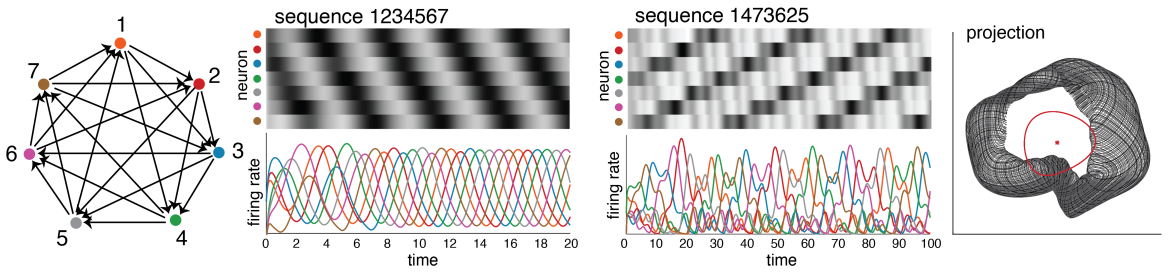

Figure 7. *A cyclically symmetric CTLN on $n = 7$ nodes. (Left) The graph is a cyclically symmetric tournament on $n = 7$ nodes. (Middle) The corresponding CTLN has two attractors: one limit cycle with sequence* 1234567 *and one quasi-periodic attractor whose sequence corresponds to a different cycle in the graph. (Right) A two-dimensional projection of the two attractors and the unique (unstable) fixed point of the network. The limit cycle and fixed point are displayed in red, while the quasi-periodic attractor is the torus-like trajectory shown in black.*

network in Figure 7 only has "forward edges" around the cycle and is thus far from symmetric. Nevertheless, due to the cyclic symmetry, one observes a bump of activity traveling around the cycle. Interestingly, the second attractor of this network has a different sequence of activity that does not respect the same geometric organization. In this solution, the activity jumps between nodes that are neighbors in the graph but not neighbors in the ring. However, using a different geometric arrangement of the nodes, so that ordering on the ring is 1473625, the activity may again be thought of as a traveling wave.

4.3. Complex rhythms. Limit cycles need not be perfectly sequential; they can also display complex rhythms, including synchronous or quasi-synchronous activity for a subset of the nodes. The network in Figure 8 has two high-activity nodes (3 and 6) that peak at different times, while nodes 2, 4, and 7 are approximately synchronous. Note that node 6 has the highest peak activity, even though it has the lowest in-degree among all nodes in the network.

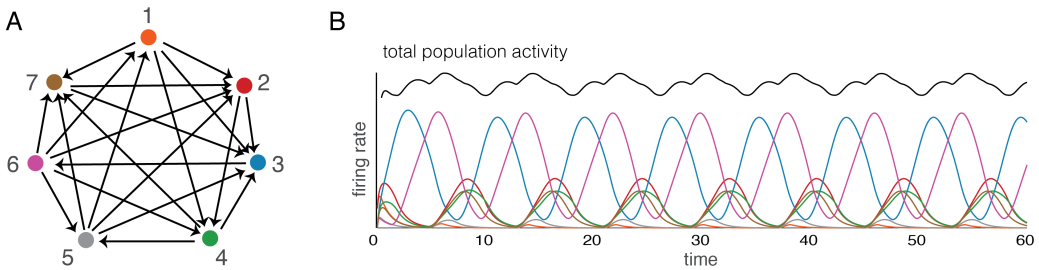

Figure 8. *Emergence of a complex rhythm. A graph on seven nodes (left) yields a CTLN model whose activity always settles into the same limit cycle (right). The activity in this limit cycle is a rhythmic, with some nodes that are quasi-synchronous rather than sequential in their activation.*

Using Theorem 3.4, we can also generate large random networks that are guaranteed never to settle into a steady state and thus possess interesting dynamic attractors. Such networks can exhibit spontaneous transitions between distinct patterns of network activity. In Figure 9, the total population activity trace has a sharp qualitative change around $t = 80$;

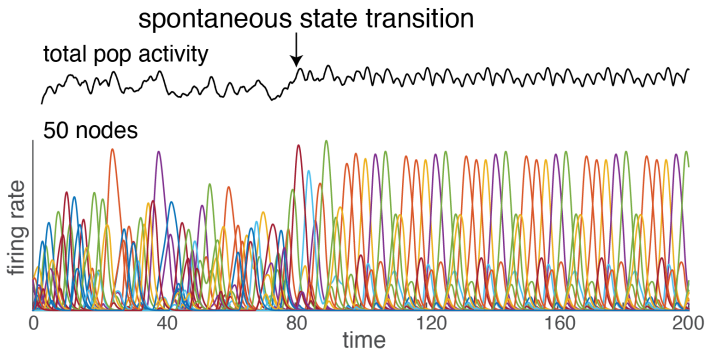

Figure 9. *Spontaneous state transition. A random network of* 50 *nodes satisfying the conditions of Theorem* 3.4 *exhibits a spontaneous state transition from irregular to periodic behavior. Total population activity (black trace) is the sum of the activity for all* 50 *nodes.*

this is reminiscent of state transitions in cortical networks observed during light anesthesia and sleep [27].

The simulations in Figures 1–9 yield attractor types that are not exclusive to CTLNs; similar behaviors have been observed in many other nonlinear systems. What is remarkable is that they can all be produced within the same model family, and with the same parameters, by varying only a simple graph of connectivity. Moreover, this model family (TLNs/CTLNs) is surprisingly mathematically tractable. Indeed, it is this tractability that enabled the discovery of such a variety of dynamics without computationally intensive searches of the TLN parameter space. For example, the chaotic attractors in Figure 3D are reminiscent of the Rössler attractor but were naturally discovered by exploring small oriented graphs with no sinks. The Gaudi attractor in Figure 6 was also discovered in this way and appears to be novel.

4.4. Network engineering. The power of Theorems 3.1, 3.4, and 3.5 is that they enable us to reason mathematically about the graph G in order to make strong yet accurate predictions about the resulting network dynamics. Such mathematical results are thus valuable tools for designing networks with prescribed dynamic properties.

Rhythmic patterns of activity, supporting locomotion and other functions, arise in central pattern generator circuits (CPGs) throughout the nervous system [28, 29]. The CTLN model provides a natural framework for CPGs. For example, Figure 10 shows that limit cycles corresponding to two different quadruped gaits, "bound" (similar to gallop) and "trot," can co-exist in the same network, with the network selecting one pattern over the other based solely on initial conditions.

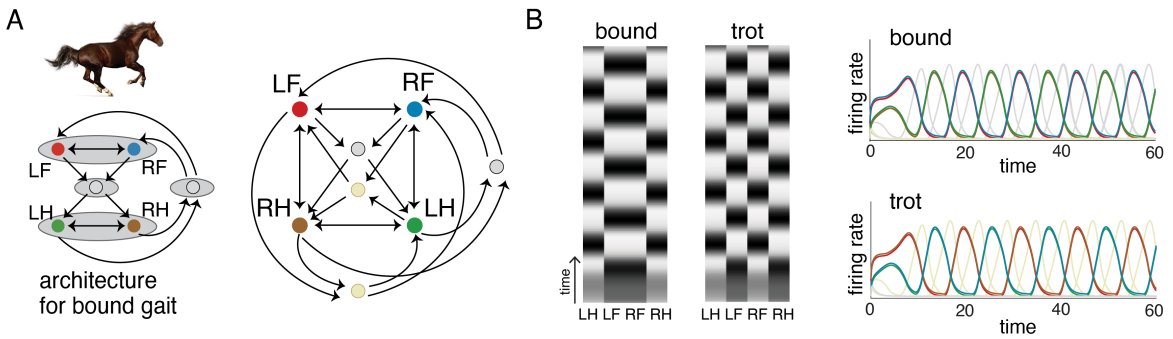

Figure 10. *A central pattern generator circuit for quadruped motion.* (A) *(Left) A* cyclic union *architecture on six nodes that produces the "bound" gait. (Right) The graph on eight nodes is formed from gluing together architectures for the individual gaits, "bound" and "trot." Note that the positions of the two hind legs (LH, RH) are flipped for ease of drawing the graph.* (B) *The network produces limit cycles corresponding to two distinct gaits. Convergence to one or the other is determined only by initial conditions.*

The network in Figure 10 was produced by essentially "gluing together" two architectures that would produce the desired gaits, identifying the graphs along the nodes corresponding to each leg. The left panel in A shows the isolated graph that produces the single gait, "bound." Notice that the clique between the left and right front legs (LF and RF) ensures that those nodes co-fire; the same is true for the clique between the left and right hind legs (LH and RH). But since each of these cliques has a target, neither of them can a support stable fixed

point (by Theorem 3.5). The activity simply flows forward from the clique to one of the (gray) auxiliary nodes and then to the next clique in the cycle. A similar network was created for the "trot" gait, with appropriate pairs of legs joined by a clique. The larger graph in Figure 10A is the result of gluing together the smaller graphs for these two gaits.

One key to the construction in Figure 10 was the use of cliques with targets to ensure a desired flow of activity. More generally, Theorem 3.5 enables us to construct networks that continually transition between cliques without getting "stuck," since we know precisely which cliques in a graph correspond to steady states. Figure 11A depicts a network with a series of overlapping cliques, only the last of which is target-free. We can also engineer networks by patching together modules that individually yield limit cycles, rather than cliques. Figure 11B depicts a network that has six overlapping limit cycles, corresponding to subsets of nodes 1–5, 4–8, 7–11, 10–14, 13–17, and 16–20. The network activity will stay in a single limit cycle

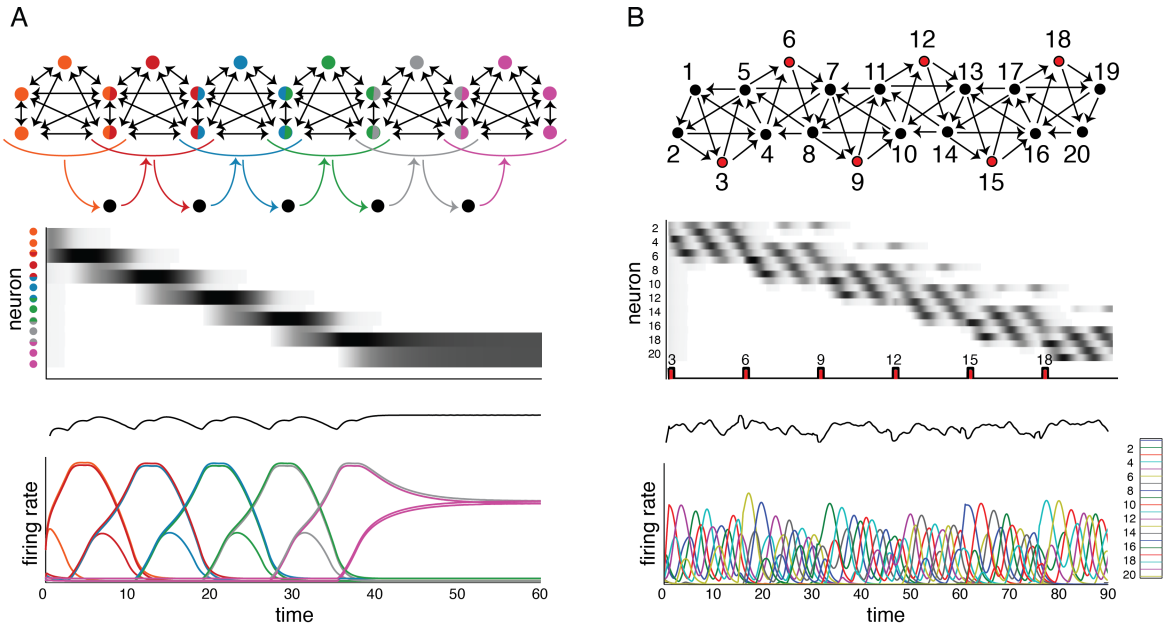

Figure 11. *Regular and irregular sequential activity from modular architectures.* (A) *(Top) A chain of six overlapping 5-clique modules, each in a different color. Nodes belonging to two adjacent 5-cliques are double-colored. Black nodes receive output edges from all nodes in one 5-clique and feed forward onto all nodes in the next 5-clique. Note that only the final 5-clique (purple) is target-free. (Middle) The solution of the network when the nodes in the first 5-clique are initialized to 0.1 and all other nodes are initialized to 0. Darker regions correspond to higher firing rates; note that the outer target nodes are omitted. (Bottom) Same solution as in middle plot. The firing rate curves are colored by the nodes in the graph, with overlap nodes receiving both colors. The activity moves slowly from one clique to the next until it stabilizes on the last clique, which is target-free. The total population activity is given by the black trace above.* (B) *(Top) A chain of six overlapping modules, where each module is a cyclic tournament on five nodes (the same graph as in Figure 6). Nodes highlighted in red receive small kicks during the simulation. (Middle) The solution of the network when nodes 1 and 2 are initialized to 0.1 and all other nodes are initialized to 0. Initially, the network activity is confined to the first module and cycles among those nodes. Every 15 time units, a small kick is given to the middle node in the next module. The timing of these kicks and the label of the affected node are shown as red pulses along the bottom of the plot. (Bottom) Same solution as in the middle plot, with individual firing rates shown in color.*

indefinitely, unless it receives an external "kick" helping it to transition to an adjacent limit cycle.

Our next example shows that a single network, with very simple architecture, can have multiple quasi-periodic attractors. The network in Figure 12A has n nodes and $n-2$ quasi-periodic attractors. Each quasi-periodic orbit selects a single node from the inner (shaded) region, which forms a dynamic sequence with the blue and gray outer nodes (Figure 12B).

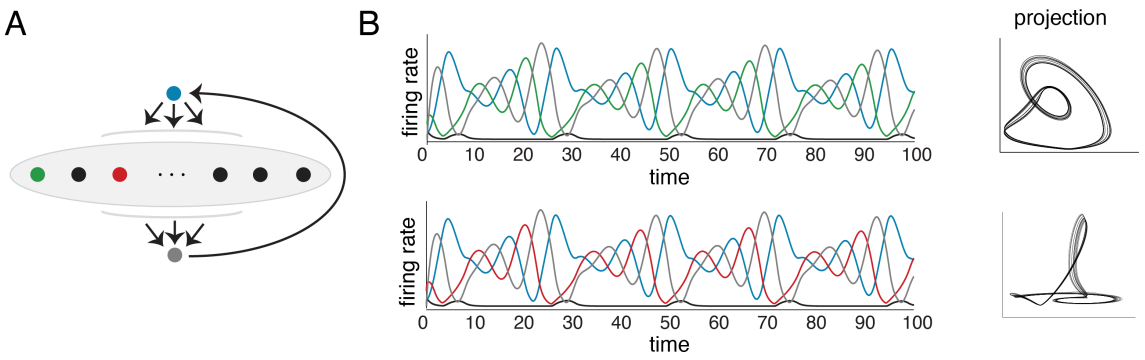

Figure 12. *Multiple quasi-periodic attractors in the same network.* (A) *A graph on n nodes, where there are $n-2$ nodes in the middle (inner) layer. The top (blue) outer node feeds onto all nodes in the middle layer, while the bottom (gray) outer node receives connections from all nodes in the middle layer and feeds back onto the top node. This architecture produces $n-2$ different quasi-periodic attractors that each involve the top node, one of the middle nodes, and the bottom node.* (B) *Two distinct quasi-periodic attractors, one involving the green middle node (top) and the other involving the red middle node (bottom). In each solution, the activity of all other middle nodes is small and synchronous and is depicted in black. Random projections of the activity (right) indicate that they are quasi-periodic trajectories and not perfect limit cycles.*

Our last example, given in Figure 13, is a generalization of the previous network: now it has m nodes in each layer, all-to-all forward connections from one layer to the next, and the last layer connecting back to the first in a cyclic fashion. For parameters $\varepsilon = 0.75$, $\delta = 4$, the CTLN yields a limit cycle with sequential firing consisting of exactly one node from each layer,

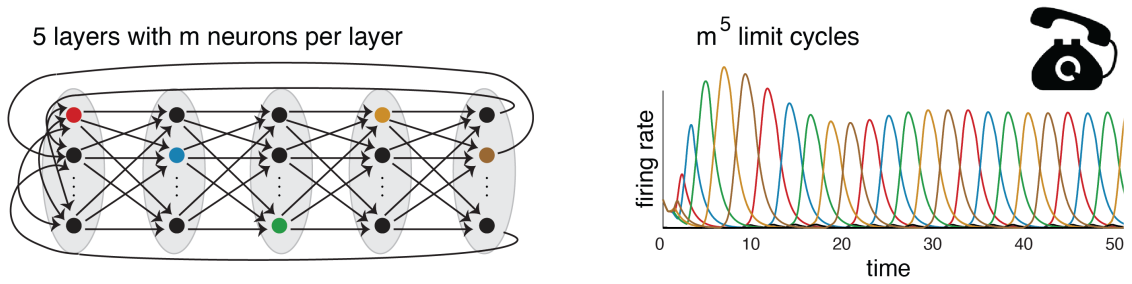

Figure 13. *The phone number network. (Left) A cyclically structured graph with m neurons per layer and all m^2 feedforward connections from one layer to the next. (Right) A limit cycle for the corresponding CTLN (with parameters $\varepsilon = 0.75$, $\delta = 4$) in which five neurons fire in a repeating sequence consisting of one neuron from each layer. Note there are m^5 such limit cycles by symmetry.*

shown on the right. By symmetry, there must be an equivalent limit cycle for every choice of five nodes, one from each layer, and thus the network is guaranteed to have m^5 limit cycles. Note that this network architecture, increased to seven layers, could serve as a mechanism for storing phone numbers in working memory (with $m = 10$ for digits $0 - 9$). The phone number is stored as a sequence that is repeated indefinitely, with different initial conditions producing different phone number sequences.

4.5. Discussion of CTLN dynamic attractors. We have seen a rich variety of emergent dynamics from CTLNs. Theorem 3.4 gave us a simple way to guarantee that a network has dynamic attractors (as opposed to stable fixed points) by ensuring the graph is oriented and has no sinks. Alternatively, we can also engineer dynamic attractors from graphs with bidirectional edges, so that groups of neurons can fire synchronously as part of an evolving sequence of activity. For example, Theorem 3.5 tells us that if a clique has a target, then there is no associated fixed point within the larger network. However, such cliques can be transiently activated. This activity is typically followed by activation of the target node(s), which can in turn activate a subsequent clique, reminiscent of cell assembly sequences (see Figure 11A). This kind of chain-like construction was also used in designing subnetworks for the quadruped gaits in Figure 10.

In fact, the bound and trot graph architectures from Figure 10A (left panel) are examples of a more general construction called a *cyclic union.* Cyclic unions are examples of "gluing" constructions, which combine subgraphs into a larger network according to certain rules. Specifically, in a cyclic union a set of disjoint graphs $G_{\tau_1}, \ldots, G_{\tau_N}$ is glued together in a cyclic manner such that every node in G_{τ_i} projects to every node in $G_{\tau_{i+1}}$ with no other edges between components. In Figure 10A, the cyclic union has subgraphs corresponding to a 2-clique, a singleton, another 2-clique, and another singleton. Interestingly, the graphs in Figures 12A and 13 are also cyclic unions. In the case of the phone number network, the cyclic union consists of five independent sets of size m. Motivated in part by these examples, cyclic unions have been studied extensively in subsequent work [9, 11]. For a summary of theorems about cyclic unions and other related "gluing" constructions, see [30].

When combining the two subnetworks for bound and trot gaits in Figure 10A (right panel), we used a different method of "gluing" two graphs. In this case, the two networks were combined by taking a union of all vertices and edges, while identifying the four common nodes (corresponding to LF, RF, LH, RH) across the two graphs. The mathematical properties of this type of gluing are not well understood. However, simulations have shown that this method often preserves the attractors of each component graph, allowing the construction of networks with multiple co-existing limit cycles. For example, in [31] a similar process is used to construct a network with five different quadruped gaits, each accessible via different choices of initial conditions.

5. Proofs. In this section, we prove Theorems 3.1, 3.3, and 3.5 and, as a consequence, Theorem 3.4 (an immediate corollary of Theorem 3.3). Our proofs build on our previous work in [5, 6, 7], as well as a useful reframing of TLNs as a patchwork of linear systems of ODEs, which we describe in section 5.1.

Throughout this section, TLNs will be assumed to be *competitive* and *nondegenerate*, as originally defined in [9].

Definition 5.1. *We say that a TLN* (W,b) *is* nondegenerate *if the following hold:*

- $\det(I - W_\sigma) \neq 0$ *for each* $\sigma \subseteq [n]$*;*
- *for each* $\sigma \subseteq [n]$ *such that* $b_i > 0$ *for all* $i \in \sigma$*, the corresponding Cramer's determinant is nonzero:* $\det((I - W_\sigma)_i; b_\sigma) \neq 0$*; and*
- $b_i > 0$ *for at least one* $i \in [n]$*.*

As in previous sections, we use the notation A_σ and b_σ to denote a matrix A and a vector b that have been truncated to include only entries with indices in σ. Here we also use the notation $(A_i; b)$ to denote a matrix A whose ith column has been replaced by the vector b, as in Cramer's rule (see, e.g., Lemma 1 in [9]).[3] In the case of a restricted matrix, $((A_\sigma)_i; b_\sigma)$ denotes the matrix A_σ where the column corresponding to the index $i \in \sigma$ has been replaced by b_σ (note that this is not typically the ith column of A_σ). Note that almost all networks of the form (1) are nondegenerate, since having a zero determinant is a highly fine-tuned condition.

The first nondegeneracy condition, $\det(I - W_\sigma) \neq 0$, ensures that all the linear systems of the TLN are nondegenerate and thus have a unique fixed point. The second condition is necessary to ensure that the fixed points of two adjacent linear systems do not coincide on a common boundary of their respective chambers in the state space. Finally, the third condition guarantees that the origin is not a fixed point of the TLN.

5.1. TLNs as a patchwork of linear systems. Recall that for a general TLN (W,b), the equations are given by

$$\frac{dx_i}{dt} = -x_i + \left[\sum_{j=1}^{n} W_{ij}x_j + b_i\right]_+ \qquad \text{for } i \in [n].$$

It is convenient to introduce the notation

$$y_i(x) \stackrel{\text{def}}{=} \sum_{j=1}^{n} W_{ij}x_j + b_i$$

and rewrite the TLN equations as

$$\frac{dx_i}{dt} = -x_i + [y_i]_+ \qquad \text{for } i \in [n].$$

From here we see that every TLN (W,b) gives rise to a hyperplane arrangement with n hyperplanes defined by $y_i(x) = 0$, subdividing $\mathbb{R}^n$ into at most 2^n regions of the form

$$R_\sigma = \{x \in \mathbb{R}^n | y_i(x) \geq 0 \quad \forall i \in \sigma \text{ and } y_k(x) \leq 0 \quad \forall k \notin \sigma\}.$$

Note that the regions cover $\mathbb{R}^n$ and overlap only on boundaries where one or more $y_i = 0$. Moreover, if x^* is a fixed point of (W,b), then

$$x_i^* = [y_i(x^*)]_+ \qquad \text{for each } i \in [n]. \tag{5}$$

[3]The use of the subscript i inside $(A_i; b)$ and $((A_\sigma)_i; b_\sigma)$ has a different meaning than the subscript σ in A_σ, because it refers to replacing the ith column by b (or b_σ), as opposed to restricting the entries of A to the index set $\{i\}$. This is an abuse of notation, but the meaning should always be clear from the context.

Through this lens, we can view the piecewise-linear TLN dynamics as a patchwork of distinct linear systems, one for each R_σ. Restricting ourselves to $x \in R_\sigma$, we see that the ODE for each node $i \in \sigma$ reduces to $dx_i/dt = -x_i + y_i$, while for $k \notin \sigma$ we obtain $dx_k/dt = -x_k$. The dynamics in each R_σ thus reduce to a purely linear system, which we denote by L_σ:

$$L_\sigma = \left\{ \frac{dx_i}{dt} = -x_i + \sum_{j=1}^{n} W_{ij}x_j + b_i \,\middle|\, i \in \sigma \right\} \bigcup \left\{ \frac{dx_k}{dt} = -x_k \,\middle|\, k \notin \sigma \right\}.$$

Assuming each linear system L_σ is nondegenerate, which is guaranteed by the nondegeneracy condition for the TLN, each has a unique fixed point x^* associated to it. Note that x^* is a fixed point of the TLN precisely when it lies inside the region R_σ where its linear system L_σ applies.

Figure 14 illustrates the hyperplanes and regions for a TLN with $n = 2$. Each region, R_σ, has its own linear system of ODEs, L_σ, for $\sigma = \emptyset, \{1\}, \{2\}$, or $\{1,2\}$. The fixed points corresponding to each linear system are denoted by x^* in matching color. Note that only one of the four regions, $R_{\{2\}}$, contains its own fixed point (in red). This fixed point, $x^* = [0, b_2]^T$, is thus the only fixed point of the TLN.

As observed above, the regions R_σ cover $\mathbb{R}^n$, and within each R_σ the TLN equations reduce to the linear system L_σ. In particular, this means that a fixed point x^* of a TLN (W, b) must be a fixed point of one of the linear systems, L_σ. Moreover, it must be an L_σ fixed point that lies within the corresponding chamber, R_σ. To determine the fixed points of a TLN, it is thus important to (1) identify the fixed points of the L_σ's and (2) have a method

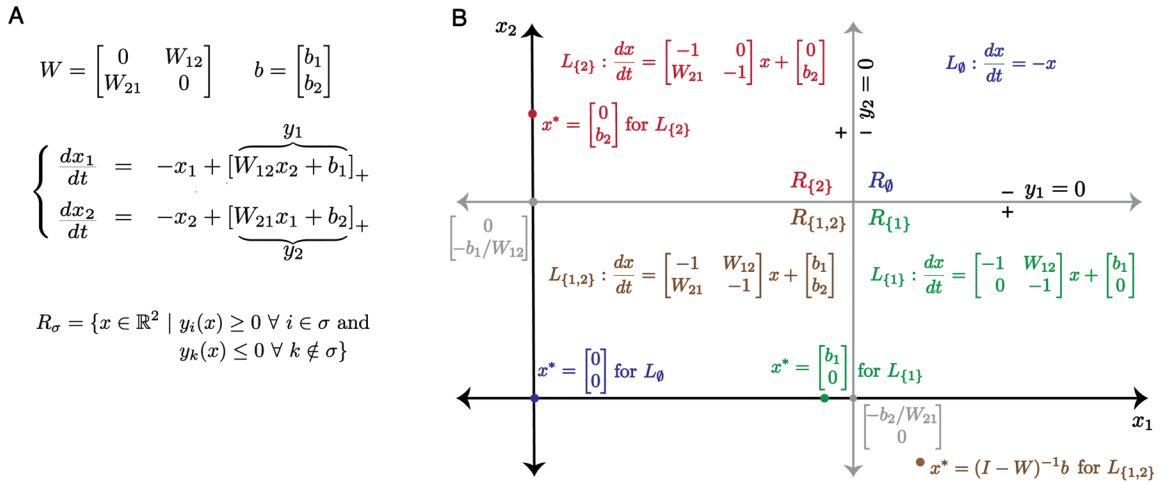

Figure 14. *Patchwork of linear systems for a TLN (W, b) on $n = 2$ nodes.* (A) *Equations for a general 2×2 TLN and corresponding R_σ regions.* (B) *An arrangement of hyperplanes $y_i = 0$ subdividing $\mathbb{R}^2$ into regions R_σ with corresponding linear systems L_σ. The fixed point of each linear system is also shown, color-coded according to its region.*

for checking whether they lie inside their respective chambers, R_σ. The following lemmas are useful for these purposes.

5.2. Fixed points of the linear systems L_σ.

Lemma 5.2. *Let (W,b) be a competitive nondegenerate TLN on n nodes. For each $\sigma \subseteq [n]$, the linear system L_σ has a unique fixed point, $x^* \in \mathbb{R}^n$, with $\operatorname{supp} x^* \subseteq \sigma$. This fixed point is stable if and only if $-I + W_\sigma$ is a stable matrix (that is, all eigenvalues have negative real part).*

Proof. Fix $\sigma \subseteq [n]$. Without loss of generality, we can reorder the nodes so that $\sigma = \{1, \dots, |\sigma|\}$. The linear system L_σ is then given by

$$L_\sigma : \ \frac{dx}{dt} = Ax + \begin{bmatrix} b_\sigma \\ 0 \end{bmatrix} \quad \text{for } A = \begin{bmatrix} -I_{|\sigma|} + W_\sigma & W_{\sigma\bar{\sigma}} \\ 0 & -I_{n-|\sigma|} \end{bmatrix}, \tag{6}$$

where $W_{\sigma\bar{\sigma}}$ is the matrix obtained by restricting W to rows indexed in σ and columns indexed by the complement, $\bar{\sigma} = [n] \setminus \sigma$. Note that $I_{|\sigma|}$ and $I_{n-|\sigma|}$ are identity matrices of sizes $|\sigma|$ and $n - |\sigma|$, respectively.

Now observe that $\det A = (-1)^{n-|\sigma|} \det(-I + W_\sigma)$, where the size of I is understood to match that of W_σ. By the nondegeneracy condition of Definition 5.1, we see that $\det A \neq 0$. It follows that L_σ is nondegenerate and has a unique fixed point given by

$$x^* = (-A)^{-1} \begin{bmatrix} b_\sigma \\ 0 \end{bmatrix}.$$

(Note that x^* may or may not lie inside the appropriate chamber R_σ.) Using Cramer's rule, we see that the entries of x^* can be written as

$$x_i^* = \frac{\det\left((-A)_i; \begin{bmatrix} b_\sigma \\ 0 \end{bmatrix}\right)}{\det(-A)} = \begin{cases} \dfrac{\det((I - W_\sigma)_i; b_\sigma)}{\det(I - W_\sigma)} & \text{for } i \in \sigma, \\ 0 & \text{for } i \notin \sigma. \end{cases}$$

Thus, $\operatorname{supp} x^* \subseteq \sigma$. (Remark: In the case where $b_i > 0$ for all $i \in \sigma$, the nondegeneracy condition guarantees that $x_i^* \neq 0$ for $i \in \sigma$. However, this does not guarantee that $x_i^* > 0$; it could be positive or negative.)

The fixed point x^* is *stable* if and only if the Jacobian of the system evaluated at x^* is stable. In a linear system of the form $dx/dt = Ax + b$, the Jacobian at every point is simply the matrix A. In our case, the matrix A for the system L_σ is block triangular, so the eigenvalues of A are the eigenvalues of $-I + W_\sigma$ together with $n - |\sigma|$ copies of -1. It follows that the fixed point x^* is stable if and only if the eigenvalues of $-I + W_\sigma$ all have negative real part. In other words, x^* is stable if and only if $-I + W_\sigma$ is a stable matrix. ∎

As an immediate corollary, we see that our TLNs have at most one fixed point per $\sigma \subseteq [n]$.

Corollary 5.3. *A competitive nondegenerate TLN on n nodes has at most one fixed point per linear system L_σ for $\sigma \subseteq [n]$. In particular, there are at most 2^n fixed points.*

5.3. From L_σ fixed points to TLN fixed points. Whether or not the fixed point x^* of the linear system L_σ is actually a fixed point of the full TLN depends on whether or not $x^* \in R_\sigma$, as this is the region where the L_σ equations apply. By definition, $x^* \in R_\sigma$ if and only if

- $y_i(x^*) \geq 0$ for all $i \in \sigma$, and
- $y_k(x^*) \leq 0$ for all $k \notin \sigma$.

These are known as the "on"-neuron and "off"-neuron conditions, respectively. Recall that for $i \in \sigma$, the L_σ equations are

$$\frac{dx_i}{dt} = -x_i + y_i(x) \quad \text{for all } i \in \sigma,$$
$$\frac{dx_k}{dt} = -x_k \quad \text{for all } k \notin \sigma.$$

It follows that at the L_σ fixed point,

$$x_i^* = y_i(x^*) \quad \text{for all} \quad i \in \sigma \quad \text{and} \quad x_k^* = 0 \quad \text{for all} \quad k \notin \sigma.$$

On the other hand, for x^* to be a fixed point of the TLN it must satisfy $x_j^* = [y_j(x^*)]_+$ for all $j \in [n]$ (see (5)). The "on" conditions thus ensure that $x_i^* \geq 0$ for all $i \in \sigma$, while the "off" conditions guarantee that $x_k^* = 0$ satisfies $x_k^* = [y_k(x^*)]_+$ for each $k \notin \sigma$. We thus have the following equivalent formulation for when a fixed point x^* of L_σ is a fixed point of the full TLN [10]:

- ("on" conditions) $x_i^* \geq 0$ for all $i \in \sigma$, and
- ("off" conditions) $y_k(x^*) \leq 0$ for all $k \notin \sigma$.

The next lemma gives precise formulas for checking the "on"- and "off"-neuron conditions. It also shows that if all $b_j > 0$, the inequalities are strict. This means that when the conditions are satisfied, $\operatorname{supp} x^* = \sigma$.

Lemma 5.4. *Consider a competitive nondegenerate TLN (W, b) on n nodes, and let x^* be the fixed point of L_σ for some $\sigma \subseteq [n]$. Then*

$$y_i(x^*) = x_i^* = \frac{\det((I - W_\sigma)_i; b_\sigma)}{\det(I - W_\sigma)} \quad \textit{for all } i \in \sigma,$$
$$y_k(x^*) = \frac{\det((I - W_{\sigma\cup\{k\}})_k; b_{\sigma\cup\{k\}})}{\det(I - W_\sigma)} \quad \textit{for all } k \notin \sigma.$$

Moreover, if $b_j > 0$ for all $j \in [n]$, then $y_i(x^*) = x_i^* \neq 0$ for all $i \in \sigma$ and $y_k(x^*) \neq 0$ for all $k \notin \sigma$.

Proof. Let x^* be the fixed point of L_σ. We have already seen that $y_i(x^*) = x_i^*$ for $i \in \sigma$. The expression for x_i^* in the statement is simply the Cramer's rule computation for x_i^* from the proof of Lemma 5.2. Moreover, as explained in the remark at the end of that proof, if $b_j > 0$ for all $j \in [n]$, then the nondegeneracy condition implies $x_i^* \neq 0$ for all $i \in \sigma$.

To see the $y_k(x^*)$ equation for $k \notin \sigma$, we use the definition of $y_k(x)$ and the formula for x_i^* in the proof of Lemma 5.2 (for both $i \in \sigma$ and $i \notin \sigma$) to compute

$$y_k(x^*) = \sum_{j=1}^{n} W_{kj} x_j^* + b_k = \sum_{j\in\sigma} W_{kj} \frac{\det((I - W_\sigma)_j; b_\sigma)}{\det(I - W_\sigma)} + b_k.$$

Now consider the Cramer's determinant, $\det((I - W_{\sigma\cup\{k\}})_k; b_{\sigma\cup\{k\}})$, where the matrix $((I - W_{\sigma\cup\{k\}})_k; b_{\sigma\cup\{k\}})$ has the vector $b_{\sigma\cup\{k\}}$ in the column corresponding to index k of $I - W_{\sigma\cup\{k\}}$. We can move the column corresponding to index k to the beginning with some number of column swaps. An equal number of row swaps brings the row corresponding to index k to the top. We thus obtain

$$\det((I - W_{\sigma\cup\{k\}})_k; b_{\sigma\cup\{k\}}) = \det\left(\begin{array}{c|c} b_k & -W_{ki_1} \cdots - W_{ki_{|\sigma|}} \\ \hline b_\sigma & I - W_\sigma \end{array}\right),$$

where $\sigma = \{i_1, \ldots, i_{|\sigma|}\}$. Applying the Laplace expansion for the determinant along the first row, we compute

$$\begin{aligned}\det((I - W_{\sigma\cup\{k\}})_k; b_{\sigma\cup\{k\}}) &= b_k \det(I - W_\sigma) + \sum_{\ell=1}^{|\sigma|} (-1)^{2+\ell}(-W_{ki_\ell})(-1)^{\ell-1}\det((I - W_\sigma)_\ell; b_\sigma) \\ &= b_k \det(I - W_\sigma) + \sum_{\ell=1}^{|\sigma|} W_{ki_\ell}\det((I - W_\sigma)_\ell; b_\sigma) \\ &= \sum_{j\in\sigma} W_{kj}\det((I - W_\sigma)_j; b_\sigma) + b_k \det(I - W_\sigma).\end{aligned}$$

Therefore, $y_k(x^*) = \frac{\det((I-W_{\sigma\cup\{k\}})_k; b_{\sigma\cup\{k\}})}{\det(I-W_\sigma)}$, as desired. ∎

An immediate corollary of Lemma 5.4 is that none of the fixed points of the linear systems L_σ lie on any of the hyperplanes defined by $y_i = 0$. In particular, we obtain the following.

Corollary 5.5. *Let (W, b) be a competitive nondegenerate TLN with $b_j > 0$ for all $j \in [n]$. If x^* is a fixed point of (W, b), then x^* lies in the interior of R_σ for $\sigma = \operatorname{supp} x^*$. In particular, the fixed points of (W, b) are all isolated.*

It is worth noting that the guarantee of isolated fixed points relies crucially on the TLN being nondegenerate. In the absence of nondegeneracy, it is possible to have two fixed points coincide with each other on a boundary $y_i = 0$.

5.4. Proof of Theorem 3.1. The proof of Theorem 3.1 is a straightforward application of the Poincaré–Hopf theorem, the famous index theorem relating the zeros of a vector field on a manifold $\mathcal{M}$ to the Euler characteristic $\chi(\mathcal{M})$. First, we define the index of a zero of a vector field, following the notation of [32].

Definition 5.6. *Let $\mathcal{M}$ be an oriented differentiable manifold. Assume that a continuous vector field $v(x)$ on $\mathcal{M}$ is differentiable in an open neighborhood of a zero, $x^* \in v^{-1}(0)$. If the vector field's Jacobian $\frac{dv}{dx}(x^*)$ has all nonzero eigenvalues, then the index $\iota(v, x^*)$ is defined as the sign of its determinant:*

$$\iota(v, x^*) = \operatorname{sgn}\det\left(\frac{dv}{dx}(x^*)\right). \tag{7}$$

Theorem 5.7 (Poincaré–Hopf theorem [32]). *Let $\mathcal{M}$ be a compact oriented differentiable manifold with boundary $\partial\mathcal{M}$, and let $v(x)$ be a continuous vector field on $\mathcal{M}$ with isolated zeros. Assume that at each point of the boundary, the vector field is pointing outside of $\mathcal{M}$. Then the following equality holds:*

$$\sum_{x^* \in v^{-1}(0)} \iota(v, x^*) = \chi(\mathcal{M}). \tag{8}$$

To apply the Poincaré–Hopf theorem in the context of competitive nondegenerate TLNs, we will consider the vector field $v(x) = x - [Wx + b]_+$, which is the negative of the vector field in (1). The zeros of $v(x)$ are thus precisely the fixed points of the TLN. The manifold of interest will be a hypercube $\mathcal{M} \subset \mathbb{R}^n$ that contains the box $\mathcal{B} = \prod_{i=1}^n [0, b_i]$ from Lemma 2.1 (see Figure 15). It will thus contain all the fixed points of the TLN (Corollary 2.2). The reason we choose $v(x)$ to have the opposite sign of the TLN vector field is so that it points *outside* of $\mathcal{M}$ on the boundary, matching the hypothesis in Theorem 5.7. In contrast, the TLN vector field, $-v(x)$, points in, yielding bounded activity with $\mathcal{B} \subset \mathcal{M}$ a globally attracting set (Lemma 2.1).

For convenience, we restate Theorem 3.1 below and then proceed to the proof. Recall that the index we defined for fixed points of TLNs is $\mathrm{idx}(\sigma) = \mathrm{sgn}\det(I - W_\sigma)$, where σ is the fixed point support. It is guaranteed to be nonzero for nondegenerate TLNs.

Theorem 3.1 (parity). *Let (W, b) be a competitive nondegenerate TLN on n nodes, with $b_i > 0$ for all $i \in [n]$. Then*

$$\sum_{\sigma \in \mathrm{FP}(W,b)} \mathrm{idx}(\sigma) = +1.$$

In particular, the total number of fixed points $|\mathrm{FP}(W, b)|$ is always odd.

Proof. Let x^* be a fixed point of a competitive nondegenerate TLN (W, b), with $b_i > 0$ for all $i \in [n]$, and let $v(x) = x - [Wx + b]_+$ be the negative of the TLN vector field. First, we will show that the index $\iota(v, x^*)$ from the Poincaré–Hopf theorem is equal to the TLN index, $\mathrm{idx}(\sigma)$, for $\sigma = \mathrm{supp}\, x^*$. Starting from Definition 5.6,

$$\iota(v, x^*) = \mathrm{sgn}\det\left(\frac{dv}{dx}(x^*)\right) = \mathrm{sgn}\det(-A) = \mathrm{sgn}\det(I - W_\sigma) = \mathrm{idx}(\sigma),$$

where A is the matrix for L_σ in (6). Because x^* lies in the interior of R_σ (Corollary 5.5), where the dynamics are governed by the linear system L_σ, the matrix A is precisely the Jacobian of the TLN vector field at x^*. It follows that $-A$ is the Jacobian of v evaluated at x^*.

Let $\mathcal{M}$ be the hypercube

$$\mathcal{M} = \prod_{i=1}^n [-\alpha, b_i + \beta]$$

for some $\alpha, \beta > 0$, and observe that $\mathcal{M}$ contains $\mathcal{B} = \prod_{i=1}^n [0, b_i]$ (see Figure 15). Thus, by Corollary 2.2, $\mathcal{M}$ contains all fixed points of (W, b) in its interior. Additionally, $\mathcal{M}$ is compact and oriented and has a boundary $\partial\mathcal{M}$. Note that although $\mathcal{M}$ is not differentiable on some parts of the boundary (it has corners), the vector field $v(x)$ is continuous on all of $\mathcal{M}$ and has isolated zeros in the interior (Corollary 5.5).[4]

[4]The boundary of $\mathcal{M}$ could easily have been chosen to be smooth, so that $\mathcal{M}$ is differentiable everywhere. However, we pick a hybercube for ease of description and showing that $v(x)$ satisfies the additional property of pointing outside of $\mathcal{M}$ along the boundary, as required by Theorem 5.7. (By continuity, $v(x)$ also points outside the boundary of a perturbed, smoothed version of $\mathcal{M}$.)

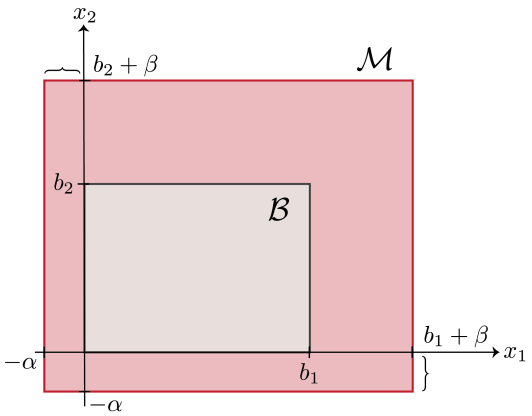

Figure 15. *A two-dimensional schematic of the hypercube $\mathcal{M} \subset \mathbb{R}^n$. By design, $\mathcal{M}$ contains the box $\mathcal{B} = \prod_{i=1}^n [0, b_i]$ from Lemma 2.1. $\mathcal{M}$ is used in the proof of Theorem 3.1.*

To show that at each point on the boundary, ∂M, $v(x)$ points outside of $\mathcal{M}$, we will now make some convenient choices for α, β. Let $w \stackrel{\text{def}}{=} \max_{i,j} |W_{ij}|$. Fix some $\alpha > 0$, and let $\beta = n\alpha w$. Note that $\mathcal{M}$ is bounded by $2n$ hyperplanes: n hyperplanes of the form $x_i = -\alpha$, with outward pointing normal vectors $-e_i$, and n hyperplanes of the form $x_i = b_i + n\alpha w$, with outward pointing normal vectors e_i. To show that $v(x)$ points outward everywhere along $\partial\mathcal{M}$, it thus suffices to show that on each of the $x_i = -\alpha$ hyperplanes we have $v_i(x) < 0$, and on each of the $x_i = b_i + n\alpha w$ hyperplanes we have $v_i(x) > 0$.

Recall that the ith coordinate of $v(x)$ is given by $v_i(x) = x_i - [\sum_{j=1}^n W_{ij} x_j + b_i]_+$. When $x_i = -\alpha$, we have $v_i(x) = -\alpha - [\sum_{j=1}^n W_{ij} x_j + b_i]_+ < 0$, as desired. On the opposite hyperplane, where $x_i = b_i + n\alpha w$, we have $v_i(x) = b_i + n\alpha w - [\sum_{j=1}^n W_{ij} x_j + b_i]_+ \geq b_i + n\alpha w - \max_{x \in \partial M} [\sum_{j=1}^n W_{ij} x_j + b_i]_+$. Since $W_{ij} \leq 0$ and $W_{ii} = 0$, the largest value of $[\sum_{j=1}^n W_{ij} x_j + b_i]_+$ on ∂M occurs when $x_j = -\alpha$ for each $j \neq i$. We thus have

$$\begin{aligned} v_i(x) &\geq b_i + n\alpha w - \left[\sum_{j=1}^n W_{ij}(-\alpha) + b_i \right]_+ \geq b_i + n\alpha w - \left[\sum_{j \neq i} (-w)(-\alpha) + b_i \right]_+ \\ &= b_i + n\alpha w - ((n-1)\alpha w + b_i) = \alpha w > 0, \end{aligned}$$

where the second inequality stems from the fact that the most negative value of W_{ij} is $-w$. We conclude that the vector field $v(x)$ points outward everywhere along the boundary of $\mathcal{M}$, and thus the Poincaré–Hopf theorem applies.

Now recall that the Euler characteristic $\chi(\mathcal{M}) = 1$ for any solid hypercube (or ball) in $\mathbb{R}^n$. It follows from Theorem 5.7 that

$$\sum_{\sigma \in \mathrm{FP}(W,b)} \mathrm{idx}(\sigma) = \sum_{x^* \in v^{-1}(0)} \iota(v, x^*) = \chi(\mathcal{M}) = +1.$$

■

5.5. Proof of Theorem 3.3. In this section, we prove Theorem 3.3. Recall from Lemma 5.2 that a fixed point x^* with support σ is *stable* if and only if $-I+W_\sigma$ is a stable matrix (i.e., all eigenvalues have strictly negative real part). We say that a matrix is *unstable* if at least one eigenvalue has positive real part. The following lemma gives a condition on W that rules out the existence of stable fixed points supported on two or more nodes.

Lemma 5.8. *Consider a TLN (W,b) where W has diagonal entries $W_{ii}=0$. If all 2×2 principal submatrices of $-I+W$ are unstable, then the network has no stable fixed points supported on more than one node.*

Proof. Suppose all 2×2 principal submatrices, $-I+W_\tau$ for $|\tau|=2$, are unstable, and note that the trace of each of these matrices is -2. In [6, Lemma 1], it was shown that if all 2×2 principal submatrices of an $n\times n$ matrix have negative trace and are unstable, then all larger principal submatrices are also unstable. It follows that $-I+W_\sigma$ is unstable for all σ of size $|\sigma|\geq 2$. This implies that (W,b) can have no stable fixed points supported on more than one node, because a fixed point with support σ can only be stable if $-I+W_\sigma$ is stable (Lemma 5.2). ∎

We are now ready to prove Theorem 3.3, which we restate below for convenience.

Theorem 3.3. *Consider a competitive nondegenerate TLN with connectivity matrix W, associated graph G_W, and uniform inputs $b_i=\theta$. Suppose that*

(i) *G_W is an oriented graph with no sinks, and*

(ii) *whenever $j\to i$ in G_W, $W_{ij}<\dfrac{1}{W_{ji}}$.*

Then the network (1) has no stable fixed points. Moreover, the network activity is bounded.

Proof. We have already seen that the activity is bounded (Lemma 2.1). To see that there are no stable fixed points, we first show that there can be no stable fixed points supported on two or more nodes. By Lemma 5.8, it suffices to show that all 2×2 principal submatrices of $-I+W$ are unstable. Each of these matrices, $\left(\begin{smallmatrix}-1 & W_{ij}\\ W_{ji} & -1\end{smallmatrix}\right)$, has negative trace and is thus stable if and only if its determinant is positive. The determinant is

$$\Delta=\det\begin{pmatrix}-1 & W_{ij}\\ W_{ji} & -1\end{pmatrix}=1-W_{ij}W_{ji},$$

which is positive if and only if $W_{ij}W_{ji}<1$.

Since the network is competitive and G_W is an oriented graph, there are two cases to consider: (a) there is a single edge $j\to i$ (or $i\to j$), so that $-1<W_{ij}\leq 0$ and $W_{ji}\leq -1$ (or vice versa), or (b) there is no edge between i and j, so that both $W_{ij}\leq -1$ and $W_{ji}\leq -1$. In case (a), hypothesis (ii) of the theorem implies $W_{ij}W_{ji}>1$ (note that the inequality reverses, since W_{ji} is negative). We thus have $\Delta<0$, and the 2×2 matrix is unstable. In case (b), we immediately see that $W_{ij}W_{ji}\geq 1$, and thus $\Delta\leq 0$. We conclude that all 2×2 principal submatrices of $-I+W$ are unstable, and thus the network has no stable fixed points with two or more active nodes.

Next, we show by contradiction that the network has no fixed points supported on a single node (i.e., there is no winner-take-all behavior). Suppose x^* is a fixed point supported on node i, so that $x_i^*>0$ and $x_j^*=0$ for all $j\neq i$, and recall that $W_{ii}=0$. It follows that

$$x_i^* = \left[\sum_{j=1}^{n} W_{ij}x_j^* + \theta\right]_+ = [W_{ii}x_i^* + \theta]_+ = \theta.$$

On the other hand, for any $k \neq i$ we must have $x_k^* = 0$ at the fixed point, and so

$$x_k^* = \left[\sum_{j=1}^{n} W_{kj}x_j^* + \theta\right]_+ = [W_{ki}x_i^* + \theta]_+ = [W_{ki}\theta + \theta]_+ = 0.$$

Now recall that G_W has no sinks (by hypothesis (i) of the theorem), and so there exists at least one vertex $\ell \neq i$ such that $i \to \ell$. This means $W_{\ell i} > -1$, and thus $x_\ell^* > 0$, contradicting the assumption that the fixed point was supported only on node i. ∎

5.6. Proof of Theorem 3.5. To prove Theorem 3.5, we make use of the fixed point conditions that were derived in [7].

Lemma 5.9. *Consider a nondegenerate CTLN model with $W = W(G, \varepsilon, \delta)$, and suppose x^* is a fixed point of (1) supported on a clique σ of G. Then x^* is a stable fixed point, and*

$$x_\sigma^* = \frac{\theta}{\varepsilon + (1-\varepsilon)|\sigma|} 1_\sigma,$$

where 1_σ is a column vector of all 1s.

Proof. If σ is a clique of G, then $-I + W_\sigma = (-1+\varepsilon)11^T - \varepsilon I_\sigma$, with eigenvalues $|\sigma|(-1+\varepsilon) - \varepsilon$ and $-\varepsilon$. Clearly, these are negative for $0 < \varepsilon < 1$, so we can conclude that $-I + W_\sigma$ is stable. It follows that any fixed point x^* with support σ is stable and unique [6, Corollary 9] (see also [7, section 1.1]). To verify the formula for x_σ^*, we simply check that it satisfies the fixed point equation $x_\sigma^* = [W_\sigma x_\sigma^* + \theta 1_\sigma]_+$. Since $x_\sigma^* > 0$, we can drop the threshold nonlinearity to obtain the equivalent constraint, $(I - W_\sigma)x_\sigma^* = \theta 1_\sigma$. Now plugging in the desired expression for x_σ^* yields

$$\begin{aligned}(I - W_\sigma)x_\sigma^* &= \left((1-\varepsilon)11^T + \varepsilon I_\sigma\right) \frac{\theta}{\varepsilon + (1-\varepsilon)|\sigma|} 1_\sigma \\ &= \frac{\theta}{\varepsilon + (1-\varepsilon)|\sigma|}\left((1-\varepsilon)|\sigma|1_\sigma + \varepsilon 1_\sigma\right) = \theta 1_\sigma.\end{aligned}$$

∎

We can now prove Theorem 3.5, which we restate below for convenience.

Theorem 3.5. *Let G be a simple directed graph, and consider an associated nondegenerate CTLN with $W = W(G, \varepsilon, \delta)$ for any choice of the parameters $\varepsilon, \delta, \theta > 0$ with $\varepsilon < 1$. If σ is a clique of G, then there exists a stable fixed point with support σ if and only if σ is target-free.*

Proof. ($\Rightarrow$) Suppose x^* is a stable fixed point with support σ, where σ is a clique of G. Then $W_{ij} = -1 + \varepsilon$ for all pairs $i, j \in \sigma$. To see that σ must be a *target-free* clique, suppose that σ has a target $k \notin \sigma$. This implies that $W_{ki} = -1 + \varepsilon$ for each $i \in \sigma$. It follows that

$$x_k^* = \left[\sum_{i\in\sigma} W_{ki}x_i^* + \theta\right]_+ = \left[(-1+\varepsilon)\sum_{i\in\sigma} x_i^* + \theta\right]_+ = \left[\frac{\varepsilon\theta}{\varepsilon + (1-\varepsilon)|\sigma|}\right]_+ > 0,$$

where we have used the expression for x^*_σ from Lemma 5.9 to obtain $\sum_{i\in\sigma} x^*_i = \dfrac{\theta|\sigma|}{\varepsilon+(1-\varepsilon)|\sigma|}$. This contradicts the fact that $x^*_k = 0$, since $k \notin \sigma$ and $\sigma = \operatorname{supp} x^*$. We thus conclude that σ must be a target-free clique.

($\Leftarrow$) Suppose σ is a target-free clique. Since σ is a clique, it follows from Lemma 5.9 that if a fixed point x^* with support σ exists, then it must be unique and stable, with $x^*_\sigma = \dfrac{\theta}{\varepsilon+(1-\varepsilon)|\sigma|}1_\sigma$. Clearly, $x^*_\sigma > 0$, and thus the "on"-neuron conditions hold. To guarantee that the fixed point of the CTLN with support σ exists, however, we must also check the "off"-neuron conditions: $y_k(x^*) \le 0$ for each $k \notin \sigma$. Since σ is a target-free clique, for any $k \notin \sigma$ there exists $i_k \in \sigma$ such that $i_k \not\to k$, and so $W_{ki_k} = -1-\delta$. We thus have

$$y_k(x^*) = \sum_{i\in\sigma} W_{ki}x^*_i + \theta \le \theta\left(\frac{-1-\delta+(|\sigma|-1)(-1+\varepsilon)}{\varepsilon+(1-\varepsilon)|\sigma|}+1\right) = \frac{-\theta\delta}{\varepsilon+(1-\varepsilon)|\sigma|} < 0,$$

showing that the "off" conditions hold. ∎

REFERENCES

[1] J. J. Hopfield, *Neural networks and physical systems with emergent collective computational abilities*, Proc. Natl. Acad. Sci. USA, 79 (1982), pp. 2554–2558.

[2] J. J. Hopfield, *Neurons with graded response have collective computational properties like those of two-state neurons*, Proc. Natl. Acad. Sci. USA, 81 (1984), pp. 3088–3092.

[3] R. H. Hahnloser, H. S. Seung, and J. J. Slotine, *Permitted and forbidden sets in symmetric threshold-linear networks*, Neural Comput., 15 (2003), pp. 621–638.

[4] X. Xie, R. H. Hahnloser, and H. S. Seung, *Selectively grouping neurons in recurrent networks of lateral inhibition*, Neural Comput., 14 (2002), pp. 2627–2646.

[5] C. Curto, A. Degeratu, and V. Itskov, *Flexible memory networks*, Bull. Math. Biol., 74 (2012), pp. 590–614.

[6] C. Curto, A. Degeratu, and V. Itskov, *Encoding binary neural codes in networks of threshold-linear neurons*, Neural Comput., 25 (2013), pp. 2858–2903.

[7] C. Curto and K. Morrison, *Pattern completion in symmetric threshold-linear networks*, Neural Comput., 28 (2016), pp. 2825–2852.

[8] K. Morrison, A. Degeratu, V. Itskov, and C. Curto, *Diversity of Emergent Dynamics in Competitive Threshold-Linear Networks: A Preliminary Report*, preprint, https://arxiv.org/abs/1605.04463v1, 2016.

[9] C. Curto, J. Geneson, and K. Morrison, *Fixed points of competitive threshold-linear networks*, Neural Comput., 31 (2019), pp. 94–155.

[10] K. Morrison and C. Curto, *Predicting neural network dynamics via graphical analysis*, in Algebraic and Combinatorial Computational Biology, R. Robeva and M. Macaulay, eds., Elsevier, Amsterdam, 2019, pp. 241–277.

[11] C. Parmelee, J. Londono Alvarez, C. Curto, and K. Morrison, *Sequential attractors in combinatorial threshold-linear networks*, SIAM J. Appl. Dyn. Syst., 21 (2022), pp. 1597–1630, https://doi.org/10.1137/21M1445120.

[12] A. Bel, R. Cobiaga, W. Reartes, and H. G. Rotstein, *Periodic solutions in threshold-linear networks and their entrainment*, SIAM J. Appl. Dyn. Syst., 20 (2021), pp. 1177–1208, https://doi.org/10.1137/20M1337831.

[13] C. Parmelee, S. Moore, K. Morrison, and C. Curto, *Core motifs predict dynamic attractors in combinatorial threshold-linear networks*, PLoS One, 17 (2022), e0264456.

[14] E. Nozari and J. Cortés, *Stability analysis of complex networks with linear-threshold rate dynamics*, in Proceedings of the Annual American Control Conference (ACC), IEEE, 2018, pp. 191–196.

[15] E. Nozari and J. Cortés, *Hierarchical selective recruitment in linear-threshold brain networks—part* I*: Single-layer dynamics and selective inhibition*, IEEE Trans. Automat. Control, 66 (2020), pp. 949–964.
[16] D. Liberzon, *Switching in Systems and Control*, Systems Control Found. Appl., Birkhäuser Boston, Boston, MA, 2003.
[17] M. Johansson, *Piecewise Linear Control Systems: A Computational Approach*, Lect. Notes Control Inf. Sci. 284, Springer-Verlag, Berlin, 2003.
[18] H. Lin and P. J. Antsaklis, *Stability and stabilizability of switched linear systems: A short survey of recent results*, in Proceedings of the 20th IEEE International Symposium on Intelligent Control (ISIC'05) and the 13th Mediterranean Conference on Control and Automation (MED'05), 2005, pp. 24–29.
[19] A. Pavlov, N. van de Wouw, and H. Nijmeijer, *Convergent piecewise affine systems: Analysis and design part* I*: Continuous case*, in Proceedings of the 44th IEEE Conference on Decision and Control, 2005, pp. 5391–5396.
[20] M. di Bernardo, C. J. Budd, A. R. Champneys, and P. Kowalczyk, *Piecewise-Smooth Dynamical Systems*: *Theory and Applications*, Appl. Math. Sci. 163, Springer-Verlag London, 2008.
[21] R. H. Hahnloser, R. Sarpeshkar, M. A. Mahowald, R. J. Douglas, and H. S. Seung, *Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit*, Nature, 405 (2000), pp. 947–951.
[22] C. Curto, J. Geneson, and K. Morrison, *Stable fixed points of combinatorial threshold-linear networks*, Adv. Appl. Math., 154 (2024), 102652.
[23] G. B. Ermentrout and D. H. Terman, *Mathematical Foundations of Neuroscience*, Interdiscip. Appl. Math. 35, Springer, New York, 2010.
[24] Y. Burak and I. R. Fiete, *Fundamental limits on persistent activity in networks of noisy neurons*, Proc. Natl. Acad. Sci. USA, 109 (2012), pp. 17645–17650.
[25] V. Itskov, C. Curto, E. Pastalkova, and G. Buzsáki, *Cell assembly sequences arising from spike threshold adaptation keep track of time in the hippocampus*, J. Neurosci., 31 (2011), pp. 2828–2834.
[26] V. Itskov, D. Hansel, and M. Tsodyks, *Short-term facilitation may stabilize parametric working memory trace*, Front. Comput. Neurosci., 5 (2011), 40.
[27] C. Curto, S. Sakata, S. Marguet, V. Itskov, and K. D. Harris, *A simple model of cortical dynamics explains variability and state dependence of sensory responses in urethane-anesthetized auditory cortex*, J. Neurosci., 29 (2009), pp. 10600–10612.
[28] E. Marder and D. Bucher, *Central pattern generators and the control of rhythmic movements*, Curr. Biol., 11 (2001), pp. R986–996.
[29] R. Yuste, J. N. MacLean, J. Smith, and A. Lansner, *The cortex as a central pattern generator*, Nat. Rev. Neurosci., 6 (2005), pp. 477–483.
[30] C. Curto and K. Morrison, *Graph Rules for Recurrent Neural Network Dynamics: Extended Version*, preprint, https://arxiv.org/abs/2301.12638, 2023.
[31] J. Londono Alvarez, K. Morrison, and C. Curto, *Sequential Control of Coexistent Dynamic Attractors in a Model of Quadruped Locomotion*, manuscript in preparation, 2024.
[32] J. W. Milnor, *Topology from the Differentiable Viewpoint*, The University Press of Virginia, Charlottesville, VA, 1965.