

Nerve Theorems for Fixed Points of Neural Networks



Daniela Egas Santander, Stefania Ebli, Alice Patania, Nicole Sanderson,
Felicia Burtcher, Katherine Morrison, and Carina Curto

Abstract Nonlinear network dynamics are notoriously difficult to understand. Here we study a class of recurrent neural networks called combinatorial threshold-linear networks (CTLNs) whose dynamics are determined by the structure of a directed graph. They are a special case of TLNs, a popular framework for modeling neural activity in computational neuroscience. In prior work, CTLNs were found to be surprisingly tractable mathematically. For small networks, the fixed points of the network dynamics can often be completely determined via a series of *graph rules* that can be applied directly to the underlying graph. For larger networks, it remains a challenge to understand how the global structure of the network interacts with

Authors Katherine Morrison and Carina Curto have equally contributed.

D. Egas Santander · S. Ebli
École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland
e-mail: daniela.egassantander@epfl.ch; stefania.ebli@epfl.ch

A. Patania
Indiana University Network Science Institute (IUNI), Bloomington, IN, USA
e-mail: apatania@iu.edu

N. Sanderson
Lawrence Berkeley National Lab, Berkeley, CA, USA
e-mail: nzs5677@psu.edu

F. Burtcher
Technische Universität Berlin, Berlin, Germany
Université du Luxembourg, Belvaux, Luxembourg
e-mail: felicia.burtcher@uni.lu

K. Morrison
University of Northern Colorado, Greeley, CO, USA
e-mail: katherine.morrison@unco.edu

C. Curto (✉)
Pennsylvania State University, University Park, PA, USA
e-mail: ccurto@psu.edu

local properties. In this work, we propose a method of covering graphs of CTLNs with a set of smaller *directional graphs* that reflect the local flow of activity. While directional graphs may or may not have a feedforward architecture, their fixed point structure is indicative of feedforward dynamics. The combinatorial structure of the graph cover is captured by the *nerve* of the cover. The nerve is a smaller, simpler graph that is more amenable to graphical analysis. We present three nerve theorems that provide strong constraints on the fixed points of the underlying network from the structure of the nerve. We then illustrate the power of these theorems with some examples. Remarkably, we find that the nerve not only constrains the fixed points of CTLNs, but also gives insight into the transient and asymptotic dynamics. This is because the flow of activity in the network tends to follow the edges of the nerve.

1 Introduction

Combinatorial threshold-linear networks (CTLNs) are a special class of threshold-linear networks (TLNs) whose dynamics are determined by the structure of a directed graph. The firing rates $x_1(t), \dots, x_n(t)$ of n recurrently-connected neurons evolve in time according to the standard TLN equations:

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

These networks derive their name from the nonlinear transfer function, $[\cdot]_+ = \max\{0, \cdot\}$, which is threshold-linear. A given TLN is specified by the choice of a connection strength matrix W and a vector of external inputs θ . TLNs have been widely used in computational neuroscience as a framework for modeling recurrent neural networks, including associative memory networks [3, 4, 6, 7, 10, 11, 18, 19].

What makes CTLNs special is that the matrix $W = W(G, \varepsilon, \delta)$ is determined by a simple¹ directed graph, as follows:

$$W_{ij} = \begin{cases} 0 & \text{if } i = j, \\ -1 + \varepsilon & \text{if } j \rightarrow i \text{ in } G, \\ -1 - \delta & \text{if } j \not\rightarrow i \text{ in } G. \end{cases} \quad (2)$$

Additionally, we fix $\theta_i = \theta > 0$ to be the same for all neurons, and we require the ε, δ parameters to satisfy $\delta > 0$, and $0 < \varepsilon < \frac{\delta}{\delta+1}$. CTLNs were first defined in [15], where the $\varepsilon < \frac{\delta}{\delta+1}$ condition was motivated by the desired property that subgraphs consisting of a single directed edge $i \rightarrow j$ should not be allowed to support stable fixed points. Note that the upper bound on ε implies $\varepsilon < 1$, rendering the W matrix effectively inhibitory. We think of the graph edges as excitatory connections in a

¹ A graph is *simple* if it does not have loops or multiple edges between a pair of vertices.




Fig. 1 (a) A neural network with excitatory pyramidal neurons (triangles) and a background network of inhibitory interneurons (gray circles) that produces a global inhibition. The corresponding graph (right) retains only the excitatory neurons and their connections. (b) TLN dynamics. (c) A graph that is a 3-cycle (left), and a solution for the corresponding CTLN showing that network activity follows the arrows in the graph (right). Peak activity occurs sequentially in the cyclic order 123. Unless otherwise noted, all simulations have parameters $\varepsilon = 0.25$, $\delta = 0.5$, and $\theta = 1$

sea of inhibition (Fig. 1a). Figure 1c shows an example solution for a CTLN whose graph is a 3-cycle.

TLNs are high-dimensional nonlinear systems whose dynamics are still poorly understood. However, in the special case of CTLNs, there appears to be a strong connection between the attractors of the network and the pattern of stable and unstable fixed points [14, 17].² Moreover, these fixed points can often be completely determined by the structure of the underlying graph. In prior work, a series of *graph rules* were proven that can be used to determine fixed points of the CTLN by analyzing G , irrespective of the choice of parameters ε , δ , and θ [8, 9]. A key observation is that for a given network, there can be at most one fixed point per support, $\sigma \subseteq [n]$, where the *support* of a fixed point is the subset of active neurons (i.e., $\text{supp } x = \{i \mid x_i > 0\}$).

For a given choice of parameters, we use the notation

$$\text{FP}(G) \stackrel{\text{def}}{=} \{\sigma \subseteq [n] \mid \sigma \text{ is a fixed point support of } W(G, \varepsilon, \delta)\},$$

where $[n] \stackrel{\text{def}}{=} \{1, \dots, n\}$. For many graphs, we find that the fixed point supports in $\text{FP}(G)$ are confined to a subset of the neurons. In other words, there is a partition $\{\omega, \tau\}$ of the vertices of G such that, for every $\sigma \in \text{FP}(G)$, we have $\sigma \subseteq \tau$ (see Fig. 2a). In these cases, we observe that solutions of the network activity $x(t)$ tend to converge to a region of the state space where the most active neurons are in τ , and those in ω are either silent or have very low firing (see Fig. 2b). In other words, the attractors live where the fixed points live.

This motivates us to define *directional graphs*. A directional graph G is a graph with a proper subset of neurons τ such that $\text{FP}(G) \subseteq \text{FP}(G|_\tau)$, where $G|_\tau$ is the induced subgraph obtained by restricting to the vertices of τ . For example, the graph

² A fixed point, x^* , of a TLN is a solution that satisfies $dx_i/dt|_{x^*} = 0$ for each $i \in [n]$.




Fig. 2 Directional graphs and covers. **(a)** A directional graph, with fixed points supported in the subset of nodes $\tau = \{3, 4\}$. **(b)** A solution of a CTLN with the graph in **(a)**. The network was initialized with the activity concentrated on the neurons in $\omega = \{1, 2\}$, but the activity flows from $\omega \rightarrow \tau$. **(c)** A graph with a partition of the nodes, each component in a different color. **(d)** A directional cover of G . Subsets of nodes, ν_i, ν_j, ν_k , and ν_l , correspond to the partition in **(c)**. The four induced subgraphs within each oval, of the form $G_{ij} = G|_{\nu_i \cup \nu_j}$, are all directional graphs with direction $\nu_i \rightarrow \nu_j$, as given by the arrows. **(e)** The nerve associated to the directional cover in **(d)**. The corresponding network can be viewed as a dimensional reduction of the one in **(c)**

in Fig. 2a is directional with a single fixed point supported in $\tau = \{3, 4\}$. We also require an additional technical condition that allows us to prove that certain natural compositions, like chaining directional graphs together, produce a new directional graph (see Definition 3.1 for the full definition). In simulations, such as the one in Fig. 2b, we have seen that directional graphs display feedforward dynamics, even if their architecture does not follow a feedforward structure. Activity that is initially concentrated on ω flows towards τ , giving the dynamics an $\omega \rightarrow \tau$ directionality. Thus, from a bird’s eye view, directional graphs behave like a single directed edge, where the activity flows from the source to the sink. These observations prompted us to ask the following question: if we cover a graph G with a collection of directional graphs, what can we say about $\text{FP}(G)$ from the combinatorial structure of the cover?

In this paper we develop tools to answer this question, inspired by the construction of the nerve of a cover of a topological space. We define a *directional cover* of G as a set of directional subgraphs that cover G and have well-behaved intersections.³ Effectively, such a cover is entirely determined by a partition of the

³ This is analogous to the definition of a “good” cover of a topological space, which also requires well-behaved intersections. Nerves of good covers reflect the topology of the underlying space [5, 13].

vertices of G , denoted $\{v_i\}$, that satisfies special properties. (See Definition 4.3 for a precise definition.) We define the *nerve* of a directional cover as a new graph $\mathcal{N}$ that has a directed edge for each directional graph in the cover, and a vertex for each component v_i of the partition. Figure 2c,d depicts a graph with a partition of the nodes (indicated by the colors), and its corresponding directional cover. The edges of the nerve reflect the local dynamics of G , and the nerve itself encodes the combinatorics of the intersection pattern of the cover: the directional graphs overlap precisely at vertices of the nerve where their edges meet. The partition of the vertices of G induces a canonical quotient map, $\pi : V_G \rightarrow V_{\mathcal{N}} := \{v_i\}$, that simply identifies all the vertices in each component v_i . Figure 2e is the nerve of the directional node in d, and the quotient map π sends each color in c to the corresponding node with the same color in e - within the nerve $\mathcal{N}$, we often label the node v_i simply as i .

As an illustration of directional covers and nerves, consider the graph in Fig. 3a. This graph is a chain of ten 5-cliques where the edges between adjacent cliques all follow the pattern shown in panel c: there are edges forward from every node in the first clique to every node in the second clique; every node in the second clique (except for the top node) sends edges back to every node in the first clique. Most edges are thus bidirectional arrows (in black), while the edges that only go




Fig. 3 Example graph with a directional cover, its nerve, and network activity that flows along the nerve. **(a)** A chain of ten 5-cliques where the edges between adjacent cliques all follow the pattern shown in panel (c). **(b)** The nerve of the graph in A induced by the partition of the vertices as $v_1, \dots, v_{10}$. **(c)** The graph G restricted to a pair of adjacent cliques. All edges in black are bidirectional, while those in green are unidirectional from the green clique to the pink clique. This restricted graph is directional, and all the graphs in the directional cover of G have this form. **(d)** A solution to the CTLN defined by G (with $\varepsilon = 0.25$ and $\delta = 0.5$), where the activity is initialized on the nodes in the first clique v_1 . The transient dynamics slowly activate each clique in sequence, following the path of the nerve, until the solution converges to the stable fixed point supported on the nodes in v_{10}

forward from clique i to clique $i+1$ are in color. The induced subgraphs $G|_{v_i \cup v_{i+1}}$ are all directional with direction $v_i \rightarrow v_{i+1}$, despite all the back edges from right to left. This means that $\text{FP}(G|_{v_i \cup v_{i+1}}) \subseteq \text{FP}(G|_{v_{i+1}})$, and we expect the activity of the neurons to flow from v_i to v_{i+1} (left to right). Figure 3b depicts the nerve of G . Figure 3d shows the solution to a CTLN defined by G , where we have initialized all the activity on the nodes in the first clique v_1 . We see that the activity eventually converges to the final component, shown in purple, where the fixed points of the network are concentrated. The transient dynamics, however, are rather slow, with each clique activated in a sequence that follows the path-like structure of the nerve. Note that this network behaves similarly to a synfire chain [1, 2, 12], despite numerous backward edges between components that completely destroy the feedforward architecture (Fig. 3c). The sequential dynamics are maintained because these backward edges do not disrupt the directionality of the graphs in the cover.

The main goal of this paper is to prove nerve theorems for CTLNs. Broadly speaking, such a *nerve theorem* is a result that gives information about the dynamics of a network from properties of the nerve. Specifically, we are interested in results that allow us to constrain the fixed points of G by analyzing structural properties of $\mathcal{N}$. Ideally, we would like to prove the following kind of result: If G has a directional cover with nerve $\mathcal{N}$, then

$$\sigma \in \text{FP}(G) \Rightarrow \pi(\sigma) \in \text{FP}(\mathcal{N}), \quad (3)$$

where π is the canonical quotient map from V_G to $V_{\mathcal{N}}$. This can be quite powerful in cases where $\mathcal{N}$ is a much smaller and simpler graph. Unfortunately, the statement (3) is not in general true. However, we do find that this holds whenever the nerve $\mathcal{N}$ is a directed acyclic graph (DAG) or $\mathcal{N}$ is a cycle (see Theorems 4.8 and 4.9). More generally, whenever $\mathcal{N}$ admits a DAG decomposition (see Definition 2.16), Theorem 4.7 gives a result similar in spirit to (3) and allows us to greatly constrain $\text{FP}(G)$.

Our nerve theorems can be used to simplify a complex network by finding a nontrivial directional cover and studying its nerve. Finding such covers is an art, however, and we do not yet have a systematic way of doing it. On the other hand, nerve theorems can also be used to engineer complex networks with prescribed dynamic properties. This is how we constructed the example in Fig. 3. We explore both kinds of applications in the last section of the paper.

The organization of this paper is as follows. In Sect. 2, we review some graph theory terminology and basic background and notation for CTLNs. We also introduce the DAG decomposition of a graph. In Sect. 3, we define directional graphs, prove that certain graph structures are always directional, and provide some other families of examples. In Sect. 4, we introduce directional covers and their associated nerves. Here we also state and prove our main results, Theorems 4.7, 4.8 and 4.9. Finally, in Sect. 5, we illustrate the power of our theorems with some applications.

2 Preliminaries

In this section we review some useful terminology from graph theory and summarize essential background and prior results about fixed points of CTLNs. We also introduce the DAG decomposition of a graph, a notion that will appear in our main nerve theorems.

2.1 Graph Theory Terminology

Definition 2.1 A *directed graph* G can be described as a tuple $G = (V_G, E_G)$, where V_G is a finite set called the set of *vertices* and $E_G \subseteq V_G \times V_G$ is the set of (*directed*) *edges*, where $(i, j) \in E_G$ means there is a directed edge $i \rightarrow j$ from i to j in G . If $(i, j) \notin E_G$, we write $i \not\rightarrow j$. A directed graph is *simple* if it has no self-loops, so that $(i, i) \notin E_G$ for all $i \in V_G$. A directed graph is *oriented* if it has no bidirectional edges.

In this paper, we restrict ourselves to simple directed graphs. Unless otherwise noted, we will use the word *graph* to refer to simple directed graphs.

Notation 2.2 Let G be a graph with vertex set V_G and edge set E_G . For any subset of vertices $\sigma \subseteq V_G$, denote by $G|_\sigma$ the *induced subgraph* obtained by restricting to the vertices σ . More precisely, $G|_\sigma = (\sigma, E|_\sigma)$ where $E|_\sigma = \{(i, j) \in E_G \mid i, j \in \sigma\}$.

Let $\sigma_1, \sigma_2 \subseteq V_G$ be two subsets of the vertices of G . We denote by $E_G(\sigma_1, \sigma_2) \subseteq E_G$ the set of directed edges *from* vertices in σ_1 *to* vertices in σ_2 in G .

Next, we define some basic notions relevant to graphs.

Definition 2.3 Let G be a graph and $v \in V_G$ be a vertex in G . The *in-degree* of v is the number of incoming edges to v . The *out-degree* of v is the number of outgoing edges from v . We say v is a *source* if v has no incoming edges, and we say v is a *proper source* if it is a source that has at least one outgoing edge. We say v is a *sink* if v has no outgoing edges. Note that a source that is not a proper source is an *isolated vertex*, and thus it is also a sink.

Definition 2.4 We say that a graph G has *uniform in-degree* if every vertex $v \in V_G$ has the same in-degree d . Note that an *independent set* is a graph with uniform in-degree $d = 0$. A *k-clique* is an all-to-all bidirectionally connected graph with uniform in-degree $d = k - 1$. And an *n-cycle* is a graph with n edges, $1 \rightarrow 2 \rightarrow \dots \rightarrow n \rightarrow 1$, which has uniform in-degree $d = 1$.

2.2 Background on Fixed Points of CTLNs

In this subsection we recall the results from [8] that are relevant for this work and include simple proofs to some of these to provide intuition to the reader.

A *fixed point* of a CTLN is simply a fixed point of the network equations (1). In other words, it is a vector $x^* \in \mathbb{R}_{\geq 0}^n$ such that $\frac{dx_i}{dt}|_{x=x^*} = 0$ for all $i \in [n]$. As explained in [8], fixed points of CTLNs can be labelled by their *supports* (i.e. the subset of active neurons), and for a given G the set of all fixed point supports is denoted $\text{FP}(G)$.

Lemma 2.5 ([8]) *Let G be a graph on n vertices, and suppose G has uniform in-degree. Then G has a full-support fixed point, $\sigma = [n] \in \text{FP}(G)$.*

In particular, this lemma says that cliques, cycles, and independent sets all have a full-support fixed point. In fact, this fixed point is symmetric, with $x_i^* = x_j^*$ for all $i, j \in [n]$. This is true even for uniform in-degree graphs that are not symmetric.

More generally, fixed points can have very different values across neurons. However, there is some level of “graphical balance” that is required of $G|_\sigma$ for any fixed point support σ . For example, if σ contains a pair of neurons j, k that have the property that all neurons mapping to j are also mapping to k , and $j \rightarrow k$ but $k \not\rightarrow j$, then σ cannot be a fixed point support. This is because k is receiving strictly more inputs than j , and this imbalance rules out their ability to coexist in the same fixed point support. To see this more rigorously, we have the following lemma.

Lemma 2.6 *Let G be a CTLN and $\sigma \subseteq V_G$. Suppose there exist vertices $j, k \in \sigma$ such that for each $i \in \sigma \setminus \{j, k\}$, if $i \rightarrow j$ then $i \rightarrow k$. Furthermore, suppose $j \rightarrow k$ but $k \not\rightarrow j$. Then $\sigma \notin \text{FP}(G)$.*

Proof To obtain a contradiction, assume $\sigma \in \text{FP}(G)$. The corresponding fixed point x satisfies $x_i > 0$ for all $i \in \sigma$, and $dx_i/dt = 0$. In particular, setting $dx_j/dt = 0$ and $dx_k/dt = 0$ (and recalling $W_{jj} = W_{kk} = 0$) we obtain:

$$\begin{aligned} x_j &= \sum_{i \in \sigma \setminus \{j, k\}} W_{ji}x_i + W_{jk}x_k + \theta, \\ x_k &= \sum_{i \in \sigma \setminus \{j, k\}} W_{ki}x_i + W_{kj}x_j + \theta. \end{aligned}$$

Now observe that for each $i \in \sigma \setminus \{j, k\}$, the fact that $i \rightarrow j$ implies $i \rightarrow k$ tells us that $W_{ji} \leq W_{ki}$, (see Eq. (2)). This means the summation term in the x_j equation above is less than or equal to the analogous term in the x_k equation. Using this fact, we see that $x_j - W_{jk}x_k \leq x_k - W_{kj}x_j$, which can be rearranged as,

$$(1 + W_{kj})x_j \leq (1 + W_{jk})x_k.$$

Now recall that $j \rightarrow k$ but $k \not\rightarrow j$, so $W_{kj} = -1 + \varepsilon$ and $W_{jk} = -1 - \delta$. The above inequality thus says that $\varepsilon x_j \leq -\delta x_k$. But this is a contradiction, because $\varepsilon x_j > 0$ and $-\delta x_k < 0$. And so no fixed point supported on σ can exist. $\square$

The conditions on $j, k \in \sigma$ used in the above lemma is an example of so-called *graphical domination*. This notion was first defined in [8], and provides a useful tool for ruling in and ruling out fixed points of CTLNs purely based on the graph structure, and independently of the ε, δ and θ parameters.

Definition 2.7 (graphical domination) Let G be a graph, $\sigma \subseteq V_G$ a subset of the vertices, and $j, k \in V_G$ such that $\{j, k\} \cap \sigma \neq \emptyset$. We say that k *graphically dominates* j with respect to σ , and write $k >_\sigma j$, if the following three conditions hold:

1. For all $i \in \sigma \setminus \{j, k\}$ if $i \rightarrow j$, then $i \rightarrow k$.
2. If $j \in \sigma$, then $j \rightarrow k$.
3. If $k \in \sigma$, then $k \not\rightarrow j$.

This definition of domination covers more cases than what we saw in Lemma 2.6. This greater generality is reflected in the main theorem about domination, which appeared as Theorem 4 in [8]. We cite a special case of this theorem below (also illustrated in Fig. 4).

Theorem 2.8 (graphical domination [8]) Let $\sigma \subseteq V_G$ be a subset of the vertices of a graph G . If there is a $j \in \sigma$ and a $k \in V_G$ such that $k >_\sigma j$ (k graphically dominates j with respect to σ), then $\sigma \notin \text{FP}(G)$.

We will furthermore use the following useful equivalence, which states that σ can only be a fixed point support if $\sigma \in \text{FP}(G|_\sigma)$ and the fixed point survives the addition of each individual $k \notin \sigma$.




Fig. 4 The two cases of graphical domination in Theorem 2.8. In each panel, k graphically dominates j with respect to σ (the outermost shaded region). The inner shaded regions illustrate the subsets of nodes that send edges to j and k . Note that the vertices sending edges to j are a subset of those sending edges to k , but this containment need not be strict. The dashed arrow indicates an optional edge between j and k

Lemma 2.9 ([8, Corollary 2]) *Consider a CTLN determined by a graph G on a set of vertices V_G , and let $\sigma \subseteq V_G$. Then*

$$\sigma \in \text{FP}(G) \Leftrightarrow \sigma \in \text{FP}(G|_{\sigma \cup \{k\}}) \text{ for all } k \in V_G.$$

In particular, $\sigma \in \text{FP}(G) \Rightarrow \sigma \in \text{FP}(G|_{\sigma})$. Moreover, $\sigma \in \text{FP}(G) \Rightarrow \sigma \in \text{FP}(G|_{\tau})$ for any τ with $\sigma \subseteq \tau$.

One simple case where graphical domination can be used to rule out a fixed point support is whenever a graph contains a proper source. This is Rule 6 in [8].

Lemma 2.10 (sources [8]) *Let G be a graph and $\sigma \subseteq V_G$. If there exists a $j \in \sigma$ such that j is a proper source in $G|_{\sigma}$ or j is a proper source in $G|_{\sigma \cup \{\ell\}}$ for some $\ell \in V_G$, then $\sigma \notin \text{FP}(G)$.*

Proof If j is a proper source in $G|_{\sigma}$, then there exists $k \in \sigma$ such that $j \rightarrow k$. Since j has no other inputs in σ , clearly $k >_{\sigma} j$. If j is not a proper source in $G|_{\sigma}$ but is a proper source in $G|_{\sigma \cup \{\ell\}}$, then $j \rightarrow \ell$, and hence $\ell >_{\sigma} j$. In either case, by Theorem 2.8 we have that $\sigma \notin \text{FP}(G)$. $\square$

The lemma above allows us to rule out fixed points of cycles that are not full support.

Lemma 2.11 (cycles) *If G is a cycle on n vertices, then G has a unique fixed point, which has full support. In other words,*

$$\text{FP}(G) = \{[n]\}.$$

Proof First observe that $[n] \in \text{FP}(G)$ by Lemma 2.5 because a cycle has uniform in-degree 1. To see that this is the only fixed point support of G , consider any proper subset $\sigma \subsetneq V_G$. Since G is a cycle, $G|_{\sigma}$ either contains a path or is an independent set. If it contains a path, then the source of that path is a proper source in $G|_{\sigma}$. If it is an independent set, then for any $i \in \sigma$, we have $i \rightarrow \ell$ in G for $\ell = i + 1$. Then i is a proper source in $G|_{\sigma \cup \{\ell\}}$. Thus by Lemma 2.10, $\sigma \notin \text{FP}(G)$. Thus, $\text{FP}(G) = \{[n]\}$. $\square$

Another simple case when graphical domination can be used to rule out fixed points is whenever σ has a target in G . For $k \in V_G$, we say that k is a *target* of σ if $i \rightarrow k$ for every $i \in \sigma \setminus \{k\}$.

Lemma 2.12 (targets [8]) *Let G be a graph and $\sigma \subseteq V_G$. Suppose $k \in V_G$ is a target of σ .*

1. *If $k \in V_G \setminus \sigma$, then $\sigma \notin \text{FP}(G)$.*
2. *If $k \in \sigma$ and there exists a $j \in \sigma$ such that $k \not\rightarrow j$, then $\sigma \notin \text{FP}(G)$.*

Proof In case 1, it is straightforward to see that $k >_{\sigma} j$ for any $j \in \sigma$. In case 2, we see that for the particular j such that $k \not\rightarrow j$, we have $k >_{\sigma} j$. In either case, by Theorem 2.8 we have that $\sigma \notin \text{FP}(G)$. $\square$

The target lemma allows us to rule out fixed points of cliques that are not full support.

Lemma 2.13 (cliques) *If G is a clique on n vertices, then G has a unique fixed point, which has full support. In other words,*

$$\text{FP}(G) = \{[n]\}.$$

Proof First observe that $[n] \in \text{FP}(G)$ by Lemma 2.5 because a clique has uniform in-degree $n - 1$. To see that this is the only fixed point support of G , consider any proper subset $\sigma \subsetneq V_G$ and let $k \in V_G \setminus \sigma$. Then k is a target of σ , and so by Lemma 2.12, $\sigma \notin \text{FP}(G)$. Thus, $\text{FP}(G) = \{[n]\}$. $\square$

Finally, using a more general form of domination defined in [8], we obtain the following *survival rule* telling us precisely when a uniform in-degree fixed point survives as a fixed point of a larger network (Theorem 5 of [8]):

Theorem 2.14 (uniform in-degree [8]) *Let G be a graph and $\sigma \subseteq V_G$ such that $G|_\sigma$ has uniform in-degree d . For $k \in V_G \setminus \sigma$, let $d_k \stackrel{\text{def}}{=} |\{i \in \sigma \mid i \rightarrow k\}|$ be the number of edges k receives from σ . Then*

$$\sigma \in \text{FP}(G) \iff d_k \leq d \text{ for every } k \in V_G \setminus \sigma.$$

Other than this theorem we will not use the more general form of domination. Therefore, in the remainder of this work when we say domination we mean graphical domination.

2.3 The DAG Decomposition

Two of our main results are nerve lemmas involving *directed acyclic graphs* (DAGs). Recall that a DAG is a graph that has no directed cycles. There is a well known characterization of DAGs in terms of a *topological ordering* of their vertices. In particular, G is a DAG if and only if there exists an ordering of the vertices such that edges in G only go from lower numbered to higher numbered vertices. In other words, if $i \rightarrow j$ then $i < j$; equivalently if $i > j$ then $i \not\rightarrow j$.

Lemma 2.15 (DAGs) *Let G be a DAG and let $\tau = \{\text{sinks of } G\}$. Then the fixed point supports of G are all the nonempty subsets of τ , i.e.*

$$\text{FP}(G) = \mathcal{P}(\tau) \setminus \{\emptyset\},$$

where $\mathcal{P}(\tau)$ denotes the power set of τ .

Proof First to see that $\mathcal{P}(\tau) \setminus \{\emptyset\} \subseteq \text{FP}(G)$, notice that any non-empty subset of τ is an independent set of sinks. An independent set has uniform in-degree 0, and thus

by Theorem 2.14, an independent set produces a fixed point when it has no outgoing edges. Since all the nodes in τ are sinks, every subset of τ has no outgoing edges, and so every subset produces a fixed point support in $\text{FP}(G)$.

Next, to see that no other sets can produce fixed points of G , consider $\sigma \subseteq V_G$ such that $\sigma \not\subseteq \tau$. Let j be the lowest number vertex in $\sigma \setminus \tau$ according to some topological ordering of G . Then j has no incoming edges from other nodes in σ since edges in a DAG can only go from lower numbered vertices to higher number vertices. Moreover, there exists some $\ell \in V_G$ such that $j \rightarrow \ell$ since otherwise j would be a sink, but $j \notin \tau$ by design, which contains all the sinks of G . Thus j is a proper source in $G|_{\sigma \cup \ell}$, and so $\sigma \notin \text{FP}(G)$ by Lemma 2.10. $\square$

Many graphs that are not DAGs nevertheless have a DAG-like structure on a subgraph. This will also be a useful concept for our nerve theorems.

Definition 2.16 (DAG decomposition) Let G be a graph. For $\omega, \tau \subseteq V_G$, we say that (ω, τ) is a *DAG decomposition* of G if $\omega \dot{\cup} \tau$ is a partition of the vertices V_G such that:

1. $G|_{\omega}$ is a DAG,
2. $G|_{\tau}$ contains all sinks of G ,
3. there are no edges from τ back to ω , i.e., $E_G(\tau, \omega) = \emptyset$.

We say a DAG decomposition is *non-trivial* if $\omega \neq \emptyset$. We say a DAG decomposition is *maximal* if ω is as large as possible. More precisely, (ω, τ) is a maximal DAG decomposition if there is no other DAG decomposition (ω', τ') with $\omega \subsetneq \omega'$.

Every graph G that has at least one proper source j has a DAG decomposition with $\omega = \{j\}$ and $\tau = V_G \setminus \{j\}$. But DAG decompositions are most valuable when τ is as small as possible. To minimize the size of τ , we'd like to “grow” ω as much as possible, as in a maximal DAG decomposition. It turns out that there is straightforward procedure for generating a maximal DAG decomposition of a graph, and moreover, the maximal DAG decomposition is in fact unique. Specifically, one can iteratively refine a DAG decomposition by moving any nodes that are proper sources in $G|_{\tau}$ to ω (see Fig. 5). This process will maintain the property that $G|_{\omega}$ is a DAG (each node that is moved to ω will be at the end of the “topological ordering” of the DAG) while also guaranteeing that there are no edges from nodes in τ back to nodes in ω . Finally, the process terminates when there are no nodes in τ that are proper sources in $G|_{\tau}$. It turns out that the τ satisfying $G|_{\tau}$ has no proper sources is both *minimal*, in the sense that $|\tau|$ is smallest and $\tau \subseteq \tau'$ for any other DAG decomposition (ω', τ') , and *unique*. As a result, this process yields the unique maximal DAG decomposition. Note in particular that in any maximal DAG decomposition (ω, τ) , τ cannot have proper sources, because if it did one could move such a vertex to ω , contradicting maximality.

Lemma 2.17 *Suppose that G contains a proper source. Then the DAG decomposition (ω, τ) of G satisfying $G|_{\tau}$ has no proper sources is a maximal DAG decomposition. In particular, G has a unique maximal DAG decomposition.*




Fig. 5 Iterative construction of DAG decompositions. (a) A DAG decomposition of a graph where ω contains only a single source. The gray highlighted node 2 is a proper source in $G|_\tau$, but not a source in the full graph. (b) A second DAG decomposition is obtained by moving node 2 to ω . Now node 5 has become a proper source in the new $G|_\tau$. (c) A third DAG decomposition is obtained by moving 5 to ω . In this decomposition, $G|_\tau$ has no proper sources. Notice that node 7 is a source, but because it has no outgoing edges, it is not a proper source so will not be moved to ω . In fact, node 7 is a sink, and thus is required to be in τ by condition 2 of DAG decompositions. We have thus arrived at the unique DAG decomposition with minimal τ and maximal DAG ω

Proof Let (ω, τ) be a DAG decomposition of G satisfying $G|_\tau$ has no proper sources, and let (ω', τ') be any other DAG decomposition of G . Suppose $\omega' \not\subseteq \omega$. Then since each DAG decomposition is a partition of the vertices, this condition on ω implies that $\tau \not\subseteq \tau'$. Then there exists a node $i_0 \in \tau \setminus \tau'$. Since $i_0 \in \tau$ and $G|_\tau$ has no proper sources, there exists some $i_1 \in \tau$ such that $i_1 \rightarrow i_0$. Since i_1 also is not a proper source in $G|_\tau$, there exists some $i_2 \in \tau$ such that $i_2 \rightarrow i_1 \rightarrow i_0$. Again i_2 is not a proper source, and so there exists $i_3 \in \tau$ such that $i_3 \rightarrow i_2 \rightarrow i_1 \rightarrow i_0$.

Note that in the other DAG decomposition (ω', τ') , since $i_0 \notin \tau'$, we must have $i_0 \in \omega'$. Moreover, by the definition of DAG decomposition, there are no edges from nodes in τ' to nodes in ω' , and so all nodes in the path $i_3 \rightarrow i_2 \rightarrow i_1 \rightarrow i_0$ must also be in ω' . We can continue to trace the path backwards in this way through $G|_\tau$ for arbitrarily many steps since it has no proper sources, but since τ is finite, at some point some node must appear twice in this path. Thus this sequence of nodes must contain a bidirectional edges and/or a directed cycle. But all the nodes in this sequence must be in ω' , and by definition, $G|_{\omega'}$ must be a DAG, thus it cannot contain any bidirectional edges or directed cycles. Thus we have a contradiction, and so $\tau \subseteq \tau'$, and thus $\omega' \subseteq \omega$.

Since any DAG decomposition (ω, τ) of G satisfying $G|_\tau$ has no proper sources must have maximal ω , it follows that there must be a unique decomposition satisfying this property. Finally, since any maximal DAG decomposition must satisfy this property, it follows there is a unique one and it is this one. $\square$

3 Directional Graphs

In this section, we focus on a special class of graphs known as *directional graphs*, first defined in [16]. The motivating heuristic behind *directional graphs* is that they

are graphs whose vertices can be partitioned into two sets ω and τ such that when the neural activity is initialized on nodes in ω , it flows to the nodes in τ . In simulations, we have seen that this flow of activity occurs whenever the fixed points of G are confined to live in τ , so that $\text{FP}(G) \subseteq \text{FP}(G|_\tau)$. In order to guarantee nice properties when we union together directional graphs, we require something slightly stronger in our definition of directional graphs, namely that the collapse of the fixed points onto the subnetwork $G|_\tau$ be the result of graphical domination.

Definition 3.1 (directional graph) We say that a graph G is *directional*, with direction $\omega \rightarrow \tau$, if $\omega \dot{\cup} \tau = V_G$ is a nontrivial partition of the vertices ($\omega, \tau \neq \emptyset$, $\omega \cap \tau = \emptyset$) such that $\text{FP}(G) \subseteq \text{FP}(G|_\tau)$ by way of graphical domination. Specifically, we require the following property: for every $\sigma \not\subseteq \tau$, there exists some $j \in \sigma \cap \omega$ and $k \in V_G$ such that k graphically dominates j with respect to σ , i.e. $k >_\sigma j$. When this is the case we say σ *dies by (graphical) domination*.

As mentioned above, we predict that directional graphs will have *feedforward dynamics*, so that activity that is initially concentrated on $G|_\omega$ should flow towards $G|_\tau$, giving the dynamics an $\omega \rightarrow \tau$ directionality. The most natural examples of directional graphs are those where G has an explicit feedforward architecture in $G|_\omega$, for example when $G|_\omega$ is a DAG, and there are no edges from τ back to ω . In this case, it seems intuitive that the dynamics will flow along this feedforward structure in ω and end up concentrated in τ .

It turns out that any DAG decomposition of a graph G immediately yields a directional partition as intuitively predicted.

Lemma 3.2 *If (ω, τ) is a DAG decomposition of G , then G is directional with direction $\omega \rightarrow \tau$.*

The key to the proof of Lemma 3.2 is the well known characterization of DAGs in terms of a *topological ordering* of their vertices. Recall that G is a DAG if and only if there exists an ordering of the vertices such that edges in G only go from lower numbered to higher numbered vertices, i.e., if $i > j$, then $i \not\rightarrow j$.

Proof To show that G is directional, we must show that any $\sigma \subseteq V_G$ that intersects ω dies by graphical domination. Suppose $\sigma \cap \omega \neq \emptyset$, and let j be the lowest numbered vertex in $\sigma \cap \omega$ with respect to some topological ordering of the DAG $G|_\omega$. Since all the sinks in G are contained in τ , $j \in \omega$ must have at least one outgoing edge in G , so $j \rightarrow k$ for some $k \in V_G$. Moreover, j has no incoming edges in $G|_{\sigma \cup k}$ because of its numbering in the topological ordering. Thus, j is a proper source in $G|_{\sigma \cup k}$ and by Lemma 2.10, $\sigma \notin \text{FP}(G)$ because $k >_\sigma j$. Thus every σ with $\sigma \cap \omega \neq \emptyset$ dies by graphical domination, and so G is directional with direction $\omega \rightarrow \tau$. $\square$

DAG decompositions are a very special case of directional graphs $\omega \rightarrow \tau$ where there are no back edges from τ to ω , and the ω component of the graph is a DAG. Neither condition needs to hold for more general directional graphs. Figure 6 shows several types of directional graphs. In panel a, there are only edges from $\omega \rightarrow \tau$ as in a DAG decomposition. In panel b, the existence of a target in τ that receives edges




Fig. 6 Three types of directional graphs. (a) A nontrivial DAG decomposition (ω, τ) is a directional graph with direction $\omega \rightarrow \tau$. (b) If τ contains a target node of ω , and there are no back edges $\tau \rightarrow \omega$, then G is directional irrespective of the structure of $G|_\omega$. (c) A more general directional graph can have a variety of forward and backward edges




Fig. 7 Directional graphs: examples and non-examples. Note that we refer to fixed point supports $\{i_1, i_2, \dots, i_k\}$ simply as $i_1 i_2 \dots i_k$. For example, 234 denotes $\{2, 3, 4\}$

from all nodes in ω guarantees that G is directional irrespective of the structure of $G|_\omega$. Finally, in panel c we see a schematic of a directional graph with both forward edges from ω to τ and backward edges from τ to ω .

In fact, directional graphs can have a surprisingly large number of back edges while still preserving their “forward” directionality. All the graphs in Fig. 7a are directional with $\omega \rightarrow \tau$, and each of the graphs in a3–a6 actually has as many back edges from τ to ω as it does forward edges. The dynamics for a3 and a6 are shown on the right, and we see that even if we initialize the activity purely on nodes in ω , the activity flows $\omega \rightarrow \tau$ as predicted by the directionality. Note that none of the graphs in panel a has a proper source, and thus none has a nontrivial DAG decomposition.

Panel b in Fig. 7 shows some example graphs that are *not* directional for any partition of the vertices. This is because every vertex is involved in at least one fixed point support, so there cannot be a collapse $\text{FP}(G) \subseteq \text{FP}(G|_\tau)$ for any $\tau \subsetneq V_G$. The graph in b2 is particularly surprising since it only has edges forward from the 2-clique 12 to 34, so we might expect this to yield a directional decomposition. But the forward edges are not sufficient to kill the 2-clique 12, and so we see from the dynamics on the right, that we are not guaranteed a directionality of flow. Instead, 12 supports a stable fixed point, and thus when we initialize activity on those nodes, it remains there, and never flows to the other stable fixed point 34.

Remark 3.3 If G is a directional graph, its directional decomposition is not unique. For example, as long as ω has more than one vertex, then vertices can always be moved from it to τ and maintain directionality. However, directional decompositions are most useful when the τ component is as small as possible, since this gives the strongest restrictions on the possible fixed point supports of the whole network. One candidate τ for such a decomposition is $\tau := \cup_{\sigma \in \text{FP}(G)} \sigma$. However, this set does not guarantee a directional decomposition since we have not guaranteed that all subsets of V_G that intersect $\omega := V_G \setminus \tau$ die by graphical domination. In order to satisfy this property, it may be necessary to add some additional vertices to τ , and doing this in a minimal way may not be unique. It is an open question if every directional graph has a unique directional decomposition with minimal τ .

4 Directional Covers and Nerve Theorems

In this section, we aim to characterize the fixed points of more complex graphs by covering the graph with directional graphs, and then analyzing a simpler associated object known as the *nerve* of the cover. The intuition is as follows. As described in Sect. 3, if G is a directional graph with direction $\omega \rightarrow \tau$, the activity of the network flows from ω to τ . Thus, from a bird's eye view, the flow of activity of such a graph can be represented by the flow of activity along a single directed edge from source to sink. Moreover, this flow of activity reflects restrictions imposed on the fixed point supports as well. With this in mind, we will take any graph G and aim to cover it with directional graphs that have appropriate pairwise intersections. From this cover, we construct a nerve, which is a simplified graph where subsets of vertices are collapsed to single points, and each directional graph of the cover is now represented by a single directed edge. These edges are glued to one another in a way representative of the intersection pattern of the cover. We will see that, with this construction, we are able to deduce certain restrictions on the fixed point supports of the original graph G by studying the fixed points of the nerve of the cover, which is in general a simpler graph.

4.1 Directional Covers and Nerves

We begin by making the notion of directional cover and its nerve precise.

Definition 4.1 (graph cover) Let G be a graph. A *graph cover* of G is a collection of induced subgraphs $\mathcal{U} = \{G_i := G|_{V_i} \mid \text{for some } V_i \subseteq V_G\}$ such that G is the union of the G_i . In other words, $V_G = \cup_{i \in I} V_i$ and $E_G = \cup_{i \in I} E_{G_i}$.

Remark 4.2 Note that every vertex and every edge of G must live in at least one G_i , but often they live in multiple G_i within the cover. In particular, since the covering graphs are induced subgraphs of G , if $u, v \in V_i$ and $u, v \in V_j$, then any edges between u and v will be in both G_i and G_j .

Next we turn to a special type of graph cover which we call *rigid directional cover*. In a rigid directional cover, we require that all the graphs of the cover are directional and that they overlap in prescribed ways that will facilitate associating a nerve to the cover and ensure that this nerve captures constraints on $\text{FP}(G)$.

The rigid condition can be informally described as follows. Consider a graph cover $\mathcal{U}$ of G , where all the covering graphs are directional. Let $G_1, G_2 \in \mathcal{U}$ be a pair of graphs in the cover, with directional decompositions $\omega_1 \rightarrow \tau_1$ and $\omega_2 \rightarrow \tau_2$, respectively. The graph cover $\mathcal{U}$ is *rigid* if for any pair G_1 and G_2 that have nontrivial intersection, their overlap is of one of the following three types:

1. The τ component of the first graph acts as the ω component of the second, i.e., $V_{G_1} \cap V_{G_2} = \tau_1 = \omega_2$. In this case we say the graphs have a *chaining overlap*. (See Fig. 8a.)
2. The two covering graphs intersect exactly at their τ component, i.e., $V_{G_1} \cap V_{G_2} = \tau_1 = \tau_2$. In this case we say the graphs have a *merging overlap*. (See Fig. 8b.)
3. The two covering graphs intersect exactly at their ω component, i.e., $V_{G_1} \cap V_{G_2} = \omega_1 = \omega_2$, and have the additional property that there are no back edges from vertices in τ to vertices in ω in either graph, i.e., $E_{G_1}(\tau_1, \omega_1) = E_{G_2}(\tau_2, \omega_2) = \emptyset$. In this case we say the graphs have a *splitting overlap*. (See Fig. 8c.)




Fig. 8 A pair of graphs G_1 and G_2 that have (a) a chaining overlap (b) a merging overlap and (c) a splitting overlap

Effectively, a rigid directional cover is always induced by a partition of the vertices of the underlying graph and this partition encodes all the information of the cover itself. Therefore, we formally define a rigid graph cover as follows.

Definition 4.3 (directional cover and its nerve) Let G be a graph. Given a partition of the vertices, $\nu = \{v_1, \dots, v_n\}$, let

$$E = E(G, \nu) := \{(i, j) \in [n] \times [n] \mid G|_{v_i \cup v_j} \text{ is directional with direction } v_i \rightarrow v_j\},$$

$$\mathcal{I} = \mathcal{I}(G, \nu) := \{i \in [n] \mid G|_{v_i} \text{ is disconnected from the rest of the graph}\}.$$

We say that the partition $\{v_1, \dots, v_n\}$ induces a *rigid directional cover* of G if:

1. For every pair (v_i, v_j) either $(i, j) \in E$ or $(j, i) \in E$ or there are no edges between v_i and v_j . In other words, the set $\mathcal{U} = \{G_{ij} := G|_{v_i \cup v_j} \mid (i, j) \in E\} \cup \{G|_{v_i} \mid i \in \mathcal{I}\}$ is a graph cover of G .
2. Whenever $G_{ij}, G_{ik} \in \mathcal{U}$, they have “splitting overlap”, meaning there are no edges from v_j to v_i and no edges from v_k to v_i , i.e., $E_G(v_j, v_i) = E_G(v_k, v_i) = \emptyset$.

We define the *nerve* of the cover, denoted by $\mathcal{N} = \mathcal{N}(G, \mathcal{U})$, to be the graph with vertex set $V_{\mathcal{N}} := [n]$ and edge set $E_{\mathcal{N}} := E$. The partition ν induces a canonical quotient map $\pi : V_G \rightarrow V_{\mathcal{N}}$ that identifies all the vertices of a component v_i , so that $\pi(v_i) = \{i\}$ for each $i \in V_{\mathcal{N}}$.

Note that an arbitrary partition will not typically induce a rigid directional cover because there will be pairs (v_i, v_j) with edges between them, but the induced subgraphs $G|_{v_i \cup v_j}$ will not be directional. In contrast, partitions that do induce a rigid directional cover must have $G|_{v_i \cup v_j}$ directional whenever there are edges between v_i and v_j . Note that we do not require the $G|_{v_i}$ for $i \in \mathcal{I}$ to be directional; these graphs are included in $\mathcal{U}$ simply to ensure that isolated components of the graph are still covered. In this paper we will only work with rigid directional covers. Therefore, in the remainder of this work we will use the term *directional cover* to refer to a rigid directional cover.

Figure 9 gives an illustration of a graph G with vertex partition $\nu = \{v_i\}$ that induces a collection of covering graphs $\{G_{ij}\}$ that are all directional. In the nerve, $\mathcal{N} = \mathcal{N}(G, \mathcal{U})$, we see a vertex i for each component v_i from G , and an edge $i \rightarrow j$ corresponding to each covering graph G_{ij} , which has direction $v_i \rightarrow v_j$.

Lemma 4.4 *Let G be a graph with nerve $\mathcal{N} = \mathcal{N}(G, \mathcal{U})$ for some directional cover $\mathcal{U}$. Then the nerve $\mathcal{N}$ is a simple directed graph that is oriented.*

Proof Recall that given a partition $\nu = \{v_1, \dots, v_n\}$ of the vertices of G , the nerve $\mathcal{N}$ is defined as a graph on n vertices with edge set E given in Definition 4.3. The edge set E is straightforward to determine from the partition ν . Specifically, if there are no edges between v_i and v_j in G , then neither (i, j) nor (j, i) are in E , since $G|_{v_i \cup v_j}$ is a disjoint union, which can never be directional. If there are any edges




Fig. 9 A graph with a directional cover and the corresponding nerve. (Left) A graph G . (Middle) A directional cover $\mathcal{U}$ of G . The sets v_i, v_j, v_k, v_l, v_m partition the vertices of G . All the edges of G are contained within some of the covering graphs $\{G_{ij}\}$ shown, each of which is directional with direction $v_i \rightarrow v_j$. (Right) The nerve $\mathcal{N} = \mathcal{N}(G, \mathcal{U})$ with a vertex for each component of the partition and a directed edge for each directional graph of the cover

between v_i and v_j , we must have either $(i, j) \in E$ or $(j, i) \in E$ in order for the $\{G_{ij}\}$ to cover G . Moreover, we can only have one of (i, j) or (j, i) in E since $G|_{v_i \cup v_j}$ can never be directional with both $v_i \rightarrow v_j$ and $v_j \rightarrow v_i$. Thus, $\mathcal{N}$ is oriented. Finally, to see that $\mathcal{N}$ is simple, notice that $(i, i) \notin E$ since G_{ii} can never be directional with $\omega = \tau = v_i$. $\square$

Given the complexity of the requirements of a directional cover, specifically that every covering graph of the form $G|_{v_i \cup v_j}$ be directional, it is natural to ask when a graph actually has such a cover. Of course, every graph has a trivial directional cover induced by the trivial partition $v_1 = V_G$; in this case, the nerve of the cover is just a single point. At the other extreme, whenever G is an *oriented* graph, the partition of singletons $v_i = \{i\}$ will induce a directional cover, whose nerve is precisely the original graph G . While these two trivial covers exist, they clearly do not provide any insight into the structure or expected dynamics of G . There is an art to finding a partition of V_G from which a cover with an informative nerve can be obtained.

It is important to note that not every graph has a directional cover induced by a nontrivial partition. For example, if G is a clique, then there is no nontrivial partition of V_G that can admit a directional cover since every G_{ij} will be a clique, and thus not directional. At the other extreme, there are graphs with multiple nontrivial partitions of V_G which induce directional covers. See for example, Fig. 10c,d, which shows two different covers of the same graph, and where the nerve of the first one is directional while the nerve of the second is a cycle.

It is an open question which graphs have at least one nontrivial partition that admits a directional cover. And unfortunately, there is currently no efficient way to find all the partitions of a graph that do induce directional covers. However, when we have a directional cover of G , obtained either by brute force search or from intuition into the original construction of the graph, the nerve of the cover can give significant insight into the collection of fixed points and consequently into the predicted dynamics of the underlying network. In particular, nerve theorems ensure that there is a provable connection between $\text{FP}(G)$ and the structure of the nerve under certain conditions on G and/or $\mathcal{N}(G, \mathcal{U})$.




Fig. 10 Counterexamples to general nerve theorems. (a) and (b) give two different graphs with partitions that induce directional covers. The nerve of each cover is shown to the right. (c) and (d) give two different partitions for the *same* graph. In (c), the partition induces a nerve that is directional, while in (d), the partition induces a nerve that is a cycle

4.2 Nerve Theorems

Ideally, we would hope for a nerve theorem that provides a strong connection between the fixed point supports of the original graph and those of the nerve. For example, we might hope that for any graph G that admits a directional cover with nerve $\mathcal{N}$ we can guarantee a condition such as

$$\sigma \in \text{FP}(G) \Rightarrow \pi(\sigma) \in \text{FP}(\mathcal{N}). \quad (4)$$

Unfortunately, though, this strong restriction on $\text{FP}(G)$ does not hold for all graphs and all directional covers.

Example 4.5 For the graph in Fig. 10a, we see that the partition $\{\nu_1, \dots, \nu_4\}$ shown induces a directional cover of G . The nerve $\mathcal{N} = \mathcal{N}(G, \mathcal{U})$ of this cover is shown on the right. Recall that a cycle is *uniform in-degree 1* and thus it supports a fixed point precisely when no external vertex receives more than one edge from it (see Theorem 2.14 in Sect. 2). Thus we have $123 \in \text{FP}(G)$ ⁴ since the external vertices 4 and 5 each receive only one edge from the cycle. However,

⁴ Recall that we write 123 to denote the fixed point support $\{1, 2, 3\}$.

$\pi(123) = 123 \notin \text{FP}(\mathcal{N})$ since in the nerve, the cycle 123 has two outgoing edges to vertex 4. Thus, $\sigma \in \text{FP}(G) \not\Rightarrow \pi(\sigma) \in \text{FP}(\mathcal{N})$ for this example graph.

An alternative style of nerve theorem would enable us to at least restrict the candidate fixed point supports of G based on the structure of the nerve. For example, we might hope that whenever the nerve $\mathcal{N}(G, \mathcal{U})$ is directional with direction $\mathcal{W} \rightarrow \mathcal{T}$ that the directionality would pullback to guarantee G is directional with direction $\omega \rightarrow \tau$ for $\omega = \pi^{-1}(\mathcal{W})$ and $\tau = \pi^{-1}(\mathcal{T})$. If this held, then we could guarantee that the fixed point supports of G were confined to $\tau = \pi^{-1}(\mathcal{T})$, and thus $\text{FP}(G) \subseteq \text{FP}(G|_{\tau})$ by Lemma 2.9. Such a result would be somewhat weaker than (4) in that the restrictions on $\sigma \in \text{FP}(G)$ would not be as strong. On other hand, it would also be somewhat stronger in a different direction, since a result like this would guarantee the presence of graphical domination relationships for ruling out fixed point supports, which is not something guaranteed by (4). Unfortunately, this alternative nerve theorem does not hold in general either, and the same graph from Fig. 10a provides a counterexample, as do the other graphs in Fig. 10.

Example 4.6 It is straightforward to see that the nerve $\mathcal{N}$ in Fig. 10a is directional for $\mathcal{W} := \{1, 2, 3\}$ and $\mathcal{T} := \{4\}$: every subset $\mathcal{S} \subseteq V_{\mathcal{N}}$ that intersects $\mathcal{W}$ either has a proper source in $G|_{\mathcal{S}}$, and thus dies from domination by Lemma 2.10, or contains 123, in which case we have $4 >_{\mathcal{S}} 1$. However, G is not directional since $12345 \in \text{FP}(G)$, so there is no collapse of the fixed point supports of G onto a proper subset τ . Figure 10c gives another counterexample where G has a full support fixed point but a directional cover whose nerve is directional.

We have seen that in general the existence of a directional relationship $\mathcal{W} \rightarrow \mathcal{T}$ of the nerve does not guarantee directionality of G . But are there certain conditions under which this holds? It turns out that we can pullback such a directionality relationship in the special case when the nerve has a nontrivial DAG decomposition (see Definition 2.16).

Theorem 4.7 (DAG decomposition of the nerve) *Let G be a graph with nerve $\mathcal{N} = \mathcal{N}(G, \mathcal{U})$ where $\mathcal{U}$ is a directional cover induced by a partition $\{v_1, \dots, v_n\}$, and let $\pi : V_G \rightarrow V_{\mathcal{N}} = [n]$ be the canonical quotient map of the partition. Then for any DAG decomposition $(\mathcal{W}, \mathcal{T})$ of the nerve $\mathcal{N}$, we have that G is directional with direction $\omega \rightarrow \tau$ for $\omega = \pi^{-1}(\mathcal{W})$ and $\tau = \pi^{-1}(\mathcal{T})$. In particular,*

$$\text{FP}(G) \subseteq \text{FP}(G|_{\tau}),$$

and so for all $\sigma \in \text{FP}(G)$, we have $\pi(\sigma) \subseteq \mathcal{T}$.

Notice that in Theorem 4.7, we can only conclude that $\sigma \in \text{FP}(G) \Rightarrow \pi(\sigma) \subseteq \mathcal{T}$, and so we do not quite have the ideal nerve theorem result that $\pi(\sigma) \in \text{FP}(\mathcal{N})$ as in (4). The conclusion in Theorem 4.7 is weaker than (4) because although the directionality of $\mathcal{N}$ guarantees that every element of $\text{FP}(\mathcal{N})$ is contained in $\mathcal{T}$ as is each $\pi(\sigma)$, we cannot guarantee that $\pi(\sigma)$ is actually a fixed point support of $\mathcal{N}$.




Fig. 11 Example directional covers and nerves. **(a1–c1)** Graphs with simple directional covers in which every pair of covering graphs have the same type of overlap (chaining overlap in **(a)**, merging overlap in **(b)**, and splitting overlap in **(c)**). **(a2–c2)** Nerves for the simple directional covers above

Next we consider when the nerve $\mathcal{N}$ is itself a DAG so that, in the maximal DAG decomposition, $\mathcal{T}$ is precisely the sinks of $\mathcal{N}$. We can immediately apply Theorem 4.7 to see that the directionality of $\mathcal{N}$ pulls back to G , but in fact we can say something stronger: the ideal nerve theorem conditions of (4) hold in this case. Moreover, it turns out that there are further restrictions on the fixed point supports of G in terms of the fixed points of the component subgraphs $G|_{v_i}$, which are prescribed by the partition.

Theorem 4.8 (DAG nerve) *Let G be a graph with nerve $\mathcal{N} = \mathcal{N}(G, \mathcal{U})$ where $\mathcal{U}$ is a directional cover induced by a partition $\{v_1, \dots, v_n\}$, and let $\pi : V_G \rightarrow V_{\mathcal{N}} = [n]$ be the canonical quotient map of the partition. Suppose that $\mathcal{N}$ is a DAG, and let $\mathcal{T} = \{\text{sinks of } \mathcal{N}\}$ and $\mathcal{W} = V_{\mathcal{N}} \setminus \mathcal{T}$. Then G is directional with direction $\omega \rightarrow \tau$ for $\omega = \pi^{-1}(\mathcal{W})$ and $\tau = \pi^{-1}(\mathcal{T})$.*

Moreover,

1. $\sigma \in \text{FP}(G) \Rightarrow \pi(\sigma) \in \text{FP}(\mathcal{N}) = \mathcal{P}(\mathcal{T}) \setminus \{\emptyset\}$, where $\mathcal{P}(\mathcal{T})$ denotes the power set of $\mathcal{T}$.
2. $\sigma \in \text{FP}(G) \Rightarrow \sigma \cap v_i \in \text{FP}(G|_{v_i}) \cup \{\emptyset\}$ for all $i \in \mathcal{T}$ and $\sigma \cap v_j = \emptyset$ for all $j \in \mathcal{W}$.

Figure 11 illustrates three special cases of directional covers $\mathcal{U}$ of a graph G with their corresponding nerves shown below. These directional covers have pairwise overlaps that are either: only chainings (Fig. 11a), only mergings (Fig. 11b), or only splittings (Fig. 11c). We refer to these types of overlaps as n -chaining, n -merging and n -splitting, respectively. We see that their corresponding nerves are DAGs where the set of sinks $\mathcal{T}$ is either a single sink, $\mathcal{T} = \{n\}$, as in panels a2 and b2, or an independent set of sinks, $\mathcal{T} = \{2, 3, \dots, n\}$, as in panel c2. Theorem 4.7 tells us that in each case the underlying graph G is directional with direction $\omega \rightarrow \tau$ for $\tau = \pi^{-1}(\mathcal{T})$. Additionally, in the case of n -splitting, Theorem 4.8 gives a stronger result. Namely,

$$\sigma \in \text{FP}(G) \Rightarrow \pi(\sigma) \in \text{FP}(\mathcal{N}) \text{ and } \sigma \cap v_i \in \text{FP}(G|_{v_i}) \cup \{\emptyset\} \text{ for } i \in \{2, 3, \dots, n\}.$$

That is: any fixed point support σ of G gets pushed forward to a fixed point support of the nerve $\mathcal{N}$, and for any $2 \leq i \leq n$ we have that if σ intersects v_i then this intersection is also a fixed point support of the induced subgraph on v_i .

Thus far, we have only seen nerve theorems in the case when $\mathcal{N}$ has a nontrivial DAG component, but it turns out that a similar nerve result holds in the case when $\mathcal{N}$ is a cycle (and thus has no DAG component).

Theorem 4.9 (cycle nerve) *Let G be a graph with nerve $\mathcal{N} = \mathcal{N}(G, \mathcal{U})$ where $\mathcal{U}$ is a directional cover induced by a partition $\{v_1, \dots, v_n\}$, and let $\pi: V_G \rightarrow V_{\mathcal{N}} = [n]$ be the canonical quotient map.*

Suppose that $\mathcal{N}$ is a cycle on n vertices. Then

1. $\sigma \in \text{FP}(G) \Rightarrow \pi(\sigma) \in \text{FP}(\mathcal{N}) = \{[n]\}$
2. *If $\{v_1, \dots, v_n\}$ is a simply-embedded partition of G , then*
 $\sigma \in \text{FP}(G) \Rightarrow \sigma \cap v_i \in \text{FP}(G|_{v_i}) \text{ for all } i \in [n].$

Theorem 4.9 is a repackaging of results from [16], which explores graphs known as *directional cycles*; in the terminology of this paper, these are precisely graphs with a directional cover whose nerve is a cycle. In [16, Theorem 1.2], it was shown that for this family of graphs, every fixed point support must nontrivially intersect every v_i , and so for every $\sigma \in \text{FP}(G)$, we have $\pi(\sigma) = [n]$. Recall that when $\mathcal{N}$ is a cycle, $\text{FP}(\mathcal{N}) = \{[n]\}$ by Lemma 2.11, and thus, combining these results, we are guaranteed that $\pi(\sigma) \in \text{FP}(\mathcal{N})$. In [16, Theorem 1.5], it was also shown that when the partition $\{v_1, \dots, v_n\}$ has a special property, known as *simply-embedded*,⁵ then every fixed point support must restrict to a fixed point in each of the component subgraphs $G|_{v_i}$, yielding the second part of Theorem 4.9.

It is worth noting that another special family of graphs with directional covers was previously studied in [8, Section 5]. That work focused on *composite graphs*, which are graphs where all the vertices in a component behave identically with respect to the rest of the graph. Consequently, the only directional covering graphs G_{ij} used in the cover are those that have all possible edges forward from v_i to v_j and no backward edges. In this context, the components of a composite graph correspond to the partition $\{v_1, \dots, v_n\}$ that induces the directional cover, and the *skeleton* of the composite graph is its nerve. With this perspective, many of the results of [8, Section 5] can be reinterpreted as nerve theorems for the special family of composite graphs.

⁵ We say that $\{v_1, \dots, v_n\}$ is a *simply-embedded partition* if every vertex in v_i is treated identically by the rest of the graph. In other words, for every $j \in V_G \setminus v_i$ if $j \rightarrow k$ for some $k \in v_i$, then $j \rightarrow \ell$ for all $\ell \in v_i$.

4.3 Proofs of Nerve Theorems

Throughout this subsection we fix the following notation: G is a graph with a partition $\{v_1, \dots, v_n\}$ that induces a directional cover $\mathcal{U}$. The nerve $\mathcal{N}(G, \mathcal{U})$ is denoted by $\mathcal{N}$, and $\pi : V_G \rightarrow V_{\mathcal{N}} = [n]$ is the canonical quotient map induced by the partition.

Before proving the nerve theorems we give an overview of the structure of the proofs. To prove Theorem 4.7 (DAG decomposition of the nerve), we first show that whenever $\mathcal{N}$ has a proper source s , we can guarantee that G is directional for $\omega = \pi^{-1}(s)$ and $\tau = V_G \setminus \omega$ (see Lemma 4.11). This gives a rather coarse directional decomposition of G . We will then consider the general case when we have a DAG decomposition $(\mathcal{W}, \mathcal{T})$ of the nerve $\mathcal{N}$. We will use the previous result and show inductively that G is directional with direction $\pi^{-1}(\mathcal{W}) \rightarrow \pi^{-1}(\mathcal{T})$. For this proof, we will use three ingredients we briefly describe now.

First, we will use a *topological ordering* on $\mathcal{W}$, which guarantees that the only possible edges in $\mathcal{N}|_{\mathcal{W}}$ are from lower numbered vertices to higher number vertices. With respect to this topological ordering, vertex 1 in $\mathcal{N}$ is a proper source; vertex 2 is a proper source in $\mathcal{N}|_{V_{\mathcal{N}} \setminus \{1\}}$; vertex 3 is a proper source in $\mathcal{N}|_{V_{\mathcal{N}} \setminus \{1,2\}}$, and so on. This ordering will allow us to induct on $|\mathcal{W}|$. The second ingredient is Lemma 4.12. This result will allow us to refine a directional decomposition of a graph in order to grow ω , and consequently shrink τ by looking at directional decompositions of the subgraph induced on the vertices in τ . The third and final ingredient is Lemma 4.13. This result will show that the construction of a directional cover and its nerve is compatible with taking subgraphs of the underlying graph corresponding to only some of the components of the partition inducing the cover.

At the core of the proofs is the process of “extending domination”. To see why this idea is central, consider a directional cover of G and $\sigma \subseteq V_G$ a subset of the vertices of G . Let μ be the intersection of σ with the vertices of one of the covering graphs such that μ is not contained in the τ component of that covering graph. Since μ is completely contained within a directional graph then it must die by domination. Under certain conditions, one can show that σ dies by domination by extending the domination relationship on μ to all of σ . Therefore, all the proofs hinge on the following technical result that determines when domination can be extended to a superset.

Lemma 4.10 (restricting and extending domination) *Let G be a graph, and let $\alpha \subset \mu \subset \sigma \subseteq V_G$ be subsets of vertices with α possibly empty (see Fig. 12). Then the following hold:*

- (a) *Restriction: Suppose $k >_{\sigma} j$ where $j \in \mu$ and $k \in V_G$. Then $k >_{\mu} j$.*
- (b) *Extension: Suppose $k >_{\mu} j$ where $j \in \mu \setminus \alpha$ and $k \in V_G$. If there are no edges from $\sigma \setminus \mu$ to $\mu \setminus \alpha$ (i.e., $E_G(\sigma \setminus \mu, \mu \setminus \alpha) = \emptyset$), then $k >_{\sigma} j$.*

Proof Recall from Definition 2.7 that $k >_{\sigma} j$ if the following three conditions hold:




Fig. 12 Figure for Lemma 4.10. A graph G with subsets of the vertices $\alpha \subset \mu \subseteq \sigma \subset V_G$ satisfying the conditions of Lemma 4.10. Specifically, the vertex $j \in \mu \setminus \alpha$, the vertex $k \in V_G$, and there are no edges from vertices in σ to those in $\mu \setminus \alpha$

1. For all $i \in \sigma \setminus \{j, k\}$ if $i \rightarrow j$, then $i \rightarrow k$.
2. If $j \in \sigma$, then $j \rightarrow k$.
3. If $k \in \sigma$, then $k \nrightarrow j$.

For (a), *Restriction*, we see that $k >_{\sigma} j$ immediately implies that $k >_{\mu} j$ since if condition 1 holds for all of σ , then it holds for the subset μ as well, and conditions 2 and 3 go through directly.

For (b), *Extension*, we see condition 2 goes through immediately since $j \in \mu \subset \sigma$. For condition 3, observe that if $k \in \sigma$ then either $k \in \mu$ or $k \in \sigma \setminus \mu$, and in both cases $k \nrightarrow j$ as required. Finally, for condition 1, notice that for $i \in \mu$ this condition holds because $k >_{\mu} j$, while for $i \in \sigma \setminus \mu$, this condition holds trivially because $j \in \mu \setminus \alpha$ and there are no edges from $\sigma \setminus \mu$ to $\mu \setminus \alpha$ by hypothesis. Thus condition 1 holds as well, and so $k >_{\sigma} j$. $\square$

We can now prove that it is possible to pull back directionality from the nerve $\mathcal{N}$ to G whenever $\mathcal{N}$ has a proper source.

Lemma 4.11 *If s is a proper source in $\mathcal{N}$, then G is directional with $\omega = \pi^{-1}(s)$ and $\tau = V_G \setminus \omega$.*

Proof Let s be a proper source in $\mathcal{N}$, $\omega := \pi^{-1}(s)$, and $\tau := V_G \setminus \omega$. To show that G is directional with direction $\omega \rightarrow \tau$, consider $\sigma \subseteq V_G$ such that $\sigma \cap \omega \neq \emptyset$. We need to show that σ dies by domination, i.e., that there exists a $j \in \sigma \cap \omega$ and a $k \in V_G$ such that $k >_{\sigma} j$.

The organization of the proof is as follows. We first find a covering graph $G_{s1} \in \mathcal{U}$ such that $\omega \subseteq V_{G_{s1}}$. Then we set $\mu \subseteq \sigma$ to be the restriction of σ to G_{s1} . Since G_{s1} is directional, μ dies by domination. We will show that by setting $\alpha \subset \mu$ to be the intersection of μ with the τ component of G_{s1} , the conditions of Lemma 4.10 hold (see Fig. 13). Thus, we can extend the domination relationship to all of σ .

To find G_{s1} , notice that since s is a proper source in $\mathcal{N}$ there exists at least one vertex in $\mathcal{N}$ that s sends an edge to; without loss of generality, label the vertices of $\mathcal{N}$ that s sends edges to as $1, \dots, m$. Since $s \rightarrow 1$ in $\mathcal{N}$, the covering graph




Fig. 13 Figure for proof of Lemma 4.11. Vertex s is a proper source in $\mathcal{N}$ (right) with outgoing edges to vertices 1, 2, and 3. In G (left), we see the corresponding components ν_s and $\{\nu_1, \nu_2, \nu_3\}$. We consider a subset $\sigma \subseteq V_G$ (outlined in red), and let $\mu := \sigma \cap (\nu_s \cup \nu_1)$ (shaded in purple). Arrows with an x through them indicate that no edges are allowed in the specified direction between the relevant components. Specifically, there are no edges from $\{\nu_1, \nu_2, \nu_3\}$ into ν_s because the graphs G_{s1} , G_{s2} , and G_{s3} must have “splitting overlap” by the definition of directional cover. There are no edges from any other vertices in G into ν_s because there are no edges in the nerve between s and any of the other vertices besides 1, 2, and 3

$G_{s1} := G|_{\nu_s \cup \nu_1}$ must be directional with direction $\nu_s \rightarrow \nu_1$. Let $\mu := \sigma \cap (\nu_s \cup \nu_1)$ (see Fig. 13). Since G_{s1} is directional, there exists a $j \in \mu \cap \nu_s$ and $k \in V_{G_{s1}}$ such that $k >_\mu j$. Following the notation of Lemma 4.10, let $\alpha := \mu \setminus \nu_s = \mu \cap \nu_1$. We will show that there are no edges in G from vertices in $\sigma \setminus \mu$ to vertices in $\mu \setminus \alpha = \mu \cap \nu_s$, enabling us to extend the domination relationship from μ to all of σ .

Note that by definition of directional cover, there can only be edges between ν_s and ν_ℓ in G if there is an edge between s and ℓ in $\mathcal{N}$ (see Definition 4.3). Thus the only candidate vertices in G that could send edges into $\mu \setminus \alpha = \mu \cap \nu_s$ are those in $\{\nu_1, \dots, \nu_m\}$ since the only edges in $\mathcal{N}$ that involve s are those from s to $1, \dots, m$. But since $(s, 1), \dots, (s, m) \in E_{\mathcal{N}}$, condition (2) of the definition of directional cover requires that the covering graphs $G_{s1}, \dots, G_{sm}$ have splitting overlap, so there are no edges from ν_ℓ to ν_s for any $1 \leq \ell \leq m$. Thus, there are no edges from $\sigma \setminus \mu$ into $\mu \cap \nu_s = \mu \setminus \alpha$, and so by Lemma 4.10, the domination relationship $k >_\mu j$ extends to give $k >_\sigma j$. Hence G is directional with direction $\omega \rightarrow \tau$. $\square$

We now give the two lemmas that allow us to inductively use the result above. First, we show that one can refine a directional decomposition of a graph by looking at possible directional decompositions of the subgraph induced on the τ component.

Lemma 4.12 *Suppose that G is a directional graph with direction $\omega_1 \rightarrow \tau_1$ and that $G|_{\tau_1}$ is also directional with direction $\omega_2 \rightarrow \tau_2$ (see Fig. 14). Then G is directional with direction $\omega_1 \cup \omega_2 \rightarrow \tau_2$.*




Fig. 14 Figure for Lemma 4.12. A directional graph G with an initial directional partition $\omega_1 \rightarrow \tau_1$. Additionally, $G|_{\tau_1}$ is directional with direction $\omega_2 \rightarrow \tau_2$. These two partitions can then be combined to show that G is directional for a larger set $\omega = \omega_1 \cup \omega_2$ (outlined in gray) and smaller set $\tau = \tau_2$

Proof Let $\omega = \omega_1 \cup \omega_2$ and $\tau = \tau_2$. To show that G is directional with direction $\omega \rightarrow \tau$, consider $\sigma \subseteq V_G$ such that $\sigma \cap \omega \neq \emptyset$. We will show that σ dies by domination.

Case 1: $\sigma \cap \omega_1 \neq \emptyset$. Then since G is directional with $\omega_1 \rightarrow \tau_1$, σ dies by domination.

Case 2: $\sigma \cap \omega_1 = \emptyset$. Then $\sigma \subseteq \tau_1$. Since $\sigma \cap \omega \neq \emptyset$, we have $\sigma \cap \omega_2 \neq \emptyset$. Then since $G|_{\tau_1}$ is directional with $\omega_2 \rightarrow \tau_2$, there exists $j \in \sigma \cap \omega_2$ and $k \in V_{G|_{\tau_1}} = \tau_1$ such that $k >_\sigma j$ in $G|_{\tau_1}$. Since $\sigma \cap \omega_1 = \emptyset$, there are no vertices in σ outside of $G|_{\tau_1}$ that could potentially subvert the domination relationship between k and j . Thus, $k >_\sigma j$ in all of G , and so G is directional for $\omega = \omega_1 \cup \omega_2$ and $\tau = \tau_2$. $\square$

We now show that the construction of a directional cover and its nerve behaves nicely for induced subgraphs of G , restricting to a subset of the partition components.

Lemma 4.13 *Let $v_I := \{v_i \mid i \in I\}$ be a subset of the components of the partition v of G , for $I \subseteq V_N$. Let $G_I := G|_{\cup_{i \in I} v_i}$ denote the induced subgraph of G on the components v_I . Then the partition v_I of the vertices of G_I induces a directional cover $\mathcal{U}_I$ whose nerve is $N(G_I, \mathcal{U}_I) = N|_I$, the restriction of the original nerve $N = N(G, \mathcal{U})$ to the vertices I .*

Proof Observe that the partition v_I yields the edge set

$$E(G_I, v_I) := \{(i, j) \in I \times I \mid G|_{v_i \cup v_j} \text{ is directional with direction } v_i \rightarrow v_j\}.$$

This edge set is clearly a subset of the edge set $E(G, v)$ for the cover $\mathcal{U}$ of G ; specifically, $E(G_I, v_I) = \{(i, j) \in E(G, v) \mid i, j \in I\}$. Moreover, the set of graphs $\{G_{ij} \mid (i, j) \in E(G_I, v_I)\}$ form a graph cover of G_I because for any $i, j \in I$ if

there are any edges between v_i and v_j in G , then G_{ij} or G_{ji} must have been in the cover $\mathcal{U}$ of G , so one of these graphs is directional, and thus (i, j) or (j, i) is in $E(G_I, v_I)$. Thus, v_I induces a directional cover $\mathcal{U}_I$ of G_I with vertex set I and edge set $E(G_I, v_I)$. Since $E(G_I, v_I)$ is precisely the edges of $E(G, v)$ among the vertices in I , we see that the nerve of the cover $\mathcal{N}(G_I, \mathcal{U}_I)$ is precisely $\mathcal{N}|_I$, the restriction of the nerve of G to the vertices I . $\square$

We are now prepared to prove Theorem 4.7 (reprinted below for convenience).

Theorem 4.7 (DAG decomposition of the nerve) For any DAG decomposition $(\mathcal{W}, \mathcal{T})$ of the nerve $\mathcal{N}$, we have that G is directional with direction $\omega \rightarrow \tau$ for $\omega = \pi^{-1}(\mathcal{W})$ and $\tau = \pi^{-1}(\mathcal{T})$. In particular,

$$\text{FP}(G) \subseteq \text{FP}(G|_{\tau}),$$

and so for all $\sigma \in \text{FP}(G)$, we have $\pi(\sigma) \subseteq \mathcal{T}$.

Proof Let $(\mathcal{W}, \mathcal{T})$ be a DAG decomposition of the nerve $\mathcal{N}$. We will show that G is directional with direction $\pi^{-1}(\mathcal{W}) \rightarrow \pi^{-1}(\mathcal{T})$ by inducting on $|\mathcal{W}|$ in the DAG decomposition of $\mathcal{N}$.

The base case of $|\mathcal{W}| = 1$ follows immediately from Lemma 4.11 since the first element of $\mathcal{W}$ must be a proper source in $\mathcal{N}$. For the inductive step, assume the inductive hypothesis holds whenever $|\mathcal{W}| < m$ and consider a DAG decomposition $(\mathcal{W}, \mathcal{T})$ of $\mathcal{N}$ where $|\mathcal{W}| = m$. Since $\mathcal{N}|_{\mathcal{W}}$ is a DAG, there is a topological ordering of the vertices such that the only edges in $\mathcal{N}|_{\mathcal{W}}$ are from lower numbered to higher numbered nodes; WLOG relabel the vertices of $\mathcal{W}$ as $1, \dots, m$ according to this ordering. Let $\mathcal{W}_1 = \mathcal{W} \setminus \{m\}$ and $\mathcal{T}_1 = \mathcal{T} \cup \{m\}$. It is straightforward to check that $(\mathcal{W}_1, \mathcal{T}_1)$ is also a DAG decomposition of $\mathcal{N}$. Since $|\mathcal{W}_1| < m$, by the inductive hypothesis, G is directional with

$$\omega_1 = \pi^{-1}(\mathcal{W}_1) \text{ and } \tau_1 = \pi^{-1}(\mathcal{T}_1) = \pi^{-1}(m) \cup \pi^{-1}(\mathcal{T}).$$

We will show that $G|_{\tau_1}$ is also directional, so that we may apply Lemma 4.12 and further refine the directional decomposition of G . Specifically, we will show that $G|_{\tau_1}$ has direction $\omega_2 \rightarrow \tau_2$ for $\omega_2 = \pi^{-1}(m)$ and $\tau_2 = \pi^{-1}(\mathcal{T}_1 \setminus \{m\}) = \pi^{-1}(\mathcal{T})$.

By Lemma 4.13, the original directional cover of G restricts to a directional cover of $G|_{\tau_1}$ and its nerve is $\mathcal{N}|_{\{m\}} \cup \mathcal{T}$. Since $m \in \mathcal{W}$ in the original DAG decomposition of $\mathcal{N}$, m is not a sink in $\mathcal{N}$, so it has at least one outgoing edge. Moreover there are no edges from $\mathcal{T}$ back to m in a DAG decomposition. Thus m is a proper source in $\mathcal{N}|_{\{m\}} \cup \mathcal{T}$. Therefore, by Lemma 4.11, $G|_{\tau_1}$ is directional with direction $\omega_2 \rightarrow \tau_2$ for

$$\omega_2 = \pi^{-1}(m) \text{ and } \tau_2 = \pi^{-1}(\mathcal{T}).$$

Finally, since G is directional with $\omega_1 = \pi^{-1}(\mathcal{W} \setminus \{m\})$, $\tau_1 = \pi^{-1}(m) \cup \pi^{-1}(\mathcal{T})$ and $G|_{\tau_1}$ is directional with $\omega_2 = \pi^{-1}(m)$ and $\tau_2 = \pi^{-1}(\mathcal{T})$, we see from Lemma 4.12, that G is directional with direction $\omega_1 \cup \omega_2 \rightarrow \tau_2$. Since $\pi^{-1}(\mathcal{W}) = \omega_1 \cup \omega_2$ and $\pi^{-1}(\mathcal{T}) = \tau_2$, we see G is directional with direction $\pi^{-1}(\mathcal{W}) \rightarrow \pi^{-1}(\mathcal{T})$ as desired. $\square$

Next we consider when the nerve $\mathcal{N}$ is itself a DAG so that in the maximal DAG decomposition $\mathcal{T}$ is precisely the sinks of $\mathcal{N}$. Theorem 4.7 guarantees that the directionality of $\mathcal{N}$ pulls back to G , but to prove the rest of the nerve theorem conditions, we must appeal to a result characterizing the fixed point supports of *disjoint unions*, proven in [8]. The disjoint union of component subgraphs is the graph consisting of those subgraphs with no edges between the components.

Theorem 4.14 ([8], Theorem 11) *Let G be the disjoint union of component subgraphs $G_1, \dots, G_N$. For any nonempty $\sigma \subseteq V_G$,*

$$\sigma \in \text{FP}(G) \iff \sigma \cap V_{G_i} \in \text{FP}(G_i) \cup \{\emptyset\} \text{ for all } i \in [N].$$

We can now prove Theorem 4.8 (reprinted below).

Theorem 4.8 (DAG nerve) Suppose that $\mathcal{N}$ is a DAG, and let $\mathcal{T} = \{\text{sinks of } \mathcal{N}\}$ and $\mathcal{W} = V_{\mathcal{N}} \setminus \mathcal{T}$. Then G is directional with direction $\omega \rightarrow \tau$ for $\omega = \pi^{-1}(\mathcal{W})$ and $\tau = \pi^{-1}(\mathcal{T})$.

Moreover,

1. $\sigma \in \text{FP}(G) \Rightarrow \pi(\sigma) \in \text{FP}(\mathcal{N}) = \mathcal{P}(\mathcal{T}) \setminus \{\emptyset\}$, where $\mathcal{P}(\mathcal{T})$ denotes the power set of $\mathcal{T}$.
2. $\sigma \in \text{FP}(G) \Rightarrow \sigma \cap v_i \in \text{FP}(G|_{v_i}) \cup \{\emptyset\}$ for all $i \in \mathcal{T}$ and $\sigma \cap v_j = \emptyset$ for all $j \in \mathcal{W}$.

Proof The fact that G is directional with direction $\omega \rightarrow \tau$ for $\omega = \pi^{-1}(\mathcal{W})$ and $\tau = \pi^{-1}(\mathcal{T})$ follows from Theorem 4.7 since the given choice of $(\mathcal{W}, \mathcal{T})$ is the maximal DAG decomposition of $\mathcal{N}$. As a consequence of this, we have $\text{FP}(G) \subseteq \text{FP}(G|_{\pi^{-1}(\mathcal{T})})$, and so we turn our attention to $G|_{\pi^{-1}(\mathcal{T})}$ to understand $\text{FP}(G)$.

Observe that since $\mathcal{T} = \{\text{sinks of } \mathcal{N}\}$, there are no edges between the vertices in $\mathcal{T}$, and so $\mathcal{N}|_{\mathcal{T}}$ is an independent set. Thus, there are no edges in G between the components v_i for $i \in \mathcal{T}$, and so $G|_{\pi^{-1}(\mathcal{T})}$ is a disjoint union of the component subgraphs $G|_{v_i}$. Applying Theorem 4.14, we see that $\sigma \in \text{FP}(G|_{\pi^{-1}(\mathcal{T})})$ precisely when $\sigma \cap v_i \in \text{FP}(G|_{v_i}) \cup \{\emptyset\}$ for all $i \in \mathcal{T}$. And since $\text{FP}(G) \subseteq \text{FP}(G|_{\pi^{-1}(\mathcal{T})})$ by the directionality of G , part (2) of the theorem statement follows immediately.

For part (1), observe that since G is directional, $\sigma \in \text{FP}(G)$ implies that $\pi(\sigma) \subseteq \mathcal{T}$. Since $\mathcal{N}$ is a DAG, by Lemma 2.15, $\text{FP}(\mathcal{N}) = \mathcal{P}(\mathcal{T})$, and so every subset of $\mathcal{T}$ is an element of $\text{FP}(\mathcal{N})$. Thus, $\sigma \in \text{FP}(G) \Rightarrow \pi(\sigma) \in \text{FP}(\mathcal{N})$ as desired. $\square$

5 Some Extensions and Applications

We now turn our attention to some examples that illustrate the power of our nerve theorems. Going back to Fig. 3a,b of the Introduction, we see that this graph and its nerve satisfy the hypotheses of Theorem 4.7 and Theorem 4.8. The nerve (panel b) is a simple path that has a maximal DAG decomposition with $\mathcal{T} = \{10\}$ (10 is the unique sink node). Theorem 4.7 thus predicts that $\text{FP}(G) \subseteq \text{FP}(G|_{v_{10}})$, since $v_{10} = \pi^{-1}(\{10\})$. In fact, $\text{FP}(G) = \{v_{10}\}$, so the network has a unique fixed point supported on v_{10} . Moreover, this fixed point is stable because it corresponds to a clique (see [8, 9]). As seen in panel d, the dynamics do indeed converge to this stable fixed point. Furthermore, for solutions with initial conditions supported on the first clique ($G|_{v_1}$), we see that the transient dynamics activate all cliques in the chain, in sequence, following the path of the nerve.

In the remainder of this section, we will discuss additional examples of networks whose graphs and nerves satisfy the hypotheses of one or more of our nerve theorems: Theorem 4.7, Theorem 4.8, and Theorem 4.9. Just as in Fig. 3, we will see that the nerve not only predicts the fixed points and asymptotic dynamics, but also provides insight into the transient dynamics of the network. In the second subsection, we will see that even when we violate a key condition of directional covers, the nerve of a network covered by directional graphs can still provide accurate predictions of the dynamics.

5.1 Iterating and Combining DAG Decomposition and Cycle Nerve Theorems

Recall that for any partition $\{v_1, \dots, v_n\}$ of the vertices of G , there is an associated quotient map $\pi : V_G \rightarrow V_N = [n]$ defined by $\pi(v_i) = \{i\}$ for each $i \in [n]$. We saw in Sect. 4 that if such a partition induces a directional cover, and the nerve $\mathcal{N}$ has a DAG decomposition $(\mathcal{W}, \mathcal{T})$, then the fixed points of the corresponding CTLN are confined to the non-DAG part $\mathcal{T}$. In other words, Theorem 4.7 tells us that

$$\text{FP}(G) \subseteq \text{FP}(G|_{\tau}),$$

where $\tau = \pi^{-1}(\mathcal{T})$. Now if we consider the restricted graph, $G' = G|_{\tau}$, we can potentially iterate this process by finding a new directional cover, with nerve $\mathcal{N}'$, DAG decomposition $(\mathcal{W}', \mathcal{T}')$, and quotient map π' . This would enable us to further restrict the fixed points of the original network to

$$\text{FP}(G) \subseteq \text{FP}(G') \subseteq \text{FP}(G''),$$

where $G'' = G|_{\tau'}$, and $\tau' = \pi'^{-1}(\mathcal{T}') \subset \tau$. Note that G'' is the original graph restricted to an even smaller subset of vertices. This kind of iteration may enable us to get more power from our nerve theorems, by further constraining $\text{FP}(G)$.




Fig. 15 Iterative application of nerve theorems

Figure 15 shows an example of a graph (a1) whose nerve (a2) is a DAG. Therefore, we can apply Theorem 4.8 to conclude that the fixed point supports of $\text{FP}(G)$ must be unions of fixed points of G restricted to the green, red, or orange components. Moreover, the activity flows towards the subnetwork corresponding to the sinks in the nerve. Indeed, for a given initial condition supported on the top gray nodes, we see the solution converge to a limit cycle supported only on green nodes (a3).

The green component, however, is itself a complex graph. Thus, we may consider the subgraph G' of G corresponding to the non-DAG part of $\mathcal{N}$. Figure 15 (right) shows G' (b1) with nodes colored according to a DAG decomposition of its own nerve $\mathcal{N}'$ (b2). (Note that G' is a disjoint union of three graphs, and the nerve $\mathcal{N}'$ is the disjoint union of nerves for each connected component of G' .) Applying Theorem 4.7 allows us to restrict the fixed points of G even further, to the vertices that map to the colored nodes in b2. We see this reflected in the dynamics as well. Panel b3 shows the same limit cycle as before, only now it's clear that the green curves in a3 correspond only to the yellow, purple, and blue neurons in b1.

We can also combine the DAG nerve theorems with the cycle nerve theorem, Theorem 4.9. Figure 16a depicts the graph of a complex network whose nodes are grouped according to a partition with 12 components (8 in gray, 4 in color). In panel b we see the nerve of the induced directional cover, with the color of each node matching those in the original graph. This nerve has a DAG decomposition with the eight gray nodes in the DAG part, $\mathcal{W}$, and the four colored nodes in the non-DAG part, $\mathcal{T}$. Theorem 4.7 thus tells us that all fixed points of the original graph G must be contained in $\tau = \pi^{-1}(\mathcal{T})$, which is the set of colored nodes in panel a. Indeed, $\text{FP}(G) \subseteq \text{FP}(G|_{\tau})$ for this graph.

Note that we can also apply another nerve theorem to $G|_{\tau}$ to say something stronger about $\text{FP}(G)$. Because $G|_{\tau}$ has a nerve that is a cycle, Theorem 4.9 tells us that any fixed point of G must intersect each of the components $\pi^{-1}(i)$ for i in the cycle. That is, fixed points of G must contain at least one vertex from each of the four colors (red, blue, green, orange) shown in panel a. This is in fact the case, as the CTLN for G has $\text{FP}(G) = \{14567\}$. Moreover, we see that even if we choose initial conditions supported only on the gray vertices, the dynamics will converge to




Fig. 16 Application of multiple nerve theorems. (a) A graph G for a CTLN. (b) A nerve $\mathcal{N}$ of G with 12 vertices. Each of the gray vertices in $\mathcal{N}$ corresponds to a cluster of gray nodes in (a), while the colored vertices in $\mathcal{N}$ correspond to the vertices with matching colors in G . (c) A solution to a CTLN with graph G and initial conditions supported on the top-most gray nodes. The activity flows down the network and converges to a limit cycle involving only the colored nodes

a part of the state space where the gray neurons are off and at least one neuron of each color is active (see Fig. 16c).

5.2 Extensions Beyond Directional Covers

In Definition 4.3 of (rigid) directional covers, we required fairly stringent conditions that enabled us to prove strong results about the fixed points of a graph in terms of the fixed points of its nerve. Here we consider some examples of “weak” directional covers, where the component graphs are all directional but one of the conditions in Definition 4.3 is violated. Nevertheless, we find that the nerve provides a remarkably accurate prediction of the network dynamics.

Our starting point for these examples is a pair of “nerve” graphs with a grid-like structure, shown in Fig. 17a. Each graph is a finite lattice with directed edges moving down and to the right across the grid, and an additional edge, $20 \rightarrow 16$, completing a cycle at the bottom. We can think of these graphs as a pair of nerves, $\mathcal{N}_1$ and $\mathcal{N}_2$, for larger networks obtained by inserting directional graphs along the edges. In this case, a vertex i corresponds to a component v_i , and an edge $i \rightarrow j$ corresponds to a directional graph $G|_{v_i \cup v_j}$ with direction $v_i \rightarrow v_j$. The only difference is that the first nerve, $\mathcal{N}_1$, includes the $15 \rightarrow 20$ edge; while the second nerve, $\mathcal{N}_2$, does not (see dotted line in Fig. 17a).

The nerves enable one to make concrete predictions about the network dynamics. Specifically, we can consider CTLNs where the graph G is chosen to be $\mathcal{N}_1$ or




Fig. 17 A grid-like nerve shapes global sequential dynamics irrespective of the component graphs. (a) Two versions of the nerve, with and without the $15 \rightarrow 20$ edge. (b–c) Solutions for CTLNs where the graph is nerve 1 or nerve 2, respectively. (d) 5-clique component graphs and a corresponding directional graph for an edge in the cover. (e) 5-star component graphs and their directional graph. (f–i) Solutions for associated CTLNs for all four combinations of nerves and component graphs

$\mathcal{N}_2$. Figure 17b shows a solution for a CTLN with $G = \mathcal{N}_1$, where the activity is initialized with $x_1(0) = 0.5$ and $x_i(0) = 0$ for $i > 1$ (we refer to this as initial condition 1). Note that the transient activity follows a hopscotch trajectory $1 \rightarrow (2, 6) \rightarrow 7 \rightarrow (8, 12) \rightarrow 13 \rightarrow (14, 18) \rightarrow 19$, where the pairs $(2, 6)$, $(8, 12)$, and $(14, 18)$ fire synchronously. Once the activity reaches the bottom row $\mathcal{T}$, the dynamics converge to a limit cycle that follows the cycle $19 \rightarrow 20 \rightarrow 16 \rightarrow 17 \rightarrow 18 \rightarrow 19$ in the graph. If we choose the same initial condition for $G = \mathcal{N}_2$, we obtain exactly the same result. This is because the activity never reaches node 15, and so the cut edge is not “seen.” Figure 17c shows the solution when we initialize instead at $x_3(0) = 0.5$, and $x_i(0) = 0$ for $i \neq 3$ (initial condition 2). In this case, the activity follows the hopscotch trajectory $3 \rightarrow (4, 8) \rightarrow 9 \rightarrow (10, 14) \rightarrow 15$ that ends in a stable fixed point at 15. This is precisely what we expect since node 15 is a sink. These basic features of the

dynamics for a CTLN defined directly on the nerve can be taken as a prediction for the dynamics of any network G that has a directional cover with $\mathcal{N}_1$ or $\mathcal{N}_2$ as its nerve.

If we consider a graph G for which $\mathcal{N}_1$ or $\mathcal{N}_2$ is the nerve of a directional cover, satisfying all the conditions of Definition 4.3, then we have additional knowledge about the fixed points of G . The first nerve, $\mathcal{N}_1$, has a maximal DAG decomposition $(\mathcal{W}, \mathcal{T})$ with nodes 1–15 in $\mathcal{W}$ and nodes 16–20 in $\mathcal{T}$. Theorem 4.7 thus predicts that all fixed points of G are supported in $\tau = \pi^{-1}(\mathcal{T})$. Moreover, since $\mathcal{N}_{|\mathcal{T}}$ forms a cycle, applying Theorem 4.9 we expect all fixed points to intersect each component of $\mathcal{T}$. In the second nerve, $\mathcal{N}_2$, $15 \not\rightarrow 20$ and thus node 15 is a sink. Any DAG decomposition for $\mathcal{N}_2$ must therefore include node 15 in $\mathcal{T}$. Applying Theorem 4.7, we expect fixed point(s) corresponding to the cycle, as before, as well as fixed point(s) supported in $v_{15} = \pi^{-1}(15)$. We may also have fixed points whose supports are unions of those from the cycle and v_{15} . These observations are all independent of the choice of components v_i of G and the subgraphs $G|_{v_i \cup v_j}$, so long as the nerve corresponds to a directional cover in accordance with Definition 4.3.

What happens if a cover by directional graphs violates one of the conditions of Definition 4.3? Figure 17d shows a directional graph that can be inserted along the edges of $\mathcal{N}_1$ and $\mathcal{N}_2$. The components are 5-cliques, and the directional graphs are the same as in Fig. 3c. Inserting these into the grid-like nerves, however, yields many nodes with splitting overlaps (unlike for the nerve in Fig. 3b). This means backwards edges within the directional graphs violate the “splitting” condition 2 of Definition 4.3, and the resulting network does not have a directional cover. Moreover, the splitting condition was essential to our nerve theorems: the fixed points of a network G obtained by inserting the Fig. 3c graph into either $\mathcal{N}_1$ or $\mathcal{N}_2$ do not satisfy the constraints given by Theorem 4.7. In fact, there are numerous fixed points supported outside $G|_{\tau}$ for $\tau = \pi^{-1}(\mathcal{T})$.

As another example, inserting the directional graph in Fig. 17e as $G|_{v_i \cup v_j}$ along the edges of $\mathcal{N}_1$ or $\mathcal{N}_2$ also produces a cover by directional graphs that violates the splitting condition. Here, each component is a cyclically symmetric oriented graph on five vertices, called the 5-star,⁶ and the directional graphs $G_{ij} = G|_{v_i \cup v_j}$ are chosen to have forward edges from v_i onto three of the nodes in v_j , and backwards edges from the remaining two nodes in v_j to all five nodes in v_i . Again, our nerve theorems fail to predict the fixed point structure of the larger network.

Nevertheless, we find that the dynamics of these networks whose graph covers violate the splitting condition are well predicted by their nerves. Figure 17f–i display the dynamics of the networks obtained from each of the four combinations: 5-clique components with nerve $\mathcal{N}_1$ (panel f), 5-clique components with nerve $\mathcal{N}_2$ (panel g), 5-star components with nerve $\mathcal{N}_1$ (panel h), and 5-star components with nerve $\mathcal{N}_2$ (panel i). In each case, the nerve dictates the global structure of the dynamics as activity flows from one component to another. Figure 17f shows the solution for a CTLN with 5-clique components and initial conditions supported in v_1 . As

⁶ Each vertex k in the 5-star has two outgoing edges: $k \rightarrow k+1$ and $k \rightarrow k+2$ (indexing mod 5).

predicted, the asymptotic behavior is of a limit cycle following the cycle $16 \rightarrow 17 \rightarrow 18 \rightarrow 19 \rightarrow 20 \rightarrow 16$ in the nerve. Figure 17g shows the activity with 5-clique components inserted into nerve $\mathcal{N}_2$, with initial conditions leading the activity to converge to a stable fixed point, as in Fig. 17c.

Inserting 5-star graphs in $\mathcal{N}_1$ and $\mathcal{N}_2$, instead of 5-cliques, also produces asymptotic behavior that settles into a repeating sequence (Fig. 17h) or a localized attractor (Fig. 17i). The structure of the inserted graphs, however, does affect the local dynamics within each component of the nerve. In Fig. 17f the transient dynamics are considerably more regular than in Fig. 17b, and the neurons within each component fire synchronously. In contrast, in Fig. 17h the transient dynamics are irregular like in Fig. 17b, and the neurons within each component do not fire synchronously. In each case, we see global aspects of the dynamics being dictated by the nerve, while local dynamics are affected by differences in the component graphs $G|_{v_i}$. For example, in Fig. 17g the activity converges to a stable fixed point corresponding to the 5-clique supported on v_{15} ; but in Fig. 17i, the network does not converge to a fixed point because the 5-star does not support a stable fixed point. Instead, the activity settles into a limit cycle typical of 5-star CTLNs, with activity confined almost entirely to the neurons in v_{15} .

Taken together, these examples suggest that nerves can be predictive of network dynamics for a broader class of directional covers, including networks where our current set of nerve theorems do not apply. It is an open question how to formalize these observations into new nerve theorems that reflect the predictions on the dynamics.

6 Conclusion

In this work, we investigated how the global structure of a network, as captured by the nerve of a directional cover, reflects the underlying dynamics. By replacing directional subgraphs with single edges, the nerve provides a significant dimensionality reduction of a network. Moreover, this reduced network is meaningful: in simulations, we have seen that the dynamics of a CTLN with a directional cover closely follows the dynamics of its nerve.

Although the observed relationship between the dynamics of a network and its nerve is so far heuristic, we have proven a number of theorems directly connecting the fixed points of a larger network G to those of its nerve $\mathcal{N}$. Specifically, we showed that whenever the nerve has a *DAG decomposition* $(\mathcal{W}, \mathcal{T})$, the nerve is directional with direction $\mathcal{W} \rightarrow \mathcal{T}$, and this guarantees that the larger network G is also directional. In particular, the fixed points of the larger network are confined to live in the pullback $\pi^{-1}(\mathcal{T})$ (Theorem 4.7). In the special case where the nerve is a DAG, we have a tighter connection between $\text{FP}(G)$ and $\text{FP}(\mathcal{N})$. Theorem 4.8 shows that every fixed point support of G projects to a fixed point of $\mathcal{N}$. In other words,

$$\sigma \in \text{FP}(G) \Rightarrow \pi(\sigma) \in \text{FP}(\mathcal{N}).$$

Moreover, every $\sigma \in \text{FP}(G)$ is a union of fixed point supports of the component subgraphs $G|_{v_i}$. Theorem 4.9 gives similar constraints on $\text{FP}(G)$ whenever the nerve is a cycle.

Due to the close relationship between fixed points and attractors in CTLNs, understanding the fixed point structure and how this is shaped by network architecture is an important step towards understanding how network connectivity shapes dynamics. The machinery developed here thus provides a useful framework for dimensionality reduction in the analysis of large networks. Moreover, it provides insight into how to engineer complex networks with desired dynamic properties from smaller building block components.

Acknowledgments This research is a product of one of the working groups at the Workshop for Women in Computational Topology (WinCompTop) in Canberra, Australia (1–5 July 2019). We thank the organizers of this workshop and the funding from NSF award CCF-1841455, the Mathematical Sciences Institute at ANU, the Australian Mathematical Sciences Institute (AMSI), and Association for Women in Mathematics that supported participants’ travel. We thank Caitlyn Parmelee for fruitful discussions that helped set the foundation for this work. We would also like to thank Joan Licata for valuable conversations at the WinCompTop workshop. CC and KM acknowledge funding from NIH R01 EB022862, NIH R01 NS120581, NSF DMS-1951165, and NSF DMS-1951599.

References

1. Abeles, M.: Local Cortical Circuits: An Electrophysiological Study. Springer, Berlin (1982)
2. Aviel, Y., Pavlov, E., Abeles, M., Horn, D.: Synfire chain in a balanced network. *Neurocomputing* **44**, 285–292 (2002)
3. Bel, A., Cobiaga, R., Reartes, W., Rotstein, H.G.: Periodic solutions in threshold-linear networks and their entrainment. *SIAM J. Appl. Dyn. Syst.* **20**(3), 1177–1208 (2021)
4. Biswas, T., Fitzgerald, J.E.: A geometric framework to predict structure from function in neural networks (2020). Available at <https://arxiv.org/abs/2010.09660>
5. Borsuk, K.: On the imbedding of systems of compacta in simplicial complexes. *Fundam. Math.* **35**(1), 217–234 (1948)
6. Curto, C., Morrison, K.: Pattern completion in symmetric threshold-linear networks. *Neural Comput.* **28**, 2825–2852 (2016)
7. Curto, C., Degeratu, A., Itskov, V.: Encoding binary neural codes in networks of threshold-linear neurons. *Neural Comput.* **25**, 2858–2903 (2013)
8. Curto, C., Geneson, J., Morrison, K.: Fixed points of competitive threshold-linear networks. *Neural Comput.* **31**(1), 94–155 (2019)
9. Curto, C., Geneson, J., Morrison, K.: Stable fixed points of combinatorial threshold-linear networks (2019). Available at <https://arxiv.org/abs/1909.02947>
10. Hahnloser, R.H., Sarpeshkar, R., Mahowald, M.A., Douglas, R.J., Seung, H.S.: Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit. *Nature* **405**, 947–951 (2000)
11. Hahnloser, R.H., Seung, H.S., Slotine, J.J.: Permitted and forbidden sets in symmetric threshold-linear networks. *Neural Comput.* **15**(3), 621–638 (2003)
12. Hayon, G., Abeles, M., Lehmann, D.: A model for representing the dynamics of a system of synfire chains. *J. Comput. Neurosci.* **18**(41–53) (2005)

13. Leray, J.: Sur la forme des espaces topologiques et sur les points fixes des représentations. *J. Math. Pures Appl.* **9**, 95–248 (1945)
14. Morrison, K., Curto, C.: Predicting neural network dynamics via graphical analysis. In: Robeva, R., Macaulay, M. (eds.) *Algebraic and Combinatorial Computational Biology*, pp. 241–277. Elsevier, Amsterdam (2018)
15. Morrison, K., Degeratu, A., Itskov, V., Curto, C.: Diversity of emergent dynamics in competitive threshold-linear networks: a preliminary report (2016). Available at <https://arxiv.org/abs/1605.04463>
16. Parmelee, C., Londono Alvarez, J., Curto, C., Morrison, K.: Sequential attractors of combinatorial threshold-linear networks. To appear in *SIAM J. Appl. Dyn. Syst.* (2022). Available at <https://arxiv.org/abs/2107.10244>
17. Parmelee, C., Moore, S., Morrison, K., Curto, C.: Core motifs predict dynamic attractors in combinatorial threshold-linear networks, *PLOS ONE*, **17**(3) (2022): e0264456 <https://doi.org/10.1371/journal.pone.0264456>
18. Seung, H.S., Yuste, R.: Principles of Neural Science. In: Appendix E: Neural Networks, pp. 1581–1600, 5th ed. McGraw-Hill Education/Medical, New York (2012)
19. Xie, X., Hahnloser, R.H., Seung, H.S.: Selectively grouping neurons in recurrent networks of lateral inhibition. *Neural Comput.* **14**, 2627–2646 (2002)

Combinatorial Conditions for Directed Collapsing



Robin Belton, Robyn Brooks, Stefania Ebli, Lisbeth Fajstrup,
Brittany Terese Fasy, Nicole Sanderson, and Elizabeth Vidaurre

Abstract While collapsibility of CW complexes dates back to the 1930s, collapsibility of *directed* Euclidean cubical complexes has not been well studied to date. The classical definition of collapsibility involves certain conditions on pairs of cells of the complex. The direction of the space can be taken into account by requiring that the past links of vertices remain homotopy equivalent after collapsing. We call this type of collapse a *link-preserving directed collapse*. In the undirected setting, pairs of cells are removed that create a deformation retract. In the directed setting, topological properties—in particular, properties of spaces of directed paths—are not always preserved. In this paper, we give computationally simple conditions for preserving the topology of past links. Furthermore, we give conditions for when link-preserving directed collapses preserve the contractability and connectedness of spaces of directed paths. Throughout, we provide illustrative examples.

R. Belton · B. T. Fasy (✉)
Montana State University, Bozeman, MT, USA
e-mail: robin.belton@montana.edu; brittany.fasy@montana.edu

R. Brooks
Boston College, Chestnut Hill, MA, USA
e-mail: robyn.brooks@bc.edu

S. Ebli
École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland
e-mail: stefania.ebli@epfl.ch

L. Fajstrup
Aalborg University, Aalborg, Denmark
e-mail: fajstrup@math.aau.dk

N. Sanderson
Lawrence Berkeley National Lab, Berkeley, CA, USA
e-mail: nzs5677@psu.edu

E. Vidaurre
Molloy College, Rockville Centre, NY, USA
e-mail: evidaurre@molloy.edu