ORIGINAL ARTICLE

Excitable networks for finite state computation with continuous time recurrent neural networks

Peter Ashwin[1] · Claire Postlethwaite[2]

Abstract
Continuous time recurrent neural networks (CTRNN) are systems of coupled ordinary differential equations that are simple enough to be insightful for describing learning and computation, from both biological and machine learning viewpoints. We describe a direct constructive method of realising finite state input-dependent computations on an arbitrary directed graph. The constructed system has an excitable network attractor whose dynamics we illustrate with a number of examples. The resulting CTRNN has intermittent dynamics: trajectories spend long periods of time close to steady-state, with rapid transitions between states. Depending on parameters, transitions between states can either be *excitable* (inputs or noise needs to exceed a threshold to induce the transition), or *spontaneous* (transitions occur without input or noise). In the excitable case, we show the threshold for excitability can be made arbitrarily sensitive.

Keywords Continuous time recurrent neural network · Nonlinear dynamics · Excitable network attractor

1 Introduction

It is natural to try to understand computational properties of neural systems through the paradigm of network dynamical systems, where a number of dynamically simple units (i.e. with attracting equilibria or periodic orbits) interact to give computation as an emergent property of the system. This is as much the case for biological models of information processing, pattern generation and decision making as it is for artificial neural networks inspired by these models. A variety of specific models have been developed to describe the dynamics and training of recurrent networks comprised of coupled neurons in biological and artificial settings. The particular challenge that we address here is the construction of arbitrarily complex, but specified, dynamical structures that enable discrete (finite-state) computation in an input-driven

✉ Peter Ashwin
  p.ashwin@exeter.ac.uk

✉ Claire Postlethwaite
  c.postlethwaite@auckland.ac.nz

1 Center for Systems, Dynamics and Control, Department of
   Mathematics, University of Exeter, Exeter EX4 4QF, UK

2 Department of Mathematics, University of Auckland,
   Auckland 1142, New Zealand

system that is nonetheless continuous in both state and time. Clearly, invariant objects of the autonomous system (such as equilibria, periodic orbits and chaotic attractors) only form part of the picture and an input-dependent (non-autonomous) approach such as Manjunath et al. (2012) is needed.

To help understand the response of systems to inputs the authors introduced in Ashwin and Postlethwaite (2016) a notion of a "network attractor", namely an invariant object in phase space that may contain several local invariant sets, but also systems of interconnections between them. This generalises the notions of "heteroclinic network" and "winnerless competition/stable heteroclinic channels" (Afraimovich et al. (2004)) which have been used to describe a range of sequence generation, computational and cognitive effects in neural systems: see for example Rabinovich et al. 2001, 2006; Afraimovich et al. 2004; Rabinovich et al. 2020; Hutt and beim Graben 2017. These models connect computational states, represented as saddles in the dynamics. In the presence of inputs or noise, the switching between saddles is a useful model for spontaneous computational processes, but there are two problems. One of these is that the states are dynamically unstable and so there are spontaneous transitions that are not noise driven. The other is that heteroclinic chains are destroyed by arbitrarily small perturbations, unless there are special structures in phase space (symmetries or invariant subspaces).

Network attractors joining stable equilibria (Ashwin and Postlethwaite 2016) overcome both of these issues. In this paper, we show that arbitrarily complex network attractors exist in an idealised class of model called a *continuous time recurrent neural network* (CTRNN) (Beer 1995). A CTRNN is a set of differential equations each with one scalar variable that represents the level of activation of a neuron, and feedback via a saturating nonlinearity or "activation function". These models have been extensively investigated in the past decades as simple neurally inspired systems that can nonetheless (without input) have complex dynamics by virtue of the nonlinearities present (Beer 1995). They can also be trained to perform arbitrarily complex tasks depending on time-varying input (Tuci et al. 2002; Yamauchi and Beer 1994). They are frequently used in investigations of evolutionary robotics (Blynel and Floreano 2003) and (in various equivalent formulations (Chow and Karimipanah 2020)) as models for neural behaviour in biological or cognitive systems. For example, the classical work of Hopfield and Tank (Hopfield and Tank 1985) considers such systems with symmetric weights to solve optimization problems, while more recently, Bhowmik et al. (2016) discusses CTRNN models for episodic memory, and several other biological and cognitive applications are discussed in Nikiforou (2019).

CTRNNs are often referred to as "universal dynamical approximators" (Funahashi and Nakamura 1993), meaning that the trajectory of a CTRNN can approximate, to arbitrary precision, any other prescribed (smooth) trajectory in $\mathbb{R}^n$. However, this does not mean that the dynamics of CTRNNs are simple to understand, or that it is easy to form the above approximation. It also raises the question of how a "trained" CTRNN performs a complex task. Gradient descent or more general evolutionary training algorithms train the network by navigating through a high dimensional landscape of possible feedback weightings and moving these towards a setting that is sufficiently optimal for the task. It may be possible to give a clear description of the resulting nonlinear dynamics of the autonomous (constant input) CTRNN, but we want to understand not only this but also how inputs affect the state of the system.

The main theoretical result in this paper focusses on "excitable network attractors": these consist of a finite set of local attracting equilibria and excitable connections between them (see Appendix A). It was demonstrated in Ashwin and Postlethwaite (2018) that an excitable network attractor can be used to embed an arbitrary Turing machine within a class of purpose-designed coupled dynamical system with two different cell types. Rather than relying on an optimization approach to design the system, that paper gave a *constructive* method for designing a realisation of any desired network attractor. However, this construction required specialist dynamical cells with quite complex nonlinear couplings between them, and a comparatively large

number of cells. It was recently shown in Ceni et al. (2020) that trained RNNs in the form of echo-state networks can realise excitable networks for certain tasks and that structural errors in the trained network can explain errors in imperfectly trained systems.

The current paper demonstrates that CTRNN dynamics is sufficiently rich to realise excitable networks with arbitrary graph topology, simply by specifying appropriate connection weights. The construction algorithm in the proof of Theorem 1 assigns one of only four values to each of the connection weights to realise an arbitrary graph on N vertices as an excitable network on N states (subject to some minor constraints on its connectivity), using a CTRNN with N cells. The CTRNN we consider in this paper (see for example (Beer 1995), which corresponds to a continuous time Hopfield model (Hopfield and Tank 1985) in the symmetric coupling case $w_{ij} = w_{ji}$) are ordinary differential equations

$$\dot{y}_i = -y_i + \sum_j w_{ij}\phi(y_j) + I_i(t), \tag{1}$$

where $\mathbf{y} = (y_1, \ldots y_N) \in \mathbb{R}^N$ is the internal state of the N cells of the system, w_{ij} is a matrix of connection weights, ϕ is a (sigmoid) activation function that introduces a saturating nonlinearity into the system and $I_i(t)$ is an input. We say system (1) is *input-free* if $I_i(t) = 0$ for all i and t.

We consider two cases for ϕ, a smooth function

$$\phi(y) = \phi_S(y) := \left[1 + \exp\left(-\frac{(y - \theta)}{\epsilon}\right)\right]^{-1}, \tag{2}$$

and a piecewise affine function

$$\phi(y) = \phi_P(y) := \begin{cases} 0 & y - \theta < -2\epsilon, \\ (y - \theta)/(4\epsilon) + 1/2 & |y - \theta| \leq 2\epsilon, \\ 1 & y - \theta > 2\epsilon, \end{cases} \tag{3}$$

In both cases, ϕ is monotonic increasing with

$$\lim_{y \to \infty} \phi(y) = 1, \quad \lim_{y \to -\infty} \phi(y) = 0, \quad \phi(\theta) = 1/2,$$

and a maximum derivative at $y = \theta$, equal to $\frac{1}{4\epsilon}$. In both cases, ϵ and θ are parameters, and we are interested in the case $0 < \epsilon \ll 1$. In general, the function ϕ need not be the same in every component of (1), but here we make a simplifying assumption that it is.

Note that both activation functions (2) and (3) have piecewise constant limits in the singular limit $\epsilon \to 0$. Such limiting systems are of Fillipov type and have been explored in various biological contexts, especially for gene regulatory dynamics. These can also have rich dynamics as discussed in

the literature (for example Gouzé and Sari 2003; Harris and Ermentrout 2015), but we do not consider this limit here.

The main contribution of Sect. 2 is to give, in Theorem 1, a construction of a connection weight matrix w_{ij} such that the dynamics of the input-free system (1) contains (or *realises*: definition given below) an excitable network attractor, as defined in Ashwin and Postlethwaite (2016). We prove this (with details in Appendix B) for the case of the piecewise affine function ϕ_P. In Sect. 3, we present evidence that this is also true for the smooth case ϕ_S for an open set of parameters. Qualitatively, this means the system will contain a number of stable equilibrium states, and small inputs (either deterministic, or noisy) will push the trajectory from one stable equilibrium into the basin of attraction of another. In this way, transitions can be made around the network, and the transition time between states tends to be much smaller than the residence times of the trajectory in neighbourhoods of the states. In particular, we can choose w_{ij} so that the network attractor has (almost) any desired topology. In Appendix A, we recall formal definitions of network attractors from (Ashwin and Postlethwaite 2016, 2018).

In Sect. 3, we consider several examples of simple graphs and demonstrate that the desired networks do indeed exist in the systems as designed. We also perform numerical bifurcation analysis to demonstrate the connection between periodic orbits in the input-free deterministic system (1) and excitable networks in the same system with additive noise, that is, the system of stochastic differential equations (SDEs):

$$dy_i = \left(-y_i + \sum_j w_{ij}\phi(y_j) \right) dt + \sigma\, dW_i(t), \qquad (4)$$

where $W_i(t)$ are independent standard Wiener processes. Here, the noise plays the role of inputs that propel the trajectory around the network, although of course this occurs in a random manner. In Sects. 3.3 and 3.4, we consider graphs that have multiple edges leading out from a single vertex and show that additional equilibria may appear in the network attractor where two or more cells are active simultaneously. We further show that the existence of these additional equilibria can be suppressed by choosing one of the parameters used in the construction of the weight matrix w_{ij} to be sufficiently large.

Section 4 concludes by relating our results to other notions of sequential computation. We also conjecture some extensions of the results shown in this paper.

2 Construction of a CTRNN with a network attractor

Let G be an arbitrary directed graph between N vertices, and let a_{ij} be the adjacency matrix of G. That is, $a_{ij} = 1$ if there
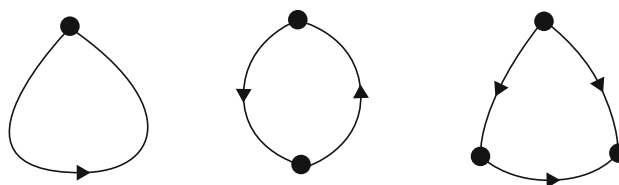

Fig. 1 From left to right, the figures show an order-one loop, an order-two loop, and a Δ-clique in a directed graph. The construction we provide realises an arbitrary graph G, as long of none of these subgraphs are present in G

is a directed edge from vertex i to vertex j, and $a_{ij} = 0$ otherwise.

Let Σ be an invariant set for a system of ordinary differential equations. We say Σ is an excitable network that *realises* a graph G for some amplitude $\delta > 0$ if for each vertex v_i in G there is a unique stable equilibrium ξ_i in Σ, and if there is an excitable connection with amplitude δ in Σ from ξ_i to ξ_j whenever there is an edge in G from v_i to v_j. The existence of an excitable connection means that there exists a trajectory with initial condition within a distance δ of ξ_i that asymptotes in forward time to ξ_j (formal definitions are given in Appendix A).

For the purposes of our construction of a network attractor, we assume that G contains no loops of order one, no loops of order two, and no Δ-cliques. Figure 1 shows each of these graph components schematically.

In terms of the adjacency matrix, for G to contain no loops of order one requires that

$$a_{ii} = 0 \text{ for all } i; \qquad (5)$$

for G to contain no loops of order two requires that

$$a_{ij}a_{ji} = 0 \text{ for all } i, j; \qquad (6)$$

and for G to contain no Δ-cliques requires that

$$a_{ij}a_{ik}a_{jk} = 0 \text{ for all } i, j, k. \qquad (7)$$

In our earlier work, we have demonstrated a network design which can admit order-two loops and Δ-cliques (Ashwin and Postlethwaite 2016), although this previous construction is not motivated by neural networks *per se* and requires a higher dimensional system of ODEs (for a given graph) than the one presented here.

Before we move into the details of the construction, we briefly discuss our terminology. A graph G has *vertices*, which correspond to (stable) *equilibria* in the phase space of the dynamical system (1). Also within the phase space, there exist *excitable connections* (sometimes abbreviated to *connections*) between the equilibria, which correspond to the *edges* of the graph. When a trajectory in the phase space

moves between neighbourhoods of equilibria along a path close to one of these connections, we say that a *transition* between the equilibria has occurred. We refer to each of the components of the dynamical system (1) as a *cell*, and say that a cell j is *active* if $\phi(y_j)$ is close to one.

2.1 Realization of arbitrary directed graphs as network attractors

We construct a weight matrix w_{ij} that depends only on the adjacency matrix a_{ij}, and on four parameters w_t, w_s, w_m and w_p. It is given by

$$w_{ij} = w_t + (w_s - w_t)\delta_{ij} + (w_p - w_t)a_{ji} + (w_m - w_t)a_{ij}, \tag{8}$$

where δ_{ij} is the Kronecker δ. This choice of w_{ij} ensures that $w_{ii} = w_s$, $w_{ij} = w_p$ if there is a directed edge in G from vertex j to i (i.e. $a_{ji} = 1$), $w_{ij} = w_m$ if there is a directed edge in G from vertex i to vertex j (i.e. $a_{ij} = 1$), and $w_{ij} = w_t$ otherwise. In later sections, we allow for different weights along different edges by allowing w_p to depend on i and j (i.e. $w_p = w_p^{ij}$). We give an overview of how each of the parameters affect the dynamics of the system in Sect. 2.2.

We write $\mathbf{w} = (\epsilon, \theta, w_s, w_m, w_t, w_p) \in \mathbb{R}^6$ to be a vector of all parameter values. The next result shows that for the piecewise affine activation function and suitable choice of parameters, there is an embedding of G as an excitable network attractor for the input-free system.

Theorem 1 *For any directed graph G with N vertices containing no loops of order one, no loops of order two, and no Δ-cliques, and any small enough $\delta > 0$, there is an open set $W_{\mathrm{ex}} \subset \mathbb{R}^6$ such that if the parameters $\mathbf{w} \in W_{\mathrm{ex}}$, then the dynamics of input-free equation (1) on N cells with piecewise affine activation (3) and w_{ij} defined by (8) contains an excitable network attractor with threshold δ, that realises the graph G.*

Recall that by *realises* we mean that all edges in the graph are present as transitions between stable equilibria using perturbations of size at most δ.

Proof We give the main ideas behind the proof here, deferring some of the details to Appendix B. We construct an excitable network attractor in $\mathbb{R}^N$ for (1) with piecewise activation function (3) and weight matrix (8). For any $\frac{1}{2} > \delta > 0$, we show there exist parameters $\mathbf{w}$ (with $\epsilon > 0$ small) and stable equilibria ξ_k ($k = 1, \ldots, N$) that are connected according to the adjacency matrix a_{ij} by excitable connections with amplitude δ. Below, we provide an explicit set of parameters that make such a realisation, and note that the realisation will hold for an open set of nearby parameters.

We show in Appendix B that the equilibria ξ_k have components (cells) that are close to one of four values Y_T, Y_D, Y_L, Y_A related to the edges attached to the corresponding vertex k in the graph G. For any $0 < \delta < \frac{1}{2}$, we use the following parameters:

$$\epsilon = \frac{\delta}{8}, \ \theta = \frac{1}{2}, \ w_s = 1, \ w_t = 0, \tag{9}$$

and then w_p and w_m are given by

$$w_p = \theta - \frac{\delta}{2}, \quad w_m = -(w_s - \theta) - \frac{\delta}{2}. \tag{10}$$

We then set

$$Y_A := w_s, \ Y_L := w_p = \theta - \frac{\delta}{2},$$
$$Y_T := w_m = -(w_s - \theta) - \frac{\delta}{2}, \ Y_D := w_t \tag{11}$$

We use square brackets and subscripts to identify the components of points in phase space, that is, $[\xi_k]_j$ is the jth component of ξ_k. Each ξ_k has:

– Exactly one cell that is *Active*: $[\xi_k]_k = Y_A$
– A number of cells that are *Leading*: $[\xi_k]_j = Y_L$ (if $a_{kj} = 1$)
– A number of cells that are *Trailing*: $[\xi_k]_j = Y_T$ (if $a_{jk} = 1$)
– All remaining cells are *Disconnected*: $[\xi_k]_j = Y_D$ ($a_{kj} = a_{jk} = 0$).

Note that the requirement of no loops of order one or two and no Δ-cliques implies that this labelling is well defined.

From equilibrium ξ_k, there is an excitable connection to any of the equilibria ξ_l with $a_{kl} = 1$, that is, any of the Leading cells can become the Active cell. During a transition, the remaining cells can be classified into six types, which are identified in Fig. 2, and depend (for each j) on the values of the four entries in the adjacency matrix a_{jk}, a_{kj}, a_{jl} and a_{lj}. We label cell k as AT (Active–Trailing) and cell l as LA (Leading–Active). Note that the lack of two cycles means that the cases with $a_{jk} = a_{kj} = 1$ or $a_{jl} = a_{lj} = 1$ (a total of seven possibilities) cannot occur, and the lack of Δ-cliques mean that the cases with $a_{jk} = a_{jl} = 1$, $a_{kj} = a_{jk} = 1$, or $a_{kj} = a_{lj} = 1$ also cannot occur (which includes the cases where a cell would switch from Leading to Trailing). The remaining six possibilities are listed below.

– Type DD: $a_{jk} = a_{kj} = a_{jl} = a_{lj} = 0$; the cell is Disconnected throughout.
– Type TD: $a_{jk} = 1, a_{kj} = a_{jl} = a_{lj} = 0$; the cell switches from Trailing to Disconnected.
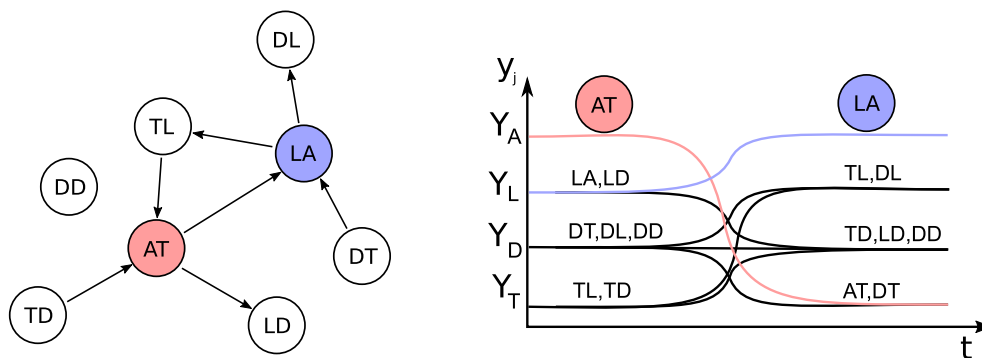
Fig. 2 Schematic diagram showing a transition in the network attractor of Theorem 1. Left: a schematic of the directed graph showing the eight distinct types of vertex. Right: schematic time series; as the Active cell changes there will be a transition in the connected Leading and Trailing cells as shown. Which cell becomes active is controlled by a small perturbation to the set of currently Leading cells

- Type LD: $a_{kj} = 1, a_{jk} = a_{jl} = a_{lj} = 0$; the cell switches from Leading to Disconnected.
- Type TL: $a_{jk} = a_{lj} = 1, a_{kj} = a_{jl} = 0$; the cell switches from Trailing to Leading.
- Type DT: $a_{jl} = 1, a_{jk} = a_{kj} = a_{lj} = 0$; the cell switches from Disconnected to Trailing.
- Type DL: $a_{lj} = 1, a_{jk} = a_{kj} = a_{jl} = 0$; the cell switches from Disconnected to Leading.

The right panel in Fig. 2 shows how a transition from cell AT active to cell LA active will occur in a general network.

To prove the existence of an excitable connection giving a realisation, we consider a perturbation from ξ_k to the point

$$\zeta_{k,l} = \xi_k + \delta e_l,$$

where e_l is the unit basis vector, and we show in Appendix B that, for small enough δ, $\zeta_{k,l}$ is in the basin of attraction of ξ_l if $a_{kl} = 1$. This means there is an excitable connection from ξ_k to ξ_l in this case. □

We believe that for small enough δ and suitable choice of weights, the realisation of G in Theorem 1 can be made *almost complete* in the sense analogous to the similar notion for heteroclinic networks (Ashwin et al. 2020), namely that the set

$$\bigcup_{\xi \in E} B_\delta(\xi) \setminus \Sigma_E$$

has zero measure. If a network realization is almost complete, then almost all trajectories starting close to some ξ_k will remain close, or will follow a connection corresponding to the realization. We do not have a proof of this, though Appendix B.2 shows that for small enough δ, $\zeta_{k,l}$ is in the basin of attraction of ξ_l if and only if $a_{kl} = 1$. This does not preclude the possibility that there exist perturbations in $B_\delta(\xi_i)$ other than $\zeta_{k,l}$ resulting in trajectories asymptotic to

ξ_l, or indeed to other attractors. Our numerical investigations suggest that by choosing w_t nonzero and large enough, any connections to other equilibria can be suppressed. This is discussed more in Sects. 3.4 and 4.

2.2 Excitable networks for smooth nonlinearities

For small enough ϵ, the smooth activation function (2) can be made arbitrarily close to the piecewise activation function (3), and so we expect Theorem 1 to also apply in the smooth case, but do not give a proof here. However, throughout Sect. 3 we use the smooth activation function (2) in our examples. We use

$$\epsilon = 0.05, \ \theta = 0.5, \ w_s = 1,$$
$$w_m = -0.7, w_p = 0.3, \ w_t = 0 \tag{12}$$

as our default parameter set (compare this choice of parameters with that given in equations (9) and (10), with $\delta = 0.4$), though there will be an open set of nearby parameters with analogous behaviour. In Sect. 3, we provide several examples of using this choice of weight matrix to realise a graph G.

From equation (11), we can see that the parameter choices directly affect the location of the equilibria ξ_k in phase space. As we will see in the following sections, the parameters also have further effects on the dynamics. In particular, the relative sizes of the parameters w_p and θ determine whether the dynamics are excitable or spontaneous: essentially, for ϵ small enough, w_p needs to be smaller than θ to observe excitable dynamics. If w_p is too large, then the equilibria ξ_k cease to exist: periodic orbits can exist instead. The parameter w_m controls how fast a Trailing cell decays, and the parameter w_t controls the suppression effects when there is more than one Leading cell. We discuss the effects of w_t in more detail in Sect. 3.3.

For a smooth activation function such as (2) that is invertible on its range, there is a useful change of coordinates to

$J_i = \phi(y_i)$ (a similar transformation is made in Beer (1995)). The input-free equations (1) then become

$$\dot{J}_i = \frac{1}{\epsilon} J_i(1 - J_i) \left(-\phi^{-1}(J_i) + \sum_j w_{ij} J_j \right), \qquad (13)$$

where each $J_i \in (0, 1)$ (which is the domain of the function ϕ^{-1}), and

$$\phi^{-1}(x) = \theta - \epsilon \ln \left(\frac{1-x}{x} \right).$$

Each vertex of the graph G is realised in the phase space of the J_k variables by the stable equilibrium with one of the J_k close to 1 and the remainder close to 0. With a slight abuse of notation, we refer to these equilibria as ξ_k. As we will see in the examples that follow, the parameters can be chosen such that the dynamics are close to a saddle-node bifurcation. In general, the system is near a degenerate bifurcation with codimension $d = \sum_{k=1}^{N} O_k$, where O_k is the out-degree of the kth vertex for the graph G. This corresponds to there being simultaneous saddle-node bifurcations (with O_k-fold degeneracy) at each equilibrium ξ_k corresponding to each of the outgoing directions. This bifurcation has global connections analogous to a saddle-node on an invariant circle (SNIC)/saddle-node homoclinic bifurcation: there are connecting orbits between the saddle nodes that reflect the network structure.

If parameters are chosen as in Theorem 1 (such that none of the saddle-node bifurcations have occurred), then for each k there will be a stable equilibrium ξ_k and a group of nearby saddles and sources. A small perturbation near ξ_k can move the trajectory out of the basin of attraction of ξ_k to effect a transition to another equilibrium: this gives an *excitable* connection from ξ_k.

If parameters are chosen such that all saddle-node bifurcations have been passed (for example, by sufficiently increasing the value of w_p), then the flow through the corresponding region of phase space near where the sink ξ_k was will be slow, and all equilibria in this region will have been destroyed. However, we will still observe intermittent dynamics as the trajectory passes through this region (Strogatz 1994, p99). In this case, we refer to a region where (a) there is a unique local minimum of $|\dot{\mathbf{y}}|$ and (b) a large subset of initial conditions in this region pass close to this minimum, as a *bottleneck region* P_k, and refer to a *spontaneous* transition past P_k. (This is also called the *ghost of a saddle-node bifurcation* in Strogatz (1994).) If all the connections corresponding to edges in the graph G are excitable, then the system contains an excitable network attractor. If all connections are spontaneous, then we typically see a periodic orbit, although we do not prove this.

For other cases (e.g. where the parameter w_p is assumed to depend on i, j in (8)), some saddle-node bifurcations will have occurred, but others will have not. In this case, we expect that some connections will be excitable, and for others, trajectories will automatically pass through bottleneck regions.

When there is a choice of two out-going connections, one of which is excitable and the other of which is spontaneous, the one chosen by the trajectory will depend on noise amplitudes and other effects: we expect there to be a rich complicated local and indeed global dynamical behaviour, the analysis of which is beyond the scope of this paper.

3 Examples for the smooth activation function

3.1 Two vertex graph

For our first example, we consider the connected graph with two vertices and a single edge joining them, that is $a_{12} = 1$, and $a_{ij} = 0$ for $(i, j) \neq (1, 2)$. We use bifurcation analysis to show that the transition between spontaneous and excitable dynamics is caused by a saddle-node bifurcation and find an approximation to the location of the saddle-node bifurcation in parameter space.

The two-dimensional system of equations is:

$$\dot{y}_1 = -y_1 + w_s \phi(y_1) + w_m \phi(y_2), \qquad (14)$$
$$\dot{y}_2 = -y_2 + w_s \phi(y_2) + w_p \phi(y_1). \qquad (15)$$

In the J_i variables, this becomes

$$\dot{J}_1 = \frac{1}{\epsilon} J_1(1 - J_1) \left(-\phi^{-1}(J_1) + w_s J_1 + w_m J_2 \right),$$
$$\dot{J}_2 = \frac{1}{\epsilon} J_2(1 - J_2) \left(-\phi^{-1}(J_2) + w_s J_2 + w_p J_1 \right), \qquad (16)$$

where $(J_1, J_2) \in (0, 1)^2$. We note the following properties of the function $g : (0, 1) \to \mathbb{R}$, with $g(x) = \phi^{-1}(x) - w_s x$:

$$g'(x) = \frac{\epsilon}{x(1-x)} - w_s, \quad g''(x) = \frac{\epsilon(2x - 1)}{x^2(1-x)^2},$$
$$\lim_{x \to 0} g(x) = -\infty, \quad \lim_{x \to 1} g(x) = \infty,$$
$$g\left(\frac{1}{2}\right) = \theta - \frac{w_s}{2}, \quad g'\left(\frac{1}{2}\right) = 4\epsilon - w_s.$$

If $w_s > 4\epsilon$, then g has local extrema at x_+ and x_-, where

$$x_\pm = \frac{1}{2} \pm \sqrt{\frac{1}{4} - \frac{\epsilon}{w_s}}. \qquad (17)$$

The J_1 and J_2 nullclines of system (16) are at, respectively

$$J_2 = \frac{1}{w_m} g(J_1) \quad \text{and} \quad J_1 = \frac{1}{w_p} g(J_2). \tag{18}$$

In Fig. 3, we show a sketch of the phase space of (16); the J_1 nullcline is shown in blue, and the J_2 nullcline in red. Solid dots show stable equilibria, open dots show unstable equilibria, and arrows show the direction of flow. Note that we do not include any nullclines at $J_i = 0$ or $J_i = 1$ because they are not in the domain of equations (16). As w_p is decreased (in the figures, moving from left to right), a saddle-node bifurcation creates a pair of equilibrium solutions.

Lemma 1 *If $\theta < w_s$, and $0 < \epsilon \ll 1$, then a saddle node bifurcation occurs in system (16) when*

$$w_p = w_p^{SN} \equiv \epsilon \log \epsilon + \theta - \epsilon(1 + \log w_s) + \frac{\epsilon^2}{w_s} + O(\epsilon^3). \tag{19}$$

To begin the proof, we note that the saddle-node bifurcation will occur when the points A and B (marked by squares in the left-hand panel of Fig. 3) coincide. These points are defined as the intersection of the nullclines with the line at $J_2 = x_- = \frac{\epsilon}{w_s} + O(\epsilon^2)$, i.e. at the local extrema of the J_2 nullcline.

Let the J_1 coordinate of A be

$$J_1^A = \frac{1}{w_p} g(x_-)$$
$$= \frac{1}{w_p}\left(\epsilon \log \epsilon + \theta - \epsilon(1 + \log w_s) + \frac{\epsilon^2}{w_s}\right) + O(\epsilon^3)$$

Let the J_1 coordinate of B be J_1^B, and write $J_1^B = 1 - \epsilon^B$, for some $\epsilon^B \ll 1$. Substituting this, along with $J_2 = x_-$, into the expression for the J_1 nullcline in (18) gives

$$x_- = \frac{1}{w_m} g(1 - \epsilon^B)$$

Expanding in terms of the small quantities ϵ and ϵ^B, this gives

$$\frac{\epsilon}{w_s} + O(\epsilon^2) =$$
$$\frac{1}{w_m}\left(\theta - \epsilon(\log \epsilon^B + \epsilon^B + O(\epsilon^{B2})) - w_s + \epsilon^B w_s\right)$$

which we rearrange to find

$$\log \epsilon^B = \frac{\theta - w_s}{\epsilon} + w_s \frac{\epsilon^B}{\epsilon} + O(1).$$

Since we are assuming $\theta < w_s$, then ϵ^B is exponentially small, that is $\epsilon^B = O(\epsilon^n)$ for all $n \in \mathbb{N}$, thus, $J_1^B = 1 + O(\epsilon^n)$.

The points A and B collide when $J_1^A = J_1^B$, that is, when

$$w_p = w_p^{SN} \equiv \theta + \epsilon \log \epsilon - \epsilon(1 + \log w_s) + \frac{\epsilon^2}{w_s} + O(\epsilon^3).$$

□

For the default parameters (12), except for w_p, we find $w_p^{SN} = 0.3027$ (4 s.f.). Note that this means for $w_p = 0.3$ we are close to saddle node and there is an excitable connection with small $\delta > 0$. More generally, note that for any fixed w_s, as $\epsilon \to 0$ we have $w_p^{SN} \to \theta$ as expected from Theorem 1.

The following result gives an approximation of the positions of the equilibria that are created in the saddle-node bifurcation. Methods similar to those used in this proof are used in later sections for larger networks.

Lemma 2 *If $\theta < w_s$, $0 < \epsilon \ll 1$, and $0 < \eta \ll \frac{\epsilon}{4}$, then if $w_p = w_p^{SN} - \eta$, the system (16) has a pair of equilibria at*

$$(J_1, J_2) = \left(1, \frac{\epsilon}{w_s} \pm \frac{\sqrt{2\eta\epsilon}}{w_s}\right) + O(\epsilon^2).$$

Recall that $x_- = \frac{\epsilon}{w_s} + O(\epsilon^2)$. Thus,

$$g''(x_-) = -\frac{w_s^2}{\epsilon}\sqrt{1 - \frac{4\epsilon}{w_s}} = -\frac{w_s^2}{\epsilon} + 2w_s + O(\epsilon).$$

Using the earlier results on the location of the J_1 nullcline, we will have equilibria when

$$\frac{g(J_2)}{w_p} = 1 + O(\epsilon^n).$$

Expanding g about $J_2 = x_-$ and writing $w_p = w_p^{SN} - \eta$ gives,

$$g(J_2) = g(x_-) + \frac{(J_2 - x_-)^2}{2} g''(x_-)$$
$$+ O((J_2 - x_-)^3),$$
$$= w_p^{SN} - \eta + O(\epsilon^n),$$
$$(J_2 - x_-)^2 = -\frac{2\eta}{g''(x_-)} + O(\epsilon^3),$$

where the final line follows because $g(x_-) = w_p^{SN}$. Substituting for $g''(x_-)$ then gives the result. □

3.2 Three vertex cycle

Our second example is the cycle between three vertices shown schematically in Fig. 4. As a heteroclinic cycle
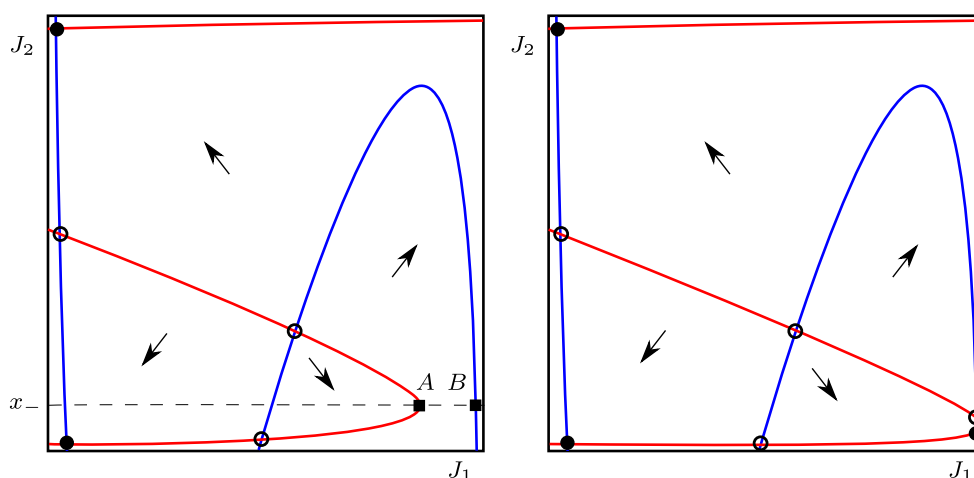
Fig. 3 Schematic illustration of the phase space and nullclines of the system (16). The blue curves are the J_1 nullclines, and the red curves are the J_2 nullclines. Solid dots show stable equilibria, and open dots show unstable equilibria. The parameter w_p is decreased from left fig-ure to right figure, creating two further equilibrium in a saddle-node bifurcation. The squares labelled A and B are referred to in the proof of lemma 1
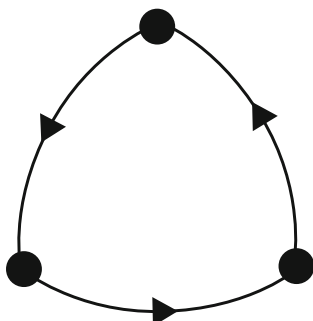

Fig. 4 Graph of the three vertex cycle

between equilibria, this system has been studied extensively in the fields of populations dynamics (May and Leonard 1975), rotating convection (Busse and Heikes 1980) and symmetric bifurcation theory (Guckenheimer and Holmes 1988).

We give some numerical examples of the dynamics of this system as realised by the CTRNN excitable network and use the continuation software AUTO (Doedel et al. 2007) to show that the transition from excitable to spontaneous dynamics occurs at a saddle-node on an invariant circle (SNIC) bifurcation generating a periodic orbit.

The deterministic equations realising this graph are:

$$\dot{y}_1 = -y_1 + w_s\phi(y_1) + w_m\phi(y_2) + w_p\phi(y_3),$$
$$\dot{y}_2 = -y_2 + w_s\phi(y_2) + w_m\phi(y_3) + w_p\phi(y_1),$$
$$\dot{y}_3 = -y_3 + w_s\phi(y_3) + w_m\phi(y_1) + w_p\phi(y_2). \quad (20)$$

We also consider the noisy case, using the setup given in equations (4).

Figure 5 shows sample time series for two different parameter sets. On the left, we show a noisy realisation with $w_p = 0.3$, and on the right, a periodic solution in the deterministic system (equations (1)) with $w_p = 0.305$.

Note that in both cases, for this system the y_k variables oscillate between three values: high ($y_k = Y_A = w_s = 1$), intermediate ($y_k = Y_L = w_p = 0.3$), and low ($y_k = Y_T = w_m = -0.7$), as the cells shift between Active, Trailing and Leading. Only the first of these corresponds to $J_k \approx 1$, as can be seen in the time series plots of the J_k variables in the lower panels of the figure, the other two correspond to $J_k \approx 0$.

We compute a bifurcation analysis of the system (20) using the continuation software AUTO (Doedel et al. 1997).

Figure 6 shows a bifurcation diagram of this system as w_p is varied. Stable solutions are shown in red. There is a saddle-node on an invariant circle (SNIC) bifurcation at $w_p = w_p^{SNIC} \approx 0.30287$. For $w_p < w_p^{SNIC}$, the diagram shows a stable equilibrium solution with $y_1 \approx Y_A = 1$ (and $y_2 \approx Y_L$, $y_3 \approx Y_T$). As w_p increases through w_p^{SNIC}, this equilibrium disappears in a SNIC bifurcation creating a stable periodic orbit. Note that the period of the periodic orbit asymptotes to ∞ as the SNIC bifurcation is approached. Due to the symmetry, there are of course two further pairs of equilibria, one pair with $y_1 \approx Y_T$, $y_2 \approx Y_A$, $y_3 \approx Y_L$, and another with $y_1 \approx Y_L$, $y_2 \approx Y_T$, $y_3 \approx Y_A$. The symmetry causes three saddle-node bifurcations to occur simultaneously, creating the periodic orbit. If we were to instead choose the w_p to be different in each of the lines in (20), the saddle-nodes would occur independently, and a periodic orbit would exist only if all three w_p's were greater than w_p^{SNIC}.

The time-series on the left-hand side of Fig. 5 has $w_p = 0.3 < w_p^{SNIC}$. Without noise, at these parameter values,
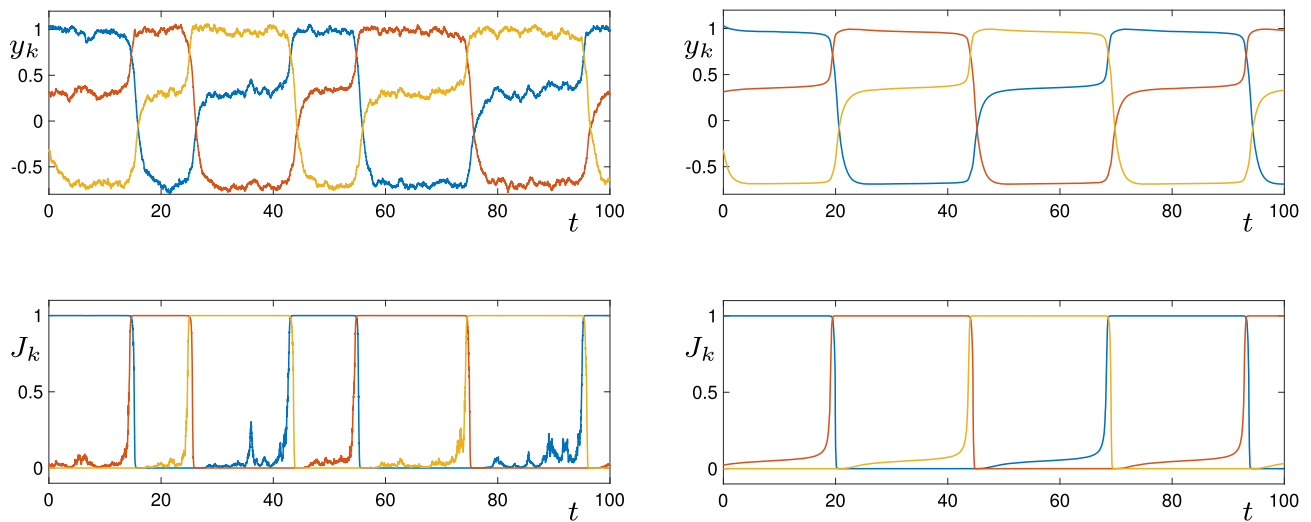
Fig. 5 For the three-vertex cycle (20), the top row of figures shows time series of the y_k variables in the excitable (left) and spontaneous (right) cases. In the left column, $w_p = 0.3$ and $\sigma = 0.05$. In the right column, $w_p = 0.305$ and $\sigma = 0$. The bottom row shows the J_k variables. Other parameters are $\epsilon = 0.05$, $\theta = 0.5$, $w_s = 1$, $w_m = -0.7$
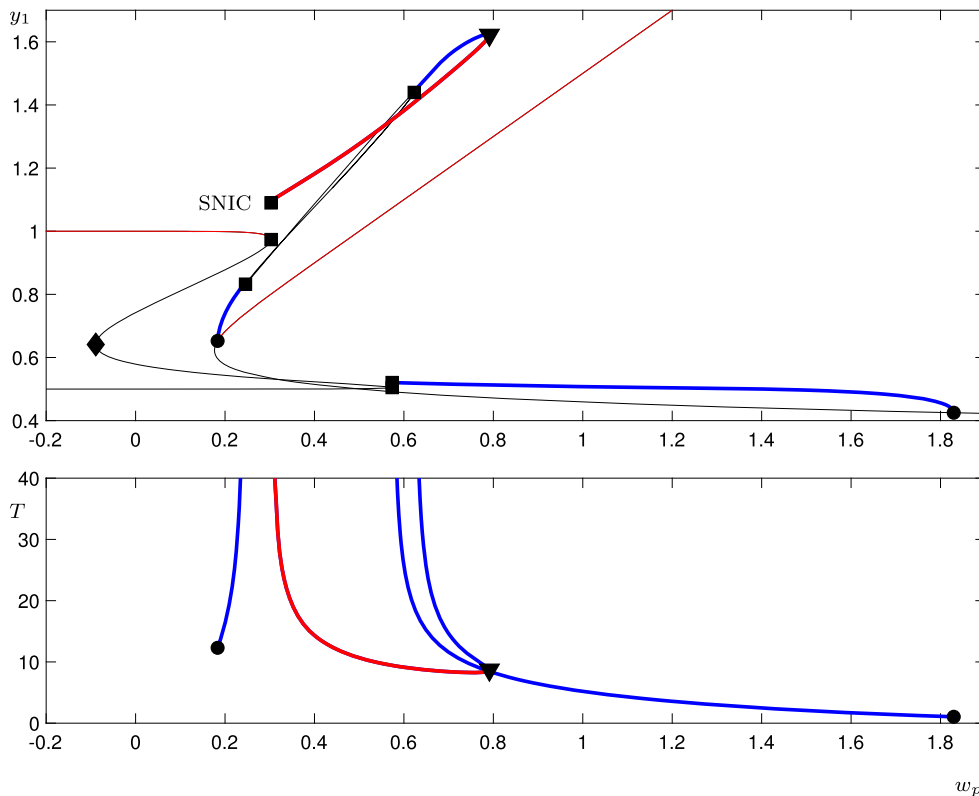

Fig. 6 For the three vertex cycle with equations (20), the figure shows a bifurcation diagram as w_p is varied. The top panel shows the y_1 coordinate of equilibrium solutions, and the maximum value of y_1 for periodic solutions. The lower panel shows the period of the periodic solutions. Equilibrium solutions are shown by a thin line, and periodic solutions by a thick line. Stable solutions are shown in red. Bifurcation points are indicated by various shapes: Hopf bifurcations by circles, saddle-node bifurcations by diamonds, saddle-node of periodic orbits by triangles, and saddle-node on invariant circles (SNIC) by squares. Note that there may be two squares for a single SNIC bifurcation because the maximum value of y_1 on the periodic orbit is not the same as the value of y_1 for the equilibria undergoing the saddle-node bifurcation. The SNIC bifurcation labelled at $w_p = w_p^{SNIC} \approx 0.30287$ is the transition between excitable and spontaneous dynamics
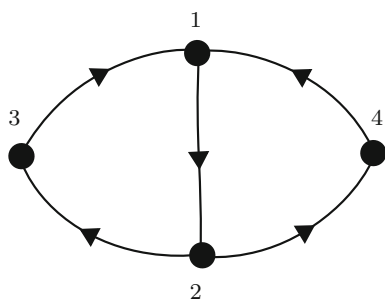
Fig. 7 Graph of the four vertex Kirk–Silber network

the system would remain at one equilibrium point indefinitely. The noise acts as inputs pushing the trajectory along the excitable connections. The time series on the right-hand side of Fig. 5 has $w_p = 0.305 > w_p^{SNIC}$ and shows the periodic orbit which has resulted from the SNIC bifurcation.

We note that this SNIC bifurcation occurs at approximately the same value of w_p as the saddle-node bifurcation found in Sect. 3.1. This is not surprising; using similar methods to those in the previous section, we can show that to lowest order in ϵ, the SNIC bifurcation occurs when $w_p = w_p^{SN}$. For $w_p < w_p^{SNIC}$, there thus exists an excitable network in the sense defined in appendix A.

3.3 Four node Kirk–Silber network

For our next example, we consider a graph with the structure of the Kirk–Silber network (Kirk and Silber 1994), shown schematically in Fig. 7. This graph has one vertex which has two outgoing edges, and the dynamics here are somewhat different to vertices with only one outgoing edge. The bulk of this section is devoted to an analysis of these differences, in particular, the possibility of an additional equilibrium in the network attractor with two active cells.

The corresponding deterministic equations for this network are (moving immediately into the J_i variables):

$$\dot{J}_1 = \frac{1}{\epsilon} J_1 (1 - J_1) \left(-\phi^{-1} J_1 + w_s J_1 + w_m J_2 + w_p J_3 + w_p J_4 \right),$$

$$\dot{J}_2 = \frac{1}{\epsilon} J_2 (1 - J_2) \left(-\phi^{-1} J_2 + w_p J_1 + w_s J_2 + w_m J_3 + w_m J_4 \right),$$

$$\dot{J}_3 = \frac{1}{\epsilon} J_3 (1 - J_3) \left(-\phi^{-1} J_3 + w_m J_1 + w_{p_3} J_2 + w_s J_3 + w_t J_4 \right),$$

$$\dot{J}_4 = \frac{1}{\epsilon} J_4 (1 - J_4) \left(-\phi^{-1} J_4 + w_m J_1 + w_{p_4} J_2 + w_t J_3 + w_s J_4 \right).$$

$$(21)$$

We can break the symmetry between J_3 and J_4 by choosing $w_{p_3} \neq w_{p_4}$. In fact, in what follows, we will frequently write $w_{p_3} = w_{p_4} + \Delta w$, for $\Delta w > 0$, and choose $w_{p_4} = w_p$ for simplicity.

We consider first the dynamics near each of the vertices that have exactly one outgoing edge (vertices 1, 3 and 4
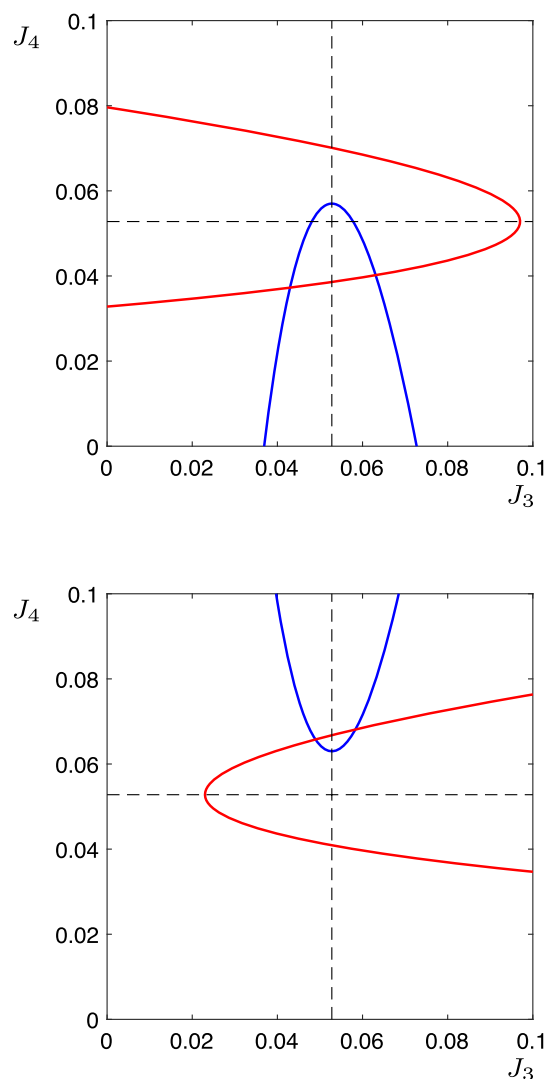

Fig. 8 Schematic illustration of the nullclines of the system (22) system along the surface $J_2 = 1$. The blue curves are the J_3 nullclines, and the red curves are the J_4 nullclines. In the upper panel, $w_t = 0.05$, $w_{p_3} = 0.30$, $w_{p_4} = 0.298$. In the lower panel, $w_t = -0.05$, $w_{p_3} = 0.306$, $w_{p_4} = 0.304$. The dashed lines show the where the nullcline would lie if $w_t = 0$

in the graph; see Fig. 7). Again using the same techniques that were used in Sect. 3.1 (lemma 2), we can show that for $w_p, w_m, w_t < \theta < w_s$, and $w_p = w_p^{SN} + \eta$, $0 < \eta < \frac{\epsilon}{4}$, there exist equilibria solutions at, for example

$$(J_1, J_2, J_3, J_4) = \left(1, \frac{\epsilon}{w_s} \pm \frac{\sqrt{2\eta\epsilon}}{w_s}, 0, 0 \right) + O(\epsilon^2).$$

That is, there is a transition from excitable to spontaneous dynamics (in this case between cells 1 and 2, but also between cells 3 and 1 and cells 4 and 1) as w_p is increased through w_p^{SN}.

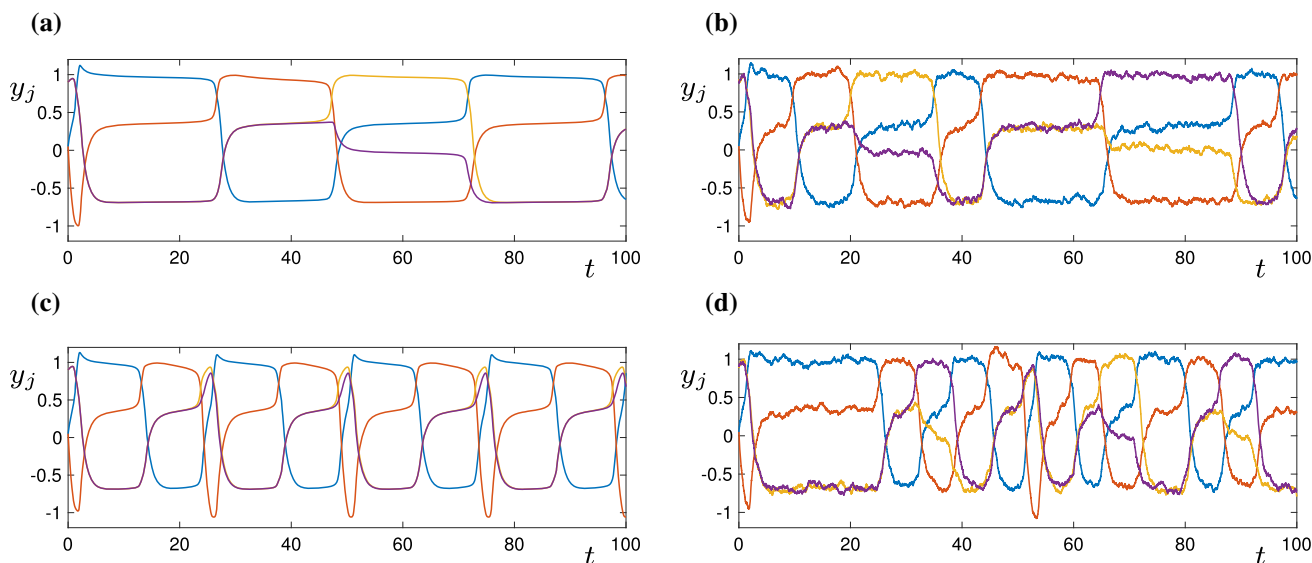The dynamics close to the vertex with two outgoing connections (vertex 2) is modified by the presence of the

(a)

(b)

(c)

(d)

Fig. 9 The figures show time series of simulations of the y_j variables in the Kirk–Silber type network (21). In (**a**), $w_p = 0.305$, and $\sigma = 0$; in (**b**), $w_p = 0.3$ and $\sigma = 0.05$; in (**c**) $w_p = 0.315$, and $\sigma = 0$; in (**d**), $w_p = 0.315$ and $\sigma = 0.05$

additional parameter w_t. Consider the three dimensional subsystem of (21) with $J_1 = 0$, that is:

$$\dot{J}_2 = \frac{1}{\epsilon} J_2 (1 - J_2) \left(-\phi^{-1} J_2 + w_s J_2 + w_m J_3 + w_m J_4 \right)$$

$$\dot{J}_3 = \frac{1}{\epsilon} J_3 (1 - J_3) \left(-\phi^{-1} J_3 + w_{p3} J_2 + w_s J_3 + w_t J_4 \right)$$

$$\dot{J}_4 = \frac{1}{\epsilon} J_4 (1 - J_4) \left(-\phi^{-1} J_4 + w_{p4} J_2 + w_t J_3 + w_s J_4 \right).$$
$$\tag{22}$$

We can perform a similar calculation to that shown in Sect. 3.1 to show that there is a section of the J_2 null-surface which lies asymptotically close to the surface $J_2 = 1$. Equilibria solutions exist on this null-surface if the J_3 and J_4 null-surfaces intersect there, that is, if there are solutions to the pair of equations

$$g(J_3) = w_{p3} + w_t J_4, \tag{23}$$
$$g(J_4) = w_{p4} + w_t J_3. \tag{24}$$

We assume without loss of generality that $w_{p3} > w_{p4}$ (i.e. $\Delta w > 0$) and then the arrangement of these curves is in one of the configurations shown in Fig. 8. If $w_t < 0$ (lower panel), equilibria solutions exist (i.e. the red and blue curves intersect) for a range of w_{p3} and w_{p4} with both larger than w_p^{SN}: that is, the transition to spontaneous dynamics happens at a larger value of w_{p_j} (than if $w_t = 0$). If $w_t > 0$ (upper panel), the opposite happens: that is, the transition to spontaneous dynamics occurs at a smaller value of w_{p_j}. Solving these equations exactly requires solving a quartic equation, and the resulting expression is not illuminating. We label

the value of w_p at which this transition from spontaneous to excitable dynamics occurs as $w_p^{SN'}$, and note that this is a function of $w_t, w_s, \epsilon, \theta$ as well as more generally, the number of Leading directions from that cell.

For the specific system (21), with $w_{p3} = w_{p4} + \Delta w = w_p + \Delta w$, we thus have two conditions. If $w_p < \min(w_p^{SN}, w_p^{SN'} - \Delta w)$, then the system is excitable along all connections. If $w_p > \max(w_p^{SN}, w_p^{SN'} - \Delta w)$, then a periodic orbit will exist. If neither of these conditions holds, then we will see excitable connections in some places and spontaneous transitions in others.

In Fig. 9, we show some example time-series of the system (21) (in the y_k coordinates). In panel (a), parameters are such that $w_p > \max(w_p^{SN}, w_p^{SN'} - \Delta w)$, so we see a periodic solution in the deterministic system. Note that the y_3 (yellow) coordinate becomes close to $Y_A = 1$ during this trajectory, but the y_4 (purple) coordinate does not: it switches between $Y_L = 0.3$, $Y_D = 0$ and $Y_T = 0.7$. In panel (b), parameters are such that $w_p < \min(w_p^{SN}, w_p^{SN'} - \Delta w)$, so without noise the trajectory would remain at a single equilibrium solution. Here, we add noise with $\sigma = 0.05$, and the trajectory can be seen exploring the network. Note that there are some transitions between ξ_2 (y_2 is red) and ξ_3 (y_3 is yellow), and some from ξ_2 to ξ_4 (y_4 is purple). In panels (c) and (d), we increase w_p further away from the saddle-node bifurcation (further into the regime of spontaneous transitions) and observe some qualitative differences in the trajectories. In the deterministic case (c), the periodic solution now transitions from near the bottleneck region P_2 to a region of phase space where y_3 and y_4 are both Active. We label this region of phase space as $P_{3,4}$. In the noisy case (d) (which is also in the spontaneous
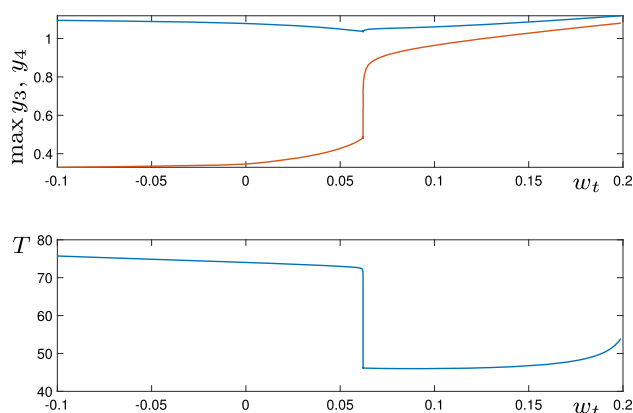
Fig. 10 Behaviour of periodic orbits in the Kirk–Silber-type four-node network (21) as the parameter w_t is varied. The top panel shows $\max(y_3)$ (blue) and $\max(y_4)$ (red) along the periodic orbit. The bottom panel shows the period of the orbit on varying w_t

regime), the trajectory makes transitions from the bottleneck region P_2 to each of P_3, P_4, and to $P_{3,4}$.

In the above simulations, we have used $w_t = 0$, but we observe that the type of qualitative behaviour observed depends both on the parameters w_p and w_t. Specifically, a sufficiently negative w_t provides a suppression effect, meaning that only a single cell y_k can be active at any one time, but the transitional value of w_t depends on w_p. In Fig. 10, we show maximum values of y_3 and y_4 along the periodic orbits as w_t is varied. It can be seen clearly here that the transition between periodic orbits which visit P_3 (where $\max(y_3)$ is significantly larger than $\max(y_4)$) and those which visit $P_{3,4}$ (where $\max(y_3) \approx \max(y_4)$) is quite sharp.

We extend these results to show the behaviour as both parameters w_p and w_t are varied in Fig. 11. The data in this figure show the observed behaviours for both noisy and deterministic systems as the parameters w_p and w_t are varied. The red lines are the curves $w_p = w_p^{SN}$ (dotted) and $w_p = w_p^{SN'}$ (dashed). If w_p is above both of these lines, then all transitions are spontaneous, and so a periodic orbit exists in the system. If w_p lies below either one (or both) of these lines, then at least one of the transitions will be excitable and so there will be no periodic solutions. The black line shows the boundary between those periodic solutions which visit P_3 (to the left of the black line) and those which visit $P_{3,4}$ (to the right of the black line), as determined by the location of the sharp transition in calculations similar to those shown in Fig. 10 for a range of w_p. The background colours are results from noisy simulations. The colour indicates the ratio of transitions to $P_{3,4}$ to the total number of transitions to P_3, P_4 and $P_{3,4}$. Interestingly, the noisy solutions require a much larger value of w_t than the deterministic ones to have a significant proportion of transitions to $P_{3,4}$.

These changes in qualitative dynamics can be explained in terms of the three-dimensional subsystem with $J_1 = 0$,

given by equations (22). In this three-dimensional system, there are stable equilibria at

$$(J_2, J_3, J_4) = (0, 1, 0) + O(\epsilon)$$
$$(J_2, J_3, J_4) = (0, 0, 1) + O(\epsilon)$$
$$(J_2, J_3, J_4) = (0, 1, 1) + O(\epsilon)$$

as well as further unstable/saddle equilibria. Recall that these equilibria are not on the boundaries of the box (which are not part of the domain). In Fig. 12, we show solutions from the full four-dimensional system (21) projected onto the three-dimensional space with $J_1 = 0$. In panel (a), we show the periodic solutions from the deterministic systems for five different values of w_p, ranging from $w_p = 0.309$ (left curve, dark purple), to $w_p = 0.3096$ (right curve, red) in increments of 0.0002. It can be seen that the first two of these trajectories approach the saddle equilibria (marked as a blue dot) from one side of its stable manifold, and the latter three from the other. The first three thus visit P_3, and the latter three visit $P_{3,4}$. It is the transition of the periodic orbit across the stable manifold which results in the rapid change in the qualitative behaviour of the periodic orbit and likely indicates that a homoclinic bifurcation to this saddle point separates these behaviours. That is, the sharp transition in T in Fig. 10 should actually extend to ∞ on both sides. In panel (b), we show both noisy and deterministic trajectories with $w_p = 0.315$. Note that only one of the noisy trajectories follows the deterministic trajectory closely: the majority visit P_3.

3.4 A ten node network

In this section, we demonstrate the method of construction described in Sect. 2 for a larger network. Specifically, we randomly generated a directed graph between 10 vertices, with the constraints that it contained no one-loops, two-loops or Δ-cliques, and such that the graph does not have feed-forward structure (i.e. you cannot get 'stuck' in a subgraph by following the arrows). The graph we consider is shown in Fig. 13.

We ran one simulation of the deterministic CTRNN system (1), and two simulations of the noisy CTRNN system (4), and in each case, randomly generated the entries for the w_p in the equation (8). For the deterministic system, the entries of w_p^{ij} were chosen independently from the uniform distribution $U(0.32, 0.34)$. For the noisy systems, the entries of w_p^{ij} were chosen independently from the uniform distribution $U(0.30, 0.32)$. The remaining parameters were set at the default parameter values given in (12) except $w_t = -0.3$ for the deterministic system, and one of the noisy systems, and $\sigma = 0.01$ for the noisy systems. Note that for the deterministic parameter values there are bottlenecks in the phase space close to the locations in phase space for the excitable

states that are present for the parameter values used in the stochastic cases.

The results of the simulations are shown in Figs. 14, 15 and 16. For the deterministic simulation (Fig. 14), we see that the system has an attracting period orbit, which visits the nodes in the order $1 \rightarrow 4 \rightarrow 7 \rightarrow 3 \rightarrow 8$. The entries of w_p^{ij} were randomly generated as described above, and we do not give them all here for space reasons, but we note that in all cases in which a vertex in this cycle has two 'choices' for which direction to leave (in the graph shown in Fig. 13), the attracting periodic orbit chooses the more unstable direction. That is, if i is a vertex in the above cycle, and if $i \rightarrow j$ and $i \rightarrow k$ are connections in the directed graph, with $i \rightarrow$

j being a part of the attracting periodic orbit, then $w_p^{ji} > w_p^{ki}$. Although we do not prove here that this will always be the case, it is intuitively what one might expect, that is, the connection from i to j is stronger than the connection from i to k.

In the noisy simulation with $w_t = -0.3$ (Fig. 15), the equilibria are not visited in a regular pattern, but random choices are made at each equilibria from which there is more than one direction in which to leave. See, for instance, the transition $3 \rightarrow 8$ at $t \approx 75$, and the transition $3 \rightarrow 10$ at $t \approx 155$. The length of time spent near each equilibria is also irregular; note for instance, the variable amount of time spent near ξ_1 and ξ_3. In this simulation, because the transverse
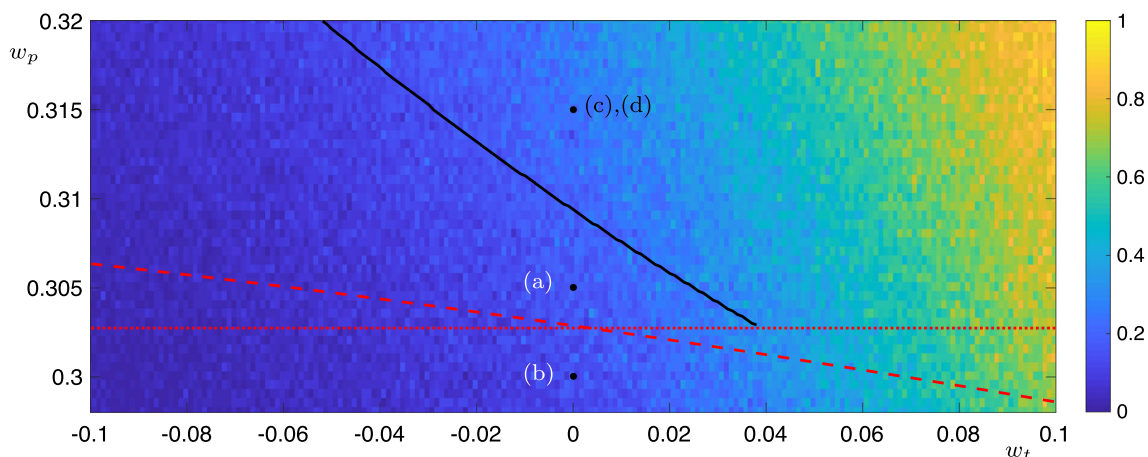

Fig. 11 Behaviour of the Kirk–Silber-type four-node network (21) as the parameters w_t and w_p are varied. The red lines are the curves $w_p = w_p^{SN}$ (dotted) and $w_p = w_p^{SN'}$ (dashed; determined numerically by solving a quartic equation). For w_p above, both of these lines periodic solutions exist in the deterministic system. To the left of the black line, these periodic solutions visit P_3, to the right they visit $P_{3,4}$. The background colours are results from noisy simulations with $\sigma = 0.05$. The colour indicates the ratio of transitions to $P_{3,4}$ to the total number of transitions to P_3, P_4 and $P_{3,4}$. The labelled dots give the parameter values of the time-series plots in Fig. 9
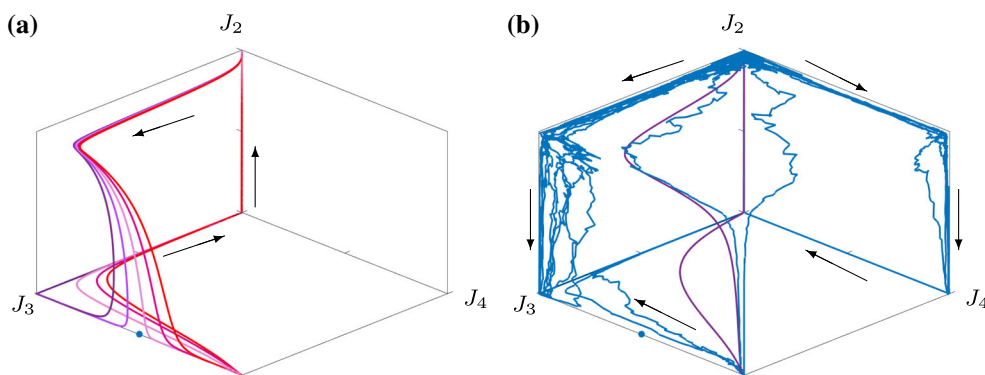

Fig. 12 The figures show trajectories for the system (21) projected onto the three-dimensional space with $J_1 = 0$. In panel (**a**), five periodic trajectories in the deterministic system are shown for different values of w_p, ranging from $w_p = 0.309$ (left curve, dark purple), to $w_p = 0.3096$ (right curve, red) in increments of 0.002. In panel (**b**), we set $w_p = 0.315$ (as in panels (**c**) and (**d**) of Fig. 9). Noisy trajectories ($\sigma = 0.05$) are shown in blue, and the periodic orbit of the deterministic system is shown in purple. In both panels, an equilibrium of the three-dimensional system (22) is shown by a blue dot. Other parameters are $\varepsilon = 0.05$, $\theta = 0.5$, $w_t = 0$, $w_s = 1$, $w_m = -0.5$. The arrows indicate the direction of flow
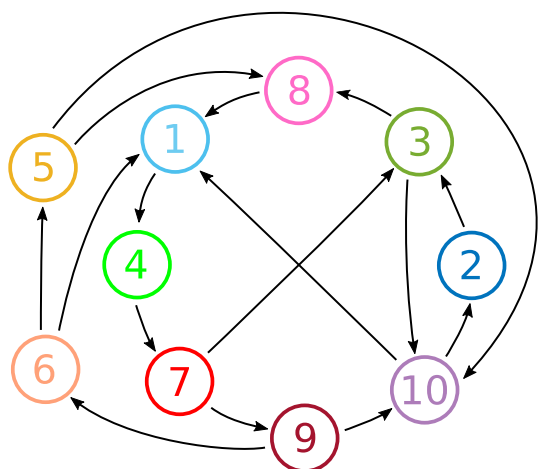
Fig. 13 A example of a directed graph between ten nodes with no loops of order one or two and no Δ-cliques. Three timeseries from realisations of this as a network using CTRNNs are shown in Figs. 14–16

parameter w_t is sufficiently negative, only one node is active at any given time.

By contrast, the simulation in Fig. 16 has $w_t = 0$. Here, the transitions again are made randomly, but without the suppression provided by the transverse parameter it is possible to have multiple cells active at once. For instance, at around $t = 55$, both cells 1 and 5 become active at the same time; they were both leading cells from the previously active cell 6. As cells 1 switches off, cell 4 becomes active, and as cell 5 switches off, cell 8 becomes active. The system continues to have two active cells around $t = 250$, at which point a third cell also becomes active. If the trajectory was to run for longer, then the number of active cells could decrease again, if an active cell suppresses more than one previously active cells.

The entire excitable network attractor for this level of noise is clearly more complicated than the design shown in Fig. 13, in that additional equilibria (with more than one active cell) are accessible to those encoded and described in Theorem 1. An interesting extension of this work would be to understand which additional equilibria appear in a network attractor generated from a given directed graph in this manner: Fig. 16 suggests that at least seven levels of cell activity are needed to uniquely describe the states that can appear when more than one cell becomes "active". In particular, when a cell is Trailing to more than one Active cell, the value of y_j for that cell is even lower than Y_T (compare the top panel of Fig. 16 with the schematic in Fig. 2).

4 Discussion

The main theoretical result of this paper is Theorem 1, which states that it is possible to design the connection weight

matrix of a CTRNN such that there exists a network attractor with a specific graph topology embedded within the phase space of the CTRNN. The graph topology is arbitrary except for minor restrictions: namely there should be no loops of order one or two, and no Δ-cliques. Theorem 1 assumes a piecewise affine activation function, but the examples in Sect. 3 suggest that the results generalise to CTRNN using any suitable smooth activation function. More generally, note that the coupled network is in some sense close to N simultaneous saddle-node bifurcations. However, the units are not weakly coupled and indeed this is necessary to ensure that when one cell becomes active, the previous active cell is turned off.

Theorem 1 proves the existence of an excitable network with threshold δ where not only the connection weights, but also ϵ and θ (properties of the activation function) may depend on δ. We believe that a stronger result will be true, namely that δ can be chosen independent of properties of the activation function, and also that this can be made an almost complete realisation by appropriate choice of parameters.

Conjecture 1 Assume the hypotheses on the directed graph G with N vertices as in Theorem 1 hold. Assume that $\epsilon > 0$ is small and $\theta > 0$. Then, there is a $\delta_c(\epsilon, \theta) > 0$ such that for any $0 < \delta < \delta_c$ there is an open set $\hat{W}_{\text{ex}} \subset \mathbb{R}^4$. If the parameters $(w_s, w_m, w_t, w_p) \in \hat{W}_{\text{ex}}$, then the dynamics of input-free equation (1) with N cells and piecewise affine activation function (3) and w_{ij} defined by (8) contains an excitable network attractor with threshold δ that gives an almost complete realisation of the graph G.

The construction in Theorem 1 uses a comparatively sparse encoding of network states—each of the N vertices in the network is associated with precisely one of the N cells being in an active state. Indeed, the connection weights (8) assign one of only four possible weights to each connection, depending on whether that cell can become active next, was active previously or neither. Other choices of weights will allow more dense encoding: and many more than N excitable states within a network of N cells. However, the combinatorial properties of the dynamics seem to be much more difficult to determine and presumably additional connection weightings will be needed, not the just four values considered in Theorem 1.

Section 3 illustrates specific examples of simple excitable networks for smooth activation function (2) on varying parameters—this requires numerical continuation to understand dependence on parameters even for fairly low dimension. For this reason, we expect that a proof of an analogous result to Theorem 1 for the smooth activation function may be a lot harder. These examples also give some insight into bifurcations that create the excitable networks.

In general, there is no reason that the realisation constructed in Theorem 1 is *almost complete* (in the sense that
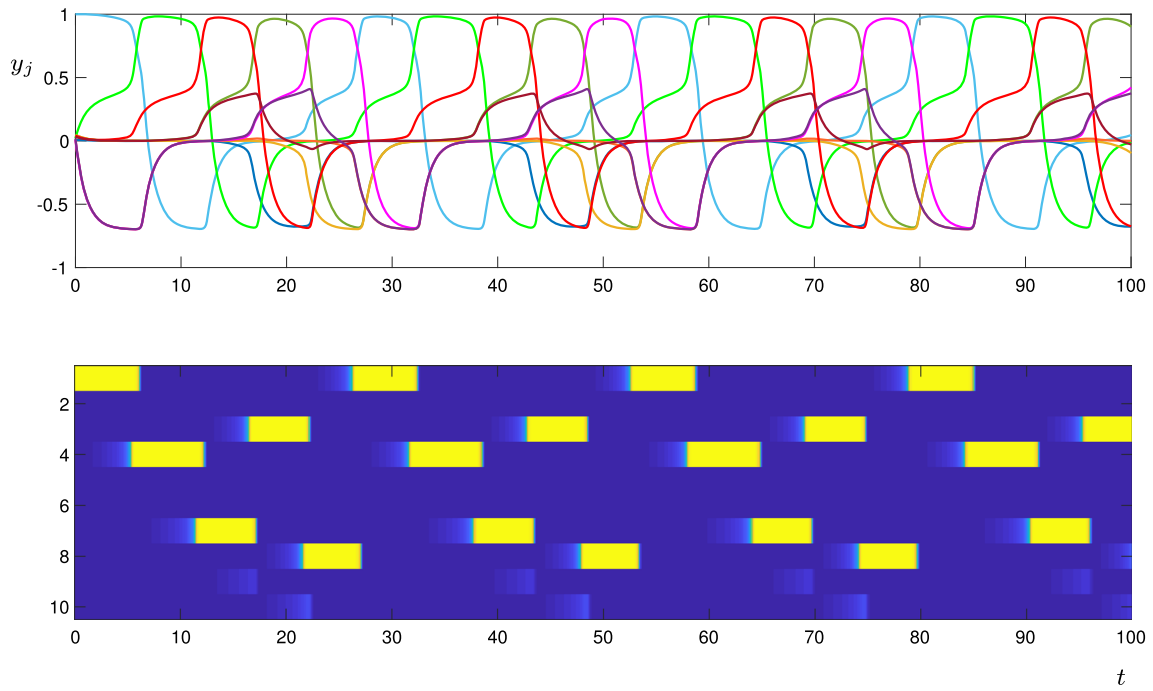
Fig. 14 Trajectories of the CTRNN network for the directed graph shown in Fig. 13, using the deterministic model (1). The top panel shows the y_j coordinates, where the colours correspond to the node colours in Fig. 13. In the lower panel, each horizontal row corresponds to one node (as labelled on the vertical axis), and the colour is blue when the corresponding J_j coordinate is close to zero, and yellow when it is close to one. That is, the yellow segments indicate when each node is active. Parameters are as described in the text
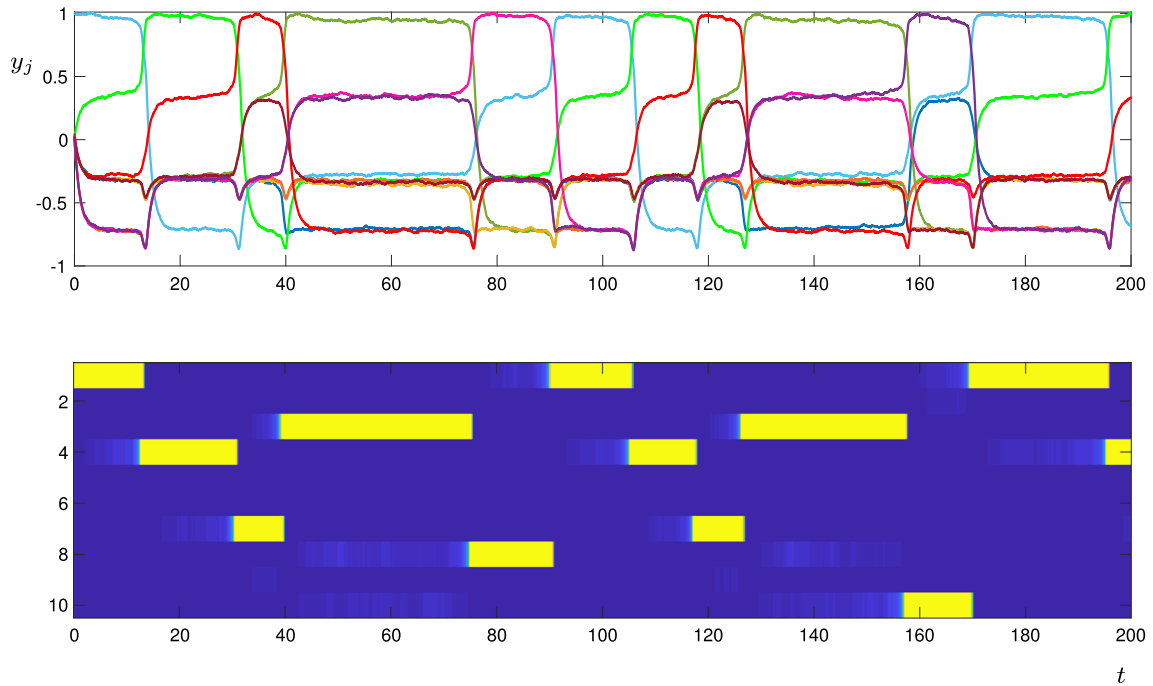
Fig. 15 Trajectories of the CTRNN network for the directed graph shown in Fig. 13, using the stochastic model (4). Lines and colours are described in Fig. 14. Parameters are as described in the text, here $w_t = -0.3$
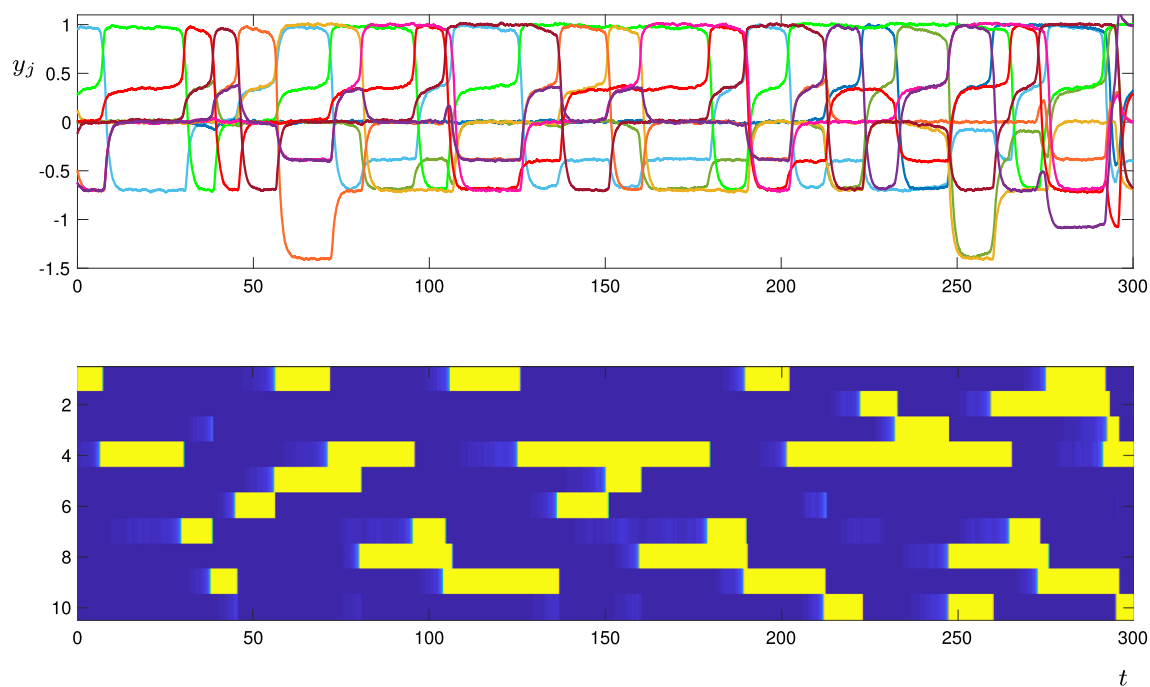
Fig. 16 Trajectories of the CTRNN network for the directed graph shown in Fig. 13, using the stochastic model (4). Lines and colours are described in Fig. 14. Parameters are as described in the text, here $w_t = 0$

almost all initial conditions in $B_\delta(\xi_k)$ evolve towards some ξ_l with $a_{kl} = 1$, analogous to (Ashwin et al. 2020, Defn 2.6)). If it is not, then other attractors may be reachable from the excitable network. Conjecture 1 suggests that the realisation can be made almost complete for small enough δ: we expect that for this we will require w_t to be sufficiently negative. If δ is too large, we cannot expect almost completeness: there are other stable equilibria (notably the origin) that can be reached with large perturbations, and the simulation in Fig. 16 shows that other equilibria may be reachable from the network. It will be a challenge to strengthen our results to show that the excitable network is an almost complete realization. However, the examples studied in Sect. 3 confirm that, at least for relatively simple graphs, this conjecture is reasonable.

Our theoretical results are for networks with excitable connections. We expect much of the behaviour described here is present in the spontaneous case, if the coupling weights are chosen such that equilibria are replaced with bottlenecks. In the absence of noise, we expect to see a deterministic switching between slow moving dynamics within bottlenecks. This dynamical behaviour is very reminiscent of the stable heteroclinic channels described, for example, in (Afraimovich et al. 2004; Rabinovich et al. 2020). However, stable heteroclinic switching models require structure in the form of multiplicative coupling or symmetries that are not present in CTRNN or related Wilson–Cowan neural models (Wilson and Cowan 1972) (see Chow and Karimipanah 2020 for a recent review of related neural models). Other models showing sequential

excitation include (Chow and Karimipanah 2020): this relies on a fast-slow decomposition to understand various different modes of sequential activation in a neural model of rhythm generation. It will be an interesting challenge to properly describe possible output dynamics of our model in the case of bottlenecks.

We remark that asymmetry of connection weights is vital for constructing a realisation as an excitable networks—indeed, the lack of two-cycles precludes $a_{jk} = a_{kj} = 1$. While this may be intuitively obvious, it was not so obvious that we also need to exclude one-cycles and Δ-cliques in the graph to make robust realisations.

Finally, although we do not consider specific natural or machine learning applications of CTRNN here, the structures found here may give insights that give improved training for CTRNN. In particular, it seems plausible that CTRNN may use excitable networks to achieve specific input-output tasks (especially those requiring internal states). For example, recent work (Ceni et al. 2020) demonstrates that echo state networks can create excitable networks in their phase space to encode input-dependent behaviour. It is also likely there are novel optimal training strategies that take advantage of excitable networks, for example, choosing connection weights that are distributed close to one of the four values we use.

A Definition of excitable network

This appendix (which extends ideas in Ashwin and Postlethwaite (2016)) gives formal definitions for excitable networks considered in this paper. We say a system has an *excitable connection for amplitude* $\delta > 0$ from one equilibrium ξ_i to another ξ_j if

$$B_\delta(\xi_i) \cap W^s(\xi_j) \neq \emptyset,$$

(where $B_\delta(\xi)$ is the open ball of radius δ centred at ξ) and this connection has *threshold* δ_{th} if

$$\delta_{th} = \inf\{\delta > 0 : B_\delta(\xi_i) \cap W^s(\xi_j) \neq \emptyset\}.$$

We define the *excitable network of amplitude* $\delta > 0$ between the equilibria $E = \{\xi_i\}$ to be the set

$$\Sigma_E = \bigcup_{i,j=1}^N \{\Phi_t(x) : x \in B_\delta(\xi_i), t > 0\} \cap W^s(\xi_j).$$

We say the excitable network Σ_E for amplitude δ *realises a graph* G if each vertex v_i in G corresponds to an equilibrium ξ_i in Σ_E and there is an edge in G from v_i to v_j if only if there is a connection in Σ_E for amplitude δ from ξ_i to ξ_j.

B Proof of Theorem 1

As in equations (9) and (10), for any $\delta < \frac{1}{2}$ we choose

$$\epsilon = \frac{\delta}{8}, \ \theta = \frac{1}{2}, \ w_s = 1, \ w_t = 0, \tag{25}$$

and then w_p and w_m are given by

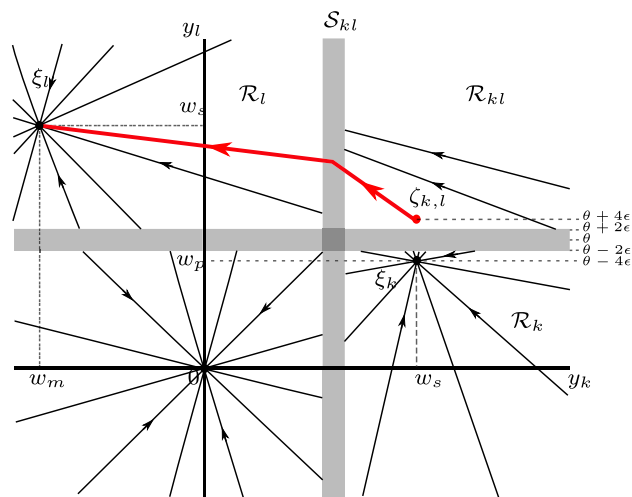$$w_p = \theta - \frac{\delta}{2}, \quad w_m = -(w_s - \theta) - \frac{\delta}{2}. \tag{26}$$

Fig. 17 Illustration of a connection (shown in red) with amplitude $\delta > 0$ from ξ_k to ξ_l, projected in the coordinates y_k and y_l. The points ξ_k, ξ_l and 0 are linear sinks. The point $\zeta_{k,l}$ is within δ of ξ_l and limits to ξ_l in forwards time. The grey areas are regions of width 4ϵ centred around $y_{k,l} = \theta$ where $\phi_P(y_k)$ and $\phi_P(y_l)$ are non-constant: these contain saddle and other equilibria (not shown)

We define equilibria ξ_k as in Sect. 2.1, where we write the jth component of the equilibrium as $[\xi_k]_j$:

$$[\xi_k]_j = \begin{cases} Y_A & \text{if } j = k, \\ Y_L & \text{if } a_{kj} = 1, \\ Y_T & \text{if } a_{jk} = 1, \\ Y_D & \text{if } a_{kj} = 0 \text{ and } a_{jk} = 0, \end{cases} \tag{27}$$

where

$$Y_A := w_s = 1, \ Y_L := w_p = \theta - \frac{\delta}{2},$$

$$Y_T := w_m = -(w_s - \theta) - \frac{\delta}{2}, \ Y_D := w_t = 0,$$

are the values of the Active, Leading, Trailing and Disconnected components, respectively. Note that

$$Y_T < Y_D < Y_L < \theta - \epsilon/2 < \theta + \epsilon/2 < Y_A. \tag{28}$$

As mentioned before, the hypotheses of Theorem 1 imply that this labelling is well defined and Fig. 2 shows how a transition from cell 1 active to cell 2 active will occur in a general network. It is simple to check that (27) is an equilibrium solution of (1) with activation function (3). Moreover, ξ_k is linearly stable with n eigenvalues -1. We define $[\mathcal{J}_k]_j := \phi_P([\xi_k]_j)$ then note that (28) implies that

$$[\mathcal{J}_k]_j := \delta_{kj},$$

in terms of the Kronecker δ_{kj}.

We consider two cases. Case 1 is where $a_{kl} = 1$ and we expect to see a connection from ξ_k to ξ_l. Case 2 is $a_{kl} = 0$ and we do not expect a connection.

B.1 Case 1

Suppose $a_{kl} = 1$ (recall that the lack of 1-cycles means that $k \neq l$), and then we define two regions of phase space, which we label $\mathcal{R}_l$ and $\mathcal{R}_{kl}$, as follows:

$$\mathcal{R}_l = \left\{ \mathbf{y} \mid y_j > \theta + \frac{\delta}{4}, \ j = l; \ y_j < \theta - \frac{\delta}{4} \text{ otherwise} \right\},$$

$$\mathcal{R}_{kl} = \left\{ \mathbf{y} \mid y_j > \theta + \frac{\delta}{4}, \ j = l, k; \ y_j < \theta - \frac{\delta}{4} \text{ otherwise} \right\}.$$

The regions and the dynamics within them are shown schematically in Fig. 17.

Within $\mathcal{R}_{kl}$, if $a_{kl} = 1$, then the dynamics are governed by the equations

$$\dot{y}_j = f_j^{(b)} := -y_j + \phi_j^{(b)},$$

where

$$\phi_j^{(b)} := \begin{cases} w_s + w_m & \text{if } j = k, \text{(AT)} \\ w_s + w_p & \text{if } j = l, \text{(LA)} \\ w_p & \text{if } (a_{kj}, a_{jk}, a_{lj}, a_{jl}) = (1, 0, 0, 0), \text{(LD)} \\ w_m & \text{if } (a_{kj}, a_{jk}, a_{lj}, a_{jl}) = (0, 1, 0, 0), \text{(TD)} \\ w_p & \text{if } (a_{kj}, a_{jk}, a_{lj}, a_{jl}) = (0, 0, 1, 0), \text{(DL)} \\ w_m & \text{if } (a_{kj}, a_{jk}, a_{lj}, a_{jl}) = (0, 0, 0, 1), \text{(DT)} \\ w_m + w_p & \text{if } (a_{kj}, a_{jk}, a_{lj}, a_{jl}) = (0, 1, 1, 0), \text{(TL)} \\ 0 & \text{if } (a_{kj}, a_{jk}, a_{lj}, a_{jl}) = (0, 0, 0, 0), \text{(DD)} \end{cases}$$
(29)

where the type of coordinate (see Fig. 2) is given in parentheses on each line. Within $\mathcal{R}_l$, the dynamics are governed by the equations

$$\dot{y}_j = f_j^{(c)} := -y_j + \phi_j^{(c)},$$
(30)

where

$$\phi_j^{(c)} := \begin{cases} w_m & \text{if } j = k, \text{(AT)} \\ w_s & \text{if } j = l, \text{(LA)} \\ w_p & \text{if } a_{lj} = 1 \text{ and } j \neq l, \text{(DL/TL)} \\ w_m & \text{if } a_{jl} = 1, \text{(DT)} \\ 0 & \text{if } a_{lj} = 0 \text{ and } a_{jl} = 0, \text{(LD/TD/DD)} \end{cases}$$
(31)

The equilbrium ξ_l lies in the interior of the region $\mathcal{R}_l$, whereas ξ_k lies in the interior of the region $\mathcal{R}_k$. We show that there is a connection of amplitude δ from ξ_k to any ξ_l with $a_{kl} = 1$, by considering a trajectory starting at

$$\zeta_{k,l} = \xi_k + \delta e_l,$$

where e_l is a unit vector in the l-direction: clearly $|\zeta_{k,l} - \xi_k| = \delta$ (see Fig. 17). Note that

$$[\zeta_{k,l}]_j = \begin{cases} Y_A & \text{if } j = k, \\ Y_L + \delta & \text{if } j = l, \\ Y_L & \text{if } a_{kj} = 1 \text{ and } j \neq l, \\ Y_T & \text{if } a_{jk} = 1, \\ Y_D & \text{if } a_{kj} = 0 \text{ and } a_{jk} = 0, \end{cases}$$
(32)

where Y_A, etc. are given in equation (11), and $Y_L + \delta = \theta + \delta/2 > \theta$. We define $[\mathcal{J}_{k,l}]_j := \phi_P([\zeta_{k,l}]_j)$, and then

$$[\mathcal{J}_{k,l}]_j = \delta_{jk} + \delta_{jl}.$$

Now, we note that $\zeta_{k,l} \in \mathcal{R}_{kl}$. We next show that a trajectory with initial condition at $\zeta_{k,l}$ will asymptotically approach ξ_l. We show first that the trajectory enters $\mathcal{R}_l$ in finite time. Then, since in $\mathcal{R}_l$, the flow is linear, with stable equilibrium ξ_l, all trajectories in $\mathcal{R}_l$ eventually approach ξ_l.

Define $\mathcal{S}_{kl}$ to be the region between $\mathcal{R}_l$ and $\mathcal{R}_{kl}$, namely,

$$\mathcal{S}_{kl} = \left\{ \mathbf{y} \mid \theta - \frac{\delta}{4} \leq y_k \leq \theta + \frac{\delta}{4}; \ y_l > \theta + \frac{\delta}{4}; \ y_j < \theta - \frac{\delta}{4} \text{ for } j \neq k, l \right\}.$$

First, consider the dynamics of all y_j, with $j \neq k, l$. This means that $y_j < \theta - \delta/4$ for all points in $\mathcal{T}_{kl} \equiv \mathcal{R}_{kl} \cup \mathcal{S}_{kl} \cup \mathcal{R}_l$. While the trajectory remains in $\mathcal{T}_{kl}$, it can be shown that $\dot{y}_j$ is negative along the line with $y_j = \theta - \delta/4$. Hence, none of the y_j will leave $\mathcal{T}_{kl}$.

Within $\mathcal{S}_{kl}$, the dynamics of y_k and y_l are governed by

$$\dot{y}_k = -y_k + w_m + \phi_P(y_k),$$
(33)
$$\dot{y}_l = -y_l + w_s + w_p \phi_P(y_k).$$
(34)

Consider the equation for $\dot{y}_k$ in $\mathcal{R}_{kl}$, $\mathcal{S}_{kl}$ and $\mathcal{R}_l$, namely:

$$\dot{y}_k = \begin{cases} -y_k + w_s + w_m & \text{if } \mathbf{y} \in \mathcal{R}_{kl}, \\ -y_k + \phi_P(y_k) + w_m & \text{if } \mathbf{y} \in \mathcal{S}_{kl}, \\ -y_k + w_m & \text{if } \mathbf{y} \in \mathcal{R}_l. \end{cases}$$

Recall that $w_m = -(w_s - \theta) - \frac{\delta}{2}$, and $0 \leq \phi_P(y_k) \leq 1$. We can use these bounds to show that

$$\dot{y}_k \leq \begin{cases} -\dfrac{3\delta}{4} & \text{if } \mathbf{y} \in \mathcal{R}_{kl}, \\ -\dfrac{\delta}{4} & \text{if } \mathbf{y} \in \mathcal{S}_{kl}, \end{cases}$$

Furthermore, if $\mathbf{y} \in \mathcal{R}_l$ and $y_k > 0$, then $\dot{y}_k < -w_m < 0$. In particular, we note that for $\mathbf{y} \in \mathcal{T}_{kl}$, with $y_k > 0$, $\dot{y}_k$ is negative and bounded below zero.

Now, consider the equation for $\dot{y}_l$ in $\mathcal{R}_{kl}$, $\mathcal{S}_{kl}$ and $\mathcal{R}_l$, namely:

$$\dot{y}_l = \begin{cases} -y_l + w_s + w_p & \text{if } \mathbf{y} \in \mathcal{R}_{kl}, \\ -y_l + w_s + w_p \phi_P(y_k) & \text{if } \mathbf{y} \in \mathcal{S}_{kl}, \\ -y_l + w_s & \text{if } \mathbf{y} \in \mathcal{R}_l, \end{cases}$$

We use this to compute $\dot{y}_l$ along the lower boundary of the three regions, $\mathcal{R}_{kl}$, $\mathcal{S}_{kl}$ and $\mathcal{R}_l$, that is, the line $y_l = \theta + \frac{\delta}{4}$, and we find

$$\dot{y}_l = \begin{cases} w_s - \frac{3\delta}{4} & \text{if } \mathbf{y} \in \mathcal{R}_{kl}, \\ w_s - \theta - \frac{\delta}{4} + \left(\theta - \frac{\delta}{2}\right)\phi_P(y_k) > w_s - \frac{3\delta}{4} & \text{if } \mathbf{y} \in \mathcal{S}_{kl}, \\ w_s - \theta - \frac{\delta}{4} & \text{if } \mathbf{y} \in \mathcal{R}_l. \end{cases}$$

Since $w_s > \frac{3\delta}{4}$ and $w_s > \theta + \frac{\delta}{4}$, we see that $\dot{y}_l > 0$ in all three cases.

Combining our knowledge of $\dot{y}_l$ and $\dot{y}_k$ tells us that a trajectory which starts in $\mathcal{R}_{kl}$, or more specifically, a trajectory starting in a small neighbourhood of $\zeta_{k,l}$ will have monotonic decreasing y_k component until (at least) $y_k = 0$. Furthermore, the y_l component cannot decrease below $y_l = \frac{1}{2} + \frac{\delta}{4}$. Thus, the trajectory will move through $\mathcal{S}_{kl}$ and into $\mathcal{R}_l$ in a bounded time.

Within $\mathcal{R}_l$, ξ_l is a linearly stable fixed point. In summary, we have demonstrated that if $a_{kl} = 1$, then there is a connection from a δ-neighbourhood of ξ_k to ξ_l. Moreover, as the equilibria are linearly stable and having a connection is an open condition, the realisation will persist for an open set of parameters.

B.2 Absence of excitable connections for edges absent from G

Now, suppose that $a_{lk} = a_{kl} = 0$. Then, the dynamics for x_k and x_l is shown schematically in Fig. 18. Equilibria are shown with dots, and all equilibria shown in this figure are linearly stable. Note that the equilibrium ξ_k has $y_k = Y_A = 1$, and $y_l = Y_D = 0$. It is clear that there are no small perturbations which allow for a connection between ξ_k and ξ_l.

For $\theta = 1/2$, if $\delta < 1/4$, then for any k and $j \neq k$ such that $a_{kj} = 0$, all trajectories starting in

$$\zeta_{k,l} = \xi_k + a e_k + b e_l,$$

with $a^2 + b^2 \leq \delta^2$. In particular, perturbations of the form (32) will return to ξ_k. This is because this set remains in the region of validity for the equivalent of (30) for which the only attractor is ξ_k. In the case, $a_{kl} = 0$ and $a_{lk} = 1$ a similar argument holds as the phase portrait corresponds to Fig. 17 reflected in the diagonal.

This shows that no perturbations of amplitude δ within the (x_k, x_l) plane that give a connection from ξ_k to ξ_l of
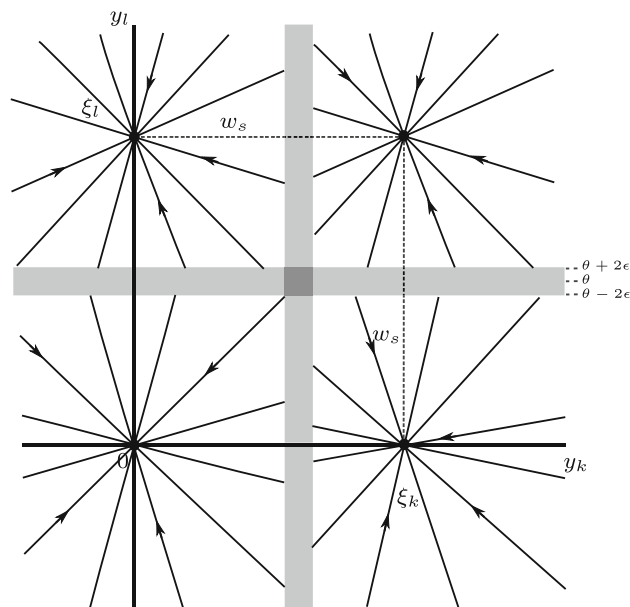

Fig. 18 Schematic diagram showing the dynamics in the y_k-y_l plane when $a_{kl} = a_{lk} = 0$. Equilibria are shown with dots, and all equilibria are linearly stable

amplitude δ. However, we cannot rule out the existence of connections from other locations in $B_\delta(\xi_k)$ to ξ_l. This would be needed to prove that the realisation is almost complete.

References

Afraimovich V, Zhigulin V, Rabinovich M (2004) On the origin of reproducible sequential activity in neural circuits. Chaos Interdiscip J Nonlinear Sci 14(4):1123–1129

Afraimovich VS, Rabinovich MI, Varona P (2004) Heteroclinic contours in neural ensembles and the winnerless competition principle. Int J Bifurc Chaos 14(04):1195–1208

Ashwin P, Castro SB, Lohse A (2020) Almost complete and equable heteroclinic networks. J Nonlinear Sci 30(1):1–22

Ashwin P, Postlethwaite C (2016) Designing heteroclinic and excitable networks in phase space using two populations of coupled cells. J Nonlinear Sci 26(2):345–364. https://doi.org/10.1007/s00332-015-9277-2

Ashwin P, Postlethwaite C (2018) Sensitive finite-state computations using a distributed network with a noisy network attractor. IEEE Trans Neural Netw Learn Syst 29(12):5847–5858

Beer RD (1995) On the dynamics of small continuous-time recurrent neural networks. Adapt Behav 3(4):469–509

Bhowmik D, Nikiforou K, Shanahan M, Maniadakis M, Trahanias P (2016) A reservoir computing model of episodic memory. In: 2016 international joint conference on neural networks (IJCNN), pp 5202–5209. IEEE

Blynel J, Floreano D (2003) Exploring the T-maze: evolving learning-like robot behaviors using CTRNNs. In: Applications of evolutionary computing, pp 593–604. Springer, Berlin

Busse FM, Heikes KE (1980) Convection in a rotating layer: a simple case of turbulence. Science 208:173–175

Ceni A, Ashwin P, Livi L (2020) Interpreting recurrent neural networks behaviour via excitable network attractors. Cogn Comput 12(2):330–356

Chow CC, Karimipanah Y (2020) Before and beyond the Wilson-Cowan equations. J Neurophysiol 123(5):1645–1656. https://doi.org/10.1152/jn.00404.2019

Doedel EJ, Champneys AR, Dercole F, Fairgrieve TF, Kuznetsov YA, Oldeman B, Paffenroth R, Sandstede B, Wang X, Zhang C (2007) AUTO-07P: continuation and bifurcation software for ordinary differential equations

Doedel EJ, Champneys AR, Fairgrieve TF, Kuznetsov YA, Sandstede B, Wang XJ (1997) AUTO97: continuation and bifurcation software for ordinary differential equations. Technical report, Department of Computer Science, Concordia University, Montreal, Canada. Available by FTP from ftp.cs.concordia.ca in directory pub/doedel/auto

Funahashi K, Nakamura Y (1993) Approximation of dynamical systems by continuous time recurrent neural networks. Neural Netw 6:801–806

Gouzé JL, Sari T (2003) A class of piecewise linear differential equations arising in biological models. Dyn Syst 17:299–316

Guckenheimer J, Holmes P (1988) Structurally stable heteroclinic cycles. Math Proc Camb Phil Soc 103:189–192

Harris J, Ermentrout B (2015) Bifurcations in the Wilson-Cowan equations with nonsmooth firing rate. SIAM J Appl Dyn Syst 14(1):43–72

Hopfield JJ, Tank DW (1985) Neural computation of decisions in optimization problems. Biol Cybern 52(3):141–152

Hutt A, Beim GP (2017) Sequences by metastable attractors: interweaving dynamical systems and experimental data. Front Appl Math Stat 3:11. https://doi.org/10.3389/fams.2017.00011

Kirk V, Silber M (1994) A competition between heteroclinic cycles. Nonlinearity 7:1605–1621

Manjunath G, Tiño P, Jaeger H (2012) Theory of input driven dynamical systems. In: ESANN 2012 proceedings

May R, Leonard W (1975) Nonlinear aspects of competition between three species. SIAM J Appl Math 29:243–253

Nikiforou K (2019) The dynamics of continuous-time recurrent neural networks and their relevance to episodic memory. Ph.D. thesis, Imperial College, London

Rabinovich M, Volkovskii A, Lecanda P, Huerta R, Abarbanel H, Laurent G (2001) Dynamical encoding by networks of competing neuron groups: winnerless competition. Phys Rev Lett 87(6):068102

Rabinovich MI, Huerta R, Varona P, Afraimovich VS (2006) Generation and reshaping of sequences in neural systems. Biol Cybern 95(6):519–536

Rabinovich MI, Zaks MA, Varona P (2020) Sequential dynamics of complex networks in mind: consciousness and creativity. Physics Reports

Strogatz SH (1994) Nonlinear dynamics and chaos: With applications to physics, biology, chemistry, and engineering. Addison-Wesley, London

Tuci E, Quinn M, Harvey I (2002) An evolutionary ecological approach to the study of learning behavior using a robot-based model. Adapt Behav 10(3–4):201–221

Wilson HR, Cowan JD (1972) Excitatory and inhibitory interactions in localized populations of model neurons. Biophys J 12(1):1–24

Yamauchi BM, Beer RD (1994) Sequential behavior and learning in evolved dynamical neural networks. Adapt Behav 2(3):219–246