

On designing heteroclinic networks from graphs

Peter Ashwin
University of Exeter,
Exeter EX4 4QF, UK

Claire Postlethwaite
University of Auckland,
Auckland, New Zealand

February 4, 2014

Abstract

Robust heteroclinic networks are invariant sets that can appear as attractors in symmetrically coupled or otherwise constrained dynamical systems. These networks may have a complicated structure determined to a large extent by the constraints and dimension of the system. As these networks are of great interest as dynamical models of biological and cognitive processes, it is useful to understand how particular directed graphs can be realised as attracting robust heteroclinic networks between states in phase space. This paper presents two methods of realizing arbitrarily complex directed graphs as robust heteroclinic networks for flows generated by ODEs—we say the ODEs *realise* the graphs as heteroclinic networks between equilibria that represent the vertices. Suppose we have a directed graph on n_v vertices with n_e edges. The “simplex realisation” embeds the graph as an invariant set of a flow on an $(n_v - 1)$ -simplex. This method realises the graph as long as it is one- and two-cycle free. The “cylinder realisation” embeds a graph as an invariant set of a flow on a $(n_e + 1)$ -dimensional space. This method realises the graph as long as it is one-cycle free. In both cases we realise the graph as an invariant set within an attractor, and discuss some illustrative examples, including the influence of noise and parameters on the dynamics. In particular we show that the resulting heteroclinic network may or may not display “memory” of the vertices visited.

1 Introduction

Heteroclinic cycle or network attractors are an interesting nontrivial form of invariant set for nonlinear dynamics where a typical trajectory recurrently approaches a number of different unstable (saddle) equilibrium states. They have been found to appear robustly in various applications [25, 27, 19, 16] and have been investigated by several groups of researchers as a way to model and understand a number of types of collective neural dynamics. In particular, the networks seem to be able to model a diverse range of systems, from (deterministic or random) sequence generation via “winnerless competition” [32] and finite-state computation [35, 5], to aspects of neural function [31, 13, 24] and, for example, binocular rivalry [7].

Close dynamical relatives are found in models of phase-resetting oscillators where “unstable attractor networks” are a limiting case of heteroclinic networks that can appear in non-invertible semiflows [11, 14, 28].

More precisely, we consider a ordinary differential equation (ODE)

$$\dot{x} = f(x)$$

with smooth $f(x)$ that defines a flow on $x \in \mathbb{R}^d$. We say an closed invariant set $X \subset \mathbb{R}^d$ is a *heteroclinic network* for this flow if it is a union of a finite set of equilibria and connecting orbits between these equilibria, such that the invariant set is chain transitive (i.e. it is a depth one heteroclinic network in the terminology of [4]). The set X is a *robust heteroclinic network* if (subject to specified constraints such as symmetries and smoothness) there is an open set of perturbations of f that have a nearby heteroclinic network that is homeomorphic to the original. We say a heteroclinic network is a *heteroclinic attractor* if it is an attractor in some sense - this could mean, for example, asymptotic stability or Milnor attraction. In the presence of additive noise, we expect trajectories starting close to such a *noisy heteroclinic attractor* to remain close to the network for long periods of time.

Until now, two of the main barriers to using heteroclinic networks to model specific neural processes are firstly, that coupled systems can possess quite complicated robust heteroclinic networks, and secondly, it can be highly nontrivial to determine whether a given coupled cell system has a heteroclinic network. Even when one can find a heteroclinic network, it may have very high complexity for only a moderate number of coupled cells. For example, heteroclinic networks with the structure of “odd graphs” [8] can be found in systems of $N = 2k + 1$ oscillators, and heteroclinic ratchets can be found in less symmetric coupled systems [20]. Up to now there seems to be no way to easily design a coupled dynamical system that realises a *given* directed network as a robust heteroclinic attractor; the closest approaches we are aware of are [16, 15, 30, 1]. The last of these specifically considers the design of heteroclinic networks using small numbers of coupled cells by a process of “inflation” of cells in a smaller network but it is not clear how to use this to design an arbitrary heteroclinic network.

The main contribution of this paper is to give two explicit methods for design of coupled cell systems that realise a given directed graph as a heteroclinic network. In doing so, we suggest possible robust ways to embed a finite-state discrete state computational system into a system of autonomous coupled dynamical systems. The method in Section 2.1 (which we call the simplex realisation) realises any graph (that is one- and two-cycle free) as a heteroclinic network for a cubic polynomial vector field on a simplex. The method in Section 2.2 (which we call the cylinder realisation) realises any graph that is one-cycle free based around a vector field on a line coupled to a number of “transition modes”. We do this by constructing a “coupled cell system” that in both cases is “all-to-all” coupled with a small number of cell types, but where the coupling is inhomogeneous. Figure 1 illustrates the realisation methods.

We analyse both of these realization methods and illustrate their application to some example graphs in Section 3. In all but the simplest cases the embedded network will be

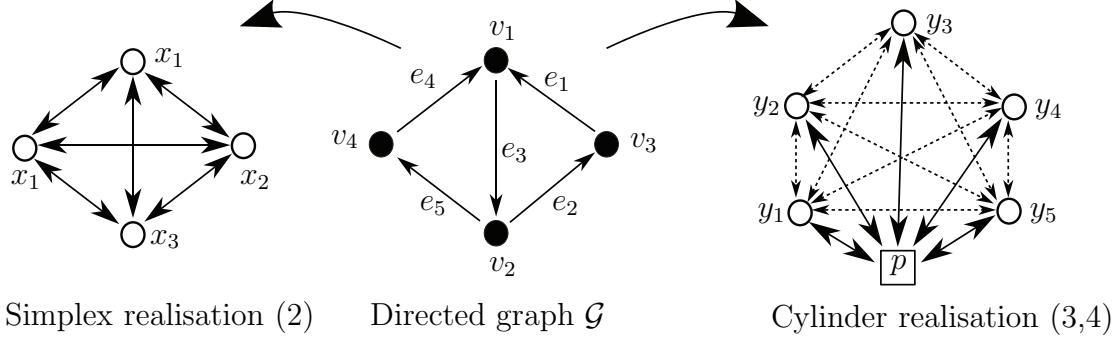


Figure 1: The two realisation methods described in this paper for a coupled cell network to realise a specific directed graph with n_v vertices and n_e edges as a heteroclinic network. The “simplex” realisation described in Section 2.1 uses n_v identical cells x_i with an inhomogeneous coupling while the “cylinder” realisation described in Section 2.2 has a special cell (p) that remembers the last vertex visited while the other n_e cells (y_i) are identical and have inhomogeneous coupling to provide dynamical connections between the states. The dashed lines indicate that the nature of coupling between the y_i is of a different nature to that between that y_i and p . The names are due to invariant structures in phase space rather than any feature of their coupling.

part of a larger chain-recurrent set that includes additional “induced vertices” and separatrices between different connection sets; nonetheless we find in numerical simulations that these induced vertices do not seem to influence the the behaviour of the network, although we have no proof of this. In Section 4 we consider long-term statistics of the behaviour near the network under the influence of noise fluctuations. We see that the transition behaviour between vertices is in certain cases well-modelled as a Markov process with transition depending only on the current vertex. The presence of “lift-off” [3] of trajectories can lead to short-term memory—i.e. cases where the transition probability depends not only on the current vertex but also on the past few vertices visited; we lose the Markov property. This is observed in numerical simulations and discussed in Section 5. Finally, Section 6 finishes with a discussion of issues around the proposed methods of realization of heteroclinic cycles, and we give general thoughts on how these constructions may be related to applications.

2 Design of heteroclinic networks realising a given directed graph

Suppose that $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a directed graph between a finite set of vertices $\mathcal{V} = \{\mathbf{v}_i\}_{i=1}^{n_v}$ with directed edges $\mathcal{E} = \{\mathbf{e}_i\}_{i=1}^{n_e}$. Let $\alpha(\mathbf{e}_i)$ denote the starting vertex and $\omega(\mathbf{e}_i)$ the finishing

vertex of an edge \mathbf{e}_i . We will write A_{jk} to be the adjacency matrix for the graph, i.e.

$$A_{jk} = \begin{cases} 1 & \text{if there is an } \ell \text{ with } \mathbf{v}_j = \alpha(\mathbf{e}_\ell), \mathbf{v}_k = \omega(\mathbf{e}_\ell) \\ 0 & \text{otherwise.} \end{cases}$$

Now consider the noise-forced vector field on \mathbb{R}^{nd} describing a set of coupled cells

$$\frac{dx_i}{dt} = f_i(x_1, \dots, x_n) + \zeta w_i(t) \quad (1)$$

where $x_i \in \mathbb{R}^d$ for $i = 1, \dots, n$ and f_i are smooth functions of their arguments. The terms w_i represent additive i.i.d. noise processes and $\zeta \geq 0$ the noise amplitude. We analyse this system by considering the noise-free ($\zeta = 0$) dynamics and inferring from this the low noise behaviour $\zeta \ll 1$, as well as through numerical simulations for $0 < \zeta$. We say the flow induced by (1) *realises* the graph \mathcal{G} as a heteroclinic network X if X is an embedding of \mathcal{G} such that each vertex \mathbf{v}_i is mapped onto an equilibrium for the flow (contained in X), and each edge \mathbf{e}_i is mapped onto (possibly one of a set of) heteroclinic connections from $\alpha(\mathbf{e}_i)$ to $\omega(\mathbf{e}_i)$, also contained in X .

It is well known that a realisation may be robust if the form of f is constrained—see for example [16, 25, 17, 6]. The network may also be attracting if certain eigenvalue conditions are satisfied. Note that (a) robustness implies necessarily that all equilibria will be hyperbolic and (b) both of our constructions give robust realisations in that there is a symmetry group generated by order two elements, and the constructions are robust to perturbations that respect this symmetry group.

In the remainder of this section, we present two methods for creating heteroclinic networks in a coupled cell network with $d = 1$ dimensional dynamics within each cell; the coupling structures are sketched in Figure 1. The first method is a cubic vector field similar to that of [15, 16] where the vertices of the graph are embedded as vertices of an $(n_v - 1)$ -simplex (note that an n -simplex has $n + 1$ vertices)—we call this the “simplex realisation”. Many examples of heteroclinic cycles and networks in the literature are of this type, especially the “winnerless competition/Lotka–Volterra” type models. Our first construction is also related to the edge and face cycle constructions given in [16, Chapter 7]. Other examples related to this first method are in [15, 18, 21, 30].

The second method, inspired by [7, Figure 11] has all vertices embedded on an invariant line on the centreline of an $n_e + 1$ dimensional cylinder—we call this the “cylinder realisation”. This method is also similar to the construction method in [1]. We will be interested in the long time (asymptotic) behaviour of the system in the presence of noise and inhomogeneities, where the heteroclinic network will not be an exact solution but it gives a clear framework in which one can discuss the dynamics of nearby solutions.

2.1 Simplex realisation

The first construction proceeds as follows. Suppose that \mathcal{G} is a directed graph with n_v vertices and consider the stochastically forced vector field on $x \in \mathbb{R}^{n_v}$ for the *simplex realisation* as

follows:

$$\frac{dx_j}{dt} = x_j \left(1 - |x|^2 + \sum_{i=1}^n a_{ij} x_i^2 \right) + \zeta w_j(t) \quad (2)$$

with $x_j \in \mathbb{R}$ for $j = 1, \dots, n_v$ and $|x|^2 = \sum_j x_j^2$. The w_j represent i.i.d. white noise modulated by $0 \leq \zeta$. In this section, we discuss the existence of the heteroclinic network X in the noise free case $\zeta = 0$. Note that in the absence of noise ($\zeta = 0$), the system has $\mathbb{Z}_2^{n_v}$ symmetry, where each of the \mathbb{Z}_2 subgroups is generated by a reflection in the x_j th coordinate. This symmetry implies that each coordinate hyperplane is invariant under the flow (for compactness, we use the convention that a set is said to be invariant if it is invariant for the noise-free system). Note that systems such as (2) are known to support a wide range of heteroclinic attractors [16].

We say a graph \mathcal{G} is *n-cycle free* if it contains no directed loops of length n . In the following proposition we will consider graphs where there are no edges \mathbf{e}_i with $\alpha(\mathbf{e}_i) = \omega(\mathbf{e}_i)$ (no one-cycles) and no pairs of edges $\mathbf{e}_{i,j}$ with $\alpha(\mathbf{e}_i) = \omega(\mathbf{e}_j)$ and $\alpha(\mathbf{e}_j) = \omega(\mathbf{e}_i)$ (no two-cycles). Note that \mathcal{G} is *n-cycle free* if and only if all diagonal entries of A^n are zero, where A is the adjacency matrix.

Proposition 1 *Suppose that \mathcal{G} is one-cycle and two-cycle free. Then a_{ij} can be chosen so that the system (2) for $\zeta = 0$ realises the graph \mathcal{G} as a heteroclinic network X . Moreover, this realisation is robust to perturbations that respect $\mathbb{Z}_2^{n_v}$ symmetry given by reflection in the coordinate planes.¹*

Proof: The system (2) has equilibria on each of the n_v coordinate axes,

$$\xi_k = (0, \dots, 0, 1, 0, \dots, 0)$$

where the 1 is in the k th position and $k = 1, \dots, n_v$. We associate ξ_k with the vertex \mathbf{v}_k of the graph \mathcal{G} . The eigenvalues of ξ_k are $-2 + 3a_{kk}$ in the radial (x_k) direction, and a_{kj} in the x_j direction ($j \neq k$). We choose the coefficients a_{ij} so that $a_{jj} = 0$, and

$$a_{ij} > 0 \text{ if } A_{ij} = 1 \text{ while } a_{ij} < 0 \text{ if } A_{ij} = 0, \quad (i \neq j).$$

If the graph \mathcal{G} contains an edge (\mathbf{e}_ℓ) with $\alpha(\mathbf{e}_\ell) = \mathbf{v}_i$ and $\omega(\mathbf{e}_\ell) = \mathbf{v}_j$ (and hence does not contain an edge from \mathbf{v}_j to \mathbf{v}_i), then $A_{ij} = 1$, hence $a_{ij} > 0$ and so the equilibrium ξ_i will be unstable in the x_j direction. Similarly, $A_{ji} = 0$, thus $a_{ji} < 0$ and so ξ_j will be stable in the x_i direction. Restricted to the two-dimensional subspace (x_i, x_j) , the resulting flow and connecting heteroclinic orbit are shown in Figure 2(a). The heteroclinic network X consists of the union of the equilibria ξ_k and the (one-dimensional) connections from ξ_i to ξ_j .

Because the network is one- and two-cycle free we have $A_{ij}A_{ji} = 0$ for all i, j . If this condition were broken, for instance if $A_{ij} = A_{ji} = 1$, then ξ_i would be unstable to perturbations in the x_j direction, and ξ_j would be unstable to perturbations in the x_i direction.

¹In particular, the robustness to perturbations implies that the construction is robust to choice of parameter values.

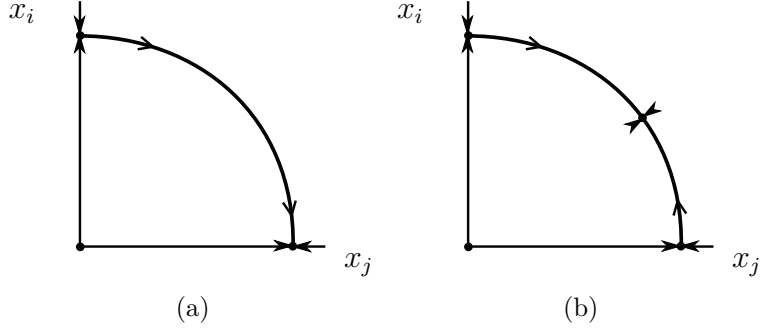


Figure 2: (a) Schematic diagram showing a connecting orbit between equilibria ξ_i and ξ_j in the (x_i, x_j) -plane for the simplex realisation. (b) The lack of connection in this invariant subspace if $a_{ij} > 0$ and $a_{ji} > 0$; there is an additional equilibrium that is a sink within the subspace. Equilibria are shown with dots; the asymptotic dynamics for the full system is on topological $(n_v - 1)$ -simplex within an attracting invariant sphere.

This would result in a flow as shown in Figure 2(b), and a new (stable) equilibrium would be created in the (x_i, x_j) plane. **QED**

This system also has a number of other equilibria; for example, the equilibrium at the origin, which is unstable (all eigenvalues are equal to 1). The coefficients a_{ij} that are non-zero can be thought of as weights on the graph; the magnitude of $a_{ij} > 0$ affects the expanding eigenvalue (in the x_j direction) at ξ_i while $a_{ji} < 0$ affects the contracting eigenvalue (in the x_i direction) at ξ_j . In later sections we will parametrize the dependence by making the a_{ij} a function of a matrix of weights $w_{ij} \geq 0$.

We conjecture that for an open set of parameter values the constructed heteroclinic network X is a subset of a larger network \tilde{X} that is asymptotically stable. In addition, we conjecture that in the limit of small noise $\zeta \rightarrow 0$, the only equilibria which typical trajectories will spend a non-zero proportion of time near will be the equilibria in X , and the only connections that will appear in non-zero frequency correspond to connections in X . Our numerical simulations in the later sections support these conjectures. Note that the constructed heteroclinic network X from Proposition 1 cannot be an asymptotically stable heteroclinic attractor if the network is more than a cycle - this is because there will be at least one vertex with an unstable manifold that is at least two dimensional. This also means that there are typically many more connecting orbits from ξ_i to ξ_j than the one in X ; this larger network \tilde{X} will contain all unstable manifolds of points in the network and hence will include not only the “edge cycles” that represent the vertices of the graph, but also “face cycles” (equilibria where two or more components are non-zero) in the terminology of [16].

Whether this larger network \tilde{X} is attracting will depend on details of the chosen parameters. Computation of the stability properties of a heteroclinic network is quite involved, see for example [22, 23] or [16] for cycles. We do not compute conditions for asymptotic

stability of the embedding heteroclinic network \tilde{X} . However, by analogy with properties of heteroclinic cycles, we conjecture that a sufficient (but not necessary) condition for asymptotic stability of the heteroclinic network is that all the contracting eigenvalues are greater in absolute value than all the expanding eigenvalues, and all transverse and radial eigenvalues are negative. Here by contracting (resp. expanding) eigenvalue we mean a negative (resp. positive) eigenvalue in a direction associated with an incoming (resp. outgoing) heteroclinic connection to an equilibrium (see [26] for cycles). This condition can presumably be achieved in this example by choosing the a_{ij} appropriately. The computation of detailed stability conditions for the network is beyond the scope of this paper, but we find in our numerical examples that the noisy heteroclinic networks can be observed to be attracting.

2.2 Cylinder network

The second construction also proceeds directly from the graph structure. Suppose that \mathcal{G} is a directed graph with n_e edges; let us define the stochastically forced vector field for the *cylinder realisation* on (y, p) with $y \in \mathbb{R}^{n_e}$ and $p \in \mathbb{R}$ by

$$\begin{aligned}\frac{dy_j}{dt} &= -y_j G_j(y, p) + \zeta w_j(t) \\ \frac{dp}{dt} &= -\sin(2\pi p) + F(y, p) + \zeta w_0(t)\end{aligned}\tag{3}$$

for $j = 1, \dots, n_e$, where F and G_j are smooth functions and are even in each of the y_j , and $F(0, \dots, 0, p) = 0$ for any p . As before, the w_j are i.i.d. white noise processes that we modulate with an amplitude $0 \leq \zeta$. Again, we analyse the system for $\zeta = 0$ while in later sections we consider numerical simulations with $0 < \zeta$.

This system with $\zeta = 0$ has an invariant line $y_1 = y_2 = \dots = y_{n_e} = 0$ parametrized by p , on which there are equilibria at $p = n$, $n \in \mathbb{Z}$, $y_j = 0$. We denote the equilibrium

$$\xi_k = (0, \dots, 0, k)$$

for $k = 1, \dots, n_v$. The invariant line is contained in each of the invariant planes $\ell = 1, \dots, n_e$:

$$P_\ell := \{(y, p) : y_j = 0 \text{ if } j \neq \ell\}.$$

We say the variable y_ℓ is activated (i.e. it is non-zero and all other y_j remain zero) in P_ℓ , on the connection from

$$\alpha(\mathbf{e}_\ell) = (0, \dots, 0, p_{\alpha\ell})$$

to

$$\omega(\mathbf{e}_\ell) = (0, \dots, 0, p_{\omega\ell})$$

where $p_{\alpha\ell}, p_{\omega\ell}$ are defined appropriately. The following result holds for any graph as long as there are no edges \mathbf{e}_i with $\alpha(\mathbf{e}_i) = \omega(\mathbf{e}_i)$ —it constructs a network that lies in a cylindrical neighbourhood of the invariant line—hence the name of the network.

Proposition 2 *Suppose that \mathcal{G} is one-cycle free. Then there are smooth functions F and G_j such that the system (3) for $\zeta = 0$ realises the graph \mathcal{G} as a heteroclinic network. Moreover, this realisation is robust to perturbations that respect $\mathbb{Z}_2^{n_e}$ symmetry given by $y_j \mapsto -y_j$.*

Proof: We consider the following specific choices for F and G_j :

$$\begin{aligned} F(y, p) &= \sum_{j=1}^{n_e} y_j^2 \tanh(p_{\omega j} - p), \\ G_j(y, p) &= \left[\left(y_j^2 - \frac{5}{4} \right)^2 - 1 - f_{\alpha j}(p) + f_{\omega j}(p) + K_i \sum_{i \neq j} y_i^2 \right], \\ f_{\alpha j}(p) &= L_{\alpha j} \operatorname{sech}^2(k_{\alpha j}(p - p_{\alpha j})), \\ f_{\omega j}(p) &= L_{\omega j} \operatorname{sech}^2(k_{\omega j}(p - p_{\omega j})). \end{aligned} \quad (4)$$

The quantities $L_{\alpha j}, L_{\omega j}, k_{\alpha j}, k_{\omega j}, K_i > 0$ are parameters that will be chosen appropriately so that there is a robust heteroclinic connection in P_ℓ corresponding to the edge \mathbf{e}_ℓ . Note that K_i is a ‘mutual inhibition’ parameter between the various y_j .

If $L_{\alpha \ell} = L_{\omega \ell} = 0$ then, within the plane P_ℓ , the system has y_ℓ -nullclines at $y_\ell = 0$, $y_\ell = \pm 1/2$ and $y_\ell = \pm 3/2$. The additional terms when $L_{\alpha \ell}, L_{\omega \ell} \neq 0$ make $\alpha(\mathbf{e}_\ell)$ unstable in the y_ℓ direction near $\alpha(\mathbf{e}_\ell)$. Constraints on the parameters can be understood by considering the geometry of the heteroclinic connection from $\alpha(\mathbf{e}_\ell)$ to $\omega(\mathbf{e}_\ell)$, which lies in P_ℓ .

- Firstly, we require that $G_\ell > 0$ (so that $\dot{y}_\ell/y_\ell < 0$) for all $y_\ell \neq 0$ at $p \approx p_{\omega \ell}$, which is achieved if

$$\left(y^2 - \frac{5}{4} \right)^2 - 1 - f_{\alpha \ell}(p_{\omega \ell}) + L_{\omega \ell} > 0, \quad \forall y$$

which is satisfied if $L_{\omega \ell} > 1$ and $f_{\alpha \ell}(p_{\omega \ell})$ is sufficiently small.

- Secondly, we require that $G_\ell < 0$ (so that $\dot{y}_\ell/y_\ell < 0$) when $p \approx p_{\alpha \ell}$ and $0 < y_\ell < y_u$ for some $y_u > 1/2$. This is achieved if

$$\left(y^2 - \frac{5}{4} \right)^2 - 1 - L_{\alpha \ell} + f_{\omega \ell}(p_{\alpha \ell}) < 0, \quad 0 < y < y_u$$

which is satisfied if $L_{\alpha \ell} > 9/16$ and $f_{\omega \ell}(p_{\alpha \ell})$ is sufficiently small.

Each of the equilibria ξ_k has eigenvalues as follows: firstly, all equilibria have a ‘radial’ eigenvalue in the p direction which we label $r_k = -2\pi$. Now, suppose there is an edge \mathbf{e}_i with $\alpha(\mathbf{e}_i) = \mathbf{v}_k$, then the eigenvalue at ξ_k in the y_i direction is

$$e_{ki} = -\frac{9}{16} + L_{\alpha i} - f_{\omega i}(k).$$

Similarly, if there is an edge \mathbf{e}_j with $\omega(\mathbf{e}_j) = \mathbf{v}_k$, then the eigenvalue at ξ_k in the y_j direction is

$$c_{kj} = -\frac{9}{16} + f_{\alpha j}(k) - L_{\omega j}.$$

If the edge \mathbf{e}_ℓ neither starts nor ends at \mathbf{v}_k (i.e. $\alpha(\mathbf{e}_\ell) \neq \mathbf{v}_k$ and $\omega(\mathbf{e}_\ell) \neq \mathbf{v}_k$) then the eigenvalue at ξ_k in the y_ℓ direction will be

$$t_{k\ell} = -\frac{9}{16} + f_{\alpha\ell}(k) - f_{\omega\ell}(k).$$

Sufficient conditions for the existence of the desired connections are that $e_{ki} > 0$, for all i for which $\alpha(\mathbf{e}_i) = \mathbf{v}_k$ (which gives the condition $L_{\alpha i} > 9/16$, as before) and $c_{kj} < 0$, for all j for which $\omega(\mathbf{e}_j) = \mathbf{v}_k$.

Hence we can choose parameters $L_{\alpha j}$, $L_{\omega j}$, $k_{\alpha j}$ and $k_{\omega j}$ so that $f_{\alpha j}(p)$ and $f_{\omega j}(p)$ are small enough when p is far from $p_{\alpha j}$ and $p_{\omega j}$ respectively. Alternatively, we can set the values of the contracting and expanding eigenvalues (subject to some constraints; see below), and then choose parameters $L_{\alpha j}$, $L_{\omega j}$, $k_{\alpha j}$ and $k_{\omega j}$ so that the network has this set of eigenvalues.

QED

Figure 3 shows the structure of the connections in phase space using this construction. One implication of the requirement that $L_{\omega j} > 1$ is that for these parameters, the contracting eigenvalues will always be less than $-25/16$. The final term in the \dot{y}_j equations, $K_i \sum_{i \neq j} y_i^2$, is a ‘‘mutual inhibition’’ term that is included so that additional stable equilibria away from the coordinate planes are suppressed. We will typically choose $K_i = 1$.

Similarly to the simplex construction, we conjecture that parameters can be chosen so that the heteroclinic network X that realises \mathcal{G} is embedded within a larger network \tilde{X} ; this larger network will not be an asymptotically stable attractor unless care is taken to ensure that the expanding eigenvalues are weak enough. More precisely, we conjecture that for an open set of parameters the heteroclinic network X is embedded within a larger asymptotically stable network \tilde{X} , and in the limit of small noise $\zeta \rightarrow 0$ the only equilibria visited with non-zero proportion of time are those in X and the only connections visited with non-zero frequency are those in X . Note that a necessary (but not sufficient) condition for asymptotic stability of the network X is that $t_{k\ell} < 0$.

Within the two-dimensional subspace P_ℓ parametrized by (p, y_ℓ) we show nullclines of the system with typical parameter values in Figure 4, in the case of $p_{\alpha\ell} = 1$ and $p_{\omega\ell} = 4$. The blue lines show the p nullclines while the red lines show the y_ℓ nullclines. Note that within the subspace $y_j = 0$, the equilibria ξ_k ($k \neq 1$) are sinks. Also, note that there are additional equilibria at $p = k + \frac{1}{2}$, which are unstable in this one-dimensional subspace, while trajectories which start near ξ_1 move towards ξ_4 .

3 Examples of networks realised as noisy heteroclinic attractors

We now consider two illustrative examples; the first (the decision graph) uses the simplex realisation and Proposition 1 while the second (the Petersen graph) uses the cylinder realisation and Proposition 2. The decision graph could be realised using the cylinder realisation,

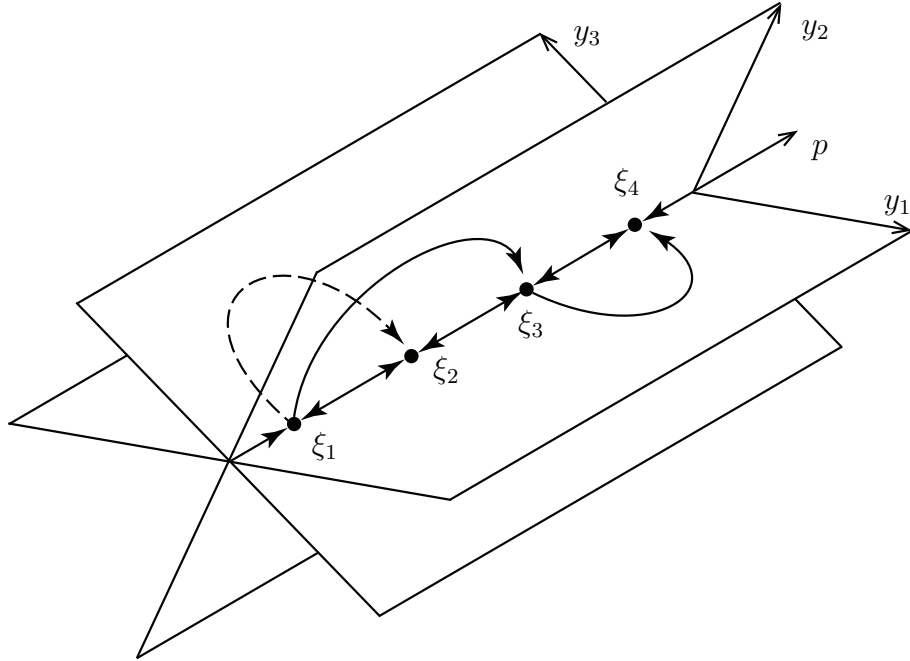


Figure 3: Schematic diagram showing the structure of part of a cylinder realisation: the vertices line up along the p -axis while the (y_j, p) planes contain connections between vertices. In this case there are connections $\xi_3 \rightarrow \xi_4$ in P_1 , $\xi_1 \rightarrow \xi_3$ in P_2 and $\xi_1 \rightarrow \xi_2$ in P_3 . Note that the planes P_ℓ are all orthogonal.

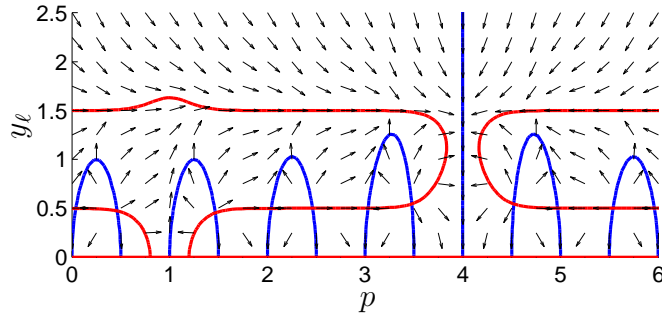


Figure 4: The figure shows the nullclines and vector field for a typical two-dimensional subspace P_ℓ parametrized by (p, y_ℓ) for the cylinder realisation; there is a connection from $p = 1$ to $p = 4$. Arrows in the vector field are all scaled to have the same length for clarity. In this case we have a connection from the saddle $\xi_1 = (1, 0)$ to the sink $\xi_4 = (4, 0)$. Note that there are additional sinks in P_ℓ at $\xi_k = (k, 0)$ for $k = 2, 3, 5$ that are surrounded by bounded basins of attraction.

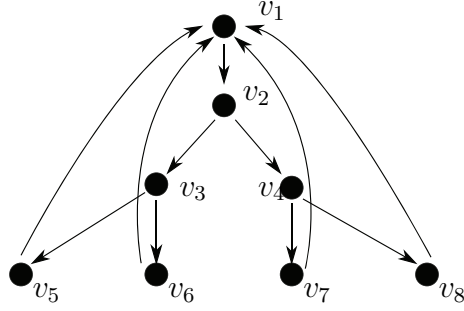


Figure 5: The “decision network” graph for the adjacency matrix corresponding to the weights in (5).

but the Petersen graph cannot be realised using the simplex realisation as it contains cycles of order two. In both cases we consider simulations that include small but non-zero levels of noise.

3.1 Decision graph

Consider the “decision tree with reset” shown in Figure 5, which for convenience we call the “decision graph”. One can think of this graph \mathcal{G} as a two-level binary decision tree followed by a reset back to the first (root) vertex of the tree. This graph has $n_v = 8$ vertices and $n_e = 11$ edges. Its adjacency matrix is $A_{ij} = 1$ for $w_{ij} \neq 0$ and $A_{ij} = 0$ for $w_{ij} = 0$, where

$$\{w_{ij}\} = \begin{bmatrix} 0 & a & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & a & a' & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & a & a' & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & a & a' \\ a & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ a & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ a & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ a & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}. \quad (5)$$

If we set parameters

$$a = 0.99, \quad a' = 0.98, \quad \sigma = 0.2, \quad \mu = 0.6, \quad \zeta = 10^{-4} \quad (6)$$

and define

$$a_{ij} = \begin{cases} -\sigma + \mu w_{ij} & \text{if } i \neq j \\ 0 & \text{otherwise} \end{cases}$$

then it is a routine calculation to check that the requirements of Proposition 1 are satisfied and (2) will realise a noisy heteroclinic network with the structure of the graph shown in Figure 5.

One can visualise the dynamics on the realisation of this graph by projecting the trajectory onto two observables that order the vertices into a ring. More precisely, for a trajectory $x(t)$ we define the complex observable

$$R(t) = \frac{1}{n_v} \sum_{k=1}^n x_k^2 \exp \left[i\pi \frac{2(k-1)}{n_v} \right] \quad (7)$$

so that the vertices of the graph are projected onto the vertices of a regular n_v -gon on the unit circle. Figure 6 shows a projection and timeseries for the dynamics of the system corresponding to the graph in Figure 5.

Proposition 1 means that the graph structure is unaffected by the exact values of the parameters used. However, the statistics of the residence times and the transition probabilities at decision points (including the possibility of memory) are strongly affected by the parameter values. We explore some of these effects in Section 4 and apply our results to investigate this graph in Section 5.

3.2 Petersen graph

As an example with a somewhat different structure (and one that is nontrivially but highly connected), we now consider a cylinder realisation of the Petersen graph (shown in Figure 7) as a noisy heteroclinic network. Each edge can be thought of as a pair of directed edges, one in each direction.

This graph has been studied extensively as one of the simplest examples of a graph with certain nontrivial colouring properties; it has also been found to organize heteroclinic networks in five globally coupled phase oscillators and in systems of five delay pulse-coupled oscillators, where it has been proposed for computational purposes [5, 29].

Figure 8 shows a simulation of a realisation of this using the cylinder realisation (3,4) and adjacency matrix

$$\{A_{ij}\} = \begin{bmatrix} 0 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 \end{bmatrix}. \quad (8)$$

The other parameters used for the simulation in Figure 8 are

$$L_{\alpha j} = 1.4376, \quad L_{\omega j} = 1.5625, \quad k_{\alpha j} = 2.017, \quad k_{\omega j} = 0.4705, \quad K_i = 1, \quad \zeta = 10^{-6}. \quad (9)$$

Note that by varying $k_{\alpha j}$ and $k_{\omega j}$ with j we can vary the expanding and contracting eigenvalues on the j th connection. For simplicity, in the example we set them to be identical.

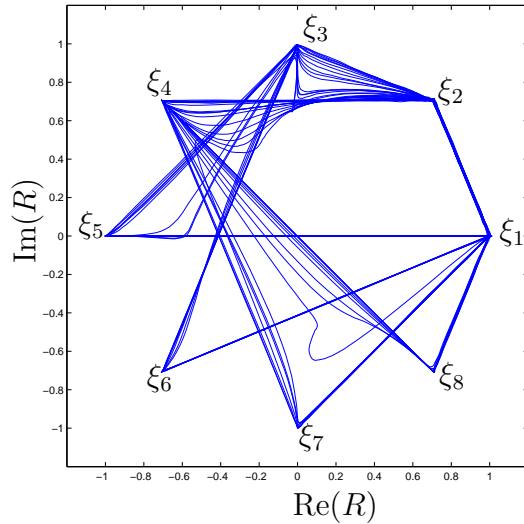
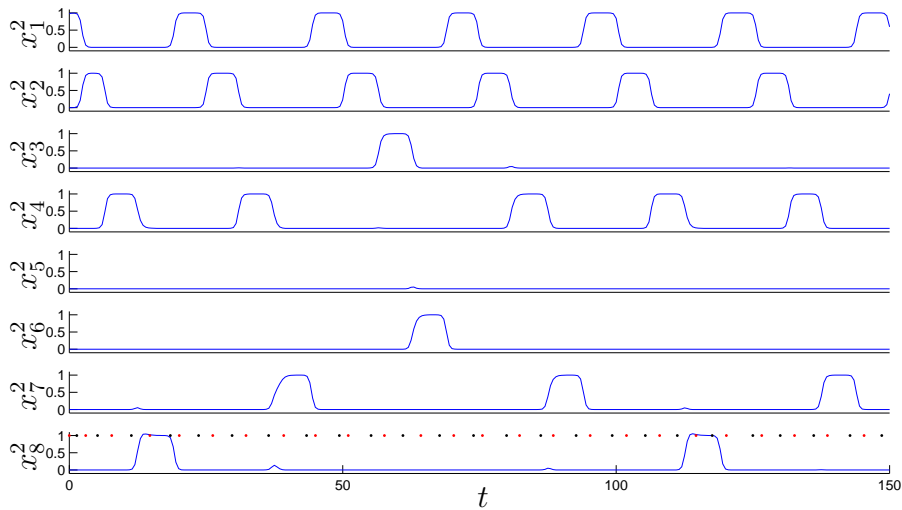


Figure 6: An example simulation of the decision graph (shown in Figure 5) using the simplex realisation (2) and weights (5). Upper: time series of x_i^2 ; the red (resp. black) dots on the lowest frame indicate times where the trajectory enters (resp. leaves) an “epoch” of being close to one of the saddles. Lower: projection onto the real and imaginary parts of order parameter for the system showing the location of the equilibria. Simulations clearly show a noisy heteroclinic attractor with the structure as in Figure 5, projected into the plane using the complex observable $R(t)$ (equation (7)). Note that presence of additional saddles can be inferred from rare excursions away from the one-dimensional connections.

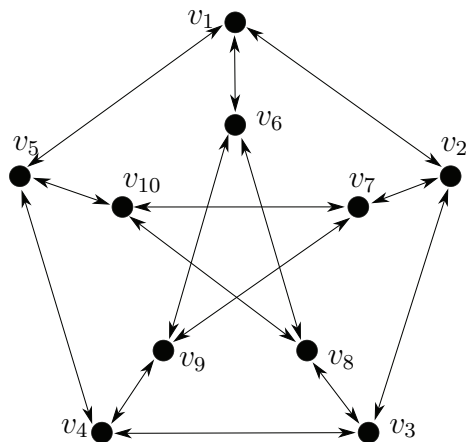


Figure 7: The Petersen graph; this can be viewed as a (one-cycle free) directed graph by considering each edge as a pair of directed edges in both directions.

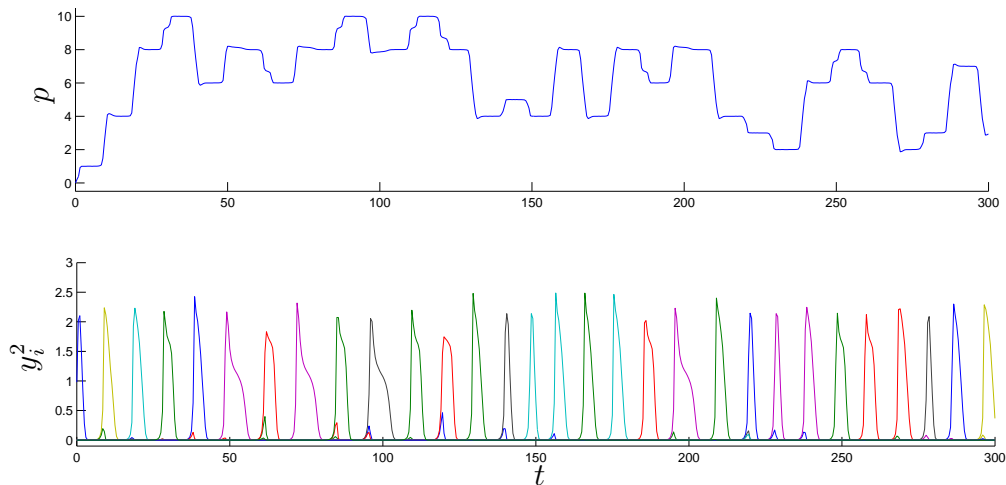


Figure 8: Time series for a cylinder realisation of the Petersen graph (Figure 7). The p -dynamics (top) can be observed to wander around a noisy heteroclinic attractor between the vertices of the Petersen graph, only making transitions corresponding to edges in the graph, while each of the components y_i for $i = 1, \dots, 15$ (bottom) become non-zero only during a transition along the i th edge. The presence of weak noise causes the dynamics to wander around the network. Note that all the saddles corresponding to vertices on the network have three dimensional unstable manifolds.

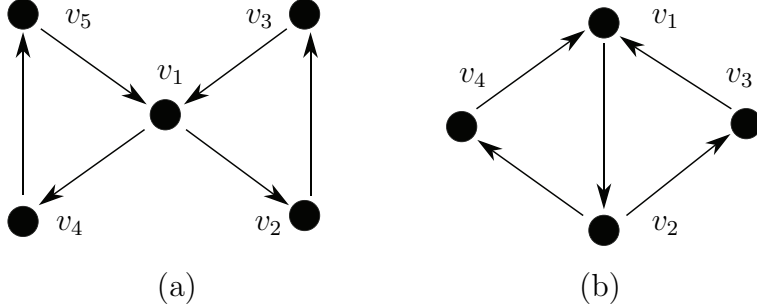


Figure 9: (a) A “Bowtie” graph of two connected cycles and (b) the “Kirk–Silber” [21] network that shows a competition of two length three cycles with a shared edge. Note that in both cases there are two cycles of length three; in (b) there is a shared edge and two shared vertices while in (a) there is only a shared vertex.

3.3 Further examples

We briefly highlight two further graphs that illustrate different ways in which vertices can connect cycles of the same type within a network. Figure 9 illustrates two networks each containing two nontrivial cycles of length three. In each case the network can in principle be realised using either “simplex” or “cylinder” realisations. Note that although the decision graph has alternative routes around the graph, all of these share at least one edge; similarly the cycles on the Kirk–Silber network share an edge, while those in the Bowtie graph do not.

4 Statistical properties of trajectories near a realised network

Propositions 1 and 2 do not give unique ways to realise a given graph \mathcal{G} as a heteroclinic attractor - rather, they give open sets of functions that give the appropriate embedding. These alternative realisations can however display different statistical properties of trajectories near the network in the presence of low noise (or more generally, time-dependent inputs) after transients have decayed. In this section, we adapt and generalise some of the analysis of Stone, Armbruster and colleagues [34, 33, 3] to understand some basic statistical properties of the noise-induced itineraries of trajectories near the network.

We fix on a size $h > 0$ for neighbourhoods of the equilibria and say a solution trajectory $x(t)$ is *close* to an equilibrium ξ_j (corresponding to a vertex of the realised graph) at time \tilde{t} if $|x(\tilde{t}) - \xi_j| < h$. We require that $|\xi_p - \xi_q| > 2h$ for any $p \neq q$ so that at any time, the trajectory is close to at most one equilibrium. We say $x(t)$ *remains close* to ξ_j during $t \in [s, s + T]$ if it satisfies

$$|x(t) - \xi_j| < h \text{ for } t \in (s, s + T), \text{ and } |x(t) - \xi_j| = h \text{ for } t \in \{s, s + T\}.$$

We divide a trajectory starting at some given initial condition \tilde{x} and given noise path \tilde{v} into

an *itinerary*, i.e. the sequence of *epochs*:

$$\{(i_k, s_k, T_k) : k \in \mathbb{N}\} \quad (10)$$

such that $x(t)$ remains close to ξ_{i_k} during $t \in [s_k, s_k + T_k]$. Note that $i_k \in \{1, \dots, n_v\}$ represent indices of the vertices (equilibria), the times of entry $0 < s_k$ are increasing and the durations $T_k > 0$ are all positive; for example, the red and black dots in Figure 5 illustrate the sequences s_k and $s_k + T_k$ respectively, for $h = 0.1$.

One can formally consider the solution of the noise-perturbed system as deterministic in the spirit of random dynamical systems [2], i.e. as a skew product evolution

$$\begin{aligned} x(t) &= \Phi_t(\tilde{x}, \tilde{\vartheta}) \\ \vartheta(t) &= \theta_t(\tilde{\vartheta}) \end{aligned} \quad (11)$$

over the evolution θ_t along the noise path. In the original system $\tilde{\vartheta}$ represents a particular Brownian path whereas in the simulation, $\tilde{\vartheta}$ can be thought of as choice of seed for a random number generator. Note that Φ_t satisfies the cocycle property $\Phi_{t+s}(\tilde{x}, \tilde{\vartheta}) = \Phi_t(\Phi_s(\tilde{x}, \tilde{\vartheta}), \theta_s(\tilde{\vartheta}))$. Then $x(t)$ remains close to ξ_j if $|\Phi_t(\tilde{x}, \tilde{\vartheta}) - \xi_j| < h$, which clearly is determined by initial condition and noise path.

If the trajectory remains close to the attractor then it will have an infinite itinerary except in the (very unlikely) case that it remains close to one of the saddle equilibria for all time. We make a *stochastic stationarity* assumption that the statistical properties of the itinerary of typical initial conditions and typical noise trajectory are stationary and independent of the initial condition and details of the noise trajectory. This means we assume that any transients associated with the initial condition will have decayed and that the initial condition is distributed according to a stationary invariant probability distribution on the space of initial conditions \tilde{x} for the noise-perturbed system, and the noise path $\tilde{\vartheta}$ is “typical”. More precisely, we make an implicit assumption that there is an ergodic probability measure for θ_t that lifts to a natural measure for the skew product flow. With respect to this distribution, we define the probability of observing a given finite sequence of vertices $\{j_k : k = 1, \dots, m\}$ as

$$\mathcal{P}(j_1, \dots, j_m) = \text{Prob}(\tilde{x}, \tilde{\vartheta} : i_\ell = j_\ell \text{ for } \ell = 1, \dots, m, \text{ for the trajectory starting at } (\tilde{x}, \tilde{\vartheta})). \quad (12)$$

The stationarity assumption implies that one can compute this from a typical trajectory simply in terms of the frequency of transitions

$$\mathcal{P}(j_1, \dots, j_m) = \lim_{k \rightarrow \infty} \frac{1}{k} \#\{0 \leq \ell < k : i_{\ell+n} = j_n \text{ for } n = 1, \dots, m\}. \quad (13)$$

This can be used to define the probability π_j of an epoch being close to ξ_j :

$$\pi_j = \mathcal{P}(j)$$

and (assuming $\pi_j > 0$) we define the transition probability between vertices ξ_{j_1} and ξ_{j_2} by:

$$\pi_{j_1, j_2} = \frac{\mathcal{P}(j_1, j_2)}{\pi_{j_1}}. \quad (14)$$

We ask the question: are the statistics of the itineraries simply those of a Markov chain whose non-zero transitions correspond to edges of the original graph \mathcal{G} ? Although this is possible, memory effects can appear (see [12]) where transition probabilities may depend not just on the current equilibrium but also on previous equilibria visited.

We say a noise perturbed attracting heteroclinic network is *memoryless* if the transition probabilities are independent of the prior itinerary, i.e. if the itineraries are Markov of order one. More precisely, a transition from ξ_p to ξ_q is said to be memoryless if

$$\mathcal{P}(j_1, \dots, j_m) = \mathcal{P}(j_1, \dots, j_{m-1})\pi_{p,q} \quad (15)$$

for any $m \geq 2$ and any sequence $\{j_k : k = 1, \dots, m\}$ with $j_{m-1} = p$ and $j_m = q$. The system is memoryless if all possible transitions are memoryless. We expect such a system will be truly memoryless only in the singular limit of low noise $\zeta \rightarrow 0$; nonetheless, later sections suggest that a system can be very close to memoryless in that corrections to (15) may be asymptotically small in ζ . Note also that a particular transition will necessarily be memoryless if $\pi_{p,q} = 1$, though this does not necessarily imply that all other transitions are memoryless.

The transition probabilities will be affected (to a greater or lesser extent) by the eigenvalues at the equilibria and by the noise level. In the following section, we outline the influence of these parameters. Finally, using the simplex realisation of the decision graph, Section 5 gives some numerical examples with and without memory.

4.1 Analysis of dynamics near heteroclinic networks: Poincaré sections and maps

Analysis of the dynamics of trajectories near heteroclinic cycles and networks is often done via construction of Poincaré maps to approximate the flow (see, e.g. [21]). The return maps are an ensemble of maps where each is a composition of two types of map: local maps past neighbourhoods of the equilibria and global map that connect the neighbourhoods. The local maps are constructed assuming no resonances and then linearising the flow near the equilibria while the global maps are truncated Taylor expansions of the maps near the connecting orbits. The lowest order truncations of these maps then approximate the flow close to the heteroclinic network.

The detailed construction of such maps for the networks (even in the noise-free case) is beyond the scope of this paper. However, in later sections we apply previous results regarding the dynamics of trajectories near heteroclinic networks in the presence of noise, so we find it useful to define the appropriate Poincaré sections here. Poincaré sections are often defined to be surfaces a distance h from an equilibrium, where h is some small constant, meaning they are spheres. However, in order to apply results of Stone, Armbruster and others [34, 33, 3] we define Poincaré sections as unions of codimension one surfaces as follows. We give explicit definitions of Poincaré sections for the simplex realisation; for the cylinder realisation they can be defined similarly.

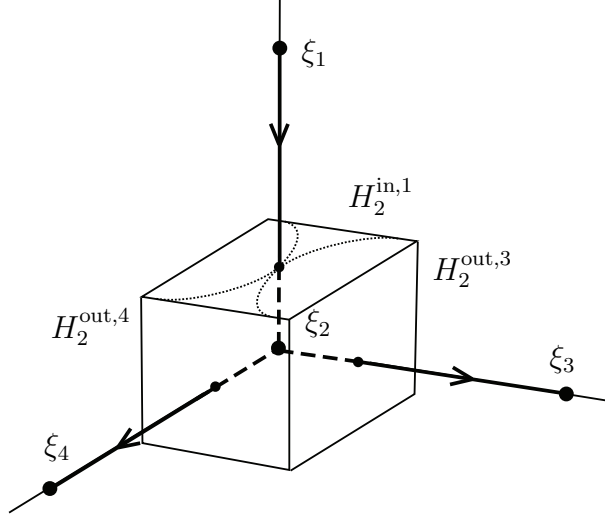


Figure 10: Schematic diagram showing a Poincaré sections $H_2^{\text{in},1}$, $H_2^{\text{out},3}$ and $H_2^{\text{out},4}$ near the equilibrium ξ_2 . Heteroclinic orbits are shown with bold lines. Large dots indicate equilibria, small dots indicate the intersection of the heteroclinic orbits with the Poincaré sections. The dotted lines on H_2^{in} indicate the dividing line between regions of trajectories which will next travel to ξ_3 or ξ_4 , in the case that $e_{23} > e_{24}$.

Consider an equilibrium ξ_k in the simplex realisation, with an incoming heteroclinic connection from equilibria ξ_j , and an outgoing heteroclinic connection towards equilibria ξ_l . We define (for a fixed small h):

$$H_k^{\text{in},j} = \{x \in \mathbb{R}^{n_v} : x_j = h, |1 - x_k| < h, |x_i| < h, i \neq j, k\} \quad (16)$$

$$H_k^{\text{out},l} = \{x \in \mathbb{R}^{n_v} : x_l = h, |1 - x_k| < h, |x_i| < h, i \neq l, k\}. \quad (17)$$

If a vertex in the graph has an incoming or outgoing degree greater than one, then there will be multiple incoming or outgoing Poincaré sections, and so we also define

$$H_k^{\text{in}} = \bigcup_j H_k^{\text{in},j}, \quad H_k^{\text{out}} = \bigcup_l H_k^{\text{out},l}$$

where the unions are taken over all incoming and outgoing directions respectively. Note that if the trajectory spends time near ξ_k it must have passed through H_k^{in} and will pass through H_k^{out} at some time thereafter. Figure 10 shows a schematic diagram of the Poincaré sections near an equilibrium ξ_2 , with an incoming connection from ξ_1 and outgoing connections to ξ_3 and ξ_4 .

In the following sections, we use the following notation for eigenvalues of equilibria. If ξ_k has a contracting direction in the x_j direction, we label the corresponding eigenvalue $-c_{kj}$. Similarly, if ξ_k has an expanding direction in the x_l direction, we label the corresponding eigenvalue e_{kl} .

4.2 Transition probabilities between equilibria

If ξ_j has a single expanding direction towards ξ_k , then clearly $\pi_{j,k} = 1$ and $\pi_{j,\ell} = 0$ for $k \neq \ell$ in the limit of low noise. The more interesting case is when the equilibrium has two expanding directions, as in the example shown schematically in Figure 10. Here, at the equilibrium ξ_2 , the trajectory makes a ‘choice’ as to whether to next visit ξ_3 or ξ_4 . We give here an outline of how to compute the probabilities of the trajectory making each choice; that is, the transition probabilities. The extension of these calculations to equilibria with three or more expanding directions is straightforward.

Armbruster et. al [3] compute the transition probabilities for an equilibrium with two expanding directions, under the assumption that the incoming distribution of coordinates of a trajectory is approximately Gaussian and centered at zero. For the schematic example in Figure 10, we show this distribution of incoming trajectories, in terms of the x_3 and x_4 coordinates on H_2^{in} , on the left-hand side of Figure 11. The proportion of trajectories which we expect to visit ξ_4 is given by the proportion of the measure in the ‘noise ellipse’ which intersects the cusp.

The computation of this area was given in [3]. The results depend on constants which come from the global part of the flow and are in general unknown. However, a scaling can be found in the limit of low noise. It was shown that if an equilibrium ξ_j has expanding directions x_k and x_m , with expanding eigenvalues $e_{jk} > e_{jm}$, then in the limit of low noise, (i.e. as $\zeta \rightarrow 0$) the transition probability from ξ_j to ξ_m scales as

$$\pi_{j,m} = O\left(\zeta^{\frac{e_{jk}}{e_{jm}} - 1}\right).$$

4.3 Lift-off and memory effects for transitions between equilibria

For some applications, it may be desirable for a network to have memory effects, i.e. where the transition probabilities depend on the recent history of the trajectory taken through the network. It is possible to create memory effects in a noisy heteroclinic network that has ‘lift-off’, an effect noted by Stone, Armbruster and Kirk [33, 3] that appears for attracting heteroclinic (but not homoclinic) networks with low amplitude additive noise.

Lift-off is a property of the distribution of the coordinates of a trajectory as it enters a neighborhood of an equilibrium such that the probability distribution of coordinates of trajectories in a Poincaré section becomes multi-modal and is a mechanism by which the transition probabilities can gain memory. We demonstrate the idea of lift-off in Figure 11, which shows the Poincaré section $H_2^{\text{in},1}$ for the schematic in Figure 10, and a representation of the distribution of the coordinates of the incoming trajectories. If no lift-off occurs, then the distribution of coordinates is approximately Gaussian and centered at zero, as shown on the left. If lift-off occurs, then the distribution of a particular coordinate is no longer Gaussian and may include several peaks, as shown on the right.

The transition probabilities can be thought of as intersections of the distribution of incoming trajectories with the set of points that progress to a given equilibrium at the next epoch. As lift-off changes this area of intersection and depends on the previously epochs,

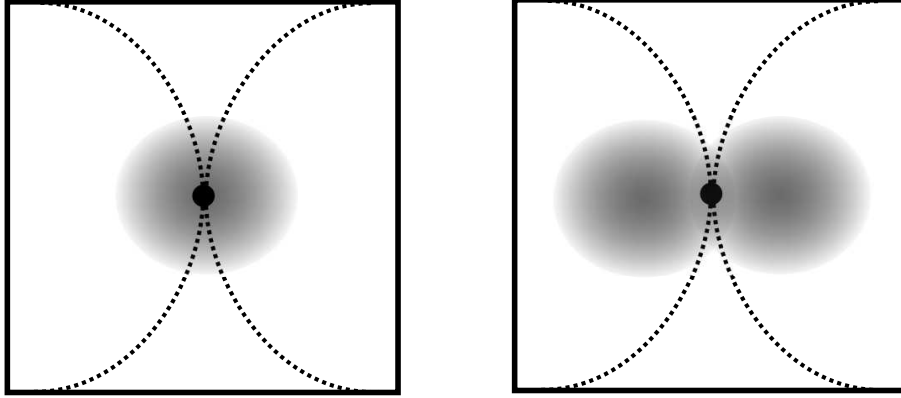


Figure 11: Schematic diagram showing the Poincaré section $H_2^{\text{in},1}$ for the connection from ξ_1 to ξ_2 shown in Figure 10, projected onto the x_3, x_4 plane. The dotted lines indicate the dividing line between regions of trajectories which travel next to ξ_3 or ξ_4 . The shaded ellipse represents the distribution of coordinates of trajectories as they pass through the section where the darker regions are visited more often, while the black dot shows the location of the noise-free connection. The left panel shows a case in which the incoming distribution of the x_3 and x_4 coordinates is approximately Gaussian and centered at zero. The right panel shows a case with lift-off in the x_3 direction and consequently a multi-modal distribution.

lift-off can affect transition probabilities and make them depend on memory. The occurrence of lift-off depends on the eigenvalues of the equilibria, and hence can be controlled. Note that a sufficient condition for lift-off *not* to occur is that all contracting eigenvalues are greater than all expanding eigenvalues.

It was shown in [33] that if the network contains a heteroclinic connection $\xi_j \rightarrow \xi_k$, then there is the possibility of creating lift-off in the x_j direction as the trajectory exits a neighbourhood of the equilibrium ξ_k . If this lift-off can be maintained until the trajectory next approaches an equilibrium with an unstable manifold which connects to ξ_j , the probability of visiting ξ_j will be higher than if the lift-off had not occurred.

Stone and Armbruster [33] compute conditions on the eigenvalues for lift-off to occur. They show that, if there exists a connection $\xi_j \rightarrow \xi_k$, lift-off will occur at ξ_k (that is, in the distribution of coordinates of trajectories leaving a neighbourhood of ξ_k) in the x_j direction if

$$\frac{c_{kj}}{e_{kl}} < 1,$$

where x_l is the expanding direction at ξ_k . That is, if this condition is satisfied, then the distribution of the x_j coordinate, as the trajectory exits a neighbourhood of ξ_k , is not centered about zero but instead is centered about a point whose distance from the connection scales as

$$\zeta \frac{c_{kj}}{e_{kl}}.$$

in the low noise limit $\zeta \rightarrow 0$. Since the distribution of the noise is symmetric about zero,

the lift-off has equal probability of occurring in either the positive or negative direction. This means that the distribution becomes multimodal. Since the global flow only scales the distribution of coordinates by an amount that is order one in ζ , the incoming distribution to the next equilibrium will scale similarly.

It is straightforward to extend the calculations of [33] to a sequence of heteroclinic connections between equilibria $\xi_j \rightarrow \xi_k \rightarrow \xi_l \rightarrow \xi_m$. That is, we already have conditions that the distribution of the x_j coordinate as it exits a neighbourhood of ξ_k is not centered about zero, and we can extend the computations to give conditions that the distribution of the x_j coordinate as the trajectory exits a neighbourhood of ξ_l is similarly not centered about zero.

These calculations are a simple extension of those given in [33], and it can be shown that

$$\frac{c_{kj}}{e_{kl}} + \frac{c_{lj}}{e_{lm}} < 1, \quad (18)$$

implies that the distribution of the x_j coordinate as the trajectory exits a neighbourhood of ξ_l is centered about a point whose distance from the connection scales as

$$\zeta \frac{c_{kj} + c_{lj}}{e_{kl} + e_{lm}}.$$

in low noise limit $\zeta \rightarrow 0$. Again, since the global flow only scales coordinates by an order one amount, the incoming distribution of the x_j coordinate at equilibrium ξ_m will then also not be centered at zero. Thus, if condition (18) on the eigenvalues is satisfied, and ξ_m has an unstable manifold (of dimension two or more) which includes the direction towards ξ_j , then the proportion of trajectories which go towards ξ_j will be higher than if lift-off had not occurred.

Note that lift-off occurs only if the contracting eigenvalues are small enough, that is, if there is not enough contraction in the x_j direction at ξ_l to ‘squash’ the lift-off. Further extension of these calculations will give conditions on lift-off being maintained over longer sequences of equilibria, and hence the possibility of longer term memory. However, as lift-off requires sufficiently small contracting eigenvalues “on average”, and stability of the network requires sufficiently large contracting eigenvalues, we conjecture that any memory effects present must appear of a sequence of epochs that is smaller than the longest cycle within the network.

5 Example: simplex realisation of the decision graph

In this section, we numerically investigate the properties described in Section 4 for the decision graph realised using the simplex realisation of Section 3.1. Note that whilst it is comparatively easy to check for the presence of memory in solution trajectories, checking for the *absence* of memory is difficult. That is, although we can perform statistical tests to show that a trajectory may be zeroth order rather than first order Markov, checking for the absence of all long-term memory effects is hard. Whilst we expect that such long term memory (i.e. over sequences of more than two equilibria) is possible, it seems unlikely to

occur in the absence of short term memory except in special cases. A detailed study of long-term memory is beyond the scope of this paper and so in the following sections we discuss only the presence or absence of short-term memory.

5.1 Transition probabilities in the decision graph

First, we consider a particular case where we predict that the network is memoryless, that is, there is no lift-off, and so the incoming distributions of all coordinates at each equilibria are Gaussians centered at zero.

We confirm the scaling given in Section 4.2 for the transition probabilities in the decision graph by performing numerical integrations at various noise levels. We use parameters

$$c_{jk} = t_{jk} = 2, e_{12} = 2, e_{24} = e_{36} = e_{48} = 1.85, \\ e_{23} = e_{35} = e_{47} = 1.99, e_{51} = e_{61} = e_{71} = e_{81} = 1.98$$

and integrate for noise levels $\zeta = 10^{-5}, 10^{-6}, 10^{-7}, 10^{-8}, 10^{-9}, 10^{-10}, 10^{-11}$. For each noise level we integrate until the number of passes the trajectory makes through ξ_1 is 5000.

At each noise level, we measure the number of visits the trajectory makes to the equilibria ξ_4, ξ_6 and ξ_8 , as a proportion of the number of visits to ξ_2, ξ_3 and ξ_4 respectively. In terms of the definitions given in Section 4, these are approximations of

$$\frac{\pi_4}{\pi_2}, \quad \frac{\pi_6}{\pi_3}, \quad \text{and} \quad \frac{\pi_8}{\pi_4}$$

respectively, using approximations (13) and $k = 5000 \times 4 = 20,000$ with standard errors. These proportions are plotted against the noise level in Figure 12 on log-log axes. Best fit linear regressions for each of these lines have slopes 0.072, 0.074 and 0.073 respectively; note that the expected slope for all three is $\frac{1.99}{1.85} - 1 = 0.076$. This means that in the low noise limit $\zeta \rightarrow 0$ we predict that almost all itineraries follow the cycle

$$\cdots \rightarrow \xi_1 \rightarrow \xi_2 \rightarrow \xi_3 \rightarrow \xi_5 \rightarrow \xi_1 \rightarrow \cdots$$

corresponding to simply selecting the most unstable expanding direction at each equilibrium.

Note that without knowing details about the global part of the flow, it is not possible to predict the precise transition probabilities for particular noise levels and eigenvalues, the best we can do is get the scaling for low noise. However, if a particular transition probability is desired, a little experimentation and alteration of noise levels can achieve this.

5.2 Numerical examples of memory effects for the decision graph

We now demonstrate an example of how lift-off as described in Section 4.3 can create memory in a network. We use the decision graph with the simplex realisation as an example, and induce lift-off in the x_3 direction at ξ_5 . We give two examples: in the first the lift-off is maintained until the trajectory reaches ξ_2 . In this example, we expect that trajectories which visit ξ_5 will then have a larger probability of visiting ξ_3 on the next circuit of the network

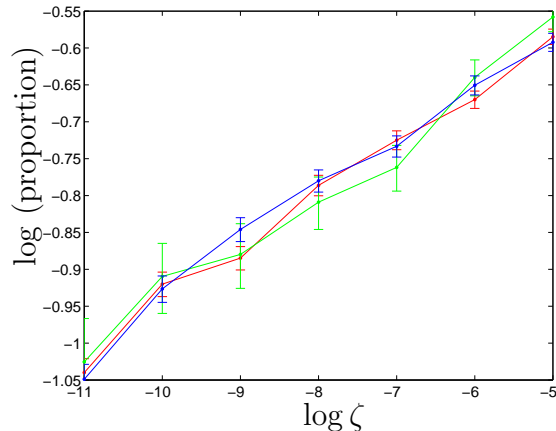


Figure 12: For the decision graph and simplex realisation, the figure shows number of visits the trajectory makes to the equilibria ξ_4 (red), ξ_6 (blue) and ξ_8 (green), as a proportion of the number of visits to ξ_2 , ξ_3 and ξ_4 respectively, for various noise levels ζ ; the bars show standard errors. For details and parameter values, see text.

than they would have done had they previously visited ξ_6 , ξ_7 or ξ_8 . In the second example, we exhibit parameters such that the lift-off is not maintained at ξ_2 , and consequently there are no memory effects.

The particular structure of the decision graph can be used to simplify our analysis. Note that the sequence of equilibria visited (that is, the sequence $i_k, k = 1, \dots$) can be deduced by only recording which of the equilibria ξ_5 , ξ_6 , ξ_7 and ξ_8 was visited by the trajectory on each ‘loop’ around the network, and we therefore focus just on transitions between this subset of equilibria.

Specifically, we consider the sub-itinerary \tilde{i}_k of (10) that consist of visits to $\{5, 6, 7, 8\}$ by choosing the minimal strictly increasing set of indices $0 < \ell_k$ such that

$$\tilde{i}_k = i_{\ell_k} \in \{5, 6, 7, 8\}.$$

Observe that from the structure of the decision graph in Figure 5 we expect that $\ell_{k+1} - \ell_k = 4$ on most occasions, and note also that we can have one-cycles (non-trivial transitions from j straight back to j) in this induced graph. Using the corresponding definitions for proportions of visits and transition probabilities between this subset of equilibria we define for any sequence of $j_k \in \{5, 6, 7, 8\}$:

$$\tilde{\mathcal{P}}(j_1, \dots, j_m) = \text{Prob}(\tilde{x}, \tilde{\vartheta}) : \tilde{i}_\ell = j_\ell \text{ for } \ell = 1, \dots, m, \text{ for the trajectory starting at } (\tilde{x}, \tilde{\vartheta}). \quad (19)$$

Analogously to (13) we can compute this from a typical trajectory as

$$\tilde{\mathcal{P}}(j_1, \dots, j_m) = \lim_{k \rightarrow \infty} \frac{1}{k} \#\{0 \leq \ell < k : \tilde{i}_{\ell+n} = j_n \text{ for } n = 1, \dots, m\} \quad (20)$$

and so define the probability $\tilde{\pi}_j$ of the first epoch being close to $j \in \{5, 6, 7, 8\}$:

$$\tilde{\pi}_j = \tilde{\mathcal{P}}(j).$$

Assuming $\tilde{\pi}_j > 0$, we define the transition probability between vertices ξ_{j_1} and ξ_{j_2} by:

$$\tilde{\pi}_{j_1, j_2} = \frac{\tilde{\mathcal{P}}(j_1, j_2)}{\tilde{\pi}_{j_1}}. \quad (21)$$

We perform numerical integrations using the Heun method (with timestep 0.01) for the cases with and without memory, and observe the effects in two ways. First, we compute the observed number of transitions between this subset of equilibria to estimate the matrix of probabilities $\tilde{\pi}_{j_1, j_2}$ for $j_1, j_2 = 5, \dots, 8$. A χ -squared test is used to determine whether or not the probabilities of visiting each of the equilibria are independent of the previous state of the trajectory. We also directly observe the distribution of coordinates as trajectories enter and leave neighbourhoods of equilibria. From the results in Section 4.3, we know that lift-off at ξ_5 in the x_3 direction occurs if $\frac{c_{53}}{e_{51}} < 1$. This lift-off will still exist after the trajectory when the trajectory enters a neighbourhood of ξ_2 if $\frac{c_{53}}{e_{51}} + \frac{c_{13}}{e_{12}} < 1$. We first perform an experiment where both conditions are satisfied, and then an experiment where only the first condition is satisfied.

5.3 An example with memory

We choose parameters so that

$$\frac{c_{53}}{e_{51}} < 1 \quad \text{and} \quad \frac{c_{53}}{e_{51}} + \frac{c_{13}}{e_{12}} < 1,$$

and integrate the equations for the decision graph simplex realisation, with noise $\zeta = 10^{-5}$, and for total time 200,000, giving a sequence of 31,568 epochs and so $k_c = 31,568/4$ cycles. The parameters used are $c_{jk} = 2$ for all j, k except for $c_{13} = 0.8$, $c_{53} = 0.8$, $t_{jk} = 2$ for all j, k , $e_{12} = 2$, $e_{24} = e_{36} = e_{48} = 1.9$, $e_{23} = e_{35} = e_{47} = 1.99$, and $e_{51} = e_{61} = e_{71} = e_{81} = 1.98$.

For a randomly chosen initial condition and noise path, the following matrix is the observed number of transitions between equilibria ξ_5, ξ_6, ξ_7 and ξ_8 in the reduced system:

$$\begin{pmatrix} 2829 & 1478 & 143 & 76 \\ 1166 & 643 & 379 & 181 \\ 355 & 158 & 88 & 57 \\ 176 & 90 & 48 & 25 \end{pmatrix}$$

which gives an approximation to the matrix $\tilde{\pi}_{j_1, j_2}$:

$$\begin{pmatrix} \tilde{\pi}_{5,5} & \tilde{\pi}_{5,6} & \tilde{\pi}_{5,7} & \tilde{\pi}_{5,8} \\ \tilde{\pi}_{6,5} & \tilde{\pi}_{6,6} & \tilde{\pi}_{6,7} & \tilde{\pi}_{6,8} \\ \tilde{\pi}_{7,5} & \tilde{\pi}_{7,6} & \tilde{\pi}_{7,7} & \tilde{\pi}_{7,8} \\ \tilde{\pi}_{8,5} & \tilde{\pi}_{8,6} & \tilde{\pi}_{8,7} & \tilde{\pi}_{8,8} \end{pmatrix} \approx \begin{pmatrix} 0.6251 & 0.3266 & 0.0316 & 0.0168 \\ 0.4922 & 0.2714 & 0.1600 & 0.0764 \\ 0.5395 & 0.2401 & 0.1337 & 0.0866 \\ 0.5192 & 0.2655 & 0.1416 & 0.0737 \end{pmatrix}$$

where standard errors are of order $1/\sqrt{k_c}$, i.e. ± 0.01 . Hence, if a trajectory visits ξ_5 , it is more likely to go via ξ_3 (and hence to ξ_5 or ξ_6 rather than ξ_7 or ξ_8) on its next circuit around the network. A χ -squared test confirms this by rejection the null hypothesis that the transition probabilities are independent of the previous equilibrium visited.

Figure 13 demonstrates the lift-off by showing distributions of the x_1 and x_3 coordinates as the trajectory passes from equilibria ξ_5 to ξ_1 . We plot the coordinates as the trajectory crosses appropriate Poincaré sections: H_5^{in} , H_5^{out} , H_1^{in} and H_1^{out} , using $h = 0.1$. On H_5^{in} , x_3 is of order h , and the coordinate of interest is x_1 . Figure 13(a) shows this to clearly be approximately Gaussian centered at zero. Figures 13(b) to (d) show distributions of the x_3 coordinate on H_5^{out} , H_1^{in} and H_1^{out} . In each of these three cases, the distribution is clearly not Gaussian, but instead is approximately the sum of two Gaussians shifted to the right and left of zero. This demonstrates the lift-off of the trajectory both in the positive and negative directions, on different loops around the network.

The conditional distributions in Figure 14 clearly show the memory effect via a scatter plot of the x_3 and x_4 coordinates of the trajectory on H_1^{out} , conditional on the previously visited cycle. The conditional distribution peaks are clearly different as only those that went past ξ_6 at the last cycle display lift-off.

5.4 An example with no memory

We now explore a second example with same parameters as in Section 5.3, except that $c_{13} = 1.3$, so

$$\frac{c_{53}}{e_{51}} < 1 \quad \text{but} \quad \frac{c_{53}}{e_{51}} + \frac{c_{13}}{e_{12}} > 1.$$

This means we expect to see lift-off in the x_3 direction at ξ_5 , but after the trajectory has passed ξ_1 this lift-off will be ‘squashed’ and hence the distribution of the x_3 coordinate as the trajectory enters a neighbourhood of ξ_2 will be Gaussian again. Thus the network will have no memory. Again, we integrate for time 200,000, which gives a sequence of $k_c = 30,892/4$ cycles.

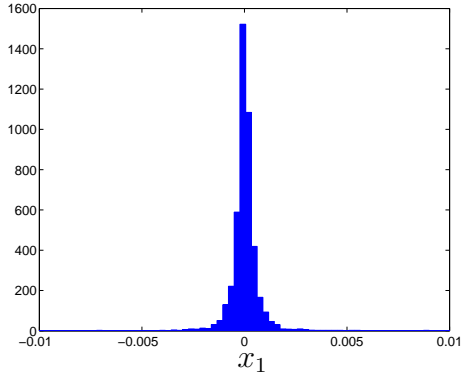
For an example run, the observed number of transitions between equilibria ξ_5 , ξ_6 , ξ_7 , and ξ_8 is:

$$\begin{pmatrix} 1686 & 957 & 660 & 348 \\ 896 & 446 & 371 & 186 \\ 697 & 328 & 254 & 142 \\ 373 & 167 & 136 & 76 \end{pmatrix}$$

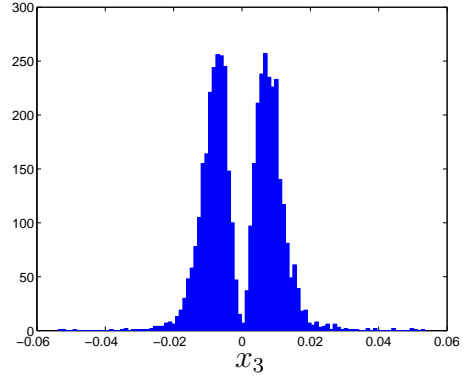
which gives the following approximation to $\tilde{\pi}_{j_1, j_2}$:

$$\begin{pmatrix} \tilde{\pi}_{5,5} & \tilde{\pi}_{5,6} & \tilde{\pi}_{5,7} & \tilde{\pi}_{5,8} \\ \tilde{\pi}_{6,5} & \tilde{\pi}_{6,6} & \tilde{\pi}_{6,7} & \tilde{\pi}_{6,8} \\ \tilde{\pi}_{7,5} & \tilde{\pi}_{7,6} & \tilde{\pi}_{7,7} & \tilde{\pi}_{7,8} \\ \tilde{\pi}_{8,5} & \tilde{\pi}_{8,6} & \tilde{\pi}_{8,7} & \tilde{\pi}_{8,8} \end{pmatrix} \approx \begin{pmatrix} 0.4618 & 0.2621 & 0.1808 & 0.0953 \\ 0.4718 & 0.2349 & 0.1954 & 0.0979 \\ 0.4905 & 0.2308 & 0.1787 & 0.0999 \\ 0.4960 & 0.2221 & 0.1809 & 0.1011 \end{pmatrix}$$

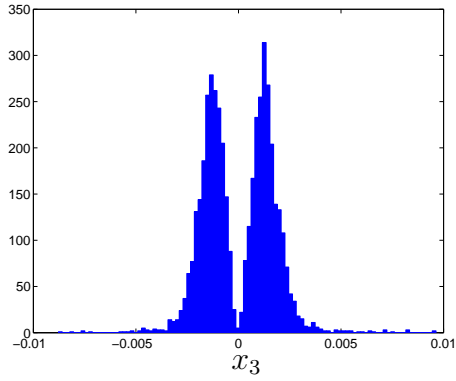
with standard errors of order ± 0.01 . In this case the χ -squared test does not reject the null hypothesis of independence for this transition matrix.



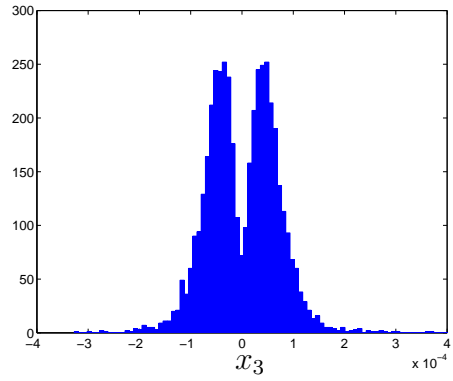
(a) Distribution of x_1 on H_5^{in}



(b) Distribution of x_3 on H_5^{out}



(c) Distribution of x_3 on H_1^{out}



(d) Distribution of x_3 on H_1^{out}

Figure 13: Distributions of coordinates for trajectories which pass from equilibrium ξ_3 to ξ_5 to ξ_1 , in the decision graph. Parameters satisfy $\frac{c_{53}}{e_{51}} + \frac{c_{13}}{e_{12}} < 1$, and so lift-off in the x_3 coordinate is expected and can be seen in the non-Gaussian distributions.

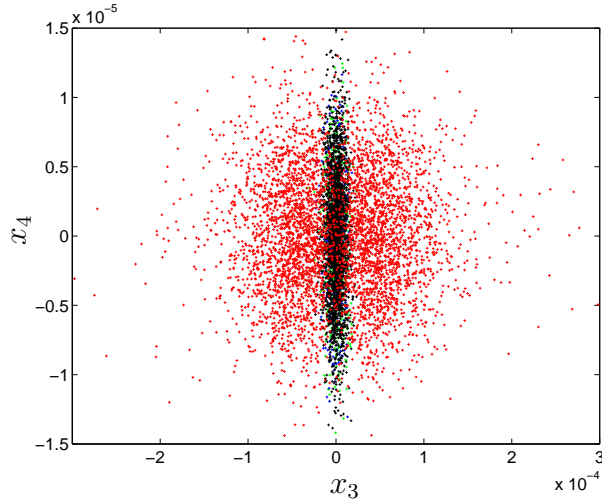


Figure 14: Scatter plot of x_3 and x_4 coordinates showing the condition distribution of trajectories as they exit from ξ_1 on H_1^{out} , for the same trajectory plotted in Figure 13. Trajectories which visited ξ_5 on the previous loop around the network are coloured red, while trajectories which previously visited ξ_6 , ξ_7 and ξ_8 are coloured black, blue and green (note that the blue and green points are plotted under the black points, making them difficult to see). Lift-off in the x_3 direction for those points that are coloured red can clearly be seen to cause memory.

Figure 15 shows the distribution of the x_3 coordinate (for those trajectories which pass through ξ_5) on H_1^{in} and H_1^{out} . It can clearly be seen that there is lift-off in the x_3 direction before ξ_1 , but on exiting, the distribution has returned to being approximately Gaussian with zero mean; this is shown as a scatter plot in Figure 16 where, in contrast to the case in Figure 14, the distributions do not appear to depend on previous epochs and so the system is memoryless.

6 Discussion

We have shown that, in principle, any finite directed graph can be realised as an embedded attracting robust heteroclinic network for a coupled cell system. In fact the method should work even if the network is not connected (respectively, if the components are not strongly connected) - in these cases the resulting network will be disconnected (respectively, the asymptotic dynamics will be contained within a subnetwork). There are clearly issues at vertices of high degree for outgoing edges - there must be nontrivial dynamics such as additional saddle points involved in separating trajectories that go to different edges. We have not been able to prove all we would like to concerning the dynamics of the constructed networks; in particular we have not been able to characterize the larger embedding network \tilde{X} for either simplex or cylinder realisations. Nonetheless, our numerical examples suggest that the dynamics is robust enough to be useful as a method for designing dynamical systems

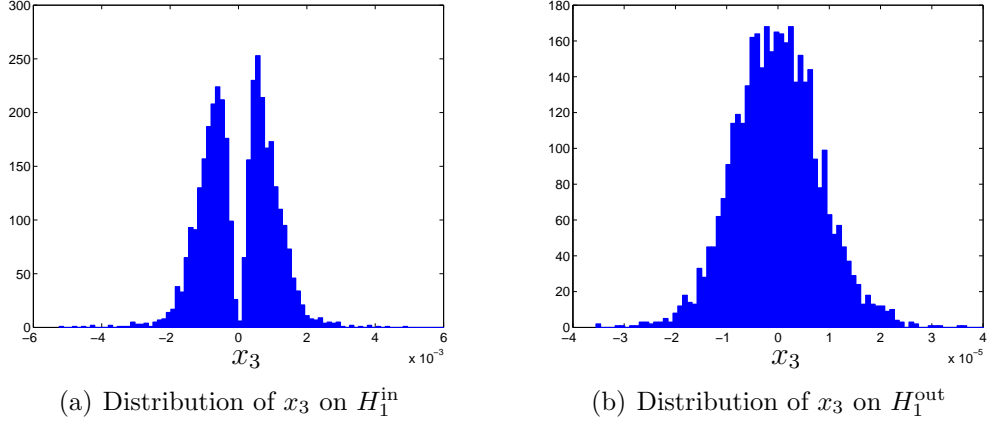


Figure 15: Distributions of coordinates for trajectories crossing H_1^{in} and H_1^{out} , having previously visited ξ_5 . Parameters in this case satisfy $\frac{c_{53}}{e_{51}} < 1$ but $\frac{c_{53}}{e_{51}} + \frac{c_{13}}{e_{12}} > 1$. Lift-off in the x_3 direction can be seen as the trajectory enters a neighbourhood of ξ_1 but it has been compressed by the time the trajectory leaves a neighbourhood of ξ_1 .

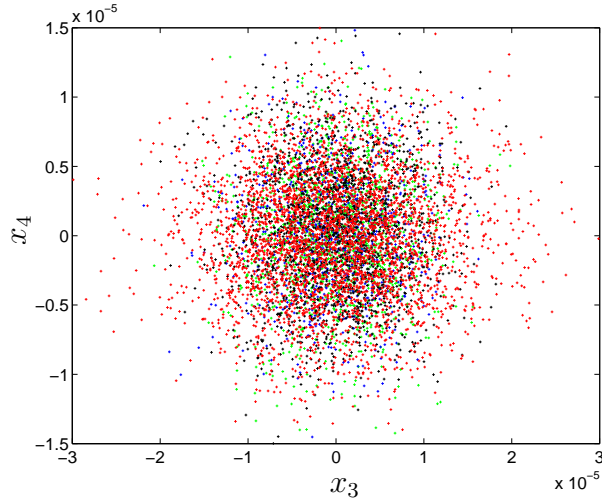


Figure 16: Scatter plot of x_3 and x_4 coordinates showing the condition distribution of trajectories as they exit from ξ_1 on H_1^{out} , for the same trajectory plotted in Figure 15. Trajectories which visited ξ_5 on the previous loop around the network are coloured red, while trajectories which previously visited ξ_6 , ξ_7 and ξ_8 are coloured black, blue and green. The distributions of all four sets of trajectories appears to be identical, indicating that there is no memory.

with specific graph attractors in phase space.

The noisy networks will be qualitatively robust to perturbations of parameters up to the point where the network undergoes one of a number of possible bifurcations. As suggested by [9, Figures 14 and 15], one bifurcation that will destroy the heteroclinic network and replace it with an excitable network is when one of the vertices becomes locally stable via a pitchfork bifurcation. At such a pitchfork bifurcation, new saddles will be created whose stable manifolds give excitability thresholds - the connections no longer form a chain recurrent set. Another bifurcation will be a resonance that causes the network to lose stability and results in bifurcation of periodic or other attractors from the network. Unlike the heteroclinic networks considered in [15] we do not expect the networks we consider to appear in primary bifurcations from a trivial state - the symmetries we require for robustness do not act transitively on the set of cells.

We do not expect it will be possible to remove the constraint of no 1-cycles that is required by both constructions, except by lifting any graph with 1-cycles to a larger graph that has no 1-cycles. Because of the constraints of transversality, it is not possible for a dissipative system to have a robust homoclinic cycle to a hyperbolic equilibrium, even if the perturbations are constrained to leave certain subspaces invariant. This is because, within the smallest invariant subspace that contains the connection, there will be a transverse intersection that will be broken by small perturbations.

6.1 Residence times near vertices

For a given typical trajectory we define $I_k(j)$ to be the probability that the trajectory is near ξ_j , that is:

$$I_k(j) = \{l \in \{1, \dots, k\} : i_l = j\} \quad \text{and define} \quad N_k(j) = \#\{I_k(j)\}$$

to be the cardinality of $I_k(j)$, i.e. the number of times a trajectory visits the equilibria ξ_j . This can be used to define the *mean residence time* of the trajectory close to the vertex ξ_j :

$$\tau_j = \lim_{k \rightarrow \infty} \frac{1}{N_k(j)} \sum_{l \in I_k(j)} T_l.$$

In this paper we do not examine properties of τ_j except to remark that this is a very interesting and nontrivial question [9]. It was shown in [33] that the mean passage time of a trajectory past an equilibrium (that is, the length of time during which the trajectory remains in a neighbourhood of an equilibrium) for a heteroclinic cycle in the limit of low noise $\zeta \rightarrow 0$ is

$$\tau_j = \frac{1}{\lambda_j} \log \left(\frac{h}{\zeta} \right) + O(1)$$

where λ_j is the largest expanding eigenvalue for equilibrium ξ_j and h is the neighbourhood size. For a heteroclinic network in the presence of noise, the passage time will be similar, though if there are several expanding eigenvalues, it is possible that the transition time may depend on the exit route, in this way the mean residence times may in fact encode the perturbations of the system to inputs with non-zero mean [36].

6.2 Relevance to computational systems

The constructions presented here should give flexible tools for designing coupled cell systems that realise finite-state computational systems. They suggest new ways to adapt coupled cell systems to “learn” networks, by modifying parameters [9]. An interesting question is whether either of the constructions can be adapted to explain neural computations that proceed on a dynamical basis. Simple network models of coupled neurons [6] can give rise to various structures of heteroclinic network in phase space and we believe the constructions here can be used to extend this to more realistic decision making systems.

In applications, it may also be desirable to have control over the statistical properties of how a trajectory moves around the heteroclinic network. We may also want to control the occurrence of memory effects in the network where transition probabilities depend on the recent history of the route taken around the network; the methods in Section 4 offers a route for doing this.

There are many further questions that one could ask about the resulting designed networks - these include, for example: Can one design a network that has not only the given network structure, but also a specified set of average residence times and/or transition probabilities (with or without “memory”)? This is likely to be an interesting and challenging problem where inclusion of anisotropic noise may be important [10]. Other questions concern the limits on memory effects within such systems and behaviour of the system at larger noise levels.

Finally, we should comment that the realisation methods here are not very efficient in terms of the dimension of the system - the phase space dimension scales linearly with the number of vertices (resp. edges) for the simplex (resp. cylinder) realisations. This could be a barrier to using these results as a paradigm for neural computation where the encoding may be very dense. By contrast, the number of vertices for “odd graph” networks is very efficient - using only $n = 2k + 1$ globally coupled phase oscillators we can find heteroclinic networks with $n!/(k!(k + 1)!)$ vertices [8]. However, as previously highlighted, the latter networks are not easily adaptable to computation because of their complex topology and large number of symmetries. Hence it is an interesting challenge to find a way to robustly realise a noisy heteroclinic attractor in a “minimal” dimension network.

Acknowledgments

This work was started during discussions at the Mathematical Biosciences Institute (MBI) in Ohio in 2011 and continued during a visit of CMP to Exeter. We thank the MBI, the Royal Society and the University of Auckland for partial support for this research and thank Marc Timme, Fabio Neves, Chris Bick, Mike Field, Vivien Kirk and Ilze Ziedins for very interesting conversations in relation to this work. Finally, we thank the referees for advice that helped us to improve the exposition in the paper.

References

- [1] Aguiar MAD, Ashwin P, Dias APS, Field M: Dynamics of coupled cell networks: Synchrony, heteroclinic cycles and inflation. *J. Nonlinear Sci.* 2011, **21**:271–323.
- [2] Arnold, L: Random Dynamical Systems Springer monographs in mathematics. Springer-Verlag, Berlin 1998.
- [3] Armbruster D, Kirk V and Stone, E: Noisy heteroclinic networks. *Chaos* 2003, **13**(1):71-79.
- [4] Ashwin P, Field M: Heteroclinic networks in coupled cell systems. *Arch. Ration. Mech. Anal.* 1999, **148**(2):107–143.
- [5] Ashwin P, Borresen J: Encoding via conjugate symmetries of slow oscillations for globally coupled oscillators. *Phys. Rev. E* 2004, **70**(2):026203.
- [6] Ashwin P, Karabacak O, Nowotny T: Criteria for robustness of heteroclinic cycles in neural microcircuits. *J. Math. Neurosci.* 2011, **1**:13.
- [7] Ashwin P, Lavric A: A low-dimensional model of binocular rivalry using winnerless competition. *Physica D* 2010, **239**:529–536.
- [8] Ashwin P, Orosz G, Borresen J: “Heteroclinic Switching in Coupled Oscillator Networks: Dynamics on Odd Graphs”, *Nonlinear Dynamics and Chaos: Advances and Perspectives, Understanding Complex Systems*, M. Thiel et al. (eds.), 2010, Springer.
- [9] Ashwin P, Orosz G, Wordsworth J, Townley S: Dynamics on networks of clustered states for globally coupled phase oscillators. *SIAM J. Appl. Dyn. Sys.* 2007, **6**(4):728–758.
- [10] Ashwin P, Podvigina O: Noise-induced switching near a depth two heteroclinic network and an application to Boussinesq convection. *Chaos* 2010, **20**(2):023133.
- [11] Ashwin P, and Timme M: Unstable attractors: existence and robustness in networks of oscillators with delayed pulse coupling. *Nonlinearity* 2005, **18**:2035–2060.
- [12] Bakhtin Y: Small noise limit for diffusions near heteroclinic networks. *Dynamical Systems* 2010, **25**:413-431.
- [13] Bick C, Rabinovich MI: On the occurrence of stable heteroclinic channels in Lotka-Volterra models. *Dyn. Syst.* 2010, **25**:97–110.
- [14] Broer H, Efstathiou K, Subramanian E: Heteroclinic cycles between unstable attractors. *Nonlinearity* 2008, **21**:1385-1410.
- [15] Dias APS, Dionne B, Stewart I: Heteroclinic cycles and wreath product symmetries. *Dyn. Stab. Syst.* 2000, **15**:353–385

- [16] Field M: *Lectures on bifurcations, dynamics and symmetry, Volume 356 of Pitman Research Notes in Mathematics Series*. Harlow: Longman 1996.
- [17] Golubitsky M, Stewart I: *The Symmetry Perspective*. Birkhäuser Verlag, Basel 2003.
- [18] Guckenheimer J and Holmes P: Structurally stable heteroclinic cycles, *Math. Proc. Camb. Phil. Soc.* 1988, **103**:189–192.
- [19] Hofbauer J, Sigmund K: *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge 1998.
- [20] Karabacak O, Ashwin P: Heteroclinic ratchets in networks of coupled oscillators. *J. Nonlinear Sci.* 2010, **20**:105–129.
- [21] Kirk V and Silber M: A competition between heteroclinic cycles, *Nonlinearity* 1994, **7**:1605–1621.
- [22] Kirk V, Lane E, Postlethwaite CM, Rucklidge AM, Silber M: A mechanism for switching near a heteroclinic cycle, *Dyn. Syst.* 2010, **25**:323–349.
- [23] Kirk V, Postlethwaite CM, Rucklidge AM, Resonance of heteroclinic networks, *To appear in SIADS*, 2013.
- [24] Komarov MA, Osipov GV, Suykens JAK: Sequentially activated groups in neural networks. *EPL (Europhysics Letters)* 2009, **86**(6):60006.
- [25] Krupa M: Robust heteroclinic cycles. *J. Nonlinear Sci.* 1997, **7**(2):129–176.
- [26] Krupa M and Melbourne I: Asymptotic stability of heteroclinic cycles in systems with symmetry, II *Proc. Roy. Soc. Edinburgh A* 2004 **134A** 1177–1197
- [27] May RM, Leonard WJ: Nonlinear Aspects of Competition Between Three Species. *SIAM J. Appl. Math.* 1975, **29**(2):243–253.
- [28] Neves FS, Timme M: Controlled perturbation-induced switching in pulse-coupled oscillator networks. *J. Phys. A* 2009, **42**:345103.
- [29] Neves FS, Timme M: Computation by Switching in Complex Networks of States. *Phys. Rev. Letts* 2012, **109**:018701.
- [30] Postlethwaite CM and Dawes JHP: Regular and irregular cycling near a heteroclinic network, *Nonlinearity* 2005, **18**:1477–1509.
- [31] Rabinovich MI, Afraimovich VS, Varona P: Heteroclinic binding. *Dynamical Systems* 2010, **25**(3):433–442.
- [32] Rabinovich MI, Volkovskii A, Lecanda P, Huerta R, Abarbanel HDI, Laurent G: Dynamical encoding by networks of competing neuron groups: Winnerless competition. *Phys. Rev. Lett.* 2001, **87**(6):068102.

- [33] Stone E, Armbruster D: Noise and $O(1)$ amplitude effects on heteroclinic cycles. *Chaos* 1999, **9**(2):499–506.
- [34] Stone E, Holmes P: Random perturbations of heteroclinic attractors. *SIAM J. App. Math.* 1990, **50**(3):726–743.
- [35] Seliger P, Tsimring LS, Rabinovich MI: Dynamics-based sequential memory: winnerless competition of patterns. *Phys. Rev. E (3)* 2003, **67**:011905.
- [36] Wordsworth J, Ashwin P: Spatiotemporal coding of inputs for a system of globally coupled phase oscillators. *Phys. Rev. E* 2008, **78**:066203.